

## Levels Indeed! A Response to Broadbent

David E. Rumelhart and James L. McClelland  
Institute for Cognitive Science, University of California, La Jolla

Although Broadbent concedes that we are probably correct in supposing that memory representations are distributed, he argues that psychological evidence is irrelevant to our argument because our point is relevant only at what Marr (1982) has called the *implementational* level of description and that psychological theory is only properly concerned with what Marr calls the *computational* level. We believe that Broadbent is wrong on both counts. First, our model is stated at a third level between the other two, Marr's *representational* and *algorithmic* level. Second, we believe that psychology is properly concerned with all three of these levels and that the information processing approach to psychology has been primarily concerned with the same level that we are, namely, the algorithmic level. Thus, our model is a competitor of the logogen model and other models of human information processing. We discuss these and other aspects of the question of levels, concluding that distributed models may ultimately provide more compelling accounts of a number of aspects of cognitive processes than other, competing algorithmic accounts.

Broadbent (1985) has generously conceded that memory is probably represented in a distributed fashion. However, he has argued that psychological evidence is irrelevant to our argument because the distributed assumption is only meaningful at the *implementation* (physiological) level and that the proper psychological level is the *computational* level.

Broadbent has raised an extremely important issue, one that has not generally received explicit attention in the psychological literature, and we applaud his attempt to bring it into focus. However, the issue is very complex and deserves very close scrutiny. Indeed, more levels must be distinguished than Broadbent acknowledges, and there are many more constraints among levels than he supposes.

We begin by pointing out that Broadbent has ignored a third level of theoretical description, the *algorithmic* level, which is a primary level at which psychological theories are stated. We then suggest that his analysis of our arguments fails to establish his claim that our model and traditional models are not competitors at the same level. We then describe other senses of levels, including one in which higher level accounts can be said to be convenient approximations to lower level accounts. This sense comes closest to capturing our view of the relation between our distributed model and

traditional information processing models of memory.

### Marr's Notion of Levels

Broadbent begins his argument by appealing to the analysis of levels proposed by David Marr (1982). Although we are not sure that we agree entirely with Marr's analysis, it is thoughtful and can serve as a starting point for seeing where Broadbent's analysis went astray. Whereas Broadbent acknowledges only two levels of theory, the computational and the implementational, Marr actually proposes three, the computational, the algorithmic, and the implementational levels. Table 1 shows Marr's three levels. We believe that our proposal is stated primarily at the algorithmic level and is primarily aimed at specifying the representation of information and the processes or procedures involved in storing and retrieving information. Furthermore, we agree with Marr's assertions that "each of these levels of description will have their place" and that they are "logically and causally related." Thus, no particular level of description is independent of the others. There is thus an implicit computational theory in our model as well as an appeal to certain implementational (physiological) considerations. We believe this to be appropriate. It is clear that different algorithms are more naturally implemented on different types of hardware, and therefore information about the implementation can inform our hypotheses at the algorithmic level.

Broadbent's failure to consider the algorithmic

---

Request for reprints should be sent to David E. Rumelhart, Institute for Cognitive Science C-015, University of California—San Diego, La Jolla, California 92093.

Table 1  
*The Three Levels at Which any Machine Carrying Out Information Processing Tasks Must be Understood (Marr, 1982)*

Computational theory	Representation and algorithm	Hardware implementation
What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?	How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation	How can the representation and algorithm be realized physically?

*Note.* From *Vision* by D. Marr. W. H. Freeman and Company © 1982. All rights reserved. Reprinted by permission.

level is crucial, we believe, because this is the very level at which information processing models (including Morton's, 1969, logogen model) have been stated. Computational models, according to Marr, are focused on a formal analysis of the problem the system is solving—not the methods by which it is solved. Thus, in linguistics, Marr suggests that Chomsky's (1965) view of a competence model for syntax maps most closely onto a computational level theory, whereas a psycholinguistic theory is more of a *performance* theory concerned with how grammatical structure might actually be computed. Such a theory is concerned with the algorithmic level of description. It is the algorithmic level at which we are concerned with such issues as efficiency, degradation of performance under noise or other adverse conditions, whether a particular problem is hard or difficult, which problems are solved quickly, which take a long time to solve, how information is represented, and so on. These are all questions to which psychological inquiry is directed and to which psychological data is relevant. Indeed, it would appear that this is the level to which psychological data primarily speaks. At the computational level, it does not matter whether the theory is stated as a program for a Turing machine, as a set of axioms, or as a set of rewrite rules. It does not matter how long the computation takes, or how performance of the computation is affected by performance factors such as memory load, problem complexity, and so on. It does not matter how the information is represented, as long as the representation is rich enough, in principle, to support computation of the required function.

The question is simply what function is being computed, not how is it being computed.

Marr recommends that a good strategy in the development of theory is to begin with a careful analysis of the goal of a particular computation and a formal analysis of the problem that the system is trying to solve. He believes that this top-down approach will suggest plausible algorithms more effectively than a more bottom-up approach. Thus, the computational level is given some priority. However, Marr certainly does not propose that a theory at the computational level of description is an adequate psychological theory.

As psychologists, we are committed to an elucidation of the algorithmic level. We have no quarrel with Marr's top-down approach as a strategy leading to the discovery of cognitive algorithms, though we have proceeded in a different way. We emphasize the view that the various levels of description are interrelated. Clearly, the algorithms must, at least roughly, compute the function specified at the computational level. Equally, the algorithms must be computable in amounts of time commensurate with human performance, using the kind and amount of hardware that humans may reasonably be assumed to possess. For example, any algorithm that would require more specific events to be stored separately than there are synapses in the brain should be given a lower plausibility rating than those that require much less storage. Similarly, in the time domain, those algorithms that would require more than one serial step every millisecond or so would seem poor candidates for implementation in the brain (Feldman & Ballard, 1982).

To summarize, Broadbent's claim that our model addresses a fundamentally different level of description than other psychological models is based on a failure to acknowledge the primary level of description to which much psychological theorizing is directed. At this level, our model should be considered as a competitor of other psychological models as a means of explaining psychological data.

#### Different Models or Different Levels?

But Broadbent offers more specific arguments aimed at establishing that distributed models are not competitors of psychological models. First, he suggests that both distributed models and local models are computationally equivalent, in that distributed systems are not capable of any computation that cannot also be performed in a system of localized storage. But this form of computational equivalence is too weak, if we are interested in specifying the representations and procedures used. Different sorting algorithms may be computation-

ally equivalent in just Broadbent's sense but may differ in whether the sorting time increases exponentially or merely log linearly with list length. Casual appeals to Turing's thesis may be sufficient to establish equivalence at the computational level, but they are generally of little use to us in psychology, precisely because we are concerned with stronger forms of equivalence of our models to psychological processes.

Broadbent then discusses whether our model could possibly offer alternatives to localist accounts of memory such as logogen models and prototype-exemplar theories of concept learning. After pointing out the insights that logogen theory captures—insights that we, of course, appreciate and have incorporated into our thinking (McClelland & Rumelhart, 1981; Rumelhart, 1977, Rumelhart & McClelland, 1981, 1982)—he goes on to point out that the proliferation of logogens that was occasioned by modality and even format specific repetition effects may not be necessary after all. One can argue that the effects of a prior stimulus pattern are not localized in the logogen itself but are distributed throughout the pathway over which the stimulus is being processed. We cannot understand why he thinks we would find this observation to be devastating to our argument. Our argument was based on an extension of the existing findings to the conclusion that virtually any alteration of a stimulus or the conditions of its presentation (differing context, etc.) that altered the pattern of activation produced internally by the stimulus (we should have stressed, in any of the several modules in which the item would give rise to a pattern of activation) would influence repetition effects. We then showed that a distributed model would automatically capture this, without encountering the need to proliferate logogens for each individual variation of presentation conditions.

All of this is not to say that we believe that it will be easy to distinguish distributed and localist models empirically; especially if, as in Broadbent's localist account, the localist model includes separate representations of both exemplars and prototypes. Actually, as we thought we made clear, there are localist models that do not require local representations of prototypes to account for Whittlesea's data (Hintzman, 1983; Whittlesea, 1983). But to say that two models make the same predictions for a given set of data is not to say that one should be seen as an implementation of the other. The models may simply be alternative constellations of assumptions about representation and process that happen to make identical or approximately equivalent predictions over some of the cases to which they can be applied.

In summary, equivalence at the computational level, the existence of ways to avoid proliferation

of logogens within logogen theory, and the possibility that local models can under some conditions account for the same data as distributed models does not prove the case that distributed models are implementations of other cognitive models.

#### Other Notions of Levels

Yet we do believe that Broadbent is partly right when he says that our distributed model (and the class of models we have come to call *parallel distributed processing* models) are at a different level than models such as the logogen model, or prototype theories, or schema theory. The reason is that there is more twixt the computational and the implementational than is dreamt of, even in Marr's philosophy. Many of our colleagues have challenged our approach with a rather different conception of levels borrowed from the notion of levels of programming languages. It might be argued that Morton's logogen model is a statement in a *higher level* language analogous, let us say, to the Pascal programming language and that our distributed model is a statement in a *lower level* language that is, let us say, analogous to the assembly code into which the Pascal program can be compiled. Both Pascal and assembler, of course, are considerably above the hardware level, though the latter may in some sense be closer to the hardware, and more machine dependent than the other.

From this point of view one might ask why we are mucking around trying to specify our algorithms at the level of assembly code when we could state them more succinctly in a high level language such as Pascal. We believe that most people who raise the levels issue with regard to our models have a relation something like this in mind. Indeed, we suspect that this notion of levels may be rather closer to Broadbent's own than the notion of levels one finds in Marr. Like Broadbent, people who adopt this notion have no objection to our models. They only believe that psychological models are more simply and easily stated in an equivalent higher level language—so why bother?

We believe that the programming language analogy is very misleading. The relation between a Pascal program and its assembly code counterpart is very special indeed. Pascal and assembly language necessarily map exactly onto one another only when the program was written in Pascal and the assembly code was compiled from the Pascal version. Had the original programming taken place in assembler, there is no guarantee that such a relation would exist. Indeed, Pascal code will, in general, compile into only a small fraction of the possible assembly code programs that could be written. Because there is presumably no compiler

to enforce the identity of our higher and lower level descriptions in science, there is no reason to suppose that there is a higher level description exactly equivalent to any particular lower level description. We may be able to capture the actual code approximately in a higher level language—and it may often be useful to do so—but this does not mean that the higher level language is an adequate characterization.

There is still another notion of levels that illustrates our view. This is the notion of levels implicit in the distinction between Newtonian mechanics on the one hand and quantum field theory on the other.<sup>1</sup> It might be argued that a model like Morton's logogen model is a macroscopic account, analogous to Newtonian mechanics, whereas our model is a more microscopic account analogous to quantum field theory. Note, over much of their range, these two theories make precisely the same predictions about behavior of objects in the world. Moreover, the Newtonian theory is often much simpler to compute with, because it involves discussions of entire objects and ignores much of their internal structure. However, in some situations the Newtonian theory breaks down. In these situations, we must rely on the microstructural account of quantum field theory. Through a thorough understanding of the relation between the Newtonian mechanics and quantum field theory we can understand that the macroscopic level of description may be only an approximation to the more microscopic theory. Moreover, in physics, we understand just when the macro theory will fail and when the micro theory must be invoked. We understand the macro theory as a useful formal tool, by virtue of its relation to the micro theory. In this sense, the objects of the macro theory can be viewed as emerging from interactions of the particles described at the micro level.

In our article (and elsewhere) we have argued that many of the constructs of macrolevel descriptions such as logogens, schemata, and prototypes can be viewed as emerging out of interactions of the microstructure of distributed models. (See Rumelhart, Smolensky, McClelland, & Hinton, in press, for a further discussion of the nature of this emergence.) We view macrolevel theories as approximations to the underlying microstructure that the distributed model presented in our article attempts to capture. As approximations they are often useful, but in some situations it will turn out that a lower level description may bring deeper insight. Note for example, that in the logogen model one has to know how many logogens there are. Are there three logogens for a word or just one, or are there perhaps two and a half? In our distributed formulation no such decision need be

made. Because the analog to a logogen is not necessarily discrete but simply something that may emerge from interactions among its distributed parts, there is no problem with having the functional equivalent of half a logogen. Thus, although we imagine that Morton's logogen model, schema theory, prototype theory, and other macrolevel theories make all more or less valid approximate macrostructural descriptions, we believe that the actual algorithms involved cannot be represented precisely in any of those macrolevel theories. They will require a distributed model of the general kind described in our article.

### Conclusion

An understanding of human memory requires an analysis at many different levels. Marr's computational and implementation levels specify the boundaries of psychological inquiry. In between, at what Marr calls the algorithmic level, lies the heart of our concerns. However, our examination of analogies to computer science and physics suggests that we may well need to consider many sublevels of analysis within the algorithmic level. At a macroscopic level, a coarse-grained analysis in which concepts, logogens, prototypes, schemata, and so forth are treated as undecomposable wholes may lead to insights; the discovery of the basic level (Rosch, Mervis, Gray, Johnson, & Boyes-Brian, 1976) would seem to be a case in point. But when we look closely, both at the hardware in which the algorithms are implemented and at the fine structure of the behavior that these algorithms are designed to capture, we begin to see reasons why it may be appropriate to formulate models that come closer to describing the microstructure of cognition. The fact that models at this level can account as easily as our model does for many of the facts about the representation of general and specific information makes us ask why we should view constructs like logogens, prototypes, and schemata as anything other than convenient approximate descriptions of the underlying structure of memory and thought. Finally, it should be said that the decision between our particular formulation and those implicit in the logogen model and other competitors will not be made on the basis of general considerations about what level of theories psychologists should or should not be interested in. Ultimately, the worth of our formulation will be determined by how useful it is at explaining the facts of memory storage and retrieval and to what degree it leads to new and fruitful insights.

<sup>1</sup> This analogy was suggested to us by Paul Smolensky.

## References

- Broadbent, D. (1985). A question of levels: Comment on McClelland and Rumelhart. *Journal of Experimental Psychology: General*, 114, 189-192.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science*, 6, 205-254.
- Hintzman, D. (1983, June). *Schema abstraction in a multiple trace memory model*. Paper presented at conference on "The priority of the specific," Elora, Ontario, Canada.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of the effect of context in perception, Part I. An account of basic findings. *Psychological Review*, 88, 375-407.
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76, 165-178.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Brian, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
- Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI*. Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., & McClelland, J. L. (1981). Interactive processing through spreading activation. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes in reading*. Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of the effect of context in perception, Part II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (in press). Models of schemata and sequential thought processes. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume II: Applications*. Cambridge, MA: Bradford Books.
- Whittlesea, B. W. A. (1983). *Representation and generalization of concepts: The abstractive and episodic perspectives evaluated*. Unpublished doctoral dissertation, MacMaster University, Hamilton, Ontario.

Received December 17, 1984

Revision received December 17, 1984 ■

---

*Psychological Documents* to be Discontinued

At its February 2-3, 1985, meeting, the Council of Representatives voted to cease publication of *Psychological Documents* (formerly the Journal Supplement Abstract Service) as of December 31, 1985, with the publication of the December 1985 issue of the catalog. Continued low submissions, decreasing usage, and rising costs for fulfillment of paper and microfiche copies of documents were reasons given for discontinuing publication of the alternative format publication, which was begun in 1971 as an "experimental" publication.

Authors who wish to submit documents for publication consideration must do so by July 1. Authors who are currently revising documents at the request of the editor should complete all revisions and submit them for final review as soon as possible, but no later than July 1. Orders for paper and microfiche copies of documents presently in the system and of those documents entered during 1985 will continue through December 31, 1986.

---