General Cognitive Principles: The Structure of Behaviour and The Sequential Imperative

William H. Edmondson

School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK

w.h.edmondson@bham.ac.uk

Submitted to

International Journal of Mind, Brain and Cognition

(submitted June 2009, revised July 2009)

Published Vol.**1**(1): 7-40, 2010.

General Cognitive Principles:
The Structure of Behaviour and The Sequential Imperative

Abstract

The paper sets out a multifaceted and multidisciplinary response to an apparently casual remark of Chomsky's to the effect that General Cognitive Principles could suffice in place of Universal Grammar. The response here takes the form of a set of General Cognitive Principles and an account of the Sequential Imperative. The discussion covers various structural issues arising from generalization of a non-linear formalism used in phonology, and a proposal for a research programme for cognitive science which focuses on the need to formulate a functional specification of the brain. The account is not limited to discussion of the human brain.

General Cognitive Principles:

The Structure of Behaviour and The Sequential Imperative

0 Introduction

This paper – a think piece, not a report of finished work – provides a multi-facetted response to Chomsky's comment in Rules and Representations (1980:29):

The study of biologically necessary properties of language is a part of natural science: its concern is to determine one aspect of human genetics, namely, the nature of the language faculty. Perhaps the effort is misguided. We might discover that there is no language faculty, but only some general modes of learning applied to language or anything else. If so, then universal grammar in my sense is vacuous, in that its questions will find no answers apart from general cognitive principles.

The response offers a set of plausible General Cognitive Principles (GCPs) and develops some ideas which illustrate the value of the GCPs as the basis for a programmatic approach to research in Cognitive Science. The paper is multidisciplinary in two senses: a) the author's work is multidisciplinary and thus the content ranges over several disciplines; and b) the implications for further work are multidisciplinary. The paper is in three parts, each subdivided into sections. The parts cover: A – discussion of the GCPs and the Sequential Imperative; B – detailed discussion of a formalism originating in phonology which, when generalised, reflects/models the material in A; C – discussion of the implications for research in Cognitive Science (using GCPs and the model from Part B).

Part A – GCPs and the Sequential Imperative

The difficulty with GCPs is not, as it might first seem, proposing or discovering them[1]. The problem is evaluating them – making GCPs do useful work not

[1] Space does not permit an account of the nearly three decades of work which has informed their formulation, but for early work on the emerging ideas see Edmondson (1986a,b; 1989a,b).

previously done, and in a compelling fashion. For work on GCPs to "catch fire" the proposals should be "blindingly plausible" – readers should see some value immediately, and further value after some work. The set of principles proposed below is presented *ex caeruleo*, and the first three then explained in detail. Following this is a discussion of the Sequential Imperative – a core insight derivable from the GCPs, but which in fact inspired early work in their formulation.

A.1 General Cognitive Principles

The set of principles is as follows:

GCP1 Sequentiality in behaviour is forced physiologically.
Corollary 1: Sequence penetrates the corporeal boundary.
Corollary 2: Sequence is semiotically free.

GCP2 Cognitive entities are:
i) Inherently atemporal
ii) Dual in nature

GCP3 Behaviour is sequencing; Perception is de-sequencing.

GCP4 Learning serves the sequential imperative.

GCP5 Attention is the management of the processes of sequencing and de-sequencing.

GCP6 Affect is an attentional mechanism.

GCP7 Cognition and affect can be distributed and projected:
i) in the environment -
space, objects, others (not just con-specifics)
ii) in time -
historical, personal

GCP8 Thought is the production of cognitive entities.

We will consider in detail the first three GCPs which taken together capture the notion of the Sequential Imperative. GCP4 returns us to our opening quotation.

A.2 Three General Cognitive Principles

The first three general principles – when taken together – provide the core insight needed to begin to address Chomsky's concern (op.cit.). This insight – that behaviour is constrained by the Sequential Imperative – will be explained below. More generally, the "big idea" the reader should have in mind is that the GCPs constitute a functional specification of the brain. The GCPs encourage us to think about what the brain is "for" in the sense that we know what the heart is "for", and the lungs, etc… We will return to this generality of perspective after reviewing the first three principles.

A.2.1 GCP1

GCP1 *Sequentiality in behaviour is forced physiologically.*
Corollary 1: Sequence penetrates the corporeal boundary.
Corollary 2: Sequence is semiotically free.

Before considering the two corollaries we look at behaviour generally. Any organism with articulators will deploy/activate/articulate them in one or another sequence – there is no choice. Clearly there will be instances of co-ordinated simultaneous activity – two hands reaching out to catch a ball, for example, or two paws holding a nut for ease of biting – but deployment of articulators will inevitably involve sequentiality. Some patterns of activity, such as the leg movements of a millipede, show some co-ordination that is both simultaneous and sequential. But inevitably, it can never make sense to try to contract all one's muscles at once, in an effort to do everything simultaneously.

So what determines the sequences? In some cases – millipede locomotion, for example – the sequencing is built-in in a rather unchallenging sense, perhaps at some "low level" in the articulator control system. In some cases sequences are provided by the specification of the activity – peeling a fruit before consuming it, for example. In yet other cases an activity may only specify sequentiality, not an actual sequence – an example might be visual search requiring active direction of gaze at a variety of targets which cannot all be gazed at simultaneously (the issue is not perception per se, but the orientation of the head/eyes). More complex sequencing still – a cat reaching up a paw to swat a passing bee, or a pianist at work – requires some internal specification of the sequencing of articulatory activity in coordination both with other articulatory activity and with sensory input. *Physiology*, and thus eventually behaviour, makes no sense without sequence. But we can get closer to understanding behaviour when we consider the two corollaries.

A.*2.1.2* *GCP1:Cor1*

This corollary states that *sequence penetrates the corporeal boundary.* What does this mean? Simply put, the corollary notes that a sequence in activity (or, as we will see, perception) is not disrupted/altered across the corporeal boundary: a sequence "within" is that same sequence externally (and vice versa). The control of muscles in the specification of performing a sequence of actions has the same sequentiality as the actual sequence – which is useful as it means that feedback can be used to check the activity has been successfully carried out. Importantly, the feedback may well involve sensory input, as when one watches one's hands perform an activity using vision to fine-tune manipulation. Without the preservation of sequence through the skin, so to speak, such feedback and control could not work (kinaesthetic feedback constitutes a form of sensory feedback which doesn't involve an external proximal stimulus – the stimulus is internally generated). Note that the preservation of sequence is the minimum temporal requirement – there is no additional requirement for preservation of actual timings.

A.*2.1.3* *GCP1:Cor2*

This corollary states that *sequence is semiotically free*. What this means is that because the sequencing of activity is a consequence of physiology it is inherently meaningless and thus semiotically available to carry meaning. Indeed, even when the activity has inherent sequence that sequence may be available for semiotic exploitation. Although arbitrariness in activity invites semiotic exploitation (say, sequences of speech sounds) it does not follow that inherent sequence (say, "wet, soap, wash, rinse, dry" as a washing routine) blocks semiotic exploitation. When this is taken with Cor1, discussed above, it becomes clear that monitoring of semiotic activity is required and also provided for in GCP1 – otherwise we could not know that our semiotic exploitation of activity is happening.

A.2.2 GCP2

> *GCP2* *Cognitive entities are:*
> *i) Inherently atemporal*
> *ii) Dual in nature*

Cognitive entities are the pre-verbal representations of thoughts, ideas, desires, beliefs, memories, intentions, and the like – but in forms not directly available to consciousness or sensation.

Cognitive entities are *inherently atemporal* in that desires, beliefs, etc., don't have a dynamic characteristic. Belief in God, desire for a boiled egg, memory of yesterday's weather, knowledge of the squares of single digit numbers, and etc., *endure* but have no temporal structure or dynamic – the desire for a boiled egg does not simmer for 4.5 mins approx. Cognitive entities are formed and also learnt – which looks like temporal activity. However, this temporal activity is not entrained by or reflective of real world temporal activity. Although there are cases where learning A must precede learning B this does not undermine the general case.

Furthermore – well established research over several decades shows that there are so many different timescales for activities within the brain that temporal structuring is surely improbable. Binaural signal differences for sound source direction identification are measured in 10s of microseconds. Evidence for temporal patterning of sensory responses to stimulation comes from studies of hearing (Geldard 1953): the patterns of discharge in the auditory nerve (of a cat) show that firing can be entrained by the stimulus up to around 850Hz. Neurons operate more slowly still: Rumelhart and McClelland (1986) cite a figure for the speed of operation of neurons as being measured in "milliseconds – perhaps 10s of milliseconds". Then there are so-called "brain-waves" – with frequencies of 1,3,5,7,10 & 13 Hz (Lenneberg 1967). The notion that one's lifetime of experiences – let alone knowledge, plans, etc. – are perpetually playing out in some vast assembly of endless loops of brain activity, ever ready to be "read out" like some old fashioned computer memory system – simply doesn't bear close scrutiny. Although they may change (learning) cognitive entities are not structured in time, they endure; they may be *about* time and sequence, but they are essentially timeless.

Cognitive entities cannot do much, if merely enduring and timeless, so need some other property to make them accessible and deployable. That property is *duality*. In the general case cognitive entities have as (sub-) components other cognitive entities, and also are components of other cognitive entities – "all the way down" (and up). This "Janus-like" componentiality – duality – is especially obvious in linguistic structures, but does not need to be construed as yielding a hierarchical structure. It may appear, for example, that there is obvious hierarchy in the relationship: book, chapter, section, paragraph, sentence, word, character. However, in speech the sounds, syllables and stress patterns do not necessarily map onto morphological structures in any obviously hierarchical manner – the English word *originality* provides a good example where morphology and syllabicity don't map one to one.

The important point about GCP2 is that it complements the sequentiality discussed in GCP1. We see this explicitly in the discussion of GCP3.

A.2.3 GCP3

GCP3   Behaviour is sequencing; Perception is de-sequencing.

GCP3 couples the first two principles. In order to execute any behaviour an organism needs to produce sequences of muscle actions from sequential specifications internally. The sequential specifications are derived from atemporal cognitive entities via successive decomposition of such entities into sub-components – "all the way down", or perhaps better – "all the way out". The entities are made sequencible by virtue of their sub-components (and sub-sub…. components) even though those sub-components may not be sequenced until "just inside the skin" or "close to the muscles". One can think of this as being the conversion of *internal distal* entities to *internal proximal* entities as sequential specifications which produce behaviour – this is behaviour as sequencing.

Conversely, perception is desequencing – beginning with the transfer "through the skin" of sequential *proximal external* stimuli to *proximal internal* stimuli in the same sequence, followed by abstraction from the sequence as the cognitive entities (*internal distal*) are recovered. Note that if the *external distal* stimuli are not sequential (compare the visual world with the acoustic world) then the sensory apparatus operates to impose sequentiality so that the transfer through the corporeal boundary can happen. Saccadic movements of the eyes have this effect in humans.

Consider avian vision in two species. Wagtails (e.g. *Motacilla alba*) eat seeds and crumbs etc., which have no internal movement. The birds walk with a jerky motion of the head which serves to impose temporal structure on their visual input and thereby provides the basis for computation of depth cues which they need to see their food items. Robins (*Erithacus rubecula*), on the other hand, which eat mobile food (e.g. worms) do not need to move their heads to create time varying sensory input. For them it is required that if they alight on a flexible/bendy twig then they have to cope with the movement of the perch whilst keeping their head still – it is impressive when you see it.

Of the remaining GCPs, GCP4 will be discussed below but GCPs 5-8 will not be explored in detail. The set of GCPs constitutes a functional specification of the brain and as such sketch a research programme for Cognitive Science. One GCP – 7 – will be considered briefly in discussion of the universality of the GCPs and the Sequential Imperative (see Part C).

A.3 The Sequential Imperative

The first three principles taken together express the Sequential Imperative (SI). It is *necessary* for any organism with a brain, and thus with atemporal cognitive entities, to map those into sequential precursors of sequential behaviour (where the sequence maps across the corporeal boundary) if behaviour is to happen: sequencing is required by the physiology and enabled by the duality of entities. Likewise, and "in reverse", perception as desequencing *necessarily* requires "stripping out" the sequence required of stimuli to map through the skin, in order to end up with atemporal cognitive entities.

The functionality of the brain is to serve this imperative: it stores entities, and their componentialities; it handles the sequencing and desequencing; it enables the semiotic exploitation of behaviour and percepts; and, via GCP4 (of which there is little to be said, it is now mainly obvious) *learning* is all about extending the materials available in the brain for sequencing and desequencing, for semiosis and other behaviour – acquisition of components and componentialities. Language, therefore, is just like everything else (species differences notwithstanding) and Chomsky's concern is, in principle, addressed (learnability is discussed below). Languages have to *conform* to the requirements of the system which does the sequencing and desequencing – the body-brain totality of the organism. The SI is not so detailed as to constrain language behaviour to a unique solution – conformity can take many forms (several thousands of languages are extant, with significant typological variations, exploiting two modalities: speech and signing). Species differences in brains, inevitably coupled to sensory and physiological differences, surface in non-trivial differences in details – that is what speciation is about – but the underpinning architecture is invariant. And, it is suggested here that the plausibility of this claim extends to claims about brains in general – anywhere. That is what they are *for* and any embodied organism, anywhere, must have a brain with such functionality.

Another aspect of the SI is that we can recognize that the mind/body "problem" is actually little more than a description of the fact of the dimensional shift into and out of the time domain. The mind will, in such a scheme, amount to the atemporal entities of intentions, goals, etc., and their re-arrangement, extension, recombination, etc. which – when done without recourse to sequencing is sub-conscious, but is conscious when done via some sort of behaviour (or precursor thereto – speech in the head), or some sort of sensational representation (sequential percepts generated internally and not via desequencing through the skin). Consciousness is more than this – as discussed later – but here it is enough to point to the temporal/atemporal dimensional disjunction/shift as expression of the mind/body "problem".

The Russian composer Gubaidulina has referred to her composing work as transforming vertical time/sound which is the musical thought into horizontal time/sound which is the musical product (the phrases/words used vary but the curious reader can find much on the web[2]). This is perhaps just one more way of talking about the dimensional shift from atemporal structures which can be about time, to sequenced structures which are temporally arranged.

Note that learning (GCP4) has consequences for cognitive entities which amount to changing them over time. But this time is different, it is simply process time, not expressed time or content time. It may take me 6 months or only 6 seconds to learn that potatoes respond well to wood-ash on the soil; depending on whether I try it for myself for a season's growing, or learn the fact from a book. The learning may well be different in some details (like the colour of the foliage or the luxuriance of the growth, or whatever) depending on the book as much as the growing season, and of course experiential knowledge is sometimes not available any other way (you cannot be taught how to ride a bicycle using a book). But this aspect of changing cognitive entities is not the time domain involved in sequencing and desequencing.

[2] See for example: http://www.csmonitor.com/durable/1997/08/27/feat/music.1.html

Learning language apparently poses special problems – the behavioural products or utterances are inherently complex and their contextualization makes both perception and production a rich mix of structures and structurings which have to be sequenced and desequenced successfully for communication to work. Extraction/recovery and learning of components and componentialities, whilst using them, is not straightforward – and one theoretical approach has given rise to the notion of "learnability" which encapsulates the difficulty facing the language learner who is exposed only to the products of a grammar (and grammars are underdetermined by their output – see Bowerman (1988), Truscott and Wexler (1989); for more recent work on learnability see Tesar and Smolensky (1998)). Whilst learning language may or may not be impressive (cf. Putnam (1971)), within a GCP framework learning any sort of behaviour is as difficult/easy as learning any other. The GCP framework addresses universality of process, not ease.

A.4 The Sequential Imperative and Language

In this section we address more explicitly Chomsky's concern (op.cit.) and consider how it is that the Sequential Imperative can be universal, in the sense that the SI is served by the functional specification for a brain – any brain; and can also be variably served in the sense that brains (species) differ, and the behaviours "accessible" to different organisms differ. Maybe only one species can manage the full complexities of human language, but others might come close, and in any case they have their own species specificities. The point is that the SI is an imperative, it is not a specification for any one sequential behaviour. Many different brains, potentiating many behaviours, must all comply with the SI.

Language structure reveals considerable complexity of data. Consider a sentence: *Stan's dog barked at him.* Let us ignore for this exercise the structural detail implied by the possessive marker on *Stan*, and note that in the diagram below it is the sounds which surface as the sequentially arranged elements. The sentence can simply be diagrammed in a conventional way (Fig. 1 below), but note that although the size of the elements in the temporal stream is shown as uniform this is merely a graphical convenience – no assumption is made of uniformity of segment duration in the speech stream.

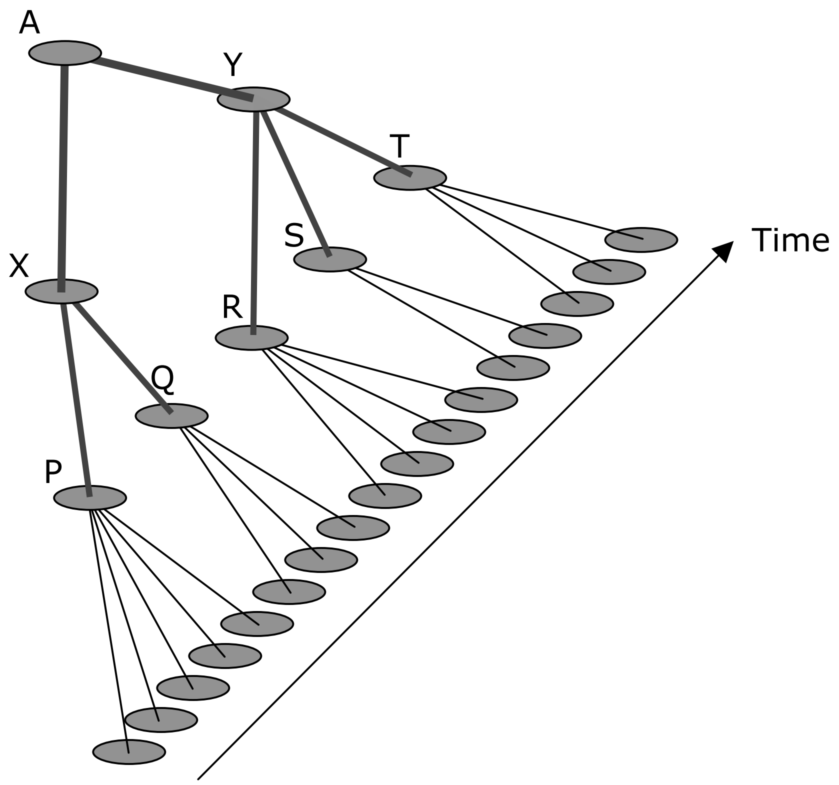

Fig. 1

The node/root A is typically labelled S for sentence, X and Y become Noun Phrase (NP) and Verb Phrase (VP) respectively, the five words map to P -> T, and the sounds of the words map onto the little blobs on the timeline. The sounds of the words are sequenced, that does not seem in doubt (duration of each sound is a separate issue, addressed later). What about the words? The diagram shows them in sequence, as one might expect from the sounds. Likewise the NP and VP are in sequence. Much is made in these "tree diagrams" of the sequential arrangement of the nodes, and lines showing connections between components on two levels may not cross – so Y could not connect to Q, with X to R, for example.

Should the hierarchical approach to structures be used without limit? When a person writes a book is the whole tree structure in place, cognitively prepared, sequentially, to be "output" as a book? How many nodes have to be organized into the right sequential slot at the right level, before outputting the text? If such sequentially complete preparedness is not required – as seems likely – the question which then arises is at what point in a complex tree with many intermediate levels between root and output sounds (if we talk about talking) is the presumption of sequence abandoned, and what is put in its place? Note that if the sequential thoroughness is presumed then accounting for errors in production becomes problematic – they shouldn't be possible. Of course, one solution is to limit the tree-structure representations to sentences, and to ignore issues of sequentiality in precursors to the sentences themselves. However, that won't work because sentences can be started which in principle have no termination:

Stan's dog barked at Fred, who stood by the door, with a basket of vegetables, grown on his plot, down by the stream, with the help of his neighbour, who supplied the seeds, which came in an envelope, slipped under the door, onto the mat, which woke the cat, who chewed on the package, …..

Or

Anne put her cup down on the table, under the window, near the door, from the kitchen, into the hall, on the ground floor, of the house, by the stream, …..

The question can be put more generally – does sequencibility *require* sequentiality in all cognitive entities at all scales prior to output as behaviours? It was claimed above that duality is all that matters, in accord with GCP2, so that the cognitive sub-components of cognitive entities, themselves cognitive entities – "all the way down"

– provide the structural resource for sequencibility, without actually demonstrating sequentiality. In this model, therefore, all that is actually sequenced is the output – the observed sequence – and the internal sequential instruction to produce the output (the same sequence, both sides of the corporeal boundary, so to speak). The cognitive precursors of the sequentially arranged instructions to muscles need not themselves be sequentially arranged – at whatever level of analysis/detail seems appropriate for labelling as the output instruction layer. The imprecision here doesn't change the overall argument – at some scale/level of analysis it will be possible to point to internal representations as unsequenced, with a finer scale of analysis showing entities which are sequenced. This argument applies as readily to cooking a meal, or designing a building, as it does to using language.

Consider Fig. 2 below which shows how it is possible to illustrate the same structural relations as in Fig. 1, but with the assumption that the time dimension (sequentiality) is reflected only in the output sequence of sounds.

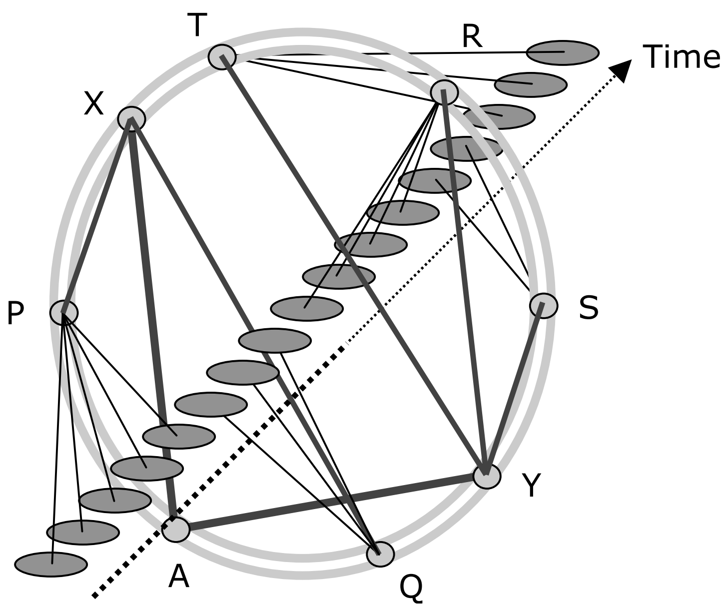

Fig. 2

Here we see that the nodes A, X, Y, P, Q, R, S, T are all arranged in a circular fashion around the sequence of output sounds – these nodes are simultaneously co-present precursors or specifications of the output. The diagram resembles a Naum Gabo sculpture with links (strings) between structures in different planes. The different planes are defined by the central "spine" of sequenced output segments running along the time axis and the sets of links to the various nodes P, Q, R, S, T so that the links are in different planes radiating out from the central spine in different directions. An alternative "visual aid" for these structures is to think of a notebook with a "spiral" binding – where each ring in the spine represents a speech segment, with pages, or planes (half-planes, strictly) radiating away from the spine. Specifications written on the pages (P, Q, R, S, T) label the linked segments in sequence but are themselves unsequenced (the pages fanned out around the binding are not sequenced, but contain specifications about the segment sequences in the output).

With this style of diagramming in mind, the reader is invited to consider the two representations – conventional and "Gabo-like" – of *discontinuous* constituents. In the phrases "… looked up the word …", and "… looked the word up …", the Verb+Particle constituent is the Verb, in a real sense, but in the second phrase the two components surface discontinuously. Conventional accounts have the Noun Phrase constituent in Object position following the Verb – which works out fine as a tree-structure representation for the first version of the phrase, but looks impossible in the second, where the links between the NP node and the words cross the link from the V to the second component "up". By contrast, in the ring-bound notebook model the specification on the page for V links to the two components wherever they are in the surface string, irrespective of the links on the page for NP to its related elements in the string. There is no crossing because the links are in different planes.

As before – the argument applies as readily to non-language behaviour as to linguistic behaviour. In one sense, therefore, Chomsky's concern is indirectly met – behaviours (including language) are served by general cognitive principles.

The discussion here has reached the conceptual core of this paper. The first three GCPs, taken together, capture the essence of the Sequential Imperative and of the functional specification of the brain (the remaining GCPs are functional corollaries and we will return to these in Part C). The illustrations above – Figs. 1 & 2 – reveal that problems expressing the relation between surface forms and cognitive precursors exist in simple formal accounts of linguistic structures and that it is possible to envisage a formalism which avoids these problems. At this point a "health warning" is appropriate – discussions of representations offered below, and of their possible computational expression, are of functionality and functional architecture. There is no claim that such representations or models are the key to understanding neurological details (although later we can use the model from Part B to discuss some details in an evolutionary context in Part C), nor is it claimed that the representation/model elaborated below is somehow the "right" solution, or the only solution to the problem of capturing and expressing the functionality of brains. The reader might form the view that a formalism based on unification would be more intellectually satisfying and more readily implemented than the rather graphical account offered. The graphical account/model is offered because it reveals rather nicely the development of a conceptual approach stretching over several decades of work in phonology.

<u>Part B – 3D models</u>

We turn now, in Part B, to consider in detail a debate about a formalism in phonology which has implications for the validity of conceptual generalization of the three dimensional structure sketched above. The general appeal of 3D diagrams of linguistic structures will be illustrated first as a way of showing why this approach is worth pursuing. The reader needs to be prepared for a sense of expository disjunction between Parts A and B, and then between B and C.

B.5 3D Models in Phonology and Syntax

In the 1980s linguistic work on signed languages included consideration of the possible role of segments in analogy with speech segments. I was active in the work (Edmondson (1986a, 1990a,b)), using the formalism of autosegmental phonology (also known as pluri-linear, or non-linear, phonology) – as were others, notably Sandler (1990, 1993). The question of segmentation in signed languages remains open, in my view, despite more recent work perpetuating analogies between activities in signing, and vowels and consonants in speech (cf. Brentari (2002) and elsewhere in Meier et al (2002)). One underlying problem with the autosegmental formalism (from which Fig. 2 above is derived) is the assumption that the units in the timeline or spine are linguistically significant units of time – proto-segments (often specified as Consonants and Vowels) which are characterised by the featural specifications provided on the pages of the notebook. This looks much more plausible as a notation for speech phenomena than it does for sign phenomena. Segmentation in signing is considered to be shown in the articulatory activity of each of the two manual articulators, leading problematically to the requirement for two formal accounts (two spines and sets of pages – see 5.3.1) (see Edmondson (1986a)).

The formalism expects units, Cs and Vs, so these have to be found in signing if the formalism is to be deployed in sign language linguistics. In one sense therefore the formalism is driving the language analysis (see 5.2.2 below). This situation is not substantially altered if the basic timing unit is considered to be the syllable (another debated phenomenon in the case of signed languages). The real issue, of course, is whether the formalism should be modified to remove any requirement for timing units in the spine. All that really matters is identifiable sequentiality – specified activities in sequence form the content to be notated.

B.5.1 3D Syntax

The foregoing is offered as a caution to the reader in relation to what follows. The use or not of timing units in the spine is properly considered as a theoretical question about properties of the formalism, not a concern for the presence or absence of phenomena in languages or modalities. We need to consider next a different technical issue with the formalism. Later, when both are discussed together, the generic value of the formal approach will become apparent. The technical detail to be considered now comes in two variants, and is quite intricate.

Consider a simple "respectively" sentence, such as: *John and Jill are, respectively, 19 and 20 years old.* The *respectively* can be placed at the beginning, or the end, if not in the middle of the sentence as shown (or even after *Jill*). It can be omitted, colloquially, to give: *John and Jill are 19 and 20*. This sentence is in fact two propositions interleaved, to be uttered "simultaneously" as shown in Fig. 3 below.

The complex structure illustrated in Fig. 3 is reminiscent of diagrams found in Van Valin's work on Role and Reference Grammar. For example, the diagram below (Fig. 4) – taken from a recent document on the RRG website – is very similar to one in an earlier publication (Van Valin (1993)). Both Fig. 3 and Fig. 4 show 3D structures with half-planes, or pages, radiating out from a central spine where surface elements are arranged in sequence. The technical question here (variant 1) is why, given the contrast between Figs. 1 and 2, is it necessary to build complex structures in just a few half-planes when simple structures on more half-planes will have the same content?

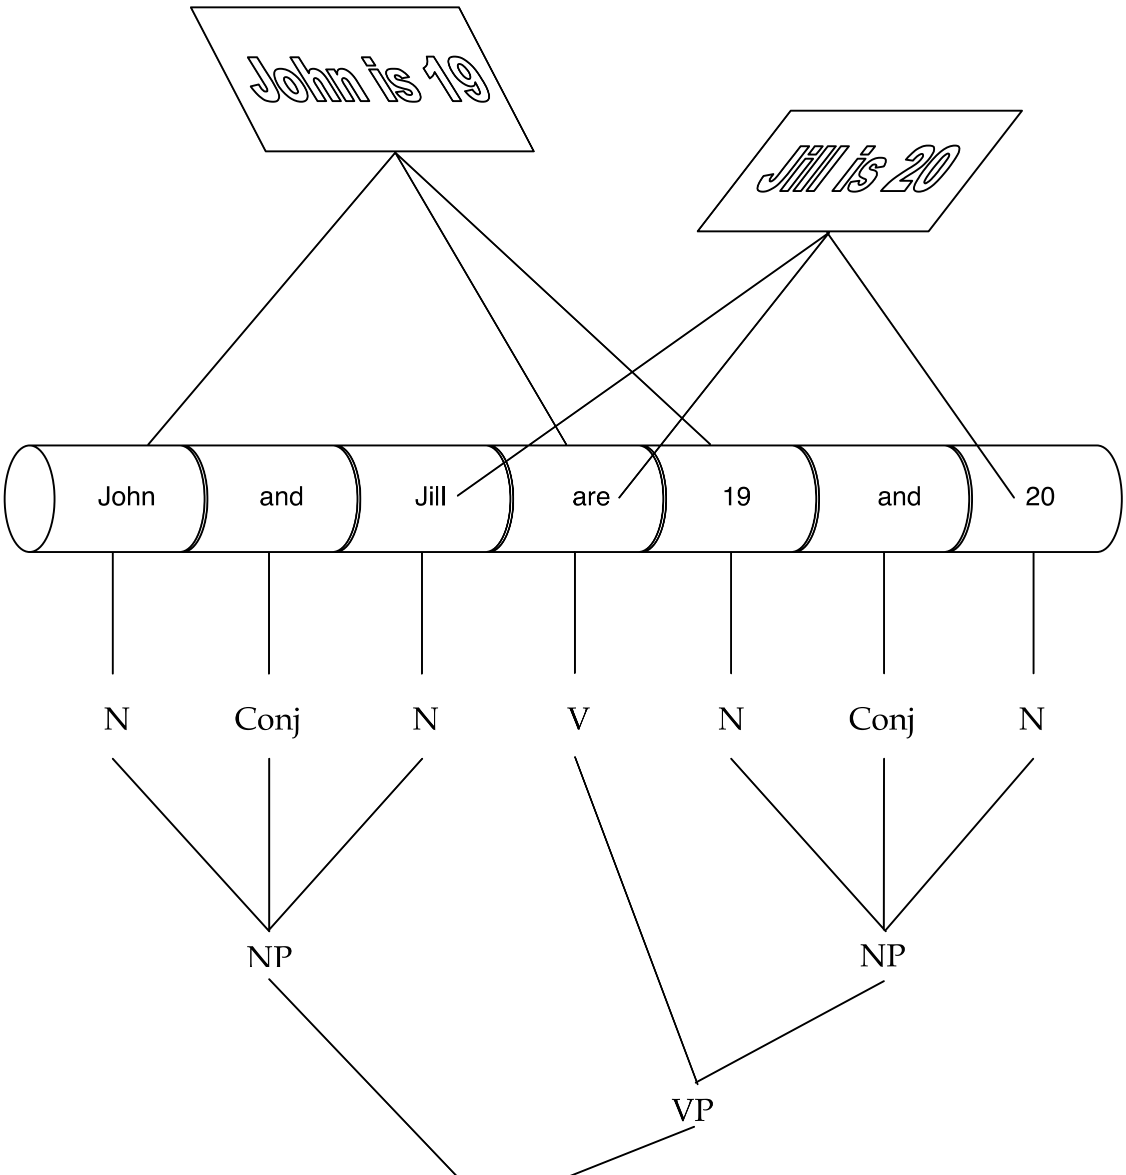

Fig. 3
This shows the ‘respectively’ sentence *John and Jill are 19 and 20* diagrammed in 3D. The syntactic structure is shown in the plane vertically below the segments in the spine labelled by the words of the sentence. Note that the plural form of the verb is required because the subject NP is plural in consequence of the conjunction. Semantically there is no plural entity, nor conjunctions of people or numbers. Semantically the two propositions are unpacked and shown by text in boxes on two different planes radiating away from the spine: one above and back, and the other above and forward.

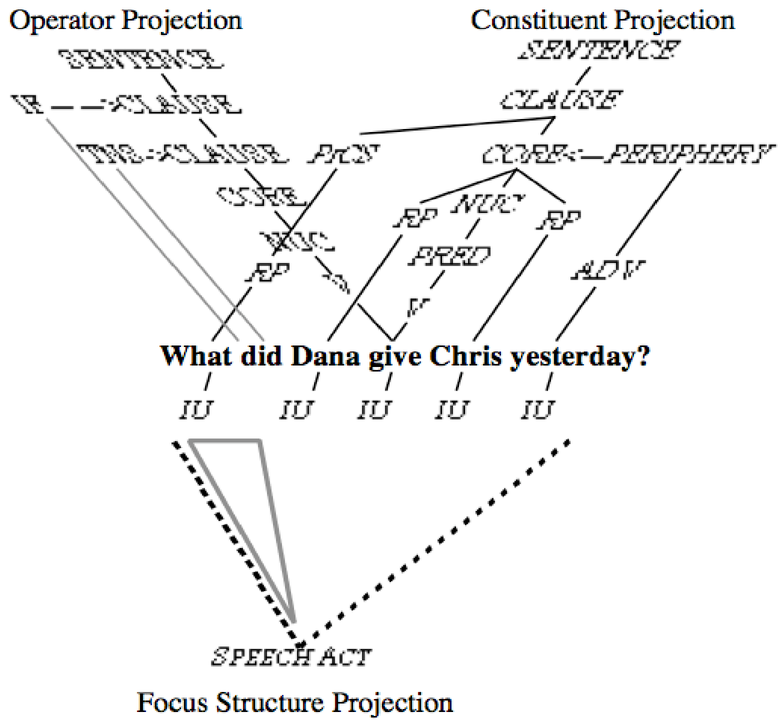

Fig. 4

See http://linguistics.buffalo.edu/people/faculty/vanvalin/rrg/RRG_overview.pdf retrieved from http://linguistics.buffalo.edu/people/faculty/vanvalin/rrg.html

From a formal point of view, more half-planes, with links between them, will always be able to reflect the details found in fewer half-planes each containing complex hierarchical structures. To be sure, the formalism requires the details of co-referencing half-planes to be specified, but in principle more and simpler half-planes will do the work, and ultimately are to be preferred because the structural details have to be spelt out, rather than being embodied in the diagrams.

A related, but subtly different point (variant 2) is made in complex and convincing detail in relation to articulatory activity, by Hayes (1990) in discussion of work by Clements (1985). In order to appreciate the details a brief discussion of basic issues in phonology is helpful. This is followed by accounts of two approaches to 3D phonology, with Hayes's simpler approach emerging as the more compelling.

B.5.2 3D Phonology

The basic temporal unit in speech is conventionally known as the segment; in articulatory and acoustic accounts this is the phone, in linguistic accounts the underlying form, or precursor, is the phoneme. In both cases, however, the assumption is that the unit is indivisible – it is the speech "atom". The segment can be thought of in several different ways, and this leads to differing theoretical perspectives.

One view is that segments are inherently one of two types – vocalic or consonantal. Vowels and consonants pattern syllabically and syllables are a necessary and independent basic building block in spoken languages. Vowels and consonants also cluster into meaningful units, but these are not syllabically constrained. Thus the basic structural units are identifiably vowel and consonant, with the further particulars or details of each segment requiring a separate account. The second, and only slightly different, approach recognizes only the notion of underlying phonetic segment – a unit of time in the speech stream, but otherwise characterless.

The third perspective is more radical, and somewhat younger than the millenia old accounts of segments familiar to phoneticians. It can be argued (Abercrombie 1965, Firth 1957, Palmer 1970, Edmondson 1990a) that the segment is a fiction, and that the behaviour of interest is simultaneous patterning of articulator activity. In this prosodic account the segment has no privileged status as a primary unit; it is useful, but is not a necessary component of every account of the production, perception or processing of speech at the phonetic level. The radical nature of the perspective lies not (just) in the denial of primacy of segmentation, but rather in the notion of many articulatory activities simultaneously organized in time so as to create *en passant* the effect of segmentation.

B.*5.2.1* *Non-linear phonology – I*

Clements (1985) introduces his discussion of the essential principles of non-linear phonology with an account of phonological features and segmentation, an account which can serve here (as for Clements) to answer our immediate question. For Clements (op cit) "one of the fundamental discoveries of modern linguistics" is that

> phonological segments, or phonemes, are not the ultimate constituents of phonological analysis but factor into smaller, simultaneous properties or features.

But, as he points out, the organization of these "ultimate constituents" is less clear. Citing Bloomfield's "well-known characterization of phonemes as 'bundles of features'" Clements argues that there are two senses in which the "inherent disorganization and lack of structure" are problematic. The first problem with the feature bundle approach arises from the fact that the individual segments (phonemes) each come with a full bundle of features which supplies their specification in isolation from the adjoining segments. Feature values are assigned to segments without regard to the fact that feature values are descriptive of the configuration of the articulatory apparatus producing the speech, and that this articulatory activity is necessarily supra-segmental. The force of the complaint can be conveyed most graphically by pointing out (as Clements does not) that the isolationist account would apply equally well to behaviour in which the configuration of the articulators always returned to some fixed inter-segmental neutral state between any pair of fully specified segments. That this does not happen (and could not without turning speech into a very different activity) reveals that segments are not readily isolatable.

The second problem, for Clements, is that irrespective of the foregoing supra-segmental issues, the features in each bundle have some internal or inter-*featural* structure. Clements (op cit) refers to

> a fundamental observation regarding the structure of the human speech-producing apparatus. The essential characteristic of speech production is that it is COMPONENTIAL in nature, involving the coordination of simultaneous and partly overlapping gestures (cf. Halle and Clements 1983). These gestures show varying degrees of mutual independence. For example one can maintain a certain oral tract configuration constant, say the one appropriate for producing the vowel [a], while varying the type of laryngeal configuration, or the position of the velum. Or one can hold the laryngeal configuration constant while varying the internal geometry of the oral tract. Following this line of thought, we can identify at least the following articulatory parameters, each of which shows a high degree of independence from the others:
>
> (4) a) laryngeal configuration
> b) degree of nasal cavity stricture (open/closed)
> c) degree and type of oral cavity stricture
> d) a pairing of an active and a passive articulator
>
> Within each of these categories, on the other hand, it is difficult, and sometimes impossible, to maintain one gesture while varying another freely.

Clements goes on to discuss the subtleties of the mutual independence and interdependence of features with the scheme just outlined. He uses a hierarchical structure which makes explicit the notion of inter-featural structural organization irrespective of segmentation. Clements develops his account of the inter-featural structural organization within a non-linear model. He does so (op.cit.) by contrasting two different 3-D non-linear models.

> According to the first of these, phonological representations involve multi-tiered structures in which all the features are assigned to their own tiers, and are linked to a common core or 'skeleton'. Such a view can be schematised as in (2):

(2)

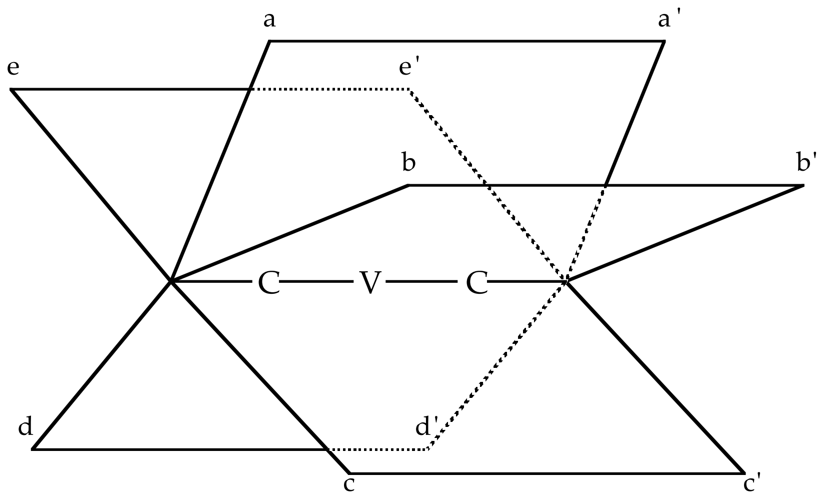

aa'= sonorant tier, bb'= continuant tier, cc'= high tier,
dd'= back tier, ee'= voiced tier

> In such a conception, a phonological representation resembles an open book, suspended horizontally from its ends and spread open so that its pages flop freely around its spine. The outer edge of each page defines a *tier*; the page itself defines a *plane*, and the spine corresponds to the *skeleton*.

Clements goes on to explain how values (binary) on each tier are associated with the "points on the spine" (the Cs and Vs). (Clements has a running example in this portion of his paper, the CVC word "pin".) He then argues that, just as with the feature bundle account, the problem with this 3-D model is the lack of clear indication of any inter-featural structure. He then introduces an alternative.

> In this conception, individual features are organised under hierarchically superordinate nodes, which I will term CLASS NODES. The class nodes themselves are dominated by a yet higher-level class node, which (following Mohanan) I will term the ROOT NODE. The root node, in turn, is directly linked to the CV tier. Under this conception, the phonetic content of a segment is arrayed on two different types of tiers, the feature tiers and the class tiers (including the root tier). As a preliminary proposal, suppose we take the view that the class tiers are exactly the following: the root tier, the laryngeal tier, the supralaryngeal tier, the 'place' tier, and the 'manner' tier (a further tier, the tonal tier, will not figure in the present discussion). These are organised as in (3):

(3)

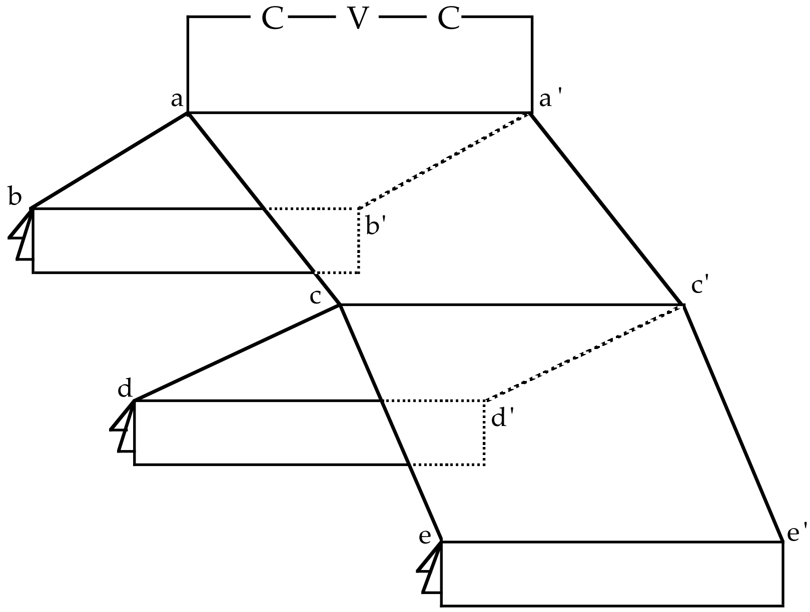

aa'= root tier, bb'=laryngeal tier, cc'= supralaryngeal tier,
dd'= manner tier, ee'= place tier

This conception resembles a construction of cut and glued paper, such that each fold is a class tier (labelled aa', etc.), the lower edges are feature tiers, and the upper edge is the CV tier. Imagine now that each element of the CV tier is individually linked by an association line to a corresponding node on the root tier, and that the first such node is linked to all the features of [p], the second to all the features of [i], and the third to all the features of [n], all placed on appropriate tiers. We then have a representation of the word [pin].

It appears that Clements is intent on developing a model of feature organization which permits him to address *both* of the problems identified earlier. This is supposed to work through recognition of the inter-featural structures of the root tier, the class tiers, and the feature tiers as the basis for inter-segmental interactions (e.g. assimilation phenomena). However, whilst the structural insight might clarify some descriptions of assimilation it remains the case that the underlying account is a segmental one (with proto-segments identified as Cs and Vs in the CV tier), and thus the supra-segmental relationship between the segments and the "coordination of simultaneous and partly overlapping gestures" needs another account of its own. In short, a solution to one of the two problems does not succeed in providing Clements with the solution to the other.

B.*5.2.2* *Non-linear Phonology – II*

The point is picked up by Hayes (1990) who has argued that the complexities of Clements' preferred model render it incapable of handling diphthongization. Hayes notes that the problem with Clements' tree structure is that it enshrines within the formalism a hierarchical organization which is inflexible (note, in passing, that this hierarchy relates to page structures – "cut and glued paper" – it is not the issue of hierarchical structures on individual pages, discussed above in 5.1). The problem which Hayes addresses he calls the "diphthongisation paradox". He illustrates this with data also used by Clements (op cit). In Icelandic, preaspiration

> converts the geminate voiceless aspirated stops /pp tt kk/ to phonetic [hp ht hk]. Articulatorily, this is a simple process, in that the aspirated voiceless stops and /h/ are both [+ spread glottis]. Because of this, the preaspirated stops can be derived simply by removing the supraglottal articulation from the first half of the geminate.

The difficulty, Hayes points out, is that there is no simple way to represent this process in the 3-D model favoured by Clements.

> In particular, consider the representation for geminate /pp/, in (5a). Since long segments are doubly linked, there is only one /p/ autosegment present on the Supralaryngeal tier. Deleting this autosegment, we obtain *[hh], as in (5b) and not the correct form [hp]:

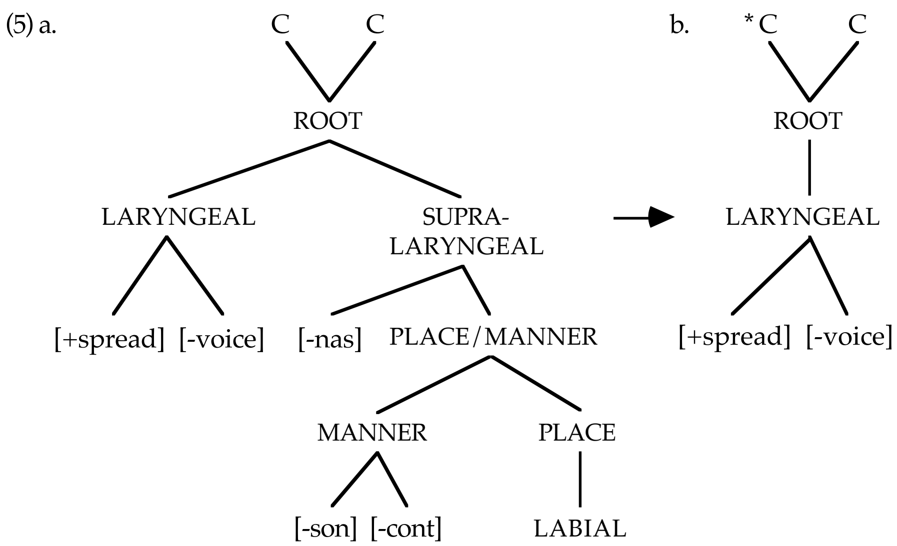

More generally, the representation provides no way to delink the supralaryngeal specifications of /pp/ only from the first half of the geminate: the only node that is doubly linked is the Root node; and if we remove its first link, we again obtain the wrong result, namely an empty C position followed by a single [p]:

(6)

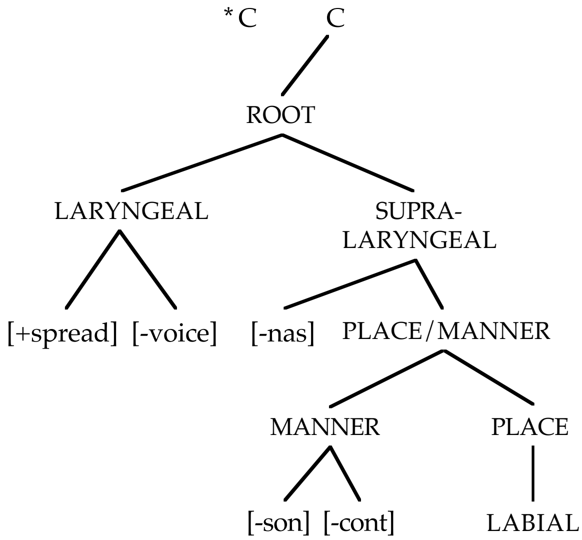

The nature of the difficulty should be clear: Icelandic Preaspiration is a diphthongisation of the supralaryngeal features, but not of the laryngeal features. If both supralaryngeal and laryngeal features are linked to the prosodic tier through a single Root node, then Preaspiration cannot be expressed as a unitary process.

Hayes goes on to clarify the nature of the problem, in general terms.

I believe the diphthongisation paradox is rooted in an ambiguity of formalism. Phonologists characteristically encode their representations graphically, as lines on a page. In principle, this is a good idea, since it usually makes the representations easier to visualise and manipulate. But it is important to have in mind a clear notion of what graphic formalisms are intended to stand for.

A line in a phonological representation can normally mean one of two things. In autosegmental theory a line indicates, roughly, simultaneity (see Sagey 1988 and Hammond 1988 for more refined interpretations). Thus (13a) indicates that H tone is pronounced together with the vowel /a/. Lines on the page have an additional meaning, however, in that they can indicate that the segments /tap/ belong to a single syllable:

(13)

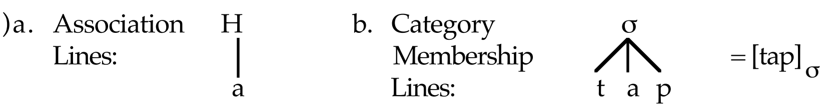

What are the lines in a feature tree? Here, the two functions are conflated: the lines depict a hierarchy of category membership, but at the same time they serve to link all nodes, directly or indirectly, to the prosodic tier. It is this ambiguity, I believe, that leads to the diphthongisation paradox. Once we separate the two functions, the paradox disappears.

Hayes separates the functions by a system of feature outlines. These can be coindexed for temporal information, giving a readable formalism in which feature structure and time are separately managed. The Icelandic preaspiration is then notated:

(34) *Icelandic Preaspiration*

	C_i C_j	CV tier
Delete i:	$[\text{-cont}]_{ij}$	Place/Manner tier
	$\begin{bmatrix}+\text{spread}\\ -\text{voice}\end{bmatrix}_{ij}$	Laryngeal tier

Hayes' observation that the formalism is ambiguous is a useful starting point for assessment of its general value. In essence the nature of his complaint (against Clements' proposal) is that category membership information, or information about the "logical grouping of features", should *not* be incorporated in the structures depicted in the representations. The graphical formalism is about characterizing temporal elements or units, and this can be done with the spine and radiating pages – for both suprasegmental details and for inter-featural details. The simplified formalism (cf. Clements (2) above) can deal with both.

B.*5.2.3* *Reformalisation*

The "reformalisation" which Hayes (op.cit.) claims is necessary "in any event" (i.e. aside from his argument against Clements' version) can only be guessed at, but nonetheless substantial reformalisation is not difficult to envisage. We need only return to earlier comments by Clements, in the light of doubts voiced by Abercrombie, to see the potential of the non-linear formalism. Having noted Hayes' concerns we can, in effect, extend his argument. The force of the argument is that the half-planes provide for individual characteristics to be associated with the spine, and that therefore the spine should be inherently characterless. The *duration* of any segment in the spine is as much a consequence of this process of characterization as any other property. The spine should not, then, consist of anything except event markers delimiting temporal segments (of whatever duration) which are characterized by half-planes or tiers as before. Segments emerge *en passant* in the flow of articulatory activity.

The reformalisation combines the insights of modern non-linear phonology with those found in the work of Abercrombie and other British phonologists working in the prosodic tradition (see Anderson 1985, Firth 1957, Palmer 1970)[3]. The technical details remain to be refined, but suffice to say here that the coindexation notion introduced by Hayes (above) provides a sketch of the way in which segments "emerge". A phonetic segment is that portion of duration in the spine which more or less simultaneously (which permits considerable phonetic detail to be captured) is characterized by many tiers or half-planes. The conclusion to be drawn from the above summary of what is in reality a long and complex debate is that the more radical approach no longer looks questionable. The formalism of non-linear phonology, it is suggested, does not need speech to be "built in" – as feature geometry in the half-planes or as (temporal) units in the spine. This is just as well, because a formalism constrained so specifically to vocal articulation cannot play a rôle in accounting for other linguistic data, linguistic data for other modalities, or behaviour more generally.

[3] The work cited in this discussion is not new – and this is intentional. Firstly, debates about formalisms and their influence on theory are part of the discourse in linguistics; secondly, work in phonology has moved away from the insights of Goldsmith (1976, 1990) and the related formal device explored here – but that does not render the device worthless; thirdly, just as Chomsky's concern was expressed in the 1980s so too was work being done around that time which could have taken the debate in a different direction – so one aspect of this presentation is "we could have seen things differently" because the ideas were around but failed to be synthesised into a new paradigm; fourthly, as noted in the text above, and in Goldsmith (1990), much earlier historical reference is reasonable from the autosegmental perspective and this sense of continuity is helpful validation of what might seem strange at first encounter.

B.5.3 Generalised 3D representation

From the “Gabo-like” diagram in Fig. 2, through the structures depicted in Figs. 3 and 4, to the complex diagrams offered by Clements – especially (2) quoted above – we find reference to a “spine” (as in the spine of a book) and pages or half-planes radiating out from the spine. Linguistic elements or segments are arranged in sequence along the spine, and the half-planes provide structural support for accounts of the elements. However, whilst this seems visually insightful, and appealing as a general scheme, the details are not self-evident.

When we look at the details it is actually helpful to broaden the general scope of the diagrams and of the account of them. Whilst this creates a bigger “big picture” it means that consideration of detail only needs to be done once, and comprehensively. Inevitably, of course, recapitulation of the scheme so far will involve restatement of some of the detail – and thus it makes sense to start there.

B.*5.3.1* *The spine and half-planes*

The spine[4] is reduced to a timeline with event markers placed to indicate “segment” boundaries where a segment is now no longer just a speech segment, but in fact any recognizable/identifiable activity with some measurable duration. The timeline has no inherent units or “ticks”, no inherent but “naked” segments such as vowels and consonants (without featural specification). The event boundaries can be coterminous so that, for example, the termination of a nasal consonant may also mark the end of a word, of a sentence, of a speech, and of a career: “… I resign.”.

[4] There is only one spine, irrespective of the number of articulators involved in behaviour. To propose two, say, for accounts of the two hands in signing, introduces formal problems of co-ordination/co-indexation of two sets of diagrams – problems which are readily solved, as in the cases discussed in 5.1 and 5.2, by reverting to a single spine with more thorough co-indexation of the half-planes. In addition, configurations of organisms, as in an ant colony for example, can also be represented in the formalism with only a single spine.

Each half-plane serves to define segmentation in the spine by supplying characterizations which "label" portions of the spine. However, just as segment boundaries can be coterminous[5] so it must be that portions of the spine can be "part of" many different segments. Segmentation is thus richly parallel and the specifications are not sequentially arranged. Note as well that structuring in the half-planes – whether like the diagrams in Figures 3 and 4, or whether like Clements's diagram in his Figure 3 – is simply not required. Equivalent detail can always be captured through the deployment of more half-planes and coindexation of the half-planes. The relationship between the spine and the half-planes is clear enough – but equally clear is the fact that 3D sketches don't constitute a formalism without additional elaboration.

B.*5.3.2* *Mechanisms – I*

The half-planes label or identify segments in the spine – but nothing has been said about how this might be done. Instead of thinking of a half-plane as a page in a notebook – blank until written on – it makes more sense to consider it as a mechanism or process which can do three things: it can read segmentation labels put "on" or "into" the spine by other half-planes; it can add labels to portions of the spine; it can communicate (bi-directionally) with back-ground processes which are "off-line" so to speak, and not temporally structured (more on these below). The labelling in the spine supplied by a half-plane process combines both a value and a type, and will be indexed to boundary events (this can be thought of as specifying them, in some cases, or simply reading them, in others).

[5] Precise time alignment of events constitutive of a segment boundary is not required.

B.*5.3.3* *Mechanisms – II*

Each half-plane process can be envisaged as a set of rules and conditions which responds as specified to values encountered in the spine and/or supplied from the back-ground processors. Without further structuring each half-plane or page would function independently of all others, and would communicate only with the back-ground processes. However, it should be noted that additional structuring can be provided very simply through the expedient of ensuring that some half-plane processes read or exploit data about others (via the type information). This permits pages to be clustered with one of the pages in a cluster serving as a sort of link or representative in other clusterings (cf. Clements' concern for articulatory feature hierarchy). All the pages or half-planes can be supplied with values from the back-ground processing, but only some will have any sort of super-ordinate role in other structures (these are in some sense *primus inter pares*).

B.*5.3.4* *Mechanisms – III*

Some of the half-planes (or clusters thereof) have special functionality. There are two such functions – crudely put these are "input" and "output". Somewhere in the system it is necessary to have some half-planes which serve as links "out" to articulators/muscles and links "in" from sensors. Importantly, these half-planes or pages have "back edges" (cf. Clements' term "outer edge" in the extract above) which are temporally coupled via the half-plane processing to the sequential detail on the spine, so that the muscle contractions, say, are "outputs" which are sequentially specified by reference to temporal detail in the spine (and, *mutatis mutandis*, temporally structured sensory input is mapped into the spine).

B.*5.3.5* *Back-Ground Processes*

Before considering the integration of these components into a functional account we need to consider some details of the back-ground processing (BGP) coupled to the "back edges" or tiers of the half-planes. The back-ground processes (also BGP) interpret values supplied from the half-planes, and also supply values to the half-planes. The processes work "off-line" and in their own time – temporal properties are not related to the time-line in the spine (see 2.2 above). Back-ground processing is the domain of cognitive entities: their formation, modification, retrieval, storage. Different components or processes can interact with one another "off-line" and thus the processing in BGP can be both "inter-process" and "intra-process", where the latter is concerned with interpretation and supply of values to half-planes and clusters thereof. BGP thereby contextualises the sequencing and desequencing effected via the spine and half-planes.

B.*5.3.6* *Mechanisms – IV*

With the exception discussed earlier, where "real-time" impinges on the spine, the segmentation in the spine, and the related half-planes, is scale independent. The account above is valid irrespective of whether a segment is 10 milliseconds or 10 seconds, or 10 hours or…. This is hugely significant. The BGP supplies values for segments regardless of duration, and interprets them similarly – but such segments are co-occurring, or simultaneous and over-lapping. This property provides for anticipatory labelling of spinal segments. For any given moment it will be the case that segmentation detail "into the future" is available and being supplied via segment labelling processes on the half-planes. In both sequencing and desequencing this delivers immense benefits.

Anticipatory behaviour in sequencing – whether “spreading” of articulatory features in speech (e.g. of lip-rounding to the consonant cluster in the second syllable of the word *construe*), or “spreading” in the sense of handling car-keys when “finishing one’s goodbyes” – is readily understood. Likewise in desequencing: some “top-down” information from BGP, triggered by desequencing thus far, will “fill-in” up-coming segmental detail in the spine providing anticipatory perception and a degree of error protection/correction, for example when listening to speech in a noisy environment. Errors in both production and perception can be explained as inappropriate segmental anticipation, and can be managed in perception by retrospective re-labelling of segments.

The flexibility afforded both by the BGP and by clustering of half-planes permits some non-determinacy so that incomplete and/or inconsistent labelling can be tolerated. This is particularly important for dealing with variability and delivering generality – close observation of the limbs of walkers reveals idiosyncracy in detail, likewise lip/mouth movements when speaking vary from speaker to speaker even when saying the same thing (just ask any deaf person reliant on lip-reading). In terms of the mechanism discussed in 5.3.4 above there needs to be tolerance of variability in many different circumstances (within and between speakers, for example). This is provided through both the “error correction” effect of anticipatory segmentation (which can sometimes go wrong, of course) and through the availability of additional half-planes which make available characterizations of segments in terms additional to core generalities such as what is being said (so that one knows something about the speaker, their age and gender, possibly their name, and so forth). Generalization and abstraction is, in this model, refinement of segmentation (see 6.2 below).

Fig. 5, below, provides a rough illustration of the whole scale-independent scheme, but note that for graphical convenience the segments in the spine are shown with a uniformity of duration which is not required or expected. To remind the reader of the autosegmental origins of this model the word "tier" is used (cf. Clements (2) above).

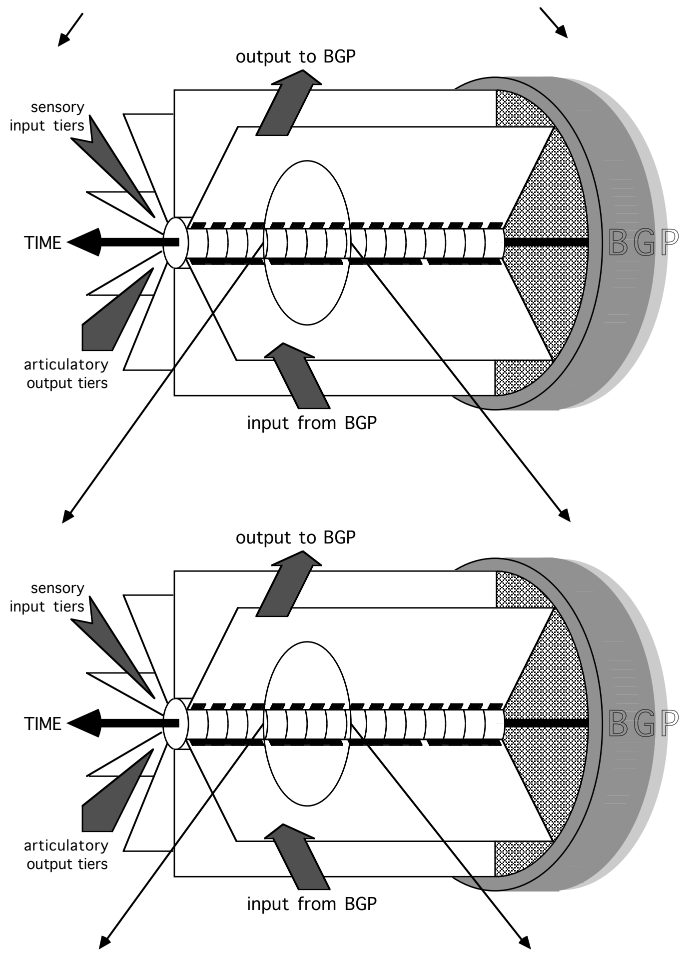

Fig. 5

B.5.4 Computational implications

The foregoing account slipped readily from discussion of a phonological formalism into discussion of a complex mechanism for assembling (and reading) sequentially organized structures in any behaviour. The mechanism is depicted as a scale-independent three dimensional configuration of central spine (segments) with radiating half-planes (containing some processes) communicating with background processes uncoupled in time from the spine. The totality manages GCP3 through reliance on GCP2 (see 2.2 above) and provides an account of GCP1 – in short the mechanism illustrates how the Sequential Imperative can be made to work, both in an abstract sense and computationally.

To be sure, there are details of co-indexation to be filled in (cf. Hayes above) but in principle the system could be instantiated computationally, and indeed a "primitive" system called *Pantome* was built and deployed for speech synthesis in the 1990s (see Edmondson and Iles (1994), Edmondson, Iles and Iskra (1996), Iles and Edmondson (1994), Iles (1995)). This is not the place to review earlier computational work; instead the focus is on the functionality of the system[6]. As such it is a model for the core functionality of the brain – any embodied brain – so we (re)turn now to consider the question – "what is the brain for?".

Part C – Cognitive Science research programme

Once again the reader needs to prepare for some expository disjunction. Part C, which follows, considers the functional specification of the brain from a number of perspectives. Some topics are discussed in terms of details of the model just described and some more generally, but in all cases the underlying concern is to show how consideration of the first three GCPs – the Sequential Imperative – illuminates topics and theories yielding a clearer understanding of "what the brain is for". Of course, the model itself may not be the "right" one – but the Sequential

[6] The system described is reminiscent of a black-board architecture in AI. It differs by not having a scheduler, in having direct communication between back-ground processes (BGP), by having "half-plane" processor coupling the spine and the BGP, and in having a "workspace" (the spine) which is time-aligned.

Imperative does have to be served somehow because it is universal, and it is motivated independently of the phonologically inspired model. The discussion does not present fully worked solutions but rather new commentary on old problems, couched in terms of a programmatic perspective on research in Cognitive Science.

C.6 Functional specification of the brain – I

The brain serves the Sequential Imperative and the first few GCPs show how this can be made to work, as noted above. The model introduced above in Part B – as a detailed interpretation of the ideas sketched in Part A – is not species specific – the configuration of articulators and sensors, along with the clustering of half-planes, is a species specific issue; the architecture is generic. In the next few sections we will use the model for descriptive illustration of concepts and details, beginning with an evolutionary perspective.

C.6.1 Evolutionary perspective on the Sequential Imperative

Within the terms of the model in Part B evolution can be thought of as producing brains with the neurological equivalent of increasing numbers of half-planes arranged in increasingly complex ways (through co-indexation and through clusterings, see 5.3 above) covering segmentation occupying increasing portions of time in the spine and/or increasingly finely divided portions of time. Each species therefore has its own instantiation of the equipment required to serve the SI and this (coupled with specific sensors and articulators) provides the behavioural repertoire and the learning capabilities for the species.

C.6.1.1 Language evolution

"Evolutionary pressure" operates on the equipment serving the SI and not on behaviours. Thus, in terms of "language evolution" and theories thereof, the issue is not how to account for the relationship between vocal language and gestural activity (for a recent account see Arbib 2005 and commentaries thereon). Rather, it is how to account for changes in the delivery of action and perception via the Sequential Imperative. In terms of the model this means looking for pressures to force the elaboration of the neurological equivalent of half-planes, the emergence of new half-planes, the realignment of clustering of half-planes, and so forth. Some of these changes may yield minimal apparent advantage, but taken together can provide a radical change in behavioural capability. A specific issue prompted by this use of the modelling concerns the emergence of a cluster of half-planes for controlling the vocal tract at much finer temporal resolution than is required for manual gesturing. If one wishes to subscribe to the notion of language emergence via gesture then this issue needs consideration.[7] Indeed it needs consideration regardless of one's view – the mistake is to assume the behaviour is evolving, not the mechanism supporting behaviour.

The sceptic might need re-assurance in more familiar terms. An analogy drawn from computing might help. Personal computers have steadily developed over time (incremental quantitative improvement without changing the architecture in any significant way). They now have faster processors, faster internal communications and more memory than 30 years ago (and different screens/input devices – of no concern here). The graphical capabilities of these machines, as evidenced in the interaction possibilities, have changed during this time in a qualitative manner. In particular, real-time graphical manipulation and rendering have become possible – where previously it was not. Incremental change in hardware (coupled perhaps with more fancy software) yields a radical step change in the user experience.

[7] See Bellugi and Studdert-Kennedy (1980), Klima and Bellugi (1979), Bellugi and Fischer (1972). The contrast between the relatively coarse-grained simultaneity of signing and the fine-grained linearity of speaking has been much discussed.

C.6.1.2 Modularity and neural networks

The model described in Part B prompts the thought that modularity, in Fodor's sense (1983), should be revisited. However, clusters of half-planes, whilst they may resemble "modules" (and perhaps offer an account for their evolution) don't actually express that conception very cleanly. Modular isolation and specificity are not strongly evidenced in *Pantome*, which in any case is closely coupled to the time domain and the requirements for sequencing and desequencing (which are not part of Fodor's conception of modularity). Likewise, neural networks as a modelling conception don't really map onto much of the model in Part B – except perhaps in the BGP. Neural networks aspire to generality of cognitive modelling but are not the sort of general cognitive principle we are concerned with here. Neither conception recognises the SI and neither addresses the issue of how to serve the SI.

C.6.2 Mirror neurons

The model deals with both output and input as linked – spinal segmentation specificational detail fills up through both routes (sequencing and desequencing), during experience of brief activity or extended contemplation. Languages are learnable because they can be produced (some aspects of language acquisition fit rather obviously in the model – e.g. successive refinement/analysis of originally holistically interpreted phrases), but this is equivalent, in an interesting sense, to saying they can be learned because they can be perceived. The segmentation issues are the same. Mirror neurons in this context look like detailed expression of the bi-directional activity involved in populating the spine with segmentational details.

C.6.2.1 Segmentation and symbolization

If mirror neurons are part of the neural expression of the close coupling between sequencing and desequencing (GCP3) then what is being mirrored must be fairly abstracted from the behavioural details (in which there is bound to be considerable variation, in perception as much as the actual behaviour (5.3.6), and from instance to instance). In terms of the non-linear model adopted here this places the neurons in the BGP. The abstraction involved looks like a step towards symbolization (cf. Bickerton (2005)), and it is therefore reasonable to suppose that behaviour produced by one member of a species can be "read" in part by another member of that species (and indeed sometimes other species).

C.6.3 Integrating linguistic theories

A major formal problem in conventional linguistics is that analyses and models focus on parts of the whole – syntax, or phonology, or semantics, or pragmatics – instead of unifying all these aspects. Part of this difficulty seems to be that each domain appears to be a significant problem in its own right. The non-linear 3D model described in Part B is scale independent and thus can apply to *and integrate* all aspects of language behaviour in a coherent way. Whereas learnability of, say, syntax or phonology, appears to be problematic when described in a formal way (see Bowerman (1988), Truscott and Wexler (1989), Tesar and Smolensky (1998)), what we can see from the generic non-linear model is that syntax is not necessarily readily isolated as a separable component and learnability is actually about the whole of language. More abstractly – the overall approach taken here accounts for language behaviour as behaviour. Other approaches (e.g. Cognitive Linguistics) attempt to broaden an essentially linguistic account to cover other relevant behaviours. The SI and the GCPs are all about behaviour, and therefore *necessarily* part of understanding language behaviour and formulating linguistic accounts. In that sense, then, Dik's notion of the model of a natural language user (M.NLU) is closer to the perspective offered here, even though his work on Functional Grammar (Dik (1989)) is restricted to just the linguistic capacity of a natural language user.

C.7 Functional specification of the brain – II

We move now from using the model from part B to provoke (re)thinking about language behaviour, behaviour generally, and cognitive architectures – demonstrating the value of the Sequential Imperative – to consideration of the General Cognitive Principles as illustrative of wider themes in a Cognitive Science research programme.

C.7.1 GCP7 and SETI

The remaining GCPs are not for detailed discussion here but, taken with the four discussed, and some sort of computational model – perhaps along the lines sketched above – they constitute a multidisciplinary programme for Cognitive Science. GCP7, in particular, has already been the subject of considerable attention over two decades (cf. Bærentsen (1989) – cited in de Léon (1999), Edmondson and Beale (2008), Hutchins (1995)). But as well as considering the GCPs in isolation the Cognitive Science programme should deal holistically with the GCPs as providing an account of the brain's functionality. One illustration of this sort of work comes from a recent proposal for searching for extra-terrestrial intelligence (Edmondson and Stevens (2003)) where it is conjectured, on the basis of understanding both the SI and GCP7, that such intelligences must have brains with functionality matching ours. It really is the case that we can say something specific about how embodied brains *must* function, wherever they are found in the universe.

C.7.2 GCPs 4-8 as functional corollaries

It was noted earlier (section 4) that whilst the first three GCPs constituted the Sequential Imperative, the remaining GCPs constituted a set of functional corollaries. The functionality of the brain – any brain, anywhere – is to serve the Sequential Imperative, but this requires more than just a statement of the imperative (whether as three GCPs or whatever). GCPs 4-8 set out what seem to be the necessary functional corollaries. It might be argued, for example, that reference to affect (GCP6) is a peculiarly terrestrial requirement, and that GCP5 could be rewritten to encompass the possibility of affect. However, affect seems to be an important aspect in much animal behaviour (from alarm calls in birds to weeping at funerals) and thus is offered a specific mention. The functional corollaries are required, notwithstanding the comment on GCP6, to ensure that the SI can be served in any but the most basic fashion. Without them, one can speculate, an organism is reduced to little more than stimulus-response behaviour, such as one can find in complex form in bees (although even this requires GCP4). GCP4, to paraphrase Chomsky, is key – if the neurological equivalent of the model proposed in Part B is supported at least by GCP4 then language learnability is unexceptional and universal grammar is vacuous. Indeed – the model is not the core issue. The SI, supported by at least GCP4, is core to language learnability. The objection that this simply restates Chomsky's comment, rather than addressing it, is refuted by the account above where the point is made that learning is obviously required for sequencing/desequencing to be possible at all (whether walking, making tea, riding a bicycle, or debating scriptures).

C.7.3 Consciousness

One remaining issue deserves further mention: consciousness. It was suggested earlier (section 3, and see also 5.3.5) that back-ground processing (BGP) is sub-conscious but that when creation/modification of cognitive entities is mediated via any sort of sequencing/desequencing this becomes conscious. Sensation is linked to the spine in the 3D model, in a time aligned way (perhaps a sort of output tier), and thus can be linked to segmentation supplied from the BGP – such as "speech in the head" which feels like thinking which could be "output" but which is not. Non-articulatory segmentation can be supplied from the BGP – "imagistic thinking" for example provides another route for sequencing/desequencing which produces no external output.

Consciousness, in the general sense, is not limited to the idea of thinking, but is rather the sensation of sequencing and desequencing, noting in passing that some of the segments may be quite long, in real time, and that attentional mechanisms such as affect (cf. GCP6) may be part of the sensational repertoire (with more or less "self-awareness" in relation to affect driven attention). It should be noted that attention does not have to be conscious (as when one attends to a vertical edge to monitor one's posture/balance – only to stagger when that edge moves because it is a train, not a building).

Because segments of various durations co-occur (and some stretch into the future in the anticipatory sense discussed above in 5.3.6, following insertion into the spine from various BGP) there are inevitably some complex co-occurring sensations which add up to the sense that consciousness is more than just "now", and even that "now" itself has a complex sense of extending from the past and into the future. Not only does this permit complex configurations of meanings to be attached to actions (the semiosis of behaviour and culture – see also Lotman's (1990) notion of the semiosphere) it yields a sense of consciousness as functional. However, it is not at all clear that the GCP-based model sketched above is missing a functional component which has to be provided by consciousness – it really isn't obvious that consciousness *does* anything (it isn't BGP, for example); it is the sensation of attention, not the driver of anything. Species differences do not obviously reside in different experiences of consciousness although these will occur – different species will have different clusterings of half-planes, different degrees of "look-ahead" in aniticipatory segment building, different sensory apparatus, and etc. So it is that the commonality of brain functionality – the underlying ubiquity of the SI and all that goes with it – makes inter-species "mind reading" less startling than might be supposed (guide dogs for the blind provide some inter-species support for notions of distributed cognition, for example, see GCP7). In short, then, consciousness is not part of any functional specification of the brain, nor need it be.

8 Conclusions

This paper provides a multidisciplinary, and potentially computationally realisable, account of the core functionality of the brain through discussion of some plausible General Cognitive Principles, and the Sequential Imperative. The account exemplifies a general cognitive approach to language learning and linguistic behaviour, and thus responds to Chomsky's comment quoted at the outset. Furthermore, the work reported illustrates the feasibility of such an approach to a Cognitive Science Research Programme. The value of the General Cognitive Principles, and the obviousness of the Sequential Imperative, combine in the observation that, without intellectual over-reaching, our understanding of cognition, and of the functionality of the brain, must apply to embodied brains elsewhere in the universe.

References

Abercrombie, D. 1965. Parameters and phonemes. In *Studies in phonetics and linguistics*. London: Oxford University Press. The paper was given, under the title 'The Contribution of Phonetics', to an International Study Group on the Development and the Disorders of Hearing, Language and Speech in Children, held in Durham in September 1963, under the auspices of the Medical Education and Information Unit of the National Spastics Society. Published in *The Child Who Does Not Talk*, Clinics in Developmental Medicine, No. 13, London, 1964.

Anderson, S.R. 1985. *Phonology in the Twentieth Century.* Chicago: The University of Chicago Press.

Arbib, M. A. 2005. From monkey-like action recognition to human language: An evolutionary framework for neurolinguistics. *Brain and Behavioural Sciences*, **28**:105-167.

Bærentsen, K. B. (1989). Mennesker og maskiner. Et virksomt liv: Udforskning af virk- somhedsteoriens praksis. M. Hedegaard, Hansen, V. R. and Thyssen, S. Aarhus, Aarhus Universitetsforlag**:** 142–187.

Bellugi, U. and Fischer, S. 1972. A comparison of sign language ad spoken language. *Cognition*, **1**:173-200.

Bellugi, U. and Studdert-Kennedy, M. 1980. *Signed and Spoken Language: Biological constraints on linguistic form*. Report on the Dahlem Workshop, Berlin, 1980. Verlag Chemie.

Bickerton, D. 2005. Beyond the mirror neuron – the smoke neuron? Commentary on Arbib (2005) in *Brain and Behavioural Sciences*, **28**:126.

Bowerman, M. 1988. The 'No Negative Evidence' Problem: How do children avoid constructing an overly general grammar? In *Explaining Language Universals,* edited by John Hawkins. Oxford: Blackwell.

Brentari, D. 2002. Modality differences in sign language phonology and morphophonemics. In *Modality and Structure in Signed and Spoken Languages,* edited by Richard P Meier, Kearsy Cormier, and David Quinto-Pozos, Cambridge University Press.

Chomsky, N. 1980. *Rules and Representations.* Oxford: Basil Blackwell.

Clements, G.N. 1985. The geometry of phonological features. *Phonology Yearbook 2.* Cambridge: Cambridge University Press.

de Léon, D. (1999). Building Thought into Things 3rd. European Conference on Cognitive Science, Siena, Italy.

Dik, S. C. 1989. *The Theory of Functional Grammar – Part 1: The structure of the Clause.* Dordrecht: Foris.

Edmondson, W. H. 1986a. Autosegmental Phonology: The syllable and the sign. In *Signs of Life*, ed B.T. Tervoort. Amsterdam: University of Amsterdam, Institute for General Linguistics, publication 50.

Edmondson, W. H. 1986b. Issues in Linearization: Prolegomena for a general theory of communication. *Language and Communication*, **6**(4):225-266.

Edmondson, W. H. 1989a. Asynchronous parallelism in the formalism of non-linear phonology. In *The Structure of Multimodal Dialogue*, edited by M. M. Taylor, F. Néel, and D. G. Bouwhuis, published by North Holland, 1989.

Edmondson, W. H. 1989b. Asynchronous parallelism in human behaviour: a cognitive science perspective on HCI. *Behaviour and Information Technology*, **8**(1): 3-12. 1989.

Edmondson, W. H. 1990a. Segments in Signed Languages: Do they exist and does it matter? In *SLR'87: Papers from the Fourth International Symposium on Sign Language Research*, eds. W. H. Edmondson and F. Karlsson. Hamburg: SIGNUM Press.

Edmondson, W. H. 1990b. A Non-Concatenative Account of Classifier Morphology in Signed and Spoken Languages. In *Current Trends in European Sign Language Research*, eds. S. Prillwitz and T. Vollhaber. Hamburg: SIGNUM Press.

Edmondson, W. H. and Beale, R. 2008. Projected Cognition – Extending Distributed Cognition for the study of human interaction with computers. *Interacting with Computers*, **20**:128-140

Edmondson, W. H. and Iles, J. P. 1994. A non-linear architecture for speech and natural language processing. Proceedings of ICSLP 94 - International Conference on Spoken Language Processing, Yokohama. Vol 1, pp. 29-32.

Edmondson, W. H., Iles, J. P., and Iskra, D. J. 1996. Pseudo-articulatory Representations in Speech Synthesis and Recognition". Proceedings of ICSLP 96 - International Conference on Spoken Language Processing, Philadelphia. pp. 2215-2218.

Edmondson, W. H. and Stevens, I.R. 2003. The utilization of pulsars as SETI beacons. *International Journal of Astrobiology*, **2**(4):231-271.

Firth, J. R. 1957. *Papers in Linguistics: 1934 - 1951.* London: Oxford University Press.

Fodor, J. A. 1983. *Modularity of Mind*. MITPress.

Geldard, F. A. 1953. *The Human Senses.* New York: John Wiley & Sons.

Goldsmith, J. 1976. *Autosegmental Phonology.* PhD dissertation. MIT

Goldsmith, J. 1990. *Autosegmental and Metrical Phonology*. Oxford: Blackwell.

Halle, M. and Clements, G. N. 1983. *Problem book in phonology.* Massachusetts: MIT Press.

Hammond, M. 1988. On deriving the well-formedness condition. *Linguistic Inquiry*, **19**:319-325.

Hutchins, E. 1995. *Cognition in the Wild*. Massachusetts: MIT Press.

Hayes, B. 1990. Diphthongisation and coindexing. *Phonology*,**7**(1):31-71.

Iles, J. P. and Edmondson, W. H. 1994. Quasi-articulatory formant synthesis. Proceedings of ICSLP 94 - International Conference on Spoken Language Processing, Yokohama. Vol 3, pp. 1663-1666.

Iles, J. P. 1995. *Text-to-Speech Conversion Using Feature-Based Formant Synthesis in a Non-Linear Framework.* PhD. School of Computer Science. University of Birmingham.

Klima, E. S. and Bellugi, U. 1979. *The Signs of Language*. Harvard University Press.

Lenneberg, E. H. 1967. *Biological Foundations of Language.* New York: John Wiley & Sons.

Lotman, Y. M. 1990. *Universe of the Mind – A semiotic Theory of Culture*. New York: Taurus.

McClelland, J. L. and Rumelhart, D. E. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models.* Massachusetts: MIT Press.

Meier, R. P., Cormier, K., and Quinto-Pozos, D., (editors) 2002. *Modality and Structure in Signed and Spoken Languages*. Cambridge University Press.

Palmer, F. R., (editor). 1970. *Prosodic Analysis.* See Palmer's introduction. London: Oxford University Press.

Putnam, H. 1971. The 'Innateness Hypothesis' and Explanatory Models in Linguistics. In *The Philosophy of Language,* edited by J. R. Searle. Oxford: Oxford University Press.

Rumelhart, D. E. and McClelland, J. L. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations.* Massachusetts: MIT Press.

Sagey, E. 1988. On the ill-formedness of crossing association lines. *Linguistic Inquiry*, **19**:109-118.

Sandler, W. 1990. Temporal Aspects and ASL Phonology. In *Theroetical Issues in Sign Language Research: Vol 1. Linguistics*, eds. S.D. Fischer & P. Siple. Chicago: The University of Chicago Press.

Sandler, W. 1993. A sonority cycle in American Sign Language. *Phonology,* **10**(2):243-279.

Tesar, B. and Smolensky, P. 1998. Learnability in Optimality Theory. *Linguistic Inquiry* **29**(2):229-268.

Truscott, J. and Wexler, K. 1989. Some Problems in the Parametric Analysis of Learnability. In *Learnability and Linguistic Theory*, edited by Robert Matthews and William Demopoulos. Dordrecht: Kluwer Academic Publishers.

Van Vallin, R. D., (editor) 1993. *Advances in Role and Reference Grammar.* Amsterdam: John Benjamins

Wilbur, R. 1990. Why Syllables? What the notion means for ASL research. In *Theoretical Issues in Sign Language Research: Vol 1. Linguistics*, eds. S.D. Fischer & P. Siple. Chicago: The University of Chicago Press.