Memory, Speak: The Computational Theory of Mind, Declarative Memory, and the Problem of Relevance

How can the role of long-term declarative memory in a computational theory of mind make tractable the computational problem of relevance: given the cognitive limitations of human mind and the massive, disparate set of information it warehouses, how does a human cognizer feasibly circumscribe the subset of information relevant for belief fixation, belief insertion, and belief revision? The content of this paper will pave a way towards answering this question by providing an answer to a less ambitious question: why does the problem of relevance ever arise for the computational theory of mind (henceforth, `CTM`)?

My answer to this question will depend on drawing attention to an unsuspected yet highly consequential assumption, the `naïve theory of memory`, that underlies the core theoretical commitments of the CTM. To facilitate the uncovering of this assumption and its purview, consider for a moment the core theoretical commitments captured under the following very rough description of the CTM: thoughts characterized as mental states are complex monadic properties – the instantiation, manipulation, and storage of which the cognizer bears to mental representations have a rich compositional syntax and semantics. A straightforward corollary of this characterization is that the structural configuration of these mental representations which serve as the cognitive grist for our reasoning and memory processes in many important ways mirrors the structural configuration of sentences of a natural language. This consequence is so central to theorizing about cognition under the CTM, it has acquired its own moniker: `the mental sentence theory`.

On the mental sentence theory, it is assumed that mental sentences are the basic units of information processing. The theorist then finds herself cornered by the following constraint: the varied and disparate mental processes of our cognitive architecture communicate and operate uniformly over inputs that have a sentential structure. It is precisely this influential constraint of

Harjeet S. Parmar   harjeetp@buffalo.edu

the mental sentence theory that has led to the collapse of a much needed distinction between the structural configuration of the basic units of processing of a theory of thought and the structural configuration of the basic units of processing of a theory of long-term declarative memory. By arguing that the mental processes that support the proper functioning of certain of our reasoning processes, do so by manipulating mental sentences, has the appeal not only of intelligibility, but also credibility.[1] However, adequately explaining how sentences are suitably stored in long-term memory in a way that characterizes their structure requires more than claiming that sentences are stored in an inert medium that itself has the requisite sentential structure. Against this naïve view of memory, the theorist must say something about the structure and operations of that medium, and how that structure and its operations determine or are related to the structure of the sentences which are a necessary condition for thought.

In this paper, I will trace the development of what I am calling `the naïve view of memory`, and how it leads to the computational problem of relevance. I will subsequently gesture at some ways that having a more sophisticated structural and organizational account of long-term memory may provide an avenue of research for ultimately solving the computational problem of relevance.

---

[1] Consider that since Frege's proof system, and Tarski's semantics for a language L, it has been customary to view the chains of inference, defined in terms of valid rules of the system, and interpretative clauses, defined in terms of truth and satisfaction, to range over the notion of `sentences of a language L`. Furthermore, this notion applied to natural language semantics provides the necessary codification for the determination of truth-conditions of sentences.