How Molecules Matter to Mental Computation

Paul Thagard†‡
University of Waterloo

Almost all computational models of the mind and brain ignore details about neurotransmitters, hormones, and other molecules. The neglect of neurochemistry in cognitive science would be appropriate if the computational properties of brains relevant to explaining mental functioning were in fact electrical rather than chemical. But there is considerable evidence that chemical complexity really does matter to brain computation, including the role of proteins in intracellular computation, the operations of synapses and neurotransmitters, and the effects of neuromodulators such as hormones. Neurochemical computation has implications for understanding emotions, cognition, and artificial intelligence.

1. Introduction. The functioning of brains in humans and other animals involves dozens of chemical messengers, including neurotransmitters, hormones, and other molecules. Yet almost all computational models of the mind and brain ignore molecular details. Symbolic models such as those based on production rules abstract entirely from neurological details (e.g. Anderson 1993; Newell 1990). Neural-network computational models typically treat neuronal processing as an electrical phenomenon in which the firing of one neuron affects the firing of all neurons connected to it by excitatory and inhibitory links (e.g. Churchland and Sejnowski 1992; Eliasmith and Anderson forthcoming; Levine 2000; Parks, Levine, and Long 1998; Rumelhart and McClelland 1986; see also such journals as *Cognitive Science, Neural Computing, Neural Networks,* and *Neurocom-*

†Send requests for reprints to the author, Philosophy Department, University of Waterloo, Waterloo, ON, Canada N2L 3G1; e-mail: pthagard@uwaterloo.ca; web: http://cogsci.uwaterloo.ca.

‡I am grateful to Baljinder Sahdra and Zhu Jing for helpful suggestions, and especially to Chris Eliasmith for skeptical comments. This research is supported by the Natural Sciences and Engineering Research Council of Canada.

Philosophy of Science, 69 (September 2002) pp. 429–446. 0031-8248/2002/6903-0000$10.00
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

puting). The role of neurotransmitters and other molecules in determining this electrical activity is rarely discussed.

The neglect of neurochemistry in cognitive science would be appropriate if the computational properties of brains relevant to explaining mental functioning were in fact electrical rather than chemical. But there is considerable evidence that chemical complexity really does matter to brain computation. I will review that evidence by discussing the role of proteins in intracellular computation, the operations of synapses and neurotransmitters, and the effects of neuromodulators such as hormones. Attending to the ways in which the brain is a chemical as well as an electrical computer provides a qualitatively different view of mental computation than is found in traditional symbolic and connectionist models. I conclude with a discussion of the implications of neurochemical computation for issues involving emotions, cognition, and artificial intelligence. First some general remarks are needed concerning the explanatory functions of computational models of mind.

2. Modeling the Mind. During the 1930s, Alan Turing and others produced rigorous mathematical accounts of computation, and in the 1940s the first general digital computers were built. The development of the theory and practice of computation had a huge impact on psychology and the philosophy of mind, because it showed how thought could plausibly be construed as mechanical. Psychologists such as George Miller and philosophers such as Hilary Putnam recognized the computational construal of mind as a powerful alternative to behaviorist ideas that had tried to make the mind go away. Allan Newell and Herbert Simon and other researchers began to produce computer programs that model intelligent behavior.

Abstract models of computation include the Turing machine, which is an imaginary device consisting of a tape with squares that contain a 0 or 1 and a head that can move from square to square. A table of very simple instructions completely determines the movements and reading and writing behavior of the head. The Turing machine, and mathematically equivalent abstractions such as recursive function theory, are very useful for clarifying what computation is. But they play no direct role in explaining particular mental functions. In order to explain particular kinds of mental abilities such as problem solving and language use, researchers develop specific kinds of computational models that posit mental representations, such as rules, and computational procedures, such as forward chaining, that operate on those rules. Rule-based systems are much better cognitive models than Turing machines because they concretely describe mechanisms that can replicate mental behavior. As far as abstract computational power goes, rule-based systems are no more powerful than a Turing ma-

https://doi.org/10.1086/342452 Published online by Cambridge University Press

chine, but they are much closer to capturing the mechanisms that underlie cognitive functions.

Besides rules, many cognitive scientists espouse alternative or complementary ways of modeling the mind, involving such representations as concepts, mental models, analogies, visual imagery, and artificial neural networks (see Thagard 1996 for a concise survey). In particular, artificial neural networks have the same abstract computational power as Turing machines and rule-based systems, but they are advocated by many researchers because they implement structures and procedures that seem to capture more closely the operations of the brain. For example, the brain uses distributed representations in which symbolic information is represented collectively by numerous simple neuronal elements, and uses massively parallel computations to draw inferences. Neural networks can be used to implement rule-based systems, but they also can support modes of computing qualitatively different from those in rule-based systems.

Most cognitive models using artificial neural networks describe the behavior of neurons by a parameter called *activation,* which represents the firing rate of the neuron, i.e. the rate at which it sends an electrical signal to other neurons. Recent models have more sophisticated dynamics, describing not only the rate of firing but the pattern of firing. Consider, for example, a neuron that fires 5 times, with the firing state represented by a 1 and the resting state represented by a 0. The firing pattern 10101 and the pattern 11100 both show the same rate of activation (firing 3 times out of 5), but they can represent very different neuronal behaviors. Neural networks that take into account such firing patterns are called *spiking* or *pulse* networks, and they have computational advantages over networks that only use rate codes. For example, there are functions that can be computed by a single spiking neuron whose computation would require many traditional rate-coding neurons (Maass and Bishop 1999, 79). Moreover, spiking neurons have psychologically important qualitative properties such as becoming synchronized with each other, and neural synchrony has been proposed as a crucial ingredient in inference, analogy, and consciousness (Engel et al. 1999; Hummel and Holyoak 1998; Shastri and Ajjanagadde 1993). Thus spiking neural networks provide a promising new approach to computational modeling of the brain.

I have gone into this brief review of cognitive modeling to indicate the form of argument that I want to develop. Just as rule-based models capture aspects of cognition that Turing machines do not address, and just as neural networks capture aspects of cognition that rule-based systems do not address, and just as spiking neural networks capture aspects of cognition that rate-coded neural networks do not address, so chemical neural networks have the potential to illuminate aspects of thinking that purely electrical neural networks do not adequately address. In order to

https://doi.org/10.1086/342452 Published online by Cambridge University Press

provide a useful supplement to existing computational accounts of mind, a new account must show that it has quantitative and qualitative advantages over old models, suggesting mechanisms of mental computing that are more powerful and more biologically and psychologically natural than in previous models. My task is to show that such advantages are to be found in chemical neural networks that explicitly recognize molecular mechanisms.

I do not mean to suggest that molecular models should supersede existing ones. Models are like maps, intended to correspond to reality but varying greatly in the level of detail that is useful for different purposes. To determine that Italy is south of Switzerland, a large scale map of the world is appropriate, whereas a much more detailed map is better for hiking in the Alps. Similarly, there are aspects of mental computing that are conveniently and accurately describable by rule-based systems and traditional electrical neural networks, but there are also aspects for which it is explanatorily useful to move down to the molecular level.

3. Proteins and Cells. How neurons and neural networks can perform computations is well understood. Each neuron receives and integrates electrical signals from other neurons, then passes its own signal on to other neurons to excite or inhibit their signaling. Neural networks are Turing complete, in that they can compute any function that a Turing machine can, and more importantly they can behave in ways that account for human cognitive functions. Only recently, however, have the computational capabilities of non-neuronal cells been appreciated.

The human body contains trillions of cells, and a typical cell contains around a billion protein molecules, with about 10,000 different kinds of protein in each cell (Lodish et al. 2000). The outer membranes of cells have *receptors,* which are proteins that bind signaling molecules circulating outside the cells. The receipt of a signaling molecule by a receptor activates *signal-transduction* proteins within the cell that initiate chemical reactions affected by *enzymes,* which are proteins that accelerate reactions involving small molecules. The chemical pathways within a cell can lead to diverse results, including cell division producing new cells, cell death, and the production of new signaling molecules that are expelled from the cell and then circulate to be bound by the receptors of other cells. For example, when the hormone epinephrine (also known as adrenaline) is produced by the adrenal gland in response to fright or heavy exercise, it circulates through the blood stream and binds to cells with appropriate receptors. These include liver cells that are stimulated to emit glucose into the blood stream, and heart muscle cells that increase the heart's contraction rate and the supply of blood to the tissues. The result is an increase in available energy for major motor muscles.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

We can think of individual cells, whether neurons or not, as computers that have inputs in the form of molecules that bind to receptor proteins, outputs in the form of molecules emitted from the cells, and internal processes carried out by chemical reactions involving proteins (Gross 1998). Proteins can function as on-off switches, for example by the process of *phosphorylation* in which proteins are modified by adding groups of atoms including phosphorus. Signals within a cell can be rapidly amplified by enzymes that can each activate hundreds of molecules in the next stage of processing. Molecular computing within the cell is massively parallel, in that many receptors can simultaneously initiate many chemical reactions that proceed concurrently in the billion or so proteins in the cell.

Multi-cellular computing also exhibits massive parallelism as cells independently receive and send signals to each other. There are three types of signaling by secreted molecules (Lodish et al. 2000, ch. 20). In *autocrine* signaling, a cell signals itself by secreting molecules that bind to its own receptors. For example, cells often secrete growth factors that stimulate their own proliferation. In *paracrine* signaling, a secretory cell signals an adjacent cell that has receptors for the secreted molecules. Neuronal signaling is paracrine, with neurotransmitters as the molecular signals, but there are also other kinds of paracrine signaling involved in cellular communication. Adjacent cells can also communicate with each other more directly than via secretions, by means of the attachment proteins that enable cells to adhere to each other and form tissues. The third type of signaling by secreted molecules is *endocrine,* in which a cell secretes a molecule, called a *hormone,* that travels through blood vessels to be received by distant target cells that may be several meters away. The computational functions of hormones are discussed in Section 5.

Is describing proteins and cells as performing computations a stretched metaphor that violates the mathematically precise notion of computation developed by Turing and others? Not at all, for there are several recent mathematical and experimental results that show that molecular processing is computational in the most rigorous sense. Magnasco (1997) proved that chemical kinetics is Turing universal in that the operations of a Turing machine can be carried out by chemical reactions. Bray (1995) showed how protein molecules can function as computational elements in living cells and can even be trained like a neural network. Adleman (1994) demonstrated that a hard combinatorial problem in computer science could be solved by molecular computation involving strands of DNA. DNA can provide cells with a kind of permanent memory, whereas protein operations serve to process information. Thus the description of cells and proteins as carrying out computations is more than metaphorical, and therefore is potentially relevant to understanding mental computation. Whether it is actually relevant requires looking more closely at the behavior of neurons.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

4. Neurotransmitters.

4.1. Properties of Neurotransmitters. The last section discussed the signaling capabilities of cells in general, but was not meant to suggest that organs such as the liver have mental properties. Human minds depend on a particular kind of organ, the brain, which has billions of cells capable of interacting with each other in special ways. A typical neuron takes input from more than a thousand neurons, and provides output to thousands of others, via special connections called synapses. Some synapses are electrical, passing ions directly from one cell to another, but most are chemical, enabling neurons to excite or inhibit each other by means of neurotransmitters that pass from the presynaptic cell to the postsynaptic cell. Neurotransmitters are not the only chemicals that allow one neuron to influence another; the next section will discuss hormones and other molecules that modulate the effects of neurotransmitters. Human brain chemistry is fundamentally the same as that found in other vertebrates.

The most important neurotransmitters include: aspartic acid and glutamic acid, (excitatory), gamma-aminobutyric acid and glycine (inhibitory), epinephrine (also a hormone), acetylcholine, dopamine, norepinephrine, serotonin, histamine, neurotensin, and endorphins. Does the abundance of different neurotransmitters used by the brain matter to mental computation? One might argue that the only computational significance is in the excitatory and inhibitory behavior of synaptic connections, and that the particular chemicals involved in excitation and inhibition are largely irrelevant to how the brain computes. I propose, however, that the array of neurotransmitters makes both qualitative and quantitative differences to mental processing, affecting both its style and speed.

The computational operation of a neural network depends on three kinds of properties of the network. The first is the internal processing capability of the neurons in the network, which may vary depending on how much the neuron can do with the various inputs coming to it and how complex its outputs can be. Most models of artificial neural networks used in cognitive science have very simple processing power, enabling them to translate input activation into output activation. Spiking neural networks have greater internal processing power in that they can respond differently to different patterns of spikes coming into them, and they produce different output patterns of spiking behavior, not just a rate of activation. The discussion of the computational power of proteins in Section 3 showed that chemical neurons have still greater internal processing power than those found in artificial spiking networks, because the chemical reactions that occur within cells are qualitatively and quantitatively different from the electrical integration and firing performed by spiking neurons.

The second key property is the topography of the network, which is the pattern of connectivity that enables one neuron to affect the firing of

https://doi.org/10.1086/342452 Published online by Cambridge University Press

another. In typical artificial neural networks, topography is determined by the excitatory and inhibitory links that connect neurons, but we shall see that chemical brains have a greatly enhanced topography. The third key kind of property is temporal. A neural network is a dynamic system that evolves over time, and how it evolves is very much affected by the order and rate of different occurrences in it. For example, artificial neural networks are sometimes synchronous, with all neurons having their activations updated at the same time, but it is more biologically natural when they are asynchronous. Real neurons are asynchronous and depend on temporal history in the form of the spike patterns that are input to them. Spiking neural networks thus have temporal properties that are different from rate-activation networks, although they are no different topographically from rate-activation networks. Chemical networks differ in all of these kinds of properties—internal processing, topographical, and temporal—from purely electrical networks. I will now discuss the topographic and temporal effects of neurotransmitters and neuromodulators.

4.2. Topographic Effects of Neurotransmitter Pathways. Neurotransmitters occur in specific nerve pathways in the brain (Brown 1994, 70). A pathway consists of connected neurons whose synapses all involve the transmission of the same chemical. For example, there are specific pathways for acetylcholine, dopamine, norepinephrine, and serotonin. Different pathways have different functions, for example the integration of movement by dopamine and the regulation of emotion by serotonin. Disruptions in these pathways can cause various mental illnesses, for example Parkinson's disease resulting from lack of dopamine, and depression resulting from lack of serotonin. Drugs can be used to treat illnesses by increasing or decreasing the amounts of neurotransmitters, as when MAO inhibitors are used to treat depression by increasing the availability of dopamine and serotonin.

The computational significance of neurotransmitter pathways is that they provide the brain with a kind of organization that is useful for accomplishing different functions. If a neuron could be connected to any other neuron, it would be difficult to orchestrate particular patterns of neuronal behavior. The brain requires cascades of activity, for example when perception of a dangerous object such as a snake leads to activation of fear centers in the amygdala and release of stress hormones. Neurotransmitters provide a coarse kind of wiring diagram, organizing general connections between areas of the brain that need to work together to produce appropriate reactions to different situations. Of course, the brain might have evolved with purely electrical pathways, but the fact is that the different kinds of neurotransmitters have served to establish patterns of connectivity that are important for its operation. Neurotransmitters

https://doi.org/10.1086/342452 Published online by Cambridge University Press

serve to restrict connectivity within the brain, but different kinds of chemical communication that enhance connectivity are discussed in Section 5.

4.3. Temporal Effects of Neurotransmitters. There are two types of synapse, the relatively rare electric synapse and the more common chemical synapse in which neurotransmitters are emitted from the vesicles of the presynaptic cell and bind to the receptors of the postsynaptic cell. The effects of chemical synapses are electrical, allowing ions to cross the membrane of the postsynaptic cell. But these effects are much slower than in an electric synapse, in which ions move directly from one neuron to another (Lodish et al. 2000, 943). Heart cells, for example, are electrically coupled, allowing groups of muscles cells to contract in synchrony. Signals are transmitted across electric synapses in a few microseconds, without the delay of .5 milliseconds found in chemical synapses.

Given the greater speed and reliability of electric synapses, it might seem puzzling why most synapses are chemical. According to Lodish et al. (2000, 942), chemical synapses have two important transmission advantages over electric ones. The first is *signal amplification,* for example when a single presynaptic neuron causes contraction of multiple muscle cells. The second is *signal computation,* in which a single neuron is affected by signals received at multiple excitatory and inhibitory synapses. "Each neuron is a tiny computer that averages all the receptor activations and electric disturbances on its membrane and makes a decision whether to trigger an action potential." (Lodish et al. 2000, 943). Thus chemical synapses, even though slower, allow for more flexible kinds of computation.

In chemical synapses, there are two classes of neurotransmitter that operate at vastly different speeds (Lodish et al. 2000, 939). Fast synapses, using receptors to which neurotransmitters bind and cause an immediate opening of ion channels, enable ions to cross the postsynaptic cell membrane in less than 2 milliseconds. In contrast, slow synapses are more indirect, requiring binding of a neurotransmitter to a receptor that initiates a chemical reaction that eventually affects ion conductance. Such postsynaptic responses are slower and longer lasting than those involving fast synapses, working on a scale of seconds rather than milliseconds.

Particular neurotransmitters can have special temporal properties. One kind of glutamate receptor, the NMDA receptor, functions as a coincidence detector (Lodish et al. 2000, 947). These receptors only open a channel if two conditions are met: glutamate must be bound and the membrane must be partly polarized by previous transmission. Thus the NMDA receptor makes possible a simple kind of learning. Galarreta and Hestrin (2001) found that networks of neurons that release gamma-aminobutyric acid (GABA) spike fast enough to be able to detect synchrony in input neurons. It used to be thought that each neuron could only release one

https://doi.org/10.1086/342452 Published online by Cambridge University Press

kind of neurotransmitter, but there is evidence that a neuron can release different transmitters and different amounts and combinations of transmitters at different times (Black 1991, 79). This complexity makes possible a degree of electrochemical encoding that has more variables than the activations and spike trains in purely electrical networks.

In sum, the different temporal properties of neurotransmitters enable them to operate on very different time scales, ranging from microseconds (electric synapses) to milliseconds (fast chemical synapses) to seconds (slow chemical synapses). We will see in the next section that even longer time effects are possible with hormones.

5. Neuromodulators. Brown (1994, 14) provides a useful taxonomy of *neuroregulators,* the chemicals that affect neuronal activity, dividing them into neurotransmitters and neuromodulators. As just described, neurotransmitters are released by neurons and act on other neurons via synapses. Neuromodulators, in contrast, can be released by non-neuronal cells as well as neuronal cells, and they act non-synaptically on both the presynaptic and postsynaptic cell to alter synthesis, storage, release, and uptake of neurotransmitters. Neuromodulators include hormones, which travel through the bloodstream, and nonhormone molecules that pass more directly between cells. The point of this section is to argue that the variety of neuromodulators used by the brain expands its computational abilities in ways that help to explain aspects of human thinking. Contrary to most computational models of neural networks, whether a neuron fires is not simply a function of its synaptic input. The influence of neuromodulators affects both the topographical and temporal properties of neural networks.

5.1. Topographical Effects of Neuromodulators. Neuromodulators dramatically change the causal structure of a neural network. Instead of having a kind of local causality, in which whether a neuron fires is determined only by the neurons that provide synaptic inputs to it, it becomes possible for neurons and other cells that are even meters away to affect firing. A neuron in one part of the brain such as the hypothalamus may fire and release a hormone that travels to a part of the body such as the adrenal glands, which stimulates the release of other hormones that then travel back to the brain and influence the firing of different neurons. Complex feedback loops can result, involving interactions between the neurotransmitter control of hormone release and the hormonal regulation of neurotransmitter release. These feedback loops can also involve the immune system, because brain cells also have receptors for cytokines, which are protein messengers produced by immune system cells such as macrophages.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

How do hormones affect neuronal firing? The internal processing of a neuron depends on a host of inputs, including neurotransmitters, hormones, and growth factors (Brown 1994, 200). All of these are *first messengers* that activate proteins to produce intracellular signals via second messengers such as the molecule cAMP, which then activate specific protein kinases (enzymes) that function as third messengers. The kinases phosphorylate proteins that act as fourth messengers to stimulate changes in membrane permeability and protein synthesis in the cell. Such changes influence the ability of the neuron to spike, and hence affect the rate and pattern with which it fires. The key point here is that whether a neuron fires and hence contributes to the computation performed by the neural network is not simply a function of neurons that provide synaptic inputs, but can also be affected by a host of other cells that produce hormones. Hence the topography of the brain is far more complex than is recognized by purely electrical models in which the inputs to artificial neurons are just activations and spike trains.

Hormonal chemical effects operate over long distances, but there are also non-synaptic connections between adjacent neurons. Cell adhesion molecules not only bind cells together to form tissues, they also carry signals between cells that can affect their development (Crossin and Krushel 2000). Song et al. (1999) discovered Neuroligin, a synaptic cell adhesion molecule that not only enables neurons to establish synaptic connections with each other, but also allows for direct signaling from the postsynaptic neuron back to the presynaptic one. Such retrograde signaling is thought to be important for learning. Other molecular mechanisms for retrograde signaling have been identified. The postsynaptic neuron can also send chemical signals back to the presynaptic neuron by means of gases such as nitric oxide and carbon monoxide, or by peptide hormones (Lodish et al. 2000, 915). Nitric oxide is a small molecule that can easily diffuse to affect many neurons, greatly expanding the computational topography of neural networks beyond synaptic connections. Koch (1999, 462) conjectures that, because of the spread of nitric oxide: "the unit of synaptic plasticity might not be individual synapses, as assumed by neural network learning algorithms, but groups of adjacent synapses, making for a more robust, albeit less specific learning rule."

Neuronal firing is also affected by glial cells, which were formally thought to function only to hold neurons together. There are 10–50 times more glial cells in the brain than neurons, and glial cells affect both the formation of connections by nerve cells and their firing. A factor released by glial cells makes transmitting neurons release their chemical messengers more readily in response to an electrical signal (Pfrieger and Barres 1997). Stimulated glial cells release calcium that trigger surrounding glia to release calcium too, producing a spreading signal (Newman and Zahs 1998).

https://doi.org/10.1086/342452 Published online by Cambridge University Press

The calcium wave releases glutamate from the glial cells, which has a direct impact on the firing of the neurons in the vicinity.

In sum, there is evidence from the behavior of hormones, nitric oxide, and glial cells that the topography of brain networks is far more complex than is captured by electrical models based only on synaptic connections. Not surprisingly, the operation of non-synaptic chemical messengers also affects the temporal patterns of neurons.

5.2. Temporal Effects of Neuromodulators. Hormones can affect the firing rate of neurons (Brown 1994, 166f.). Gonadal hormones increase the electrical activity of some neurons and inhibit the activity of other neurons. For example, estrogen can modulate the release of dopamine and serotonin. Thus hormones can slow down or speed up neuronal firing.

Many neurons secrete neuropeptides such as endorphins and oxytocin. Unlike classical neurotransmitters, these molecules are released outside the synaptic zone and can have effects that last for hours or days (Lodish et al. 2000, 936). Thus the temporal effects of neuropeptides operate on a very different scale from the much briefer effects of neurotransmitter emission described in Section 4.2.

Thus a computational system that involves neuromodulators can be expected to have different temporal behaviors than one with neurotransmitters only, and we already saw in Section 4.2 that different neurotransmitters give rise to different temporal properties. Hence molecules matter for the temporal behavior of neural networks.

6. Emotional Cognition. My general argument to this point has been that there are reasons to expect that neurochemistry should matter to mental computation, but I have not shown any particular kinds of mental computation that are affected. There is little direct evidence that the highest-level mental computations involved in problem solving are tied to the influences of specific neurotransmitters and neuromodulators. However, there is substantial evidence that these neuroregulators are important for emotions, and there is also evidence that emotions greatly affect problem solving and learning. I will review these two bodies of evidence and conclude that even the most cognitive of mental functions are subject to neurochemical understanding. Chemistry has both positive and negative effects on emotions and problem solving.

6.1. Emotions and Neurochemistry. Panksepp (1993) provides a concise review of the neurochemical control of moods and emotions, including examples of how neurotransmitters are linked to particular emotions. Adminstration of glutamate, the most common excitatory neurotransmitter in the brain, can precipitate aggressive rage and fear responses. NMDA

https://doi.org/10.1086/342452 Published online by Cambridge University Press

receptor blockage in the amygdala can modulate extinction of fear behaviors. The inhibitory neurotransmitter GABA figures in the control of anxiety. Norepinephrine influences sensory arousal and becomes prominent in high-affect situations such as threat. Dopamine is associated with positive emotionality, and adenosine is a natural soporific that is blocked by weak mood enhancers such as caffeine.

Neuroregulators also play prominent roles in specific emotions. Corticotropin-releasing factor instigates a stress response that has a major impact on fear and anxiety. Oxytocin enhances maternal behavior as well as feelings of acceptance and social bonding, and it contributes to sexual gratification. Arginine vasopressin is under testosterone control and can provoke male aggression. Estrogen receptors in the brain are involved in female sexual behavior, aggression, and emotionality (Brown 1994, 154). Many other peptides also affect emotional behavior.

Additional evidence concerning neurochemical influences on mood and emotion comes from the medical effectiveness of drugs that target particular neurotransmitters (Panksepp 1998, 117). Depression can be treated both by drugs like Prozac that prolong the synaptic availability of neurotransmitters such as serotonin and dopamine and by drugs that inhibit the enzyme monoamine oxidase (MAO) that normally helps degrade neurotransmitters following release. Antipsychotic drugs used to treat schizophrenia generally dampen dopamine activity. Most antianxiety agents interact with a specific receptor that can facilitate GABA activity, whereas newer drugs reduce anxiety by interacting with serotonin receptors. A new generation of psychiatric medicines is being developed to deal with problems such as bulimia that may arise from imbalances in particular neuropeptides.

There is thus abundant reason to believe that understanding of human emotions will require attention to the effects of neuroregulators on thinking. It follows immediately that neurochemistry is relevant to understanding the nature of emotional consciousness. Feelings of happiness, sadness, fear, anger, disgust and so on emerge from brain activity by mechanisms not yet understood, but the diverse ways in which neurochemicals influence emotion suggest that it is unlikely that emotional consciousness emerges only from the electrical activities of the brain. I return to this topic in the discussion of artificial intelligence in Section 7.

6.2. Cognition. It might be argued that, even though chemical explanations are relevant to emotion, they have no bearing on central cognitive processes such as problem solving, learning, and decision making. However, there is increasing evidence in psychology and neuroscience that cognition and emotion are not separate systems and that emotion is an intrinsic part of human cognition (Dalgleish and Power 1999). Reviewing

https://doi.org/10.1086/342452 Published online by Cambridge University Press

this evidence would take a book in itself, but here I will only report a few salient examples of the cognitive impact of emotions.

Isen (1993) reviews an extensive literature on the impact of positive affect on decision making. The presence of positive feelings can cue positive materials in memory, making access to such thoughts easier. Positive but not negative emotion provides retrieval cues for situations relevant to a current problem. Positive affect also promotes creativity in problem solving and negotiation, and efficiency and thoroughness in decision making. People in whom positive affect have been induced are able to categorize material more flexibly and to see more similarities among items. Kunda (1999, 248) reports that mood manipulations by small gifts or pleasant music have been shown to influence a host of judgments, including assessment of one's own competence, one's general satisfaction of life, and evaluations of the quality of political leaders. Affect may also influence our cognitive strategies: people in a bad mood are more likely to use elaborate, systematic processing strategies. Happiness has been found to increase our reliance on social stereotypes, whereas sad people have reduced reliance on negative stereotypes. Thus basic cognitive functions such as categorization, problem solving and decision making are under emotional influence.

Ashby, Isen, and Turken (1999) have developed a neuropsychological theory of how positive affect influences cognition. They propose that positive affect is associated with increased brain dopamine levels that improve cognitive flexibility. Many readers of this article are familiar with the enhancement in problem solving ability brought about by caffeine, which blocks the inhibitory neurotransmitter adenosine (Brown 1996). In contrast, alcohol can disrupt mental functioning by inhibiting receptors for the excitatory neurotransmitter glutamate, including NMDA receptors important for learning. (More pleasantly, alcohol reduces anxiety by binding to GABA receptors and increasing their inhibitory function, while inducing euphoria through increased dopamine levels in the brain's reward centers and released endorphins.)

It might be thought that decision making would improve if emotions were removed from decisions, but the neurophysiological research of Damasio (1994) and his colleagues suggests that this is emphatically not the case. People who have brain damage that severs links between the most cognitive areas of the brain in the neocortex and the most emotionally important areas in the amygdala are very ineffective decision makers, even though their verbal and mathematical abilities are unaffected. Their problem is that they have lost the emotion-driven ability to make decisions on the basis of what really matters to them. Bechara et al. (1997) found that this disability also made it difficult for patients to learn a card playing task in which normal subjects unconsciously learned strategies that enabled them to avoid bad outcomes.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

This neurological research on the role of emotions in decision making fits well with recent psychological theories that find deficiencies in purely cognitive accounts of decision making. Loewenstein, Weber, Hsee, and Welch (2001) show that many psychological phenomena involving judgment and decision making under uncertainty can be accounted for by understanding peoples estimates of risk as inherently emotional. Similarly, Finucane et al. (2000) propose that human decisions are heavily affected by what they call the "affect heuristic." Legal and scientific thinking are also inherently emotional (Thagard 2002, forthcoming).

I have mentioned only a small part of the evidence that challenges the traditional psychological division between cognition and emotion and the ancient philosophical distinction between reason and passion. But it suffices for the purpose at hand, to show that the demonstrable relevance of neurochemistry to emotions carries over to cognition in general. If human cognition is mental computation, it is a kind of computation determined by the chemical as well as the electrical aspects of the brain. This conclusion has important implications for the prospects of developing intelligence in non-human computers.

7. Artificial Intelligence. Kurzweil (1999) and Moravec (1998) have predicted that artificial intelligence will be able to match human intelligence within a few decades. Their prediction is based on the exponential increase in processing speed of computer chips, which continues to double every 12–18 months as it has for decades. Kurzweil estimates the computing speed of the human brain to be around twenty million billion calculations per second, based on 100 billion neurons each with a thousand connections and the slow firing rate of 200 calculations per second. Assuming continued exponential increase in chip speed, digital computers will reach the twenty million billion calculations (on the magnitude of 10^{15}) per second mark around the year 2020.

However, the molecular chemistry of the brain suggests that this estimate of its computational power may be very misleading, both quantitatively and qualitatively. If we count the number of processors in the brain as not just the number of neurons in the brain, but the number of *proteins* in the brain, we get a figure of around a billion times 100 billion, or 10^{17}. Even if it is not legitimate to count each protein as a processor all by itself, it is still evident from the discussion in Section 3 that the number of computational elements in the brain is more than the 10^{11} or 10^{12} neurons. Moreover, the discussion of hormones and other neuroregulators discussed in Section 5 shows that the number of computationally relevant causal connections is far greater than the thousand or so synaptic connections per neuron. I do not know how to estimate the number of neurons with hormonal receptors that can be influenced by a single neuron that

https://doi.org/10.1086/342452 Published online by Cambridge University Press

secretes hormones or that activates glands which secrete hormones, but the number must be huge. If it is a million, and if every brain protein is viewed as a mini-processor, then the computational speed of the brain is on the order of 10^{23} calculations per second, far larger than the 10^{15} calculations per second that Kurzweil expects to be available by 2020, although less than where he expects computers to be by 2060. Thus quantitatively it appears that digital computers are much farther away from reaching the raw computational power of the human brain than Kurzweil and Moravec estimate.

Moreover, intelligence is not merely a matter of raw computational power, but requires that the computer have a sufficiently powerful program to produce the desired task. My Macintosh G4 laptop computer can calculate $2^{100,000}$ in a couple of seconds, the same amount of time in which I can only calculate 2^5, but the computer lacks the programming to be able to understand language and solve complex problems. Kurzweil and Moravec are aware that it is a daunting task to write the billions or trillions of lines of software that would be needed to enable the superfast computers of the future to approach human cognitive capabilities, but they blithely assume that evolutionary algorithms will allow computers to develop their own intelligent software. Evolutionary computation, which uses algorithms modeled in part on human genetics, is indeed a powerful means of developing new software (Koza 1992), but it is currently limited by the need for humans to provide the evolving programs with a criterion of fitness that the genetic algorithms serve to maximize. In humans, the evaluation of different states is provided by emotions, which direct us to what matters for our learning and problem solving. Computers currently lack such intrinsic, biologically-provided motivation, and so can be expected to have difficulties directing their problem solving in non-routine directions.

Perhaps software will be developed that does for computers what emotions do for us, but current computational research on emotions is very limited compared to the complexity of the human emotional system based on numerous neurotransmitters and neuromodulators. There is a current resurgence in AI of interest in emotion, which is however treated by researchers as a symbolic or electrical rather than a chemical phenomenon. The complexity of human emotions, based on looping interactions among neural, hormonal, and immune systems, may be too complex for people to figure out how to program and also too complex for a program created by humans to evolve.

This does not mean that computers of great intelligence in special areas will not be developed. It may be quantitatively and qualitatively difficult for AI to duplicate the human brain, but intelligent computers may be developed by other means, just as IBM managed to build the world's best

https://doi.org/10.1086/342452 Published online by Cambridge University Press

chess player by combining clever software with extraordinarily fast computer chips. But we should not expect a computer developed in this way to have all the mental capacities of humans, and we certainly should not expect it to have anything like human consciousness, which Section 6.1 suggested is intrinsically tied to human emotions and hence to our peculiar brain chemistry.

8. Conclusion. My arguments that neurochemistry matters to mental computation are not meant to show that computational models of the mind have to be at the molecular level. As I stated at the end of Section 2, models are like maps in that various levels of detail are useful for different purposes. Symbolic models of high-level inference and neural network models with and without spiking neurons have proven very useful in explaining many facets of cognition, and I have no doubt that they will continue to be useful. Cognitive science benefits from a combination of many different fields and methodologies, with different researchers attacking the problem of understanding mind and intelligence at different levels.

Without recommending abandonment of the techniques of computational modeling that have served cognitive science well, it is nevertheless evident that there are new possibilities for enhancing understanding of mind by working more at the molecular level. Consider, for example, the computational study of emergent properties of chemical pathways conducted by Bhalla and Iyengar (1999), including integration of signals across multiple time scales and self-sustaining feedback loops. It is possible that computational modeling of brain activity at the molecular level will discover additional emergent properties that are important for understanding some of the most currently intractable problems in cognitive science, such as the origins of emotional consciousness. Hence without abandoning traditional concerns and methods, it may be time for psychology and the philosophy of mind to become, like current biology and medicine, molecular.

REFERENCES

Adleman, Leonard. M. (1994), "Molecular Computation of Solutions to Combinatorial Problems", *Science* 266: 1021–1024.

Anderson, John. R. (1993), *Rules of the Mind.* Hillsdale, N.J.: Erlbaum.

Ashby, F. Gregory, Alice. M. Isen, and And U. Turken, (1999), "A Neuropsychological Theory of Positive Affect and Its Influence on Cognition", *Psychological Review* 106: 529–550.

Bechara, Antoine, Hanna Damasio, Daniel Tranel, and Antonio R. Damasio (1997), "Deciding Advantageously before Knowing the Advantageous Strategy", *Science* 275: 1293–1295.

Bhalla, Upindar S. and Ravi Iyengar (1999), "Emergent Properties of Networks of Biological Signaling Pathways", *Science* 283: 381–387.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

Black, Ira. B, (1991). *Information in the Brain: A Molecular Perspective.* Cambridge, Mass.: MIT Press.
Bray, Dennis (1995), "Protein Molecules as Computational Elements in Living Cells", *Nature* 376: 307–312.
Brown, Richard. E. (1994), *An Introduction to Neuroendocrinology.* Cambridge: Cambridge University Press.
Brown, Stephen, (1996), *Buzz: The Science and Lore of Alcohol and Caffeine.* New York: Penguin.
Churchland, Patricia. S. and Terrence Sejnowski (1992), *The Computational Brain.* Cambridge, Mass.: MIT Press.
Crossin, Kathryn. L. and Leslie A. Krushel (2000), "Cellular Signaling by Neural Cell Adhesion Molecules of the Immunoglobulin Superfamily", *Developmental Dynamics* 218: 260–279.
Dalgleish, Tim and Power, Michael. J. (eds.) (1999), *Handbook of Cognition and Emotion.* New York: Wiley.
Damasio, Antonio R. (1994), *Descartes' Error.* New York: G. P. Putnam's Sons.
Eliasmith, Chris and Charles H. Anderson (forthcoming), *Neural Engineering: Computation, Representation and Dynamics in Neurobiological Systems.* Cambridge, Mass.: MIT Press.
Engel, Andreas K., Pascal Fries, Pieter R. Roelfsema, Peter König, and Wolf Singer (1999), "Temporal Binding, Binocular Rivalry, and Consciousness", *Consciousness and Cognition* 8: 128–151.
Finucane, Melissa L., Ali Alhakami, Paul Slovic, and Stephen M. Johnson (2000), "The Affect Heuristic in Judgments of Risks and Benefits", *Journal of Behavioral Decision Making* 13: 1–17.
Galarreta, Mario and Shaul Hestrin (2001), "Spike Transmission and Synchrony Detection in Networks of GABAergic Interneurons", *Science* 292: 2295–2299.
Gross, Michael (1998), "Molecular Computation", in Tino Gramss et al. (eds.), *Non-standard Computation.* Weinheim: Wiley-VCH, 15–58.
Hummel, John E. and Keith J. Holyoak (1998), "Distributed Representations of Structure: A Theory of Analogical Access and Mapping", *Psychological Review* 104: 427–466.
Isen, Alice M. (1993), "Positive Affect and Decision Making", in Michael Lewis and Jeannette M. Haviland (eds.), *Handbook of Emotions.* New York: Guilford Press, 261–277.
Koch, Christof (1999), *Biophysics of Computation: Information Processing in Single Neurons.* New York: Oxford University Press.
Koza, John R. (1992), *Genetic Programming.* Cambridge, Mass.: MIT Press.
Kunda, Ziva (1999), *Social Cognition.* Cambridge, Mass.: MIT Press.
Kurzweil, Ray (1999), *The Age of Spiritual Machines.* New York: Viking.
Levine, Daniel S. (2000), *Introduction to Neural and Cognitive Modeling,* 2nd ed. Mahwah, NJ: Erlbaum.
Lodish, Harvey et. al. (2000), *Molecular Cell Biology,* 4th ed. New York: W. H. Freeman.
Loewenstein, George F., Elke U. Weber, Christopher K. Hsee, and Ned Welch (2001), "Risk as Feelings", *Psychological Bulletin* 127: 267–286.
Maass, Wolfgang and Christopher. M. Bishop (eds.) (1999), *Pulsed Neural Networks.* Cambridge, Mass.: MIT Press.
Magnasco, Marcelo O. (1997), "Chemical Kinetics Is Turing Universal", *Physical Review Letters* 78: 1190–1193.
Moravec, Hans (1998), *Robot: Mere Machine to Transcendent Mind.* Oxford: Oxford University Press.
Newell, Allen (1990), *Unified Theories of Cognition.* Cambridge, Mass.: Harvard University Press.
Newman, Eric A. and Kathleen R. Zahs (1998), "Modulation of Neuronal Activity by Glial Cells in the Retina", *Journal of Neuroscience* 18: 4022–4028.
Panksepp, Jaak (1993), "Neurochemical Control of Moods and Emotions: Amino Acids to Neuropeptides", in Michael Lewis and Jeannette M. Haviland (eds.), *Handbook of Emotions.* New York: Guilford Press, 87–107.
——— (1998), *Affective Neuroscience: The Foundations of Human and Animal Emotions.* Oxford: Oxford University Press.

https://doi.org/10.1086/342452 Published online by Cambridge University Press

Parks, Randolph W., Daniel S. Levine, and Debra L. Long (eds.) (1998), *Fundamentals of Neural Network Modeling.* Cambridge, Mass.: MIT Press.

Pfrieger, Frank. W. and Ben A. Barres (1997), "Synaptic Efficacy Enhanced by Glial Cells in Vitro", *Science* 277: 1684–1687.

Rumelhart, David E. and Jay L. McClelland (eds.) (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Cambridge Mass.: MIT Press/Bradford Books.

Shastri, Lokendra and Venkat Ajjanagadde (1993), "From Simple Associations to Systematic Reasoning: A Connectionist Representation of Rules, Variables, and Dynamic Bindings", *Behavioral and Brain Sciences* 16: 417–494.

Song, Ji-Ying et. al. (1999), "Neuroligin 1 is a Postsynaptic Cell-adhesion Molecule of Excitatory Synapses", *Proceedings of the National Academy of Sciences* 96: 1100–1105.

Thagard, P. (1996), *Mind: Introduction to Cognitive Science.* Cambridge, Mass.: MIT Press.

——— (2002), "The Passionate Scientist: Emotion in Scientific Cognition", in Peter Carruthers, Stephen Stich, and Michael Siegal (eds.), *The Cognitive Basis of Science.* Cambridge: Cambridge University Press 235–250.

——— (forthcoming), "Why Wasn't O. J. Convicted? Emotional Coherence in Legal Inference", *Cognition and Emotion.*

https://doi.org/10.1086/342452 Published online by Cambridge University Press