recognition procedure. But then that has to be specified. There are other problems as well. Both analogies suggest that the pieces of information in this model are discrete – one piece gets put on top of another piece. But how, in this system, can we account for the distortions that are so common in memory? In many instances distortions arise because of a confusion between two items; however, in some instances distortions result from a combination of elements from different sources. For example, in one study (Loftus 1979), subjects saw slides of people going about routine activities. One person in the slides was seen reading a book with a green cover. Later, a leading question suggested that the book had actually been blue. When subjects were finally asked to select the remembered color on a color wheel their choices tended to be a compromise between what they actually saw and what they were told later on their questionnaire. How would "the man at the desk" make these sorts of errors? He could not simply file a letter in the wrong file folder. An analogous error might be one in which the man at the desk recalled a letter received by Mary Smith as one that had been signed by Jane Doe.

In any event, Broadbent's Maltese cross focuses its emphasis toward the shorter-term memories and away from long-term memories and the well-known and interesting distortions of the latter. In fact, Broadbent has rather little to say about long-term memory, except to stress its associative nature. Other recent models take the opposite tack (e.g. CHARM; Eich 1982). They can successfully account for a number of long-term memory phenomena. We see no reason why one or more of these long-term memory models could not be successfully combined with the Maltese cross, perhaps by some operation like "convolution" (Eich 1982), to provide a more complete explanation of the human information-processing system.

ACKNOWLEDGMENTS
The writing of this commentary was supported by National Science Foundation grants to E. Loftus and G. Loftus, and by an Office of Naval Research contract to E. Hunt.

## The homunculus as bureaucrat

Alan K. Mackworth
*Department of Computer Science, University of British Columbia, Vancouver, B.C., Canada V6T 1W5*

The Maltese cross metaphor offered as a replacement for the outmoded telephone exchange and pipeline metaphors might better be called "the homunculus as bureaucrat." The image of a gnomish, desk-bound homunculus is indeed appealing. As Broadbent points out, introducing a homunculus does not lead to an infinite regress if he is confined to executing a set of well-defined rules. The interpreter for that set of rules may well be a second rule-governed homunculus defining the virtual machine for the first; however, that recursion eventually bottoms out on the real machinery of the brain. This view of the structure of intelligent systems is the communal metaphor of artificial intelligence. The primary scientific task now is the determination of the actual computational architecture of those levels.

There are many universal computing devices; they are all equivalent in theoretical power but differ markedly in their expressive power for building cognitive systems and in their suitability for implementing on neural hardware. The two poles of the cognitive task and the real machine exist in a space of possible architectures, creating a tension within it. That space can and must be spanned by a hierarchy of machines, but workers tend to focus their efforts near one pole or the other. Those clustered around the task pole favor production systems (Newell & Simon 1972) or schemata (Havens & Mackworth 1983) while those nearer the real machine pole favor relaxation-based cooperative computation (Marr 1982) and connectionist schemes (Hinton & Anderson 1981). The research strategies of theoretical and experimental psychologists must be designed to identify the most appropriate architecture for the top-level cognitive machine and to specify a necessary and sufficient set of primitives for that machine. Broadbent's target article should be judged by its success in defining and achieving that goal.

Although Broadbent tries to disarm his potential critics by modestly presenting a simplistic theory, in broad brush strokes, some of the details that are present can be questioned. The system architecture is problematic. Broadbent favors production systems but does not pay sufficient attention to their structure. His proposal is in part modeled on the von Neumann model of a stored program computer with separated processing system and memory. In that model the instructions are stored in memory with the memory location of the next instruction stored in the processor's program counter. Since location-addressable memory is assumed, retrieving the next instruction from memory is a single fetch operation. However, Post production systems require a different architecture. Broadbent is not clear on the contents of his long-term associative memory. He says it "retains a running total of the number of times Event A and Event B have occurred together." But later he says, "Although the rules on which it acts are held in long-term memory, it is left to the central control system to recognise the pattern of conditions appropriate to each of the possible rules." Hence the central processing unit must permanently store all the condition predicates in order to carry out the function of the program counter in the stored program model and retrieve the applicable rule(s) from memory. Since the set of rules has no apparent internal structure, the processor must evaluate a very large number of predicates in a very short time to decide what to do next. Moreover, Broadbent states categorically that the learned equivalence of "2" and "two" is *not* stored in the associative long-term memory, but in the processor itself. All this suggests that the serial processor and separate memory architecture proposed is simply inappropriate for the cognitive machine.

Broadbent provides an excellent summary of some of the discontent with the pipeline or linear stage model, but one is still left with a feeling that the new model is a hybrid of the telephone exchange and pipeline models dressed up in new computational clothes. The confusion over the nature and location of "memory" seems to originate from the associationist exchange model. Most of the experimental data provided derive from a school of psychology firmly embedded in the pipeline model. Many of those experiments ask questions that are simply not meaningful in a computational model of perception and cognition. For example, the "logogen" pipeline theory of word recognition is offered as the model for "the pattern recognition system," an otherwise unexplained subhomunculus living inside the central processing system. The notion of a sensory store of "relatively raw input" paradoxically indexed, for sound, by speaker ("the nature of the voice") is another example. One can also discern an attempt to smuggle in the pipeline theory of attention ("the selection of one part of sensory store for further transmission"). Perhaps, as Neisser (1976, p. 84) says, "when perception is treated as something we do rather than as something thrust upon us, no internal mechanisms of selection are required at all."

However appealing the bureaucrat at his desk may be, the metaphor still conveys too passive an image of the perceiver. His filing cabinet stores statistical correlations between events; his many in-baskets fill up with paper, independently of his own purposes, plans, and actions, which he must then attend to or ignore. Requirements for an active, schema-based cycle of perception theory have been outlined elsewhere for psychology (Neisser 1976) and artificial intelligence (Mackworth 1978). The psychology of mind must undergo a thoroughgoing revision more radical than we see here if it is to adopt the computational model.