VALERIE GRAY HARDCASTLE

14 Neurobiology

Most of the issues found in traditional philosophy of science are recapitulated in the philosophy of neurobiology. In particular, philosophers of neurobiology worry about what counts as appropriate empirical justification for a theoretical claim, how to determine which level of organization is the correct one for a scientific explanation, what explanations should look like, whether all explanations will or should reduce to some primitives, and how what we learn about the mind/brain should affect larger social, economic, and political decisions.

In addition, philosophers of neurobiology concern themselves with some traditional aspects of philosophy of mind, including worrying how it is a brain can represent, if it does, and how and whether this representation ties to other notions of representation in cognitive science and beyond. It is difficult to focus on only one of these concerns to the exclusion of the rest. Most likely, as we come to understand some particular aspect of the practice of neurobiology, we will also understand others as well. In what follows, I discuss these areas of concern as they *differ* from traditional arguments. This discussion therefore should be laid on top of and be seen to complement the very rich literature in traditional philosophy of science and philosophy of mind.

1. THEORIES IN NEUROBIOLOGY

Brains are complicated and messy affairs; theories about brains share these same traits. The difficulty is that in order to make a simple generalization about how some aspect of the brain functions, scientists have to retreat to such a broad level of abstraction that their assertions become almost empirically meaningless. In order to

make their claims testable in a laboratory, neurobiologists have to confine their ideas to particular animals, to particular experimental tasks, or to both. As a result, scientists end up with neurobiological ''theories'' that contain two distinct parts: a broad statement of theoretical principle and a set of detailed descriptions of how that principle plays out across different animal models and experimental tasks. Though the detailed descriptions fall under the general principle, they are not immediately derivable from it. Moreover, the detailed descriptions can be incompatible with one another, though each will maintain a family resemblance with the others. (See Hardcastle 1995, Schaffner 1993, Suppe 1989 for similar approaches to understanding theories in the biological sciences.)

At a gross level, mammalian brains are remarkably similar to one another. Indeed, the central nervous system (CNS) in invertebrates is not all that different from the mammalian CNS either. There are innumerable homologous areas, cell types, neurotransmitters, peptides, chemical interactions, and so forth. However, once we scratch the surface of different animal brains, we do find important differences.

For example, consider the semicircular canal. All mammals have roughly the same five end organs in their ears to support their auditory and vestibular systems, and they all work to keep their lateral semicircular canals in their ears parallel to the horizontal plane relative to the Earth, for keeping it in that position allows them to get the best possible information about head position in space. (The lateral canal is maximally excitatory to a yaw (left-to-right) head motion; keeping the canal in line with the horizontal plane allows the organ to detect this motion with the greatest accuracy.) But rodents ambulate with their necks extended, which keeps their heads in an extreme dorsal position, while humans incline their heads about twenty degrees when walking naturally. In general, scientists can correlate the differences in the shape of the semicircular canals in the ear with skull shape and the position that an animal's head is normally in. (It is an unanswered but intriguing question whether scientists find the canal structures they do because different heads evolved to be oriented in different directions or whether animals naturally hold their heads in different positions because their semicircular canals evolved differently.)

For another example, consider the retina. There are striking differences between herbivores and predators in brain structure, for creatures who munch on grasses and trees require much less precise environmental information than those who hunt moving targets in order to survive. As a result, rodents have no foveae. To maintain visual fixation on a point, they move their necks, using what is known as the vestibular-colic response. The vestibular system in their ears tells them how their head is oriented and they use that information to reorient their heads in order to keep whatever object currently fascinates them in their line of sight.

In contrast, primates have foveae and they move their eyeballs to keep their target within the foveal area, using the vestibular-ocular response. This is a much more precise orienting mechanism, which allows them to move their eyes to compensate for changes in head position such that they can keep objects foveated for as long as they wish. For some indication of how important computing horizontal eye motion is to our brains, consider that the abducens (or sixth) nerve in humans, which controls horizontal eye abduction, feeds into one of the biggest motor nuclei in the brain stem. This ocular nucleus, which controls only one very tiny muscle, is only slightly smaller than the nucleus that controls all of our twenty or so facial muscles.

In more striking contrast still, bats do not maintain ocular position in the same fashion as the rest of the mammals. Because they fly and so have greater freedom to move in three-dimensional space, maintaining body position relative to horizontal is not an easy option. As a result, they use other sense organs, primarily hearing (the other half of the eighth nerve), to determine how their eyes should be oriented. Consequently, they need not rely on vestibular-ocular responses as we do, even though their bodies are equipped with such reflex machinery.

All of these anatomical and physiological differences are important when neurobiologists want to investigate something like the way the brain learns to compensate for damage to the vestibular pathways. What may seem as small and insignificant differences from a broad mammalian perspective become hugely important as scientists seek to understand the particular mechanisms of brain plasticity. Can they use animals with no foveae and a vestibular-colic response to learn about how foveated mammals recover their

vestibular-ocular response? More generally, they need to know how well particular animal models translate across the animal kingdom. Should they be allowed to generalize from experiments on a single species (or set of species) to the way nature functions?

In all vertebrates, a unilateral labyrinthectomy (UL), or a lesion of the labyrinthine structure in one ear, gives rise to two types of ocular motor disorders. There are static deficits, such as a bias toward looking toward the lesioned side when the head is not moving, and dynamic deficits, such as abnormal vestibular-ocular reflexes (VORs), which occur in response to head movements. In only two or three days after the UL procedure, the brain starts to compensate for its loss and the static deficits disappear. Since labyrinthine structures do not regenerate, and peripheral neurons continue to fire abnormally, whatever the brain is doing to recover has to be a central effect (Shaefer and Meyer 1974). Single neuron recordings from a variety of animals indicate that the vestibular nuclei (VN) on the same side of the brain as the lesion start to show normal resting rate activity as the brain learns to compensate for its injury. Scientists do believe that whatever the mechanism is, it is also likely to be a general procedure the brain uses for recovery, for they find similar resting rate recoveries of the sort they see with the ipsilateral vestibular nuclei after denervation in the lateral cuneate nucleus, the trigeminal nucleus, and the dorsal horn, among other areas. Exactly how an argument to defend these convictions is supposed to run, though, is unclear, since it is fairly easy to find significant differences in the ways organisms recover and compensate for vestibular damage across the animal kingdom. Frogs, for example, appear to rely on input from the intact labyrinth to regulate the resting activity of the vestibular nuclei. Mammals, however, do not. The recovery of their vestibular nuclei occurs independently of transcommissural inputs (Flohr et al. 1981). In addition, static symptoms follow different time courses in different animals. In rats, spontaneous nystagmus disappears within hours after UL, while in the rabbit and guinea pig, it persists for several weeks (Baarsma and Collewjin 1975, Sirkin, Precht, and Courjon 1984). In humans, it may continue in one form or another for several years (Fisch 1973).

There is a fundamental tension in neurobiology between the big picture story and what is found in particular instances. All sciences

strip away features of the real world when they devise their generalizations. Physicists neglect friction; economists neglect altruism; chemists neglect impurities, and so on. However, what neurobiologists are doing is not analogous to what physicists, economists, and chemists are doing. In each of the other cases, the scientists are simplifying the number of parameters they must consider in order to make useful and usable generalizations. In contrast, if neurobiologists were to ignore the differences they find across species, then they would have no data left to build a theory with. There is not anything left over, as it were, once neurobiologists neglect the anatomical and physiological differences found in the brain across the animal kingdom. There is much left over when physicists neglect friction; most of classical mechanics is left, in fact. In distinction to the other sciences, in neurobiology we find a tension between the general rules one hopes to find that describe all brains and the particular cases neurobiologists happen to study.

What should the scope and degree of generalization for neurobiological theories be? It appears scientists are confronted with an unpleasant choice. Either they settle for large-scale abstract generalizations, which gloss over what may be important differences, or they focus on the differences themselves, at the expense of what may be useful generalizations. However, despite appearances, they do not have an either-or proposition that they have to resolve before they can move ahead, for a proper neurobiological theory contains both general (and fairly vague) abstractions and detailed comments on specific anatomies and physiologies. The paradigm theories for physics are simple elegant equations with universal scope. Theories in neurobiology read more like a list of general principles plus detailed commentaries. One feels the tug of the dilemma posed above only if one is operating with a restricted notion of what a scientific theory is. Some theories are pithy and succinct; some are not. Neurobiological theories are of the latter sort.

In neurobiology, scientists start with a theoretical description at the most general level; it is what we might call the ''theoretical framework'' – the most general component in a neurobiological theory. Once they adopt the framework, they can make more precise hypotheses as a way of filling out their theoretical proposal. These claims can be local to particular phyla or species; hence, they are not intended to be a more detailed specification of the general

framework. Instead, they can be thought of as instances or examples of how the framework might be cashed out in particular cases.

However, it is not the case that all "fillings out" fail to generalize. For example, the dynamic symptoms of UL recover by using a different mechanism (probably). One hypothesis is that brains use a form of sensory substitution to compensate for the vestibular-ocular reflex (Berthoz 1988, Miles and Lisberger 1981). In this case, the brain uses internally generated signals from the visual or somatosensory systems to compensate for the vestibular loss. It may substitute computations from the saccadic or a visual pursuit system, both of which (probably) reconstruct head velocity internally, for vestibular throughputs. Data drawn from experiments on frogs, cats, and humans indicate that they all apparently use the same mechanism, though it remains to be seen whether this proposal will be applicable to all creatures and whether it can be generalized much beyond vestibular reflexes.

There are different degrees of abstraction one might use once some theoretical framework is adopted. Some discussions are going to be restricted to a single species, or maybe even one developmental stage within a species; others will include several unrelated species or phyla. Both are legitimate ways of cashing out the framework in particular instances, and neither is to be preferred to the other. The data will dictate the scope of subhypotheses, and scope can vary dramatically.

And this is how theories in neurobiology are built and structured. Detailed conclusions regarding a single animal model give rise to general theoretical principles. These principles inspire new experiments done with other animal models, which in turn give us new (and probably incompatible) details but also new general principles. These new principles then connect to other detailed studies using different protocols on still other animals, and so it goes.

At the end of the day, we have a set of related theoretical principles that jointly compose a general theoretical framework. And these principles are held together by the detailed data from a wide variety of animal studies. Neurobiology continually moves between two different ways of understanding the nervous system, first in broad and sweeping strokes and second by submergence in the minutiae. General theoretical principles arise out of and then feed back into particular animal experiments done on different animal

models. Because physiology differs across species, specific experimental protocols are appropriate only for specific models. Sometimes the data arising out of the different animal models and different experimental procedures overlap, but largely they do not. Hence, sometimes the detailed conclusions are consistent, but sometimes – a lot of the time – they are not. Neurobiologists weave a story through their animal models and experimental protocols united by a common guiding theoretical thread. They both find commonalties and define differences. And this entire exercise, taken together, fashions the theoretical structure of neurobiology.

2. THEORY-LADEN OBSERVATIONS AND SINGLE-CELL RECORDINGS

It is almost a truism in philosophy of science that there is no unproblematic distinction between observation and theory. That is, any scientific observations we make are filtered through and by a prior theoretical framework. Raw data become observations as we interpret the ways they either fit or belie our hypotheses (Woodward 1989). In short: what counts as an observation and how that observation functions in the business of science are heavily mediated by theory. In neurobiology in particular, it is easy to change the fundamental nature of our observations using accepted methodological techniques for manipulating raw data.

Good data allow scientists to discriminate among competing claims about phenomena (Suppe 1989). The particular practices of the scientific subfield tell us how to judge whether data are good. Sometimes these practices involve explicit calculations and formal derivations; sometimes they involve matters of personal judgment and skill. The cases in neurobiology involve both. In particular, it is a matter of personal judgment in the world of single-cell recording when to employ certain computational procedures. Different sorting techniques give rise to different data, so which techniques to employ is an important question. But that is also a question for which no easy or accepted answer exists.

It has only been during the last decade or so that neurobiologists have been able to record from the extracellular space of a large number of neurons from awake and behaving animals. When they record with an electrode near a single cell, they do pick up the cells'

action potentials, which are commonly believed to be the means through which neurons communicate. But they also record things that look like action potentials, but are instead voltages generated by axonal bundles or the field potentials from parallel sets of dendrites. Moreover – and especially if the microelectrode has a relatively low impedence – extracellular electrodes pick up signals from several neurons at the same time, recording from all the cells in a nearby area.

The problem is how to differentiate the contributions of the different cells and cell parts from a single lump recording. In many cases scientists only care about one particular action potential; the rest, from their perspective, is background noise. The challenge is how to separate what they want from all the electrical signals they do not want. The challenge is how to move from the recordings of the electrode's output to genuine, reliable, and informative data.

This challenge is compounded by the noisy nature of the recordings themselves. Some of the noise is mechanical and arises from the amplifiers themselves, but some is biological and comes from the neurons. Brain cells jitter around constantly (cf. Connors and Gutnick 1990). Neurons are not quiet until they fire off a spike, as some might think. Instead, they are always producing some activity or other. All in all, scientists have to cull their data from quite a din.

Finally, because the components in a recording are not constant, it is difficult to get a theoretical hook into the waveform. Spike shapes can change over time; electrodes can drift during recording session, changing position relative to the cells, which would also alter the spike amplitudes; and the electrical properties of electrodes vary with changes in tip condition or background impedance. Gathering data from single unit activity presents neurobiologists with a serious technical challenge.

In order to get usable data – to get genuine observations – out of what the electrode transmits, scientists must isolate each neuron's contributions to the recorded waveform. They first need to ascertain exactly how many neurons the recorded waveform reflects. How can they do this if they have a mess of overlapping action potentials and field potentials from a variety of cells at different and unknown distances from the electrode? This question becomes particularly vexing if other neurons in the same area have spikes of the same or a similar shape and amplitude.

There are several decomposition algorithms; however, each is imperfect (see Lewicki 1998). Each represents a different way to move from raw output to interpreted and interpretable data, giving scientists different ways of refining the waveforms they have recorded so that they can later interpret them. Each is what philosophers are thinking about when they talk about the theory-ladenness of data. Scientists have to choose what to do with their measurements in order to get something that can be scientifically useful. And the way they choose is determined by previously accepted theories.

But even with all these advanced sorting techniques, it is still hard to predict the number of neurons eliciting the data. Ideally, scientists would like to claim that one neuron generates each cluster of spikes we have identified, but if the cells are firing in complex bursts, or if there is nonstationary noise, or if the spike trains overlap one another, they cannot get accurate classifications at all. It is simply an unsolved problem how to decompose coincident action potentials with variable spike shapes. The best scientists can do at this point is guess. Their guesses are informed by their years of experience, but they are guesses nonetheless.

Guessing is not quite what philosophers of science have in mind when they talk about the theory-ladenness of observation. Their vision of creating data is one of more "scientific method." That is, to pull data out of the dial movements or changes in color or squiggles on the page, philosophers generally hold that there is some explicit background theory, devised in some other scientific inquiry, that scientists learn and then use to interpret what they are seeing or measuring as something useful for their studies. But there is a theoretical gap, as it were, in the move from raw recordings to genuine data, a gap scientists cannot fill with any sort of decision-making algorithm. The best scientists can do at this point is simply leap across the gap, on blind faith, with an eye to where they want to go.

Neurophysiology travels in a cognitive circle; scientists use what they know to cull data that support what they believe to be the case. Nevertheless, progress is not stymied. Knowledge accrues in small increments, with each set of single-cell recordings altering the face of what is known a wee bit at a time. Because neurobiological sorting techniques rely so heavily on previously accepted neurobiological hypotheses, there will likely never be an abrupt or dramatic conceptual revolution. But what is known can evolve slowly but

surely until the final resting position is quite far removed from the place where the investigation began.

3. LOCALIZATION AND REDUCTION

When scientists do single-unit recordings from a set of neurons they assume that they are examining a discrete system. They have been wildly successful using this strategy, identifying at least thirty-six different topographical visual processing areas in cortex (De Gelder 2000), differentiating the ''what'' from the ''where'' object processing streams (DeYoe and Van Essen 1988; Mishkin, Ungerleider, and Macko 1983), and distinguishing motion detection from contour calculations (Barinaga 1995), to name but a few examples. Maps of brain function are getting more and more complicated as more and more is learned about the processing capacities of individual cells. And all these projects are founded on the belief that brains have discrete processing streams that feed into one another.

Yet the most neurons scientists have ever been able to record from simultaneously are a few hundred; the most cells they can ever see summed local field potential activity over are a few thousand. But brain areas have hundreds of thousands of neurons, several orders of magnitude more than can be accessed at any given time. And these neurons are of different types, with different response properties and different interconnections with other cells, including other similar neurons, neurons with significantly different response properties, and cells of completely different types. Any conclusions scientists draw about the behavior of whatever cells they are recording from are going to be limited to very basic stimulus-response and correlation analyses of whatever neuronal subtype they are currently examining. Hence, the functionality they ascribe on the basis of these relatively meager sorts of experiments might be much more restricted than what the cells are actually doing.

They insert an electrode in or near a cell and then record what it does as they stimulate the animal in some fashion. They record from a cell in a vestibular nucleus and then move the animal's head about to see whether doing so changes the activity of the neuron. If it does, then they move it more or they move it differently and see how that changes the neuronal output. If it does not, then they either try another nearby cell or try some other stimulus. But what they

cannot do is record from all the neurons in some isolated area, even if the area is very small. And what they cannot do is test any given cell for all the known functional contributions of brain cells in general. So, what they conclude about any cell will only reflect the cells they have actually recorded from using stimuli they have actually used. This research strategy systematically underestimates when neurons actually respond and under what conditions.

Unit studies attempt to combine scores, hundreds, or even thousands of single-unit recordings to try to analyze the population. Theoretically, scientists could perhaps, in principle, delineate a nervous system region stereotaxically if it had reproducible correlations between afferent and efferent connections such that they could ultimately articulate the neurobiological function of the defined region. However, the likelihood of success for this type of study decreases as the complexity of the organism increases. Scientists can draw functional conclusions regarding the activities of neurons in the abdominal ganglia of *Aplysia*, or the segmental ganglia of the leech. But the architecture of these organisms' central nervous system is so different from mammals' that the probability of successfully using similar techniques for understanding humans is very low to zero.

In addition, the actual processing of information that goes on in those cells involves lots of different kinds of excitatory and inhibitory inputs from other areas in the brain stem, cerebellum, and cerebral cortex. The dorsal horn is supposed to integrate afferent nociceptive information from the periphery and pass it onto the motor system (among other things), but it does not do that segregated from the rest of the brain and what the brain is trying to do. It is integrating and passing as the organism is trying to pursue prey or flee from an enemy. Moreover, the brain regions that perform these tasks are often connected to the very area scientists are recording from. The motor system feeds back down into the dorsal horn, as do the thalamus and significant parts of cortex.

The impact on cognitive processing of such rampant feedback connections in the brain is only just now starting to be explored in neurobiological research, though exactly how to do this is a difficult question to answer. Of course, neurobiologists design their experiments keeping in mind the known anatomic connections between and among the relevant structures. At the same time, any actual

experimental observations of all the remote influences on the dorsal horn, for example, are impossible, despite however many individual neurons scientists record from. They simply do not have any way of conducting such extensive, invasive tests on live animals. At best, the particular influences assumed in any particular recording series are a matter of previously accepted gospel, dogma, and faith.

Ideally, neurobiologists try to conjoin single-cell studies with some sort of lesion experiment. Once scientists construct a general flowchart of the relevant structures based on anatomy experiments, and they have estimated normal unit behavior from a series of single-cell studies, they then try to knock out the hypothesized functions by placing lesions in otherwise normal animals. They run their experiments on the basis of the assumption that these lesions, placed in regions known to be important, will change the unit behavior of cells they are studying in a consistent fashion. If they witness such a change, they use that information to explain the relative functional contributions of the lesioned region to the cells under scrutiny. In other words, they are using lesion studies to try to derive a functional boxology for the brain, just as cognitive psychologists use reaction time distributions and error measurements to find one for the mind.

But there is a larger theoretical concern. What neurobiologists know, but generally ignore, is that any functional change in the central nervous system will lead to compensatory changes elsewhere (e.g., Merzenich et al. 1983). Because the brain is highly plastic, lesioning it in one place will provoke it to react in some fashion in some other place. Usually these other places are not components in the system or region being studied. But even if they are, neurobiologists ignore plasticity of the brain in favor of assuming a consistent functional alteration as caused by the lesion and nothing more. How are investigators supposed to evaluate some observed functional change when the difference they see might have been evoked by the brain's attempt to compensate for its loss and not by any specific deficit induced by the lesion?

The short answer is that they cannot if they are restricted to single-cell recordings and lesion studies. To answer this question we need to be able to see the activity of the entire brain at once and over time. The excitement over functional magnetic resonance imaging (fMRI) and other imaging techniques concerns exactly this

point: there is a way of looking at the activity of the whole brain at one time as tied to some cognitive activity or other. But magnetic resonance imaging, the best noninvasive recording device we currently have, only has a spatial resolution of about 0.1 millimeter and each scan samples a few seconds of activity. This imprecision forecloses the possibility of directly connecting single-cell activity – which operates three to four orders of magnitude smaller and faster – with larger brain activation patterns.

Methodological difficulties with current imaging techniques are now well known (Bechtel 2000, Cabeza and Nyberg 1997, 2000). Most center around the fact that MRI is a blood-oxygen-level-dependent (BOLD) measure, which can only be imperfectly correlated to brain activity. That is, MRI measures changes in the oxygenation level of blood; it does not directly measure anything about actual neuronal activity. Others are tied to the fact that the measure cannot differentiate between inhibitory and excitatory activation, and that can confound the way the images are interpreted. An area might be ''read'' as being part of the processing stream for some input, even though what is showing up in the MR analysis is that area actively damping down activity. A third set of limitations is tied to the sparse distribution of some processing systems. If a system – nociception in somatosensory cortex might be one example – is widely but sparsely distributed, then its activity level might never reach what is required for a BOLD measure to notice, given that cells surrounding the system are not activated by the particular stimulus in question.

The final set of concerns revolves around the subtraction method used in imaging studies to cull data. In brief, here is how that technique works. The experimenter picks two experimental conditions that she believes differ along only one dimension: they differ only with respect to the cognitive or perceptual process she wants to investigate. She then compares brain activity recorded under one condition with what happens in the second condition, looking for regions whose activity levels differ significantly across the two. These areas, she believes, constitute the neural substrates of the task under scrutiny. By subtracting one set of scans from the other, the hope is that one has removed activity not specifically relevant to the task at hand.

Let us set aside the fact that this method has no way of determining whether the differences found are actually tied to the

cognitive process and not to something else occurring concurrently but coincidentally. Let us also set aside the fact that some activity might be both relevant to the task at hand and relevant to the baseline task. Notice that how well the subtraction method will work depends upon the sensitivity of the measuring devices such that the worse the instrument is, the better the method seems to be for localization studies. Low signal-to-noise ratio (SNR) means that scientists will find only a few statistically significant differences across conditions. And these are the sorts of results neurobiologists need in order to bolster any claims identifying particular cognitive processes with discrete brain regions.

But as the imaging technology improves and the SNR increases, scientists see more and more sites that differ across trials. The more sites they get, the more it appears that essentially the entire brain is involved in each cognitive computation. And the more it appears that the entire brain is involved in each thought, the less it is they can justify any assumption of functional specificity in the brain. If we extrapolate from what scientists might learn with more sensitive measures, we can easily see that there will be a time when this whole approach just will not work anymore. Put in the harshest terms, brain imaging seems to support reductionism because the science is not very good yet.

For example, Brodman area 6 appears significantly active after subtraction in studies of phonetic speech processing, voluntary hand and arm movements, sight-reading of music, spatial working memory, recognizing facial emotions, binocular disparity, sequence learning, idiopathic dystonia, pain, itch, delayed response alternation, and category-specific knowledge, to list only a subset of activities in which it is significantly and differentially active. It could be the case that if scientists keep on doing the sort of subtraction studies that they currently are doing, then eventually they will find a unifying and pithy way to describe what premotor cortex is doing in humans. In this instance, neurobiology would be on the right track to determining brain function, but they still have a long way to go. But it could also be true that how a region functions depends heavily on the ''neural context.'' Its functional role in a cognitive economy depends on how it is connected to other areas and how those other areas are responding. (The function of these areas would also be dependent on their particular connectivity and

the current patterns of activation. And so it would go.) If this is correct, then searching for ''the'' function of particular areas is misguided, for different brain regions play different roles depending upon the cognitive tasks at hand.

4. NEUROETHICS

As progress is made into understanding how the brain works and how to influence brain functioning, serious ethical questions arise concerning how the medical, insurance, and governmental leaders should react to new information and possibilities (Marcus 2004). Neuroethics is a newly burgeoning area of research, with national attention only now being focused on the issues. Particular questions that philosophers of neurobiology will have to answer concern how and whether we should alter normal functioning brains, how and whether we should use brain technology to track individuals' social behavior, and how and whether what we learn about the brain changes the way we think of ourselves as human.

We know a lot about how memory works, and, more importantly, how it fails us. Seven basic ways in which memory can fail are decreasing accessibility to memories over time, lapses in attention, temporary inability to access stored information, false recognition of something, false memory of something, contamination of stored information by current beliefs, and remembering of items at inappropriate times. All of these processes are perfectly normal and occur in all of us at some time or another. Suppose we have some way of correcting some or all of these deficits. Should we? Or should we accept less-than-perfect memories as the way we are?

Neurobiologists are already tracking where and how moral decisions are made in the brain; they are also looking at brain differences between normal and sociopathic, psychopathic, and violently impulsive individuals. We know that such individuals respond to violent or otherwise disturbing situations with increased activity in the amygdala and decreased activity in the frontal lobes relative to normal individuals. We can now identify such trends in individuals before they actually commit any crime. Should we? And what should we do with such information once we have it?

If we come to believe that violence is biologically based, as are all other behavioral decisions, then what does this say about notions of

self or free will? How might this alter our court systems, since they operate under the assumption that one is guilty if one could have done otherwise in a situation but chose not to? Similar questions arise with gender differences in the brains. We know that female brains differ from males'. What effect, if any, should this fact have on our educational systems, our social expectations of gendered behavior, or men's and women's professional lives?

We are only beginning to confront these sorts of questions, as our technology is only beginning to allow us to understand and change the brain to any significant degree. As our knowledge of the mind/brain continues to increase exponentially, these and other similar questions will only become more pressing.