Evolving the linguistic mind [1]

Keith Frankish

ABSTRACT. It is sometimes suggested that we can think "in" natural language. According to this "cognitive" conception of language, we have a linguistic mind, or level of mentality, which operates by manipulating representations of natural language sentences. This paper outlines two evolutionary questions that the cognitive conception must address and looks at some versions of it to see which provides the best answers to them. The most plausible version, I argue, is the view that the linguistic mind is a virtual system (a "supermind"), which arose when early humans learned to engage in private speech and to regulate it using metacognitive skills originally developed for use in public argumentation.

1. Introduction

It is sometimes suggested that we can think "in" natural language –that natural language not only serves as an input-output system for cognition but is also centrally involved in cognition itself. According to this view, sometimes called the *cognitive conception* of language, we have a linguistic mind, or level of mentality, which operates by accessing and manipulating representations of natural language sentences (e.g., Bickerton, 1990, 1995; Carruthers, 1996, 1998b; Dennett, 1991; Harman, 1973). This view has attractions; there are theoretical reasons for endorsing it and introspection supports it, too. However, I shall not be defending the cognitive conception here. Instead, I shall be asking how a linguistic mind might have evolved. I begin by outlining two evolutionary questions that a satisfactory version of the cognitive conception must answer. I then look at some versions of the doctrine and ask which provides the best answers to these questions. This will give us a reason for preferring that version, should we decide to endorse the cognitive conception in the first place.

2. An evolutionary perspective

Many questions arise once we adopt an evolutionary perspective. For the present, however, I shall concentrate on just two.

How? How did natural language become involved in central cognition? It is widely accepted that the language faculty is a modularized, peripheral system, which is relatively encapsulated from the rest of cognition and which originally evolved for

---

[1] This is the author's version of an article published in *Linguistic and Philosophical Investigations* 9: 206-14. (2010). It may differ in minor ways from the published version. The article itself is a revised version of a paper delivered at the 3rd International Evolution of Language Conference (Paris, 3–6 April, 2000) and published in the conference *Proceedings*, edited by J-L. Dessalles and L. Ghadakpour (Ecole Nationale Supérieure des Télécommunications, 2000, pp.104-8).

communicative purposes. How did such a system come to play a role in *central* cognition (that is, in flexible, intelligent, nonencapsulated, conscious thought)?

When? When did language-based cognition evolve? Most writers agree that fully fledged grammatical language evolved some time within the last quarter of a million years, and some place its emergence as late as 50,000 years ago. This does not leave much time for the subsequent development of language-based cognitive mechanisms – not, at any rate, if this would have involved substantial alterations to neural anatomy.

Any satisfactory version of the cognitive conception must address these questions, and any version that conspicuously lacks the resources to do so can be ruled out in advance. So, for example, I think we can rule out the hypothesis that a separate language-based general-purpose reasoning system evolved subsequently to the emergence of language. There would simply have been no time for such a system to develop. There remain, however, a number of viable candidates.

3. Language and the off-line mind

According to Derek Bickerton, language and central cognition co-evolved (Bickerton, 1995). Bickerton argues that structured non-demonstrative thought ("off-line thought" as he calls it) requires a system of schematic representations and a set of combinatorial principles defined over those representations. Language, too, requires a set of schematic representations and a combinatorial syntax, and parsimony suggests that the same neural resources play both roles.

This suggestion has attractions. If Bickerton is right, then there is no problem of how language got involved in central cognition; it *is* central cognition – at least, in so far as central cognition involves off-line thinking. The "when" question is also dispelled. There was no need for further adaptation after the development of language; the development of fully grammatical language was also the development of structured off-line thought.

The proposal has some serious drawbacks, however. I shall mention two. First, it is unlikely that language possessed cognitive functions from the very start. For it would have been simpler to build a purely communicative language system than to build one that had both communicative and cognitive functions. A communicative language system requires only syntax, phonology, and comprehension systems, together with a lexicon. And while these could be adapted to play a role in cognition, they would not in themselves constitute a cognitive system. Additional subsystems would be required, in particular, some sort of central processor. (Syntax alone might give you structured *thoughts,* but not structure-sensitive *thought-processing*.) But if it would have been easier to construct a purely communicative language system than one that also had cognitive functions, then we should expect the former to emerge before the latter – assuming (as seems plausible) that purely communicative language would itself have carried significant benefits.

Second, the proposal overlooks evidence for modular structure within central cognition. Bickerton tends to view the whole of human central cognition as language-based, and thus to suppose that it is a fairly recent system with little inherited structure. Yet there is mounting evidence that we have lots of innate cognitive competences, realized in functionally distinct, partially encapsulated modules, which have developed gradually over the last million years or so and in many cases predate language (e.g., Barkow, Cosmides, and Tooby, 1992). Such evidence tends to undermine Bickerton's solution to the "how" problem. If the linguistic mind is not the whole of central cognition, then how is it related to the rest of it?

4. Language and the modular mind

Is the cognitive conception compatible with a modularist view of central cognition? Peter Carruthers has argued that it is, building on suggestions by Steven Mithen (Carruthers, 1998a; Mithen, 1996). According to Mithen, the human mind developed in three phases. In phase 1, it consisted of a rudimentary general-purpose problem-solving system. In phase 2, this was supplemented by a number of self-contained domain-specific modules, which were fast but inflexible and did not communicate with each other or with general intelligence. Finally, in phase 3, there was a growth of "cognitive fluidity." The previously isolated central intelligences began to communicate with each other and with general intelligence, either through direct channels or through the mediation of a metarepresentation supermodule. Carruthers argues that this picture naturally supports a version of the cognitive conception. Even while remaining internally isolated, he points out, the central modules would have formed input-output links with the language faculty. Natural language would then have been the obvious vehicle for inter-modular information transfer, once the internal barriers started to come down. In this way, natural language would have come to serve as a cognitive lingua franca.[2]

I have two worries about this proposal. First, it does not amount to a full-blooded vindication of the cognitive conception – not, at least, if we take that doctrine to involve the claim that language can act as a medium of *inference* as well as *thought*. In Carruthers's scenario all the real inferential work is done within modules, using their own internal representational media, and natural language serves merely as a conduit between them. Second, the proposal does not fully answer the "how" question. A neural lingua franca may be a necessary condition for inter-modular co-operation, but it is not a sufficient one. Coherent trains of thought do not just spring into existence spontaneously; a problem has to be identified and the various modular resources deployed intelligently to its solution. Some kind of executive would be needed to marshal the problem-solving resources of the different modules and to co-ordinate their outputs. Moreover, this system would need to process sentences in a way that was sensitive to their semantic properties. But then it starts to look like the sort of hard-wired language-based cognitive processor whose existence we have already ruled out.

---

[2] Carruthers develops these and related ideas at much more length in his 2002, 2006, and 2009.

5. Language and the virtual mind

We have been thinking of the linguistic mind and the processing mechanisms that support it as *part* of the brain. But perhaps this is wrong. Perhaps it is more like a *program* running on the brain, a feature of our mental software rather than our neural hardware. Such a view has been defended by Daniel Dennett (Dennett, 1991). The modern human mind, Dennett claims, is not a biological system at all, but a *virtual machine*, the product of learned behaviours (Dennett calls them "good tricks" or "memes"), which have reprogrammed our biological brains. The behaviours in question, Dennett suggests, are linguistic. We acquire virtual minds by talking to ourselves – producing, rehearsing, and rearranging sentences in overt or silent soliloquy. This stream of private verbalization transforms the activity of the biological brain, causing its parallel, multi-track hardware to simulate the behaviour of a serial, single-track processor, operating upon natural language sentences. Dennett calls this softwired system the *Joycean machine*.[3]

This story is particularly attractive from our current perspective. There is no special problem about how or when such a virtual mind could have evolved. Its development would have involved a process of memetic, cultural evolution, rather than the emergence of new neural structures, and there would have been ample time for it to occur after the development of language.

Still, the story will not do as it stands. The problem lies in the way the Joycean machine is supposed to work. According to Dennett, the key mechanism is one of *self-stimulation*. Inner speech is channelled through a feedback loop from speech production to speech comprehension. Internally generated sentences are then processed by the comprehension system just like externally produced ones, often evoking similar responses. So, for example, questioning yourself may prompt an instinctive verbal reply, just as a question from another person might. This reply will then itself be processed by speech comprehension, like any other heard utterance, giving global neural publicity to the information it carries. In this way, information that was previously held by just one neural subsystem can be made available to them all. Dennett suggests that subsystems routinely compete for control of the vocal system and the self-stimulatory mechanisms it supports. As a result, the Joycean machine comes to act both as a bulletin board, where locally stored information is made globally available, and also as a sort of virtual executive, focusing attention, marshalling resources, and co-ordinating the activities of the different subsystems.

It is likely that inner verbalization does have a self-stimulatory function of this sort (Diaz and Berk, 1992). But this cannot be all there is to the linguistic mind. For one thing, it is doubtful that self-stimulation could generate sustained trains of intelligent thought. It might help to produce some regularity and consistency in one's inner

[3] This is in tribute, of course, to James Joyce's depictions of the stream of consciousness in his novel *Ulysses*.

verbalizations, but it is hard to see how it could give rise to coherent inferential sequences of the kind involved in dedicated problem-solving. In such cases, it seems, our inferential subsystems are not *competing* for vocal control, each shouting out its favoured (and not particularly bright) solution. Rather they are *co-operating*, each subordinating its activity to a global objective. And it is hard to see how they could be induced to do this without executive supervision of some sort. Second, self-stimulations will not have the cognitive role typical of linguistic thought. Consider the sort of cases that lend intuitive support to the idea that we can think in language. I notice that the steering on my car is uneven and say to myself, "The wheel alignment needs checking." Here, it seems, I am not instructing or encouraging myself to think that the wheel alignment needs checking; I am *judging* that it does. And this judgement may have long-term effects, such as getting me to the garage the following day, which a transient self-stimulation would not have.

6. Language and the supermind

Dennett gives us part of the story, then, but important features of the linguistic mind remain unaccounted for. Can we complicate his picture in order accommodate these features? I think so.[4] The trick is to think of linguistic reasoning as, to some extent, under personal control – as something we *do*, rather than something that happens in us. The linguistic mind, I suggest, is indeed a virtual one, developed through the discovery and transmission of good tricks. However, these include, not only inner verbalization, but also various metacognitive and metalinguistic skills. We do not only speak to ourselves, I suggest; we also adopt attitudes towards our inner verbalizations and perform explicit inferential operations upon them. In particular, we adopt some of our private utterances as expressions of premises and goals, and manipulate them so as to construct chains of explicit reasoning, using learned inferential skills. I have suggested elsewhere that these activities constitute a distinct level of mentality, which is intentionally formed and sustained and which constitutively involves natural language. I call this the *supermind*. (The term is intended to capture the idea that the states and processes that constitute this level of mentality supervene on intentional states and processes at a more basic level.)[5]

This proposal retains all the advantages of Dennett's. Like the Joycean machine, the supermind is the product of memetic and cultural evolution, rather than changes in neural anatomy. Indeed, there is a plausible story to tell about how it emerged. The metalinguistic and metacognitive skills needed to develop a supermind – the ability to think about one's thoughts and words and to articulate cogent trains of argument – are just the skills needed for engaging in reasoned argument with one's peers. Such

---

[4] The following sketch draws on my 1998a and 1998b. The ideas in question have been developed at more length in my 2004 and 2009.

[5] I have argued elsewhere that the supermind corresponds to the "System 2" posited by dual-system theorists in cognitive and social psychology (Frankish, 2009).

skills would have carried many benefits in early human society – in securing social influence, resources, and mates – and there would have been strong independent pressure for their development. And once they were in place, supermental abilities would have followed naturally. Humans would have begun to develop linguistic minds as soon as they started to internalize their skills in interpersonal argument, reasoning and debating with themselves. (The development of these metacognitive and metalinguistic abilities might also have involved some neural changes, of course, but these would have been minor in comparison to those involved in constructing a distinct language-based general-purpose reasoning system.[6])

This view can, I believe, resolve the problems facing Dennett's account. Take the question of executive control. For self-conscious agents, equipped with metacongitive skills, problem solving will assume a dual aspect. They will be able to think, not only about the first-order problem of what to do or think next, but also about the metaproblem of how to solve that problem. Suppose they have some general ideas about how to solve this metaproblem. So, they want to evaluate candidate hypotheses as they occur to them, preferring those that harmonize well with premises and goals they have previously endorsed, and rejecting those that conflict with them. And as a subgoal they want to trace out the implications of each hypothesis, searching for data that might confirm it or refute it or indicate how it should be revised. These desires then drive their subsequent attempts to tackle the first-order problem. As various candidate hypotheses occur to them (thrown up, let us suppose, by modules), they set to work evaluating them, in line with their metacognitive goals. They persist in this, refining and complicating their hypotheses, until they reach a solution that satisfies those goals. In this way, a person's metacognitive attitudes can *regulate* their first-order problem-solving activities. In effect, people can act as their own central executives, marshalling and directing their low-level cognitive resources.

This view also offers an explanation of how inner verbalizations can assume a direct cognitive role, rather than a merely self-stimulatory one. We can *decide* that an inner verbalization will have the role of a thought by deciding to adopt it as a premise and use it to guide inference and action. Executing such a decision will involve using one's explicit reasoning skills to make sure that the verbalized proposition has the appropriate inferential role – for example, by taking it as a premise in one's explicitly constructed syllogisms. Language-based reasoning will thus be genuinely computational, though the computations in question will be carried out at an explicit, personal level.

Of the candidates reviewed, then, I suggest that the last is the best placed to provide a full-blooded and neurologically plausible defence of the cognitive conception. If we have a linguistic mind, then it is most likely to take the form of a supermind.

---

[6] Note that I am not suggesting that metacognitive processes are themselves language-based; the claim is that they *shape* linguistic thinking, not that they *constitute* it.

References

Barkow, J.H., Cosmides, L., and Tooby, J. (1992), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.

Bickerton, D. (1990), *Language and Species*. Chicago: University of Chicago Press.

Bickerton, D. (1995), *Language and Human Behavior*. Seattle: University of Washington Press.

Carruthers, P. (1996), *Language, Thought, and Consciousness: An Essay in Philosophical Psychology*. Cambridge: Cambridge University Press.

Carruthers, P. (1998a), "Thinking in Language? Evolution and a Modularist Possibility," in Carruthers, P. and Boucher, J. (eds.), *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press, 94–119.

Carruthers, P. (1998b), "Conscious Thinking: Language or Elimination?" *Mind & Language* 13(4): 457–476.

Carruthers, P. (2002), "The Cognitive Functions of Language," in *Behavioral and Brain Sciences* 25(6): 657–674.

Carruthers, P. (2006), *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Oxford University Press.

Carruthers, P. (2009), "An Architecture for Dual Reasoning," in Evans, J. St. B. T. and Frankish, K. (eds.), *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press, 109–127.

Dennett, D.C. (1991), *Consciousness Explained*. Boston: Little, Brown and Co.

Diaz, R.M. and Berk, L.E. (eds.) (1992), *Private Speech: From Social Interaction to Self-regulation*. Hillsdale, NJ: Lawrence Erlbaum.

Frankish, K. (1998a), "A Matter of Opinion," *Philosophical Psychology* 11(4): 423–442.

Frankish, K. (1998b), "Natural Language and Virtual Belief," in Carruthers, P. and Boucher, J. (eds.), *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press, 248–269.

Frankish, K. (2004), *Mind and Supermind*. Cambridge: Cambridge University Press.

Frankish, K. (2009), "Systems and Levels: Dual-system Theories and the Personal-Subpersonal Distinction," in Evans, J. St. B. T. and Frankish, K. (eds.), *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press, 89–107.

Harman, G. (1973), *Thought*. Princeton, NJ: Princeton University Press.

Mithen, S.J. (1996), *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. London: Thames and Hudson.