

Lessons for cognitive science from neurogenomics

- 1. From developmental molecular biology to neurogenomics**
- 2. More than you wanted to know about short term and long term implicit memory**
- 3. How are explicit memories stored?**
- 4. How the brain recalls memories**
- 5. Each explicit memory is just a lot of implicit memories**
- 6. Is 'knowledge how' computable?**
- 7. Computationalism and neuroscience**

Assume that cognitive science is the research program of computationalism, acting on a wager about how the mind works inspired largely by philosophical arguments. The wager is roughly that thinking shows powers of creativity, systematicity and productivity that can only be explained as the purely syntactical processing of symbolic representations. Optimists about the wager hope that computational neuroscience will make good this hypotheses about thought that cognitive science advances to explain the powers of thought, by showing how the wet stuff substantiates, explains, systematizes and improves these hypotheses. Pessimists about cognitive science will expect cognitive neuroscience to show why findings about the brain limit the prospects of the research program computationalism motivates. In this paper I argue that it is not too early to begin to take sides on whether the optimists or the pessimists are going to turn out to be correct. I will argue that what is already known in cognitive neuroscience, and has been known for a decade or so, is enough to significantly effect our expectations of the shape that cognitive science will take, and that shape will be significantly different from what computationalism leads us to expect.

In particular, I will argue that work by some of the leading figures in the field, employing the resources of neurogenomics, have already provided good grounds to be pessimistic about the representations to which a computational theory of mind is committed, but to be optimistic about the syntactical character of processes of thinking and reasoning in the brain. Many philosophers will hold that the most pressing philosophical problem of computationalism is to explain how propositional

representation in the brain is even possible, let alone actual, as much cognitive neuroscience presumes. There is no general agreement on how physical states can embody representations of distinct and different propositions as we individuate them and express their differences in speech and writing—that is, in media with “derived intentionality” [Searle, 1980]. The failure to provide an adequate account of representation is however easily explained by findings in neuroscience which I shall expound. These findings strongly suggest that brain states cannot represent distinct propositions as we apparently individuate them in spoken and written language, but that brain states can come asymptotically close to doing so. If this is correct then propositional content is explain away as a natural and perhaps useful heuristic foisted upon us in ordinary life by the closeness of the approximation to it that the brain attains. Whether this heuristic assumption is one that a cognitive science should also exploit is another matter on which pessimism may be in order. On the other hand, if the study of the brain reveals the same sorts of processes that molecular biology has uncovered elsewhere in development and operation of bodily processes, computationalism’s commitment to discovering a syntax of thought may already be guaranteed. The only question that remains will be whether it will be much like a “language of thought” that some computationalists urge cognitive science to look for.

1. From developmental molecular biology to neurogenomics

The work by neuroscientists expounded here deals with neural mechanisms of memory-storage and memory recall that originated in the work of Eric Kandel but now includes a large number of other laboratories. It has accelerated over the last decade largely due to advances in genomics and proteonomics¹, methods that enabled experimenters first to identify the genes whose products are differentially expressed in the various tissues of the body, then to switch them on and off, or even selectively increase and decrease the levels of their expression, and eventually to assay for individual proteins, to locate their concentrations in parts of cells, identify their structure and quantify their enzymatic roles in chemical reactions,

¹ The study of the proteins the genes encode and of their interrelation, on the model of genomics—the study of the genome; whence proteonomics—the study of the proteome

To understand its significance requires a quick run-through of developments in molecular developmental biology. This discipline has now provided the details of embryological development across the phylogenetic spectrum from the *Drosophila* to mammals. What the details show is that across this spectrum, development is the result of the switching on and off of pretty much the same structural genes at different times by a diversity of regulatory genes, and that it is differences among these regulatory genes that results in the obvious differentiation on organismal body-plans, and behavioral capacities. This discovery reduces the surprise associated with the discovery that humans and mice have roughly the same number of genes, and only small multiples of the number of genes in much simpler systems such as fruit flies, the worm, *C.elegans*, and the sea slug, *Aplasia*. Indeed, so far as development is concerned, one set of genes of central importance, the so-called *Hox* genes, differ across this spectrum largely in their copy numbers, so that increasing complexity of animals is the result of successive duplication and mutation of the same original set of regulatory genes. In the case of *Drosophila* molecular biologists have been able to identify a literal program of development (one which can also be implemented on a desk-top or lap-top computer) that involves some 30 genes or other and their gene products, and by the repeated implementation of a structure of sub-programs, builds the *Drosophila* embryo.² Biologists have been able to do this because advances in gene sequencing and in computational biology have enabled them to locate both the structural genes and the regulatory ones that switch them on and off.

What these breakthroughs first provided for developmental biology, they soon began to provide for the understanding of the working of normal adult cells and groups of them. Having located genes on the chromosome, it became possible to precisely knock out these genes (and only them) in the germ-line of an organism so the resulting deficits among offspring will enable the experimenter to identify the gene's somatic function. Subsequent breakthroughs—RNAi (i for interference) and double-stranded RNA—now enable experimenters to temporarily silence individual genes in individual cells of the body, or to simply reduce the number of protein products they express. Since everything that goes on in and at the surface of a cell is a matter of the switching on and off of genes

² See Rosenberg, 2007, chapter 2 and 3, for an introduction and Davidson, 2006 for an advanced treatment.

and the chemical reactions—catalytic or transformational—of their protein products, this technology will eventually enable the molecular biologist to learn all the details of cell-physiology down to the level of the individual molecule.

This research has already had three sorts of results relevant to the present matter. First, as in development, it is now evident that many of the processes of cellular physiology are also well understood as the operation of structured digital programs instead of analog mass-action chemical reactions. Second, the diverse programs realized within and among cells are often structured iterations of subprograms to be found ubiquitously in many different processes in many different tissues of the body. These subprograms have come to be called ‘motifs’ owing to their repetition in diverse molecular processes and the search for these motifs, and larger programs built out of these motifs has become a recognized research program under the label “systems biology.” [See Alon, 2006] Third, a class of genes has been identified as Immediate Early Genes. These genes, already implicated in development, physiology and learning are, as the label implies, switched on and off in quick response to molecular stimuli, and synthesize “ Among the immediate early genes already identified are ones with profound roles in behavior. The fos-B gene for example when knocked out seems to destroy normal nurturance behavior while leaving no other detectable deficits in mice [Brown, et. al. 1996, Gingrich and Hen 2000]. Similarly, knock-outs for selective olfactory protein genes in male mice can by themselves inhibit males learning of appropriate courting behavior in these males [Swaney, et. a., 2006]. The discovery of such genes made it clear how somatic cells, and neurons, especially can respond in quantitatively and qualitatively profound ways, to stimuli, rapidly expressing large quantities of catalytic molecules that quickly effect behavior.

2. More than you wanted to know about short term and long term implicit memory

In the introductory paragraph of “Towards a molecular definition of long-term memory storage” Kandel and his coauthors, Bailey and Bartsch, [Kandel, et. al.1996] ask the question, “Can molecular biology provide novel insights into the mind?”. Then the paper turns technical and few specific claims about human cognition are broached

therein. But the paper had important and unnoticed consequences for the nature of human cognition.

Kandel et al begin by distinguishing ‘explicit’ or what they call declarative memory, “the conscious recall of knowledge about people, places, and things” well developed in the vertebrate brain, from ‘implicit’ nondeclarative nonconscious recall of motor skills, classical conditioning, and sensitization (p. 13445). The seats of these two types of memories appear to be separated in the brain. Explicit memory is subserved by structures in the temporal lobe of the cerebrum (and especially the hippocampus), while implicit learning involves learning processes in the sensory-motor pathways of organisms, including invertebrates like the *Aplysia* that do not have anything like a cerebrum. Both explicit and implicit memory are characterized by distinct short term and long term capacities, and both of these are dependent on the number of training trials to which the neural circuits are exposed. Owing to advances in neurogenomics—the use of knock-out and gene silencing techniques in the study of neurons—the macromolecular differences between short and long term implicit memory in *Aplysia*, *C.elegans*, and *Drosophila* are now relatively easy to study. They reveal unsurprisingly enough that the difference between short-term and long term is a fairly obvious difference between establishing temporary co-valent bonding relationships between molecules that degrade quickly and building more lasting anatomical structures.

Short term implicit learning results from conditioning in which a chain of molecular signals and ambient catalytic molecules produce a short-lived modification in the concentration and the conformation (secondary and tertiary structure or shape which changes binding and/or catalytic activity) of neurotransmitter-molecules in preexisting synapses. Specifically, most neuroscientists hold, the conditioning begins when dopamine or (in invertebrates) serotonin molecules are released at the synapse between neurons in a pathway from a stimulated sensory neuron. The serotonin molecules bind receptor molecules at the sensory neuron which in turn activate molecules of the enzyme adenylyl cyclase that catalyze the formation of cyclic AMP, and it in turn activates a kinase protein molecule, **PKA**. This kinase catalyzes the release of another transmission molecule at the synapse between the sensory neuron and a motor neuron. The chain thus laid down will result in the same motor neuron firing again later when another weaker instance of the

same stimulus (i.e. lower concentrations of the same molecules at every prior node) provided it occurs soon enough after the initial stimulation. The neural pathway has ‘remembered’ how to respond to the stimulus. The sea slug, *Aplysia* was important in Kandel’s original research on the neurophysiology of learning because the relevant pathway for a given learned behavior is large enough to have been located and examined even before the tools of neurogenomics became available.

Long term implicit memory appear to be mainly the result of the stimulation of somatic genes to orchestrate the production of *new* synapses connecting sensory and motor neurons. In long term implicit memory the initial steps are the same as in short term learning: the repeated stimulus causes increases in and persistence of serotonin at the synapses of interneurons and sensory neurons, where receptor molecule’s activation of the adenylyl cyclase repeatedly catalyzes the production of more cyclic AMP and thus activates the kinase, **PKA**. But at this point something different from the short term pathway’s chain of events happens: some of the larger number of **PKA** molecules (that result from repeated stimulation) diffuse to the sensory neuron’s nucleus where they activate a protein that binds to the DNA and switches on a number of genes whose molecular products form new synaptic connections between the sensory neurons and the motor neurons. Long term implicit memory is realized by an anatomical change at the cellular level that involves a protein the binds to gene-promotor sequences. The cyclic AMP or **cAMP Response Element Binding protein of type 1 (CREB-1** for short) and the **cAMP Response element**, a gene that is a promoter or on-switch for a suite of genes. Each of the new synaptic connections work in the same way as the smaller number of connections laid down for short term implicit memory, but their larger number means that the learned response will be manifested even if a significant number of the synaptic connections degrade, as happens over time. Thus, the new construction of additional synaptic connections provides for long term implicit memory.

What is stored in implicit memory, long- or short-term? Kandel and neuroscientists in general describe what is stored in implicit memory as learned behavioral dispositions or capacities. This is what philosophers at least as far back as Ryle [1949, chapter 5] have called “knowledge how”. It is nonpropositional and there is no questions of its requiring to be “represented” in the neural structures that realize it. To

the extent that a neural structure realizes knowledge of *how to respond*, for example, to aversive stimuli, an account of its “content” appears to be unproblematic. The content of the system of interneurons, sensory and motor neurons that intervene between painful stimuli and withdrawal of the body-part painfully stimulated, can be characterized variously in our language. For example, we may accord the pathway or some part of it the content “potential or actual tissue-damage at bodily location a.” But we do so without any temptation to attribute concepts such as ‘potential’ or ‘tissue’ or even ‘bodily location’, still less some propositional attitude to a sea slug or a fruit fly. We accept that the experimenter’s attribution of ‘knowledge how’ to the simple pathway must take some propositional form, but we don’t take the particular proposition chosen “seriously” as *the* content; we don’t suppose that any experiment could narrow down the actual complete content to any particular unique proposition, nor is there any need to do so in order fully to understand the nature of implicit memory, short or long term. It would not be difficult for neuroscientists to narrow down the set of propositions that more completely characterize what is stored in implicit memory by varying stimuli and/or stimulating peripheral neurons that intervene between stimulus and the sensory-motor neurons . Such experiments enable the experimenter more narrowly to identify what features of the stimuli or peripheral neurons are “detected” (i.e. responded to) by the path way.

2. How are explicit memories stored?

Explicit memories, recall, are those which Kandel et al [1996, p. 13445] identified as “declarative:” “the conscious recall”, or better, storage “of knowledge about people, places and things” that is particularly well developed in the vertebrate brain. Storage of course is hardly even conscious, and it’s clear that what Kandel is interested in is information that *can be recalled* to consciousness, expressed propositionally and that is presumed to be represented in the brain. (I will have nothing more to say about consciousness hereinafter.)

Explicit memories storage is localized to the temporal lobe, and especially the hippocampus³, structures unknown in the Aplasia and far from the sensory pathways in

³ There is some evidence that remote early memories in vertebrates are moved from the hippocampus to storage in the neocortex. But there is also evidence that the

vertebrates that are homologous to the ones storing implicit memories—short or long term—in the sea slug. Studies of neural processing in these regions of the temporal lobe began with the determination that neural pathway there are subject to long term potentiation, LTP, the process in which synapses become much more sensitive when they are repeatedly stimulated. In particular repeated stimulation by the same concentrations of neurotransmitters of the neurons in the hippocampus results in much higher production by them of neurotransmitters that stimulate down-stream neurons. This increased sensitivity can endure for long periods. LTP occurs in three different pathways in the hippocampus (Mossy fiber, Schaffer collateral, medial perforant) and goes through three different stages, short term, early, and late LTP. There is of course a great deal of evidence that LTP in the hippocampus constitutes information storage in laboratory animals, especially mice.

Kandel et al.'s [1996] studies of these stages of LTP in these distinct pathways showed that the *same* molecular mechanisms, involving the same somatic genes that build new synaptic connections in the *Aplasia* implicit long term sensory motor memory, are responsible for all the forms of LTP in all the hippocampal pathways that subserve explicit memory in vertebrates. They write,

Similar to the presynaptic facilitation in *Aplasia*, both mossy fiber and Schaffer collateral LTP [two of the three types of LTP] have distinct temporal phases... The early phase is produced by a single titanic stimulation [release of neurotransmitters], last 1-3 hours, and requires only covalent modification of preexisting proteins. By contrast, the late phase is induced by repeated titanic stimulations, and is dependent on new proteins and RNA synthesis. As is the case with long-term memory in *Aplasia*, on the cellular level there is a consolidation switch, and the requirement for [gene] transcription in LTP has a critical time window. In addition, the late transcription-dependent phase of TLP is blocked by inhibitors of **PKA** ... Recent studies by Nguyen and Kandel now indicate that these features of LTP also apply to a third major hippocampal pathway, the

macromolecular processes obtaining in this area of the brain are substantially the same—involve the same genetic programs and the same neurotransmitters as those obtaining in the hippocampus and the sensory motor cortex. See Silva 2004.

medial performant pathway... Thus, as in *Aplasia* presynaptic facilitation, cAMP-mediated transcription appears to be the common mechanism for the late form of LTP in all three pathways within the hippocampus. [p. 13452]

Note the role of **PKA** and cAMP which has been demonstrated to be the same in both implicit and explicit memory. Moreover, the presence of RNA synthesis implicates gene expression in the manufacture of new synapses. One thing missing from Kandel et al [199] account of the mechanism of LTP in the hippocampus was the identification of the same genes and their transcription factors—**CREB1**, **CREB-2** (cyclic AMP responsive element binding protein) or analogs-- already known to be at work in implicit learning-- as responsible for the RNA synthesis Kandel reports in explicit memory LTP. This would demonstrate that the same genes build the same new anatomical structures in explicit as in implicit memory. Kandel et al observed, “it will be of particular interest to investigate whether cAMP-dependent transcription factors...are also required for long term synaptic modification in mammals.” [13452] Many studies subsequent to the Kandel paper vindicated the expression of interest in this question. The answer appears to be a resounding yes.⁴

With understandable caution, Kandel et al concluded: “The apparent similarity of the molecular steps that underlie learning-related synaptic plasticity may reflect the fact that long-term memory for both implicit and explicit storage is associated with structural changes.” [Kandel, et al., 1996, p. 13452] Not just changes, but the very same structural changes, down to the same genes, transcription factors, and proteins, most of them identical at the molecular level.

What happened in the next ten years or so turned this speculation about the identity of explicit and implicit memory storage into well grounded theory. To understand the basic outlines of how Kandel’s picture has been filled in, we need to introduce the **NMDA receptors** (so named because they can be stimulated by the drug **N-methyl-D-aspartate**, by contrast to other receptors which cannot be so stimulated). These receptors are to be found elsewhere in the brain, including in the sensory-motor cortex

⁴ See Bailey, C., Kandel, E., Si, H., 2004. Another recent example is to be found in Yang Zhou, et al., 2006.

where implicit knowledge is stored.⁵ When NMDA or the amino acid glutamine binds to these receptors, their shape changes and opens a channel into the neuron through which Calcium ions flow, increasing the sensitivity of the neuron to subsequent electrical impulses. Work in one important region of the hippocampus, the CA1 region, has shown that when NMDA receptors are bound by a drug so that they cannot respond to the NMDA or glutamate molecules, LTP does not occur; when genes in the neurons produce below normal quantities of NMDA, no LTP occurs and when these genes produced above normal quantities of NMDA, the neuron's susceptibility to LTP and animals' abilities to store explicit information is enhanced. This research further substantiates Kandel's hypothesis that long term implicit and explicit memories are realized at the individual synaptic level by the same mechanisms of somatic gene regulated multiplication of anatomical structures.

Although there remains a substantial possibility that something entirely new in kind will be reported, several decades of work now strongly point to the conclusion that the molecular mechanisms of implicit and explicit memory, both short and long term differ only in kind. Before drawing several of the obvious morals of this story for the nature of representation in the cerebrum, it is worth summarizing what is now known about the mechanism of information recall, to go along with what we know about information storage.⁶

3. How the brain recalls memories

Recall is a capacity associated only with explicit or declarative memory, of course, and if it is a matter of accessing representation laid down in explicit memory, it must take place in the hippocampus.. Research on recall of stored information has focused on both the CA1 and the CA3 region which project on to the CA1 cells on which

⁵ Such receptors have recently been identified in the *Aplasia* ganglia. See Ha, Kohn, Bobkova, Moroz, 2006.

⁶ It is probably also worth pointing out that the similarity of molecular mechanism in these four different kinds of memory processes is entirely neutral with respect to any "multiple realizability" claims about the neuroanatomy of memory. Similarity of mechanism is not the same as identity of it. Moreover, none of the inferences drawn here from the similarity of mechanisms to the general nature of cognitive processes turn on whether there is substantial, some, slight or no multiple realizability in the relationships in question.

the work reported in section 2 was focused. The strategy for studying how recall is effected unsurprisingly involves modulating the presence or the behavior of the NMDA receptor proteins on the surface of neurons. Tonegawa et al (2002) were able to provide an account of the molecular mechanism of how stored information is recalled in response to environmental clues by knocking out genes that code for an important component of the NMDA receptor proteins, NR1 (for NDMA receptor). The work also shows how information is degraded in the brain, and how degraded information can be recalled and used effectively.

Normal mice trained in a Morris water maze to recall the location of an underwater platform (on which they can rest) continue to do so without difficulty even when most of the environmental clues to the platform's location have been removed from the maze. Tonegawa et al called this behavior 'pattern completion.' Mice in which the genes for NR1 have been knocked out showed normal behavior and normal protein distributions through development and maturation. But by week 18 the NR1 protein had disappeared from the CA3 region, while remaining at normal levels elsewhere (including CA1 region) in the brains of these knock out mice. The cytoarchitecture of the mutant hippocampus was otherwise normal, and LTP elsewhere in the hippocampus (even LTP requiring NR1 in CA1) was unaffected by the gene deletion. Knock-out and normal mice showed no difference in ability to acquire and retrieve explicit spatial memories of the location of the underwater platform in a Morris water maze with 4 cues on the surrounding walls. But when 3 of the 4 cues were then removed and the animals retested, knock-out mice were significantly worse at finding the platform. The experimenters excluded the hypothesis that the difference was due to memory-loss by showing that the knock-out mice did as well as controls when all 4 cues were restored. "In summary, under the full-cue conditions, both mutant and control mice exhibited robust memory recall, When three out of the four major extra-maze cues were removed, control mice still exhibited the same level of recall, whereas the mutants' recall capacity was severely impaired." (p. 215). What neural mechanisms underlie this deficit?

Recall what the research described in section 2 showed about explicit memory storage in the CA1 region of the hippocampus. Assuming that neural processing in this region stores memory, a comparison between cellular activity at the CA1 regions of

normal and knockout mice subjected to the same water maze training regime and then tested with and without cues would show whether knocking out the NR1 gene impairs memory storage in CA1. The result of this experiment strongly suggests that mutant mice memory storage in CA1 is *unaffected* by the knock outs. It is retrieval that is affected. Since the absence of NR1 proteins effects storage in CA3, Tonegawa advanced a two stage theory of memory retrieval involving connections between the CA1 and CA3 neurons. Information, in this case spatial information about cues, that indicate location of the under-water platform is stored in neurons in both CA1 and CA3. A normal CA3 circuit stores full cue information and in the presence of full cue inputs stimulates the CA1 neurons to produce the hippocampal output that leads to rapid arrival at the platform. In mutant animals, the NR1 deficit impairs memory storage in the CA3 areas but has no effect on the CA1 storage, which in the presence of all 4 cues produces behavior indistinguishable from that of normal mice. Partial cue removal reduces the input stimulus to the CA3, but in normal mice these circuits are sufficiently strengthened by previous stimulation that they drive the full output pattern of the CA1 neurons. In mutant mice the number of CA3 neurons providing input to CA1 neurons is significantly reduced and the corresponding CA1 neurons are not driven to produce the full output pattern of normal mice. Tonegawa et al conclude that their “results reflect a primary deficit in NR dependent memory formation in CA3 that is then revealed as a deficit in recall under limited cue conditions.” (p. 218). Tonegawa and his collaborators were not reluctant to draw inferences from this research about CA3 cells in humans. They connect these results to the previously established neurochemical alternation in CA3 cells of some Alzheimer’s patients.

The data from these experiments along with the work of Kandel and others led Tonegawa to propose a schematic probabilistic synaptic program for memory-storage and recall realized by the CA1 and CA3 cells, along with the rest of the architecture of the region of the hippocampus (the entorhinal cortex, the dentate gyrus, the Schaffer and recurrent collateral neurons, mossy fiber neurons and the perforant pathway) [p. 217]. The model suggests multiple memory location of the same information; its probabilistic output mirrors the fact that memories are imperfect even in normal animals, and that the deficit produced by knocking out the gene coding for NR1 protein in CA3 cells reduces

the probability of correct recall from a very high level to a much lower one. But the take home lesson for our purposes is the distinctive recurrence of the same neurogenomic “motifs” across different parts of the brain, and different organisms, engaged in different memory-expressing behaviors.

4. Each explicit memory is just a lot of implicit memories

It may not be safe to assume that human memory is realized in human brains by the same mechanism that realizes it in rat and mice brains, but there seems to be great deal of evidence in favor of it, and none against it. It seems equally reasonable to adopt as a working hypothesis that the differences between implicit and explicit or declarative memory is a matter of location and degree. The former is realized by the number and sensitivity of neuronal connections in the sensory motor areas of the vertebrate brain, and the latter is realized in the cerebrum and in particular in the hippocampus. More significantly, the molecular biology of long term implicit memory and explicit memory appears also to be substantially the same, indeed identical except for the particular molecular configuration of the neurotransmitters and the nucleic acid sequence difference of the genes and RNAs that regulate changes in the microarchitecture of synaptic connections.

If the difference between the details of the neural connections that constitute long term storage of implicit memory—storage of knowledge how—differs only by number of connections from long-term storage of explicit memory—knowledge that—then it is reasonable to consider whether long-term explicit memory storage differs only in degree from long term implicit memory storage. Or in other words, it’s worth considering whether propositional knowledge is nothing but a large number of synaptic connections each one of which is a bit of associative learning, a neural circuit that realizes a conditional disposition to respond to stimuli in an environmentally appropriate way, a little bit of knowledge how.

Recall the point that the sensory-motor pathway produced by classical conditioning in the *Aplasia*, constitutes the stored disposition to respond to noxious or positively rewarding stimuli in an environmentally appropriate manner. In these cases it is natural to attribute content to the circuits owing to their functional role. Nevertheless

there is no temptation to attribute specific propositional content to such circuits, still less to identify a “language of thought” in which sentences expressing propositions about the presence of noxious or rewarding stimuli are written in the arrangements of neurotransmitters. The intentional attributions are instrumental, short hand, heuristic. When, as we have seen, a large number of these circuits are “grown” owing to somatic gene expression involved in long term memory, an organism like the *Aplysia* persistently responds distinctively to stimuli. Under these conditions we may be tempted to accord the relevant circuits, when ‘switched on’, content *derived* from our own “original” intentionality, our representations to ourselves of the presence of aversive stimuli in the sea slug’s immediate environment. But derived intentionality is not original intentionality, and it is no more problematical that the intentionality of the pixels or ink marks now reflecting light on to your retinas.

Move the same circuits to the hippocampus, multiply their numbers by several order of magnitude, and the result is long term explicit memory. Recall that Kandel called this explicit memory “declarative” and labeled it as conscious, meaning presumably that in humans the information stored can often be recalled “at will,” that when it is recalled can be verbalized and be the subject of conscious awareness, events of successful recall “at will,” and verbalization in humans at least. Now, mere change of location cannot turn individual circuits that do not represent and lack intentionality into ones that do so and have it. Nor can it turn large packages of circuits that might have derived intentionality into ones with original intentionality. *Natura non facit saltum*. At least change in location and numbers cannot change structures that don’t represent and are not intentional into ones that do and are, unless this change is merely a matter of degree, something no one should grant.

So, how do the circuits at CA1 or CA3 or the circuits that include synaptic connections from CA1 to CA3 and all the other subregions of the hippocampus, acquire the original intentionality that we believe reposes in the human brain? If the storage of information in the hippocampus is not a matter of original intentionality, then there seem to be only two options:

- 1) Representational states in the human (or primate or mammalian or vertebrate, take your pick) brain, the ones with true, real, original intentionality, are not to be found in

exclusively in the hippocampus; other parts of the brain are necessary and it is they that are responsible for the original intentionality.

2) Some (much, all?) representations are fully located in the hippocampus, and there is no original intentionality to be found anywhere in these representational states; intentionality is an illusion foisted on us by the sheer numbers of synaptic connections which produce such exquisitely fine grained responses that mistakenly attributing specific propositional content is overwhelmingly natural to creatures like us, and heuristically invaluable.

Available neuroscience does not recommend the first alternative. As I have said, there are always surprises in science. Cross-phylogenetic inference from experiments on rodent brains to ours is always to some degree fraught. Deficit, lesion, fMRI, PET-scan studies in humans may tomorrow show that memories accessible to consciousness are also stored elsewhere [modulo footnote 3 above] or that their storage spans regions that include more than just the hippocampus. But the betting must for the moment be against these alternatives.

Let's consider alternative 2) somewhat further. *Aplasia* is a wonderful model for understanding how the neuronal network realizes conditioning, but it is limited by the narrow repertoire of behavioral responses that the *Aplasia* makes to a range of stimuli. This narrow range prevents us from attributing much discrimination of its environment to the sea slug. The more nuanced the differences in behavior as a function of the presence of different stimuli, or differing quantities of the same stimuli, the more content it is convenient to attribute to the neural networks intervening between stimuli and behavior. Though of course, in the case of sensory motor networks in these organisms, the attributions of content are merely heuristic (like those familiar from molecular biology's description of molecules as, for instance, "recognizing second messengers"). The number of different kinds such synaptic pathways multiplies as the peripheral sensory neurology of organisms becomes more differentiated and the number of central neural synaptic connections that are possible increases. As this process proceeds in phylogeny, attributions of content will become increasingly specific. Once we get to *Drosophila* it will be natural to attribute information about self-location, local temperature, the direction and nature of nearby vegetation and food-sources, the presence of fertile

females or sexually aroused males, etc. And of course when we get to vertebrates, our attributions will become very fine grained indeed. At this point indeed, common-sense and cognitive science will begin to attribute content not just heuristically but literally. However once we recognize that what is going on, neurophysiologically, in these cases is just more of the same as what happens all the way down to the *Aplasia* we should want to resist such attributions and explain them as heuristic devices fostered by the wonderfully sophisticated behavioral responses that these trained-up neural pathways give rise to.

When we get to the human case we are treating systems with hippocampuses composed of millions of cells. The rat CA1 area is estimated to contain upwards of 355,000 pyramidal cells, each making synaptic contact with hundreds or thousands of interneuron cells. Though the difference between what transpires in the rat hippocampus and in the sea slug sensory motor system differs only by degree, it is a difference of at least 5 orders of magnitude. This is a difference in degree large enough to so finely narrow the range of propositions heuristically attributable as content to sets of synaptic connections in the hippocampus that literal ascription becomes overwhelmingly tempting.

In the human case we may be able to add a further order of magnitude difference in degree from the rodent (even though the number and structure of the relevant somatic structural and regulatory genes seem to be similar). What this vast multiplication of synaptic structures does for the human is asymptotically to approach the literal content ascriptions common sense and cognitive science make. Our use verbal and inscriptional sentential expressions to express out thoughts strongly suggests that we have particular propositions ‘in mind’ that are expressed by the voice or the hand. Agreement and disagreement with others suggests a common acceptance of the same propositions. But of course experiences with ambiguity and differences in the meaning of sentences, expressed and inscribed, long ago taught all of us that each person may have a different proposition in mind when they assert or assent to the same sentence inscribed or vocalized in a natural language.

The approach to content inspired by the continuity of implicit and explicit memory storage takes us only a little way beyond this lesson. What it shows is that the set of circuits in the hippocampus and/or elsewhere in the brain that realize some information about the world are equally well described by more than one particular proposition from a

set of only very slightly different ones, because the neural architecture does not in fact literally realize *any* propositional content at all. Rather the circuitry is connected to a set of behavioral dispositions that will equally well support any one of the members of this set of extremely similar propositions, and there are circumstances we can arrange which will enable us further to narrow down this set of equally-well supported propositions. But they won't narrow it down to a unique proposition that fully and completely expresses the representational content of the synaptic connections in question. For, like the much smaller number of otherwise similar synaptic connections in the *Aplysia*, they are not capable of representing propositional content. They are not capable of original intentionality. Note that if the neural architecture cannot represent unique single propositions, it cannot represent precise (or fuzzy) sets of them either, for to begin with if it did so it would also have to represent the unique precise single (vague or nonvague) proposition that is constituted by the disjunction of the member of the set, an even harder task than literally representing one of the members of the set.

5. Farewell to original intentionality

But wait, if there is derived intentionality, as surely there is (just look at the words on this page), there must be original intentionality somewhere in the universe from which the derivation comes. And surely this original intentionality must be realized in the human hippocampus if it is instantiated anywhere? These rhetorical questions ride roughshod over several possibilities that need to be ruled out before we accept the indicative claims they really express.

Does an inscription's having any derived content at all require that there be something else—a brain state-- with some original intentional content or other on which the inscription's intentional content is dependent? Assume the answer is yes. By *modus tollens*, therefore, it follows from the denial of original intentionality in the brain that there is no real derived intentionality, that inscriptions, vocalizations and other signs are not really *symbols*. This means of course that our attribution to them of the literal status of 'symbol' is mistaken. But what if it turns out that "treating them" as symbols is a heuristic device we employ owing to the same considerations which lead us heuristically

to ascribe intentional states to the synaptic circuits of the *Aplysia*.⁷ If it is owing to the asymptotic approach of our neural circuitry to representing particular propositions that we mistake our heuristic attribution of derived intentionality for the literal attribution to signs of symbolizing—of representing specific propositions. Derived intentionality, if any, is dependent on original intentionality. *Mutatis mutandis*, the asymptotic approach of signs to derived intentionality is dependent on the asymptotic approach of neural architecture to original intentionality.

It is worth noticing how the present account of explicit memory as on a continuum with implicit memory deals with one of the most persistent recent objections to a teleosemantic account of intentionality: the so-called disjunction problem [Fodor, 1991]. This is the challenge to teleosemantic theories' attribution of representations to neural pathways that they cannot distinguish cases in which such pathways mistakenly realize a *false* propositions from cases in which they correctly realize a *disjunctive* proposition. An illustration will help. Suppose some neural pathway is said to realize "Raven here now" owing to the environmental appropriateness of the behavior to which it leads in the circumstances in which it occurs. Now suppose that the same pathway is instantiated by a treepie (a black magpie, a bird of the same family as the raven, with much the same appearance). Any account of neural representation (for example a teleosemantic account) must provide a principled distinction between the contents, "Raven here now" and 'Raven or a treepie here now.' Tokening the former content is a natural mistake creatures can easily make. Tokening the latter is presumably rare, especially among non-linguistic creatures. If a theory cannot distinguish between them it does not have the resources to account for unique specific representational content. Or so the argument goes.

The view that explicit memory is a huge number of implicit memories asymptotically approaching representation deals with this problem naturally and easily. To begin with, it holds that all representations are disjunctive, in the sense that there is no unique proposition realized by a set of synaptic connections; rather there is always a

⁷ Here of course one must be careful not to lapse into literal attribution of intentionality while explaining why there can be no such thing. "Treating signs as symbols" cannot mean bringing them under descriptions or any other activity whose description is referentially opaque. But this is a philosopher's problem, not a neuroscientist's.

disjunction of them each equally suited to a heuristic attribution of content. Does a synaptic network's connections make it natural to attribute 'raven here now' v 'raven or treepie here now'? That depends on a previous history of discriminations and failures to discriminate ravens and treepies. If there is no such history the two propositions are on a par, along with indefinitely many other propositions, all of which it makes sense to attribute in the light of stimuli and subsequent behavior. The most complex neural architecture never represents a unique single proposition, because it never represents any propositions at all. And only rarely will it be the case that there is a unique nondisjunctive proposition which it is unavoidable heuristically to attribute.⁸

Fodor, who famously taxed teleosemantics with the disjunction problem has advocated a theory of (original) intentional content which he claimed is not held hostage by the problem. This is his theory that intentional states are ones that are involved in relations of asymmetrical causal dependence. Consider a synaptic network, call it SN₁. Adapting Fodor's theory, the conditions SN₁ must meet for it to represent "Raven here now." is roughly this:

1. There are occasions on which the presence of a treepie causes a synaptic network to produce raven-appropriate behavior. Call this synaptic network SN₂.
2. When SN₂ obtains, it does so because SN₁ is in place, but not vice versa.

While the theory of heuristic attribution advanced here does not actually involve literal attribution of original or any other intentionality, it explains perfectly well why such a theory as Fodor's should look attractive. If a synaptic network SN₁ is "trained up" to produce long-term raven-appropriate behavior in the presence of ravens, this will be due to the building of a large number of new synaptic connections, as we have seen. Such connections will facilitate occurrent behavior of the sort that leads from treepies to raven-present behavior that otherwise would not obtain. After all, under some conditions some treepies look very much like some ravens. But the reverse dependence does not obtain. Heuristically attributing "Raven here now" to SN₁, as opposed to "Black colored bird or black colored prey or black colored moving object here now", requires that the behavior

⁸ It is perhaps only when we think of a mathematical truth such as $2 + 2 = 4$ or first person truths such as "I am here now," that it will be reasonable to narrow down the class of propositions attributed to a unique one.

SN₁ produces be rather specific, specific enough to exclude “Treepie here now” as a heuristically useful attribution. It is evident that on its initial occurrence, SN₂ --the “Raven here now” synaptic system caused by a treepie and resulting in raven-appropriate behavior—does not causally contribute to the initial appearance of, presence of, or long term build up of SN₁ synaptic connections. Whence the asymmetrical causal dependence Fodor thinks is the essence of original intentionality.⁹

6. Is ‘knowledge how’ computable?

The claim that explicit declarative memory is just large quantities of implicit memory realizing behavioral dispositions, but not inscribed with propositions, may be resisted, owing to an attachment to the computational theory of the mind—the thesis that cognition is at least often a process of syntactical manipulation of representational states with semantic content. To a first approximation this theory holds that thinking is a mathematical or inferential process familiar, to logicians and computer scientists, in decidable procedures and algorithmic programs, that operates syntactically on symbols—inscriptions with semantic meaning--to produce further representations and ultimately behavior. Insofar as the view defended here is inhospitable to the literal attribution of representational content to synaptic networks, it may be viewed as hostile to a computational approach to cognition. But there are important aspects or components of computationalism to which this approach is quite amenable. Indeed, it holds out the promise of actually providing a concrete program of research that will substantiate an important component, indeed the most important component of computationalism.

Begin by separating the computational theory of the mind into two well recognized components: a) the claim that cognition is literally intentional, that brain states are symbols with (original or derived) or derived intentionality, and b) the claim that cognition is a form of reasoning that consists in the formal or purely syntactic

⁹ Suppose after SN₁ has been established, treepies begin to show up regularly and that SN₂ becomes established. That would lead us to revise our heuristic attribution to SN₁ of the proposition “Raven here now” to “Raven or treepie or jay or other member of the *Corvidea* family of birds here now” with no temptation to credit the bearer of SN₁ or SN₂ with acquaintance with Linnaean binomial nomenclature.

manipulation of these brain states. It is the second of these two claims that neurogenomics holds out the promise of vindicating.

It is of course obvious that synaptic connections and the somatic genes that construct and regulate them can realize the truth-functional connectives familiar in syntactical symbol processing. But that is only the tip of the computational iceberg whose proportions are now becoming familiar in neurogenomics and proteonomics. To begin with the computational character of development has become well understood in the early stages of the embryological development of many species [see Rosenberg 2007 for examples], and it appears to embody many of the same subprograms across the phylogenetic spectrum from insects to humans. These subprograms are called “motifs” and the term has become common in neurogenomics as well. The aims of “systems biology” are to uncover these motifs and indeed higher order motifs composed of repeated iterations of a small number of such motifs in much the way that structured programs in computer science are composed. The identification of computational programs realized by sets of structural and regulatory genes and the protein products they produce is the major current focus of molecular biology. This is equally true in neurobiology as it is in developmental biology. Illustrations of this interest are easy to find. One recent example is provided by the elucidation of sensory information processing in *C.elegans*. [Chasalani et al. 2007] In the presence of certain food related odors this worm will move towards their sources, and in the absence of them engage in a characteristic random turning-and-moving behavior. Neurobiologists have now identified the molecular steps in the neuronal pathways from the worm’s nose through the sensory motor neuron pathway. In particular they have shown the existence of a circuit of switches, called AIB and AIY that probabilistically start and stop turning-behavior.¹⁰

¹⁰ Why should the control be probabilistic: “Why not just respond in an all-or-none manner? The reason is that a circuit that generates probabilistic behavior leaves itself open to further modulation. The circuit involving AWC, AIY and AIB neurons is only a sub-circuit within a much larger, interconnected neuronal network.’ Chasalani, et al. 2007, p. 35] Inputs from this network could modulate the output of smaller sub-circuits, allowing the organism to integrate information from several environmental sources and accordingly alter the probability of a response. Thus, a probabilistic network might cope better with uncertainty and unpredictability — characteristics of the real-world environment.

subject to parallel but opposite regulation by a pair of initial peripheral neuron, AWC^{on} and AWC^{off}. One thing they have shown is how a quite sudden removal of stimulation in the olfactory neurons produces a sustained level of excitation in downstream neurons that start and keep the worm in its search-behavior mode, thereby accounting for “how ‘working memory’ of a transient stimulus is maintained.” [p. 35] It is no surprise that the molecular signals realizing this program—glutamate molecules-- are of the same general type as in implicit and explicit memory across the phylogenetic spectrum. More significantly, the olfactory pathway thus elucidated by Chalasani and colleagues is described by a commentator as “a circuit motif that seems to be used in several contexts in different organisms. In vertebrate vision, information flow through parallel channels of opposite signs is crucial for contrast sensitivity¹³. The use of a similar functional motif in *C. elegans* olfaction emphasizes the efficiency of this circuit in processing sensory information, and supports the idea that systems of very different complexities may nevertheless use shared strategies to perform similar tasks¹⁴. In other words, much within man might still be worm.”¹¹

Examples can be multiplied and in effect have been in section 2 above. Thus the process of memory recall elucidated there constitutes a computation program which when

¹¹ Chalasani et al [2007, p. 68] write:

The sensory responses of AWC neurons are similar to those of vertebrate rod and cone photoreceptors, which have tonic activity in the dark, are hyperpolarized by light, and are depolarized by the removal of light³⁷. Like AWC neurons, photoreceptors are non-spiking and have graded glutamate release. Molecular analogies also link AWC neurons with vertebrate photoreceptors: their sensory transduction pathways rely on G-protein-coupled receptors, G_i-like proteins, receptor-type guanylate cyclases and cyclic GMP-gated channels, and their differentiation is controlled by Otx homeodomain proteins. The synaptic connections between AWC, AIY and AIB neurons are also reminiscent of those between vertebrate photoreceptors and their targets, the ON and OFF bipolar cells: In the retina, glutamate from photoreceptors is sensed by AMPA-type receptors on the OFF bipolar cell, a connection that is functionally and molecularly analogous to the AWC-to-AIB connection. The vertebrate ON bipolar cell is inhibited by glutamate, as are AIY neurons. In mammals, the ON bipolar cell is inhibited by a G-protein-coupled glutamate receptor and, in fish, by a glutamate-gated chloride channel that is functionally although not molecularly similar to *C. elegans* GLC-3. The parallel ON and OFF streams enhance contrast sensitivity in vertebrate vision; it is possible that the parallel AIY and AIB neurons have analogous functions in odour detection.

operating normally, takes reduced clue data as input and using stored memories, computes a probable location for the underwater platform in the Morton water maze. When certain somatic genes are knocked out the computation cannot be effected. The sub-programs of this computation established by LTP are identical to those computed in the interneuron-sensory-motor pathway of the *Aplasia*. It is early days in the search for the motifs that phylogenetic and genomic continuity lead one to expect across the metazoan sensory and neural equipment.

It seems safe to say that current research in neurogenomics is identifying the physical systems that realize data structures and the physical processes that employ them to compute outputs of various kinds, following a variety of nested, parallel and independent programs, and that the programs are not just “curve-fitting overlays” but reflect what is really going on in the brain. In fact the kind of programs neurogenomic methods are enabling neuroscientists to identify as realized in regions of the brain, hold out the prospects of ***computational compositionality, systematicity and productivity***. And these are the very features of thought which constitute the best arguments for computationalism in cognitive science. Accordingly, these neural programs identified by genomic/proteonomic methods have the very features we should look for in processes that vindicate a significant part of computationalism about the mind.

Consider first the compositionality of synaptic structures in the hippocampus. If each such structure embodies a specific schedule of learning from inputs, output and feedback, and if there are upwards of a million or so of them in the three regions of the hippocampus about which we already now something, then as genomic and proteonomic programming adds and subtracts, or just temporarily blocks or opens connections between them, the combinatorial combinations of these synaptic patterns will add up to unimaginably large exponential powers of the numbers of basic synaptic structures that are to be found in a normal hippocampus. This gives us probably as much productivity as we require to explain the fact that actual patterns of synaptic structures are only a small proportion of possible ones, and to show why the actual behaviors they eventuate in are only a small proportion of the repertoire of behaviors brain states can give rise to. A huge number of such basic synaptic structures operated on by even a small number of computational programs, encoded in somatic genomes and their protein products, will be

just as powerful as a more complex computational program operating on a smaller number of highly structured propositional representations. Indeed, the large number of basic dispositional “knowledge how” synaptic structures operated on by many iterations of the same relatively limited package of subprograms can produce the same output, indeed can be the very realization of the higher level program operating on syntactically richer structures. This is just what goes on in the realization of higher level programs assembly level programs, by machine-language programs at the basement-level computations realized in the or-, and-, and not-gates of the microprocessors in a computer.

Because every synaptic network comes equipped with exactly the same set of structural and regulatory genes, prepared to respond to various signals, the synaptic structures will also show the systematicity that computationalism credits to the mind. The same syntactic structure can figure in a large number of different networks depending on how changes in the neural inputs to it effect the switching on and off of genes and the neurotransmitters they produce which leave the initial structure intact but simply allow or establish connections to other such structures. It is as if each basic structure comes with a set of connectors that can be plugged into the connectors of other members of a large but not infinite set of structures that have the same or coordinated connectors. These connectors can be used to change the direction in which a structure sends signals, can coordinate the sending of signals with other structures, or when pulled apart stop sending signals in the same direction, etc. What does the changing, coordinating, disaggregating? The obvious candidate is the general neurogenomic/proteonomic program laid down in the genome, carried by all the cells in the body but only switched on in the brain tissue cells by the operation of regulatory genes and environmental inputs in those very cells. Most of these programs, exploiting immediate early genes, are epigenetically switched on in all the brain cells, while other more specialized neurogenetic programs operate in each of the distinct regions of the brain, such as the CA1 or CA3 regions of the hippocampus.

7. Computationalism and neuroscience

It requires little argument to show that cognitive psychologists cannot uncover the programs that implement cognition by reverse engineering from behavior. Just trying to

infer the programs a much simpler system such as a lap top's *central processing unit* is too much to demand. Of course there has been some success reverse engineering the program which the visual cortex employs to solve the so-called inverse optics problem of inferring a 3-dimensional environment from the 2-dimensional array on the retina, and some advance in identifying the problem-solving heuristics employed by human cognitive agents in stead of less convenient algorithms that assure correct answers. But for many reason neither of these lines of research will vindicate computationalism. The comparative ease of discovering programs operated in the sensory modules is a reflection of their modularity and of the relatively simple problems they have to solve along with the significant constraints that make the range of alternative solutions relatively small. And the problem-solving heuristics that have been uncovered are more in the nature of data that a computational theory would need to explain than they are advances in a computational research program for cognitive science.

The other main lines of research in cognitive neuroscience, especially those devoted to localization of distinctive cognitive function to particular areas of the brain, to the timing, temporal sequence and duration of cognitive processes, that functional MRI or other techniques provides is one much too coarse-grained to tell us anything much about how the regions they identify do what they do, let alone the programs which are realized by processes in these regions. If anything, such studies are the necessary precursors to a neuro-genomic approach which first seeks differences in the timing and rates of gene expression between areas mapped by the coarse grained anatomical studies to vindicate the taxonomies they suggest.¹² Then if neurogenomic studies finds these differences in gene- and protein expression, it begins knock out-studies to reverse engineer from deficits to the genomic/proteonomic programs and subprograms to these programs' computational structure.

What are the chances that the resulting computational neuroscience will vindicate a computational cognitive science? Computationalism as traditionally understood is the "bet" that top-down cognitive science can uncover programs operating over such

¹² Timing and rates are everything in such studies, as there is good reason to believe that almost all the neural genes are expressed at some time or other in almost all the regions of the brain, See the on-line brain gene expression map at <http://www.stjudebgem.org/web/mainPage/mainPage>.

semantically characterizable states in the brain. If we could do this, the next step in the research program would be cognitive neuroscience—the search for the physical nature of such computational processes and the physical composition of the semantically characterized states which these processes syntactically manipulate. Were computationalism to succeed, we should expect that it would meet bottom-up neuroscience in one or another smooth “reduction”.¹³

It’s hard to say whether bottom up neuroscience will meet top-down cognitive science, since there really isn’t any of the latter yet. In fact the paucity of successful computationally motivated theory in cognitive science, in spite of its comparative antiquity and the success of computational neuroscience in spite of its youth, may suggest that there is nothing much for a bottom up science to meet. At most so far computationally motivated cognitive science has identified the behavioral capacities a cognitive theory needs to explain and has placed some general constraints on any theory that does so: it should explain the productivity, systematicity and creativity of thought as evinced in behavior.

The best case scenario for a smooth link up between some future cognitive science and neuroscience will be one in which a theory committed to representations is vindicated as a model to which the real nonrepresentational neural processes approximate or approach asymptotically. This would require two things: first, that representations required by computational cognitive science be approximated by independently individuated sets of synaptic structures which never represent any distinct propositions; second, that the “high level” computations that cognitive science identifies be shown to be constituted or closely approximated by sets of “lower level” computations over these sets of synaptic structures or sub sets of them. The worst case scenario for a cognitive science motivated by computationalism would be the conclusion that neither its representational states nor computations that operate over them are anything like sets of synaptic structures and the computations which the genome/proteome subjects them to.

¹³ Using the term loosely enough so that it includes ones that may leave important residual philosophical difficulties, such as those involved in the “reduction” of thermodynamics to statistical mechanics, and excludes eliminative explanations like the latter day redescription of phlogiston chemistry’s isolation of oxygen as a case of reduction.

In the absence of any significant success of computationally inspired cognitive science, philosophers and cognitive scientists will offer a range of arguments familiar from other disciplines in favor of continuing a top-down research program. Such arguments conclude that if there are any interesting theories or models of computation at the level which approximates to semantically evaluable representations, then a bottom-up approach will miss them. When pressed to give examples of such higher-level regularities from elsewhere in biology or the special sciences, exponents of this claim can at best cite some *ceteris paribus* generalizations that in fact describe the *explananda* of the science in question, not the *explanantia*, or else they cite mathematical models—necessary truths—whose explanatory power is limited, and which a bottom-up approach is unlikely to miss in any case. On the one hand, exponents of a bottom-up approach can claim a long history of vindication in the discovery of regularities; equally important the bottom up results can explain both those explananda-describing *ceteris paribus* generalizations and the circumstances under which their *ceteris paribus* clauses must be invoked.

On the other hand, the best relevant example adduced in favor of bottom-up approaches is both powerful and highly relevant to the top-down/bottom up dialectic in cognitive science. It is to be found in developmental molecular biology and its “evo-devo” applications. Here in the last 20 years more than one Nobel Prize has been awarded for the uncovering the genetic program that actually explains more than a century of descriptive regularities in development of many species. Until these achievements of genomics in developmental biology there were no real explanations whatever, and a good deal of agreement about what such explanations had to explain, in particular developmental recapitulation, serial homology, and phylogenetic novelty, features that bear important similarities to the productivity, systematicity, and creativity that computationalism demands any cognitive theory explain. The bottom-up approach of molecular developmental genetics has provided the required explanations in terms of computationally realizable programs of development. And it has shown that there no generalizations to be uncovered at the level of organismal development, only descriptions of phenomena that needed to be explained.

The precedent of molecular developmental biology and the conspicuous initial success of a computational neurogenomics are part of a strong case that the most

promising line of research in cognitive science is bottom up. If the best case scenario is on the cards, then a bottom-up approach should vindicate it as the systems biology of the nervous system discovers motifs at successively higher levels of aggregation of synaptic structures. Such successes will suggest experiments in humans and infrahuman species that could narrow down the representational states to which the sets of synaptic structures approximate. If on the other hand, no such vindication of a computationally inspired cognitive science is on the cards, time, money and genius will not have been wasted trying to develop a top-down theory that does not exist.

Alex Rosenberg

Center for Philosophy of Biology

Duke University

References

- Alon, U., 2006, *Introduction to Systems Biology: Design Principles Of Biological Circuits*, Boca Raton, CRC.
- Bailey, C., Kandel, E., Si, H., 2004, "The persistence of long term memory; a molecular approach to self-sustaining changes in learning-induced synaptic growth," *Neuron*, 44, pp. 49-57.
- Brown. J.R., H. Ye, R.T. Bronson, P. Dikkes and M. Greenberg, 1996, "A defect in nurturing in mice lacking the immediate early gene fosB." *Cell* 86, pp. 297–309.
- Chalasanani, S.H., Chronis, N., Tsunozaki, M., Gray, J.M., Ramot, D., Goodman M.B., and Bargmann, C.I., 2007, "Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*," *Nature* 450: pp. 63-70
- Davidson, E. H., 2006, *The Regulatory Genome*, New York, Academic Press
- Fodor, J, 1991, *The Theory of Content*, Cambridge, MIT Press
- Gingrich, J., and Hen, R., 2000, "Commentary: The broken mouse: the role of development, plasticity and environment in the interpretation of phenotypic changes in knockout mice," *Current Opinion in Neurobiology*, 10, pp. 146-152
- Ha, T.J., Kohn, A.B., Bobkova, Y.V., Moroz, L.L. 2006, "Molecular characterization of NMDA-like receptors in *Aplysia* and *Lymnaea*: relevance to memory mechanisms," *Biological Bulletin* 210(3): 255-270
- Kandel, E., Bailey, C., Bartsch, S., 1996, "Toward a molecular definition of long-term memory storage," *Proceedings of the National Academy of Science*, 93, pp. 13445-13452
- Rosenberg, A., 2006, *Darwinian Reductionism or How to Stop Worrying and Love Molecular Biology*, Chicago, University of Chicago Press
- Ryle, G., 1949, *The Concept of Mind*, Hammersmith, Penguins.
- Searle, J., 1980, "Minds, Brains, and Program", *The Behavioral and Brain Sciences*.3, pp. 417-424.
- Silva, A., Brown, R., Talton, A., Wiltgen, 2007, "New Circuits for Old Memories: The Role of the Neocortex in Consolidation. *Neuron*, 44, 101-108

William T. Swaney W. , Curley, J.P., Champagne, F.A. and Keverne, E.B. 2007, “Genomic imprinting mediates **sexual** experience-dependent olfactory learning in male **mice,**” *Proceedings of the National Academy of Science*, 104, pp. 6084-6089

Yang Zhou, Hao Wu, Shuai Li, Qian Chen, Xue-Wen Cheng, Jing Zheng, Hiroshi Takemori, and Zhi-Qi Xiong 2006, “Requirement of TORC1 for Late-Phase Long-Term Potentiation in the Hippocampus,” *PLoS ONE*. 2006; 1(1): e16. Published online 2006 December 20. doi: 10.1371/journal.pone.0000016