COGNITION AND BRAIN THEORY, 1983, 6(2), 237-249

CRITICAL DISCUSSIONS

Fodor's Representations

RADU J. BOGDAN
Department of Philosophy
Tulane University

I. OVERVIEW

Mental processes can be best understood as being symbolic and formal, that is, computational, in the sense that they engage and operate on representations in virtue of their shape or form. This presupposes a code or language in which mental representations get expressed and a set of formation and transformation rules for the formulae of the language. This language of representation (LR) extends from peripheral processes such as perception to central ones such as thought. The pervasiveness of LR explains not only how organisms handle information at various levels (by computing it) but also how these levels interact (by exchanging results of computation in a common language). This is the *computational* picture of cognition. There is also a parallel *physical* picture which assigns to any formula in LR a corresponding physical state of the hardware, such that interformulaic derivations are matched by causal relations among the relevant physical states. A particular cognitive or behavioral feat can thus be seen either as the output of a causal sequence of neural states under causal laws or as a last line in a derivation under rules of computation. The computational picture belongs to psychology, the physical to the neurosciences.

The computational theory of mind (CTM) that emerges from the psychological picture comes in a philosophical package. If you buy CTM, you may also have to buy: (1) materialism in its weaker, token-physicalist form, because only mental particulars or tokens, but not mental types or kinds, are physical and obey causal laws; hence (2) the autonomy of psychological explanation, because an explanation essentially involves statements about types and the types you get are irreducibly psychological; (3) anti-behaviorism, because there is a mental inside, the computational mechanics, which does a lot of work independently of, and prior to, any behavior;

Requests for reprints should be sent to Radu J. Bogdan, Department of Philosophy, Tulane University, New Orleans; La., 70118.

(4) a restricted (qualia excluding) functionalism, because computations are abstracted from hardware (neural) realization and are also treated relationally, in terms of input (initial formula/cause)—output (derived formula/effect) and (5) innatism, since any and every cognitive process, being computational, requires a prior LR and the learning of any such language, itself a cognitive hence computational process, requires a still prior LR, and so forth, in a regress that only innatism can stop.

Historically, the two pillars of CTM may well be Chomsky and the computer. Yet CTM's intellectual debts reach further in the philosophical space and time: to rationalists, for innateness and a computational view of logic and language; to empiricists, for the combinatorial mechanics of concept formation and the inductive logic of belief fixation; to skeptics for indicating the futility of semantic questions about mental representations, and to solipsists, for showing that such questions cannot be answered; to materialists, for making mind matter (metaphysically); and to nonreductionists, for minding matter (explanatorily).

This, more than roughly, is a CTM with its philosophical pedigree and critical horizon in what I take to be its best articulated and most imaginatively defended version, Jerry Fodor's. A detailed exposition and motivation of CTM can be found in Fodor's 1975 book, *The Language of Thought*. The volume under discussion, *Representations*[1], is not a new book but rather a collection of eleven essays, two new (an introductory survey and a 60 page discussion of innate concepts), the rest written and published in the last 15 years or so. Most essays revolve around some of the main topics of the earlier book: the critique of Wittgenstein, behaviorism, reductionism, and procedural semantics; the explanatory irreducibility of psychology; and the problem of functionalism. I will have very little if anything to say about these issues. Instead I will focus on two major themes of *Representations:* the narrow, nonsemantic, solipsistic scope of cognitive psychology delineated by the formality condition, and the determination of mental content and belief. To my ear these are dissonant themes: A narrow, solipsistic psychology cannot by itself determine and explain content and belief. The rest of this paper attempts to explain why.

II. THE PROBLEM

Let us have first and briefly Fodor's treatment of the two themes. A belief is a relation to a formula in LR. That formula expresses the content of the belief. Content determination, then, must be computational. Since contents do produce other contents, in thought and inference, as well as action, an

---

[1]The MIT Press, Cambridge, Mass. (A Bradford Book), 1981.

explanation of cognition and behavior must invoke contents and be computational. This is the only form of explanation that the cognitive psychologist can and must engage in; it is also the only one which ensures the autonomy and legitimacy of his discipline (Chapters 5 and 6). But there is a price to be paid: Computational explanation has a narrow scope. This is where the second theme comes in (Chapter 9). It runs like this. Mental states and processes are computational in the sense that they encode and apply to contents only as forms. The *formality condition* (FC) protects this characterization of, and approach to, cognition by stipulating that content be exclusively treated as form. Consequently, FC entails that mental states and processes should be construed as not having access to the *semantic* properties of their contents. Since the cognitive psychologist has to go computational if he is to do his explanatory job, it follows that, in complying with FC, he has to subscribe to a form of *methodological solipsism* (MS) which says that to explain cognition one must blackbox its semantic contact with the world. Naturalist philosophy and psychology take the opposite line and thus become implausible if not impossible projects.

We need to pause here for clarification. The anti-semantic force of FC, I suggest, can be read in a strong or weak sense. The *strong* reading completely rules out any concern whatsoever with the semantics of mental representation. The *weak* one allows only those semantic aspects which can be structurally captured in the form of a mental content. What complicates the picture, though, is the prevalent and mistaken assumption, shared by Fodor, his naturalist opponents, and countless other theorists, that the semantics of mental representation amount essentially to questions of truth, reference, indexicals, existence, and the like. If this were so, then the strong reading of FC would be legitimate. Indeed, as Fodor correctly argues, thus joining a distinguished tradition going back to Hume and Descartes, these semantic parameters are excluded by FC precisely because they cannot be represented in the form of a mental content. The reason is simple. I will put it in my terms. Truth, reference and the rest are *extrinsic* (and second order), not intrinsic, properties of mental representations. Like seeing and knowledge, they cannot, and are not meant to, characterize something that a representation intrinsically possesses or is capable of doing. Instead the job of the semantically extrinsic parameters is to ascertain whether (and to what extent, if probabilities are included) mental representations discharge their intrinsic semantic task. In this sense, the job of the extrinsic parameters is *evaluative* rather than descriptive of cognition. From an external vantage point, as it were, they measure an organism's informational dealings to see whether they are successful or not. This is no different from measuring the goodness of an action. No intrinsic aspect of the action itself can possibly embody goodness just as no intrinsic aspect of a representation can embody truth or reference. No wonder, then, that the latter are necessarily excluded by FC.

The above analysis presupposes a distinction I make between semantic *interpretation,* which is a theory of the intrinsic semantics of mental representation, and semantic *evaluation,* which is a theory of the extrinsic semantic success of mental representation.[2] We can now restate Fodor's programmatic position. If he takes FC to exclude semantic evaluation, then he is right. If this is what MS amounts to, then it is a most sensible research strategy. Since, moreover, psychological and philosophical naturalism are generally concerned with semantic evaluation, Fodor is also right in holding that naturalism is none of the cognitive psychology's business. If, however, FC is read strongly, and MS is therefore construed, as excluding semantic interpretation too, then Fodor is wrong.

Now Fodor's position, as I understand it, appears to fall in between. This creates a problem at the very heart of his enterprise. Yet, because of its centrality, the problem also creates an opportunity for its dialectical solution. I will first identify the problem and its centrality and then anticipate the dialectical solution.

The *problem* is this. Fodor wants a weak FC in and semantic evaluation out. So far so good. But, as earlier indicated, he also appears to assume that extrinsic semantics is the only semantics there is for mental representation. This is not so good. The two views cannot coexist. One must go. The reason is this. If there is only extrinsic semantics, then the semantic aspects which, according to a weak FC, can and should be treated as form are extrinsic. But this, as we have seen, does not make any sense. If, then, the extrinsic aspects are the only semantic aspects there are, a weak FC is contentless. Fodor appears to assume the antecedent while denying the consequent. This is the problem.

If a weak FC requires semantic aspects, as it must, and these aspects cannot possibly be extrinsic, then there must be *other* semantic aspects for CTM to account for. Therein lies the solution: A weak FC must presuppose

---

[2]The distinction between the intrinsic and extrinsic semantics does not match, although it may overlap with, the standard distinction between intensional and extensional semantics. Georges Rey, for example, criticizes Fodor's reading of FC as exclusively extensional. See *The Brain and Behavioral Sciences* (1980), *3*, p. 91. I find the standard distinction unsatisfactory for understanding mental representation, although useful in other contexts. Firstly, it labels rather than explains the distinction intended. Secondly, it is not specifically geared to mental representation but to languages in general. This fosters the impression, thirdly, that some intensional aspects (the ones I take to be intrinsic such as information or meaning) are properties of a language as such. This is a mistake. They are rather properties of a language instantiating a cognitive performance. Fourthly, the standard distinction *either* assumes that it differentiates between two perspectives on the *same* phenomena, whence the attempts either to reduce one perspective to the other, or simply to explain one away, thus paralleling the materialism-dualism debate; *or* it restates in semantical terms the difference between concrete particulars and abstract entities, thus paralleling the nominalism-platonism debate. Needless to say, either alternative has nothing to do with what I take to be the contrast between the intrinsic and extrinsic semantics of mental representation.

only intrinsic semantic aspects, and semantic interpretation is the theory of these aspects. Or so I will argue.

What is *central* about the problem I think Fodor faces is this. The main overall objectives Fodor assigns cognitive psychology are identifying and explaining the mental contents and propositional attitudes of an organism. He thinks that only *as CTM* can psychology pursue these objectives. But the problem I have identified for Fodor can be shown to prevent psychology-as-CTM from doing its assigned job. The reason is that CTM cannot possibly identify contents and associated attitudes, let alone explain them, if it does not consider the intrinsic semantics of mental representation. This means reliance on semantic interpretation. What can be *dialectical* about this solution is that it complies with Fodor's requirements for psychological theorizing: semantic interpretation satisfies, and in fact legitimizes, a weak FC; it also satisfies the opacity condition; and it does not presuppose semantic evaluation. But there is another dialectical aspect involved. The case for semantic interpretation, and a parallel approach to belief, can and will be made from within a framework that Fodor himself has used to make his case for CTM. This dialectical concern should indicate my basic sympathy for the main thrust of Fodor's CTM. It is just that I do not see it as self-contained as Fodor appears to see it, for I do not see how it can take off without the map provided by semantic interpretation for mental contents and by performance interpretation for propositional attitudes.

My next task, in the section to follow, is to document the diagnosis of the problem of content determination and then show why semantic interpretation is needed. A parallel tactic will be used, in the section after next, to outline the problem of belief facing CTM and show why performance interpretation is needed.

III. CONTENT

Documenting the diagnosis: We have, on the one hand, the weak reading of FC and the semanticity of mental content. A mental representation, Fodor writes, is viewed as a "semantically interpreted object (p. 30)." Propositional attitudes inherit their semantic properties from those of mental representations (p. 31). Indeed, content itself is a "semantic notion par excellence (p. 277)." We are also told that CTM "cannot get off the ground without having a look at the organism-environment relations that (presumably) determine the semantic interpretation of thoughts." Fodor indicates that such a look is compatible with FC.[3] So content is semantic and FC must be weak.

---

[3]In his answers to comments on the original version of Chapter 9 in *The Brain and Behavioral Sciences* (1980), *3*, pp. 102, 103, and 106.

Yet Fodor's view of semantics, on the other hand, appears to be consistently extrinsic. The evidence is both direct and indirect. Directly, as mentioned earlier, the key semantic relations Fodor almost always discusses are extrinsic. Also, he appears to equate semanticity with transparency (p. 239), and the latter is indeed an extrinsic notion. Finally, there is Putnam. A good part of Chapter 9 is a reaction to Putnam's semantic naturalism, a doctrine which programmatically construes semantic relations as extrinsic. This is called the *wide* construal of mental states: it combines what is in the head with what, extrinsically, is out there to determine the content of those states. Since the wide or extrinsic construal now defines semanticity, it is no surprise that a narrow reaction to it, such as Fodor's, will exclude semanticity altogether and declare that "semantic notions aren't psychological categories (p. 253)." But the wide-narrow contrast need not be exhaustive. There is, I will suggest, a third alternative, semantic interpretation, which posits a narrow psychological basis but a wide semantic horizon.

The indirect evidence is of the dog-which-did-not-bark sort. There are some notions which I take to characterize the intrinsic semantics of mental representation. Chief among them are information and meaning—the latter understood not as a concept, that is, a computational device (p. 206), but rather as the mental act of a speaker. Fodor studiously avoids talking about them. I think this is indirectly significant. It is also significant in another sense, to be amplified later. Information, meaning as well as belief have intrinsic semantic aspects whose identification requires consideration of the overall cognitive and behavioral performance of an organism. This global view of performance is not characteristic of CTM. But so much for the diagnosis.

My other task is to argue for semantic interpretation. I take the intrinsic semantic function of a mental representation to be that of coding information in some form. Information is an intrinsic notion because its determination relies essentially on data about the organism's cognitive competence, its various mental functions such as beliefs, goals, and prior knowledge, and the context of its current performance. Since, in an organism, information and its flow materialize as, or literally take the form of, computational structures and transformations, the determination of information, while semantic, can be accommodated by a computational account of cognition and thus comply with FC.

This is how we should look at semantic interpretation. It consists, roughly, in figuring out an organism's informational transactions, both external (with the world) and internal (inferential, generative of behavior). This is, in other words, trying to figure out what an organism perceives, remembers, thinks. Semantic interpretation is *opaque* in that it takes the organism's perspective on, and has no independent or transparent access to, what is represented. It is only opaquely that the interpreter can determine

what counts as information-for-an-organism in a certain context and how that information is represented, acted upon, and so on. It is only semantic evaluation that is transparent.

Semantic interpretation is a *hypothetical* enterprise. On the basis of data and various constraints such as intelligibility, coherence, and so on, the interpreter speculates about what an organism extracts or uses as relevant information in a given situation. In so doing he need not assume specific semantic relations of the extrinsic sort, nor need he evaluate the semantic worth of the contents he is trying to identify. This does not mean that the interpreter may not and does not, in general, rely on data and hypotheses about the organism-environment set up. He had better do so if the interpretation is to make any sense. The point rather is that a particular interpretation need not rely on a semantic evaluation of the very contents being figured out. This is my reading of FC.

Also, the data and hypotheses generally brought to bear on various interpretations do not have to depend on the ultimate truth of science. Fodor believes that Putnam's semantic naturalism requires this dependence (pp. 247–250). I, on the other hand, take semantic evaluation and the wide construal of content to assume an extrinsically natural (causal, lawlike) connection to the facts and things out there *independently* of how we determine and explain them. In other words, the assumption is ontological, not epistemological. But this is something semantic interpretation need not worry about anyway, which is a virtue.

A principled argument for semantic interpretation, as being indispensable for content determination, can be mounted from a perspective which Fodor has been for years so instrumental in promoting. This is the other dialectical aspect mentioned earlier. In a number of works, some reprinted in this book (Chapters 5 and 6; also pp. 6–11), Fodor has argued for the autonomy of psychological explanation and its irreducibility to neurological explanation. He sees the computational turn in our understanding of mind as essential in securing this autonomy. His argument is that such explanatory constructs of psychology as content and computations (pp. 114, 164) delineate kinds that cannot be directly captured by, because they have no counterparts in, the reductive explanations of neurology. Unless content is posited in a vocabulary and network of facts and regularities specific to psychology, the neurologist may have all possible data about an organism and still be unable to figure out contents. He would not know what to look for. The neural patterns he has theoretical access to may vary a good deal and yet, from a psychological angle, realize the same content relations. If, however, the neurologist is told in computational terms how to identify content, he may then *redescribe* it in neurological terms. But token redescription is not type explanation. Content, being an offspring of explanation, is lost in the redescription. This is the moral of Chapter 6.

A *parallel* argument can now be mounted against CTM if, as construed earlier, it excludes semantic interpretation. In a rough outline, the argument goes like this. Just as neurological theorizing needs higher level, irreducible constructs and explanations of the computational sort in order to map and redescribe the mental syntax in neurological terms, so does the computational theorizing itself need the higher level, irreducible constructs and explanations of semantic interpretation in order to be able to treat computational states and processes as semantically significant or contentful. When the methodologically minded neurologist urges that computational form be reconstructed as neural pattern, he cannot mean that, by itself, the latter would tell us what the former is. Likewise, when Fodor urges that content be reconstructed as form, he cannot mean that, by itself, the latter can identify and explain the former. In both cases one needs to know antecedently *what* to reconstruct at the lower level.

What this parallel implies, then, is that just as content-as-form "evaporates" when neurologically redescribed, so semantic content must "evaporate" when computationally redescribed. What really "evaporates" is the explanatory, type identity of the constellation of facts subsumed under the concept being redescribed. What often obscures the latter "evaporation," of semantic content when computationally reconstructed, is the tendency that many philosophers and psychologists alike have to take the simplest units of content-as-form, say, a formula in LR made of object and property, as being the real format of the content of one's cognition. As we will see in the very next paragraph, even such a simple formula (like 'a is F') needs a prior story to deliver its content. But the point I want to make now is that cognitive content is typically organized in macrounits and increments in information which an atomistic, unit-by-unit, bottom up, compositional analysis that CTM is capable of producing cannot capture by itself. Think, in the visual domain, of entire scenes and of informational increments geared to the organism's areas of uncertainty and interest. Think also, in the linguistic domain, of such macrounits of speech as topic and presupposition, various increments within a discourse or text. All CTM can do, nere as in vision, is to give the story of their syntactic microrealization. Semantic interpretation, on the other hand, chases content from the simplest units to the more complex ones while always aware that the "phenomenological reality" of content, that is, the story of content commensurate with how the organism represents what it does, comes typically in macro, not micro, units.[4] But, as anticipated, CTM is in trouble even with the microunits. To see this we have to go back to FC.

---

[4]If, perhaps understandably, people think that the informational units and increments I take semantic interpretation to identify are not genuinely semantic but rather pragmatic or something of that sort, then I will restate my discussion accordingly but insist that my critical position remains intact. I stick with the adjective 'semantic,' meant intrinsically, because many theorists do think that such notions as information, content and meaning are semantic.

FC is not only a methodological injunction about how to treat content. It is also, Fodor tells us, a method of telling contents apart. It is not, however, a very successful one. As Fodor notes, two propositional attitudes differ in content only if they differ in form (p. 227, etc.). This is much too weak. We can have same form and different contents if the form is semantically interpreted in different ways, just as we can have distinct forms with the same content. Examples in computer or natural languages illustrate both cases, as the reader can verify. So, unless an interpretation is provided, FC cannot help determine content. The latter, Fodor also argues, can be identified by functional role too. This is a consequence of FC. The argument is this. Mental states cause other mental or behavioral states in virtue of their content; but content is form, so formal identity/difference will amount to functional identity/difference; content, therefore, can be determined by its functional role. But the earlier objection still applies. One can imagine and describe a situation, at a certain level of generality consistent with how several agents represent it, in which they have, content-wise, similar beliefs and goals whose causal, formal and hence functional redescriptions are quite different. Again, one needs *prior* knowledge of what makes various behaviors and experiences functionally equivalent. It takes a pre-formal characterization of content to specify functional equivalence.

In short: Semantic interpretation without CTM may well be empty, but CTM without semantic interpretation is certainly blind. A similar story can be told about belief. To it we now turn.

IV. BELIEF

Mental contents cannot be determined abstractly, in isolation. As Fodor's FC implies, contents are construed as such precisely to the extent to which they play some *role* in the cognitive and behavioral economy of an organism. If a content is the information on which an organism acts or which it relies on to seek or infer further information, then we say that the content functions as a belief. If a content is information used to anticipate a contingency, we say that the content operates as a plan. And so forth. Thus, to determine a content is to determine an attitude an organism has toward it, whence the generic concept of propositional attitude. It is then natural, as Fodor insists, for a psychology of computational cognition to be in fact a psychology of propositional attitudes. Its goal is to systematize, explain and state generalizations about these attitudes (pp. 18–19, 25, 102–103, 158, etc.). Belief is the central attitude and Fodor devotes most attention and space to it: three full Chapters (4, 6, and 7) and many other fragments (e.g., pp. 18–20, 24–30, 236–240, 242–244). But there is an intriguing fact here. Fodor tells us that cognitive psychology is essentially a study of propositional attitudes, of belief in particular. Yet an inspection of the literature,

from textbooks to articles in the specialized journals, yields nothing of the sort. There are no chapters on belief in the textbooks or major monographs, no entries for it in the indexes either. The inspection also fails to locate articles on belief in the mainstream work on cognition. I do not discount exceptions yet it seems obvious that current cognitive psychology does not think highly of belief. This cannot be accidental. An explanation is needed.

To understand mine we should contrast the position of cognitive psychology with that of folk psychology, and then both with Fodor's. Unlike cognitive psychology, folk or common sense psychology is up to its neck in speculations and generalizations about belief. Why this difference? A very crude answer goes like this: Cognitive psychology is a study of competence, whereas folk psychology is concerned with performance; belief is a performance concept, so we have our explanation. But this is not literally true. Cognitive psychology is necessarily interested in cognitive performance but in a very indirect and reductive sort of way. The typical experiments show this too well. They often strike an outsider as utterly artificial, that is, far removed from actual performance. There is an explanation for this. Performance matters in these experiments only to the extent to which it exhibits and illuminates a specific capability, that is, only to the extent to which it is competence realized, tokened. Without *some* performance one cannot get at competence. The eye doctor needs some minimal performance to check visual competence but he is seldom interested in what the patient really sees. More theoretically, the same is true, at different levels, of the neuroscientist and cognitive psychologist.

Perhaps a more general metaphor may sharpen the distinction I am after. We may think of cognitive performance in a *local* and *vertical* way if we take, rather artificially, a limited and isolated unit of performance (say, seeing colors or rotating a mental image) and treat it as a sort of probe taking us vertically to the underground level of competence to illuminate a specific disposition or mechanism. This is the position I am attributing to various scientists of cognition and behavior. Or, with the folk psychologist and the rest of us, we can think of cognitive performance in a *global* (or integrated) and *horizontal* way if we relate and align a macrounit of performance (say, an action) to several others which are deemed responsible for it (say, beliefs and goals) in an attempt to figure out the organism's overall performance, irrespective of the neural structures or the computations which underly the performance in question. The latter, folk strategy casts a net over the cognitive "surface" with the aim of delineating a centrally initiated, overall performance; the former, scientific strategy probes locally into the cognitive "depth" where competence is at work.

How does Fodor's position fit into all this? Not too neatly. He occasionally waves in the direction of folk psychology (pp. 19, 24, 184) but his

methodological commitments pull him in the other, scientific direction. He borrows the performance concept of belief from folk psychology, which understands it in a central, global, and horizontal manner, but he gives it a local and vertical twist in the direction of competence. It is this twist which ensures that any cognitive performance, belief in particular, can be understood and explained in the vocabulary and network of generalizations appropriate to computational competence. Belief, according to Fodor, is an organism's attitude toward a formula in LR. That formula expresses the content of the belief. The attitude itself is to be thought of as empirically identical to a computational relation. Here comes the twist: The computational relation is understood in terms of specific operations on formulae such as storing, recalling, accepting, and the like. The emphasis now is on competence or software (pp. 26, 166–168, 200, etc.).

This is the right way to go. Perhaps more than anybody else in this business, Fodor has boldly and insightfully exploded that vague, opaque and rather lazy talk of belief as mental state or disposition, period, and has offered instead a concrete and fruitful proposal. But it is not the right way to *start*. Fodor's definition provides the "computational reality" of a belief but not its identity. Only when the latter is antecedently determined, as a type, can the former redescribe it. Once this reductive move is made, a story can be told in terms of memory storage, motor commands, verbal responses, and the like. The computational mechanics takes over. This is precisely where cognitive psychology comes into the picture. But this is no longer an explanatory story of belief although it is, at a lower level, an explanatory story of what makes belief computationally possible. The explanatory integrity of belief, recognized only by folk psychology or some sort of performance interpretation, is now dissolved, since redescribed token-wise, into the local computational workings of the mind. No need, therefore, to advert to the initial, interpretational insights or concepts. Whence the absence of 'belief' from the professional talk. What all this boils down to is a *FC for belief* which cognitive psychology follows both in spirit and letter: Reconstruct (and talk of) aspects of believing only as aspects of computations.

If you sense a *déjà vu,* you sense right. This is the story of content all over again. No wonder. A belief *is* a content in a functional capacity. Content, as Fodor reminds us, is in general to be characterized functionally, in terms of its role in cognition and behavior. It is the specification of this role that makes the difference. If the specification is central and global, with an eye to the overall performance of the organism as centrally initiated, then we say that the content functions as a belief. Such a specification, which is the business of an interpretation of performance, comes in a conceptual package concerned with connecting belief to the other mental processes and

propositional attitudes such as plans, goals, actions, values, decisions, intentions, and the like. This is a familiar point often made by philosophers. It may be clarified and strengthened by another familiar argument concerning human action.

Like a mental state, an action can be viewed from several angles, say, neurological, physical, possibly computational. An action can be counted as *intentional.* We say, for instance, that his pushing the man in front of him was intentional (he wanted to see what was going on) although his pushing the dean was not (he did not know the man was the dean and would not have pushed him had he known that) just as presumably various muscle motions and neural firings involved cannot be seen as intentional. Same action, or series of physical events, different interpretations only one of which yields intentionality. Obviously a physical, neurological, or even computational story cannot be the one to tell us where and how to attribute intentionality to an action. Rather, it is an alignment of action to central agency and control, to awareness and hence the agent's representation of the situation, to his beliefs and goals, that accounts for the attribution. It takes a particular interpretation to produce this conceptual alignment. The latter, to pursue my earlier metaphor, is a horizontal enterprise, quite oblivious to its vertical tokenings.

The same is true of belief. We attribute believingness to contents the same way we attribute intentionality to actions: by aligning content to the relevant parameters (attitudes, states, processes) in the horizontal package offered by an interpretation of performance. There will be contents whose function may be so local that no alignment can make them into beliefs, just as there are actions which fail to engage the parameters that will make them intentional. Yet in the cases of both belief and action, their computational and neural tokenings cannot tell the difference. Believingness as well as intentionality of action emerge elsewhere. It takes a performance interpretation to get there.

V. CONCLUSION

It is now time to put things in perspective. It may turn out, as I think it will, that mental content can be attributed only when, simultaneously, a batch of interacting propositional attitudes are also attributed. If content *can* be attributed independently, then people will likely cease calling it 'content' or else will cease attributing it to an organism as such. In that case, semantic interpretation and performance interpretation will turn out to be two sides of the same enterprise. We now call the enterprise 'folk psychology' or think of it along its lines. We should neither exaggerate nor minimize its importance. The right understanding of Fodor's pioneering work can tell us why.

Before Fodor and the recent work in the foundations of computational cognition, the dominant contrast has been that between a populist, fluid, intuitive, externalist understanding of mind and cognition, something like folk psychology, and the reductionist, brute-facts approach of neuroscience. This had been a simplistic contrast which deserved simplistic reactions such as behaviorism, dualism, and type physicalism. What cognitive psychology has produced, and what computational metareflections such as Fodor's have justified, is the insertion of an autonomous level of theorizing about the mind between the two extremes just mentioned, a level which makes the mental inside intelligible for the first time while closing the gap, since mediating, between the folk psychological or interpretational attribution of content and propositional attitudes, on the one hand, and their material tokenings, on the other. This computational insertion is a remarkable development which Fodor documents and rationalizes so well and vigorously. If on the right track, my critical remarks indicate that the new computational gambit cannot fully absorb the attributional strategies of semantic and performance interpretation. In particular, they indicate that the latter are needed to determine the formally and functionally equivalent classes which, computationally, token contents and beliefs. In brief, the thrust of my remarks has been that, in creating a new level of theorizing, CTM further relativizes the distinction between type and token, between explanation and description. Indeed, the computations CTM posits provide new types for neural tokens but in turn become tokens relative to the higher level types posited by semantic and performance interpretation. Somehow Fodor appears to think that the latter, upward relativization does not affect computationalism because, presumably, it can absorb the interpretational enterprises. He has not shown this. Interpretation remains independently important. Yet, on the other hand, its role should not be exaggerated. The general lesson to be learned from Fodor's work is that *new* theoretical insertions may occur between the computational and the interpretational levels of understanding the mind, thus eating into the latter's once vast territories and regimenting its insights into more rigorous patterns of explanation. How far this will go it is hard to predict. But one thing is sure: Such new insertions do not amount to reductions. It is the merit of Fodor's work, of computationalism in general, to have shown that a better explanation of mind is often measured by its distance from, not proximity to, neurological and physical theorizing. Computationalism itself must be wise and learn modesty from this lesson.

Representations: Philosophical Foundations of Cognitive Science. By Jerry A. Fodor. Cambridge: MIT Press, 1981, 356 pp.