

Renewing the link between cognitive archeology and cognitive science

Article (Draft Version)

Thornton, Chris (2012) Renewing the link between cognitive archeology and cognitive science. *Journal of Archaeological Science*, 39 (7). pp. 2036-2041. ISSN 0305-4403

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/41778/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Renewing the Link between Cognitive Archeology and Cognitive Science

Chris Thornton
Informatics
University of Sussex
Brighton
BN1 9QH
UK
c.thornton@sussex.ac.uk

October 31, 2012

Abstract

In cognitive archeology, theories of cognition are used to guide interpretation of archeological evidence. This process provides useful feedback on the theories themselves. The attempt to accommodate archeological data helps shape ideas about how human cognition has evolved and thus—by extension—how the modern form functions. But the implications that archeology has for cognitive science particularly relate to traditional proposals from the field involving modular decomposition, symbolic thought and the mediating role of language. There is a need to make a connection with more recent approaches, which more strongly emphasize information, probabilistic reasoning and exploitation of embodiment. Proposals from cognitive archeology, in which evolution of cognition is seen to involve a transition to symbolic thought need to be realigned with theories from cognitive science that no longer give symbolic reasoning a central role. The present paper develops an informational approach, in which the transition is understood to involve cumulative development of information-rich generalizations.

1 Introduction

Our commitment to the Darwinian theory of evolution means we normally hope to understand the evolution of a species in terms of well-evidenced ‘survival of the fittest’ events. In the case of human evolution, unfortunately, there are significant difficulties in achieving this goal. Part of the problem is the fact that, in this context, it is behavioural changes that are particularly significant. Given these may create little or no trace in the archeological record, the difficulty of inferring relevant transitions is inevitably increased.

Coming to grips with this problem, archeologists have deployed the approach of *cognitive archeology* (e.g. Plotkin, 1982; Donald, 1991; Mellars, 1991; Mellars and Gibson, 1996; Dunbar, 1996; Renfrew and Zubrow, 1993; Mithen, 1996a; Renfrew, 2007). In this approach, theories of cognitive functionality come to guide interpretation of physical evidence. The adoption of a lithic technology in a particular context, for example, might be explained by demonstrating that the cognitive skills involved are enabled by activities pursued in an immediately preceding context. Application of cognitive theory becomes a way of constraining and shaping the interpretation placed on the emergence of particular tool-use.

The approach has proved useful as a means of explaining curiosities in the evolutionary record, e.g., the fact that ‘the most dramatic developments in human cognition seem to have occurred without any concomitant increase in brain size.’ (Bickerton, 1996, p. 44). It has also been particularly effective for explaining the interval between emergence of anatomically modern humans around 200,000 years ago, and the emergence, somewhat later, of significant levels of characteristically modern human behaviour. The earliest known artistic artefact—the incised slab of shale from the Blombos Cave—is dated to more than 70,000 years ago (Henshilwood et al., 2002). But the torrent of art, technology, ritual and symbolism that is deemed the distinctive signature of *Homo sapiens* is seen to develop momentum somewhat later, with the change being particularly dramatic at the Middle/Upper Paleolithic transition in western Europe (Henshilwood and Marean, 2003). The question is then what explains the delay? Renfrew poses it thus. ‘If the genetic basis of the new species is different from that of earlier hominids, and of decisive significance, why is that new inherent genetic capacity not more rapidly visible in its effects, in what is seen in the archeological record?’ (Renfrew, 2008, p. 84-85). Renfrew terms this the ‘sapient paradox’.

2 Approaches to the sapient paradox

One way to deal with the paradox is to posit the occurrence of a mutation in human DNA, that had the effect of establishing a more sophisticated cognitive engagement with the world, but without any accompanying change in anatomy (e.g. Klein, 1999, 2001). In proposals of this type, language may be seen to play an important, mediating role (e.g. Bickerton, 1996). But such schemes remain controversial, given the relatively sparse evidence for symbol-oriented behaviour prior to 50,000 years ago and the lack of evidence for any ‘highly advantageous neurological change’ after that date (Henshilwood and Marean, 2003, p. 630). Cognitive archeology is able to produce more graduated accounts, however. Rather than assuming the occurrence of some critical event which had the effect of ‘switching on’ modern cognition, structured processes of development can be envisaged that had the same effect over a longer period.

Taking the transition to modern cognition to entail the adoption of an increasingly *symbolizing* mode of thought (Renfrew, 2007), such accounts often explain the ‘hold-up’ in terms of sequential constraints. Domain integration can

be the critical concept here. If development of symbolic thought facilitates (or is facilitated by) integration of domain-specific competences, it is clear the latter must precede the former. The necessity of passing through two distinct phases of development can then be used to explain why emergence of modern cognitive sophistication seems to have been delayed (Rozin and Schull, 01988).

Prominent among accounts of this type is Mithen's 'cathedral' model. In this theory (Mithen, 1993), the domain-specific entities are understood to be specialized intelligences, along the lines of (Gardner, 1993). Seeing these as broadly analogous to the chapels of a cathedral, Mithen equates general intelligence with the cathedral's nave. It is a central area through which the multiple intelligences come to be connected. The ultimate effect is a unified, cognitively fluid system of general intelligence (Mithen, 1996b, p. 72).

Mithen's account particularly invokes Tooby and Cosmides' evolutionary model, in which specialized cognitive modules are analogized with the tools of a Swiss army knife (Tooby and Cosmides, 1989, 1992). It also references Fodor's (1983) 'Modularity of Mind'. This proposal envisages mental architecture to comprise a large number of informationally-encapsulated 'input systems' under the management of an integrative reasoning system. Fodor commits to this medium being an inner, compositional language—the so-called 'language of thought'. Mithen is more agnostic on this point, although stressing the degree to which integration must involve processes of analogy and metaphor (Mithen, 2006).

A key advantage of Mithen's account is its ability to address the sapient paradox. It becomes possible to see why a critical development in the evolution of modern human cognition might have involved no gross change in anatomy. The transition can be seen to have been more a change in 'software' than 'hardware'. Delays in cognitive emergence are then more easily comprehended. Indeed, by developing more structured accounts of the changes involved, it is possible to reach a point where the time allowed seems almost too short.

In Donald's (1991) account, for example, the evolution of modern cognition is understood to involve a progression through four distinct *cultures* of representation. The initial culture is *episodic* representation. The main entity here is a kind of situational snapshot; reliance on it significantly limits possibilities for engaging with temporally-extended and otherwise relational phenomena. Episodic representations are then seen to be superseded by *mimetic* representations. These are language-like generalizations, but mediated by non-linguistic forms of expression such as mime and body language. Exploitation of temporal and relational phenomena becomes a possibility.

In the subsequent culture of *mythic* representation, we see emergence of language itself, with expression in mythical entities and traditions. The final phase in the sequence is then characterized by use of external symbol storage (e.g., written representation) and *theoretic/scientific* culture. Mapping these four cultures onto the archeological record then addresses the sapient paradox in a more fine-grained way. The effect is accounted for in terms of the progression through mimetic and mythic stages of representation, within an overarching journey from episodic to theoretic/scientific culture.

It is an important advantage of these cognitively-informed theories that they make it easier to understand why there may have been a delay between appearance of modern human anatomy and modern cognitive sophistication. But they are not without their problems. Allowing cognitive evolution to proceed in a way that is largely disconnected from anatomical change deals with the sapient paradox; but it also tends to eliminate constraints on absolute timing. The question arises of how long we should expect such progressions to take. Cognitive science's reliance on computer simulation means it is not well equipped to give an answer (Boden, 2006). Indeed, the time-scale of most simulation work allows that Donald's four-stage progression might be accomplished in a relatively modest number of generations. Mithen's might even be completed in a single lifetime. A question mark remains hanging over the issue of timing, therefore.

From the present point of view, the more pressing problem with these accounts relates to their terms of reference. Evolution of modern cognition is seen to be a process through which domain-specific functionalities come to be integrated through the operations of a centralized system. This is understood to be either dependent on, or somehow constituent of processes of symbolic reasoning. Development of this style of reasoning is generally assumed to interact closely with evolution of language. But the connection is difficult to discern, partly because language seems somewhat overpowered with regard to its initial application (Dunbar, 1996), and partly because it is extremely hard to disentangle cause and effect (Hauser et al., 2002).

As McBrearty and Brooks note, 'Abstract and symbolic behaviors imply language, but it is doubtful that the point at which they can first be detected coincides with the birth of language' (McBrearty and Brooks, 2000, p. 486). Henshilwood and Marean suggest the latter is likely to have come first, noting the 'capacity for language probably existed in humans well before it was manifested in material culture' (Henshilwood and Marean, 2003, p. 635). But the degree of integration between evolution of symbolic reasoning and evolution of language is not presently of concern. It is the implication of symbolic reasoning being *fundamental* in modern cognition that is more significant.

The difficulty is that this way of conceiving cognition is increasingly out of step with developments in cognitive science. The field has changed significantly in the last two decades. Commitments from earlier years that have recently been revised (and in some cases overturned) include some of those that particularly inform proposals from cognitive archeology. Where cognitive science once emphasized factors of modular decomposition (e.g. Newell and Simon, 1972; Johnson-Laird, 1983; Fodor, 1983; Haugeland, 1985), it now more strongly stresses efficient coding and information use (e.g. Eliasmith, 2007; Griffiths, 2009; Friston, 2010). Where it once emphasized the importance of representational multiplicity and centralized integration (e.g. Anderson, 1983; Gregory, 1984) it now gives as much weight to exploitation of scaffolding and embodiment (e.g. Wheeler, 1994; Beer, 2000). And where it once committed to symbolic reasoning being the medium of high-level integration (e.g. Marr, 1977; Boden, 1977; Winston, 1984) it increasingly recognizes the greater potential (and neural

plausibility) of probabilistic forms (e.g. Doya et al., 2007; Chater and Oaksford, 2009; Clark, 2008).

Conceptions of cognition in which symbolic reasoning takes charge are increasingly questioned (Thelen and Smith, 1993; Ballard, 1991). Indeed, they are often seen to be philosophically flawed (Wheeler, 2005). The essence of the charge is that they are *ideomorphic*, i.e., dependent on conceptual projection. Where a set of ideas are found to deal effectively with the behaviour of a system, there is the temptation to assume they must pick out physical or functional constituents. In practice, there may be no correlation whatsoever. Yet any such scheme will seem to be explanatory purely as a result of the deployment of familiar concepts.

A familiar illustration of the effect is provided by neuroscientific work on language. In the history of this area, we see how conceptions of functionality came to shape ideas about mechanisms of implementation. Dominating the process is the idea of language as a distinct cognitive functionality. Investigations naturally came to focus on the ‘faculty’ responsible. With the actions of this faculty conceptualized to connect language comprehension with speech production, it seemed particular areas of the brain must be involved. Investigations proceeding accordingly then generated evidence which seemed to place the faculties in distinct areas of the brain: comprehension came to be located in Broca’s area, speech production in Wernicke’s area. Ongoing work, however, has revealed a far less simplified arrangement. In fact, these areas have come to be recognized as ‘convenient fictions’ (Donald, 1991, p. 55). In Pinker’s view ‘no one really knows what either Broca’s area or Wernicke’s area is for.’ (Pinker, 1994, p. 311).

Much the same story can be told in the case of cognitive science’s involvement with symbolic reasoning. This way of conceptualizing cognition has proved effective over more than two millenia. It naturally came to shape ideas of implementation. The human cognitive system came to be viewed as a machine for executing symbolic reasoning (cf. Newell and Simon, 1976; Boden, 2006). But with no neurological evidence to give support, the account lacked a firm foundation. In the event, the problems with it came rapidly to light. Research in artificial intelligence demonstrated convincingly that symbolic reasoning machines cannot replicate the power and fluidity of human cognition (Clark, 1997; Beer, 2000; Wheeler, 2005). Most researchers in cognitive science have then moved on to other proposals. Almost all now emphatically reject the possibility of there being any ‘executive center’ for performance of symbolic reasoning in the brain (e.g. Clark, 1997, p. xiii). As Donald observed two decades ago, ‘there do not appear to be any good neuroanatomical candidates for an unconstrained central processor’ (Donald, 1991, p. 54).

Other ways of understanding this outcome can also be framed. The difficulty of obtaining cognitive functionality from mechanised symbolic reasoning can be seen to arise from the ‘frame problem’, for example. This particularly highlights the effects of increasing complexity. As reasoning processes becomes more complex, interactions between separate thread of reasoning build up to the point where inferential functionality breaks down completely. Symbolic rea-

soning machines can reproduce cognitive functionality in simple ‘toy’ situations. But the fluidity of functionality we associate with human cognition remains out of reach.

A close connection also exists between ideomorphic and homuncular explanation. The effectiveness of conceptualizing system behaviour in terms of the actions of an intelligent homunculus become the basis for seeing the homunculus as a real constituent of the system. But if the homunculus is imaginary (as it invariably is) such explanations ultimately collapse. A third way of understanding the failure of accounts invoking symbolic reasoning is to draw attention to the degree to which they are implicitly homuncular. Donald particularly emphasizes this in connection with Fodor’s ‘Modularity of Mind’. Fodor envisages a centralized reasoning system taking steps to integrate information emerging from specialized input systems. But as Donald warns, ‘there is a not-so-covert homunculus lurking in all such models’ (Donald, 1991, p. 364).

The difficulty confronting the relationship between cognitive science and cognitive archeology then comes into focus. There is divergence on the assumed substrate of human cognition. Taking symbolic reasoning to be the means of integrating domain-specific intelligences entails that this process is functionally critical. But for cognitive scientists, that assumption has been the path to manifestly dysfunctional models of mechanism. For these researchers, symbolic reasoning *cannot* be a functionally critical constituent of human cognition.

Is there any way to bridge this gulf? The present paper argues that it can be done by revising definitions a little. In particular, we need to adopt a broader interpretation for the process that cognitive archeology deems to mediate the transition to a ‘symbolic style of thought’. While it seems this must involve acquiring the ability to *execute* processes of symbolic reasoning, that is not the case. In fact, there is an interpretation that eliminates the assumption altogether. The point can be illustrated using an example.

A Tale of Two Forecasters

Imagine we have two individuals F1 and F2 that compete to produce the best weather forecasts for a certain country. F1 relies on perceptual cues obtained from direct observation of the environment. These forecasts exploit simple associations between perceptual patterns and meteorological outcomes. A grey cloud cresting the horizon directly upwind is taken to be an indication that it is likely to rain in the next 15 minutes. Obtaining this perceptual cue, F1 forecasts a strong probability of rain in the immediate vicinity within 15 minutes. Other forecasts are generated in much the same way.

Wanting to achieve a greater level of precision than F1, F2 decides to deploy an approach based on symbolic reasoning. Borrowing the standard symbols for representing centres of high and low pressure, F2 develops a way of classifying warm and cold fronts, such that the location of any front can be predicted to reasonable accuracy up to 24 hours ahead of time. Using this system, F2 is able to produce highly accurate forecasts covering the entire country purely through inferential reasoning.

Not to be outdone, F1 decides that a change of strategy is in order. Understanding that F2's remarkable forecasts are obtained through symbolic reasoning, F1 takes steps to adopt the same strategy. Symbols are introduced for purposes of representing different cloud colourings. F1 establishes one symbol for threatening grey clouds, another for fluffy white ones, and so on. Further symbols are introduced to represent different locations and types of outcome: rain, sun, fog etc. With all of this in place, F1 is able to produce forecasts purely through symbolic reasoning, just like F2. Unfortunately, F1's forecasts turn out to be no better than they were before.

The moral of the story is probably self-evident. Adopting the strategy of symbolic reasoning does not *itself* guarantee any improvement in efficacy. What makes a difference is the precision and power of the generalizations that are brought into play. F2's forecasts are better than F1's due to deploying generalizations that are both more predictive and more precisely defined. F1's adoption of symbolic reasoning makes no difference because there is no change in the generalizations referenced. Examining the reasons for F2's supremacy over F1, we then arrive at a different understanding of what the transition to symbolic thought may involve. We do not have to assume it involves increasing utilization of symbol *processing*. We can see it has involving use of increasingly predictive, increasingly well-defined and increasingly broad generalizations.

In a sense, this is just another way of saying what is meant by 'more symbolic', of course. Generalizations with these characteristics are at a higher level of abstraction, and thus closer to the ideal of formal symbolism, in which signs are fully divorced from content (Casson, 1981). But it does get around a major difficulty with the idea of symbol usage. As Donald notes (1991), if we assume there is *literal* symbol usage in human cognition, we need to explain how referential meaning is established. It cannot be done on the basis of more symbolic processing, on pain of infinite regress. Allowing symbol usage to be identified with deployment of predictively powerful generalizations avoids the problem.

This way of conceptualizing the transition to a symbolic style of thought has a close relation with Bickerton's proposal for a transition from 'on-line' to 'off-line' thinking (Bickerton, 1996), and also to Deacon's (1997) proposal for a progression from indexical to symbolic representation. But the proposed interpretation has the advantage of getting away from any implication of the denouement being emergence of explicit symbolic reasoning.

3 Information-rich generalization

The framework of information theory (Shannon, 1948; Wiener, 1948; Shannon and Weaver, 1949) can be used to give a more formal statement of the proposal. The key principle of this framework is that the information content of an event is the amount of uncertainty it resolves. Given a way of measuring uncertainty (e.g., the entropy formula), we can then *quantify* the amount of information contained in a message, signal or event. The approach can also be used to explain the advantage that is acquired from use of more abstract generalizations,

however. Generalizations which predict a broader range of phenomena with a greater level of precision necessarily resolve more uncertainty. They can thus be seen to *generate* more information. This is one part of the benefit obtained.

There is a second aspect. More abstract generalizations have fewer and more precisely defined criteria. If we take satisfaction of these to be the events that convey the relevant information, it is clear that the fewer there are, the greater must be their average information value. More abstract generalizations thus benefit from making broader and more precise predictions, but also from having more specific criteria. The transition to a more symbolic style of thought can then be seen to be the adoption of increasingly *information-rich* generalizations, where these are understood to have relatively greater breadth, precision and criterial specificity.

Recognizing this triad of informational benefits, we obtain a more formal conception of what the transition to symbolic thought may involve. Rather than seeing it as the innovation of symbolic reasoning, we can see it as the cumulative adoption of increasingly information-rich generalizations. The main advantage of the idea is that it more easily accommodates modern frameworks of cognitive science, particularly where they have an informational basis.

The relation between emergence of symbolic styles of thought and integration of domain-specific intelligences is then called into question, however. The former process is often held to be critically implicated in the latter, although usually with some flexibility over flow of control, i.e., acceptance that it might be integration facilitating the transition, or the other way around.¹ Bringing the informational interpretation to bear, it becomes possible to see how the two processes might be one and the same. As noted, the pursuit of information-rich generalizations can mediate the transition to a symbolic style of thought. But it is no less plausibly the vehicle for forming cross-domain generalizations. It is, after all, a generalization mechanism. Domain integration may thus be progressed by the same underlying process that produces the transition to a more symbolic style of thought.

How well this unification of processes can be reconciled with other accounts for the symbolic transition, such as (Bickerton, 1996) and (Deacon, 1997), needs further investigation. It does seem reasonably consistent with Donald's representation cultures, however. As noted, Donald proposes evolution of modern cognition passes through four distinct cultures of representation: episodic, mimetic, mythic and theoretic/scientific. But these form a sequence in which later cultures must deploy constructs built on previously constructed representations. Since what is represented in later culture must then generalize what is represented in earlier cultures, it is possible that the process is also mediated by generalization. To this degree, Donald's proposal is broadly consistent with the notion of an underlying, information-seeking generalization mechanism.

One question then arising involves the oscillations in cognitive orientation

¹Relevant to this is evidence favouring the significance of cognitive structure in evolution of language. As Dunbar notes, it is 'more likely that language is parasitic on thought, than it has the kind of grammatical structure we give it (the subject-verb-object form) because that is how we naturally think (Dunbar, 1996, p. 105).

that are seen in different phases of human evolution. Cognitive development (in human ancestry) between 100 million and 10 million years ago seems to move in the direction of a more generalized type of intelligence. A period of approximately 10 million years then ensues in which developments seem to move in the opposite direction—towards a more specialized type of intelligence. Finally, in the era of *Homo sapiens*, the trend is reversed again, with developments seeming to move in the direction of more generalized intelligence (Mithen, 1996b, p. 242). How can this alternation be squared with the assumption of a single process?

Again, informational notions can be of use. Recognizing generalization to be a cumulative operation—subsequent constructions must build on and reference prior constructions—we see how the pursuit of informational reward might produce the *effect* of increasing specialization at certain stages. Where we see construction of a generalization that is relatively more specific than the context allows, preference for the construction might be based on it yielding greater reward. Nothing rules this out. Indeed, adoption of more *specific* criteria will generally increase average information among criterial events. Pursuit of informational reward thus has the potential to produce movement in the direction of either increasing specificity or increasing generality. Cognitive developments which seem to prioritize generality in one phase, and specificity in another might then arise from contextual factors rather than from any fundamental ‘polarity shift’ in the process itself.

The informational proposal does have its problems, of course. One drawback is its inability to address the issue of absolute timing. Proposals from cognitive archeology seek to explain the delay in emergence of modern human sophistication in terms of scheduling requirements and constraints, e.g., the necessity of domain-specific learning preceding integration of the knowledge acquired. But there remains the question of absolute duration. There seems no way to pin down the amount of time a particular sequence should take. The informational approach does no better on this point. Taking the underlying mechanism to be information-seeking generalization, the amount of time required to complete any particular operation remains largely unconstrained. On the other hand, by eschewing commitments to symbolic reasoning and modular decomposition, we increase the possibility of bringing other theories to bear in resolving the issue. Further work is required to get more insight on this.

4 Concluding comments

In Wynn’s (1985) view, archeology has ‘the potential to make serious contributions to the study of intelligence’ (p. 32). Mithen also stresses the point, observing that ‘if you wish to know about the mind, do not ask only psychologists and philosophers: make sure you also ask an archeologist’ (Mithen, 1996b, p. 259). The sentiment surely applies in spades to cognitive science, which has long suffered from lack of empirical data to constrain an abundance of theory. But interactions between cognitive archeology and cognitive science face the problem

of increasingly divergent terms of reference. Cognitive archeologists particularly emphasize a gradual transition to symbolic styles of thought, and the mediating role symbolic reasoning can play in cross-connecting domain-specific modules of intelligence. Recent approaches in cognitive science, on the other hand, stress non-symbolic (i.e., probabilistic) forms of reasoning and move *away* from notions of modular decomposition based on abstract notions of mechanism. The problem is how to bring these diverging trends together.

Behind the differing terms of reference there is substantial continuity of conceptualization, fortunately. The transition to symbolic thought need not involve wholesale adoption of explicit, symbolic processing. It may be mediated by the pursuit of informationally rewarding generalizations. On this basis, evolution may have been able to cross what Donald describes as the ‘abyss from non-symbolic to symbolic representation’ (Donald, 1991, p. 368) in relatively easy, information-seeking steps. Ideas relating to processes of module integration (as envisaged by Mithen), or about cultures of representation (as envisaged by Donald), become more easily related to current concepts of probabilistic representation and efficient coding (Friston, 2010). It also becomes possible to see how theories from different domains of investigation might ultimately come be unified under the mathematical principles of information theory.

The proposal does not dispute the salience of a symbolic style of thought. It contends only that symbolism is on a continuum. Instead of there being a rigid distinction between ‘what is symbolic’ and ‘what is not symbolic’, concepts are seen to become relatively more symbolic, as they become more abstract, more precisely defined and more broad in generalization. The proposal also does not intend to downplay the potential significance of language. As Tattersall comments ‘it’s as certain as anything inferential can be that language and the mental abilities directly associated with it loom large indeed behind the capacity to think, on which our species’ success is founded (Tattersall, 1998, p. 227). Nothing in the present proposal contradicts that idea. Taking the transition to symbolic thought to involve information-oriented generalization is in no way inconsistent with evolution of language being intertwined. (This might be in the manner envisaged by Carruthers (2002).) It is more a functional reduction of the mechanisms assumed to mediate the process.

Finally, there is no particular commitment on chronology. The present evidence seems to favour emergence of some symbolic activity after 70,000 years ago, with a significant eruption of art, technology and other forms of symbolically-mediated activity after 40,000 years ago. But the accuracy or otherwise of these dates has no implications for what is presently proposed. The envisaged mediation for the transition is consistent with it having got underway somewhat later, or perhaps considerably earlier. Indeed, the ‘presumptive evidence that boats were being built by *Homo erectus* some 500,000 years ago’ for purposes of making sea voyages around Indonesia (Renfrew, 2007, p. 139) can only suggest (to anyone who has embarked on a similar undertaking) application of fairly sophisticated styles of thought within that era.

References

- Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, Mass: Harvard University Press.
- Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48 (pp. 57-86).
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4 (pp. 91-99).
- Bickerton, D. (1996). *Language and Human Behavior*. London: UCL Press.
- Boden, M. (1977). *Artificial Intelligence and Natural Man (1st edition)*. Harsocks: Harvester Press.
- Boden, M. A. (2006). *Mind as Machine: A History of Cognitive Science (Vol. 1 and 2)*. Oxford: Oxford University Press.
- Carruthers, P. (2002). The cognitive functions of language. *Brain and Behavioral Sciences*, 25 (pp. 652-726).
- Casson, E. (1981). *Language, Culture, and Cognition*. New York: Maxmillan.
- Chater, N. and Oaksford, M. (2009). Precis of bayesian rationality: the probabilistic approach to human reasoning. *Behavioral and Brain Sciences*, 32 (pp. 69-120).
- Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action and Cognitive Extension*. Oxford: Oxford University Press.
- Deacon, T. W. (1997). *The Symbolic Species: The Co-evolution of Language and the Brain*. New York: W.W. Norton & Company, Inc.
- Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, Mass: Harvard University Press.
- Doya, K., Ishii, S., Rao, R. P. N. and Pouget, A. (2007). *The Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press.
- Dunbar, R. (1996). *Grooming, Gossip, and the Evolution of Language*. Cambridge, MA: Harvard University Press.
- Eliasmith, C. (2007). How to build a brain: from function to implementation. *Synthese*, 159 (pp. 373-388).
- Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: The MIT Press.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci*, 11, No. 2 (pp. 127-138).

- Gardner, H. (1993). *Multiple Intelligences: The Theory in Practice*. New York: Basic Books.
- Gregory, R. L. (1984). *Mind in Science*. Harmondsworth, Middlesex: Peregrine Books.
- Griffiths, T. L. (2009). The strengths of - and some of the challenges for bayesian models of cognition (commentary on oaksford and chater, 2009, precis of bayesian rationality). *Behavioral and Brain Sciences*, 32 (pp. 89-90).
- Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. Cambridge, Mass: MIT Press.
- Hauser, M. D., Chomsky, N. and Fitch, W. T. (2002). The faculty of language: what is it, who has it, and how did it evolve? *Science*, 198 (pp. 1569-1579).
- Henshilwood, C. S. and Marean, C. W. (2003). The origin of modern human behavior: critique of the models and their test implications. *Current Anthropology*, 44, No. 5.
- Henshilwood, C. S., D'Errico, F., Yates, R., Jacobs, Z., Tribillo, C., Duller, G. A. T., Mercier, N., Sealy, J. C., Valladas, H., Watts, I. and Wintle, A. G. (2002). Emergence of modern human behaviour: middle stone age engravings from south africa. *Science*, 295 (pp. 1278-80).
- Johnson-Laird, P. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. London: Cambridge University Press.
- Klein, R. G. (1999). *The Human Career: Human Biological and Cultural Origins*. Chicago: University of Chicago Press.
- Klein, R. G. (2001). Fully modern humans. In Feinman and Price (Eds.), *Archeology at the Millenium* (pp. 109-135). New York: Plenum.
- Marr, D. (1977). Artificial intelligence: a personal view. *Artificial Intelligence*, 9 (pp. 37-48).
- McBrearty, S. and Brooks, A. S. (2000). The revolution that wasn't: a new interpretation of the origin of modern human behavior. *Journal of Human Evolution*, 39 (pp. 453-563).
- Mellars, P. and Gibson, K. (1996). *Modelling the Early Human Mind Symbolism, Language and the Neanderthal Mind*. Cambridge UK: McDonald Inst. Archeol. Res.
- Mellars, P. M. (1991). Cognitive changes and the emergence of modern humans in europe. *Cambridge Archeological Journal*, 1, No. 1 (pp. 63-76).
- Mithen, S. (1993). From domain specific to generalized intelligence: a cognitive interpretation of the middle/upper paleolithic transition. In Renfrew and Zubrow (Eds.), *The Ancient Mind: Elements of Cognitive Archeology* (pp. 29-39). Cambridge: Cambridge University Press.

- Mithen, S. (1996a). *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. Thames and Hudson.
- Mithen, S. (1996b). *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science (Phoenix paperback edition)*. Phoenix.
- Mithen, S. (2006). *The Singing Neanderthals: The Origins of Music, Language, Mind and Body*. Phoenix.
- Newell, A. and Simon, H. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Newell, A. and Simon, H. (1976). Computer science as empirical inquiry: symbols and search. *Communications of the ACM*, 19, No. 3 (pp. 113-126).
- Pinker, S. (1994). *The Language Instinct: The new Science of Language and Mind*. The Penguin Press.
- Plotkin, H. C. (1982). Evolutionary epistemology and evolutionary theory. In Plotkin (Ed.), *Learning, Development and Culture: Essays in Evolutionary Epistemology*. Chichester: Wiley.
- Renfrew, C. and Zubrow, E. B. W. (1993). *The Ancient Mind: Elements of Cognitive Archeology* (pp. 29-39). Cambridge: Cambridge University Press.
- Renfrew, C. (2007). *Prehistory: The Making of the Human Mind*. Weidenfeld & Nicolson.
- Renfrew, C. (2008). *Prehistory: The Making of the Human Mind (paperback edition)*. Phoenix.
- Shannon, C. and Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, Illinois: University of Illinois Press.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27 (pp. 379-423 and 623-656).
- Tattersall, I. (1998). *Becoming Human: Evolution and Human Uniqueness*. New York: Harcourt Brace.
- Thelen, E. and Smith, L. (1993). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Mass: MIT Press.
- Tooby, J. and Cosmides, L. (1989). Evolutionary psychology and the generation of culture, part I theoretical considerations. *Ethology and Sociobiology*, 10 (pp. 29-49).
- Tooby, J. and Cosmides, L. (1992). The psychological foundations of culture. In Barkow, Cosmides and Tooby (Eds.), *The Adapted Mind* (pp. 19-136). New York: Oxford University Press.

- Wheeler, M. (1994). *For Whom The Bell Tolls? The Roles of Representation and Computation in the Study of Situated Agents*. School of Cognitive and Computing Sciences, University of Sussex.
- Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. MIT Press.
- Wiener, N. (1948). *Cybernetics*. New York: Wiley.
- Winston, P. (1984). *Artificial Intelligence (2nd edn)*. Reading, Mass: Addison-Wesley.
- Wynn, T. (1985). Piaget, stone tools and the evolution of human intelligence. *World Archaeology*, 17, No. 1 (pp. 32-43).