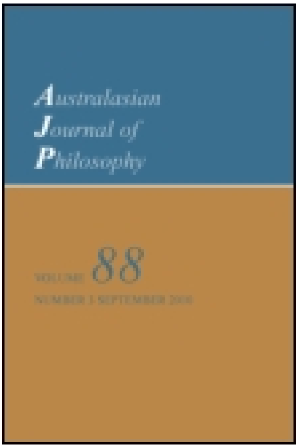

Australasian Journal of Philosophy
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/rajp20

Critical notice

D.D. Gamble [a]
[a] University of Adelaide
Version of record first published: 02 Jun 2006.

Australasian Journal of Philosophy
Vol. 70, No. 3, September 1992

CRITICAL NOTICE

Cummins, Robert, *Meaning and Mental Representation*, Cambridge, MA: Bradford Book/MIT Press, 1989. A$17.85 (paper).

Many philosophers, Robert Cummins argues in *Meaning and Mental Representation*, abuse the computational model of mind, driven by a desire to exploit the computational notion of representation to explain intentionality in belief-desire psychology. Resulting theories, Cummins maintains, are incompatible with computationalism, or non-explanatory or both. Either way, the very legitimacy of the 'computer hypothesis' is undermined.

J.A. Fodor is a philosopher deemed guilty of this wrong approach to mental representation. He is criticised for treating representations as the neutral givens of computation, the problem of intentionality just being that of attaching intentional contents to these givens.

Philosophers guilty of the wrong approach attempt to slide from the constraints of legitimate computation *qua* function-simulation to illegitimate 'computational belief-desire psychology', exploiting the computational notion of representation and riding on the occurrence in both contexts of rationality and epistemological constraints. Desiring to explain intentionality philosophers grasp at a scientifically respectable notion of representation quite blind to its foundational incompatibility with their theoretical requirements. A good dose of philosophy of science is needed to cure philosophers of such erroneous ways.

The recommended cure goes like this. If you want a scientific theory of 'x' then find some well established scientific framework where 'x' is used and analyse its use and nature in that context. The results must constrain any future theorizing about 'x'. Specificity of context is important so for the present case of representation Cummins restricts consideration to empirical theories of cognition which are based on classically computational principles (not connectionist). The Computational Theory of Cognition or CTC is used as an umbrella term to denote actual theories (e.g. of perception, psycholinguistics etc.) in use by practicing cognitive scientists. Hence philosophers must pursue representation via an analysis of the construct in use in CTC theories.

The question of mental representation is held to be a question in the philosophy of science analogous to the question of physics 'What must we suppose the nature of space to be . . . if General Relativity is to turn out to be true and explanatory?' (p. 13). Analogously, 'What must mental representation be like if CTC theories are true and explanatory?' CTC assumes cognition can be explained computationally. That is, its approach is determined by supposing that

(i) computational objects are symbols representing the arguments and values of cognitive functions; (ii) systems have cognitive capacities in virtue of being able to compute the functions specifying those capacities, (p. 148). Since on CTC cognition is computation, CTC is explanatorily *grounded* in computation. Hence the answer to the methodologically sound question is a *computer science* answer derived from the fundamentals of computation in noncognitive contexts. The answer is *S-representation*. This is what philosophy of science reveals as the basis of a science of computation and so as the basis of a scientifically respectable theory of mental representation.

Cummins' dialectic is as follows:

(1) Computation explains cognition, (CTC).
(2) Computation is essentially a form of function simulation/instantiation.
(3) To explain computation is to explain *how* its an event that can be understood as function simulation/instantiation.
(4) S-representation is integral to understanding *qua* (3).
(5) The central assumption of CTC is that the *same* notion of explanation carries over from calculators to cognition.
(6) Thus the same notion of representation and semantics applies to calculators and cognition: 'It is an absolutely central thesis of the CTC that representation in cognitive systems is *exactly the same thing* as representation in computational systems generally' (p. 118, my emphasis).
(7) S-representation is not equipped to explain intentionality.
(8) Any alternative semantics imposed on CTC is incompatible with CTC, nonexplanatory, or both.

Little will be said here about (7) and (8) though Cummins' arguments are not very convincing on that score. Instead I will indicate how unattractive S-representation is to many philosophers of mind and take issue with Cummins' improper mandate to a style of Artificial Intelligence (AI) to dictate theories of mental representation to philosophers. Such a mandate is what (6) above comes to.

Introducing S-Representation

'The CTC proposes to explain cognitive capacities by appeal to representation and computation in exactly the way that arithmetical capacities of calculators . . . are standardly explained . . . '

CTC according to Cummins considers cognition to be a special class of function simulation. It is special only because cognitive functions are special. Following the rationalist tradition, cognitive functions are essentially rationally-epistemologically constrained functions, i.e. input-output patterns, or arguments and values, which individuate specific capacities or competences, are normatively and semantically evaluable in a way that makes sense with respect to the competence manifested. Normative-semantic evaluation implies assessment as rational, truth-bearing etc. relative to some domain. Hence cognitive functions are the subclass of functions marked by propositional, inferential specifications. Inferential relations between propositionally specified values and arguments being the only suitable vehicles for normative-semantic evaluation. Such specifications are thus the distinctive feature of cognitive as opposed to noncognitive functions.

A calculator is a simple computer artifact. Calculators *qua* computers belong to the broad class of function-simulators. The functions they simulate are arithmetical, not cognitive. On Cummins' (5) and (6) CTC assumes univocal notions of explanation and representation from calculators to cognition. The crucial notion of explanation is the 'Tower Bridge' picture. S-representation emerges as a construct from this picture.

A calculator's capacity to add is explained by the Tower Bridge picture. Addition is a function from numerals to numerals. A calculator can be described physically. Externally one can specify a button-pressing-to-display function; (observable inputs and outputs). Internally one can specify states correlated with and mediating button-pressing-to-display functions. Call a button-pressing-to-display function a 'g' function. A *g*-function is *satisfied by* the internal states which mediate it: 'the arguments and values of *g* are literally physical states of the physical system' (p. 90). 'Bridging' is effected — between physical description and mathematical function — by interpretation. Bridging occurs when components in virtue of which the *g* function is satisfied can be viewed under interpretation as representations of numbers which are arguments and values of addition. Cummins calls this 'S' (for simulation) representation. Bridging allows one to view the bottom span (physical) as simulating or instantiating, hence representing, the top span (mathematical). The two 'spans' are joined forming a bridge by interpretation. Interpretation or S-representation is thus integral to grasping a physical event as the instantiation of an abstract function, which is to grasp the event *as* a computation, which is to give a computational explanation of the event (in this case, addition).

S-representation emerges from the Tower Bridge story via interpretation. 'We call an object of computation a representation when it is important to see the computation as an instantiation of something else' (p. 110). Explanation is incomplete until it is shown *why* a physical system satisfies *g*. It does so because it is programmed to. That is, it's designed so that patterns of its structured state transformations under interpretation parallel, are isomorphic to, or track, patterns of transformations identifying a target function. That is, in the calculator's case, the causal structure of the system is devised to mirror the arithmetical structure of the target function.

Now consider the Tower Bridge extended to cognitive functions. The Tower Bridge allows us to see physical structure as computational in virtue of mirroring cognition understood as rational or inferential structure. S-representation in the context of cognitive, hence epistemologically constrained hence semantically evaluated functions is a version of Interpretational Semantics. 'The thesis of Interpretational Semantics is that S-representation is the kind of representation the CTC requires to ground its explanatory appeals to representation' (p. 88). Components in virtue of which cognitive *g* functions are satisfied may be viewed under interpretation as representations of semantic contents.

It is important that we have a clear picture of S-representation. First, something's being a representation is *constituted by* the fact that it is an aspect of a structure grasped as a simulation: 'representation is just a name for the relation induced by the interpretation mapping . . . ' (p. 93). Anchored thus in the pragmatics of modelling theory S-representation is a pragmatic and relative phenomenon. What a structure S-represents (to some degree or other) depends on what an interpreter sees it as representing which is shaped by what (s)he's aiming to get a structure to simulate and how successful (s)he's being at the task. Second, S-representation is global and procedural in nature. It is evaluated, modified, successful or unsuccessful on a global basis. It's tracking over time in a coherent manner that counts: 'successful representation is not simply a matter between a symbol and its interpretation (p. 98) . . . What is S-represented is essentially a matter of the processes . . . ' (p. 100). Finally, global coherence does not dictate unique interpretation. The one dynamic physical structure or event may successfully be viewed as simulating indefinitely many functions, under equally acceptable interpretations.

The initial problem with S-representation should be clear. Since structures *qua* computational objects can be mapped onto indefinitely many abstract functions, how can S-representation *constitute* representation in anything but an utterly trivial sense?

Cummins makes various types of moves to counter this and related familiar objections. First he advances the principle that interpretations themselves and not interpretation users should do the work of interpretation. This rules out various kinds of *indirectness* of interpretation. Thus, although a multiplication function can be mapped onto an addition function we can refuse to see the one computational object as simulating both because only the latter (in the case discussed) is *directly* simulated. To view this structure as simulating the former the interpreter must calculate the transpositions and reinterpret display units. The complexity of interpretation is taken over by the interpreter. It is wrong to tie interpretation to what interpreters must do, Cummins maintains. Rather, interpretation should be tied 'to the complexity (or something) of the proposed interpretation function' (p. 104).

Requiring 'internalized' and completely tracking isomorphisms is one way to rule out degenerate because indirect interpretations. But there are other ways of being degenerate. For example, interpreting a certain structure as 'cat and beloved by the ancient Egyptians' or '5' except after '9' when 'George Bush'. What is required is not just direct but direct and 'proper' — i.e. perspicuous and relevant — interpretations. Against indirectness Cummins stressed the complexity of the interpretation function against the cleverness of its user. But in striving for 'direct and proper' interpretations the triviality of the interpretation function is emphasised: 'the system has to do the work, not the interpretation function' (p. 105).

But what kind of complexity can guarantee the elimination of unwanted interpretations? It is a given that there are indefinitely many ways of consistently describing representational structure whose terms equally succeed in being fully isomorphic with causal structure. This problem may be termed 'the problem of cheap/unwanted contents'.

A further possible option for dealing with the problem of cheap/unwanted contents is raised by Cummins but rejected. Quite rightly so; he is not, consistently, entitled to it. This is the idea that 'proper' interpretations could be explicated in terms of selected functions or selection-by-function. Two ideas seem to be run together in Cummins' treatment of this option. One, that only some interpretations actually matter to systems by virtue of being 'used by' systems. Two, that some interpretations *get* to matter to systems by virtue of being derived through specific environmental interaction. Environmental interactions, the supposition is, determine or select structures with specific functional roles, hence representational content. Cummins is aware that there is no solution here for him. For him selection-by-function comes down to no more than program execution on the one hand and environmental input specifications on the other, (i.e. cognitive function specification). Thus it reduces to more of the same interpretation, hence S-representation. So it introduces no new constraints on interpretation other than those already present (p. 107).

Hence Cummins confesses to no adequate account of direct and proper interpretation. Thus there is no principled solution to the problem of cheap/unwanted contents. Cummins closes however on an article of faith. The Tower Bridge story presupposes some nontrivial notion of interpretation, so there must be one: '*Something* must account for the fact that instantiating *f* isn't enough to instantiate every function isomorphic to *f*' (p. 105).

Against S-Representation

Cummins' aim is to re-educate philosophers such as Fodor about limitations imposed on theorizing by constraints arising from the computer hypothesis itself. The computer hypothesis, Cummins urges, yields S-representation as its central explanatory construct. Cummins concurs with the conclusions of many that S-representation is an unsuitable vehicle for the vindication of intentional causation. It is his view that the entire vindication program, not S-representation, should be called in question.

It may be useful to rehearse the reasons for the unsuitability of S-representation as a vehicle for intentional realism. Standard intentional realism proposes that the contents of mental states are integrally involved in the causal explanation of behaviour. It assumes that there is some determinate, objective matter of fact about the contents of states. Vindication is also assumed to require a naturalistic story about intentional contents and their role in the causal order. S-representation, it seems proper to say, is non-naturalistic, non-causal and non extendable to intentional psychology. Cummins indeed agrees with the third of these claims but not the prior two. The philosopher of science in him wants to defend the 'scientific realist' status of S-representation. To do this he seeks — I think unsuccessfully — to establish that S-representation is both naturalistic and causal. Insofar as Cummins could manage to justify some scientific realist status for S-representation, it would not be of the standard sort since it would be one on which the *causal* and the *explanatory* component of a criterion of scientific realism come unstuck. In other words, although S-representation is deemed 'causal', its explanatory role is not. Its explanatory role is exhausted, remember, in enabling the act of *grasping* a structure as a simulation of something, hence as a computation.

Not Naturalistic

In spite of Cummins' claims to the contrary, S-representation is not a naturalistic construct. Naturalistic accounts of intentional or semantic phenomena are accounts which explain the phenomena in non-intentional, non-semantic terms. S-representation is essentially embedded in modelling pragmatics. Hence it would appear to be essentially dependent on intentional interpreters. How then can S-representation be naturalistic?

S-representation's naturalism Cummins claims derives from its merely being an application of the ontologically benign Tower Bridge story. Thus S-representation specifications *don't invoke interpreters* and actual interpreters *don't 'endow' states with contents.* Specifications are simply hypotheses which quantify over physical structures, abstract structures and interpretation mappings:

> *R* represents *x* in S iff there are functions *g*, *f*, and interpretation I such that S satisfies *g*, *g* simulates *f* under *I*, and *I(R)* = *x*. (p. 129)

However a naturalist is not concerned only with the status of S-representation specifications. Insofar as S-representation is a semantic notion the naturalist wants an *explanation of content determination* which those specifications express. This explanation must itself be naturalistic. ie. in terms of more basic, nonintentional facts. It seems likely that Cummins cannot provide such an explanation. In any case he makes no real attempt to.

Cummins in fact sidesteps the issue with the disclaimer that interpreters don't 'endow' states with content. This claim however draws attention to further problems. Discussion slips between two distinct types of semantic theory. So long as S-representation remains consistently a semantics which does not endow states with content it will go with a non-reductively intentional (hence nonnaturalistic) account of content determination. Where it is at risk of turning into another kind of semantic theory it is possible it could have a naturalistic explanation. But on this other theory states *do* get endowed with content.

To elaborate. We may distinguish, among semantic theories, between those that are reductive and those that are nonreductive and between functional role theories and model-theoretic or indexing theories. S-representation should properly be classed as non-reductive model-theoretic semantics. Yet Cummins is periodically in danger of turning it into reductive functional role semantics. In evidence for this claim I cite the following.

We see in Cummins' discussion a peculiar line up of the notions of Individualism, Functionalism and functional role semantics (FRS). FRS supposes a structure's content to be determined by its functional, causal or computational role. Some theorists extend the determination of functional roles to environmental causes. Computationalism seems to be committed to Individualism. Individualism is sometimes held to require that only those aspects of functional role which occur within the system can be relevant to a psychological notion of content. Now Cummins views FRS as merely an extension of Functionalism; (a common assumption Fodor often complains about). Thus,

> Mental states are individuated functionally . . . . But some states — representations — . . . are individuated by their contents. Hence functionalism implies that mental contents are individuated functionally. (p. 114)

Not only does Cummins appear to run FRS together with Functionalism but he appears also to collapse FRS into Individualism. Individualism supposes that computational equivalents must be cognitive equivalents, hence representational equivalents, regardless of historical-environmental differences. Cummins construes Individualism in such a manner to imply that cognitive equivalents are representational equivalents *because* cognitive-computational role *determines* content. Thus Individualism is expressed as the assumption that:

> current ahistorical state . . . determines current cognitive capacities, and hence must determine current representational content. (p. 81)

The impression that there is a collapse is further strengthened by the observation that Individualism is the 'existence condition for AI':

> if we *give* a system the same data structures that a natural system must acquire by learning, then . . . we have a cognitively equivalent (hence representationally equivalent) . . . system. (p. 84)

The picture here is one where states are endowed with contents in virtue of functional roles. Thus the 'Individualism' of S-representation threatens to turn it into FRS. The picture is one where CTC theorists *do* attempt to *endow* states with contents — at least when they are the states of artificial systems. And they think they are able to make the attempt because nature (via functional roles) has endowed natural minds with contents. In either case, the endowment depends on the idea that content is a function of functional role which is a function of structure which is endowed.

Cummins of course does not want to collapse S-representation into FRS. S-representation is held to be explanatory where FRS is not. FRS gives us 'no hint as to why being a node in a network of computational relations should make something a representation or endow it with a particular content' (p. 122). S-representation, on the other hand, presupposes 'some antecedently specified function the realisation of which will be a representational system . . . '. S-representation as a form of model-theoretic semantics is thus contrasted with and made explanatorily superior to FRS.

In seeking to distinguish S-representation from FRS Cummins passes up for S-representation the kind of naturalistic explanation possibly open for FRS. However S-representation's vaunted explanatoriness surely remains trivial if there is no principled solution to the problem of cheap/unwanted contents, hence no definitive theory of content determination. There is only a pragmatic solution. Interpreters simply ignore nonperspicuous interpretations. That there is only an intentional solution is borne out by Cummins' own remarks. Interpretational semantics is likened to science in general. Science is problem solving. It proceeds by devising good conceptual schemes. Good conceptual schemes 'screen out' (ignore) nonperspicuous concepts (pp. 110-111). Thus insofar as S-representation's explanatoriness depends on interpreters restricting themselves to *hypotheses stating* and not *content endowment*, the hypotheses must be perspicuous. But the perspicuity of interpreters' hypotheses depends essentially on the perspicuity of interpreters themselves (presumably a matter for further hypotheses). I conclude therefore that Cummins has given us insufficient reason to accept S-representation as a naturalistic notion. (Unless he has an independent and noncircular account of perspicuity or rationality to offer.)

Not Causal

For an entity or property to be causally efficacious it must be a causal power or be causally relevant or determine the causal powers of the entity that possesses it. Cummins argues for the causal relevance of S-representation. I find the arguments both puzzling and unconvincing. A 'received view' according to which CTC contents are held to be epiphenomenal is rejected by Cummins on behalf of S-representation. This view allows that states of a system with contents have causal powers and that appeal to causal powers is central to psychological explanation. It asserts however that such states have their causal powers in virtue of their physical, not their content properties (content properties not being reducible to their underlying physical bases). The 'received view' thus concludes that contents are epiphenomenal, hence not scientifically real or explanatory. Now Cummins definitely wants to deny the final conclusion on behalf of S-representation. It appears, but is not certain, that he does this by denying that S-representation is epiphenomenal and resting that denial on a rejection of the claim that states have their causal relevance or causal powers in virtue of physical and not content properties. That this latter proposition is rejected seems to follow from assertions such as that CTC content can be accorded 'whatever 'reality' or membership in the 'natural order' goes with having a causal role' (pp. 134-135), and that CTC is 'in a position to avail itself of the usual reasons for supposing that states of cognitive systems have effects in virtue of their contents' (p. 136).

Unfortunately, as the intentional realist knows, the 'usual reasons' for justifying the causal status of contents in the face of epiphenomenalist challenges arising from the physicalist assumptions of CTC are not all that easy to establish. Any argument for the causal efficacy of content must come to terms with the locally based physical basis of causal powers in a way that does not neutralize the efficacy of content properties in their own right. Moreover one expects naturalism to come into play again in that it is usual to look for a connection between the properties in virtue of which content is naturalistic and the properties in virtue of which it is causal-explanatory.

Cummins approaches the causal status of S-representation via the question of generalizations. Settling the question of (G) is advanced as the criterion for causal efficacy. To settle the question of (G) is to determine whether (G) is confirmed in a specific context.

(G) In a true singular content-ascribing causal statement having the content ascribed is (at least sometimes) the 'causally relevant factor'.

(G) can be settled either by the 'Way of Generalization' or the 'Way of Counterfactuals'. Since the latter is supposed to collapse into the former, the Way of Generalization it is. That is, (G) is true 'if singular content-ascribing causal statements generalize along lines traced by content' (p. 131).

Following good philosophy of science, Cummins addresses this question initially in the context of calculators. For example, a causal statement describing the incrementation of an accumulator will be generalized along lines of content, for the reason that computational functions are involved that need to be 'hardware independent' hence need to be described in function-representational terms. Thus the Way of Generalization confirms (G) in this instance via multiple-realizability. Basically the same argument is run for the context of CTC but with a peculiarly *a priori* emphasis. It is termed 'the argument from the existence of the CTC'. It holds that wherever CTC exists, autonomous content generalizations *must* exist. This is because epistemological, hence semantic, constraints mark the generalizations in virtue of which functions *count* as cognitive and are individuated. Feed in multiple-realizability of computational instantiations. Derive the irreplaceability of content generalizations. Hence (G) is confirmed for CTC. Hence content is causal in CTC.

With respect, all Cummins has established, as others before him, is that one loses interesting generality if one eliminates content generalizations over computational systems. Mere physical specifications won't indicate what a number of physical systems have in common. Namely, that they are all instantiations of the same cognitive function. But in what sense does establishing this fact establish the causal efficacy of contents in computational or psychological *processes*?

Intentional realism depends on more than the autonomy of content generalizations. Although there are many ways of running an argument needed to establish causal status for representational content, most exploit some notion of supervenience of contents on naturalistic properties. The supervenience relation must be strong enough to guarantee at least equal causal efficacy for content properties as the physical basis they supervene on. For this it would seem to be required that contents are at least locally 'realised' in physical systems. Fodor, for example, advances a naturalistic nomological (causal covariation) theory of ('narrow') content, along with a nomic subsumptive view of intentional causation resting on the postulation of intentional laws. There are problems with

Fodor's and other intentional realists' theories. But the point is, if one wants to go from a 'way of generalization' to content efficacy it must be via *some* such substantive argument. Cummins makes no attempt to supply such an argument. Furthermore, for S-representation, no such argument could be provided.

Of course, Cummins is not interested in establishing intentional realism where this entails intentional explanation as causal subsumptive explanation. The causal status of S-representation is required only to guarantee its scientific realist status. But the notion of 'causal status' that emerges is most unclear. S-representation seems best construed as epiphenomenal. Given Cummins' interpretationist view of science its not clear why he is not satisfied with epiphenomenalism. S-representation supervenes too weakly and on the wrong sorts of (intentional, non-local) facts for causal efficacy. It's doubtful, given the normative, pragmatic dimension of S-representation, that it even meets the criteria of contents being *realised* in systems. Causal status seems to require at least that an entity have the causal role it does because it has the property in question — in this case a specific content. If it would have had the causal role or causal powers regardless of having the property then the property did not determine causal powers or was not causally efficacious. Yet this is exactly the situation with S-representation. Structures' causal powers can remain invariant under various interpretations, i.e. when theorists change their minds about what target functions (amongst infinitely many) are being simulated. (Cummins may try to retreat into global physicalism but I don't think he will find a solution.)

Cummins will object to the foregoing argument. He maintains that it is like thinking you can knock off the top span of the Tower Bridge and still have the bottom span. This is illegitimate: the top span's (ie. content's) 'presence is *entailed* by the structure in the lower span . . . Once you have built the lower span you have built the contents; there is nothing more to S-representation than instantiation' (p. 135). Without any proper solution to the problem of cheap/ unwanted contents the appropriate reply should be we have 'built' too potentially many contents to have actually built any at all. Except, that is, if we defer to the structure the theorist *intends* him/herself to have built. But such deferment is not grounds for establishing the causal efficacy of content.

I have argued from the essentially pragmatic, non-unique, model-relativity of S-representation to the conclusion that it is non-naturalistic and non-causal. Cummins concludes only that it cannot provide a vehicle for belief-desire psychology. The intentional contents of beliefs and desires, Cummins recognizes, must be objective, determinate matters of fact and not artefacts essentially relative to interpreters. Nevertheless, he contends, S-representation does provide the *scientific* construct for the explanation of cognitive functions. Cognitive functions need have nothing to do with intentional psychology. Therefore CTC should stick to S-representation and cognitive functions and drop intentional psychology, philosophers of mind lowering their sights accordingly. I suggest that if Cummins is right that CTC or the Computational Theory of Cognition is committed to S-representation, we might as well give up CTC. But *is* Cummins right about this?

Philosophy, Science and Methodology

Cummins holds that S-representation is the construct philosophy of science entitles us to. We are as entitled to it as we are to the theories of space delivered by General Relativity. What, we need ask, do all these entitlements come down to on the present view?

Cummins advances a three-termed analogy holding between (1) doing science, (2) cognizing, and (3) doing S-representation or interpretational semantics. To quote at some length:

> The goal of science is . . . to develop concepts . . . What a good conceptual scheme does is force us to characterize the system under study in a way that screens out all information that is irrelevant to the problem at hand . . . a kind of filter, like a pair of conceptual glasses . . . Analogously, the proper semantics — i.e. the right interpretation will allow you to see the right computations . . . as cognizing. (pp. 110-111) The CTC . . . is just the idea that cognizing is generally like doing science. (p. 112)

A key feature in the above is the notion of problem-solving. A conception of explanation as understanding *qua* interpretation is correlated with it. In effect, the three termed analogy reduces to a trilogy: science is problem-solving, interpretation is problem-solving and cognizing is problem-solving. Crucial presuppositions underlie the emphasis on problem-solving, viz., that (i) it requires domains which can be 'prepackaged', and (ii) explanation afforded is not primarily causal and is governed more by coherence than correspondence principles. Let's begin with science. Its alleged purpose is to enable us to 'grasp' aspects of nature in ways which seem perspicuous, coherent and autonomous. Structure in nature must be significantly compartmentalizable. Science maps nature's structure onto symbolic conceptual structures. The latter have their intrinsic rules of consistency. Mappings are effected so respective structures 'track' each other. Symbolic structures are essentially instruments. We measure or index nature's structure in terms of them. Thus we make nature intelligible to us. There will always be many legitimate ways of rendering nature intelligible, i.e. conceptual glasses we can wear.

The situation is more complicated in contemplating cognition as problem-solving. Clearly for Cummins CTC , though supposedly a broad umbrella term covering empirical theories in psychology as well as AI, in fact takes on the distinctive complexion of standard AI. We have this situation: (a) cognizing is problem-solving, (b) explaining cognizing is problem-solving, and (c) CTC as empirical theory *embodies* assumptions (a) and (b). However what actually embodies these assumptions is primarily a well-established form of AI. If Cummins insists that CTC in general does this amounts to reducing CTC to that form of AI.

The predominant AI methodology is to simulate semi-autonomous capacities on general-purpose computers. AI's top-down functional analyses of capacities into functions become explanations when reconstructed in bottom-up computational instantiations of the components of analyses. Components of analysis, now functional components of simulations, are constructed from repertoires of known instantiations already made available and embedded in developed programming languages. This is what explanation *qua* function simulation in this form of AI is all about. Problem-solving occurs and is instantiated along a number of dimensions. On one, S-representation enables us to grasp computational instantiations as function simulations which are arrived at by breaking a target capacity or problem down into its component capacities or sub-problems. Meanwhile the functions simulated may themselves, in virtue of their inferential input-output specifications, be viewed as instantiated 'homunculi' or mini problem-solvers (cf. John Haugeland's 'black boxes'). The comprehensive picture then is one where both science in general, and CTC as AI in particular, integrally involve S-representation, ie. rendering structure intelligible via some preconceived system of intelligibility.

Science then, legitimizes S-representation for CTC. But CTC as a theory of cognizing is a theory of doing S-representation. CTC theories emerge as theories which both reconstruct and depend on cognizing as instantiation of S-representation. In understanding how it is possible to understand a structure *qua* mind, we explain the mind. So interpretation explains interpretation by means of interpretation. S-representation is explained by S-representation.

Cummins' view of science should be contrasted with a genuinely scientific realist one. On this view the universe is a largely causal structure populated with objectively existing entities which participate in objective causal structure by virtue of the distinctive properties and relations by which they are classified into kinds and particulars. Science's purpose is to discover and describe. It cannot be known prior to investigation what 'conceptual scheme' or interpretation best fits the facts. On interpretationism, explanation derives from the rationality of interpreters as creators of interpretation schemas. Such schemas are in danger of intruding between theorizers and reality. For realists rationality is *itself* a fact to be discovered and described. Evidence suggests minds participate in the causal order of things. Realists want an explanation of minds and agency as part of the causal order of things. Interpretationists provide an *application* and not an explanation of rationality.

If we accept Cummins' convergence of science, cognition and semantics the results are circular and hence not very satisfying. But why should we accept the convergence? Scientific realism of the standard sort grows out of commonsense realism and unless given powerful reasons to think otherwise it is the rational approach to science. Cummins' general arguments for the sole legitimacy or usefulness of S-representation then are tied up with a view of science many will see little reason to accept.

Physics and CTC as introduced don't seem to be on a par. Physicists aim at a comprehensive theory. If current theories of physics are true and a certain notion of space is integral to them, that notion is justified not because of its coherence or elegance but because it's part and parcel of an accurate description of reality. Standard AI by contrast is committed to no true-by-virtue-of-correspondence theory of cognition of which a notion of representation is an integral part. AI is simply committed to results any way it can get them.

Cummins then has offered philosophy of science as the cure for philosophers led astray in their theories of mental representation. I have questioned this philosophy of science. For Cummins, S-representation and CTC-cum-AI are conceived as microcosms of science. If this were so, science would be as uninteresting as, ultimately, S-representation and CTC-AI are.

Philosophers are not to be yoked to some 'canonical' notion of representation allegedly grounding the science of cognition in the science of computation. Where do we find a science of computation? All we have are the largely *a priori* Turing machine principles and their empirical outworkings in the developing of programming languages that has shaped AI. You can see these developments as holding the key to the inner nature of intelligence if you like, but it's probably not very wise to do so. Even if you do you need not find that the key is S-representation. Allen Newell, for example, is one who viewed developments in computer science as revealing the secrets of intelligence. He too has advanced a canonical notion of representation and it is not S-representation ('Physical Symbol Systems', *Cognitive Science* 4 (1980) pp. 135-183). Newell argued that universal machines both provide the necessary and sufficient conditions for intelligence and contain a definitive notion of symbol. He called this notion *designation*. It is essentially the notion of action at a distance. That is, one computer object designates, hence represents, another because the system behaves when it operates on the former the way it would have behaved if it had operated on the latter. (The first computer object, in the most fundamental kind of designation, is 'assigned' to the second.) Newell held designation (one computer object designating another computer object) to be *generalizable* to real world designation (a symbol designating an external object). Although I cannot argue it here, the most plausible generalization would probably be in terms of covariation or functional role semantics, both semantics Cummins blackballs as being incompatible with computationalism.

In AI practice designation, though operationally distinct, is perhaps not a genuine alternative to S-representation. What symbols end up 'designating' would depend on what programmers aim at simulating, hence S-representing. However Newell in effect built real-world designation into his empirical symbol systems theory from the start as a potential theory of real-world intelligence. Newell's option should show us that 'grounding' cognitive theoretical constructs in a 'science of computation' need not mandate S-representation as central construct.

Cummins has supposed that science-respecting philosophers would drop the intentional realism program in the context of computationalism if they approached computationalism as philosophy of science. Fodor is a primary target. Yet it is by no means clear that Fodor's philosophy needs correction by 'philosophy of science'. Fodor is a scientific realist of the standard sort. He believes in no 'science' of computation. Empirical computational theories he holds to be neutral on the question of representation. The fact that typical AI can only proceed at all on the basis of S-representation is a fact of minor interest to philosophers. Fodor locates the scientific credentials of his theory in the status of *psychology* as a (special) science. Computation supplies merely the mechanism for the implementation of intentional laws. Psychological data drives theorizing and construct articulation. Psychological constructs must be compatible with but not driven by computational theory. AI (though by no means all cognitive scientists) typically aims at models which are computationally complete; hence the emphasis on universal machines and general problem-solvers. But results in this branch of AI suggest the assumptions are psychologically implausible. Fodor's assumptions and approach are more psychologically plausible — though some would disagree — while radically computationally incomplete. Whether Cummins or Fodor has chosen the right constraints on theorizing depends on whether there is more likely to be a fruitful science of computation or science of psychology.

To view computation as supplying mechanisms for psychology while not driving psychological theorizing, in particular remaining neutral on the question of representation, is to construe computation primarily as a *kind of process*, and not essentially as function simulation. What kind of process? Well, the sort where if some computer objects are semantically interpreted then causal processes of the system guarantee systematic relations between semantically interpreted symbols. Science (psychology) reveals patterns of connections (which may or may not exhibit rationality) between contentful states. Computation is a *process whereby physical systems can bring about connections of contentful states.* That's *all* the 'computer hypothesis' for philosophers of mind need consist in. The question of what it is for computer objects to be semantically interpreted is left an independent issue. Thus representation is left an open question. Cummins pre-empts the question by making grasp of structure by intelligent designers in terms of preconceived, pre-articulated functions *constitutive* of computation. Thus his construal of computation virtually builds S-representation in by definition. But I maintain that this is too closely bound by contingencies of most AI practice up to now. (If AI were in the business of building real-life, environment-interacting, replicas of us, S-representation need *not* be the central explanatory construct of representation. But this has all been said before.) Only insofar as CTC is equated with an AI committed to constructing simulations to instantiate pre-articulated abstract functions is CTC committed to S-representation. But CTC need not be handcuffed in this way. We may understand CTC broadly as the 'theory' (empirical assumption) that cognition is implemented computationally. We need only suppose that it is conceivable that God or nature, in designing structures whose causal relations preserve semantic relations, also determined that the semantics of symbols are themselves objectively given by nature. We can speak of cognition satisfying functions in some sense but not the sense Cummins tries to impose on us.

Cummins' arguments against alternative semantics do focus on severe difficulties facing those programs. I have been concerned only with the 'ideological' aspects of his discussion. That is, the idea that the foundations of the computer hypothesis dictate S-representation as the basis of any scientifically respectable theory of mental representation, I have taken issue in effect with Cummins' thesis that calculators and cognizers must be in the same boat.

Philosophers of mind are entitled to think otherwise.

D.D. Gamble University of Adelaide