CHAPTER 3

CONSTRUCTIVISM: AN ALTERNATIVE THEORETICAL FOUNDATION

> The way we segment the flow of our experience, and the way we relate the pieces we have isolated, is and necessarily remains an essentially subjective matter. Hence, when we intend to stimulate and enhance a student's learning, we cannot afford to forget that knowledge does not exist outside a person's mind. (von Glasersfeld, 1996, p. 5)

> Engaging in practice, rather than being its object, may well be a condition for the effectiveness of learning.
> (Lave & Wenger, 1991, p. 93)

In the previous chapter, I illustrated that the extent to which learners are actively involved in the visualization process has the most significant impact on successful learning outcomes. A logical starting point in the search for an alternative theoretical foundation, then, is to consider pedagogical alternatives that get students *even more involved* in the visualization process.

One obvious alternative is to have students construct their own AVs, rather than having them interact with AVs constructed by someone else. While not motivated by learning theory, a few computer science educators have already explored so-called "animation assignments," in which students are asked to construct their own AV of a given algorithm. For example, Brown (1988, Appendix A) describes animation assignments used in past offerings of Brown University's undergraduate algorithms course. For such assignments, students were provided with a collection of animation routines that they could insert into their programs in order to animate them. The animation routines allowed students to generate animations of their programs that resembled the ones presented in the lectures.

In contrast, Stasko (1997) has developed a high-level interpreted language called *Samba*, which is designed specifically to support students in their algorithm animation-building efforts. In past offerings of his third-year undergraduate algorithms course, Stasko used Samba as the basis for animation assignments. In these assignments, students developed animations of fundamental sorting and graph algorithms "that would help explain the algorithm[s] to a person not familiar with [them]" (p. 6). In his assessment of animation assignments, Stasko is quite positive. He notes that students appear not only to enjoy building their own animations, but also to gain increased understanding as a result.

Recall that the main argument of this thesis is that, in order to harness the pedagogical promise of algorithm visualization, we need to rethink our underlying theory of effectiveness. What theoretical foundation would predict the pedagogical shift both suggested by the meta-analysis of the last chapter, and recently explored by a few computer science educators?

In this chapter, I propose *constructivism* as the theoretical position that best accounts for the pedagogical shift toward getting students more involved in the visualization process. As with EF Theory, there are multiple versions of this theory. I begin by introducing the cognitive version, which, because it focuses on individual knowledge, bears some resemblance to EF Theory. However, I argue that the target skills taught in an algorithms course are inherently *social*. Consequently, an alternative version of the theory, *sociocultural constructivism*, is actually more appropriate because it situates learning in the social environment in which it occurs. In the second part of the chapter, I discuss sociocultural constructivism and identify the implications it has for the design, evaluation, and pedagogical use of algorithm visualization technology. As we shall see, if we are to take sociocultural constructivism seriously, we need to go beyond having students construct their own AVs: We need to encourage students to participate within a community that *practices algorithms*, which fundamentally entails the construction, presentation, and discussion of AVs. To explore and validate the potential of this theoretical shift requires a naturalistic field study—an endeavor I take up in Chapter 4.

3.1 Cognitive Constructivism

Inspired by Piaget's research into childhood development (see, e.g., Gruber & Voneche, 1977), cognitive constructivism encompasses both an epistemological framework and learning theory that differ markedly from those of the EF view. Like the EF view, cognitive constructivism views knowledge at the level of the individual learner. However, rather than regarding knowledge as representations of an objective reality that people carry around in their heads, cognitive constructivism asserts that there is no absolute knowledge. Instead, cognitive constructivism holds that individuals *construct* their own individual knowledge "out of the bewildering array of sensations which have no order or structure besides the explanations. . .which [humans] fabricate for them" (Hein, 1991).

This epistemological position gives rise to a theory of learning that differs fundamentally from that of EF Theory. In particular, it implies that knowledge cannot merely be transferred, in the EF sense, from person to person by means of visual representations. Instead, knowledge must be individually (re-)constructed by the individual. Learning as a knowledge construction process, as well as the way in which this process changes as a learner grows, form key assumptions of the cognitive constructivist view. Below, I discuss these assumptions, as well as their implications for the design, evaluation, and pedagogical use of AV technology.

3.1.1 Key Assumptions

3.1.1.1 Learners Actively Construct Their Own Knowledge

The first key assumption of the cognitive constructivist view is an epistemological one, and is perhaps best understood by way of contrast with EF Theory. On the EF View, there exists a deterministic relationship between an objective reality and individual knowledge about that reality. On the cognitive constructivist view, in contrast, knowledge is an individual creation; it is constructed by each learner for herself. There is thus no notion of absolute knowledge that is somehow represented in people's heads.

How do learners construct their own knowledge? According to cognitive constructivism, in the process of knowledge construction, learners assimilate their experiences in the world into their existing understandings. In other words, learners actively construct new understandings by interpreting new experiences within the context of what they already know (see, e.g., Resnick, 1989).

3.1.1.2 Stages of Intellectual Development: Concrete to Abstract

A second key assumption of cognitive constructivism follows from Piaget's theory of childhood development (see, e.g., Gruber & Voneche, 1977). Piaget's theory postulates that children's intellectual development progresses through different stages, beginning with concrete thinking, and ultimately advancing to formal thinking. These stages are held to be genetic or innate; they are not formally taught. Moreover, they have implications for the ways in which learners construct meanings. When learners are in the concrete stage of thinking, which is well under way at age 6, their conceptual structures—that is to say, the knowledge into which they assimilate their experiences—are firmly grounded in the physical world. Later on, in the formal stages of thinking (which begins around age 12), learners' conceptual structures are more abstract, enabling them to construct knowledge through more abstract ways of thinking. For example, they may use the concept of variables to reason about the possible states of a physical system, or they may use meta-knowledge (knowledge about knowledge) to reflect on their own thought processes (see, e.g., Papert, 1980, ch. 1).

3.1.2 Implications for Pedagogy

What implications do these assumptions have for the ways in which computer science educators ought to use AV as a pedagogical aid? The first assumption implies that instructors should not use AV technology to facilitate lectures or demonstrations in which students are members of a passive audience. Instead, the theory suggests that computer science educators ought to use AV technology as a means of getting students actively involved in their own meaning-construction processes; AV technology should be a basis for *active learning*.

Active learning, on the cognitive constructivist view, fundamentally entails self-directed engagement with one's environment. Here, environment is not meant in the literal sense, as one's physical surroundings. Rather, in the cognitive constructivist sense, environment is quite subjective, encompassing the "cognitive and perceptual structures" that one has abstracted and constructed from one's experiential world (von Glasersfeld, 1996, pp. 4–5). Active learners choose their own problems based on their interests and motivation; construct conceptual and physical objects to represent the problems; manipulate the objects; observe the effects; and, through this process, construct new understandings. Along the way, learners are bound to formulate false theories—that is, incorrect understandings of what is to be learned. However, according to cognitive constructivism, the development of such false theories is an important step in the knowledge-construction process. In fact, the theory suggests that it is foolish to force-feed learners the correct theories; students must ultimately *invent* the correct theories for themselves.

If computer science educators are to take active learning seriously, they should encourage students to use AV technology as a tool for discovering for themselves how algorithms work. In the previous chapter, I showed a trend, in AV-based pedagogical approaches, toward "active viewer involvement": having learners explore a visualization by selecting their own input data sets and observing the execution of the animation for those data sets. While this learning paradigm is certainly a step in the right direction, cognitive constructivism would hold that it does not go far enough, because it fails to involve learners in the process of defining and constructing the objects of the visualization. Indeed, the cognitive constructivist view would maintain that, in order to reap the most learning benefits from an AV, students must construct the AV for themselves. Note that this is precisely the case in the AV construction assignments described above.

The second assumption, which follows from Piaget's child development theory, establishes a theoretical basis for the pedagogical use of AVs. Papert (1980), among others, has observed that the concrete and formal stages of learning of which Piaget speaks can be bridged through computer-implemented models that stand in analogy to the abstract concepts to be learned. As Papert puts it, "computer models [seem]. . .able to give concrete form to areas of knowledge that. ..[appear]. . .so intangible and abstract" (Papert, 1980, p. 23). On the cognitive constructivist view, then, AVs are pedagogically valuable precisely because they model an abstract process in terms of concrete objects, thus establishing a basis on which students can progress from the construction of concrete knowledge (AVs as physical objects) to the construction of abstract knowledge (AVs as general models of algorithms).

3.1.3 Implications for Technology Design

As discussed in Chapter 2, the dominant user model for AV software encourages a division of labor between instructors and students. Instructors are generally the ones who design and implement AVs, and learners are generally the ones who view and interact with AVs. Given its position that learners actively construct their own knowledge, cognitive constructivism suggests a break from this model: AV technology should be designed for learners, not instructors. In other words, AV technology should be designed such that learners are the ones who design, implement, view, and interact with AVs.

What, specifically, does this implication have to say about design? The fact that learners are the ones who design and implement AVs suggests that the models that AV systems define for AV specification must be carefully adapted to the needs and abilities of learners. In many extant AV systems (e.g., Brown, 1988; Roman, Cox, Wilcox, & Plun, 1992; Stasko, 1989), AV creation requires the use of relatively low-level graphics packages. While one would expect an expert instructor to be able to handle low-level graphics, it would be unreasonable to expect learners to be able to program an AV in terms of low-level graphics.

To empower learners to construct their own AVs, AV systems should support AV creation environments that consists of objects—both physical and conceptual—with which learners are likely to be familiar. Indeed, according to the cognitive constructivist view, unless learners are able to represent algorithms in terms of familiar objects, they will be unable to bridge the gap between the concrete and the abstract. As a result, AV technology will lose its value as a tool with which learners can think, experiment, and construct new understandings.

3.1.4 Implications for Effectiveness Evaluation

According to cognitive constructivism, in order to set up the best possible conditions for learning, effective learning technology should make available the right conceptual and physical tools—those that will make it possible for learners to build their own understandings. In the case of algorithms, what would the desired understandings be? That is, what knowledge and skills would one expect learners to take away from their experiences with AV technology?

Recall that EF Theory sees knowledge as symbolic structures in the heads of individuals. It is common, in fact, for EF-minded psychologists to distinguish between two kinds of symbolic knowledge: "declarative" knowledge (facts), and "procedural" knowledge (skills). As we saw in Chapter 2, the post-tests used in past studies of AV effectiveness were designed specifically to measure learners' acquisition of these two forms of knowledge. In the case of an algorithm, declarative knowledge might include general statements about the algorithm's best-case or worst-case behavior, and procedural knowledge might equate to an ability to provide the correct output of the algorithm for a given input.

Like EF Theory, cognitive constructivism situates knowledge squarely in the head of individuals. However, cognitive constructivism suggests that a qualitatively different form of knowledge arises out of effective learning environments—a form of knowledge not reducible to knowing a fact or skill. Papert (1980) cites the Logo programming environment as a prime example of a learning environment that supports the process of "getting to know an idea" (p. 136). In the Logo environment, students explore ideas by programming the Turtle—the equivalent of a pen with a heading—to draw. Papert argues that, in contrast to traditional forms of learning, in which learners merely memorize facts and practice skills, learners in the Logo environment "get to know the Turtle" by "exploring what a Turtle can and cannot do" (p. 136). This activity, Papert contends, "is similar to the child's everyday activities, such as making mudpies and testing the limits of parental authority—all of which have a component of 'getting to know'" (p. 136).

Given that constructivist learning encompasses more than the acquisition of knowledge and skills, what are educators interested in evaluating the effectiveness of AV techology left with? Is it possible to measure the extent to which a learner has "gotten to know" an algorithm, in the cognitive constructivist sense? By closely examining a learner's progress over time, a skilled evaluator certainly could get a good sense of the learner's feel for an algorithm. However, it is important to point out two shortcomings of such an evaluation technique. First, it would most likely require significant amounts of time—more time than most evaluators are willing to spend. Second, such an evaluation might be criticized both because it is subjective, and because it is not replicable. Of course, one could allay some critics by making one's evaluation criteria public, and by using multiple evaluators and demonstrating a high level of agreement among them. However, these efforts would only serve to complicate the evaluation, adding to its already costly time requirements.

In the interest of practicality, then, it seems reasonable to consider less complicated, less time-intensive evaluation techniques. Here, the evaluation techniques adopted by EF Theory might serve as crude but acceptable alternatives. Indeed, notice that cognitive constructivism does not preclude the possibility that learners who have "gotten to know" an algorithm, in the constructivist sense, would perform well on conventional tests of facts and skills. As Papert (1980) puts it, learners who work in constructivist learning environments "certainly do discover facts, make propositional generalizations, and learn skills" (p. 136). Hence, if one hopes to perform a practical, replicable effectiveness evaluation of AV technology with respect to the assumptions of cognitive constructivism, conventional tests of skills and knowledge may be a suboptimal, but viable alternative.

3.2 Discussion: Algorithms as a Social Practice

The preceding section sketched out a learning theory that, like EF Theory, is squarely focused on individual knowledge. Learning, according to cognitive constructivism, involves learners' constructing their own knowledge by reorganizing and restructuring their existing knowledge. At first, concrete physical objects are involved in this knowledge construction process. Later on, learners build understandings out of more abstract conceptual objects.

The individualistic understandings emphasized by cognitive constructivism are certainly part of the competence that an algorithms student would ideally develop. On closer inspection of the activities and learning objectives of an algorithms course, however, I argue that "knowing" algorithms encompasses much more than cognitive skills. Rather, I suggest that the "schoolbook algorithms" taught in an undergraduate algorithms course is a fundamentally *social* endeavor; at its heart, it demands interaction with the social world of schooled algorithmaticians. Consequently, I contend that cognitive constructivism proves inadequate as a theoretical framework for interpreting the effectiveness of AV technology. What is required is a theoretical framework that views knowledge, and hence learning, at the level of *communities*.

To defend this position, let me first review the normative learning objectives of a typical undergraduate algorithms course. Upon completing an algorithms course, students would ideally be able to do such things as (a) write down a procedural description of a given algorithm; (b) argue convincingly that an algorithm is correct; (c) analyze the efficiency of an algorithm; and, ultimately, (d) apply the appropriate algorithm to solve a given problem.

Notice that a student's level of competence in these skills cannot be objectively determined by the student's performance on a test. Rather, a student's level of competence is ultimately a matter of agreement among members of a community of expert "algorithmaticians" who are concerned with the conceptual foundations of algorithms. Consider, for example, how one might determine the "goodness" of a procedural description of an algorithm. In the context of a programming environment, a sufficient description is one that, when entered into a computer, actually compiles and executes as expected. In contrast, within the scope of an algorithms course, no programming environment is enlisted as the ultimate judge of the description. Instead, whether a given procedural description is sufficient depends largely on whether those who read or listen to it can understand it—that is, on whether the description communicates the algorithm in understandable terms. This, in turn, depends on whether the description makes use of an established language for describing algorithms—most often some sort of pseudocode.

Likewise, whether a given proof or efficiency analysis is convincing and correct depends largely on whether its audience is convinced and deems it correct; indeed, it has often been said that “proof is a social interaction.” Once again, success depends not only on the way in which a student presents the proof or efficiency analysis, but also on the extent to which the proof or efficiency analysis appropriately makes use of established notation, conceptual tools, and techniques, including loop invariants, induction, and Big-O notation.

Finally, although it may not be obvious, whether a student appropriately applies an algorithm to solve a given problem is also largely a matter of social consensus. Indeed, applying an algorithm is not done in a vacuum, but rather within a social environment in which others must be convinced that the algorithm is the right one for the problem. And without some sort of established conventions with respect to what “an appropriate algorithm” is, it would be impossible to determine what it would mean to choose the “appropriate” algorithm for a given problem.

In short, while the cognitive skills required to supply written answers to test problems are certainly important in an undergraduate algorithms course, I suggest that the competence that algorithms teachers would like students to acquire goes beyond an ability to answer written test problems adeptly. What teachers would really like their students to learn *is an ability to engage with competence in the practice of algorithms*. Here, the notions of “competence” and “practice of algorithms” have distinct meanings; let me briefly clarify them.

On the one hand, competence indicates a command of the conceptual tools, representations, language, and techniques that expert algorithmaticians use and recognize as appropriate. On the other hand, it indicates an ability to use those tools, representations, languages, and techniques in practice—that is, in a way that established algorithmaticians recognize and find convincing. By “practice of algorithms,” I mean the typical activities in which those who are interested in the conceptual foundations of algorithms engage. Key among such activities is *algorithms discourse*—that is, conversations and dialog about algorithms, including how an algorithm works, how efficiently it runs, and why it is correct.

If one accepts the premise that an undergraduate algorithms course is really about training students to become competent members of a community of practice, one realizes the inadequacy of cognitive constructivism’s purely individualistic view of knowledge and learning. In the case of an undergraduate algorithms course, an adequate theoretical framework must account for students’ development as social members of a community that recognizes and values competence at such activities as algorithm explanation, analysis, and proof, all of which have a rich heritage of established techniques, tools, and language. Below, I introduce a more recent version of constructivist theory that, because it views knowledge and learning at the level of the community, would appear to be a more appropriate theoretical framework for analyzing the effectiveness of algorithm visualization technology.

3.3 Sociocultural Constructivism

Spurred by a “growing disillusionment with the individualistic focus” of cognitive theories of learning (Cobb, 1996, p. 34), sociocultural constructivism proposes a fundamental reinterpretation of the constructivist ideas discussed above.[11] Like cognitive constructivism, sociocultural constructivism

---

[11]Sociocultural constructivism has evolved out of the work of a number of social scientists working within different disciplines, including social psychology, anthropology, sociology, and educational research. Not surprisingly, these social scientists have had different agendas, have used different terminology, and have emphasized different aspects of the position. The account of the theory I offer here draws primarily on the influential work of anthropologist Jean Lave and her Ph.D. students (Lave & Wenger, 1991; see also Lave, 1988, Lave, 1993; Lave, 1997; Wenger, 1998). Note that Lave and colleagues label their position *situated learning theory*, not sociocultural constructivism.

views knowledge not as cognitively-instantiated representations of an objective reality, but rather as a constructed entity. However, whereas cognitive constructivism focuses on knowledge at the level of the *individual*, sociocultural constructivism focuses on knowledge at the level of the *community*. According to the former, individuals construct their own knowledge by assimilating their experiences into their existing knowledge. According to the latter, knowledge is socially constructed; it is the collaborative achievement of persons engaged in the practices of a community.

In shifting its epistemological focus from the individual to the community, sociocultural constructivism understandably stresses a markedly different set of assumptions regarding learning. Most significantly, for sociocultural constructivism, learning only makes sense within the context of participating in the ongoing activities of a particular community of practice (Lave & Wenger, 1991); skill and knowledge do not exist independently of such a context. Moreover, sociocultural constructivism views learning in terms of changing participation; to learn is to participate, in increasingly competent ways, within a community. Here, the concept of *identity*—the way one views oneself, and is viewed by others with respect to the community—plays an important role; one's identity undergoes significant changes as one's level of participation changes within the community. Finally, the sociocultural constructivist position emphasizes that *access* and *opportunities for participation* are essential to the learning process. Without access to the practices, resources, and members of the community, and without opportunities to participate in the practices of the community, learning would be impossible.

Below, I elaborate further on sociocultural constructivism. I begin by relating an alternative version of the Allegory of Musica introduced in Chapter 2. The intention of this alternative allegory is not only to highlight the key assumptions of sociocultural constructivism, but also to contrast it with EF Theory. I then move to a discussion of the implications of the position for design, pedagogy, and evaluation. As I shall illustrate, sociocultural constructivism does not contradict cognitive constructivism. Rather, it reinterprets the benefits of so-called active learning, placing greater emphasis on the importance of authentic activities and collaboration.

3.3.1 The Allegory of Musiphonia

Recall that the Allegory of Musica (see Chapter 2) tells of a singing mammal, the Musibeest, that once roamed the land of Musica, and of the Musican school system's commitment to teaching students to imitate the Musibeest's song as accurately as possible. In the neighboring land of Musiphonia, citizens also enjoy the songs of the Musibeest. However, their Musican counterparts, in a historic invasion, captured and took away all of the Musibeests living in Musiphonia. Shortly after the Musican invasion, a few Musiphonians set about the task of preserving the heritage of the Musibeest song. With the help of a few surviving Musibeest recordings, they developed a musical brass ensemble instrument called the Musiphone. Its inventors dedicated themselves to organizing Musiphone ensembles, whose charter was to celebrate and share the song of the Musibeest.

Unlike the Musicans, the Musiphonians were not concerned with the accuracy of their Musibeest reproductions. (Even if they had been interested in accurately reproducing Musibeest sounds, they no longer had in their midst Musibeests against whom to check their reproductions.) Instead, the Musiphonians turned their interest to building cohesive, tightly-knit Musiphone ensembles to experience—and, indeed, to define—what it means to know Musibeest music.

Within the Musiphone ensembles, a system of educating newcomers gradually evolved that was radically different from that of the Musicans. Musiphonians who decided they wanted to dedicate themselves to the music of the Musibeest were taken in by Musiphone ensemble masters. When they joined the ensemble, those newcomers were given simplified versions of the Musiphone and small, simple parts in the concerts. As they gradually gained experience, they were given increasingly sophisticated versions of the Musiphone, and they were called on to play increasingly difficult parts in concerts. In addition, they might be asked to help the Musiphone masters out with other administrative responsibilities entailed in Musiphone ensemble membership—for example, booking concerts, maintaining instruments, and scheduling consultations with interested newcomers. Eventually, a newcomer would gain so much responsibility in the ensemble that her role would be hard to distinguish from that of the ensemble master. Indeed, she would gradually come to be recognized as a master in her own right.

3.3.2 Key Assumptions

3.3.2.1 Knowledge and Learning within a Community

As the Allegory of Musiphonia indicates, sociocultural constructivism differs markedly from both EF and cognitive constructivist theory with respect to the status of knowledge and the process of learning. Whereas the EF and cognitive constructivist perspectives situate knowledge in the head of the individual, the sociocultural constructivist view reconceptualizes knowledge as the collaborative achievement of a community. As Lave and Wenger (1991) note, "[a] community of practice is an intrinsic condition for the existence of knowledge" (p. 98).

A contrast between the Allegory of Musica (Chapter 2) and the Allegory of Musiphonia (above) serves to elucidate this view. For both the Musicans and the Musiphonians, knowledge is intimately entwined in the songs of the Musibeest. The relationship between knowledge and Musibeest songs, however, is markedly different in the two allegories. In Musica, knowledge fundamentally entails replicating a coveted absolute: the song of the Musibeest. The closer the match with that absolute (a close match can be measured objectively), the truer the knowledge. In Musiphonia, on the other hand, knowledge—that is, the song of the Musibeest—does not exist except in relation to people engaging in the practices of a community; it is the collaborative achievement of Musiphonian ensembles making music together.

This view of knowledge implies a fundamentally different position with respect to what it means to learn. As we have seen, both the EF and cognitive constructivist views focus on learning at the level of the *individual*. These views hold that individual knowledge is somehow transformed through the learning process, and they aim to explain just how such changes in individual knowledge come about. By contrast, sociocultural constructivism interprets learning at the level of the *community*. Instead of asking "What knowledge does one gain through learning activity *x*," sociocultural constructivism asks "What form of participation in community *y* is enabled through learning activity *x*?" As Lave and Wenger (1991) emphasize, "participation in the cultural practice. . .[of a community]. . .is an epistemological principle of learning" (p. 98).

In the Allegory of Musiphonia, for example, the focus is not on individual learning, as evidenced by an ability to replicate Musibeest songs with increasingly high accuracy. Rather, the allegory portrays learning in terms of participating more fully within a Musiphonian ensemble. Fuller participation, on this view, entails gaining the kinds of competence necessary to take on the more central practices of the community, such as playing more important parts in concerts with more complicated instruments, and arranging meetings with newcomers.

3.3.2.2 Learning Is Changing Participation and Identity

Whereas EF and cognitive constructivist theories focus on changes in the mind of the learner, sociocultural constructivist theory focuses on changes in the way in which the learner *participates* in the practices of a community. According to this view, a community implicitly defines a participation structure, which encompasses various ways of participating in the community and thereby contributing to its reproduction. When newcomers enter a community, they participate at its periphery; their role in contributing to the ongoing practices of the community are minimal. According to Lave and Wenger (1991), such peripheral participation, which is legitimized by the community, defines the learning process. It is only through the process of *legitimate peripheral* *participation* that newcomers are able to learn—that is, to participate in more "expert" ways within a community.

For instance, in the Allegory of Musiphonia, newcomers start out by playing small parts on simple versions of the Musiphone. In this early stage in their Musiphone Ensemble career, they are participating on the periphery. Gradually, as they participate for a longer time, they may be given the opportunity to participate in more central ways. For example, they may play more important parts in concerts, or they may take on more administrative responsibility. In so doing, they move toward fuller participation in the community. If they stay with a Musiphone Ensemble long enough, they may gain a level of participation that is indistinguishable from that of a Musiphone master.

A key aspect of a community member's advancement toward fuller participation is the member's *identity*: how the member views herself, and is viewed by others, within and with respect to the community. On the sociocultural constructivist view, one's status within a community depends not only one's level of competence and participatory role in the practices of the community, but also on one's own perceptions of one's place in the community, as well as those of others. As Lave and Wenger (1991) put it,

> Activities, tasks, functions and understandings do not exist in isolation; they are part of broader systems of relations in which they have meaning. The person is defined by as well as defines these relations. Learning thus implies becoming a different person with respect to the possibilities enabled by these systems of relations. To ignore this aspect of learning is to overlook the fact that learning involves the construction of identities. (p. 53)

The Allegory of Musiphonia can be used to illustrate the role of identity in learning. At first, an ensemble newcomer participates in ways that indicate she is a newcomer. She plays simple parts on simple instruments, and, for the most part, only observes fuller members as they take care of the business of the ensemble. She regards herself as a beginner in the ensemble, and others also recognize her as such. As she moves through the ranks of the ensemble, however, she constructs a new identity for herself. As she is given the opportunity to play more important parts in concerts, and as she is granted more administrative responsibilities, she comes to see herself as a more central player in the ensemble. Likewise, others regard her as such, and, as a result, begin to treat her differently. For instance, they might call her by different names or greet her differently. The most significant consequence of her change in identity, however, is her change in level of participation: She takes on a more central role in the community, and does so with the support and blessing of other community members.

3.3.2.3 Access Is Crucial

As discussed above, the sociocultural constructivist view maintains that learning can only take place through participation in the authentic activities of a community of practice. An important implication of this position is that *access* to the community is absolutely essential to the learning process. As Lave and Wenger (1991) explain, through access to the community, learners are able to become more central members of the community:

> From a broadly peripheral perspective, [learners] gradually assemble a general idea of what constitutes the practice of community. This uneven sketch of the enterprise (available if there is legitimate access) might include who is involved; what they do; what everyday life is like; how masters talk, walk, work, and generally conduct their lives; how people who are not part of the community of practice interact with it; what other learners are doing; and what learners need to learn to become full practitioners. It includes an increasing understanding of how, when, and what they enjoy, dislike, respect, and admire. In particular, if offers exemplars (which are grounds and motivation for learning activity), including masters, finished products, and more advanced [learners] in the process of becoming full practitioners. (p. 95)

This quote suggests that learners require access to a broad range *of community resources*, including the ongoing activities of the community, central members of the community, other learners, and the tools, artifacts, and technology of the community. In Musiphonian Ensembles, for example, newcomers require access to the ensemble leader, other ensemble members, sheet music, Musiphones, performances, and rehearsals.

In addition to community resources, sociocultural constructivism emphasizes that learners require access to *participation* in the authentic activities of the community—both peripheral forms, and more central forms. For instance, in the Musiphonian ensembles, newcomers are initially allowed to participate in the ensemble peripherally; they play simplified instruments and small parts. Later on, as they mature as ensemble members, they gain access to more central forms of participation. For example, they play more important parts in concerts, and help to recruit newcomers. In sum, access to resources and participation in a community is crucial if learners are to learn—that is, if they are to gain fuller membership in the community.

3.3.3 Implications for Pedagogy

Sociocultural constructivism, as articulated by the assumptions just discussed, has important implications for how algorithms instructors ought to use AV as a pedagogical aid. As we have seen, sociocultural constructivism views learning in terms of increasingly central participation within a community of practice. I call the community that is in the process of reproducing itself through an undergraduate algorithms course the "Community of Schooled Algorithmaticians" (COSA). This community consists of people who "practice" algorithms with a shared sense of such things as (a) what is important and interesting about algorithms (e.g., procedural behavior, efficiency, correctness); (b) how to communicate about algorithms with others (e.g., comparisons of naïve and efficient algorithms, pseudocode, graphical notations, correctness proofs, Big-O analyses); and (c) how to do algorithmic problem solving.

As we have seen, sociocultural constructivism assumes that learning algorithms fundamentally entails participating more centrally in the COSA, and that learners can only advance toward more central participation if they have access to the community. With respect the pedagogical use of AV technology, these assumptions have a clear message: namely, that algorithms instructors would do well to set up situations in which students have access to (a) the *resources* of the COSA (e.g., ongoing activities, old-timers, learners, AV technology), and (b) *opportunities for participation* in the authentic activities of the COSA. The important question for instructors thus becomes, "How might AV technology be used to provide opportunities for participation in the authentic activities of the COSA?"

I argue that AV technology, whether high-tech graphical simulations or low-tech sketches, provides access to at least three key forms of COSA participation:

1. meaningfully interpreting visual representations of algorithms (*AV reading*);
2. constructing visual representations of algorithms (*AV writing*); and
3. presenting visual representations of algorithms to, and discussing them with, members of the community (*AV presentation/discussion*).

Notice that these three activities are listed roughly in order of increasing centrality. Whereas newcomers in the community might be expected to read an algorithms text and meaningfully interpret the graphical representations contained therein, newcomers are seldom expected to construct their own visual representations of algorithms, much less to engage in presentations and discussions involving such representations. In other words, AV reading is a peripheral form of participation in the community; AV writing, and AV presentation/discussion are more central. These layers of AV technology-mediated participation, which define an important part of the COSA participation framework, are illustrated in Figure 14.

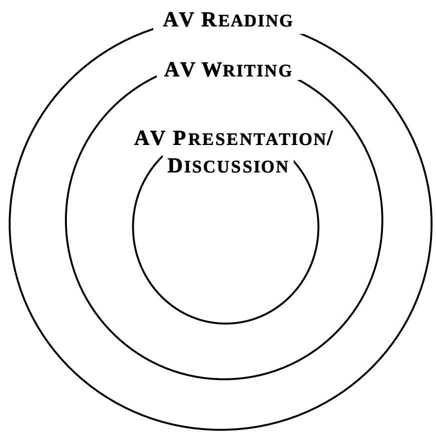

Figure 14**.** Layers of AV technology-mediated participation

As I illustrated in Chapter 2, the traditional pedagogical approach to using AV technology has been to have students view and interact with graphical representations of algorithms. On the sociocultural constructivist view, this approach grants students access to the activity of AV reading. However, if it is true, as sociocultural constructivism assumes, that access to *increasingly central* participation is key to learning, then algorithms instructors would do well to provide students with access to the more central forms of AV technology-mediated participation: AV writing and AV presentation/discussion.[12] Granting such access would involve setting up situations in which students construct their own AVs, and then present them to, and discuss them with, a group of COSA members—both old-timers and newcomers.

3.3.4 Implications for Technology Design

As discussed in Chapter 2, EF Theory holds that the key to successful learning is a close denotational match between an expert's mental model and its graphical representation in an AV. As I argued, this view implies that AV technology must support high fidelity representations, which enable an algorithms expert to portray an algorithm with a high degree of accuracy. Extant AV technology, as I illustrated, has taken this implication of EF Theory to heart through its consistent support of input-general, high-typeset fidelity AVs.

In contrast, sociocultural constructivism asserts that AV technology's pedagogical value lies in its ability to provide access to increasingly central participation in the COSA. Clearly, the implications of this position for AV technology design differ markedly from those of EF Theory. Indeed, sociocultural constructivism, like cognitive constructivism, rejects the importance of a close denotational match between an expert's mental model and an AV. Instead, sociocultural constructivism implies that AV technology designers ought to concentrate on building AV technology that supports both *participation* and *communication*.

Sociocultural constructivism suggests that, to support increasingly central forms of participation in the COSA, AV technologists must jettison the traditional dyadic user model discussed in Chapter 2, as this dyadic model fosters the student-instructor dichotomy, and denies students access to more expert forms of participation. But what user model should be used in its place? Like cognitive

[12]As Lave and Wenger (1991) put it, "the understanding to be gained from engagement with technology can be extremely varied, depending on the form of participation enabled by its use" (p. 101).

constructivism, sociocultural constructivism suggests that AV technology should enable students to participate in AV construction—that is, it should provide an environment in which students can construct their own AVs. However, as we have seen, sociocultural constructivism goes beyond cognitive constructivism in asserting that AV technology must also enable students to participate in the most central AV technology-mediated activities: AV presentation and discussion.

According to sociocultural constructivism, it is in the activities of AV presentation and discussion that AV technology holds promise in facilitating effective communication about algorithms. In this capacity, sociocultural constructivism does not see AVs as merely "transferring" knowledge from the AV presenter to the audience, as EF Theory holds. Rather, according to the sociocultural constructivist view, AV technology serves to *mediate* interaction between students and instructors. In the words of Roschelle (1990), AV technology serves to manage the inherent "uncertainty of meaning in conversations" by functioning as "a common ground for shared activity and talk among learners and experts" (p. 3).[13]

The Allegories of Musica and Musiphonia well juxtapose the communicative role of AV technology according to EF Theory, on the one hand, and according to sociocultural constructivism, on the other. In Musica, AV technology is analogous to a tape recorder; it conveys with precise accuracy the song of the Musibeest, which symbolizes absolute knowledge. In Musiphonia, on the other hand, AV technology is analogous to the Musiphone; it enables members of Musiphonian ensembles to negotiate shared meaning—the socially agreed-upon song of the Musibeest—through their concerts. Clearly, to play the role of a Musiphone, AV technology must support conversations about algorithms. Just what specific design features this implies is a question that only an empirical study can answer; I take the question up in the next chapter.

3.3.5 Implications for Effectiveness Evaluation

According to sociocultural constructivism, effective pedagogical exercises involving AV technology provide access to increasingly central forms of participation in the COSA, including AV construction and AV presentation/discussion. Thus, one immediate measure of effectiveness suggested by sociocultural constructivism is *activity relevance*: to what extent does a pedagogical activity involving AV technology involve learners in the authentic activities of the COSA? The more relevant and authentic is a learner's experience with respect to the COSA, the better.

The sociocultural constructivist position further suggests that, through participating in increasingly central ways, learners gradually become fuller members of the COSA. Yet, how should one measure something as ostensibly vague as membership in a community? Recall that one's level of membership, on the sociocultural constructivist view, is intimately connected to one's *level of participation* and *identity*. As discussed above, the roles, responsibilities, and tasks that one takes on within a community fit into a participation framework that resembles an onion: Some of these roles, responsibilities, and tasks are on the periphery (the skin of the onion); others are more central (near the core of the onion). Thus, the roles, responsibilities, and tasks that one takes on within a community provide a rough indication of one's level of participation, and hence the extent to which one has gained full membership. By conducting ethnographic fieldwork within a community of practice, one can observe longitudinal changes in people's levels of participation within the community.

As discussed above, a key component of level of membership is *identity*: how one perceives oneself, and is perceived as others within and respect to the community. Like level of participation, identity is difficult to measure in a quantitative sense. However, just as it may be used to assess one's level

---

[13]Roschelle (1990) uses the term *symbolic mediation* to refer to the use of visual representations in this way—that is, as a mediator of conversations.

of participation, ethnographic fieldwork can be used to gain a qualitative sense of the way in which one's identity with respect to a particular community of practice changes over time.

3.4 Summary and Research Questions

Inspired by the empirical trend, noted in the last chapter, that suggests that increased student involvement in the visualization process leads to higher learning outcomes, this chapter has proposed constructivism as an alternative, more appropriate guiding theory for the pedagogical use of AV technology. Cognitive constructivism, the first brand of constructivism I considered, views knowledge as a highly subjective, individually-constructed entity. In contrast to EF Theory, the cognitive constructivist position would predict the empirical finding that active viewer involvement is important. In particular, cognitive constructivism implies that, since learners must actively construct their own understandings, they will learn best by constructing their own visual representations of algorithms.

This chapter has suggested, however, that understanding algorithms entails far more than an ability to recount accurately procedural behavior, that understanding algorithms fundamentally involves an ability to engage in the *practice of algorithms*. Since that practice is inherently social, I have argued that cognitive constructivism, with its narrow focus on individual knowledge, proves inadequate as a guiding theory of effectiveness. As a consequence, I turned to sociocultural constructivism, a more recent version of the constructivist position that situates knowledge within communities of practice, and recasts learning in terms of changing levels of participation in such communities. I argued that sociocultural constructivism, with its broader analytical scope, proves more appropriate as a guiding theory, and I outlined its implications for the pedagogical use, design, and evaluation of AV technology.

While many of its implications were similar to those of cognitive constructivism, two key differences did arise: first, that sociocultural constructivism emphasizes the importance of providing students with opportunities to participate in the social, AV technology-mediated activities of AV presentation and discussion; and second, that sociocultural constructivism stresses the importance of designing AV artifacts as *mediational resources* for such presentations and discussions.

As we have seen, sociocultural constructivism offers plenty of guidance with respect to how to design, implement, and gauge the benefits of AV technology-based pedagogical exercises. Yet, until we actually try out its ideas within the context of an actual algorithms course, we cannot know whether sociocultural constructivism is appropriate as a guiding theory of effectiveness for AV technology. As we have seen, in order to observe the outcomes of the sociocultural constructivist ideas, one must view the algorithms course in which they are implemented as a distinct community. And one way to proceed to do this is to conduct an ethnographic field study. In the next chapter, I report on a series of exploratory ethnographic studies that I conducted in two different algorithms courses that incorporated AV technology using a sociocultural constructivist approach.