Charting an evolutionary path to syntax, semantics and consciousness

Donald M. O’Malley*, Department of Biology, Northeastern University, Boston, MA, USA

Correspondence:

Dr. Donald M. O’Malley
Department of Biology
Northeastern University
134 Mugar Hall
Boston, MA, 02115
USA
d.omalley@neu.edu
617-373-2284

running title: Evolution of Syntax and Semantics

key words: symbols, neural networks, evolution, neocortex, representations, sleep, Turing test, grammar

3272 words
no figures

Abstract

It is a challenge to understand how humans and other beings derive meaning from the world, i.e. from the bits of information being routed through their information processing systems. One way to segment distinct forms of information is as *syntax* and *semantics*, nominally the structure and meaning of a message. While more complex forms exist in modern languages, this perspective looks in the opposite direction, towards the deepest evolutionary roots of neuronal information processing. The most ancient forms of world knowledge are hard-wired into nervous systems, wherein neuronal representations might be likened to a basic semantics of objects and other items. In addition, the world can be viewed as possessing a kind of syntactic structure: actions, consequences, space, time and object attributes constitute a universal physics that shapes the survival of individuals and thus evolution. While grammatical structures are not universal across languages, modern human languages encode the syntactic complexity of the world. The path to such capabilities begins with ancient neuronal algorithms whose continued evolution, under representational constraints, led to mammalian neocortex and its associated conscious states.

Revised Manuscript

Introduction

How do indistinguishable blips of electricity, running about in our brains, convey useful information to us? At a slightly higher level, we can ask how it is that we understand the syntax and semantics of messages. This issue has been formulated in computer science as the *symbol grounding problem*: How do we get computers to understand the meaning of arbitrary symbols? In 1990, an influential report by Stevan Harnad proposed a means by which arbitrary symbols could come to have meaning. He argued that symbols should be grounded in more fundamental iconic and categorical representations, thus linking symbolic operations with lower-level (neural-net style) connectionist systems (Harnad, 1990). There are important parallels between computer and neuronal processors, with both showing successively more complex levels of organization: the *symbol grounding problem* applies to the highest levels of human cognitive activity.

Classical work on symbols, linguistics and meaning that emerged from the domains of artificial intelligence and cognitive science (see e.g. Turing, 1964; Chomsky, 1965; Fodor, 1975; Haugeland, 1981; Jackendoff, 1985; Deacon, 1997), is now being complemented by an intense and more biological effort (Dikker and Pylkkanen, 2011; Fitch, 2011; Kim and Sikos, 2011; Isel and Shen, 2011; Wilson et al., 2011; Cappa, 2012; Griffiths et al., 2012; Grainger et al., 2012). But such high-level processes ultimately come down to basal units of information: neuronal action potentials vs. the 1's and 0's of digital computers, elements that further the identity between biological and digital computers. Both codes are tasked with the same core problem, that of representing the items and relationships of the world. Of course, neuronal information systems (brains) exhibit far more complex representational and communicative capabilities than current computers. A more fundamental difference, however, is that brains evolved from extremely simple neuronal information processors. We should therefore take an evolutionary journey deep into primitive information processing systems to explore the neuronal precursors to syntax and semantics. Ultimately, these led to modern biological communications.

The Structure of Messages

Messaging today is dominated by computer and linguistic communications. At their core, these systems rely on syntax and semantics. From a most stripped down viewpoint, syntax and semantics can be described as the structure and meaning of a message (the complexities of modern language notwithstanding). We would like to understand how the precursors to syntax and semantics came to exist in nervous systems. As central nervous systems grew larger and more complex, early computing algorithms of adaptive value would have been conserved and so may reside to some extent in modern neuronal architectures.

Syntax and semantics should also be considered in the context of their underlying content. For example, “the cat bit the dog” has three semantic elements, while the syntax defines who bit whom. Such a message is meaningful only because the brain knows (via prior experience) something about *cat*, *dog* and *bite*. But must a “message” be deliberately sent to be imbued with semantic and/or syntactic attributes? With this question in mind, we begin with quite ancient computations that might be associated with semantic information processing.

How did Semantic Systems Arise?

Among the simplest of vertebrate brains is that of the larval zebrafish, which has innate capabilities to identify, track and consume prey items (Borla et al., 2002). The visual firing pattern, i.e. a small object with a certain movement trajectory, is a structured input that is operated upon by brainstem circuits to produce a precisely targeted and powerful strike on a hapless paramecium (McElligott and O’Malley, 2005; Bianco et al., 2011). But is this a “semantic” result? At its most basic, semantics entails a sign or symbol that represents something other than itself. In this example, the structure of the transduced visual stimulus conveys more than the electrical blips of retinal ganglion cells: it conveys actionable information about something else, about an item in the world. This signal, however, lacks the arbitrary symbolism of the letters “p-r-e-y i-t-e-m”, and would not be considered a semantic or symbolic result by Harnad (1990). This contradiction might (with some reluctance) be considered a matter of “semantics”, but it does raise a number of real issues.

The immediate problem with calling a structured pattern of action potentials a “semantic” result is that all neuronally-processed items detected in the world would be considered semantic results, not just animal communications (as e.g. in Seyfarth et al., 1980, 2010; Pepperberg, 2002, 2009; Rendall et al., 2009; Fischer and Hammerschmidt, 2011). Would such a view produce incongruous results? Deliberate animal signals might have their origins in “motivated signs” (Burling, 2007) where e.g. the non-symbolic act of biting evolves into the baring of teeth plus growling. This requires evolutionary changes in both the senders and receivers of the message, in that evolution is learning a sequence [mouth proximity ==> teeth ==> bite/harm] with predictive value to the receiver (avoid harm) and additional value to the sender (avoid conflict). The baring of teeth constitutes an iconic symbol, whereas the sound of a growl is indexical—correlated to an impending bite (see e.g. Deacon, 1997). At the origin of communications, such signals were not fully arbitrary symbols, nor were they processed by pure symbol-manipulating systems. Instead, general auditory and visual processing systems adaptively processed such items and they, at some point, were enhanced by dedicated temporal and parietal lobe systems for memory and spatial analyses.

Predating the earliest signals or signs, myriad items in the world, as processed by ancestral neuronal circuits, would seem to have an *intrinsic* semantic aspect. Consider the operation of a simple autoassociative network in the brain of an animal or computer. This network receives visual input of a furry animal that could be a cat or dog. Such networks can take a noisy input and via recursive synaptic connections produce an output that is a canonical representation of either a dog or cat, depending on which stored pattern in the network best matches the incoming visual pattern (see e.g. Trappenberg, 2010). Here, the “syntax” (of sorts, see below) of the message was the structure of the pattern of incoming action potentials, which emanated from lower-level visual processing networks. Via the winner-take-all property of autoassociative networks, an output pattern is computed that is purely CAT or purely DOG, in terms of the set of output features (e.g. size, shape, gait). The output of this network can feed into the linguistic network of the brain and elicit e.g. the word CAT. The arbitrary symbol c-a-t conveys, without a doubt, semantic information. But this semantic information did not make a miraculous leap from the world directly into a linguistic processing center. Semantic information must have been conveyed via the structured results of visual information processors (aided in all likelihood by experiences previously stored in the brain).

These results do not preclude assigning different terms to different classes of representation, e.g. *linguistic* semantics (of multiple kinds) vs. an ancient world-driven *intrinsic* semantics. For example, a message from optic tectum to the brainstem motor network might be viewed as a deliberate internal message. Such definitions, however, do not resolve the biggest questions: what is the nature of systems that process arbitrary symbols and how do linguistic systems relate to the processors used by non-linguistic animals to understand the world? This is an opportune point to transition to a second aspect of communication systems, syntax, wherein neuronal antecedents to syntax were likely co-evolving with the antecedents of semantics.

Nature and Origins of Syntax

Syntax is evident in the communications of both humans and computers. It might also be extended to certain structured communications of other animals (Savage-Rumbaugh, 1990; Esser et al., 1997; Mercado et al., 2005; Ouattara et al., 2009; Seyfarth et al., 2010; Berwick et al., 2011; Clay and Zuberbühler, 2011; but see Nowak et al., 2000). In the internal communications of the zebrafish brain, the sequential movements of an appropriately sized item across the optic tectum symbolize the presence of a consumable item. But this signal does not have the structural complexity of even a rudimentary grammar or computer language and since even simpler animal responses are evident, it reasons that syntax appeared only after the instantiation of intrinsic semantic meaning in the neural code (also see Zhang et al., 2010). This might seem strange since syntax is integral to semantics in many information-processing systems, but that is because syntax is required to generate more complex kinds of integrated information, as e.g. language.

Three main elements of syntax are (1) sequential elements, (2) order-dependency and (3) the use of distinct syntactic elements to construct higher-order representations. While the image of a paramecium moving across a zebrafish's retina contains sequential elements and the order dependency needed to initiate predation, it is just one repeated element. Similarly a "language" that consisted entirely of nouns would not meet the criteria of even a rudimentary syntax. While a proto-lexicon along these lines is a plausible step towards human language (Bickerton, 2009), the addition of basic syntax could have resulted in the first human protolanguage. Human language appeared at least within the past 2 or 3 million years, but by most estimates much more recently (Burling et al., 2007; Bickerton, 2009; Dunbar, 2009). From the perspective of neuronal computing, however, the precursors to syntax have far deeper roots.

Syntax comes from The World

Jackendoff (1985) has argued that modern grammar reveals much about higher-level brain operations (thought), but this can be true only if underlying brain operations give rise to the structure of grammar. Antelopes, via neocortical representations, likely recognize a variety of predators and their actions. Consider e.g. the antelope's interpretation of two similar items: a stalking lioness vs. a strolling lioness (often with tail held high, an indicator that she is not actively hunting). Antelopes will largely ignore the strolling lioness, while responding with agitation or flight to one spotted in stalking mode. The antelope's behavior is thus a consequence of the neuronal coupling of action and item, with different couplings producing opposite responses. Animals couple a great many things, such as "my cubs are hidden in the grass" and "I have to stay in the tree because I am surrounded by hyenas". While we cannot directly know what it is like to be a leopard high up in a tree (Nagel, 1974), they know something about gravity and many other things that correspond to nouns, verbs, adjectives, prepositions and relationships: hyena, lion, antelope, run, jump, climb, hot, cold, tasty, rancid, thirsty, up, down, inside, near, far, friendly, hostile, their offspring and their parents (to name a few).

Essential relationships between such sundry items are embedded in the knowledge that many animals have of the world—an ageless universal physics that perhaps gave rise to a vastly younger universal grammar (Chomsky, 1970). Such complexities of the world must reside in the neocortical representations that underlie grammatical operations, such as those explored by Jackendoff and others (see e.g.: Lakoff and Johnson, 1999; Margoliash and Nusbaum, 2009).

Transition to Language

The beauty of language is its creative and limitless ability to use a small variety of syntactic operators to manipulate a vast catalog of semantic items. At present we cannot answer a most fundamental question regarding the immediate precursors to syntax and semantics: did these ancestral neuronal circuits (or genetic duplicates of them) *become* linguistic circuits or did they retain their ancestral character, continuing to serve mainly as a source of *content* for modern language circuits? This limitation stems, at least in part, from an even more fundamental limitation in that we do not know how basic items are represented in neocortex. Possible scenarios range from grandmother cells to sparsely-distributed representations (Dehaene et al., 2005; Quiroga et al., 2008; Bowers, 2009; Plaut and McClelland, 2010). The actual form of such representations might constitute a major constraint on the operation of fully-symbolic linguistic operations.

There is intense debate on the origins of language, so this brief perspective could not do justice to many interesting theories, e.g. whether protolanguage was gestural in nature, an exaptation, a direct descendant of vocal animal communication or perhaps a *de novo* capability emerging from the improved symbolic and recursive abilities of *Homo's* neocortex (Savage-Rumbaugh et al., 2001; Hauser et al., 2002; Christiansen and Kirby, 2003; Rizzolatti and Craighero, 2004; Arbib, 2008; Rilling et al., 2008; Corballis, 2010; Sole, 2010; Fitch, 2011; Enard, 2011; Aboitiz, 2012; Steele et al., 2012). One thing, however, seems certain: the nouns, verbs, prepositions and adjectives of human language are deeply entangled with neuronal representations of the world (objects, actions, attributes, etc.) that animals have understood (colloquially speaking) for hundreds of millions of years (Allen et al., 2010; Yamazaki et al., 2010; Battaglia and Pennartz, 2011). Their conjoining via the universal physics of the world parallels the (non-linguistic) ability of mammals to understand myriad relationships amongst a world of varied items.

To Neocortex and Beyond

The most profound question is one not yet broached: how do neuronal representations, be they visual, auditory or linguistic, come to have meaning to us? The intrinsic semantics derived from the visual image of a paramecium “means” something to a larval zebrafish in some sense of the word: the larva detects and responds. Some might think the larva conscious of its world, but if it is a mindless automaton, then it is not computationally distinct from a thermostat responding to a change in temperature: both engage in detection, decision making and action. But the question arises as to whether linguistic computations are materially different from simpler ones. This is where the rubber of the symbol grounding problem meets the road. If it is just assumed that any device that receives inputs and produces outputs (including calculators) “understands” what it is doing, then there is no symbol grounding problem. But it is only when conscious that humans can generate and understand language and do myriad other things including storing and interrogating day-long memory records (O'Malley, 2011a, b). Several physiological and computational aspects of consciousness are of relevance to the understanding of symbols.

First and foremost, we know not why a collection of neurons should cause physical beings to consciously experience the world. Most would agree that a few cultured, synaptically-connected neurons in a dish are not conscious, but when you add more neurons and organize them into a sophisticated central nervous system and instantiate them in the body of say an ant, now you have a being that explores and interacts with the world. If we suppose the ant is conscious, then some crucial capacity emerged in between the culture dish and the ant. But alas there is NO aspect of the ant's neuronal operations that requires consciousness. Everything the ant CNS does, including teaching (Franklin et al., 2012), can be defined in terms of inputs, outputs, action potentials and internal computing circuits: consciousness need not apply. One could set the bar higher and claim that only mammals with sleep-wake cycles and 5 stages of sleep are the truly conscious beings, especially given our perceived free will, but when we look into the human CNS it's just more axons, synapses and action potentials. At no point is consciousness required to explain any operation of human neocortex.

A second key point is that the sleeping, non-dreaming human brain is a massively powerful computer performing a multitude of operations, including housekeeping stuff like preventing us from rolling off the bed or awakening us to a fire alarm. So it is not the case that as soon as you have a certain number of neurons working together (in an ant, fish or mammal brain), that you are conscious. You might be—or you might not. But can a non-conscious or unconscious being understand the meaning of events in the way that we can? Consider that your son comes into your bedroom and says, "Mommy, I'm sick". This does not seem to mean anything if you remain asleep. While the words might reverberate in working memory circuits long enough to be perceived were you to awaken, if you remain asleep the words are lost forever. Here *intrinsic semantic* processing seems to diverge from *perceived meaning,* because there might be semantic processing of the message in sleeping mommy's brain. Neuroscience may lack good data on this point, but it is possible that the words *mommy*, *I'm* and *sick* are semantically processed during sleep by e.g. causing the firing of neurons associated with the memories/representations of these items. Moreover, the sentence structure might activate circuits that interpret grammar. But absent consciousness, there is no perceived meaning of the sentence.

It seems plausible that consciousness is a property of integrated neocortical activity, given that the functional disconnection of neocortical columns via the fragmenting effects of anesthetics renders humans unconscious (Mashour, 2006). This fits with the integrated information theory of consciousness, which argues that a threshold level of complexity is needed for the emergence of consciousness (Tononi, 2008; Rosanova et al., 2012). There is some means of transforming neuronal (including linguistic) operations into perceived experience, but despite much effort the neuronal machinery and processing steps remain elusive (Posner, 1994; Llinas et al., 1998; Baars and Franklin, 2003, 2007; Dehaene and Changeux, 2004; Roy and Llinas, 2008). Dogs exhibit sleep-wake states (Nishino and Mignot, 1997; Bush et al., 2004) and the similar states of zebrafish (Zhdanova, 2011) suggest an even earlier origin of conscious states, perhaps entailing very different mechanisms. But regardless of phylogenetic origins, consciousness is a real experience and its continued presence in brain-damaged individuals a matter of dire importance (Schiff et al., 2009; Rosanova et al., 2012).

The Symbol Grounding Problem Reprised

Alan Turing (1964) proposed that one could decide if one was speaking to another sentient (i.e. conscious) being by holding a conversation across a barrier. If the other entity could obviously be recognized as a computer (like the telephone agent *Siri*), then it failed the test. But if you could not reliably distinguish its performance from that of your fellow, like-minded humans, then it was granted the status of a conscious being. This seems tangential to the problem, but we have no better test. Koch and Tononi (2011) proposed a new Turing-style test where only “conscious” beings can determine if certain kinds of images are “possible” or not. However, as with Turing’s test, all mammals other than humans fail this test. Many mammals exhibit sleep-wake cycles and multiple sleep stages including REM sleep. Behaviorally, these animals seem “obviously conscious” when awake (Griffin, 2001) yet cannot pass the Turing or Koch tests. Since neither test can sort mindless automatons from apparently conscious beings, such tests are neither necessary nor sufficient to determine consciousness (McGinn, 1999).

In the case of Harnad’s *Symbol Grounding Problem*, where a goal is to enable machines to perform at higher cognitive levels, the work is a clear success. But it does not address the ability of entities to perceptually understand symbols in the way that we do. If you (in 2012) can generally understand what is written here, then you are not a *Google* text-searching algorithm. Whether or not non-biological devices will ever understand the meaning of symbols, in the sense that humans do, boils down to one basic question: will they become conscious? This might be viewed as an untestable hypothesis, but perhaps one day, after humans better intuit the 11 dimensions of string theory, we might be able to explain why consciousness appears when we wake and disappears when we sleep. We might be able to discern with rigor which beings are conscious and which are not. Until then, consciousness should be seen as the core of our existence for at least one practical reason: we can pull the plug on any human body that we know, with judicial certainty, can never regain consciousness.

Reference List:

Aboitiz F (2012) Gestures, vocalizations, and memory in language origins. Front. Evol. Neurosci. 4:2, 15pp.

Allen K, Ibara S, Seymour A, Cordova N, Botvinick M (2010) Abstract structural representations of goal-directed behavior. Psychol. Sci. 21:1518-1524.

Arbib MA (2008) From grasp to language: embodied concepts and the challenge of abstraction. J. Physiol Paris 102:4-20.

Baars BJ, Franklin S (2003) How conscious experience and working memory interact. Trends Cogn. Sci. 7:166-172.

Baars BJ, Franklin S (2007) An architectural model of conscious and unconscious brain functions: Global Workspace Theory and IDA. Neural Networks 20:955-961.

Battaglia FP, Pennartz CM (2011) The construction of semantic memory: grammar-based representations learned from relational episodic information. Front. Comput. Neurosci. 5:36.

Berwick RC, Okanoya K, Beckers GJ, Bolhuis JJ (2011) Songs to syntax: the linguistics of birdsong. Trends Cogn. Sci. 15:113-121.

Bianco IH, Kampff AR, Engert F (2011) Prey capture behavior evoked by simple visual stimuli in larval zebrafish. Front. Syst. Neurosci. 5:101.

Bickerton D (2009) Adam’s Tongue: How humans made language, how language made humans. Hill and Wang Publishers, New York, 286 pp.

Borla MA, Palecek B, Budick SA and O’Malley DM. (2002) Prey capture by larval zebrafish: evidence for fine axial motor control. Brain Behavior & Evolution, 60:207-229.

Bowers JS (2009) On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychol. Rev.* **116**:220-251.

Burling R (2007) The talking ape: how language evolved. Oxford University Press, USA, 298 pp.

Bush WW, Barr CS, Stecker MM, Overall KL, Bernier NM, Darrin EW, Morrison AR (2004) Diagnosis of rapid eye movement sleep disorder with electroencephalography and treatment with tricyclic antidepressants in a dog. J. Am. Anim. Hosp. Assoc. 40:495-500.

Cappa SF (2012) Imaging semantics and syntax. Neuroimage 61:427-431.

Chomsky N (1965) Aspects of the Theory of Syntax. MIT Press, Cambridge MA, 251pp.

Chomsky N (1970) Remarks on Normalization. In: Readings in English transformational grammar. L. Jacobs and P. Rosenbaum (Eds.), Ginn and Company, Waltham MA.

Christiansen MH, Kirby S (2003) Language evolution: consensus and controversies. Trends Cogn. Sci. 7:300-307.

Clay Z, Zuberbühler K (2011) Bonobos extract meaning from call sequences. PLoS One 6:e18786, 10pp.

Corballis MC (2010) Mirror neurons and the evolution of language. Brain Lang. 112:25-35.

Deacon TW (1997) The symbolic species. WW Norton & Company Inc., New York, NY, 527 pp.

Dehaene S, Changeux J-P (2004) Neural mechanisms for access to consciousness. In Gazzaniga MS (ed.) *The Cognitive Neurosciences*, pp. 1145-1157. MIT Press, Cambridge, MA.

Dehaene S, Cohen L, Sigman M, Vinckier F (2005) The neural code for written words: a proposal. Trends Cogn. Sci. 9:335-341.

Dikker S, Pylkkanen L (2011) Before the N400: effects of lexical-semantic violations in visual cortex. Brain Lang. 118:23-28. .

Dunbar R (2009) Why only humans have language. In: R.Botha & C.Knight (eds), The Prehistory of Language. Oxford University Press.

Enard W (2011) FOXP2 and the role of cortico-basal ganglia circuits in speech and language evolution. Curr Opin Neurobiol. 21:415-424.

Esser KH, Condon CJ, Suga N, Kanwal JS (1997) Syntax processing by auditory cortical neurons in the FM-FM area of the mustached bat Pteronotus parnellii. Proc. Natl. Acad. Sci. USA 94:14019-14024.

Fischer J, Hammerschmidt K (2011) Ultrasonic vocalizations in mouse models for speech and socio-cognitive disorders: insights into the evolution of vocal communication. Genes Brain Behav. 10:17-27.

Fitch WT (2011) The evolution of syntax: an exaptationist perspective. Front. Evol. Neurosci. 3:9.

Fodor (1975) The language of thought. Harvard University Press, Cambridge MA.

Franklin EL, Robinson EJ, Marshall JA, Sendova-Franks AB, Franks NR (2012) Do ants need to be old and experienced to teach? J. Exp. Biol. 215:1287-1292.

Grainger J, Dufau S, Montant M, Ziegler JC, Fagot J (2012) Orthographic processing in baboons (*Papio papio*). Science 336:245-248.

Griffin DR (2001) Animal minds: beyond cognition to consciousness. University of Chicago Press, Chicago, IL, 376 pp.

Griffiths JD, Marslen-Wilson WD, Stamatakis EA, Tyler LK (2012) Functional organization of the neural language system: dorsal and ventral pathways are critical for syntax. Cereb. Cortex 2012 Feb 10. [Epub ahead of print].

Harnad S (1990) The symbol grounding problem. Physica D 42:335-346.

Haugeland J (1981) “Semantic Engines: an Introduction to Mind Design” in Mind Design, Cambridge MA: MIT Press, pp. 1-34.

Hauser MD, Chomsky N, Fitch WT (2002) The faculty of language: what is it, who has it, and how did it evolve? Science 298:1569-1579.

Isel F, Shen W (2011) Modulation of semantic integration as a function of syntactic expectations: event-related brain potential evidence. Neuroreport 22:195-199.

Jackendoff RS (1985) Semantics and Cognition. MIT Press, 304 pp.

Kim A, Sikos L (2011) Conflict and surrender during sentence processing: an ERP study of syntax-semantics interaction. Brain Lang. 118:15-22.

Koch C, Tononi G (2011) A test for consciousness. Scientific American 304:44-47.

Lakoff GP, Johnson M (1999) Philosophy in the Flesh. Basic Books, 640 pp.

Llinás R, Ribary U, Contreras D, Pedroarena C (1998) The neuronal basis for consciousness. Philos. Trans. R. Soc. Lond. B 353:1841-1849.

Margoliash D, Nusbaum HC (2009) Language: the perspective from organismal biology. Trends Cogn Sci. 13:505-510.

Mashour GA (2006) Integrating the science of consciousness and anesthesia. Anesth. Analg. 103:975-982.

McElligott MB and O’Malley DM (2005) Prey tracking by larval zebrafish: axial kinematics and visual control. Brain Behav. Evol. 66:177-196.

McGinn C (1999) The mysterious flame: conscious minds in a material world. Basic Books, New York, NY, 242 pp.

Mercado E 3rd, Herman LM, Pack AA (2005) Song copying by humpback whales: themes and variations. Anim. Cogn. 8:93-102.

Nagel T (1974) What is it like to be a bat? Philosophical Review 83:435-450.

Nishino S, Mignot E (1997) Pharmacological aspects of human and canine narcolepsy. Prog Neurobiol. 52:27-78.

Nowak MA, Plotkin JB, Jansen VA (2000) The evolution of syntactic communication. Nature 404:495-498.

O’Malley DM (2011a). Did flash memory, a non-linguistic symbolic stream, play a role in the emergence of language? Soc. Neurosci. Abstracts #401.18.

O’Malley DM (2011b) Neural coding, flash memory stores and symbolic representation. 4th INCF Congress of Neuroinformatics. Boston MA, Sept. 4-6.

Ouattara K, Lemasson A, Zuberbühler K (2009) Campbell's monkeys concatenate vocalizations into context-specific call sequences. Proc. Natl. Acad. Sci. USA 106:22026-22031.

Pepperberg IM (2002) The Alex Studies: Cognitive and Communicative Abilities of Grey Parrot. Harvard University Press, Cambridge MA, 434 pp.

Pepperberg IM (2009) Alex & Me: How a Scientist and a Parrot Discovered a Hidden World of Animal Intelligence--and Formed a Deep Bond in the Process. Harper Collins Publisher. 288 pp.

Petersson KM, Folia V, Hagoort P (2012) What artificial grammar learning reveals about the neurobiology of syntax. Brain Lang. 120:83–95.

Plaut DC, McClelland JL (2010) Locating object knowledge in the brain: comment on Bowers's (2009) attempt to revive the grandmother cell hypothesis. Psychol. Rev. 117:284-288.

Posner MI (1994) Attention: the mechanisms of consciousness. Proc. Natl. Acad. Sci. 91:7398-7403.

Quiroga RQ, Kreiman G, Koch C, Fried I (2008) Sparse but not 'grandmother-cell' coding in the medial temporal lobe. *Trends Cogn. Sci.* **12**:87-91.

Rendall D, Owren MJ, Ryan MJ (2009) What do animal signals mean? Animal Behaviour 78:233-240.

Rilling JK, Glasser MF, Preuss TM, Ma X, Zhao T, Hu X, Behrens TE (2008) The evolution of the arcuate fasciculus revealed with comparative DTI. Nat. Neurosci. 11:426-428.

Rizzolatti G, Craighero L (2004) The mirror-neuron system. Ann. Rev. Neurosci. 27:169–192.

Rosanova M, Gosseries O, Casarotto S, Boly M, Casali AG, Bruno MA, Mariotti M, Boveroux P, Tononi G, Laureys S, Massimini M (2012) Recovery of cortical effective connectivity and recovery of consciousness in vegetative patients. Brain 135:1308-1320.

Roy S, Llinás R (2008) Dynamic geometry, brain function modeling, and consciousness. Prog. Brain Res. 168:133-144.

Savage-Rumbaugh ES (1990) Language acquisition in a nonhuman species: implications for the innateness debate. Dev. Psychobiol. 23:599-620.

Savage-Rumbaugh S, Fields WM, Taglialatela JP (2001) Language, speech, tools and writing: a cultural imperative. J. Consciousness Studies 8:273-292.

Schiff ND, Giacino JT, Fins JJ (2009) Deep brain stimulation, neuroethics, and the minimally conscious state: moving beyond proof of principle. Arch. Neurol. 66:697-702.

Seyfarth, RM, Cheney DL, Marler P (1980) Vervet monkey alarm calls: Semantic communication in a free-ranging primate. Animal Behaviour 28:1070-1094.

Seyfarth R, Cheney DL, Bergman T, Fischer J, Zuberbühler K, Hammerschmidt K (2010) The central importance of information in studies of animal communication. Animal Behav. 80:3–8.

Sole RV, Corominas B, Valverde S, Steels L (2010) Language Networks: their structure, function and evolution. Complexity 15:20-26.

Steele J, Ferrari PF, Fogassi L (2012) From action to language: comparative perspectives on primate tool use, gesture and the evolution of human language. Philos. Trans. R. Soc. Lond. B 367:4-9.

Tononi G (2008) Consciousness as integrated information: a provisional manifesto. Biol. Bull. 215:216-242.

Trappenberg TP (2010) Fundamentals of Computational Neuroscience, 2nd edition. Oxford University Press.

Turing AM (1964) Computing machinery and intelligence. In: *Mind and Machines*. AR Anderson, Ed., Prentice Hall, Englewood Cliffs, NJ.

Wilson SM, Galantucci S, Tartaglia MC, Rising K, Patterson DK, Henry ML, Ogar JM, DeLeon J, Miller BL, Gorno-Tempini ML (2011) Syntactic processing depends on dorsal language tracts. Neuron 72:397-403.

Yamazaki Y, Yokochi H, Tanaka M, Okanoya K, Iriki A (2010) Potential role of monkey inferior parietal neurons coding action semantic equivalences as precursors of parts of speech. Social Neurosci. 5:105-117.

Zhang Y, Yu J, Boland JE (2010) Semantics does not need a processing license from syntax in reading Chinese. J. Exp. Psychol. Learn Mem. Cogn. 36:765-781.

Zhdanova IV (2011) Sleep and its regulation in zebrafish. Rev. Neurosci. 22:27-36.