

Thoughts from Chris Reynolds for

AIT Forum on the Histories of the Internet - 9th January 2024

“Artificial Intelligence research has involved chasing one overhyped paradigm after another, with intervening winters, and I am worried that the current large language model paradigm, while very powerful, is a black box approach which does not properly reflect human intelligence or how the brain evolved, and I agree with most of what John Handby has said.

In my first job, in 1961, I was employed as a 100% human chatbox. It involved reading research and development correspondence in an international company, setting up an indexed archive, and then using this archive to provide management reports and answer questions. This gives me an interesting insight of what the latest chatboxes are doing. Following the ideas in Bush's 1946 prediction "As We may Think" I decided to move to the data processing department of a nearby major oil marketing company, and later I was headhunted to work on the future requirements for the next generation of large commercial computers which would have terminals and hopefully integrated management information systems. This research resulted in a proposal for an experimental human-friendly electronic clerk using an interactive language called CODIL. The main requirement was that the system should be self-documenting and 100% transparent and could work symbiotically with human staff. However, the approach ran into difficulties because "everyone at the time assumed that computers must be black boxes." In particular the Turing Test involves a black box - so it is not surprising that systems designed to pass the Turing test end up as black boxes!!!

I have recently been re-examining the CODIL project archive in the light of modern AI research, and the result suggests that the CODIL approach unconsciously involved (perhaps because I am neurodiverse) reverse engineering the way the human short term memory works. A key feature was that the language involved describing the task in term of sets, where each set was a node in an associatively addressed network. This is radically different to normal programming where the task is split the task into rules and data, which are mapped onto a digitally addressed array of numbers. The project ended after an experimental educational version had got very favourable reviews simply because mu head of department decided that the approach did not agree with the AI paradigm which was currently favoured during the 1980s. This was a time when AI network models were out of fashion.

The underlying story is that the project archives show that there is a possibility that modern AI is going up a blind alley by neglecting brain modelling. The archive also contains much information relating to the lack of adequate support for creative blue-sky research in a rapidly developing field where there is a ferocious rat race to get funding.”