Chapter 4

What Matters

Getting Started

After my son James was born, we lived for a year in Kyoto, famous for its grand Buddhist temples. We lived near the Kyoto Zoo, but James had little interest in the animals. However, he loved the statues of the Buddhas, austere, serene gold faces dimly lit by votive candles. Set him down on the wooden walkway, and he would immediately crawl back to peer in. The elderly Japanese women who had come to pray were delighted by the manifest devotion of the little blond-haired foreign baby who surely had taken a wrong turn somewhere in transmigration.

What indeed did attract him to those statues? Perhaps some form of karmic connection: good deeds in his former life? Karma—not a serious answer—at least has the virtue of explaining the specificity of a highly individuated response: my son's inborn disposition simply reflects the transmitted *atman*, or soul. In a specifically Buddhist context, however, this answer cannot work so easily, for there is no *atman*, no essential self to pass from life to life.[1] To make matters more difficult, karma itself has no independent existence. This Buddhist quandary, created by categories emptied through relentless dialectical argument, in fact

[1] The concepts of *atman* and *karma* come from the Vedic culture of early India out of which Jainism, Buddhism and Hinduism all grew.

resembles our own as we seek a neuroscientifically reasonable account of my son's behavior. Easy postulates concerning innate preferences, predispositions, and the like that might have worked in a descriptive psychology are no longer adequate. Like karma and the Buddhist self, they are not independent entities. Preferences need a neural substrate. As in the Buddhist accounts, we will need to find our answers in the substructures of desire, perception, and intention. Moreover, the version of the self arising out of such an approach will be closer to the Buddhist than to the Cartesian model, a point to which I shall return at the end of the chapter.

Paying Attention

The primary perception systems in the brain organize themselves in response to patterns detected in the environment. Higher order systems then use these codings to attend more carefully to the environment, to construct and retrieve memories, and to guide responses. So far, so good. Central to these processes, however, are the related ideas of selective attention and biological salience. Deciding what matters is vital to making the system work well. In most cognitive research, this decision is not especially difficult. The chimpanzee wants the food, the human subject follows instructions, and the rats hate getting wet. Often the question will be *where* in the brain the decision happens and how that region then influences other systems. The particular decision itself, however, is simply given. For most purposes, this approach is entirely reasonable: experiments attempt to probe one question at a time, and why *this* rather than *that* is rarely the question.[2]

The logic and mechanisms behind preference is the domain of affective neuroscience. One small example: give a baby the chance to look at a bright blue ball for ten minutes. Take the ball away, then bring back both the blue ball and a red one: the baby reliably will look longer at the new red ball. (Reversing the colors doesn't change the pattern.) This preference for the new seems to be built into the baby's mental hardware.

This rather mild bias toward novelty does not seem to fit our usual understanding of emotion. Emotion is something big—joy, anger, fear, disgust—that sets the heart racing. Nonetheless, we are distracted from more subtle and important patterns if we look only at the big emotions. Emotion is both broader and deeper. As Joseph LeDoux argues,

> I think starting with universal behavioral functions is a better way of producing a list of basic emotions than the more standard ways—facial expressions, emotion words in different languages, or conscious introspection.[3]

---

[2] There are many experiments that explore priming effects, how exposure to one stimulus affects later choices. Subliminally showing the picture of an angry face before the picture of a neutral object, for example, makes the viewer biased against the object in a later preference test. The tester, however, usually simply assumes that the angry face is somehow a negative image without worrying too much about the details of what that "negative" is. The exceptions here are the experiments that explore the inability to judge angry faces that derives from damage to the amygdala, the subcortical structure that will loom large in this chapter.

[3] Joseph Le Doux, *The Emotional Brain: the Mysterious Undepinnings of Emotional Life* (New York: Simon and Schuster, 1996), p. 126

Emotions, that is, ensure the performance of “fundamental life tasks.” This approach to emotion as a mechanism to guide basic behavior and ensure survival reappears throughout neurobiological accounts of affect:[4]

> All mammals, indeed all organisms, come into the world with a variety of abilities that do not require previous learning, but which provide immediate opportunities for learning to occur. The influence of these systems varies as a function of the life span in each species. Analysis of the emotional systems that control behavior is complicated by the fact that the intrinsic arousability of underlying brain systems may change in many ways as organisms age. Still, the present premise will be that emotional abilities initially emerge from “instinctual” operating systems of the brain, which allow animals to begin gathering food, information, and other resources needed to sustain life.[5]

Jaak Panksepp concludes more abstractly:

> …emotions are the psychoneural processes that are especially influential in controlling the vigor and patterning of actions in the dynamic flow of intense behavioral interchanges between animals, as well as with certain objects during certain circumstances that are especially important for survival.[6]

---

[4] Although there are many variations in the terminology for emotional systems in the brain, most approaches consider “emotion” to refer to the built-in responses. Since the brain captures information about these responses in cortical memory/processing structures, these representations become available (in some form) to consciousness and are part of the system for responding to the world. Antonio Damasio, for example, considers “feelings” to be “the private, mental experience of an emotion,” while “affect” refers to the “entire topic of emotion and feeling.” My concern in this chapter is for the impact of primary emotional mechanisms on the large-scale representational structures in the cortex, and I use Damasio’s most general term—affect—to point to this broader area of inquiry. In this sense, “an affect” is a cortically articulated and mediated emotional response. See Antonio R. Damasio, “A Second Chance for Emotion,” in Richard D. Lane and Lynn Nadel, *Cognitive Neuroscience of Emotion* (Oxford: Oxford University Press, 2000), pp. 12-23.

[5] Jaak Panksepp, *Affective Neuroscience: the Foundations of Human and Animal Emotions* (Oxford: Oxford University Press, 1998) , p. 25.

[6] Jaak Panksepp, *Affective Neuroscience*, p. 48.

Working primarily with the clusterings of affect in rats, Panksepp proposes four major categories of emotion to encompass the most important survival activities:

1. Seeking (including food, water, warmth, sex, social contact)
2. Fear
3. Rage (including hate, anger, and indignation)
4. Panic (including loneliness, grief, and separation distress)

Dividing emotions into categories can be very simple or highly elaborate. Edward Rolls, for example, keeps his assumptions about the basic logic of emotion as simple and general as possible: "emotions are states elicited by rewards and punishers, including changes in rewards and punishments." Although he suggests that "different emotions could be produced and classified in terms of rewards and punishers received, omitted or terminated," he is much more interested in the shared structures and processes that signal, evaluate, and record these rewards and punishers than in the logic of differentiation within the system.[7] In sharp contrast, some neuroscientists have followed the model of appraisal theory in cognitive science and have derived highly articulated categories. Gerald Clore and Andrew Ortony, for instance, argue that emotional responses rely on three cognitive criteria of "goals, standards and attitudes." Their proposed taxonomy of emotions based on these criteria is shown in the diagram below.

[7] Edmund T. Rolls, *The Brain and Emotion* (Oxford: Oxford University Press, 1999), pp. 60-61.

Global Structure of Emotion Types

(from Gerald L. Clore and Andrew Ortony, "Cognition in Emotion: Always, Sometimes, or Never?" in Lane and Nadel, *Cognitive Neuroscience of Emotion*, p. 30)

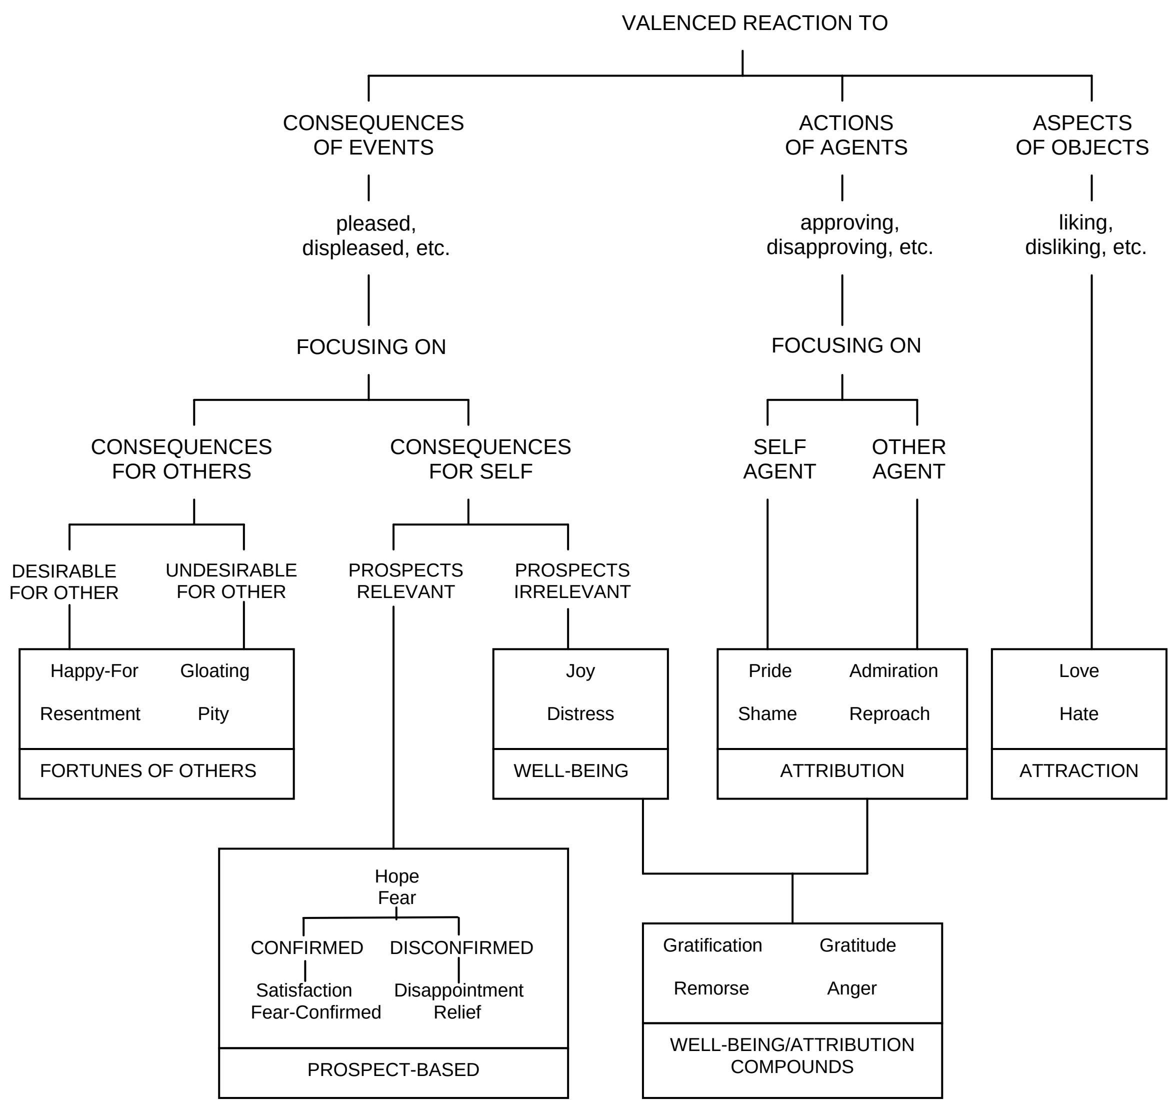

Panksepp's comparatively simple four-fold scheme is interesting because he attempts to discern common elements between behavioral studies and information on discrete neurotransmitter pathways and activation patterns. His stress on "Seeking" as a major emotional system also provides an account for infants and their preference for the unfamiliar red ball as well as for my son's odd predilection for Buddhas. Panksepp's recasting of seeking—a form of cognitive and purposive activity—as an *emotion* reflects the ways in which neuroscientists have begun to rethink the fundamental relationships between perception, cognition and emotion.[8] Basic emotions are genetically encoded response systems that present biological salience to the brain. Yet these response systems do not work in isolation but as guides that strongly direct activity in the brain. Their function is precisely to focus perception, thought, and indeed the entire representational system of the brain on what matters. They do so through the gating mechanisms discussed in Chapter 2 and through the sort of top-down attentional feed-back presented in Chapter 3. The pervasiveness of emotional biasing throughout the cortex has become increasingly clear in the past decade as neuroscientists explore in detail the connections between cortical regions. The interactions between "cognitive" and "emotional" centers have forced researchers to conclude that the old distinction between cognition as ideally objective and disinterested and emotion as the importuning of subjectivity and bodily desire simply no longer fits our understanding of the brain.[9] Antonio Damasio, initially basing his work on the

---

[8] Panksepp's SEEKING system centers on the dopamine circuits of the frontal cortex, basal ganglia, and ventral tegmental area (VTA) which he argues "has long been misconceptualized as a 'reward or reinforcement.'" Instead, he suggests, "it appears to be a general purpose neuronal system that helps coax animals and humans to move energetically from where they are presently situated to the places where they can find and consume the fruits of this world." Jaak Panksepp, *Affective Neuroscience*, pp. 53-54.

[9] For an account of a similar transformation in thinking in social psychology, see Joseph P. Forgas' introduction to his edited volume, *Feeling and Thinking: the Role of Affect in Social Cognition* (Cambridge: Cambridge University Press, 2000), pp. 1-7.

functions of the prefrontal cortex rather than on emotion per se, announced his rejection of this model in the title of his important book *Descartes' Error*.[10] The brain as a *res cogitans* is, in its deep structure, a *res extensa*.

Emotional information influences processing throughout the brain, but two cortical regions in particular are important for the shaping of that information. First is the so-called limbic system, and second is the orbitofrontal cortex. These two cortical regions have strong connections to the subcortical units that control the visceral responses we usually associate with strong emotion, and they are equally connected to the sensory cortex, the medial temporal lobe (MTL) memory structures, and the attentional system. The limbic system, however, has become close to an anachronism, and Joseph LeDoux has argued that we should drop the term because it invokes an outmoded model of brain function.[11] Indeed, although the two key cortical components of the limbic system—the amygdala and the hippocampus—are important in emotion, they have profoundly different roles in emotional experience. However, I introduce the term (and acknowledge Le Doux's critique) primarily because it appears extensively in the neuroscientific literature, though in general I use the term simply as a short-hand reference to the amygdala and hippocampus. In any case, within the limbic system, the amygdala is the more central structure for emotional experience. The hippocampus responds to emotional information in its role as facilitator of memory, but the amygdala

---

[10] Antonio Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Avon Books, 1994).

generates important cortical components of the emotional information to which the hippocampus responds. Because the amygdala more directly participates in emotional response and because the orbitofrontal cortex is not well developed in that prime experimental subject—the rat—the amygdala's structure and role in emotional experience have been extensively studied. The orbitofrontal cortex, in contrast, usually has been considered as part of the working-memory decision-making cognitive system. Its functional regions and cortical connections, moreover, remain poorly understood. However, with the growing sense of emotion's integration into cognitive function, the orbitofrontal cortex has come to be seen as perhaps the most important site for this integration.

Because the story becomes complex—much more complex than the relatively self-contained account of sensory processing presented in Chapter 3—sketching an overview of current approaches to emotional processing requires the separate delineation of the two interrelated systems. The amygdala-centered system comes first both in terms of evolution and of processing logic: it is the primary cortical pathway for emotion in non-primate mammals and remains important for emotional experience even in humans.[12] The orbitofrontal system, coming second,

---

[11] Joseph Le Doux, *The Emotional Brain*, pp. 98-102.

[12] It has been long noted that removal or dysfunction of the amygdala in humans has a far less catastrophic effect on affective experience than in other primates. For a discussion of the issues, see John P. Aggleton and Andrew W. Young, "The Enigma of the Amygdala: On its contribution to Human Emotion," in Lane and Nadel, *Cognitive Neuroscience of Emotion*, pp.106-28.

is essentially a neocortical elaboration of the amygdala network that allows representations, plans, and responses of much greater complexity.

Although the amygdala appears to mediate the experience of many emotions, fear has been the most intensely studied because it is the most amenable to systematic manipulation in the laboratory. Many rats and psychology students have encountered torment, but their sacrifices have allowed us significant insight into the fear system as a model for emotion and its systemic role in brain function.

The Traces of Fear

A fearless animal is, in all probability, a dead animal. Fear is not good in itself, but is far better than being caught flat-footed by a creature with sharp teeth. What concerns humans, however, is not *appropriate* fear, but the inappropriate. Anxiety disorders are the most prevalent of modern mental discomforts. So we as a society spend considerable resources trying to clarify the mechanisms that shade from fear of present real danger into nameless dread. The key seems to be memory: how frightening objects and events are categorized, remembered, and reinvoked.

As with much research in cognitive science, the first subjects were rats. They were zapped with electric foot shock and learned to fear all sorts of cues that the shock was coming again. Scientists studied how long it takes to learn the cues, how cues mixed with other events complicate learning, and how long it takes to persuade a rat that the old cue is harmless now. The researchers began to probe into the brain. They damaged specific regions to discover what was needed to make fear-driven learning possible. Some lesions destroyed the fear; some destroyed the ability to learn cues, while others destroyed the ability to remember context ("Oh no! Not the cage with the two water bottles again!")

This work with rats, combined with studies of the impact of brain damage on fear responses in human patients, helped delineate a neurobiology of fear. Knowing where to look, researchers at present are using recently developed imaging techniques to more fully elucidate the interacting neural pathways that subserve our experience of fear.

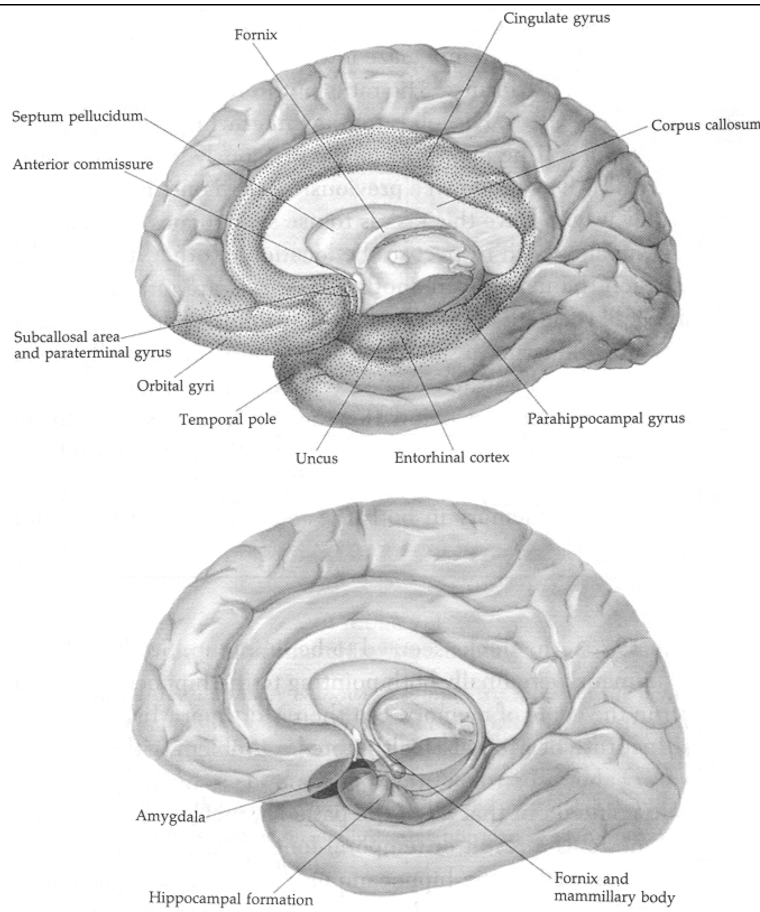

From Joseph LeDoux, *The Emotional Brain*, p. 190.

The center of the fear system is the amygdala. “Amygdala” is Latin for “almond” and refers to a small but complex almond-shaped structure next to the front of the hippocampus. The front part of the amygdala (the basolateral amygdala) receives input from the sensory areas of the thalamus (LGN for vision, MGN for hearing), from the sensory cortex, as well as from those cortical areas around the hippocampus that seem to serve the short-term episodic memory

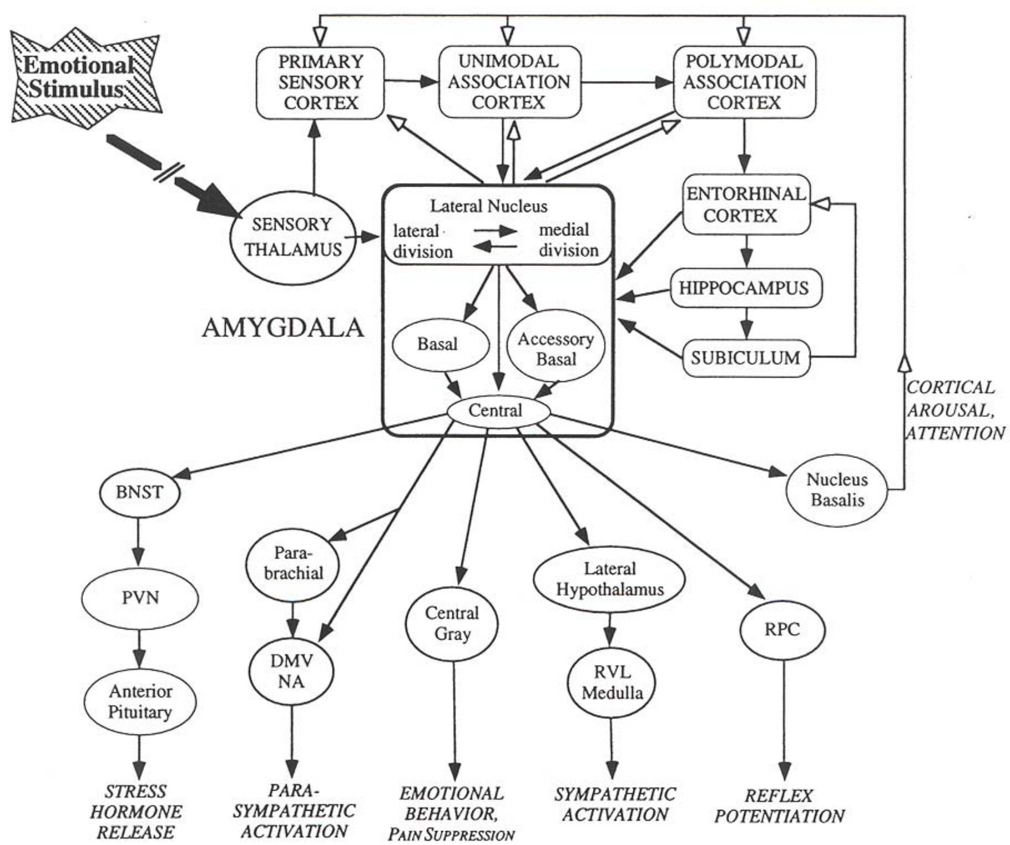

From: LeDoux, JE; Muller, J. Emotional memory and psychopathology. *Philosophical Transactions of the Royal Society of London*. Series B: Biological Sciences, 1997 Nov 29, 352(1362):1719-26.

system. The front part of the amygdala then connects to the middle region, which in turn connects to the hypothalamus. The hypothalamus in concert with the pituitary and adrenal glands controls the autonomic nervous system responses and releases various neuromodulator hormones into the blood stream. The amygdala has connections via the stria terminalis to the nucleus basalis that largely controls the release of acetylcholine in the brain and thus shapes cortical arousal states. The amygdala also has direct, often reciprocal connections to many cortical regions. Tracing these circuits in a bit more detail clarifies the functioning of this basic fear system.

The Path from Thalamus to Amygdala

This route is very fast but not overly refined. In the visual system, for example, the shortest path to the amygdala is from the retina to the superior colliculus to the thalamus and amygdala. The synapses that come from the retina, however, largely pass information from the rods. They respond strongly to movement in the periphery and to curved lines but not to color. The neurons of the thalamus subject the retinal activations to some processing, but the result is a form of feature detection rather than the more sophisticated object recognition that requires the visual cortex. Nonetheless, the information made available via this route is sufficient—in the absence of cortical processing—to persuade a rat to fear a light or a pure tone.[13]

---

[13] For debates around the learning of conditioned stimuli, see page 120.

The Routes from the Cortex and Hippocampus to the Amygdala

These are the "usual" routes for fear generated by higher processes. The startle reflex elicited by an unidentified "something" buzzing loudly past my ear took the shortcut, but cautiousness around the paper wasps as they mind their own business comes from a more structured recollection. I recently stepped on one in my bare feet when going to pick up the morning paper. I cannot recall the pain, but I now slow down and look carefully when I pass that spot. This form of conditioning appears to be mediated—like all declarative memory—by the hippocampus and the cortical areas around it. Disable the hippocampus within a short time after the encounter, and the conditioning fails. Removing the hippocampus later has no effect: the memory has been assimilated.

The Routes from the Amygdala

What sears the context of trauma into the traumatic memory is part of what happens when activation spreads to the basolateral amygdala. The basolateral amygdala passes the activation on to the central nucleus of the amygdala, which then activates a large array of responses. The central nucleus innervates the nucleus basalis, which activates acetylcholine-producing synapses that selectively control cortical arousal. The amygdala also activates the hypothalamus, which controls bodily hormonal responses and peripheral nervous system arousal. We flush or grow pale; our heart pounds; our breath freezes, our mind focuses, or freezes. The central nucleus of the amygdala sets off this complex repertoire of reactions to help us marshal our resources for action.

The arousal system of the nucleus basalis and the hormonal activation of the hypothalamic-pituitary-adrenal system are particularly important in exploring how the brain decides what matters. The attentional system focuses resources on specific sensory systems: hearing sharpens remarkably when one walks home alone at night. Moreover, the hippocampus receives cholinergic arousal and reacts both to the peripheral release of hormones and to the stimulation of more specific intra-cranial stress hormone pathways: these stimuli strongly activate the hippocampus' role in committing the focused cortical information to long-term memory. *This*—the body says—matters.[14]

Orbitofrontal Elaborations

The question is, What exactly is the *this* that matters? And why precisely does it matter? The amygdala is good at activating visceral and cortical systems to take note of events that hurt, but it must rely on other brain regions to integrate the fact of exigency into larger patterns for response and remembering. As Rolls notes,

> [T]he amygdala has direct projections back to many areas of the temporal, orbitofrontal, and insular cortices from which it receives inputs.... [T]he function of these back projections includes the guidance of information representation and storage in the neocortex, and recall (when this is related to reinforcing stimuli). Another interesting set of output pathways

---

[14] The brain works by complex indirect routes here that are still not well understood. See B. Setlow, B. Roozendaal, and J. L. McGaugh, "Involvement of a basolateral amygdala complex–nucleus accumbens pathway in glucocorticoid-induced modulation of memory consolidation" *European Journal of Neuroscience* 12 (2000):367-375.

> of the amygdala are the projections to the entorhinal cortex and dentate gyrus, and ventral subiculum, which provides a major output of the hippocampus.[15]

As discussed in Chapter 3, the latter structures—the hippocampus, etc.—are the major components of the MTL systems for episodic and semantic memory. As Rolls further argues,

> we humans and other animals do not generally want to learn that a particular pure tone is associated with reward or punishment [i.e. the sort of conditioning possible through the thalamic-amygdalar route]. Instead, it might be a particular complex pattern of sounds such as a vocalization (or, for example, in vision, a face expression) that carries the reinforcement signal, and this may be independent of the exact pitch at which it is uttered. Thus cases in which some modulation of neuronal responses to pure tones in parts of the brain such as the medial geniculate (the thalamic relay for hearing) may be rather special *model* systems (i.e., simplified systems on which to perform experiments), and not reflect the way in which auditory-to-reinforcement pattern associations are normally learned. (For discrimination of more complex sounds, the auditory cortex is required.) [Emphasis in the original.][16]

Rolls, among many others, proposes the orbitofrontal cortex as the brain region which performs the more complex forms of integration of emotional valences with cognitive information (i.e. high order perceptual information and its associated memory structures).[17] First, the orbitofrontal cortex has the right connections, and secondly, lesion studies point to a corresponding integrative role for it in emotional responses.

---

[15] Edmund T. Rolls, *The Brain and Emotion* (Oxford: Oxford University Press, 1999), pp. 97-98.

[16] Rolls, *The Brain and Emotion*, p. 104.

[17] Rolls, *The Brain and Emotion*, p. 127.

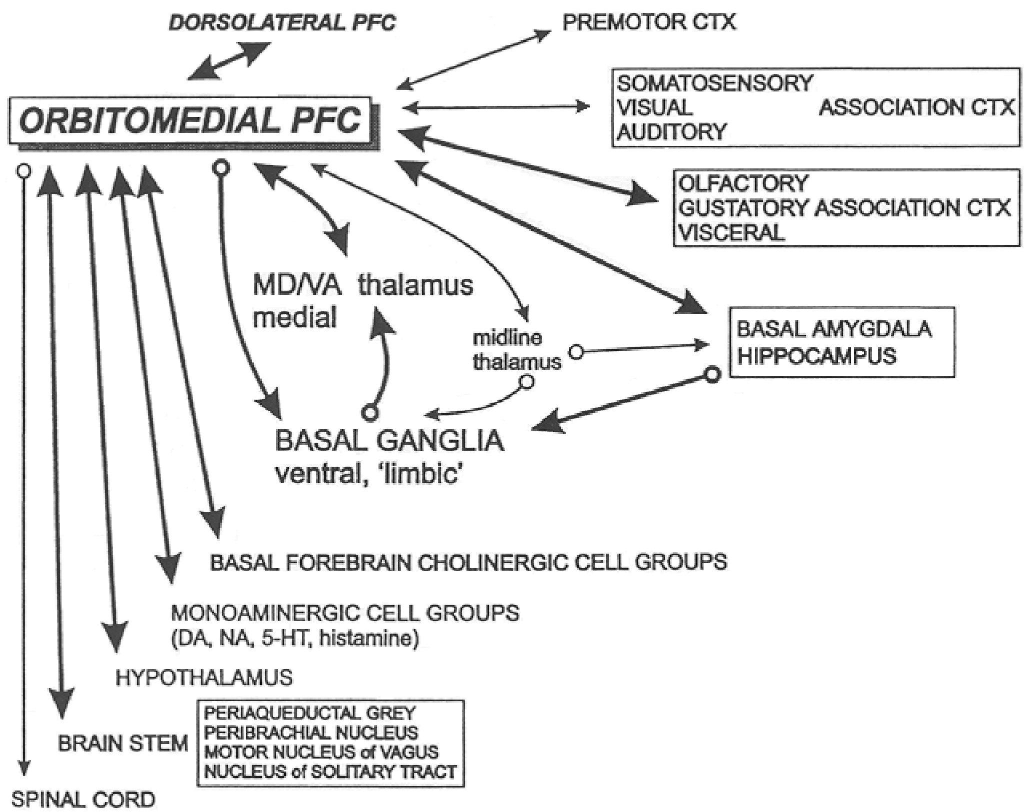

As the above diagram indicates, the orbitomedial prefrontal cortex (i.e. the orbitofrontal cortex) is strongly and reciprocally connected with the working-memory structures of the dorsolateral prefrontal cortex, the sensory cortices, the amygdala, hippocampus, and the various subcortical response systems.[18] It moreover has a strong link to the basal ganglia motor-planning system that in turn has backprojections mediated by the thalamus.

Perhaps the most extensively studied set of connections is the dopamine (DA) system linking the orbitofrontal cortex, the midbrain dopamine-producing

---

[18] The diagram is from Henk J. Groenewegen and Harry B.M. Uylings, "The prefrontal cortex and the integration of sensory, limbic, and autonomic information," *Progress in Brain Research* 126 (2000):19.

cells and the ventral striatum in the basal ganglia. These three units jointly represent a “reward system” for the brain:

1. The **midbrain dopamine** cells become strongly but briefly active either when a stimulus associated with a reward is detected or a reward appears that had not been predicted, but not when an anticipated reward appears. However, they become depressed when the anticipated reward does *not* appear.
2. In a delayed-task experiment, various populations of cells in a rat’s **striatum** become active when (a) an instruction appears telling the rat to anticipate a signal for a reward, (b) while the rat awaits the signal to act (to get the reward), (c) when performing the action, (d) while awaiting the reward, and (e) when the reward actually arrives.
3. In the **orbitofrontal cortex**, however, cells in the dopamine circuit become briefly active only in response to the instruction, while waiting for the reward (after the action), and when the reward arrives. (In contrast, populations of cells in the dorsolateral prefrontal cortex remain active during the delay between the instruction and the signal to act: these reflect working memory and planning for movement.)

These three different types of responses allow the brain to first capture reward information (the dopamine cells’ response to unanticipated rewards), remember associations that can predict the reward next time, and learn to act on the prediction. In the orbitofrontal cortex, populations of cells discriminate between different types of rewards and different types of stimuli, and they can quickly reverse their response patterns if the associations between stimulus and reward changes.[19] Edmund Rolls points out that when, for example, responses to visual stimuli reverse, this reversal appears in the orbitofrontal neurons and not one synapse earlier in the inferotemporal (IT) cortex where the object's visual identity presumably is stored. That is, the prefrontal cortex is the site of the integration of visual and reward information rather than earlier in the visual system itself. Moreover, a monkey's orbitofrontal neurons respond to a particular food only when it is hungry rather than sated: these cells also integrate visceral information into their interpretation of visual data.

The orbitofrontal cortex not only integrates emotional, sensory and visceral information with memory and action patterns. It also makes these integrations available both to other cortical systems and—through encoding into memory structures—to its own future integrating activity. This accessibility of emotionally valenced response patterns seems to be a key feature of the prefrontal cortical system. Many researchers agree in particular that both consciously available emotions and the ability to reflect on and (to some extent) control emotional

---

[19] This account of the dopamine system relies primarily on the work of Wolfram Schultz and his lab. See such article as Jeffrey R. Hollerman, Leon Tremblay and Wolfram Schultz, "Involvement of basal ganglia and orbitofrontal cortex in goal-directed behavior," *Prorgess in Brain Research* 126 (2000): 193-215; Wolfram Schultz, Leo Tremblay, and Jeffrey R. Hollerman, "Reward Processing in Primate Orbitofrontal Cortex and Basal Ganglia," *Cerebral Cortex* 10 (Mar 2000):272-83; and Wolfram Schultz, "The Reward Signal of Midbrain Dopamine Neurons," *News in Physiological Science* 14 (Dec 1999)249-55.

reactions arise out of the prefrontal cortex's creation of representations of the associations between emotional responses and their contextual information.[20]

Antonio Damasio is willing to go further. He argues that "structures in the ventromedial prefrontal cortex provide the substrate for learning an association between certain classes of complex situations, on the one hand, and the type of bioregulatory state (including emotional state) usually associated with that class of situation. The ventromedial sector holds linkages between the facts that compose a given situation, and the emotion previously paired with it in an individual's contingent experience."[21] This linkage of emotional state with episodic context is a form of second-order mapping that serves a pivotal function in Damasio's model for the central role of representations of emotion in consciousness. Damasio begins with "core consciousness," which occurs "when the brain's representational devices generate an imaged, nonverbal account of how an organism's own state is affected by the organism's processing of an object and when this process enhances the image

---

[20] See, for example, Geoffrey Schoenbaum, Andrea A. Chiba, and Michela Gallagher, "Changes in Functional Connectivity in Orbitofrontal Cortex and Basolateral Amygdala during Learning and Reversal Training," *Journal of Neuroscience* 20.13(July 1, 2000):5179-89. In the quite different context of exploring the basis for conscious emotions, Richard D. Lane proposes a role for orbitofrontal cortex as providing the first level of abstraction, the gateway between the automatic responses of the amygdala and the conscious control of higher levels of emotional representation. See Richard D. Lane, "The Neural Neural Correlates of Conscious Emotional Experience" in Lane and Nadel, *Cognitive Neuroscience of Emotion*, pp 345-70.

[21] Antoine Bechara, Hanna Damasio, and Antonio R. Damasio, "Emotion, Decision Making and the Orbitofrontal Cortex," *Cerebral Cortex* 10 (March 2000):296.

of the causative object, thus placing it in a spatial and temporal context."[22] This is a complex idea that requires some elaboration. First, notice that the process of "enhanc[ing] the image" ties well into the arguments about consciousness entailing the participation of attentional, sensory and object-memory systems in adaptive resonant loops. To this sensory-object-memory loop Damasio adds the representing of the sort of contextual information associated with episodic memory. Finally, he binds all these to cortical representations of the impact of the event on the person or animal, i.e to representations of emotions. Note that this core consciousness is explicitly non-verbal. What we tend to talk to ourselves about in our unending inner dialogue is a part of "extended consciousness," an elaboration of the representations made available through the functioning of core consciousness.

The orbitofrontal cortex, in Damasio's view, is not necessary for core consciousness, but it does become important for synthesizing the information made available *through* the representational activities of core consciousness. Damasio describes the processes of assimilation that begin when an object or event—for whatever reason—triggers basic emotional responses:

> [E]motion induction sites trigger a number of responses toward the body and toward other brain sites and unleash the full range of body and brain responses that constitute emotion. First-order neural maps in both subcortical and cortical regions represent changes in body state.... Feelings emerge. The pattern of neural activity at the emotion-induction site is mapped in second-order neural structures. The proto-self [see below] is altered because of these events. The changes in proto-self are also mapped in second-order neural structures. An account of the foregoing events,

---

[22] Antonio R. Damasio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (New York: Harcourt, Brace & Co., 1999), p. 169.

> depicting a relationship between the "emotional object" (the activity at the emotion-induction site) and the proto-self, is thus organized in second-order structures.[23]

As already seen, Damasio argues that the site of the second-order structures that map the relationship between an event and the emotional response is precisely the ventromedial prefrontal cortex. The "proto-self" that Damasio invokes is the brain's slowly constructed internal representation of the organism of which it is a part: "As the representations of the body grow in complexity and coordination, they come to constitute an integrated representation of the organism, a proto-self."[24] These mappings, both first and second order, all become available for incorporation into the ever more complex and integrated structures through which we mediate our encounter with the world, especially as we know it through the reflections of extended consciousness.

Damasio's account synthesizes much of what we know about the brain and emotion, even if it is still preliminary because there remains so much we do *not* know. Damasio proposes mechanisms for the encoding of experience that accommodate both the rawest of emotional encounters as well as forms of reflections where affective valences are but the faintest of echoes. I find the model a very compelling beginning that corresponds with—and gives us ways to systematically explore—our lived engagement with the world. It offers a vision of how, as our experience grows, the structures that inform our manner of engaging

[23] Damasio, *The Feeling of What Happens*, p. 283.

[24] Damasio, *The Feeling of What Happens*, p. 284.

the world grow correspondingly complex. These structures form a map that charts the particular history as well as the shared creaturely logic within the intersection of the self and the world. There remains much *terra incognita* here, but the exploration of the systemic function of emotion as the trace of the body in the neural organization of the mind seems well under way.

PRIMING THE PUMP

Tracing affective pathways in adults is deeply complex because we have much history behind us. We have not remembered everything that has happened to us—just events that somehow mattered. Nor have we remembered every detail of events we found meaningful: we remember the aspects that proved significant from our particular perspective at the time. These events shade into ever more articulated maps of the possibilities for emotion, for experience, and for our selves as actors in the world. Lesion studies help clarify what happens when parts of these maps become inaccessible, and imaging studies can help us watch the brain invoking those maps. Yet the higher cortical systems that participate in evaluating what matters are hard to read precisely because they are well established.[25] Another approach to exploring the way in which the brain organizes what it finds important in experience is to pay attention to developmental issues. In watching an infant's growth, we see the mind constructing itself. We see the movement from innate

---

[25] Well established cortical neural networks tend to have sparse connections with minimized activation. Rolls points out that therefore current techniques probably are

responses to complex behavior mediated by the infant's cumulative experience of the world.

For example, consider the baby who prefers the new red ball to the previously seen blue one. This response is just the opposite of the adult so-called Mere Exposure Effect. For adults, and even older infants, briefly seeing one of two random images is sufficient to make one prefer that image over the other. Infants, in contrast, like novelty. This interest in the new takes many forms. One of my favorite examples is an experiment in impossible movement. A hidden assistant slowly moves a doll along a display until it disappears behind an obstruction. Then the assistant moves an identical doll from behind a paired obstruction a foot away after a delay equal to the time it would have taken the doll to traverse the distance. Infants stare in wonder at this discontinuity. (Having subjected my own children to cognitive dissonance at an early age, I imagine how much fun these experiments must have been.) Researchers use this experiment to argue for certain types of conceptual understanding, but I suspect that can also reflect a higher-order bottom-up prediction based on the types of calculations used to make smooth tracking of objects possible. In any case, the doll clearly violates the regularity of normal experience, and the infant pays close attention.

The unexpected draws the infant's attention, but does it generate an emotional response? In particular, does it cause the release of any of the hormones—from acetylcholine to cortisol, dopamine and oxytocin—that modulate

---

inadequate to capture much of the crucial activity in such regions as the orbitofrontal cortex. Edmund T. Rolls, *The Brain and Emotion*, p. 121.

hippocampal activation and activate gating in the temporal lobes and prefrontal cortex? The event is new, but is it important enough to be preferentially registered and remembered? Inborn mechanisms for generating attention, like the response to novelty, are a form of bet that the general category of inputs that share the attention-getting feature is important.[26] "The New" is a clearly good, if extremely general category of this sort. Among the most studied of more specific mechanisms—and among the most important—is "face recognition."

Young infants pay attention and selectively look at faces. Older babies continue to spend a lot of time looking at faces. They learn how to "read" faces (expressions that convey "Yes." "No!" or merely a mild "OK"), especially that of the primary care-giver. Adult brains have a cortical region specifically dedicated to identifying faces.[27] Faces, that is, prove to be an important part of an infant's early environment. Thus the neonate's seemingly simple built-in attraction to faces plays a pivotal role in shaping the infant's emotional and social life. In reflecting on this preference, however, it is important to recall that the cortical circuits in a neonate are not very functional. The superior colliculus, interacting with the thalamus, controls early eye movement. The superior colliculus, moreover, receives direct input from only the ganglial cells of the retina that take their input from rods (good

---

[26] See, for example, the discussion in Jeffrey L. Elman et al., *Rethinking Innateness: a Connectionist Perspective on Development* (Cambridge, MA: MIT Press, 1996), pp, 107-118.

[27] Current research suggests that this "dedicated" region in fact accomplishes face recognition as one among other high-order perceptual tasks. See, for example, Isabel

in dim light, poor resolution, no color) rather than the higher acuity cones of the foveal region.

The neonate is attending to particular features shared by faces rather than to faces as "faces." It will look just as happily at approximately face-like blobs. The new-born prefer clearly marked, curved line segments, movement, and objects that appear in front of them: faces (especially hair-lines and chins) usually meet these criteria. It also turns out that visual search (the scanning with the eye) is poorly developed in young infants. While older infants (starting at 2 months) will move their focus from place to place within an object marked by borders (like a face), young infants even at 1 month remain fixated on the borders themselves. These limitations in the perceptual and cognitive mechanisms do not seem to offer a very promising beginning, but they also show the real power of small systemic biases. The face is almost always there when biologically important things happen: feeding, being wrapped or unwrapped, cuddling, and so on.) As the infant's cortex matures, faces come to matter. From 2 months, infants start to make eye contact, and at 5 months, the infant strongly fixates on the face. After 5 months, however, the baby starts to squirm to look around. Life gets more complicated and varied: the baby learns to crawl, to direct not only its gaze but its movement. As the infant expands its range, it still checks up on the person it knows best. Interestingly, "[w]hen infants from 6 to 9 months looked toward the parent while playing at a distance, they were as likely to look at the parent's body as the parent's face. In

---

Gautier, Marlene Behrmann, and Michael J. Tarr, "Can Face Recognition Really be Dissociated from Object Recognition?" *Journal of Cognitive Neuroscience* 11.4(1999)349-70.

contrast, children from 14 to 22 months not only looked more frequently at the parent but looked almost exclusively toward the face."[28]

The nature of the internally represented relationship to the parent (from which the child seeks assurance and response in looking) has a profound impact on the baby's explorations. Developmental psychologists have correlated an infant's ability to cope with stress to the quality of its attachment to its parent (usually, its mother). The categories they use are broad, and the scoring criteria clearly reflect cultural biases. Nonetheless, the measurable effects are there: whether the mother is in a room makes no difference to a 5 month old baby when a stranger approaches: the baby is interested, and heart beat drops (a sign of focused, aroused attention). A 9 month old is usually a bit distressed, with an elevated heartbeat, when a stranger approaches. When the mother is in the room, distress lessens. The better the attachment (as parameterized), the less a baby shows stress.

The usual argument for why this correlation appears is because of coping. Initially, affective responses, here measured by the release of the stress hormone cortisol, mark biologically important input. However, with experience (and an increasingly functional cortex), other, more complicated contextual responses become available: these usually are motor sequences associated with the input; i.e., infants develop coping strategies. If such cortical structures exist, they mediate and can inhibit the emotional response. As a result, cortisol production drops significantly from the third to the sixth month. At six months, a baby in the

[28] Holly A. Ruff and Mary K. Rothbart, *Attention in Early Development: Themes and*

doctor's office for a check-up may fuss and cry as much as a 3 month old, but cortisol levels are significantly lower. This dissociation between crying (behavioral distress) and cortisol release has surprised researchers: by perhaps 6 months, a baby has discovered that crying can be a very effective coping strategy.

Cortical structures can suppress emotional responses. Access to coping responses is one learned structure. Simple habituation is another; the infant concludes: been there, done that, it's not so bad. Distraction is a third approach. Affective responses focus attention on salient input. After about 3 months, one can interrupt a baby's distress by drawing its attention to something new.[29] When the distraction loses its charm, however, the infant returns focus to its former complaint. (Although I can offer only anecdotal evidence, I always moved my children from the scene of the offense so that the conspecifics which might reactivate the former state all changed.) Between 9 and 18 months, as the prefrontal cortex develops, an infant gains increasing control over his or her attentional system. Posner and Rothbart, who have explored the role of the cingulate cortex in this process, write:

> It is not so much the stimulus of pain or distress that activates cingulate, but the feeling of distress related to pain or efforts to cope with or control these feelings. Thus amygdala-cingulate interaction might be a reasonable candidate for the earliest form of self-regulation in the infant. In the infant, control of orienting is partly in the hands of care-givers' presentation of relevant information. However, the infant is clearly involved in soliciting attention from the adults (Stern 1985). During the first years of life, more direct control of attention passes from care-givers

---

Variations (New York and Oxford: Oxford University Press, 1996), p. 87.

[29] Michael I. Posner and Mary K. Rothbart, "Attention, self-regulation and consciousness," *Philosophical Transactions of the Royal Society of London,* Series B (1998):1915-1927.

> to infant. It seems likely that the same mechanisms used to cope with self-regulation of emotion are then transferred to issues of control of cognition during later infancy and childhood.[30]

Studying the mental life of babies suggests that for them, affect and cognition are closely bound together. Only as emotional responses diminish in intensity and frequency do the two systems seem to disaggregate.

EMOTION AND COGNITION

The last section explored several aspects of the relationship between affect and cortical systems in infants. One major point to be stressed is that there is not much in the way of higher cortical structure at the beginning and that, instead, nascent cognitive systems are strongly driven by emotional responses, that is, the built-in brain responses to crucial environmental and somatic cues. Secondly, we know that activation of the amygdala cues the hippocampus to assimilate current input into long term memory. Third, the hippocampus has a dense population of receptors for most of the emotion-related hormonal neuroregulators: epinephrine, cortisol, oxytocin, etc. And finally, the prefrontal cortices for higher-order structuring are still developmentally immature.

These facts suggest an important conclusion. In an infant, the hippocampus and the MTL memory system capture the episodic content of emotional experience in particular. This makes sense (though reasonableness is admittedly not a very

---

[30] Michael I. Posner and Mary K. Rothbart, "Attention, self-regulation and consciousness," p. 1922.

good argument for truth): the infant needs to know fastest what matters most. However, the hippocampal memory system does not recall discrete "things:" it is a slowly trained neural network, but it is a neural network nonetheless. Like the increasingly higher-order bottom-up systems that feed it activations, it captures the strongest regularities in its input. Because of the gating processes that mark biological salience, emotional events in particular serve as the key early data as it sorts out an implicitly structured input domain. It is difficult to imagine the dimensional logic of such a space of built-in biological (emotional) cues, particularly given the early input: thumbs, toes, tongues, hunger, thirst, nameless anxiety, breasts, faces, bowel-discomfort, blankets, the sounds of rattles, wind-up music boxes and voices, and on and on, all perceived through poorly articulated sensory cortices. The sensory conspecifics surely do not appear in object memory in what we as adults take to be their "objective" perceptual form. In whatever form their salient features are coded, however, they establish links to cortical and sub-cortical information on the recurring internal states to form a system of patterns.

The dimensionality of the space for semantic representation set down through the first year of a person's life is that of affect and bodily need. This is a remarkable conclusion: the memory at the heart of our cognitive capabilities is deeply human. What exactly does this mean? It means that the mutually defined systemic relationships that bind together associations in semantic memory can be described as vectors in affective experience space, whatever that may be. The sensory material in "Roses are red, violets are blue"—the shapes, colors, and smells, as well as the letters and sounds—are all components of the higher regions of the sensory cortices that provide the elements to be linked in semantic coding, which, being at the top, can link any modality with almost any other.

The notion of affective commitments within all remembered information may seem odd. "Water is wet." Who cares? Actually, babies happily splashing in a tub enjoy that fact quite a bit. "Water is H_2O (a famously belabored philosophical statement):" the affective commitment here is admittedly more attenuated. For most people, the sentence is a bit of rote learning they were compelled to memorize sometime in school. Fear of failure or a desire to please is sufficient to account for its place, and indeed its meaning ramifies no further than perhaps a teacher's face or a parent's smile. Under any circumstances, the criteria for the cortical calculations that marked the fact as worthy of commitment to memory—the criteria used to bind together the disparate parts of the sentence in memory—were part of the affectively driven semantic system. The system can draw on itself, remaining consistent with its earlier dimensional self-structuring, to articulate itself as it expands. It is important to recall here the representational power of even the fairly simple coarse coding of the retinal color information. Systems of coding built upon affective patterns can cover an immense breadth of experience as the infant grows to a toddler, child, teen-ager, and adult. The system grows highly articulated so that direct links to primary emotional experience (the sort of information perhaps stored in the retrosplenial cortex) can become very attenuated indeed. The semantic memory system, when fully elaborated, must be capable of representing all of human experience from the most immediate of bodily pleasures to the most abstract mathematical musings. However, all these memories remain vectors within affective experience space. It is hard for a mathematician to explain why homological algebra matters, or for me to explain why 13th century classical Chinese poetry matters in a deep, viscerally emotional way, but they do.

Converging Views on the Role of Affect in Cognitive Structure: Perspectives from "Classical" Cognitive Science and Situated Robotics

In the previous section, I combined the basic self-structuring feature of neural networks with what we know of the neuroscience of emotions to argue that semantic memory is built upon an affective structure. Viewing semantic memory in this way—as a space determined by the dimensionality of emotional experience—is consistent with Damasio's account of the role of emotion in the higher-order representations that underlie consciousness. It also confronts the inadequacy of the naive notion of "objects" independent of the sorts of self-defining structural maps that the mathematics of neural net modeling has introduced into the neuroscientific account. Finally, it also removes the little homunculus who controls decisions about attention and action: salience (importance) automatically is a direct component of all objects and action plans that draw on semantic information to compete for final selection.

Other disciplines outside of neuroscience that deal with issues of modes of extracting information from the world have independently arrived at models that incorporate a fundamental emotional dimension to cognitive data as a way to assess physical saliency. The affective organization of semantic memory, that is, solves problems of decision-making and the allocation of resources that are part of any effort to usefully organize the data of experience: this is as true for robots that must learn to operate on their own as it is for a new-born infant who must quickly sort out what matters.

Recent work in cognitive science on the interaction of emotion, perception, and attention, for example, has come to stress the structural role of emotion in cognitive processes. Various groups have looked at such phenomena as semantic priming for emotions, perceptual biasing, and the role of affect in shaping attention. Paula Niedenthal and her associates in particular have come to an understanding of emotion and its role in categorization that approaches my more radical argument for affective coarse coding at the heart of semantic memory.[31] Consider, for example, her discussion of the internal structure of emotion (citations omitted):

> Although the notion of specific emotions is appealing to many emotion theorists who do not endorse a strict dimensional account of the structure of emotional experience, the notion of biologically basic emotions is not appealing to some. However, one can posit the existence of discrete, nondecomposable emotional states without assuming that the states have a biological basis. In this *discrete emotions* [her emphasis] approach, emotional states are differentiated, defined, and labeled through experience; represented as schemas or organized units of information (perhaps corresponding to clusters of mid-level emotion terms

[31] Indeed, the work done in the cognitive science of affect pushes affective information deep down into the cognitive structures (i.e. memory) through which perceptual input is processed. These recent models of the interactions suggest that the larger arguments of this chapter and the next stand even if the more radical (and, I think, more elegantly simple) version does not hold up.

in the individual's native language); and bound by culture.[32]

Niedenthal had previously observed in a footnote that "a set of emotional states that are not meaningfully reducible to a smaller number of common dimensions" is "not meaningfully distinguishable from a model that conceptualizes specific emotions as locations in a high-dimensional space" (p. 339). Studying the effect of affective priming on category judgments, she concludes that "along with their perceptual features and conceptual associations, objects and events are represented in terms of the emotional responses they have evoked" (p. 351). The categories for encoding these responses, furthermore, "might be thought of as schematic memories for specific emotion-eliciting episodes and stimuli" (p. 357).

Niedenthal is one of the editors of *The Heart's Eye: Emotional Influences in Perception and Attention*.[33] This volume includes many studies that stress the structural rather than merely adventitious interaction between emotion, perception, and cognition. As Shinobu Kitayama and Susan Howard explain about their paper, "Amplification and Semantic Priming:"

> In this chapter we review recent research on affective regulation of perception and comprehension, especially focusing on our own experimental work. First, we present two theoretical principles of affective regulation of cognition that can be identified in the literature, namely amplification and semantic priming. *Amplification* refers to affect's function to energize some cognitive or behavioral responses. We believe that this function of affect can have pervasive influence on

---

[32] Paula Niedenthal, Jamin B. Halberstadt, and Åse H. Innes-Ker, "Emotional Response Categorization," *Psychological Review* 106 (1999): 340.

[33] Paula Niedenthal and Shinobu Kitayama, ed., *The Heart's Eye: Emotional Influences in Perception and Attention* (San Diego: Academic Press, 1994).

> cognitive processes before the affect is semantically encoded. Affect however, can be semantically encoded as "good, "bad," "happy," "sad," and the like; and once so encoded, it can activate a variety of associated thoughts, images, and meanings. This more semantic or cognitive influence of affect is called *semantic priming*. [Their emphases.][34]

Kitayama and Howard take the early processing of objects and events in the environment to be a bottom-up analysis and therefore below semantic coding. If, however, we assume that object identification is a more complicated negotiation between higher and lower structures, then semantic coding can be a source of the affective information that leads to amplification. Top level structures easily fit within the process that Kitayama and Howard describe:

> The perceptual code corresponding to an impinging stimulus may be called the *target* code. Once the target perceptual code has been preconsciously activated, two additional processes come into play. First, although this activation can happen preconsciously, it is sufficient to evoke some associated affective response. Perhaps the target code thus activated preattentively summons affective circuits of the brain located in the limbic or subcortical regions. Second, simultaneously with the elicitation of the associated affect, a set of additional cognitive processes that operate in a more top-down and serial fashion, called *attentive processing*, is engaged. Attentive processing is selective, limited solely to the code to which it is directed. It, therefore, must first be directed to a target code and, once so directed, it furthers the processing of the information represented in the attended code and thus of the impinging stimulus. Notice that conscious perception is not merely a function of sensory (i.e. externally caused) activation. The external activation must be combined with attentional (i.e. internally caused) activation to produce a full-fledged conscious percept (pp. 45-46).

The model, like Damasio's, provides a continuum between affect, attention, and cognition. The preconscious attentional calculus is based upon the affective weight

---

[34] Shinobu Kitayama and Susan Howard, "Affective Regulation of Perception and Comprehension: Amplification and Semantic Priming," pp. 42-43.

of the input activation patterns. These "affective weights," to apply Niedenthal's analysis, are the emotional response codes. The structure of these codes that bind together the activation patterns as systemically significant are, in turn, what I suggest we call semantic memory.

The issue of strategies for extracting significance from a field of input to allow one to make appropriate decisions is crucial here, for it proves to be extremely difficult to create information systems that do the job flexibly and well. It does not matter whether the memory system that develops internal models to aid its parsing of the world is made of neurons and encased in flesh or is made of silicon and stored on magnetic media: the fundamental problems of building an internal system of representations that captures what matters in experience is the same. The approach of viewing semantic space as fundamentally affective in structure thus reflects a direction that researchers in artificial intelligence and robotics are heading. The question always has been how to make an independent robot properly assess its input: the device must integrate goals, constraints and—in the best of all worlds—past experience into its reaction to present conditions.[35] Internal representations as affectively valenced structures can accomplish this task

[35] See the discussion of autonomous agents in Margaret Baden, "Autonomy and Artificiality" in Andy Clark and Josefa Toribio, ed., *Cognitive Architectures in Artificial Intelligence*, p. 301.

and are the sort of embodiment that Andy Clark and—from a different perspective—researchers in "situated robotics" are exploring.[36]

In the parlance of computer science, "situatedness" is when "an agent's controller interacts directly with the environment of the agent rather than in a manner mediated by an internal formal description of that environment."[37] J. C. T. Hallam and C. A. Malcolm argue that a situated agent (i.e. robot) therefore "must be able to recover from its environment by appropriate sensing whatever it needs to know to determine its course of action, and thus must be able to participate in its environment." The next step beyond situatedness is embodiment, where "a physically embodied agent ... interacts directly with the world by means of physical sensors and physical activators, i.e. by means of direct physical causation" (p. 36)

Researchers in situated robotics both worry about the nature of robot semantic systems and stress that physical causality is central to any meaningful semantics. Discussing embodiment, Hallam and Malcolm argue that "an embodied agent engages in the causal pathways of which the real world is made, and this physical engagement with reality provides the ultimate grounding for any semantic machinery inside the agent. Without such physical grounding, any semantic entities found in the agent exist only in the eye of the beholder and are completely devoid

---

[36] See Andy Clark, *Being There*: *putting brain, body, and world together again* (Cambridge, Mass. : MIT Press, 1997).

[37] J. C. T. Hallam and C. A. Malcolm, "Behavior: perception, action and intelligence—the view from situated robotics" *Philosophical Transactions of the Royal Society of London,* Series A (1994):34. I rely on the account in this interesing article for the comments that follow.

of any intrinsic meaning" (p. 36). They distinguish between ascribed semantics and intrinsic semantics. Ascribed semantics, as the name suggests, is a system in which an external programmer creates the "semantic entities," that is, ascribes meanings and action routines to define the data objects to be encountered by the robot. In contrast, "[i]ntrinsic semantics is processed by the properly grounded symbol system of an autonomous robot.... [I]ntrinsic semantics involves the autonomous operation of an *active* symbol system and its connection to external referents via causal links used in its interaction with the world" (p. 39).

What then is the nature of the interaction that generates the grounded symbol system?

> It is characteristic of all robots so far constructed by roboticists, and of all the animals so far analyzed by biologists, that at the lower levels of sensorimotor interaction with the world, great use is made of goal-seeking mechanisms, i.e. of control. In autonomous agents, whether biological or artificial, the purposes and parameters of these goal-seeking mechanisms can provide a very convenient ready made source of grounded symbols. While there must also be other mechanisms involved (e.g. to cover the case of an autonomous agent which has the purpose of finding an algorithm for generating the *n*th prime number) there is no doubt about the importance of these goal-seeking mechanisms as a source of symbols, because they can neatly encapsulate both many of the capabilities of the agent, and the perceptions relevant to these capabilities.... Symbols which are grounded via autoteleological mechanisms will naturally thereby have intrinsic meaning for the agent so hosting them, i.e. autosemantics.... (p. 40-41)

Semantic memory, even for robots, should be affectively grounded.

There is a cost, however, associated with this solution. First, training—the building of the structure of grounded symbols—takes a long time: 18 years in a human is almost enough. Even if a robot could speed up the process, however, there is a clear conflict between speed of acquisition and appropriateness of internal structure. The fantasy of speed-reading all the great works to be all-knowing and wise cannot work: the hippocampal two-step assimilation along with prefrontal synthetic analysis is needed to provide both gating and structure. Secondly, the "goals" that mark the initial state would have to be very general: there can be no mention of explicit objects of any sort. Instead, any specific planning would have to be part of the training process. The Mars rover might get to Mars and just want to watch the moons rise. Third, emotions would not be an option (unlike Data's "emotion-chip") but built into both the structure of the hardware and into the emergent cognitive system. (This could give future informatic ethicists a nightmare: when do you have the right to turn off the machine?) Finally, such a machine would have the sorts of problems of reliability from which we suffer. Its choices would be at the mercy of its unique and perhaps idiosyncratic experiences. And its conclusions, based on that experience and its own internal commitments, might not be our own.

Other Perspectives on Affective Structuring

Not only is this fundamental inscription of the human body into the cognitive realm consistent with (small branches of) neuroscience, cognitive science, and computer science, it also reflects important currents in humanistic thought. A theme running through this book has been that the humanities and neuroscience share more common ground than either suspects. As explored in Chapter 2, the general model of a microstructure of parallel distributed representation leads into the macrostructures of semiology. Here, the role of affect in structuring memory provides a neuroscientific corollary to three important traditions in philosophical speculation on the nature of our relationship to the world. The oldest and most influential is part of the Buddhist concept of dependent co-arising. The next is a strain of Chinese thought that found clear articulation in the writings of Wang Yangming in the 15th century. Finally, neuroscience leads us to Heidegger.[38]

In the story of the night of Gautama's Awakening, Gautama progresses to increasingly deep levels of meditation and insight.[39] First he becomes aware of all his former incarnations. Then he becomes aware of all the former incarnations of all sentient beings. Finally he sees into the grand pattern revealed in the universal process of *karma* ("action"), and he sees a way out. In his enlightenment, he comes to understand the cycle of "dependent co-arising" (*paṭicca samuppāda*), the logic behind the so-called Wheel of Life. This is a twelve-step cycle of ignorance and craving through which one comes to believe in the existence of the self and of things. The account is based on earlier Indian technical classifications of the nature and structure of experience. Conditioned experience belongs to the realm of *nāma-rūpa*, Name and Form, mental and physical phenomena. This realm can be

---

[38] R. A. Brooks has a great sub-head, "It isn't German philosophy," for explaining the implications of his model for intelligent robots in "Intelligence without representation" in Andy Clark and Josefa Toribio, ed., *Cognitive Architectures in Artificial Intelligence*, p. 251.

[39] Many scholars discuss Gautama's enlightenment and the role of dependent co-arising in early Buddhist thought. A particularly accessible account is in Richard H. Robinson and

analyzed both in terms of the media of perceptions and in terms of the self that perceives. The media of perception are the six sense-bases: the usual five (belonging to Form) plus their objects and a "mind-organ" (belonging to Name) that perceives mental objects like memory, thought, imagination, and the input from the five senses plus all the objects that this mind-organ can perceive. The "self" is a collection of Five Aggregates (*khanda,* Sanskrit *skanda*): form, feelings (of pain or pleasure), perception (classifying and labeling), constructing activities (thoughts, intentions, etc.), and discriminative consciousness. (The last four aggregates belong to the category of Name.) Gautama's concern was why we take the fleeting aggregations of experience to have real substance and how we come to believe the Five Aggregates to be a permanent self. His analysis goes:

1. Aging, dying, sorrow, and pain depend upon birth.
2. Birth depends upon becoming.
3. Becoming depends upon sustenance, a drawing out of the momentary from a clinging to the Five Aggregates.
4. Sustenance depends on craving for the end of pain and continuation of pleasure.
5. Craving depends on feelings of pain and pleasure.
6. Feelings depend on contact with objects of the phenomenal realm.
7. Contact depends on the six sense-bases (senses + objects).
8. The six sense-bases depend on "Name and Form" (mind + body).
9. "Name and Form" depend on discriminative consciousness of the senses.
10. Consciousness depends on mental Constructing Activities.
11. Constructing Activities depend on Ignorance (of the Four Noble Truths).

Willard L. Johnson, *The Buddhist Religion: A Historical Introduction* 4th ed. (Belmont, CA: Wadsworth, 1997), pp. 9-29.

12. Ignorance is the root.

The arguments and terminology are complex. However, the point to be made here is the idea of objects coming into being *as objects* of mental awareness—rather than as aspects of the flow of sensation—through cravings associated with feelings of pain and pleasure. This Buddhist account has been successful for the simple yet profound reason that it is true to human experience. Its analysis of human consciousness has engendered long-enduring meditative traditions throughout South, Southeast, and East Asia. Buddhist practice has soteriological commitments that leave the phenomenology of experience behind and to which we can remain indifferent. Similarly, one can remain agnostic about the existence of the *dharmakaya* ("body of the Law") or *tathagathagarbha* ("womb of suchness") that ground these traditions. Still, Gautama's insight into desire, selfhood, perception and consciousness—so very different from Western formulations—have been the guides that led Buddhists through the ages to accumulate vast experience with exploring the mind's engagement with the world and with selfhood.

This Buddhist philosophical psychology also influenced a strand of Chinese Confucian speculation that culminated in the teachings of Wang Yangming. Wang concluded that phenomena 物 are where intention 意 is lodged. In his youth, he tried to understand the Principle 理 that underlies all formed objects by meditating on bamboo. The argument, derived from Zhu Xi, the great synthesizer of the Learning of the Way, was that one Principle inheres in all things. One comes to apprehend this all-pervading Principle by first understanding part of it in one object, then more in another object, and so on until one achieves a breakthrough of realizing the coherence running through all the particulars. Wang's intense meditation on the bamboo only made him fall ill, so he abandoned the model. His response to Zhu Xi's metaphysics parallels Xun Qing's reaction to Zhuang Zi: what humans know is the human realm. Whatever else there may be in the world, objects become objects for *human* experience only through their participation in the logic of human intentionality.

Heidegger introduces into Western European thought a similar approach of fundamentally linking objects—in their very existence—with human intentionality. Heidegger built upon Husserl's focus on intentionality to make readiness-to-hand central to *dasein*'s engagement with the world. Seeking to escape the Kantian impasse of the inaccessibility of the world, Heidegger sought to reframe the question of being. He recast phenomenology as an analysis of the manner in which authentic existence discloses itself. This most authentic existence is *dasein*, "being there." Heidegger's exploration of *dasein* in *Being and Time* is long and intricate. What concerns us here is not *dasein* in itself but Heidegger's analysis of *dasein* as a form of being-in-the-world. He seeks to understand how things in the world are fundamentally accessible for self-disclosure to *dasein*. His answer is in the idea of entities being ready-to-hand. They exist in their status as equipment for human use. I remember being appalled at this formulation when I first read *Being and Time*. Not only is this relationship of instrumentality deeply constricting, but Heidegger further ascribes this readiness-to-hand as the basic nature of "things in themselves:" phenomenology here pushes into ontology. After studying the

Chinese tradition, I can appreciate the elegance and insight of making readiness-to-hand central to the "world." In both the Chinese and Buddhist traditions, however, the phenomenology of experience serves to displace and preclude ontological questions. Foregrounding problems of access to experience usually explicitly makes existence that is outside the possibility of experience radically inscrutable, as it is for Hume, who insists that all we have outside the experiential is revelation and faith.

Despite such questions about the claims Heidegger makes, his account of the dynamics of engagement driven by concern remarkably parallels the model of cognitive structures built upon affective encounters. For example, Heidegger argues:

> When we concern ourselves with something, the entities which are most closely ready-to-hand may be met as something unusable, not properly adapted for the use we have decided upon. The tool turns out to be damaged, or the material unsuitable. In each of these cases *equipment* is here, ready-to-hand. We discover its unsuitability, however, not by looking at it and establishing its properties but rather by the circumspection of the dealings in which we use it. When its unusability is thus discovered, equipment becomes conspicuous.... (p. 102)
>
> In our concernful dealings, however, we not only come up against unusable things *within* what is ready-to-hand already, we also find things which are missing—which not only are not 'handy' but are not 'to hand' at all. Again, to miss something in this way amounts to coming across something un-ready-to-hand. when we notice what is un-ready-to-hand, that which IS ready to hand enters the mode of *obtrusiveness*.... (p. 103)
>
> The structure of the Being of what is ready-to-hand as equipment is determined by references and assignments. In a peculiar and obvious manner, the 'Things' which are closest to us are 'in themselves;' and they are encountered as 'in themselves' in the concern which makes use of them without ever noticing them explicitly—the concern which can come up against something unusable. When equipment cannot be used, this

> implies that the constitutive assignment of the "in-order-to" to a "toward-this" has been disturbed. The assignments themselves are not observed; they are rather 'there' when we concernfully submit ourselves to them. But when an assignment has been disturbed—when Something is unusable for some purpose—then the assignment becomes explicit. (p. 105)[40]

Heidegger's argument is that "concern" is fundamental to our Being-in-the-world. This concern underlies the manner of our engagement with things in the world and structures our attentiveness to the world. Heidegger's analysis of the ready-to-hand here is finally empirical, based on his experience in the world, rather than aprioristic. One can conceive of other modes of engagement with things: one could judge all objects by their aesthetic qualities, whether they form harmonious wholes within their particular contexts. One can also imagine the creature thus encountering the world dying very quickly, but this would be a fact of experience: it doesn't work. One of the central tasks of the integration of neuroscientific accounts of meaning with those in the humanities will be to tease out which claims asserted within the humanities are in fact empirical and which have the status of transcendental postulates about the conditions for the possibility of experience. Those that we can acknowledge as transcendentally given—like temporality, causality, the being-in-the-world of this particular body—are outside the realm of confirmation or disconfirmation by neuroscientific exploration. *Empirical* claims that just don't work will have to be abandoned as perhaps interesting but wrong. Hence it becomes important that Heidegger's analysis of the fundamental

---

[40] All citations are from Martin Heidegger, *Being and Time*, translated by John Macquarrie and Edward Robinson (New York: Harper and Rowe, 1962).

engagement with objects as ready-to-hand survives. Large-scale theorizing in the humanities will not simply evaporate like dawn mists. Yet where theories in the humanities attempt to account for empirical facts of experience, they will be challenged.

Turning Outward

Heidegger turns to the question of signs and reference in his analysis of readiness-to-hand and Being-in-the-world. This insight is important: if concern underlies our engagement with the world, it also shapes the final content of all systems of signification. If my reading of the neuroscientific data is right, the affective structure of memory—of all our knowledge—creates deep commitments in our understanding of the broader world. Heidegger aggregated these as "concern." Freud, in a different framework, called them forms of object cathexis. For both, however, the interplay of knowledge and perception cannot be neutral, for our commitments must be honored. This question of the emotional price of thought ramifies in all aspects of our lives: our intimate relationships, our neighborhoods, our jobs, our entire built environment of roadways and cities, as well as our political commitments and religious beliefs. We may try to keep these domains separate. In the end, however, they intersect in the structural nature of memory. They cannot help but impinge on one another. This human world, in which each of us has our own, often conflicting affective stakes, is not simple. It is not harmonious. Because the world "out there" is not made to our measure, it can threaten to pull apart the fragile self woven within the fabric of memory. We need to understand the tensions and fault lines. More deeply, we must explore the nature of the affective commitments we need to have answered by our world if we are to live well within it.