

Defending Information-Free Genocentrism

In this paper I defend genocentrism, the thesis that the genes have a special causal/explanatory role in embryological development I do so while granting that the genes do not carry information about development in any interesting sense. But I also show that the supposition that they do is not required by genocentrism. What genocentrism requires is that the genes program the embryo in the same way that the hardware in an automobile assembly plant's robots' central processing units realize a program that enable them to weld chassis without human supervision and more accurately than any human can perform this task. This latter thesis requires no philosophical defense. It is obvious to any one who examines the available accounts of the embryological development of, say, the *Drosophila*. The mistake is to suppose that programming requires information.

The debate between genocentrist and their opponents has been focused on the question of whether the genes bear information. The opponents argue that nothing in the biosphere has the status of representing development or hereditary-guiding information, including the genes. *Ergo*, claims based on its unique informational role are unwarranted. At most, all the talk about genes as information, as programming development, as blue-prints or recipes for the body, as sending signals, or being proof read, decoded, and for that matter as 'master-controls' is so much mere metaphor, and seriously misleading metaphor at that.

Here the exponents of the genetic program explanation of development have responded that calling the DNA sequence a coded message is no metaphor, that the genes carry information in literal sense not shared by other causally necessary conditions for development.. Genocentrist will grant that there is an attenuated sense of information in which many other things besides the genome carry information about development and hereditary traits. This is the sense of information as it figures in mathematical information

theory. It is true that the mathematical theory of information due to Shannon and Weaver applies to any causal chain, including ones in the environment transmitting epigenetic “hereditary” transmission, just as well as it does to genetic transmission of hereditary information. This is simply because the Shannon-Weaver formalism is a way of measuring the quantity of information and the reliability of transmission which any causal chain can be employed to transmit. But the transmission of information requires more than a causal chain. Otherwise the casual chain which transmits a scrambled TV signal would count as transmitting information (beyond the “information” that the signal was scrambled). And of course the opponents of genocentrism will grant that genes bear Shannon-Weaver information as well.

One reason genocentrists maintain that genes bear real (and not just Shannon-Weaver) information is that they express protein structure in an arbitrary code just as sentences express facts about the world in an arbitrary code [Maynard Smith, 2000, following Monod, 1971]. The nucleotide sequence of Adenine, Thymine, Guanine and Cytosine (hereafter A,T,G,C) constitute a code in which the structure of proteins is expressed, and from which that structure can be read off. That it is literally a code is reflected in several facts about it: for one the genetic code is redundant. Though with four units one can send 2^6 (i.e. 64) different bits of information, only 20 bits—the 20 amino acids are actually encoded, and some amino acids can be signaled by any one of three different messages (three nucleotide “codons”). Histidine’s codons are CAT and CAC, for example. Thus, the genetic code is informational in a sense nothing else in development is informational, in large measure because, like a signal system (say Morse code) it is arbitrary. It is held to be arbitrary in the following sense: as a matter of physical possibility, any particular triplet, say CAT, which codes for histidine, could just as well have coded for the amino acid Glutamate. That is, we

can imagine a process of protein synthesis that attaches a Glutamate molecule to a transfer RNA with the CAT codon instead of a Histidine molecule. There is of course an explanation for why CAT codes for Histidine and not Glutamate, and that explanation involves natural selection and drift operating on the initial conditions obtaining at the time the coevolution of nucleic acid/amino acid began. It will presumably show that the actual coding of 20 amino acids by a particular redundant pattern of 64 nucleic acid codons, is a “frozen accident” [Crick, 1968]: i.e. it could, consistent with the laws of chemistry and physics, have turned out that a different coding emerged. Can epigenetic codings make the same claim of physical arbitrariness? That is, would it be the case that, holding the laws of nature constant and merely changing the initial conditions at the time that an epigenetic hereditary mechanism kicked in, have resulted in the epigenetic transmission of the same traits? Almost certainly not.

Unfortunately, it may be difficult to show positively that the code which translates CAT into histidine is really arbitrary. What exactly does ‘arbitrary’ mean in the present connection? Compare the English word ‘cat’ to the English word ‘ouch’. There is a full and complete explanation of why ‘cat’ names cats, in English. But, consistent with all the laws of nature, ‘cat’ could have carried information about hats or bats or cots or cabs or casts, etc. That in general it carries information about cats is presumably fixed by some laws or other operating on some set of initial conditions, facts about fairly local conditions, that obtained on the earth at some time in the philological past. By contrast its onomopoeitic sound suggests that ‘ouch’ carries information about pain somewhat non- arbitrarily. Consistent with the laws of nature, it could have meant what ‘ah’ means. That it carries information about pain is still a matter of initial conditions, but perhaps less local ones than the ones that resulted in ‘cat’ meaning cats. By contrast clouds carry information about rain independent

of any local conditions: the general connection is explained by laws alone, and initial conditions are not involved. Roughly, then, a generalization about meaning is arbitrary, in the present context, since its explanation depends in part the role of initial conditions in the explanation of the connection. The more local the conditions, the more arbitrary the connection expressed by the generalization.

The trouble is that among the competing explanations for the genetic code several either do not require initial conditions at all, and others make only the most minimal reliance to such conditions. At least five or six accounts of the origin of the genetic code have been offered. One of them, the stereochemical theory of the origin of the code requires as explanans only the laws of chemistry and physics to bring about the particular distribution of codons to amino acids that figure in the genetic code as we know it. Indeed, this theory even provides a scenario according to which there is a natural sequence of chemical events that builds derived amino acids out of basic ones and changes existing codon-sequences in step with the amino acid syntheses to build up the genetic code. If the theory is right, it gives the code a non-arbitrariness and inevitability which approaches the periodic table of the element's association between atomic structure and chemical properties.

Rival evolutionary accounts trade on features of the code noticed soon after its discovery: in general structurally similar amino acids share more similar codons, and amino acids with pyrimidine (T and C) nucleic acid based codons are more similar to one another than ones coded for by purine (A and G) bases. Accounts trading on these facts suggest that the code would develop in the particular order it did owing to the operation of natural selection wherever the set of 20 amino acids, and 4 nucleic acids interact. [Jukes 1985.] The selectionist explanations for the code have been strengthened recently their ability to explain

why certain exceptions to the code have appeared, for example the shift in the codon CUG from coding for leucine to serine in certain *Candida* species [Santos, 1999].

If the distribution of the chemical substrate on which natural selection operates is almost everywhere in the universe the same, and if there is a purely chemical affinity between nucleic acid codons and amino acid molecules, then we can expect the same genetic code to develop ubiquitously. Even if chemical affinities do not result in just one code, natural selection operating at the level of macromolecules could choose among competing codes. The fact that the code is nearly ubiquitous among all biological systems on Earth is good evidence that if it is the result of evolution, it must have won the competition among alternative codes so completely that no competitor was left in the field. Since natural selection among vast numbers of molecules over evolutionary time scales allows only vanishingly small scope for drift, the conclusion that the code is more than a frozen accident is not surprising, and neither is the persistent search among molecular biologists for an explanation of the code that reveals its non-arbitrary character.

The explanation of the code is a question on which much current research is devoted among molecular biologists, and for all we know the “frozen accident” [Crick, 1968] view which would vindicate arbitrariness may be correct. But pending the outcome of this debate the exponents of DNA’s distinctive informational role cannot help themselves to arbitrariness as evidence of the genetic code’s informational character.

Moreover, there is another more significant requirement that must be satisfied by the genes if they are literally to be said to carry information in a way that the rest of the causally necessary conditions for development and heredity do not carry information. At a minimum the three nucleotides of a codon must carry information about a particular amino acid. The codon has to be “about” the amino acid in the way that ‘cat’ is “about” felines. A codon’s

being about an amino acid, is a matter of the codon's containing information that, for example, the amino acids transfer RNA's anticodon should have a certain sequence. The 'aboutness' or 'content of a codon is the requirement that the informational state of the nucleic acid sequence must be, in the philosopher's argot, "intentional". Molecular biology is of course riddled with intentional expressions: we attribute properties such as being a *messenger* ("second messenger") or a *recognition* site, we ascribe "*proof reading*" and '*editing*' capabilities, and say that enzymes can *discriminate* among substrates (as when "synthetase *avoids* hydrolyzing isoleucine-AMP, a *desired* intermediate"). Even more tellingly, as we have seen, molecular developmental biology describes cells as having "positional information", meaning that they *know* where they are relative to other cells and gradients. The naturalness of the intentional idiom in molecular biology presents a problem. The naturalness of this idiom in molecular biology cannot be written off as mere metaphor, at least when it comes to information in the genome. The claim manifestly cannot be merely metaphorical, not at any rate if the special role of the gene is to turn on its informational content. But to have a real informational role, the genome must have intentional states. So at any rate both the genocentrists and their opponents agree [compare Griffiths, 2001 and Maynard Smith, 2000]

Exactly what is intentionality? Intentionality is best introduced by examining a paradigm case of informational content, such as belief-states. If Lois Lane believes that Superman was born on Krypton, then his state of belief "contains" the proposition that Superman was born on Krypton; the belief is "about" Superman, and attributes to him a property, that of being born on Krypton. The example uses Superman in order to reflect the fact that what a belief is about need not exist, nor need it attribute a property that anything actually does have. What most clearly reflects the intentionality of beliefs, their containing

propositions and being about or “directed at” objects which may or may not exist is an extremely interesting logical feature they all share. Suppose that Lois Lane believes that Superman was born on Krypton. Now substitute for ‘Superman’ the words ‘Clark Kent’. Now consider the claim that Lois Lane believes that Clark Kent was born on Krypton. This claim, derived by substituting ‘equals for equals’—Clark Kent for Superman—is presumably false. Here is another way a statement about what Lois believes can be changed from true to false just by substitution of equivalents. Since Krypton is the only planet which blew up just after Superman’s departure from it as an infant, we can substitute for ‘was born on Krypton’ ‘was born on the only planet which blew up just after Superman’s departure as an infant’. Since Lois presumably doesn’t know anything about the history of Krypton, the resulting statement that she believes superman was born on a planet that blew up just after his departure is false. By making an innocent substitution in the contained statement (Superman was born on Krypton) that preserves its truth we have changed the containing statement (Lois believes that Superman was borne on a planet which exploded...) into a falsehood. (Note the fact that Superman, Lois Lane and Krypton are all fictional items, is irrelevant to the points made here. The same conclusions are illustrated by Cicero and Tully, the Morning and Evening Stars, and other co-referring terms familiar to philosophers since the work of Frege and Russell). Beliefs of course are not the only kinds of things that bear intentionality. Among other things, inscriptions do as well, including those which represent other things, and thereby convey (semantic) information. Moreover, this feature of intentional states is not the whole story about them, but for our purposes it may be enough. For now we have a useful test of whether a state is intentional or not.

If the state, of one thing is *about* another thing, or *represents* it, or has informational content, then the truth or falsity of a description of the representing state should be sensitive

to the way its content—the represented state-- is described. In the case of beliefs, desires, hopes, fears, plans, actions, and other human psychological states and their effects in action, the source of intentionality and of the sensitivity of these states and actions to how they are described is of course due to their all being states that rely on thought. Since thinkers are not omniscient, there will be lots of ways of describing their objects of thought that thinkers don't recognize as true of their objects of thought, and so we get the sensitivity to substitutions of descriptions that is the hallmark of intentionality.

Now, for nucleotide sequences to carry intentional information, as those who accord it an informational role in development require, the descriptions of the information the sequences contain must be sensitive to the terms in which we describe that information. When we say that CAT means histidine, or refers to histidine, or carries information about histidine, it will have to be the case that there are some ways of describing histidine which we could substitute for the words 'histidine' in the statement that 'CAT means histidine' or 'CAT represents histidine' or 'CAT is about histidine' that would convert the statement from true to false. But there are no such descriptions: consider the following attempts to produce such a falsehood. They will be bizarre, perhaps even funny, certainly without scientific interest. But none of them will be false. 'CAT' means the only amino acid spelled with an initial 'h' in English; 'CAT' represents Francis Crick's favorite molecule; 'CAT's' informational content is about molecule whose chemical structure is symbolized like a baseball home plate with a Christian cross in the right handed batter's box.

There is a more general way to see the problem if we consider how advocates of the informational character of the code argue for its having intentional content. The distribution of black marks on the white paper 'cat' means the domesticated feline owing to our endowing the marks with meaning, *interpreting* them as a sign for cats. Roughly speaking,

the intentionality of the marks on paper 'cat' is *derived*: it gets its intentionality from our beliefs about the references of particular English words and our desires to communicate our thoughts about cats to others. What gives the codon CAT its meaning? The standard argument is that natural selection does so [Maynard smith, 2000]. That is, CAT has been selected to mean histidine by a schedule of variation and selection which lead to the ubiquity of the code in which CAT is histidine codon. Notice this claim assumes that several theories of the origin of the code are false, e.g. the stereochemical theory, and it may be incompatible with the claim that the code is arbitrary. Let's leave these problems aside in what follows. There are more serious problems for this proposal about where 'CAT' gets to mean 'histidine.'

So, CAT's functional role in protein synthesis, which is the result of natural selection, confers upon individual CAT codons their intentionality. This intentionality is presumably not derived, but *original* intentionality. For presumably 'mother nature' doesn't have its own "free-floating" desires, beliefs, or other intentional states that could interpret CAT as meaning histidine. To suppose otherwise is to accord to nature, or whatever it is that natural selection depends on, the very sort of mentality, purpose or design that Bishop Paley hoped for and Charles Darwin expunged from nature. So, the question that needs to be addressed is whether and how natural selection can produce intentionality. And on this question, the jury is, so to speak, *in*. It cannot. And the reason is that the functions which natural selection accords to structures—molecules, organelles, cells, tissues, and for that matter organs and organisms, etc —don't produce the sensitivity to alternative descriptions of their functional "content" which intentionality requires.

To see why consider a case where content seems uncontroversially attributable. The brain of the frog which has been programmed by natural selection to cause the frog to flick

its tongue out in exactly the right direction and exactly the right time to trap a fly at location (x,y,z,t) . Can we attribute to the state of that brain, or some component of it, perhaps a few thousand of its neurons, the intentional content 'fly at x,y,z,t '? One would think so. After all, that brain state was selected for, phylogenetically in the frog lineage, and ontogenetically in this particular frog because it led to environmentally appropriate fly-flicking. But consider, as a result of the same history of selection, the same brain state will also cause the frog to flick its tongue out at a bee bee at x,y,z,t . So, the brain state's content must be 'fly or bee bee at x,y,z,t '. And of course it will also stick out its tongue at a black current at x,y,z, t . So, is the brain state's intentional content about the disjunction, fly or bee bee or black current? Of course our imagination is the only limit on our ability to further and indeed endlessly expand this disjunctive list of what the content of the frog's brain state is. In other words, by making truth preserving substitutions for 'fly' in the sentence 'there is a fly at x,y,z,t ' we cannot convert to a falsehood the statement that 'the frog's brain state contains "fly at (x,y,z,t) ." And of course there is another problem, fly means the same as 'mouche' in French, *Drosophila* in Latin, etc. So we can substitute the word for fly in any human language into the sentence which gives the content of the frog's brain state, without changing the claim that the brain state contains the statement that there is a frog at x,y,z,t from a truth to a falsehood (compare Lois Lane's beliefs about Superman). The upshot is that if natural selection does accord the frog's brain state with content, it is not *intentional* content. And of course the same argument can be advanced for the claim that the codon 'CAT' is about, means or represents histidine. The codon's content is not intentional. So, it's not informational in the required sense.

There is another way to see that the genetic program that builds the *Drosophila* embryo does not bear information in the way required to differentiate it from other factors necessary

for development. This way trades on Searle's [1979] "Chinese Room" argument alleging to show that "Strong Artificial intelligence's" conception of cognition as following a program is a puerile mistake. Searle invites us to examine a fluent English speaker sealed in a room with a mail slot through which ink blots are slid. The prisoner follows instructions written in English about which ink blot to put through the mail slot for any given entering ink-blot. Unbeknownst to the prisoner, what he thinks are ink blots are Chinese pictograms and Chinese speakers outside the room interpreting them as such, conclude that there is a fluent Chinese speaker in the room participating in a written conversational exchange with them. As Searle points out however, there is no one in the room who understands Chinese, who thinks he is in a conversation, or who has any interest in communicating anything to anyone by means of the cards he slides through the mail slot. All he is doing is following a program. Therefore following a program is not sufficient for or even interestingly necessary for understanding Chinese.

There are many responses in the philosophical literature to Searle's argument against the claim that intentional cognition is a matter of following a program. And the defenders of the informational role of the genetic code had better concern themselves with these replies. For they will need them. Searle's argument certainly shows that following a program is not sufficient for intentionality in the sense required for the transmission of information.

However, opponents of genocentrism should not rejoice prematurely. It turns out that for their purposes Searle's argument may have proved too much. In fact gives us a way to see the genetic program's special role in development even while it has at most derived intentionality, derived that is from our interpretation, and does not transmit information, at least not independently of us and our interpretations. The genocentrist, it runs out, can claim

for the genes the special role in development they require without needing to argue that the genes bear information at all.

Searle rightly notes that we compute, using computers do so, and we accord content to their inputs, outputs and the programs they run by interpreting them as representing cognitive states of our own which have underived intentionality. The soft ware/hardware distinction familiar to us reflects the fact that many physically different pieces of hardware can run the same programs. Apples and PCs, for example, both run Word for Windows programs, but “realize”, “instantiate”, “implement” these programs on physically different types of microprocessors, and other pieces of hardware. They do so in part because higher level programs like Word for Windows get rewritten in the proprietary assembly languages that differed from manufacturer to manufacturer of the PCs we employ. Indeed, for the purposes of high speed intensive calculation of the sort required to deal with the most difficult problems a high level program can be implemented on, say a single Cray super computer and on a system of PCs running in parallel. The architecture and the components of the hardware will be quite different, even as some of the programs running on them are the same.

In recent years computers have been used to help complete proofs in mathematics for theorems that otherwise would have remained in doubt: the four color theorem and Fermat’s theorem are two examples that spring immediately to mind. Here, of course, following Searle, we cannot say that the computers proved lemmas or intermediate theorems needed by the mathematicians proving the theorems. Rather, the mathematicians used the computer programs, interpreted the computer inputs and outputs in such a way as to themselves prove the theorems (theorem proving is an intentional act). One interesting fact about even the most powerful of current computers is that there are some computational

problems which are too complex for them to be employed to solve. Among these are calculations involving the so-called NP-hard problems (NP from Nondeterministic Polynomial). Among the most famous of these NP-hard problems is the traveling sales man problem: for any finite number of cities to be visited by a traveling salesman, given the distances between them, find the route that minimizes travel distance, while enabling the sales man to visit them all and return to his starting point. There is no general solution to this problem of which is the shortest route. Once the number of cities grows to more than several dozen, finding the answer to this question will require more computer power and time than is available on even the fastest super-computers of the largest network of parallel computers employing deterministic programs. Approximation methods on non-deterministic (parallel) computers must be employed to solve this problem and there is no implementable algorithm for establishing that the solution derived is optimal.

What does all this have to do with the genetic code? Well, it has been shown that when the number of cities to be visited grows beyond 50 or so, a computer composed of strands of nucleic acids can provide a better approximate solution to this problem faster and cheaper than any silicon chip microprocessor based computer. Of course both computers will run the same high-level program (though their “assembly languages” will implement the higher level program differently), but the DNA computer produces a result faster and more cheaply. The high level program consists in the following five steps: Suppose the number of cities to visit is n . First one generates a large number of random paths through some or all of the n cities. Then, screen the set of paths, eliminating any path that does not return to the starting point, then eliminate all sets of paths which do not go through some cities n times, and then eliminate all those which miss any city, and finally eliminate those paths which do not enter all cities at least once. Any remaining paths will be solutions,

though perhaps not optimal solutions to the problem. It sounds simple, but the combinatorial possibilities rapidly become unmanageably great. The fastest supercomputers can perform 10^9 calculations per second and this is too slow to provide an answer when the number of cities begins to exceed 70. But a DNA computer can be constructed out of strands of molecules which can be combined, preferentially anneal with one another, amplified by polymerase chain reaction, cut by restriction enzymes, ligated at their sticky ends, and detected by electrophoretic techniques. By the mid-nineties, molecular biologists and computer scientists had designed DNA computers which can solve simple versions of the traveling sales man problem and can in theory solve complex ones that are beyond the powers of the silicon based super computer. This should of course not be surprising. Molecular interactions enable the DNA computer to perform 10^{14} calculations per second, a five fold improvement on conventional microchip computers, at an energetic cost which is a staggering 10^{10} lower than that of a contemporary silicon based computer, while storing data at a density of approximately 1 bit per cubic nanometer, compared to existing storage media which record 1 bit per 10^{12} cubic nm. It's no surprise therefore that a DNA computer can provide results of equal accuracy with extraordinarily faster, and more economical means than conventional computers (see Adleman, 1994).

The upshot is obvious. If a DNA computer can implement the same program that a silicon-chip based computer can, then in whatever sense a computer follows the program, the DNA computer does too. If a DNA computer can more effectively and efficiently deal with large number instance of an NP-hard problem than any silicon chip based computer—series or parallel, super or not, then surely it can deal with a simpler problem such as implementing the program that builds the *Drosophila* embryo. How can we be confident of this? Well, the program that the DNA implements can also be run as a Pascal Program on an

Apple Mac or a PC lap or desk top. And of course, if it is the same program that the Apple runs to produce a print out of successive stages of the embryo that the embryo's genome itself runs to build successive stages of the *Drosophila*, then they must literally both be running the same software. Recall, that many physically different computers hooked up to quite different in-put and output peripheral devices can run the same programs.

But the silicon-chip based computer is an information storage and processing device, owing to our interpreting its inputs, outputs and internal states as having derived intentionality, like any other artifact. The *Drosophila* genome was not constructed by us, but by evolution, and evolution, I have argued, cannot confer intentionality, derived or underived. Ergo, the objection concludes, the *Drosophila* genome does not bear information, whether it is literally a program or not.

This objection misses several points. First, it needs an argument to show that the causal origin of a computer is relevant to its functioning as one, i.e. to our according the program it runs as having derived intentionality. Second, we may even grant that before the revolutionary developments in molecular biology no cognitive agent knew the program which the *Drosophila* genome runs, and it is only in the last decade or so that we have been able to accord the program derived intentionality. This means however, that long before its first interpretation by molecular biologists, the program the genome implements could have been accorded a kind of possible derived intentionality, or capacity for derived intentionality, and this is a counterfactual property that apparently nothing else in the developmental life-cycle of the *Drosophila* could have had.

For those, like me, to whom counterfactual properties provide cold comfort, there are more powerful considerations to mount against the argument that the genomes recent acquisition of derived intentionality disqualifies it from a unique role as the program of

development. First, notice that derived intentionality requires original intentionality in us. And if natural selection is the process that made our brains, it is hard to see where they got their original intentionality. The problem of naturalistically explaining the original intentionality of the human (and infra-human) brain is perhaps the most serious fundamental challenge facing neuroscience and its philosophy. No one has yet solved it. Pending the solution it would be unwarrantably complacent to help oneself to the problematical notion of a mysterious underived original intentionality in us in order to argue that the *Drosophila* genome cannot have the derived intentionality of our artifacts. Indeed, if human intentionality turns out to be derived from some evolutionary process as yet unimagined (and it will have to be unimagined so far if it is to prove “unmysterious”), it will turn out that both artifacts and genomes will be on a par, neither deriving their intentionality directly from something with underived intentionality, and both tracing their intentionality back to evolution by natural selection. Of course, as I indicated above, I am dubious that natural selection can actually produce original intentionality in the brain or anywhere else, and so it cannot produce derived intentionality either. Both will, on my view, turn out to be illusions, like the purposes we overlay on nature and which natural selection has dispelled. But this is another story in the philosophy of psychology and we need not pursue it further for present purposes.

The crucial question is not intentionality but programming! Whether intentional or not, what seems hard to deny is that the *Drosophila* genome programs the embryo. The only real issue is whether it is unique in doing so. Opponents of the genocentric approach to development will want to deny this. The argument given above, that the genome has no privileged informational role, will not support the denial that it programs the embryo, since

an (underived) informational role is evidently not required for the genome to program the embryo.

Of course we can exploit almost any regular natural process to compute some function or other. Thus, by employing Ohm's law, we can compute products and quotients given readings of voltage, wattage and resistance. We can use the turning of leaves from green to red in the autumn and the rings the tree trunks form to compute seasons and years. Accordingly, it is not beyond the bounds of human ingenuity to employ regular biological processes besides those implicating the genes to solve some computational problems. These processes too will therefore turn out to be programs. If there are regularities relating non-genetic factors to developmental outcomes, then presumably these too may be construed as or actually employed as programs. And of course we are familiar with such non-genetic factors—the epigenetic ones such as methylation or the finch song on which widowbirds imprint. Will epigenetic regularities we are ingenious enough to employ to effect certain computations, almost always ones that we could make much more easily, suffice to deny the genome its claim to uniqueness as the program for development? Well, it will certainly be sufficient for the developmental systems' theorist, or others who wish to deny the genome an in-principle uniqueness. But I dare say it will not move the scientist, either the computer scientist or the molecular biologist. The reason is obvious. DNA computation is a practical reality, the basis of a device with profound advantages over silicon-chip based computers, which can be programmed and reprogrammed, to realize any algorithm silicon based computers can implement.

Even if we grant that there are generalizations about some biological systems regular enough to realize some mathematical function or other, and thus to be employed by us as computers, no one could reasonably describe these systems as capable of implementing a

program of any real flexibility, power, and computational utility. Indeed, it could be argued that merely realizing a mathematical function that we could (though never would actually) employ to compute some ordered pairs of numbers is insufficient to mark something out as implementing a program. For some physical configuration of matter to constitute a systems that implements a program, it should non-trivially and literally instantiate the “hardware/software” distinction. That is, we should be able to identify other physical configurations that will implement the same mathematical function, reflecting the multiple realizability of programs by physically different processors; and we should be able to identify other programs which changes in the original physical configuration will enable it to realize. Notice that the program of *Drosophila* development for example, as we now know its details, satisfies both of these requirements. Doubtless, clever philosophers will be able to identify a small number of cases elsewhere in the biosphere which can also do so. There are two things worth bearing in mind about such cases. In order to be reliable and useful in real computation they will have to reflect biological regularities that operate with close to invariable regularity and at fairly high speeds. *Ergo*, they will be about molecular interactions, not the sort of factors that opponents of genocentrism are likely to attempt to put on a par with the genome (e.g. Beaver dams or finch songs). They will not be anything like as ubiquitous in the biosphere as the hardware/software configuration of the genome and the genetic code. The genocentrist can accept such examples of other programs with equanimity, for the assertion that the genome is unique in its programming of all development cannot be undercut by the recognition that there may be other programs operating at the molecular level, and even one or two such molecular programs also involved, along with the genome, in the programming of development.

References

- Crick, F., 1968, "The origins of the genetic code", Journal of Molecular Biology, 38: 367-379.
- Griffiths, P., 2001, "Genetic information: a metaphor in search of a theory," Philosophy of Science, 67: 26-44.
- Jukes, T. 1985 , "A change in the genetic code in *Mycoplasma capricolum*", Journal of Molecular Evolution, 22, 361-362
- Maynard Smith, J. 2000, "The concept of information in biology," Philosophy of Science, 67: 177-194
- Monod, J. 1974, Chance and Necessity, London, Fontana
- Santos, M, .et. al , 1999, "elective advantages created by codon ambiguity allowed for the evolution of an alternative code in *Candida* spp.", Molecular Microbiology, 31, 937-947
- Searle, J., 1980 "Minds, brains, and programs," Behavioral and Brain Sciences, 3: 417-457.