Clancey, W. J. 1998. Interactive coordination processes: How the brain accomplishes what we take for granted in computer languages—and then does it better. In Pylyshyn, Z. (Ed.) Constraining Cognitive Theories: Issues and Options (pp. 165–190). Greenwich: Ablex.

Interactive Coordination Processes: How the brain accomplishes what we take for granted in computer languages—and then does it better

William J. Clancey

Institute for Research on Learning
2550 Hanover Street
Palo Alto, CA 94304
Bill_Clancey@IRL.ORG

Abstract

How can children and intelligent animals such as ravens and chimpanzees conceptualize and *perform* procedures without a *theoretical* language for *describing* what they do? Situated cognition theory suggests that we look for a mechanism that doesn't consist entirely of stored descriptions, or equivalently, a resequencing and composing process that doesn't require explicit instruction or reflective reasoning. Edelman's theory of neuronal group selection suggests that multimodal, sensorimotor coupling is the foundation. Chomsky emphasizes that this capability goes beyond physical habit formation—grammatical regularities are more than path finding (such as a bee's navigation) requires. This approach to developing an artificial intelligence seeks to relate neural processes to the best cognitive models, applying a means-ends analysis. To introduce this perspective, I analyze an example of sending two messages in an e-mail program. The example reveals a fundamental sequence-construction mechanism by which perceptual categories and motor schema are automatically generalized. By this mechanism, the human brain accomplishes more flexibly what we take for granted in stored-program computers—ordered steps (a sequence of operators in a problem space), variable bindings, conditional statements, and subgoaling. Remembering and acting are not literally problems of matching (indexing via labels) or searching, but reconstructing relations between processes. Conceptualization coordination is categorizing internal relations across modalities—verbal, gestural, visual, rhythmic—sequentially or simultaneously. These neural processes only exist as physical organizations when they are active and must be contrasted with a parametrized description of how the world or behavior appears—the textual representations of "symbolic models."

> ...it was a long time before I fully realized the importance, for many psychological experiments, of putting the situations which are used to produce response into sequential form.
>
> Sir Frederic Bartlett, *Thinking*, 1958, p. 141

Background

By the mid-1980s, many AI researchers began to realize that the mainstream approaches to developing intelligent machines were fruitful, but incomplete. For example, knowledge engineering—the process of codifying human knowledge in expert systems—makes a solid contribution to computer modeling, but such programs neither replicate in all important aspects nor explain the full capability of human intelligence. A means-ends approach to engineering an artificial intelligence machine now suggests that we focus on the differences between human capabilities and the best computer programs. These differences suggest two basic limitations in the "symbolic" approach of game-playing programs, expert systems, natural language processing, planning, etc. First, human memory is much more than a storehouse where structures are put away, indexed, and rotely retrieved. Second, human reasoning involves more than searching, matching, and recombining previously stored descriptions of situations and action plans. Indeed, these hypotheses are related: Remembering and reasoning both involve *reconceptualization*. That is, human learning occurs with every thought; learning is not just a reflective process of formalizing knowledge after deeds are done.

For many computer scientists the phenomenon to be replicated—human intelligence—and knowledge are placed in new relation. Previously, we explained intelligent behavior in terms of knowledge stored in memory: word definitions, facts and rules modeling complex causal processes in the world, plans and procedures for doing things. Now, we view these descriptions as *descriptive models*—which people use like maps, handbooks, and standards for guiding their actions. These descriptions, like the rules in an expert system, are not knowledge itself but *representations of knowledge*—a step removed from what is occurring in the brain. Like an article in an encyclopedia, such descriptions require interpretation in the context of use, which in people involves not only generating additional descriptions, but conceptualizing in non-verbal modes. As Newell (1982) put it, “Knowledge is never actually in hand...Knowledge can only be created dynamically in time.”

The change in perspective can be bewildering: no longer are symbols, representations, text, concepts, knowledge, and formal models equated. Symbols in the brain are not immutable patterns, but structural couplings between modular systems, continuously adapted in use by a generalization process that dynamically recategorizes and sequences different modalities of sensation, conception, and motion (Merzenich, 1983). Representations are not just coded artifacts such as diagrams and text, but in the brain are active processes that categorize and sequence behavior in time (Edelman, 1992; Rosenfield, 1992; Bickhard, 1995). Conceptualization is not just verbal, but may be rhythmic, visual, gestural—ways in which the neural system categorizes sequences of multimodal categorizations in time. In the view of Howard Gardner (1985), people have multiple intelligences, with different ways of making sense and coordinating action. The neuropsychological studies of Oliver Sacks (1987) reveal how people who cannot verbally abstract their experience may nevertheless dance and speak poetically; and people who speak in the particulars of geometric forms and rules may be unable to see the forest for the trees. Indeed, our textually-rooted robots, like these patients, once loaded with a knowledge base and shoved out into the world of shopping malls, doctor’s offices, and national parks, may appear to be dysfunctional morons or *idiot savants.*

With these observations in mind, we seek to invent a new kind of machine with the modularity and temporal dependence of neural coordination. Some connectionists are exploring computational recognition systems based on dynamic coupling, in which subsystems exist only in mutual reinforcement (Freeman, 1991). Structures in the brain form during action itself, like a car whose engine’s parts and causal linkages change under the load of steeper hill or colder wind. Such mechanisms are “situated” because the repertoire of actions becomes organized by what is perceived—and what is perceived depends on the relations of previous coordinations now reactivating.

In retrospect, we view symbolic models (e.g., natural language grammars) as descriptions of gross patterns that “subsymbolic” systems produce in the agent’s interactive behavior. The “symbolic” AI of the 1950s through the early 1980s is not wrong, but it may be a bit inside out. People do articulate rules, scripts, and definitions and refer to them. These descriptive representations are not in their brains, but on their desks and in their libraries. The representations in the brain are different in kind—dynamically formed, changing in use, categorizations of categorizations *over time,* and multimodal (not just verbal), as means of *coordination*. Meaning can be written down (described), but words are not concepts; the map is not the territory.

This paper illustrates how viewing human experience in a fresh way, applying a means-ends approach, may be fruitful for suggesting new kinds of physical coordination mechanisms. A simple example is analyzed in great detail to bring out properties of sequencing in practiced behavior. These properties are described by a notation called an ICP diagram, representing the "interactive coordination processes" of neural maps. This diagram allows us to contrast neural processes with aspects of stored computer programs, which we take for granted in our symbolic cognitive models: ordered steps, variable bindings, conditional statements, and subgoaling. The analysis begins to reveal how the brain accomplishes these structures and relations more flexibly, and hence constitutes a new kind of computational device—one we don't fully understand and can't yet replicate. Finally, the example reveals how talk about "representations" in the symbolic approach confused stored text and diagrams ("symbol structures") with active, always adapted processes, namely perceptual categorization, conceptual categorization of sequences of activation, and non-verbal forms of conceptualization (especially rhythm and imagery). The advice for AI research is to investigate how the brain works, especially to explore how "multiple intelligences" (especially non-verbal) are interactively formed and coordinated in everyday behavior.

Sending two messages: Forms organized by active processes

I wish to examine in some detail how I typically used the Xerox-Lafite electronic mail system in the early 1980s (for simplicity I will use present tense). The process is as follows: I decide to send a message to a particular person, and select "Send Mail" in the Lafite control window (Figure 1, top center). Then, perhaps before typing anything in the message window (W1), I remember that I want to send a message to someone else, so I button "Send Mail" in the control window a second time, in order to produce a second message window template (W2).

Why do I create a second window rather than using the first one? Why do I find it so easy to later return to the first window and send the message I first intended, when it may be unlabeled and I have no notes about what I intended to do with it? Why does it feel difficult to hold the first message in mind and reassign the first window to the second message?

The impression I have is that the first window, W1, is assigned; it embodies my intent; it has meaning already. I see W1 as "a message to person 1." Using W1 to send a message to the second person means seeing it in a new way. My feeling is that this disrupts my thoughts. It is much easier and obvious to create a second window.

Using the terminology of Bamberger and Schön (1991),the physical structures visible on the screen, W1 and W2, are *reference entities*.

> A reference entity serves to single out, externalize, hold for current attention some emergent object or relation. A reference entity serves the function of on-the-spot naming within, and as part of, the making process.

Related to these reference entities are activities, which I describe as "the process of sending a message to person 1" and "the process of sending a message to person 2." Crucially, I see W1 and W2 as being related to these activities, what I am currently doing and planning to do. Somehow, through neural processes inaccessible to me, W1 and W2 are bound or viewed as part of different activities.

To reassign W1 is to see it another way, to reperceive it. But because W1 embodies the process of sending a message to person 1, to see W1 in a different way is to *disrupt that process*. This is a crucial observation: Seeing materials in a different way may require disrupting an ongoing neural process, an active coordination between what we are seeing and what we are doing, that otherwise has no observable manifestation. Significantly, "doing" has the larger sense of what I am doing during this current electronic mail session. In some sense, I am *still sending* that message to person 1, even though I have shifted my focus to person 2; the neural aspect of seeing and sending that message is still active. In cognitive jargon, we would say that I am still "intending" to send message 1.

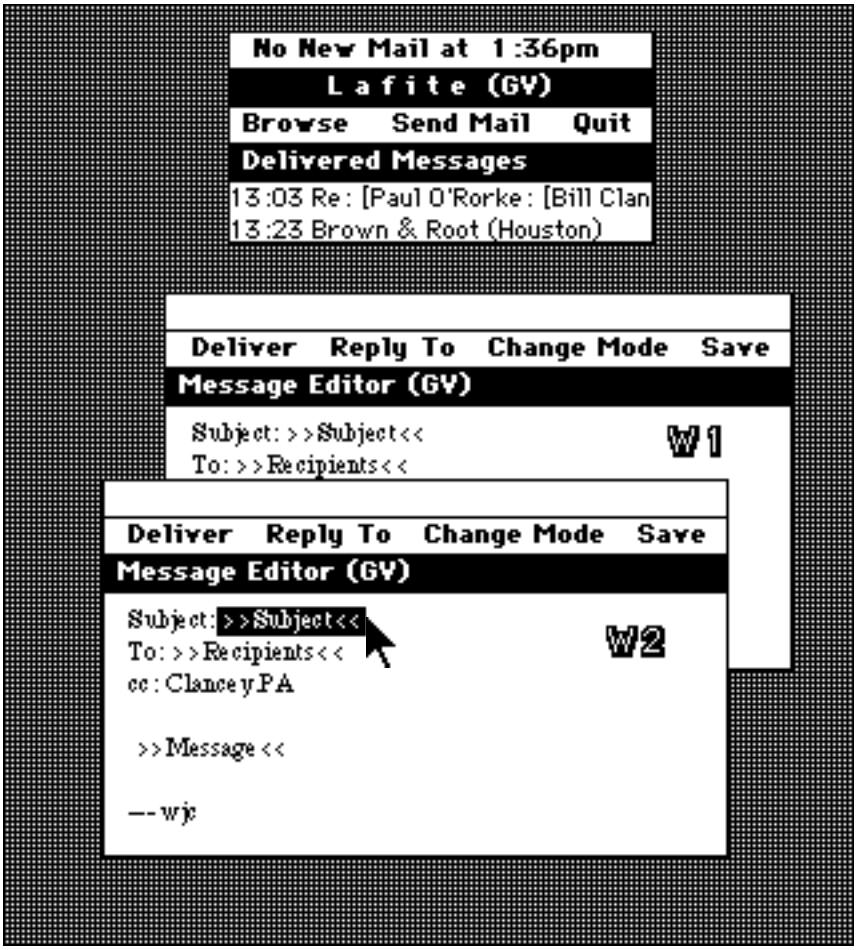
No New Mail at 1:36pm

13:03 Re: [Paul O'Rorke: [Bill Clan
13:23 Brown & Root (Houston)

Subject: >>Subject<<
To: >>Recipients<<

Subject: >>Subject<<
To: >>Recipients<<
cc: ClanceyPA

>>Message<<

--wjc

Figure 1. Example of two active "Send Message" windows, labeled W1 and W2.

Instead of disrupting the first process, I *hold it aside* and start up a second process, embodied by W2. I send the second message and then return to W1. The windows on the screen serve as my external representation of messages I want to send. I don't even need to type in something in the TO: or SUBJECT: fields in W1. I can *remember how to see* W1. *I know what it means.*

The following aspects of this coordination are striking:

1. Visible spatial reminders on the screen allow me to coordinate multiple activities. I can interrupt what I am doing, both physically and mentally setting it aside, and pick up the activity later where I left off. I don't disrupt the interaction, I hold it active, keeping it current but making it peripheral.
2. The capability to create multiple processes in the electronic mail program facilitates my mental coordination. Activating the second "Send Mail" process in the mail program parallels and facilitates my allowing a second activity (sending a message to person 2) to take form. Indeed, the idea of setting aside message 1 and beginning

message 2 is afforded by the concreteness of W1 now in place, and the “Send Mail” button, allowing a repeated operation to occur. Selecting back and forth between W1 and W2, by buttoning the mouse in each window in turn, feels like shifting my attention back and forth between the activities of sending the two messages. The effortless flow and context shift as windows shift forward and backward parallels my experience in thinking about the messages. The window I am seeing is the activity I am doing.

3. Completing the second message *first* is a bit odd. The impression is of a stack: I need to get rid of the second message before I can return to the first. The feeling is that the second message is more pressing. If I don't do it now, I will forget it. The idea of sending the second message arose effortlessly from the process of sending the first; but somehow the second message, as an interruption that only came to mind as a second thought, needs to be handled immediately.

Discussion: Figure-ground priorities

To develop these points further, consider that many alternative ways of using the Lafite system when you realize that you want to send two messages. Table 1 summarizes first how one might reassign the windows or reorder the sending of messages.

Table 1. Alternative ways of sending two messages.

	Keep interpretations	**Reassign windows**
Keep order	W1 <- msg1 W2 <- msg2	W2 <- msg1 W1 <- msg2
Reorder messages	W2 <- msg2 W1 <- msg1	W1 <- msg2 W2 <- msg1

For example, one might keep the order (sending the first message first), but reassign the windows, so the first message is typed into W2, and one returns to W1 to send the second message. In the two cases where one uses W2 second, one might create W2 before or after entering a message into W1. For example, one might reassign W1 to the second message, then create a second process for the first message *after* sending off the second message. This gives six basic combinations. But there are other possibilities:

- Type in reminders about the contents of the first message before creating W2.
- Type in reminders about the contents of the first message after creating W2.
- Immediately type in reminders for the second message in W1, then create W2 and use it for sending the first message.

The coordination I developed was to reorder messages, but keep interpretations without reminders. Assuming that this selection bears some relation to neural processes, what conjectures can we make about how neural processes are activated, held active, coordinated, and reactivated?

First, in my experience, *it is difficult to reassign* W1. The impression is that I would lose the process of sending the first message. There is “no way to remember it,” aside from reaching for a pad and making a note. It is difficult to switch contexts, to see W1 as being about (or for) the process of sending a message to person 2. The meaning persists; perceived forms are coordinated with activities. To destroy or reassign perceived objects is to lose the coordination: To lose a way of seeing is to lose an activity.

Second, *it is difficult to do message 1 before message 2*. Message 2 is active now, after "Send Mail" was selected a second time. I'm doing *that* now. I don't want to disrupt message 2 now. There would then be *two disruptions*. I need to keep going forward. (Later, we relate this observation to the inability to "move noun phrase references," a key pattern in Chomsky's Universal Grammar.)

It feels difficult to *go back* to message 1, given that message 2 is in process already. Hence, I can hold in abeyance message 1 and begin message 2 easily, but I resist going back to message 1! My sense is that I am finishing what's active now (and visually, it's the foreground window). Consistently, I sometimes (but rarely) start up a third message process and do it first. I have a sense of "the current process"—what I'm doing now.

Oddly, it seems easier to remember what I intend to do with W1 than to make the effort of typing in some TO: or SUBJECT: information as a reminder. Typing anything in W1 after already pulling down W2 feels like actually doing the first message—going back to it. Crucially, I immediately pull down a second window by buttoning "Send Mail" at the very moment that I realize that I want to send a message to a second person.

Thus, we have the curious and fortunate balance: It is difficult to reinterpret forms that are part of a sequence of activity ("a chain of thought"). But it is relatively easy to recall how materials were previously perceived and engage in that process again. These possibilities work together here because interrupting the first process is perceived as a "holding aside," not a disruption, by the representation on the screen. The first message activity is held, not reorganized or ignored; it is visible peripherally on the screen. The activity is set aside like an object in the physical world, and can be picked up again at the click of the mouse.

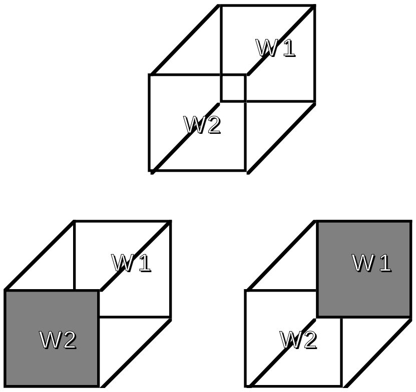

Figure 2. Figure-ground relations between W1 and W2 shown as a Necker Cube.

The effect is like illusions that reorganize the visual field, such as the Necker Cube, the duck-rabbit, old-young woman, etc. (Figure 2). Buttoning one of the windows makes that activity the *figure* and the other activity part of the *ground*. Fixing a figure-ground relationship on the screen by bringing W1 or W2 to the fore corresponds to making that activity *"what I am doing now"* and every other process in which I am engaged *peripheral*. The figure-ground relation appears to be basic to the idea of focus of attention. Referring to the discrete manner in which we see forms as one thing and can shift back in forth, Arbib (1980, p. 33) suggests that "The inhibition between duck schema and rabbit schema that would seem to underlie our perception of the duck-rabbit is not so much 'wired in' as it is based on the restriction of the low-level features to activate only one of several schemas."

Why did the reminder occur?

To better understand how the neural processes are created and held relative to each other, we will introduce and use the ICP notation to show how the W2 process developed in the context of W1. Why did I experience the idea of sending a message to the second person, just when I was about to write to the first? When such an interruption occurs, I experience an image or name relating to the second person, a phenomenon which Bartlett identifies with the process of recollection. The activity of sending a message apparently enables an effortless beginning of the process of sending a different message. Building on this clue, and Bartlett's analysis I illustrate the relation of the two processes (Figure 3).

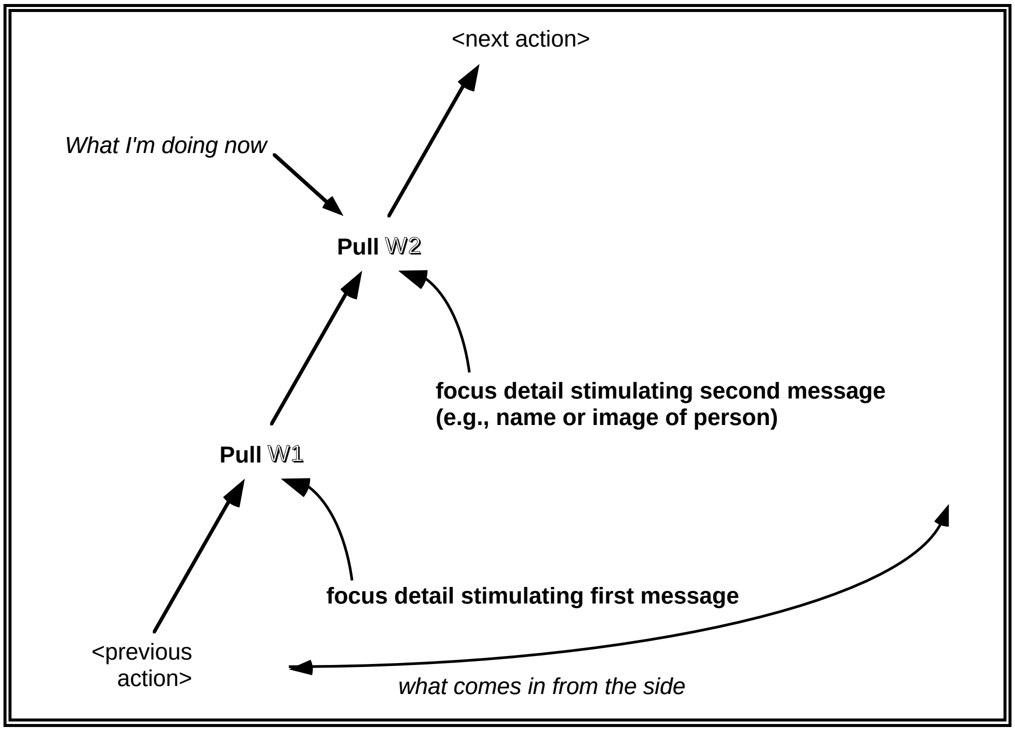

Figure 3. An ICP diagram: Process of sending second message viewed as a generalization of the active process of sending a message.

Figure 3 shows the coordination just after I have pulled down W2. Just after pulling down W1 a detail came to the foreground (e.g., the idea of another person or message). The diagram shows this second "focus detail" coming in "from the side." Because of the materials available to me, I didn't experience a sense of disruption (an impasse). For example, without the multiple-process capability of Lafite, I might have been forced to turn aside and write something down about message two on a sheet of paper. Instead, I immediately repeated the current activity of writing a message, pulling down W2. By this coordination, I incorporated the detail; I made it "what I am doing now." The relation is non-symmetric: I can button back and forth to *see the windows as different activities*, but I resist *actually doing* the first message activity after I have pulled down W2.

Notice that the constructive process is not a matter of seeing the first activity in a different way (W1 is not reinterpreted). Instead, the second activity is built around the spontaneously created idea of another message (and commentary about it, "Oh yes, I intended to send a note to Mike"). Since there is in general no content relation between the two messages, it is apparently the process of sending the first message itself—the active neural processes after W1 was pulled down—that brings the new focus detail to the foreground. The effect is that I "recall" the need to send a second message.

The two activities are parallel and coordinated, as if they indeed integrated in my understanding. Figure 3 is intended to represent this sequence; in some sense *what I am doing now includes sending the first message*. Narrowly, my focus is on the second message, but broadly, my coordinated activity includes the first message. Furthermore, the activity of describing the second message or pulling down W2 is inseparable from the process of including the new focus detail coming in from the side in my activity. In terms of the ICP notation, the node labeled “pull W2” represents a neural process that is activated just as the new focus detail comes to my attention. That is, there is a dialectic relation between the idea of sending a second message and the activity of doing so; in my experience, they come into being together. The new composition, represented by the node relating “pull W1” and the focus of the second message, is created by the activity of sending the second message.

In contrast, a serial or parallel architecture would involve a different timing and different kinds of representation: The idea of sending a second message would be described internally by data structures that included some kind of reference to the general description of the procedure of sending a message, plus a reference to the topic of the second message. Reaching to button “Send Mail” would be interpreted as executing this procedure, and Figure 3 would be viewed as a representation of a stack, showing the order in which procedures and “arguments” were “bound” and “invoked.”

Neural map activation over time

Figure 3 is not a representation of a stack, but of a sequence of links between neural maps. The ICP notation shows the relation of neural processes constructed over time by activating and composing maps. I conjecture that “Pull W2” literally reuses the neural maps involved in “Pull W1,” but coordinates them with different perceptual figures. Just as perceptual forms can be “reused” by different visual organizations (Figure 2), neural maps can be “reused” by different neural organizations. Just as I can slip back and forth between how I see the display or how I conceive of my activity, a given set of neural maps can be simultaneously reorganized into a different, ongoing coordination.

Crucially, the production of one neural organization from the previous one is embodied in the timing of neural activations, so the next-next-next sequence is biased to be reconstructed. That is, the activation of a perceptual-motor map will activate the next map in sequence by virtue of having neuronal groups that are physically activated by groups in the previously active map. In the example of sending two messages, a given coordination is repeated, but composed with different perceptual and conceptual maps. This might be the most simple kind of sequencing: Doing something again, but acting on different objects in the world.

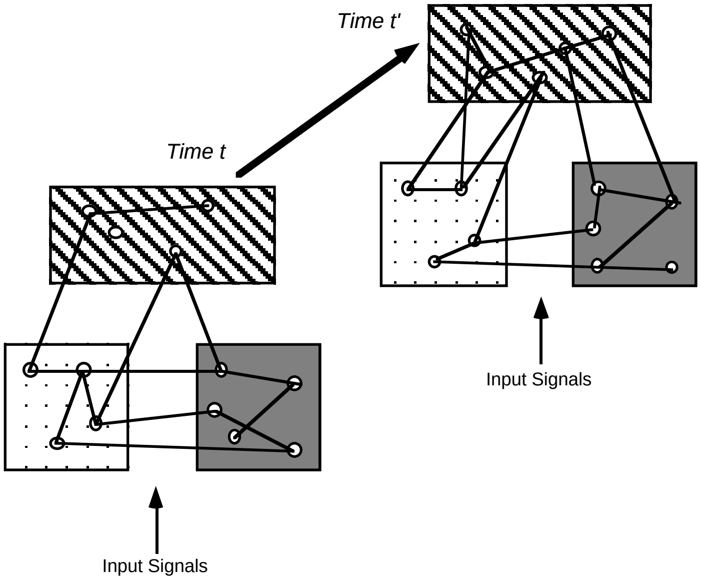

Figure 4. Memory as recategorization: Inputs categorized similarly produce similar “output” relations. Circles are neuronal groups (perhaps thousands of neurons); lines indicate bi-directional activation; squares are “maps”; rectangles are maps of maps (adapted from Edelman, 1992; Figure 10-1, p. 103).

In this analysis, we are going beyond Edelman’s focus on multimodal coupling of neural maps to consider how map activation occurs over time. Figure 4 illustrates how a given set of maps are reactivated at a later time. With similar inputs, there is similar output. The groups and the reentrant links that are activated are always prone to change, because of changes in the external environment and changes in the correlations activated throughout the neural system. Broadly speaking, TNGS claims that every activation is a *generalization* of past activation relations. Put another way, every activation is a *recategorization*, as opposed to a literal match or retrieval operation.

This recategorization is always occurring simultaneously within a larger sensorimotor coordination, involving conceptualization (at the level of maps of types of maps) and, often in people, verbalization. By conjecture, *memory for sequence of activation* involves one map of maps activating another over time (Figure 5). This corresponds to Bartlett’s (1932) notion of a schema as a *coordinating process*.

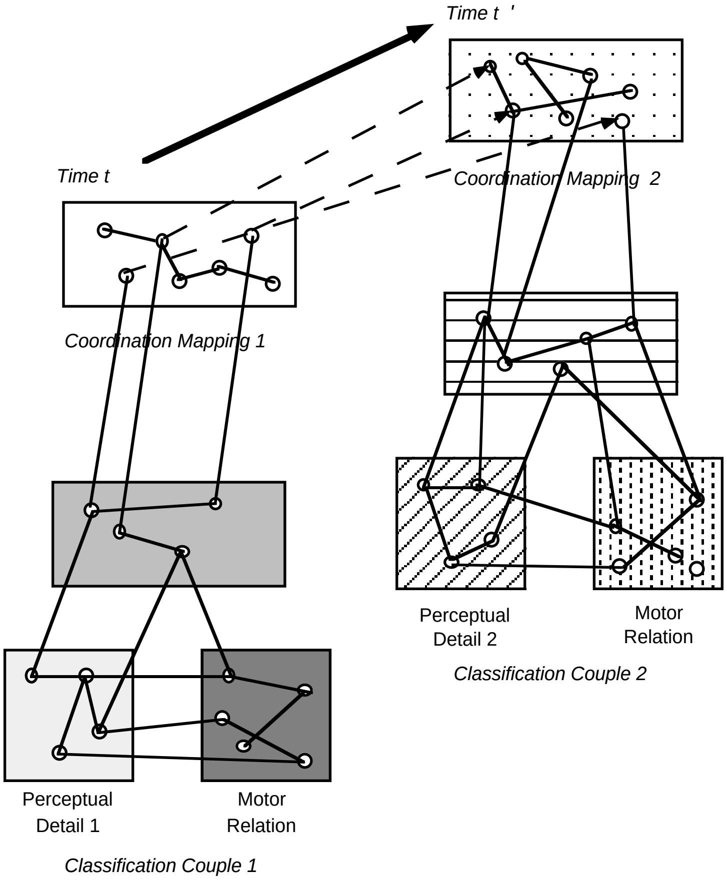

Figure 5. Higher-order maps, including categorizations of global (cross-modal) maps, physically activate in sequence over time, serving as the memory of how perceptual details are coordinated within an activity. Solid lines indicate co-activation in a circuit; dotted lines indicate activation over time.

In Figure 5, we see *different* maps of maps linked to *different* high-order maps. A perceptual detail may be a classification of a sound, a color, an object, a word, etc. The higher-order maps are physically linked by reentrant connections. However, the second higher-order map (labeled *time t'*) only becomes activated after the first map has become active (labeled *time t*). In this manner, it is conjectured that *maps of maps are sequenced over time via physical subsumption*. By conjecture, the number of possible connections in the brain is so large that there are sufficient links to establish a link between any sets of maps, as required by activity. Furthermore, all learning is categorization over time. This means that all learning is activation "in place" of pre-existing links, which are composed over time into, and always within, sequences of activation. As Edelman states, there must be some way of holding such activations so they can be compared and further coordinated, allowing for continuity in our experience.

This analysis is driven by two key observations about memory: First, Bartlett describes new conceptions as being *manufactured* out of previously active coordinations. How ways of seeing, moving, talking, etc. have worked together in the past will bias their future coordination, in the sense that they are physically connected again by the same relations. Second, I am building on Vygotsky's observation that every coordinated action (or descriptive claim, image, conception) is a *generalization* of what we are already doing (saying, seeing, categorizing). Putting these together, we tentatively view the ICP notation as showing a composition process over time (corresponding to inclusion of neural maps of the previous active coordination), such that each composing along the way is a generalization, qualified by details that specialize the neural construction, making it particular (focused on a thing) and unique.

Intention: Incorporating processes in higher-order categorizations

The coordination of sending two messages can be compared to a deliberate attempt to recall messages I need to send. Again, phrasing this in terms of Bartlett's model, I would seek and construct first a detail (e.g., an image of a journal paper might come to my attention) and then knit it together with "main stream of interest" through a story relation (e.g., "I need to get back to Wendy"). I would probably create a list on paper of these activities.

In sending two messages, a generalization occurs immediately, without a sense of disruption (impasse) or need for a rationalizing story, because the new construction involves a *reactivation* of what I am currently doing. No seeking or manipulation of materials and reperceiving is necessary. Coordinating occurs when I say to myself, "Oh yes, I intended to send a note to Mike." By this statement, I mark what I am doing now. The process of sending a second message is constructed directly from my currently active processes, my coordinated activity of looking at the screen and manipulating the mouse. I can compose the detail with my main stream of interest (the activity of sending messages) by simply doing again what I am already doing; I simply repeat the action of bringing down another message window.

By this analysis, the process of sending a second message is not a recursive procedure call, in the sense of storing a description of the procedure and its previous arguments (the topic of the first message). Instead, the neural links create physical structure "in place," such that a sequence of maps is constructed that includes the different ways of viewing the screen and my current activity. To do this with computer programs, we would have one copy of a procedure, *temporally* associated with different arguments (data). That is, the focus of a procedure is dynamically activated within an embracing physical structure, such that the containing structure *conceptually* relates the activity and the focus. In this way, I view the ICP diagram as representing a sequence of conceptual compositions over time, with each new node in the sequence representing the concept of "what I am doing now." This conceptualization (in these examples) simultaneously incorporates a particular focus, while broadening the previous conceptualization. Here, the broadening is such that "what I am doing now" (namely, sending messages) now includes two messages, but has a specific focus of the second message.

Significantly, in this simple example, the new focus detail *bears the same relation* to "Pulling down W1 " as the first focus detail. I am able to see the second focus detail as *an instance* of "what I am already doing" (sending a message). The process of "sending the first message" (involving buttoning and following hand-eye, typing, information-entering processes) is reactivated, but I do it with respect to a different figure (the name or image related to Mike that came to my attention). We can say that the two focus details (essentially the topics of the two messages) are *correlated* by their composition with the conceptualization of sending a message. But holding the two details together, the activity of deliberately relating the two topics (a crucial capability constituting higher-order consciousness), is very different from being in the activity of sending two messages. Grounded in the activity, the details are mutually exclusive. In *reasoning* about alternative actions they are both active and related as details within an encompassing conceptualization of ordering, correspondence, cause, negation, identity, etc. (the fundamental reasoning processes emphasized by Piaget).

Shifting from one message activity to another, the previous focus becomes part of the background; the associated neural maps are still active, but I am not conceiving "sending W1" as what I am doing now. This is analogous to looking at the cube in Figure 2: "I am looking at the cube and seeing it as W1" or "I am looking at the cube and seeing it as

W2." An activity, a way of coordinating perception and action, is reenacted, but incorporating different perceptual categorizations.

In other circumstances, I might realize that the second focus detail might be another thing to mention in my first message, or something I can ask the first person to relay to someone else. In this case, the result would be a new composition of "What I am telling person one now," with the second detail included. I conjecture that we build up "next-next-next" relations in this way—by a composition process—and have the capability to return to an earlier process-state and "read out the details" (going forward) by reactivating containing generalizations. In this respect, our use of stories and how we construct stories—in the sequence of naming, framing, history-telling, and design (Schön, 1979)—has an underlying basis in this neural composition mechanism.

The tree form of the ICP notation (Figure 3) should not suggest that there is a data structure that is being traversed, inspected, or otherwise manipulated by some other process in the brain. Just as in perceptual reorganizations, in which lines get reorganized to form different angles and objects, neural maps get reorganized to form different figure-ground relations to foci and conceptualizations of what I am doing now. When I button back and forth between W1 and W2 I am recoordinating my activity. I conjecture that each buttoning back and forth results in *a new composition* higher on the path shown in Figure 3. Indeed, the conceptualization of what I'm doing now becomes, "I'm buttoning back and forth on the screen" *(a categorization of the sequence)*. That is, every activity must incorporate what is currently active as either figure or ground. The mechanism is not a process of moving pointers up and down a tree or restoring a previous organization by writing it into a buffer for execution. We are always moving forward, creating a new composition, albeit often a generalization of what we have done before.

Distinguishing between awareness and representations

The example of sending two messages illustrates that statements like "how the subject is representing the windows to himself," must be distinguished from "how the subject is *describing* the windows to himself." Indeed, like the fly coordinating his tongue with a passing fly or the blind man using the cane as an extension of his arm, a practiced message sender is perhaps not *representing* the windows any more than I am representing the location of the mouse on the desk, which I cannot see.

The experienced message sender is detecting previously categorized forms while coordinating the cursor and typing, with an activity of composing a message. Our experience indicates that more and more of the screen becomes peripheral through practice. We are *apprehending relations*, and no longer *seeing forms*. The distinction between conscious and subconscious is then reframed as a distinction between figures and contrasts and peripheral processes.

Contrast my experience in sending two messages with my asking an assistant to finish my work: He is given two identical, blank windows, already open on the screen. He is told to send two messages to different people. Window position has no representational significance, except he will probably use the front one first. Now suppose I require that he assign the windows in the same manner as I saw them when they were created. Now I must represent these windows for him: "This is the window for Mark's message; here is the one for Wendy." Without a descriptive set of instructions, my friend can't replicate the activity I was engaged in: He needs to know what the two open windows represent.

But the windows weren't representations to me, they were just part of my activity of sending messages. Commonly, we would say that in sending the two messages I knew what the windows represent (or similarly, "what they mean"). But this means that if asked, *I could tell you what I am doing*, not that I view the windows *as being representations.* For me, they are places in which I am composing messages—not representations of messages, but the real thing. I am not perceiving the forms as *windows*, but as *spaces* in which to type.

Given that the windows are identical except for spatial location, observers will say that they are examples of indexical representations (like the word “now”). Meaning is embedded in use; the process of seeing-as “registers” each of the windows as different activities. But for me they aren’t representations at all! They are reference *entities*, not references. They *are* the messages I am in the process of sending, not representations of them. Engaged in my activity, unlike the message program I don’t have a map or registration table that associates the windows with my message-sending processes. Their meaning is embodied in my activity, and held by active processes of categorizing.

The conceptual relation of my understanding “what I am doing now” isn’t *between descriptions*. Nor does my understanding relate neural processes and the world. Instead, seeing places on the screen is inseparable from interacting with them, what I am *doing* with them. As Ryle said, demonstrating know-how (doing) is not two things: knowing-that (articulating facts) *plus* doing. What I am perceiving, my way of categorizing my experience, is literally part of the construction of what I am doing. I don’t describe what I am doing when I return to W1; I simply reconstitute that process from the ground, making that place and the topic the figure of my conception.

Relation of the ICP notation to stored-text procedural models

What does this analysis reveal about stored-program computer languages? We start with the null hypothesis that there is fairly direct mapping between hierarchicalization of subgoals, impasses, operators, and chunking in symbolic cognitive models and the activation relations of neural structures. Very likely, the architecture we seek will not be constructed out of structures that cleanly map onto a modeler’s domain terms and relations. We take care to distinguish categorized details (especially words) from categorizing processes that incorporate such details. In general, there is no simple mapping between words and conceptualizations. Nevertheless, we assume that conceptual coordinating involves recurrent relations between physical (neural) structures, which are stable recategorizations of past neural activations, at some level corresponding to words, clichés, and the patterns of logical thought. By this, we shift from a stored-structure *manipulation* view to a process *construction* view.

The ICP notation is useful for comparing grammatical models to the flexibility we claim is actually present in human conceptualization (Figure 6). Essentially, the primitive components of computational models—ordered steps (a sequence of operators in a problem space), variable bindings, conditional statements, and subgoaling—can be contrasted with classification, sequencing, and composition of neural maps. Furthermore, coordinated sequences constructed multidimensionally (from existing sequences, corresponding to grammatical orders) would be like remembering stacks (a sequence of procedural calls and variable bindings). In procedural terms, an ordering of steps and binding constraints could be reestablished without running the procedures themselves; this is what Soar’s chunking provides. In terms of the ICP notation, a sequence of perceptual classifications (a next-next-next path) could be reestablished without recalling the specific actions that led to these perceptions in the past.

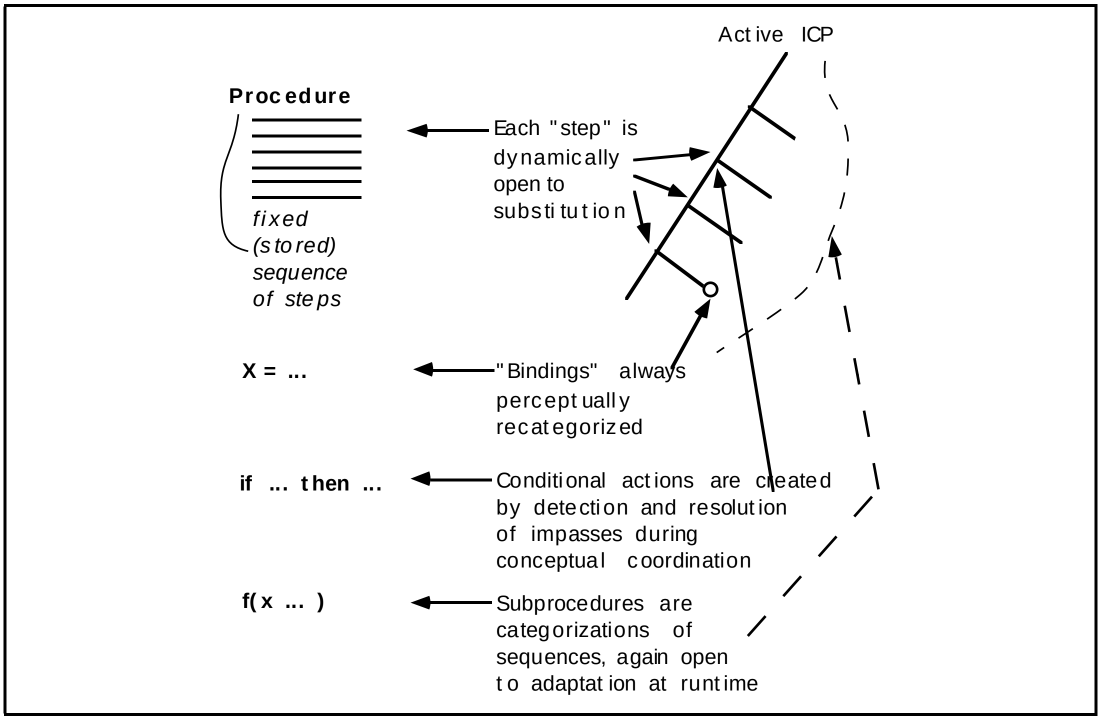

Figure 6. Relation of stored procedures and rules to ICP construction

The ICP notation is a tool for visualizing where stored descriptive models are inflexible relative to the TNGS. It also reveals what Soar assigns to the "cognitive architecture," but doesn't explain. Fundamentally, we need the ICP account to explain not only where the knowledge in the lower problem spaces comes from, but *why an operator hierarchy is constructed at all* and how operators—situation-action coordinations occurring at different points in time—can be *held active and related.*

Diagrams of classification couples (Figures 4 and 5) are a bit unwieldy. The ICP notation combines the maps of maps (rectangles) and secondary classification (e.g., coordinated motor relations) into a single node (Figure 7). Emphasis is thereby given to the sequence of perceptual details in the person's experience over time.

Each sequence node (e.g., A1) in the ICP diagram represents a conceptualization of global mapping. Below is shown the perceptual detail (e.g., P1) included in the associated global-mapping activation. This detail is the focus of attention (illustrated by the example of sending two messages). The next node (e.g., A2) is activated in sequence, over time. This global mapping is partially activated by the new environment, here represented exclusively by P2, another perceptual detail. Crucially, A2 and often P1 itself are activated by physical, reentrant links from the preceding global map, A1 (compare to Figure 5). In people, this coordination is often associated with linguistic processes, so the person could say "what I'm doing now." In using the ICP notation to represent problem-solving protocols, we think of each transition (e.g., from A1 to A2) as being a shift in "what I'm doing now," accompanied by a shift in perceptual detail. More generally, the ICP notation shows simply *a sequence of global mapping activations*, which in practiced form would not be accompanied by verbalizations (e.g., as in playing a musical instrument). Of special interest is how perceptual details in a sequence are related.

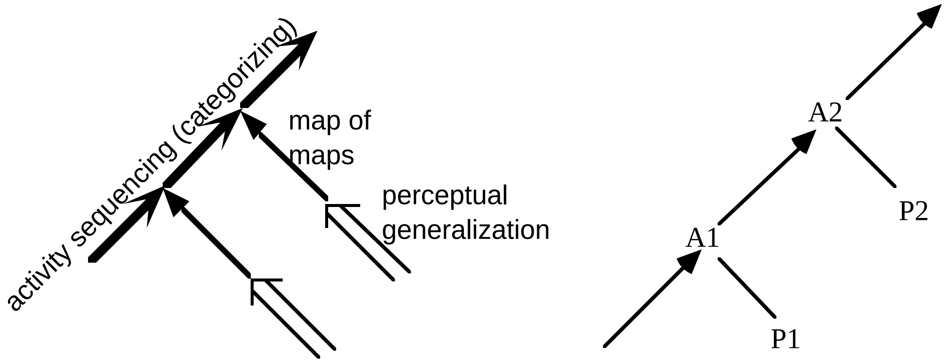

Figure 7. Relation of ICP diagram to TNGS: Ai represent global maps or categorizations of global maps; Pi represent perceptual details.

Categorizing sequences of activation (e.g., A1-A2-) is a form of *conceptualization*, corresponding to creation of a subgoal hierarchy, such as Neomycin's task hierarchy (Figure 9). The categorization itself enables the capability to "activate or reconstruct portions of past activities of global mappings of different types—for example, those involving different sensory modalities. They must also be able to recombine or compare them." (Edelman, 1992, p. 109). For example, my conceptualization of "Pursue-Hypothesis" includes speaking a name for the idea (itself a complex perceptual-motor coordination), a visualization of the name in a tree diagram, and the first step of checking for exposure to a disease (that is, my understanding is now tied to my formal representation of the process).

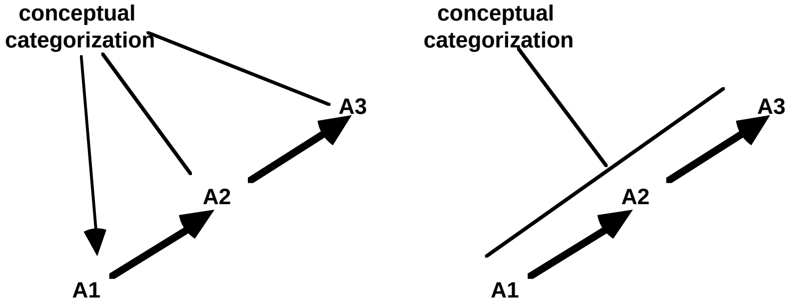

Figure 8. Conceptualization as categorization of sequences of activation. Left diagram shows possible "feed-forward" activation to the head of the sequence, with couplings between the global map of the conceptual categorization and subsequent global map categorizations. We use the equivalent and simpler notation on the right.

Mapping a symbolic program's hierarchies onto neural processes is an attempt to explain how subgoal hierarchies are learned, or equivalently, in Soar's terms how new problem spaces are learned. An key insight is that construction needn't operate on textual representations, as both the dominant inductive and deductive models of machine learning suppose: Compositions may develop bottom-up from perceptual details—*in the very process of constraining what details are perceived.* In terms of the ICP notation, the sequence isn't included necessarily under something else that preexists, which is "running" or "executing" and monitoring every step. Instead, the brain is constantly creating a generalization of "what I am doing now," and hence constructing orderly, coherent sequences. An orderly, and eventually multi-leveled, pattern develops by the bias to redo a previous sequence, resee a kind of detail, and reestablish a cross-modal coordination.

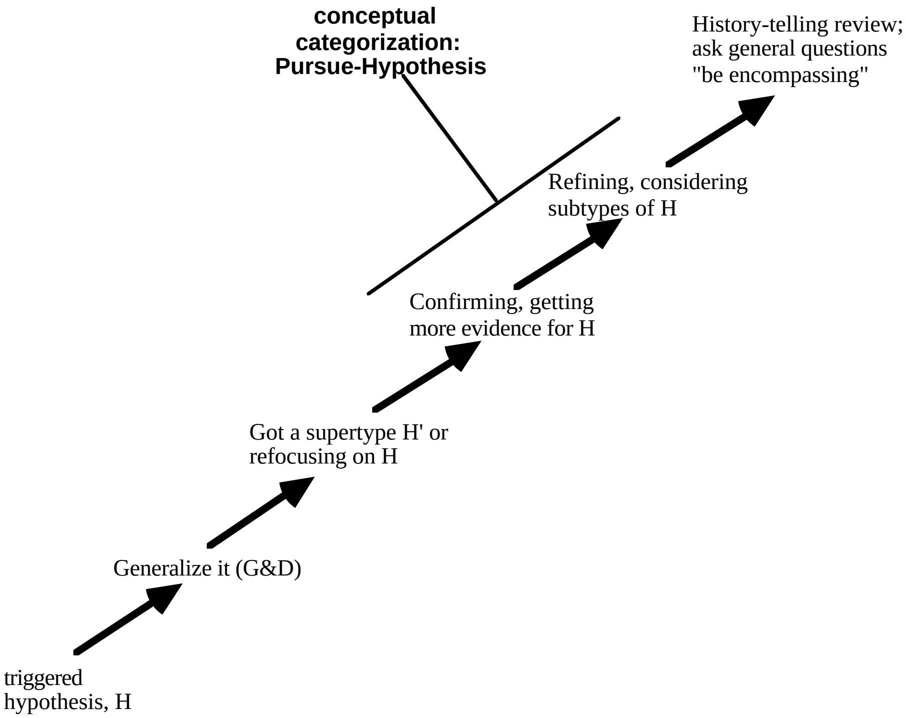

Figure 9. Habitual path in diagnosis, with a subsequence categorized as a new operator

To be specific, Figure 9 shows a plausible sequence of diagnostic activities corresponding to Neomycin's subtasks. Subsequences become segmented (chunked) and then categorized because "steps" always occur together and *they have a common kind of focus detail* (here the category "hypothesis"). These segments are now "conceptual coordinations"; activation is possible in a top-down manner by the inclusion of the higher-order category (e.g., "pursue-hypothesis") in another recurrent sequence. Goal-directed control then becomes possible if such a conceptualization can be held, biasing the activation of further sequences. A collection of diagrams like Figure 9 has the same embedded structure as a conventional hierarchical notation. But instead of merely describing the hierarchical relations, the ICP notation argues that the inherent formative process is *categorization over time*; and these nodes are always-adapted neural maps, not stored texts and *pointers* to locations in memory.

To summarize:

- *Chunks* are habitual sequences. The effect of practice is to establish reentrant links between perceptual details themselves, so attention to conceptualization drops out. In other words, a sequence of bindings is established from higher-order maps down to perceptual details and motor actions. (Note: Visual chunks, such as the arrangement of pieces on the chessboard, are non-verbal and non-sequential; they are *simultaneous categorizations*.)
- *Goals* are categorizations of sequences (aka conceptualized coordinations) with a common kind of focus.
- *A subgoal hierarchy* is composed of categorizations of sequences (conceptualizations). The ordering dependency represented by the ICP notation is *temporal*. The hierarchical order of Figure 9 represents temporal dependency of activation, not a simple *spatial* property of physical inclusion of neural maps.

Relating the innate and the learned: Evidence from linguistics

The ICP perspective reifies two dimensions of organizers: On the one hand, relationships conceptually constructed in the past are organizing new utterances—previous sequences are being reactivated, generalized, and are operating on different levels of abstraction. On the other hand, the kinds of conceptual constructions that are possible *de novo* are determined by the process of perceptual categorization, sequencing, reactivating sequences, categorizing sequences (conceptualizing), and constructively composing sequences on different levels. This distinction has been primary in the linguistic theory of transformational grammar and is much misunderstood.

Berwick's (1983) summary of modern linguistic theory recapitulates the argument against mediating descriptions: 1) parse trees are not explicitly built, 2) the units of information are associated with individual words, and 3) surface structure is the product of many constraints acting in parallel. The idea of abstract rules describing classes of words and how they interact is replaced by a mechanism by which how words interact to produce sequences is determined by ordering and "thematic" constraints. Hence syntax and semantics are effectively integrated. The phrase structure rules of early transformational theory are reinterpreted as being an observer's abstractions over a set of sentences, not part of the mechanism itself.

Berwick's description still refers to dictionary entries and encoding of features. On the other hand, connectionism does not satisfactorily explain how orderings are learned and how this relates to other kinds of sequential learning.[1] The ICP notation provides a perspective that supports and improves upon Berwick's thesis. For example, Figure 10 shows Berwick's representation for a phrase structure rule in which a verb is indicated as having an inanimate object, as in the full sentence, "Sincerity frightens John." "Nanim" (corresponding to the word "sincerity") is a feature that "percolates to a higher node to carry out a compatibility check." (Berwick, p. 393).

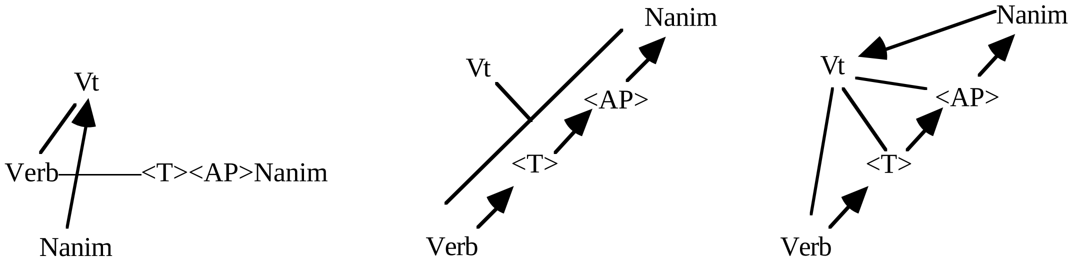

Figure 10. Representations of a phrase structure rule (left) shown by two forms of ICP notation, with a Verb Phrase (Vt) represented as a categorized sequence.

The middle of Figure 10 shows the ICP notation for this sequence, in which Vt is a conceptualization of a kind of verb phrase. The relation of the ICP notation to Edelman's TNGS (see Figures 4 and 5) suggests that each of the nodes have a reentrant link to the categorization of the sequence (Vt).[2] Furthermore, as shown on the right side of Figure 1,

---

[1] (Clancey, in preparation) shows how the ICP notation could be implemented as a generalization of Simple Recursive Networks (Elman, 1989; Servan-Schreiber, et al, 1991).

[2] Reentrant links would allow a conceptual categorization to "get back" to the head of a sequence (Figure 8, node A1) and begin to activate subsequent categorizations (A2, A3). As the sequence is constructed, the overarching conceptual categorization is generalized and strengthened. Possibly *the conceptual categorization is the first node in the sequence*, generalized to include the subsequent global maps, including them both temporally and by reentrant links. In this way, the sequence is subsumed within a global map, but is the global map itself, which "operates" by activating included maps as a process over time. The unit is a coordinated sequence over time. It is unclear how the end is marked; "closure" might

the nodes at the head and the tail may have a special relation within the categorization, allowing reconstructing the sequence from the end or activating the head via feed-forward activation. A conceptualization—a categorized sequence—unlike a felt path, may be reconstructed from the end because the subsuming categorization provides an activation link back to the head. Here we have Berwick's representation, with no added feature tagged onto the Vt categorization—it is inherent in the sequence itself that the "kind of verb" is qualified by the object, Nanim, the last category in the sequence.[1]

The ICP representation in Figure 10 suggests that there may be a special relation between the categorization of the sequence and the head of the sequence. Indeed, Chomsky's X-bar theory of 1970 claims:

> A Verb Phrase *must* be expanded with a verbal "Head"—that is what it means to be a Verb Phrase. Similarly, a Noun Phrase is always expanded as ...N..... In short, every phrase of type "X" is grounded on a word of type X. Phrases, then, are merely projections of lexical items. ... There is no specific phrase structure rule needed to describe the required constraints... (Berwick, p. 396)

But the relation between the phrase categorization and the root is even more general—the ordering is the same for all kinds of phrases!

> Chomsky observed that not only does every phrase of type X have an X-based root, but also that the order of arguments of each phrase is fixed across all phrase types. For instance, in English, Verb Phrases are headed by Verbs, followed by optional Noun Phrases, Prepositional Phrases, and Sentential Phrases. Prepositional Phrases are headed by Prepositions, and are also followed by phrases of the same type. Similarly for Adjectival Phrases . That is, once we have fixed the phrasal order XP->X NP PP, it is fixed for all phrasal types. (Berwick, p. 396)

This linguistic data suggests that a general sequencing "coordinator" is learned and applied systematically whenever speaking a given language. In ICP terms, a conceptualization which orders sequences within a phrase applies to all kinds of phrases. Only one kind of sequence categorization is applying throughout. The constraints referred to by Berwick are the activations in sequence and the formation of new phrase conceptualizations in accord with previous phrase conceptualizations. This suggests that learning a new kind of phrase conceptualization is generalizing an existing one; thus constructing the categorization XP of the sequence X->NP->PP (where here "->" denotes the ICP sequence link). Or perhaps better, to allow for the fact that kinds of phrases are not necessarily learned separately, learning phrase ordering in a given language involves conceptualizing *what ordering means in that language*, which is one categorization, not many. To repeat, the constraints of transformational grammar are revealed in the ICP notation as being processes of sequential activation and recategorization of sequences, not stored dictionary entries as in a descriptive cognitive model. These observations fit Berwick's theoretical claims in relating modern transformational grammar to connectionism, and further support the ICP (*process construction*) perspective.

The "inability to move Noun Phrases" over multiple phrases reveals from the other direction a limitation in our ability to coordinate multiple conceptualizations. That is, the

---

be conceived simply by the absence of further maps in sequence. But some sequencing and timing mechanism is required to prevent looping or stuttering. All of this is beyond Edelman's speculations, and I cannot go further here.

[1] This interpretation fits Small's theory, upon which Berwick draws, that a thematic argument such as a kind of NP is a "concept structure," which is a "grouping of words." (Berwick, p. 410) The difference is that we replacing "encodings in a dictionary" by categorizations of process sequences.

activation and sequencing process imposes a constraint that a subsumed phrase sequence can be moved "into the next higher sentence, but no farther.." (Berwick, p. 405). So we can understand "Who do I think that Mary likes" but not "Who do I believe the claim that Mary likes?" Again, in phrase structure terms this "locality principle" is merely structural, based on *proximity*. The ICP notation suggests that a subsumed process may be *incorporated again* in the next sentence categorization, but not after that—a constraint of *activation* concerning physical construction of processes. But this appears to merely restate (and support) the claim that each conceptualization is reusing or reactivating its previous construction. Again, there is no need for buffers or pointers: Constructing the next sentence is a process of *reactivating the previously active processes of the Sentence conceptualization*—hence a "who phrase" previously uttered is now available to be categorized within the new sequences "Mary likes ..." "We do not actually 'move' NPs around." (Berwick, p. 406). In short, Chomsky's locality principle should be considered as a constraint of process reconstruction, rather than serial order of the surface structure which is its manifestation.

When linguistic observations are reformulated by the ICP notation we find the kind of mechanism that Chomsky has argued for—a physiological process that structures conceptual content. The flexibility of "working from the head" or "working from the tail" suggests again that physiologically, conceptualizations of phrase utterances are *circuits*, activating and available to be worked with as single processes. Even though the head and tail are manifested in behavior at different points in time, they are incorporated in one categorization "couple" (Figure 4). Further discussion of TNGS and consciousness (Clancey, in preparation) considers how the human brain allows us to hold active and coordinate multiple, competing sequences as "a soup of strings."

Conclusions

My interest throughout has been architectural, that is, to describe the process mechanisms by which behavior is sequenced and composed. Following Bartlett's lead, I began by assuming that learning and memory should be viewed as the dynamic construction of coordinated sequences from *previous coordinations*. That is, *physical structures* are reactivated and reused in accord with how they categorically and temporally worked together in the past. With further examples from neuropsychology (Clancey, in preparation), I have generalized the notion of coordination to multiple sensory systems and, especially in humans, multiple modes of conceiving. Fitting the original hypothesis that knowledge is not a body of text, this perspective recognizes most of all that verbalizations are not isomorphic to conceptualizations, and that scenes, rhythms, and melodies are other ways of organizing behavior. Hence, "gestalt" understanding, which appeared to our text-based mentality to be a pre-scientific concept like "intuition," can now be understood as a *non-descriptive form of conceptualization*.

At the first order, we now make no distinction between "memory," "learning" and "coordinating." The control regime of symbolic models, such as Neomycin's strategy rules, is replaced by temporal-sequential conceptualizing: reactivating, generalizing, and recoordinating behavior *sequences* (whether occurring as outward movement or private imagination). In this respect, "knowledge" is constructed as processes "in-line" and integrated with sensorimotor circuits. Cognition is *situated* because perception, conception, and action are physically coupled. By this reformulation, cognition is not only physically and temporally situated ("knowledge is in the environment"), but *conceptually* situated as a tacit understanding of "what I am doing now." In people, this conceptualization is pervasively social in content; and hence we see that the proponents of situated action are making psychological claims about both the learning mechanism and the content of knowledge.

In accepting the broad patterns the schema theories of cognitive psychology describe, we are able to go beyond other models as well, such as Lashley's serial organization and Edelman's perceptual categorization. Nevertheless, for the purpose of psychological theory (though perhaps not pedagogical design) we reject the idea of a "body" of knowledge, a store that is indexed, compared, instantiated, and applied. Putting such descriptions aside, we ask What *inherently developmental* process or architecture could Neomycin's strategy rules and MOPS scripts be describing?

This exploration presents a new research program and the methods of analysis point to new synthetic approaches. The Interactive Coordination Processes (ICP) notation is a means for inventing new learning mechanisms. The rules of a descriptive cognitive model are more convenient for *classifying* behavior, but the ICP notation is better for understanding development of sequences. By eliminating the mediating role of descriptions from the mechanism, we can show intuitively how "proceduralization" is not a secondary compilation step, but a constructive process that began in evolution as sensorimotor coordination and became the process we call *conceptualizing,* a means of holding active, ordering, and composing sequences. The ICP notation, although valuable for coarsely diagramming the phenomenology of cognition, imperfectly describes multimodal coordinations and must be assumed to still misconstrue the spatial and temporal properties of the neural mechanism.

The result is a sketch of an architecture for "process memory," based on notions of coupling, activation sequencing, composing, figure-ground shifting—the neurological origins of the transformations and organizing principles studied by Lashley, Piaget, Chomsky, and Bateson. A more complete analysis (Clancey, in preparation) relates the verbal and visual to temporal-sequential, manipulo-spatial, and scene conceptualization. In analyses of dysfunctions and dreaming, we discover further evidence for specialized brain "modules" for conceptual organizing—hypothesizing that these are always mutually configured, neither parallel or serially-related in conventional terms, and have no existence as structures or capabilities apart from some integrated coordination in time with other processes.

References

Arbib, M. (1980). Visuomotor Coordination: From Neural nets to Schema theory. *Cognition and Brain Theory* 4(1) 23-39.

Bamberger, J. and Schön, D.A. 1991. Learning as reflective conversation with materials. In F. Steier (Editor), *Research and Reflectivity.* London: Sage Publications.

Bartlett, F. C. (1932/1977). *Remembering--A Study in Experimental and Social Psychology.* Cambridge: Cambridge University Press. Reprint.

Berwick, R.C. (1983). Transformational grammar and artificial intelligence: A contemporary view. *Cognition and Brain Theory,* 6(4) 383-416.

Bickhard, M. H. & Terveen, L. (1995). *The Impasse of Artificial Intelligence and Cognitive Science.* Elsevier Publishers.

Clancey, W. J. (1993). Situated action: A neuropsychological interpretation (Response to Vera and Simon). *Cognitive Science, 17*(1), 87-107.

Clancey, W.J. (in preparation) *Situated Cognition: On Human Knowledge and Computer Representations.* Cambridge University Press.

Edelman, G. (1992). *Bright Air, Brilliant Fire: On the Matter of the Mind.* New York: Basic Books.

Elman, J. L. (1989). Structured representation and connectionist models. *Proceedings of the Annual Conference of the Cognitive Science Society,* Lawrence Erlbaum Publishers, pp. 17-23.

Freeman, W. J. (1991). The Physiology of Perception. *Scientific American* (February), 78-85.

Gardner, H. (1985). *Frames of Mind: The Theory of Multiple Intelligences*. New York: Basic Books.

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), 33-55.

Newell, A. (1982). The Knowledge Level. *Artificial Intelligence, 18*(1), 87-127.

Polk, T. A., & Newell, A. (In press). Deduction as verbal reasoning. *Psychological Review*.

Rosenfield, I. (1992). *The Strange, Familiar, and Forgotten*. New York: Vintage Books.

Sacks, O. (1987). *The Man Who Mistook His Wife for a Hat*. New York: Harper & Row.

Schön DA. Generative metaphor: A perspective on problem-setting in social policy. In: Ortony A ed. *Metaphor and Thought*, Cambridge: Cambridge University Press, 1979;254-83.

Servan-Schreiber, D., Cleeremans, A., McClelland, J.L. (1991). Graded State Machines: The representation of temporal contingencies in simple recurrent networks. *Machine Learning* 7(2/3), Sept ‘91, 161-193.

Steels, L., & Brooks, R. (Eds.). (1995). *The "Artificial Life" Route to "Artificial Intelligence": Building Situated Embodied Agents*. Hillsdale, NJ: Lawrence Erlbaum Associates.