Seeking a Computational Brain

Paul Buchheit
Harold Washington College
30 E. Lake St.
Chicago, IL 60601

Abstract

This paper compares the various theories of mental processing, and suggests that some form of computationalism drives our thought processes. Since we know little of the actual neurophysiology behind the workings of the mind, we have to hypothesize the likely structures based on higher-level functions. Propositional representation, word concept localization, the abstraction of general concepts, and the reification of parts into a unified whole are posited as high-level functions that might help reveal some of the details of brain organization. The INFANT System is based on such representational techniques. It is shown in this paper how these techniques can be used to simulate the learning environment of a small child, and perhaps to stimulate thinking about the lower levels of mental processing.

1. Introduction

The idea of computationalism in the brain is a very controversial issue, having implications for the existence of mind, self, and consciousness. For centuries learned philosophers and scientists have disagreed on matters of the mind, and today's leading thinkers, despite -- or perhaps because of -- recent developments in neurophysiology and computer modelling, are equally disunited. Figure 1 attempts to classify the various philosophical theories of the mind. Compiled from a number of sources [1,2,3,4,5,6,7,8,9], it depicts the relationships among theories and many of their best-known proponents:

A summary of the main theories:

Dualism (anti-materialism): The mind and brain are entirely separate concepts

Functionalism: The brain is a computational device that is functionally independent of neural matter, and can therefore be expressed in any of a number of different substrates, including perhaps the silicon of a computer.

Computationalism: The workings of the mind can be explained in algorithmic terms, as on a computer (specifically a Turing machine). It is usually synonymous with *functionalism*, but may also refer to the neural computations of connectionist systems.

Materialism (anti-dualism): The brain is dependent on its neurobiological architecture, and possibly on its substrate. It represents information as distributed patterns of analog activity. Although in theory the brain may be definable, at least in part, as a computational device in a different substrate, in practice it cannot work without the biological matter that allows it to function in real time.

Behaviorism: an extreme materialist viewpoint that contests the existence of internal mental processes, and claims that all human activity can be explained by input stimuli and output responses.

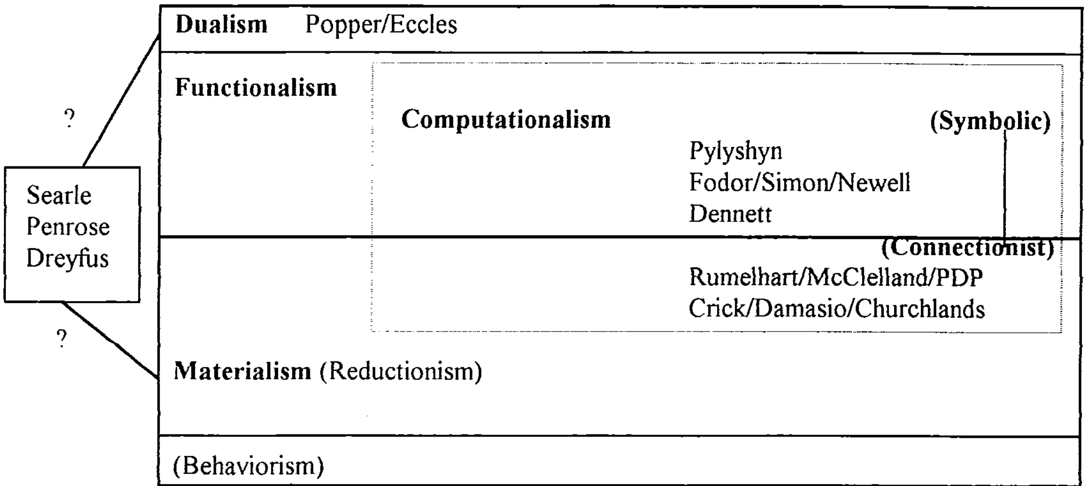

Figure 1 -- Theories of the Mind

2. Toward Computationalism

It is recognized that such categorizations are themselves subject to dispute, and may not reflect the subtleties and nonconformities of opinion of the people named in the chart. But the purpose of the graphic is to show the wide range of viewpoints about mental processes. In particular, it depicts the ample support for computationalism among both philosophers and neurobiologists. There seems to be a general agreement in the literature (or at least a lack of disagreement) that some brain processes will eventually be represented in neural networks, or in a hybrid symbolic/connectionist device.

The modern-day computer, as powerful as it seems to us, is still a crude imitation of the machinery of the mind: certainly faster, but lacking the plasticity and massive parallelism that characterize the brain. If and when we develop machines with parallel, cooperative neural networks, depth- and orientation-sensitive pattern recognizers, and sensory input-output devices, we might be prepared to assemble a model of humanlike intelligence.

But according to the non-computationalists, we might not. Searle's Chinese Room argument [10] suggests that an automated process, even if performing an apparently intelligent task, never really "understands" what it is doing. But the argument is circular. The task used to demonstrate his point -- automaton-like language translation -- has never been done, and if ever accomplished it will be difficult to argue that the requisite blending of syntactic, semantic, and pragmatic knowledge is anything short of an "understanding" process. Yet Searle uses this currently unperformable understanding process as part of his argument against any such task ever being performed.

The role of the cognitive scientist is to consider the range of opinions, examine any new evidence, form a hypothesis about the nature of the computational method behind the brain's representationalism, and test it with our constantly expanding computer power. Whether the neurobiological substrate is a necessary part of an architecture for intelligence, or if instead the algorithms of intelligence are independent of neural matter, the symbols of human communication are encoded in the brain in a manner that may be discernible at an abstract, non-neural level.

Unfortunately, we do not yet have the cooperative networks, pattern recognizers, and sensory I/O devices needed for the job. But we have considerable knowledge of the higher-level functions of the brain. We can state with some certainty that:

(1) *Propositional* representation plays a significant role in comprehension [11].

(2) Parts of speech (nouns and verbs) are *localized* in the brain [3].

(3) Children learn language by *abstracting* general patterns and word concepts from spoken input [6].

(4) The brain uses *recursion* to process sentences with nested clauses [12]. Since short-term memory holds only 4-7 chunks of information [2,4], long-term memory must use a *reification* technique to combine sentence parts into a simpler form.

3. The INFANT Approach

Recent work on the INFANT System [13,14,15,16] has used these hypotheses and observations to facilitate language processing on a computer. INFANT has been implemented as a symbol-based system in three phases: (1) as a conversational system for communication at the level of a child; (2) as a NLP interface to a single-user operating system; and (3) as a working cognitive model of language comprehension. The earliest application provided the foundation for current work with the original *propositional network* concept. For the second application, the knowledge domain was restricted to typical data management activities such as file movement and disk searching. An extended lab experiment was conducted in which students accomplished a variety of such tasks through the use of the INFANT Talk Program. In the third and ongoing phase, INFANT defines a model for language understanding through the massive interaction of standardized and semantic primitive-based propositional forms. A syntactic/semantic translation procedure converts sentences into a tree structure of propositions called *Hierarchical Logical Form* (HLF), which lends itself to an algorithmic compositional analysis of similarly-structured sentence parts. For example, the sentence
The boy with the little dog is in the house
can be mapped to the HLF:

BOY_x IN HOUSE

BOY_x WITH DOG_y

DOG_y LITTLE

The propositional segments, which in the INFANT System are likened to collections of neural units in a connectionist network, are linked through symbolic, *reified* representations of subordinate propositional parts. In the above example the subordinate proposition [DOG_x LITTLE] is reified to DOG1, and the subordinate proposition [BOY_x WITH DOG1] is reified to BOY1, resulting in an overall sentence representation of [BOY1 IN HOUSE]. All possible generalized forms of the sentence are maintained in the knowledge base. These would include:

(BOY1 IN PLACE)
(BOY IN HOUSE)
(BOY IN PLACE)
(PERSON1 IN HOUSE)
(PERSON1 IN PLACE)
(PERSON IN HOUSE)
(PERSON IN PLACE)

Reification, inheritance, generalization, and the standardization of the propositional form work together to facilitate understanding. Any subsequent information about a person who is *in the house* will generate a link to the reification term PERSON1, and from there to the possibility of that person being "with a dog." Similarly, any further reference to a person *with a dog* will link to the possibility of that person being "in the house." In order for this process to work, a vast degree of *implicit* forward chaining is built into the declarative structures. Each Subject-Predicate-Object proposition can be linked to many others as a result of reification and deductive inferencing. The introduction of new sentence information triggers a connectionist-like flurry of activity among related propositions. As a result of this activity, which involves mainly the interaction of new and old memory-based inferences, the propositional knowledge base reaches a state of equilibrium in which new assertions are in accordance with past beliefs. Thus, over time, knowledge is gained only through the competing influences of very many propositional units of varying strength and specificity.

4. Computational Processing with INFANT

As an example of the INFANT method, the following sentences were used as a training set for a sample run:

The baby is in the playpen.
(BABY IN PLAYPEN)
The big dog is in the doghouse.
(DOG1 IN DOGHOUSE)
{DOG1: (DOG BIG _)}
The big dog barks.
(DOG1 BARK _)
The boy with the little dog is in the house.
(BOY1 IN HOUSE)
{BOY1: (BOY WITH DOG2)}
{DOG2: (DOG LITTLE _)}
The baby sees the little dog.
(BABY SEE DOG2)
The little dog barks.
(DOG2 BARK _)
The little dog is near the playpen.
(DOG2 NEAR PLAYPEN)
The baby is in the playpen.
(BABY IN PLAYPEN)
The big dog is near the playpen.
(DOG1 NEAR PLAYPEN)
The baby cries
(BABY CRY _)

The sentences were entered into the system repeatedly, and in random order except for an intentional bias toward the likelihood of the baby crying in the presence of the big dog. After a number of passes through the data, INFANT was able to acquire a great deal of information through propositional association, generalization, and inductive inferencing. All propositions and their generalized forms were linked to subsequent propositions and evaluated for possible dependencies. For example, the following data was recorded after the processing of 55 sentences:

(BABY IN PLAYPEN) *{sentences 5 7 11 13 18 23 27 32 37 43 44 47 48 51 52}*

→ *(baby cry _)(54 38 33 25 19)* **(A)**

→ *(dog near playpen) (53 28 24 17 12)* **(B)**

→ *(dog1 near playpen)(53 24)* **(C)**

(DOG1 NEAR PLAYPEN) *{sentences 10 17 24 31 36 40 53}*

→ *(baby in playpen)(37 32 18 11)* **(D)**

AND (baby in playpen)
→ *(baby cry _)(54 38 33 25 19)* **(E)**

(DOG NEAR PLAYPEN) *{sentences 4 10 12 17 24 28 31 34 35 36 40 50 53}*

→ *(baby in playpen) (37 32 18 13 11 5)* **(F)**

AND (baby in playpen)
→ *(baby cry _)(54 38 33 25 19)* **(G)**

As a result, it *might* be determined from (E) above that the simultaneous presence of the big dog (DOG1) and the BABY in/near the playpen caused the baby to cry at 5 different times. It can *definitely* be determined from (C) and (D) that DOG1 and BABY were simultaneously in/near the playpen at 6 different times. Also, it can *definitely* be seen from the frequency of (BABY IN PLAYPEN) and from (A) that just being in the playpen did not cause the baby to cry. Thus it can be concluded that (E) is a valid inductive inference. If BABY and DOG1 are both in/near the playpen, BABY will cry.

On the other hand, it can be determined from (B) and (F) that DOG (ie, *any* dog) and BABY were simultaneously in/near the playpen at 11 different times. But BABY only cried the 5 times DOG1 was nearby. Thus it *cannot* be more generally inferred that BABY will cry when *any* DOG is near the playpen.

In this example, the various simulated computational aspects of high-level brain functionality (propositions, localization, abstraction, reification) contribute to ease and efficiency of usage:

(1) the *propositional* form: it is more efficient to work with the single proposition (DOG1 NEAR PLAYPEN) than with the compound proposition (*the big dog is near the playpen*).

(2) the *localization* of inheritable word concepts such as IN, WITH, SEE, and NEAR:

(BABY IN PLAYPEN)
→
(BABY NEAR PLAYPEN)

and, to a lesser extent,

(BABY SEE DOG)
→
(BABY NEAR DOG) .

(3) *abstracting* general concepts: in the example, the baby is able to inductively learn that (DOG BARK _) -- that is, dogs in general bark -- based on the gradual accumulation of data such as *The little dog barks* and *The big dog barks.*

(4) *reification*: as explained above, the baby learns through an inductive process that proximity to the big dog is undesirable:

(DOG1 NEAR PLAYPEN)
AND
(BABY IN PLAYPEN)
→
(BABY CRY _)

Thus the reification method facilitates the inferencing process by clearly differentiating among the word concepts DOG1, DOG2, and DOG.

5. Conclusion

Rule-based representations cannot substitute for the temporally-managed internal images that probably drive our mental processes. But we can use them to simulate the computational processing within the biological structure. By modeling the higher-level brain functions on a computer, we may, in a type of reverse engineering process, be able to speculate on the nature of the lower-level details that so far have escaped us. Current plans for the INFANT System include a comprehensive training session with input similar to that described above. This will incorporate the conversational system of Phases 1 and 2 with the more connectionist-like approach of Phase 3, and will make it possible to evaluate the effects of a long-term learning process.

Bibliography

1. Peter Baumgartner and Sabine Payr (1995). *Speaking Minds: Interviews with Twenty Eminent Cognitive Scientists* (Princeton University Press).

2. Patricia S. Churchland and Terrence J. Sejnowski (1992). *The Computational Brain* (The MIT Press, Cambridge MA).

3. Paul M. Churchland (1995). *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain* (The MIT Press, Cambridge MA).

4. Francis Crick (1994). *The Astonishing Hypothesis: The Scientific Search for the Soul* (Touchstone, Simon & Schuster, New York).

5. Antonio R. Damasio (1994). *Descartes' Error: Emotion, reason, and the Human Brain* (Avon Books, New York).

6. John C. Eccles (1989). *Evolution of the Brain: Creation of the Self* (Routledge, London).

7. Rodolfo Llinas and Patricia S. Churchland (1996). *The Mind-Brain Continuum: Sensory Processes* (The MIT Press, Cambridge MA).

8. Alwyn Scott (1995). *Stairway to the Mind* (Springer-Verlag, new York).

9. Neil A. Stillings, Steven E. Weisler, Christopher H. Chase, Mark H. Feinstein, Jay L. Garfield, and Edwina L. Rissland (1995). *Cognitive Science: An Introduction* (The MIT Press, Cambridge, MA).

10. J. R. Searle (1980). "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3: 417-457.

11. George W. Smith (1991). *Computers and Human Language* (Oxford University Press).

12. N. Chomsky (1957). *Syntactic Structures* (The Hague: Mouton).

13. Buchheit, Paul (1996). 'A General-Purpose Natural Language Interface to an Operating System.' *The International Journal of Intelligent Systems*, Volume 11, Number 7, July 1996.

14. Buchheit, Paul (1995). 'A Conversational Interface for File Management in an Operating System.' *Proceedings of the 1995 Midwest Artificial Intelligence and Cognitive Science Society Conference*, Carbondale, IL., pp. 43-47.

15. Buchheit, Paul (1993). 'INFANT: A Modular Approach to Natural Language Processing.' *Proceedings of the 21st ACM Computer Science Conference*, Indiana Convention Center, Indianapolis, IN, February 16-18, 1993, pp. 410-417.

16. Buchheit, Paul (1991). *INFANT: A Connectionist-Like Knowledge Base and Natural Language Processing System*. Ph.D. Thesis, University of Illinois at Chicago.