



biolinguistics

volume 6
issue 2

spring
2012

XP

ZP

...

γ

α

...

XP

ISSN 1450-3417

The Biolinguistics of Autism: Emergent Perspectives

Nicolas Bourguignon, Aparna Nadig & Daniel Valois

This contribution attempts to import the study of autism into the biolinguistics program by reviewing the current state of knowledge on its neurobiology, physiology, and verbal phenotypes from a comparative vantage point. A closer look at alternative approaches to the primacy of social cognition impairments in autism spectrum disorders suggests fundamental differences in every aspect of language comprehension and production, suggesting productive directions of research in auditory and visual speech processing as well as executive control. Strong emphasis is put on the great heterogeneity of autism phenotypes, raising important caveats towards an all-or-nothing classification of autism. The study of autism brings interesting clues about the nature and evolution of language, in particular its ontological connections with musical and visual perception as well as executive functions and generativity. Success in this endeavor hinges upon expanding beyond the received wisdom of autism as a purely social disorder and favoring a 'cognitive style'-approach increasingly called for both inside and outside the autistic community.

Keywords: autism spectrum disorders; executive functions; language processing; music; vision

Saying "person with autism" suggests that the autism can be separated from the person. But this is not the case. I can be separated from things that are not part of me, and I am still the same person. I am usually a "person with a blue shirt" one day, and a "person with a yellow shirt" the next day and I would still be the same person, because my clothing is not part of me. But autism is part of me. Autism is hard-wired into the ways my brain works. I am autistic because I **cannot** be separated from how my brain works.

(from J. Sinclair, 1999, "Why I dislike person first language"¹)

We are grateful to Kleantes Grohmann, Cedric Boeckx, Antonio Benítez-Burraco, and another anonymous reviewer for very fruitful exchanges in preparing this manuscript.

¹ The full version of this text is available under <http://autismmythbusters.com/general-public/autistic-vs-people-with-autism/jim-sinclair-why-i-dislike-person-first-language>. Mention of this reference to justify the use of the word 'autistic' rather than 'person with autism' was first made in Dawson *et al.* (2007). The term *autistic* will be used accordingly throughout the present article.



1. Introduction

The present article aims to make the study of Autism Spectrum Disorders (ASD) a chapter of the biolinguistic program, i.e. the study of language as an internal system of human biology (Jenkins 2000). It is argued that a cognitive neuroscience of ASD, in light of recent advances in neurolinguistics and cognitive psychology, can deepen our knowledge of the constitutive features of language and its evolution.

This paper has two explicit motivations. The first is to raise awareness of a view of ASD within the framework of ‘cognitive styles’ (Happé 1999; see also Mottron 2003) defined by strengths and weaknesses equally worthy of investigation. The specific strength–weakness fraction to be dwelled upon in this discussion is that of enhanced auditory and visual perception contrasted with decreased integration of perception into higher-order representations. The existence of different cognitive styles within the human species, notably as a result of variations in genetic and neurobiological underpinnings, holds promise for refining the comparative work integral to biolinguistics and cognitive science (Hauser *et al.* 2002, de Waal & Ferrari 2010). Accordingly, the second motivation is to provide an alternative to the common view of ASD as deficits mainly affecting the socio-cognitive aspects of language, specifically ‘theory of mind’, or the ability to infer from a person’s behavior their mental states, including beliefs, desires and emotions (Baron-Cohen 1995). Theory of mind and its precursor skills are taken to be important prerequisites for the acquisition and proper use of language in context (e.g., Bloom 2002). As a result, most early research on language in autistics focused on their striking *pragmatic* impairments, sometimes driven by the theory of mind model (Baltaxe 1977, Tager-Flusberg 1992, Surian *et al.* 1996), without undertaking — or paying full attention to — investigations of every aspect of language structure. Yet, despite its widespread success in the cognitive science culture and its recognized importance for early stages of language acquisition, theory of mind falls short as an explanatory account of ASD phenotypes (Frith & Happé 1994). ASD also involve symptoms and characteristics outside the realm of social cognition, which are addressed by alternative, domain-general and bottom-up approaches to ASD such as enhanced perceptual functioning (Mottron *et al.* 2006), weak central coherence (Happé & Frith 2006) and disruptions of executive functions (Ozonoff *et al.* 1991, Russo *et al.* 2007).

We argue that these theories reveal novel and important facts about language in ASD, in particular a generally different mode of language development possibly encompassing all levels of linguistic representation (e.g., phonology, semantics, syntax, in addition to pragmatics), rooted in important differences in neurobiological architecture. We present a synthesis of findings evaluating these alternative models, with a focus on the various neural discrepancies affecting perceptual functioning, central coherence, and executive function in ASD. We provide a discussion of their implications for the study of language structure and development in autism and hope to demonstrate how the rich, neurophysiologically grounded science of ASD can contribute to intrinsic developmental–evolutionary questions of biolinguistics.

2. Autism and Biolinguistics: Advantages and Challenges

Importing the study of ASD into the province of biolinguistics may further the advancement of comparative models of language development and evolution, principally their genetic and neurophysiological aspects. The main challenge to be faced in this enterprise, however, resides in the large genetic and neurophysiological heterogeneity of the autistic spectrum itself.

2.1. *Advantages: Intra-Species Variability*

From a genetic and neurobiological vantage point, the study of ASD has allowed for significant forays into the 'emergence hypothesis' (Casanova & Tillquist 2008), whereby the advent of language is thought to have endowed human populations with the cognitive armamentarium to ignite their dramatic social and cultural development (Tattersall 2004, Chomsky 2006, 2007). In the wake of seminal approaches put forth to study language evolution despite the paucity of reliable biological artifacts, cognitive biologists ventured to compare human and animal cognition as a means of inferring which of the building blocks of language may be shared between humans and animals on the one hand (Hauser *et al.* 2002, deWaal & Ferrari 2010) and between language and social cognition on the other (Fitch *et al.* 2010). Nevertheless, while cross-species comparisons and animal models certainly are useful in tracing back the "foundational abstractions" of human language and intelligence (Gallistel 2009), comparative work would be incomplete without consideration of the differences emerging *from within* the human species. As the Human Genome Project reached its first significant milestones, it has become incontrovertible that genetic variations, and the interaction thereof with the organism's environment, lie at the source of many psychiatric conditions, including autism (Cowan *et al.* 2002). It follows that genetically-based conditions affecting the neural building blocks of language constitute a promising means to explore its nature and origins, along with the ontological connections between language and other constituents of the human mind (Fisher & Marcus 2006, Marcus & Rabagliati 2006). Given the co-occurrence of the linguistic and social atypicalities that characterize autistic phenotypes, the study of ASD has long been considered a candidate of choice. Although the question of autism as a proxy to investigate the relationship between language and social cognition is not excluded, a central goal of the present article is to show that social cognition is not the only aspect of language in autism that deserves consideration.

2.2. *Challenges: Different Routes to the Same Outcome*

Despite the aforementioned merits of studying autism as part of biolinguistics, the most likely challenge to be faced in that enterprise is the large genotypic and phenotypic heterogeneity observed in the autistic spectrum, which leads one to expect great variability at the neurophysiological level as well. Textbook descriptions of

autism (DSM-IV; APA 1997) as a triad of *reduced social interactions, delayed or atypical language, and repetitive and restricted interests and behavior* portray only in broad strokes a highly heterogeneous set of symptoms and degrees of severity that often goes beyond the large unevenness in verbal and nonverbal performance across autistic individuals, how it comes to reorganize itself differently from individual to individual in the course of development, and how this reorganization should be explained at the neurobiological level (Joseph *et al.* 2002). A description of the functioning and abilities of autistics needs to incorporate many dimensions such as age, verbal and nonverbal intelligence, and the settings in which behavior takes place (e.g., experimental vs. natural settings; Klin *et al.* 2003).

This patchwork-like picture of autism brings about several caveats and empirical hurdles: First, any investigation of cognitive abilities in ASD must ideally discriminate the broad categories of *high-functioning* autism (which characterizes a substantial 45–60% of individuals with ASD in recent reports (Newschaffer *et al.* 2007; see also Steiman *et al.* 2011), or individuals without intellectual delay, as measured by standardized intelligence tests, and with functional or fluent language abilities, from autism accompanied by mild or severe intellectual delay and minimal or generally non-functional language. Yet, surveying current evidence in both high- and low-functioning autism may provide important information about the potential endophenotypes of ASD as a whole.

Second, many of the neurophysiological studies to date test individuals with a very broad age range and there is little comparability across tasks employed. Focusing on tighter age spans but testing hypotheses over the course of development, and selecting tasks and methods that complement prior findings would provide a clearer picture of how and why language may or may not develop in subpopulations of the autism spectrum.

Third, a careful understanding of language design in autism requires that one consider the distinction between autistics *with* and those *without* formal language impairment. To that effect, while the former may have genetic overlap with specific language impairment (Kjelgaard & Tager-Flusberg 2001; but see Whitehouse *et al.* 2007 for a counterargument²), the forthcoming review of neurophysiological data suggests that autistics without behaviorally-defined language impairment may also display patterns of language acquisition and processing that depart from that of typical populations.

This third point highlights that a complete understanding of individual differences in language acquisition and processing demands comparisons across language

² With regard to the debate on the genetic relationship between autism and SLI, a series of genetic analyses have broadened the focus of attention from the well-known FOXP2 gene to the neurexin-encoding gene CNTNAP2 by suggesting that mutations affecting the former, while not being a major susceptibility gene for autism or language impairment (Newbury *et al.* 2002), may nevertheless have upstream consequences on the latter's regulation (Vernes *et al.* 2008). By bringing in autism together with other common types of language disorders, this type of evidence suggests that language development (and evolution) might result from a cascade-like interaction of different genetic factors. See also Benítez-Burraco (in press) for discussion.

disorders to determine which aspects (beyond decreased pre-verbal social communication in early development, Tager-Flusberg *et al.* 2005) are ASD-specific, rather than common to individuals with language impairment more generally.

In fact, the heterogeneity of ASD phenotypes yields a vexing tension for scientists keen on developing a generalized model of autism. After intensive efforts to formulate a unitary explanation of these complex phenotypic characteristics, the current state of knowledge has converged on a more fragmented etiology of autism (Happé *et al.* 2006), notably for reasons including its very intricate and still incompletely understood genetic and neural underpinnings. Indeed, existing evidence points to several dozen different genetic mutations associated with autistic behavior (Geschwind 2008, Walsh *et al.* 2008). This, along with the behavioral diversity of ASD (Volkmar & Klin 2005), calls for an approach to autism as a collection of multiple genotypic and phenotypic traits and subgroups rather than a unitary cognitive disorder or condition. Yet, we must still account for the aforementioned triad of features that define ASD. Neuroanatomically, a possible explanation for this is that initially distinct genetic mutations hold analogous consequences for general cortical design or the development of neural networks (Geschwind & Levitt 2007, Walsh *et al.* 2008). In the next section we review findings on brain structure in ASD populations at the levels of minicolumns, hemispheric lateralization and functional connectivity. This overview will serve as a basis upon which the various linguistic discrepancies of ASD can be introduced in light of nonsocial approaches to autism.

3. Brain Architecture in ASD

Discrepancies have been observed at various levels of neurobiological architecture in autistic populations, in particular minicolumnar organization, hemispheric lateralization and connectivity. Although these levels have been studied independently, unified models of autistic neurobiology are beginning to emerge.

3.1. Minicolumns

Casanova *et al.*'s (2002) postmortem morphometric studies on the columnar architecture of the superior and middle temporal gyri in nine autistic patients revealed that their minicolumns were more numerous, smaller and less compact (i.e. more dispersed) than in non autistic individuals. The dorsal and middle portions of these areas typically support the spectro-temporal analyses of speech sounds, while more posterior and ventral parts are involved in accessing lexical representations (Hickok & Poeppel 2007). Minicolumns are vertical bundles of approximately 100 neurons that constitute the basic units of information processing in the brain (Mountcastle 1997). Among other mechanisms³, these assemblies bind their temporal activity via

³ For reasons of space, we do not address the issue of columnar functioning at a molecular level, although evidence points to the impact of columnar disorganization on several neurotrans-

different levels of oscillatory coherence, allowing for top-down sensory integration across distant cortical areas (cf. Senkowski *et al.* 2008, Gray *et al.* 1989).

Studies on cortical oscillatory rhythms during sound and speech processing report an asymmetric and hierarchical temporal sensitivity of auditory cortices, with increased left temporal and premotor sensitivity to segmental (i.e. phonemic) information (~40 ms, the duration of the gamma-band), but greater tuning to suprasegmental (i.e. syllabic) information in the right temporal auditory and premotor cortices, correlated with the duration of the theta-band (~200 ms; Luo & Poeppel 2007, Giraud *et al.* 2007). Other studies show that neurons in the right hemisphere are preferentially sensitive to more basic features of auditory processing such as pitch (Belin *et al.* 1998) and slower modulations of sounds typical of musical and prosodic phrases (Belin *et al.* 2002). This hemispheric asymmetry is presumably attributable to differences in the structure and physiology of neuronal assemblies in the left and right hemispheres (Giraud *et al.* 2007).

Under normal circumstances, minicolumns in the left hemisphere contain a greater number of large pyramidal neurons than those in the right (Hutsler 2003). These large neurons typically fire at higher temporal frequencies than the smaller neurons on the right. However, in line with Casanova *et al.*'s findings, several studies report significantly reduced cell size in autistic adults' brains (Kemper & Bauman 1998), including in the hippocampus (Raymond *et al.* 1996), the main source of theta oscillations (Vertes 2005, in Giraud *et al.* 2007). These data suggest that decreased cell size might mostly be detrimental to the phonemic perceptual functions of the left hemisphere, while preserving the right hemisphere's tuning to the syllabic and prosodic characteristics of speech. The 'left-ear' dominance hypothesis of auditory perception in autism (formulated as early as Blackstock 1978) is explored in section 5.

3.2. Hemispheric Lateralization

Given the close links existing between columnar development and brain lateralization (Stephan *et al.* 2007), the features of columnar organization in autism outlined above are likely to impact hemispheric lateralization generally, affecting particularly the large cortical network of language processing (Chugani 2008). Using an MRI regional cortical volume analysis in 16 autistic boys, Herbert *et al.* (2002) reported reversed brain asymmetry in anterior cortical areas traditionally linked to language processing. A region included in Broca's area (pars opercularis), active during syntactic processing (Embick *et al.* 2000) and verbal working memory (Smith & Jonides 1999), appeared 27% larger in the right hemisphere in the ASD group relative to 17% larger in the left hemisphere in controls.

Another study by De Fossé *et al.* (2004) comparing ASD children with or with-

mitters putatively involved in regulating important aspects of language development and brain plasticity, in particular the influence of GABA-ergic transmission during the critical period (Hensch 2005). Specific hypotheses on the correlates of minicolumnar disruption on GABA transmission in autism are formulated in Casanova *et al.* (2003).

out language impairments, children with specific language impairments (SLI), and typically developing children, suggests that reversed lateralization of frontal language areas is related to language impairments rather than autistic disorders per se. Herbert *et al.*'s (2004) comparison between ASD, language impaired children and typical controls reports that language impaired and autistic children had proportionally greater right hemisphere volume relative to typically developing and language-impaired participants, but that this right hemisphere bias was more pronounced in the autistic than the language impaired group. Detailed investigation of a shared rightward lateralization between ASD and SLI individuals is beyond the scope of this paper; based on neuroimaging and phenotypic data, Whitehouse and colleagues proposed that the brain asymmetry in SLI and ASD constitutes the same expression of different neurobiological etiologies (Whitehouse *et al.* 2007, 2008).

The lateralization of temporal regions implicated in the auditory and lexical processing of speech is less clear and probably depends in great part on variabilities in the exact anatomy and function of these areas as well as on methodological considerations. In Herbert *et al.*'s (2002) *a priori* analysis, a region corresponding to the Planum Temporale appeared 25% larger on the left in the autistic group relative to 5% larger on the left in the control group, but this difference was much less extreme than that observed in Broca's area. Post-hoc analyses revealed that the leftward lateralization in the autistic group was actually strongest and reached statistical significance in the posterior temporal fusiform gyrus, a region implicated in picture naming and lexical processing (cf. Indefrey & Levelt 2004 for review), which was 20% larger in the left in autistic subjects relative to 6% larger in the right in controls. Adjacent regions, however, showed a trend towards rightward lateralization in the ASD group, including the inferior fusiform gyrus implicated in face processing (Kanwisher *et al.* 1997). However, Jou *et al.* (2010) report significantly enhanced rightward cortical volume in the posterior superior temporal gyrus of ASD adolescents, and normal cortical volumes have been observed in the right Planum Temporale in ASD adults (Rojas *et al.* 2002) and children and adolescents (Rojas *et al.* 2005). Contrary to Herbert *et al.* (2002), Rojas *et al.*'s studies revealed decreased cortical volumes in the left Planum Temporale. Further research is needed to better establish the degrees of lateralization in Wernicke's area and the Planum Temporale in ASD, but existing evidence points to aberrant patterns of hemispheric lateralization in the cortical network of language in ASD populations.

3.3. *Functional Connectivity*

Besides its impact on hemispheric lateralization, atypical columnar development also has significant consequences on cortical connectivity (Casanova & Trippe 2009), in particular those that characterize large associative areas engaged in complex cognitive and linguistic functions. The large pyramidal cells of the left hemisphere mentioned earlier are thought to form the long-range connections between anterior and posterior language areas (Hutsler 2003). Accordingly, decreased amounts of magno-

pyramidal cells and correspondingly smaller minicolumns are likely to disrupt long-range connectivity. This was observed in fronto-parietal and parieto-temporal networks using structural and functional MRI (McAlonan *et al.* 2005, Just *et al.* 2007), as well as in central subcortical fiber structures such as the arcuate fasciculus using diffusion tensor imaging (Fletcher *et al.* 2010). By contrast, locally normal or enhanced short-range connectivity has been reported in posterior primary sensory cortices (occipital visual areas, cf. Belmonte & Yurgelun-Todd 2003; see also Buxhoeveden *et al.* 2004) and regions contained in Wernicke's area (Just *et al.* 2004).

Thus, studies on connectivity in autism distinguish between underconnectivity over large association areas and normal or enhanced connectivity of primary visual and posterior temporal areas. This distinction led several researchers to suggest that local overconnectivity might compensate for large-scale underconnectivity in the successful completion of specific cognitive tasks (Mottron *et al.* 2006, Just *et al.* 2004, Bertone *et al.* 2005, Williams & Casanova 2010). Interestingly, microstructural studies in typical brains indicate that the amount of large pyramidal cells in temporal language areas decreases as one moves posteriorly (Hutsler 2003), possibly making posterior areas less vulnerable to dysconnectivity and impaired developmental trajectories compared to more anterior brain regions (Carper *et al.* 2002). Also, the spacing of columnar assemblies in posterior language areas is greater in the left hemisphere than in the right in normal brains — an anatomical pattern similar to that observed in the visual cortex and suggesting stronger modular organization in the posterior parts of the left hemisphere (Galluske *et al.* 2000). Given the increased number and greater-than-normal dispersion of minicolumns observed in autistic brains by Casanova *et al.* (*op. cit.*), the hypothesis has emerged that autistic brains might be characterized by more numerous and hyperactive cortical modules, which may account for specific features of autistic behavior (Williams & Casanova 2010).

3.4. *Hopes and Hurdles for Unification*

Although the various discrepancies documented in the investigation of brain anatomy in autism have to a large extent been studied separately, one cannot afford to ignore the strong interdependencies between them. Attempts to integrate these observations in a single framework will prove useful, and necessary, in formulating empirically testable hypotheses on the distinctive cognitive processes that define autism (Coleman 2005). Geschwind (2008) expresses this expectation while also allowing for possible divergences in neural architecture within the autistic spectrum itself. Beyond the many developmental routes potentially related to multiple and divergent cases of autism, current integrated neurobiological hypotheses to date (e.g., Markram *et al.* 2007, Williams & Casanova 2010) managed to emphasize the following dichotomy to describe autistic cognition generally: On the one hand, skills requiring multimodal integration of information, for example language and social cognition, will likely be more vulnerable to dysfunction. For example, Damasio & Maurer (1978: 779) noticed that “the verbal defects of autism [...] are seen only in a set of [...]

transcortical aphasias that result from a more or less complete anatomical isolation of speech areas". On the other hand, principles of economy in wiring (Cherniak 1994; mentioned in Williams & Casanova 2010) may compensate for this large-scale underconnectivity with a local overconnectivity and hyper-functioning of modular cortical systems reacting to psychophysically 'simple' environmental features.

It is important at this point to clarify the particular meaning of the terms 'simple' or 'complex' as they are understood in our discussion. As in Samson *et al.* (2005), and in line with hierarchical cortical models of perception and learning (e.g., Friston 2005), we consider a neurocognitive system as 'complex' if it is organized into elemental but hierarchically nested units that encode correspondingly complex information. Accordingly, a decrease in the hierarchical organization of processing systems in autism may lead to the processing of narrower, possibly non-hierarchical units. In this sense, 'complexity' at the neurocognitive level should not be confounded with complexity at the level of a particular *task*, in that complex tasks may involve the manipulation of simple stimuli.

This propensity for complex manipulation of simple material is now often assumed to be a characteristic trait of autistic cognition. In its extreme form, it gives rise to special splinter skills (e.g., letter decoding, calculation, list memory, 2D- and 3D-drawing, and music) before functional language is attained at the cost of long, deliberate efforts in some individuals. Special talents are far from the rule in ASD, but are nonetheless particularly informative to the extent that they magnify cognitive trends that might be generally distributed across the autistic spectrum (Mottron *et al.* 2006), and provide important clues on the neuronal systems that may define autism as a whole. If such hypothesis holds, a crucial question arises for language — a prime example of hierarchical complexity at all levels of structure and use. In particular, individuals with ASD might extend their initial cognitive strengths in processing simple/unimodal stimuli to the learning and processing of higher-order and hierarchically complex cues over the course of their development, including those characterizing speech and natural syntax (Mottron *et al.* 2006). Yet, the dearth of longitudinal studies of neural development in autism makes it unclear if neuroanatomical differences reflect the end-state of years of living with a different phenotype and consequent differences in interaction with the environment, or a relative continuity of differences present in the 'initial state' of ASD. A crucial focus of current work in the neuroscience of autism should thus be to determine if these anatomical and functional differences are similarly observed in young children with ASD. In this scenario much work lies ahead in specifying how neuroanatomical differences modify the mechanisms of language acquisition, and, in turn, unraveling how atypical brain development determines language processing in autism.

4. Alternatives to Socio-Cognitive Models of Autism

The unifying hypotheses presented above echo several cognitive psychological mo-

dels of autism that do not consider social communication as its prime domain of deficit. To varying degrees, these models have accounted for autistic language processing in terms of the simple-complex dichotomy developed earlier: The models of enhanced perceptual functioning (EPF; Mottron *et al.* 2006) and weak central coherence (Happé & Frith 2006) have prominently shifted the focus of autism research to the *positive* impacts of autistics' processing bias towards simple, *non-hierarchical* cues. By contrast, models dwelling on autistics' weaknesses in executive functions (see Hill 2004, Russo *et al.* 2007) emphasize the possible difficulties autistics experience as a result of their limitations in processing and producing hierarchically complex stimuli, including sentences (Just *et al.* 2004).

In the remainder of this paper, we take each of these approaches as an illustration of how language in autism could be studied outside of its socio-cognitive aspects: Perceptual functioning in phonology, central coherence in word and sentence processing via visual imagery, and executive functions in the relation between language, thought, and action. We also endeavor to map these observations to those made in neurobiology. But before we proceed, we wish to emphasize that we do not treat these approaches as mutually exclusive in the sense that one (say, perceptual functioning) fares better than the other (say, central coherence) in accounting for a particular aspect of language (say, phonology). Given the theoretical proximity between some of these approaches, there is good reason to believe that they might end up complementing each other in explaining the same aspect of autistics' speech processing abilities. Nor do we claim that a particular discrepancy found at one level of language processing in autism necessarily entails a similar discrepancy at another level. Finally the great phenotypic variability so characteristic of ASD forces us to interpret any observed discrepancies as applying to the tested subgroup of individuals with ASD, without assuming that they should be found uniformly in all autistics. Resolving these issues will depend on the success of our predictions, on a better delineation of the various autistic phenotypes observed, and on how the aforementioned models of autism develop in the future.

5. Phonological Processing: Enhanced Perception of Local Auditory Features

Neurobiological and cognitive psychological evidence suggests a 'left-ear' preference of speech processing in autism as a result of smaller minicolumns, rightward hemispheric lateralization and decreased connectivity in left-hemispheric language areas. This might account for autistics' enhanced perception of phonological primitives processed preferentially in the right hemisphere and shorter neuronal assemblies, namely syllables and prosody, and suggests decreased hierarchical processing of phonemic within syllabic information. Developmental evidence shows that this pattern occurs early. Putative links with preserved or enhanced musical abilities in autism are discussed.

5.1. *Neurophysiological Evidence for Rightward Dominance of Speech Processing in Autism*

Beginning with adult data, decreased left-lateralization during auditory language processing was reported in a positron emission tomography (PET) study by Müller *et al.* (1999) with five high-functioning participants, and in an fMRI study with 26 young adults by Anderson *et al.* (2010).⁴ In another PET study on the processing of 200 ms steady-state synthetic CVC speech-like sounds in five autistic adults, Boddaert *et al.* (2003) observed both significantly lower activity in the left superior temporal cortex and increased activation of the right superior temporal and frontal areas.

Directly addressing the question of when such pattern occurs in development, a follow-up study with intellectually delayed autistic children (Boddaert *et al.* 2004) reported decreased left-hemispheric activity but failed to replicate any right hemispheric effect, suggesting that rightward lateralization of speech processing might occur as a function of age, IQ, and/or verbal ability. ERP and MEG research on sound-related cortical components (in particular the N/M100 cortical response reflecting early auditory processing) and fMRI studies on speech processing in ASD children have begun to refine the relationship between rightward lateralization and development in autistics: Delays in the right hemispheric N/M100 responses to subtle tone contrasts in ASD children are taken as evidence for atypical maturational development of the auditory system in autism generally (Gage *et al.* 2003a, 2003b, Roberts *et al.* 2010).

Beyond these potential delays, other evidence goes along Boddaert *et al.*'s (2004) assumption that the development of autistics' speech recognition system might also follow distinctive maturational trajectories. Compared to the well-established route towards increased left-lateralization in typical children's cortical activation to speech, Flagg *et al.* (2005) found a significant, age-related rightward lateralization in ASD children. Bruneau *et al.*'s (1999) study with intellectually delayed children with autism, normal and intellectually delayed controls reported tone intensity effects on the N/M100 amplitude in the right hemisphere in the ASD group only. Bruneau *et al.* (2003) replicated these results and showed that the ampli-

⁴ Interestingly, the reversed lateralization observed by Müller *et al.* (1999) in ASD participants was related only to speech perception, suggesting a dissociation between production and perception systems and lateralization in ASD. Subsequent imaging research on language production in ASD individuals remains scarce and offers mixed and oftentimes surprising results. In a response-naming fMRI study with ASD adolescents, Knaus *et al.* (2008) reported less left-lateralization but greater activation of Broca's area in the ASD relative to the control group. In a functional transcranial Doppler ultrasonography study on language production in adults with autism, adults with a history of SLI, language-impaired adults, and typical adults, Whitehouse *et al.* (2008) reported that the ASD group, like the typical and SLI-history group, had significant activation in the left hemisphere, while right-hemispheric or bilateral activation was mostly significant in the non-ASD language impaired groups. These results led the authors to suggest (in line with Whitehouse 2007) that the aberrant lateralization patterns shared between ASD and SLI individuals might be the similar expression of different neurobio-logical causes.

tude of the right temporal N/M100 was larger as participants' verbal and non-verbal communication abilities increased.

Along the same line, Redcay & Courchesne (2008) report that 2- to 3-year-old toddlers with provisional diagnosis of ASD showed greater rightward activity when presented with auditory bedtime stories during natural sleep (see also Eyler *et al.* 2010). Again, correlations showed that right-hemispheric activation was positively linked to verbal abilities and negatively correlated with autism severity. Interestingly, Wilson *et al.*'s (2007) MEG study reports reduced left-hemispheric steady state gamma-responses to non-speech sounds in autistic adolescents, while frequency power in the right hemisphere did not differ from controls. By contrast, Murias *et al.* (2007) observed significantly increased resting state theta rhythms in autistic relative to controls subjects. This increase in theta oscillations, most detectable in left temporal and frontal regions, is argued by the authors to reflect a decrease in long-range connectivity. The implications of these factors to autistics' language processing will be considered in turn.

5.2. 'Left-Ear' Bias in Speech Processing: Syllables and Prosody

Samson *et al.*'s (2005) review of the literature on auditory processing in ASD points out autistic populations' enhanced performance in tasks involving spectrally and temporally simple material, accounting for their superiority in identifying pitch changes (i.e. absolute pitch, Heaton *et al.* 1999), pure tone discrimination (Bonnell *et al.* 2003, Heaton *et al.* 1998), detection of local changes in contour-preserved melodies (Mottron *et al.* 2000), or — more occasionally — exquisite musical talent (Miller 1999). Other research has applied this hypothesis directly to language processing.

In a study comparing the perception and comprehension, by fluent autistic adolescents and non-autistic controls, of simple sentences with specific prosodic modulations and analogous musical sequences, Järvinen-Pasley *et al.* (2008, Study 1) observed that autistic adolescents performed significantly better than the control group in perceiving prosodic variations in both the linguistic and non-linguistic perceptual samples.⁵ Enhanced perceptual processing in autistics has also been found at the word and syllable levels. Mottron *et al.*'s (2001) study of word recall comparing high-functioning autistic and typical individuals reported that whereas typical individuals benefited more from semantic cueing in word recall, the autistic group was equally biased by semantic and syllabic cueing, suggesting that autistics "benefit equally from superficial (syllabic) and deep (semantic) recall cues" (p. 258).

⁵ Enhanced perception of prosody may appear as a striking contrast to reports of aberrant expressive prosody produced by autistic speakers (Nadig & Shaw 2012, Peppé *et al.* 2007, Shriberg *et al.* 2001). Global pitch production as well as different functional types of prosody (affective, grammatical, pragmatic) appear to be more disregulated than comprehension of prosody in ASD. Recent work documents atypical production of pitch and duration in non-social situations as well (e.g., Bonneh *et al.* 2011, naming; Diehl *et al.* 2011, imitation), suggesting that basic motor planning or production-perception feedback mechanisms (Russo *et al.* 2008) contribute to differences in prosodic production in ASD.

Using slightly larger groups and narrower age-ranges, Järvinen-Pasley *et al.* (2008, Study 2) compared typical and high-functioning ASD children's perception and comprehension of short sentences displaying specific syllabic rhythms. The autistic group performed significantly better than controls in perceiving syllabic rhythmicity, while the control group showed higher levels of sentence comprehension. Although these data point to enhanced perception of syllabic and prosodic patterns in autistics, it is difficult for now to know whether this pattern might ultimately be detrimental to language comprehension (see McCleery *et al.* 2010 for potential neurophysiological effects of auditory processing on the N400 ERP component in autistic children).

5.3. *Neurophysiological Evidence for Decreased Hemispheric Synchronization*

As neurophysiological research on phonological processing suggests that large neurons in the left hemisphere show increased sensitivity to phonemic variations (Giraud *et al.*, 2007), reports of long-range connectivity disruption (Fletcher *et al.* 2010) and smaller columnar units in auditory cortices (Casanova *et al.* 2002) in autism lead one to predict that autistics may show reduced sensitivity to subtle phonemic variations within syllabic tiers, as in the detection of consonant (e.g., /dɪp/ vs. /tɪp/) or vowel changes (e.g., /ɑ/ vs. /æ/). A recent fMRI study by Dinstein *et al.* (2011) comparing brain activation in autistic, language-delayed, and typically developing toddlers during verbal and non-verbal auditory stimuli presentation in natural sleep found significant evidence of hemispheric desynchronization in the ASD group.⁶

At a more fine-grained level, Event Related Brain Potentials (ERPs) studies provide evidence of decreased sensitivity to phonemic modulations, including those embedded in syllabic units. Ceponienė *et al.*'s (2003) ERP study on autistic participants' sensory and attentional integration of deviances involving simple tones, complex tones, and natural speech vowels in an 'oddball' paradigm (i.e. the detection of unpredictable events in otherwise consistent auditory sequences; cf. Näätänen *et al.* 1978, 1990) reports intact *sensory* processing of all sound categories but no *attentional* processing of vowel modulation, confirming ASD participants' atypical processing of phonemic variations but intact processing of non-speech sounds. Subsequent neurophysiological research corroborates atypicalities in attentional processing of phonemic changes contrasted with greater sensitivity to pitch (Lepistö *et al.* 2005, 2008) but decreased tuning to phonemic changes within syllables (discriminating /taa/ from /kaa/, for example; cf. Jansson-Verkasalo *et al.* 2003).

5.4. *Summary and Prospective Research Questions*

Atypical right-hemispheric dominance in auditory speech processing in autism has come to be increasingly consensual (see Haesen *et al.* 2011 for another review). Coupling such observations to those made on hemispheric specialization for speech pro-

⁶ It is important to note here that Dinstein's study did not allow the authors to determine the directionality of lateralization between the groups.

cessing leads us to formulate the following predictions: Autistics might show a 'left-ear' bias towards syllabic and prosodic patterns, a feature possibly shared in their preserved or enhanced processing of rhythmic and melodic patterns. By contrast, evidence suggests decreased sensitivity to primitives typically subserved by the left hemisphere, namely subtle phonemic variations, whether or not nested in syllabic constituents. This pattern appears to occur early in development, but the extent to which it is compensatory or detrimental to speech perception remains an open question. Beyond possible maturational delays in cortical activity of the right-hemisphere in autistic children without intellectual impairments (Roberts 2010), positive correlations between rightward lateralization of speech/non-speech sound perception and age (Flagg 2005) or verbal abilities in autistic children with intellectual delay (Bruneau *et al.* 2003) suggest that right hemisphere processing of speech is a compensatory mechanism in at least some subgroups of autistic participants.

Answers to the question as to how auditory language processing functions in autism might contribute a good deal to our understanding of how the evolution of complex auditory abilities could have furthered communication, hence social interactions. As Siegal & Blades (2003) point out, discrepancies in complex sound processing in autism, and their impact on autistics' social abilities, may well be more adequately accounted for through investigations of brain structures supporting human voice processing than by appeal to social-cognitive models of autism (see also Gervais *et al.* 2004). On the other hand, autistics' peculiar strengths in auditory perception and their link to language ability might appear quite valuable in studying the relationship between spoken language and cognitive capacities relying on the right hemisphere such as music (Levitin & Tirovolas 2009).

Detailed investigations of the link between musical capacities or enhanced perception of rhythmic/melodic patterns in autistics and their potential ability to exploit these skills in the perception of speech (syllabic vocalization, rhythm and prosody) could shed significant light on the evolutionary connection between these domains of human cognition. In any event, approaches to phonological perception in autism based on discrepancies at the structural and functional levels of neuronal assemblies seem to be gaining promising speed (Giraud & Poeppel 2012).

6. Word and Sentence-Level Processing: Greater Reliance on Visual Imagery in Lexical and Sentential Processing

Evidence shows that some autistics' visual processing is atypically active during performance in tasks of higher cognition, including language comprehension. Increased visual imagery might be particularly important, if not compensatory, in their integration of verbal material, in particular at the levels of words and sentences. Parallels with savant visual abilities and implications for language comprehension are addressed.

6.1. Behavioral and Neurophysiological Evidence for Enhanced Visual Imagery

Early reports of some autistics' strengths in visual processing were based on their enhanced performance on measures of visual intelligence such as the Embedded Figure Task (EFT; Shah & Frith 1983, Joliffe & Baron-Cohen 1997), whereby participants must detect geometric figures contained in more complex visual patterns. In particular, their success on the EFT indexes a tendency to ignore the global properties of images to the benefit of their local features. This local bias in visual integration contrasts radically from typical visual perception, which rather proceeds from global features to hierarchically organized subparts (Navon 1977). Interestingly, autistics' performance in the EFT is correlated with greater cortical activity in occipital areas relative to comparison participants (Ring *et al.* 1999), providing the neurophysiological basis for a 'visual imagery' approach to problem solving.

On a more general basis, several studies demonstrated that ASD individuals' level of intellectual functioning reached significantly higher results when measured through minimally verbal visual tasks such as the Wechsler Block Design subtest or the Raven's Progressive Matrices than through verbal subtests (Happé 1994, Dawson *et al.* 2007). Soulières *et al.* (2009) also demonstrated that autistics' performance in the Raven's matrices was linked to higher activation of occipital regions, while performance in the control group was linked to increased activity of prefrontal areas supporting working memory (Postle *et al.* 1999, Smith & Jonides 1999). A patent example of autism as a visual cognitive style nevertheless comes from autistic draftsmen able to reproduce scenes and objects with exquisite fidelity (Mottron & Belleville 1993) but evidence also shows that autistics' visual integration abilities decrease whenever second-order visual information is involved (Bertone *et al.* 2003), indicating that visual strengths in autism are restricted to simple, non-hierarchical visual material. This latter observation may explain autistic individuals' impaired perception of hierarchically-organized stimuli such as biological motion (Blake *et al.* 2003) or facial masks (Deruelle *et al.* 2010).

It must be reiterated yet again, however, that cognitive peaks in visual abilities are not always found in ASD. Higher verbal than visual abilities are found as well and these profiles may in fact specify different subgroups of autistic individuals (Black *et al.* 2009). Several studies using EFT did not replicate visual facilitation in autistic children, and researchers have recently come to criticize this task and its application to autism on a number of counts (see White & Saldaña 2011). Although neural imaging confirms enhanced activity of the visual cortex in autistics, careful replication of visual processing tasks in ASD individuals is needed to strengthen this argument.

In the late 1980s, autistics' islets of visual abilities figured as evidence for the development of the central coherence approach to autism (Frith 1989, Frith & Happé 2006). On a par with EPF, this approach also stresses the prevalence of simple over complex perception and derives from this perceptual hallmark autistic populations' typical attraction for small, isolated features of the environment and obsessive drive

for sameness. Extended to general cognitive processes (including auditory processing; see Frith & Happé 2006 for a synthesis), this perspective thus emphasizes that autistic perceptual processes are primarily not hierarchical, favoring fragmentary over holistic processing.

Here we focus on the primary findings that spawned the development of weak central coherence, namely peculiarities in visuo-spatial tasks, but findings of decreased hierarchical configuration and enhanced visual imagery have had ramifications in the description of language phenotypes in ASD (see Happé 1999 for review). Specifically, they predict that ASD individuals should show intact processing of isolated lexical items and would be inferior in processing hierarchically structured sentential constituents (see Frith & Snowling 1983 for early evidence).

An ancillary prediction linking facilitated lexical access and enhanced first-order visual processing is that people with autism should show near intact, even enhanced lexical access *via* visual imagery. Neuroanatomically, this phenomenon may find its roots in the greater activation of vision-related areas of the brain during the EFT, Block Design, or Raven's tasks mentioned above, but also in reports of aberrant lateralization of posterior temporal regions (Herbert *et al.* 2002), which are engaged in picture-naming tasks (Indefrey & Levelt 2003), mental image generation (D'Esposito *et al.* 1997), and reading (Dehaene & Cohen 2007) on the left, and in face processing on the right (Kanwisher *et al.* 1997), including during audio-visual speech processing in degraded auditory environments (Kawase *et al.* 1997). Interestingly, face-processing areas in autism show remarkably weak activation during face scanning (Pierce *et al.* 2001), suggesting the possibility that audio-visual perception of speech might be problematic in ASD (see section 6.4 below).

6.2. *Visual Imagery Enhances Lexical Access*

Existing behavioral and neurophysiological evidence with autistic participants supports the prediction that lexical access and visual imagery can be intact or superior in autism. Autistics appear to show relative strengths in lexical acquisition relative to other aspects of language (Tager-Flusberg *et al.* 2005) and are advantaged in word access in the pictorial (Kamio & Toichi 2000) and orthographic modalities (Toichi & Kamio 2002). Interestingly, Walenski *et al.*'s (2008) picture-naming study comparing high-functioning autistic and typical children report faster naming performance in the ASD compared to the typically developing group, providing evidence for more efficient lexical access in autism.

Current imaging research also suggests that facilitation in lexical access in autistics is related to increased activation of posterior temporal and occipital areas, even in the absence of pictorial prompts. In an fMRI study on word classification in ASD adults, Harris *et al.* (2006) observed increased activation of left posterior temporal areas (Wernicke's area) in the ASD group compared to the control group. Gaffrey *et al.*'s (2007) fMRI study on word classification in ASD participants and typical controls reported significantly increased bilateral activation in the visual cortex in

the ASD compared to the control group. Finally, in their fMRI study comparing performance in a pictorial reasoning task in 12 children with high-functioning autism and 12 age- and IQ-matched controls, Sahyoun *et al.* (2009) showed that although the two groups displayed similar activation in the typical language areas when verbal mediation was necessary, the autism group had substantially greater activation of occipital and ventro-temporal areas in the tasks requiring verbal mediation, while greater activation was found in temporo-frontal language regions in the typical group. The authors suggest that enhanced engagement of posterior regions across tasks in the autistic group indicates greater “reliance on visual mediation [...] in tasks of higher cognition”.

6.3. *Visual Imagery at the Sentence Level*

While current evidence supports the view that visual imagery might be linked to greater performance at the word level in ASD, evidence for decreased integration of words in hierarchically structured expressions is mixed, and questions remain unresolved as to whether autistic populations may achieve similar performance as typical, yet through different strategies. Early claims of weak central coherence effects in sentence processing come from studies reporting autistics’ decreased ability to choose the appropriate pronunciation of homographs according to their sentential context (e.g., *In her eyes/dress there was a big tear*; Frith & Snowling 1983, Happé 1997, Jolliffe & Baron-Cohen 1999, Lopez & Leekam 2003).

However, these claims have been challenged and/or refined on a number of counts. In a disambiguation study comparing children with autism and concomitant language impairment, children with autism but without language impairment, language-impaired children, and typically developing children using a picture selection paradigm, Norbury (2005) reported that both the autism group with language impairment and the language-impaired group performed equally worse than the ASD group without language impairments and the typically developing group, indicating that decreased ability to use context for disambiguation may stem from language impairment rather than autism *per se*. This effect was replicated in a lexical ambiguity resolution study by Nadig (2011), where children with high-functioning autism did not differ from typically developing peers matched on language level in being able to use a sentential context to disambiguate a homophone (e.g., *fan, bank, cell*) when pictures of each versions of the homophone were presented, as reflected by their anticipatory eye-movements.

Brock *et al.*'s (2008) findings from an eye-tracking study of sentence processing in 24 ASD adolescents and 24 controls brings fine-grained evidence that impairments in the use of sentential context to identify a particular word might be attributable to language impairment irrespective of whether or not participants are autistic. In one condition, a visual display accompanying an auditory sentence (e.g., *He stroked the hamster*) presented only the picture of a phonological competitor for the object noun (e.g., *hammer*) and unrelated pictures. Importantly, these sentences were semantically

constraining, such that the phonological competitor (*hammer*) was not a viable object for the verb *stroke*. ASD participants without language impairment and the language unimpaired control group inhibited looks to the hammer following constraining versus neutral verbs such as *chose*, demonstrating online use of sentential context. However, for constraining sentences both autistics with poor language skills and language-impaired controls continued to look at the hammer as candidate based on its phonological onset, despite the lack of fit with the semantics of the verb.

Taken together, these findings are at odds with the prediction of local, piecemeal processing of words in autism, and the consequent prediction of insensitivity to global sentential context. However the question remains as to whether underlying processing strategies are similar between autistics and typicals. Notably, given autistics' putatively intact or enhanced visual processing abilities, it is possible that the use of visual stimuli in lexical disambiguation or phonological competition tasks would have advantaged or facilitated processing in the autism groups.⁷ Earlier homograph disambiguation studies (e.g., Happé 1997) that found poorer performance in ASD groups did not present pictorial stimuli. Importantly, other research suggests that superior visual processing might not be sufficient for the comprehension of complex hierarchical structures and operations such as c-command or A-movement. For example, Perovic *et al.* (2007) tested autistic children's comprehension of actional vs. non-actional passives (e.g., *Mary was pushed by Thom; Mary was loved by Thom*) and anaphora vs. pronoun structures (e.g., identifying the antecedent in *Bart_i's dad_j is washing himself_f/him_i*) using a sentence-picture matching task. Autistics' poor performance at these tasks despite the use of pictorial material indicates that visual imagery may not be sufficient to compensate for core aspects of (Reuland 2001), at least in the early stages of language development.

Nevertheless, neural imaging has brought significant evidence that the use of visual imagery and enhanced lexical access still seems to constitute a key factor in autistics' sentence interpretation. For example, Kana *et al.*'s (2006) fMRI study compared brain activation between high-functioning autistic individuals and normal adults in processing sentences with high-imagery (e.g., *The number eight when rotated 90 degrees looks like a pair of eyeglasses*) vs. low-imagery (e.g., *Addition, subtraction, and multiplications are all math skills*) semantic content. In typical individuals, the processing of high-imagery sentences had already been shown to simultaneously engage areas typically activated during language comprehension and posterior areas subserving visuo-spatial processing, while processing low-imagery sentences activates language-related areas only (Just *et al.* 2004a), suggesting that large-scale integration of visual and verbal information is required when sentences have high imageability

⁷ By design, the majority of the target-competitor word pairs in Brock *et al.*'s study began with the same syllable (e.g., *bucket – butter; medal – medicine*), while Happé's (1997) stimuli contained phonemic variations within syllables (e.g., *There was a big tear in her eye/dress*). According to the hypotheses formulated in section 4, the fact that the ASD group performed as well as the control group in Brock *et al.*'s study but not in Happé's may be explained by their presumably intact perception of syllabic patterns but reduced perception of phonemic variations within hierarchically larger units.

content. In Kana *et al.*'s study, by contrast, whereas the simultaneous activation of language- and vision-related areas was triggered only by high-imagery sentences in the control group, ASD participants had increased activation of occipital and parietal areas for *both* high- and low-imagery sentences, while the language network was significantly less activated.

Based on these findings, the authors suggested that "there is a tendency in people with autism to use more visuo-spatial processing by recruiting posterior brain regions in accomplishing even language tasks" (p. 2485). Importantly, they propose to consider this effect as "an adaptation to the underconnectivity in autism, making greater use of parietal and occipital areas and relying less on frontal regions for linguistic processing" (p. 2492). A lexically- (and perhaps visual imagery-) rather than syntactically-based account of sentence processing in autism was also provided in an earlier fMRI study by the same group (Just *et al.* 2004b), in which enhanced activity in the posterior parts of the left superior and middle temporal gyri (i.e. Wernicke's area) in the ASD group contrasted with significantly increased activity of frontal areas in the control group. These results suggest that, "autistic participants may rely more on an enhanced word-processing ability and less on integrating processes that bring the words of a sentence together into an integrated syntactic and semantic structure".

Similar hypotheses on language processing in autism have already been formulated within the framework of other research agendas (e.g., Ullman 2004), but open questions persist as to the proper characterization of autistics' visually/lexically-based sentence processing strategies. First, we must still determine *what* particular visual representations are indeed activated in autistics' processing of verbal material, namely images of *words* or other, more abstract representations (if not both). Many of the studies described above involved reading written sentences or watching pictorial representations. As such, it is difficult to tell if the activation of visual and multimodal language areas reflected activation of graphemes or images with transparent semantic content. Also, warnings about heterogeneity in visual processing across the autistic spectrum must dampen the claim that *all* autistics profit from enhanced visual imagery to process language. In effect, these two issues might at some point end up confronting each other: If the hypothesis that activation of visual cortices in sentence processing actually reflects enhanced grapheme decoding turns out to be correct, then it must readily take into account the great heterogeneity of reading skills in autistics, ranging from floor to ceiling (Nation *et al.* 2006).

6.4. Summary and Prospective Research

Many questions remain open with regard to the place vision occupies in language design. These questions have often been the centre of much attention in language sciences, from lexical semantics (Jackendoff 1983) to language acquisition (Gleitman 1990) or speech processing (van Wassenhove *et al.* 2005) and language evolution generally (Corballis 2009). Studying the nature and use of visual imagery during

speech integration in ASD individuals may thus prove valuable on several counts. Notably, could autistic individuals' greater reliance on neural areas subserving visual processing to extract the meaning of words and sentences tell us anything about the mechanisms by which lexical concepts are acquired, processed and combined over time? Does there exist a correspondence between levels of visual complexity and particular levels of linguistic representation, and is it necessary, or even correct, to explain this correspondence by appealing to autistics' social deficits instead of the core mechanisms underlying their visual abilities?

From a computational point of view, the study of autism may help enlighten many grey areas regarding the computational origins of speech and language, in particular when these are assumed to have emerged from the 'social experience' of visually presented information (Gallese 2008). For example, autistic individuals seem to show resistance to McGurk effects (McGurk & McDonald 1976), involving cross-sensory integration of speech and facial articulatory movements (e.g., Mongillo *et al.* 2008). Should this phenomenon be explained in terms of autistic individuals' impaired social comprehension of facial masks, by their putatively deficient 'mirror neuron' detector (Williams *et al.* 2004) or rather by their decreased ability to use facial movements as hierarchical predictors of the speech input? While theory of mind may limit the explanation of this phenomenon to a failure to sense the social significance of face perception, an account centered on the levels of visual complexity in autism would allow for an exploration of the possible connections between visual intelligence and the underlying computational principles of natural languages. Naturally, exploring this territory will necessarily involve a deeper understanding of the computations of audio-visual speech. Luckily, evidence in this domain grows at a rather fast rate (Arnal *et al.* 2011).

On another line of thinking about the significance of graphical evidence in the evolution of language and mind, autistic draftsmen's accurate reproductions of visual scenes have led several authors to note that sophistication in human graphic feats may not necessarily be the sign of verbal intelligence as it is characterized in typical individuals today (Humphrey 1998 contra Tattersall 1998),⁸ sparking both new ideas and new doubts about early artistic artifacts as tokens of full-fledged human intelligence. In this respect, autism presents an undeniable comparative advantage. Importantly, one can view the study of autism as an opportunity to identify the distinctive roles that vision and language might have (had) with regard to internal thought processes, and what their respective benefits or disadvantages could be for human consciousness (Dennett 1992: Chap. 7).

That language and vision constitute initially independent but complementary

⁸ Among the most suggestive parallels drawn by Humphrey (1998) between cave art and savant drawings is the striking lack of symbolism, which puts into question interpretations of cave art as evidence for the emergence of a symbolic, hence possibly computational mind. It is also worth pointing out, as Humphrey does, that these parallels serve as arguments on what "we should not assume about the mental capacities of the cave artists" (p. 171) and constitute in no case the basis for speculations about common clinical phenotypes between modern autistic populations and cave artists.

tools for thought is reflected in anecdotes from autistic savant artists. For example, Lorna Selfe (1995) tells us the story of Nadia, a gifted autistic child born in 1967, whose drawing abilities ultimately waned following her first steps in actual linguistic communication at the age of eight. Temple Grandin's (1996) book *Thinking in Pictures*, by emphasizing the primacy of visual over verbal information in her daily stream of consciousness, has a similar sort of flavor. If these personal stories turn out to be correct, we believe that certain types of autism as being at one extreme of the 'verbalizer-visualizer' cognitive continuum, where the cognitive functions of 'inner speech' (Carruthers 2002) could be compared to those of 'private diagram-drawing' (Dennett 1992), set the stage for a direct investigation of their respective advantages and weaknesses.

Empirical research in this area is obviously challenging, and therefore scant (see Hulburt *et al.* 1994 for an early attempt with ASD individuals), but the issues at stake have begun to emerge along with an adequate research framework. Two questions deserve consideration: First, if private speech allows for cognitive functions that private diagram-drawing does not, autistics' performance should be decreased in tasks tapping the former, but not the latter. Second, if private diagram-drawing allows for roundabout strategies to solve problems typically hinging upon inner speech, as seems to be the case for sentence processing, neural imaging should provide ways to discover how this happens in autism. As for the particular research framework within which these questions can be addressed, Hinzen's (2008: 355) mention of the "systems of *executive control* that both human and non-human animals exercise when planning a sequence of actions so as to achieve a particular goal" (italics ours) provides an ideal entry into the problem. In the last section of this paper we sketch out how an Executive Function (EF) approach to autism might serve the purposes of biolinguistics. This section is admittedly the most speculative part of our discussion, so we will limit ourselves to a brief description the areas of EF in autism that we think merit close attention.

7. Executive Functions in Autism: Connectivity and the Prefrontal Cortex

Aberrant neural organization in the prefrontal cortex in autism is linked to weaknesses in higher-order executive control of thought and action, with possible ramifications for several aspects of language comprehension and production, specifically the role of inner speech in complex planning, monitoring of verbal information along its various dimensions, and generativity.

7.1. *Neurophysiological and Behavioral Evidence for Executive Function Discrepancies in Autistic Speech*

The most striking patterns of aberrant developmental trajectories and cortical architecture in autism appear in the prefrontal cortex (Carper *et al.* 2002). Among other

discrepancies, Courchesne & Pierce (2005) point out excessive and disorganized connectivity *within* the frontal lobes and poor connectivity *between* the frontal lobes and other cortical areas. The importance of the prefrontal cortex and the long-range connections it shares with virtually all other regions of the brain has long been acknowledged in subserving complex EF such as problem solving, language, decision, attention, planning, and goal-directed behavior (Fuster 2008). It is therefore unsurprising that autistic populations show several deficits in mental flexibility and planning, or perseveration (Hill 2004). Regions of the prefrontal cortex for which aberrant lateralization has been reported, such as Broca's area, are not only tonically active in processing language-like hierarchical structures (Musso *et al.* 2003) but also seem to play a critical role in the hierarchical organization of human behavior generally, leading to the conjecture that language may share the same hierarchical properties as those underlying complex human activities (Koechlin & Jubault 2006, Fuster 2008).

Hypotheses of EF as the 'private speech' underlying human thought and behavior (Vygotsky 1962, Luria 1979) not only echo the linguists' suggestions that language may constitute the very "skeleton of thought" (Hinzen 2009), but also conflate the ideas of EF and language as workspace of human planning and decision-making (Hinzen 2008). The rapprochement appears equally well as language and EF have both been assumed to constitute the basis of human creativity, in particular the generative properties so typical of natural languages (Goldberg 2009, see also Fuster 2008: Chap. 5). The proposed limited use of inner speech in autistic populations (Whitehouse *et al.* 2006) resulting from their EF impairments therefore raises at least three questions: Do autistics' "deficits in planning and discourse processing" (Hinzen 2008) tell us anything about the role of language in regulating human thought? (2) Do autistics' superior skills in visual processing lead them to manipulate verbal information in peculiar ways? And (3) do autistics EF impairments have connections to language generativity? We will briefly touch on these points in turn.⁹

7.2. *Inner Speech and Planning*

Regarding question (1), if inner speech has a role to play in an individual's decision-making ability, autistics should show specific impairments in planning as a result of limited use of inner speech. Poor performance on the Wisconsin Card Sorting Task (WCST), tapping into participants' rule and set-shifting ability, was part of the first evidence to have motivated the development of executive theories of autism (Ozonoff *et al.* 1991). Impaired performance on WCST is believed to reflect an inability to establish goal hierarchies and flexibly shift attention from one set of rules to another. Interestingly, neuropsychological studies suggest that WCST performance is verbally

⁹ It is important to note that there are multiple components of executive function and that atypical EF profiles are present in neurodevelopmental disorders more generally (cf. Happé *et al.* 2006, Ozonoff & Jensen 1999). Future work should pinpoint more clearly the profile specific to ASD, and how this set of EF strengths may be related to enhanced performance on visual imagery tasks (cf. Eigsti 2011).

mediated and depends on the integrity of crucial language brain regions (Baldo *et al.* 2005, but see Konishi 1998). It is intriguing to note from Baldo *et al.*'s (2005) study that inner speech impairments in aphasic patients provoked perseverations, or repetitive responses not related to the changing problem presented, not only in WCST, but also in the Raven's, even though both tasks initially tap into visual processing.

However, a proportion of high-functioning autistic individuals are impaired in the former, but unimpaired or superior in the latter, suggesting that enhanced visual processing could compensate or successfully replace weaker inner speech in solving certain visual problems but not others (Kunda & Goel 2011). A possible answer lies in the fact that whereas WCST requires fluctuant application of different rules to the same input, the Raven's Matrices do not. If this turns out to be the critical factor, one could infer that inner speech (or lack thereof) specifically supports (or impairs) the ability to flexibly switch from one task to the other (see Emerson & Miyake 2000 for experimental evidence). Further research is needed to explore this question.

Another EF task possibly requiring covert vocalization and for which individuals with autism show particular impairments is the Tower of London task or its variants (Ozonoff & McEvoy 1991).¹⁰ It is possible that the Tower of London and WCST both necessitate inner speech to a greater extent than the Raven's matrices as a result of requiring more complex planning abilities. If so, this would support the hypothesis that language is an important tool for setting long-term goals. Along similar lines, Carruthers (2002) proposes that EF and inner speech have an important part to play in perceiving and planning the behavior of other people, making them important components of theory of mind (Carruthers 2002, Newton & deVilliers 2007, but see Forgeot d'Arc & Ramus 2011).

These hypotheses parallel those of studies attempting to link autistics' ability to pass false-belief tasks and their acquisition of complement syntax (Tager-Flusberg & Joseph 2005; for an original argument on the relationship between complementation and theory of mind, see de Villiers & Pyers 2002) or other striking reports of autistics' success at false-belief tasks after achieving a certain verbal mental age (Happé 1995). Regarding social cognition generally, authors have observed that autistics' level of social functioning was significantly linked to their verbal abilities (Joseph *et al.* 2002), possibly making linguistic competence a crucial compensatory mechanism of their deficit in social cognition, perhaps more so than in typical children, strengthening further the link between language and social cognition.

7.3. *Monitoring Verbal Information across its Various Dimensions*

With regard to question (2), EF and the prefrontal cortex are important for the flexible selection of stimuli according to their nature, context and cross-temporal contingencies (Koechlin *et al.* 2003), for example when subjects are asked to judge the same verbal item along its different levels of representation, e.g., orthography, phon-

¹⁰ For an application of the Tower of London to prefrontal functions, see Shallice (1982). A variant of this task is the Tower of Hanoi.

ology, and meaning. Research on working memory and EF also shows hemispheric selectivity between left and right prefrontal regions, with the left frontal cortex subserving verbal information, and the right visuo-spatial stimuli (Smith & Jonides 1999). Accordingly, autistics' enhanced perceptual bias towards the visual features of words along with their rightward bias in Broca's area might lead them to perseverate on their orthographic rather than phonological or semantic aspects. This was shown in Toichi & Kamio (2002), who compared autistic and learning-disabled adults and adolescents' discrimination of words based on their orthographic properties, pronunciation, or meaning.

Results indicated not only that the autistic group had no level-of-processing effect compared to the control group, but also that the autistic group performed better than the control group in the orthographic task, suggesting a processing perseverance at the orthographic relative to phonological and semantic level. Interestingly, Harris *et al.*'s (2006) fMRI study on levels-of-processing effects in autistic and control participants reports that while activation of Broca's area was significant for the semantic relative to the orthographic task in the control group, its activation was undifferentiated between the two conditions in the ASD group. Koshino *et al.*'s (2005) fMRI study on verbal working memory comparing high-functioning and control participants provides even more compelling evidence. The authors observed that the control group had substantially more activation in the left and right prefrontal regions, while the autistic group had significant activation in right prefrontal and parietal regions, suggesting that autistic participants would have used a "visual-graphical approach [...] in which they coded the shape of the alphabet letters without naming them" (p. 818).

Such conclusions are interesting but raise a few parallel issues to be worked through. First, the link between right prefrontal regions and 'letter decoding' must be checked against neurophysiological theories that locate letter decoding in left inferior temporal regions (Dehaene & Cohen 2007), which — interestingly enough — also showed signs of significantly greater activation in the ASD relative to the control group (see also hypotheses on visual imagery sketched in section 6). Second, that visuo-spatial strategies could somewhat supplant manipulation of verbal information does not entail that inner speech is totally absent in ASD populations (Williams *et al.* 2008), nor that visuo-spatial working memory capacity is exempt from impairments as a function of stimulus complexity (Williams *et al.* 2005). Further research will be needed to refine this question, taking into account age, functioning, task demands, and neurophysiological factors.

7.4. *Generativity*

We wish to end this section with a brief mention of the studies that have investigated *generativity* in ASD populations. The notion of EF as an important contributing factor of creativity (Shallice 1988, Goldberg 2009, Fuster 2008) has been used to account for autistics' impaired ideational fluency in play (Lewis & Boucher 1995) and, more

recently, language production (Turner 1999). These characteristics might be visible at varying degrees in the development of intellectually unimpaired and impaired individuals. Tager-Flusberg *et al.*'s (1990) longitudinal study on language development between autistic children and children with Down syndrome remarks that "autistic children [...] tend to rely on a narrower range of grammatical structure in their spontaneous speech" (p. 17), despite similar levels of syntactic development as, and higher IQ levels than, children with Down syndrome.

Other research points to autistics' lack of flexibility in structural levels of linguistic representation, as reflected in extreme forms of echolalia¹¹ (Roberts 1989), 'stereotyped language', and "gestalt language learning patterns exhibited by autistic individuals who, unlike unimpaired children, may not develop a truly flexible syntactic rule system" (Landa 2000: 127). Interestingly, cases of limited syntactic flexibility must also be contrasted with instances of exaggerated lexical creativity such as the production of neologisms and idiosyncratic language (Volden & Lord 1991). Facts such as these are difficult to accommodate within a socio-cognitive account of autism but certainly deserve closer inspection from a 'generative' perspective.

8. Spreading the Net: Conceptual Payoffs for the Bilingual Program

Granted some consensus emerges on the topics we have discussed, we believe that the perspective advocated in the present article might help advance some of the core theoretical work in biolinguistics in a more concrete and observable way. In particular, more light could eventually be shed on the definition and the relative contribution of the conceptual divide between the Faculty of Language in the broad and narrow sense (FLB vs. FLN; Hauser *et al.* 2002) as well as on the relationship between language and other facets of cognition. Importantly, the constructs brought forth by alternative models of autism — *central coherence* in auditory and visual perception; *visual imagery* in concept acquisition and audiovisual language; *generativity* and *monitoring* in executive functions — may all in our view be part of the infrastructure of FL. We see two significant advantages to their introduction into biolinguistics: one related to the constituents of cognition that could have served and interacted as precursors to this faculty altogether; the other related to the importance of embedding ASD and its features into solid computational theories of neural functioning. We will briefly exemplify them in turn.

First, the cognitive phenomena highlighted throughout this text turn out to be necessary for other cognitive abilities likely to form part of FL broadly and narrowly defined. For example, prefrontal executive functions are necessary components of

¹¹ One must note that echolalia takes different forms in autism, with different levels of severity and functional roles as a result of different levels of development or functioning. Early studies on echolalia in autism have proposed interesting ways of using autistic echolalia as an indicator of propositional speech development (Baltaxe & Simmons 1977). Accordingly, we speculate that various forms of echolalia could be related to different levels of sophistication in grammatical generative power.

meta-representation (Stuss *et al.* 2001, Ozonoff *et al.* 1991), which is in turn deemed to be an important requisite for sophisticated intra-species communication. Furthermore, the view that executive functions are the “generative capability of the frontal lobes that made complex propositional structures possible” (Goldberg 2009) points to yet new bases for complex recursive thinking. At a lower level, central coherence could be analogous to the temporal *binding* of sensorimotor information necessary to construct higher-order representational hierarchies across neural networks, be it for auditory language or for other cognitive abilities (Engel *et al.* 2001). As a case in point, our discussion of the possible impacts of underconnectivity on cortical oscillations and phonological processing is only part of the broader discussion on the role of endogenous cortical cycles in perception and cognition (Fries *et al.* 2007), providing strong empirical and theoretical extensions of central coherence in autism. Similarly, we mentioned that impairments in the hierarchical integration of audiovisual information could contribute to autistics’ resistance to McGurk illusions. Rather than appealing to socio-cognitive explanations of this phenomenon, our understanding of this impairment would gain significant depth through hierarchical cortical models of perception (Friston 2005, Rao & Ballard 1999), especially if it is confirmed that cortical hierarchies are precisely what may be jeopardized in ASD. One advantage for taking these factors into account in characterizing FL is to understand not only what the precursors to language are (e.g., vision, central coherence, generativity, etc.), but how they interface hierarchically with one another within the constraints of neural architecture to eventually give rise to a full-fledged capability for language structure and use.

The second advantage follows directly from the first: A very exciting move in the study of language in ASD would be to look at central coherence, enhanced perceptual functioning, and executive function in light of existing computational theories. For example, the study of central coherence could be embedded within fine-grained and biologically realistic models of binding, asymmetric sampling and predictive coding (Engel *et al.* 2001, Bever & Poeppel 2011, Giraud & Poeppel 2012).¹² The same is arguably true for the computational principles underlying EF, which have received much support both from a theoretical (Dehaene & Changeux 1997) and empirical point of view (Koechlin *et al.* 2003) but remain largely absent from the literature on autism. In effect, the reason why autism research has been so hard to reconcile with contemporary language science beyond its socio-cognitive considerations is possibly the failure to appreciate that autism, much like social cognition or language, is a collection of different perceptual and cognitive factors, each of which is altered in its own computational and neurobiological machinery. If, by contrast, the multiple perceptual and cognitive facets of autism — and, for that matter, of *every* developmental disorder implicating language — are understood and specified through grounded explanatory theories of neural computation, biolinguistics could go a long way into the reverse-engineering agenda it has set out to pursue.

¹² We are grateful to one of the reviewers for bringing this point to our attention.

9. Conclusion

The present article was an attempt to integrate the study of autism within the framework of the biolinguistic program along two interconnected perspectives, namely that of autism as a cognitive style, on the one hand, and of autism as a heterogeneous set of verbal and nonverbal behaviors outside the realm of social cognition, on the other. These perspectives have led us to consider three alternative approaches of autistic cognition that focus on differences in perception and cognition (driven by differences in neural architecture), and their application to linguistic traits observed in autism. We propose that these traits hold promise for understanding individual linguistic differences if they are explored in the neurosciences of language: brain lateralization in auditory language processing, the role of visual intelligence in defining the nature and trajectories of language design and evolution, and the parallel between language and executive functions.

Importantly, we emphasize that our paper should be construed less as a discussion *on autism* than as a review of the *ways in which autism can feed the research program pursued in biolinguistics*. It is therefore neither comprehensive, nor integrative. Its primary goal is to show that the use of comparisons with autism to elucidate only pragmatic aspects of language is an insufficient and unnecessarily limited approach, and that this should be complemented with bottom-up, alternative, and empirically testable hypotheses that do not necessarily appeal to social cognition. In short, we hope to have shown that there is more to study about language in autistic populations than their assumed “blindness to Gricean Maxims” (Surian *et al.* 1996) and that thorough understanding of linguistic phenotypes in autism requires domain-general, neuroscientifically explainable, and ultimately *computational* hypotheses encompassing *every* level of linguistic representation.

This is not to say, however, that studying the interface between language and social cognition through autism is no longer worthwhile. To the contrary, we argue that the perspective defended here might bring pending research questions back to the forefront: Where are the links, both biological and psychological, between social cognition and language to be found? *Are* there any such links? Are these links a “spandrel” or otherwise characterized “cultural recycling” of the brain (Dehaene & Cohen 2007)? More particularly, did the computational complexity of social cognition, if any, feed into language or vice versa (Fitch 2005)? Addressing these issues also requires recognizing that to fully understand the social phenotype in autism, one must strive to tease apart aspects of autistics’ social cognition that do present deficits from those that don’t. As Sinclair’s epigraph expresses quite clearly, a growing number of people within the autistic community struggle for their recognition within society as ‘another intelligence’, where their preoccupations and interests deserve as much heed as our common habits of verbal interchanges (Wollman 2008). As in any other fields of science, this paradox certainly summarizes how complex the problem becomes when looked at carefully, but comes yet again with novel and exciting questions about the place of language within human nature and society.

References

- Anderson, Jeffrey S., Nicholas Lange, Alyson Froehlich, Molly DuBray B., Jason Druzgal, Michael P. Froimowitz, Andrew L. Alexander, Erin D. Bigler & Janet E. Lainhart. 2010. Decreased left posterior insular activity during auditory language in autism. *American Journal of Neuroradiology* 31, 131–139.
- Arnal, Luc H., Valentin Wyart & Anne-Lise Giraud. 2011. Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience* 14, 797–803.
- Baldo, Juliana V., Nina F. Dronkers, David Wilkins, Carl Ludy, Patricia Raskin & Jiye Kim. 2005. Is problem solving dependent on language? *Brain and Language* 92, 240–250.
- Baltaxe, Christiane A. 1977. Pragmatic deficits in the language of autistic adolescents. *Journal of Pediatric Psychology* 2, 176–180.
- Baltaxe, Christiane A. & James Q. Simmons. 1977. Bedtime soliloquies and linguistic competence in autism. *Journal of Speech and Hearing Disorders* 42, 476–393.
- Baron-Cohen, Simon. 1995. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press.
- Belin, Pascal, Stephen McAdams, Bennet Smith, Sophie Savel, Lionel Thivard, Séverine Samson & Yves Samson. 1998. The functional anatomy of sound intensity discrimination. *The Journal of Neuroscience* 18, 6388–6394.
- Belin, Pascal, Stephen McAdams, Lionel Thivard, Bennet Smith, Sophie Savel, Mônica Zilbovicius, Séverine Samson & Yves Samson. 2002. The neuroanatomical substrate of sound duration discrimination. *Neuropsychologia* 40, 1956–1964.
- Belmonte, Matthew K. & Deborah Yurgelun-Todd. 2003. Functional anatomy of impaired selective attention and compensatory processing in autism. *Cognitive Brain Research* 17, 651–664.
- Bertone, Armando, Laurent Mottron, Patricia Jelenic & Jocelyn Faubert. 2003. Motion perception in autism: A “complex” issue. *Journal of Cognitive Neuroscience* 15, 219–225.
- Bever, Thomas G. & David Poeppel. 2010. Analysis by synthesis: A (re-)emerging program of research for language and vision. *Biolinguistics* 4, 174–200.
- Bishop, Somer L., Whitney Guthrie, Mia Coffing & Catherine Lord. 2011. Convergent validity of the Mullen scales of early learning and the differential ability scales in children with autism spectrum disorders. *American Journal of Intellectual and Developmental Disorders* 116, 331–343.
- Black, David O., Gregory L. Wallace, Jennifer L. Sokoloff & Lauren Kenworthy. 2009. Brief report: IQ split predicts social symptoms and communication abilities in high-functioning children with autism spectrum disorders. *Journal of Autism and Developmental Disorders* 39, 1613–1619.
- Blackstock, Edward G. 1978. Cerebral asymmetry and the development of early infantile autism. *Journal of Autism and Childhood Schizophrenia* 8, 339–353.
- Blake, Randolph, Lauren M. Turner, Moria J. Smoski, Stacie L. Pozdol & Wendy L.

- Stone. 2003. Visual recognition of biological motion is impaired in children with autism. *Psychological Science* 14, 151–157.
- Bloom, Paul. 2002. *How Children Acquire the Meaning of Words*. Cambridge, MA: MIT Press.
- Boddaert, Nathalie, Pascal Belin, Nadia Chabane, Jean-Baptiste Poline, Catherine Barthélémy, Francis Brunelle, Marie-Christine Mouren-Simeoni, Yves Samson & Mônica Zilbovicius. 2003. Perception of complex sounds: Abnormal pattern of cortical activation in autism. *American Journal of Psychiatry* 160, 2057–2060.
- Boddaert, Nathalie, Nadia Chabane, Pascal Belin, Marie Bougeois, Vincent Royer, Catherine Barthélémy, Marie-Christine Mouren-Simeoni, Anne Philippe, Francis Brunelle, Yves Samson & Mônica Zilbovicius. 2004. Perception of complex sounds in autism: Abnormal auditory cortical processing in children. *American Journal of Psychiatry* 161, 2117–2120.
- Bonneh, Yoram S., Yoram Levanon, Omrit Dean-Pardo, Lan Lossos & Yael Adini. 2011. Abnormal speech spectrum and increased pitch variability in young autistic children. *Frontiers in Human Neuroscience* 4, 1–7.
- Bonnel, Anna, Laurent Mottron, Isabelle Peretz, Manon Trudel, Erick Gallun & Anne-Marie Bonnel. 2003. Enhanced pitch sensitivity in individuals with autism: A signal detection analysis. *Journal of Cognitive Neuroscience* 15, 226–235.
- Brock, Jon, Courtenay F. Norbury, Einav Shiri & Kate Nation. 2008. Do individuals with autism process words in context? Evidence from language-mediated eye-movements. *Cognition* 108, 896–904.
- Bruneau, Nicole, Frédérique Bonnet-Bribault, Marie Gomot, Jean-Louis Adrien & Catherine Barthélémy. 2003. Cortical auditory processing and communication in children with autism: Electrophysiological/behavioral relations. *International Journal of Psychophysiology* 51, 17–25.
- Bruneau, Nicole, Sylvie Roux, Jean-Louis Adrien & Catherine Barthélémy. 1999. Auditory associative cortex dysfunction in children with autism: Evidence from late auditory evoked potentials. *Clinical Neurophysiology* 110, 1927–1934.
- Buxhoeveden, Daniel, Katerina Semendeferi, Natalie Schenker & Eric Courchesne. 2004. Decreased cell column spacing in autism. *Society for Neuroscience* 582, 6.
- Carper, Ruth A., Pamela Moses, Zachary D. Tigue & Eric Courchesne. 2002. Cerebral lobes in autism: Early hyperplasia and abnormal age effects. *NeuroImage* 16, 1038–1051.
- Carruthers, Peter. 2002. The cognitive functions of language. *Behavioral and Brain Sciences* 25, 657–726.
- Casanova, Manuel F., Daniel P. Buxhoeveden & Juan Carlos Gomez. 2003. Disruption in the inhibitory architecture of the cell minicolumn: Implications for autism. *The Neuroscientist* 9, 496–507.
- Casanova, Manuel F., Daniel P. Buxhoeveden, Andrew Switala & Emil Roy. 2002. Minicolumnar pathology in autism. *Neurology* 58, 428–432.
- Casanova, Manuel F. & Christopher R. Tillquist. 2007. Encephalization, emergent properties, and psychiatry: A minicolumnar perspective. *The Neuroscientist* 14,

101–118.

- Casanova, Manuel F. & Juan Trippe. 2009. Radial cytoarchitecture and patterns of cortical connectivity in autism. *Philosophical Transactions of the Royal Society* 364, 1433–1436.
- Ceponienė, Rita, Tuulia Lepistö, Shestakova, A., Vanhala, R., Alku, P. & Risto Näätänen. 2003. Speech-sound selective auditory impairment in children with autism: They can perceive but do not attend. *Proceedings of the National Academy of Sciences of the United States of America* 100, 5567–5572.
- Cherniak, Christopher. 1994. Component placement optimization in the brain. *The Journal of Neuroscience* 14, 2418–2427.
- Chomsky, Noam. 2005. Three factors in language design. *Linguistic Inquiry* 36, 1–22.
- Chomsky, Noam. 2006. *Language and Mind*. Cambridge, MA: Cambridge University Press.
- Chomsky, Noam. 2007. Approaching UG from below. In Uli Sauerland & Hans-Martin Gärtner (eds.), *Interfaces + Recursion = Language: Chomsky's Minimalism and the View from Syntax-Semantics*, 1–29. Berlin: DeGruyter.
- Chugani, Diane C. 2004. Serotonin in autism and pediatric epilepsies. *Mental Retardation and Developmental Disabilities Research Reviews* 10, 112–116.
- Coleman, Mary. 2005. A neurological framework. In Mary Coleman (ed.), *The Neurology of Autism*, 40–74. Oxford: Oxford University Press.
- Corballis, Michael C. 2009. Language as gesture. *Human Movement Science* 28, 556–565.
- Courchesne, Eric & Karen Pierce. 2005. Brain overgrowth in autism during a critical time in development: Implication for frontal pyramidal neuron and interneuron development and connectivity. *International Journal of Developmental Neuroscience* 23, 153–170.
- Cowan, Maxwell, Kathy L. Kopinsky & Steven E. Hyman. 2002. The human genome project and its impact on psychiatry. *Annual Reviews in Neurosciences* 25, 1–50.
- Craik, Fergus I. & Robert S. Lockhart. 1972. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior* 11, 671–684.
- D'Esposito, Mark, John A. Detre, Geoffrey K. Aguirre, Matthew Stallcup, David C. Alsop, Lynette J. Tippet & Martha J. Farah. 1997. A functional MRI study of mental image generation. *Neuropsychologia* 35, 725–730.
- Damasio, Antonio R. & Ralph G. Maurer. 1978. A neurological model for childhood autism. *Archives of Neurology* 35, 777–786.
- Dawson, Geraldine, Karen Toth, Robert Abbott, Julie Osterling, Jeff Munson, Annette Estes & Jane Liaw. 2004. Early social attention impairments in autism: Social orienting, joint attention and attention to distress. *Developmental Psychology* 40, 271–283.
- Dawson, Michele, Isabelle Soulières, Morton A. Gernsbacher & Laurent Mottron. 2007. The level and nature of autistic intelligence. *Psychological Science* 18, 657–662.
- Dehaene, Stanislas & Jean-Pierre Changeux. 1997. A hierarchical framework for

- planning behavior. *Proceedings of the National Academy of Sciences of the United States of America* 94, 13293–13298.
- Dehaene, Stanislas & Laurent Cohen. 2007. Cultural recycling of cortical maps. *Neuron* 56, 384–398.
- Dennett, Daniel C. 1992. *Consciousness Explained*. London: Penguin.
- Deruelle, Christine, Cecile Rondan, Bruno Gepner & Carole Tardif. 2004. Spatial frequency and face processing in children with autism and Asperger syndrome. *Journal of Autism and Developmental Disorders* 34, 199–210.
- deWaal, Frans B.M. & Pier Francesco Ferrari. 2010. Towards a bottom-up perspective on animal and human cognition. *Trends in Cognitive Sciences* 14, 201–207.
- Diehl, Joshua John & Rhea Paul. 2011. Acoustic differences in the imitation of prosodic patterns in children with autism spectrum disorders. *Research in Autism Spectrum Disorders* 6, 123–134.
- Dinstein, Ilan, Karen Pierce, Lisa Eyler, Stephanie Solso, Raphael Malach, Marine Behrmann & Eric Courchesne. 2011. Disrupted neural synchronization in toddlers with autism. *Neuron* 70, 1218–1225.
- Eigsti, Inge-Marie. Executive functions in ASD. In Deborah A. Fein (ed.), *The Neuropsychology of Autism*, 185–203. Oxford: Oxford University Press.
- Embick, David, Alec Marantz, Yasushi Miyashita, Wayne O’Neil & Kuniyoshi L. Sakai. 2000. A syntactic specialization for Broca’s area. *Proceedings of the National Academy of Sciences of the United States of America* 97, 6150–6154.
- Engel, Andreas, Pascal Fries & Wolf Singer. 2001. Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience* 2, 704–716.
- Eyler, Lisa T., Karen Pierce & Eric Courchesne. 2010. Abnormal brain response to language stimuli in sleeping infants and toddlers with ASD. Paper presented at the 9th Annual International Meeting for Autism Research, Philadelphia, PA (20–22 May 2010).
- Fisher, Simon E. & Gary F. Marcus. 2006. The eloquent ape: Genes, brains and the evolution of language. *Nature Reviews Genetics* 7, 9–20.
- Fitch, W. Tecumseh. 2005. Computation and cognition: Four distinctions and their implications. In Anne Cutler (ed.), *Twenty-First Century Psycholinguistics*, 381–400. New York: Lawrence Erlbaum.
- Fitch, W. Tecumseh, Ludwig Huber & Thomas Bugnyar. 2010. Social cognition and the evolution of language: Constructing cognitive phylogenies. *Neuron* 366, 376–388.
- Flagg, Elissa J., Janis E. Oram Cardy, Wendy Roberts, Timothy P.L. Roberts. 2005. Language lateralization development in children with autism: Insights from the late field magnetoencephalogram. *Neuroscience Letters* 386, 82–87.
- Fletcher, Thomas, Ross T. Whitaker, Ran Tao, Molly B. DuBray, Alyson Froehlich, Caitlin Ravichandran, Andrew L. Alexander, Erin D. Bigler, Nicholas Lange, Janet E. Linhart. 2010. Microstructural connectivity of the arcuate fasciculus in adolescents with high-functioning autism. *NeuroImage* 51, 1117–1125.
- Forgeot d’Arc, Baudouin & Frank Ramus. 2011. Belief attribution despite verbal

- interference. *Journal of Experimental Psychology* 64, 975–990.
- De Fossé, Lies, Steven M. Hodge, Nikos Makris, David N. Kennedy, Vernes S. Caviness, Lauren McGrath, Shelley Steele, David A. Ziegler, Marthia R. Herbert, Jean A. Frazier, Helen Tager-Flusberg & Gordon J. Harris. 2004. Language-association cortex asymmetry in autism and specific language impairment. *Annals of Neurology* 56, 757–766.
- Fries, Pascal, Danko Nikolic & Wolf Singer. 2007. The gamma cycle. *Trends in Neurosciences* 30, 309–316.
- Friston, Karl. 2005. A theory of cortical responses. *Philosophical Transactions of the Royal Society* 360, 815–836.
- Frith, Uta. 1989. *Autism: Explaining the Enigma*. Hoboken: Wiley-Blackwell.
- Frith, Uta & Francesca Happé. 1994. Autism: Beyond 'Theory of Mind'. *Cognition* 50, 115–132.
- Frith, Uta & Margaret Snowling. 1983. Reading for meaning and reading for sound in autistic and dyslexic children. *British Journal of Developmental Psychology* 1, 329–342.
- Fuster, Joaquin. 2008. *The Prefrontal Cortex*, 4th edn. London: Academic Press.
- Gaffrey, Michael S., Natalia M. Kleinhans, Frank Hais, Natacha Akshoomoff, Ashley Campbell, Eric Courchesne & Ralph-Axel Müller. 2007. Atypical participation of visual cortex during word processing in autism: An fMRI study of semantic decision. *Neuropsychologia* 45, 1672–1684.
- Gage, Nicole M., Bryna Siegel, Melanie Callen & Timothy P.L. Roberts. 2003a. Cortical sound processing in children with autism disorder: An MEG investigation. *NeuroReport* 14, 2047–2051.
- Gage, Nicole M., Bryna Siegel & Timothy P.L. Roberts. 2003b. Cortical auditory system maturational abnormalities in children with autism disorder: An MEG investigation. *Developmental Brain Research* 144, 201–209.
- Gallese, Vittorio. 2008. Mirror neurons and the social nature of language: The neural exploitation hypothesis. *Social Neuroscience* 3, 317–333.
- Gallistel, Randy C.H. 2009. The foundational abstractions. In Massimo Piattelli-Palmarini, Juan Uriagereka & Pello Salaburu (eds.), *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country*, 58–73. Oxford: Oxford University Press.
- Gervais, Hélène, Pascal Belin, Natalie Boddaert, Marion Leboyer, Arnaud Coez, Ignacio Sfaello, Catherine Barthélémy, Francis Brunelle, Yves Samson & Mônica Zilbovicius. 2004. Abnormal cortical voice processing in autism. *Nature Neuroscience* 7, 801–802.
- Geschwind, Daniel H. 2008. Autism: Many genes, common pathways? *Cell* 135, 391–395.
- Geschwind, Daniel H. & Pat Levitt. 2007. Autism spectrum disorders: Developmental disconnection syndromes. *Current Opinions in Neurobiology* 17, 103–111.
- Giraud, Anne-Lise, Andreas Kleinschmidt, David Poeppel, Torben E. Lund, Richard

- Frackowiack & Helmut Laufs. 2007. Endogenous cortical rhythms determine cerebral specialization for speech perception and production. *Neuron* 56, 1127–1134.
- Giraud, Anne-Lise & David Poeppel. 2012. Cortical oscillations and speech processing: Emerging computational principles and operations. *Nature Neuroscience* 15, 511–517.
- Gleitman, Lila. 1990. The structural sources of verb meanings. *Language Acquisition* 1, 3–55.
- Goldberg, Ekholon. 2009. *The New Executive Brain: Frontal Lobes in a Complex World*. New York: Oxford University Press.
- Grandin, Temple. 1996. *Thinking in Pictures: And Other Reports from My Life with Autism*. New York: Vintage Books.
- Gray, Charles M., Peter König, Andreas K. Engel & Wolf Singer. 1989. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338, 334–337.
- Haesen, Brigitt, Bart Boets & Johan Wagemans. 2011. A review of behavioral and electrophysiological studies on auditory processing and speech perception in autism spectrum disorders. *Research in Autism Spectrum Disorders* 5, 701–714.
- Happé, Francesca G.E. 1994. Wechsler IQ profile and Theory of Mind in autism: A research note. *Journal of Child Psychology and Psychiatry* 35, 1461–1471.
- Happé, Francesca G.E. 1995. The role of age and verbal ability in the Theory of Mind task performance of subjects with autism. *Child Development* 68, 843–855.
- Happé, Francesca G.E. 1999. Autism: Cognitive deficit or cognitive style? *Trends in Cognitive Sciences* 3, 216–222.
- Happé, Francesca G.E., Rhonda Booth, Rebecca Charlton & Claire Hughes. 2006. Executive function deficits in autism spectrum disorders and attention-deficit/hyperactivity disorder: Examining profiles across domains and ages. *Brain and Cognition* 61, 25–39.
- Happé, Francesca G.E. & Uta Frith. 2006. The weak coherence account: Detail-focused cognitive style in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 36, 6–25.
- Happé, Francesca G.E., Angelica Ronald & Robert Plomin. 2006. Time to give up on a single explanation for autism. *Nature Neuroscience* 9, 1218–1220.
- Harris, Gordon J., Christopher F. Chabris, Jill Clark, Trinity Urban, Aharon Itzhak, Shelley Steele, Lauren McGrath, Karen Condouris & Helen Tager-Flusberg. 2006. Brain activation during semantic processing in autism spectrum disorders via functional magnetic resonance imaging. *Brain and Cognition* 61, 54–68.
- Hauser, Marc D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 22, 1569–1579.
- Heaton, Pamela, Beate Hermelin & Linda Pring. 1998. Autism and pitch processing: A precursor for savant musical ability? *Music Perception* 15, 291–305.
- Heaton, Pamela, Linda Pring & Beate Hermelin. 1999. A pseudo-savant: A case of exceptional musical splinter skills. *Neurocase* 5, 503–509.

- Hensch, Takao K. 2005. Critical period plasticity in local cortical circuits. *Nature Reviews Neuroscience* 6, 877–888.
- Herbert, Martha R., Gordon J. Harris, Adrien Kristen, David A. Ziegler, Nikos Makris, David N. Kennedy, Nicholas T. Lange, Christopher F. Chabris, Anna Bakardjiev, James Hodgson, Takeoka Masanori, Helen Tager-Flusberg & Verne S. Caviness. 2002. Abnormal asymmetry in language association cortex in autism. *Annals of Neurology* 52, 588–596.
- Herbert, Martha R., David A. Ziegler, Curtis Deutsch, Liam M. O'Brien, David N. Kennedy, Fillipek, P.A., Anna Bakardjiev, James Hodgson, Takeoka Masanori, Nikos Makris, Caviness & Verne S. 2004. Brain asymmetries in autism and developmental language disorders: A nested whole-brain analysis. *Brain* 128, 213–226.
- Hickok, Gregory & David Poeppel. 2007. The cortical organization of speech processing. *Nature Reviews Neuroscience* 8, 393–402.
- Hill, Elisabeth L. 2004. Executive dysfunction in autism. *Trends in Cognitive Sciences* 8, 26–32.
- Hinzen, Wolfram. 2008. Prospects for an explanatory theory of semantics. *Biolinguistics* 2, 348–363.
- Hinzen, Wolfram. 2009. Hierarchy, merge and truth. In Massimo Piattelli-Palmarini, Juan Uriagereka & Pello Salaburu (eds.), *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country*, 123–140. Oxford: Oxford University Press.
- Hoffman, Donald. 1998. *Visual Intelligence*. New York: Norton.
- Hulburt, Russ, T., Happé, Francesca G.E. & Frith, Uta. 1994. Sampling the form of inner experience in three adults with Asperger syndrome. *Psychological Medicine* 24, 385–395.
- Humphrey, Nicholas. 1998. Cave art, autism and the evolution of the human mind. *Cambridge Archaeological Journal* 8, 165–191.
- Hutsler, Jeffrey. 2003. The specialized structure of human language cortex: Pyramidal cell size asymmetries within auditory and language-associated regions of the temporal lobes. *Brain and Language* 86, 226–242.
- Hutsler, Jeffrey & Ralf A.W. Galluske. 2003. Hemispheric asymmetries in cerebral cortical networks. *Trends in Neuroscience* 36, 429–435.
- Indefrey, Peter & Willem J. M. Levelt. 2004. The spatial and temporal signatures of word production components. *Cognition* 92, 101–144.
- Jackendoff, Ray. 1983. *Semantics and Cognition*. Cambridge, MA: MIT Press.
- Jansson-Verkasalo, Eira, Rita Ceponienè, Marko Kielinen, Kalervo Suominen, Ville Jäntti, Sirkka-Liisa Linna, Irma Moilanen & Risto Näätänen. 2003. Deficient auditory processing in children with Asperger syndrome, as indexed by event-related potentials. *Neuroscience Letters* 338, 197–200.
- Järvinen-Palsey, Anna, Gregory L. Wallace, Franck Ramus, Francesca Happé & Pamela Heaton. 2008. Enhanced perceptual processing of speech in autism. *Developmental Science* 11, 109–121.

- Jenkins, Lyle. 2000. *Biolinguistics: Exploring the Biology of Language*. Cambridge: Cambridge University Press.
- Jolliffe, Therese & Simon Baron-Cohen. 1997. Are people with autism and Asperger syndrome faster than normal on the Embedded Figure Test. *Journal of Child Psychology and Psychiatry* 38, 527–534.
- Joseph, Robert M., Helen Tager-Flusberg & Catherine Lord. 2002. Cognitive profiles and social-communicative functioning in children with autism spectrum disorder. *Journal of Child Psychology and Psychiatry* 43, 807–821.
- Jou, Roger J., Nancy J. Minshew, Matcheri, S. Keshavan, Matthew P. Vitale & Antonio Y. Hardan. 2010. Enlarged right superior temporal gyrus in children and adolescent with autism. *Brain Research* 1360, 205–210.
- Just, Marcel A., Vladimir L. Cherkassky, Timothy A. Keller, Rajesh K. Kana & Nancy J. Minshew. 2005. Functional and anatomical cortical underconnectivity in autism: Evidence from an fMRI Study of an executive function task and corpus callosum morphometry. *Cerebral Cortex* 17, 951–961.
- Just Marcel A. Vladimir L. Cherkassky, Timothy A. Keller & Nancy J. Minshew. 2004b. Cortical activation and synchronization during sentence comprehension in high-functioning autism: evidence of underconnectivity. *Brain* 127, 1811–1821.
- Just, Marcel A., Sharlene D. Newman, Timothy A. Keller, Alice McEleney & Patricia A. Carpenter. 2004a. Imagery in sentence comprehension: An fMRI study. *NeuroImage* 21, 112–124.
- Kamio Yoko & Motomi Toichi. 2000. Dual access to semantics in autism: Is pictorial access superior to verbal access? *Journal of Child Psychology and Psychiatry* 41, 859–867.
- Kana, Rajesh K., Timothy Keller, Vladimir L. Cherkassky, Nancy J. Minshew & Just, Marcel A. 2006. Sentence comprehension in autism: Thinking in pictures with decreased functional connectivity. *Brain* 129, 2484–2493.
- Kemper, Thomas L. & Margaret L. Bauman. 1998. Neuropathology in infantile autism. *Journal of Neuropathology and Experimental Neurology* 57, 645–652.
- Kjelgaard, Margaret M. & Helen Tager-Flusberg. 2001. An investigation of language impairment in autism: Implications for genetic subgroups. *Language and Cognitive Processes* 16, 287–308.
- Koechlin, Etienne, Chrystèle Ody & Frédérique Kouneiher. 2003. The architecture of cognitive control in the human prefrontal cortex. *Science* 302, 1181–1185.
- Koechlin, Etienne & Thomas Jubault. 2006. Broca's area and the hierarchical organization of human behavior. *Neuron* 50, 963–974.
- Konishi, Seiki, Kyochi Nakajima, Idai Uchida, Masashi Kamevama, Kiyoshi Nakahara, Kensuke Sekihara & Yasushi Miyashita. 1998. Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nature Neuroscience* 1, 89–94.
- Koshino, Hideya, Patricia A. Carpenter, Nancy J. Minshew, Vladimir L. Cherkassky, Timothy A. Keller & Marcel A. Just. 2005. Functional connectivity in an fMRI

- working memory task in high-functioning autism. *NeuroImage* 24, 810–821.
- Kunda, Maithilee & Ashok K. Goel. 2011. Thinking in pictures as a cognitive account of autism. *Journal of Autism and Developmental Disorders* 41, 1157–1177.
- Landa, Rebecca. 2000. Social Language use in Asperger syndrome and high-functioning autism. In Ami Klin, Fred R. Volkmar & Sara S. Sparrow (eds.), *Asperger Syndrome*, 159–171. New York: The Guilford Press.
- Lepistö, Tuulia, Marika Kajander, Raija Vanhala, Paavo Alku, Minna Huotilainen, Risto Näätänen & Teija Kujala. 2008. The perception of invariant speech features in children with autism. *Biological Psychology* 77, 25–31.
- Lepistö, Tuulia, Teija Kujala, Raija Vanhala, Paavo Alku, Minna Huotilainen & Risto Näätänen. 2005. The discrimination of and orienting to speech and non-speech sounds in children with autism. *Brain Research* 1066, 147–157.
- Levitin, Daniel J. & Anna K. Tirovolas. 2009. Current advances in the cognitive neuroscience of music. *Annals of the New York Academy of Sciences* 1156, 211–231.
- Lewis, Vicky & Jill Boucher. 1995. Generativity in the play of young people with autism. *Journal of Autism and Developmental Disorders* 25, 105–121.
- Lopez, Beatriz & Susan R. Leekam. 2003. Do children with autism fail to process information in context? *Journal of Child Psychology and Psychiatry* 44, 285–300.
- Luo, Huan, Zuxiang Liu & David Poeppel. 2010. Auditory cortex tracks both auditory and visual stimulus dynamics using low-frequency neuronal phase modulation. *PLoS Biology* 8, e1000445.
- Luo, Huan & David Poeppel. 2007. Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron* 54, 1001–1010.
- Marcus, Gary F. & Hugh Rabagliati. 2006. What developmental disorders can tell us about the nature and origins of language. *Nature Neuroscience* 9, 1226–1229.
- Markram, Henry, Tania Rinaldi & Kamila Markram. 2007. The Intense World Syndrome — an alternative hypothesis for autism. *Frontiers in Neuroscience* 1, 77–96.
- McAlonan, Grainne M., Vinci Cheung, Charlotte Cheung, John Suckling, Lam, Grace Y., K.S. Tai, L. Yip, Declan G.M. Murphy & Siew E. Chua. 2005. Mapping the brain in autism. A voxel-based MRI study of volumetric differences and inter-correlations in autism. *Brain* 128, 268–276.
- McCleery, Joseph, Rita Ceponienè, Karen M. Burner, Jeanne Townsend, Mikaela Kinnear & Laura Schreibman. 2010. Neural correlates of verbal and nonverbal semantic integration in children with autism spectrum disorders. *Journal of Child Psychology and Psychiatry* 51, 277–286.
- McGurk, Harry & John McDonald. 1976. Hearing lips and seeing voices. *Nature* 264, 746–748.
- Miller, Leon K. 1989. *Musical Savants: Exceptional Skills in the Mentally Retarded*. Hillsdale, NJ: Erlbaum.
- Minschew, Nancy J., Gerald Goldstein & Don J. Siegel. 1997. Neuropsychologic functioning in autism: Profile of a complex information processing disorder. *Journal*

- of the *International Neuropsychological Society* 3, 303–316.
- Mongillo, Elizabeth A., Julia R. Irwin, Doug H. Whalen, Cheryl Klaiman, Alice S. Carter & Robert T. Schultz. 2008. Audiovisual processing in children with and without Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 38, 1349–1358.
- Mottron, Laurent. 2003. *L'autisme: Une autre intelligence*. Bruxelles: Mardagas.
- Mottron, Laurent & Sylvie Belleville. 1993. A study of perceptual analysis in a high-level autistic subject with exceptional graphic abilities. *Brain and Cognition* 23, 279–309.
- Mottron Laurent, Michèle Dawson, Isabelle Soulières, Bénédicte Hubert & Jake Burack. 2006. Enhanced perceptual functioning in autism: An update and eight principles of autistic perception. *Journal of Autism and Developmental Disorders* 36, 23–47.
- Mottron, Laurent, Karine Morasse & Sylvie Belleville. 2001. A study of memory functioning in individuals with autism. *Journal of Child Psychology and Psychiatry* 42, 253–260.
- Mottron, Laurent, Isabelle Peretz & Élise Ménard. 2000. Local and global processing of music in high-functioning persons with autism: Beyond central coherence? *Journal of Child Psychology and Psychiatry* 41, 1057–1065.
- Mountcastle, Vernon B. 1997. The columnar organization of the neocortex. *Brain* 120, 701–722.
- Müller, Ralph-Axel, M.E. Behen, R.D. Rothermel, D.C. Chugani, O. Musik, T.J. Mangner, H.T. Chugani. 1999. Brain mapping of language and auditory perception in high-functioning autistic adults: A PET study. *Journal of Autism and Developmental Disorders* 29, 19–31.
- Murias, Michael, Sara J. Webb, Jessica Greenson & Geraldine Dawson. 2007. Resting state cortical connectivity reflected in EEG coherence in individuals with autism. *Biological Psychiatry* 62, 270–273.
- Musso, Mariacristina, Andrea Moro, Volkmar Glauche, Michel Rijntjes, Jürgen Reichenbach, Christian Büchel & Cornelius Weiller. 2003. Broca's area and the language instinct. *Nature Neuroscience* 6, 774–780.
- Näätänen, Risto & Anthony W.K. Gaillard & Sirkka Mäntysalo. 1978. Early selective attention effect on evoked potential reinterpreted. *Acta Psychologica* 42, 313–329.
- Näätänen, Risto. 1990. The role of attention in auditory information processing as revealed by event-related potentials and other brain measures of cognitive function. *Behavioral and Brain Sciences* 13, 201–288.
- Nadig, Aparna. 2011. Use of sentence context for lexical ambiguity resolution in high functioning autism. Paper presented at the *12th International Congress for the Study of Child Language*, Montreal (19–23 July 2011).
- Nadig, Aparna & Holly Shaw. 2012. Acoustic and perceptual measurement of expressive prosody in high-functioning autism: Increased pitch range and what it means to listeners. *Journal of Autism and Developmental Disorders* 42, 499–511.

- Nation, Kate, Paula Clarke, Barry Wright & Christine Williams. 2006. Patterns of reading ability in children with Autism Spectrum Disorder. *Journal of Autism and Developmental Disorders* 36, 911–919.
- Navon, David. 1977. Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology* 9, 353–383.
- Newbury, Dianne F., E. Bonora, J. Lamb, S.E Fisher, C.S.L. Lai, G. Baird, L. Jannoun, V. Slonims, C.M. Stott, M.J. Merricks, P.F. Bolton, A.J. Bailey, A.P. Monaco & the International Molecular Genetic Study of Autism Consortium. 2002. *FOXP2* is not a major susceptibility gene for autism or Specific Language Impairment. *American Journal of Human Genetics* 70, 1318–1327.
- Newschaffer Craig J., Lisa A. Croen, Julie Daniels, Ellen Giarelli, Judith K. Grether, Susan E. Levy, David S. Mandell, Lisa A. Miller, Jennifer Pinto-Martin, Judy Reaven, Ann M. Reynolds, Catherine E. Rice, Diana Schendel & Gayle Windham. 2007. The epidemiology of Autism Spectrum Disorders. *Annual Reviews of Public Health* 28, 21–24.
- Newton, Ashley & Jill G. deVilliers. 2007. Thinking while talking: Adults fail false belief attribution. *Psychological Science* 18, 574–579.
- Norbury, Courtenay F. 2005. Barking up the wrong tree? Lexical ambiguity resolution in children with language impairments and autistic spectrum disorders. *Journal of Experimental Child Psychology* 90, 142–171.
- Ozonoff, Sally & Jenise Jensen. 1999. Specific executive function profiles in three neurodevelopmental disorders. *Journal of Autism and Developmental Disorders* 29, 171–177.
- Ozonoff, Sally & Ryan E. McEvoy. 1994. A longitudinal study of executive function and theory of mind development in autism. *Development and Psychopathology* 6, 415–431.
- Ozonoff, Sally, Bruce F. Pennington & Sally Rogers. 1991. Executive function deficits in high-functioning autistic individuals: Relationship to Theory of Mind. *Journal of Child Psychology and Psychiatry* 32, 1081–1105.
- Paul, Rhea, Amy Augustyn, Ami Klin & Fred R. Volkmar. 2005. Perception and production of prosody by speakers with Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 35, 205–220.
- Peppé, Susan, Joanne McCann, Fiona Gibbon, Anne O'Hare & Mel Rutherford. 2007. Receptive and expressive prosodic ability in school-age children with high-functioning autism. *Journal of Speech, Language and Hearing Research* 50, 1015–1028.
- Perovic, Alexandra, Nadia Modyanova & Kenneth Wexler. 2007. Knowledge of c-command and A-movement in children and adolescents with autism and with Asperger syndrome. Paper presented at *GALA 2007: Generative Approaches to Language Acquisition*, Barcelona (6–8 September 2007).
- Pierce, Karen, Ralph-Axel Müller, J. Ambrose, Greg Allen, Eric Courchesne. 2001. Face processing occurs outside the fusiform 'face area' in autism: Evidence from functional MRI. *Brain* 124, 2053–2073.

- Postle, Bradley R., Jeffrey S. Berger, Mark D'Esposito. 1999. Functional neuroanatomical double dissociation of mnemonic and executive control processes contributing to working memory performance. *Proceedings of the National Academy of Sciences of the United States of America* 96, 12959–12964.
- Rao, Rajesh P. N. & Dana H. Ballard. 1999. Predictive coding in the visual cortex: A functional interpretation of some extra-classical fields effects. *Nature Neuroscience* 2, 77–87.
- Raymond, Gerald V., Margaret L. Bauman & Thomas L. Kemper. 1996. Hippocampus in autism: A Golgi analysis. *Acta Neuropathologica* 91, 117–119.
- Redcay, Elizabeth, Eric Courchesne. 2008. Deviant functional magnetic resonance imaging patterns of brain activity to speech in 2-3-year-old-children with Autism Spectrum Disorders. *Biological Psychiatry* 64, 589–598.
- Reuland, Eric. 2001. Primitives of binding. *Linguistic Inquiry* 32, 439–492.
- Ring, Howard A., Simon Baron-Cohen, Sally Wheelwright, Steve C.R. Williams, Mick Brammer, Chris Andrew & Edward T. Bullmore. 1999. Cerebral correlates of preserved cognitive skills in autism. *Brain* 122, 1305–1315.
- Roberts, Timothy P.L., Sarah Y. Khan, Michael M. Rey, Justin F. Monroe, Katelyn Cannon, Lisa Blaskey, Sarah Woldoff, Saba Quasmieh, Mike Gandal, Gwen L. Schmidt, Deborah Zarnow, Susan E. Levy & Christopher J. Edgar. 2010. MEG detection of delayed auditory evoked responses in autism spectrum disorders: Towards an imaging biomarker for autism. *Autism Research* 3, 8–18.
- Rojas, D. C., S. D. Bawn, T. L. Benkers, M. L. Reite & S. J. Rogers. 2002. Smaller left hemisphere planum temporale in adults with autistic disorder. *Neuroscience Letters* 328, 237–240.
- Rojas, Donald C., Suzanne L. Camou, Martine L. Reite & Sally J. Rogers. 2005. Planum temporale volume in children and adolescents with autism. *Journal of Autism and Developmental Disorders* 35, 479–486.
- Russo, Natalie, Tara Flanagan, Grace Iarocci, Darlene Berringer, Philip D. Zelazo & Jacob A. Burack. 2007. Deconstructing executive deficits among persons with autism: Implications for cognitive neuroscience. *Brain and Cognition* 65, 77–86.
- Russo, Nicole, Charles Larson & Nina Krauss. 2008. Audio-vocal system regulation in children with autism spectrum disorders. *Experimental Brain Research* 188, 111–124.
- Sahyoun, Chérif P., John W. Belliveau, Isabelle Soulières, Shira Schwartz & Maria Mody. 2010. Neuroimaging of the functional and structural networks underlying visuospatial vs. linguistic reasoning in high-functional autism. *Neuropsychologia* 48, 86–95.
- Samson, Fabienne, Laurent Mottron, Jemel Boutheina, Pascal Belin, Valter Ciocca. 2006. Can spectro-temporal complexity explain the autistic patterns of performance on auditory tasks? *Journal of Autism and Developmental Disorders* 36, 65–75.
- Selge, Lorna. 1995. *Nadia Revisited: A Longitudinal Study of an Autistic Savant*. New York: Taylor & Francis.

- Senkowski, Daniel, Till R. Schneider, John J. Foxe & Andreas K. Engel. 2008. Cross-modal binding through neural coherence: Implication for multisensory processing. *Trends in Neuroscience* 31, 401–409.
- Shah, Amita & Uta Frith. 1983. An islet of ability in autistic children: A research note. *Journal of Child Psychology and Psychiatry* 24, 613–620.
- Shallice, Tim. 1982. Specific impairments of planning. *Philosophical Transaction of the Royal Society of London* 298, 199–209.
- Shallice, Tim. 1988. *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Shriberg, Lawrence, Rhea Paul, Jane L. McSweeny, Ami Klin, Donald J. Cohen & Fred R. Volkmar. 2001. Speech and prosody characteristics of adolescents and adults with high-functioning autism and Asperger syndrome. *Journal of Speech, Language and Hearing Research* 44, 1097–1115.
- Siegal, Michael & Mark Blades. 2003. Language and auditory processing in autism. *Trends in Cognitive Sciences* 7, 378–380.
- Smith, Edward E & John Jonides. 1999. Storage and executive processes in frontal lobes. *Science* 283, 1657–1661.
- Soulières, Isabelle, Michelle Dawson, Fabienne Samson, Elise B. Barbeau, Cherif P. Sahyoun, Gary E. Strangman, Thomas A. Zeffiro & Laurent Mottron. 2009. Enhanced visual processing contributes to matrix reasoning in autism. *Human Brain Mapping* 30, 4082–4107.
- Steiman, Mandy, Rebecca Simon, Lisa Reisinger & Eric Fombonne. 2011. Trends in autism rates — Is there an epidemic? In M. Elena Garralda & Jean-Philippe Raynaud (eds.), *Increasing Awareness of Child and Adolescent Mental Health*, 162–192. Lanham: Jason Aronson.
- Stephan, Klaas E., Gereon R. Fink & John C. Marshall. 2007. Mechanisms of hemispheric specialization: Insights from analyses of connectivity. *Neuropsychologia* 45, 209–228.
- Stuss, Donald T., Gordon G. Gallup Jr. & Michael P. Alexander. 2001. The frontal lobes are necessary for 'Theory of Mind'. *Brain* 124, 279–286.
- Surian, Luca, Simon Baron-Cohen & Heather van der Lely. 1996. Are children with autism deaf to Gricean Maxims? *Cognitive Neuropsychiatry* 1, 55–71.
- Tager-Flusberg, Helen. 1992. Autistic children's talk about psychological states: Deficits in the early acquisition of a Theory of Mind. *Child Development* 63, 161–172.
- Tager-Flusberg, Helen, Susan Calkins, Tina Nolin, Therese Baumberger, Marcia Anderson & Ann Chadwick-Dias. 1990. A longitudinal study of language acquisition in autistic and Down syndrome children. *Journal of Autism and Developmental Disorders* 20, 1–21.
- Tager-Flusberg, Helen & Robert M. Joseph. 2003. Identifying neurocognitive phenotypes in autism. *Philosophical Transaction of the Royal Society* 358, 303–314.
- Tager-Flusberg, Helen & Robert M. Joseph. 2005. How language facilitates the acquisition of false-belief understanding in children with autism. In Janet

- Wilde & Jodie A. Baird (eds.), *Why Language Matters for Theory of Mind*, 298–318. New York: Oxford University Press.
- Tager-Flusberg, Helen, Rhea Paul & Catherine Lord. 2005. Language and communication in autism. In Fred R. Volkmar, Rhea Paul, Ami Klin & Donald L. Cohen (eds.), *Handbook of Autism and Pervasive Developmental Disorders*, 3rd edn., 335–364. Hoboken, NJ: Wiley & Sons.
- Tattersall, Ian. 1998. *Becoming Human: Evolution of Human Uniqueness*. New York: Harcourt Brace.
- Tattersall, Ian. 2004. What happened in the origin of human consciousness? *The Anatomical Records* 276B, 19–26.
- Toichi, Motomi & Yoko Kamio. 2002. Long-term memory and levels of processing in autism. *Neuropsychologia* 40, 964–969.
- Turner, Michelle A. 1999. Generating novel ideas: Fluency performance in high-functioning and learning disabled individuals with autism. *Journal of Child Psychology and Psychiatry* 40, 189–201.
- Ullman, Michael T. 2004. Contributions of memory circuits to language: The declarative/procedural model. *Cognition* 92, 231–270.
- Vernes Sonja C., Dianne F. Newbury, Brett S. Abrahams, Laura Winchester, Jérôme Nicod, Matthias Groszer, Marciela Alarcon, Peter L. Olivier, Kay E. Davies, Daniel H. Geschwind, Anthony P. Monaco & Simon E. Fisher. 2008. A functional genetic link between distinct developmental language disorders. *The New England Journal of Medicine* 359, 2337–2345.
- Vertes, Robert P. 2005. Hippocampal theta rhythm: A tag for short-term memory. *Hippocampus* 15, 923–935.
- de Villiers, Jill G. & Jennie E. Pyers. 2002. Complements to cognition: A longitudinal study of the relationship between complex syntax and false-belief-understanding. *Cognitive Development* 17, 1037–1060.
- Volden, Joanne & Catherine Lord. 1991. Neologisms and idiosyncratic language in autistic speakers. *Journal of Autism and Developmental Disorders* 21, 109–130.
- Volkmar, Fred R. & Ami Klin. 2005. Issues in the classification of autism and related conditions. In Fred R. Volkmar, Rhea Paul, Ami Klin & Donald L. Cohen (eds.), *Handbook of Autism and Pervasive Developmental Disorders*, 3rd edn., 5–41. Hoboken, NJ: Wiley & Sons.
- Vygotsky, Lev. 1962. *Thought and Language*. Cambridge, MA: MIT Press.
- Walenski, Matthew, Stewart H. Mostofsky, Jennifer C. Gidley-Larson & Michael T. Ullman. 2008. Brief report: Enhanced picture naming in autism. *Journal of Autism and Developmental Disorders* 38, 1395–1399.
- Walsh, Christopher A., Eric M. Morrow & John L. R. Rubenstein. 2008. Autism and brain development. *Cell* 135, 396–400.
- van Wassenhove, Virginie, Ken Grant & David Poeppel. 2005. Visual speech speeds up the neural processing of auditory speech. *Proceedings of the National Academy of Sciences of the United States of America* 102, 1181–1186.
- White, Sarah J. & David Saldaña. 2011. Performance of children with autism on the

- embedded figure test: A closer look at a popular task. *Journal of Autism and Developmental Disorders* 41, 1565–1572.
- Whitehouse, Andrew J.O. & Dorothy V.M. Bishop. 2008. Cerebral dominance for language function in adults with specific language impairment or autism. *Brain* 131, 3293–3200.
- Whitehouse, Andrew J.O., Johanna G. Barry & Dorothy V.M. Bishop. 2007. The broader language phenotype of autism: A comparison with specific language impairment. *Journal of Child Psychology and Psychiatry* 48, 822–830.
- Williams, David, Francesca G.E. Happé & Christopher Jarrold. 2008. Intact inner speech use in autism spectrum disorder: Evidence from a short-term memory task. *Journal of Child Psychology and Psychiatry* 49, 51–58.
- Williams, Diane L., Gerald Goldstein, Patricia A. Carpenter & Nancy J. Minshev. 2005. Verbal and spatial working memory in autism. *Journal of Autism and Developmental Disorders* 35, 747–756.
- Williams, Emily L. & Manuel F. Casanova. 2010. Autism and dyslexia: A spectrum of cognitive styles as defined by minicolumnar morphometry. *Medical Hypotheses* 74, 59–62.
- Williams, Justin H. G., Dominic W. Massaro, Natalie J. Peel, Alexis Bosseler & Thomas Suddendorf. 2004. Visual-auditory integration during speech imitation in autism. *Research in Developmental Disabilities* 25, 559–575.
- Wilson, Tony W., Donald C. Rojas, Martin L. Reite, Peter D. Teale & Sally J. Rogers. 2007. Children and adolescents with autism exhibit reduced MEG steady-state gamma responses. *Biological Psychiatry* 62, 192–197.
- Wollman, David. 2008. Yeah, I'm autistic. You got a problem with that? *Wired Magazine* (16 March 2008).

Nicolas Bourguignon & Daniel Valois
University of Montreal
Department of Linguistics
C.P. 6128, succursale Centre-ville
Montreal, QC H3C 3J7
Canada

Aparna S. Nadig
McGill University
School of Communication Sciences & Disorders
1266, Pine Av. West
Montreal, QC H3G 1A8
Canada

nicolas.bourguignon@umontreal.ca, daniel.valois@umontreal.ca, aparna.nadig@mcgill.ca

Ever Since Dennett: On the Origins of Biolinguistics

Hans-Martin Gärtner

In their ‘*Biolinguistics Manifesto*’, Boeckx & Grohmann (2007: 3) rightly point out “that the recent resurgence of interest in ‘biolinguistics’ is due in large part to the advent of the minimalist program in linguistic theory”. Various reasons have been given for the necessity of moving from (some kind of) GB-style grammar research toward minimalism, some more conceptual (cf. e.g. Chomsky 2007: 19), some more empirical (cf. e.g. Holmberg 2000). However, arguably one of the motivating sources neglected so far is a remark by Daniel Dennett, which this very brief note is meant to bring to everyone’s attention (again).

In reflecting upon the explanatory burden put on UG by Chomsky (1980b, 1980c), Dennett gives vent to an uneasy feeling about “passing the buck to biology”. He therefore — constructively, I believe (cf. Dennett 1995: 388) — “challenges Chomsky” as follows:

Perhaps no one supposes there is a larger innate contribution than Chomsky does, and perhaps the facts will eventually bear out a position close to his, but his polemics sometimes ignore the perfectly reasonable motivation behind the contrary perspective — what we might call the *minimalist* research strategy. (Dennett 1980: 19)

Given the context of the early GB era, Chomsky (1980a: 44) defended his position as the most promising way of “developing what Dennett calls a ‘realistic picture’ of the basis in innate endowment for cognitive growth”. Thirty years on, however, it seems that some credit is due to the challenger. Thus, ironically, although Dennett’s (neo-)Darwinian adaptationist outlook may be considered incompatible with (certain strands of) biolinguistics (cf. Hinzen 2006), his visionary postulation of a “minimalist research strategy” could well be taken to have contributed a pebble — be it ever so small — to the biolinguistic edifice. *Honni soit qui mal y pense.*

References

- Boeckx, Cedric & Kleanthes K. Grohmann. 2007. The *Biolinguistics* manifesto. *Biolinguistics* 1, 1–8.
- Chomsky, Noam. 1980a. The new organology. *Behavioral and Brain Sciences* 3, 42–61.
- Chomsky, Noam. 1980b. *Rules and Representations*. Oxford: Blackwell.
- Chomsky, Noam. 1980c. Rules and Representations. *Behavioral and Brain Sciences* 3, 1–61.



- Chomsky, Noam. 2007. Of minds and language. *Biolinguistics* 1, 9–27.
- Dennett, Daniel. 1980. Passing the buck to biology. *Behavioral and Brain Sciences* 3, 19.
- Dennett, Daniel. 1995. *Darwin's Dangerous Idea*. London: Allen Lane.
- Hinzen, Wolfram. 2006. *Mind Design and Minimalist Syntax*. Oxford: Oxford University Press.
- Holmberg, Anders. 2000. Am I unscientific? A reply to Lappin, Levine, and Johnson. *Natural Language and Linguistic Theory* 18, 837–842.

Hans-Martin Gärtner
Hungarian Academy of Sciences
Research Institute for Linguistics
Benczúr u. 33
1068 Budapest
Hungary
gaertner@nytud.hu

Lights and Shadows in the Evolution of Language

Hurford, James R. 2007. *The Origins of Meaning: Language in the Light of Evolution I*. New York: Oxford University Press.

Hurford, James R. 2012. *The Origins of Grammar: Language in the Light of Evolution II*. New York: Oxford University Press.

by Evelina Leivada & Ana M. Suárez

James Hurford's *Language in the Light of Evolution* two-volume project aims to bring together an up-to-date account of the synthesis of language in the light of evolution, going from meaningful mental representations, widespread in the animal kingdom, to the emergence of the first words and their grammatical combination. The then upcoming successor of volume I alongside with its contents is already sketched out in the preface of the first volume. The two parts are organized to cover different aspects of human language and its precursors, although the length of the second volume makes some degree of overlap unavoidable. Volume I deals with the content of meaning (i.e. semantics) and its interpersonal use (i.e. pragmatics), while volume II focuses on core notions of grammar and discusses the ins and outs of the evolution of language in a three-step travel: a first shared lexicon, a two-word stage, and grammaticalization procedures. Albeit their different objects of study, the two volumes complement each other and are unified under an evolutionary approach. It is worth mentioning for the sake of completeness that the duology could be a trilogy as Hurford notes in the preface of the volume II. Having covered the origins of meaning in volume I and the origins of grammar in volume II, the origins of speech (i.e. phonetics and phonology) should be addressed next. However, as Hurford notes, such accounts exist in the literature; among the most recent ones is Fitch (2010).

Starting off with *The Origins of Meaning*, the standard gradualist (neo-) Darwinian perspective through which Hurford approaches the topic of language origins is made explicit from the beginning. This contrasts with the relative lack of clarity when it comes to the linguistic framework against which the discussion unfolds — contrary to what happens in volume II, where Minimalism together with other frameworks (e.g., Construction Grammar, Formal Language Theory) is often put under (at times, comparative) scrutiny. Hurford frequently refers to

We are grateful to Cedric Boeckx and Luis Eguren for valuable comments and suggestions. We also acknowledge support through grant FFI2010-20634 from the Spanish Ministerio de Ciencia e Innovación and a scholarship from the Cyprus State Scholarship Foundation (Leivada) and through grants EDU2008-01268 and FFI2011-23829 from the Spanish Ministerio de Ciencia e Innovación (Suárez).



Tomasello's and Call's work in volume I; a fact that led to the claim that his linguistic stance is probably cognitivist (Edwardes 2009: 191). The Minimalist Program is not brought into the discussion up until chapter 8, and when this happens, the passage from pre-Minimalism days (before Chomsky 1995) to what Hurford refers to as the "current Chomskyan position [...] that the only possibly distinctive property of the human language faculty in the narrow sense (i.e. not shared with either animal communication or with human non-linguistic cognition) is recursion" (p. 285) is mentioned only in passing.¹

In the first part of volume I, Hurford presents the elements that comprise animal thought of the world when thought and perception was yet untrimmed by the need to communicate as humans do. Special emphasis is laid on defining concepts, also in terms of an evolutionary continuum: Proto-concepts were succeeded by pre-linguistic concepts which were followed by linguistic concepts (p. 12). Hurford's understanding of the term 'concept' is partly in line with Fodor's (1998) treatment of concepts as indicating subparts of states of mind. It is argued that Fodor sets five conditions for having concepts and Hurford embraces viewing concepts as (i) mental particulars, (ii) categories, (iii) compositional, and (iv) often learned. He claims that he parts company with Fodor only on the fifth requirement: (v) that concepts are public; "they're the sorts of things that lots of people can, and do, *share*" (Fodor 1998: 28). This deviance on the fifth condition is the result of Hurford talking about animal, pre-linguistic concepts that exist before communicative and social needs arise, whereas Fodor describes human, post-linguistic concepts. However there is some (yet not well known) degree of difference between what is ascribed to Fodor and what Fodor has really claimed. Therefore, Fodor and Hurford part company in point (iv) already, since Fodor does not assume concepts to be learned.

The first volume is a very welcome contribution in that it brings together various experimental findings that pertain to the evolution of animal (i.e. pre-linguistic) cognition. It surveys many aspects of non-human cognitive functions in a range of primates, but also other (marine) mammals and birds. It also touches upon a recently-evolved domain of cognition that has been argued to be unique to humans (Tulving 1999): episodic memory. Hurford rightly juxtaposes the argument for episodic memory being a uniquely human trait to experimental findings coming from a variety of (food-storing) species such as squirrels, American scrub jays, honeybees, and great apes. The conclusion he draws is that there is evidence from experiments with animal food-storing and neurophysiological studies of rat dreaming, for assuming a kind of episodic memory in non-humans that is less domain-general than the one humans have, but still one that could be viewed as a seed from which the human capacity evolved (p. 83).²

¹ Notice here that this is not the then current Chomskyan position (it is, in fact, Hauser *et al.*'s 2002) to the extent that Fitch *et al.* (2005), to whom Hurford makes no reference in the first volume, are more perspicuous than Hauser *et al.* on the possibility that the contents of the faculty of language in the narrow sense (FLN) are subject to empirical determination and FLN might turn out to be empty, resulting in a claim that only "language as a whole is unique to our species" (Fitch *et al.* 2005: 181).

² However, human episodic memory, although better than the one observed in other species, can be proven fairly poor. Hurford cites the case of eyewitness testimony (Wells & Olson 2003); more recent experiments suggest that poor performance might be attributed to episodic memory undergoing a reconsolidation process after recalling an event (Chan *et al.* 2009).

Discussing proto-propositions, Hurford argues against a drastic jump in the course of evolution and contrasts his view that animals are capable of having a proposition-like cognition with Dummett's (1993) view that animals show proto-thought but are not in a position to entertain propositions. Discussing numerical limits on the size of propositions (an issue re-introduced in volume II), Chomsky's distinction between competence and performance is introduced and an argument is made for memory limitations on the size of simple propositions being consistent with the generative stance that a predicate could potentially take an unlimited number of arguments. In this context, the limit of maximum four arguments per predicate is viewed as a matter of performance, since "U[niversal] G[rammar] imposes no constraint on the number of arguments a predicate may take" (p. 90). However first, performance-imposed restrictions are not necessarily UG-derived constraints and second, the 'maximum four arguments per predicate' constraint is not absolute. Consider (1), for example:

- (1) X hit *y* with *z* at place *p* and time *t*.

In the second part of the first book, the focus is on communication which translates into interactions between animals of the same species, since when animals communicate with animals belonging to different species (e.g., the example of a shepherd whistling commands to a sheepdog that Hurford lists), this communication is not reciprocal in the sense that even if there is a response, this will be in a different code. Communicative acts in the systems of various primates are brought together in Hurford's review of a broad amount of literature in animal communication research which does justice to most current findings that primatology and biology report. Of course, evolution of such devices only makes sense if animals are biologically disposed to share information with their peers by communicating it to each other. The origins of this need to communicate are sketched out in relation to the phenomenon of niche-construction in biology — a correlation more developed in the last part of the second volume — which is argued to have facilitated rapid changes that in turn gave rise to new domains of social interaction and co-operation. The answer Hurford gives to the crucial question of why such a principle came to exist boils down to the aspect of the Darwinian theory that emphasizes the importance of selected traits being of some benefit to the individuals that make use of them. Hurford concludes the second part of this first volume by admitting that no single theory (such as Kin Selection or Sexual Selection) can on its own "adequately explain the unique human characteristic of freely giving information in such structurally complex ways as we do every day with language" (p. 333).

Maintaining the perspective on the "no fundamental difference" (Darwin 1871: 35) between non-humans' and humans' mental life, Hurford deals in *The Origins of Grammar* with the core components of language and the consequences of its breakthrough in communicating thoughts, putting in parallel alike communicative traits between both genera. Within the domain of what he calls pre-grammar, where he meticulously describes the disposition of rhythmic patterns in animal songs, he does not go past some cognitive factors, such as memory limitations constraining communication. He presents several striking

constants in the messages of other species, which resemble constraints one observes in human language: parallel chaffinch song has a median of seven apparent ‘phrases’; bird songs are roughly the same length as typical spoken human sentences, between one and ten seconds. If humans also possess a limit on the *selected* arguments a predicate can take (‘the magical number 4’) — a fact to which he gave, probably for the first time in (bio)linguistic research, a psychological explanation (i.e. ‘subitization’, similar in monkeys and humans) — these computational limitations, Hurford argues, should not be bypassed in the acquisition and development of language; that is, they should not be regarded as a matter of performance, according to generative thinking, but comprised — from the very beginning — into what Chomsky (1965 *et seq.*) called ‘competence’. Departing, thus, from the mainstream generative view, Hurford introduces the term ‘competence-plus’ with which he spans grammatical specifications (what *competence* stands for) and numerical constraints (the *plus*) into one: “[N]o organism learns or acquires competence immune from the quantitative constraints of its body” (p. 56); “the interaction happens not after competence has been formed, but while competence is being formed” (p. 247).³ According to his view, the mechanisms constraining language processing are central to the human language faculty (UG+), and, accordingly, the human language faculty should not keep being considered “unaffected by such grammatically irrelevant conditions as memory limitations” (Chomsky 1965: 3) anymore. A different question is whether this *re-orientation* (the *plus*) on such linguistic presupposition (*competence*) will be included at some point in future linguistic research — and how so.

When it comes to the comparative issue of cognitive differences, the conviction with which in *The Origins of Meaning* non-human animals were endowed with concepts seems now diminished in *The Origins of Grammar* by an important factor: the enhancement of Thought. Hurford acknowledges that “those who deny that animals can have full concepts do have a point [...]. There is a difference between pre-linguistic concepts, or proto-concepts, such as I have *freely postulated* in the earlier book, and fully-fledged human concepts” (p. 154, emphasis ours). The genuine ‘feedback loop’ which enriches in tandem language and thought proves to be a crucial reason for the terminological distinction. Hurford alludes to Sapir’s (1921) and Bickerton’s (2009) claims on the reciprocity between the ‘instrument’ and the ‘product’ or the ‘species’ and the ‘niche’ respectively in order to show that this view of the Language–Thought relationship can receive support by observing the parallels it has in other domains.

But what did exactly enable this very reciprocal relation? Hurford depicts some coherent scenarios on the matter. The public labeling of mental representations — our ‘first shared lexicon’ — was probably started up, according to his exercise of retrospection, by a three-step process: First, a hominin should have produced a random noise attempting to convey some idea; second, some hearer guessed from the context what idea was meant; and third, the random noise became arbitrarily associated with that idea. Subsequently, by a combination of

³ Physical patterns/laws even precede living beings’ nature, as it is put forward in Burge’s theory of perceptual representation (2010: 521): “the rhythms of the environment are encoded in an organism’s physical rhythms. Objective intunement precedes objective representation”.

social ingredients, such as trust and cooperation, humans became to realize meaningful signals were advantageous, and a voluntary control over their use increased too. It is in this way that Hurford conceives Deacon's (2010) concept of *symbolic niche* as a point of no return towards our current complex languages. Nevertheless, there still is 'something' missing at a deeper level, which must have favored this new habitat: If we assume that only humans make use of a symbolic (triadic) communication and a developed thought, there must be an important *mental* transition between core cognitive levels,⁴ a transition that did not occur in other animals which are also inhabitants of niches.⁵

Of course, this terminological issue has something to do with Hurford's adscription to the gradual continuity in syntactic evolution particularly, and in general animal cognition. He defends the continuum path from the 'one-word stage' to our current state of complex syntax, in the same way as he promoted in volume I the graduality of meaning in our wide animal kingdom. While pleading the case for an 'evolutionary journey', left-hemisphere specialization for auditory attention to conspecific calls links monkeys and humans in a thread of continuous evolution; by contrast, the same cannot be said regarding continuity from primate calls to words, where facts go against a straightforward continuity: "[T]here is no inconsistency in arguing for (1) continuity in production of simple calls and auditory attention to the calls of conspecifics, and (2) discontinuity in vocal learning and in production of complex signals" (p. 111). However, the underlying mechanisms which triggered both processes were essential for human spoken languages to develop; what is more, appealing to the lowest common denominator — i.e. 'voicing' — would include not only primates, but all mammals (Tallerman 2011: 486).

Once the symbolic niche was settled, speech began to evolve. Going through several sections about common properties of languages, anthropology, particularities in syntax, and genotype changes, Hurford reaches a primitive linguistic stage where units had "an internal coherence distinguishing them from any kind of looser discourse-level organization. At its simplest, this coherence is marked by pauses at the boundaries of the units" (p. 608). In Bickerton's view of protolanguage, many two-word utterances at that stage should have contained words denoting actions and objects, although non-syntactically determined. Hurford explains this phenomenon in pragmatic terms, namely, the bipartite organization responds to a distinction between 'constant' (what you are talking about: topic) and 'changing' (new information: comment); moreover, if a single word is not enough to identify the topic, another object-denoting word may be placed together (and will eventually become an adjective, if grammaticalized).

⁴ This alludes to cognitive levels such as perception, conceptual spaces, propositional thought, (core) knowledge, etc., which imply computation and mental representations. If there was a significant change, one might legitimately wonder from where to where exactly this transition took place. Moreover, if non-human animals possess all the mentioned representational 'stages' in a primitive (/proto-)way, what is the nature of the spark that made humans, but no other animals, evolve in particular way, given that non-human animals form their own niches as well?

⁵ By way of illustration: "If we do not rule by fiat that the term *concept* be reserved only for linguistic creatures such as adult humans, the categories in terms of which an animal segregates its experiences can reasonably be called *proto-concepts* at least" (Hurford 2011: 371).

In an insightful link in volume I, Hurford puts forward the connection between the ventral and dorsal neural pathways, and predicate–argument structure respectively (a connection already suggested by Jackendoff & Landau 1992: 121–123). In a nutshell, the dorsal pathway (‘where-stream’) is said to identify the location of an object; the ventral pathway (‘what-stream’) gives all the properties necessary to identify it. External objects delivered by the dorsal stream are given individual variables (x, y, z); categorical judgments about objects’ properties are delivered by the ventral stream and considered predicates (*red, cat, Mary*). Additionally, two types of psychological attention take part in this process: global attention delivers one-place judgments about the whole scene, while local attention delivers one-place judgments about the objects within each scene; the two processes, global and local, operate in parallel. The notation Hurford proposes follows that of the Discourse Representation Theory (Kamp & Reyle 1993) regarding the use of boxes; he depicts each one containing the conflated information of both the dorsal and the ventral stream through the two attentional processes; e.g., $\boxed{\text{FLY-SMALL}}$ (local attention), or $\boxed{\text{FLY} \boxed{\text{FLY-SMALL}}}$ (global and local). However, the public arrangement of this bipartite information is not developed any further in neuroanatomical terms but in linguistic ones and the neural streams make way for pragmatic distinctions. Despite this, recent parallel research indirectly underpin this interdisciplinary enterprise of Hurford’s, whose importance lies in the appealing connection between neuroanatomy, cognition, and language. The set of linkages that the Complex Systems Theory proposes allows envisaging a primitive mechanism of this word-word coherent association; as Solé (2005: 289) argues, “two words will be linked if they share at least one object of reference”. Similarly, in the field of semantics, Pietroski (to appear) advocates that phrasal meanings are instructions for how to build concepts; in his view, for instance, the lexical expression *brown cow* assembles the concepts BROWN and COW to form the monadic concept BROWN-COW (i.e. no brown cows can be thought as not being cows). Solé’s and Pietroski’s associations seem to resemble Hurford’s neural linkage, and their two-word/concept conjunction susceptible of being originally identified by the dorsal pathway (which would be tracking one object, assigned a variable x), and categorized twice by the ventral pathway — according to his view. In this sense, the external shape may be somehow mirroring the internal one, or following similar (neurological) principles; a subtle issue, that Hurford has preferred not to go into in volume II.

In an overall assessment and comparison of both volumes, the first one deals with issues on the evolution of communication in a somewhat clearer manner than the second one deals with aspects of grammar again in the light of evolution. This could be partly due to the size of the second volume — which at times makes it hard to retain the kind of thematic continuity that one enjoys throughout the first volume — and partly due to the fact that the first volume is more succinct in placing the discussion of whatever notions unwoven in the appropriate (i.e. as informative as possible) perspective through forming fruitful connections with a broad amount of literature from primatology, neuroscience, philosophy, and linguistics.

The second volume is slightly lacking in this respect; consider, for instance, the argument about non-compositionality in complex signs in animal communi-

cation: Hurford correctly points out that the syntax of these sequences is not semantically compositional, however this claim should ideally be connected to pre-existing related claims (e.g., Mirolli & Nolfi 2010, Byrd & Mintz 2010) that one finds in the literature on this topic. Another reason that the first volume is more coherent and overall making a more significant contribution is that the content of this volume (i.e. the evolution of semantics and pragmatics) offers fertile ground for identifying parallels of (sub-aspects of) key functions of human cognition in the cognitive systems of other species, therefore it nicely supports an argument against a drastic, single-step advancement in the course of evolution. On the contrary, the subject matter of the second volume is the evolution of grammar. Once the point of inquiry shifts from semantics and pragmatics to syntax, the task of finding the relevant parallels in other species that would allude to the gradual character of evolution becomes more burdensome because it is precisely in this level of linguistic analysis that humans differ greatly from other species (take, for instance, the ability to — compositionally and hierarchically — combine lexicalized concepts into larger sequences).

The aforementioned concerns notwithstanding, Hurford's duology provides in its totality one of the broadest syntheses of a variety of topics in the area of the evolution of communication and grammar in humans. As such, this work is of great interest especially to those who wish to acquire a background in biolinguistics that is truly interdisciplinarily informed by recent developments in many interfacing fields.

References

- Bickerton, Derek. 2009. *Adam's Tongue: How Humans Made Language, How Language Made Humans*. New York: Hill and Wang.
- Burge, Tyler. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrd, Dani & Toben H. Mintz. 2010. *Discovering Speech, Words, and Mind*. Malden, MA: Wiley-Blackwell.
- Chan, Jason C. K., Ayanna K. Thomas & John B. Bulevich. 2009. Recalling a witnessed event increases eyewitness suggestibility: The reversed testing effect. *Psychological Science* 20, 66–73.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Darwin, Charles. 1871 [1981]. *The Descent of Man*. Princeton, NJ: Princeton University Press.
- Deacon, Terrence W. 2010. On the human: Rethinking the natural selection of human language. *On the Human*. [Online forum provided by the National Humanities Center at <http://onthehuman.org/2010/02/on-the-human-rethinking-the-natural-selection-of-human-language>.]
- Dummett, Michael E. A. 1993. *The Origins of Analytical Philosophy*. London: Duckworth.
- Edwardes, Martin. 2009. A Darwinian approach to meaning. *Journal of Evolutionary Psychology* 7, 191–194.

- Fitch, W. Tecumseh. 2010. *The Evolution of Language*. Cambridge: Cambridge University Press.
- Fitch, W. Tecumseh, Marc D. Hauser & Noam Chomsky. 2005. The evolution of the language faculty: Clarifications and implications. *Cognition* 97, 179–210.
- Fodor, Jerry A. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
- Hauser, Marc D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298, 1569–1579.
- Hurford, James R. 2011. The origins of meaning. In Maggie Tallerman & Kathleen R. Gibson (eds.), *The Oxford Handbook of Language Evolution*, 370–381. Oxford: Oxford University Press.
- Jackendoff, Ray & Barbara Landau. 1992. Spatial language and spatial cognition. In Ray Jackendoff, *Languages of the Mind*, 99–124. Cambridge, MA: MIT Press.
- Kamp, Hans & Uwe Reyle. 1993. *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer.
- Mirolli, Marco & Stefano Nolfi. 2010. Evolving communication in embodied agents: Theory, methods, and evaluation. In Stefano Nolfi & Marco Mirolli (eds.), *Evolution of Communication and Language in Embodied Agents*, 105–122. Berlin: Springer.
- Pietroski, Paul M. To appear. *Semantics without Truth Values*. Oxford: Oxford University Press.
- Sapir, Edward. 1921. *Language: An Introduction to the Study of Speech*. New York: Harcourt, Brace & Co.
- Solé, Ricard. 2005. Syntax for free? *Nature* 434, 289.
- Tallerman, Maggie. 2011. Protolanguage. In Maggie Tallerman & Kathleen R. Gibson (eds.), *The Oxford Handbook of Language Evolution*, 479–491. Oxford: Oxford University Press.
- Tulving, Endel. 1999. Episodic vs semantic memory. In Robert A. Wilson & Frank C. Keil (eds.), *The MIT Encyclopedia of the Cognitive Sciences*, 278–280. Cambridge, MA: MIT Press.
- Wells, Garry L. & Elizabeth A. Olson. 2003. Eyewitness testimony. *Annual Review of Psychology* 54, 277–295.

Evelina Leivada
Universitat de Barcelona
Department of Linguistics
Gran Via de les Corts Catalanes, 585
08007 Barcelona
Spain
evelina@biolinguistics.eu

Ana M. Suárez
Universidad Autónoma de Madrid
Department of Linguistics
Francisco Tomás y Valiente, 1
28049 Madrid
Spain
am.suarez@uam.es

Knots, Language, and Computation: More Bermuda than Love

David J. Lobina & Mark Brenchley

[T]he greater part of what is said and written upon it is mere windy talk, the assertion of subjective views which commend themselves to no mind save the one that produces them, and which are apt to be offered with a confidence, and defended with a tenacity, that are in inverse ratio to their acceptableness. This has given the whole question a bad repute among sober-minded philologists [...].

W.D. Whitney, *On the Present State of the Question as to the Origin of Language*

[t]he relevance of Ristad's results, regardless now of their accuracy [...]

S. Balari et al. (2012: 94)

1. Introduction

We commend Balari *et al.* (2012; BEA2 henceforth) for managing to wring a 33-page response from a 9-page critique (Lobina 2012a) of the arguments put forth in Camps & Uriagereka (2006; C&U) and Balari *et al.* (2011; BEA1).¹ And we certainly welcome the detail and slow pace; good attributes when the understanding of eclectic minds is at stake. Nevertheless, despite being impressed by the 'form' of their response, we find ourselves unmoved by its 'substance'. In particular, we find no reason to abandon the main conclusions reached by Lobina (2012a), namely that Knot Theory (Knott) has nothing to say about the knot-tying abilities of humans, and perhaps even less about the general nature of human cognition.

Be that as it may, we feel a further response is in order, not least because the argument outlined in C&U and BEA1 has now become something of a moving target. To this end, we begin by reminding BEA2 of the chronology of the arguments they purport to defend, pointing up these arguments as they actually appeared in C&U and BEA1. Section 3 then presents a critical analysis of the specifics of the proposal put forth in BEA2 by way of response to Lobina's criti-

The first author thanks the second for his collaboration, his very few professional commitments having allowed him to participate in this essay. Each author would also like to take this opportunity to state, categorically, that any errors to be found herein are unequivocally the fault of the other. Part of this research was funded by an AGAUR grant 2009SGR-401 awarded to the first author and by an ESRC studentship grant ES/I017224/1 to the second.

¹ BEA2 also discuss Lobina (2012b), a longer version that remains unpublished. We will only mention the former on a number of occasions here, befitting BEA2's own emphasis.



cisms; a proposal we show to be entirely novel with respect to those originally presented in C&U and BEA1, and just as unsatisfactory. This demonstrated, section 4 switches tack and offers a broader diagnostic of the deeper structural problem that we take to mark our contending authors, arguing the conceptual framework of BEA2 to rest on an unfortunate and fundamental equivocation. By way of conclusion, section 5 offers some cautionary remarks regarding the moral of the BEA2 story for the biolinguistic enterprise, at least as it relates to evolutionary speculation.

2. To Recap

In Figure 7a of their paper, C&U (p. 47) presented the reader with a simple knot and its reputed transformation from a loose string. In this graphic, C&U assigned symbols to different segments of the knot — the “implied relations” within a knot (*ibid.*) — and suggested that these symbols can be manipulated by a production (string-rewriting) system, thereby reducing knot-tying to a succession of grammatical rules in the technical sense of formal language theory. In rather intuitive terms, C&U contended that one has “to ‘hold’ and ‘skip’ [the internal elements of a knot] to be able to lace the knot back into place” (*ibid.*), a process that, according to C&U, cannot be the “consequence of a Markovian process of sequencing adjacent elements” (*ibid.*). In order to formally ground this intuition, C&U drew a link between knot-tying and Knott, a sub-field of mathematical topology, relying on but a single reference to do so: an unpublished software manual meant solely for the studying of Knott (Mount 1985). In this booklet, Mount mused that the Unknotting problem — a particular Knot recognition problem of Knott — could perhaps be modelled as a context-sensitive formal grammar problem.² From this, C&U (p. 63) concluded that knots were not describable by a generative procedure of less than context-sensitive power, a conclusion they categorically stated to be beyond “rational debate” (*ibid.*; see below). BEA2 (p. 104) make much of the fact that Lobina (2012a) ridicules this statement, but given the evidence C&U adduced in its favour — the musings of a software manual about a Knot recognition problem that has nothing to say about how a human being actually transforms an actual string into an actual knot — it is hard to know how else a reasonable reader could be expected to respond.³

Note that we are here abstaining from judging the general validity of the statement per se; we are simply pointing out that the evidence that was actually provided in its favour was ludicrous. As we will show later on, BEA2 go to great lengths to justify C&U and admonish Lobina (2012a), but they seem oblivious to the fact that this is all bit *post facto*. Put simply, Lobina (2012a) had ample reason to query such a conclusion, there being nothing in C&U to support it.

BEA1, in turn, offered a very similar argument. Taking their heed from C&U, and attempting to find domains that might “presuppose a ‘technical intelli-

² In fact, C&U claimed that Mount (1985) had shown that “we need a context-sensitive system to *create* a knot” (p.47, our emphasis), a blatant misinterpretation of the reference they quote. We follow BEA2 in writing ‘Knot’ for mathematical knots.

³ Unsurprisingly, BEA2 abstain from mentioning (or indeed justifying) the employment of Mount (1985) by C&U.

gence' that could well be [...] parasitic on the Faculty of Language" (BEA1: 11), the authors drew the reader's attention to the "complexity of *knot-tying*" (*ibid.*; our emphasis), claiming that when making a knot one must "relate a portion in the knot with the background 'figure'" (*ibid.*). This, they "intuitively" continued, involves an operation that implicates "grouping and long distance-like relations" (*ibid.*). Naturally, the mention of such relations is clearly intended to reference those features that make natural language mildly context-sensitive, and, as pointed out in Lobina (2012a), this is then connected to the computational complexity of '(un)tying knots', which "seems to require an underlying computational system of Type 1" (that is, a context-sensitive system; BEA1: 11). It is in this context that BEA1 referenced Hass *et al.* (1999) and claimed that the computational complexity involved in "determining whether any string is knotted is known to have a complexity comparable to the one needed to process linguistic expressions" (BEA1: 11).⁴

Note two things then: (a) Both C&U and BEA1 focused on knot-tying, that is, their arguments centred on how you go from a loose string to a knot; and (b) knots/Knots were claimed to be only describable by a context-sensitive system, even though no actual proof of this was provided; it was merely stated (and, considering the reference C&U used, simply imagined).

It was this overall argument, briefly recapped here, that Lobina (2012a) undertook to critique, and it was a recurrent point of that paper that Knott has nothing to say about how a string is converted into a knot/Knot, for the critical and substantive reason that Knott quite simply doesn't consider how a loose string becomes tied into a knot — which was without a doubt the focus of both C&U and BEA1.

Now, BEA2 (p. 98) seem quite agitated by Lobina's (2012a) suggestion that Knott takes 'tied knots' as a starting point, a statement they consider downright false. Admittedly, the formulation chosen in Lobina (2012a) is perhaps a little clumsy, but we point out that 'tied' is there used as a synonym of *bound*, and not of *knotted* (the antonym would be *loose*). In other words, this definition merely stated that Knott studies closed Knots, together with their relation to the Unknot — and *not* the conversion of strings into Knots/knots. It seems to us that such an interpretation ought to be obvious to anyone capable of a careful analytical reading of either the published Lobina (2012a) or the unpublished Lobina (2012b) given the emphasis therein placed on the irrelevance of Knott for real-life knot-tying; this would certainly have saved us from the irrelevant barrage of random quotes that BEA2 (pp. 98–99) grace us with.⁵ Thus, the contradiction BEA2 see between employing the locution 'tied knots' and the Unknotting problem of

⁴ However, BEA1 never actually offered any details or references regarding the relevant complexity needed to process linguistic expressions.

⁵ BEA2 engage in a lot of 'proof quoting' to make their points, but this is not always accompanied by proper interpretation of the material they cite. Consider, for example, BEA2's (pp. 95–96) insistence on the relevance of Knott for the study of real knots, contra an individual statement they select from Lobina (2012a), namely that "the knots that Knott studies have nothing to do with real knots" (p. 74). We, with Lobina (2012a), don't deny that Knots can be regarded as abstract, mathematical models of real knots; rather, the point is that the Knot recognition problem "has nothing to do with the computational complexity or expressive power of (un)tying a knot in *real* life" (Lobina 2012a: 76; our emphasis).

Knott as described in Lobina (2012a) is, for us, the result of misinterpretation, a failure to substantively engage with the issue under discussion.

Be that as it may, let us accommodate BEA2 and employ their definition of the Unknotting problem; the overall point made in Lobina (2012a) still follows. BEA2 describe the Unknotting problem — recall, a Knot recognition problem — in terms of an explicit question, namely: “Is this projection of a Knot a projection of the unknot [sic]?” (p. 99), a problem that involves “telling apart the unknot [sic] and any knot [sic] projection” (p. 100; see also the relevant figures in Lobina 2012a and BEA2). Clearly, on this or anyone else’s definition, the Unknotting problem is not *prima facie* related to the conversion of a loose string into a knot/Knot in the sense in which either C&U or BEA1 seem to imply; moreover, neither set of original authors offered any plausible reasons to relate the computational complexity of the Unknotting problem with that of natural language processing. This, at least, was the main point that Lobina (2012a) tried to convey and which BEA2 seem to have entirely missed.

3. The “All Tied in Knots” Recognition Problem

So much, we hazard, for the original claims critiqued in Lobina (2012a, 2012b). As noted above, however, whilst BEA2 apparently set out to defend and justify the ways of C&U and BEA1 to man, they primarily do so by presenting a new means for relating the Unknotting problem to real-world human knot-tying. That is, whilst the argument presented in C&U and BEA1 is clearly centred on the possibility of modeling the conversion of a string into a knot as a specific sort of formal grammar problem, this is not quite the case within BEA2, where an entirely novel argument, based on their presumption that human artefacts are cognitively transparent with respect to the “cognitive resources” that produced them (p. 79), is poured forth as if it were smoothly related to that which went before. Specifically, they argue that, in order to tie a knot, you have to first visualise the particular knot you are going to create, and such a visualisation is an instance of visual object recognition, a process BEA2 regard as analogous to the Knot recognition problem of Knott. We turn to this argument now, but forewarn its ultimate failure to offer a coherent response to the original criticisms. Indeed, as best we can tell, the actual defence mounted by BEA2 seems to rest on two evidentially dubious steps, and a rather persistent confusion between the computational complexity of string recognition and the parsing complexity of language processing.

With this in mind, the initial step in their argument involves the claim that tying a knot requires a prior act of visualisation, ‘knot production’ instantiating “at least a particular case of the more general problem of object recognition” (BEA2: 97); a claim that follows from observations which BEA2 themselves note to be “admittedly scant” (*ibid.*). The authors adduce two such observations, in fact, opening with a reference to some empirical evidence which they claim indicates human beings to be particularly poor at learning to produce knots by simply observing motor sequences. BEA2 then connect this evidence to their own personal experience of knot-tying, suspecting that knowledge of complex knottings is taught “by resorting to mnemonic techniques” (*ibid.*) which involve the

working out of the number and direction of the relevant crossings.

Unfortunately, neither set of observations seems to withstand much scrutiny. Take, first, their supposed experimental ‘evidence’, at least with respect to the specific citation from Michel & Harkins (1985; BEA2: 97). Having taken the liberty of following this study up, we are more than a little disconcerted to find it more than a little misrepresented. So, whilst it is true that only 37% of the subjects successfully learned all three knots, the figure BEA2 would have us focus on, a full 97% still managed to learn at least one knot (Michel & Harkins 1985: 598), the subjects doing so based on observations of a mere five demonstrations per knot, each such demonstration lasting a mere 15–20 seconds each (*ibid.*). To our mind, such unconvincing learning conditions hardly constitute evidence against learning by observing motor sequences. Even more disconcertingly, it seems that when handedness was taken into account, *the actual object of the study*, a full 90% of subjects somehow managed to learn at least two of the knots (*ibid.*); that is, observational learning was greatly enhanced when teacher and student were either *both* left-handed *or* both right-handed. At the very least, this suggests that any of the evinced difficulties may actually have resulted from having to observe a teacher who utilised what was, for the differently-handed subjects, a non-natural set of hand co-ordinations.⁶ Quite how BEA2 take this to be evidence for their proposal is beyond us. Indeed, according to the logic of their own argument, motor sequence internalisation supposedly goes hand-in-hand with successful observational learning; since, therefore, we would actually seem to have evidence of successful observational learning, what BEA2 actually offer up is evidence of successful motor sequence internalisation.⁷ We are much obliged.

This leaves their citation of personal experience; a particularly strange sort of evidence to present and have taken at face value. Nevertheless, since this is the kind of evidence we are apparently allowed to present, we dutifully note that, though neither of us is, has been, or likely intends to be a boy scout, at least one of us used to sail in their youth, and distinctly recalls learning knots sometimes by copying the movements they observed others making, sometimes by being explicitly taught (usually in terms of confusingly metaphorical rabbits, trees, and holes), and sometimes by both methods. Such, indeed, is the value of scant observations. Indeed, even taken together, we doubt that these two sets of observations, at least as presented, can seriously lend *any* kind of substantive support to the claim that “knotting abilities have little (if anything) to do with the accurate internalization [sic] of a motor sequence” (BEA2: 97), or that “to make a

⁶ And this is not even to point out that three different types of knot were demonstrated, with the ‘magic’ and ‘butterfly’ knots each being learned faster than the ‘sheepshank’ knots (Michel & Harkins 1985: 598–599), clearly allowing for the possibility that the final knot was of a completely different order of difficulty.

⁷ Just to be clear, the issue we point out here is entirely framed according to the apparent internal logic of BEA2 itself, based on the argument and evidence they present. It would clearly be injudicious of us to read overmuch into Michel & Harkins (1985) given our non-intimate familiarity with that specific line of research. Unlike BEA2, who seem able to make sweeping claims based on a modicum of evidence, we genuinely worry about issues of interpretation. So, for example, we wouldn’t dream of using a couple of arbitrary papers to claim that a complex phenomenon such as dyslexia can be neatly reduced to a “visuo-constructive deficit” (BEA2: 85).

knot, one needs first to represent it and to represent it one needs to figure out its topology" (*ibid.*). As far as we are concerned, therefore, this part of the argument is little more than strained supposition; a sort of proof by speculation.

Having failed to establish the plausibility of knot recognition with respect to actual knot-tying, the next step in BEA2's argument involves their assertion that this assumed process of knot recognition — by which they mean the explicit mental recognition of a knot's formal properties — can be substantively related to the Knot recognition problem of Knott. In this case, BEA2 literally offer no evidence for connecting the two. Rather, BEA2 merely admit the possibility that "knot recognition in humans is a totally different thing from Knot recognition" (p. 102), and then proceed to claim the connection between the two to be a "hypothesis that is as informed as it could possibly be" (*ibid.*). Informed by what? We are not being purposely confrontational here, the authors really do offer zero grounds for connecting these two *prima facie* disparate phenomena.⁸ Recall that, for BEA2 (p. 99), the Unknotting problem (a Knot recognition problem) is explicitly characterised as a yes-or-no answer to the question "Is this projection of a Knot a projection of the unknot [sic]?". How can *that* have anything to do with the visual recognition of a knot? Is the suggestion seriously to be, for example, that boy scouts check the success of their knot-tying activities by mentally visualising an Unknot and then determining whether their particular knot is, indeed, such a projection?

This leaves us with the final key issue in BEA2's novel argument, namely the supposed relation between parsing and recognition, together with their respective computational complexities. That is, having failed to plausibly argue for the involvement of knot recognition in knot-tying, and having failed to make any kind of case for a *prima facie* link between knot recognition and Knot recognition, BEA2 next attempt to dismiss Lobina's (2012a) claim that *parsing* a sentence is not quite the same thing as *recognising* a string of a formal language and, since the computational complexity measures we do have refer to the latter and not the former, there are quite simply no grounds for comparison between Knot recognition and language processing.⁹ To this BEA2 (p. 90) respond with a somewhat arbitrary quotation in which parsing and recognition are equated, enabling them to carry on without an apparent second thought for the actual nature of the matter under consideration. Unfortunately, BEA2 show a remarkable ignorance of basic issues in psycholinguistics.

Given the referential pyrotechnics of BEA2, we assume they would at least be moderately aware of Berwick & Weinberg (1989), for example, a publication that treats the relationship between parsing and recognition more carefully. As Berwick & Weinberg clearly state, sentence processing "involves associating a meaning with a phonological string", which "demands parsing, not just recognition" (p. 252, n. 13). That is, parsing means the recovering and assigning of structural descriptions to the linguistic input (p. 264, n. 55), and from this

⁸ At this point, BEA2 write that their hypothesis is not a possibility that "Lobina has been able to really call into question" (p. 102). That is quite right, but also rather insincere, considering that this is an entirely novel proposition, not present in either C&U or BEA1.

⁹ Recall, nevertheless, that neither C&U and BEA1 provided any evidence or references regarding the computational complexity of processing linguistic sentences.

Berwick & Weinberg conclude that parsing is harder, computationally speaking, than recognition. Whilst this is certainly true, we would nevertheless like to point out that parsing and recognition actually divorce in a much clearer and more principled way.

Pointedly, for instance, a central result of psycholinguistics has it that sentence processing proceeds incrementally, by which it is meant that partial meaning interpretations are computed during parsing (i.e. before the end of a clause). Consequently, the sentence processor carries out many valid parses that would have no string recognition equivalent; indeed, many of these parses relate to incomplete ‘chunks’ rather than full expressions *per se*. And there are many more interesting cases, each pointing to the clear divorce between parsing and formal language recognition: such as ungrammatical sentences which are nonetheless successfully parsed (or at least provided with some sort of interpretation), as in the missing verb effect (Frazier 1985); such as grammatical sentences that are nonetheless unparsable, as is the case with reduced relatives (Bever 1970); and such as the clear existence of grammatical illusions (Phillips *et al.* 2011).

In other words, though they have clearly read some of the formal grammar literature, BEA2 seem completely unable to distinguish the interests of psycholinguists from scholars working within the discipline of computational linguistics. Traditionally, the latter have been more interested in the computational complexity of language recognition, whilst the former have recently started to investigate the parsing complexity involved in the cognitive processing of linguistic utterances; a very different type of investigation, despite the superficial similarity that might otherwise be implied by their mutual use of the word ‘complexity’. As such, BEA2 simply cannot rely on the ‘complexity’ of language recognition as a measure of the ‘complexity’ involved in the parsing of linguistic structures, and they are just hopelessly confused when they state:

[I]t is the task of psycholinguists to incorporate these [computational complexity] results when building their performance models, that whatever memory limitations they postulate, whatever parsing strategies they propose, etc, should take into account the inherent structural/computational complexity of natural language. [BEA2: 94]

This is simply false: The memory limitations and strategies implicated in sentence processing are a matter of *sui generis* discovery, and they are clearly independent of the structural and computational complexity of formal languages.¹⁰ Rather, the role of the psycholinguist is to find out how the performance systems *cope* with the linguistic input they receive, which is to say that psycholinguistics aims to discover what strategies and memory limitations human psychology exhibits when undertaking the business of actually processing a sentence. In this sense, the results of formal language theory have *prima facie* very little to do with such an investigation.

¹⁰ BEA2 use the expression “structural complexity” for what Lobina (2012a) termed the “expressive power” of a grammar (that is, the set of strings a grammar can generate). We think the former formulation is somewhat misleading, but won’t dwell on this point here. In what follows, we’ll use both terms interchangeably.

This is a matter of principle, in fact, for the products of linguistic cognition are not solely the result of the underlying grammatical system, they are produced by a much larger complex; a conglomerate that likely includes a parser (as distinct from the grammar), working memory, and perhaps other systems. Consequently, there is little reason to believe that a proper analysis of linguistic productions will be able to rely solely on the results of formal language theory. On the contrary, if we are to take seriously the proposal that cognitive artefacts are transparent with respect to the “cognitive resources” that produced them, we must also include processing considerations in our models, these considerations distinct from (and perhaps bearing *at best* only an indirect relationship to) such theory. What is certainly not the case, despite BEA2’s misconstrual, is that the expressive power of natural language must be built *into* performance models.¹¹

This point was already clear as far back as Miller & Chomsky (1963) where the relationship between the grammar and the parser was explored, two linguistic constructs that need not be, and quite probably aren’t, isomorphic. So, for example, among the many things Miller & Chomsky discussed was the working memory capacity involved in the processing of centre-embedded sentences.¹² After reviewing some experimental results, Miller & Chomsky concluded that subjects could successfully parse up to seven centre-embedded clauses, a limit they linked to the ‘magical’ number 7, in reference to Miller’s (1956) now-classic study of working memory’s capacity to recode information into manageable chunks. Note, then, that the memory capacity Miller & Chomsky assigned to the processor bears no direct relation to the expressive capacity of language. That this is so is demonstrated by the human ability to actually produce the centre-embedded construction itself; the capacity of human memory relating only to the *number* of central embeddings it can cope with. Thus, the clear implication is that the memory capacity of humans can only be determined by the measuring devices and experimental paradigms the psycholinguist has at their disposal.

A similar point applies with respect to parsing strategies. Naturally, of course, it is true that the grammar must be somehow related to the parser, for if this weren’t the case, the parser would be unable to assign linguistic structure to the linguistic input; that is, the parser needs to have access to a grammatical ‘knowledge’ base if it is to be properly operative. However, what does not follow from this is that the parser implements the rules of the grammar in a direct and transparent manner; in fact the computations that the parser implements need not be isomorphic to those of the grammar at all (see Bever 1970 for a relevant discussion of this last point). We would, then, be talking about two very different types of ‘computations’: those underlying the sound–meaning pairs the language faculty generates in the technical sense employed within generative grammar, and those implicated in the operations of the parser during real-time processing. This is, again, clearly so in the case of nested structures: Whilst the grammar in principle allows for unbounded centre-embedded structures, the ability of the

¹¹ Indeed it is perfectly possible, at least in principle, that human language considered in terms of a specific knowledge base demonstrates any number of expressive properties that are never realised due to a mismatch with the capacity of the parser to implement these properties.

¹² Naturally, the data Miller & Chomsky (1963) analysed are now somewhat dated; however, our interest here focuses on the underlying idea borne out by their paper.

parser to process structures like these is hampered by processing constraints. These constraints determine, to a certain extent, the character of the computations the parser effects.

Some of the processing constraints we have in mind here would at least include perceptual strategies, memory limitations, the design of mental architecture, the role of context, frequency effects, and so on. Now, whilst it is at present very uncertain how all these factors conspire into an overall model of language comprehension (not to mention production), we can nevertheless highlight one theoretical perspective in order to bring the general point home, that of Townsend & Bever (2001). Therein, the authors present an analysis-by-synthesis approach according to which the processor undertakes, first of all, a preliminary analysis of the signal by imposing a Noun–Verb–Noun (NVN) template onto the input. This stage is then followed by the application of the rules of grammar, in some fashion or other. What interests us here is the postulation of the NVN template, a *sui generis* perceptual strategy based on the high frequency of such configurations (at least in English). Critically, for our purposes, this specific parsing strategy is proposed in order to explain data produced by experimental studies of online language comprehension and, contra BEA2, has no direct basis in formal language theory.

Simply put, psycholinguists will do well, as they already do, to ignore the misguided advice BEA2 bestows upon them, focusing instead on the effectively *sui generis* complexity involved in the actual parsing of linguistic products (and quite independently of the abstract formal complexity these products might have). So, it is entirely in the spirit of Miller & Chomsky (1963), for instance, that the psycholinguistics literature has recently provided a number of studies which have sought to discern the *parsing* complexity involved in recovering and assigning the right structural description to the linguistic input. To name but two examples, Gibson (1998) calculated parsing complexity in terms of the number of new discourse referents that are introduced in a sentence, whilst Hawkins (2004) focused on the number of syntactic nodes required to handle a particular piece of syntax.¹³ Of course, it is far too early to settle on a specific measure or to favour a given proposal, but if the computational complexity involved in processing language is to be related to that of another cognitive phenomenon, we believe that the focus should lie on the sort of approach that Gibson (1998) and Hawkins (2004) advocate, and not on formal language recognition as understood within the technical sense of the theory of computation.

In any case, the manner in which BEA2 treat the issue of formal language recognition has its own problems, and it may be worth at least mentioning them here. Thus, BEA2 take up ample space in showing that whilst the structural complexity of a formal language can only be assessed directly via “the devices that are capable of specifying” a language (p. 89), and then tell us that the computational complexity involved in recognising a formal language can be so determined; that is, independently of any formalism (pp. 92–94). Yet they hardly argue in favour of such a position, merely offering an off-hand reference to Ristad (1993; cited therein). We do not, of course, doubt the value of Ristad’s

¹³ There are other possible variables, of course, such as number of words or simply time sequences.

results, at least for formal language theory, but it seems to us that they are of very little application to either linguistics or psycholinguistics. After all, to abstract away from the device that specifies a language is to focus on an infinite set of strings (the formal language), and the latter has no cognitive relevance whatsoever. We should perhaps remind BEA2 that the important concept in linguistics is the ‘grammar’ and not a formal language as a set of strings whose recognition is, presumably, between P and PSPACE (BEA2: 93). That is, since it is the grammar that is postulated to be mentally represented in the minds of speakers and hearers, the computational complexity of processing language ought to be closely related to this construct. As a matter of fact, Lobina (2012a) made the point of referencing a recent summary of formal language theory (viz. Pratt-Hartmann 2010), wherein it is stated that different grammatical formalisms are in different computational complexity classes. Further, Lobina (2012a) pointed out that none of the complexity measures he was aware of matched the computational complexity of the Unknotting problem, as determined in the single reference BEA1 had employed (namely, Hass *et al.* 1999). So why are BEA2 so sanguine about Ristad (1993)?

It seems to us that it has nothing to do with Ristad (1993) *per se*; it is merely based on the fact that Ristad’s conclusions are simply convenient for BEA2, as Ristad determined, according to BEA2 (p. 94), that the computations underlying natural language are NP-complete, which is precisely the complexity class of the Unknotting problem BEA2 would have us focus on. Had BEA2 decided to follow Pratt-Hartmann (2010) and settle on a specific grammatical formalism — such as tree-adjoining grammar, for instance —, they would not have been able to connect the computational complexities of such disparate phenomena, a major point that we return to in our final section. In other words, if we postulate that it is the grammar and not the set of strings that is mentally represented (and the latter couldn’t be because we quite simply can’t represent infinite sets of strings), then the computational complexity of processing language *has* to relate to the mental reality of the grammar — and therefore its formalism.

Taking all of the above into account, therefore, the lesson for BEA2 would seem to be that the actual “cognitive resources” involved in the processing of a human artefact cannot be so easily discerned through a formal mathematical analysis of these artefacts. And there is a very simple reason for this: In order to determine the appropriate cognitive resources, what we actually need to postulate is a plausible cognitive model. Such a model, if we are right, will necessarily involve properties which go beyond an artefact’s formal properties and which cannot be assumed to bear any kind of transparent relationship to these properties. This is a critical claim in the context of our reply to BEA2, and something to which we return in section 4. In order to provide some further support for this claim within the context of the present section, however, let us make a final point and close with an important caveat.

Regarding the former, take the case of the structural complexity underlying language. As is well known, Chomsky (1956, 1963) was able to show that the expressive power of language had to be at least context-free because a finite-state grammar couldn’t account for unbounded nested structures. Critically, however, Chomsky did so by relying on *both* theoretical argumentation *and* grammatical

judgements over this sort of sentences, and grammatical judgements are each a type of experimentation *and* performance data. Naturally, had Chomsky only had access to linguistic artefacts in isolation from any grammatical judgements, as in printed material or corpora transcripts, a finite-state system would have sufficed to account for the data, as corpora and the like do not even hint at the possibility of unbounded nested structures. *Mutatis mutandis*, we here claim, with respect to the analysis of the fossil record.

As for the caveat, imagine that the formal equivalence of knots/Knots and linguistic expressions had, in fact, been demonstrated. What should we take this to actually mean? If we were to specify the structural complexity of a language — that is, the set of strings a grammar can generate — we could, of course, employ either a grammar or an automaton for this purpose; but do we have an analogous analysis for knot-tying? The answer is a clear “no”; or at least it is not something that C&U, BEA1, or BEA2 provide. If, on the other hand, we were to determine the computational complexity of language recognition, we would be probing the rate of growth of space and memory resources as manifested in an automaton; but is there a similar study in the case of knot-tying? The answer is “not quite”. That is, whilst we do have a Knot recognition problem for which a complexity measure can be calculated, BEA2 have nowhere demonstrated that knot recognition can plausibly be construed as Knot recognition. Furthermore, we have here insisted that formal language recognition has very little to do with the actual processing of natural language. That is, even though a formal language may be recognised by an automaton, the expressions of a natural language must be parsed by humans, and the latter is crucially a very different matter. In this respect, and once more, what of knot-tying? Well, we have argued that no reasons have been provided to warrant the supposition that knot-tying is preceded by knot recognition. And, be *that* as it may, it seems to us that knot-tying is surely a case of knot *production*, meaning, surely, that the relevant line of comparison to pursue would more likely be language ‘production’, rather than language ‘recognition’. Now, whilst we won’t pursue the last speculation in this paper, our admonition here is simply that the computational properties of language and those of knots/Knots/knot-tying do not seem to match up at *any* level.

Relatedly, BEA2 are certainly proceeding too fast (and too loose). After all, the fact that two computational problems happen to be in the same complexity class does not mean that they are actually related; or that they share the same “cognitive resources”. We can certainly recognise that C&U and BEA1 attempted to model knot-tying as a grammatical production system, even though no demonstration was in fact provided. We also recognise that it is one of the aims of BEA2 to show that the Knot recognition problem can be modelled as a grammatical problem, but it is telling that the closest they come to achieving this is by using one of the references included in Lobina (2012a, 2012b), namely Turing (1954). In fact, Lobina (2012a) was not sceptical of the possibility of modelling the Unknotting problem as a formal grammar problem — that is why the reference to Turing (1954) was included in the first place. What’s more, we are not even sceptical of the possibility of modelling, among other things, *knot-tying* as a grammar problem. Our point is twofold; firstly, the key notion is *knot-tying*, and it is this that is the phenomenon that needs to be modelled as a grammatical pro-

blem, but so far no such thing has been provided; and secondly, even if the latter were to be successfully undertaken, it is not the case that by describing two *prima facie* un-related problems with the same formal machinery (i.e. the tools of formal language theory), the common features of a postulated underlying computational system are *ipso facto* unearthed. Formal analysis, after all, is no substitute for substantive argumentation. (The latter is an important issue, to which we turn in the next section.)

As a final point, let us mention that towards the end of their paper (p. 104), BEA2 remind us of the precise claim which C&U considered to not be subject to rational debate, namely that “knots are *not describable* by any generative procedure that does not have enough operational memory to count as context-sensitive” (C&U: 63). BEA2 take this statement to mean, rather banally and with no insignificant liberty of interpretation, that the “inherent complexity” (BEA2: 104) of knots is not controversial, a complexity that can apparently be related to the structural and computational complexity of natural language; a torchlight for the biolinguistics enterprise, it seems to them (*ibid.*). This, according to BEA2, is the message C&U tried to convey, and it is their message too.

This would be all good and proper, if it wasn't for the blatant disingenuousness. To recap once again: C&U suggested, albeit incompetently, that a string can be converted into a knot by following a series of grammar rules — which, we suppose, is what is meant by a knot being ‘describable’ — but their final conclusion, cited *supra*, was based on a preposterous reference to an unpublished software manual in which it was mused that the Unknotting problem, and not the conversion of a string into a knot, could be so modelled. Clearly, C&U were not entitled to hold such a belief; *a fortiori*, they were not entitled to affirm that it was not subject to rational debate to discuss such a conclusion, the latter a ridiculous claim rightly ridiculed in both Lobina (2012a) and Lobina (2012b). BEA2's ‘description’ of a knot is instead based on visualising a knot first, a process they claim to be related to the Knot recognition problem of Knott. Note that such a ‘description’ of a knot is completely different from that of C&U; note further that it is also different from the conceptualisation of knot-tying advanced in BEA1 in terms of relating a knot to its background figure, which includes a number of the grouping and long-distance relations that arise thereby (in any case, a ‘description’ that BEA1 didn't justify either). Recognise, then, that BEA2's argument is very different indeed.¹⁴ Clearly, too, BEA2 are not entitled to use their novel argument to justify either C&U or BEA1, as this can only be regarded as an entirely *post facto* justification. The overall result is nothing more than a gigantic increase in the volume of fog that no amount of ‘new ordering’ will fix. It is no wonder, then, that C&U, BEA1, and BEA2 have got themselves so tangled up; and no wonder that they are seemingly incapable of *recognising* their predicament.

All in all, then, BEA2 fail to make a reasonable case for relating the subject-matter of Knott with formal language theory, linguistics or psycholinguistics. In particular, Knott has nothing to say about how one produces a knot from a

¹⁴ In fact, given BEA2's insistence on the need to visualise a knot before tying it, and their dismissal of the role of motor sequencing in learning how to tie a knot, they are unsurprisingly pretty much indifferent to the actual action of tying a knot!

string. There is also no reason to believe that a knot is visually recognised or represented before tying a string into a knot; and even if this were the case, we have not been offered any grounds to relate such ability to the Knot recognition problem of Knott. Finally, the computational complexity of the Knot recognition problem is not directly relatable to the computational complexity implicated in linguistic processing; if the latter is to be defended, it needs to be substantively related, not analogically hinted at. In short, the charges levied at C&U and BEA1 by Lobina (2012a) remain basically untouched: Knott is quite simply misapplied and unrelatable to the study of the faculty of language. If we were to be charitable, we would simply point to the rather glaring errors that we have thus far sought to highlight, and leave it at that; but, in the end, it would seem that we are not that charitable after all.

4. A Natural Computational System by any Other Name

Thus far the present paper has essentially been addressing BEA2 in terms of their specific responses to Lobina (2012a, 2012b), arguing a consistent failure of these rejoinders to hit their respective marks. Approached at a more general level, however, we believe this consistency to be, if not willful, then, at the very least, no accident; rather, it is the inevitable outcome of the conceptual muddle that BEA2 seem to have gotten themselves into. It is to this muddle that we now turn, beginning with a question which may well have been bothering those readers inclined to have suffered us so far: why all the lather about knots?

The answer, perhaps unsurprisingly, is that the authors aren't really interested in knots *per se*. On the contrary, knots are considered to be a particularly instructive case-in-point regarding their core "thesis". This thesis we take to consist in the following set of interrelated claims:

(A) The animal mind is at least partially constructed in terms of a "natural computational system (NCS)" (BEA2: 80). This NCS is a "core engine [...] subserving some (but not necessarily all) of the main cognitive functions" (*ibid.*), and "which may be modeled by an abstract machine or automaton in the sense of the mathematical theory of formal languages and automata" (*ibid.*). Each such NCS represents a "computational phenotype" (p. 83) that "one can associate to certain specific neuroanatomical configurations" (p. 80), and which is "functionally unspecific" (p. 84).

(B) Artefacts, those objects which are the material products of animal minds, instantiate a direct relationship between their formal properties and the cognitive biology of the particular animal mind that happened to produce them; that is, "manufactured objects are transparent with respect to the biological structures underlying the processes necessary to produce them" (BEA2: 79).

(C) Considered as artefacts, a formal analysis of knots and natural languages suggests them to be strikingly similar in terms of their computational properties (BEA2: 102–104).

(D) Given this closeness, it is plausible to suppose that human cognition does, in fact, represent the implementation of an NCS (BEA2: 94).

To the best of our apprehension, we believe the above to be a fair summary of the thesis underlying BEA2.¹⁵ So summarised, the centrality of knots to the authors is immediately apparent. For, taken at face value, the supposed equivalence in computational complexity transforms the somewhat trivial observation that humans happen to produce both language and knots into substantive support for the authors' notion of a "natural computational system" (NCS).¹⁶ That is, the very reason human beings are able to produce *both* knots *and* language is that we have evolved a particular cognitive setup that represents the biological implementation of an NCS of a specific computational phenotype. This is no doubt a novel thesis, one of which BEA2 seem suitably proud: Not only does it apparently open out a new perspective on the nature of cognition, it also opens up a way for researchers to productively mine the archaeological record, thereby bringing new evidence to bear on both the biolinguistics enterprise and the evolution of cognition. Much is, indeed, ado about knotting.

A central problem as we see it, however, is that the only rationale BEA2 actually seem to have for taking the supposed mathematical equivalency of knots and language seriously as substantive evidence for the cognitive equivalency of knots and language is their claim that the mathematical properties of artefacts are also their cognitive properties. Unfortunately, the only reason BEA2 seem to provide for taking *this* claim seriously is their own notion of an NCS, which they *define* outright to be something both cognitive and mathematical; something, that is, which simultaneously underlies "some (but not necessarily all) of the main *cognitive* functions of an animal mind" (BEA2: 80, our emphasis) and which is also a "*computational* phenotype" (BEA2: 83, our emphasis). In other words, BEA2 nowhere offer substantive reasons for linking the two domains; instead, the authors merely equate, and thereafter interpret all evidence accordingly.¹⁷ As such, we believe the edifice of BEA2 to be founded, and foundered, on an enviable equivocation, one which has fooled its authors into thinking they can swap out theory of cognition talk for theory of computation talk as if they were one and the same, never stopping to think either what it might mean or how it might be for the two domains to be related in the very real world of the mind; perhaps not so difficult a thing to do when the world can be predefined to suit one's fancy. Put bluntly, we're afraid they've been jumping too fast. Let's proceed at a slower pace.

¹⁵ This does, of course, assume such a summary to be genuinely possible, something of which we are not entirely convinced given the rather scattered remarks from which we have attempted to reconstruct the thesis and given that the actual thesis itself appears something of a tangled web in which premises also seem to serve as their own substantive conclusions; see, for example, BEA2 (pp. 79–111).

¹⁶ Hence, presumably, their bizarre, and somewhat random, statement that "the archaeological record [...] shows a strong correlation between the presence of language and knotting". (BEA2: 94).

¹⁷ Hence, no doubt, their bizarre, and somewhat random, statement that "the archaeological record [...] shows a strong correlation between the presence of language and knotting" (BEA2: 94).

4.1. *The Theory of Computation and Theories of Cognition*

Our initial reason for supposing equivocation to be at the heart of BEA2 is simply the sheer amount of space devoted, firstly, to recounting some essentials of the theory of computation, and, secondly, to discussing some of the computational properties of knots and language (pp. 89–104); as if this were somehow enough in itself to refute the concerns of Lobina (2012a, 2012b). Yet, so devoted, BEA2 entirely miss the thrust of the original criticisms, which never doubted the possibility of *formally* equating knots and natural language.¹⁸ Rather, what was queried was the legitimacy of moving from the possibility of formal equivalence to the plausibility of substantive equivalence; that is, of meaningfully moving from the domain of the theory of computation to the domain of cognition.

To see that any such movement cannot consist in simple switches, swapping out talk at the theory of computation level for talk at the theory of cognition level, we need only note the *prima facie* distinct concerns of the two domains when they come to consider real-world objects, whether or not said objects happen to have actually been manufactured. Thus, the theory of cognition deals with the properties of entities as actually instantiated within, and realised by, animal minds; that is, its concern lies with the properties objects have as *cognitive* objects. The theory of computation, on the other hand, deals with the properties of entities considered in the mathematical abstract; that is, its fundamental concern lies with the properties of objects modelled as *formal* objects. As such, the theory of computation is clearly under no obligation to give any thought whatsoever either as to how these formal ‘objects’ might actually be, or as to how they might actually be related to each other, within the specific confines of an animal mind. In this sense, at least, the theory of computation is somewhat like statistical analysis: You give the chosen statistical technique a particular set of numbers, and it outputs a set of results without a thought for the actual interpretation of these numbers and whether they really measure what the researcher believes them to be measuring.

By way of illustration, let us take two objects chosen at random; knots and human language, say. Handily, both objects turn out to be humanly-produced, thereby presenting themselves as reasonable candidates for explanation within the theory of cognition. Accordingly, it should be possible to assign each certain properties which serve to ground them cognitively; call these, for complete want of imagination, *cog-properties*. Equally handily, both objects turn out to be formally analysable within the framework of the theory of computation (the caveats identified in section 3 notwithstanding). Accordingly, they can each be assigned various formal symbols that enable them to be treated computationally; call these, again for complete want of imagination, *comp-properties*. It so turns out that, assigned certain *comp-properties*, knots and human language evince a certain similarity with respect to these properties, at least according to BEA2 (pp.

¹⁸ As already noted above, such a possibility was clearly acknowledged through the original referencing of Turing (1954), Turing therein showing that knots can, indeed, be modelled in terms of the theory of computation. In point of fact, we are really quite comfortable with the idea that *any* real world object can be modelled mathematically, though we do think it perhaps sensible to draw the line at Italian cuisine (Hildebrand & Kenedy 2010).

102–104). The question that can now be asked is what this computational equivalence might have to tell us about their cognitive equivalence? Very little, we suggest, and quite likely nothing, for the simple reason that to speak of knots and language computationally is first and foremost to speak of these objects as already having been couched in terms of the formal symbols that are the province of the theory of computation. As such, it is only at this level of description that any formal results are defined and at which any equivalence directly holds. Hence, there is no *necessary* reason to believe that this formal equivalence will hold when we come to consider knots and language within the domain of the theory of cognition, even though both may well be cognitively-derived objects, because there is no *necessary* reason to believe that the comp-properties theory of computation researchers might legitimately choose to assign knots and language bear any relation at all to their cog-properties; that is, exactly those properties which the human mind takes into account with respect to its competence for knots and language. *Prima facie*, being mind-external physical objects, knots and natural language utterances may have any number of properties in common, any number of which the human mind neither notices nor processes, and which are therefore irrelevant in terms of a theory of cognition.

Now, to be fair to BEA2, there are moments where the authors acknowledge this point. So, for instance, they state that “from the fact that Knots are perfectly good models for knots it does not follow that they are good models for *cognitive* representations of knots” (BEA2: 96). Yet, despite the occasional glimmer of recognition, the general thrust of BEA2 rather suggests this point to be one the authors all too easily forget. Indeed, as best we can tell, their core thesis apparently reduces to the claim that you can infer back *from* the computational equivalence of artefacts *to* the cognitive equivalence of these artefacts:

In a series of papers we have been developing a proposal for a novel methodology to ‘read’ the archeological record [...]. Our proposal is based on the idea that a formal analysis of the material remains left by our ancestors may prove useful in determining the kinds and amount of cognitive resources deployed to produce such objects, in other words, that manufactured objects are transparent with respect to the biological structures underlying the processes necessary to produce them. By performing such an analysis [...] one is capable of inferring the computational complexity of the said cognitive tasks and to advance hypotheses concerning the architecture of the mind capable of performing them. (BEA2: 79–80)

According to the position we have been outlining in this section, however, this is precisely what researchers are not capable of inferring. After all, given that the relevant ‘manufacturers’ have long since passed, to do so would be to take the comp-properties of an artefact as strong evidence for their cog-properties *independently of the possibility of any direct cognitive analysis*. But, since the comp-properties of objects need bear no relation to their cog-properties, independently of any direct cognitive analysis there are scant grounds for considering the comp-property equivalence of objects to be much kind of evidence at all. As such, when it comes to the archaeological record, there is simply no transparency to be found at all; and to claim otherwise could only be to assume that one can switch back

and forth between the two domains without supplying the specific reason needed to justify each specific switch. To equivocate, in other words.

Perhaps most emblematic of this equivocation, however, for us at least, is the euphoric note on which BEA2 end, the authors presenting an analysis whereby “Knot recognition is reduced to a language recognition problem as required by computational complexity theory” (BEA2: 103). Strikingly, it is following on from the apparent success of this analysis, that they state:

On the whole, and considering the different kinds of data we have presented here, it seems likely that natural computational systems, knots, and language do not define such a bizarre love triangle after all as pretended by Lobina. (BEA2: 104)

Unfortunately, what seems to have escaped BEA2’s attention is the fact that the string language analysis preceding this bullish conclusion is still entirely couched within the terms of formal language theory; that is, it deals only with ‘Knots’ and ‘Language’ rather than ‘knots’ and ‘language’.¹⁹ As such, it directly relates only to certain comp-properties of knots and language, having no necessary bearing at all on their cog-properties and, hence, their substantive relationship according to a theory of cognition. The latter still remains something that must always be independently demonstrated to hold at the *cognitive* level. Unless, of course, you happen to have some concept that allows you to arbitrarily inter-define the two domains within the same sentence; a natural computational system, say.

4.2. *The Problem of the Missing Link*

If the previous section’s line of reasoning is in any way proceeding along the correct path, then, BEA2 would, in fact, seem to be operating under the mistaken belief that theory of computation equivalencies between artefacts are *prima facie* grounds for considering these artefacts to also be equivalent at the cognitive level, this the result of an unfortunate equivocation between the two domains that we take to be embodied in the authors’ own notion of an NCS. As further support for our argument, however, we believe this equivocation can yet be highlighted another way. For, if comp-properties and cog-properties are in principle distinct,²⁰ it becomes criterial that proper consideration be given to the manner in which they can actually be substantively related. As such, it cannot be enough to show that certain artefacts have certain comp-properties in common and that they also have certain cog-properties in common. Rather, it must also be demonstrated that it is *because* of these cog-properties that the comp-property equivalencies hold.

We believe this to be an especially important point for BEA2 to grasp, at least if their thesis is to actually go through, since it is surely not the case that knots and natural language utterances have whatever cog-properties they do be-

¹⁹ The capitalisation of ‘Language’ here is simply intended to reference language as defined in theory of computation terms, on the analogy of ‘Knots’ and ‘knots’.

²⁰ Something, recall, that BEA2 would themselves seem to believe: “Of course, from the fact that Knots are perfectly good models for knots it does not follow that they are good models for *cognitive* representations of knots” (BEA2: 96).

cause of the specific comp-properties they do, but that they have whatever comp-properties they do because of the specific cog-properties they do. After all, theory of computation accounts of cognitive phenomena are, ultimately, *models* of such phenomena and, hence, entirely dependent on the actually existing cog-properties for their own reality. Taking the particular case at hand, therefore, what this means is that even if knots and language do turn out to have some comp-properties in common, and even if the artefacts' comp-properties are genuinely relatable to the artefacts' cog-properties, the latter could nevertheless still be entirely distinct at the cognitive level, thereby rendering any computational equivalence completely irrelevant from the perspective of a substantive theory of cognition: They would simply be different things, only coincidentally equivalent at the theory of computation level.

Now, given that BEA2 (p. 80) apparently acknowledge the fact that theory of computation accounts are, indeed, mathematical *models*, this is intuitively something for which they ought to display serious concern. Yet, despite the token acknowledgement, we find little evidence of any genuine concern for this state-of-affairs in the actual paper itself. Hence, once again, the authors' apparent belief that researchers can uncover cognitive facts by a simple perusal of the archaeological record. Hence, also, the switching back and forth between cognitive findings and computational findings, without the authors ever really providing any direct arguments that would serve to justify these highly general switches.²¹ And hence, in particular, their seeking to argue for the cognitive equivalence of knots and language on the basis of their supposed computational equivalence (BEA2: 102–104),²² without actually providing any reason to believe that these equivalencies have the same cognitive base (except, of course, for that handily provided by their own notion of the NCS). Something of a topsy-turvy state-of-affairs, to be sure.

To more clearly demonstrate what we mean by this point, as well as how it specifically relates to BEA2, let us return to the case of language, here considered apart from knots, and taking for granted that the human linguistic competence has an expressive power that is mildly context-sensitive; a comp-property which BEA2 persistently return to (pp. 83, 84, 90, 93, 103). What can we conclude from this? Arguably, two things. The first is that, couched within the terms of the theory of computation, natural language syntax is mildly context-sensitive. The second is that, whatever particular grammatical frameworks linguists are seeking to develop, were we to formalise any of these frameworks and thereby take advantage of the extra precision such formalisms afford, these frameworks should plausibly manifest an expressive power that is mildly context-sensitive.

²¹ Indeed, near as we can tell, what BEA2 effectively present is a somewhat convoluted argument by analogy, the authors providing two entirely separate evidence bases, one relating solely to the cognitive level and the other solely to the computational level, it being left to the reader to magically join up what BEA2 only ever present as distinct dots.

²² Strictly speaking, BEA2 do not even do this, since they never really argue that knots and language are exactly equivalent, as they surely must if they are to give their notion of an NCS any kind of plausibility. Instead, they seem quite content with the mere belief that "Knots (and knots) are complex objects but no [sic] too complex, perhaps sitting in a region of complexity space similar or not too far away from that of language" (p. 103); which seems nothing more than a roundabout way of admitting knots and language to not, in fact, be equivalent.

The first thing, of course, is little more than a truism (eggs *is* eggs, after all). The second, on the other hand, is potentially quite a useful thing to know, since it enables linguists to further evaluate any proposed grammatical framework according to the comp-properties it evinces, thereby constraining the range of frameworks that can be taken as reasonable candidates for modelling the linguistic competence which human cognition instantiates. So, based on such an evaluatory approach, for example, it would seem to be the case that the comp-properties of Head-Driven Phrase Structure Grammar and Lexical Functional Grammar mark them out as apparently too powerful a framework for adequately describing natural language syntax, whereas those of Combinatory Categorical Grammar, Minimalism, or Tree Adjoining Grammar apparently mark them out as more likely ‘just right’ (see Stabler 2010, together with the references therein). Accordingly, linguists attempting to ascertain the nature of human grammatical competence might therefore wish to use this state-of-affairs as reason for focusing less on the former two frameworks, and more on the latter three (or, indeed, any other framework that can be shown to be mildly context-sensitive from a theory of computation perspective).²³

On this account, then, there would indeed seem to be at least one sense in which the comp-properties that are assignable to natural languages can be held to substantively relate to some of the cog-properties that we take to be characteristic of human linguistic competence. This state of affairs, however, only holds in this case because it is of the nature of grammatical formalisms that they bridge comp-properties and cog-properties. That is, whilst the comp-properties follow from the formalism being a formalism, and therefore ripe for treatment in theory of computation terms, the cog-properties follow from the fact that each formalism constitutes a formalisation of a specific grammatical framework, these frameworks specifically motivated in order to account for the criterial features of human linguistic competence. As such, to formalise a particular grammatical framework is perforce to provide a formalisation of the cog-properties of human language and, thereby, to provide a substantive mechanism for linking some of the comp-properties of natural languages to their cog-properties.

Unfortunately for BEA2, however, it is genuinely hard to conceive of any means by which the two sets of properties can be meaningfully bridged other than via the formalisation of a particular grammatical framework. Thus, for example, it is presumably not the case that the mind literally instantiates some infinite store of utterances; rather, what it instantiates is some form of productive competence on the basis of which this particular set of utterances can be generated. Similarly, whilst this competence can plausibly be characterised as generating some infinite set of expressions, what it actually, or “strongly”, generates is not some infinite set of symbol *strings*, so much as a set of structured representations over which the appropriate string set can be abstracted; a set of structured representations, furthermore, which BEA2 (p. 89) themselves admit to be beyond the scope of the theory of computation. In other words, viewed from

²³ Note, of course, that this is still very much only a might, there really being no *requirement* that linguists must take such properties into account, comp-properties being only one of a number of factors that might make a particular grammatical framework attractive from a linguistic point of view.

the perspective of the theory of cognition, the cognitively-relevant comp-properties of human language would seem to be entirely *derivative*. That is, they have no direct cognitive reality in-and-of-themselves, being essentially by-products of the underlying competence, it being this competence which *is* cognitively real and which is actually responsible for the cog-properties that ground the comp-properties of human language. Accordingly, if we are to genuinely establish the kind of NCS link so critical for BEA2, this can really only be accomplished by first establishing a direct account of the aforementioned competence. And since such an account cannot apparently be provided by the theory of computation (something which, just to reiterate, BEA2 (p. 89) themselves seemingly admit to be the case), our only recourse would seem to be the grammatical frameworks that are the focus of professional linguists because it is these frameworks which are expressly developed in order to directly account for the human grammatical competence and these frameworks which are capable of providing the kind of structural descriptions needed to properly model said competence.

Assuming the above-argued state of affairs to be in any way accurate, therefore, it would seem to be just these grammatical frameworks which constitute the requisite locus of description necessary for substantively bridging the comp- and cog-properties which are instantiated by the human competence with respect to natural language; call this level of description, for ever-persistent want of imagination, the *grammatical* level. As such, it seems clear that there is no real sense in which BEA2 can plausibly seek to legitimate the cog-property equivalence of knots and language on the basis of theory of computation results, even if such results may actually turn out to have a genuine cognitive basis. Rather, for the authors to actually make good on their claim regarding the NCS equivalence of knots and language, what they must be doing is demonstrating that this claim holds at the all important grammatical level. Unfortunately, this sort of argument is quite clearly a very different one from that which BEA2 seem interested in providing.

Indeed, to see this, one need only consider some of the grammatical frameworks on which various linguists are currently working; those of Culicover & Jackendoff (2005), Sag *et al.* (2003), and Steedman (2000), to name but three. For what even a moment's such consideration amply demonstrates is that these frameworks are directly motivated by the need to account for such criterial and highly specific features of natural language syntax as agreement, binding, constituency, dependencies, displacement, grammatical functions, scope, the selectional properties of lexical items, etc; not to mention the rather obvious requirement that linguistic expressions must somehow be such as to be both semantically interpretable and phonologically expressible. Now, since it is these properties which the various frameworks are expressly designed to capture, on our account it can only be these properties which result in the cognitively-relevant comp-properties of human linguistic competence. Hence, in order to even begin arguing for the existence of any kind of substantive NCS link between natural language and another human artefact such as knots, what BEA2 would actually have to argue is either that knots demonstrate similar criterial features to those of natural language syntax or that any of the highly specific (and, prefer-

ably, mildly context-sensitive) grammatical frameworks which linguists posit in order to account for these linguistic features can also be plausibly thought appropriate for handling the criterial features of the human competence with respect to knots (whatever these criterial features might actually turn out to be). So, taking the case of Combinatory Categorical Grammar as a particular case in point, BEA2 would have to show either that it makes sense to model knot-competence using the combinatorial framework postulated by, for example, Steedman (2000), or that the linguistic features which motivate this account have any directly equivalence with respect to those features which serve to cognitively ground said knot-competence.

Perhaps needless to say, we remain rather sceptical about even the principled possibility of such a demonstration. After all, it seems hard to conceive of any meaningful way in which human knot-competence could legitimate a treatment in terms of bluebirds, starlings, and thrushes. Regardless of the outright difficulty of such an approach, however, what is important for our purposes is simply that BEA2 nowhere attempt to mount any kind of argument at what we have called the “grammatical” level. Rather, the only comp-property accounts of knots and natural language they do provide are precisely those framed in either time/space terms (pp. 91–101) or language recognition terms (pp. 102–104), pure theory of computation accounts which make no reference at all to the “grammatical” relationships which might exist between knots and language and which would provide a genuine bridge between the cognitive and the computational. In fact, BEA2 are quite open about their beliefs here, explicitly following a claim about the complexity of natural language, for example, with the statement that “this inference is entirely independent from any consideration concerning parsing, *choice of grammatical formalism*, or any other architectural or formal consideration about performance models” (p. 93, our emphasis). What they seem to be arguing, in other words, is that you don’t need any kind of grammatical framework to explore the computational properties of natural language.

Viewed purely from the perspective of the theory of computation, of course, this is quite possibly right (though even here this is perhaps still a rather limiting position to take). Viewed from the perspective of the theory of cognition, at least as we see it, however, this statement is highly misleading, and emblematic of the equivocation we take to underlie BEA2. For what the statement ignores is the critical fact that, to demonstrate any kind of substantive link between the computational and cognitive properties of human artefacts, this demonstration can only be unpacked at the grammatical level. For it is grammatical frameworks that serve to directly model human linguistic competence, and ultimately these frameworks which serve to ground any cognitively-relevant computational properties that human artefacts might have. Otherwise, all you are left with are some cog-properties and some comp-properties of one type of artefact and some cog-properties and some comp-properties of another type of artefact, with no means for substantively relating these properties. Which is perhaps just another way of noting that all BEA2 actually offer are hypothesised comp-property relationships which, as presented, are purely coincidental when approached in terms of the substantive reality of human cognition. Unless, of course, you happen to have some notion that

enables you to arbitrarily inter-define two domains within a single concept, allowing you to move from the computational to the cognitive as if it were the former that grounded the latter; a natural computational system, say.

4.3. *As Sure as Eggs is X*

And were a second strain of reasoning not enough, BEA2 further oblige us to consider a third means by which the equivocation at work in the paper can be highlighted. Thus, it is a central point of their paper that computational systems can be distinguished in terms of their ‘computational power’, the latter understood as the memory resources “a computational device has at its disposal” (BEA: 82). In fact, their NCS is described as being composed of a “very conservative core engine” (*ibid.*), of which we are told very little indeed, and the all-important working memory device (*ibid.*). As such, according to BEA2, structural differences among computational systems follow “from the constraints on the working memory space the device has as its disposal to perform the computation” (p. 81). This is true enough for the means BEA2 have chosen to model their computational system — namely, those automata that formal language theory studies, and which constitute instances of a Turing Machine — but it is striking that the structural differences they seem so intent on outlining have so little to do with the underlying mechanism (or any of the operations) with which a computational system is usually identified. One could, after all, draw a distinction between, say, Kleene’s (1943) partial recursive functions and untyped versions of the lambda calculus, or between any of the latter two and a Turing Machine, or between any of the latter three and Post’s (1943) production systems; and doing so is to focus on the *intensional* differences among these systems.

BEA2 are sure to remonstrate that all these systems are extensionally equivalent — that is, they can all generate the same input-output pairs — and so their internal differences are not that important. That is a fact about their comp-properties, but as stated earlier, it is not the infinite set of input-output pairs that should preoccupy the cognitive scientist, but the intensional properties of a computational system — its cog-properties. After all, to determine that this or that automaton can recognise this or that formal language is to *specify* this or that formal language, and this can just as well be done with a string-rewriting system, thereby downplaying (actually, eliminating) any role memory resources may have.

So why are BEA2 so keen on the memory resources a computational system has access to? This is in fact hard to determine, but Lobina (2012b) did point out that both C&U and BEA1 made reference to Uriagereka (2008: Chap. 7), wherein it was defended that the Chomsky Hierarchy, qua a ranking of production systems, had so far only modelled the weak generativity (string generation) of grammars. It was further supposed by that author that ‘re-interpreting’ the Hierarchy in terms of automata provided an account of strong generativity (structure generation), the only construct of some relevance for linguists. The connection between automata (including their memory resources) and strong generativity is of course a false one, something which Uriagereka (2008) seems to not have fully

grasped given his apparent confusion of the memory resources of an automaton with psychological models of memory (and, by extension, with the structural properties of the representations so manipulated). Nevertheless, though BEA2 (p. 89) explicitly state themselves to be fully aware of the falsity of such connection, is it really the case that they are themselves so free of Uriagereka's (2008) equivocation?

BEA2 (p. 83) assure us that their NCS is neither a psychological model of memory nor a performance model; rather, it is an abstract characterisation of a model of computation, a formulation they consider to be similar to that of the faculty of language in the narrow sense of Hauser *et al.* (2002; cited therein).²⁴ However, when it comes to listing the evidence for their NCS — and this evidence is of two kinds, either clinical (*viz.* to do with cognitive disorders) or neurological, all to be found in pp. 85–86 — BEA2 seem oblivious to the fact that the data they provide are the result, as we have stressed above, of cognitive resources that include both the underlying computational system *and* whatever systems partake in performance, including, naturally, real-time memory.

The equivocation between formal and psychological models of memory is clearest when BEA2 consider the neural substrate of their NCS, as they reference the respective models of Lieberman & Ullman (pp. 85–86; cited therein), both of which quite explicitly outline a psychological, rather than a formal, model of memory. According to BEA2, Lieberman proposes a sequencer (perhaps the conservative core engine BEA2 advertise?) and a *working* memory in order to account for our ability to *process* symbolic elements (p. 85), while Ullman hypothesises about the location of *procedural* memory (p. 86). Apparently, both accounts are “ultimately conceived as to subserve the *learning* and *execution* of diverse tasks” (p. 86, our emphasis).

Can BEA2, therefore, really believe that their ‘abstract’ working memory construct is analogous to the working or procedural memory hypothesised in most of cognitive psychology? Are they really unaware that the cognitive data they selectively reference are informative of cognitive resources that must include, surely, a psychological model of memory in addition to whatever computational system underlies whatever cognitive skill? That by employing automata theory one is merely specifying formal languages? That any supposed distinction among computational systems in terms of memory access is no more than a result of the chosen formalism, and therefore not a genuine distinction? That, in any case, automata can only model weak generative power and therefore are irrelevant for the study of cognition? If BEA2 are indeed privy to all this, how is it they have managed to cleave so persistently to their cognitive tale? Is it because they have found themselves able to postulate the state of affairs to be thus? In terms of an NCS, say?

4.4. Coda

Equivocation, then; and persistently so. According to our analysis, therefore, it is

²⁴ BEA2 also draw a connection between their NCS and whatever computational system Fitch & Hauser (2004; also cited therein) were in fact studying. We won't engage this issue here, but it seems to us that these similarities are more than a little exaggerated.

ultimately to this equivocation which can be traced BEA2's thoroughgoing (and otherwise perplexing) failure to engage with the nub of Lobina's (2012a) original criticisms. Indeed, taking the accuracy of our analysis at face value, this failure would hardly seem surprising, the very possibility of genuine and substantive engagement essentially ruled out *ab initio*, the entire edifice of BEA2 resting on a flawed conceptual foundation which has enabled its authors to arbitrarily confound distinct conceptual domains and present a grab-bag of disparate information as if it constituted a substantive evidence base. As such, BEA2 (pp. 84, 89) would seem critically mistaken when they assert their thesis to be one that is empirically testable. On the contrary, there is really nothing there with the kind of conceptual coherence necessary to even begin getting a proper purchase on the world.

Having paused for summary and breath, we find ourselves in something of a quandary, feeling strongly that there is yet more to be said by way of response to BEA2. Thus, we could take further issue with the specific pieces of evidence provided by our contending authors. So, for example, we could discuss the fact, pointed out to us by Mark Steedman (p.c.), that their example 3 (p. 103) is actually in the linear context-free rewriting class rather than in the triple copy class, and that there is scant reason to believe it bears any relation to their example 4 (p. 103). Or we could discuss the Herzfeld & Lestel (2005) study that BEA2 reference (p. 97); a study which, contrary to BEA2's interpretation, actually offers up evidence of knot-tying in apes, raising the inevitable question as to why, if apes *can* knot and *if* knotting has the expressive power of natural languages, apes have not so far been found to have a capacity for natural language syntax. Or we could discuss the apparent contradiction in using knot-learning evidence by way of support for their thesis (p. 97), when they themselves attempt to immunise this thesis from criticism by explicitly stating that it involves no "focusing on learning capabilities" (p. 83).

Thus, we could take issue with BEA2's (p. 88) mentioning of a "thesis" that was supposedly attributed to them by Lobina (2012b), namely that there is a "processing competence" which composes meaningful expressions, and which is additionally connected to "rich, contentful, language-like thoughts" (p.88). This is a thesis which they are apparently able to doubt on account of some supposed problems with adaptationist explanations of the theory of evolution. So, we could note that we are not sure what the term "processing competence" is supposed to refer to, but that we are certainly sure that Lobina (2012b) didn't ascribe any such thesis to them (unsurprisingly, they don't offer a page reference). We could also note that, in the greater scheme of things (there is life beyond C&U and BEA1, after all), Lobina (2012b) was seeking to discuss the relationship between language and thought, and it was therein argued that a rich conceptual representational system must be postulated if the acquisition of language is to be at all possible, a belief the present authors actually hold themselves (pp. 87–88). And we could note that however this might pan out for a theory of evolution is something we neither know nor care much about, but that we are definitely amused that BEA2 feel able to reference Fodor & Piattelli-Palmarini (2010; cited therein) as an authoritative critique of adaptationism.

Thus, we could also take time to more fully discuss the further implications

of our reasoning for BEA2's claims regarding the particular NCS that human cognition supposedly instantiates. So, for instance, we could point out that if, as we have essentially argued in section 4.2, language has whatever expressive power (or comp-properties) it does because of the syntactic structures (or cog-properties) it does, then these (let us assume) mildly context-sensitive properties are ultimately the result of the human linguistic competence taken *as a whole*, a competence which is as much an artefact of the linguistic representations that are operated on as it is of the underlying system that does the operating. Hence, it really makes little sense to speak of the human NCS, in the sense of a "core engine" (BEA2: 80), as being itself mildly context-sensitive (p. 83), since it is not this engine, considered in isolation, which gives language its overall expressive power.²⁵ As such, should human cognition genuinely turn out to instantiate some domain-general computational system, which is presumably what BEA2 mean by their "functionally unspecific" device (BEA2: 84), then it is more than likely that this computational system will not behave uniformly with respect to the various cognitive domains over which it operates. After all, these domains will essentially comprise different systems and data sets, which will in turn require the computational system to carry out distinct computations according to each particular domain's own particular requirements, computations which should thereby result in distinct types of cognitive artefact, each with their particular expressive power. As such, it is perfectly conceivable that the (let us assume) mildly context-sensitive expressive power of human language will be of no import whatsoever for the study of other cognitive domains. Indeed, in such a situation, it would be quite literally meaningless to speak of human cognition as instantiating a natural computational system characterised by a particular expressive power, because such a system could not be meaningfully said to have any such particular power of its own.

And so on, and so forth. Suffice it to say, there is a lot more that we would like to have said.²⁶ Nevertheless, being aware that a proper treatment of these points and issues would require more space than is likely reasonable in terms of the present paper, and being unwilling to try the patience of the reader any further, such a treatment is no doubt best left for a more appropriate context; the addendum to a certain forthcoming book, perhaps...

5. An Old-World Apology, or Thereby Hangs a Tale

Almost all being said and done, there remains one final point that we feel does need addressing; namely the *ad hominem*, levelled against at least one of us, of "formal bullying" (BEA2: 104).²⁷ This is a strong claim, one that at least both of us

²⁵ To put this another way, it is language, not some domain-general NCS, that gives language the particular expressive power it seems to have.

²⁶ So, for example, we haven't even bothered to mention BEA2's (p. 98) below-the-belt charge of behaviourism, a mischaracterisation that again only serves to highlight the general inadequacy of their own conceptual framework (which apparently mismeasures 'trial-and-error' learning with 'stimulus-response' learning, something which is likely the result of their inability to distinguish between claims regarding mental architecture and claims regarding performance).

²⁷ An accusation we take to have been levelled with a certain amount of irony, originating in a

feels to be essentially unwarranted; for the simple reason that it is, essentially, unwarranted.

True, the spirited form of Lobina (2012a, 2012b) may not have suited all tastes; but, read thoughtfully, there really is nothing in the actual substance of these papers tantamount to “bullying”. In particular, neither Lobina constituted some churlish refusal to engage with the matter at hand. Rather, sustained criticism was put forth in an attempt to substantively address the specific points and general claims of C&U and BEA1. Surely, to so criticise is not to bully. Nor, surely, is it bullying to point out lack of understanding if either (a) in general, there does indeed seem to be such lack of understanding, or (b) as specifically written, C&U and BEA1 can reasonably be argued to demonstrate such lack of understanding. This is what Lobina (2012a, 2012b) undertook to argue, and nothing in BEA2 suggests either that the original criticisms were ill-founded or that the situation has been substantively improved; at least not to us. Hence, the present paper.

And, just to be absolutely clear, no opposition to the *principle* of the endeavours represented by BEA2 *et al.* is to be found anywhere herein.²⁸ Indeed, in this sense, we are entirely in agreement with the spirit of the authors’ attempts to advance new methodologies, methodologies that would enable fresh evidence to be uncovered and productively pursued; such undertakings are commendable. We simply disagree with the substance of their specific proposal, in its present form, and fail to see how querying this proposal is in anyway unproductive. Unless, of course, there is some sense to be had in cleaving to something that cannot claim to do what it sets out to do.

Here it is worth pointing out the assumption on the part of BEA2, and it is purely an *assumption*, that their framework advances our understanding; or at least has the potential to advance our understanding (p. 104). What Lobina (2012a, 2012b) was at moderate pains to point out, however, and what we have sought to reargue here, is that it is far from clear that the proposed framework actually does advance our understanding, being apparently based on false analogies, tenuous evidence, and dubious interpretation. Most fundamentally, there is clearly no virtue, in-and-of-itself, to a priori define one distinct conceptual domain in terms of another, whether or not said definition strikes some mysterious chord, and whether or not a particular set of authors are able to dress up their argument in superficially impressive formalities. After all, the history of linguistics is littered with such dead ends, false prophecies that muddy more than clarify and entangle more than merge. These prophecies are hardly surprising, there being no doubt an immense aesthetic appeal to find that language originated in the croaking of frogs or the thunder of Jove (Brisset 2001, Vico 1744/1948); and, well, you know, it sort of kind of looks like it does, you know, assuming, of course, that you are able to look at it in the ‘right’ way. Unfortunately, what stands to ‘right’ reason does not always stand to reality.

paper composed by five established academics for the specific purpose of critiquing not more than one of us. After all, for someone to be able to bully, one would have thought that they would first need to be in some actual position of power...

²⁸ Nor, of course, is there to be found anywhere herein any general opposition to, or dismissal of, the theory of computation as taken on its own terms.

And so it was, for example, that the statutes of the Linguistic Society of Paris included the well-known 1866 moratorium on evolutionary talk, explicitly recognising the highly speculative nature of that particular enterprise as it stood at the time.

If there is one underlying motivation with respect to the present paper and its two forebears, therefore, it is perhaps that the 1866 moratorium was issued with good sense, and that certain linguistic work with a biolinguistic/evolutionary flavour ought to take the spirit of that moratorium very much to heart, explicitly recognising the highly speculative nature of the enterprise as it currently stands. Not that we wish to dismiss outright any such work or demand the literal issuing of any such moratorium. After all, the familiar history of early twentieth century research into language and cognition demonstrates the pitfalls that easily arise through such *a priori* dictats. Rather, we make the simple point, easily forgotten in all the speculative excitement, that if linguists are to genuinely establish and cash out an apt biological/evolutionary framework for understanding human language, it will be important to proceed both thoughtfully and *critically*; not least because it is not especially clear what or where the relevant evidence base will turn out to be, or even what such a framework might itself actually mean given the highly interdisciplinary requirements of the task. As such, we take the three critiqued papers to represent a telling cautionary tale. For, if it is true, as has been remarked, that “[t]here is no end to plausible storytelling” (Lewontin 1998: 129), then surely it is even more true that we first have a plausible story to tell.

References

- Balari, Sergio, Antonio Benítez-Burraco, Marta Camps, Víctor M. Longa & Guillermo Lorenzo. 2012. Knots, language, and computation: A bizarre love triangle? Replies to objections. *Biolinguistics* 6, 79–111.
- Balari, Sergio, Antonio Benítez-Burraco, Marta Camps, Víctor M. Longa, Guillermo Lorenzo & Juan Uriagereka. 2011. The archaeological record speaks: Bridging anthropology and linguistics. *International Journal of Evolutionary Biology* 2011, doi:10.4061/2011/382679.
- Berwick, Robert C. & Amy S. Weinberg. 1989. *The Grammatical Basis of Linguistic Performance: Language Use and Acquisition*. Cambridge, MA: MIT Press.
- Bever, Thomas G. 1970. The cognitive basis for linguistic structures. In John R. Hayes (ed.), *Cognition and Language Development*, 279–362. New York: Wiley-Blackwell.
- Brisset, Jean-Pierre. 2001. *Œuvres complètes*. Dijon: Les Presses du Réel.
- Camps, Marta & Juan Uriagereka. 2006. The Gordian knot of linguistic fossils. In Joana Rosselló & Jesús Martín (eds.), *The Biolinguistic Turn: Issues in Language and Biology*, 34–65. Barcelona: Publications of the University of Barcelona.
- Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions of Information Theory*, IT-2, 113–124.
- Chomsky, Noam. 1963. Formal properties of grammars. In R. Duncan Luce,

- Robert R. Bush & Eugene Galanter (eds.), *Handbook of Mathematical Psychology*, vol. 2, 323–418. New York: John Wiley.
- Culicover, Peter J. & Ray Jackendoff. 2005. *Simpler Syntax*. Oxford: Oxford University Press.
- Frazier, Lynn. 1985. Syntactic complexity. In David R. Dowty, Lauri Karttunen & Arnold M. Zwicky (eds.), *Natural Language Processing: Psychological, Computational, and Theoretical Perspectives*, 129–189. Cambridge: Cambridge University Press.
- Gibson, Edward. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68, 1–76.
- Hass, Joel, Jeffrey C. Lagarias & Nicholas Pippenger. 1999. The computational complexity of knot and link problems. *Journal of the ACM* 46, 185–211.
- Hawkins, John A. 2004. *Efficiency and Complexity in Grammars*. Oxford, England: Oxford University Press.
- Herzfeld, Chris & Dominique Lestel. 2005. Knot tying in great apes: Ethnology of an unusual tool behaviour. *Social Science Information* 44, 621–653.
- Hildebrand, Caz & Jacob Kenedy. 2010. *The Geometry of Pasta*. London: Boxtree.
- Kleene, Stephen Cole. 1943. Recursive predicates and quantifiers. In Martin Davis (ed.), *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvability Problems, and Computational Functions*, 254–287. Mineola, NY: Dover.
- Lewontin, Richard C. 1998. The evolution of cognition: Questions we will never answer. In Don Scarborough & Saul Sternberg (eds.), *An Invitation to Cognitive Science*, vol. 4: *Methods, Models, and Conceptual Issues*, 107–132. Cambridge, MA: MIT Press.
- Lobina, David James. 2012a. All tied in knots. *Biolinguistics* 6, 70–78.
- Lobina, David James. 2012b. Conceptual structure and the emergence of the language faculty: Much ado about knotting. Ms., Tarragona: Universitat Rovira i Virgili. [<http://ling.auf.net/lingBuzz/001397>]
- Michel, George F. & Debra Harkins. 1985. Concordance of handedness between teacher and student facilitates learning manual skills. *Journal of Human Evolution* 14, 597–601.
- Miller, George A. & Noam Chomsky. 1963. Finitary models of language users. In R. Duncan Luce, Robert R. Bush & Eugene Galanter (eds.), *Handbook of Mathematical Psychology*, vol. 2, 419–492. New York: John Wiley.
- Mount, John. 1985. *KnotEd: A Program for Studying Knot Theory*. [<http://mzllabs.com/JohnMount>] (27 January 2012)
- Phillips, Colin, Matthew W. Wager & Ellen F. Lau. 2011. Grammatical illusions and selective fallibility in real-time language comprehension. In Jeffrey T. Runner (ed.), *Experiments at the Interfaces*, 147–180. Bingley: Emerald Group Publishing.
- Post, Emil Leon. 1943. Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics* 65, 197–215.
- Pratt-Hartmann, Ian. 2010. Computational complexity in natural language processing. In Alexander Clark, Chris Fox & Shalom Lappin (eds.), *The Handbook of Computational Linguistics and Natural Language Processing*, 43–73. Oxford: Blackwell.

- Sag, Ivan A., Thomas Wasow & Emily M. Bender. 2003. *Syntactic Theory: A Formal Introduction*, 2nd edn. Stanford, CA: Centre for the Study of Language and Information.
- Stabler, Edward P. 2010. Computational perspectives on minimalism. In Cedric Boeckx (ed.), *Oxford Handbook of Linguistic Minimalism*, 616–641. Oxford: Oxford University Press.
- Steedman, Mark. 2000. *The Syntactic Process*. Cambridge, MA: MIT Press.
- Townsend, David & Thomas G. Bever. 2001. *Sentence Comprehension: The Integration of Habits and Rules*. Cambridge, MA: MIT Press.
- Turing, Alan M. 1954. Solvable and unsolvable problems. In B. Jack Copeland (ed.), *The Essential Turing*, 576–595. Oxford: Oxford University Press.
- Uriagereka, Juan. 2008. *Syntactic Anchors*. Cambridge: Cambridge University Press.
- Vico, G. 1744 [1948]. *The New Science* (trans. T. G. Bergin & M. H. Fisch). Ithaca, NY: Cornell University Press.

David J. Lobina
Universitat Rovira i Virgili
Departament de Psicologia
Centre de Recerca en Avaluació i Mesura
de la Conducta
Ctra. de Valls s/n
43007 Tarragona
Spain
davidjames.lobina@urv.cat

Mark Brenchley
University of Exeter
Graduate School of Education
St Luke's Campus
Heavitree Road
Exeter
EX1 2LU
United Kingdom
mdtb201@exeter.ac.uk

On the Feasibility of Biolinguistics: Koster's Word-Based Challenge and Our 'Natural Computation' Alternative

Sergio Balari, Cedric Boeckx & Guillermo Lorenzo

1. Introducing the Challenge and Our Motivation

The present paper examines Jan Koster's "skeptical view on Biolinguistics and linguistic internalism" (Koster 2009: 61),¹ and concludes on a far more positive note than he does regarding the prospect of this emerging field. Examining Koster's challenge also gives us the opportunity to point that while it is important to remember, as Jackendoff (2010) stresses, that thinking about the biology of language (e.g., its evolution) 'depends on one's view of language', it is equally important to bear in mind that thinking about the biology of language also depends on one's view of biology. We think that this point is worth emphasizing at a time when both modern linguistics and biology are re-examining their foundations.

Let us begin by stating Koster's argument in a nutshell. In order for languages to be acquired and used, Koster (hereafter, K) agrees, certain uniquely human biological requirements are required. However, so K's argument goes, in as much as human biology is not 'transparent' with respect to its role in language, the idea of translating these biological underpinnings into a distinctive mental faculty (often called 'the faculty of language', FL, or 'the language organ') makes no more sense than positing distinctive faculties for human activities such as trumpet playing or bicycle riding, or (to invoke distinctively human anatomical sites) hat-wearing or glasses-supporting systems. It is K's contention that this 'Panglossian' drift of modern linguistics (we have a language faculty in order to support languages, much like we have a nose to support glasses) is the historical consequence of generative grammarians having uncritically adopted a series of conceptually problematic and empirically unwarranted compromises since Chomsky (1957), the latest outcome of which is the chimerical discipline now

This work has been partially supported by the Spanish Ministerio de Ciencia e Innovación through research grants FFI2010-14955 (with FEDER funds) (Balari and Lorenzo) and FFI2010-20634 (Balari and Boeckx), by the Generalitat de Catalunya through research grant 2009SGR1079 to the Centre de Lingüística Teòrica of the UAB (Balari), and by the European Union through the Marie Curie International Reintegration Grant PIRG-GA-2009-256413 (Boeckx), as well as research funds from the Fundació Bosch i Gimpera (Boeckx).

¹ For a complete presentation of his argument, see Koster (2009). Partial accounts of it are also found in a number of shorter papers from 2005 to 2010, accessible from Koster's webpage (<http://odur.let.rug.nl/~koster>).



known as Biolinguistics, devoted to uncovering the biological foundations of said language faculty.

We, as advocates of Biolinguistics, think that is important to meet K's challenge. After all, K has had a distinguished career as a theoretical linguist, and is deeply familiar with the issues and practice of theoretical, 'Chomskyan' linguistics. In addition, K's is a sophisticated argument constructed upon several interesting premises that are worth thinking about, and to which we now turn.

2. Language is Words, Words, and Words

K's first premise is that language is not within our brains, but within our words. This is not an in itself invalidating argument against Biolinguistics, as words could still happen to be natural, biological objects (within our brains), but it nevertheless is, according to K, a first necessary step in order to dismantle the whole biolinguistic project. Let's see why.

K argues that Biolinguistics is constructed upon the belief that the human brain contains a system specifically devoted to computing linguistic expressions, which can be abstracted away from other peripheral components of FL and which is furthermore considered the site of one of the most distinguishing features of language: namely, 'recursion' (unbounded nested-embedding) — i.e. as in Hauser *et al.* (2002). But such a belief is, according to K, nothing more than a historical residue of Chomsky's (1957) thesis on the 'autonomy of syntax', an idea that K takes to have been *de facto* rejected in Chomsky's subsequent works with the adoption of the 'lexicalist hypothesis' (Chomsky 1970), and to which we will come back.

K's own contention regarding this issue reads as follows: Linguistic recursion is not a property of an autonomous system of computation, but a property of words, because, according to the lexicalist position that K endorses, it is in the words' content that instructions are encoded as to how they are to be combined — for example, by means of a structure headed by a word of the same type.

It is obvious that in order for the resulting structures to be full-fledged, a rather powerful computational space is required. However, according to K, such a space is just a biological substrate that 'facilitates' the completion of the properties of words (such as linguistic recursion). Moreover, K takes such a space not be 'transparent' with respect to these properties, meaning that the biological substrate would remain 'ignorant' of the facilitated properties were it not for the fact that humans have historically given it such a linguistic functionality. So K's conclusion is that the human brain does not incorporate a system in charge of computing linguistic expressions, but a general-purpose computational space, on which linguistic functionality simply rides, thanks to words. Because K views words themselves as inherently cultural symbols, as opposed to natural, biological units, there cannot be any proper field of study devoted to the biology of language.

3. Decomposing and Deconstructing K's Arguments

Let us note that K's view is an updated version of an idea with a long-standing

pedigree that K traces back to Sapir (1921), but as a matter of fact it is explicitly articulated as well in Whitney (1875), where the contention is already made that language is just a matter of having (culturally) discovered how to put into a derived or secondary use natural resources inherently unrelated to it. It also was Saussure's (1916) idea, who expressed in very similar terms the relation of language to Broca's region, as well as the official position of European functionalism — as witnessed, for example, in Martinet (1960). For K, departing from this venerable tradition would require more compelling arguments than those so far adduced by biologically oriented linguists.

Historical and traditional considerations aside, the substantive part of K's argument actually splits into two different theses, the first one having to do with the centrality of words, and the second one with the non-specificity of the system subserving the computation of linguistic expressions. We will now try to show that none of these arguments is compelling enough as to support K's anti-naturalist stance on language.

3.1. *Lexicalism vs. Lexicocentrism*

For purposes of K's argument, we can define lexicalism as the position according to which grammars to a great extent have the forms that they do thanks to the instructions encoded in words, contained in their lexicons. As we already stated above, the reason why K believes that lexicalism is such a problematic, indeed lethal aspect for any biolinguistic project is that words are inherently cultural, not natural/biological entities: Words are “man-made, public cultural objects” (Koster 2009: 66). Accordingly, if lexicalism is assumed and the combinatorial properties of language are taken to depend on properties of words, such combinatorial properties will have to be traced back to cultural, not natural/biological attributes.

But how true is lexicalism? And how essential are words?

For K, as we saw, they are pretty much everything. As he writes on his website (see fn. 1), where he summarizes his view: “Invented words rather than syntax are at the essence of language in this view, while recursive syntax is seen as a successful extension of the properties of the cultural objects in question (‘words’). In other words, for K, “[s]yntactic structures are not generated by lexicon-independent rules (like phrase structure rules or Merge) but as the spelling out of the contextual properties of lexical items (‘valency’)”.

Other linguists, too, ascribe an essential role to words (though they do not conclude from this that Biolinguistics is a doomed enterprise). Here is a representative quote from Pinker and Jackendoff (2005):

We now come to an aspect of language that is utterly essential to it: the word. In the minimal case, a word is an arbitrary association of a chunk of phonology and a chunk of conceptual structure, stored in speakers' long-term memory (the lexicon) [...]. [W]ords have several properties that appear to be uniquely human [...]. Our assessment of the situation is that words, as shared, organized linkages of phonological, conceptual, and grammatical structures, are a distinctive language-specific part of human knowledge [...]. [A] good portion of people's knowledge of words (especially verbs and functional morphemes) consists of exactly the kind of information that is

manipulated by recursive syntax, the component held to make up the narrow language faculty. (Pinker & Jackendoff 2005: 213–215)

But a growing number of linguists are coming to the conclusion that words² are not distinguished building blocks in syntax or morphology or semantics, and that lexicalism, or as one of us has come to call it, ‘lexicocentrism’ is not the right model for FL. Here are a few representative quotes:³

In this work I have claimed that a word is a morpheme sequence that shows internal cohesion and has independent contribution relative to other morphemes. I have argued that these properties stem from the syntax: although each morpheme is inserted separately into syntactic structure, syntax may cause some groups of morphemes to show the behavior characteristic of words [...]. If two morphemes form a distributional unit of this kind [one that cannot be interrupted — SB, CB & GL] every time they appear together, the two morphemes in question will be perceived as one word. This means that ‘word’ in the non-phonological sense is a distributional concept [...]. Crucially, there is not one single syntactic configuration that underlies all complex words [...]. On this approach, words are not necessarily syntactic constituents [...]. The consequence of these claims is that words do not really have a place in grammar at all. From the point of view of grammar, ‘word’ is an epiphenomenon, and not a basic concept. (Julien 2002: 321–322)

A Word, as conventionally conceived, is a syntactic constituent which (happens) to correspond to a phonological unit of a given size (e.g., for the assignment of primary stress). While it is likely that there are some universal constraints on what syntactic constituents can correspond to such phonological units, beyond that, the mapping is language specific, and syntactic constituents of equal complexity may or may not be phonological-stress units. Crucially, then, Words are not syntactic primitives or atomic in any meaningful sense. (Borer 2005: 1)

Marantz (2000) adds the following relevant observation:

It’s somehow intuitive to think that knowing a language involves knowing the words of the language. Linguists that start with this notion quickly get into trouble by not being clear about what a ‘word’ is such that a speaker might know it or what ‘know’ is such that a speaker might ‘know’ a word. Jackendoff (1997) argues that the ‘lexicon’ should be extended to include units larger than phrases. But doesn’t the Wheel of Fortune corpus rather argue against the correlation between ‘memorized’ and ‘special linguistic properties’? We know we’ve encountered [Any friend of yours is a friend of mine] just as we know we’ve encountered ‘nationalization’ (with a certain measurable degree of certainty). That means, in some sense, we’ve stored these items — in some way or other. But does ‘storage’ necessarily imply ‘storage in a special linguistic Lexicon’? Jackendoff’s observations call into question the notion that we don’t store information about structures unless the structures have special linguistic properties. None of the examples [he provides] have special structure — none involve special connections between sound and meaning. Rather than arguing for an extended lexicon, Jackendoff is actually arguing that we should abandon the notion of a

² Or even morphemes; see Starke (2010), Boeckx (2010a).

³ We quote the relevant passages in full, as we don’t want to give the impression of constructing a strawman.

'lexicon' (of items with internal structure) entirely.

Jackendoff pulls a fast one on us. He suggests that anyone trying to keep 'fixed expressions' out of the lexicon is trying to keep them out of the language. But, since fixed expressions are made of words (phrases, phonology, etc.), they are clearly part of language. What he fails to argue successfully is that fixed expressions have the sorts of meanings that need to be negotiated by the linguistic system. Knowledge about 'any friend of yours is a friend of mine' is clearly knowledge about a linguistic object — but that linguistic object is constructed via the generative system of the language. (Marantz 2000: 1–2)

True, as Marantz (1997: 201) points out, most contemporary theories of grammar assume a general organization in which elementary constituents are drawn from a place called the 'Lexicon' for composition in the syntax. But when linguistic practice is scrutinized, as Boeckx (2010a) has done, far less than the full array of properties traditionally ascribed to words turns out to be needed. In fact, Boeckx goes so far as to argue that no notion more than the 'edge feature', as defined by Chomsky in the following quote, is needed to reconstruct the essential properties ascribed to the faculty of language in the narrow sense:

For a L[exical] I[tem] to be able to enter into a computation, merging with some [syntactic object], it must have some property permitting this operation. A property of an LI is called a feature, so an LI has a feature that permits it to be merged. Call this the edge-feature (EF) of the LI.

(Chomsky 2008: 139)

If it is indeed true that the edge property is the only lexical property needed to jump start (Narrow) Syntax, words lose the centrality they have in K's argument.

What we have pointed out just now is in fact an old observation, already made by Otero (1976). Consider the following quotes:⁴

Given the theoretical framework Chomsky had developed in [Chomsky (1965)], it is somewhat surprising that he did not go on to draw what, from a generative perspective, appears to be a very natural, if not inescapable, conclusion, namely that morphemic representations play no role in the (syntagmatic) derivation of a sentence.

Otero goes on to formulate the 'Dual Hypothesis', according to which "a grammatical system consists of two major modules: (i) a syntagmatic grammar; (ii) a paradigmatic grammar". Otero notes that this "yields a much improved theory of generative grammar" — "a form of grammar that is conceptually simpler":

[O]ne with fully differentiated but internally homogeneous components. The syntagmatic subsystem consists of a central component (the syntax) and two interpretive components (the phonetics and the semantics). The syntactic component consists of a recursive set of context-free phrase-structure rules and a transformational subcomponent with root transformations, one nonlocal transformation ('move C') and a set of local transformations in the sense of Emonds (to a great extent language particular), which together generate what might be called 'construction forms' (cf. LSLT

⁴ Otero's important study remains unpublished, and the transcript of the talk hard to gain access to. For this reason, we reproduce significant portion of his argument here.

[Chomsky 1975 — SB, CB & GL], §33.1), that is, abstract phrase markers including only syntactic category and subcategory feature specifications [...] The ‘construction forms’ will presumably be enough to derive a ‘logical form’ [...]; a full interpretation can only be derived after the insertion of phonological matrices of words (in the extended sense) from the paradigmatic subsystem.

Otero further notes that:

A syntagmatic grammar is essentially universal (biologically given in essence), while a paradigmatic grammar is, to a considerable extent, a historically evolving subsystem, burdened with the weight of the past, like other cultural systems. Only a paradigmatic grammar can be fossiliferous. This brings to mind the distinction between ‘core grammar’ and a ‘periphery’ of ‘borrowings, historical residues, inventions, and so on’, which we can hardly expect to — and indeed would not want to — incorporate within a principled theory of UG.

Every paradigmatic grammar is, to a considerable extent, language particular, and to some extent fossilized, while the syntagmatic grammar can be assumed to be a fairly direct reflection of the language faculty of the mind/brain [...]. No student of human language ever dreamed of a universal dictionary.

Otero concludes that at the syntagmatic level “there is only one language, as the evolutionary biologist would expect”.

This, we submit, is what makes Biolinguistics possible, a point to which we will return presently when we deal with the issue of transparency. But let us first expand a bit more on the question of lexicalism with an additional observation we believe it may be relevant to complement the ones put forth so far.

To be sure, K could argue that what we have just shown is that lexicalism is a debated issue — something that linguists have been unable to agree upon — and that this doesn’t necessarily invalidate the contention that if lexicalism is true, then Biolinguistics is impossible. To which we could retort that K’s faith in lexicalism doesn’t make it true either and that a closer look at the fate of what we could call ‘radical lexicalist frameworks’ casts serious doubts, if not over the ‘truth’ of the lexicalist hypothesis (‘truth’ being a rather strong word), perhaps over the general viability of the whole program, at least in the form K appears to interpret it.

It is important to note from the outset that K’s main reference for lexicalism is the Government & Binding (GB) framework of Chomsky (1981) and subsequent work, a grammatical theory that always kept a substantial bulk of its grammatical principles away from the lexicon and that was often seen from other quarters (most of them based in California) as relying too much on configurational notions to account for certain grammatical phenomena. A case in point is, for example, Binding Theory (by the way, one of K’s favorites; see, for example, Koster 1987: Chaps. 3 and 6), which was mostly based on the notion of c-command, a structural relation between nodes in a tree which need not both be within the domain of a head word and its dependants, and which was ‘lexicalized’ by, for example, Pollard & Sag (1992, 1994) in an illustrative attempt to remove such principles from the syntax and place them directly in the lexicon, in the internal structure of words.

Our reference to HPSG is not casual, as it is perhaps the grammatical framework that most clearly illustrates the point we want to make here, although similar points could be made with respect to LFG or the various versions of Categorical Grammar (CG). The point is that LFG, HPSG, CG never, ever, assumed that phrase structure could or had to be lexicalized in some way or another — witness LFG's c-structures, HPSG's ID rules (or its two types of signs: words and phrases) and CG operations of functional application and functional composition (plus type raising, which is 'syntactic' not lexical). Now, from the fact that all these frameworks have always needed something more than just words to account for the context-free backbone is, in our opinion, the demonstration that it is not enough with the combinatorial properties of words to get some structure. Structure doesn't come for free and, in fact, assuming this strikes us as a fallacy equivalent to that of claiming that genes are 'replicators', which they aren't, since no replication is possible without all the cellular machinery in charge of actually doing that; in a similar vein, words are not 'combinators', but combinable building blocks in (desperate) need of a combinatorial operation. It is perhaps impossible to lexicalize structure and this perhaps explains, for example, HPSG's steady development in the direction of becoming a version of construction grammar (Ginzburg & Sag 2000, Sag, in press), where the traditional distinction between lexicon and grammar is blurred into a continuum consisting of pure lexical constructions at one extreme of the spectrum and multiword (or 'combinatorial', to use Sag's, in press, terminology) constructions at the other end.⁵

Summarizing, then, perhaps words and groupings of words are some kind of cultural objects, but as such *they are transparent with respect to their underlying biological structures* (call them Unification, Merge, Phrase Structure Rules or whatever operation is responsible of building the structures). This takes us to the issue of transparency and domain-specificity.

3.2. *Transparency and Domain-Specificity*

As for the second part of K's argument, it reduces to the idea that the system that computes linguistic expressions is not transparent with respect to the properties of words *because* it is a system with no inherent specialization, subserving and facilitating the tasks of a wide array of domains — language among them. In K's own words: "My argument is not against innateness but against the idea that biological structures are transparent with respect to their cultural functions, including their role in language" (Koster 2009: 66).

K's argument is subtle and deserves careful examination. Its initial premise is that biological structures are not transparent with respect to their functions, indeed that biological structures are all functionally unspecific. Thus, the

⁵ This is not to be taken as a criticism of construction grammar, as, for the purposes of this paper, we would like to remain agnostic as to what is the best approach to grammatical description. We just find it symptomatic that linguists with a historical strong commitment with lexicalism are abandoning it in favor of other clearly non-lexicalist models. Thus, in addition to HPSG, we could cite the case of Ray Jackendoff, also coming closer to construction grammar (e.g., Goldberg & Jackendoff 2004) or Joan Bresnan, now favoring probabilistic approaches to grammar (e.g., Bresnan 2007).

computational system underlying language, being biological, did not evolve ‘for language’ (remember, words) but only acquired its linguistic functionality once language (i.e. words) was invented; from this we can only conclude that there isn’t anything internal specific to language (which is external), since its current functionality was imposed from the outside, and, thus, Biolinguistics, being concerned with the internal biological structures underlying language, is impossible.

The problem here is that Koster has it backwards. It is certainly true that biological structures are never transparent with respect to their functions (cultural or otherwise), expecting the contrary would constitute a natural theological assumption that was untenable for most biologists even before Darwin (e.g., Owen 1849). Indeed, the process of acquiring one (or more) functions is a historical one, a dialectic between the formal properties of the biological structure and several environmental factors. From this, however, it doesn’t follow that form and function are entirely decoupled, as K seems to suggest, but rather that functions, behaviors, etc., *are transparent with respect to the biological structures underlying them*, and that from the analysis of behaviors, cultural objects, etc., *independently of their function*, one can infer important properties of the said biological structures. In other words, we contend that from the lack of specificity and the lack of functional transparency of structure — from which K derives the thesis that the system that computes words in combinations is inherently unfamiliar to language — nothing of interest can be said, since transparency works in the other direction, from function to structure. For us then, a system of computation can be unspecific and, at the same time, an inherent component of FL — as well as the other faculties that it subserves, as actually witnessed by the fact that it is transparent to the properties manifested by words — or other symbols in different domains (say, music or arithmetic). This position, we think, deserves to be carefully explained and contrasted with K’s opposing views.

K’s position is that there does not exist such a thing as a language-specific system of computation, and that in the absence of such a system it makes no sense to postulate the existence of FL — i.e. a naturally evolved cognitive system in charge of linguistic tasks. The rationale underlying K’s contention is that organic systems acquire their functional specializations by two different means: a) as a result of ‘natural selection’, in which case they can be properly deemed ‘adaptations for doing X’ — as in the case of the lungs and breathing; or b) as a consequence of ‘intentional decision’, in which case they become ‘instruments for doing X’ — as in the case of the lungs and playing the trumpet. According to K, language belongs to the second category, as it is the cumulative outcome of particular intentional agents having historically decided to give a secondary use to systems — including a computational system—naturally evolved for other purposes. So K’s idea is that in as much as it makes sense to speak of a language-dedicated computational system, it is just as an instrumentally adapted apparatus to a non-natural function. As a consequence, no natural language-dedicated computational system can be said to exist — and, concomitantly, no such a thing as FL actually exists.

It is our impression, however, that there is a flaw in K’s rationale that

compromises this chain of deductions. In a few words: the idea that every single organic system has a ‘proper function’ corresponds to a narrow-minded, old-fashioned, and probably wrong biology of sorts. Let’s explain why.

Contrary to common wisdom, organic systems are not inherently adapted to fulfill particular functions. They naturally evolve certain structural properties that endow them with the capacity of performing some activities, while others fall completely outside of their dynamic potential. So the activity that a given structure normally or most prominently runs — the one that we are tempted to attribute it as its ‘proper function’ — is better to be understood as a contingent effect of that structure’s connections to other organic systems and to a particular environment. What it is truly inherent to natural systems is their potential to perform a more or less open array of activities, were their organic or environmental context to change — in Reid’s (2007) terminology, their ‘adaptability’; see also Balari & Lorenzo (2010a, 2010b).⁶ Based on this, our claim is that it makes perfect sense to speak of an organic system as inherently devoid of a specific function, while acquiring different specializations as it naturally evolves certain connections to other systems and starts to be sensitive to certain environmental inputs. For us, this is a very suitable description for the system of computation underlying linguistic brain activity, and one capable of legitimately inspiring the biological study of FL.

Curiously enough, K’s reasoning is to a certain extent parallel to our own and he even makes use of a notion of ‘recycling’ — adopted from Dehaene (2009) — that can be seen as the cultural counterpart of the idea of ‘adaptability’ referred to above. K is in apparent agreement with us when he contends that “there is no such a thing as an intrinsic function of a physical structure” (Koster 2009: 69). However, while we defend that this is the case even when a structure seems to fix some practical specialization within a certain context — internal, external, or both, for K this is a state of affairs that applies only up to the point at which either natural selection ‘adapts’ (or ‘exapts’) it for a natural function — as in the case of breathing — or human invention ‘recycles’ it for non-natural tasks — as in the case of language.

Before closing the topic of the domain-unspecific character of the cognitive resources dedicated to language, let’s observe that K’s argument against Biolinguistics contrasts with another current of opinion according to which for linguistics to fulfill the project of becoming a branch of the natural sciences, a relaxation of the degree of specificity of the said resources is a crucial requisite. Otherwise, no true convergence with standard biological disciplines as neuroscience or genetics can reasonably be expected (Boeckx 2010b, Hornstein 2009). The logic underlying the idea is clear: the more specific the mechanisms put into use in a certain domain — as it is routinely assumed by most descriptive approaches in the case of language, the more difficult it becomes to connect them with their putative variants in other organisms — and, consequently, the less plausible any evolutionary explanation for their emergence. Thus, far from

⁶ An evolutionary corollary of this idea is that highly specialized structures — ‘adaptations’ — are more a risk than an advantage in the long run, given both the plasticity of organisms and the instability of environments. This kind of considerations is not, however, our main focus of interest here.

putting at risk the biolinguistic enterprise, the task of decomposing previously thought language-specific mechanisms into domain-general ones is for some, including ourselves, an urgent necessity in order to frame linguistic explanations within normal biological practice.

In other words, what K takes to be a lethal path for Biolinguistics — the path that Chomsky (2007) has called “approaching Universal Grammar from below”, we take to be a desideratum for a *rapprochement* between linguistics and biology. What makes these two opposite interpretations possible is the fact that what counts as biology is not fixed once and for all. Jackendoff (2010) was certainly right when he said (adapting his statement slightly) that one’s view of the biology of language depends on one’s view of language, but we wish to stress that it also depends on one’s view of biology.⁷ It is indeed important to bear in mind that biology is far from a simple field. Many are the biologists who have argued for a pluralist conception of the life sciences (note the plural!) (see, e.g., Gould 2002, Pigliucci & Mueller 2010), and even strong advocates of narrow, pan-adaptationist conceptions of biology such as the late Ernst Mayr (“the Darwin of the 20th century”) recognized the need to distinguish between two kinds of science, cutting across traditional disciplines like biology, for instance. Mayr (2004: 13, 24) leans toward attaching what he calls functional/‘mechanistic’ biology (i.e. molecular biology) to the natural sciences, and what he calls evolutionary biology to the historical sciences, and notes that each science has its own methodology and principles.

The same distinction may be necessary in the context of the language sciences, with one part of the field devoted to more cultural aspects of language, the languages — call this part (theoretical) philology —, and another devoted to the more natural aspects of language — call this Biolinguistics.

Ironically, K himself once pointed out (Koster 2003) that Chomskyan linguistics pursued along minimalist lines was “not philology by other means” (p. 171). We think this is exactly right, and moreover we think that this vindicates Biolinguistics. What is true of largely cultural entities like languages need not be true of the language faculty: whereas few would deny that the morphosyntax or grammar of particular languages is largely determined by their lexicons, we

⁷ Ironically enough, a conception of words almost identical to that of K is not seen as an obstacle, according to some authors, in order to approaching language with biologically informed lenses. It is the case, for example, of Millikan’s (2005) self-styled ‘biological model’ on linguistic conventions, according to which words are individually created items that replicate, proliferate and eventually become fixed for their coordinative benefits — relatively to similar units — within a community of users. Thus lexical inventories — with their grammatical, semantic, and pragmatic associated rules — are historical outcomes of ordinary Darwinian processes of differential reproduction and survival of the fittest.

Despite appearances, Millikan’s and K’s are not incongruous models. On the contrary, K’s anti-biologist stance on words is fully compatible with Millikan’s biological model, as Millikan’s is just a variant of Universal Darwinism, which means that the ‘model’ put forward in order to explain linguistic units is biological, but not the ‘object’ to which the model applies. Besides, Millikan is explicit in declaring that the objects to which her model applies have nothing to do with individual psychology and, therefore, with Biolinguistics as properly understood. So Millikan belongs to the same anti-biological quarters as K, even if contending that some laws exist that universally hold in both the biological and the cultural realms. So the question remains whether words are inherently foreign to individual psychology — and whether they justify a biological approach to language.

submit that the syntactic principles of Universal Grammar are completely independent of the cultural constructs we call words (or morphemes). Such principles continue to depend on atomic units, but these units only consist of natural instructions.

4. A Note on the Proposition that Words Are Tools and Language Is the Technology Embodied by Them

In this section we would like to briefly examine the contrast K introduces between the conception of 'language-as-an-organ' (FL) and his conception of 'language-as-a-technology' (TL).

The latter conception is not completely new, as it is vaguely suggested in McLuhan (1964) and developed in Logan (2007), where the contention is explicitly made that language belongs to a series of human inventions comprising spoken language, writing, mathematics, science, computing, and the Internet. Logan's thesis is that all these practices are technical improvements connected to the human necessity of representing and transmitting knowledge, each one historically emerging at points of informational overload that made insufficient the pre-existing technologies. A shortcoming of the idea — and one of which Logan is not unaware — is that improving an existing technique is a thing very different from creating it from scratch, so the question remains of how something like a TL could be created. This is by the way a question that, in slightly different contexts, also worried Humboldt (1836) and Rousseau (1781), and to which both responded by appealing to an instinctual basis for language.

K's ingenious alternative is a different one — and one that deserves to be carefully scrutinized: language (TL) is a human creation resulting from the prior invention of words. In other terms, TL results from the impact of words on the human brain — a source of extremely powerful cognitive resources, but otherwise a linguistic blank slate. Towards the end of this section we'll return to the issue of transparency in order to argue, among other things, that it is not an easy task — if a feasible one at all — to explain how properties such as compositionality and productivity could be added to the pack of inventions associated to words. But problems with K's conjecture are more serious than that. It is K's opinion that his view can comfortably be framed within the 'extended mind' paradigm (Clark 1997, Clark & Chalmers 1998), as it purports that language-associated mental activity results from the recruitment of external inputs (words) by general purpose and linguistically opaque cognitive systems. It is not clear, however, that such an assumption is so congenial with the extended mind framework as normally envisioned by its advocates. Let's see why.

Proponents of the extended mind model are not committed to any particular cognitive architecture — and this obviously includes the idea of mind as a blank slate of sorts prior to its embedding in the world. They just defend that (some) cognitive systems incorporate elements of the environment, so they comprise both internal and external components, the role of which in the normal execution of the system's activities is seen as functionally equivalent. The question is however orthogonal to that concerning the specificity — or lack thereof — of the relevant systems. So, in principle, the 'extended' thesis is

compatible with cognitive architectures of any degree of modularity. However, there exist strong arguments — put forward in Rupert (2009) with great detail — suggesting that ‘extended’ systems are only operative if their internal components are highly articulated and robustly constructed modules, in the absence of which it is not to be expected any particularly useful sensitivity to external inputs. As a matter of fact, Rupert’s (2009) conclusion is that the extended mind rhetoric can be dispensed with entirely without great harm to our understanding of cognition, saying instead that minds comprise an integrated set of mechanisms and capacities in the functioning of which certain environmental inputs may exert an important causal impact.

It is not particularly important for our argument whether the question is settled in favor of or against preserving the ‘extended’ idea and its vocabulary. The substantive part of the question is that words — understood as ‘man-made, public cultural objects’ — seem to be of little help to cognition in the absence of an associated set of internal mechanisms and capacities that, as we have argued at length, happens to be ‘transparent’ with respect to the properties of these external objects — and thus deserve the name of FL.

At this point, we would like to stress what we take to be another important aspect of words, in the context of K’s challenge. A particular lexical inventory is obviously a cultural phenomenon — there is no point in discussing this. However, it is far more contentious whether the units composing them belong to a same category with the symbols of other non-linguistic cultural inventories (for the ease of discussion, we will refer to the former as ‘words’ and to the latter as ‘symbols’). In our opinion, there exist good reasons to believe that they do not (Balari *et al.* 2011, for a detailed argument). The crucial point is that the information encoded in symbols is ‘opaque’ in a sense in which the information encoded in words is not, in that in order to be a competent user of symbols one needs to be familiar with the contexts in which they show up and how they relate with each other in each particular context of use (Eco 1975). This kind of acquaintance is not however a requisite in order to be fully competent as a word-system user, as once one knows the information encoded in given words, she gains access to the information encoded in combinations thereof, even without previous familiarity with the contexts in which these words’ use is appropriate and even if no such contexts happen to exist. These are well-known facts, of which philosophers and linguists have been aware for a long time — but of which no clear explanations have been traditionally offered.

Indeed, we think that it is one of the strengths of Biolinguistics — under the guise criticized by K — that it comes with the only reasonable explanation hitherto offered to this recalcitrant problem: Word-systems are inextricably connected to a system of computation that — returning to K’s terms — is ‘transparent’ with respect to the properties of words — namely, words are used compositionally and productively. Such an explanation vanishes as soon as this mind internal connection is severed and words are reduced to the same condition as other non-linguistic symbols — i.e. man-made, public cultural objects. It is worth remembering that Wittgenstein, the most conspicuous and influential defender of this ‘words-as-external-symbols’ view (Wittgenstein 1953), assumed — coherently with the model — that linguistic meanings were not compositional.

An idea that he defended ingeniously and enthusiastically (Wittgenstein 1958), but that only seems to really work in the case of phrases used (quasi-idiomatically) in ritualistic or other highly stereotyped situations — in which words are actually reduced to the condition of symbols.

Wittgenstein's is not for sure K's strategy to deal with this problem. It is not however completely clear what his strategy is. We think that he is forced to admit that compositionality and productivity are inventions added at a certain point — if not from the very beginning — to the way words behave. But once this is admitted, and given that these are properties unfamiliar to other 'man-made, public cultural' symbols, the burden of the proof is clearly on the side of the defender of the 'words-as-external-symbols' view, as it seems extremely counter-intuitive that the said properties are imprinted via words on an system of computation inherently opaque with respect to them.

Let us conclude this section by emphasizing that nothing thus far said purports to deny the evidence that words have external counterparts (to which we can refer as 'E-words'), or even the admission that parts of the information encoded in words have external origins, meaning that FL is not transparent with respect to this particular pieces of information ('E-features'),⁸ but as we have argued none of this actually undermines a natural, biological study of linguistic computations. Note only that, by assuming a position like K's according to which words/language are cultural objects with, of course, a biological basis, but one that is inaccessible/irrelevant for the study of language, we run the risk of falling into the trap of the thesis of the ontological autonomy of culture held by many anthropologists and which Dan Sperber (1996) has cogently criticized as being blatantly contradictory. Sperber's point is that ontological autonomy is untenable because it is a form of cryptodualism: if you are a materialist, your cultural ontology has to be grounded on a physical/material ontology (cultural objects must have a material basis if we are not going to accept an 'irrational' account of causality), otherwise, your cultural ontology is vacuous (see also Jablonka & Lamb 2005 for a congenial criticism). An alternative, perhaps closer to Koster's position, is to just say 'of course there's a material basis, who ever denied that?', and leaving it there. This is empty materialism (Sperber 1996: 11), a position totally incapable of justifying an ontology of cultural objects, which, once its material basis is investigated, may turn out to be false.

⁸ As a matter of fact, this is a suitable way of making sense of Chomsky's concept of linguistic 'imperfection' (Chomsky 1995 and subsequent works), an idea thought to capture those aspects of grammars (case and agreement features being two conspicuous examples) that seem to lack any motivation from the point of view of the cognitive systems (sensory-motor and conceptual) that the computational machine of FL accesses. A reasonable conjecture, worth being empirically tested, is that these features work as external devices that stimulate the development of the computational system and ease its normal functioning (Lorenzo & Longa 2003). This idea would justify preserving to a certain extent K's instrumental view on words. It does not justify, however, the strong instrumentalist thesis (see our section 4) according to which words are the tools that create language (TL), as in any event they are tools that clearly presuppose the existence of a robust cognitive system (FL) devoted to dealing with them.

5. Conclusion: What a Shame It Would Be to Abandon Biolinguistics so Soon!

The alternative to the various shortcomings of K's theses is, of course, to stay firmly in the idea of FL as an organ — a part of the nervous system that deserves the dedication of a special branch of Biology. Questions routinely directed to other aspects of the biological realm make also perfect sense when aimed at this particular object: What are its component parts and how they compound a coherent unit of activity? How can this activity be described in the abstract and how is it physically realized? How does it become developmentally assembled and how did it evolve this developmental pattern? Such a research program cannot be seen, however, as the denial that systems of grammatical conventions also exist, the historical creation and transformations of which also deserve a scientific branch of specialization. For us it is nonetheless clear that the existence of historical grammatical systems is only possible against the background of a natural system (FL), the study of which seems mandatory in order to understand how they emerge within speaking communities and how they are acquired and used by individuals, as well as to establish the putative role of these systems in the opposite direction — i.e. as agents with a causal impact in the early development of FL in the individual and even in the evolutionary (or co-evolutionary) process of the faculty in the species.

We can give up doing Biolinguistics. Granted. But is it really worth the price of renouncing to understand questions like these?

References

- Balari, Sergio & Guillermo Lorenzo. 2010a. ¿Para qué sirve un ballestrinque? Reflexiones sobre el funcionamiento de artefactos y organismos en un mundo sin funciones. *Teorema. Revista Internacional de Filosofía* 29, 56–76.
- Balari, Sergio & Guillermo Lorenzo. 2010b. Communication. Where evolutionary linguistics went wrong. *Biological Theory* 5, 228–239.
- Balari, Sergio, Antonio Benítez-Burraco, Marta Camps, Víctor M. Longa, Guillermo Lorenzo & Juan Uriagereka. 2011. The archaeological record speaks: Bridging anthropology and linguistics. *International Journal of Evolutionary Biology*, 2011 (Special Issue: Key Evolutionary Transformations in Homo Sapiens), Article ID 382679, 17 pages, doi:10.4061/2011/382679.
- Boeckx, Cedric. 2010a. Defeating lexicocentrism. Ms., ICREA/Universitat Autònoma de Barcelona. Available at <http://ling.auf.net/lingBuzz/001130>.
- Boeckx, Cedric. 2010b. *Language in Cognition. Uncovering Mental Structures and the Rules behind Them*. Malden, MA: Wiley-Blackwell.
- Borer, Hagit. 2005. *Structuring Sense* (2 vols.). Oxford: Oxford University Press.
- Bresnan, Joan. 2007. A few lessons from typology. *Linguistic Typology*, 11, 297–306.
- Chomsky, Noam. 1957. *Syntactic Structures*. Mouton: The Hague: Mouton.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

- Chomsky, Noam. 1970. Remarks on nominalization. In Roderick A. Jacobs & Peter S. Rosenbaum (eds.), *Readings in English Transformational Grammar*. Waltham, MA: Ginn.
- Chomsky, Noam. 1975. *The Logical Structure of Linguistic Theory*. New York: Plenum Press.
- Chomsky, Noam. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 2007. Approaching UG from below. In Uli Sauerland & Hans-Martin Gärtner (eds.), *Interfaces + Recursion = Language? Chomsky's Minimalism and the View from Syntax–Semantics*. Berlin and New York: Mouton de Gruyter.
- Clark, Andy. 1997. *Being There. Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, Andy & David J. Chalmers. 1998. The extended mind. *Analysis* 58, 7–19.
- Dehaene, Stanislas. 2009. *Reading in the Brain*. New York: Penguin.
- Eco, Umberto. 1975. *Trattato di semiotica generale*. Milano: Bompiani.
- Ginzburg, Jonathan & Ivan A. Sag. 2000. *Interrogative Investigations: The Form, Meaning, and Use of English Interrogatives*. Stanford, CA: CSLI Publications.
- Goldberg, Adele E. & Ray Jackendoff. 2004. The English resultative as a family of constructions. *Language* 80, 532–568.
- Gould, Stephen J. 2002. *The Structure of Evolutionary Theory*. Cambridge, MA: Belknap Press of Harvard University Press.
- Hauser, Mark D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 289, 1569–1579.
- Hornstein, Norbert. 2009. *A Theory of Syntax: Minimal Operations and Universal Grammar*. Cambridge: Cambridge University Press.
- Humboldt, Wilhem von. 1836. *Über die Verschiedenheit des menschlichen Sprachbaues und ihren Einfluss auf die geistige Entwicklung des Menschengeschlechts*. Berlin: F. Dümmler.
- Jablonka, Eva & Marion J. Lamb. 2005. *Evolution in Four Dimensions. Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, MA: MIT Press.
- Jackendoff, Ray 1997. *The Architecture of the Language Faculty*. Cambridge, MA: MIT Press.
- Jackendoff, Ray. 2010. Your theory of language evolution depends on your theory of language. In Richard Larson, Viviane Déprez & Hiroko Yamakido (eds.), *The Evolution of Human Language. Biolinguistic Perspectives*, 63–72. Cambridge: Cambridge University Press.
- Julien, Marit. 2002. *Syntactic Heads and Word Formation*. Oxford: Oxford University Press.
- Koster, Jan. 1987. *Domains and Dynasties*. Dordrecht: Foris.
- Koster, Jan. 2003. Not philology by other means. *Glott International* 7, 171–172.
- Koster, Jan. 2009. Ceaseless, unpredictable creativity: Language as technology. *Biolinguistics* 3, 61–92.
- Logan, Robert K. 2007. *The Extended Mind: The Emergence of Language, the Human Mind, and Culture*. Toronto: University of Toronto Press.

- Lorenzo, Guillermo & Víctor M. Longa. 2003. *Homo Loquens: Biología y Evolución del Lenguaje*. Lugo: Tris Tram.
- Martinet, André. 1960. *Éléments de linguistique générale*. Paris: Armand Colin.
- McLuhan, Marshall. 1964. *Understanding Media: The Extensions of Man*. Cambridge, MA: MIT Press.
- Marantz, Alec. 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. *Proceedings of the 21st Annual Penn Linguistics Colloquium — University of Pennsylvania Working Papers in Linguistics* 4.2, 201–225.
- Marantz, Alec. 2000. Words. Ms., Massachusetts Institute of Technology, Cambridge, MA.
- Mayr, Ernst. 2004. *What Makes Biology Unique?* Cambridge: Cambridge University Press.
- Millikan, Ruth G. 2005. *Language: A Biological Model*. Oxford: Oxford University Press.
- Otero, Carlos. 1976. The dictionary in a generative grammar. Paper presented at the 91st Annual Convention of the Modern Language Association, New York (26–29 December 1976).
- Owen, Richard. 1849. *On the Nature of Limbs. A Discourse*. London: John Van Voorst.
- Pigliucci, Massimo & Gerd B. Mueller (eds.). 2010. *Evolution: The Extended Synthesis*. Cambridge, MA: MIT Press.
- Pinker, Steven & Ray Jackendoff. 2005. The faculty of language: What is special about it? *Cognition* 95, 201–236.
- Pollard, Carl J. & Ivan A. Sag. 1992. Anaphors in English and the scope of Binding Theory. *Linguistic Inquiry* 23, 261–303.
- Pollard, Carl J. & Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Chicago: University of Chicago Press.
- Reid, Robert G. B. 2007. *Biological Emergences. Evolution by Natural Experiment*. Cambridge, MA: MIT Press.
- Rousseau, Jean-Jacques. 1781. *Essai sur l'origine des langues* (ed. by Jean Starobinski). Paris: Gallimard.
- Rupert, Robert D. 2009. *Cognitive Systems and the Extended Mind*. Oxford: Oxford University Press.
- Sag, Ivan A. In press. Sign-Based Construction Grammar: An informal synopsis. In Hans C. Boas & Ivan A. Sag (eds.), *Sign-Based Construction Grammar*. Stanford, CA: CSLI Publications.
- Sapir, Edward 1921. *Language: An Introduction to the Study of Speech*. New York: Harcourt, Brace and Co.
- de Saussure, Ferdinand. 1916. *Cours de linguistique générale*. Paris: Payot.
- Sperber, Dan 1996. *Explaining Culture: A Naturalistic Approach*. Oxford: Blackwell.
- Starke, Michal. 2010. Nanosyntax: A short primer to a new approach to language. *Nordlyd*, 36, 1–6.
- Whitney, William Dwight. 1875. *The Life and Growth of Language: An Outline of Linguistic Science*. New York: D. Appleton and Company.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations* (translated by G.E.M. Anscombe). Oxford: Basil Blackwell.

Wittgenstein, Ludwig. 1958. *The Blue and Brown Books*. Oxford: Basil Blackwell.

Sergio Balari
Universitat Autònoma de Barcelona
Departament de Filologia Catalana
& Centre de Lingüística Teòrica
Edifici B, Campus UAB
08193 Bellaterra (Barcelona)
Spain

Sergi.Balari@uab.cat

Cedric Boeckx
ICREA & Universitat de Barcelona
Department of Linguistics
Gran Via de les Corts Catalanes, 585
08007 Barcelona
Spain

cedric.boeckx@ub.edu

Guillermo Lorenzo
Universidad de Oviedo
Departamento de Filología Española
Facultad de Filosofía y Letras
Campus El Milán 33011 Oviedo
Spain

glorenzo@uniovi.es

A Turing Program for Linguistic Theory

Jeffrey Watumull

1. The Big Question

Universal to systems so various and complex as the foundations of mathematics, cryptography, computer science, artificial intelligence, and morphogenesis is the “the concept of ‘mechanical procedure’ (alias ‘algorithm’ or ‘computation procedure’ or ‘finite combinatorial procedure’). This concept is shown to be equivalent to that of a ‘Turing machine.’ [In fact,] due to A.M. Turing’s work, a precise and unquestionably adequate definition of the general concept of formal system can now be given” (Gödel in Davis 1965: 71–72,). For this triumph, *inter alia*, the world is now celebrating the centenary of the mathematician Alan Mathison Turing (1912–1954).¹ The mathematical universality of the Turing machine — the abstract system of computation — implies that it is not only *relevant* to biolinguistic research, but *intrinsic* to linguistic — indeed any cognitive-neurobiological — computation. To demonstrate this is the desideratum of my proposed *Turing Program for Linguistic Theory* (TPLT). The proposal is very summary and very sketchy; I proffer no answers, only approaches to questions.

One of the “Big Questions” of the Turing Centenary (see <http://www.turingcentenary.eu>) is whether there exists “a successful mathematical model of intelligent thought”. The answer is surely not yet, but I am sure there will be. Cognition is clearly computational (see Gallistel & King 2009) and computation is mathematical by definition: Procedures are run to determine the symbolic outputs (values) of functions given symbolic inputs (arguments); in the domain of the brain, the running of procedures is referred to informally as ‘thinking’. In a successful model of this process, functions would be “*completely* determined” (Turing 1936: 232, emphasis original) by rules and representations so “perfectly explicit” (Chomsky 1965: 4) as to be automatable.²

For comments and criticisms, my many thanks to Bob Berwick, Noam Chomsky, Randy Gallistel, Marc Hauser, Steve Pinker, and Ian Roberts.

¹ From the editors of the “Turing at 100” special issue of *Nature* (2012, 482: 440):

Nature invites its readers to embrace and celebrate [Turing,] one of the brightest minds of all time [...]. The scope of Turing’s achievements is extraordinary. Mathematicians will honour the man who cracked David Hilbert’s *Entscheidungsproblem* or ‘decision problem,’ and cryptographers and historians will remember him as the man who broke Nazi Germany’s Enigma code and helped to shorten the Second World War. Engineers will hail the founder of the digital age and artificial intelligence. Biologists will pay homage to the theoretician of morphogenesis, and physicists will raise a glass to the pioneer of nonlinear dynamics.

² For “*if you can’t program it, you haven’t understood it*” (Deutsch 2011: 146, emphasis original).



The successful model would be *descriptively* and *explanatorily adequate* (see Chomsky 1964): It would completely describe *what* the system of rules and representations is (see Roberts 2011) and completely explain *how* it develops — how internal and external factors determine the genotype-to-phenotype expression of the cognitive system (see Pinker 1984). The model might even go *beyond* explanatory adequacy (see Chomsky 2004a) to answer the biggest question of all: *Why* does the system assume this one form out of the infinity of conceivable forms? The answer might ultimately derive from computational constraints reducible to mathematical laws. This model would thus succeed in explaining not only *something of the specific nature of intelligence*, but ultimately *something of the general nature of reality*.³

These *what*, *how*, and *why* questions are obviously too big to answer as framed, and must therefore be decomposed into smaller solvable problems. Hence I propose the TPLT, a research program based on Turing's mathematics, to discover and model mathematically important aspects of intelligent thought in the domain of language.

2. Language

One strategy for answering the general question whether there could be a successful mathematical model of intelligent thought is to reformulate it as a mathematical modeling problem for a specific form of human intelligence. The form I propose to consider is that of language.

Why study language? [One reason is to] discover abstract principles that govern its structure and use, principles that are universal by biological necessity and not mere historical accident, that derive from mental characteristics of the species [...]. By studying the properties of natural languages, their structure, organization, and use, we may hope to gain some understanding of the specific characteristics of human intelligence. We may hope to learn something about human nature; something significant, if it is true that human cognitive capacity is the truly distinctive and most remarkable characteristic of the species. (Chomsky 1975: 4)

As an explanandum of scientific inquiry, 'intelligence' is so polysemous that any theory not formalizing it is "too meaningless to deserve discussion" (Turing 1950: 442).^{4,5} *Mutatis mutandis* for 'language', commonly conflated with

³ I am but one of many to have conjectured that fundamentally the universe *is* mathematical (computational/informational).

⁴ Polysemy is not the only (or even the main) problem. As Noam Chomsky (p.c.) observes, the question whether a machine *really* thinks is equally meaningless to the question whether a submarine *really* swims: "These questions have to do with choice of metaphor. They are not substantive. Questions about swimming, thinking, [etc.] are about the meanings of linguistic elements/concepts, and how they are used" (see Chomsky 2009).

⁵ Essentialism is comparably meaningless: "[S]ooner or later we will be able to assemble programs of great problem-solving ability from complex combinations of heuristic devices — multiple optimizers, pattern-recognition tricks, planning algebras, recursive administration procedures, and the like. In no one of these will we find the seat of intelligence. Should we ask what intelligence 'really is'?" (Minsky 1963: 446–447). Like an onion — that imperishable if slightly stale staple of analogies — we can peel away the purely mechanical functions of the mind in search of the 'real' mind. If "we eventually come to the skin which

any system of communication or representation (perhaps quasi-symbolic); or interpreted as a social-political construct governed by shared norms and conventions of usage; or classified as “the totality of utterances made in a speech community” (Bloomfield 1926: 155); or reduced to a Skinnerian repertoire of habits, abilities, and dispositions to respond to verbal stimuli; or stipulated to be a set of essentially Quinean well-formed formulae; or obfuscated in a Wittgensteinian game or “form of life”. By contrast, modern linguistics — *generative linguistics* (subsuming *biolinguistics*) — established and expounded by Noam Chomsky, in a Turing-style consilience of the formal and natural sciences, adopts a rigorous and empirical definition of ‘language’ as an *I-language*: A cognitive computational system — a function in *intension* — *internal* to an *individual* of the species *Homo sapiens sapiens*.^{6,7} The function recursively generates syntactic structures mappable via formal semantics and rule-based morphology–phonology to interfaces with conceptual-intentional and sensory-motor systems, respectively. *I-language* is thus a system of *d-infinity* (discrete/denumerable/digital infinity) analogous to the natural numbers: A finite system that in principle can generate an infinite set of hierarchically structured expressions by recursively combining discrete elements.

A generative system *strongly generates* structures/sets (material to linguistic cognition) and *weakly generates* strings/sequences (marginal to linguistic cognition). The structures, not the strings, represent the grammatical information mappable to representations of semantic and morphological-phonological information, as evidenced by the fact that one string can correspond to many structures (in a many-one function). Consider that the one string *the boy saw the man with binoculars* is two-ways ambiguous because it corresponds to two possible structures representing two possible interpretations: (i) $\{\{the, boy\}, \{saw, \{the, \{man, \{with, binoculars\}\}\}\}\}$; (ii) $\{\{the, boy\}, \{\{saw, \{the, man\}\}, \{with, binoculars\}\}\}$.

The strong generative capacity of the language faculty was probably exapted in human evolution to connect interfaces with systems necessary for general intelligence (see Hauser 2009).⁸ For instance, Minsky (1963) argues that the sophistication of human pattern-recognition necessitates “provisions for recursive, or at least hierarchical use of previous results” as in the “articulation” of a scene into descriptions of “elementary figures” and “subexpressions [...] designating complex subfigures” with a “figure [...] first divided into two parts; [and] then [with] each part [...] described using the same machinery” (pp. 434, 423).

has nothing in it”, we cannot but conclude that “the whole mind is mechanical” (Turing 1950: 454–455).

⁶ To define the function in *extension* is to define the set of syntactic objects (strings/structures) it generates (see Church 1941). For linguistic theory, the function needs to be defined in *intension*: The definition of the function qua procedure for generating sets of structures; the properties of the structures derive, in part, from the properties of the function.

⁷ Of course, a universal — species-typical — *I-language* can be abstracted from particular *I-languages* for formal and empirical inquiry. And if Platonically inclined, we may abstract a universal *I-language* as a mathematical object for metaphysical analysis.

⁸ Updated with technology from generative linguistics, the general problem solving program T.O.T.E. of Miller, Galanter, and Pribram (1960) has been argued to be active in linguistic and extralinguistic domains (see Jackendoff 2007).

And as McCarthy (1956) explained, recursively hierarchical generativity is necessary for complex planning methods because inefficient if not technically intractable problems can be solved only by decomposition into tractable and efficient sub-problems; composition of the functions and solutions for the latter can solve the former for infinite sets of problems.⁹ Accordingly, recursively generated “Hierarchy [...] is one of the central structural schemes that the architect of complexity uses” (Simon 1962: 468) (see Turing 1947 on the central importance of ‘subsidiary tables’, i.e. recursive subroutines).

3. Language and the Brain

I submit that the object of linguistic inquiry is, or can be regarded as, “the thing in itself”, a computational — ergo mathematical — system abstracted away from spatiotemporal contingencies, as a Turing machine is with its memory space and operating time unlimited: “With this will come a mathematical characterization of a class of [...] functions, the functions ‘computed’ by these Turing machines. These functions will be called *computable functions*, [identifiable with] the intuitive concept of effectively calculable function” (Davis 1958: 3). In short, “[s]omething is computable if it can be computed by a Turing machine” (Gallistel & King 2009: 105)¹⁰ and “[a]ny Turing machine is completely described by a *machine table*” (Putnam 1975: 365), a *functional organization*, specifying its mathematical-logical rules and representations, not its physical implementation (if it even has one): “[T]he ‘logical description’ of a Turing machine does not include any specification of the *physical nature* of [its rules and representations] — or indeed, of the physical nature of the whole machine [...]. In other words, a given ‘Turing machine’ is an *abstract* machine which may be physically realized in an almost infinite number of different ways” (Putnam 1975: 371, *emphases original*), *if at all*.¹¹ So it is mere “superstition” to attach “[i]mportance [...] to the fact that modern digital computers are electrical, and that the nervous system also is electrical. Since Babbage’s machine was not electrical, and since all digital computers are in a sense equivalent, we see that this use of electricity cannot be of theoretical importance [...]. If we wish to find [...] similarities, we should look rather for *mathematical analogies of function*” (Turing 1950: 439, *emphasis added*).¹²

Consistent with this reasoning, the object of linguistic inquiry can be defined as a form of *mathematical-functional competence*: the “underlying system of

⁹ Behaviorist psychology assumed nonhierarchical chaining theories (Markov models), but these were determined to be inadequate by Lashley (1951) in a paper (unappreciated until Chomsky 1959b) on the necessity of hierarchical planning in solving the problem of serially-ordered behavior.

¹⁰ In this proposal I do not address the limits of the computable (e.g., the relation of artificial neural networks, so-called ‘Super-Turing machines’, etc. to classical Turing machines). That is a task for the TPLT.

¹¹ The general and necessary and sufficient conditions an object must satisfy for it to be defined as a (type of) Turing machine — a computer — are purely mathematical-functional (see Carnap 1955 on such intensional definitions).

¹² Charles Babbage (1791–1871) was a Cambridge mathematician and designer of the Difference Engine and the Analytical Engine, which “had all the essential ideas” (i.e. the mathematical-functional components and procedures) of Turing’s computers (Turing 1950: 439), but were mechanical.

rules" (Chomsky 1965: 4) in the mind that "represents the information concerning sentence structure that is available, in principle, to one who has acquired the language" (Chomsky 1963: 326–327). This information is represented as an "idealization [...] leaving out any limitations [...] of memory, time, and access" (Chomsky 1965: 4, 10). Idealization of the linguistic system reveals the mathematical-functional components and procedures necessary and sufficient to define it as a subtype of Turing machine. Such idealization is part and parcel of the methodology and the *metaphysics* of normal science, which proceeds by the "making of abstract mathematical models of the universe to which at least the physicists give a *higher degree of reality* than they accord the ordinary world of sensation" (Weinberg 1976: 28, emphasis added).¹³ For those working in the Turing Program for Linguistic Theory (TPLT), it would be natural to give the highest degree of reality to the mathematical form of I-language.

In the "world of sensation", things in themselves, often abstract, are confounded by arbitrary constraints, often physical. For computational systems, confounding the abstract with the physical can conflate the truistic yet only lip serviced distinction between software and hardware; thus this important distinction remains unassimilated, preventing recognition of the fact that "[a]s our knowledge increases, the abstract mathematical world becomes farther removed from the world of sensation" (Weinberg 1976: 28). For instance:

You know that if your computer beats you at chess, it is really the *program* that has beaten you, not the silicon atoms or the computer as such. The abstract program is instantiated physically as a high-level behaviour of vast numbers of atoms, but the *explanation* of why it has beaten you cannot be expressed without also referring to the program in its own right. That program has also been instantiated, unchanged, in a long chain of different physical substrates, including neurons in the brains of the programmers and radio waves when you downloaded the program via wireless networking, and finally as states of long- and short-term memory banks in your computer. The specifics of that chain of instantiations may be relevant to explaining how the program reached you, but it is irrelevant to why it beat you: there, the content of the knowledge (in it, and in you) is the whole story. That story is an explanation that refers ineluctably to abstractions; and therefore those abstractions exist, and really do affect physical objects in the way required by the explanation.

(Deutsch 2011: 114–115, emphases original)¹⁴

Consistent with this reasoning, it is not unreasonable to "give a higher degree of reality" to an "abstract mathematical model" of linguistic computation than to its implementation in the "ordinary world of sensation." But this poses a problem for the "Mind, Mechanism and Mathematics" theme of the Turing Centenary:

The joint study of brain and language [...] has achieved some basic results correlating linguistic phenomena with brain responses, but has not advanced to any explanatory theory that identifies the nature of linguistic computation in the brain [...]. The absence of an explanatory theory of this type is the result of the conceptual granularity mismatch and the ontological incommensurability between the foundational concepts of linguistics and

¹³ This can be construed as a restatement of the Platonic theory of forms.

¹⁴ This is to restate the Aristotelean distinction of matter and form.

those of neurobiology [...]. Consequently, there is an absence of reasonable linking hypotheses by which one can explore how brain mechanisms form the basis for linguistic computation. (Poeppel & Embick 2005: 14–15)

The ‘conceptual granularity’ in theories of I-languages is measured on ‘higher’ computational and algorithmic levels whereas the primitives of neuroscience are posited on the ‘lower’ implementational level (see Marr 1982). *Prima facie*, the ontologies of these levels are incommensurable: set-formation, phrases, and so on in linguistics; action potentials, neurons, and so forth in neuroscience. If a mathematical model of intelligent thought is to be formulated — and *a fortiori* realized as artificial intelligence — a novel and nontrivial unification of the levels of analysis, the levels “at which any machine carrying out an information-processing task must be understood” (Marr 1982: 25), is imperative, requiring interdisciplinary research. The concept that unifies the research is *computation* — essential to which is *information* — and the concept that unifies computation is the Turing machine. Indeed, the beauty of the Turing machine is that in its abstractness it subsumes and thereby relates all computational primitives; in principle therefore it renders commensurable the computational ontologies of linguistics and neuroscience — or so I would endeavor to prove in the TPLT.

A Turing machine is a mathematical abstraction, not a physical device, but my theory is that the information it specifies in the form of I-language must be encoded in the human genetic program — and/or derived from the mathematical laws of nature (‘third factors’ in the sense of Chomsky 2005) — and expressed in the brain. Central to the machine is a generative procedure for d-infinity; however, “[a]lthough the characterizations of what might be the most basic linguistic operations must be considered one of the deepest and most pressing in experimental language research, we know virtually nothing about the neuronal implementation of the putative primitives of linguistic computation” (Poeppel & Omaki 2008: 246). So is presented the great challenge for the TPLT: To precisify (formalize) the definitions of linguistic primitives in order that ‘linking hypotheses’ (not mere correlations) to as yet undiscovered neurobiological primitives can be formed.

4. Generative Systems

It was in the theory of computability and its equivalent formalisms that the infinite generative capacity of a finite system was formalized and abstracted and thereby made available to theories of natural language (see Chomsky 1955 for a discussion of the intellectual zeitgeist and the influence of mathematical logic, computability theory, etc. at the time generative linguistics emerged in the 1950s). In particular, a *generative grammar*¹⁵ was defined as a set of rules that recursively generate (enumerate/specify) the sentences of a language in the form of a *production system* as defined by Post (1944) and exapted by Chomsky (1951):

$$(1) \quad \phi_1, \dots, \phi_n \rightarrow \phi_{n+1}$$

¹⁵ Linguists use the term with systematic ambiguity to refer to the explananda of linguistic theory (i.e. I-languages) and to the explanantia (i.e. theories of I-languages).

“[E]ach of the ϕ_i is a structure of some sort and [...] the relation \rightarrow is to be interpreted as expressing the fact that if our process of recursive specification generates the structures ϕ_1, \dots, ϕ_n then it also generates the structure ϕ_{n+1} ” (Chomsky & Miller 1963: 284); the inductive (recursive) definition derives infinite sets of structures. The objective of this formalization was analogous to “[t]he objective of formalizing a mathematical theory a la Hilbert, [i.e.] to remove all uncertainty about what constitutes a proof in the theory, [...] to establish an algorithm for the notion of proof” (Kleene 1981: 47) (see Davis 2012 on Hilbert’s program). Chomsky (1956: 117) observed that a derivation as in (1) is analogous to a proof with ϕ_1, \dots, ϕ_n as the set of axioms, the rewrite rule (production) \rightarrow as the rule of inference, and the derived structure ϕ_{n+1} as the lemma/theorem. For a toy model, let (2) be a simplified phrase structure grammar with S = Start symbol *Sentence*, $\widehat{}$ = concatenation, $\#$ = boundary symbol, $N[P]$ = Noun [Phrase], $V[P]$ = Verb [Phrase]:

- (2) $\# \widehat{S} \widehat{\#}$
 $S \rightarrow NP \widehat{VP}$
 $VP \rightarrow V \widehat{NP}$
 $NP \rightarrow the \widehat{man}, the \widehat{book}$
 $V \rightarrow took$

(3) is one possible derivation given the grammar (production system) in (2):

- (3) $\# \widehat{S} \widehat{\#}$
 $\# \widehat{NP} \widehat{VP} \widehat{\#}$
 $\# \widehat{the} \widehat{man} \widehat{VP} \widehat{\#}$
 $\# \widehat{the} \widehat{man} \widehat{V} \widehat{NP} \widehat{\#}$
 $\# \widehat{the} \widehat{man} \widehat{took} \widehat{NP} \widehat{\#}$
 $\# \widehat{the} \widehat{man} \widehat{took} \widehat{the} \widehat{book} \widehat{\#}$

The derivation proceeds deterministically, stepwise, “remov[ing] all uncertainty about what constitutes a [derivation] in the theory”.

Restricted and unrestricted production systems were proved by Post (1944, 1947) to be formally equivalent to the effectively calculable functions — and by extension the λ -calculus and by extension Herbrand-Gödel general recursion (Gödel 1934, Church 1936, Kleene 1936) — proved by Turing (1936) to be equivalent to the computable functions.^{16,17} These equivalences necessitated additional restrictions on generative grammars (or seemed to): “[A]n arbitrary Turing machine, or an unrestricted rewriting system, is too unstructured to serve as a grammar [...]. Obviously, a computer program that succeeded in generating

¹⁶ Independent of Turing, Post (1936) formulated a computational (mathematical) machine equivalent to Turing machines and Gödel-Church recursiveness — and inferred from it provocative psychological conclusions (contradicting those of Penrose 1989).

¹⁷ A Turing machine M can be described by a set Σ of rewrite rules that convert a string $\#S_0\phi\#$ to $\#S\#$ just in case M accepts ϕ . (Given the determinism of M , Σ is monogenic.) The set of rewrite rules Σ' containing $\psi \rightarrow \chi$ just in case $\chi \rightarrow \psi$ is in Σ and containing a Stop rule $\#S_0 \rightarrow \#$ is an unrestricted rewriting system.

sentences of a language would be, in itself, of no scientific interest unless it also shed some light on the kinds of structural features that distinguish languages from arbitrary, recursively enumerable sets" (Chomsky 1963: 359–360).¹⁸

Thus the Chomsky hierarchy of grammar types — type 0 (unrestricted \approx Turing machines) \supset type 1 (context sensitive \approx linear bounded automata) \supset type 2 (context free \approx pushdown automata) \supset type 3 (finite state \approx finite automata) — was formulated to define the generative capacities of grammars and corresponding automata (see Chomsky 1956). The objective for theoretical linguistics was then — *and should be again* — to discover the type of grammar/automaton for natural language; language defined as a cognitive system realizing a Turing machine subtype. But this research program failed because of a preoccupation with weakly generated strings (sentences) rather than strongly generated structures. In the original conception of the hierarchy (see Chomsky 1959a), but not since (see Boden 2006), it was understood that merely enumerating sentences was of no interest to the empirical science of natural language: "Along with a specification of the class *F* of grammars, a theory of language must indicate how, in general, relevant *structural information* can be obtained for a particular sentence generated by a particular grammar" (Chomsky 1959a: 138, emphasis added). Such a theory of the hierarchy has *yet* to be formulated. A novel mathematical model — a hierarchy revamped for strong generation — is necessary; hence the proposed TPLT.

With such mathematical formalism, I would not be "play[ing] mathematical games", but rather "describ[ing] reality" (Chomsky 1955: 81), which is mathematical: That is, I would not be precisifying mathematically a theory of a non-mathematical system, but formulating a mathematically precise theory of a cognitive system that is mathematical; the theory needs to be mathematical because the phenomenon is mathematical (see Turing 1954). With such apparatus, generative grammar substantializes the romantic intuition of language as "the infinite use of finite means" (von Humboldt 1836: 122).

Intuition is to be *explained* by a theory of linguistic computation; it is to be derived (as an effect), not stipulated (as a cause).¹⁹ This "move from the intuitive hints and examples of traditional grammar to explicit generative procedures" (Chomsky 1995: 24) was but a subroutine in the general research program starting with the demonstration that "the computable numbers include all numbers which could *naturally* be regarded as computable" (Turing 1936: 230, emphasis added). Indeed, in constructing abstract machines, Turing and Chomsky — like all theorists of algorithms — were exorcising ghosts: "If the

¹⁸ See Pinker (1979) for a discussion on the (un)decidability and (un)learnability of recursively enumerable sets.

¹⁹ Explanations must 'go all the way down', a nontrivial truism: "In imagining that there is a machine whose construction would enable it to think, to sense, and to have perception, one could conceive it enlarged while retaining the same proportions, so that one could enter into it, just like into a windmill. Supposing this, one should, when visiting within it, find only parts pushing one another, never anything by which to explain a perception" (Leibniz 1714: 83). If intuition were 'in control' of the cognitive-neurobiological windmill, an explanation for *it* would be necessary, so the explanation must go all the way down, but where it 'bottoms out' is a *mystery* — perhaps one of "[nature's] ultimate secrets [stored in] that obscurity, in which they ever did and ever will remain" (Hume 1763: 323).

grammar is [...] perfectly explicit — in other words, if it does not rely on the intelligence of the understanding reader but rather provides an explicit analysis of his contribution — we may [...] call it a *generative grammar*” (Chomsky 1965: 4); “[w]e may compare a man in the process of computing a real number to a machine which is only capable of a finite number of [configurations]. If at each stage the motion of the machine [...] is *completely* determined by the configuration, we shall call the machine an ‘automatic machine’” (Turing 1936: 231–232, emphasis original), now called a Turing machine.

With the ghost exorcised, the elementary components and procedures of generative grammar must reduce to the elementary components and procedures of the Turing machine, consistent with the thesis that if a function f is effectively calculable, then f is computable by a Turing machine, hence by the Universal Turing machine — the computer which can assume the form of any Turing machine M given the mathematical-functional description (machine table) of M as input.²⁰ The formulation of a generative grammar as a Turing machine would be novel to the TPLT and necessary for progress in discovering — and realizing artificially — the system for cognitive-biological computation.

Generative grammar was initially formulated as a Post-production system, but that project in “mathematical linguistics is [now] in a plateau”, having “got about as far as it could from the point of view of any impact on [empirical] linguistics” (Chomsky 2004b: 68–69) because now “virtually the entire subject [of mathematical linguistics] deals with weak generation; strong generation, while definable, is [now] too complex for much in the way of mathematical inquiry” (Chomsky 2007: 1097). Only strong generation encodes grammatical information, and therefore mathematical linguistics must be formulated in its terms to impact empirical linguistics.

One problem is that a formulation of strong generative capacity in Post-production systems is infeasible. However there is a different and deeper problem: Even strong generation is insufficient for empirical adequacy. As with weak generation, a grammar defined only in terms of strong generative capacity is a function defined in extension (i.e. defined in terms of expressions generated) when the desideratum of linguistic theory is a function defined in intension (i.e. defined as the generator of expressions). The reasons for this desideratum are self-evident. First, the enumeration of the extensional set is possible only by running the procedure for the intensional function; thus the logical priority of the latter over the former.²¹ Second, the ontological status of the sets of expressions is debated (see Chomsky 1986, 2001, Postal 2004, 2009), but the reality of the generator is manifest in the behavior of any normal human engaged in linguistic creativity, the ‘infinite use of finite means’. It is this function in intension that evolved in the brain; the complexity of Post-production systems cannot be posited with plausibility in theories of neurobiology and its evolution (see Ber-

²⁰ Turing discovered that a program can be a form of data. The profundity of this equation has not yet been appreciated in biolinguistics.

²¹ Gödel recognized the extensional equivalence of Herbrand-Gödel recursiveness, λ -calculus, and Turing machines, but by intensional analysis was convinced that, in his words, “the correct definition of mechanical computability was established beyond any doubt by Turing” (see Soare 2009).

wick 2011).²²

A novel theory of generative capacity is necessary, and here “Turing’s computability is intrinsically persuasive in the sense that the ideas embodied in it directly support the thesis that the functions encompassed are all for which there are algorithms” (Kleene 1981: 49). If the project is to unify the computational, algorithmic, and implementational levels of language — and ultimately intelligence — then the Turing machine is ideal, for it defines the mathematical-functional components and procedures that any computational system *must* realize: “[Turing’s] concepts underlying the design of computing machines arise out of a kind of *conceptual necessity*. [I]f one analyzes any computing machine that is powerful, fast, and efficient, one *will* find these concepts realized in its functional structure. [And yet Turing’s concepts] have been largely ignored in contemporary efforts to imagine how the brain might carry out the computations that the behavioral data imply it does carry out” (Gallistel & King 2009: 125, 144, emphases added).²³

Linguistic computation is demonstrably powerful, fast, and efficient, and yet contemporary biolinguistics — indeed all of theoretical linguistics and neurolinguistics — has largely ignored Turing’s concepts. If I-language were precisified as a form of Turing machine, our imagination for how I-language computes in mathematical abstraction and how it relates to concrete computation in a brain or machine would be profoundly expanded — so as, perhaps, to compass the truth. Such an expansion of the imagination would be the goal of the TPLT.

[Imagination] is that which penetrates into the unseen worlds around us, the worlds of Science [...]. Those who have learned to walk on the threshold of the unknown worlds [...] may then with the fair white wings of Imagination hope to soar further into the unexplored amidst which we live.

(Lady Lovelace, Babbage’s programmer, in Gleick 2011)

Research at the abstract computational level constrains research at the concrete implementational level: For instance, if research on the former level demonstrates that the mind *does* generate d-infinity, then it is incumbent upon research at the latter level of the brain to demonstrate *how*.^{24,25} Failure to concede

²² Some neurobiologists (e.g., Zylberberg *et al.* 2011), do posit ‘production rules’ in neuronal computation, but these are not Post-productions.

²³ I cannot overemphasize the fact that the Turing machine represents the *multiply realizable mathematical-functional structure* of a computer: “Since Babbage’s machine was not electrical, and since all digital computers are in a sense equivalent, we see that this use of electricity [in brains and computers] cannot be of theoretical importance [...]. If we wish to find [...] similarities we should look rather for *mathematical analogies of function*” (Turing 1950: 439, emphasis added); “[t]he brain is a purely physical object made up of completely sterile and inanimate components, all of which obey exactly the same laws as those that govern the rest of the universe. The key is not the *stuff* out of which brains are made, but the *patterns* that can come to exist inside the stuff of the brain. Brains are *media* that support complex patterns” (Hofstadter 1979: P-4, emphases original).

²⁴ Analogous logic has proved successful elsewhere in science and needs to be assumed in the sciences of mind and brain: “[Computational-Representational] theories [of the mind] will provide guidelines for the search for [neurophysiological] mechanisms, much as nineteenth-century chemistry provided crucial empirical conditions for radical revision of fundamental physics. The common slogan that ‘the mental is the neurophysiological at a higher level’ —

this top-down logic has retarded research, as with “[c]onnectionists draw[ing] their computational conclusions from architectural commitments”, which is fallacious given our limited understanding of the mathematical-functional meaning of neuronal architecture, “whereas computationalists draw their architectural conclusions from their computational commitments” (Gallistel & King 2009: ix), which is sound because of our greater understanding of computation, due to Turing (see Vaux & Watumull 2012 for a critique of connectionism as implemented in Optimality Theory phonology). The understanding of mental software thus precedes and conditions the understanding of neurophysiological hardware: “[I]t is the mentalistic studies that will ultimately be of greatest value for the investigation of neurophysiological mechanisms, since they alone are concerned with determining abstractly the properties that such mechanisms must exhibit and the functions they must perform” (Chomsky 1965: 193).

Hence it is necessary, as neurobiologists Zylberberg *et al.* (2011: 294) argue, for research at the lower level to “investigate which neural architecture could implement a Brain Turing Machine”, which neural architecture could implement “the concept of a production [...] essentially equivalent to the action performed by a Turing machine in a single step”, necessitating a “selection of productions [...] determined by the contents of working memory, which plays the role of the tape in the Turing machine [...]. Iteration of the cycle of production selection and action constitutes the basis of Turing-like programs”.²⁶ For the ‘Brain Turing Machine’ research program to succeed, the objects of inquiry must be identified by novel and precise higher level definitions of Turing’s components and procedures (e.g., ‘tape’, ‘production’, etc.). To formulate such definitions, defining objectives to unify the sciences of mind and brain would be my ambition for the TPLT. As of yet, I can proffer merely quasi-formal definitions and simplified (simplistic) prototypes of theories and models that in future research could answer big questions of minds, machines, and mathematics.

where C-R theories are placed within the ‘mental’ — has matters backwards. It should be rephrased as the speculation that the neurophysiological may turn out to be ‘the mental at a lower level’” (Chomsky 2000: 25–26).

²⁵ Given the evidence for cognitive computation, it is in my judgment logically necessary to assume the language faculty to be a form of Turing machine with addressable read/write memory. This logic applies to the brain generally (Gallistel & King 2009: 125, 105, i): If “the brain is an organ of computation[,] then to understand the brain one must understand computation”, which necessitates formalization; “[Turing] created a formalization that defined a class of machines”, with mathematical-functional components and procedures so elementary as to be multiply physically realizable. And “[b]y mathematically specifying the nature of these machines, and demonstrating their far-reaching capabilities, [Turing] laid a rigorous foundation for our understanding of what it means to say something is computable”. A formalization can in this way define conditions of adequacy that any theory in a particular domain of inquiry must satisfy to be true. Therefore, if it is demonstrated by research on higher levels of analysis that “brains are powerful organs of computation”, and that a formally definable “[addressable read/write] memory mechanism is indispensable in powerful computing devices”, then we must demand of neuroscience an implementational theory of an addressable read/write memory.

²⁶ Interesting and unorthodox work in neurobiology (e.g., Gallistel & King 2009, Hameroff & Penrose 1996) speculates that subcellular systems (e.g., in microtubules) could be efficiently implementing components and procedures of the Brain Turing Machine.

5. The Linguistic Turing Machine

A Turing machine is not a physical device — nor an intended model/simulation thereof — but rather a mathematical abstraction representing the functional conditions necessary and sufficient for any system — including the brain — to be computational. Let the components of the linguistic Turing machine L be a control unit, a read/write head, and a tape.

(4) *The linguistic Turing machine L is the 5-tuple $(Q, \Gamma, \delta, \#_s, \#_H)$* ^{27,28}

Q : Set of states/instructions (i.e., principles and parameters)

Γ : Set of symbols for syntactic objects (e.g., lexical items, phrases, etc.)

δ : $Q \times \Gamma \rightarrow Q \times \Gamma$ (i.e. transition function from state/symbol to state/symbol by search/merge)

$\#_s (\in \Gamma)$: Start (boundary) symbol

$\#_H (\in \Gamma)$: Halt (boundary) symbol

The potentially infinite bidirectional tape represents the inputs and outputs (memory) of L and stores its program (which can thereby be modified in language acquisition); for empirical adequacy (e.g., syntactic transformations), the symbolic memory of the tape can be structured as a stack with (a restricted form of) random access (NB: L is not a classical Turing machine; it is domain-specific and could be equipped with multiple tapes/stacks, etc.). The control unit is a transition function (machine table) defining a finite set of states and a finite set of instructions for the operation of L given a state and an input represented on the tape and/or possibly a symbol on the stack memory; the states and instructions are defined by universal linguistic principles and parameterization for particular natural languages (see Roberts 2011 for a computational theory of parametric variation with principled constraints). The read/write head is a search/merge procedure: Effectively, a rewrite rule of the form in (1); it reads (searches for) an input symbol (or symbols) on the tape/stack and, as commanded by the control unit, writes (merges) an output symbol on the tape (equivalently, pushes a

²⁷ Different but equivalent definitions could be formulated. A formulation of the search/merge procedure in Polish notation could precisify the efficiency of its execution (Randy Gallistel, p.c.).

²⁸ States for acceptance and rejection are not specified here, *tentatively*, because the acceptance (as grammatical) and rejection (as ungrammatical) of expressions are, *arguably*, functions of interpretability conditions at the interfaces with the conceptual-intentional and sensory-motor systems; the generation and interpretation of an expression are, *arguably*, independent processes in principle (but in practice the latter can condition the former). And at the interfaces, an approximately continuous scale of deviance — not discrete states of grammaticality (acceptance/rejection) — would probably be required: Some technically grammatical sentences are intuitively rejectionable and some intuitively acceptable sentences are technically ungrammatical and some sentences are, on some passes, of indeterminate status. As Putnam (1961) argues, and I concur, grammaticality judgments could be formally equivalent to decidability problems; perhaps, therefore, some judgments are technically undecidable, but this must be investigated in the TPLT. In addition to these complexities are Kripke's (1982) 'Wittgensteinian problems' (see Chomsky 1986), such as the issue of ungrammaticalities being 'mistakes', cases of not 'following the rules' of grammaticality; we can wonder whether a linguistic Turing machine could even 'make a mistake' (see Turing 1950). These would be interesting problems to address in the TPLT.

symbol onto the stack) in a stepwise process; and the rules themselves, stored as a program, can be written-to/read-from — a fact with profound implications for theories of the evolution of I-language, language change, language growth/acquisition, syntactic computations, *inter alia*, to be investigated in the TPLT.

To recapitulate, a Turing machine is a mathematical abstraction, not a physical device, but my theory is that the information it specifies in the form of I-language (4) must be encoded in the human genetic program — and/or given by mathematical law — and expressed in the brain. Central to the machine is the search/merge procedure, and yet “[a]lthough the characterizations of what might be the most basic linguistic operations must be considered one of the deepest and most pressing in experimental language research, we know virtually nothing about the neuronal implementation of the putative primitives of linguistic computation” (Poeppel & Omaki 2008: 246). Thus is presented the great challenge for the TPLT: To precisify (formalize) the definitions of linguistic primitives so that ‘linking hypotheses’ (not mere correlations) to neurobiological primitives — constitutive of a Brain Turing Machine — can be formed. The beauty of the Turing machine is that in its abstractness it subsumes and thereby relates all computational primitives; in this way it could render commensurable the computational ontologies of linguistics and neuroscience. And the key to linking these ontologies is the generative procedure.

Rewrite rules of the form in (1) were defined in Chomsky normal form (5); this can be updated for L , with the search/merge procedure — call it *Merge* — formulated as in (6).

- (5) Chomsky normal form: Uppercases represent nonterminals (e.g., S , VP , NP , etc.), lowercases represent terminals (e.g., *the*, *man*, *took*, etc.), ϵ represents the empty string.

$$\begin{aligned} A &\rightarrow BC \\ A &\rightarrow a \\ S &\rightarrow \epsilon \end{aligned}$$

- (6) Merge is a set-formation function: The syntactic objects α , β , γ can be simple (e.g., lexical items) or complex (e.g., phrases) such that the nonterminal/terminal distinction is unformulable; consequently, the system is simplified and generalized.^{29,30}

$$\begin{aligned} \#_s &\rightarrow \{\alpha, \#_s\} \\ \{\alpha, \#_s\} &\rightarrow \{\beta, \{\alpha, \#_s\}\} \\ \{\gamma, \#_s\} &\rightarrow \{\#_{Hr}, \{\gamma, \#_s\}\} \end{aligned}$$

²⁹ As a component of a Turing machine, the Merge procedure is necessary and sufficient to implement the elementary operations of arithmetic (consistent with the composition of functions and set-theoretic definitions of the natural numbers); and interestingly, with Merge defined as a binary operator on the set of syntactic objects, L can be formulated as a free magma in universal algebra.

³⁰ Boundary (Start/Halt) symbols are logically necessary for any computational system and can be demonstrated (see Watumull 2012) to solve problems with syntactic structures (e.g., labeling and linearizing the first merged elements, transferring/spelling-out the final phase, structuring inter-sentential coordination).

To simplify: $\text{Merge}(X, Y) = \{X, Y\}$.³¹ A simplified derivation by Merge of *the man took the book* in (3) assumes the form in (7): The predicate and subject are generated by Merge in parallel and then merged:

- (7) $\#_s$ $\#_s$
 $\{book, \#_s\}$ $\{man, \#_s\}$
 $\{the, \{book, \#_s\}\}$ $\{the, \{man, \#_s\}\}$
 $\{took, \{the, \{book, \#_s\}\}\}$ $\{\#_H, \{the, \{man, \#_s\}\}\}$
- $\{\{\#_H, \{the, \{man, \#_s\}\}\}, \{took, \{the, \{book, \#_s\}\}\}\}$
 $\{\#_H, \{\{\#_H, \{the, \{man, \#_s\}\}\}, \{took, \{the, \{book, \#_s\}\}\}\}\}$

It is obvious from (7) that strongly generative — structure (set) forming — recursion is necessary to generate any nontrivial expression: Information from step i must be carried forward for merger at step $i + 1$; as a Turing machine, L “carries symbolized information forward [...], making it accessible to computational operations” (Gallistel & King 2009: 122).³² A finite state automaton, which “cannot write to tape [or push to a stack] cannot store the results of its computations in memory for use in subsequent computations” (Gallistel & King 2009: 122), is provably inadequate (see Chomsky 1956). Nor is a connectionist network possible (probably), because a look-up table architecture — to which connectionist networks reduce (arguably) — cannot (or can only inefficiently) implement numerous linguistic phenomena such as d-infinity and phonological rules (see Vaux & Watumull 2012).

Given the generality of Merge, it is not obvious that I-language is *not* as powerful as an unrestricted Turing machine. Nor is it obvious that I-language is computable in general. The implications of incomputability would obviously be profound for our understanding not only of language, but of nature generally (see Cooper 2012) — implications to be explored in the TPLT. It is intriguing though that in the history of generative linguistics, solutions to empirical problems reducible to but usually unrecognized as incomputability problems unfailingly but unnoticeably reestablished the computability of I-language.

³¹ Merge entails ‘external’ and ‘internal’ forms: $EM(X, Y \mid X \notin Y \wedge Y \notin X) = \{X, Y\}$; $IM(X, Y \mid X \in Y \vee Y \in X) = \{X, Y\}$. EM corresponds to phrase structure rules (and argument structure semantics), IM to transformation/movement rules (and discourse semantics).

³² Recursively generated hierarchy can be super-sentential (i.e. Merge can structure discourse representations) and realized in pushdown automata:

In conversational storytelling, [...] embedded stories are normally flashbacks or flashaheads. The system [in state q] recognizes flash markers on the incoming clause, normally pluperfect marking for flashbacks (accompanied perhaps by a temporal deictic comment like ‘before this...’) and present progressive or future tense plus comments for flashaheads. The flash marker effects a PUSH to the embedded story [state q'] which parses subsequent clauses as [stative] or [eventive] and stores them in [...] registers. When a clause arrives for parsing which indicates the end of the flash with a ‘so,’ ‘anyway,’ ‘and then,’ [etc.], which act as POP markers, the parse of the embedded unit is complete and the parser returns to [q] to accept further input.

(Polanyi & Scha 1983: 164–165)

The recursive generation of discourse would be a research priority in the TPLT as it has not been in generative linguistics hitherto.

Indeed, I would argue — and expound in the TPLT — that the computability of I-language would be strong evidence against ‘Super-Turing machines’ and ‘hyper-computation’ (in I-language and beyond); I would concur that the existence of non-Turing processes is a “myth” (Davis 2004). So I assume that with additional work, I could prove Merge (as formulated in (6)) to be computable and tractable (i.e. efficient in being upper-bounded by a polynomial function) and optimal.³³

The emergence of optimality — even beauty — fascinated Turing (1952); he thus would have been interested in the fact that *binary* Merge is optimal in minimizing formal representations and spatiotemporal resources — measurable by Kolmogorov complexity and similar concepts — in the process of maximizing the strong generation of syntactic structures, consistent with the minimax theorem (see von Neumann 1928, Watumull 2010).^{34,35}

It has been demonstrated mathematically (see Turing 1952 on morphogenesis, Mandelbrot 1982 on fractals, Chaitin 2012 on natural selection) and simulated computationally (see Minsky 1985 on Turing machines, Wolfram 2002 on cellular automata) that evolution does tend to converge on simple procedures capable of generating infinite complexity. These demonstrations and simulations need to be applied in biolinguistics and unified in explanation; this would be a project in the TPLT.

Watumull *et al.* (in preparation) conduct a rigorous mathematical and biological analysis of the hypothesis that a “computational mechanism of recursion”, such as Merge, “is recently evolved and unique to our species” and unique to language (Hauser *et al.* 2002: 1573).³⁶ It has been assumed, dubiously, that this controversial hypothesis can be adjudicated empirically with counter-evidence that some language does not run recursive computation (see Everett 2005) or that linguistic universals do not exist at all (see Evans & Levinson 2009). This assumption is false (as Watumull *et al.* explain), but granting it *arguendo*, my formalization of I-language as the linguistic Turing machine *L* entails the properties of the humanly unique “computational mechanism of recursion” as one of many linguistic universals; this is satisfying because “[t]he most interesting contribution a generative grammar can make to the search for universals of language is specify formal systems that have putative universals as *consequences*, as opposed to merely providing a technical vocabulary in terms of which autonomously stipulated universals can be expressed” (Gazdar *et al.* 1985: 2, emphasis original).

As formulated in (6), Merge is necessarily recursive, automatically generative of hierarchically structured expressions over which set-theoretic relations relevant to linguistic cognition (e.g., c-command, probe-goal) are defined, not stipulated. These properties derive as logical consequences from the formulation of Merge as the simplest read/write mechanism — logically necessary in any pow-

³³ See Watumull (2012) for examples of (in)computable linguistic phenomena, arguments against Super-Turing machines, and the optimality/efficiency of syntactic computations.

³⁴ Syntactic structures are even fractal (see Watumull & Eglash 2011).

³⁵ I argue (Watumull 2010) that Merge is a *minimax function*: Effectively unary (compact) but binary (combinatorial).

³⁶ In brief, the sense of *recursion* in Hauser *et al.* (2002) is that of an effective procedure fixing a computable function that enumerates and maps to interfacing systems an infinite set of possible expressions.

erful/productive and economical/efficient formal system: Recursively generated hierarchy defined by structural (grammatical) relations is *deduced* — not merely predicted — to be a linguistic universal. To formalize this informal deduction, deriving additional linguistic universals, would be at the core of the TPLT.^{37,38}

6. Acquisition

One of the biggest questions for the TPLT is the derivation of knowledge, a question embedded in the even bigger question “whether there is a characteristic human psychological phenotype (‘human nature’ in earlier editions) that can be attributed to a characteristic human genetic endowment” (Fodor 2001: 102). If the thesis of the TPLT that I-language is a type of Turing machine is correct, then it is evident that the elementary components and procedures of the linguistic Turing machine *L* must be preinstalled in — genetically specified for — the brain and/or given by mathematical law, and thus characteristic of humans. Equally evident is the fact that variation exists in its implementation and execution: I-language is a species-typical property of *individuals*, so technically there exist as many languages as there exist (and have existed) individuals, and if a sufficient number of individuals converge (approximately) on a specific implementation and execution of *L*, they are classified as knowing the ‘same’ language (e.g., ‘English’, ‘Swahili’, etc.). This variation is a function of the fact that some parts of I-languages must be acquired (‘learned’) in an astonishing process that is not yet understood by science:

A normal child acquires [linguistic] knowledge on relatively slight exposure and without specific training. He can then quite effortlessly make use of an intricate structure of specific rules and guiding principles to convey his thoughts and feelings to others, arousing in them novel ideas and subtle perceptions and judgments. For the conscious mind [of the scientist], not specially designed for the purpose, it remains a distant goal to reconstruct and comprehend what the child has done intuitively and with minimal effort. Thus language is a mirror of mind in a deep and significant sense. It is a product of human intelligence, created anew in each individual by operations that lie far beyond the reach of will or consciousness.

(Chomsky 1975: 4)

This is but a particular articulation of the general question posed by Bertrand Russell (1948: v, emphasis original): “*How comes it that human beings, whose contacts with the world are brief and personal and limited, are nevertheless able to know as*

³⁷ A profound question in computational linguistics and computational neuroscience is whether linguistic rules are structure-dependent (defined over sets) or linear-dependent (defined over strings) and why (see Perfors *et al.* 2011, Berwick *et al.* 2011). My answer (i.e. rules are structure-dependent because of the set-theoretic relations entailed by the running of Merge), can be deduced if the Turing machine *L* does in fact represent linguistic computation. I would show this in the TPLT.

³⁸ An additional property of language predicted — deduced — to be universal on my model is that of binarity: Merge as defined for *L* is a binary function such that the expressions it generates are composed of (embedded) 2-sets. I can demonstrate that Merge of any arity *n*, *n* > 2, is either incomputable or intractable (lower bounded by an exponential function) (see Watumull 2012). This is interesting because many natural languages are argued to contain non-binary structures.

much as they do know?". Any theory of mind and machine needs to answer this question. And by dint of logic — let alone empirical observation — the answer needs to be *specific* to the cognitive domain:

From a computational point of view, the notion of a *general purpose* learning process (for example, associative learning), makes no more sense than the notion of a general purpose sensing organ. [P]icking up information from different kinds of stimuli — light, sound, mechanical, and so on — requires organs with structures shaped by the specific properties of the stimuli they possess [...]. Learning different things about the world from different kinds of experience requires computations tailored both to what is to be learned and to the kind of experience from which it is to be learned [...]. No one would suppose that [a path integration learning organ] would be of any use in the learning of a language. Applying this learning organ to that learning problem would be like trying to hear with an eye or breath with the liver.

(Gallistel 2010: 194–196, emphasis added)

Language acquisition is patently a computational process necessitating procedures for the recognition and analysis of linguistic data — out of the “blooming, buzzing confusion” (James 1890: 488) of sensory stimuli — so as to map the information into the linguistic Turing machine, which must, within a “critical period” (see Lenneberg 1967, Pinker 1984, 1994), parameterize the initial state of its control unit — interpolating language-particular instructions between (genetically-installed/mathematically-given) language-universal principles — so as to generate outputs consistent with the inputs and ultimately attain (grow) a steady state of linguistic competence. To discover the mechanisms of this process, its complexity can be abstracted into the simplicity of a Turing machine, but as of yet, the abstraction has not been conducted — another job for the TPLT.

Reinforcing the classic argument from the poverty of the stimulus (see Chomsky 1980) — that is, an innate language acquisition device LAD is necessary because the linguistic competence the child displays is underdetermined by the linguistic data to which it is exposed — is an argument from the complexity of the computation (see Aaronson 2011). For instance, let C be a child of generation g unequipped with the domain-specific LAD and exposed to a set of n -bit strings (sentences) weakly generated by paths in a nondeterministic finite automaton M (the grammar of generation $g-1$) with a number of states $< n$.³⁹ It has been proved mathematically (see Kearns & Valiant 1994) that if C can reconstruct M — if the child can form the grammar — then C can break the RSA cryptosystem, which is technically intractable.⁴⁰ “The grammar of any real human language is much too rich and complicated to be captured by a finite automaton. So this result is saying that even learning the least expressive, unrealistically simple language is already as hard as breaking RSA” (Aaronson 2011). Neither C nor any efficient algorithm

³⁹ The primary linguistic data to which the child is exposed are strings from which the strongly generated structures must be reconstructed. But this reconstruction is technically an ‘ill-posed problem’ if no LAD is assumed — indeed incomputable in general as stated in Rice’s Theorem (see Batchelder & Wexler 1979 for a related discussion).

⁴⁰ RSA encodes messages on the Internet by raising the encoded message to an exponent modulo a composite (semiprime) number p with neither of its prime factors given. To decode the message, it is necessary (though not sufficient) to factor p into two primes.

yet formulated by computer scientists can break RSA,⁴¹ but C (assuming it to be a normal child) attains/grows an I-language. It follows that C must be genetically preprogrammed with some type of LAD.⁴²

The LAD could run the ‘semantic bootstrapping’ algorithm rigorously and empirically defined and expounded by Steven Pinker (1984). The LAD (genetically) en-codes a set of semantic primitives that recognize syntactic categories (genetically specified) given in the linguistic input. With the categories recognized, the child can infer some of the rules particular to the language expressed in the input; the child can then apply these rules to additional input, recognizing additional categories and acquiring additional rules, etc., so as ultimately to attain (grow) the steady state. This bootstrapping logic was formalized (Pinker 1984), but needs to be updated into a Turing machine to prove its success in satisfying conditions any theory of the LAD must satisfy (Pinker 1979): (i) *learnability* (languages are acquired); (ii) *equipotentiality* (the LAD must run for all natural languages); (iii) *time* (languages are acquired in efficient time); (iv) *input* (the input information in the theory must be realistic); (v) *development* (the stages of acquisition are predict-able); (vi) *cognitive* (the theory of the LAD must be compatible with theories of extralinguistic cognitive faculties). These conditions could be efficiently satisfied in a PAC algorithm (see Valiant 1984) running in a hierarchy of linguistic parameters (see Roberts 2011).

A PAC (“probably approximately correct”) model is an efficient, polynomial (not exponential) bounded algorithm presented with a set of (random) points x_1, \dots, x_n from some set S with the points classified $f(x_1), \dots, f(x_n)$. The model succeeds if the function f is inferred such that $f(x)$ can be predicted for the majority of future points $x \in S$. The function f is a member of a hypothesis class H (the set of possible hypotheses). The form of H is important:

[T]he PAC approach has the advantage [over a Bayesian approach] of requiring only a *qualitative* decision about which hypotheses one wants to consider, rather than a *quantitative* prior over hypotheses.

(Aronson 2011, emphases original)

For language (see de Wolf 1997), let the domain S be the set of all possible sentences (which in the TPLT I would define in terms of strong generation); a function f is a grammar such that H is the set of possible grammars (genetically specified and mathematically governed); the child needs to succeed in converging on f given a set of sentences x_1, \dots, x_n . The algorithm can *query* an *oracle* device (Turing 1939) as to the membership of a given x in S for a given f ; an oracle is a ‘black box’ connected to a Turing machine to which queries can be entered and from which answers are returned. For language, the child does not query persons in its external environment (obviously), but (unconsciously) of internal representations of syntactic structure; those representations genetically specified constitute the ‘oracular information’. (‘Universal Grammar’ is a theory of the genetically installed linguistic oracle.) Polynomial PAC acquisition of type

⁴¹ I set aside Shor’s quantum algorithm (see Shor 1997).

⁴² See Pinker (2002) on the problems with presuming the child brain to be equipped with “little mechanism and lots of blank sheets” such that a learning machine can be “easily programmed” (Turing 1950: 456).

3 languages is possible only with membership queries (see Angluin 1987); it has not been proved — *yet* — that type 2 or type 1 languages can be acquired even with membership queries.⁴³

A proof could be possible if these grammars are reformulated in the linguistic Turing machine because the structures Merge generates in parsing (i.e. the reverse derivations the child performs in recognizing and analyzing linguistic data) necessitate querying to be interpretable (i.e. mappable into the initial (genetically determined) state of I-language). Some answers are given in the structure (e.g., those defining arrangements of phrases, etc.), obviating oracle queries; residual answers need to be given in the genetic endowment as oracular information (e.g., the semantic and syntactic primitives of Pinker 1984) and/or derived from mathematical law. So designed, acquisition is the process of running a bootstrapping PAC algorithm on a hierarchy of queries, the answers to which specify the setting of language-particular parameters. Ian Roberts has established and is expounding a rigorous research program on parameter hierarchies that derives linguistic universals. Thus, it could be that the “innate [linguistic system] contains little more than the single combinatorial operation Merge and a schema for syntactic categories and features” (Roberts 2011).

In sum, I-language could reduce to the Turing machine *L*: To prove this — and ergo prove the possibility of a unification via mathematics of mind and machine — would be the triumph of the TPLT.

7. The Big Answers

The success of the TPLT, however small, could constitute a proof of concept for a strategy to answer the Big Question of the Turing Centenary as to whether there could be “a successful mathematical model of intelligent thought”. I-language is an organ of human intelligence, and I submit that a successful mathematical model of linguistic cognition is within reach.

[O]ut of a lot of the very interesting work now going in [biolinguistics] there may come a sharpening and clarification of notions which would make it possible to undertake a new mathematical study, and maybe a lot of the questions that were asked about phrase-structure grammars [in the 1950s–1960s] could now be re-asked and many new questions could be asked about these new [...] systems. At that point one would hope that there would develop a new upsurge of mathematical linguistics and at that point I think [...] mathematical logic might very well be relevant. Ideas from recursive function theory and model theory and the theory of more elaborate logics seem at least to be in the same neighborhood as the kind of formal questions suggested by language study and it might turn out that a new field would develop out of this convergence. (Chomsky 2004b: 69)

To work for such a convergence would be to share and substantiate Turing’s dream that not only can we model nature mathematically, but we may even discover that, fundamentally, nature (reality) *is* mathematical.

⁴³ Until the Chomsky hierarchy is updated in the TPLT in terms of strong generation, it is not possible to accept or reject these proofs as conclusive.

References

- Aaronson, Scott. 2011. Why philosophers should care about computational complexity. Ms., MIT, Cambridge, MA.
- Angluin, Dana. 1987. Learning regular sets from queries and counterexamples. *Information and Computation* 75, 87–106.
- Batchelder, William H. & Kenneth Wexler. 1979. Suppes' work in the foundations of psychology. In Radu Bogdan (ed.), *Patrick Suppes*, 149–186. Dordrecht: Reidel.
- Berwick, Robert C. 2011. Syntax *facit saltum redux*: Biolinguistics and the leap to syntax. In Anna Maria Di Sciullo & Cedric Boeckx (eds.), *The Biolinguistic Enterprise*, 65–99. Oxford: Oxford University Press.
- Berwick, Robert C., Paul Pietroski, Beracah Yankama & Noam Chomsky. 2011. Poverty of the stimulus revisited. *Cognitive Science* 35, 1207–1242.
- Bloomfield, Leonard. 1926. A set of postulates for the science of language. *Language* 2, 153–164.
- Boden, Margaret. 2006. *Mind as Machine: A History of Cognitive Science*. Oxford: Oxford University Press.
- Carnap, Rudolf. 1955. Meaning and synonymy in natural languages. *Philosophical Studies* 6, 33–47.
- Chaitin, Gregory. 2012. *Proving Darwin: Making Biology Mathematical*. New York: Pantheon.
- Chomsky, Noam. 1951. The morphophonemics of modern Hebrew. Master's thesis, University of Pennsylvania, Philadelphia.
- Chomsky, Noam. 1955. The logical structure of linguistic theory. Ms., Harvard University/Massachusetts Institute of Technology, Cambridge, MA.
- Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions of Information Theory* IT-2: 113–124.
- Chomsky, Noam. 1959a. On certain formal properties of grammars. *Information and Control* 2, 137–167.
- Chomsky, Noam. 1959b. Review of *Verbal Behavior*, by B.F. Skinner. *Language* 35, 26–57.
- Chomsky, Noam. 1963. Formal properties of grammars. In R. Duncan Luce, Robert R. Bush & Eugene Galanter (eds.), *Handbook of Mathematical Psychology II*, 323–418. New York: Wiley and Sons.
- Chomsky, Noam. 1964. *Current Issues in Linguistic Theory*. The Hague: Mouton.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 1975. *Reflections on Language*. New York: Pantheon.
- Chomsky, Noam. 1980. *Rules and Representations*. New York: Columbia University Press.
- Chomsky, Noam. 1986. *Knowledge of Language: Its Nature, Origins and Use*. New York: Praeger.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, Noam. 2001. Derivation by phase. In Michael Kenstowicz (ed.), *Ken*

- Hale: A Life in Language*, 1–52. Cambridge, MA: MIT Press.
- Chomsky, Noam. 2002. *On Nature and Language*. Cambridge: Cambridge University Press.
- Chomsky, Noam. 2004a. Beyond explanatory adequacy. In Adriana Belletti (ed.), *Structures and Beyond: The Cartography of Syntactic Structures*, vol. 3, 104–131. Oxford: Oxford University Press.
- Chomsky, Noam. 2004b. *Language and Politics* (ed. by Carlos Otero). Oakland, CA: AK Press.
- Chomsky, Noam. 2005. Three factors in language design. *Linguistic Inquiry* 36, 1–22.
- Chomsky, Noam. 2007. Review of Margaret Boden's *Mind as Machine: A History of Cognitive Science*. *Artificial Intelligence* 171, 1094–1103.
- Chomsky, Noam. 2009. Turing on the 'Imitation Game'. In Robert Epstein, Gary Roberts & Grace Beber (eds.), *Parsing the Turing Test*, 103–106. Amsterdam: Springer.
- Chomsky, Noam & George A. Miller. 1963. Introduction to the formal analysis of natural languages. In R. Duncan Luce, Robert R. Bush & Eugene Galanter (eds.), *Handbook of Mathematical Psychology II*, 269–321. New York: Wiley and Sons.
- Church, Alonzo. 1936. An unsolvable problem of elementary number theory. *American Journal of Mathematics* 58, 345–363.
- Church, Alonzo. 1941. *The Calculi of Lambda Conversion*. Princeton, NJ: Princeton University Press.
- Cooper, Barry. 2012. The incomputable reality. *Nautre* 482, 465.
- Davis, Martin. 1958. *Computability and Unsolvability*. New York: McGraw-Hill.
- Davis, Martin (ed.). 1965. *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvability Problems and Computable Functions*. New York: Raven Press.
- Davis, Martin. 2004. The myth of hypercomputation. In Christof Teuscher (ed.), *Alan Turing: Life and Legacy of a Great Thinker*, 195–212. Berlin: Springer.
- Davis, Martin. 2012. *The Universal Computer: The Road from Leibniz to Turing*. New York: Taylor & Francis Group.
- Deutsch, David. 2011. *The Beginning of Infinity*. London: Allen Lane.
- Evans, Nicholas & Stephen C. Levinson. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences* 32, 429–492.
- Everett, Daniel L. 2005. Cultural constraints on grammar and cognition in Pirahá: Another look at the design features of human language. *Current Anthropology* 46, 621–646.
- Fodor, Jerry A. 2001. Doing without what's within: Fiona Cowie's critique of nativism. *Mind* 110, 99–148.
- Gallistel, C. Randy. 2010. Learning organs. In Jean Bricmont & Julie Franck (eds.), *Chomsky Notebook*, 193–202. New York: Columbia University Press.
- Gallistel, C. Randy & Adam Philp King. 2009. *Memory and the Computational Brain: Why Cognitive Science Will Revolutionize Neuroscience*. New York: Wiley-Blackwell.
- Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum & Ivan A. Sag. 1985. *Generalized Phrase Structure Grammar*. Cambridge, MA: Harvard University Press.

- Gleick, James. 2011. *The Information*. New York: Pantheon.
- Gödel, Kurt. 1934 [1965]. On undecidable propositions of formal mathematical systems. In Martin Davis (ed.), *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvability Problems and Computable Functions*, 39–74. New York: Raven Press.
- Hammeroff, Stuart R. & Roger Penrose. 1996. Orchestrated reduction of quantum coherence in brain microtubules: A model for consciousness. In Stuart R. Hameroff, Alfred W. Kaszniak & Alwyn C. Scott (eds.), *Toward a Science of Consciousness: The First Tucson Discussions and Debates*, 507–540. Cambridge, MA: MIT Press.
- Hauser, Marc D. 2009. The possibility of impossible cultures. *Nature* 460, 190–196.
- Hauser, Marc D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298, 1569–1579.
- Hofstadter, Douglas R. 1979. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.
- von Humboldt, Wilhelm. 1836 [1989]. *On Language: The Diversity of Human Language-Structure and its Influence on the Mental Development of Mankind*. Cambridge: Cambridge University Press.
- Hume, David. 1763. *The History of England: From the Invasion of Julius Caesar to the Revolution in 1688*. Google eBook.
- Jackendoff, Ray. 2007. *Language, Consciousness, Culture: Essays on Mental Structure*. Cambridge, MA: MIT Press.
- James, William. 1890. *The Principles of Psychology*. New York: Holt.
- Kearns, Michael J. & Leslie G. Valiant. 1994. Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of American Academy of Mathematics* 41, 67–95.
- Kleene, Stephen C. 1936. General recursive functions of natural numbers. *Mathematische Annalen* 112, 727–742.
- Kleene, Stephen C. 1980. The theory of recursive functions, approaching its centennial. *Bulletin of the American Mathematical Society* 5, 43–61.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- Lashley, Karl S. 1951. The problem of serial order in behavior. In Lloyd A. Jeffress (ed.), *Cerebral Mechanisms of Behavior*, 112–146. London: Chapman & Hall, Limited.
- Leibniz, Gottfried Wilhelm. 1714 [1992]. Monadology. In Nicholas Rescher (ed.), *G.W. Leibniz's Monadology*, 45–308. Pittsburgh, PA: University of Pittsburgh Press.
- Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.
- Mandelbrot, Benoit B. 1982. *The Fractal Geometry of Nature*. New York: W.H. Freeman.
- McCarthy, John. 1956. Inversion of functions defined by Turing machines. In Claude E. Shannon & John McCarthy (eds.), *Automata Studies*, 177–182. Princeton, NJ: Princeton University Press.
- Miller, George A., Eugene Galanter & Karl H. Pribram. 1960. *Plans and the Struc-*

- ture of Behavior*. New York: Holt, Rhinehart, & Winston.
- Minsky, Marvin. 1963. Steps toward artificial intelligence. In Edward A. Feigenbaum & Julian Feldman (eds.), *Computers and Thought*, 406–450. New York: McGraw-Hill.
- Minsky, Marvin. 1985. Why intelligent aliens will be intelligible. In Edward Regis (ed.), *Extraterrestrials*, 117–128. Cambridge, MA: MIT Press.
- von Neumann, John. 1928. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* 100, 295–320.
- Penrose, Roger. 1989. *The Emperor's New Mind: Concerning Computers, Minds and the Laws of Physics*. Oxford: Oxford University Press.
- Pinker, Steven. 1979. Formal models of language learning. *Cognition* 7, 217–283.
- Pinker, Steven. 1984. *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.
- Pinker, Steven. 1994. *The Language Instinct: The New Science of Language and Mind*. New York: Penguin.
- Pinker, Steven. 2002. *The Blank Slate: The Modern Denial of Human Nature*. New York: Penguin.
- Poeppl, David & David Embick. 2005. Defining the relation between linguistics and neuroscience. In Anne Cutler (ed.), *Twenty-First Century Psycholinguistics: Four Cornerstones*, 103–118. Mahwah, NJ: Lawrence Erlbaum.
- Poeppl, David & Akira Omaki. 2008. Language acquisition and ERP approaches: Prospects and challenges. In Angela D. Friederici & Guillaume Thierry (eds.), *Early Language Development*, 231–253. Amsterdam: John Benjamins.
- Polanyi, Livia & Remko Scha. 1983. On the recursive structure of discourse. In K. Ehlich & Henk van Riemsdijk (eds.), *Connectedness in Sentence, Discourse and Text*, 141–178. Tilburg: Tilburg University.
- Post, Emil. 1936. Finite combinatory processes: Formulation 1. *Journal of Symbolic Logic* 1, 103–105.
- Post, Emil. 1944. Recursively enumerable sets of positive integers and their decision problems. *Bulletin of the American Mathematical Society* 50, 284–316.
- Post, Emil. 1947. Recursive unsolvability of a problem of Thue. *The Journal of Symbolic Logic* 12, 1–11.
- Postal, Paul M. 2004. *Skeptical Linguistic Essays*. Oxford: Oxford University Press.
- Postal, Paul M. 2009. The incoherence of Chomsky's 'biolinguistic' ontology. *Biolinguistics* 3, 104–123.
- Putnam, Hilary. 1961. Some issues in the theory of grammar. *Structure of Language and Its Mathematical Aspects: Proceedings of the Symposium in Applied Mathematics: American Mathematical Society* 12, 25–42.
- Putnam, Hilary. 1975. *Mind, Language and Reality: Philosophical Papers*, vol. 2. Cambridge: Cambridge University Press.
- Roberts, Ian. 2011. Parametric hierarchies. Ms., University of Cambridge.
- Russell, Bertrand. 1948. *Human Knowledge: Its Scope and Limits*. London: George Allen & Unwin.
- Simon, Herbert A. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* 106, 467–482.
- Shor, Peter W. 1997. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* 26,

- 1484–1509.
- Soare, Robert I. 2009. Turing oracle machines, online computing, and three displacements in computability theory. *Annals of Pure and Applied Logic* 160, 368–399.
- Turing, Alan M. 1936. On computable numbers, with an application to the *Entscheidungsproblem*. *Proceedings of the London Mathematical Society* 42, 230–265.
- Turing, Alan M. 1939. Systems of logic based on ordinals. *Proceedings of the London Mathematical Society* 45, 161–228.
- Turing, Alan M. 1947 [2004]. Lecture on the automatic computing engine. In B. Jack Copeland (ed.), *The Essential Turing*, 378–394. Oxford: Oxford University Press.
- Turing, Alan M. 1950. Computing machinery and intelligence. *Mind* 59, 433–460.
- Turing, Alan M. 1952. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences* 237, 37–72.
- Turing, Alan M. 1954. Solvable and unsolvable problems. *Science News* 31, 7–23.
- Valiant, Leslie G. 1984. A theory of the learnable. *Communications of the American Academy of Mathematics* 27, 1134–1142.
- Vaux, Bert & Jeffrey Watumull. 2012. The phonological Turing machine. Ms., University of Cambridge & MIT, Cambridge, MA.
- Watumull, Jeffrey. 2010. Merge as a minimax solution to the optimization problem of generativity. MPhil Thesis, University of Cambridge.
- Watumull, Jeffrey. 2012. The computability and complexity of generativity. Ms., MIT, Cambridge, MA.
- Watumull, Jeffrey & Ron Eglash. 2011. Fractal syntax. Ms., MIT, Cambridge, MA & Rensselaer Polytechnic Institute, Troy, NY.
- Watumull, Jeffrey, Marc D. Hauser & Robert C. Berwick. In preparation. Problems and solutions for comparative research on the evolution of syntactic computations.
- Weinberg, Steven. 1976. The forces of nature. *Bulletin of the American Academy of Arts and Sciences* 29, 13–29.
- Wolf, Ronald de. 1997. Philosophical applications of computational learning theory: Chomskyan innateness and Occam's razor. Master's thesis, Erasmus University Rotterdam.
- Wolfram, Stephen. 2002. *A New Kind of Science*. Champagne, IL: Wolfram Media.
- Zylberberg, Ariel, Stanislas Dehaene, Pieter R. Roelfsema & Mariano Sigman. 2011. The human Turing machine: A neural framework for mental programs. *Trends in Cognitive Science* 15, 293–300.

Jeffrey Watumull
University of Cambridge
Department of Theoretical and Applied Linguistics
Sidgwick Avenue
Cambridge CB3 9DA
UK
jw647@cam.ac.uk

Massachusetts Institute of Technology
Department of Linguistics and Philosophy
77 Massachusetts Avenue, Building 32
Cambridge, MA 02139
USA
watumull@mit.edu

Biolinguistics

Volume 6, Issue 2

Spring 2012

TABLE OF CONTENTS

- 124 The Biolinguistics of Autism: Emergent Perspectives
Nicolas Bourguignon
University of Montreal
Aparna Nadig
McGill University
Daniel Valois
University of Montreal
- ★ BRIEFS ★
- 166 Ever Since Dennett: On the Origins of Biolinguistics
Hans-Martin Gärtner
Hungarian Academy of Sciences
- ★ REVIEWS ★
- 168 Lights and Shadows in the Evolution of Language
Evelina Leivada
Universitat de Barcelona
Ana M. Suárez
Universidad Autónoma de Madrid
- ★ FORUM ★
- 176 Knots, Language, and Computation: More Bermuda than Love
David J. Lobina
Universitat Rovira i Virgili
Mark Brenchley
University of Exeter
- 205 On the Feasibility of Biolinguistics: Koster's Word-Based Challenge and Our 'Natural Computation' Alternative
Sergio Balari
Universitat Autònoma de Barcelona
Cedric Boeckx
Universitat de Barcelona
Guillermo Lorenzo
Universidad de Oviedo
- 222 A Turing Program for Linguistic Theory
Jeffrey Watumull
University of Cambridge
& *Massachusetts Institute of Technology*

