Clancey, W. J. 1997. International Journal of Educational Research, 27(1) 5-19.

CHAPTER 2

CONCEPTUAL COORDINATION: ABSTRACTION WITHOUT DESCRIPTION

WILLIAM J. CLANCEY

Institute for Research on Learning, 66 Willow Place, Menlo Park, CA 94025, U.S.A.

Abstract

Conceptual coordination is a learning process that relates multiple perceptual-motor modalities (verbal, visual, gestural, etc.) in time. Lower-order categorizations are thus related by sequence and simultaneity, as shown by neurological dysfunctions. Heretofore, many theories of abstraction have only considered verbal behavior and assumed that the neural mechanism itself consists of manipulation of descriptions (linguistic models of the world and behavior). This broader view better relates physical and intellectual skills.

Information processing theories of cognition explain thinking, learning, and action in terms of symbol structures and the processes operating on them (Newell & Simon, 1972). These theories are expressed as computer programs, called *descriptive cognitive models*, in which the symbol structures and processes have characteristics very similar to how we communicate with each other and encode knowledge in external media: language, special-purpose notations, formulas, and so on. Such representations are viewed by many cognitive scientists as corresponding to human knowledge that is explicit and consciously available. Similarly, the inferential processes traditionally defined over such representational structures are voluntary, fully-deliberate "steps" — decisions made with attention to alternatives and consequences. Such steps can be described as fitting "rules", such as the backward-chaining of symptom–disease heuristics in a medical expert system (Buchanan & Shortliffe, 1984). By assumption, these steps (heuristics and inferential processes) can be rendered fully explicit.

However, a significant part of human cognition is different from deliberate reasoning (Polanyi, 1966). There is little awareness of where actions and thoughts "come from"; they emerge spontaneously in response to a situation or in the course of acting, speaking, writing; people do not always plan in advance what to say; they just speak; sometimes they even surprise themselves with what they say, write, or do. Sometimes people plan what to do, but such plans are not as neatly laid out and thoroughly controlling of behavior as plans in descriptive cognitive models. Often people do not know how they accomplished something until they reflect afterwards on what they did.

Some everyday examples of conceptual coordination are quite striking. People can imitate an accent without describing it theoretically; they can visually project events in their imagination (such as whether a spilled cup of coffee will reach the end of the table); and we find in many physical activities, as in dance, sports, and music, that words often get in the way. All of this points to cognitive structures and processes that are implicit, tacit, and organized on different sensory and temporal dimensions. Apparently such "steps" and actions are *generated, generalized, and coordinated along different dimensions in the course of acting.* When the cognition of children and animals who patently do not act by first *describing* their world and alternative behaviors is considered, this point appears even stronger.

The view of knowledge as tacit and generative, rather than explicit and programmatically applied, implies that alternative notions of representation must be sought — symbol structures and processes with different characteristics than those employed in descriptive cognitive models. Inspiration can be found in the neural network models of Edelman (1992) and Freeman (1991), related work in contextualist (Hoffman & Nead, 1983) and ecological psychology (Gibson, 1979), and in anthropology (Suchman, 1987; Lave, 1988). The correct explanation is unlikely to be identical to today's best models, and some researchers do not offer alternatives to the descriptive approach, but in general these efforts point in the right direction.

Within the emerging understanding, often called "situated cognition" (Clancey, in press), there is room for abstraction. Indeed, it makes sense that tacit knowledge structures are abstract; this is part of their generative power. Hence, there is no contradiction between searching for new mechanisms that lie outside the realm of the descriptive models explored hitherto and the goal of explaining the higher-order intellectual accomplishments of human beings. On the contrary, one goal of situated cognition is to explain higher-order cognitive accomplishments, "abstract thought" and the like, *in terms of* tacit and implicit but generative processes.

A confusion has developed in the scientific community because all kinds of concepts have been equated with *descriptions* (e.g., equations, heuristic rules, semantic nets, and diagrams). "Symbolic reasoning" has been viewed as the foundation of all cognition, such that any sensory input can be mapped to descriptions by some form of *encoding* (Bickhard & Terveen, 1995), and all intelligent action requires and only depends on manipulation of descriptions of the world and behavior. Often other modalities of conceptualization (rhythm, accent, imagery, and gestures) are viewed as merely the input or output of such manipulations. The nature of *conceptual coordination* has been misconstrued by viewing verbal conceptualization as the landing place and controller of all thought. Hence abstract thinking has been misconstrued as being fundamentally verbal, and other modalities of abstraction, on which verbal thinking often depends, are inadequately understood.

Subsequent sections provide an overview of major distinctions and discuss the ways they are sometimes misconceived. To illustrate the nature of conceptual coordination, two examples of neurological dysfunctions are analyzed. Conceptualization and descriptions are differentiated and a neural perspective for understanding the distinction is suggested. Finally, the ways these ideas can help resolve debates about the rule-like nature of knowledge are described.

Background: Representations and Models

To understand abstraction, the reader needs first to understand *representation*. Because the descriptive cognitive modeling approach has dominated cognitive science, the very ideas of "symbol", "representation", and even the research discipline itself are bound up in this one theoretical approach. "Symbolic reasoning" has been equated with cognition to the point that knowledge is equated with descriptions and symbols with words:

> It is sufficient to think of symbols as strings of characters and of symbol structures as a type of data structure . . . The following are examples of symbols: Apple, Transistor-13, Running, Five, 3.14159. And the following are examples of symbol structures: (On Block1 Block2) (Plus 5 X) (Same-as (Father-of Pete) (Father-of (Brother-of Pete))) (Hayes-Roth, Waterman, & Lenat, 1983, p. 61).

In moving forward, terms need to be generalized and varieties of "symbols" and "representations" distinguished. Perhaps the most frequent cause of misunderstanding is interpreting technical terms and existing models in different ways.

In improving and building on the descriptive cognitive modeling approach, situated cognition researchers are *not*:

- attempting to abolish the idea of abstraction, (e.g., suggesting that it is a kind of illusion);
- claiming that "transfer" of knowledge across situations is impossible or unimportant (cf. Anderson, Reder, & Simon, 1996); or,
- claiming that parts of cognition are "non-representational" or "non-symbolic" (in a broad sense) (cf. Vera & Simon, 1993).

Instead, situated cognition research, broadly construed, is:

- *trying to highlight differences* between what existing computer models capture and naturally-occurring cognition. For example, what is rejected is the *transfer metaphor* for explaining: *how memory works* (i.e., categorizations are generalized and recoordinated in use) (Bransford, McCarrell, Franks, & Nitsch, 1977; Iran-Nejad, 1987; Jenkins, 1974; Rosenfield, 1988; Schön, 1979) and *how descriptive generalizations are applied* (i.e., everyday design and policy interpretation typically requires reconceptualization, not using written rules as if they were recipes; (Suchman, 1987; Wynn, 1991)).
- *focusing on non-linguistic aspects of human coordination*, (e.g., visualization, social relations, and rhythm). By hypothesis, symbolic reasoning and language depend on these (Edelman, 1992).
- emphasizing the importance to psychology of *understanding how the brain actually works*, to the point of having an engineering specification that would enable replicating its functions for the broad range of human intelligence (Effken & Shaw, 1992; Freeman, 1991). For example, not all "symbol structures" in an animal's brain are entities for *referring* to the external world (Maturana & Varela, 1987).

Overall, this is a means-ends analysis: What is the gap between what psychologists and robot designers have accomplished and where they are trying to go?

The clarifications provided here are important because what some researchers take for granted others might never have believed. For example, one cognitive psychologist wrote to me, "'Storage' is only a metaphor; nobody in the symbolic cognition tradition thinks of it as actually storing something in a space, whatever that would mean". But one can easily find the opposing view throughout the artifical intelligence (AI) literature. For example, Pylyshyn provided the following commentary at the 22nd Carnegie Symposium on Cognition:

> The choice of both notation and architecture are central empirical issues in cognitive science, and for reasons that go right to the heart of the computational view of mind. It's true that in the physical sciences,

> theoretical notation is not an empirical issue. But in cognitive science our choice of notation is critical precisely because the theories claim that representations are written in the mind in the postulated notation: that at least some of the knowledge is explicitly represented and encoded in the notation proposed by the theory. The architecture is likewise important because the claim is that these are literally the operations that are applied to the representations . . . In cognitive science, theories claim that the mind works the way the model does, complete with notation and architecture. What is sometimes not appreciated is that computational models are models of what literally goes on in the mind (Pylyshyn, 1991, p. 221).

Newell and Simon, in their landmark work, *Human Problem Solving* also expressed this point of view:

> The theory posits a set of processes or mechanisms that produce the behavior of the thinking human. Thus the theory is reductionistic; it does not simply provide a set of relations or laws about behavior from which one can often conclude what behavior must be. (The elementary processes and their organization, of course, are not explained: reduction is always relative.) Thus, the theory purports to explain behavior—and not just to describe it, however parsimoniously. (We are aware that some would dispute such a distinction, viewing all causal explanations as simply descriptions.) (Newell & Simon, 1972, p. 9)

Thus, the expressions in cognitive models, according to this view, are not merely metaphorical descriptions. These structures and processes equivalently "produce the behavior of a thinking human". As Pylyshyn put it 25 years later, the "representations are written in the mind in the postulated notation". Not only is this not a "straw man view" of cognitive theory, it has been the dominant view, which successfully drove the development of theories of natural language processing, novice–expert differences, learning, and the like (Gardner, 1985b).

Dysfunctions Reveal Non-Descriptive Cognitive Processes

To illustrate the idea of conceptual coordination, two cases of neurological dysfunctions presented by Oliver Sacks (1987) are analyzed. These examples lead us to reconsider the nature and role of abstract thinking. Is the realm of abstract thinking to understand, create, and apply scientific theories? To design buildings and cities? To interpret company policies when responding to a customer? Surely abstract thinking includes all of this. But what about finding one's way around a block? Does putting on a shirt and buttoning it involve abstract thinking? Does tying one's shoes? Oliver Sacks argues that everyday actions involve a form of thinking that has a non-verbal, but necessarily "abstract" aspect. At the same time, dysfunctions reveal otherwise hidden processes on which verbal behavior normally depends.

Consider for example Rebecca, whom Sacks characterizes as having two modes of being. The first, a mode of thought measured by formal testing, requiring pattern-seeing and problem solving, revealed her as defective, lacking basic human capability.

> A poor thing . . . a multitude of apraxias and agnosias, a mass of sensorimotor impairments and breakdowns, limitations of intellectual schemata and concepts . . . a mere mosaic of higher cortical functions . . .
>
> She had done appallingly in the testing—which, in a sense, was designed, like all neurological and psychological testing, not merely to uncover, to bring out deficits, but to decompose her into functions and deficits. She had come apart, horribly, in formal testing . . . (Sacks, 1987, pp. 180–181).

The second mode of being was revealed as Rebecca sat on a garden bench, enjoying a spring day. She gestured to the foliage and spoke poetically, in spurts, "'spring', 'birth', 'growing', 'stirring', 'coming to life', 'seasons', 'everything in its time'." Rebecca cannot coordinate a schematic, spatial view of her behavior, as is required in finding her way around the block. She can metaphorically relate two images, but cannot sequence their concrete relation, as in fitting a hand to a glove or a key in a keyhole. Her inability to *compose behavior sequences* is manifest again in her speech. She can follow a meaning metaphorically, but she can't conceptually coordinate her own narratives.

Rebecca apparently experiences ways of seeing *directly* (without narrating her experience) and she can pattern herself after an ongoing *concrete form* in the environment in which she embeds her activity. She says, "I'm sort of like a living carpet. I need a pattern, a design, like you have on that carpet. I come apart, I unravel, unless there's a design" (pp. 184–185). She needs to be supplied a narrative structure, some pattern-rhythm to interact with directly in the environment. She can't compose scenes of her own conception, but she could be "composed by a natural scene", which presents itself to her as a dramatic unity, with aesthetic sense. Attempting to achieve coordinated action by her own spatial–temporal constructions, she becomes lost, appearing moronic and spastic. Top-down, internally driven organizers of verbal sequencing and ordering of scenes into imagined plans appear to be impaired.

The "abstract versus concrete" dichotomy takes on new meaning when we consider conceptual coordination as involving different sensory modalities, as illustrated by an example of a contrasting dysfunction. Dr P is the famous "man who mistook his wife for a hat". Dr P lives in the opposite of the autistic world; he lives in the world of abstract conceptions which he cannot appropriately relate to concrete things in the scene around him (pp. 7, 20, 229):

> Not only did Dr P increasingly fail to see faces, but he saw faces when there were no faces to see . . . in the street he might pat the heads of water hydrants and parking meters, taking these to be the heads of children; he would amiably address carved knobs on furniture and be astounded when they did not reply (p. 8).

Dr P proceeds as if he were mimicking a symbolic computer program, and consequently something is wrong with him:

> There was something radically wrong with the whole way he proceeded. For he approaches these faces—even of those near and dear—as if they were abstract puzzles or tests. He did not relate to them, he did not behold (p. 13).

Given a single rose, Sacks reports that Dr P spoke like an expert system program manipulating descriptions:

> He took it like a botanist or morphologist given a specimen, not like a person given a flower. "About six inches in length", he commented. "A convoluted red form with a linear green attachment . . . It lacks the simple symmetry of the Platonic solids, although it may have a higher symmetry of its own . . . " (pp. 13–14).

Visual imagination and memory are impaired, too. Asked to visualize and describe a familiar street, Dr P doesn't mention buildings on the left side. Apparently, there is a problem within the right brain, which processes the left visual field and is generally attributed with recognition of images as wholes. Asked to recall a novel, "he had an undiminished grasp of the plot, but completely omitted visual characteristics, visual narrative, and scenes" (p. 22).

The impairments in knowledge explored by Sacks are not scientific misconceptions or the kind of failures in high school physics tests. Rather, he explores *the abstract nature of conceptual* *coordination*, where the "abstracting" process is not only verbal, but includes other kinds of organizers in time: visual, rhythmic, manipulospatial.

Teasing apart the lessons from neuropsychological dysfunctions is complicated because the patients illustrate, even in their dysfunctions, what computer models cannot do. For instance, Dr P can only coordinate his eating and dressing by continuing to hear a song in his mind (p. 17)! On the one hand, these patients illustrate the *reality* (Dr P) and *necessity* (Rebecca) of abstraction by descriptive modeling in everyday life. However, their experiences suggest that more might be going on; something which is also integral to human cognition.

To sort out these organizers, the mechanisms of symbolic reasoning (descriptive modeling) can be contrasted with sequential and simultaneous relating in human coordination (Table 2.1). Of the two patients, Dr P is more like a computer program operating on descriptive models. But taken together, the examples illustrate that abstraction via description manipulation is insufficient — descriptive cognitive models do not adequately capture or replicate everything that people can do.

To elaborate a bit more, "concrete" in Rebecca's understanding means especially a *direct coupling* between perceptual patterns and action. She lacks a kind of non-verbal abstraction required for hand–eye coordination or constructing an imaged plan of motion over time, as required in walking around the block. She compensates by embedding her action in a narrative conceived over her visual-auditory space. But her *metaphoric understanding*, a form of abstraction, cannot be explained by descriptive cognitive models, which postulate that metaphor is a process of matching feature descriptions (Schön, 1987; Hofstadter, 1995) — which she cannot do at all.

"Concrete" in Dr P's understanding means describing and relating properties of objects. He lacks a kind of visual abstraction required for relating details within a simultaneously perceived configuration. He compensates by manipulating descriptions, like an expert system. But his *visual problems* cannot be explained by descriptive cognitive models which postulate that visual abstraction is a process of matching feature descriptions — which he can do very well indeed.

On the surface, the conventional formal definition of abstract thought appears reasonable: "...Thinking that goes beyond immediate experience is regulated by knowledge structures called *abstract schemas*" (Ohlsson, 1993, p. 51). Conventionally, such abstract schemas are taken to be descriptions of objects, properties, and events, as in explanations of Dr P's reasoning. But the following forms of conceptual organization also go beyond immediate experience: hearing a tune in one's head, visualizing a planting border around a lawn, placing an arm in a sleeve. Such

Table 2.1
Dimensional Analysis of Human Experience and Models

Abstract cognitive processes	Symbolic reasoning	Ability to compose sequential relations	Ability to simultaneously relate image and sound in coordinated action
Example	manipulating descriptive model of world, e.g., deductive inference	physically aligning objects, grammatical speech, projecting ordered events, e.g., hearing a song in one's head	dancing, speaking metaphorically
Rebecca			✓
Dr P	✓	✓	
Descriptive cognitive model	✓	(✓)	

behaviors are conceptually coordinated and do not all involve verbal descriptions in their form or regulation. Indeed, most studies of abstract thought focus on scientific theories, not how to find one's way around the block. If we take *thinking* to involve any organizing performed by the brain in relating and ordering of actions in time, then a more general notion of abstracting might be called *conceptual coordination.*

From this broader perspective of conceptualization the concrete and the abstract can begin to be placed in a different relation—not just an ordering of *descriptions* from specific to general or implicit to explicit. For example, Varela (1995), says that "the proper units of knowledge are primarily concrete, embodied, lived . . . The concrete is not a step toward anything: It is how we arrive and where we stay" (pp. 11–12). He continues:

> In a nutshell, the enactive approach [to cognition] consists of two key points: (a) that perception consists of perceptually guided action; and (b) that cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided (p. 16).

Here is the "mechanistic" aspect of situated cognition theory. *Perception, conception, and action can be coupled; they may develop and be reactivated as an ensemble.* Hence, Rebecca can couple her movements to a perceived pattern in her environment, and Dr P apparently couples his movements (in dressing or eating) to the reconstructed perceptual experience of music. Such behavior is "direct" (in Gibson's sense) and "concrete" (Varela, Luria) because it is not mediated by description and inferential chaining. But the organizers are nevertheless "abstract" because they are generative, recurrent ways of coordinating behavior.

The *coupling relation* between perceptual and motor systems, even those involving conceptual organizers, was emphasized by Dewey in his famous critique of early stimulus–response theory (Dewey, 1896). The relation of sensation, perception, and motor processes is dynamic, as part of a circuit, such that the momentary interactions within the system are sustained and directed as one developing ensemble, each momentary organization leading to the next, and (in Bartlett's terms) each organization is literally built out of the components that have worked together in the past (Bartlett, 1932). Again, this includes conceptual processes, as in Rebecca's conception of narratives and metaphors — there are other ways to form and relate concepts than by inferential chaining of descriptions. Understanding musical intelligence, spatial reasoning, visual recognition, and their relations to symbolic reasoning (Gardner, 1985a) is enhanced by this shift to a multimodal, coordination view of cognition.

Towards a Better Formalization: The Role of Descriptions

To better distinguish the abstract from the concrete, descriptions must be distinguished from conceptualization. Figure 2.1 shows how these terms might be ordered. Notice the distinctions among describing, conceiving, reasoning, and calculating.

- *Descriptive representations* are linguistic expressions (e.g., texts and diagrams), including numeric models and computer datastructures called "symbolic representations".
- *Conceptualization (representing in brain)* refers to modalities for coordinating experience and action, includes verbal and non-verbal, the simultaneous and the sequential.
- *Reasoning* consists of categorizing, ordering, and comparing conceptualizations to construct a "mental model" or "action plan"; it includes inner speech and is never strictly formal because it always involves conceptualization.

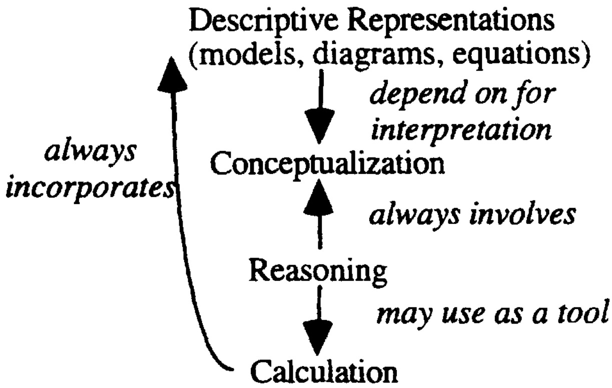

Figure 2.1. Relation of descriptions to thought and calculation.

- *Calculation* is the formal manipulation of descriptions, as in mathematics, expert systems, and other descriptive cognitive models (e.g., Soar in Newell, 1990).

Reasoning may include verbalization, projective causal envisioning (imagery), metaphorical conception, rhythmic ordering, or some combination of these. Calculation is a tool sometimes used as part of reasoning to reify conceptions and record them as a configuration of objects and relations. Calculation always involves descriptions, such as equations, statements, and notations (e.g., geometry diagrams). Computer programs, constructed from descriptions alone, may replicate the "formal" aspect of calculations (in this sense, they are often called "purely syntactic").

Within the "situated view" there is a fundamental difference in kind between descriptions/calculation, and conceptions/reasoning. But the descriptive modeling view equates conceptualization with descriptive representation and reasoning with calculation. This misses the point that abstract schemas described by Bartlett (what I call conceptualizations) are not descriptions, but *neural categorizations* coordinating different modalities. The exclusively descriptive models of natural language processing omit the non-verbal aspects of comprehension and suppose instead a mechanism built entirely out of words (above the phonemic level). In contrast, conceptualizations are processes of representing. Unlike texts and diagrams, they exist only in the time of use. They do not encode and are not stored or perceived as things. Conceptualizations involve aspects of perceiving and — by virtue of how the categorizing mechanism works — are inherently integrated with physical activity.

Conception of Activity: A Fundamental Kind of Abstract Schema

To further characterize conceptualization and reasoning in terms of situated cognition theory, the social-functional meaning of "situated" must be considered: Reasoning, perceiving, and motion are all occurring within the person's conceptualization of what he or she is doing. That is, what is seen, how it is understood, and what people think to do are all constructed within the *conceptual context* of an activity. Too often, AI researchers have interpreted "situated" to mean "physical", "in a setting" or "interactive", again reducing conceptualization to data or an objectively known world (cf. Anderson et al., 1996). The point is that people are always conceiving "what I am doing now" and this serves as a context for the satisficing orientation of everyday problem solving.

For example, Sacks (1987) emphasizes that Dr P's problem is not just a loss of the isolated function of visual processing. Perhaps more important, to lose the visual capability to conceive scenes, that is to identify objects as wholes, is to lose an aspect of human *subjectivity*. The holistic, conceptual grasp of a face *as being a face* and personal judgment are related. Without a holistic visual grasp, a *means* of experiencing feeling, of relating personally to the world is lost. When Dr P did manage to *infer* the name of the person or thing, he could experience an emotional attachment. Otherwise, things that other people found significant (like his wife sitting on the chair to his side) just went past him.

More generally, being "socially situated" means appropriately choreographing activities—"ways of being", roles, ways of spending time, "things we do when". Examples of such activities are reading the Sunday paper, going to the movies, being the clinic physician-in-charge, being on a business trip, attending a workshop, staying in a hotel, living in California, taking a vacation, and writing a book. Activities are always temporally extended "things we do", often restricted to a certain time and place, with conventions for when we do them, what we wear, how we talk, and what value we place on events. Activities are always socially constructed, in the sense that they are negotiated (by action and feedback) within a community.

Activities are abstractions, like all conceptualizations. But activities must not be identified with their descriptions. It would be easy to slide into calling every activity a task and specifying a goal description and rules or procedures for carrying them out. This is how the exclusively task-oriented view of work leaves out people's conceptions of who they are, how they allocate their time, their allegiances, and their career trajectories. Activities are known by human behaviors; they are what people do (Frake, 1980; Rommetveit, 1987; Wynn, 1991. As conceptions they constitute part of the context in which goals become defined and tasks are assigned and carried out. The real world is part of this context too, but it is the mental conceptualization of role, community, practice, and the like—the choreography of action—that shapes how people think of something to do and how they think about how to do it.

From this perspective, knowledge does not consist of theories and models per se, but comprises our conceptualizations and our perceptual categorizations that coordinate what we see and do. Activity conceptualizations are adaptively activated in different physical and social contexts. In this sense, they are general: A person may go to a restaurant in a different country and understand what is happening around him or her and how to behave, even though the menu and money may be incomprehensible.

Activity conceptualizations were described by Schank and Abelson's (1977) formalization called "scripts". However, in practice human knowledge doesn't consist of a single "restaurant script" per se, but a different *conceptualization* for each actor: the chef, the owner, the waiter, the patron, and the guest. Instead of *universal descriptions* that are shared, conceptualizations involve an inherently subjective point of view. Conceptualizations are alike not just because of the language we use during the activity. More importantly, as categorical relations between roles, stuff in the world, and conventional actions, conceptualizations develop within and sustain a *coordinated practice* of *behavioral interactions*. The similarity, what is shared, lies in *interactive compatibility*, not isomorphism of stored descriptions. For example, in the restaurant, different players "hand-off" their work and interpret materials in compatible ways.

In summary, there are *different kinds of generalization* in practiced motor skills, conceptualizations, and descriptive models. All are abstracted, but in different ways. There is no one-to-one correspondence between them; in particular, we cannot exhaustively describe the "definition" or "meaning" of everyday conceptualizations and cannot functionally replace the neural conceptualization process with an engine that manipulates and controls behavior on the basis of descriptions alone. The evidence for this is provided by neurological dysfunctions. Dr P resembles an inferential engine, unable to see the forest for the trees. On the other hand, a fully "poetic" coordination process like Rebecca's, lacking a "symbolic" organizer, reveals an inability to deliberately relate categorizations in space and time — but she can dance.

Relating Descriptions and Neural Models

In summary there are two domains in which to talk about abstraction (see Figure 2.2). The first is a comparison of descriptions in which classes with features serve as variables to generalize a description (e.g., "give aspirin to a patient with a fever" is more abstract than "give aspirin to Mary"). The second is temporal activation and configuration of sensorimotor processes, including the phenomena of perceptual categorizing, the practice effect (categorizing sequences, conceiving chunks), and multimodal coordination (singing while dressing).

The first is the domain of an observer's descriptions in text, speech, and diagrams. The second is the domain of sensorimotor coupling (what Maturana & Varela (1987) call *embodied action*). The *historical* relation of abstract and concrete are different in these two domains, though both emphasize a causal relation: the relatively abstract is constructed from the concrete. But in the purely descriptive domain of computer models, abstract descriptions are created by examining concrete descriptions (e.g., cases or input examples) and generalizing them, a process of finding patterns and stating rules with variables (e.g., see Buchanan, Sutherland, & Feigenbaum, 1969). In the domain of embodied action, there are at least three kinds of abstraction relations:

- Perceptual categorization from sensory signals;

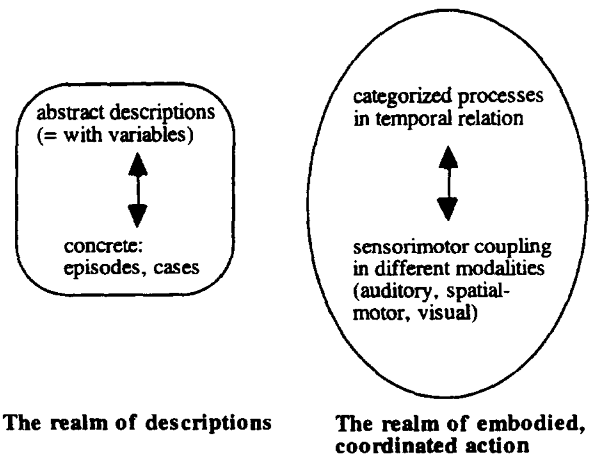

Figure 2.2. Two domains for using the term "abstraction": a comparison of descriptions or a characterization of recursive neural categorizations.

- All categorizations are biased by (built from) previously formed categorizations; and
- Verbal chunks are categorizations of *sequences* of categories forming in time, and a visual chunk involves a *concurrence* of categorizations (what Luria, 1968 calls "simultaneity").

When people create descriptive models (such as formulating a simple rule about life), they engage in embodied action in order to manipulate descriptions. A key characteristic is that neural categorizations forming now embody *relations that developed historically*. These include especially the contextual relations of multiple modalities, such that emotion, smell, visualization, hearing, and motion are coupled in memory by the coordinations previously constructed. Such "remembered" categorizations are not stored, indexed, and/or matched at a primary level, but directly activated "in-line" within circuits (Edelman, 1992; Freeman, 1991; Merzenich, Kaas, Wall, Nelson, Sur, & Felleman, 1983; Rosenfield, 1988).

Categorizations are neural structures that activate and hence constitute structures that are forming at the time of experience itself. This is to be contrasted with a *stored* memory. Neural memory is "content-addressable", but without retrieval as independently existing things. Associations are in some sense direct, involving what Edelman (1992) calls "classification couples" and "reentry" (mutual excitation) between levels of categorization. By this view, the "seven plus or minus two" size of short-term memory is a limit on *how many processes* we can sequentially chain together and *hold active* at one time. That is, we construct an activation sequence by which one global neural map feeds forward into the next and do this for 7± 2 maps. It is not a constraint on *space* (a buffer size) but on *time* (with respect to sustaining activations).

Other key properties of this emerging understanding of neural processes of representing include the following:

- Neural categorizations are *not immutable forms* (unlike tokens in descriptive models) (Merzenich et al., 1983);
- Neural categorizations are stable *relations* between processes (Rosenfield, 1988);
- Neural categorizations may be *non-verbal*, as in imagery and music (Gardner, 1985a); and,
- *Categorizing sequences of categorizations over time* may correspond to the inferential chunking mechanism of descriptive cognitive models (as in Soar; Newell, 1990).

Figure 2.3 shows one simplistic way of visualizing how words in descriptive cognitive models might be related to different neural processes: (1) perceptual-motor categorization (e.g., looking

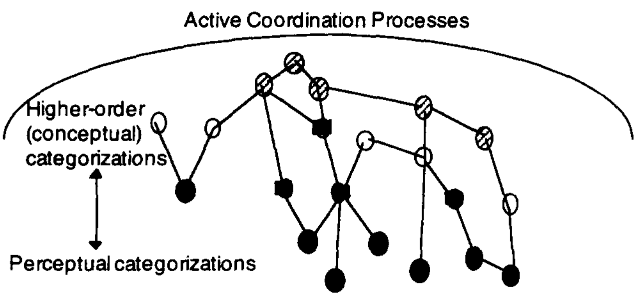

Figure 2.3. Sketch of categorical relations at a certain moment in time.

at a shiny surface) the lower-most nodes, (2) categorizations of words (auditory sequences, shown as black nodes), and (3) categorization of activities ("what I'm doing now" shown as nodes with diagonals). Each node represents an active neural process (in Edelman's (1992) terms, a global neural map). The arc at the top signifies "what's active now".

The point is that words probably correspond to recurrent neural categorizations, but there are more abstract, subsuming conceptualizations organizing the person's experience (activity conceptualizations), as well as more concrete, subsumed conceptualizations of perceptual-motor experience (Lakoff, 1987). Speaking and interpretation of descriptions occurs within this *conceptual activity and perceptual-motor context.* The higher-order understanding affects how people move and where they look as they interact in the world-what perceptual categorizations are of interest to them and how their interpretations are biased. In some activities, such as computer programming or mathematics, people's actions are strongly conventional and regulated; in others, such as spending an afternoon sailing or spending an evening on the town, their actions are still conventional, but there is more room for improvisation.

In short, the idea that conceptualization should be contrasted with descriptions is quite complex, involving not only how words relate to neural processes ("where are the symbols in the brain?"), but how activity is coordinated over time, including how we regulate our choice of words and schedule the tasks of the day, and in the large, how our sense of identity is constructed as social actors. By contrast, descriptive cognitive models are relatively flat, construing all the nodes in Figure 2.3 as words or networks of *encodings*, and viewing all the constructive relationships as processes of indexing, retrieving, matching, and instantiating. An alternative view claims that *mechanisms we do not yet understand* are involved, accounting for such phenomena as rapid figure–ground shifts, musical and rhythmic memory, visualization, silent speech, and projection of imagined movements in space.

Conclusions: Relating Conceptualizations, Descriptions, and the "Operating Principles"

In summary, coming from diverse directions of neuroscience, psychology, and computer science, researchers are converging on the idea that non-verbal phenomena must be considered if we are to understand the nature of consciousness, and thus understand the relation of human intelligence to other forms of cognition. This broader view of cognition embraces the phenomena of non-verbal modalities, neurological dysfunction, and animal cognition, leading to a theory of the evolution of representing (Donald, 1991; Barresi & Moore, 1996). As a first step, there is a need to distinguish among:

1. *Learned neural coordinations*, which are generalized into conceptualizations resembling grammars (e.g., subject-verb order in English);
2. *Descriptions of recurrent neural processes*, such as formal grammars, expert system rules, and other cognitive models; and
3. *The "operating principles" of the hardware*, that is, how new categorizations and sequences are constructed from previous coordinations, especially how reconstructing and holding active multiple categorizations allows us to categorize relations of identity, negation, causality, and correspondence.

Within the cognitive psychology community, these three aspects of coordination are often framed as dichotomies: Is thinking driven by rules or something else? (Smith et al., 1992)

Sometimes the debate concerns #1 versus #2 (implicit versus explicit) and at other times it concerns #1 versus #3 (learned versus innate). But such arguments often make the same assumptions about knowledge — identifying memory with encodings — as if innate rules are just compiled versions of learned rules, like the difference between the hardware logic of the CPU and the software logic of the programs. A contrasting argument is that the operating principles are different in kind from descriptions; they are not rules per se, but "things that happen". Piaget believed that such rules could not be taught explicitly because they are inherently embodied in actions (Nisbett, 1993, p. 3); they are processes that actively organize experience.

Once this perspective is included, many examples of such non-descriptive conceptualization can be found. For example, understanding irony or a pun involves apprehending an relation that is not in itself verbal, and may only with hesitation be expressed (with a sense of frustration at making the ineffable *a thing*, stating it in words). Another example is the imitation of an accent. Americans may conceive the patterning of British English and mimic it. This is done by apprehending the relations of pronunciation and stressed sequences in a coordinated "coupling" of perceptual-motor categorizations; people conceive the accent as *a way of performing*, without having to first describe the accent formally and carry it out instructively as a procedure.

In contrast, creating and using descriptions involve modeling a situation in some language or notation and using inferential steps to derive valid implications and new questions — performed either in our head by inferential conceptualizations or by a symbolic calculus as in an expert system. Nisbett's (1993) report about teaching statistical reasoning examines such descriptive modeling at work. But the problems he encountered highlight the different kinds of regulators. Of special interest are the logical patterns of thought close to the limits of the neural processor. It is difficult, for example, for some people to juggle the equivalence of "if p then q" and "if not-q then not-p" in their heads. Instead, a conceptualization such as "the semantic notion of obligation" allows holding the details of a problem active (as neural processes) and ordering them appropriately. These "pragmatic reasoning schemas" may exist without the person's articulation of the relations in formal terms (as stated in p's and q's) or even without an ability to execute on paper a logic proof requiring modus tolens. Instead, the person engages in a form of *concrete thinking*, arranging the elements of the situation in a mental model according to a *conceptual* scheme (Wu, 1995).

Furthermore, concrete thinking of this form can be taught by *describing* the conceptual schema and providing examples of how to use it. In this respect, the rule description is a sign post, which may or may not be discarded in practice. "The rules can be made more accessible by teaching examples of their use, and especially by teaching people how to decode the world in ways that make it more accessible to the rule system" (Nisbett, 1993, p. 11). Nisbett's terminology must be used advisedly here — "decoding" must be viewed as moving from a description to a situation conceptualization.

Nisbett's conclusion that "it is a mistake to try to found a theory of mental life on mere associations or connections between concretely-defined elements" (p. 12) can be viewed with Rebecca's experience in mind as affirming the idea of conceptual coordination. But it might be turned the other way. It is a mistake to try to found a theory of mental life on mere associations or connections between verbally-defined elements. For then all people would be like Dr P and expert systems. Indeed, the descriptive modeling approach has attempted to embrace all aspects of cognition within discrete, sequential, and often exclusively verbal conceptualization. This view has dominated how cognitive science itself is pursued, constraining what constitutes data, what kinds of mechanisms are considered, and what kinds of partial understandings are recognized as reportable and professional. In teasing apart the map (our descriptions) from the territory (our conceptualizing) and asking what remains to be done, we are challenged to recognize that we know more than we can describe. Consequently, models based on encodings will always be impoverished relative to the neurological processes we seek to replicate. Abstract descriptions may be the epitome of scholarly thought, but they are mere shadows of our concrete understanding.

Acknowledgements — The introduction section is adapted from a summary written by Stellan Ohlsson. Andrew Brook also provided many helpful comments. This work is funded in part by NYNEX Science and Technology, Inc.

References

Anderson, J. R., Reder, L. M., & Simon, H. A. (1996). Situated learning and education. *Educational Researcher, 4*(25), 5–11.

Barresi, J., & Moore, C. (1996). Intentional relations and social understanding. *Behavioral and Brain Sciences, 19*, 107–154.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press. (Reprinted in 1977.)

Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Amsterdam: Elsevier Science.

Bransford, J. D., McCarrell, N. S., Franks, J. J., & Nitsch, K. E. (1977). Toward unexplaining memory. In R. E. Shaw & J. D. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 431–466). Hillsdale, NJ: Erlbaum.

Buchanan, B. G., Sutherland, G., & Feigenbaum, E. A. (1969). Heuristic Dendral: A program for generating explanatory hypotheses in organic chemistry. In B. Meltzer & D. Michie (Ed.), *Machine intelligence* (pp. 209–254). Edinburgh: Edinburgh University Press.

Buchanan, B. G., & Shortliffe, E. H. (1984). *Rule-based expert systems: The MYCIN experiments of the heuristic programming project*. Reading, MA: Addison Wesley.

Clancey, W. J. (in press). *Situated cognition: On human knowledge and computer representations*. New York: Cambridge University Press.

Dewey, J. (1896). The reflex arc concept in psychology. *Psychological Review, III* (July), 357–70. (Reprinted in 1981.)

Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, MA: Harvard University Press.

Edelman, G. (1992). *Bright air, brilliant fire: On the matter of the mind*. New York: Basic Books.

Effken, J. A., & Shaw, R. E. (1992). Ecological perspectives on the new artificial intelligence. *Ecological Psychology, 4*(4), 247–270.

Frake, C. O. (1980). Plying frames can be dangerous: Some reflections on methodology in cognitive anthropology. In C. O. Frake (Ed.), *Language and cultural description* (pp. 45–60). Stanford, CA: Stanford University Press.

Freeman, W. J. (1991). The physiology of perception. *Scientific American, 264*(2), 78–87.

Gardner, H. (1985a). *Frames of mind: The theory of multiple intelligences*. New York: Basic Books.

Gardner, H. (1985b). *The mind's new science: A history of the cognitive revolution*. New York: Basic Books.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Hayes-Roth, F., Waterman, D., & Lenat, D. (1983). *Building expert systems*. New York: Addison-Wesley.

Hoffman, R., & Nead, J. (1983). General contextualism, ecological science and cognitive research. *The Journal of Mind and Behavior, 4*(4), 507–560.

Hofstadter, D. (1995). A review of mental leaps: Analogy in creative thought. *AI Magazine, 16*(3), 75–80.

Iran-Nejad, A. (1987). The schema: A long-term memory structure or a transient functional pattern. In R. J. Tierney, P. L. Anders, & J. N. Mitchell (Eds.), *Understanding readers' understanding: Theory and practice* (pp. 109–127). Hillsdale, NJ: Erlbaum.

Jenkins, J. J. (1974). Remember that old theory of memory? Well, forget it! *American Psychologiot, November*,785–795.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.

Lave, J. (1988). *Cognition in practice*. Cambridge: Cambridge University Press.

Luria, A. (1968). *The mind of a mnemonist*. Cambridge, MA: Harvard University Press.

Maturana, H., & Varela, F. (1987). *The tree of knowledge: The biological roots of human understanding*. Boston, MA: New Science Library.

Merzenich, M., Kaas, J., Wall, J., Nelson, R., Sur, M., & Felleman, D. (1983). Topographic reorganization of somatosensory cortical area 3B and 1 in adult monkeys following restricted deafferentation. *Neuroscience, 8*(1), 33–56.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard Press.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R. E. (1993). Reasoning, abstraction, and the prejudices of 20th century psychology. In R. E. Nisbett (Ed.), *Rules for reasoning* (pp. 1–12). Hillsdale, NJ: Erlbaum.

Ohlsson, S. A. (1993). Abstract schemas. *Educational Psychologist, 28*(1), 51–66.

Polanyi, M. (1966). *The tacit dimension*. New York: Doubleday Anchor.

Pylyshyn, Z. W. (1991). The role of cognitive architecture in theories of cognition. In K. VanLehn (Ed.), *Architectures for intelligence* (pp. 189–223). Hillsdale, NJ: Erlbaum.

Rommetveit, R. (1987). Meaning, context, and control: Convergent trends and controversial issues in current social-scientific research on human cognition and communication. *Inquiry, 30*, 77–79.

Rosenfield, I. (1988). *The invention of memory: A new view of the brain*. New York: Basic Books.

Sacks, O. (1987). *The man who mistook his wife for a hat*. New York: Harper & Row.

Schank, R.C., & Abelson, R. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum.

Schön, D. A. (1979). Generative metaphor: A perspective on problem-setting in social polity. In A. Ortony (Ed.), *Metaphor and thought* (pp. 254–283). Cambridge: Cambridge University Press.

Smith, E. E., Langston, C. & Nisbett, R. (1992). The case for rules in reasoning. *Cognitive Science, 16*, 1–40.

Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge: Cambridge University Press.

Varela, F. J. (1995). The re-enchantment of the concrete. In L. Steels & R. Brooks (Eds.), *The artificial life route to artificial intelligence* (pp. 11–22). Hillsdale, NJ: Earlbaum.

Vera, A. H., & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science, 17*(1), 7–49.

Wu, Kuang-Ming. (1995). Chinese concrete thinking in classical Taoism. *Chinese Studies* (June). History Department, National Chung-Cheng University, Min-Hsiung, Chia-I Taiwan, ROC.

Wynn, E. (1991). Taking practice seriously. In J. Greenbaum & M. Kyng (Eds.), *Design at work: Cooperative design of computer systems* (pp. 45–64). Hillsdale, NJ: Erlbaum.

Biography

William J. Clancey is a senior research scientist at the Institute for Research on Learning. His work has spanned the development of early medical expert systems, automated explanation, and instructional programs. He explores situated cognition in studies of "situated robots", neuropsychological dysfunctions, and social construction of policies and plans, emphasizing the differences between people and computer programs. He is developing a multi-agent simulation of work practice for use in work systems design and instructional settings, which models people's identities and activities.