Posthumanism and Instrumental Eliminativism

Abstract:

I distinguish transhumanism from posthumanism in both speculative and critical forms. I take the term 'posthuman' to refer to hypothetical descendants of current humans that are no longer human in consequence of some history of technical augmentation, claiming that this might occur via technically induced changes in the structure of human cognition. In the remainder of this article I sketch one such scenario for posthuman technogenesis. It assumes what I call the 'linguistic constitutivity thesis': the claim that language is a cognitive tool necessary for the possession of structured propositional attitudes and their associated concepts. Given linguistic constitutivity, an augmentation that replaced public language with a powerful non-symbolic medium would remove the preconditions for propositional and conceptual thought. Lacking these preconditions, human minds would cease to exist. They would be replaced by posthuman minds with characteristic repertoire of non-propositional attitudes exploiting non-linguaformal media for mental representation.

Keywords:

Transhumanism, posthumanism, propositional attitudes, the extended mind, eliminativism, continuous computation.

1) Posthumanism and Transhumanism

Contemporary transhumanists argue that human and non-human natures ought to be redesigned where such alterations are conducive to sentient beings leading better lives. Transhumanism, then, is an ethical position combining the progressivism of radical Enlightenment thought with advocacy of the transformative potential of powerful new technologies: in particular, Nanotechnology, Biotechnology, Information Technology, and Cognitive Science – the so-called "NBIC" suite.

Transhumanists are particularly obsessed with using NBIC's for cognitive enhancement – the targeted improvement or extension of mental aptitudes such as reasoning or memory–since more intelligent humans may be better placed to develop and oversee other technological improvements in our collective lot.

However, some transhumanists hope (and their bioconservative critics' fear) that converging 'nano, bio or info' tech may not only raise intelligence within the current human range but also raise mentation to unprecedented levels. The philosopher David Chalmers argues that this would be possible were cognitive enhancers to become *extendible.* A method is extendible if significantly improving it yields a correlatively smarter system (Chalmers 2009).[1] An extendible method that results in a very smart human equivalent AI, for example, could be improved by the AI itself to produce a still greater intelligence (AI+), which in turn

[1] Our only current means of producing human-equivalent intelligence is non-extendible. As Chalmers writes: "If we have better sex . . . it does not follow that our babies will be geniuses" (Chalmers 2009).

could design a still superior intelligence (AI++), leading to an explosion in mentation that, so the story goes, would leave humans far behind.[2]

In an article from 1993 computer scientist Virnor Vinge dubbed this point of accelerating recursive improvement as "the technological singularity". For Vinge, a singularity would be an unprecedented 'transcendental event', transforming the world in ways that pre-singularity flat footers cannot envisage. Current humans might no more understand a post-singularity dispensation than goldfishes could understand justice or public transportation.

Whether or not a runaway intelligence explosion of this kind is possible or likely, Vinge's idea of a post-singularity intelligence is a powerful inducement to philosophical speculation. For it requires us to consider the possibility of a *posthuman* mode of being significantly alien or 'other' to the human. Whereas transhumans are humans who (by current standards) are exceptionally gifted, a posthuman is a being whose augmentation history renders it inhuman.

Elsewhere I have proposed that posthumanism should be viewed as a metaphysical position logically independent of any transhumanist ethic. I have referred to this as 'speculative posthumanism' to distinguish it from a position sometimes referred to as 'critical posthumanism' in the theoretical humanities.

Critical posthumanists like Katherine Hayles and Donna Haraway claim that current technoscientific change has 'deconstructed' the transcendental centrality of the human subject (Hayles 1999). Speculative posthumanism likewise rejects the claim that the conditions of possibility of cognition or value are constituted by human subjectivity or cognition (transcendental humanism). However, it assumes that a human-posthuman gulf could be as profound as the cognitive and social disconnect between humans and non-human primates. So speculative transhumanism rejects transcendental humanism but retains anthropological humanism: the assumption that humans are distinguished from non-humans by human-specific traits or sorting mechanisms.[3]

The key thesis of speculative posthumanism can be formalized thus:

> *There could be (wide) 'descendants' of current humans that are no longer human in consequence of some history of technical alteration* **(SP).**

This schema captures the idea that posthumans would not be like superior aliens, gods or angels. They would have 'wide human lineage'. Their genesis would be a *technogenesis*.

Linguistic Constitutivity

[2] The extendible method in question might result from some targeted biological alteration in human beings, from the use of "human/computer interfaces" of the kind anticipated by Kurzweil, an emergent property of large information systems, or from convergences in robotics, synthetic life and artificial neural networks.

[3]One can be a transhumanist while denying that a posthuman disconnection from the human is possible or likely. Similarly, speculative posthumanism is consistent with the rejection of transhumanism. One could hold that a posthuman disconnect is a significant possibility but not a desirable one. Although some hold that the singularity is 'beyond good or evil', one might hold that certain posthumans would be worse off than even the most miserable human; a possibility that could warrant anti-transhumanist policies such as technological relinquishment or pre-emptive species suicide.

In this paper I wish to consider one scenario for posthuman technogenesis. This hinges on two claims:

> 1) Humans are distinguished from non-humans by their capacity for mental states with structured propositional contents such as beliefs and desires (**Propositional Distinctiveness**).
>
> 2) The human capacity for propositional attitudes depends upon continued participation in linguistic cultural practices (**Linguistic Constitutivity**)

According to the second claim, language is a cognitive tool which equips us with distinctively human minds. I refer to 2) as the *linguistic constitutivity* claim.

From PD and LC it follows that were our technologically advanced descendants to produce cognitive tools that superseded natural language their minds would cease to be human and would become posthuman in accordance with the SP schema. I will refer to this narrative, for reasons which will become clear, as the 'instrumental eliminativism scenario'.

To put bones on this technogenesis scenario, it would help to sketch out a philosophical framework for linguistic constitutivity. I have chosen to articulate it from within the Extended Mind Thesis (EMT) – an independently plausible position in cognitive science with a particular ease of integration within an instrumental eliminativist scenario.

Proponents of EMT, like Andy Clark, argue for 'parity' between cognitive processes in the head and functionally equivalent processes beyond the skull.[4] The parity principle implies that mental processes need not occur only in biological nervous systems but in the environments and tools of embodied thinkers. EMT coheres with the linguistic constitutivity claim because EMT theorists hold that language does not merely express pre-formed propositional attitudes but provides a shared suite of inferential and monitoring tools which renders propositional thinking of any grain and sophistication possible.

To better appreciate how language might extend cognition and the implications of eliminating it, we need to specify in sufficiently broad terms the features of language that might plausibly warrant attributing a human-constitutive role to it:

> L1) Linguistic behaviour consists in the production of discrete symbols. Symbols are articulated (or finitely differentiated). If it is not possible to determine which symbol type a given mark belongs to, it can't be a symbol.
>
> L2) Symbols are ideal, repeatable entities. A given symbol can be employed by different users in different times and places for different purposes.
>
> L3) Linguistic symbols exhibit compositional structure. Complex symbols like sentences are composed according to the grammatical rules of the language. Rules governing logical connectives or embedded clauses, for example, allow symbols of

[4] 'Parity Principle. If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is (for that time) part of the cognitive process'.(from Clark and Chalmers (1998) p.XX)

arbitrary complexity to be composed recursively by reapplying them to their own output.

Compositionality plus recursion intuitively furnishes a huge representational gain. Not only does it allow the rule-governed generation of arbitrarily complex grammatical strings, it allows the dissemination of culturally transmissible 'meta-tools' such as analytical methods and models of inference.

Donald Davidson's work in semantics provides a suggestive schema for a *strong linguistic constitutivity thesis*. Davidson argues that the ability to have beliefs (and hence other propositional attitudes such as desires) *requires a grasp of what belief is* since to believe is also to understand 'the possibility of being mistaken'. This in turn requires us to grasp that others might have true or false beliefs about the same topic. Thus no belief can be adopted by someone not already involved in evaluating her own and others' attitudes on common terms (Davidson 1984, p. 170). Public language provides the intersubjective framework in which beliefs can be at once attributed, assessed and their fined-grained differences marked as differences in compositional structure and inferential consequence.[5]

The ideal repeatability of symbols arguably plays a key role in attributing and expressing attitudes. Davidson calls this repeatability 'the autonomy of meaning'. Linguistic symbols are 'autonomous' insofar as they retain a semantic identity across divergent uses. Thus while the sentence 'There is a tiger in the woods' can be used for blunt assertion, it can be embedded in an utterance that attributes the belief that there is a tiger in the woods or used sarcastically or ironically (Ibid, pp. 164-165).[6] So, repeatable and compositionally-structured linguistic symbols provide the meta-tool for evaluating propositional attitudes in the course of interpreting them. Given Davidson's assumption that having attitudes requires evaluating them in inter-subjective coin, we could not qualify as true believers without the reflexivity afforded by complex linguistic structures: in particular the recursive power of language to apply sentence forming operations to sentences.

One might object that even granting the tie between evaluating attitudes and having them, it is as plausible to suppose that belief contents are meta-represented in mentalese or some other neural medium of representation (Bermudez 2003, p. 158).[7]. However, our use of hybrid

[5] "Our manner of attributing attitudes ensures that all the expressive power of language can be used to make such distinctions. One can believe that Scott is not the author of *Waverley* while not doubting that Scott is Scott; one can want to be the discoverer of a creature with a heart without wanting to be the discoverer of a creature with a kidney. One can intend to bite into the apple in the hand without intending to bite into the only apple with a worm in it; and so forth. The intensionality we make so much of in the attribution of thoughts is very hard to make much of when speech is not present. The dog, we say, knows that its master is home. But does it know that Mr Smith (who is his master), or that the president of the bank (who is that same master), is home? We have no real idea how to settle, or make sense of, these questions" (Ibid. 163).

[6] Davidson's paratactic analysis of direct and indirect discourse actually denies that s in attributive sentences like 'Fred believes that s' is a semantic constituent of the sentence attributing the belief. Davidson prefers treating 'that' as an indexical device. However, this is still consistent with compositionality so the details of his theory of direct and indirect speech need not concern us.

[7] This issue is, in part, an empirical one and has received considerable attention from psychologists concerned with the role of language in children's development of folk-psychological competences such as the ability to infer behaviour from false-belief attributions (Pyers and Senghas 2009) Rules governing sentential embedded clauses allow language users to form sentences about linguistic or propositional objects such as 'Joan believed

mental representations with both linguistic and non-linguistic parts arguably confers benefits that hypothetical 'inner' media such as mentalese or Churchland-style 'neuralese' could not provide. José Bermúdez argues that cognitive hybridity renders structured thoughts accessible for conscious or personal recapitulation and reflection, a feature absent in the case of mentalese or neuralese (Bermúdez 2003). If so, language *appears* in a unique position to supply what Clark refers to as 'second-order cognitive dynamics' (Clark 1996, 177): a suit of self-monitoring capacities which furnish the 'distinctly human capacity' to think reflexively about our thoughts.

This opens up weak and strong routes to linguistic constitutivity. Firstly, as emphasized by Bermúdez, it allows us to represent complex concepts via the articulate structures of public language. This would suffice for a considerable portion of human conceptual-space to be of a hybrid nature.[8]
The Davidsonian position that the attitudes are, as such, linguistically constituted is stronger. Given weak constitutivity alone, there could be languageless beings possessing a propositional attitude psychology but lacking access to the refine conceptual space up by the external representation of conceptual and discursive relationships.

3) The Technogenesis of the Posthuman

It should be emphasized that even the stronger form of linguistic constitutivity does not entail that the attitudes are mere *abstracta* fashioned for their predictive efficacy. Nor does it entail they are posits of some more or less moribund folk psychology. To say that propositional attitudes constitutively depend on linguistic practices is not to impugn their reality.

However, if linguistic hybridity is constitutive of propositional thought (strong constitutivity) or refined conceptual thought (weak constitutivity) *we don't have to theoretically defeat folk psychology to eliminate the mental. We can eliminate the mental by eliminating (the need for) language*.

Could this occur as a consequence of some or other augmentation technology?

Well, for any such technology to count as an augmentation it would have to render language vestigial while providing equivalent cognitive scope. For example, we would expect our hypothetical posthumans to be capable of non-symbolic analogs of the 'second-order cognitive dynamics' Clark attributes to linguistically-mediated cognition: e.g. the capacity to use vehicles of content to represent and 'reflect' upon their own contents.

In 'Eliminativism Materialism and the Propositional Attitudes', Paul Churchland discusses an inter-cranial commissure thought on analogy with the corpus-callosum – the thick trunk of nerve fibers linking right and left cerebral hemispheres. Churchland supposes that a high bandwidth connection between brains might allow two or more people to 'coordinate their behaviour with the same intimacy and virtuosity displayed by your own cerebral

---

that Bill is the culprit' or 'Nick said "Bill is falsely accused"'. These rules can be applied recursively - as in 'Nick hoped that Joan's belief that Bill is the culprit was mistaken' (Bermúdez 2002).

[8] Working within a modular approach to cognitive architecture, Peter Carruthers has argued that competence in natural language furnishes the means to integrate the output of domain-specific modules (e.g. those for tracking color or geometric relationships) using abstract representations with no appurtenance to a cognitive or sensory modality (Carruthers 2002).

hemispheres'. Were the commissure to become widely available, Churchland speculates, 'language of any kind might well disappear completely, a victim of the "why crawl when you can fly?" principle. Libraries become filled not with books, but with long recordings of exemplary bouts of neural activity' (Churchland 1981).

However, while this vision of microwaved bliss was a nice rhetorical flourish; there are technical problems integrating it into the vector-coding model that Paul and Patricia Churchland later adopted. If the vectors corresponding to activation states in neural networks are the sole determinants of content identity or similarity, then two networks can only be in states with the same content if their state spaces are mathematically equivalent (Churchland 1998; Garzon 2000). As Jerry Fodor and Ernest Lepore argue, given reasonable assumptions about human neural idiosyncrasy, this proposal has the counter-intuitive consequence that no two individuals could share the same concepts (Fodor and Lepore 1992; Garzon).

Churchland has since argued that the idiosyncrasy problem could be obviated by defining conceptual similarity using a measure of the ordering of Euclidean distances between prototype states independently of their dimensionality (Churchland 1998). However, if the linguistic constitutivity thesis is right, coupled brains could lack the capacity for second-order cognitive dynamics furnished by language unless it were *also* possible to represent conceptual relationships in a public medium. Assuming linguistic constitutivity, this resource is not yet available to bare human brains.

Still, it seems to me that we can propose a cognitive prosthesis in the context of the EMT that would have a similarly devastating effect on libraries. We can envisage the lineaments of a cognitive technology which adds enough functionality to be worth adopting, but which employs non-symbolic vehicles for thinking, communicating, interpreting, etc. in preference to symbolic vehicles such as sentences and names.

To ensure that these are non-symbolic we can specify initially that they should have a non-symbolic syntax. One of the syntactic requirements of a symbol system such as a language is that it is possible to determine unambiguously which disjoint symbol a given token types (finite differentiation). Contrastingly, we can specify that our proposed augmentation should use vehicles which do not have to belong to definite representational types and which are subject to a far wider class of computational processes than symbols.

Now, if this hypothetical technology were to count as an augmentation one would expect it to offer at least the functionality of language and then some. For example, language properties like ideality and compositionality afford a shared workspace in which inferential processes can be represented and thoughts made objects of higher order reflection. Might a *non-symbolic workspace* (NSW) mimic or exceed this representational power?

No such technology exists at present, so the only way in which to begin to evaluate this possibility is by considering how the properties of non-symbolic media might furnish this cognitive potential.

Some of these formal properties have been catalogued in computer scientist Brian MacLennan's work on Continuous Formal System (CFS) or 'simulacra'. A simulacrum is opposed to Discrete Formal System (DFS) like a formal language. In a DFS every token belongs unambiguously to one and only one type. Any real symbol system that approximates

to a DFS is thus finitely differentiated since it is possible to determine for a given mark unambiguously whether it is a token of some type.

Whereas a DFS is made out of symbols, MacLennan calls the basic representational vehicles of CFS's 'images'. Whereas a given token either does or does not belong to a syntactic type, an image's membership of a correlative 'type' is a matter of degree (MacLennan 1995, p. 2). So we can understand images as 'regions of bounded variation' rather than discrete syntactic atoms.

Similarly, whereas interpretability is all or nothing in DFS's (only grammatically well-formed strings can be interpreted) the meaning of an image 'varies continuously' with changes to its form (Ibid., p. 4).

For example, if we use analogs of traditional semantic concepts to describe the domain of interpretation, a given image could be more or less *true* or *false* the closer it is to some prototypical point in a region of syntactic variation. Here we can think of the domain as a unit interval with a maximally undefined point in the middle and the values true and false at either extreme.

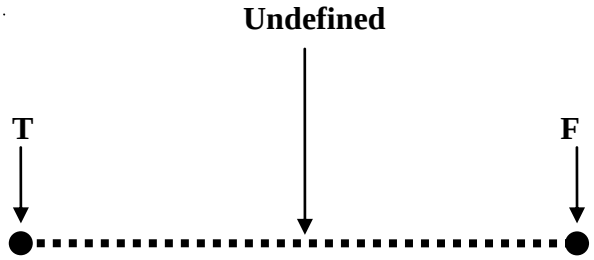

(Ibid., p.6). MacLennan also shows how we can construct a domain of interpretation for a CFS based from the domain of a DFS by mapping each value in the discrete domain (e.g. a referent or a class) onto an ordered pair, the first member of which corresponds to the degree to which a given image belongs to the fuzzy set of a given symbol, then second to the referent from the discrete domain. So a propositional image might have the interpretation:

[0.3, True]

Whereas a nominal image might have the interpretation

[0.76, Achilles]

Here the second component is just the referent of 'Achilles' in a DFS whereas the first corresponds to the interpretability of the symbol in the CFS (Ibid., p.7). Clearly, there are various formal options for describing the semantics of a CFS, depending on how wedded we are to the familiar semantics of discrete systems.

It should be recalled that semantic notions like ‘true’ or ‘false’ are derived from our understanding of formal languages modeled on natural languages. It is an open question how useful such concepts are liable to be for understanding continuous formal systems in the wild. An instantiation of a CFS like our proposed *non-symbolic workspace* might be experienced by its users as a dense plenum with quasi-symbolic nodes, ridges or vortices – more like pictorial than symbolic representation - since translations of an image would yield continuous variations in ‘significance’.

Computation in simulacra would involve the continuous evolution of images rather than the discontinuous transitions between states involved in symbolic computation. Thus a computation could involve a trajectory of a point in state space or an operation that gradually transforms a higher-dimensional surface into another. A discrete symbol can be understood as any symbol in syntax space with well-formedness equal to 1. Thus a continuous computer could accomplish Turing-style operations on discrete symbols. However, for any primitive action of a Turing machine such as ‘write’ or ‘move’, there is a real infinity of actions in an ideally continuous system.

Importantly, MacLennan observes that if a program consists of a trajectory of an image, then the path followed that image is representable as a ‘guiding image’.[9] Thus just as language users can represent rules for valid inference in symbols, so inferences in a continuous computer could be represented by guiding images. Thus a non-symbolic workspace would allow its users to broadcast and evaluate useful guiding images in a manner analogous to the dissemination of cultural ‘good tricks’ through human natural languages.

A similar principle can be applied to the idea of grammatical transformation. Images can be syntactically transformed into terminal images by guiding images (Ibid, 11-13). Thus it ought to be possible to use images and guiding images to form and parse complex images.

It follows that there are formal analogies to be drawn between the structures available within natural languages understood as compositionally structured systems of discrete, repeatable symbols and those available in principle within a continuous system. However, the range of transformations possible in a continuous system is *indefinitely larger*, while its vehicles of content can express intensities and nuances that are arguably only implicit in verbal thought and discourse. MacLennan's theory of simulacra thus allows us to envisage a representational format which is a) *non-symbolic* and b) *has computational resources unavailable to symbolic systems* and c) *capable of representing its own computational procedures and grammatical structures in terms of its own imagistic resources*.

Let us suppose, then, that a radical cognitive enhancement becomes available that implements or closely approximates styles of continuous computation associated with simulacra, which includes not only imagistic representations of non-images (entities in the world outside the system) but images of programs and images which 'parse' other images by representing their grammatical derivation. Presumably, this non-symbolic workspace (NSW) would have to be represented for its users in some form or other to allow them to use it as a transmission

[9] For example, the method of gradient descent (employed in variants of the Delta Rule in the theory of Neural Networks) can be conceptualized as the use of an error surface (the 'program' image) to define the rate and direction of change in some continuous variable corresponding to the computed image (Ibid, 9-10).

medium for their culture. Perhaps it would use a form of virtual reality technology – like the consensual hallucination of cyberspace predicted in William Gibson in his novel *Neuromancer*.

Is there any reason to expect that the NSW could instrumentally eliminate typically human forms of propositional and conceptual thinking? For this to happen, both **propositional distinctiveness** and **linguistic constitutivity** would have to hold. That is, propositional thinking would, firstly, have to a feature or sorting mechanism that distinguishes humans from non-humans and, secondly, be *constituted by* our linguistic practices. Finally, the NSW would have to provide enough functionality to eventually displace public language.

The first two assumptions are independently plausible, if obviously contestable. Moreover, we have seen that there are features of CFS that suggest ways in which non-symbolic media might rival the cognitive prowess of symbolic systems. So, given linguistic constitutivity the successful displacement of public language by a powerful non-symbolic medium would remove the conditions that make propositional attitudes possible. Given that propositional attitudes are human-distinctive, in the way described, human minds would cease to exist. They would be replaced by posthuman minds with characteristic repertoire of non-propositional attitudes exploiting non-linguaformal media for mental representation.

These holders of non-linguaformal attitudes would have ceased to be human through their own technological agency. *And since their inhumanity would have resulted from a technological alteration process, they would satisfy the requirement for posthumanity introduced in the* **SP** *schema*.

Bibliography

Bermúdez, José (2002), L 'Domain Generality and the Relative Pronoun', *Behavioral and Brain Sciences* 25, 676-677.
Bermúdez, José (2003), *Thinking without Words*, Oxford: OUP.
Bostrom, Nick (2005), 'A History of Transhumanist Thought', *Journal of Evolution and Technology*, 14 (1).
Bostrom N, Sandberg A (2006), 'Converging Cognitive Enhancements', Ann. N.Y. Acad. Sci. 1093: 201–227.
Carruthers, Peter (2002), 'The cognitive functions of language', *Behavioral and Brain Sciences* 25, 657-726.
Clark Andy, Chalmers David (1998), 'The Extended Mind', *Analysis* 58(1), pp. 7-19.
Clark, Andy (2003), *Natural Born Cyborgs*, (Oxford OUP).
Clark, Andy (2006), Language, Embodiment and the Cognitive Niche', *Trends in Cognitive Science* 10(8), pp. 370-374)
Clark, Andy, (1993) *Associative Engines*, (MIT Bradford).
Clark, Andy, (2006) 'Material Symbols', Philosophical Psychology Vol. 19, No. 3, June 2006, pp. 291–307
Churchland, Paul (1998), 'Conceptual Similarity Across Sensory and Neural Diversity: The Fodor/LePore Challenge Answered', *Journal of Philosophy* , XCV, No.1, pp. 5-32.
Churchland, Paul (1995), *The Engine of Reason, The Seat of the Soul* (Cambridge Mass.: MIT Press.
Churchland, Paul (1989),'Folk Psychology and the Explanation of Human Behaviour,' *Philosophical Perspectives* 3, pp. 225–241.

Churchland, Paul (1981), 'Eliminative Materialism and the Propositional Attitudes', *Journal of Philosophy* 78(2), 67-90.
Davidson, Donald (1984) 'Thought and Talk', in *Inquiries into Truth and Interpretation* (Oxford, Clarendon Press), pp. 155-170.
Jeffery Elman (1995), 'Language as a Dynamical System,' in R Port and T Van Gelder (eds), *Mind as Motion: Explorations in the Dynamics of Cognition*, (MIT: Cambridge Mass.), pp. 208–222.
Fodor, Jerry (1987) *Psychosemantics* (Cambridge Mass.: MIT 1987).
Fodor, Jerry (1990) *A Theory of Content and Other Essays*, (Cambridge Mass.: MIT 1990).
Fodor, Jerry and Lepore, Ernst (1992), *Holism: A Shoppers Guide* (Oxford, Blackwell).
Garzon, Francisco C (2000), 'State space semantics and conceptual similarity: reply to Churchland', Philosophical Psychology 13(1), pp. 77-95.
Hayles, N Katharine (1999), *How We Became Posthuman: Virtual bodies in Cybernetics, Literature and Informatics* (Chicago: University of Chicago Press).
Horgan, T, Tienson J (1994), 'A Nonclassical Framework for Cognitive Science' , *Synthese*, 101:3, p.305-345.
Horgan, T, Tienson J (1996, *Connectionism and the Philosophy of Psychology*, (Cambridge, MA: MIT Press).
Jones, Richard (2009), 'Brain Interfacing with Kurzweil', http://www.softmachines.org/wordpress/?p=450, Accessed 08.09.2009.
Kurzweil, Ray (2005), *The Singularity is Near* (New York Viking).
MacLennan, B.J. (1995), 'Continuous Formal Systems: A Unifying Model in Language and Cognition' in *Proceedings of the IEEE Workshop on Architectures for Semiotic Modeling and Situation Analysis in Large Complex Systems*, Monterey, CA, pp. 161–172, ftp://cs.utk.edu/pub.maclennan/CFS-UMLC.pdf. Cited 20 December 2009.
MacLennan, B.J (2002), Transcending Turing Computability, *Minds and Machines* **13:** 3–22.
O'Brien G, Opie J (2002), 'Radical connectionism: thinking with (not in) language', Language & *Communication*, 313-329.
Pyers J and Senghas A (2009), 'Language Promotes False-Belief Understanding Evidence From Learners of a New Sign Language', *Psychological Science*, 20(7), pp. 805-812.
Vinge, Vernor (1993), 'The Coming Technological Singularity: How to Survive in the Post-Human Era', http://manybooks.net/titles/vingevother05singularity.html, Accessed 2008.24.04.