Chapter 1

Attention and Cognition

Attention

Of all the tasks the human brain performs, perhaps none is more consequential for the performance of other tasks than attention. When human beings attend, they perceive. When human beings attend and perceive, they remember. When human beings attend, perceive, and remember, they learn. When human beings learn, they can act deliberately and with forethought.[1] When performing a task, one must, conversely, reduce the need for constant attention to some of its specific components, allowing those components to be carried out automatically, yet the very act of pushing these elements into the background of consciousness occurs only because one must attend to something else. In short, perceiving, thinking, learning, deciding, and acting require that we "budget" our attention.

So, what is attention? A search for a concise definition need go no farther than this famous quotation from William James: attention is "the taking possession by the mind, in a clear and vivid form, of one out of what seem several simultaneously present objects or trains of thought" (1910: 403-404). Current thinking among many cognitive scientists and neuroscientists is that this "taking possession by the mind" is not a single entity or mechanism, but rather the name given to a finite cadre of modal-specific brain processes that interact mutually with other brain processes involved in the performance of perceptual, motor, and conceptual tasks.

This "taking possession by the mind", however, is not currently considered the purview of a single entity or mechanism, at least not for neuroscientists such as John Groeger (2000) Michael Posner (1995), Raja Parasuraman (1998), or Kimron Shapiro (2001). Rather, attention falls under the purview of a finite cadre

[1] My use of the rhetorical *gradatio* to describe attention was inspired by Raja Parasuraman's own usage in his introduction to *The Attentive Brain* (1998: 3).

of modal-specific brain processes that interact mutually with other brain processes involved in the performance of perceptual, motor, and conceptual tasks. Groeger (2000) provides a general neurological description of attention as diffuse patterns of activation among distinct cortices during the performance of actual and imagined sensori-motor tasks like driving a car. According to Groeger, a motivation to attend and act begins in the cingulate cortex of the limbic system and passes to the prefrontal cortex to inaugurate planning before marshaling the resources of premotor, motor-sensori-cortices and sensori-motor cortices. As the difficulty or novelty of the task increases, so does activity in the prefrontal cortex. Attention narrows as cognitive load increases.

Attention is so important to human cognition because it places limits on what we think about at the same time that it helps determine what our thoughts, words, beliefs, and deeds are "about" at any given time. Concerning the issue of limits, consider a specific study involving visual attention.

Dees & Frith (1999: 84) observed that unattended items could still be perceived but only under conditions of “low-cognitive load”. Thus, fMRI (functional Magnetic Resonance Imaging) results show significant blood flow to auditory cortices even as the subject is attending to an item in the visual field. The researchers concluded that the peripheral item is perceptually available but unattended. Under conditions of “high cognitive load”, however, blood flow to cortices associated with the unattended or distracting item (i.e., auditory stimulus) was practically undetectable. This suggests that subjects do not even perceive the moving item. Thus, what we attend and how we attend is task specific, and the complexity of the primary task will influence the likelihood of distraction.

Concerning the issue of intentionality (to be addressed extensively in the next chapter), consider another study conducted by Driver and Spence (1999: 130). When subjects anticipated a target item in one modality (e.g., vision), their judgments about other items from other modalities (e.g., sound or touch) would improve dramatically if and only if these diverse items appeared in the same location. Driver and Spence's findings suggest to me (but perhaps not to them) that human beings pre-attentively create multi-

modal spatial representations that integrate perceptual items. When an item becomes an "intentional object", other items in the close proximity to the perceptual or conceptual "space" take on significance in relation to the intentional stance exhibited toward the initial item. For instance, suppose your intention is to find a smooth, marble statue in a particular room. It may be that you will remember, say, what music was playing in the background, not because you were paying close attention to the music but because it happened to be playing as you were looking for the statue. A purely coincidental relation can now take on new significance for you. While attention limits the specific nature of a task (physical or mental) at one moment, it, quite conversely, expands the resources made available to memory at subsequent moments of recall and reflection.

This and other studies underscore the point that all mental events require commitment of limited cognitive resources at key moments, for we just cannot "take it all in at one time". Therefore, we attend by spending the resources necessary to prime the relevant cognitive processes, such as memory. We all must operate within a limited attentional budget.

The economics of attentional budgeting dictate that human beings can do multiple tasks simultaneously only if the requisite cognitive load is within their budget, but that performance will suffer (or break down entirely) if cognitive load exceeds the attentional budget. This explains why individuals cannot understand two verbal messages at the same time: the verbal stimulus entering an unattended channel will not be understood, because the resources needed to process it are already fully engaged. This despite the fact that the individual will remain aware that the unattended message is a message. It also explains why individuals can use the attended items to remember unattended items if those items do not interfere with modal-specific processes.

In defining the limits of cognition, attention initializes the elaborate chains of inferences we use to construct “readings” of the world. As asserted previously, attention also adds a directional component to behavior, modulating responses to the environment by focusing the mind on specific objects, locations,

persons while suppressing or attenuating surrounding irrelevancies, particularly when the cognitive load is greatest. It is a family of effects promoting the processing of one set of items over another.

The Theory

The attentional system I wish to describe has three principal components: *explicit attention, subsidiary awareness* and *shared attention.*[2] Taken completely, the system possesses scalar dimensions. Items within the field of attention can occupy a place on a scale from *inactive* to *active* to *salient,* with inactive items occupying subsidiary awareness and active and salient items occupying conscious attention.[3] Salient items readily play determining roles in thought and action, for they are immediately accessible with little or no effort; active items also play a conscious role in thought and action but require slightly more effort to bring them into the foreground; and inactive items play a preconscious role in thought and action, constituting the background from which one can extract salient items. Bringing inactive items into full conscious attention requires greater effort, greater cognitive load, and greater resources from long-term memory. An item can become salient and active by two routes: directly through the perception of external prompts (often referred to as *material anchors*) or indirectly through *spreading activation.* If an item impinges directly on my visual, auditory, tactile, olfactory, or gustatory systems, I then place it momentarily in focal attention for further processing. An item can become salient as a byproduct of direct focus should further processing effort connect the focal concept to a closely connected concept. Direct perception of a bumble bee may activate through spreading activation the concept *hive* or it may direct attention to other buzzing creatures, such as flies, and mosquitoes.

The first two components of attention consist of six (six or five ?)secondary components, each specifying functional characteristics. The two sub-components of awareness are *alertness* and *orienting,*

[2] I am using this term in the manner described by Polanyi (1962: 58).

[3] These variable dimensions of attention form the cornerstone of Anderson's (1983) theory of attention.

and the three components of attention are *selection*, *sustain*, and *control*.(here there are five secondary components)

A third component of attention needs mention before proceeding further: *shared attention*. The components of individual attention can only be fully appreciated in terms of shared attention. Complex human behaviors and abilities never occur in a vacuum; in fact, they will not even get off the ground without shared attention. One fundamental condition of the human infant seems to be that she comes into the world expertly prepared to coopt her caretakers' attentional budget. She spends nearly all of her precious mental and bodily resources attending to the caretaker as the caretaker, in turn, attends to her. Together, they engage in shared attention. Extrapolating from this "primordial scene," one can also say that although attentional budgeting is a process of individual brains, the very motives behind attention tend to be intersubjective. I save a fuller treatment of shared attention for chapter 3.

Exposition of the components of awareness and attention will proceed through a description of my journey from my campus (Guilford House) office across to the university library (Kelvin Smith) to borrow a copy of James Boswell's famous biography, *The Life of Samuel Johnson, L.L.D.*

The processes of subsidiary awareness are *alerting* and *orienting*.[4] Altering refers to the process by which I maintain my general readiness to process novel stimuli, while orienting refers to my disposition to select particular kinds of input over others.

As I walk to Kelvin Smith library, my body and brain remain sensitive to happenings in the immediate environment. Walking through an environment of other pedestrians, drivers, cyclists, roller-bladers and skateboarders (not to mention construction workers operating heavy machinery) places my visual, auditory, and motor receptors on highest alert, thus "clamping" attention to the campus environment.[5] In contrast, sitting at my desk editing a draft of this chapter places my visual receptors for recognizing print and handwriting on highest alert, thus clamping attention to the immediate textual

[4] This account follows from Posner's (1988, 1992) exposition of visual awareness and attention.

[5] This term has gained currency among cognitive scientists to refer to the fixation of attention, or to the biases of an

environment. As I walk along Ireland Way (the name given to the avenue running from my office building to the library and student center), my senses are predisposed to attend to items within the immediate campus environment. By virtue of my location, I am predisposed to be on the lookout for diverse forms of external stimulation. By virtue of my bodily orientation relative to the environment, I am predisposed to focus attention on the renovation project at Severance Hall taking place some sixty meters ahead of me.

As I sit in my office and edit this text, I am less predisposed to attend to Severance Hall, even though I can see it from my office window. By virtue of my location, I am predisposed to be on the lookout for a much narrower band of sensory stimulation. By virtue of my bodily orientation, I am predisposed to focus attention on the paper laying on my desk, pushing to the background other objects not immediately related to the task set before me.

In either scenario, my present location (sidewalk or office) and bodily disposition (moving or sitting still) play a determining role in what, how, and why I attend. Alerting and orientation are two components of awareness that always factor in the higher-level attention, a description of which I now turn to.

Selection (or *setting*) of attention, the component of attention described in James' oft-quoted definition, is perhaps the most widely researched of all the components. Selectivity of processing is required because, as previously stated, human cognition must operate within a finite budget. The primate brain evolved mechanisms for data extraction based on selective attention for the purpose of coping with information processing, and information processing is really a matter of making readings of present, past, future, or imagined happenings in the world. As I walk to the library, I am presented with multiple and simultaneous items of which I can process only a small portion; therefore, I must select one and discard others.[6] Selection of attention depends on bottom-up processing to the extent that it is at the mercy of the

---

AI program for some data over another. I am using it the same way Glenberg (1997) uses it.

[6] My use of the term "discard" may be problematic, for it may imply conscious deliberation, when in fact discarding irrelevant stimuli may be automatic. I use the term simply to illustrate that some kind of filtering process is going

spatial origin and intensity of immediate perception. During my trek to the library, I am confronted by many items vying for attention, such as the sound of heavy machinery at a nearby construction site, the sight of a colleague walking toward me, the feel of the warm, humid air against my skin, the smell of hot pizza from the student center, the residual taste of balsamic vinegar lingering in my mouth from lunch, and the dull ache of my sore knee, and so forth. The advantage of selective attention is that I can concentrate on one kind of activity without interference from other activities by selectively attending to those stimuli or messages that reinforce my primary task. The disadvantage of selective attention is that I cannot accomplish many tasks simultaneously. While it is certainly possible for me to appreciate the aroma of pizza and walk to the library simultaneously, it is not within my attentional budget to look for two different books simultaneously, even if these books appear on the same shelf. (I have to search for them sequentially.) In summary, selective attention facilitates mental processing of one task while inhibiting the completion of others tasks. It accounts for the fact that I must choose which book to look for first. Without selective attention cognitively modern human beings would be ill-equipped to act coherently in the face of competing and distracting sensory information.

While selective attention supports the choice of goal-directed tasks of all sorts, *sustained* (or *focused*)attention ensures a task's completion by taking up the lion's share of the attentional budget. The need for focused attention defines a component of attention distinct from selection in that it involves concentration. While selective attention is subject to the contingencies of bottom-up perception, sustained attention depends on top-down framing of a situation or scene. Memory of a task and knowledge of how to accomplish it (i.e., going and getting a book from the library) will influence how I interpret all that I see, hear, smell, touch, and taste. If my desire to borrow a copy of Boswell's biography uses up most of my attentional budget, all my bodily movements and interactions with the campus environment are likely to be influenced by this desire. For instance, suppose the colleague walking toward me is holding a book in his hand. Memory of my goal might very well make that book stand-out from a background of

---

on, but probably one working at temporal intervals too quick to invade conscious thought.

competing information. In this case, concentrating on a task will influence selective attention or awareness of objects encountered along the way.

While selection and sustained attention can function as mutually reinforcing processes, they can also oppose one another, most notably in rich sensory environments where the alerting and orienting mechanisms are prone to respond to any sensory cue from above and below, front and back, and to the left and to the right. In this respect, it is perhaps best to think of selection and sustain as opposing processes that ensure attentional balance. That is, a high rate of stimulus presentation induces iterations of selective attention, thence decreasing sustained attention. A sudden sound of an explosion will force me to reckon a different attentional budget to deal with a possible threat. The stimulus and its aftermath may be so intense or consequential as to supplant my original plan and becomes the new focus of attention.

Cognitive psychologists specify two mental activities associated with sustained attention[7] that bear mention here: *vigilance* and *search*.

Individuals engage in vigilance tasks when they detect signals presented to them only infrequently over a long time span (psychologists usually study vigilance tasks that last more than an hour) in unpredictable intervals but in a predictable location. A parade example of vigilance task would be a postal worker whose job is to look for incoming mail with the zip code 44122 and place it in a special bin marked “local mail”. With such tasks, the number of letters and the instances are unpredictable, but what is predictable is that every target letter will be found in a specific location in the mail room. Another example would be driving a car on a freeway. While driving you remain vigilant to the task of taking a certain exit. You know it is somewhere on this stretch of highway, but you do not know exactly where, so you cannot calculate exactly when you are to turn off. Because you know that a situation will arise requiring you to turn off, you remain in a state of alertness, even as you must switch to other immediate tasks, such as breaking, shifting gears, passing slower drivers, talking with your

[7] See Davies et. al. 1984 and Parasuraman 1986 for detailed studies of search and vigilance tasks; see Matlin 1987 for an accessible overview of this research.

passenger, and so on. Let us move inside Kelvin Smith Library, so I can describe a relevant instance.

Librarians assigned to work the circulation desk must wait on patrons who wish to borrow or return books. The number and intervals of each request is not, strictly speaking, predictable, but the location of those requests is highly predictable. In fact, librarians may sit for hours at a time without waiting on a single patron. This kind of task is location-based[8]: Where an event is likely to occur has a great influence on how the attentional budget will be spent. When a person walks between the two ropes that create a "service cue," the librarian assumes that he or she wants to make a specific request; thus the librarian attends to the movements, words, and gestures of that person as if she were a patron. Top down processing with respect to location influences, if not determines, how one will respond to bottom-up perception.

Individuals engage in search tasks when they detect intermittent stimuli; whereas uncertainty persists with respect to *when* and *what* signal will be detected with vigilant tasks, uncertainty persists with respect to *where* a signal will be detected with respect to search tasks. Consider as an example how I use the Library of Congress numbering system when searching for *The Life of Samuel Johnson, L.L.D*, catalogued as item **PR3533.B5**. Once I proceed to the shelf corresponding to the call letters **PR** and number **3533**, my search task becomes more efficient. Once in that very region, I no longer focus attention on the call letters and first numbers, but on the letter and number combination appearing to the right of the decimal point. In this case, I am looking for the letter **B** and number **5**, the indicators to the right of the decimal point corresponding to an edition of Boswell's biography. I do not know exactly where the book will be (even though I have narrowed the possibilities down significantly), so I single-mindedly ignore the **PR3533** and search exclusively for **B5** to the right of the first decimal point. In sum, these two search tasks underscore the point that sustained attention depends on figure and ground relations, with some tasks emphasizing location (ground) and other tasks emphasizing objects or object-

[8] In many respects, the process of orienting determines vigilance tasks; the mere fact that a person is standing or sitting in a carefully circumscribed location frames the scene and constrains greatly the kinds of interpretations one is likely to construct of events and actions occurring therein.

token relations (figure). Importantly, these tasks are not mutually exclusive; the librarian's continued vigilance, although emphasizing location, certainly shares characteristics of search and vice versa.

Sustaining information-processing over time in the face of distraction is one means of maintaining goal-directed behavior. The activity may need to be stopped (in order to respond to some other contingency) and then be resumed; there may be other concurrent activities and their future fulfillment must be coordinated with meeting the primary task. The punctuated nature of goal-directed behavior coupled with the ability to coordinate several strands of information at one time, keeping them in their proper order, is known as *control* of attention. Attentional control is particularly critical in theories of working memory and planning. This is the least understood component of attention, but it is probably the most crucial. Most of the time, I can accomplish my goal of going to the library and checking out a book and seamlessly integrate perceptual, conceptual, and motor tasks with other competing, distracting activities, such as admiring at a distance architectural features of Severance Hall. I do so by budgeting my attention. In other words, I can be distracted by other signs of life incompatible with my present tasks, but not so distracted that it derails the completion of that task.

I chose the quotidian event of going to the library to borrow Boswell's biography because an accurate narrative of this incident would reveal a break down in the *control* of attention. This shift of attention, however, was nothing so dramatic as an explosion, but a result of my own absent mindedness. As I walked along Ireland Way I directed my line of sight to the architectural detail of the newly renovated Severance Hall adjacent to Kelvin Smith Library. I suspect that I was aware of Severance Hall as a landmark for mentally positioning the location of the library relative to my body. As I approached the library, however, the edifice of Severance Hall flooded my senses. Captivated by the details of new rear facade, I walked past the library doors. Only when I stood twenty-five feet away from Severance Hall did I realize I had missed my target. Initially selection of Severance Hall was a useful relative landmark for accomplishing my initial intended goal; subsequently it flooded selective attention so that I could no

longer remain focused on my quest to borrow a book from Kelvin-Smith Library. The unwarranted shift of attention occurred because I switched the figure ground relationship of these two building.

In its most ecumenical sense, attention is simply the global activity that initiates and maintains goal-directed behavior in the face of multiple, competing distractions.

The general account of attention I submit provides a proper grounding for modeling how human beings construct meanings about anything and everything relevant for individual and collective existence. Therefore, the processes alerting, orienting, selecting, sustaining, controlling, and sharing attention may count as the basic cognitive scaffolding of complete semiotic theory. Of course, all models of meaningfulness are in some respects reductive—which is what makes them models—because it is not possible to capture on paper every nuance of even the simplest task. My claim is that a satisfactory model needs to start with the components of attention and work outward to study acts of meaning making that occur nonverbally, verbally, and communally. The theory I develop shares many features with Peircean semiotics, cognitive linguistics, and rhetorical theory, each the subject of the next three chapters.

Before proceeding any farther, two tasks need our attention: to explore the relation between attention and a trio of related cognitive processes, and to define some basic concepts and terms on which this study depends.

Related cognitive processes

The ability to share, select, sustain, and control attention is the fundamental component of semiotic processes, both in guiding the physical, mental, social, and communicative world of the individual and in guiding individuals' attempts to influence the social, mental, and communicative worlds of others. The individual and social processes of attending determine, and in turn, are determined by, three interrelated cognitive processes: *memory*, *categorization*, and *valuation*.

Memory

Rarely has the question "what is memory for?" been explicitly asked in the cognitive sciences. The unfortunate consequence of not asking this question is that ever more sophisticated probes into the nature of human memory proceed as if memory exists for the purpose of "memorizing". Approaching what is clearly one of the most fundamental of cognitive processes from this tautologous starting point has not been particularly useful for building a theory of meaning. For the purposes of this study, I will embrace Glenberg's tentative answer to this fundamental question: human memory "evolved in service of perception and action in a three-dimensional world to facilitate interaction with the environment" (1997:1). I favor this basic approach, because it makes memory a service of interpreting the environment, and gives us a reasonable starting point for abstract thought. According to Glenberg, the purpose of memory is

> . . .to mesh the embodied conceptualization of projectable properties [i.e., sensations & perceptions] of the environment (e.g., a path or a cup) with embodied experiences that provide nonprojectable properties [i.e., conceptualizations & abstractions]. Thus the path becomes the path home and the cup becomes my cup. This meshed conceptualization, the meaning, is in the service of control of action in a three-dimensional environment. (1997:4)

This view of memory means that we match past instances of interaction, social or individual, with present or future interactions as similar or different for the purposes of adjusting anticipations, expectations, and behavior. The memory of a place or situation guides behavior because memory *is* mental simulations of places and situations. Our memory of something is entrenched to the extent that simulation holds; our memories change to the extent that new projectable properties of the environment require us to alter our mental maps.

As a process, memory is the capacity to reconstruct important scenes, and the reason one wants to

reconstruct complex scenes is, at base, “to repeat a performance" (Edelman 1992:102). Hence, completing my simple task of borrowing a book from Kelvin Smith Library entails a prodigious reconstruction of many related scenes. This study assumes the capacity to construct and repeat a performance, understood as a complex cognitive routine in which certain elements are afforded a great deal of attention in working memory, while other elements are severely attenuated from working memory. Together, these capacities constitute our ability to construct elaborate scenes.

Another way to define memory is to regard it as an *achievement*. Human beings especially marshal and reconstruct complex scenes in order to integrate novel combinations of items. Remembering the location of the library and the elaborate code of behavior associated with university libraries serves as material for reconstructing in a scene that guides my present and future behavior. If attention can be thought of as the selection, sustain, and control of information processing, then memory can be thought of as the means by which we activate diverse items with the aim of integrating them into more-or-less coherent scenes.

Before discussing specific components of memory, I would like to distinguish different memory tasks. The most fundamental and pervasive memory task is *recognition*. Recognition of something as something (see categorization) is largely an automatic process (perceptual categorization, for instance, usually occurs in 200 millisecond intervals), but it is the most pervasive. Recognizing that the colleague walking toward me is carrying a book is an example of recognition, both of the person and of the artifact he is carrying. Selective attention depends on recognition; navigating the environment depends on recognition. Thus, any theory of online meaning construction has to understand the importance and influence of recognition tasks in virtually all areas of cognition. Recognition tasks are distinct, however, from tasks of *recollection*. It is when we engage in recollection tasks that most people consider remembering. If recognition is an automatic process, recollection is a deliberate process, often understood as a means of combating forgetfulness. Suppose that I walked to the library and forgot what book I

wanted to borrow. The task of reconstructing my past intentions and desires is preeminently a task of recollection. Glenberg (1997) notes that one peculiar feature of recollection tasks is the tendency on the part of the individual to close her eyes and hold her head either up or down when in the act of recollecting something, apparently to stave off the intrusion of external, distracting signals. Attention and memory, then, are intimately connected with recognition aiding an individual's fluid and safe navigation through a potentially dangerous three-dimensional world, and recollection aiding an individual's effortful attempts to stave off the intrusion of the three-dimensional world in a single-minded effort to reconstruct an idea, object, or scene not present in the here-and-now.

Now that we have a theory of what memory is for, we can discuss the components or kinds of memory in terms of their differing functional capacities. The following classification follows closely that of Barsalou (1993). When psychologists and cognitive scientists speak of memory, they often do so by subdividing it as *sensory*, *short-term*, *working*, and *long-term*. For the purposes of this study, I will distinguish between working and long-term memory, with the proviso that short-term memory is simply an aspect of working memory with a shelf life of approximately 20 seconds.

Working memory consists of a set of mechanisms that work together to perform strategic processing. So we will define working memory as a combination of procedural and declarative information processing. Thus, when I see the window on the facade of Severance Hall, recognition of it permits the recruitment of long-term memory "stores" associated with my interactions with it, such as seeing out of it, opening it, or cleaning the individual panes. I think that working memory and long-term memory are part and parcel of the same system, and are always working in dialogue with each other.

Any theory of temporary memory must account for the strong effects of long-term memory on working memory and for the similarities between them.

Following Barsalou (1993), I adopt a unitary view of working memory. I assume that long-term memory can contain many types of information, including executive productions, auditory cues, visual

images, motoric patterns, and so forth. At any given point in time, this information can be salient, laying in the temporary memory state. Working memory then is information from long-term memory that is currently in the temporary state of activation. Thus, the memory system is not a set of distinct "warehouses" (in fact it is not a store at all), but a unified system with different activation states. Thus, the vocabulary of explicit attention—selection, sustain, and control—may be understood as the perceptual and conceptual mechanisms that relate the immediate present stimuli, the sensory store, with working memory. In summary, memory connects the here-and-now with the there-and-then.

Selective attention picks out an element from the environment, or an element from long-term memory (by means of some sign) and relates it to information in working memory. Working memory involves integrating diverse strands of information and knowledge into coherent scenes for guiding thought and action, the products of which can be conventional and mundane or novel and fantastic or somewhere in between. Sustain of attention takes a particular set of information from the sensory store and working memory and links it to conceptual and cultural schemas (see section 0.3.2) in the long-term memory store. Long term-memory, then, comes into play whenever we need to sustain attention.

Long-term memory, in turn, influences working memory and even perceptual categorization. An individual can clamp onto a search image, what we expect to find, and not clamp onto the immediate environment so that we do not register subtle changes to the environment. Uexküll (1957 [1934]: 62) speaks of the "search image" (by which I assume he means memory of an object) annihilating the "perceptual image" (by which I assume he means an external signal of something), and does so by recounting an instance from his own social life. Every time he went to a friend's house for dinner, a large, clay pitcher of water would be in the middle of the dining room table. He would always reach for the pitcher and pour himself a glass of water. One time he entered the dining room but could not find the pitcher. The hostess pointed to the pitcher standing in the same spot in the middle of the table. Uexküll initially could not see it; then he saw it. It was no longer clay, but glass. His map of the room included an

opaque pitcher of water and that element so strongly determined his mental map of the room that he could not register the change from one artifact to another. He was looking for an opaque item because memory told him that was what to search for. Memory can, in turn, annihilate immediate perceptual categorization (but only for so long).

Control of attention often entails shifting to different elements in working memory, or attending to an entirely new set of elements in working memory as they match the sensory store.

Categorization

To live a life, one must continuously engage in perceptual and conceptual categorization. Interpreting the world involves putting objects, events, beings, and ideas into categories, taking them out of other categories, and transferring emotions and attitudes onto them as a result of either placing them in or taking them out of a preregistered category. It should be clear by now that attention, categorization, and memory are co-dependent processes. In fact, categorization is memorization to the extent that perceptual categorization is a form of recognition and conceptual categorization often needs a great deal of time and effort.

Categories, however, should not be considered a priori, fixed, or context-free. They are, instead, the result of dynamic, changing, and context-dependent cognitive routines. One of the great achievements of cognitive psychology, supported by work in anthropology, is a near complete revision of what it means to construct categories (Rosch 1978). People do not categorize based on necessary and sufficient conditions; instead, human categories reflect what Wittgenstein (1953) called *family resemblances*, a phenomenon in which category members can be related to one another even if its members share none of the properties that would define them classically.

Categories also have degrees of membership, or what is usually known in the literature as *prototype effects*, some being easier to recall in context neutral situations. Thus, most westerners would

easily categorize the pair of *windsor chairs* in my office as good examples of the category *chair*, whereas a *bar stool* would be a less good example, but an example nevertheless. Most categories have "fuzzy" boundaries. In specific contexts, a large rock with a flat surface can be categorized functionally as a chair because it can function that way. In fact, a speaker can even felicitously refer to it as "my chair". These ad hoc categories are not an exotic extra process separate from other categorizing processes, but are likely the core process itself, thus studying instances of ad hoc categorization may be the key to understanding the nature of human categorization generally. In addition to family resemblances and prototype effects, category members can be organized around a *basic level*, a level at which individuals can determine membership faster, can remember the actions associated with the category at greater frequency, and, generally, can name them most easily. Thus, *chair* is a basic level category, whereas *furniture* or *chaise* are superordinate and subordinate categories. According to Lakoff, the basic level categories typically represent the level at which "people function most efficiently and successfully at dealing with discontinuities in the environment" (1987: 269). With lower level categories (subordinate) people usually find distinguishing between members more difficult, whereas with higher level categories (superordinate) people usually find it difficult to generating concrete imagery associated with the category. With respect to linguistic structure, basic level categories tend to be morphologically simple. With respect to language acquisition, basic level categories tend to be among the first items learned by children (Lakoff 1987: 46). Linguistically, the basic level is represented by names for basic colors, qualities, plants, animals, substances, objects and actions, tall, short, hard soft, rose, lily, tree, dog, cat, horse, running, walking, jumping, and eating. Basic level categories, however, are not confined to the real world or physical experience (although interpreting the immediate physical environment is a fundamental use of categorization that makes all other uses possible), since cultural and fantasy categories like mother, father, brother, sister, ghost, unicorn, priest, gang, and so forth are basic as well. They seem to be indispensable for negotiating an immediate, culturally defined environment.

The concept *book* is a basic level category because: 1) members of this category have similar shapes—square or oblong; 2) a single mental image can reflect the entire category—an image of a paperback; 3) a person uses similar motor programs for interacting with category members—I can hold it, carry it, open and close it; 4) subjects can identify it faster than superordinate and subordinate counterparts—I will call it a book before I call it a "publication" or "folio"; 5) members commonly use this label when referring to subordinate members instead of using the subordinate level label—a diary counts as a book; 6) children understand and learn this category before its superordinate and subordinate members—my oldest son used the term "book" to refer to dictionaries, magazines, newspapers, and pamphlets; 7) the label for this category is morphologically and phonologically 'shorter' than its higher and lower level counterparts—notice that "book" is monosyllabic, whereas "publication," and "monograph" are polysyllabic; and 8) the category can be evoked in context neutral situations—the term "monograph" is specific to academic situations, whereas "book" is not (Lakoff 1987:133; Rosch et. al.1976:382-439).

These newer perspective on categories and categorization (see Lakoff 1987) do not replace completely classical categorization, however. Nevertheless, their robust influence on cognition in general does indicate that how we typically construct, maintain, and alter categories does not follow the classical model of necessary and sufficient constraints. Classical categorization is a real mental phenomenon, but it is also an artificial and normative one, in great need of elaborate systems of knowledge (e.g., law, disciplines, taxonomy) to sustain it.

My remarks on the nature of categories should suggest that I regard categorization as fundamentally a non-linguistic process—even though much of our conscious categorizing proceeds by linguistic symbolization. Categorization plays a defining role in the mental life of all living organisms, because, at base, all organisms need to construct mental maps of their immediate environment and niche, and part of that process involves categorizing things *as kinds of things* (however rudimentary) so as to

create a set of anticipatory responses to contingencies in the world.[9]

Categorization is crucial to conscious mental life because it links our present environment with past instances; it is a tool for constructing mental maps for dealing with contingencies in the world. I categorize my trek to the library as a kind of activity so that I know what to expect and how to behave. To the extent that the Kelvin Smith Library fits the prototype category for libraries, the more likely I am to apply a predetermined set of anticipations, expectations and actions I associate with libraries when I enter it. On the other hand, the book shelves in my office also count as a library, but a private one. Since these shelves and books do not fit the prototype very well, I am less likely to apply the same set of anticipations, expectations, and behaviors to it. [10] I need to abstract it as a kind of activity related to other activities in which the present instantiation fits. Doing so helps me retain it in memory and, I would argue, provides me with the means of focusing attention on the present task in the face of multiple competing distractions. Without categories, all appearances would be novel, new, and important, nothing would be differentiated, and the budget for our mental resources would be constantly depleted.

Value

In order for human beings to attend to something, it must have value. Attentional dispositions use up

[9] Organisms vary greatly in the detail and flexibility of these mental maps. For instance, a common wasp can be said to categorize a dead cricket as *food*. Once applied to a particular object in the world, the category activates a response routine that involves taking it back to the nest, putting it down by the entrance, inspecting the entrance for suitability, picking the food up again and depositing down the tunnel entrance. The slightest disturbance in the routine—for instance, an experimenter moving the piece of food a millimeter from where wasp originally laid it when inspecting the entrance—and the wasp will reposition the piece of food and reinspect the tunnel. It will do so *ad infinitum*, never skipping what you and I would consider to be pointless extra placements and inspections. Although the prey does not lose its categorical designation, it nevertheless instantiates a categorical designation of cognitive routine that cannot be adjusted. Nonhuman primates, on the other hand, develop much richer and flexible cognitive maps with categories that are not so rigidly applied to routines. That is, they categorize certain fruits, plants, and vegetables as *food* and have general maps of where to find them, but will maintain that category if found in an unusual location. For instance orangutans will categorize a piece of fruit as *food* regardless of location (see Tamasello & Call 1997: 8-9).

[10] I do not need to be quiet when looking for a book on my shelf, as I would if I were in Kelvin Smith Library; I do not ?borrow it,' since I can take it out of my office and return it as I please. I can drink coffee and eat lunch while browsing through my shelves, something I am not permitted to do in the library; I usually write and underline passages in my own books, but doing so would not be considered unacceptable if done to a library book.

precious energy. Accordingly, whatever occupies an individual's attention must be worthwhile. Survival depends on paying attention to the right things. Thriving in an environment and culture usually entails an implicit adherence to certain value hierarchies, the most basic and probably universal of which would be the value claims *it is better to live than to die*, and *it is better to see than not see* (even though one can conceive of extraordinary contexts in which *it is preferable to die* or *it is preferable to be blind*). Culture specific value hierarchies might include *it is better to be rich than poor*, or *it is better to be literate than illiterate*. Valuation is a fundamental component of human existence. We are constantly engaging in value claims and acting accordingly. If attention is, as Reisberg (1997: 122) suggests, an *achievement* in which we decide to select certain kinds of information and ignore others, then these achievements are acts of valuation, of deciding (even implicitly) moment-by-moment what is important and what is not, or at least less so. The question of what we value individually and collectively varies (probably less so than many would like to think) depending on cultural and sub-cultural constraints.

What a whole culture values is reflected most clearly in the basic categories comprising its language. As Barsalou (1993: 271-272) recounts in detail, English-speaking cultures have elaborate color terms, such as *red*, *blue*, and *green*, while many languages of non-Western cultures, such as Dugum Dani, have an "impoverished" color terminology (only encoding terms for *black* and *white*). Conversely, languages of plant gathering cultures have extensive terminologies for flora, whereas industrialized peoples do not. If linguistic categories develop to serve everyday thought, then these two cultures develop categories that reflect what is most important to them. In industrialized cultures, Barsalou tells us, it is more important for individuals to extract colors from objects so they can isolate colors in paints and dyes, or match and coordinate colors. In plant gathering cultures, it is more important for individuals to recognize immediately different species of plants so they can either cultivate them or avoid them. This is not to say that individuals in these two cultures perceive color and plant life differently (each are capable of acquiring or even creating subordinate categories for colors and plants). Rather, this is to say that by

virtue of acquiring a language, an individual acquires certain valuations reflected in the structure of language and what is designed to do. Basic-level categories for color in industrial countries would count as an extensive subordinate set of categories in Dugum Dani. The basic-level categories for plants in plant gathering cultures would count as an extensive subordinate set of botanical categories in English and other Indo-European languages.

The concept of valuation has had little treatment in the cognitive sciences. It has, however, been treated extensively by rhetorical theorists (see chapter 3). Two Belgian rhetoricians, whose theory of argumentation I will describe later, stress the importance of valuation with respect to inducing agreement. Chaïm Perelman and Lucie Olbrechts-Tyteca write that

> agreement with regard to value means an admission that the object, a being or an ideal, must have a specific influence on action and a disposition toward action and that one can make use of their influence in an argument . . . (1969:74)

Their remarks capture an important component of meaning construction: the capacity and predilection to judge the relative importance of an object, event, or idea.

This capacity to judge is reflected in our experience of basic emotions, moods, and temperament. Theorists of emotion (e.g., Ekman, Oatley, and Tooby & Cosmides) argue that basic emotions, such as fear, anger, pity, anxiety, disgust, sadness, and happiness, are adaptive responses to recurrent and fundamental situations, like fighting, escaping predators and capturing prey, and falling in love. Universal among humans and primates, basic emotions are observable in others by facial expressions (the preponderance of evidence suggests expressions for basic emotions are cross culturally invariant). Emotions are distinct from other psychological phenomena in that their quick onset and limited duration offer very fast appraisals of current events, which has led theorists like Ekman (1994), Oatley (1992), and Tooby & Cosmides (1990) to conclude that basic emotions are innate capacities reflecting our ancestral past. Regardless of their evolutionary status, basic emotions are valuable to online meaning

construction because they offer quick appraisals of situations, leading individuals to act as if certain things are true of an unfolding event on the basis of what was true about a past event. The strength of being able to attach an emotional valence to a situation is that it offers a quick evaluation in a world where quick interpretation and response can mean the difference between life and death. The weakness of being able to attach an emotional valence to a situation is it can be wrong, because such responses tend to be automatic (occurring before conscious attention can catch up) and incapable of being sensitive to fine grained variations among similar situations. The control of emotions is especially valuable in complex societies where most of our days are spent in situations sheltered from the kind of basic, life and death situations that characterized the existence of our ancestors.

Although deeply entrenched and automatic, emotions are intimately connected to higher order concepts. All events and situations acquire emotional valences; they have to in order to properly alert and orient the individual. But not all events acquire the same emotional valences or exhibit the same degree of intensity. My walk to the library reflects a very diffuse or neutral emotional intensity, since it so frequently makes up quotidian reality. Encountering a colleague I have not seen in several weeks will confer a memorable, probably pleasurable, emotional color to this event. In most respects, representations of going and getting a library book are not fraught with intense negative emotions; however, my felt need to borrow a book may be a response to a situation based on emotions, such as anxiousness about my research and writing.

Since languages encode emotional terms, one aspect of human emotion is that it can be represented deliberately, i.e. selected. Emotions can, therefore, be manipulated discursively.

Human affect is greatly influence by mood. If we experience emotions as automatic in their onset and brief in duration. (Ekman argues that enduring emotions are really a series of briefer emotional episodes.) While emotions are accompanied by distinctive facial expressions, moods are not. But the real difference between emotion and mood, as proposed by Davidson (1994), is that emotions relate to

situations where quick reaction is necessary, thus modulating or biasing *action*. Mood, on the other hand, functions in situations that call for considerably more deliberation, and, therefore, function to modulate or bias *thought*. As such, mood will bias the kinds of interpretations we construct and, concomitantly, what ways we represent situations to others. For instance, if I were in a sad mood, I would have increased access to sad memories and decreased accessability to happy memories. This suggests that mood affects spreading activation, or what conceptualizations become salient at any particular time. In addition, Davidson reports that positive moods facilitate cognitive flexibility, whereas negative moods inhibit cognitive flexibility. Mood directly influences subsidiary awareness (alerting and orienting), which in turn determines explicit attention during acts of sign production and consumption. If I am extremely sad as I walk to Kelvin Smith Library, everything I see and hear is more likely to elicit dysphoric interpretations, and vice versa if I am extremely happy.

Everyone is capable of feeling intense emotions, and everyone has so-called mood swings. But everyone differs in the distribution and duration of emotions and moods. Davidson uses the term "affective style" for describing variations among individuals in how they react to events. Temperament, writes Davidson, "refers to early consistent differences that are assumed to be at least in part under genetic control. . ." leading to "systematic biases in emotional reactivity" (1994:54). Differences in affective style appear early in development and may have a determining effect on learning. Individuals who grow up in stressful environments are more likely to develop affective styles that are pathological in nature, i.e, subject to various mood disorders. The challenge for semiotics is to account for both the invariant and variant factors influencing online meaning construction. While emotions are relatively invariant, moods and temperaments are highly variant. Therefore, it may be the case that one of the key sources of variable interpretations is affective style or temperament. The question remains open, however, as to whether to treat temperament as a purely individual trait or whether to treat it demographically.

Attention, memory, categorization, and value comprise the quartet of perceptual and conceptual

processes that make us cognitively modern human beings.

Basic concepts

Before the separate discussions of semiotics, linguistics, and rhetoric begin, I offer extended remarks on specific terms common in cognitive and human sciences. The purpose of these definitions is twofold: to provide a theoretical orientation for the readers before proceeding to subsequent chapters; and to offer a reference point for readers to consult as they read on, in case they are not sure how I am using a particular term.

The tour of definitions begins with a few words about two of the most amorphous terms in the cognitive and human sciences, cognition and meaning, then proceeding to define more specific and tractable concepts, such as schema, role, imagery, mental space, and so on, before ending my tour with some thoughts on the relationship between cognition and culture.

Cognition

According to the *Oxford English Dictionary*, cognition refers to

> the action or faculty of knowing taken in its widest sense, including sensation, perception, conception, etc., as distinguished from feeling and volition; also, more specifically, the action of cognizing an object in perception proper.

The editors also define cognition as an achievement, a product of such actions. This study follows broadly this catholic definition of cognition and emphasizes its sensory-perceptual underpinnings as sketched out by Barsalou (1999). Cognition as an act or faculty is commonly thought to be *internal to* and *an achievement of* individuals. I do not quarrel with this definition; however, I wish only to stipulate that a theory of cognition cannot be achieved without a theory of culture as well, and, as I shall argue shortly, a

model of purely individuated cognition will not always (in fact rarely) yield a satisfactory account of intelligent human behavior. That said, when I use the term cognition, I am committing to the notion of an individual, intentional subject who experiences body sensations, assembles perceptions and thoughts idiosyncratic to her alone. I am also committed to the notion that individuals engage in extended reasoning about a topic, to make inferences about aspects of that topic, to employ *modus pollens* and *modus tollens* logical strategies, to set up pragmatic scales such that an established set of inferences and values can apply to a new, contested set of inferences and values of immediate concern. Additionally, I agree that language in particular and discourse in general provide the necessary scaffolding for extended thinking and reasoning; however, it would be a mistake to simply equate cognition with language, or linguistic intent, or the meaning of words, because cognition and meaning making depend equally on mental imagery, schemas, and domains.

Meaning

As advertised in the introduction, this study attempts to assemble a general theory of meaning; therefore, it behooves me to say a few words about the meaning of meaning, lest readers think that I use the term as philosophers of language have traditionally used it either to describe properties of words, or the property of words in sentences that are true or false, or the linguistic encoding of a speaker's intention. Paradoxically, my definition of meaning simultaneously comprehends all of these senses and none of them: it is all of these insofar as the term can be felicitously used to stipulate the senses of words, of asserting truth or falsity, and of displaying a speaker's intentional stance; it is none of these insofar as my definition of meaning emphasizes the fact that meaning is the ephemeral product of an activity not an enduring thing, that sense of words are continually being constructed and reconstructed and not fixed “in” the signs themselves, that truth or falsity are rarely an overriding concern for human semiosis, that a speaker's intentions are not necessarily “encoded,” and perhaps most important, that meaning does not fall under the proprietary control of language proper, but rather is a function of semiosis and communication

in general. Therefore, the base elements of meaning are not words and sentences, per se, but information.

The informational nature of meaning can be further defined as a piecemeal process with results that are obtained only within particular semiotic events. In their study of written communication, Kaufer and Carley offer a definition of meaning that is particularly apt for developing a general semiotic. For them, meaning is the resultant product of local recognition and interpretation, deriving "from the relationships between . . . discrete pieces of information, built on pieces that are known and how they are interlinked" (1993:106-107). This piecemeal view of meaning suggests further that any given datum can have multiple meanings. What is more, the same individual may assign a different meaning to the same datum in later moments, depending on the precise interpretive task required of her and on the set of presuppositions available to her. In the most ecumenical sense, meaning is the result of conceptualization, or the process by which individuals construct specific ways of thinking about something.

From an analytic perspective, something becomes meaningful when interpreted with respect to an activated schema within a conceptual domain that, when made manifest in language, is further interpreted with respect to particular semantic frames and semantic roles.

Schemas, frames, roles, and domains

Schemas. Perhaps no concept has enjoyed more widespread currency among anthropologists, cognitive scientists, and linguists than the notion of "schema". Although experts differ greatly in precise definitions of a schema and how it differs from a "script," a "scene," or a "scenario" (see Mandler 1984 and Rumelhart 1980), a definition that I believe would elicit widespread agreement and one that best captures my use of the term is that offered by Kemmer and Barlow. For them, a schema is ". . . a cognitive representation comprising a generalization over perceived similarities among instances of usage"(2000: xviii). By repeatedly "activating" a set of co-occurring properties in a particular way, individuals develop "top-down" conceptualizations of different kinds of experiences, with each repeated instance becoming "an organized framework of objects and relations which have yet to be filled in with

concrete detail" (D'Andrade 1995: 122). Entering Kelvin Smith Library activates my schema for LIBRARY that includes slots for such roles as "librarian," "patron," "student," "faculty member," any of which can be filled with specific values. When I imagine myself standing in line at the circulation desk, I am instantiating that schema with myself as a filler for the role "patron" or "faculty patron". One can legitimately say that the event I have just imagined is a sub-schema of the larger library schema; in fact, one can legitimately call it a "script": it has a clearly delineated series of actions that include searching for the book in the catalog, going to the precise stack where the book is to be found, pulling it from the shelf, carrying it to the circulation desk, handing it over with my library card to the circulation librarian, who processes it and hands it back to me before I exit the building. In order to avoid an explosion of terms, I will simply use the term schema throughout.

Frame. When interpreting the words and phrases of others, we create what Fillmore (1975) calls a text model, a transitory model constructed from schemas for the purpose of interpretation. The key component of this text model is the notion of a frame, which he defines as ". . . the specific lexico-grammatical provision in a given language for naming and describing the categories and relations found in schemas" (1975: 127). Fillmore uses this term when discussing the lexical and grammatical means of eliciting associated schemas. Lexical and grammatical items are always interpreted in relation to a system of categories whose structure issues from a specific motivating context. For example, the verb *climb* evokes a frame involving animate, up or down motion. Thus, the speaker strongly implies gravitational resistance from the immediate environment when using this verb. The verb phrases *climbs up the stairs* and *climbs into the bathtub* implicitly code gravitational resistance on the part of the stairs and the bathtub, whereas the verb phrases *ascend the stairs* and *collapse into the tub* do not imply such gravitational resistance. I will use the term frame when discussing the contrasting of implications of different (but synonymous) linguistic items. Otherwise I will use the term schema.

Semantic roles. At a lower level of linguistic organization is the semantic role, as descriptive

systems developed by many schools of linguistics used to specify the grammatical relations between predicates (events or states) and arguments (entities). There are two main classes of semantic roles: participant and non-participant (cf. Frawley 1992: 197-250). The class of participant roles can further be divided into actor roles and spatial roles. Actor roles include *agent*, *effector*, *instrument*, *patient*, *experiencer* and *benefactor*. An agent is an intentional, volitional, potent instigator of an event; an effector is an executor of a predicated event but is not a direct cause of the event; an instrument refers to the means by which an event is carried out; a patient specifies the argument directly affected by the event; an experiencer represents an argument whose internal state or constitution has been affected by the event; a benefactor specifies an argument that derives something from the actions of another. Spatial roles include *mover*, *source*, and *goal*, where a mover specifies an argument that is displaced along a trajectory, a source specifies an argument designating the origin of a displaced entity, and where goal specifies an argument designating the destination of a displaced entity.

For example, the sentence *I grabbed the book from the shelf and gave it to the librarian* can be analyzed thus: *I* is the agent, *book* is the mover, *shelf* is the source, and *librarian* is a patient (who can also be interpreted as goal in some contexts). The sentence, *The librarian suggested I look for the book on the third floor* casts the librarian in the role of the effector, since she would be only indirectly responsible for my going to the third floor. The sentence, *The student uses his physics textbook as a door stop*, casts the direct object, *physics textbook*, in the role of the instrument. The sentence, *I could not help admiring the Severance Hall's new facade* cast the pronoun, *I*, in the role of the experiencer, since the subject's internal state is altered by the building's presence. And the sentence, *I borrowed a book for my son* cast the object of the preposition, *son*, in the role of the benefactor.

A non-participant role includes *location*. This role specifies the argument for a fixed site of predication. The sentence, *The book is in the third floor stacks under call letter* ***P***, casts the object of the preposition, *Third-Floor Stacks*, in the role of the location.

Domain. A domain refers to broad areas of conceptualization. Langacker (1999: ch.1, ch.6) uses the term basic domain to refer to instances entailing conceptualizations of **three-dimensional space, visual images, sounds, smells**, and **haptic sensations**. Complex domains, such as conversation and kinship, still index broad areas of human knowledge and experience but are secondary in that they depend on basic domains for coherence. In any case, domains are superordinate to schemas and frames, which are blueprints guiding thought and action in specific situations. We can further analyze schemas and frames according to the primary and secondary domains that comprehend them: the WRITING schema, for instance, encompasses the conceptual domains for **human body, seeing, touching, movement**, and **conversation.** In this study, the term "domain" refers to basic and complex areas of conceptualization without respect to specific information regarding participant roles and perspective.

Imagery and image schemas

Imagery. Much of cognition is forming mental images, and the term *imagery* emphasizes the perceptual origins of concepts; that is, concepts (even abstract concepts) develop from representations of sensory experience—a conglomeration of visual, auditory, haptic, motoric, olfactory, and gustatory sensations. My definition of imagery follows that of Palmer (1996: 47), who writes:

> Images are mental representations that begin as conceptual analogs of immediate, perceptual experience from the peripheral sensory organs. Because they are analogs of peripheral experience, they are also, therefore, indirect conceptual analogs of the environment, broadly construed to include society, natural phenomena, our own bodies and their organic (mental) processes, and the rest of what is called "reality" or "the world out there".

Once registered in the mind, the immediate perceptual experiences thus defined can be abstracted and replicated so that we can make sense of our environment, reason about contingencies therein, and act .

While immediate perceptions form the basis of mental imagery, the images themselves are abstractions in which the individual can fill in details. They become schemas, or templates for framing new experiences. Before tackling these next terms, it should be noted that mental imagery can be (and, I would argue, always are at some level) activated by peripheral sense organs, or they can be activated autonomously from long-term memory, evoking the corresponding imagery in absence of the relevant sensory stimulus (cf. Langacker 1987: 112). This is especially important with respect to language. What language can do is evoke mental images *in absentia*.

Cognition depends on attention to these mental images which are partially activated representations from memory achieved when the attentional mechanisms group eliciting stimuli into organized networks. These organized networks depend on schemas and domains.

Image Schemas. Living and dwelling in this world depends on acquiring patterns for arranging information, what has been called schemas. The number of schemas needed to engage in purposeful behavior is staggering. Every piece of information must be placed in a schema or schemas to be interpreted. Without a schema, a phenomenon cannot be categorized and interpreted.

A class of schemas that has proven indispensable for understanding human thought and reason is the image schema. Initially developed by Lakoff and Johnson (1980) and elaborated into the areas of philosophy by Johnson (1987), into human categorization by Lakoff (1987), into poetic metaphor by Lakoff and Turner (1989), into literary criticism by Turner (1987, 1991), and into formal linguistics by Langacker (1987, 1991) and Talmy (1985), image schemas are thought to make possible the mind's ability to map spatial structure onto conceptual structure. An image-schema is a condensed redescription of perceptual experience. When fully developed in a conceptual system, an image schema operates as “a dynamic pattern that functions somewhat like the abstract structure of an image, and thereby connects up a vast range of different experiences that manifest the same recurring structure” (Lakoff 1987: 113-114). (Missing subject) describes many of these schemas and their transformations that provide the ground for

cognitive models—be they propositional, metaphoric, metonymic, or perceptual analogues. The most common image-schemas used in everyday thought include: TRAJECTOR-LANDMARK; CONTAINMENT; PART-WHOLE; CENTER-PERIPHERY; SOURCE-PATH-GOAL; LINK; REFLEXIVE; FRONT-BACK; FORCE; COMPULSION; MOMENTUM; BARRIER and UP-DOWN.

These schemas often combine to give basic structure to both concrete and abstract concepts. The schemas we employ as we think involve relations among these image schemas; these relations are known as image-schema transformations (Gibbs & Colston 1995: 347-378; Johnson 1987: 25-27; Lakoff 1987:440-444; Palmer 1995:68-74; Turner 1991:177). Thus, conceptualizations of *walking over to the library* involves a path-focus to end-point-focus transformation where conceptualization follows a moving object along a path and then shifts focus to the point where it comes to rest (Johnson 1987:26). Conceptualizing the abstract notion of *going into debt* entails building a mental model of debt by metaphorically transforming an abstraction into a concrete location (again via the path-focus to end-point focus transformation where the end-point) transforms into a containment image. Conceptualizing the act of selecting a book from a large shelf of books can proceed by a *mass to multiplex* transformation. In this case, the shelf appears to the viewer from a distance as a single homogenous mass, but once the viewer "zooms in" the mass turns into a cluster of individual items. With respect to linguistic structure, image-schemas are activated by closed-class items, i.e., prepositions, particles, deictics, affixes, and so forth.

Image-schemas are conceptual primitives because they are *topological*. Along with Talmy (2000: 25-31) and D'Andrade (1995:133), I use the term topology as it is used in mathematics as "spaces" sectioned into areas without specifying actual magnitude, shape, or material. Thus, the schema for *over*, which involves a image schematic transformation of PATH and UP relative to a landmark, does not specify the magnitude of spatial gap between trajector (up) and landmark (path); the trajector can be construed as making contact with the landmark, as in *walking over to the library*, such that the magnitude of the gap is perceptually negligible, or the trajector can be construed as above the landmark, as in *the balloon flew*

over the library, such that the magnitude of the gap is perceptually salient.

To summarize, image schemas represent the regularity experiences, and although they are both meaningful and structured, they are not richly meaningful, which makes them highly useful as an apparatus capable of describing a wide array of semiotic phenomena. In short, image-schemas are said to lie at the core of an individual's understanding of objects, events, and ideas. Under this theory they are not primarily linguistic. They are pre-verbal, making them viable candidates as the proposed link between thought and language. I also take the position that the repertoire of image schemas is not innately (i.e., genetically) specified, but arises very early in ontogeny; image-schema transformations come online as individuals learn, elaborate, and relate more concrete schemas, frames, and models.

Cognitive models

To be human, to engage in purposeful behavior means to model or simulate events, actions, objects and relations in the world. Kenneth Craik introduced the notion of a cognitive model in his study *The Nature of Explanation.* He writes:

> If the organism carries a “small-scale model” of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilize the knowledge of past events in dealing with the present and future, and in every way to react in a much fuller, safer, and more competent manner to the emergencies which face it. (1943: 13)

It is now virtually a matter of strong agreement among cognitive scientists, cognitive linguists, semioticians, and anthropologists that human beings build imagistic mental models in order to think, talk, listen, and act. Mental models can be said to govern the production and consumption of signs, verbal and nonverbal. Given the centrality of this concept, a precise description of cognitive models is now required.

To begin, a cognitive model is the same for me as an idealized cognitive model (ICM) is for

Lakoff (1987:68-76), in that each model establishes and, in turn, is further developed into a world view. Cognitive models then are also cultural models.

As mental simulations of the "same relational structure as the phenomena they represent"(Johnson-Laird 1983: 11), a cognitive model is a multi-modal set of images organized in schemas. These models can also include propositional structures, image-schematic transformations, and metaphoric and metonymic relations. For example, the cognitive model I have for borrowing a book from Kelvin Smith Library includes a body moving through space (as do most models) through a specific trajectory that can include concrete walkways, grass, hallways and doors. It includes an embedded model of the library itself, which consists of an entrance, great foyer with a large spiral staircase that leads to the third floor stacks. It includes a series of moveable shelves. It includes an elevator which I can take down to the first floor where the circulation desk stands and is populated with one or two people, and a line of patrons waiting to be served. This entire set of representations (now cast in propositional form) depends on the event schema that includes a borrower, a lender, and an object. Once in place, the sight of a book on my shelf can activate the whole mental model for going to Kelvin Smith Library because of the metonymic link between book and library. I can create propositions based on this model, such as *The third floor houses moveable stacks*, which organizes in symbolic form mental images that are perceptually, spatially structurally, and transformationally equivalent to the library itself.[11]

To elaborate, my mental model of the library is perceptually equivalent to the extent that the imagery activates primary and secondary sensory cortices (visual, auditory, motor, etc.) in my brain as it

[11] Finke (1989) offers these four terms as principles of mental imagery. The principle of *perceptual equivalence* states that "imagery is functionally equivalent to perception to the extent that similar mechanisms in the visual system are activated when objects or events are imagined as when the same objects are actually perceived" (41). The principle of *spatial equivalence* states that "the spatial arrangement of the elements of a mental image correspond to the way objects or their parts are arranged on the actual physical surfaces or in the actual physical space" (61). The principle of *structural equivalence* states that "The structure of mental images corresponds to that of actual perceived objects, in the sense that the structure is coherent, well organized, and can be reorganized and reinterpreted" (120). The principle of *transformational equivalence* states that "Imagined transformations and physical transformations exhibit corresponding dynamic characteristics and are governed by the same laws of motion" (93).

would if I were perceiving the library at that moment. My mental model of the library is spatially equivalent to the extent that memory of the arrangement of floors, desks, chairs, books, paths, doors, etc., in mental space are similar to the actual arrangement of these elements in physical space. My model of the library is structurally equivalent to the extent that the mental objects are coherent, well-organized and can be reorganized and reinterpreted so that the mental space of the library can be updated as to fit changes (real or imagined) in the physical space of the library. My model of the library is transformationally equivalent to the extent that the mental space of the library is subject to the same laws of motion as the physical space of the library so that I mentally ascend stairs in the same way that I physically ascend stairs. (My model of the library does not typically include people floating from one floor to the next, even though I can imagine it happening.)

What makes a cognitive model different from a schema is its concreteness. It is an instantiation of a schema. In addition to the present cognitive model of the library, I have a cognitive model of several other libraries, both public and private, that overlap schematically in that they all have slots for "borrower" and "lender" for "book" and "publication"; however, they differ in several specific details. For instance, my model of the United States Library of Congress includes "borrowing" and "lending" but does not include images of me browsing the stacks or taking the book home with me. So, I have several mental models of libraries all of which cluster around a culture-specific schema for "libraries".

Integration, coherence, mental spaces & blends

Mental models, schemas, frames, and domains reflect human beings' attempts to produce conceptual integration, the products of which produce what can only be called coherence. When we try to make sense of anything—an event unfolding before our eyes, a person's facial expression, a text, a picture, and so on—we develop integrated interpretations that best match the available information. According to Thagard, a proponent of a multi-coherence theory of thought and action,

> the best interpretation is one that provides the most coherent account of what we want to understand, considering both pieces of information that fit with each other and pieces of information that do not fit with each other. (2000:16)

Unlike its products, the processes of coherence is not one thing; rather, coherence relations exist along six dimensions: perception, conception, analogy, deduction, explanation, and deliberation, each of which is heavily influenced by what Thagard calls "metacoherence emotions," such as contentment, anxiety, happiness, surprise, sadness, anger, fear, pity, empathy, disgust (2000:193). Each dimension deserves its own description.

Perceptual coherence depends on the balancing of what Ruthrof (2000: ch. 5, 6) calls *intersemiotic* and *heterosemiotic* readings of the sensory world.[12] A scene makes sense to use to the extent that our interpretations cohere with low-level representations from sensory input as interpreted through the gestalt principles of proximity, contiguity, and similarity (see Koffka 1934). Differing signs attain a level of rapport with one another. For instance, my ability to perceptually categorize an entity as a *book* can occur through visual semiosis alone; however, my acquiring and using as part of my mental tool kit the concept for *book* depends on a converging interaction and integration of different semiotic sources ranging from the visual and haptic sensations (in the sense that a book looks and feels a certain way) to the olfactory (in the sense that I associate old books and new books with distinct smells) to auditory sensations (such as the sound of pages being turned) to symbolic associations (such as with academic and the leisure classes). Incompatible sensations fail to cohere with one another and can produce heterosemiotic readings, often resulting from contrasting conceptual and emotional associations produced by the concurrent presence of contrasting signs. An extreme example would be a situation in which a holographic image of a book sitting on a table baits me to pick it up. My hand passes right through it. In

[12] For Ruthrof, higher order symbolic structures like language only make sense in terms of the non-verbal interpretations of real, imagined, past, present, and future worlds that affect us as we effect them.

this instance, there is an extreme heterosemiotic relation between visual sensation and haptic sensation—the visual mode leads me to expect structural integrity, weight, and gravitational pull not forthcoming in the haptic mode. Different signs can be said to conflict with one another. The concurrent presence of incompatible information puts pressure on the comprehender to partition representations. The heterosemiosis of perception and sensation pressures me to revise my overall mental model of present reality. The *prima facie* coherence of a scene depends on the integration of sensory-perception with background knowledge.

To illustrate, suppose you notice hanging on the wall a three-dimensional portrait of a human face made out of dishware mounted on a flat canvas, with a demitasse cup as a mouth, matching saucers as eyes, and spoon as nose. We make sense of it on the basis of perceptual similarity and contiguity relations; that the demitasse forms an aperture means it can serve as a rough approximation of a human mouth; that the cup is placed in relation to other items whose similarity functions approximate two eyes and a nose means it coheres with respect to the known part-whole relations of a perceived human face. However, if someone has in her mind the daily functions of using a coffee cup, then the vertical orientation of the cup on the canvas means that it conflicts with the background knowledge that the cup will hold coffee, since, *ceteris paribus*, it violates the coherence conditions under which we understand gravitational forces and containment. So, drinking coffee from this cup would be heterosemiotic to the present situation; the functional and aesthetic dimensions do not cohere, and may compete with each other.

Analogical coherence depends on mapping one set of coherence relations onto another. The example of the tea service portrait is a good example of analogical coherence based on perceptual coherence, where the formal properties of the dishware map onto contour features known about eyes, ears, noses, and mouths; the analogical relations cohere because each artifact from the coffee service can be positioned on a plane that maps onto the contiguity features of the human face. One can introduce an

element of analogical incoherence, for instance, by rearranging the cups, saucers, and spoons so that the contiguity mappings between dishware and facial features do not cohere. Selective modification--as would happen if I preserve every contiguity mapping but alter the contiguity mapping between spoon and nose--can produce a piece that still resembles a human face but with the nose on the left side of the face. My tea-service portrait may take on an additional analogical mapping in which the whole portrait resembles some of Picasso's notorious female cubist portraits.

Very briefly, a situation attains deductive coherence when all the propositions attributed to that situation are compatible, whereas a situation fails to achieve deductive coherence when propositions attributed to it contradict one another. A situation attains explanatory coherence when hypotheses and evidence correlate positively with one another, whereas a situation fails to achieve explanatory coherence when hypotheses contradict one another or when hypotheses fail to correlate with the known evidence. A situation achieves deliberative coherence when it matches our desired goals and outcomes, whereas a situation fails to achieve deliberative coherence when it contradicts our desired goals and outcomes. Carefully observing the architectural features of Severance Hall is incoherent with respect to my task of going to the library insofar as the aim of each task is incompatible. A simple matter of attentional control can make these two activities cohere simply by sequencing them: I examine the architectural features of Severance Hall and *then* I walk into the library to borrow the book.

Coherence theory is still quite vague with respect to modeling online meaning construction. Enter Fauconnier and Turner's theoretical framework known as conceptual blending or conceptual integration; it is a promising model of online meaning construction closely based on a coherence theory of understanding. As a theoretical framework for exploring human information integration, blending theory is largely consistent with Thagard's principal notion that human thought and action depends on rendering coherent mental representations across a variety of dimensions. Achieving coherence involves a set of operations for combining dynamic cognitive models in a network of *mental spaces* (Fauconnier, 1994), or

partitions of speakers' referential representations. Fauconnier & Turner (1998) suggest that a small set of partially compositional processes operate in the creative construction of meaning in analogy, metaphor, counterfactuals, concept combination, and even the comprehension of grammatical constructions. Blending processes depend centrally on projection mapping and dynamic simulation to develop emergent structure, and to promote novel conceptualizations, involving the generation of inferences, emotional reactions, and rhetorical force.

Mental spaces contain partial representations of entities and relations of any given scenario as perceived, imagined, remembered, or otherwise understood by a discourse participant. Elements represent each of the discourse entities, and simple frames represent the relationships that exist between them. Because the same scenario can be construed in multiple ways, mental spaces are frequently used to partition incoming information about elements in the referential representation. For example, the sentence, *When I was twelve, my parents took me to a concert at Severance Hall* prompts the addressee to construct at least two mental spaces, one for the present enunciation space, and one for the time when the speaker was twelve years old, the event space. The correspondence between the focal participant in the utterance and the event space is represented via an identity connector between the two. Partitioning the information into two mental spaces allows the reader to understand that while the speaker saw a concert at Severance Hall when he was twelve, he need not have seen one since. The virtue of mental spaces is that they allow the addressee to divide information at the referential level into concepts relevant to different aspects of the scenario.

The mental spaces and blending framework constitutes the primary semiotic model used in this study, and as such deserves an extended description in that (this?)chapter. The principal notion here is that mental space networks produce blended mental spaces for the purpose of allocating attention to specific tightly organized scenes for satisfying perceptual, conceptual, analogical, deductive, explanatory, deliberative, and emotional purposes. A conspicuous characteristic of many blends is that they tend to

achieve coherence by compressing relations of *identity*, *time*, *space*, *change*, *causation*, and *part-whole* into tightly organized scenes. According to Fauconnier & Turner (2000; In press) an optimally coherent blend tends to compress what is diffuse in order to achieve a human scale and produce a global insight. Blends also strengthen relations, generate stories, and shrink from many to one. For example, the sentence *Roosevelt dragged the United States out of the depression*, cited by Lakoff & Johnson (1999: xx) as an embodiment metaphor, can be analyzed as a blend generating a story that compresses diffuse information to achieve a human scale and strengthen vital relations by providing a very short, image-rich force dynamic story involving one man. In the blend, relations of identity fuse into one person (the many people in Roosevelt's cabinet, congress, the judiciary). The relation of time compresses from years to minutes; relations of space compress from diffuse regions of the country and world to an unspecified path for dragging; and relations of cause are compressed into one action for which we have rich imagery. The final product is a blend that forms a tight, perceptually and conceptually coherent scene, even though, *sui generis*, it is far from veridical.

Culture and cognition

The reader may note that I have been using the term culture throughout these pages, but have yet to provide a precise definition of this slippery term. For the purposes of this study, I will rely on Goodenough's famous definition of culture as "whatever it is one has to know or believe in order to operate in a manner acceptable to its members" (1957: xx). Although not very precise, this definition suggests, perhaps paradoxically, that culture is simultaneously an individual achievement and a collective set of constraints on the individual.

On the individual side, a cognitivist view has it that discernable patterns exist because the evolved conceptual system of individuals permits them to exist. Talmy (2000: vol. 2; ch. 7) proposes that individuals have a "cognitive culture system" designed to evaluate the patterns of behavior and affect

observed in others, and budgets attentional resources for the purpose of being instructed in such patterns, making them a part of one's identity. Talmy's main point is that enculturation is a highly structured process that begins at birth and continues unabated throughout the course of one's life. According to Talmy, a cognitive cultural system is designed to ensure individuals perform the following tasks: determine groups most relevant to the self; abstract across members of each group; selectively attend to phenomena manifested by its members; resolve conflicts among patterns of different groups. (2000: 374).

On the collective side, conceptual schemas, frames, and models should be regarded as issuing from a culture or subculture, and what defines a sub-culture is its differences in the degree to which individuals comprising it *internalize* certain schemas (an implication of Talmy's fourth task). For instance, I am a member of an academic subculture within a greater American-European culture. As such, I have internalized the library schema to such a degree that it provides a defining cultural representation of what I take to be true, correct, and right. That is, I believe in the importance of libraries generally and in the Kelvin Smith Library in particular. On the other hand, I have not internalized cultural representations of Jesus as my lord and savior. As a member of a culture that is, by most measures, religious, I repeatedly observe cultural patterns gathering around the figure of Jesus and am familiar with his theological importance, but I have not assented to its descriptive or normative claims. Since library-going is a part of my daily practice, it shapes my identity. Since I rarely if ever attend Christian services, theological Jesus does not shape my identity. As a consequence, I regard the existence and support of public libraries as true, correct, and right. In fact, I am likely to vehemently and emotionally defend the interests of my local library should it's funding decline. Contrastingly, I do not regard the existence of Jesus, the son of God as true, correct, and right; hence, I would vehemently oppose measures that would coerce worship.

Thinking of culture in the ways just suggested has consequences for how we think about cognition itself. Although I certainly believe that individuals do think and reason, I also think what we call cognition is not just a property of individuals. In many circumstances, cognition can refer to the

mental work that goes on among individuals in a specific setting. In some respects, it is desirable to speak of cognition as *distributed*. The emergence of distributed cognition as a legitimate theoretical and methodological approach to this study is significant and deserves extended quotation from one of its most distinguished proponents, Edwin Hutchins. Hutchins (1994) argues that

> . . . the cognitive properties of groups are produced by an interaction between structures internal to individuals and structures external to individuals. All human societies face cognitive tasks that are beyond the capabilities of any individual member. For example, even the simplest culture contains more information than could be learned by an individual in a lifetime, so the tasks of learning, remembering, and transmitting are inevitably distributed. The performance of cognitive tasks that exceed individual abilities is always shaped by a social organization of distributed cognition. (quoted in D'Andrade 1995: 210)

Since individual cognition is always grounded in situations and settings, and since those situations and settings are the products of cultures, and since one infallible sign that a culture exists is the presence of artifacts in specific setting, it is plausible that artifacts themselves have cognitive properties. So, let us define an artifact as a fabricated physical structure which helps a humans perform some cognitive operation. They range from complex computers to a pencil and paper.

The main stacks on the third floor of Kelvin Smith Library are movable, as to allow for maximum storage space within a limited space. Learning to work these stacks certainly alters my mental model for going to the library and getting a book. The movable stacks helps me inasmuch as all the books are located on one floor, eliminating the need to move between floors when looking for books tagged with call letters on opposite ends of the alphabet. They constrain me inasmuch as I am less inclined to browse the stacks, for fear of annoying a patron who is waiting to retrieve a book from an adjacent stack. The existence of movable stacks betokens a cultural representation, prominent among western, industrialized

countries, of “efficiency” and “progress”. The efficient use of space competes for and wins out over other concerns, such as the need for leisurely browsing.

Conclusion

Cognitively modern human beings create meaning only when a quartet of mental processes lead by many attentional processes operate as an ensemble. Attention leads memory, categorization, and valuation in an attempt to orchestrate locally coherent mental scenes regarding past, present, future, or imagined scenarios or situations. The purpose of this chapter was to lay out a theory of attention in six parts, relate these parts to these other processes, and offer descriptions and definitions of key terms likely to appear in the forthcoming chapters.