The Abstraction Engine

Extracting patterns in language, mind and brain

Michael Fortescue

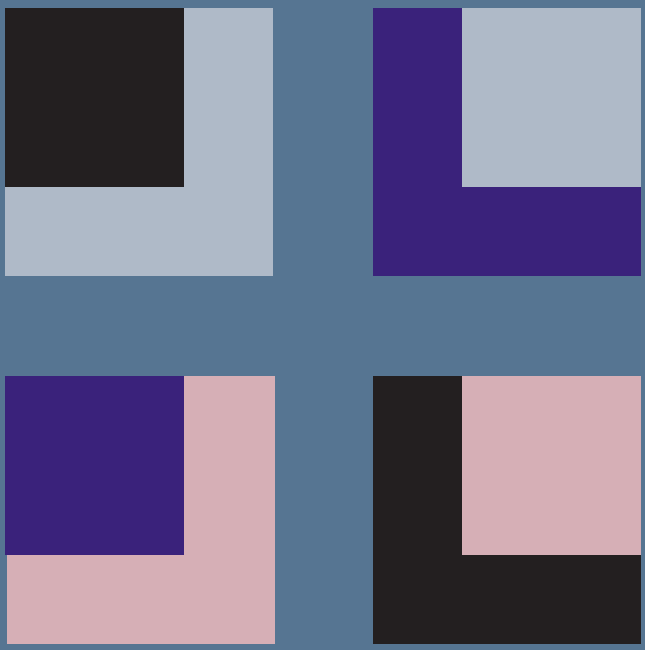

Advances in Consciousness Research 94

John Benjamins Publishing Company

The Abstraction Engine

Advances in Consciousness Research (AiCR)

ISSN 1381-589X

Provides a forum for scholars from different scientific disciplines and fields of knowledge who study consciousness in its multifaceted aspects. Thus the Series includes (but is not limited to) the various areas of cognitive science, including cognitive psychology, brain science, philosophy and linguistics. The orientation of the series is toward developing new interdisciplinary and integrative approaches for the investigation, description and theory of consciousness, as well as the practical consequences of this research for the individual in society.
From 1999 the Series consists of two subseries that cover the most important types of contributions to consciousness studies:
Series A: Theory and Method. Contributions to the development of theory and method in the study of consciousness; Series B: Research in Progress. Experimental, descriptive and clinical research in consciousness.
This book is a contribution to Series A.

For an overview of all books published in this series, please see
http://benjamins.com/catalog/aicr

Volume 94

The Abstraction Engine. Extracting patterns in language, mind and brain
by Michael Fortescue

The Abstraction Engine

Extracting patterns in language, mind and brain

Michael Fortescue
University of Copenhagen and St Hugh's College, Oxford

John Benjamins Publishing Company
Amsterdam / Philadelphia

DOI 10.1075/aicr.94

Cataloging-in-Publication Data available from Library of Congress.

ISBN 978 90 272 1361 7 (HB)
ISBN 978 90 272 6584 5 (E-BOOK)

John Benjamins Publishing Company · https://benjamins.com

Table of contents

CHAPTER 1

Introduction

Abstraction all the way down

A principal claim of this book is that the brain is not so much an organ of computation – the currently prevailing view – as one of abstraction, which in turn provides the basis for all forms of subsequent projection. This is so, I shall argue, not just as regards language – the most obvious manifestation of abstract cognition in human beings – but also as regards sensory perception, memory, emotional and aesthetic experience, and the social scenarios on which most action and planning is based. Abstraction can moreover be layered: abstractions over abstractions are possible at successively higher levels, both ontogenetically and phylogenetically. The function of abstraction – according to the definition with which I shall finally provide it – is one shared by all creatures with brains, however primitive. The main thrust of the book will be a search for the deepest roots of abstraction, specifically in the activities of the human brain. It will draw upon many areas in which processes of abstraction can be seen, in effect presenting a tapestry of ideas woven from many different threads. Sometimes the coverage will be quite superficial, sometimes quite detailed, the purpose being essentially to present a coherent synoptic overview of a far-reaching subject that will lead up to a novel conclusion.

I shall at times be using the word 'abstract' (literally 'drawn out from') in a more dynamic and general sense than normal in daily usage, one close to Whitehead's notion of 'transmutation', of which I shall have a good deal to say later. It has in fact at least two different senses that have often been confused. On the one hand there is the more traditional sense of 'disembodied from specific sensory features' (let us call this 'abstract$_1$'), and on the other there is the sense of 'simplified or generalized across instances' (let us call this 'abstract$_2$'). A typical dictionary entry, corresponding to our everyday use of the word, such as the following from Chamber's 10th edition (2006), does not resolve the confusion: "Apart from actual material instances, existing only as a mental concept, opposite of concrete, away from practice, theoretical, denoting a quality of a thing apart from the thing, e.g. 'redness'." A point that such definitions miss is that 'abstract' words in the traditional dictionary sense usually have more complex meanings than concrete ones in so far as they refer to complex contexts of a social or moral nature that require elaborate definitions, and this, as we shall see, is reflected in the way in which they

are distributed in the mental lexicon. Such words are only 'simpler' in the sense of applying holistically to less easily visualized meanings whereas concrete ones are 'simpler' in being more easily grasped through discrete sensory images. Another misleading use of 'abstract' is in contrasts made between 'abstract' – i.e. arbitrary – symbols and more sensorially motivated 'icons' and 'indices' (I shall return to this Peircean distinction between types of sign in what follows). It remains to be seen whether anything at all can be salvaged of the vague distinction between 'abstract' and 'concrete'. More importantly, I shall attempt in what follows to disentangle the various senses of the word 'abstract' and to see if some broader perspective on abstraction can be revealed that will cover them all.

Intimately linked to the concept of 'abstraction' is that of 'projection', which is dependent on preceding abstraction. It is close to the notion of 'mapping' in cognitive science, but again I take the term in a broader procedural sense than when cognitive linguists discuss the relationship between schemas in one domain 'mapping' onto another. As I use the term it is operative on multiple levels of perception and action, including acts of linguistic 'referring'. It corresponds better to Whitehead's use of the term to refer to the projection of sensory input in the mode of 'presentational immediacy', in other words to the creation of a perceived world through the combination of an inner past abstracted from experience and the present given to the senses. This is very different from the Humean input-output view of perception still shared by adherents of the digital computer model of the brain. Perhaps the analogy of a movie 'projected' through a complex, hidden apparatus onto a screen is not entirely out of place here (although the notion of an observing homunculus certainly is). However, 'projection' should not be understood as confined to perceptual experience and bodily actions based on them – it also covers the products of imagination and the deployment of language both internally and externally. The externalized results of projection on the environment will in turn need to be abstracted in order to be assessed as successful or not.

Another way of looking at the overall functioning of the brain is in terms of prediction, of the statistical calculation of probabilities: it is continually concerned with perceptual expectations and appropriate reactions based on accumulated experience. The title of the present book could in fact have been "The prediction engine", but I prefer the one I have chosen since it suggests an active process pointing both forward and backward in time.[1] It better captures the extent of this fundamental process beyond purely logical calculation, extending into the more fluid realms of perceptual and emotional experience. Abstraction and prediction are in fact two sides of the same coin since any prediction affecting future behaviour must perforce be based on abstraction from previous experience, just as what is abstracted

1. See Clark (2013) for a related notion of the "prediction machine".

from on-going experience is determined by whatever is predicted to be relevant by the experiencer, consciously or otherwise. I use the word 'engine' in a somewhat tongue-in-cheek reference to Charles Babbage's 'difference engine', the precursor to the modern computer, designed but never completed in the early 19th century.[2] What I have in mind is more of a 'similarity engine'. Pattern recognition, generalization, analogy, intuitive leaps to probabilistic conclusions as well as the learnt patterns of logical inference, all of these rely upon abstraction in its broadest sense.

Throughout the book I shall attempt to stay as close as possible to what contemporary neuroscience tells us of the functioning of the human brain in terms of the continual updating of connections between neural assemblies in networks of neural associations. At the same time I shall maintain a bird's eye view over the purposes for the individual that the fundamental process of abstraction serves. I shall take as starting point that quintessential realm of abstraction, language, then proceed downwards through successively more basic layers of cognition towards its neural underpinnings. Finally I shall pull back to the broader picture and what is perhaps the most complex level of abstraction of which the human brain is capable – the level at which it becomes possible for the brain to enquire reflexively into the roots of its own capacity for abstraction.

Of course the brain does other things beside derive abstractions, it also manipulates them in various ways in different sensory-based modalities. These may cross over, be correlated and linked to produce still further abstractions, hierarchically or sequentially organized in higher centres of the cortex. It performs remarkable feats of 'cognitive integration' of information from multiple sensory sources. It may deconstruct and reassemble sensory experience, synthesizing new combinations in the 'seemingly seamless' products of imagination, but the 'stuff' worked upon must itself have been abstracted from experience with the help of the innate biological machinery with which we are all born. The endpoint of abstractive processes is often further concrete or mental action, but such actions are themselves the projection of stored abstractions – plans and routines that subserve goals. In realizing them, cascades of neural activity will flow from high level intentions (potentially conscious) to lower level implementations far too rapidly for conscious experience of the distinct role of the chained processes to be possible.

In Chapter 2 I shall deal with language as an integrated higher level abstraction-projection function, operating largely at the level of recognizing or generating abstractions from more basic sensory patterns, themselves abstracted from environmental experience. The abstractions we will be concerned with here cover the

2. Occasionally one may meet the term 'abstraction engine' used for search programs in the context of Information Technology – I should stress that 'abstraction' in the present title is to be taken in a much broader sense than this.

recognition of phonemes and of the hierarchically organized patterns of syntax, as well as the meanings of individual words and phrases – including the function of purely grammatical words and constructions. This will lead to consideration of their combination for communicative ends as speech acts fine-tuned by appropriate prosodic moulding. From this perspective even the function of referring can be seen as based on abstraction-plus-projection. Some recent psycholinguistic models for the basic processes involved will be sketched. This kind of abstractness needs to be carefully distinguished from the purely descriptive abstractness of Chomsky's influential brand of linguistics, which searches for 'explanatory' universals. Its historical stages will be illustrated briefly. Abstract lexicon in both familiar Indo-European and less familiar Inuit languages will be touched upon, as well as the special kind of abstractness of Chinese logograms. It will be argued that the distinction often loosely made between concrete and abstract words is quite misleading.

In Chapter 3 I shall review the perspective on abstraction provided by a particular school of linguistic thought that has been gaining momentum as an alternative to Chomskyan generative linguistics in recent decades, one that does not rely on assuming the innateness of grammatical abstractions. In contrast to the latter approach, Cognitive Linguistics emphasizes the anchoring of language in sensorimotor 'image schemas' – it is 'embodied', and its surface abstractions are projections, principally via metaphor, from these pre-linguistic sources. This kind of linguistics (which has much in common with functional and neurolinguistic approaches) has a good deal to say on the subject of abstraction (see for instance Langacker 2000), although, as will be seen, it is not entirely free of the terminological confusion around the term 'abstract'. I shall argue that while much of language is indeed 'embodied', there is also much that is truly 'disembodied' and nothing to do with metaphor, a view that is acknowledged by some within the ranks of Cognitive Linguistics itself.

The following chapter, Chapter 4, takes a look at the role of cultural and social practice in supplementing the embodied view of language. Abstractions at this level – in terms of 'frames' and 'scenarios' – are an essential ingredient in both individual word meanings and in the projection (or deployment) of language in communicative situations. Again, this is an area that has drawn the attention of practitioners of Cognitive Linguistics more and more in recent years. It leads to the more general question of the role of context not only in language but in all forms of cognition. It will be seen that there are various layers of context, something we usually take for granted (we are so used to assuming them). An excursion into the nature of metaphor in hunter-gatherer societies will illustrate the important role of cultural and physical context.

The next three chapters start to leave the realm of language altogether in order to investigate the role of abstraction and projection at the level of event structures

and the way in which they are retained in memory (Chapter 5); of perception, imagination and dream (Chapter 6); and of abstract emotion (Chapter 7). These are vast subjects that have extensive literatures within the different branches of psychology and psycholinguistics that can only be summarized in these brief coverages. The purpose is to show that the fundamental processes of abstraction and projection that we principally associate with language have much deeper roots. At the forefront of these chapters is the relationship of perceptual experience and memory to processes of abstraction. Chapter 5 is concerned with our simplified structuring of events in both narrative and direct experience as retained for longer or shorter periods in memory, for instance in the form of the 'mental models' of Johnson-Laird. Paivio's theory of dual coding, which posits distinct linguistic and perceptual codes, will be touched upon in Chapter 6, as will Globus's creative theory of dream projection. The special kind of kinetic abstraction involved in gesture will also be discussed (drawing upon McNeill's theory of the 'growth point'). Chapter 7 will consider the distinction between what Damasio calls 'feeling' and 'emotion', illustrated with a sketch of the expression of emotion in Inuit languages, and concludes with a look at the abstract expression of emotion in art and music.

In Chapter 8 I come to terms with the neurological underpinnings of abstraction and projection. The relevance of connectionism, a centrepiece of modern cognitive science, is discussed. Its 'hidden layers' throws some light on the ontological cleft separating the external world of objects and our conscious experience of it. The coverage here is not technical – the mathematical abstractness of this approach to neural networks tends to obscure the highly concrete biological reality it reflects. The main thrust of the chapter is to illustrate how meanings (whether 'abstract' or 'concrete') can be represented in distributed neural networks spread over myriads of individual neural units and their inter-connections, the statistical, probabilistic weightings of which are constantly being updated by new input. The nature of the processes whereby such cortical traces accumulate will again be discussed in terms of abstraction, i.e. generalizations across experience bearing upon the concepts concerned. What is known of the hierarchical organization of the cortex and its pathways of feedback and feed-forward will be drawn upon in an attempt to clarify these deep layers of abstraction.

Chapter 9 focuses on the working of abstraction through the early stages of the life of the individual, as the child learns about the world around it and comes to master the basics of adult language use. The process of course continues throughout life as the adult adds to and refines their knowledge of the world and of themselves, but it is at this early period that the groundwork is laid for all further development. The acquisition of language is the paramount exemplar of pattern-extraction through abstraction from sensory input (irrespective of the role genetic predispositions play). The successive phases of this process as passed through by

children worldwide has been shown in numerous studies to be largely the same everywhere, with relatively minor and predictable variations according to the nature of the target adult language. This presupposes earlier, still more universal prelinguistic development, such as the sensorimotor and functional stages of concrete and formal operations described by Piaget, or the more socially determined stages described by Vygotsky.

Chapter 10 takes a further step back to observe the wider philosophical picture. The nature of abstraction (whether as ideas, concepts or universals) has occupied philosophers of all shades throughout the ages. After a brief characterization of some of the main perspectives that have been applied to this subject from Plato to Wittgenstein I shall concentrate on one particular approach, namely that of Alfred North Whitehead. This is not an arbitrary choice, for Whitehead had much to say about abstraction ('transmutation') as a general process applicable not just to language but to perception and emotion and much else besides. For him perception covers two distinct modes, that of 'causal efficacy' (vague but powerful subjective echoes of bodily engagement with objects in the past) and that of 'presentational immediacy' (the distinct 'here and now' of sensory experience). The continual interplay between causal efficacy and presentational immediacy is the usual way in which we experience the perceptual world. 'Transmutation' is part and parcel of this – it is the process whereby aggregates of related *sensa* are summated and experienced as unitary objects displaying relevant attributes. The relevance of this for the mental lexicon will be highlighted. The chapter is rounded off with a consideration of the concepts crucial to abstraction of 'similarity' and 'context' which figured prominently in the later Wittgenstein's thought.

Finally, in Chapter 11, I bring together the various threads of the story by showing the essential role of abstraction in the evolution of our species. The question addressed is this: what is the purpose of abstraction, what end has it served so relentlessly, accompanying every stage of our development through countless millennia from unthinking automata to conscious beings? The starting point is with the lowliest form of life displaying a nervous system, and the endpoint is with the highest forms of abstraction of which we as cultural and imaginative beings are capable ('the future', 'the world', 'the Good', etc.), a level only attainable with the help of language. The answer proposed will tie in both with Jackendoff's 'Intermediate Level' theory of consciousness and with Popper's notion of problem-solving as the guiding force behind evolution in general and of 'World 3' abstract cultural artefacts as the highest human extension of this. An even broader perspective will be suggested in terms of the maximizing of negative entropy ('useful order') by natural processes that tend towards ever more fine-grained informational discriminations in the face of universal thermodynamic decay. In the process 'abstract$_1$' and 'abstract$_2$' will be generalized to a more comprehensive 'abstract$_0$'.

CHAPTER 2

Language, abstract and less abstract

2.1 Levels of linguistic abstraction

The word 'abstract' is used by linguists and dictionary makers alike with reference to certain types of word meanings. Of the following English words (all of which will be touched upon in what follows) which would *you* consider 'abstract'?

> –accommodation; anger; animal; be; beauty; big; cat; conquer; cooking; entropy; farm; free; game; holiday; house; love; marriage; mercy; murder; open; primate; redness; same; sincerity; soul; thing; truth; up; weather

But wherein exactly does the abstractness of the words you select reside? How does it relate to the general abstractness of language as a symbolic system? Trying to answer these questions is not so simple as it might appear – indeed, the more you look into it the more it you begin to feel a little like Alice when confronted by the curious behaviour of the Cheshire Cat:

> *"I wish you wouldn't keep appearing and vanishing so suddenly: you make one quite dizzy."*
> *"All right," said the Cat; and this time it vanished quite slowly, beginning with the end of the tail, and ending with the grin, which remained some time after the rest of it had gone.*

Are the meanings of abstract words like the grin of the famous cat? Can one peel away the concrete trappings of instances to which they apply and be left with their 'abstract$_1$' essence ('disembodied from specific sensory features')? Of course this particular sensory feature (the cat's grin) exists only in the imaginative world of Lewis Carroll, but one might say that it stands emblematically for the whole creature he invented. In so far as the appearance of the grin will henceforth always herald the potential appearance of the whole creature could it be abstract in the other sense of the word, 'simplified or generalized across instances' ('abstract$_2$')? One might say that the first kind of abstraction is a matter of 'extraction' but also of 'attraction' (of complex associations), the second of 'subtraction' or 'contraction' (to simplified schemas). But don't be distracted by the word play here – I do not wish to detract from a serious point, which is that the '-tract-' part of all these words is in some way analogous to the Cheshire Cat's grin. In English (but not in

the original Latin source) the verbal root is decidedly 'abstract', although not hard to guess. You probably use it every day in one form or another without realizing the concrete meaning behind it. (To quote the Cheshire cat again: "You can't help that, we're all mad here.")[1]

I shall not attempt to follow up this somewhat cryptic line of enquiry here – I shall be returning to a closer consideration of abstract lexical items in section 2.4. Right now I wish to take up instead the question of the role that abstraction of one sort or another plays in our comprehension of language at a number of different levels. I concentrate here on language perception – I shall have more to say on language production later. Each level presupposes a specific kind of context, and these are summarized in Table 1. All of them may be simultaneously operative in any given situation of language use.

Table 1. Linguistic contexts

a)	phonological context – the recognition of successive phonemes
b)	individual words – the semantic fields to which they belong
c)	syntactic (and morphological) constructions in which words are embedded
d)	textual/ discourse embedding
e)	the physical setting of the communicative situation
f)	speech act conditions
g)	facial expressions and gestures of one's interlocutor
h)	attitude or intention towards one's interlocutor and/or the subject of discourse
i)	stylistic context or 'genre'

First, consider how we abstract strings of phonemes, the building blocks of words, from the stream of acoustic information reaching our ears. A good deal is known about this process, both as regards adult perception of the sounds of their language(s) and as regards the acquisition by children of the ability to distinguish them. The abstraction of categorical units from an acoustic signal is a complex matter but well charted in terms of phonetic 'clues' – frequency formants and transitions and specific kinds of 'noise'. It clearly involves not only the ability to transpose between different voices and speeds of delivery, etc., but also crucially depends on phonological context. By this I mean that a particular vowel or consonant will sound differently according to other sounds preceding or following it. Depending on the language concerned, such subtle adjustments may correspond

1. For the 'Cheshire Cat effect' in psychology see Crick (1994: 219). It involves the deleting of selective parts of a perceived whole by a 'trick' setup and in a sense is the converse of Gestalt theory 'filling in' of missing information.

to two distinct phonemes or just represent acceptable variation of one.[2] This is of course what defines a 'phoneme': contrast with other units within a system of choices such that its substitution by another may produce a different word. In languages with few vowels (like Greenlandic Eskimo) a single vowel phoneme may be realized in a wide range of variants depending on flanking consonants, and thus correspond to more than one English vowel.

In fact even for speakers of the same language, English say, phonological context (as well as speed and sloppiness of articulation) may obscure whole phonemes. But the human brain has a remarkable ability to recreate missing information, to calculate the most likely words being uttered (and their component phonemes) despite all manner of signal noise. This recognition of holistic 'gestalts' from imperfect input is of course what speech recognition is all about, a clear example of abstraction as a fundamental process in the human cortex – one of greater generality than its application to the meaning of words alone.

However, this process is more than simply a linear process of filtering the input in order to abstract successive phonemes that match the patterns embodied in the speech centres of the brain. Already at the phonological level it involves some parallel or hierarchical organisation into syllables and still higher 'metrical' units. This is where the 'double articulation' of language enters the picture: the string of phonemes that has been analysed can in turn be analysed at a higher level of abstraction into a succession of meaningful units, morphemes. This is a first example of a type of abstraction I shall call 'abstraction by matching' as opposed to 'abstraction by emergence' – the matching is of course with stored representations of morphemes (including affixes and whole lexical words). There is information in the stream of speech that provides clues as to where individual words begin and end (aided in English at least by patterns of stress assignment), and there are internal expectations as to which sounds are likely to follow a given sequence or are precluded from doing so. In some languages the choice of a vowel early in a word will determine a narrow choice of which vowels may follow within the same word (vowel harmony).[3] In languages like Welsh the grammatical context of a particular word may determine its phonological shape (by 'mutation' of initial consonants), and expectations generated at these higher levels (as well as statistical expecta-

2. Think of the English voiceless /t/, which is aspirated word initially but not following an /s/ – in languages like Nivkh these variants can contrast in the same position.

3. In a language like Chukchi the determining vowel may even be later in the word, so that earlier ones must adjust to it. Parallel or suspended processing is more likely than forward peeking. Just knowing the word concerned as a whole is not enough since a productive suffix may contain the determining vowel. Thus for example *ewlə-ɣərɣən* 'length' from *iwlə-* 'long' – the initial 'dominant' /e/ is due to dominant suffix *-ɣərɣən*.

tions at the level of individual words) will doubtless expedite the recognition of the phonemes themselves, only initiating further processing when expectations are not immediately met.

As regards the level of recognizing individual words, there are various psycholinguistic models for what this entails – I shall return to some of these in the following section. The point here is simply that this is indeed a further matter in which abstraction is involved – a honing down of the potential multiple purport of sensory input to allow a match with specific entrenched memory traces. It should be clear by now that abstraction in the broadest sense I am using for the word is not just a matter of 'abstraction from input', but also of 'abstraction from memory'. This latter side of the process will become more and more significant as we proceed. Ultimately the two must come together in the process of matching input with memory – mismatches will be acted upon, matches can be ignored if expected.

Whatever the details of how individual words are recognized, access to their meaning requires activating widespread networks of associations whose potential ramifications are curtailed by contexts of a different kind. These include the semantic context – the 'mental model' being abstracted from an on-going speech situation – and the parts of speech expected in the grammatical context. Thus if someone asks you if you have read Lewis Carrol's famous tale (pronounced /teyl/) you will first and foremost think of the Alice story, since a 'tail' (pronounced the same way) would not be activated in the context of the verb 'read'.[4] In the context of hearing someone read aloud the little dialogue given above, on the other hand, you would hardly interpret /teyl/ as referring to anything other than the cat's tail – and not say the tail of an aeroplane or the tail of a long dress.

Once activated, any given content word will have multiple associations to other specific words or constructions, including frequent conventional collocations. Intrinsically the word will also have its own position within systematic 'semantic fields' of relationships analogous to phonological systems but more complex and open-ended (as we shall see in the following chapter on prototype theory). It will also belong to a particular abstract (i.e. general) 'part of speech'. Nouns will often have hypernyms (more general 'kinds' which they belong to, like the relation of 'cat' to 'animal'), adjectives may have antonyms (like 'right' and 'wrong'), and verbs may have both near-equivalents ('disappearing' and 'vanishing') and opposites ('appearing' and either of the two verbs just mentioned). Verbs in grammatical context will reflect different kinds of 'states of affairs' (situation types such as actions with goals as opposed to activities without), in turn generating further expectations for a continuation of the text. These potential expectations may or may

4. But earlier in the Alice story there is indeed a 'tail' that can be read: the mouse's tale about Fury, which is presented on the page in the shape of a mouse's tail, a lovely visual pun.

not be relevant to the on-going construction of a mental model. Only those that are will be integrated with it. The degree to which abstracted meanings of words are truly holistic or decomposable into underlying conceptual components will be discussed in following chapters.

As regards grammatical context, each successive word must find a place within a syntactic structure suitable to its part of speech as well as being compatible in meaning to the listener's accreting mental model. These structures are themselves highly abstract, a fact which the application of convenient names to such things as 'noun phrase' or 'indirect object' may obscure. These labels probably do not have any representation as such in the brain, they are 'virtual' reflections of processes of grammatical abstraction, the nature of which I shall return to in later chapters. 'Verbs' are words that have certain grammatical expectations associated with them – their need to associate themselves with other abstract units such as objects or complement clauses (their 'case frames'). Thus 'keep' in the Alice dialogue expects either a noun phrase or – as here – a following gerund (an '-ing' form of another verb). The organization of grammatical templates is hierarchical, not just a matter of fixed linear succession, and the construction exemplified by 'keep appearing and vanishing' – a 'verb phrase' – is highly general. Thus 'keep' could also be interpreted in a more concrete sense in such collocations as 'keep your hat on' – still a verb phrase, but one consisting of a nominal object following the verb and a closely associated adverbial particle 'suspended' until the phrase is closed.

Then there is the context of the whole on-going discourse or text. The precise meaning of successive words will depend to some degree on the 'mental model' being built up in the mind of the interlocutor or reader – and also on the physical setting of the discourse. One linguistic function linked to both these kinds of context is 'reference' (one of the most basic of all). This may either be in terms of the environment in which communication is taking place, by verbal pointing as it were ('deixis'), or purely textual, referring back and forth to people, places or actions already mentioned in a text or dialogue. This is typically done by pronouns and adverbs, but there are other means as well (cf. Halliday and Hasan 1976). On the face of it, this is less obviously a matter of abstraction, though clearly one of projection. Actually, even the ordinary use of 'he' is abstract in the sense of being 'under-specified' and presupposing context to interpret. Reference in sign language is more obviously abstract – the signer assigns the referents he or she is discussing at specific 'loci' in their 'frame of reference' and points or turns in that direction to refer to them. As Engberg-Pedersen (1996: 41–2) explains with respect to Danish Sign Language, the choice of loci is free but not entirely arbitrary – empathy with a referent may result in choosing a locus in close proximity to the signer's own body, for instance.

Also in using spoken language we may project 'frames of reference' to help anchor our utterances in the specific spatial context in which we are situated. Apart from gestures like pointing and eye gaze to aid communication in shared spatial situations, there are choices of demonstratives and directional terms like 'to the left' and 'to the right'. Levinson (2013) discusses the alternative frames here as 'absolute', 'intrinsic' or 'relative' – to the speaker's 'left' or 'right' belongs to the last of these, whereas to the 'north' and to the 'south' are 'absolute', and the natural 'front' or 'back' of an object are intrinsic. These are indeed abstractions, though very deep ones with pre-linguistic roots, namely abstractions of bodily orientation in space.[5]

This leads us on to the context of 'speech acts'. The interpretation of an utterance will depend largely on the type of speech act or 'illocution' it is performing. These are best seen in terms of conventionalized 'conditions' (cf. Searle 1969). Thus the Cheshire Cat's "All right" can only be understood in the given context as an expression of compliance to Alice's request (another speech act) for him to stop making her dizzy. This is a rather specific but very simple response fulfilling that function. Alice's request is more complex and indirect – it is formulated as a wish (with a 'reason' tacked on). The cat understands the 'uptake' (or point) of this speech act, partly because of its internal constitution – using a verb and a construction conventionally used to express a wish that something might come about which the recipient is in a situation to do something about – but also at a higher level, because it is considered polite to respond to a wish of this kind, whether to desist or to refuse to do so, etc.

But beyond the conventionalized – or at least deducible – speech act situation, there are other contextual factors that have to do with the speaker's ultimate intentions, e.g. to get someone to do something other than what is overtly expressed, to influence them or indicate one's attitude towards them – the so-called 'perlocutionary' as opposed to 'illocutionary' purport of an utterance. A speaker may wish to express irony, honesty, jocularity, solidarity, or to confide in or comfort or mock their interlocutor. Such attitudes can also concern the subject of discourse rather than the interlocutor (or indeed be aimed towards oneself, expressing humility or its opposite). Intonation (not to mention innuendo) may play an important role here.[6] Facial expression and gesture is also relevant.

5. Linguistic means for referring to location in space may be further extended to orientation in time (Levinson op. cit.: 16) – cognitive linguists typically see the phenomenon as metaphorical.

6. But in some languages discourse particles play the major role. These have been analysed in detail by Davidsen-Nielsen (1996) for Danish. He distinguishes *sgu* (emphatic certainty), *nok* (unemphatic assumption), *jo* (unemphatic attachment), *da* (emphatic separation), *skam* (positive emphatic separation), *da* (emphatic presupposing separation), *dog* (negative emphatic separation), and *nu* (unemphatic separation).

The whole speaker-hearer-subject context needs to be abstracted and continually updated by both parts for successful communication to take place. The Cheshire Cat's grin may well be a non-linguistic expression of mockery, a perlocutionary intent expressed in a purely behavioural manner. Alice, in turn, shows a polite but reproachful (and somewhat confused) attitude towards the know-it-all creature. The relation of emotional expression to abstraction will be dealt with in Chapter 7.

Finally in this catalogue of contexts relevant to language, there is the stylistic context of the spoken utterance or written sentence. This is a matter of 'genre' as well as of the level of speech appropriate to a given situation as it affects the choice of individual lexemes. In the case of Lewis Carroll's tale, this is both a matter of literary genre (a story intended mainly for children, albeit an unusual and fantastic one, akin to dream) and of the level of speech used by Alice and the cat respectively to address each other (here familiar or conversational). This is a complex matter, and the recognition of a given stretch of text as belonging to a specific genre or style (as holistic concepts) obviously requires considerable abstraction across distributed linguistic features that must be matched with relevant experience.

In sum: language is much more than a string of words hierarchically organized in syntactic 'tree' structures to which the meaning of the individual words can be attached. In all instances of actual use a 'richer' meaning has to be abstracted from successive surface utterances (or 'sentences' if the text is written), relating to the various layers of context in which they are embedded. The surface utterance is just the tip of the iceberg. Thus to understand fully the meaning of Alice's remark "I wish you wouldn't keep appearing and vanishing so suddenly" you need to know not only about the forms and meanings of the individual words and constructions within it (both affected by immediate context), and to recognize that it expresses a wish and that the wish is for a change to the repetitive behaviour on the part of the cat she is addressing, you also need to understand that the cat is capable of complying with that wish if it wants. You can also deduce that such behaviour on its part is something Alice has not expected (hence the phrase 'so suddenly'). It is confusing to her (as her ensuing words explain), although she has actually seen it doing this vanishing trick in front of her previously more than once – the 'keep…-ing' construction here refers to an alternation over the short time she first encountered the cat, not, say, to a habitual occurrence once a day at the same time. The 'appearing' and the 'vanishing' are moreover conceptually linked and one can assume that they refer to the same object (the cat).

2.2 Some psycholinguistic models

At this point let me introduce a brief sketch of some of the actual processes involved in the brain according to psycholinguistic models of the abstraction of meaning of words and sentences from auditory input. An influential model for the kind of process involved in word recognition is the so-called 'Cohort' model (cf. Marslen-Wilson 1989 for the revised version), whereby the initial sound of a word as it is analysed out from the stream of speech triggers semi-activation of all words starting with it in the mental lexicon. As successive phonemes are analysed the 'search set' is rapidly reduced until there is a perfect match with just one word in memory, which is then fully activated. Again, this process is in reality not to be understood as entirely linear: there is ample scope for parallel processing and short-cuts (e.g. from higher level expectations of the endpoints), or by 'incremental' rather than simple linear processing (cf. Levelt 1993: 23–28). To give an example from the exchange between Alice and the Cheshire Cat, the word 'appearing' would (if we were somehow the cat) first semi-activate all words beginning with unstressed 'a' (phonemically /ə/), then limit the search to those beginning with /əp/ ('appeal', 'approach', 'appear', etc.), then those starting with /əpi:/, ('appeal', 'appear', etc.) and so on. But analysis in parallel of the prosodic shape of the word, its 'metric template', will probably speed up the process, limiting it to words consisting of a sequence of unstressed-stressed-unstressed syllables. The holistic shape of the word (affecting syllables boundaries and stress levels among other things) would also exclude possible misinterpretations such as 'a peer ring'.

Friederici (2002) proposed a four-phase model to account for ERP (event-related potential) data recorded during the early stages of comprehension. Each phase can be related to at least one particular region of the brain and to one known ERP phenomenon (specific electromagnetic indicators picked up by electrodes on the scalp). Friederici' s model of sentence comprehension assumes an initial syntactic analysis based on word category recognition (left hemisphere only) followed by one in which semantics plays an increasingly important role, but allows for very early intervention of prosodic information from the right hemisphere. The model starts its time course with primary acoustic analysis of the incoming signal in the auditory cortex, as described by Hickok & Poeppel (2007). This takes place in the dorsal superior temporal lobe prior to the splitting of the signal into two parallel 'streams', the 'ventral' one essentially for semantic processing and the 'dorsal' one for the sensorimotor interface to speech production. At the immediately ensuing 'phase 0', phonemes are identified and shunted to Wernicke's area in the superior temporal lobe, where during 'phase 1' word forms are identified, followed by the identification of the word's category (verb, noun, etc.), information which is further shunted to Broca's area where syntactic structure building is probably

focused. The ventral stream continues on to the mid-temporal lobe where information about the word's meaning and morphological behaviour is stored (or at least anchored). In 'phase 2' the results of the semantic and morphosyntactic analysis, entering distinct regions of Broca's area, is integrated again. ('Phase '3 is concerned with what happens when an anomaly requires reanalysis and repair.)

In a later review, following advances in fMRI scanning which have facilitated the mapping of actual white matter fibres connecting specific cortical area, Friederici & Gierhan (2013) pinpointed more than one 'ventral' and 'dorsal' streams (two of each, connecting somewhat different regions of the known speech centres), but the overall picture is the same: the language network is organized so as to facilitate the abstraction of phonemes, then words, then whole prosodically encapsulated syntactic constructions at the clause level from the acoustic signal (visual if for sequences of written words). This is done by matching input patterns with the patterns stored in long-term memory, which in turn can be projected through the anterior, motor parts of the network as speech. Perception and production are intimately bound in one tightly coherent system.

Various computer simulations of word recognition in a neural network have been carried out in recent decades by 'connectionist' modellers. One of the earliest was TRACE (McClelland & Elman 1986). This was capable of simulating basic co-articulation effects by adjacent speech segments. It operated on pre-analysed spoken word input coded for eleven acoustic phonetic features. Words were represented as patterns of activation on feature, phoneme and word 'nodes' (simulating neurons), the connections between which were updated in successive time frames in a step by step approach to matching the words contained in the input array. Phoneme units were detected and distinguished from each other by feature-level activation patterns and their sequences fed on to the word level units. Co-articulation effects were handled in a top-down fashion that allowed phoneme nodes to modulate connection strengths between phonemes and feature values in adjacent time frames. I shall have more to say on such models in Chapter 8.

In actual brains, what lies below the surface of a perceived utterance is largely tacit, accessed unconsciously and at high speed without having to 'think it through'. Much of the processing involved is simplified by shortcuts based on the probabilities accumulated through experience and 'natural' inferences accruing from them. If the meaning (extraction of which is the ultimate goal) already fits into the overall contextual mould, much can be ignored – what is essential is for deviations from it and additions to it to be registered in at least short-term memory. To be retained in long-term memory some further kind of abstraction or simplification will probably be required – the subject of Chapter 5.

2.3 Generative linguistics and abstract theory

The different kinds of linguistic abstractness seen above from a processual or functional point of view can be contrasted with another kind of abstractness, that of linguistic theory as such. Compared to the prevailing historical linguistics of the 19th century the structuralism of de Saussure and his followers can be characterized as a leap forward in theoretical abstraction. The linguistic system could now be seen in terms of internal relationships, mainly discreet binary oppositions, with the focus on 'form' as opposed to 'substance' (whether phonetic or semantic). With Hjelmslev, structuralism reached a new peak of abstraction away from concrete speech phenomena, but this was soon eclipsed by Noam Chomsky's equally abstract but terminologically more transparent transformational grammar. The continued impetus of Chomsky's theoretical approach towards further abstraction can be traced through the stages of development of the generative school that sprang from his seminal work. His own philosophical allegiance – his stand on innateness and universals – has always been to Descartes, whom I shall briefly discuss in Chapter 10. His unwavering goal has been to define what is universal behind the grammars of individual languages ever more precisely. In what follows I can only give a very superficial sketch of the successive stages of his linguistic thinking.

The initial generative stage, presented in Chomsky (1957), was largely formulated in mathematical terms and was strictly limited to syntactic phenomena. It developed into a 'standard model' following the appearance of Chomsky (1965), which posited a large number of underlying 'transformational' rules that relate invisible 'deep structure' with 'surface structure' in various syntactic configurations (mainly for English). The possibility of a distinct semantic component (which he did not develop far himself), interpreting 'meaning' directly from deep structure, now complemented the earlier purely structural model. It should be stressed that the terms 'transformation' and 'generative' are to be understood in terms of logical, quasi-mathematical relationships, not of actual on-line speech production or comprehension.[7] Chomsky's theory evolved thereafter through successive stages into a collection of sub-theories under the rubric of 'Principles and Parameters' (Chomsky 1981), of which 'Government-Binding' (GB) formed the core (and still does for many practitioners). At this stage the transformational component was simplified at the expense of a considerable complexification of underlying phrase structure. Attention was turned to incorporating phenomena from more languages than just English. The acquisition by children of different languages was now

7. The latter is the domain of 'performance' or, more recently, external 'E-language' as opposed to 'competence' or internal 'I-language', which is what Chomsky confines his interest to.

explained in terms of early choices of 'parameter settings' taken from an innate set as conforming with input from the ambient language that they are exposed to. These 'switches' determine the range of structural choices the child faces and according to Chomsky's arguments from the 'poverty of the stimulus' must be innate (Chomsky 1965: 31–33). Linguistic input is regarded as being both 'degenerate' (words blurred, not clearly separated, full of hesitations, etc.) and insufficient to allow the learning of a grammar, since children will not hear enough explicit examples of different constructions to reach the abstract syntactic principles involved.

As an example of an early transformational rule consider the 'affix-hopping' rule which would account for the position of affixes of tense and modality within the verb phrase – specifically for English. Figure 1 illustrates the 'deep structure' of the sentence "It seems that John is unhappy" expressed as a tree structure or 'phrase marker'.

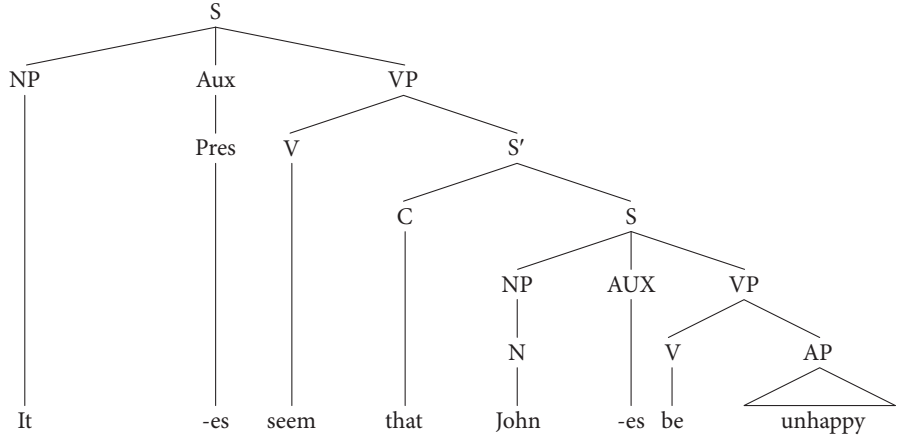

Figure 1. A transformational grammar 'deep structure'

The 'affix hopping' transformation states that if you scan a phrase marker and encounter an affix immediately followed by an auxiliary or verbal element, attach the affix to the immediately following element. Applied to this structure (and after looking up the results of combining present tense affix -es to 'seem' and 'be' respectively), you get the correct surface sentence. The abstractness here is twofold: first the transformational rule, which acts on a phrase marker generated by phrase structure rules, here S → NP+Aux+VP (twice), VP → V + S′, S′ → C+S, and VP → V + AP. (S is sentence, NP noun phrase, VP verb phrase, and AP adjective phrase.) Whether the rule as such has any cognitive reality in speaker's heads is

another matter (the link to sensory/perceptual reality is very tenuous).[8] Second, the phrase marker itself, which constitutes a Peircean 'diagram' (a species of 'icon') representing "the relations, mainly dyadic, or so regarded, of the parts of one thing by analogous relations in their own parts" (Peirce 1932: 2.277). I shall return to Peirce's theory of signs in Chapter 4. The 'parts' of the diagram here (the branching lines) represent the hierarchical organization of the syntactic units of the sentence. The abbreviations for the different constituents (e.g. 'C' for complementizer) are themselves abstractions defined by their function in the overall structure, not symbols 'in speakers' heads' – a point Jackendoff (2002: 24) stresses.

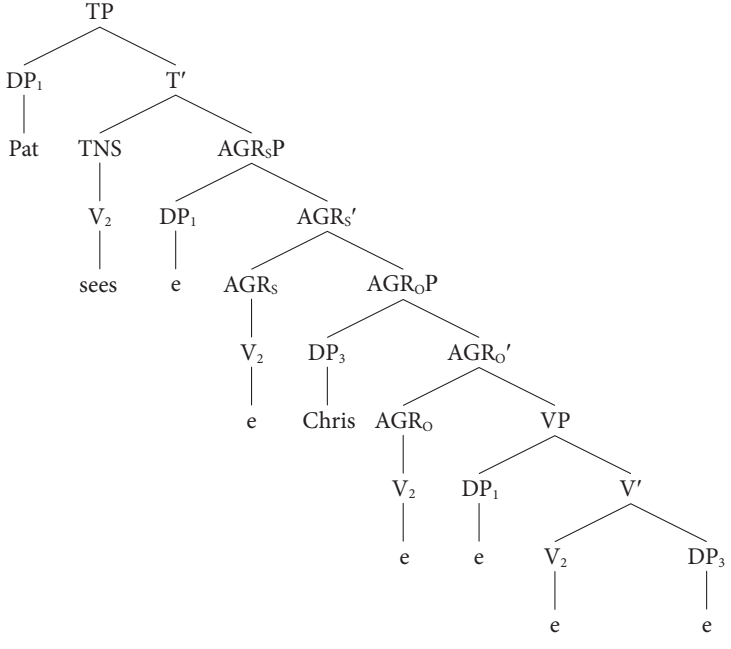

Figure 2. A 'Government-Binding' phrase marker (Van Valin 2001)

Figure 2 exemplifies the much more complex tree diagrams of the next, GB stage, here for the simple sentence "Pat sees Chris". It incorporates numerous levels of 'functional projection' like TP ('tense phrase') and various types of AGR ('agreement phrase'); 'e' indicates an 'empty' trace left by a moved constituent. All these additional levels are taken to be universal. The question as to whether they correspond to anything inside speakers' heads is deemed quite irrelevant. The theory is now unreservedly 'abstract$_1$' (as well as 'abstract$_2$').

8. Attempts to demonstrate the psychological reality of transformations rapidly foundered when no discernible correlation between transformational complexity and reaction times for comprehending sentences could be found (Harley 2008: 291–3).

Chomsky himself moved on to what he termed the 'Minimalist Program' (Chomsky 1995), perhaps in reaction to the over-complexity of underlying structural levels and movement processes to which GB had led, none of which are empirically 'provable' – not that Chomsky was particularly concerned with that. What he *was* concerned with was abstracting the most pared down, idealized aspects of his earlier theories that would explain the variety of human languages solely in terms of differences in lexicon. Syntax would now be derived from morphology. The simplified grammatical apparatus now posited was still to be taken as universal and innate. Transformations, deep structure vs. surface structure, phrase structure trees incorporating 'hidden' functional projections, and many other aspects of GB are done away with, leaving just two generalized rules ('Move' and 'Merge') operating on lexical items whose morphological and syntactic requirements need to be matched in a complex process of 'spell out' to produce paired phonetic and logical forms as output as in Figure 3.

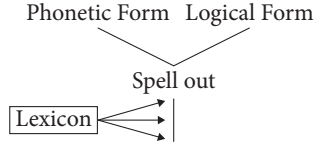

Figure 3. Basic structure of the Minimalist Program

Many of Chomsky's erstwhile most fervent followers have been unwilling to align themselves with this position of extreme theoretical simplification – in fact it brings Chomsky's thinking somewhat closer to the perspective of empirically minded cognitive linguists in so far as it attempts in effect to describe an autonomous linguistic module in the brain ('module' in the sense of Fodor 1975).[9] By opening up the possibility of a cognitive interpretation of this organization in actual processes in the brain, one might claim that the Minimalist Program is potentially somewhat less 'abstract$_1$' than its predecessors – but it is far more 'abstract$_2$', i.e. generalized, than them. The details of this sophisticated theoretical framework are not relevant here.

Generative linguistics in general has always displayed abstractness in the sense of 'simplified or generalized across instances', but at the same time it has also been 'abstract$_1$' in the sense of 'disembodied from specific sensory features' – even in the Minimalist Program the level of abstraction is still far from the sensory world. But it is not just generative theory that is abstract – so are most formalised linguistic

9. These are 'encapsulated', i.e. autonomous and innate, with their own specialized internal categories and structure. Thus the 'language module', in line with Chomsky's thinking, is not impinged upon by sensory or general-purpose problem-solving modules. Jackendoff (2002: 218–221) has a more nuanced perspective involving 'interface' modules. I shall return to this in 11.2.

theories that contain rules. Thus Simon Dik's Functional Grammar employs (at least some of the time) a quite elaborate formalism involving labelled brackets rather than tree structures. Thus in Dik (1997: 153) we find the following structural analysis of the sentence "John is easy to please":[10]

Decl E: X: Pres e_1:
{*easy* [A] (e_j: [*please* [V] (gx_1)$_{Ag}$ (Ax_j)$_{Go}$])}($d1x_j$: *John*)$_{ø}$

There is no need to go into the details of this representation (which is clearly semantic rather than syntactic) , but note that it is a layered structure, with a predication (e) contained within a proposition (X), within a 'clause' (E), and that it expresses that an entity x_j (here 'John') has the property of being 'easy' in relation to the state of affairs of anyone pleasing x_j. The essential rule this analysis presupposes is a 'predicate formation rule' for combining an adjectival predicate ('easy to please') and an embedded construction into a complex predicate of the required type. Such rules are of course abstract, but like phrase structure rules, have a rather direct relationship to surface structures.

Not all kinds of linguistic theory are abstract in the 'abstract$_1$' sense (though all involve necessarily the 'abstract$_2$' sense). Cognitive Linguistics is a case in point, one I shall focus on in Chapter 3, since unlike Chomskyan theory, which purposely ignores the deployment of language in usage to focus on structural configurations and relationships, it precisely takes usage and the relation between language and the sensory and behavioural world as central. This rival approach grew out of the Generative Semantics of the early 1970s, which attempted to replace the abstract syntactic deep structure of transformational grammar with equally abstract underlying semantic structure – with mixed success.

2.4 Abstract morphology and lexicon

Before taking the step to Cognitive Linguistics it is important to get a handle on the 'abstract$_2$' nature of language that *any* linguistic theory has to deal with in describing any or all of the world's many languages. This is best illustrated with items which in pre-theoretical terms lie, as anyone would surely agree, at the extreme end of the scale of abstractness. As regards English, these items include de-verbal and de-adjectival derivational affixes like -ity, -ation, -ment and -ness. Such morphemes are found in most if not all of the languages of the world – it is

10. This sentence is one of the examples generativists like to give of structures that are unlearnable from usage alone (owing to 'poverty of the stimulus') and require some innate knowledge of deep principles (including 'object-to-subject raising'). Functionalists disagree.

not as if the languages of 'primitive' cultures lacked such highly abstract items. Lexical items derived with these affixes are themselves amongst the most 'abstract' items in the lexicon, referring to abstract states, qualities, results, causes or actions such as 'sincerity' or 'condemnation' or 'redness'. The exact meaning will depend on the semantics of the verb or adjective they are attached to, often in unpredictable ways. Some of them (like those just mentioned) are fully lexicalized, the affixes no longer (or barely) productive today, though there is generally some leeway for neologisms (consider cases like 'pushiness'). Similar items are found in West Greenlandic (Eskimo), where for example *ilisimassuseq* 'wisdom' is derived from *ilisima-* 'know', and *naalanneq* 'obedience' from *naalag-* 'obey'. These are 'abstract$_2$' meanings in so far as all content words in any language are general 'types' applicable to multiple 'tokens' or instances. But are they also 'abstract$_1$' ('disembodied from specific sensory features')? This seems an odd thing to ask. What of such a word as 'redness' for instance? It is clearly associated with and evokes a very specific (if rather fuzzy) kind of sensory experience.

Let's leave that for a moment and consider a still more 'abstract' kind of morpheme, namely those that have a purely grammatical function, like the '-s' in English required on present tense verbs in the 3rd person singular. English would still be perfectly comprehensible if that rule was dropped (as indeed it is in pidgins). Nevertheless, the presence of this morpheme is not devoid of all meaning: it indicates precisely a 3rd person subject, and can furthermore indicate present tense (as in 'he hits' vs. 'he hit'). Its principal function is to cue the hearer or reader to look for a suitable 3rd person referent to correlate with the subject of the verb.[11] The same kind of thing can be argued for syntactic categories that on the face of it are pure abstractions in both of our senses, such as NP ('noun phrase'). The 'meaning' of this category is that here is a site for a potentially referential term. It is a functional meaning, as was that of '-s' above; both are high on the scale of 'abstract$_1$' abstractness, i.e. as regards reference to the world of sensory experience.

When we turn to lexical items, it is clear that some words have 'concrete' meanings close to the sensory world ('cow', 'house', 'hammer', etc.), and some have highly 'abstract$_1$' meanings like 'sincerity' and the other abstract derivations mentioned above, but there is a large grey area in between. There are also abstract (or rather 'abstracted') lexical elements like '-tract-' in 'abstract' that I mentioned earlier – the meaning is actually quite concrete ('draw' or 'pull') but it is not fully transparent unless one is aware of its source in Latin. Consider the following

11. Though actual searching may be dispensed with if there is just one suitable candidate held in working memory at the moment of its utterance – the brain is attuned to all manner of inferential shortcuts.

arguably abstract words not from English but from a hunter-gatherer language (Polar Eskimo), taken from Fortescue (1991):

Table 2. Abstract words in Polar Eskimo

Qualities and shortcomings: *iriarnaqtuq* 'beautiful', *naalaanngitsuq* 'disobedient', *nutaaq* 'new'
Mental faculties: *ihumaliuqtuq* 'thinks', *paahigaa* 'understands it', *ulirukkaqa* 'mistakes it (for another)'
Kinship and social relations: *iqqarliq* 'relative', *nuliaqtaaqtuq* 'gets married (man)', *atsiaq* 'someone named after a dead person'.
Beliefs: *tarniq* 'soul', *agliqtuq* 'observe taboo rules', *ilihairaa* 'casts a spell on him'

Why would we want to call them 'abstract'? Words expressing 'qualities and shortcomings' such as those given above are participial forms used for modifying head nouns. They are clearly 'abstract' in the sense of being generalized across instances, but they can hardly be characterized as abstract$_1$ in the sense 'disembodied from specific sensory features'. Recognizing an object or person as 'beautiful' or 'new', etc., presupposes a valuation of such features displayed by it. With *naalaanngitsuq* 'disobedient' the features concerned are at a higher level of behavioural traits, but they still presuppose a complex of sensory features displayed by a person acting in this manner – in this case embedded in a valuated model of 'normal' social behaviour. As regards words of mental activity, any sensory activity is internalized and 'hidden' to an outside observer, and that is what makes us think of them as abstract – but in Fortescue (2001b) I analysed the relations between words in this semantic field in a wide variety of languages in terms of metonymy, which may in turn link them to external manifestations. Thus 'knowing' something is the result of 'learning' something, and what one 'knows' one 'believes' to be true; 'remembering' something is a matter of 'experiencing again' what has previously been learnt or experienced; one 'understands' the meaning or intention of words and situations and partial understanding of something may be deepened by further 'thinking' about it (whether in words, pictures or just by 'felt' conclusions). Our innate propensity for empathy plus our understanding of behavioural externalizations of inner thoughts and feelings allows us to intuit – at least probabilistically – the private mental activity of others. There is no absolute disjunction from sensory experience here.

When we come to words (both verbal and nominal) for kinship and social relations the social nexus is to the fore, and sensory features are of lesser significance, although *nuliaqtaaqtuq* 'gets married' obviously presupposes features indicative of a male subject (the meaning is literally 'wife-get', here lexicalized). It

is a matter of sociocultural background defining what 'counts as' marriage.[12] As regards highly culturally specific words like *atsiaq* 'someone named after a dead person' there is presupposed knowledge of how to name the living individual (a social convention) and knowledge of the dead person he or she is named after (by another accepted convention). The concept is abstract but also complex. It also has specific behavioural consequences: the act of naming or calling the individual concerned. The same can be said of some of the religious/cultic concepts listed in the final category above, such as *ilihairaa* 'casts a spell on him', which combines an invisible 'spiritual' component with a physical component (the action of the witch/evil shaman in expediting the spell). Perhaps what to us seems the most abstract term of all, *tarniq* 'soul', is (or was) surely conceptualized by speakers of these languages as something almost physical and certainly 'real'.

To sum up so far, the meanings of many such 'abstract' words are clearly anchored in basic bodily actions or perceptual categories, but a good many of them are relatively divorced from such things and reflect not bodily experience as such but rather social conventions or beliefs. This in itself is not surprising. More to the point is whether the more abstract lexical meanings are somehow projected via metaphor (or metonymy) from embodied schemas in the manner claimed as general by cognitive linguistics such as Lakoff. The general position of metaphor within Cognitive Linguistics I shall return to in Chapter 4. But within that paradigm Barsalou, taking the example of 'anger', emphasizes that abstract words cannot be entirely explained in terms of metaphor. Just as indirect speech acts can 'bypass' literal interpretation, familiar metaphors (like 'explode in anger') can 'bypass' metaphorical mappings (Barsalou 1999: 600–602). For him three mechanisms relating to perception – framing (context), selectivity (or highlighting), and introspective (including proprioceptive) symbols – are all available for representing abstract concepts. He illustrates this with the abstract concept 'truth'. In fact Barsalou & Wiemer-Hastings (2005: 129) see words like 'farming' and 'cooking' in English as lying on a continuum of abstractness somewhere between the two extremes of 'concrete' and 'abstract' (in the 'abstract$_1$' sense) as regards complexity and variability of meaning. Complexity refers for them to the complex situations to which abstract words refer, and by variability they mean that abstract words can apply to a wider variety of situations than concrete ones.[13] I shall return to the

12. For what this might once have involved in this culture see the meaning of **qətəg-* 'jump' in Siberian Yupik in 4.3 below!

13. For a general discussion of complexity in language see Dahl (2004). One kind of complexity relevant to abstract$_1$ words is "algorithmic information content" whereby the complexity of an entity can be understood as the amount of information needed to re-create or specify it, identifiable with the length of the shortest possible complete description of it (op. cit.: 21).

important issue of the 'complexity' of abstract words in the following chapter, and I shall expand on Barsalou's theory of 'perceptual symbols' covering both concrete and abstract concepts in 6.2. In 10.2 I shall have more to say on the formal semantic approach to abstractness.

What of words at the extreme of abstractness, beyond even 'sincerity' and 'truth'? The verb to 'be' and its equivalents in other languages is a good case to consider, since it is both maximally general ('abstract$_2$') and very far removed from sensory reality ('abstract$_1$'). But even here there are some uses of 'be' that are more 'abstract$_1$' than others: it has besides its existential sense purely grammatical uses as a copula, also in constructions such as the progressive tense ('he is talking') or passives ('he was seen'). Despite the importance philosophers have attached to this verb in its existential sense, 'being' would seem to be so abstract a concept that it is not even treated by Wierzbicka in her 'Natural Semantic Metalanguage' as a universal 'primitive' alongside such items as 'something', 'this', 'good', 'big', 'do' or 'happen' (Wierzbicka 1992). In Simon Dik's Functional Grammar copular verbs like 'be' are introduced by grammatical rule, namely 'copula support' and do not contain independent meaning at all (Dik 1989: 165–169). Many languages such as Russian or Maltese Arabic dispense with an overt copula (in the present tense at least) in existential, equational and locational constructions, merely juxtaposing nominal or locative phrases.[14] At the other extreme, Athabaskan languages famously display a whole array of (full) verbs in a locational sense, depending on the nature (size or consistency, etc.) of the subject or object of the verb. In Chipewyan, for example, there are ten, including verbs for compact round objects, animate objects, flat and flexible objects, granular objects, objects contained in an open container, objects contained in a closed container, and stick-like objects. Other languages fall in between, such as Gaelic, with distinct verbs for existential, equative and locative 'being', and Spanish with its well-known distinction between *ser* and *estar*, the latter indicating temporary or contingent being, as opposed to more permanent *ser*. It is an open question whether metaphor is involved in the origin of copular words in general. Buck gives as the source of the various Indo-European verbs a number of different etymological sources: besides 'colourless' **es*-, and **bheu*- (originally 'become'), also 'remain, stay', 'stand' (as in 'exist' itself), and 'sit' (Buck 1988: 635–6), some of which are clearly metaphorically extended. In the following chapter we shall see how 'BE' is treated as a universal semantic primitive in the work of Jackendoff.

14. They may have separate full verbs meaning 'exist' or 'be positioned', but these tend to be more stylistically marked variants adding nothing to the basic meaning. Compare also Lakoff (1987: 462–470) on deictic as opposed to existential uses of the 'there is' construction in English.

We can agree, I think, that the verb to 'be' is abstract in both our senses, and yet in the existential sense it also covers 'be alive' (the opposite of Hamlet's 'not to be'), which is a state perfectly evident to the senses. The overall conclusion so far should by now be evident: there is no clear cut-off point between 'concrete' and 'abstract' meanings, only a continuum from more to less 'abstract$_1$' (disembodied from sensory reality), cross-cut by widely differing degrees of generality or complexity. The complexity is a matter of the situational context which may range from culturally specific to universally human. 'Being' is maximally general but not particularly complex, whereas the Polar Eskimo for 'casting a spell', for example, is culturally specific, referring to a complex belief system of quite indeterminate generality.

2.5 Logograms and 'abstract' sign language

Finally, before leaving this matter of abstract$_1$ words, it will be instructive to look at the orthographic representation of such words in a logogramic language like Chinese. These are in a sense doubly abstract since all logograms (which represent meaning in pictorial schemas of varying degrees of transparency) contain both pictogramic abstraction and semantic abstraction. Thus the word best corresponding to 'be' in Indo-European languages is (in Mandarin Chinese *pinyin* form) *shì*, written [是]. This is glossed by Wieger (1965: 267), the standard work on the etymology of Chinese characters, as 'truthfulness, reality, existence' and constitutes a ubiquitous existential verb 'be' in the modern language alongside a secondary (retained) meaning of 'right, correct'. The logogram consists of [正] *zhèng* 'right, correct' and [日] *rì* 'sun, light' and is explained as 'what was controlled at sun's light'. The 'sun' part is a pictogram of the sun rising over the horizon, rendered square rather than round due to a historical shift in writing instrument. The first part consists in turn of 'stop' and 'line, limit', so 'stop at the limit, where one had to reach, without going astray'. Another extension of this part is [定] *dìng* 'order, decide, definite, stable', which combines it with a 'primitive' semantic element meaning 'house', so 'order in the house, tranquility'. Wieger points to another logogram sharing the meaning of 'right' with existential *shì*, namely [直] *zhí* 'perfectly right, straight', explained (op. cit.: 36) as 'ten eyes having looked at s.th. did not discover any deviation'. The latter, in turn, is combined with the 'heart' logogram *xīn* (reflecting in its lobed shape its pictogramic origin) and the 'radical' meaning 'take a step, walk' (here 'conduct') to produce [德] *dé* 'rectitude of heart, virtue, moral conduct directed by a righteous heart'. So the result is a thoroughly abstract meaning (one requiring a whole body of doctrinal elaboration to understand fully), but one reflecting its pictogramic origin in rather concrete imagery.

Let's take just a few more examples. Wieger (op. cit.: 362) explains [物] *wù* 'thing', matter' as consisting of a left-most phonetic element (roughly indicating the pronunciation) and a semantic element meaning 'ox', because 'the ox is the largest of things' (it was certainly among the most valuable possessions a man could have in ancient times, as Wieger comments). Another *wu* (with rising tone), *wú* 'want, defect, negation' [無], is today a general negative marker 'have not, be not, nothing, void' (op. cit.: 36). It apparently indicated originally a multitude of men acting upon a forest, clearing wood from a tract of land.[15]

Wieger (op. cit.: 179) further gives [義] *yì* (his i^4) 'harmony' as consisting of 'good' (originally indicating a sheep, a peaceful creature), peace restored after a 'conflict' – the latter consisting of two weapons in conflict.[16] Another word for 'harmony', 'agreement' or 'union' is *hé* [合] (Wieger's ho^2), explained (op. cit.: 45) as consisting of elements meaning 'union, assemblage, many' (symbolized by three lines forming a triangle) and 'mouth', i.e. 'many mouths speaking together, good understanding'. The ubiquitous modern word for 'good, love' is *hǎo* [好], transparently consisting of 'woman' and 'child'. Older forms of the logogram were clearly pictogramic, the form for woman on the left, for example, showing the woman (headless) with her arms folded in front in correct ritual posture (Wieger op. cit.: 169).

Two more words indicating somewhat different kinds of abstractness, temporal and epistemic respectively, are [世] *shì* 'an age, generation, duration of a life' and [必] *pí* 'certainly, necessarily'. The former indicated originally a period of thirty years (the duration of a man's active life) by extending the left-most vertical line of 'thirty' (three times 'ten') to express 'the prolongation and duration of life' (Wieger op. cit.: 71). As for *pí,* this looks today like the heart symbol given above crossed by a curving line, but according to Wieger (op. cit.: 59) it originally represented an arrow dividing a doubt or dilemma, "a kind of interjection pointing out a strict order".

What kind of abstraction do such logograms contain? As linguistic symbols they are just as abstract as English written words, but to some degree they reflect their meaning, which a strictly alphabetic script does not. Perhaps they can be characterized as 'semi-abstract', in so far as their pictogramic origins still show through to some degree. Clearly they reflect some kind of metaphoric or metonymic basis but (apart from cases like is *hǎo* above) they are generally far from transparent and sometimes (like with *zhí* 'perfectly right, straight') involve a whole

15. The original form once contained a further element that indicated entering into a hiding place, hence the meaning 'vanish'.

16. Interestingly enough this is now the 1st person pronoun 'I' *wǒ,* a meaning developed via 'two opposing rights' and then 'my right'.

imaginary – and not very likely – perceptual scenario which it is unlikely that contemporary readers would recognize. In fact once suspects that folk (i.e. scholarly) etymology lie behind certain traditional explanations, and this has certainly resulted sometimes in orthographic changes that reflect such reanalyses (as with *pí* above). The most one can say with certainty is that logograms are schematic simplifications of pictogramic representations from which they are abstracted to various degrees.

A somewhat similar phenomenon is characteristic of sign languages: although the majority of signs are either iconic (in Peirce's sense) or arbitrary symbols, some seem to fall somewhere in between, being both abstract and partially iconic. Take for example the sign meaning 'EXPLAIN' in Danish Sign Language discussed by Engberg-Pedersen 1996: 43). As can be seen in Figure 4, the signer starts with the splayed fingers of both hands almost touching and finishes with them opened out in the direction of the recipient (at a certain 'locus'). The sign is repeated, indicating a durative action. (This is itself iconic but also abstract in the same sense as an English morphological affix is abstract.) The direction of the motion would have been the reverse if the speaker were the recipient of the action. The enclosing position of the fingers at the outset indicate something that is being held within, out of sight to the recipient – here an explanation. The movement itself, with fingers ending up pointing towards the recipient is shared by a number of verb signs indicating transferral. Here what is being transferred is explanatory information, which is a rather abstract notion in its own right. Whether this counts as the sub-type

Figure 4. The initial and final position of the hands in the sign EXPLAIN in Danish Sign Language (model Eva Abildgaard)

of icon that Peirce calls a 'diagram' (based on analogy rather than direct imagery) is an open question. I shall return in 10.34 to Peirce's distinctions. The EXPLAIN sign is at all events certainly 'schematic'. In the following chapter I shall look more closely at the nature of 'schemas' generally.

CHAPTER 3

Cognitive linguistics

Language embodied

3.1 Image schemas – abstract or not?

The sense of 'simplified or generalized across instances' ('abstract$_2$') is specifically applicable within Cognitive Linguistics (CL) to the 'image schema', literally a schematic version of a sensory image, a representation of 'embodied' experience. This covers such things as our knowledge of containers, paths and forces specific to the various sensory domains that structure our bodily experience. One of the foundational conceptions of Cognitive Linguistics and its embodied approach to meaning is the notion of the metaphorical projection from bodily image schemas (as source) to more abstract domains (as target). Croft and Cruse (2004: 44) have pointed out the ambiguity of the term 'abstract' when applied to image schemas: they are abstract in one sense ('schematic') but not in another (they are 'embodied' rather than 'disembodied').

One of the chief architects of CL, Langacker with his 'cognitive grammar', sees grammar as corresponding closely to conceptual structure, so the productive combination of words and parts of speech correspond to the productive combination of perceptual elements such as processes and properties. For Langacker 'abstraction' is one of several basic cognitive processes[1] – its function is 'for conceiving of situations with varying degrees of specificity and detail' (Langacker 2000: 2). Note that this preliminary definition – though useful – does not distinguish in any obvious way between 'abstract$_1$' vs. 'abstract$_2$'. For him syntactic rules, for instance, are simply the most highly abstracted/ generalized construction schemas, those that are most firmly 'entrenched' through repetition. They are clearly 'abstract$_2$', but they compete with lower level, more lexically specified constructions instantiating them – these are probably more directly accessible through association with specific event structures (Langacker 1998: 25), and thus are lower on the 'abstract$_1$' scale than the more general schemas. Langacker focusses on a particular

1. These further include the directing and focusing of attention, the imposition of figure/ground organization, the capacity for mental scanning and the creation of abstract 'things' by conceptual reification, all presumably innate (Langacker (2000:171).

process of abstraction, namely 'schematization'. Let me cite him at greater length (Langacker 2000: 93):

> …*abstraction* is the emergence of a structure through reinforcement of the commonality inherent in simple experiences. By its very nature, this abstractive process "filters out" those facets of the individual experiences which do not recur. We will be most concerned with a special case of abstraction, namely *schematization*, involving our capacity to operate at varying degrees of 'granularity' (or 'resolution'). Structures that appear very different when examined in fine-grained detail may nonetheless be quite comparable in a coarse-grained view. A *schema* is the commonality that emerges from distinct structures when one abstracts away from their points of reference by portraying them with less precision and specificity.

A 'schema' here should not be too literally taken as 'embodied' in any specific sensory modality – the 'point of reference' may already lie at a fairly abstract$_1$ level. Thus Langacker discusses the abstraction of phonological segments from syllables, and of morphemes from words within a 'usage-based' framework (as in Bybee 2010, which I return to in Chapter 9). Recurrent pairings between particular conceptual and phonological structures are recognized in this way, for example the -er affix, which occurs in words like 'propeller', 'ruler' and 'stretcher'. He further illustrates how the abstract consonantal templates in Semitic languages like Aramaic can be abstracted from specific inflected instances of words (Langacker 2000: 132–134). Note that this is not exactly what I mean by emergent abstraction (to which I shall return in Chapters 8 and 9) – if the language concerned is already known to the hearer this is more like the 'abstraction by matching' mentioned above in connection with the 'double articulation of language' since the patterns abstracted must already be 'there' at some level of memory storage. But Langacker is talking here as a linguist rather than a hearer, abstracting successively schematized levels of analysis. Only if one is talking of the *initial* formation of schemas would the process actually be emergent.

The strong embodiment position within CL, as represented by Lakoff (1987), is that all linguistic meaning is ultimately anchored in simple physical, sensorimotor schemas ('abstract$_2$' and at the same time experientially concrete) – it is extendable from these by metaphor into more abstract$_1$ realms. As I claimed in the preceding chapter, words that can be called abstract$_1$ are not necessarily simple, in fact their meaning is often much more complicated than that of concrete words in so far as they refer to a wide range of contexts and scenarios of a social, moral or introspective nature (e.g. 'marriage'). These are by no means 'simpler' than more basic perceptual notions even when peripheral connotations are stripped away. They may require lengthy definitions of a culturally determined nature.

Barsalou & Wiemer-Hastings (2005: 150), from within the CL paradigm, point out that the conceptual content of abstract concepts is indeed more complex than the content of concrete concepts. Focusing on the situational embedding of abstract words they attempt in effect to reconcile – or at least to correlate – the two senses of 'abstract' above. They view word meanings in general (concepts) as large collections of situational representations rather than as being organized around simplified prototypes ('core knowledge'): "Learning a concept establishes the ability to completely represent a wide variety of situations relevant to interacting with the concept's instances" (op. cit.: 156). In their word-association investigation of abstract concepts, they found that abstract words were more closely associated with words expressing social aspects of situations such as people, communication and social institutions, as well as introspection and belief, than concrete ones. Associations were also more hierarchically complex, with greater depth of nested clusters of categories. Metaphorical extension from more concrete concepts is not necessarily involved (though it may play a role), and 'abstract' concepts are not necessarily divorced from experience – we have direct experience of such introspective concepts as 'anger', for example, which is directly 'simulatable'. Among the concepts for complex social and cultural phenomena that they mention is 'postmodernism' (op. cit.: 159), which presupposes verbal definition as well as potential situations to which the word might appropriately be applied. Such words are surely beyond the limit of 'embodiment'. In fact the authors readily admit that their study says little about embodiment.

Other researchers within the CL paradigm have discussed still more abstract concepts in terms of embodiment. Prinz (2005) discusses moral concepts as grounded in emotion ('good' vs. 'bad'). Talmy (2000) has discussed in considerable detail under the rubric of 'Force Dynamics' the way whereby abstract relations of causality can be projected from basic spatiotemporal schemas of oppositions between 'agonistic' and 'antagonistic' physical forces. He has also analysed locative prepositions such as 'across' as involving abstraction within the embodiment framework (Talmy 1996). Langacker sees grammatical rules too as 'extracted regularities with some kind of enduring cognitive presence', templatic schemas immanent in their token instantiations (Langacker 2000: 144), and this also characterizes Construction Grammar within the broader CL paradigm (cf. Goldberg 1995). Zwaan & Madden (2005: 240–41) suggest the theoretical possibility of an embodiment treatment even of such highly abstract notions as negation. Then there are abstract concepts like 'time', 'location', and 'cause', of which we surely have direct – if diffuse – 'embodied' experience, just as we have of 'redness' and 'triangle'. It could be argued in fact from this point of view that all type concepts like 'father', or event concepts such as 'flood' or 'storm', are also abstract generalizations,

though the tokens they summate and refer to are complex, not simple concrete entities or phenomena.

3.2 Lexical concepts and conceptual schemas

How does this relate to my distinction between 'abstract$_1$' and 'abstract$_2$'? If all content words are abstract in the sense of being 'simplified or generalized across instances' and abstractness in the sense of 'disembodied from specific sensory features' is always going to remain a matter of degree, does maintaining the distinction gain us any mileage at all when considering lexical meaning? Perhaps so, if it helps us decide what is truly embodied and what is not. One useful way of looking at this is through the distinction Evans (2009) makes between 'lexical concepts' and 'cognitive models' in his LCCM ('Lexical Concepts and Cognitive Models') variant of CL theory. The former are defined as "bundles of various types of schematic knowledge conventionally associated with a unique phonological vehicle in a symbolic unit" (Evans op. cit.: 349). They can be 'open-class' (e.g. the content of the words 'waiter', 'serve', or 'customer') or 'close-class' (e.g. (in)definite articles, plural markers, lexical class, or specific grammatical relations – op. cit.: 103). 'Cognitive models' are on the other hand "a coherent body of multimodal knowledge of any kind to which a lexical concept can facilitate access and which can give rise to a simulation." They consist of 'frames' relating to individuals, types, or situations (op. cit.: 345). The cognitive models associated with words – especially those that can be considered to some degree abstract$_1$ – are complex and diffuse, drawing upon episodic as well as general semantic memory. Lexical concepts are underspecified for their situated interpretation, in agreement with Langacker's (1987) notion of 'schematic' as opposed to 'rich' content – they are abstractions over multimodal content and do not directly encode multimodal experience themselves but provide access to the far richer cognitive models that do. Individual words corresponding to lexical concepts (their 'vehicles'), deprived of a specific context in which their meaning can be 'simulated', are thus all in a sense abstract$_1$ (i.e. distanced from sensory experience) as well as abstract$_2$ (i.e. schematized). Context and content as such, however, is a matter of cognitive models, which are in principle abstract$_2$ (i.e. schematized) but not abstract$_1$.

Evans' 'lexical concepts' correspond to bundles of what I called in Fortescue (2009) 'micro-functional affordances'. (They are 'affordances' in the sense of being higher-level clues to interpretation contained within the speech stream and to which the hearer is sensitive.) These adhere to individual phonological words as relevant to the grammar – both paradigmatically as regards categories and syntagmatically in terms of their combinatoriality. They may include (for action

words) schematic event and/or force dynamic features and correspond generally to Levelt's 'lemmas' – the meaning plus syntax of lexemes minus their phonological form (Levelt 1993:188–9).[2] 'Cognitive models' on the other hand corresponded in my model to the diffuse 'macro-functional affordances' associated with a word (scenarios, situations, frames, etc. – I shall return to this in 8.4). Since this is a matter of accumulated (complex) situations and contexts, they are *a priori* going to be more complex than 'micro-functional' lexical concepts. Much of what is pre-theoretically conceived of as a difference between 'concrete' and 'abstract' words boils down to a difference in complexity of their semantics. However, this is a very unreliable distinction, since some 'concrete' words refer to very complex objects (like 'automobile'), and some abstract words are rather simple in essence (if extremely vague), like 'be' discussed in 2.4.

For Evans, lexical concepts can be simple or complex (op. cit.: 97–98), and for a lexical concept to be 'complex' means that it may have part-whole organization. LCCM incorporates a construction grammar approach to abstractness and complexity, in other words it operates with 'symbolic units' ranging from single words to completely general syntactic constructions, and 'internally complex lexical concepts' are typically complex constructions, such as 'THING X CAUSES THING Y TO RECEIVE THING Z', which is associated with the linguistic 'vehicle' NP FINITE VERB NP NP (the 'ditransitive' construction', which could be realized as, for instance, 'The man gave the book to the girl'). As an example of an 'internally simple' concept Evans gives FRANCE associated with the 'vehicle' 'France' (the word). All the complexity adhering to that word is a matter of the associated 'cognitive model' rather. Croft with his Radical Construction Grammar presents a taxonomy of symbolic units ranging in complexity from 'complex and schematic' syntactic constructions, through 'complex and specific' idioms, 'complex but bound' morphology, 'atomic and schematic' word classes, to 'atomic and specific' lexical items (Croft 2002: 17). This approach completely eliminates grammatical rules and categories, reducing grammar and lexicon to part-whole relationships between symbolic units.

Most adherents of Chomsky's generative theory (apart from the rebellious 'generative semanticists' who split away from him)[3] relegated semantics to a secondary position within linguistics. Jackendoff, however, has developed with his 'conceptual semantics' an approach to semantic matters relevant to both generativists and functionalist/ cognitivists alike. He points out in connection with the verb 'be' that

2. In his revised model there is in fact a level of lexical concepts added between communicative intentions and lexical lemmas.

3. Several of whom, including George Lakoff, later becoming key promulgators of Cognitive Linguistics.

there is a parallelism (albeit highly restricted) between sentences containing the same words in a concrete spatial sense of location and motion and extensions to the more abstract domains of possession, the ascription of properties, and scheduling activities (Jackendoff 2002: 356–9). Thus "The messenger is in Istanbul" (location) compared with "The money is Fred's" (possession), "The light is red" (ascription), and "The meeting is on Monday" (scheduling). Similarly with the verb 'go' for change of location and 'keep' for 'caused stasis', thus "The light went from green to red" (ascription). Unlike cognitive linguists, he sees this not so much in terms of metaphorical projection from a basic meaning, but rather as representing "precise abstract underlying conceptual patterns that can be applied to many different semantic fields… the basic machinery that permits complex thought to be formulated and basic entailments to be derived in any domain". These 'underlying conceptual patterns' involve semantic primitives that correspond to the elements making up Evans' 'lexical concepts'. In both frameworks the analysis is in terms of underlying abstractions like 'location' generalized over different domains as well as situation types. Jackendoff specifically represents this by BE(X,Y), which is "of the ontological category *State*: it is the conceptualization of a static configuration that can be localized at a point in time or throughout an interval of time", where X is an object and Y a location (op. cit.: 360–1). It can be applied to different domains as described above. Recall the discussion of the 'abstract' surface verb 'be' in English in the preceding chapter: according to Jackendoff BE(X,Y) is the underlying conceptual abstraction beneath this virtually 'empty' surface item.

Another central concern of CL is polysemy and prototype theory, as first fully developed by Rosch (1978). A prototype is an abstraction that is the best example of a category (for example, a typical 'bird' in our culture might be the blackbird). Category boundaries are fuzzy and members of it are identified by how closely they resemble the prototype 'average' along a number of parameters (e.g. 'flies', 'has feathers and wings', etc.). Rosch further argues that there is generally a 'basic level' of categorization, a default which is most salient psychologically, by being the most distinctive and economic and indeed the most easily imaginable (thus 'table' rather than superordinate 'furniture', or particular lower level types of table). This is the level at which children first learn to start using words. A prototype is a special kind of schema, and Lakoff has identified a number of different schema types that show prototype effects. One, the so-called 'radial category', he illustrates with the word 'mother', which besides its 'central' meaning can refer to a number of special, culturally specific instances such as 'biological mother' or 'surrogate mother' (Lakoff 1987: 83–84). These are conventional, culturally specific, and cannot be predicted by rule.

One problem with this approach is that it is difficult to fit abstract$_1$ words like 'truth' into such a category. Moreover, typicality judgements vary systematically

depending on context. This is precisely the starting point for Barsalou & Wiemer-Hastings (2005) in their attempt to fit abstract$_1$ words into CL theory. In general, the complexity of such words is more to do with their widespread applicability in somewhat differing senses according to specific situation or context types, including 'folk theoretical' models. This would appear to have little to do with 'underlying conceptual primitives' (which is more a matter of 'abstract$_2$' schemas).

In fact the perspective I shall adopt on the nature of abstract lexicon and 'embodiment' is not the 'strong embodiment' of Lakoff and colleagues but the 'embodied abstraction' perspective of Binder & Desai (2011), which is based on evidence indicating that the brain – and in particular semantic memory – operates at different levels of abstraction, from the modally fully 'embodied' level of individual sensory and emotional modalities through modal convergence zones (correlating two modalities) up to completely amodal or supramodal zones. More of this in subsequent chapters.

3.3 Compositionality and blending

Let us push the question of the compositionality of abstract underlying word meanings ('lexical concepts') a little further: to what degree (if any) can the meaning of words in fact be reduced to a handful of universal semantic ingredients? Take as an example the verb 'remain' from the Cheshire Cat dialogue ("…*the grin, which remained some time*…"). Its meaning is analysed by Jackendoff (1985: 172) as:

$$[_{\text{Event}}\ \text{STAY}\ ([_{\text{Thing}}\ x], [_{\text{Place}}\ y])]$$

In this context the 'thing' would be the cat's grin, the static 'event' the remaining after the rest of the body had gone, and the 'place' where the cat happened to be at that time (on a bough). The universal semantic element is STAY, an abstract meaning that underlies other words like 'persist' and 'persevere' and 'keep' besides 'stay' itself. The whole point of the analysis being that STAY has a number of features (shared with other surface words) that allow us to infer other semantic features of the word 'remain', for instance that it refers to a state and thus is compatible with such expressions as 'for some time' but not others like 'at 2.30 precisely'. It also presupposes that 'thing' x is at 'place' y already. There is good reason to think that this particular highly general meaning (STAY) is universal (all languages surely have some means of expressing it) and that it is also 'embodied' in sensory experience, but is it exhaustive? 'Remaining' surely means a bit more than just continuing to be at place x for some indefinite time. At the very least there is a stylistic difference between it and 'staying'. Jackendoff would call such a language-specific

difference a 'completer', a secondary or 'typicality' feature rather than a necessary one (Jackendoff 2002: 349f.).[4]

As a more complex example compare the following analysis of "Amy kept the doll" (op. cit.:192), where 'Poss' indicates the field of possession (as opposed to a spatial or temporal field):

$$[\text{CAUSE}\ ([\text{AMY}],\ [\text{STAY}_{\text{Poss}}\ ([\text{DOLL}],\ [_{\text{Place}}\ \text{AT}_{\text{Poss}}\ ([\text{AMY}])])])]$$

Levinson (2003: 297ff) describes two extreme positions on compositionality, both associated with the generative linguistic camp (not his own), namely the holistic perspective advocated by Fodor et al. (1975), which denies the possibility of decomposition entirely, and the more atomistic perspective of Jackendoff, which allows all words – in principle – to be decomposable. A more thorough-going version of this position is represented by Wierzbicka with her theory of universal semantic primitives (e.g. Wierzbicka 1992). Actually Jackendoff's position is one of limited (non-exhaustive) compositionality, not really so different from Levinson's.

Levinson's own position (which I share), lies between these extremes. As he puts it (op. cit.: 298): "There is a way to have both the non-compositional cake and to eat it with compositional relish. We simply need to assume that there is a level where the decomposition can and has taken place, and another level at which we habitually operate – where we cease to decompose." He proposes a 'dual level' model, which combines a 'molar' and an 'atomic' level, both of them cognitively real. The 'molar' level provides unitary meanings of lexemes for cognitive processing, while the 'atomic' level provides decompositional meanings for learning purposes. The latter level is as least partially universal and can presumably be related to pre-linguistic 'image schemas' (though Levinson is wary about equating semantic and conceptual structure directly), whereas the 'molar' level is culturally relative. This perspective is nevertheless germane to the CL one and is close to the position of Langacker on 'partial compositionality' (2000: 152), also to Evans' LCCM approach to compositionality (2009: 23–4). In fact most cognitive linguists today appear to accept the view that there is considerable redundancy in the organization of the mental lexicon, but that some compositionality must nevertheless be allowed for. It provides the basis for the almost unlimited potential of paraphrase and translation. Words at any level of complexity (apart from function words) are associated with their own conceptual content and defining contexts,

4. For a somewhat sceptical view on the usefulness of the notion 'completer' see Fortescue (2010a:188ff) where I considered the case of the words 'open' and 'door' and their collocation 'open (a) door'. The kinetic 'schema' underlying opening a door is far removed from the schemas relevant to the opening of a book or a letter or one's eyes, for instance. There is a 'family resemblance' here, and it is difficult to see this in terms of some feature added to a core meaning of 'open' when it is used in a specific context like these – the context itself is the 'completer'.

and these may arguably be analysed in terms of their own quasi-universal, perceptually based primitives. They are just not necessarily the primitives of language, which has its own, grammatically relevant combinatorial categories

Compositionality – in the broad sense of the term – not only pervades language at all levels (phonological, morphological, lexical and syntactic), it also plays a central role in imagination and dream imagery, as we shall see in Chapter 6. Within the CL purview, the work of Fauconnier and colleagues on Mental Space theory represents a somewhat more 'analog' approach than the 'digital' approaches to compositionality of Jackendoff and Wierzbicka. A mental space is a kind of 'projected' context, originally applied to counterfactual, belief and 'want' constructions in which presuppositions can be 'floated' from a 'base' space to some other, projected mental space. In more recent blending theory the core relationships of more than one input spaces can be combined in a novel 'blended' space. This covers the compounding of nouns, whole sentences, and novel metaphors.[5] As an example of the latter consider a statement like "Washington is the father of our country", where two input spaces – one for the generic father-son relationship, and one for the founder-nation relationship (here between Washington and the USA) – are involved in producing the novel blend (Fauconnier & Turner 2002: 141). In fact Fauconnier & Turner see blending as pervading all human reasoning.

But what has this to do with abstraction? Well, as regards semantic primitives, these can be viewed in two ways: whereas both generativists and cognitive linguists would doubtless agree that these are 'abstract$_2$' ('simplified or generalized across instances'), they would disagree as to whether they are also completely abstract in the 'abstract$_1$' sense ('disembodied from specific sensory features'). Whereas generativists who believe in an 'encapsulated' language module must by necessity regard them as abstract also in this sense (or as at best lying at the interface with other modules), this is not so with cognitivists who see the meaning of words (apart from purely grammatical 'function words') as ultimately anchored in bodily 'image schemas', albeit projected from them via metaphor. As regards 'blending' theory, this depends on a process of 'compression', which is a kind of simplification that abstracts the core relationships involved in the input spaces concerned. It brings to the forefront the question of context, which in CL is mainly handled in terms of 'frames', a matter I shall return to in the next chapter.

We might summarize the perspective of cognitive linguists on abstractness by stating that while they embrace schematic 'abstract$_2$' abstraction (as in Langacker's definition above), they eschew the 'abstract$_1$' kind characterizing Chomskyan Universal Grammar. Though problems remain for them as regards how exactly

5. For an example of the compositionality of novel noun compounds such as 'jar lid' see Langacker (2000: 109–110). He treats this in terms of 'constructional schemas'.

to take 'abstract$_1$' words on board, their usage-based approach, which takes context of use seriously and relates directly to the way in which children learn such words, offers a more realistic approach to the problem than any hand-waving about 'innateness' (more of this in Chapter 9). Moreover, their perspective on image schemas and 'deep' metaphor brings us closer to understanding the interface between sensorimotor imagery and linguistic expression, which will be the main subject matter of Chapters 5 and 6. In later chapters we will eventually reach a perspective from which all the different kinds of abstraction we have touched upon so far will be seen to be related within a broad procedural framework. This will necessitate excursions beyond language as such, in particular into neuroscience and philosophy.

CHAPTER 4

Abstraction, metaphor and cultural context

4.1 Metaphor and analogy

At this point we need to look at the relationship between abstraction, metaphor and context more closely. Metaphor (along with metonymy) is a key issue in Cognitive Linguistics. Supposedly universal 'conceptual metaphors' such as 'UP IS HAPPY' or 'HEAVY IS DIFFICULT' proposed by Lakoff and his associates are in themselves prime examples of abstraction from experience which may be projected upon individual languages. To what extent such things are actually universal is an empirical question. In this chapter I shall contrast two language families which differ greatly in the degree to which they employ metaphor, either overtly or 'concealed' in their basic lexicons. The first is Eskimo-Aleut, which displays very little obvious metaphor, and the second is Na-Dene, which displays a great deal. This is somewhat at odds with the opinion of cognitive linguists who predict that all languages are more or less equally suffused with metaphors based on basic physiological experience (cf. for instance Lakoff 1987: 407–8 on metaphors of anger). My own position is closer to the cultural relativism of Gumperz & Levinson (1996). Actually some researchers working close to the Cognitive Linguistics paradigm also share this position. Thus Sinha & Jensen de López (2000) argue for 'extended embodiment' to include cultural experiences. This deemphasizing of bodily image schemas and focusing on social situations may be the sensible way to reconcile universal with language-specific metaphors in a manner that does not presuppose deriving the one from the other.

Some abstract$_1$ terms are difficult to see as 'embodied' in the sense of Lakoff & Johnson (1980), depending as they do on complex social relations and belief systems rather than on metaphorical extension from image schemas. The degree of abstractness and the degree of complexity of such terms is not the same thing, however. Thus the Japanese word *giri,* roughly 'duty, obligation', is both more specific (less general) and more semantically complex than these English terms, which are highly abstract and general but not as complex as the Japanese one, around which whole books have been written (cf. for example Benedict 1947). Essentially it involves serving one's superior with a self-sacrificing devotion, but is manifest in very specific forms of conduct within a highly complex social network. In a single language abstract hypernyms may cover a number of more complex but

they express are not – the dimensions of metaphoricity and abstractness here are orthogonal to each other.

An influential distinction within the metaphor theory of cognitive linguists has been that between primary (universal) and compound (language-specific) metaphor, as discussed in some detail by Grady (2005). Primary metaphors are generic patterns or schemas rather than fleshed-out instantiations. Compound metaphors are further language-specific elaborations of primary ones. Doubt, however, has been cast on the cross-linguistic validity of the claim of universality for primary metaphor by Sinha et al. (2011), specifically as regards the 'TIME-from-SPACE' relationship. Grady himself (op. cit.: 1610–11) concedes that primary metaphors may not be entirely universal, only 'likely' across languages. For him the source-target relationship in metaphors is less one of abstraction than of the greater subjectivity of the latter (op. cit.: 1607–8). The source is close to the physical world as input to the senses and the target is the correlation of the source configuration with subjective experience, e.g. of a heavy object with the experience of difficulty moving it as regards the conceptual metaphor 'HEAVY IS DIFFICULT'. This is the basis for the blend that is the primary metaphor which may be further elaborated in language-specific compound metaphors. However, the universality of primary metaphors can be questioned. Circularity lurks: the temptation is to declare any 'deep' conceptual metaphor reflected in a given language as evidence for some universal schema.

Thus there is not much evidence of either kind of metaphor in at least one large linguistic family with which I am familiar, namely Eskimo-Aleut (EA), where metaphor in 'normal' (non-shamanic/ esoteric) usage is limited to a handful of frozen and culturally rather specific ones. In the following two sections I shall contrast the situation in EA with that in neighbouring Northern Athabaskan languages (NA), where there is indeed a good deal of language-specific metaphor – although it is not at all clear that this is compounded or elaborated in the majority of cases from universal primary metaphors. Both language families are (or were until recent times) spoken by hunter-gatherer communities, so it can hardly be the way of life as such that determines the degree to which metaphor is exploited by a given language. In both cases one is led to wonder whether a language can display 'compound' metaphors without the 'primary' ones they are purported to arise from.

4.2 Metaphor amongst the Inuit

As regards Eskimo-Aleut, I shall limit myself largely to discussing West Greenlandic (WG), though what I have to say applies equally well to other EA dialects. It is generally regarded as a particularly 'concrete' language (compared say to Danish),

although as we saw in 2.4 Greenlandic dialects do have a good many abstract$_1$ words, both lexical stems and derivations – it is really just its sparsity of metaphor that lies behind such judgements. I shall not say anything about 'poetic' metaphor that has come into the language in recent times through direct glossing of Danish metaphors. Greenlanders are aware of some of these influences and may joke about them – e.g. the use of *tigu-* 'take' (literal only) in equivalents of typical Danish expressions for 'taking a taxi', 'taking one's holidays', etc. The fact is that West Greenlandic has little in the way of native idioms – i.e. conventionalized non-literal or metaphorical expressions. Nevertheless, one can point out some extensions of terms for basic physical reality to more abstract domains. Thus spatial (directional/locative) expressions may have temporal usage – a very common phenomenon amongst the world's languages, e.g. stem *siu-* 'in front' which also means 'before' in time. Ancestors are also *siuaasat* (lit. 'ones who came first'), vs. *kinguaat* 'descendants' (lit. 'ones who come behind/after'). The allative case form *siumut* 'forward' can be extended (in literary WG at least) to progress in the attainment of social/political and personal goals, but this may well be under Danish influence. An isolated 'heat' metaphor (HEAT IS ANGER) can be seen in *uumigaa* 'hates him', from Proto-Eskimo **ugu-* 'be heated up' (via **ugumi-* 'be infuriated'). All such forms are to be found in Fortescue et al. (2010).

Also the up/down dimension, so prominent in Lakoff & Johnson's approach to metaphor, is represented in a few expressions, e.g. *sursunnirsuup ataani*, 'during (literally under) the war', *qummukalirpuq* 'things are looking up', literally 'it is going upwards', and *ammut isigaa* 'he disdains him' (lit. 'looks down at him') – all probably glossed from the Danish ('*under krigen*', '*ser ned på ham*' and '*det går opad*' respectively). I am not aware of these extensions in other forms of EA. Some general verbs of motion such as *ingirla-* 'go' may occasionally be found with time expressions as subject (e.g. a nominal meaning 'day' or 'one's life'). This may just reflect the liberality with which WG allows inanimate subjects of verbs in general, including in transitive meteorological expressions like *anurliuppaa* 'he was overtaken by a gale' (from intransitive *anurlir-* 'blow (wind)') in which the unexpressed ergative subject is understood vaguely as 'the weather'. At all events there is no whole-scale extension of the meaning of verbs of motion beyond their basic spatial sense.

There is no doubt that metaphorical extensions do occur in EA languages then (in particular the 'time as space' one) , but it is much less in evidence than in European – or indeed some other native American – languages, as will be seen below in the case of Athabaskan ones. The best place to look for its presence is in 'frozen' etymological metaphors, of which there are a number of clear examples. Thus the WG derived verbal stem *itigartit-* 'reject, refuse, turn down' is the causativized form of a stem *itigar-* which when used transitively means 'hit ineffectively, be

deflected (harpoon aimed at sea mammal), not go off (bullet when gun is fired), not effect (a person)' and when intransitive in reflexive use means 'receive a rejection or refusal, be disappointed'. Similar polysemy is found in other Eskimo dialects in Canada and Alaska, so clearly this is an old case of metaphor.

Another example of a frozen metaphor is the WG verbal stem *paasi-* 'understand' (the Polar Eskimo form *paahi-* was given in 2.4). Historically this is from *paa(q)* 'entrance, mouth (of fjord or river)' plus *-si-* 'get, find, come across', although this would hardly be recognized by speakers today since it is completely lexicalized. The image evoked is of a hunter in his kayak out at sea aiming for the entrance to a fjord, perhaps on his way home. This particular form is not found outside of Greenland, but parallel derivations with the same affix *-si-* are found across the Eskimo world. Thus many Canadian and all Alaskan forms of the language have *kangiqsi-* in this sense, literally 'find the source of s.th./ the bottom of a fjord', while some Canadian dialects have *tukisi-*, literally 'find the longitudinal axis of s.th. (e.g. a fjord)'. A very culturally specific metaphor then. Another ancient example involving kayaking is WG *assuari-* 'have s.th. against, reprove', the transitive form corresponding to Proto-Eskimo **aðgura-* 'be stubborn or unwilling', in turn from **aðgur-* 'go against the current or wind'. In Alaskan dialects the intransitive form means 'be stubborn' or 'be bold, undaunted' and in the Copper dialect 'be proud', while the transitive form (*ažguagi-* in Alaska) means 'be undaunted by'. Finally note WG *kappia(sug)-* 'be worried, anxious', ultimately from Proto-Eskimo **kapət-* 'fit tightly, be narrow', and (as a kind of culturally specific euphemism?) *kannat* 'male sexual organs', literally 'bowsprit, projecting end of skin boat railing'.

Other metaphorical expressions are restricted to West Greenlandic, where one suspects Danish influence, as for instance in *putu(aa)* 'stare at someone to make them confused' (lit. 'bore into him' – cf. Danish '*gennemborende (blik)*' 'a piercing look'), further extended to intransitive *putu(vuq)* 'he is in a confused state, crazy' (lit. 'a hole is bored in him') – cf. Danish '*det er hul i hovedet*' (i.e. 'it's crazy'). However the extension of *qapug-* 'foam, froth up' to 'talk loudly/angrily, quarrel' in Greenlandic seems to be a purely native development.

Often the direction of extension in such metaphors in EA is not from the concrete to the abstract (the direction focused on by Lakoff & Johnson), but from one concrete domain to another – sideways as it were. Thus the Proto-Eskimo stem **uvit-* 'open one's eyes' has come to mean 'bloom (of flower)' in some dialects, and 'break away from shore ice to form open water (pack ice)' in others, although the original meaning has also been retained everywhere. And Proto-Eskimo *iva-* 'sit on eggs, brood (birds)' has been extended in the Utku dialect of Netsilik according to Briggs (1970: 314) to refer to lying next to and cuddling someone else in bed. In Polar Eskimo it applies to warming one's hands against

another's body.[2] A West Greenlandic example is the extension of the stem *avalag-* 'sail out (from shore)' to 'go out onto the dance floor', and its reverse *tulag-* 'come ashore (from boat)' to 'come back from the dance floor'. More widespread is the extension of Proto-Eskimo body term **təmə* 'body' via 'main part of s.th. (to which s.th. else is attached)' to 'body of boot (above sole)' and 'land up from the sea, inland, innermost part of house' (in Central Siberian Yupik also 'close relative' and in WG also 'water between boat and land'). See Fortescue (1988) for a lexical area where terms referring to the interior of the human dwelling have been extended out to geographical orientation around the dwelling and beyond.

Other kinds of figurative language can be found in earlier shamanistic vocabulary such as WG *punnguuq* 'dog', literally 'one going with its head bent forward', also recorded for Iglulik Inuktitut and Polar Eskimo (cf. Dorais 2010: 132–134). Metaphorical and metonymic neologisms caused by 'death taboo' against naming the dead and important or dangerous natural phenomena have also resulted in a partial re-lexicalization of the nominal lexicon of isolated East Greenlandic in ways sporadically to be found elsewhere.[3] Thus East Greenlandic *itturmiilaq* 'edible root' (lit. 'which lives in the soil') and *tusaat* 'ear' (lit. 'which is used for hearing'). The usual words for these things in other Inuit dialects are simply opaque nouns (Dorais op. cit.: 133). That such processes are still at work in modern West Greenlandic can be seen in Petersen's (1976) article on WG neologisms for new cultural/ technological items introduced from Europe. Many of these involve the nominal affix *-usaq* 'thing like a –', as in *issuusaq* 'turnip' from *issuk* 'testicle'. These rarely refer to abstract concepts but again project from one concrete domain to another.

An example of varied metaphorical extension from the same native source across dialects is seen in the development of the meanings of the Proto-Eskimo stem **qətəg-* 'jump' as reconstructed in Fortescue et al. (2010). In Alaskan Yupik this meaning is preserved, while it is extended in Central Siberian Yupik to 'attack, endeavour to consummate marriage (man claiming bride at parents' house)', in North Slope Inupiaq to 'play games', in Netsilik 'amuse oneself with athletics', in Caribou 'play cards', in North Baffin Inuktitut 'play (of children)', in Labrador 'be joky, giddy', and in Greenlandic 'dance' (also frequentative extension *qittar-* 'be choppy (sea)'). Whether metaphor or metonymy is involved in these extensions of meaning is debatable. A clearer example of metonymy is WG *maaryug-*,

2. In WG one also finds it used in the sense 'brood over (a thought)', undoubtedly based on Danish '*ruge over*', which covers both the literal and the metaphorical meaning.

3. Human names are usually etymologically transparent in Eskimo languages, containing everyday lexical items. Taboo on naming the dead would normally only last while the memory of the deceased was alive, but in East Greenland the situation became at some stage more permanent.

another's body.[2] A West Greenlandic example is the extension of the stem *avalag-* 'sail out (from shore)' to 'go out onto the dance floor', and its reverse *tulag-* 'come ashore (from boat)' to 'come back from the dance floor'. More widespread is the extension of Proto-Eskimo body term **təmə* 'body' via 'main part of s.th. (to which s.th. else is attached)' to 'body of boot (above sole)' and 'land up from the sea, inland, innermost part of house' (in Central Siberian Yupik also 'close relative' and in WG also 'water between boat and land'). See Fortescue (1988) for a lexical area where terms referring to the interior of the human dwelling have been extended out to geographical orientation around the dwelling and beyond.

Other kinds of figurative language can be found in earlier shamanistic vocabulary such as WG *punnguuq* 'dog', literally 'one going with its head bent forward', also recorded for Iglulik Inuktitut and Polar Eskimo (cf. Dorais 2010: 132–134). Metaphorical and metonymic neologisms caused by 'death taboo' against naming the dead and important or dangerous natural phenomena have also resulted in a partial re-lexicalization of the nominal lexicon of isolated East Greenlandic in ways sporadically to be found elsewhere.[3] Thus East Greenlandic *itturmiilaq* 'edible root' (lit. 'which lives in the soil') and *tusaat* 'ear' (lit. 'which is used for hearing'). The usual words for these things in other Inuit dialects are simply opaque nouns (Dorais op. cit.: 133). That such processes are still at work in modern West Greenlandic can be seen in Petersen's (1976) article on WG neologisms for new cultural/ technological items introduced from Europe. Many of these involve the nominal affix *-usaq* 'thing like a –', as in *issuusaq* 'turnip' from *issuk* 'testicle'. These rarely refer to abstract concepts but again project from one concrete domain to another.

An example of varied metaphorical extension from the same native source across dialects is seen in the development of the meanings of the Proto-Eskimo stem **qətəg-* 'jump' as reconstructed in Fortescue et al. (2010). In Alaskan Yupik this meaning is preserved, while it is extended in Central Siberian Yupik to 'attack, endeavour to consummate marriage (man claiming bride at parents' house)', in North Slope Inupiaq to 'play games', in Netsilik 'amuse oneself with athletics', in Caribou 'play cards', in North Baffin Inuktitut 'play (of children)', in Labrador 'be joky, giddy', and in Greenlandic 'dance' (also frequentative extension *qittar-* 'be choppy (sea)'). Whether metaphor or metonymy is involved in these extensions of meaning is debatable. A clearer example of metonymy is WG *maaryug-*,

2. In WG one also finds it used in the sense 'brood over (a thought)', undoubtedly based on Danish '*ruge over*', which covers both the literal and the metaphorical meaning.

3. Human names are usually etymologically transparent in Eskimo languages, containing everyday lexical items. Taboo on naming the dead would normally only last while the memory of the deceased was alive, but in East Greenland the situation became at some stage more permanent.

Before illustrating this, certain structural traits of Athabaskan languages should be pointed out that may well have a bearing on the prevalence of figurative language in them. They have still fewer verb stems than in EA, and the most widespread and derivationally productive of these – its 'classificatory' stems – are also more abstract$_2$ than any lexical stems in EA.[5] Depending on the accompanying morphology they may indicate either such and such a type of object being in a location, or someone handling such and such a type of object. In Chipewyan, as mentioned earlier, there are ten. One of these can be seen in the following example (from Rice 1998: 101). The actual basic form of a lexical verb is the so-called 'theme', consisting of a stem (and attached 'classifier')[6] plus one or more preceding prefix (often discontinuous) – in this case an incorporated postpositional phrase meaning 'in(to) her palm', positioned before the pronominal object marker, so literally 'he transferred a closed container into her palm'.

ye-tł'aghe-ye-į-ł-tą
3SG.POSS-palm-in-3SG.SUBJ-CLASS-PERF.handle.a.closed.container
'He gave/handed it to her.'

The same stem (for handling a closed container) is involved in various derived forms including ones meaning 'taking' or 'seizing' among many other possibilities. Compare thus the following, with the same verb stem but an extra 'thematic' (probably once directional) prefix *-r(i)-*:

se-tł'aghe-ha-ye-r-ę-l-tą
1SG.POSS-palm-out-3SG.SUBJ-CLASS-PERF.handle.a.closed.container
'He took it from me.'

Such generality/abstractness of meaning is not restricted to classificatory stems alone. Amongst the stems of very general meaning are some that may be used to express emotion and cognition, comparable to the morphological category of emotional roots in EA (which I shall discuss in 7.1) but much more loosely organized and dependent on specific prefix strings to derive the meanings involved. In themselves, note, these formulae are abstract schemas which typically have other prefixes slotted in between their elements. Also unlike in comparable EA expressions, they have evidently been subject to frequent metaphorical extension.[7] For

5. Apart from the 'pro-verb' stem *pi-* 'do, see, think, say', also as a nominal 'thing', in all EA.

6. This traditional term does not refer to a classificatory morpheme but to a set of old voice markers, now largely lexicalized in close connection with the following verbal stem.

7. Rice (1998: 128–9) discusses the still active metaphor and word-play that coupling classificatory stems with particular objects allows – including with abstract objects like 'word', 'idea', or 'plan', which normally require the default THING classificatory stem (if used at all with them).

example, one common stem in Ahtna from which emotional 'themes' are derived is *-niic* 'move the hand, grab, feel' (also in classificatory use 'handle fabric-like object'), as in *u+n+niic* 'love, like' (Kari 1990: 310 – slightly simplified). This is one of the most ubiquitous stems in the language, with many distinct meanings, perhaps originally 'feel s.th. with the hand', made more specific by prefix strings such as the directive sequence *u+n-* here. The meaning appears to be metaphorical, from reaching out and feeling someone. Some of the other themes formed by this stem with preceding derivational prefix strings refer to being happy, being alive, being lucky, hearing, expecting, knowing and believing.

Another stem of broad (abstract$_2$) meaning with a number of positive emotional themes based on it is *-'aa^o* 'linear object extends', as in the following: *su+ko+d+'aa^o* 'be happy, glad' (Kari 1990: 77), where *su* is a thematic element ('happiness, enough'). The same stem can also appear in themes with a negative meaning such as, on the same page: *c'+u+d+'aa^o* 'feel a sharp internal pain', with different prefixes.

A further very general stem (one exhibiting considerable suppletion) is *yaak/laak* 'happen, affect', themes formed from which range in meaning from 'fix, make, hurt, kill, change, say' to 'make a noise'; a particularly complex derivation is *k'e+na+O+l/D+yaak* 'get even with, take revenge on O' (where *k'e* indicates 'following after'). Athabaskan languages generally express non-physical, culturally anchored abstract$_1$ meanings through derivations of (general) abstract$_2$ stems such as this. Another example: *yaniteł aa* 'he is faithful to her' from *-'aa^o* introduced above; the intermediate template here is O+G+ł+*'aa^o* 'keep linear, elongated O in position' (where G is a word class/gender marker). There are a few morphologically simplex abstract$_1$ words, but these are typically nouns or predicate adjectives like *'engii* 'taboo, forbidden, evil', *gheli* 'good, real, true', and *iine'* 'thought, mind'.

As claimed at the outset, the lexicons of NA languages differ from those of EA in the relative prevalence of overt metaphor and metonymy. The way is partly 'prepared' by the small number and semantic generality of verbal stems illustrated above, and in particular by the fluidity of classificatory verb stems, ripe for figurative extension through their rich derivational potential, but there are no doubt cultural aspects to this too. Rice (2012: 66–68) discusses some of these under the rubric 'creativity under constraint'. She also gives examples of the descriptive (metonymic) terms for body parts and (most) animals in Dene Sųłiné (Chipewyan), reminiscent of taboo replacements in East Greenlandic. She also lists some of the conceptual metaphors displayed by this language, namely orientation metaphors like 'IN IS GOOD/OUT IS BAD', ontological ones like 'PEOPLE ARE PLANTS/ANIMALS' and 'BEING AT IS DOING', and structural ones like 'PARTS ARE OTHER PARTS', as well as metonymies like 'CONTAINER FOR CONTENTS' (Rice 2012: 67). Some of the most interesting of these involve times of day, where

in Cognitive Linguistics would be covered by an 'action schema'). At the more abstract (and more complex) end of the scale we find socio-cultural ceremonies or 'scenarios', such as 'getting married' or 'being baptized'. To see what I mean by calling something as simple as opening a door a 'cultural context' think of what is involved in that scenario as we in a western suburban environment typically understand it. This clearly differs greatly from what is involved in some other cultures by 'opening a door'. Thus the underlying schema for Eskimos living in traditional houses involved entering through a sunken tunnel that traps the heat inside but where no door covers the entrance. There is a word for 'door' in West Greenlandic, where people live in modern wooden houses today (*matu*), but the original meaning was 'lid, cover', which (as in other Eskimo dialects) it can indeed still mean. What about Aleut, where the entrance to the house was by a ladder down through the smokehole in the roof? (*galRiX* 'door', now used of a door in a modern house, is etymologically 'smokehole'). And what of the Chukchi *yaranga*, the tent-like structure used by nomadic reindeer-herders on the Siberian tundra which one enters by pulling aside two flaps, but which again lacks a door as such? The noun *tətəl* can refer either to a modern door or to the flap of the interior, tent-like sleeping compartment. Obviously the meronymy of the 'house' is different in these cultures, and the nearest equivalents to 'opening a door' involve very different sensorimotor schemas.

Does OPEN represent a universal primitive then? It does not figure on the list in Wierzbicka (1992), though one suspects that a similar word exists in most languages. What it can be broken down into at an 'atomic' level is an array of immediate causal actions serving an analogous (common) functional goal, which varies in sensorimotor realization from context to context and from culture to culture. The ingredients in these possible combinations with various objects-to-be-opened may be universal, but they are rather far removed from the semantic fields and features defining individual items in the mental lexicon. What is common to all the uses of the verb 'open' (and its near-equivalents) is its underlying 'logic'. It is an essentially 'molar', language-specific entity, neither itself a primitive nor decomposable into atomic primitives other than CAUSE-BE.OPEN, where CAUSE is some fairly direct causal action and what constitutes being OPEN is a functional matter, highly dependent on the type of object of the action and on context. It is therefore a poor candidate for itself being a universal primitive (hence, no doubt, its absence from Wierzbicka's list).

As regards the semantics of words involving cultural contexts at the more abstract$_1$ end of the scale, consider the word 'holiday'. In Fortescue (2009: 56) this word is analysed and contrasted with comparable French words. This required not just a single holistic scenario or frame but arguably five, all associated either negatively or positively with other overlapping frames, namely a 'days off work'

frame (related to a contrasting 'work' one), a (typical) 'holiday places' one, a 'summer holiday' one (related to a 'no school' frame which contrasts with one for 'term time', presupposing a child at school), and a 'weekend' one contrasting with a 'week day' one (both in terms of the calendar). The word is associated with various collocations ('take a –', 'go on a –', etc.) and is related to a number of other near-synonyms. The different types of 'holiday' concerned (each with their own associations) is clearly culturally determined – consider the English bank holiday and the protracted summer holiday from university, not to mention specific 'holy days' in the church calendar. Comparable (but overlapping) words in French include *congé* (leave from work or military service); *vacances* (summer holiday, from work or school, i.e. '*les grandes vacances*'); *jour férié* (day officially free from work); *fête* (religious holy day); and *villégiature* (a stay in the country or at a resort). There are clearly differences as regards the official or informal status of holidays in the two different cultures. I shall return to the way in which the various 'scenarios' associated with the English words can be represented in a neural network model in 8.4.

The way in which Cognitive Linguistics handles contexts of this sort is in terms of 'frames'. These are the stuff of Frame Semantics. Fillmore & Baker (2000) lists hundreds of such entities without (as yet) much systematization or discussion of relevant levels of abstraction. This approach hugs the linguistic ground rather closely, with essentially every word being associated with at least one frame, and does not attempt to account for the dynamic extension and merging of scenarios independently of verbal expression. Some of these 'frames' correspond to very general semantic features or categories – e.g. 'SENTIENT', 'CONTAINER', 'PERCEPTION' – rather than to experiential scenarios as such. For a detailed discussion of one particular frame, that for 'RISK', see Fillmore & Atkins (1992). The authors analyse this frame in terms of the components: Chance, Harm, Victim, Valued Object, Situation, Deed, Gain, Purpose, Beneficiary and Motivation. Any given use of the verb 'risk' focuses on only some of these elements, but all of them are evoked potentially.

Many of the metaphorical extensions in Eskimo-Aleut and – especially – Northern Athabaskan languages that we have seen presuppose frames of a particularly complex nature, namely 'folk-theoretical' models of the kind described by D'Andrade (1987). These include the way that speakers conceive of the mind (op. cit.:127) – in 7.1 I discuss the way in which speakers of Eskimo-Aleut languages appear to have a different understanding from English speakers of the boundary between 'thinking' and 'feeling'. I shall return to the nature of abstract cultural 'artefacts' in general in 11.3.

In sum, metaphor is largely relativistic: an essential ingredient in discussing it is cultural context, the background against which the 'figure' of the core meaning is contrasted and from which it is abstracted. This is in fact the nub of Barsalou

& Wiemer-Hastings' conclusions as regards 'abstract$_1$' words (op. cit.: 157): such words are complex in so far as they take on somewhat differing senses according to specific situation or context types. This would appear to apply equally to metaphor. This is not to put in doubt the powerful notion of 'deep metaphor', but one wonders whether what lies behind for example Talmy's 'Force Dynamics' is not a general ability to recognize and utilize analogy (anchored in physical experience) rather than a specific set of universal conceptual metaphors. At all events, analogy, like metaphor, presupposes abstraction: the basis for the similarity, the features that are comparable, have first to be abstracted both from the source and from the target of the comparison.

CHAPTER 5

The abstraction of events in narrative and memory

5.1 Event structures

What we need to do now is to explore more closely what kind of real-life situations, what kind of contexts, are stored in the speaker's mind against which metaphor and abstract thinking in general can play out. This entails investigating the way in which larger scale entities than the meanings of individual words can be abstracted from daily experience and potentially stored in memory, whether short-term or long-term. On what basis can we convert a continuous flow of perceptual experience so effortlessly into a sequence of events? How can we represent them verbally in a sequence of clause-sized propositional entities? And is this really the way in which we store these narratives in memory, whether from verbal or purely perceptual input?

Within the Cognitive Linguistics paradigm Talmy, in discussing the typology of what he calls 'event integration', introduced the notion of the 'macro-event', the basic unit of conceptualization of happenings in the world (Talmy 2000, vol. 2, Ch. 3). A macro-event represents the 'conceptual integration of a complex event' and consists of one main or 'framing' event defining the abstract core of the event plus one or more subordinate or ancillary event ('co-event') filling in the more specific and perceptually vivid content of the macro-event, plus the relationship between the two, as in Figure 5 (from Talmy op. cit.: 221). The framing event, whether expressed as verb or 'satellite' (e.g. prepositional phrase or adverb), constitutes the core of the event, it is what articulates the activity of the co-event according to general categories of 'temporal contouring' (aspect), 'change of state', 'realization', 'coactivity', and – the basic dimension, from which the others analogically project – 'motion'.[1] It determines the 'upshot' of the whole macro-event and is,

1. Languages differ greatly as to how fine-grained distinctions of 'temporal contouring' are articulated. Thus in Nuuchahnulth (previously 'Nootka') there is not only a basic distinction between 'perfective' (completed) and 'imperfective' (non-completed) aspect (as in Slavic languages), but the latter category is divided into a number of morphologically types distinguished by largely iconic processes of reduplication, lengthening, shortening, affixation or all of these at once. Davidson (2002: 222) lists 'graduative', 'durative', 'continuative' (for dynamic situations),

for example, what is denied under negation and asked about under interrogation (Talmy op. cit.: 282–3). Under the co-event types 'constitutiveness' refers to the support relation of filling in the conceptual 'body' of the temporal contour provided by the framing event. This could, for instance, be an agent's action in bringing about the macro-event.

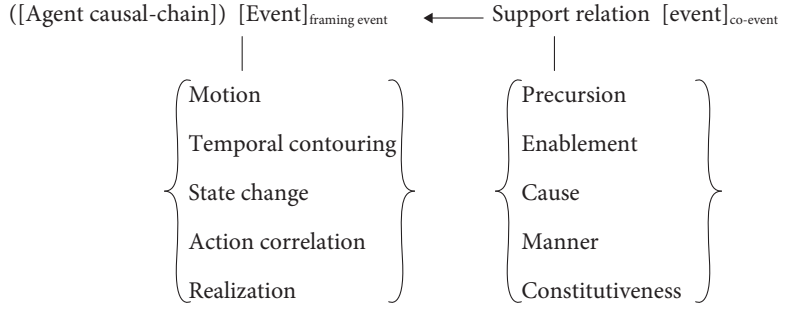

Figure 5. The conceptual structure of the macro-event (Talmy 2000)

A macro-event may be compacted into a single clause or expressed as more than one syntactic clause. In analytic languages like English this is typically expressed in several clauses (and several independent words within each clause), whereas in polysynthetic languages like Mohawk or West Greenlandic this integration is maximal, potentially compacting the macro-event into a single word (although there is usually a choice between maximally integrated and less integrated constructions even in these languages). Just as much as individual clauses in English, polysynthetic 'word-sentences' display – and are constrained by – what Bohnemeyer et al. (2011: 48) call the ´macro-event property', by which they mean that sub-parts of a complex clause cannot be individually modified by temporal operators (for tense and/or aspect), such operators by necessity taking scope over the whole clause. Note that they regard this as a matter of 'language design' at the interface between semantics and syntax and not of underlying cognitive structure – it depends on the lexical and morpho-syntactic nature of the individual language how 'compressed' event representation can be. Amongst other types of construction that they discuss is the serial verb construction. This is maximally compacted in polysynthetic languages like Yimas, which can be contrasted with more analytical, typical serial verb languages like Kalam, another Papuan language described by Pawley (2011). Thus compare the following sentence from Yimas (Foley 1986: 179):

'repetitive' (at regular intervals), and 'iterative' (at irregular intervals). Aspect can moreover be combined, e.g. 'inceptive + graduative + perfective'. These are all clearly abstract$_2$ generalizations across event types.

impa-mpu-yakal-irim-tay-ntut
3DL U-3PL A-EXT-stand-see-RM PAST
'They stood watching them two'

with the following from Kalam (Pawley 2011: 16):

kik am mon pu-wk d ap agi kn-ya-k
they go wood hit-smash get come ignite sleep-3PL-PAST
'They gathered firewood for the night'

Both these sentences correspond arguably to a single macro-event, as do the English glosses below them. Pawley countered Givón's claim that the analytic sequence of serial verbs in a typical Kalam sentence do not represent distinct verbs but jointly reflect a single complex verb constrained by the paucity of lexical verbs in the language (Givón 1990). He pointed out that there is a distinction between 'compact' and 'narrative' serial verb sequences. The former consist of conventionalized/ lexicalized sequences combining a more content-rich verb with a very general one for perception, transfer, transporting or resultative state, and can be said to correspond to single verbs in English. 'Narrative' serial sequences (distinguished by their morphosyntactic behaviour) are much more diverse and *ad hoc*, as in the example above which breaks up the event into distinct stages minimally elaborated – if further elaboration were involved they could indeed be treated as individual clauses taking their own aspectual morphemes for example. The choice is one of degree of compression expressed in the narrative sequence, in turn related to what is felt as a culturally expected scenario (or deviation from it). Note that there are also types of single clauses which express two tightly knit sub-events also in English (e.g. 'she wiped the table clean', 'he made me angry', etc.) – these represent a combination of a framing event (the resultative state) and a co-event (the causing action) in Talmy's scheme. There is some fuzziness to the 'macro-event' notion depending on what is regarded as a typical complex sequence in a given linguistic culture. However, the basic notion seems to be a useful one for characterizing the mapping from pre-linguistic event to linguistic clause structure.

Vendler, in his influential analysis of situation types distinguished states, activities, accomplishments, and achievements (Vendler 1967: 97ff). The last two are both telic (resulting in a goal) but accomplishments take place through stages that can be segmented in time while achievements are 'all or nothing'. One could exemplify these categories from our Alice excerpt: 'wish' and 'remained' refer to states; 'keep appearing and vanishing' is an (iterated) activity; 'vanished' is an achievement, whereas 'vanished quite slowly' (not the standard use of this verb) could be considered an accomplishment.

Zwaan (2008), discusses the categories within an 'Event indexing' variant of Situation Semantics (cf. Dijk & Kintsch 1983 and Barsalou 1999), which analyses on-line language processing in terms of 'simulatable' sensorimotor schemas and operates with the five dimensions of time, space, entity, causation, and intentionality. Focusing on the temporal dimension, he argues that the simulation of accomplishments ("He painted the wall") is easier than that of activities ("He was painting the wall") since telic actions (such as accomplishments) emphasize the resultant state rather than the action, so are more visual than kinetic. The degree of 'motor resonance' (kinetic simulation) varies according to situation type – habitual states or activities and iterated actions are more abstracted away from currently observed or described situations than accomplishments and achievements. (There is no focal object, as is also the case with abstract words in Barsalou & Wiemer-Hastings' analysis discussed above). Zwaan argues that all these categories are nevertheless simulatable. Activities or accomplishments stretching over real time, for example, can be simulated by just brief exemplars of the action/activity types concerned (Zwaan op. cit.: 19–22).

Other converging approaches to the 'event' can be mentioned. Thus for Jackendoff (2002: 364–369) the 'event' is a high level node over purported cognitive primitives (ontological categories) such as OBJECT, STATE, GOAL, PATH, CAUSE or another EVENT, as in Figure 6 for the sentence "Bill buttered the bread". The underlying event is homologous with a surface sentence or clause (decontextualized) and does not indicate any constraint on the nature of the information conveyed.

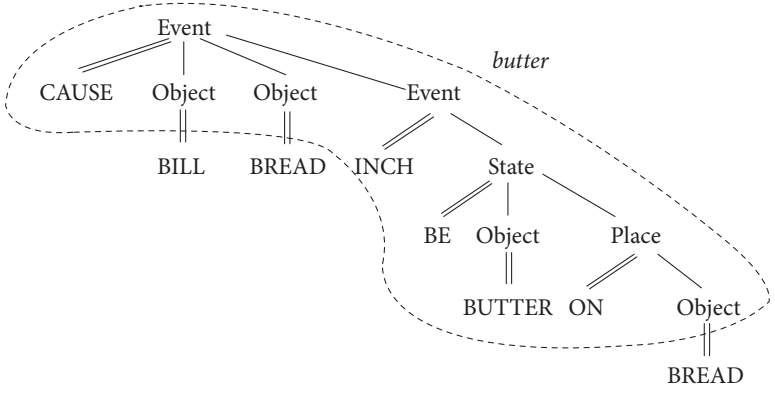

Figure 6. An event structure in Jackendoff (2002)

I should interject at this point a general caveat about cause and effect between events. Both causes and effects are in reality rarely simple, more typically they are multiple except at the most basic level of physical actions – and even then they

may be complex, depending on the level you look at them and (if a human agent is involved) how direct the chain is between intention and result – something some languages, including West Greenlandic, can distinguish. Just consider the chain of events (some perceptual, some involving muscular control) between Bill's intention to butter the bread and his achieving that result. Intentions may themselves be complex and are often hierarchically organized (like sentences), involving discrete steps, each with their own internal sequential structure, that must be completed in a certain order. But the human brain is inclined to abstract simple patterns of cause and effect in all that it experiences – this is often a deceptive form of abstraction in that it can grossly oversimplify causal relations by giving purely statistical probabilities the status of rigid axioms. Consider especially historical 'events' like World War 1, which involved extremely complex chains of cause and effect at different levels, some intentional, others not. But the same is true as regards events in one's personal life, in which multiple factors may come to bear on one's actions (only some of which are accessible to consciousness). At the quantum level, chains of cause and effect are indeed replaced by probabilities. At the level of the events of everyday life there is also going to be a modicum of randomness involved – or at least the involvement of contextual effects so complex that it seems so. This of course does not deny that utterances of a causal nature can be interpreted (at least for some languages) in terms of a single underlying primitive CAUSE linking event representations. In fact the abstraction of such a feature by children from individual unanalysed verbs has been plotted by Bowerman (1974).

Linguists of a more functional orientation like Chafe (1994) are concerned with constraints on the information that can be conveyed by a single clause or utterance. His 'intonation unit' is defined in terms of what can be held within a single 'focus of consciousness'. The principle he proposes is 'one new item per unit', i.e. one 'newsworthy' item per utterance (op. cit.: 288ff.). This is similar to Halliday's 'information unit', consisting of an obligatory new element plus an optional given one (Halliday 1985: 275). Though Chafe's intonation unit is by and large coterminous with a written sentence, he distinguishes the new item it contains from the on-going 'topic', which is hierarchically organized within the broader discourse and may cover several individual sentences – he defines it as the semantic content most active/ highlighted for conscious perusal by the reader/hearer at the moment. (There is also semi-activated background material.) It determines the perspective of the hearer/reader at that moment (Chafe 1994: 137ff.). Since what is 'newsworthy' depends on the on-going context we see here how event and context are inter-dependent. Abstraction by a speaker of what is most 'newsworthy' in the context will result in he or she modulating the information structure of their ensuing utterances for pragmatic ends (e.g. by prosodic or syntactic means of focusing or topicalizing).

One early advocate of the importance of context for meaning is Vygotsky, appreciation for whose seminal work on the acquisition of language (as well as the nature of inner speech in adults) has been waxing in recent decades. Corresponding to the cognitive 'event' discussed above, his 'psychological predicate' is an entity in inner speech consisting of a foregrounded predicate as content (i.e. a proposition minus its subject, since that is generally presupposed) against a background context of what one is thinking about (Chafe's 'topic'). For him the meaning of words is only "one of the zones of sense, the most stable and precise zone", sense being a dynamic fluid whole, embracing all the psychological events aroused by the word in the context in which it appears, and "in inner speech, the predominance of sense over meaning, of sentence over word, and of context over sentence is the rule" (Vygotsky 1962: 146–7).

Vygotsky's position has influenced McNeill, who adds to the 'psychological predicate' the notion of 'growth point', a unit of speech/gesture synchrony marking a newsworthy departure in the immediate context, which itself consists of meaningful oppositions (McNeill 2012: 33). For him gesture provides largely subconscious non-verbal suppletion of meaning that words alone can't supply and orchestrates the verbal expression that is 'unpacked' from the growth point. But are the gesticulations of growth points 'abstract' in the same sense as the image schemas of Cognitive Linguistics? They are certainly simplified schemas but they are also *ad hoc*, unlike sign language signs in which hand forms, for example, can be conventionalized analogues of human or animal bodies, etc. (as are hand movements like walking a hand, which also bear categorical meaning). Sign language (like spoken language) is segmentable into proposition-like units (between successive movements or positions), and this seems questionable as regards *ad hoc* gestures. However, McNeill's point is that gestures form synchronized unities with verbal expressions (i.e. form 'growth points'), which are indeed segmentable into propositions that reflect underlying cognitive entities and are ultimately anchored in sensorimotor schemas (and thus not abstract$_1$). Growth points may be unpacked as event-sized 'propositions' abstracted from temporally extended perceptual experience (potentially orchestrated by gesture). But the gestural components themselves can hardly be characterized as 'abstract$_2$' ('simplified or generalized across instances') or 'abstract$_1$' ('disembodied from specific sensory features') – they are clearly 'embodied', although their purport may not be immediately obvious to an observer.

Finally, an important distinction was made by Lyons between 'first order entities' (referring to physical objects), 'second order entities' (referring to events, processes and 'states of affairs' that can be witnessed), and 'third order entities', namely propositions outside of space and time (Lyons 1977: 142–5). The latter category is clearly more abstract than the other two, for example the complement clause in

"I know that John is in London", which may be true or false (though 'John' and 'London' are referential, the object clause is only a belief of the speaker). As Lyons points out, second order events can nevertheless be referred to by 'abstract' nouns in English (thus "I saw John's departure" as opposed to "I saw John leaving") although both refer to a (concrete) second order entity. In some languages (like West Greenlandic) there is a clear difference of construction with the 'objects' of transitive verbs in the three cases, third order entities being expressed by nominalizations as opposed to second order ones that are expressed by finite verbal clauses (and first order ones by simple noun phrases).[2] It is at all events no coincidence that third order entities can be expressed in clause-length units analogous to second order ones in most languages – the basic subject-predicate structure of the 'event' underlies both.

All these approaches to underlying cognitive 'events' converge on something richer and more experientially meaningful than the amodal representations of propositional meaning of the kind formal truth-functional semantics deals with. They derive their richer meaning from context – context taken in the broad, multi-levelled sense introduced in 2.1. Even 'third order entities' (the objects of belief or other 'epistemic stances') are subject to the same contextual constraints on the use of words at other levels. We are evidently equipped innately to recognize and recall events as such – and not just simple physical ones but also specific social situations of an abstract kind such as those behind idioms like 'he let the cat out of the bag' or 'when the cat's away the mice will play'. What we must consider next is how successive events abstracted from written or spoken 'texts' can be built up into larger entities in memory.

5.2 Mental models

These larger entities have been dubbed 'mental models' (Johnson-Laird 1983; for further developments see Garnham & Oakhill 1992). Mental models may be temporary or permanent schemas in memory, but typically they do not directly retain the actual linguistic input for long but rather store the result of its analysis in the form of mental representations of the situations described by the verbal input. These are image-schematic representations, structural analogues of real or imaginary worlds. Their organization may nevertheless reflect the propositional structure of the input, in which case they have been called 'situation models', with

2. Thus *nalu-aa Hansi-p tikin-ner-a* (ignore-3s/3s.INDIC HANS-POSS arrive-NOM-his) 'he was unaware of Hans' arrival' (where the subject = nominal possessor) as opposed to *taku-aa Hansi tikit-soq* (see-3s/3s Hans arrive-3s.PART) 'he saw Hans arriving'.

each proposition corresponding to an event structure element (Dijk & Kintsch 1983). They are akin to written 'abstracts' of articles or books that highlight in a few key sentences the central ideas to be expressed. Kintsch (1988) in his extension of the situation model to a 'Construction-integration' model separates two phases, a 'construction phase' in which word meanings are activated, propositions formed and inferences made, and an 'integration' phase in which a network of interrelated items is integrated into a coherent structure drawing upon schemas already stored in long term memory and any contradictions resolved. Mental models are presumed to be largely spatially organized, though they may also contain causal and temporal information about people's goals and characteristics (Gernsbacher 1990). Presumably they must further contain referential 'nodes' to which Vygotskyan psychological predicates can be pinned, allowing rapid anaphoric reference to participants in the text. Ferstl et al. (2005: 724) sum up the relationship between verbatim text and mental models in the following way:[3]

> On a level more removed from the verbatim input, comprehenders form a global representation of a mental model of text or discourse information. This so-called situation model integrates the current language input with both general knowledge and the prior discourse context. In contrast to bridging inferences, elaborative and predictive inferences are considered part of the situation model… The situation model of narrative texts has been shown to contain information about several important dimensions of the story, including the where, when, who, what, and why of events. Its representation is not necessarily verbal or propositional but flexibly tailored to the specific contents.

As an example of a 'bridging inference' consider the kind of intuitive jumps one makes between the statement "He returned home at 9 o'clock" and "The door was open," allowed by our implicit knowledge of houses (and homes) having doors. Such immediately obvious information filling in gaps in the logical train of events presumably does not need to be added to accreting mental modals as one reads one's way through a text.

One theme common to later variants of mental model theory is the multimodal nature of such representations – Dijk & Kintsch's 'situational model' for example is multi-modal rather than just consisting of non-modal propositions in a neutral 'Mentalese' in the manner of Fodor (1975) or Pinker (1984). This will be discussed further in 6.2. Mental models could even consist of imagery-plus-language units with some verbal material (names or 'key words') retained along with spatiotemporal imagery as such, reminiscent of McNeill's syncretic 'growth

3. The authors point out that evidence from brain damage cases indicates that the right hemisphere of the brain plays a crucial role here (I shall have more on this in Chapter 8).

points' or of Barsalou's 'perceptual symbols'.[4] In fact that is what was suggested by my investigation of the retention of a short anecdote by multiple speakers of English (Fortescue 2007). In this experimental investigation 24 native speakers were asked to read an anecdote then, following distraction, to write it down as they remembered it – and then once more at a later date. The result was (both times) 24 different versions in which the verbatim surface was clearly largely lost or decayed to various degrees (though near-verbatim retention was achieved in a single rare case). A common core of propositions was elicited corresponding to the core of the anecdote (just enough to preserve its point). These propositions were analysed as forming Whiteheadian 'strong nexūs' of concepts corresponding to situations/events in the real world, as opposed to the 'weak nexūs' of links between forms and functions typical of symbolic systems.[5] (I shall return in Chapter 10 to the relevance of Whitehead, for whom 'propositions' are anchored in perception.) The results generally pointed towards a coherent retention of the causal sequence of events (including the intentions of the protagonists) leading up to a succinct version of the 'punch line' of the anecdote, but with much variation in detail apart from the naming of certain key concepts (usually verbatim). Individual stylistic variation reflected differences in familiarity with conventional anecdotal forms (and whether the individual was a habitual teller of such stories himself/ herself). The overall nature of the various renditions was clearly a subjective matter.

Zwaan (2008: 14), discussing mental simulations in general, claimed that a new event is easier to integrate into a model and more strongly connected to it in long term memory if it shares at least one 'index' with the active part of the representation (i.e. the same entities, time frame, spatial location, causal chain or intentionality). Thus, bearing in mind his caveat on differences between directly observing and reading about events (op. cit.: 17ff.), one may assume that in reading the passage from Alice in Wonderland given in Chapter 2, it is easy to integrate Alice's reply with the mental model being built up by the reader at that point:

> *"I wish you wouldn't keep appearing and vanishing so suddenly: you make one quite dizzy."*

From the preceding context the reader knows full well that 'you' refers to the Cheshire Cat and 'one' to Alice herself. Moreover, both the time frame and the location form a seamless continuation of the preceding event (Alice first noticing the cat on a bough), so this new information (however bizarre) can easily be integrated with one's on-going construction of a mental model of the story.

---

4. Note, however, that for Barsalou a 'mental model' has a somewhat different, broader sense as a specific simulation of a 'frame', e.g. for object types like 'car' (Barsalou 1999: 591).

5. Note that the plural of Greek *nexus* is *nexūs* – it is a term I shall be referring to again.

The reader of the passage will also be familiar with narrative conventions in story-telling (though Alice in Wonderland hardly adheres to all of them), and certain 'top-down' large-scale schemas may assist as a template for both comprehending and for production. In fact highly schematic suggestions for 'story grammars' have been proposed in the past. Thus Labov (1973: 363) suggested the following generalized components of narratives in English: 1) Abstract (announces the story and indicates what it is about); 2) Orientation (identifies the initial context, e.g. time, place and participants); 3) Complicating action (answers the question: what happened?); 4) Resolution (reveals the outcome of the complicating action); and 5) Coda (summary remark signalling that the narrative is finished). It is 3) and 4) that are typically made up of narrative clauses representing individual events in sequence. The problem with such endeavours is that they fail to account for how stories are actually produced or understood, and what the level of generality and the relevant categories actually are. They also fail to account for individual differences in experience and expectation.

Consider how an individual reader might in fact have built up an internal representation of the Alice story up to the point where he or she reaches the end of Chapter 6 ("Pig and Pepper"), just after the exchange with the Cheshire Cat that we have seen. This will also provide a handy illustration of the hierarchical nature of abstraction. The written narrative is of course already divided up within the overall book into sentences consisting of individual words, and of paragraphs grouping together a number of sentences, and of chapters containing a certain number of paragraphs. But this does not necessarily reflect the hierarchical arrangement of information stored by the reader – we can only really guess at that, although recall techniques such as those in Fortescue (2007) can bring us closer to accessing that. As time passes after the initial building of a mental model at first reading we can envisage further stages of generalization coming to the fore, i.e. abstractions over abstractions, though in fact the higher levels of generalization may have been building up gradually already during the first reading and not require much later reformulation. Let us assume as a minimum that over the course of time the reader's memory of the content of the book will become blurred or overlaid, with the level closest to the verbatim text largely discarded and the most superordinate level (the gist of what the book is about) retained intact the longest.

At the most local and immediate level the reader will presumably have analysed each successive sentence and added to their mental model a simplified version of the content either in clause-sized propositions or perceptual images, each containing a single event or state (including individual quoted utterances and 'thoughts'). That is, in so far as that content is new and does not merely repeat what is already stored. What he or she will minimally have retained after finishing Chapter 6 will not be a verbatim sequence of words nor even a linear chain of

propositions corresponding to each successive sentence, but some representation of the kernel sequence of events (minus all secondary detail, including most of the dialogue and its punning quips). In theory this could be in the form of a sequence of 'predicates' (whether modal or amodal) attached to 'nodes' representing the different characters (Alice in particular), in the manner of Vygotskyan inner speech. But without committing ourselves to any particular format we might imagine that what is retained is something like the following. At the start Alice is standing outside a house which at the end of the last chapter was small, but now that she has nibbled one of the pieces of mushroom given her by the Caterpillar she has grown small herself (the dashes to the left distinguish episodes, i.e. changes of scene):

– Alice observes a fish-footman come running out of the wood and rap at the door – the door is opened from within by a frog-footman – the fish-footman hands over a letter to the frog-footman – he says it is an invitation from the Queen to play croquet – they bow to each other and their wigs become entangled – Alice laughs and runs back to the wood until the fish-footman leaves – she comes out again and knocks at the door – the footman explains why there is no answer – there's so much noise within they can't hear – she gets no further sense out of the footman and opens the door herself –
– Alice sees the Duchess nursing a howling baby and the Cook stirring a cauldron – there's so much pepper in the air she sneezes – she notices a grinning cat on the floor – the Duchess says it's a Cheshire Cat – she is rude to Alice – the ill-tempered cook starts throwing pots and pans at the Duchess – the Duchess ignores her and starts singing a lullaby to the baby – she starts tossing the baby violently up and down – she flings it to Alice – Alice catches it as the Duchess hurries away to get ready to join the Queen –
– Alice carries the baby out of the house as she tries to nurse it – the baby turns gradually into a pig as she proceeds into the wood – she lets the pig down and it trots off – she thinks of other pig-like children she knows –
– Suddenly she sees the grinning Cheshire Cat on the bough of a tree – she addresses it, asking directions – it indicates the directions of the houses of the March Hare and the Mad Hatter – they have a discussion about who exactly is mad – the cat says it will see her again at the croquet game – it vanishes – then appears again and enquires about the baby – Alice explains that it has turned into a pig – the cat vanishes again – she walks on heading for the March Hare's house – the cat appears again and wants to know if she said 'pig' or 'fig' – Alice replies [and here comes the short section of dialogue given in 2.1] – the cat disappears more slowly this time, its grin going last – Alice thinks to herself: how curious! –

– She walks on until she reaches the house of the March Hare – because it is so big she nibbles a piece of the other mushroom – and grows big again herself –

At a later point in time (after finishing the whole book) the reader may have forgotten a good deal of this content, recalling only the following still further reduced version which he or she is able to reconstruct upon request:

> "A fish-footman has an invitation for the Duchess, which he delivers to a frog-footman at the door. Alice observes this transaction and, after a perplexing conversation with the frog, lets herself into the house. The Duchess's cook is throwing dishes and making soup with too much pepper, which causes Alice to sneeze. Alice is given the baby that the Duchess is nursing and to her surprise the baby turns into a pig. The Cheshire Cat appears in a tree, directing her to the March Hare's house. He disappears, but his grin remains behind floating in the air for a while, which confuses Alice."

All the person in question might recall many years later from having read the whole book when young might be something like this:

> "It's about this nice little Victorian girl Alice who falls asleep at a picnic and dreams of a white rabbit she sees hurrying past on two legs holding up its watch. Following him she falls down a rabbit hole and after falling a long way she lands, unhurt, in Wonderland. Curious things start to happen to her – like growing bigger or smaller after eating bits of mushroom given her by a caterpillar smoking a hookah. She meets all sorts of funny creatures – like the Cheshire Cat, which has a trick of disappearing gradually, leaving just its grin, and the Mad Hatter and the March Hare who invite her to a crazy tea party – and the Mock Turtle that can't stop crying – and let's see, there's an angry queen always threatening to cut people's heads off – and she's organized a game of croquet where the mallets are live flamingos and the balls hedgehogs – and Alice is herself put on trial for something or other she says that displeases the queen. But it all ends when she wakes up back at the picnic again…"

Note that Chapter 6 of the book is reduced here to remembering the Cheshire Cat. Clearly this final version is more 'abstract$_2$' than the first one. Perhaps more of the first mental model of the story created at the time of the first reading does remain in the reader's memory but is simply no longer accessible, or perhaps it has gradually become degraded to the point of non-retrievability. Either way, right after the first reading the reader could presumably have given accounts at any of the three levels of 'granularity' represented above if asked to do so – he or she could have abstracted from their mental model just that level of detail they were asked about (perhaps including whole chunks of dialogue remembered verbatim), so several levels of abstraction were potentially available.

Figure 7. Tenniel's illustration to Chapter 6 of "Alice's Adventures in Wonderland" (Carroll 1865)

As to the modal or amodal nature of the mental model initially retained by our hypothetical reader, it may well be that what was retained constitutes (in large part at least) a coherent spatio-temporal aggregate of propositional units in the manner of Dijk & Kintsch's situational models, with each proposition corresponding to an event structure element. It may alternatively be retained as a more explicitly perceptual representation made up of chained 'perceptual symbols' in the manner of Barsalou. It is reasonable to suppose that some information at least is stored not in propositional format at all, but as rough pictorial representations – think of the grotesque illustration by Sir John Tenniel (Figure 7) that accompanies the original text and shows Alice observing the duchess nursing her baby with the cook in the background. Surely the reader will remember this *as well as* the linear flow of events. Details such as the grinning cat curled up at the bottom left will probably not be retained, however, the Cheshire Cat's image being more likely to be associated with the bough where most of its dialogues with Alice take place. The emotional response evoked by the picture and the overall atmospheric 'feel' of the events of the chapter must be part and parcel of the 'glue' that holds the

reader's mental model together as a distinct entity. The retention of certain key items in the form of words (verbatim) should also not be ruled out – especially the names of the characters. Clearly these cannot be represented by anything other than, well, their names.

We may then agree, I hope, that the end result of event representation in mental models is typically a much simplified amalgam of coherent elements of both a perceptual and a verbal nature (in varying proportions depending on subject and context), probably organized to some degree in 'proposition'-length units, but with much individual variation. We bring to bear on our understanding of texts and discourse, besides knowledge of the language system, a wide variety of schemas that we have abstracted from our experience in the past. We may further utilize such models to shape our own narrative productions – to ourselves or others. Whether the kind of abstraction from written texts illustrated above can be transferred directly to our memory of everyday events is an open question, but there is reason to believe that the organization of episodic memory is also hierarchical. Thus I can remember the overall order of things that I have done during the day – for example, the way I slotted in the writing of the preceding passage between other tasks that needed doing of both a routine and an *ad hoc* nature. But I can also recall some of the more fine-grained decisions I made and some of the actions I undertook in writing it, and in what order – right down to a fairly fine-grained memory of how I moved the caption of Figure 7 into its present position below the illustration from the position to which it had been dislodged when I expanded the text above it. Whether I retain this information in memory much longer than today is another matter – in another year or so I may just remember the writing of this book as a holistic undertaking, with the detail of its writing lost in a blur.

5.3 Kinds of memory

I have spoken until now of memory in a rather undifferentiated manner, but several kinds have been distinguished. Tulving (1972) introduced the distinction between semantic (general, context-independent) and episodic (subjective, context-dependent) memory, a distinction that has been revised in more dynamic, procedural terms in recent times (cf. Zelinsky-Wibbelt 2011). A further distinction was introduced by Ullman (2004): declarative vs. procedural memory, each with a distinct role to play in language and involving distinct parts of the brain. Declarative (or explicit) memory involves especially the temporal lobes of the brain and is organized in a distributed, perhaps connectionist manner. It supports the mental lexicon – including irregular morphological forms. Procedural (or implicit) memory, instantiated in frontal cortex and the basal ganglia, is taken to handle automatic

'how to' knowledge such as how to ride a bicycle but also the symbol-based sequencing of grammar (syntax), including regularly affixed morphological forms. As Bourtchouladze points out, subconscious memory 'without remembering' (as with 'priming') is probably broader than the 'procedural' distinction was meant originally to cover – just think of the 'tip-of-the-tongue' experience, a matter of incomplete activation of a memory trace. Remembering (recall) is not the same as recognition – the latter can be unimpaired while the former is severely compromised (Bourtchouladze 2002: 62).

Then there is emotional memory, which can according to LeDoux be either implicit, involving the amygdala, or explicit (of an emotional situation), involving the hippocampal system (LeDoux 1998: 200–204). Loss of the amygdala can result in loss of the ability to recognize emotion in faces (and of 'social' cognition as such – Bourtchouladze op. cit.: 83). There is also memory for faces in general and for other specific kinds of content (such as names or musical pieces), which can be distinguished by impairment affecting them but not other forms of memory, e.g. 'prosopagnosia' for the inability to recognize faces (Bourtchouladze 2002: 31–2).[6] For the role of 'long term potentiation' (LTP) affecting the efficiency of synaptic transmission in the hippocampus, the key organ in memory formation, see Reiser (1993:119ff), also Bourtchouladze (op. cit.: 132–4).[7] Bourtchouladze goes into considerable detail as to the complex neurotransmitter pathways involved (op. cit.: 112–144).

The mental lexicon is different again, a form of memory generally accepted to be highly distributed across the cortex although 'anchored' in the left hemisphere speech centres – I shall return to this in 8.3 in connection with 'abstract' lexemes. Of particular interest in this general examination of memory is the question as to how mental lexicons for more than one language are located in the same bilingual brain. Are they separate or integrated? In Fortescue (2014) I gave an account of recent hypotheses in this area and attempted my own version of an integrated model.

Dong et al. (2005) discuss various theories that have been proposed for the organization of the bilingual mental lexicon – both shared and separate storage variants. They present experimental evidence indicating that there is a shared (distributed) conceptual system for the bilingual speaker's two languages, with

6. The recognition of faces may in fact involve several stages, any of which may be individually compromised, thus, curiously, a subject may be shown to recognize a face implicitly without being aware of having done so (Bourtchouladze op. cit.: 57–8).

7. The function of the hippocampus in mice is for spatial orientation and mapping (e.g. finding the way through a maze with the help of landmarks), and some of this original function probably remains in humans. Finding one's way through a maze and finding one's way through long-term memory have something fundamental in common.

asymmetrical links to the first language (L1) and the second language (L2), the former priming elements of the latter more strongly than the reverse. They argue, further, that conceptual representations may contain both L1- and L2-specific information, converging. Most recent neuroimaging investigations support the idea that a single common neural network is involved in the processing of both languages for proficient bilinguals, late as well as early.

In my model I proposed that the different languages of a bilingual speaker may share the same overall 'mental lexicon' (including combined or converging experiential 'scenarios' associated with individual words), but that there are also language-specific 'subsets' within it, in the manner of Paradis, who has proposed a 'Three Store' model. According to this there are separate lexical sub-systems intimately intertwined but separately activateable or inhibitable as wholes (Paradis 2004: 196–198). The distinction between fluent/ balanced bilinguals and those who learn their second language late in life is, according to Paradis, a matter of the greater on-line reliance on consciously accessible 'declarative' memory in the case of late learners – they also draw more heavily on pragmatics and context than on automatic 'procedural' memory. Early, fluent bilingualism may, he suggests, result in a more integrated network in frontal grammar cortex for the two languages concerned (not a matter of lexicon as such). Late or less fluent bilingualism, involving second language learning after the full acquisition of the dominant one, may on the other hand require more widespread cerebral activation. What L2 lacks for most speakers may be the deep limbic (i.e. emotional) anchoring of L1, on which the learning of a second language must piggyback (Paradis 2004: 24–25). This approach is also in line with de Bot's (2004) multiple processing model, which contains language-specific sub-sets in each of three stores, conceptual, syntactic and formal (phonological). He adds a 'language node' switch to Levelt's (1993) model of speech production.

In the same paper I further suggested a principle of cortical proximity which places 'mediatory columns' for near-equivalent words in the bilingual speaker's two languages in proximity within the same association areas of the cortex. (I shall unpack the term 'mediatory word column' in 8.3.) This principle could facilitate the operation of intertwined language-specific networks by ensuring, among other things, that rapid, quasi-automatic code-switching from one language to the other can take place. Such networks could be distinguished by differential phonological anchoring: word forms presented as sequences of the phonemes of one of the languages will act as activating input to neural representations (word columns) for words in that language but will not 'resonate' with those for the other – it does

not matter so much how physically closely they are situated.[8] This meshes with my discussion of Evans' 'lexical concepts' in 3.2. since this may be the level at which bilinguals can translate back and forth rapidly between words and phrases in their separate languages, bypassing both having to correlate the more diffuse cognitive models associated with them and having to translate lexeme by lexeme between them.

More radically distributional models of the bilingual lexicon have been proposed, for instance de Groot's (1992) Distributed Feature Model in which translation from one language to the other is bi-directional via the shared feature-based conceptual system. This is taken to account for the relative ease of translating concrete words (sharing many as opposed to few semantic features) as opposed to abstract ones. It allows for any degree of semantic overlap between words. What most of these models reflect is the compatibility of an integrated, intertwined mental lexicon for bilingual speakers with widespread distribution of the whole integrated system. The consensus appears to be that this integration is likely to be more compact in early rather than late bilinguals.

All the forms of memory addressed in this section so far are relatively long term, although episodic memory is supposed to be transformed gradually from temporary storage in the hippocampus to longer term storage in the cortex where long-term mental models are stored. Contrasting with them is working memory, i.e. short-term memory, consisting according to Baddeley & Hitch (1974) of a visuo-spatial sketch pad, a phonological loop and a central executive. As is well known from studies of dementia, short-term memory can be dissociated from long term memory, with the former compromised while the latter remains intact. Also the opposite is possible: long-term memory can be lost while working memory is intact (Sacks reports on a number of such cases, e.g. in Sacks 1995: 49–50).

How do these various forms of memory (if they are indeed distinct) determine how we abstract the most relevant elements from our environment to enter into our personal memory store? How are these integrated into our ongoing 'stories' from day to day, from month to month and year to year? It is well known that our memories are not fixed but continually undergoing changes owing to new information and shifting contexts. Perhaps memory traces are not just 'there', ready made, but are so only in potential states of activation and have to be abstracted in

8. This may be a matter of how early in relation to the onset of myelination the languages were acquired. One might speculate that proximity between near-equivalent lexical items must apply to the mental lexicon of the early fluent bilingual, since these items would have been laid down before myelination of the relevant cortical areas had proceeded very far and when the proximity to the relevant sensory/sensorimotor region would have been crucial in the acquisition of basic vocabulary in either language.

context whenever summoned by active processes, with missing links filled in by inference from past experience and general knowledge. Answers to similar questions to those above have been proposed from the perspective of cognitive science (and Cognitive Linguistics in particular) by Glenberg (1997). For him memory is a 'meshed' set of patterns of action "integrated by virtue of their analogical shapes" (op. cit.: 3). Memory records are updated whenever there is a change in conceptualization of the environment. This may occur either in automatic mode, meshing directly with past experience or, when the environmental input is supressed and conceptualization is guided by conscious, intentional 'trajectories'.[9] The purpose of memory, he states, is to serve perception and action, "to mesh the embodied conceptualization of projectable properties of the environment with embodied experience that provides non-projectable properties" – 'non-projectable' meaning that which cannot be abstracted simply from the 'optical flow' of experience (op. cit.: 3–4). Glenberg sees initial coding not in terms of discrete digital 'features' (as in connectionist models) but as analog, in terms of patterns of possible action. From this perspective there is no deep distinction between episodic and semantic memory: "Episodic recollection is a type of pattern completion via meshed bodily constraints on action" (op. cit.: 8), a view already adumbrated by Miller & Johnson-Laird (1976), for whom there is essentially a single encyclopaedic store with varying relations to initial context.

In a similar vein, from the perspective of 'perceptual symbols' and simulation semantics, Barsalou (1999: 604–5) argues that working memory (involving buffers for all of the senses) is the system that runs perceptual simulations, both 'on line' in perception and 'off line' in thought. Whereas long-term memory contains 'simulators', working memory implements specific simulations (I shall discuss these terms further in Chapter 6). He regards memory retrieval as another form of perceptual simulation, either unconscious in 'implicit memory' or conscious in 'explicit memory'. The activated simulation in remembering may differ from the original perception owing to various kinds of restructuring and/or forgetting (and possibly to less 'bottom up' constraints).

From a more detailed neurological perspective, Binder & Desai (2011) present a multiple level model of memory that offers a useful framework for explaining the workings of abstraction in memory, one compatible as regards the relationship between perceptual and conceptual representation with the position of 'embodied abstraction'. This posits variable degrees of sensorimotor embodiment with progressive abstraction of conceptual knowledge from perceptual experience (op. cit.: 534). Semantic memory thus covers both modally specific and supramodal

9. The notion of 'trajectory' (how one situation flows into another) is reminiscent of Whitehead's 'perception in the mode of causal efficacy', which I shall discuss in 10.3.

representations (op. cit.: 528). Episodic scene construction (also prospective planning) presupposes semantic/ conceptual knowledge, and common to all memory is "learning through generalization across individual exemplars and gradual abstraction from perceptual detail, used over a wide range of domain-specific tasks" (op. cit.: 534).[10] I shall return to their model in Chapter 8.

10. It should not be thought, however, that memory alone is sufficient for abstraction: consider the immune system, which is indeed an active and highly ramified form of molecular memory built up during the lifetime of any living organism, but it lacks a subject or 'prehender'. The vehicle of such a system, say a human body, has no way of accessing what is going on with the defense potential of white cells in its circulatory system (except through gross effects such as illness).

CHAPTER 6

Pre-linguistic abstraction and projection

Perception and imagination

6.1 Visual abstraction

The question now is whether abstraction of the kind we have dealt with up till now is also reflected in domains other than the immediately language-related. It is time for a first excursion away from language. The role of abstraction in pre-linguistic perception is on the face of it rather harder to bring out than is the case with language since the processes involved are far from simple and involve much parallel processing. However, the end result (recognizing an object or face, say, behind a complex array of visual sensory data) is clearly a matter of initial abstracting followed by matching of what is abstracted with the contents of memory. As with the abstraction of meaning from speech this is a matter of 'abstraction by matching'. The following much simplified account can be compared with the sketch of the abstraction of words from auditory input in 2.1.

Much is known today about the human visual system. Light from a visual stimulus hits the retina at the back of the eye, where light-sensitive cells convert it into electrical impulses which are carried by the optic nerve via the geniculate body and the thalamus to the visual cortex at the rear of the brain, left-of-field input from both eyes carried to the right hemisphere, right-of-field to the left. Here the processing is passed on to different areas, each specialized for a particular type of analysis: general scanning, stereo vision, depth and distance, colour, motion, and position in space. These are cross-linked but divide into two principal streams, a dorsal and a ventral one, the former (going from the occipital to the parietal to the frontal lobes) determining 'where' the object perceived is located or moving, the latter (going from the occipital to the posterior temporal to the anterior temporal lobes) determining 'what' kind of object it is. Colour, for example, is on the 'what' path, depth and distance on the 'where' one. Actually a more recent view of the dorsal stream sees it as involved with more general visuo-motor coordination (cf. Hickok & Poeppel 2007: 396). One could say (somewhat oversimplifying) that abstraction largely involves the ventral and posterior parts of the cortex, projection the dorsal and anterior parts. The first area the input hits (before the paths split) is that which is responsible for general scanning and since only a small part

of the input is in good focus (from the foveal centre of the retina), the eye must jump around in small rapid 'saccades' to scan the visual field in detail. The patterns projected onto this part of the visual cortex is already distorted from a straightforward map, with different neural assemblies responding to lines and angles and other features of the array.

Consider as a concrete example how Alice might recognize the Cheshire Cat perched on its bough: first her eyes take in the whole visual field encompassing the cat and its immediate surroundings, this is then shunted from the retinae to the visual cortex where the 'picture' is analysed into shapes and sizes and colours and position, etc., the combination of which leads to the activation of areas elsewhere in her cortex responsible for matching with her experience of animals in the past. Of the information stored here only the part commensurate with the presence of a cat in her visual field (shape, colouring, whiskers, tail, etc.) will be held active, the rest suppressed or ignored. But it is not just any old cat – she has already met the Cheshire Cat and remembers it, not only by its typical feline features but also by at least one unique feature – its un-catlike grin (well, also its ability to talk). The 'raw' sensory data entering Alice's eyes will not make sense (if it does!) except through matching not only with her previous experience with cats in general (stored as memory traces capable of reconstructing the visual and other sensory experience) but also of having met this particular cat before. Her emotional response to recognizing it, involving the deeper 'limbic' system of her brain, is also relevant to her renewed confrontation with the cat. Knowing now which cat it is she has certain expectations of the way it will behave (if her brief memory of it is to be relied upon), though given the enigmatic pronouncements of the creature this is open to surprises. Whatever the details of the process, it is hard to say where 'abstraction from sensory input' ends and 'matching with the contents of memory' begins. Of course you can't – it is a seamless process from one end to the other, albeit one that proceeds through various stages ending in higher levels of cognition (in hierarchically superior cortex). And beyond recognition there is further projection of the recognized object back onto the 'screen' of the spatial world that she experiences herself as being within at the moment. She perceives the cat in a particular location – on a bough of a tree not too far from (and somewhat above) her. The parallel 'where' stream analysis of the sensory information she has received is no doubt part of this process.

Of course the processes involved are much more complicated than this sketch suggests – a truer picture is presented in Felleman & Van Essen (1991), which investigates the hierarchical organization of neural pathways between the visual sub-regions of the macaque cortex (not much different from the corresponding regions of the human cortex). Referring to the authors' complex circuitry diagram (Fig. 4, p. 30), Crick (1994: 168) explains : "As one ascends the visual hierarchy…

one goes from cortical area VI, which deals with rather simple visual features (such as oriented lines) that occur all the time, to areas that deal with objects like faces, which occur much less frequently, until one reaches the cortex associated with the hippocampus where the combination of signals it responds to (both visual and others) corresponds largely to unique events." It is not until that point that one can speak of a specific face having been abstracted from sensory input. Elsewhere Crick discusses the influential computational account of vision of Marr (1983), who describes the successive stages from the 'primal sketch' of the input at the retina to a '2 ½ D 'sketch' and finally to a 3D visual image, becoming more and more abstract as distance from the retinal input increases (Crick op. cit.: 76).

All of this involves initial simplification of the input in various ways, but the big question is how it all gets put back together again in the process of recognition and whether this still involves 'abstraction' in some sense. There are theories around as to how it is done informed by neuro-scanning techniques of ever more precise spatial and temporal precision, involving temporal synchrony in specific synthesizing regions of the anterior temporal lobes. I shall return to this in 8.2. Clearly some kind of matching with the contents of memory is involved. This can involve a number of discrete steps (mostly unconscious) – for instance recognizing the general kind of object concerned (e.g. a face), whether or not it is familiar, whose face it is and what is known of that person, and finally perhaps attaching a name to the face. This is a matter of narrowing down possibilities (somewhat like the case of the 'Cohort' model for language mentioned in 2.2), so characterizing the overall process as one of abstraction is not an unwarranted extension of the usual sense of the word. The way to final recognition may be complex – it involves parallel pathways as well as top-down expectations that intervene and speed the process at all levels, also the automatic Gestalt principle of 'filling in' missing visual information (when lines are broken by an obstacle in the foreground, for example). It is nevertheless a continuation of the analysis of the input, using the outcome of the purely sensory analysis as the 'way in' to abstracting the specific content of memory for which a match is sought. The brain may be fooled both by optical illusions (a function of its inner workings) or by top-down imposition of expectations from prior experience – hence the first artists' impressions of kangaroos as giant hares in the early paintings of flora and fauna from Australia. In integrating novel sensory input with the contents of memory the brain automatically applies 'close enough' analogy if exact matches are elusive.[1]

1. There are similar effects in language, such as misleading 'garden path' sentences (of the "The horse raced past the barn fell" kind) and the relationship between idiomatic and literal meanings of sentences. These complement the anatomical parallels between vision and language circuitry.

How can information abstracted from the environment in the visual mode be transposed into a verbal mode of expression? This is the interface problem, and whether or not one regards language as encapsulated in its own 'module' in the manner of Fodor (1975) it needs to be addressed. One way of doing so (one close to my own perspective on the matter) is the dual coding hypothesis of Paivio (1986). This posits a perceptual code that encodes the perceptual characteristics of a concept and a verbal code that encodes the abstract, non-sensory aspects of a concept, either of which – or both – may be relevant to a particular task. Representations in the mental lexicon are modally specific and multimodal. (There is some confirmation for this from the way in which concrete and abstract words are processed differently, a matter I shall return to in 8.3). Jackendoff's approach to the relation of linguistic conceptual structure (CS) to spatial structure (SpS) is similar but somewhat more abstract (Jackendoff 2002: 347–350). In fact his SpS structures are – like image schemas – also rather abstract (e.g. the position of one star above another). They encode imagistic percepts but are not themselves images. Some things can only be encoded in the one or the other way (thus predicate-argument structure for example is only encodable in CS). However, there is overlap allowing the two systems to communicate, e.g. as regards part/whole and causal relationships. SpS structures are thus part of the meaning of many words (but not the most abstract$_1$ ones).

Much of what has been said above about visual abstraction and projection applies also to the other senses – the auditory mode is of particular interest since it is time-extended and has much in common with the analysis of the stream of speech. Larsson (2015) reviews recent research on the recognition of 'auditory objects', for example by the characteristic vibrations made when different things are struck (as with a tool). Such objects require sequential integration (as opposed to the essentially simultaneous integration binding the perception of visual objects), which involves grouping temporally separated sounds emanating from the same source (notes, words, footsteps, etc.) into a coherent auditory stream. The auditory mode does also have a spatial dimension (the location of the source of sounds in space), as the visual mode also has a temporal dimension (for example the streaming of the environment as one moves through it). The suppression or 'masking' of self-generated signals (sounds) to distinguish and isolate external signals is particularly important in the auditory mode. Larsson sees the integration of information in the auditory and visual mode (as well as the kinaesthetic mode in the rhythmic action of tool use and self-locomotion) as crucial in the evolution of both music and language.

6.2 Thinking with and without words

As we have seen, Vygotsky described 'inner speech' as an abbreviated form of external speech, essentially involving predicates with subjects (and other referents) understood from context. This is perhaps not very different from 'Mentalese', described by Pinker (1994: 81–82) as corresponding to a simplified but semantically explicit, universal human code not greatly removed from actual spoken 'surface' languages (though not reflecting their phonological form). Another view – one corresponding perhaps better to experience – would be that inner speech is largely based on well-rehearsed fragments of usage. But is thought possible entirely without being shaped by speech? If so, does it use the elements abstracted from visual experience such as Paivio's perceptual code? Can inferences and predictions, the stuff of rational thought, be based on such a non-verbal code?

One writer who is convinced that non-verbal 'rational' thought is not only possible but widespread in the animal kingdom is Bermúdez. He distinguishes several kinds of rationality that he calls respectively level 0, 1 and 2 rationality, on all of which adult humans may operate (Bermúdez 2007: 117–128). Level 0 refers to simple tropistic, instinctive behavioural dispositions, but it nevertheless represents adaptive 'choices', albeit generic, not particular ones. Level 1 rationality does not extend to full-blown decision making but is applicable to behavioural tokens – 'seeing' Gibsonian affordances (possibilities for re-action) in the environment, leading to a particular course of action as opposed to another in a given situation at a particular point in time (cf. Gibson 1972).[2] Such simple comparisons of alternative behaviours available at the moment may be driven by short or long term advantages. An example would be the way my cat Jasper sometimes early in the morning throws his front paws at the door of the bedroom, from which he's excluded at night. He knows the loud thump will wake the occupants and bring breakfast somewhat closer – however, this is not a reflex action at seeing any old closed door, it is a certain door at a certain time when he is spurred on by hunger (or boredom). Annoying but effective.

Level 2 rationality, true conceptual thought or 'second order reflexion', requires language and what he calls 'intentional ascent', i.e. the ability to hold thoughts in mind and think reflexively about them. However, the use of language isn't in itself criterial. Jasper does a recognizable if approximate imitation of the human word 'hallo', which he uses to greet his owners or generally call attention to himself – this was not taught him by conditioning but simply copied from human use of that word that he has heard again and again when popping up unexpectedly. It

2. Bermúdez gives a detailed example in the 'roaring' competitions of rutting red deer stags (op. cit.: 135–8).

probably counts as another instance of Level 1 rationality, since he can choose when to use it and to whom. He surely abstracted the type of situation in which to use it from situations in which we his owners have habitually used it ourselves. His usage can be contrasted with the Level 2 rationality presupposed by ordinary human use of the same sequence of sounds in a much broader range of situations (with considerable prosodic modulation of meaning). In at least some of these there is reflexivity in so far as we are aware of the response our use of the word will arouse in the addressee and of our own reasons for using it in the given situation.[3] Jasper cannot (alas) be said to display Level 2 rationality, which requires the implementation of consequence-sensitive 'instrumental' belief-desire couplings, for instance knowledge of how to achieve a goal using the intermediary of tools and the ability to think reflexively about it.

Level 2 thought is compositional – its elements can be chained and infinitely modulated according to circumstance and goal – but Bermúdez claims that other animals (mammals at least) can have thoughts of a compositional nature with determinate content already at Level 1. It is possible to understand how they think about their environment in terms of 'proto-inferences' that keep track of regularities between states of affairs. Jasper's 'thought' before thumping the bedroom door with a particular goal in mind probably fits the bill. This kind of rudimentary causal cognition is widespread in the animal kingdom and among very young infants (op. cit.: 149). Bermúdez suggests that it is possible to distinguish between the two levels experimentally by investigating for example whether a creature persists in its behavioural response despite the cessation of the relevant 'instrumental contingency' (the situation supporting the linkage between action and goal). For him there is no inconsistency in envisaging forms of inference associated with each of these levels (cf. op. cit.: 140–149 on 'proto-logic'). Quite apart from level 2 'logical' inferencing, adult humans apply 'natural' level 1 inferences all the time in everyday life – using analogy, intuitive guesswork and the 'feel' of probabilities. What Davidson (1980) terms 'practical reasoning' based on particular sets of beliefs and attitudes that can lead causally to particular intentions to act may belong here.[4]

---

3. Compare a sincere 'hallo' spoken in a situation of greeting someone you have just been introduced to, or a breezy 'hallo' accompanying a wave as you passes a stranger while out walking, or an exaggeratedly drawn out low-to-high 'hallo…' when encountering an unexpected situation – or even the irritating female usage (spread no doubt through American films) of a high level 'hallo', with a sharp rise-fall on the last syllable meant to draw an abstracted male's attention back to the serious business at hand…

4. Though Bermúdez is not happy with the degree of reflection – typical of Level 2 reasoning – his analysis presupposes (op. cit.: 148).

Particularly relevant as regards our main theme of abstraction is the study of the African Grey parrot Alex by Irene Pepperberg (1999), since she presents convincing evidence for the ability of at least one intelligent non-human for forming and handling abstract information. Throughout her work with Alex she emphasized the importance of social context, the learning of new items being greatly facilitated by the inserting of new items to be learnt in a framework already established (e.g. of a particular phonological word or a particular type of question), much as with human children (op. cit.: 227, 237). Alex was thus able – to an impressive degree – to form and label categorical classes with human words, i.e. to "respond similarly to discriminated stimuli". This was arguably not just by 'stimuli generalization', i.e. by "memorizing a specific reference stimulus or set of stimuli followed by responses to new stimuli based on a failure to discriminate between the reference stimuli and the new stimuli". By teaching Alex to apply a symbolic label like the English word 'red' correctly she was teaching him not just to refer to a class consisting of a concrete attribute but to form a separate class/ category referred to by a colour label (an arbitrary sound pattern) which indeed related to a physical attribute but could be applied to a wide variety of objects, and not, say, to a shared shape (op. cit.: 52–4). Of these phonological labels he learnt several hundred during his lifetime. He was capable of grasping hypernyms like 'colour' or 'shape' as well as hyponyms of specific colours or shapes (correctly answering questions like 'what colour is X', where object X was characterized by shape as well as colour).

Through further carefully designed experiments that capitalized on the essential factor of 'joint attention' between Alex and his trainers, Pepperberg could demonstrate Alex's ability to recognize abstract 'sameness' or its absence when comparing objects, and to discriminate *which* attributes were the same or different, getting him to answer questions like "What is the same?", referring to items differing in all features except one, for example 'colour' (op. cit. 62–3, 89–93). Alex also showed an ability to use abstract labels for (small) numerical quantities correctly and to reply to 'recursive-conjunctive' sentences like "What colour is the 3-cornered wood (object)?" He clearly displayed 'referential mapping', i.e. the ability to actively apply his categories to new items, not just labelling ones already taught (op. cit.: 248). We are naturally dealing here with 'abstract$_2$' generalizations rather than 'abstract$_1$' concepts detached from specific sensory experience. And of course parrots like Alex in the wild do not use human words in this symbolic fashion – their skill at imitation serves other purposes there.

Within the Cognitive Linguistics paradigm, Barsalou (1999) presents a well-developed theory of 'perceptual symbols'. These are modally specific schemas capable of being manipulated in a variety of cognitive functions. They are as 'abstract' as any kind of symbol, but unlike the hypothetical 'amodal' symbols of 'Mentalese',

encapsulated in a language 'module' impenetrable to perception, they maintain an analog relation to perceptual input (much like image schemas) from which they are 'extracted' (i.e. abstracted) rather than 'transduced' to arbitrary symbols of the amodal kind. Perceptual symbols are seen by Barsalou as implemented through a combination of bottom-up activation of sensory-motor areas and top-down activation of association areas to become organized as 'simulators'. These are generalized concepts around a common frame in memory capable of producing infinite simulations, multiple conceptualizations or 'construals' in specific contexts (op. cit.: 587). A 'frame' for him is an integrated system of perceptual symbols used to construct specific simulations of a category, i.e. mental models, and the simulation emerging reflects the strongest attractor in the frame's state space (op. cit.: 591). Frames function as contexts for specific meanings (e.g. 'foot' when referring to a human body as opposed to that of a horse or a tree). Perceptual symbols are componential (some aspects may be activated while others not) and need not stand for a particular individual – they can be indeterminate and generic (op. cit.: 584). Words associated with simulators provide control over their manipulation (e.g. in linguistic constructions built up from their combinations). They further allow us to control each other's simulations in the absence of actual referents (op. cit.: 594). Simulators may combine physical and introspective elements to produce abstract concepts since perceptual symbols can represent sequences of events and allow the necessary framing, selectivity and introspection as well as sensory experience (op. cit.: 600).

For Barsalou perception and cognition are not watertight systems – they are in part penetrable. Bottom-up input (from the 'clamped' perceptual environment) takes precedence over top-down schemas when conflict arises, but in the absence of bottom up input processing is possible in pure imagery and 'conception' (thought), and with various possibilities in between, in which filling in, anticipation, implicit memory, etc., may affect perception (op. cit.: 590). Perceptual symbols being modal and thus analog, they are realized by associative patterns of neurons with dynamic properties, with different contexts distorting activation of the original pattern, much like connectionist 'attractors' – they are not just images, but neural patterns 'standing for' them. Most connectionist models, however, are constructed on the basis of amodal features (op. cit.: 579).

Perceptual symbols form categories (prototypes and exemplars) that allow categorical inferences from actual sensorimotor experiences (op. cit.: 587) and involve them in decision-making and planning for the future. They may implement propositions (binding type predicates to token individuals) that represent the 'gist' of sentences. Though non-referential themselves they may play a heuristic role in referring through complex simulations, adding specific contexts that correspond to Russell's 'definite descriptions' (Russell 1905). Barsalou further claims that they

can form the basis of formal symbol manipulation (op. cit.: 605–6). Compare also Langacker (2005) on logic based on transformations of image schemas.

Other approaches that converge on a perspective of 'thinking without (necessarily) words' – but also without positing a separate Fodorian 'language of thought' – include that of Prinz (1997) on the coordinated perception of events and action planning (the 'common coding' of events and actions). He posits an 'action-effect principle' whereby actions are planned and controlled in terms of their effects. This is relatable to the close relationship between abstraction and projection argued for in the present book. McNeill (2012) allows for narrative expression in terms of gesture alone (cf. the experiment described on p. 79 ff.). An intermediate position is proposed by Slobin (1996). In coming to terms with Whorfian relativism in a palatable contemporary form he introduced the notion of 'thinking for speaking' whereby one's language influences the way one formulates one's thoughts prior to expression by constraining it (normally) to the habitual pathways afforded by one's language.[5] This presupposes a degree of choice as regards realizing concepts as linguistic output.

In the following section we shall look a little more closely at the projective side of cognition in imagination and dream, both of which depend on schemas abstracted from experience and stored in memory, but recombined in ways that are not 'clamped' by direct perception of the environment. Though we are essentially dealing with non-verbal forms of cognition both of them may also contain verbal material.

6.3 Dreaming and imagining

Coming from a different direction than that of the psycholinguists and cognitivists discussed above, Globus (1987) delves into dream imagery as a psychiatrist and philosopher with phenomenological interests. He argues against the 'tinkering with memory traces' view of dreaming, and refers instead to the 'abstract specifications' behind the imagery of both dreams and memories. An example might be 'the movement of one's left arm', which could be satisfied by many different variant concrete movements (op. cit.: 40). This is a sense of abstractness more akin to that of image schemas ('abstract$_2$'), i.e. generalized across instances, than to that of abstract$_1$ words.[6] His term covers primitive types of intentional act (like kicking

5. This is a contemporary take on the original so-called Sapir-Whorf hypothesis, which in its strong form sees language as directly affecting and constraining thought.

6. And to be distinguished from yet another kind of 'abstraction', colloquially synonymous with 'absence of mind'. What is such an 'abstracted' state of daydreaming abstracted from?

or hitting) as well as higher-level perceptual schemas relatable by family resemblance to specific instantiations (such as his examples 'a dome-like structure' and 'luminous filaments'). They are potential vehicles of subjective meaning for the experiencer. Dreaming imagery is for him created *ab ovo* by mental acts, not by simple rearranging of mnemonic copies of waking experience. The relationship to dream life imagery and certain episodes of one's 'wake life' is abstract, with the same specifications bridging the two worlds. Both are in a sense 'projected'. In the waking state, "sensory input influences the world generated by picking out certain tunings of filters on input… In dreaming, certain tunings from waking become reoperative and determine the life-world of dreams." (op. cit.: 109).

Globus attempts to answer the question as to what sort of machine could be 'formatively creative' – not any kind of computer conceivable today. He proposes a 'wet Leibniz machine' capable of dealing with a vast number of possible worlds, creating like the human immune system a kind of 'anti-world' as it goes. This he sees as feasibly generated by the interconnected 'hyperneurons' of the brain which integrate both sensory and 'tuning' inputs. They act as filters with satisfaction conditions activated via 'good enough' matches (op. cit.: 133–37). He relates their hologram-like creativity to the 'unfolding of an 'implicate order' of possible worlds in the manner of Bohm (1980). Be that as it may, the intimate link between abstraction and projection is clearly germane to the perspective on the brain's essential function that I have been promoting. I would just add a caveat: dream imagery obviously involves considerable 'morphing' and blending of existing memory traces but this presupposes that we store neural representations not just of individual features but also of whole analog simulations, in particular of human faces. It is surely whole neural instantiations of faces that 'morph', not bundles of features.

Abstraction in general has in fact a particular relevance for dreaming (and *vice versa*). This has been obscured in the past by excessive focus on the question of whether or not dreaming is 'symbolic'. What it certainly does seem to involve is the superimposing of abstract 'schemas' on sensory imagery which may or may not have anything to do with them. Such schemas (including goal-oriented ones) may have been active during the day in connection with a particular problem or task one was engaged with – or they may be of a more long-standing nature. The result is 'hidden meaning' (that of the schemas 'working themselves out') in a tenuous or arbitrary relationship with more or less randomly generated (and hence sometimes bizarre) association-linked imagery experienced simultaneously.

A more mainstream neurological approach to dreaming and imagining is that of Reiser (1993). For him the retrieval of memory traces in general involves reactivation of the processing patterns whereby they were registered and these patterns

---

Presumably one's total sensory environment.

"may be stored as fragmented or part patterns of originally registered stimuli". He believes that neural networks must somehow contain "every last perceptual pattern fragment of perceptual information that has ever been stored" (op. cit.: 141–2). This presumably includes whole gestalts of, for example, faces and voices, as well as 'fragments' into which they may have been analysed. In the waking state perceptual residua can be reassembled and visualized 'in the mind's eye' (not hallucinated 'as in the dream'). There is neurological evidence that analogous modality-specific regions are activated in mental imagery as in ordinary waking state perception. Thus according to Kosslyn et al. (2001) mental imagery can be shown to utilize the same neural machinery as external perception and motor control – the primary visual cortex in particular can be shown to be activated during mental imagery. In the waking state the parts are "reassembled to provide a very close approximation of the original stimulus pattern". However (and this is where Reiser begins to differ from Globus), in dreaming sleep, with external sensory input blocked, the brain is functioning "in such a way that the parts may be rearranged and combined to form composite images, even from patterns that were originally registered on different occasions." His preferred explanation for what happens during dream-filled REM sleep is that stimulation by the PGO (pons-lateral geniculate-occipital cortex circuitry) could activate stored perceptual *residua* in association cortex and that these activated patterns would then follow a backward path, finally stimulating area 17 (the primary visual cortex), which in turn could initiate ordinary forward-moving programming in dream consciousness (op. cit.: 144–5). Presumably something similar underlies hallucinations, whether visual or aural. Thus Crick (1994: 147) points out that neurons have projective as well as receptive fields. He also discusses 'blindness denial' (projecting imagined objects), and the confabulations of split brain patients as well as hallucinations in general (op. cit.: p. 167).

From a Cognitive Linguistics perspective, Fauconnier & Turner (2002) discuss 'the construction of the unreal' by verbal means – counterfactual statements, lying and pretending (op. cit.: 217–24). This is at a level of conscious intention rather far removed from dream imagery, but is nevertheless germane to the kind of abstraction presupposed by any form of imagination. Elsewhere they characterize creative imagination in terms of 'blending', which they see as furnishing the 'compression' of time, space and causality involved in the establishment of 'mental spaces' (op. cit.: 92–3). The sudden temporal accelerations in dream may be regarded as instances of mental 'compression'. These may in fact be influenced by language, our internal narratives being shaped into proposition-sized units. Whereas verbal narratives handle temporal shifts with the help of conjunctions, dream sequences may involve the simple speeding up or curtailing of whole episodes, leading to the typically obscure or missing causal links in dreams.

Lewis Carol's Wonderland world is a particularly delightful product of the imagination, one packed with playful verbal puns of the kind dreaming occasionally throws up. One can but speculate on what its imaginary creatures owe to the author's own perceptual experiences, but it is clear that he never encountered a dodo or a Cheshire Cat in real life. However, he must have had access to the necessary conceptual content to create such plausible imaginary referents, whether *ab ovo* or by 'tinkering' with fragments of past experience. "Alice's Adventures in Wonderland" as a whole can be seen as a plausible dream projection displaying its own dream logic, with its illogical mergings and fluid associative shifts highly reminiscent of REM sleep dreaming. Perhaps it is just a little too logical, too literary, however: dreaming, it seems to me, is largely random and uncontrolled. We just try our best to make sense of it, like Alice. It is probably good training for dealing with the unpredictable, chaotic side of life.

You have read how Alice might have observed the Cheshire Cat. Now I propose a little exercise in imagination that you can try for yourself. Close your eyes and try to imagine how the cat looked to her on the bough (not on the floor as in the picture reproduced in Figure 7)… How easy was that? How detailed a picture could you form? Was the grin prominent? Was it grinning straight at you? What colour was the cat? What position was it in – were its paws tucked under it? Did you see the bough as well? Perhaps you saw the outline of a cat-like creature. Did you at least feel the vague presence of a cat in front of you? Was the background dark? Think what it was that you did (even if you failed to summon up much of a picture at all). Were you influenced by images of the Cheshire Cat you have seen before? Or was it just some kind of average cat such as the word 'cat' evokes for you onto which you somehow superimposed a grinning mouth? Unless you are eidetic and have retained a precise version of some reproduction of the Cheshire Cat you have seen, the image that you produced (if you managed to do so) would no doubt have been much simplified and schematized. Could one call such an image 'abstract' then? In what sense? I'll leave that up to you to decide.

To sum up the content of this chapter: just as abstraction and projection (comprehension and production) are tightly knit within the language system, visual perception and projection – mediated by memory – are two sides of the same coin and neurologically inter-dependent. This is particularly evident in the case of perceptual hallucinations but is also manifest in the products of imagination and dreaming, and indeed extends into our ordinary perception of the world.

CHAPTER 7

The feel of things

Abstract emotion

7.1 Basic emotions: Abstract and universal?

If perception, just like language, is based on abstraction, on 'image schemas' extracted from bodily experience, what about emotion? Hasn't this been left out of the equation? We first need to agree on what constitutes an emotion, and more specifically on what constitutes a basic emotion common to all who possess a human body. Psychologists discussing the 'basic' array of human emotions often refer to studies of universal facial expressions (especially to Ekman 1984). But deciding what is to count as 'basic' is a controversial matter, since different lists are arrived at depending on one's definition of 'emotion', and on what basis they are to be determined – by overt behaviour, by measurable somatic reactions, or by verbal report. Most importantly, the distance between unconscious 'emotions' and conscious 'feelings', verbally labelled, has somehow to be taken into account (cf. Damasio 2000: 35ff.). There is clearly a long way between autonomous physical responses to situations (heart rate, sweating, temperature changes, etc.) and the conscious registration of emotion – in fact there is still considerable debate about 'which comes first'.

In Fortescue (2016b) I investigate emotion words in the Eskimo-Aleut language family, a particularly useful source of such coverage since there is a morpho-syntactic category of 'emotional roots' recognizable by the special set of affixes that must be attached to them. On Table 3 I list the 24 types that can be reconstructed for proto-Eskimo stems (the starred forms). Two of them (in parentheses) are not reconstructable to Proto-Eskimo but have reflexes in many contemporary dialects. The more subtle differences between the forms given are not relevant here.

Most of the 24 types distinguished here on semantic grounds can be grouped under the six basic emotions described by Johnson-Laird & Oatley (1992), namely: happiness (defined as the perception of improving progress towards a goal); sadness (when a goal is lost); anger (when a plan is blocked); fear (when a goal conflict or a threat to self-preservation occurs); disgust (of something to reject); and desire (of something to approach). The authors argued against the compositional theory of emotion that denies the existence of basic emotions altogether. They suggested

Table 3. Emotions expressed by emotional roots in Proto-Eskimo

1.	Fear (**ira-*, **alikə-*, **naŋyar-*)
2.	Loneliness and sad feelings (**aliga-*,**ar(ə)yu-*,**nəka-*)
3.	Frustration (**capir-*)
4.	Anxiety (**kappəya-*)
5.	Shame (**kayŋu-*)
6.	Timidity (**qikə-*,**əgtug-*)
7.	Apprehension (**paqu(mi)-*, **kama-*)
8.	Worry (**pəŋəg-*)
9.	Disgust (**maruyug-*, **əplər-*)
10.	Not feeling like doing s.th. (**əq(ə)ya-*)
11.	Anger (**nəŋ(ŋ)ar-*,**qənər-*,**ugumi-*)
12.	Regret (**qivru-*)
13.	Jealousy or envy (**cikna-*,**tucu-*)
14.	Anguish or suffering (**ikvig-*)
15.	Happiness (**quvya(yug)-*)
16.	Amusement (**təmci-*)
(17.	Feeling at ease)
18.	Thankfulness (**quya-*)
19.	Lust (**əkli-*)
20.	Longing (**qi(C)əlir-*)
21.	Pity (**naŋłəg-*)
(22.	Feeling protective/loving towards)
23.	Uncertainty (**nəryu(g)-*, **ukvər-*,**nału-*)
24.	Feeling or noticing (generally) (**əlpəkə-*)

that an examination of the words used by languages to describe emotions should be taken seriously and not rejected as mere folk etymology. Theirs is a kind of prototype perspective on emotion, with the individual emotion words of a given language being 'more or less' like the basic, universal set, but subject to culturally specific variation. Their study lies close to the linguistic surface, and is specifically concerned with the emotional reaction to the perception of general categories of event by speakers of English. The distinction between different emotions within each of these types can be accounted for in terms of different stimulus situations, different strengths of emotion, and specific distinctions of a cultural/social nature. Applying this to Eskimo languages it will be noted that 'believing', 'doubting', and 'noticing' are also expressed by emotional roots. This is because 'feeling' in this

culture (unlike in English) covers any internal mental phenomenon that is tinged by emotion – in fact feelings that are principally manifested by outer behaviour (such as violent anger) are not expressed by emotional roots at all.

An important question arises here: what is the cognitive status of the Johnson-

for culturally specific categories. This they do by defining contexts of use in terms of Wierzbicka's purportedly universal metalanguage. They are sceptical about the notion of 'basic emotion' and claim that socially acceptable ways of thinking about the kind of events that provoke anger, etc., and the kind of behaviour that results from these feelings, are integral parts of the emotion itself (op. cit.: 7). Certainly the social contexts in which emotional words are seen to be applicable cannot be ignored.

7.2 Secondary emotion

Damasio (2000: 51) makes the distinction between primary (universal) and secondary or social emotions. He also distinguishes what he calls 'background' emotions. These encompass such internal, visceral or somatic states as fatigue, excitement, sickness, relaxation, stability, and discord. They are "an index of momentary parameters of inner organism state" (op. cit.: 286). Secondary emotions encompass such things as embarrassment, jealousy, guilt and pride, which are subject to socio-cultural variability. Typically they are complex, shame for example involving elements of both (primary) fear (e.g. of social ostracizing or mockery) and disgust (at oneself), also perhaps hope (of getting away with something unnoticed). The development in terms of evolution is, according to Damasio, from background to primary to secondary emotions (op. cit.: 342). Drives and motivations and states of pain or pleasure generally fall outside of these categories (though all have been labelled 'emotions' by some investigator or another).

The general picture emerging from Damasio's work is of a universal set of subconscious basic emotions, neurologically hard-wired in their own sub-systems and serving evolutionary survival functions, overlaid by a wider array of explicit types of 'feelings' coupled to long-term memory and thus available to conscious awareness and assessment. We recognize their applicability to other people by their behaviour, similar to our own when we are subject to situations of the same type. The meaning of emotion words can be expected to display a degree of social modulation and conventionalization by the very fact of their being labelled lexically. We can thus say that primary 'happiness', for example, lies behind a number of more specific secondary emotions depending on circumstance – at least those of protective tenderness (towards one's offspring), of love (towards one's mate or a close friend), of amusement (in a convivial situation, e.g. of verbal banter), or of thankfulness (for receipt of assistance in attaining some goal).

Still more culturally specific uses of words related to a basic emotion may then develop further under the aegis of traditional beliefs, including organized religion. Even those Eskimo emotional roots that lie very close to their 'primary' source,

such as 'disgust', refer not just to the basic emotion concerned but to the type of stimulus evoking it, as modulated by the particular cultural and physical environment lived in by the Inuit. In fact all emotional roots in these languages have a transitive counterpart (with a particular affix) referring to the type of stimulus evoking the feeling concerned, glossable as for example 'be disgusted by –'. In other words, recognizing the type of stimulus (object or situation) is essential to the feeling. So characterizing as 'secondary' all the emotions referred to by the 24 categories of emotional root analysed in Fortescue (2016b) may be justified, whatever their ultimate relationship to Johnson-Laird & Oatley's six 'basic' emotions.

The distinction between different sub-types of emotion is thus rather diffuse, and indeed the divide between emotion and cognition (as emotion-neutral thought processes) is itself somewhat fuzzy. What is perhaps special about the Eskimo category of 'emotional root' is that it distinguishes morphologically a relatively clear-cut semantic domain of concepts felt to be primarily emotional. That the division between emotion and cognition is not made at quite the same place as we in Europe would make it is not so important: the clear-cut opposition of the two in the western world is perhaps due to the long-standing history of dualistic thinking in this domain – in 'folk' as well as philosophical theories.

A more universally relevant aspect of my investigation of Eskimo emotional terms lies in its diachronic aspect, Proto-Eskimo being eminently reconstructable by comparing the contemporary languages and dialects. It supports the idea that the semantic development has always been predominantly from the physical/visceral correlates of emotion to more abstract or general emotion and finally to more culturally determined meanings, in the manner Damasio proposes. Metonymy coupled with contextual modulation appears to have been the principal driving force. This is not surprising, given that the relationship between an emotion and its outward behavioural manifestation can itself be characterized as metonymic. A similar point as regards the relationship between emotions and their physiological effects is made by Lakoff (1987: 382) in his analysis of anger. What this all suggests is that the processing of bodily emotions that results in conscious feelings involves successive stages of abstraction, paralleled by the development of their verbal expression. And it is worth adding: abstraction is not in itself an unemotional matter!

7.3 Aesthetic abstraction in art and music

We have seen how abstraction is involved in the construction and understanding of narrative, but how about in forms of aesthetic expression in which abstraction is directly tied in with emotion? The two are most obviously linked in abstract art, in particular abstract expressionism, where the dynamic externalization of

emotion is of the essence and pictorial representation reduced more or less to nil. In fact the history of contemporary art is (or was for most of the last century) one principally of increasing abstraction and simplification, away from its anchoring in the world of external sensory perception – from the impressionists through Cezanne and cubism, the surrealists, Matisse, Kandinsky, Miró, and Mondrian, leading up to fully abstract forms of art. According to Lucie-Smith's account of the post-war period there were two strands to abstract expressionism: first the energetic and gestural one in the works of for example Jackson Pollock, whose own guiding emotions are reported to have been 'self-doubt and anxiety' (Lucy-Smith 1969: 36). Pollock himself is reported to have seen his own paintings as having a life of their own which he tried to 'let come through'. Then the more abstract and tranquil strand in for example the plain (but 'fuzzy') blocks of colour of Mark Rothko, for example.

A subsequent development was towards 'post-painterly abstraction' (e.g. in the coloured squares of Josef Albers, with his link to Bauhaus functionalism). This was 'hard-edge' abstraction, coming to the fore mainly in America (op. cit.: 94ff.). It was now entirely divorced from emotion and reflected the modernist architecture and design of the times – simplified, functional, cool. Lucie-Smith (op. cit.: 100) likens it to Logical Positivism in the philosophy of the time (linguistic and analytic, disembodied).

Perhaps as a reaction against the inhuman geometricization of art and architecture of this period public interest was roused in the early seventies (initially in scientific circles) by a quite new form of visual art, one based on remarkably simple mathematical formulae infinitely reiterated – the Mandelbrot set, and other self-similar 'fractals'. As Gleick described the phenomenon, the swirling organic shapes of fractal art (as in Figure 8) resonate with the way nature organizes itself and the way human perception sees the world – "the harmonious arrangement of order and disorder as it occurs in natural objects… dynamical processes jelled into physical forms" (Gleick 1987: 116–7). If a short mathematical formula (e.g. $z_{n+1} = z_n^2 + c$) can reflect the infinite variety of nature's 'chaotic' forms when expanded into a dense array of points and areas (and subsequently given colour by the artist), cannot the formula itself be said to be a pure abstraction waiting to be projected back into the sensory world which it somehow reflects? I shall return to the Platonic world of mathematics and the relationship between mathematical abstractions and the perceptual world in 10.1.

On a somewhat different aesthetic plane, the cartoon is also worth mentioning here – both the comic drawing and the animated film kind are abstract, but in what sense? Both the figures and the backgrounds in a cartoon are simplified compared to real life but are they in that case abstract$_2$ (generalized across instances)? This is more obviously so as regards backgrounds (stereotyped settings)

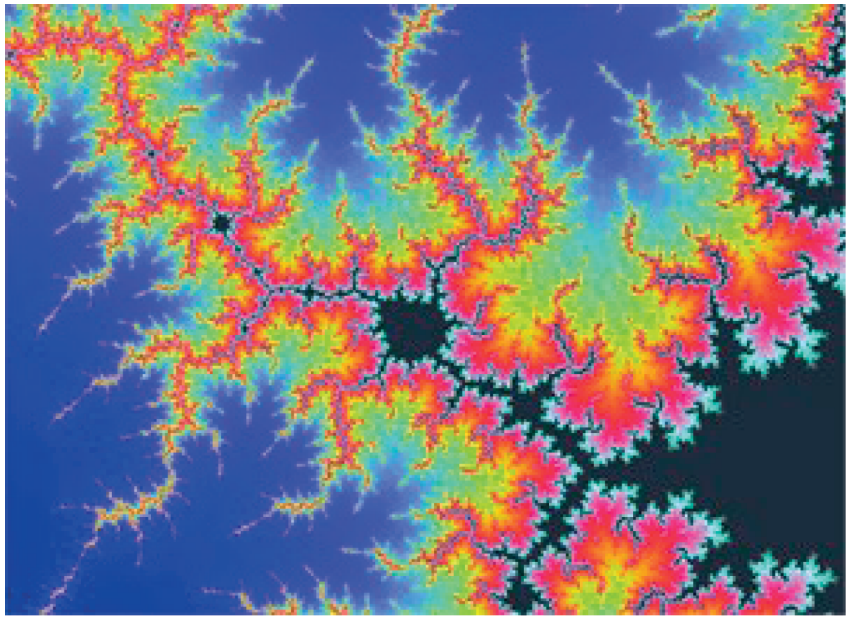

Figure 8. Detail of a Mandelbrot set fractal

and situation types rather than figures – in a satirical cartoon the 'model' may be a single public figure, but it may also be a 'typical' Englishman or Frenchman, or it may reflect a known person as standing for the latter. Clearly this is a scalar matter. The same goes for the degree to which a cartoon figure is abstract$_1$ (disembodied from specific sensory features): this may range from detailed, realistic depictions of actual people to blobby children's cartoon figures barely recognizable as human (or as some particular animal). It is a matter of granularity. This is not true only of static cartoons but also of animated ones approximating extension in real time like a Walt Disney film version of a fairy tale, or a comic book series like "Superman" in which the action is broken up into successive cross-sections, usually of particularly significant or dramatic events. Here the 'granularity' may have little to do with the closeness to sensory reality, since the figures and backgrounds and actions are typically pure fantasy. Nowadays there are even clever blends of 'real' human actors and actions and purely imaginary cartoon characters, backgrounds and animated actions as in the film "Avatar". In fact, through the media we are all our own 'avatars' as we empathize with the characters in the films and documentaries we see on the screen. But what is shown to us is already 'abstracted' to some degree from reality and at the same time simplified – but not necessarily 'generalized across instances' (though episodes of say a crime drama series may match some standard generic 'template').

A television film based on a well-known book, for example, is certainly not anything like an image schema (though it may evoke many in the observer trying to comprehend it), nor is it abstract in anything like the sense in which an abstract word or an abstract painting is abstract. It is nevertheless abstracted from reality in a general sense. One might say in fact that it is abstracted in a hierarchical manner: first the book itself was abstracted from reality (the author's experiences or his imagination, or both), then the script writer adapted this for television by further steps of simplification, and the actors under the direction of the film director converted this into physical actions and spoken words against a background based loosely on the descriptions in the book that the cameraman and/or set designer are responsible for.[1] Finally, the viewer has to abstract his or her own coherent mental model of the story from the succession of images across the screen and the accompanying words. The intermediate level of abstraction, forming meaningful images from a series of rapidly oscillating 'lines' of photons across the screen, is far too rapid for conscious awareness but has nevertheless to be undertaken by the visual system of the brain. But is the final result necessarily far removed from sensory reality (abstract$_1$)? Not so – the viewer's emotional responses at least may be as sharp as if they were themselves directly involved. The truth is, much of our daily experience is vicarious but nonetheless emotionally vivid and capable of evoking quite detailed sensori-somatic 'replays'. This raises a broader question: are our embodied 'simulations' really abstract in either of the two principal senses I have introduced? Somehow we need a broader characterization of 'abstraction' if all the phenomena discussed so far are to be captured by it. In Chapter 11 I shall endeavour to do just that.

Turning to music, this is an art form that is generally understood as going straight to emotional responses, bypassing referential meaning altogether.[2] Music has both form and content like language – but it is in that case a 'language of emotion'. In fact the analogy between musical and linguistic form has been studied in some detail by Lerdahl & Jackendoff (1983), and there are indeed certain striking similarities between especially musical rhythm and linguistic prosody – most obviously so in the case of poetry. Thus the setting of linguistic texts to music involves the alignment of two kinds of abstract metrical grids, both built

1. Of course different film versions can be made from the same literary source – and this is obviously the case with "Alice in Wonderland", some versions being purely cartoons, others with human actors (however fancifully dressed up), others a mixture of the two. Some (the purely cartoon ones?) could be called more abstract than others. I doubt if children experience them as 'abstract' though.

2. Except in forms like opera, in which the referentiality is carried largely by language and/or Wagnerian leitmotivs tangentially standing for people or objects.

of patterns of stress and meter (regular beats); that of the linguistic text is typically distorted by the musical one. Of course language lacks musical pitch/ tonality and spoken language is also much more irregular than music as regards meter. Nevertheless, 'phrasing' is a common feature of both language and at least classical and popular music.

But it is more the content of music that concerns us here. The 'phrasing' of both speech and music is linked to the expression of emotion, and it is not by chance that the same area of the right hemisphere can be shown to be activated in the perception and production of music as in the emotional modulation of speech by prosodic patterns (intonation and stress). Brust (2003: 186–7) discusses the right hemisphere advantage for melody and pitch perception (the 'global' aspects of music). The analytic left hemisphere is typically more active for professional musicians than for laymen (perhaps partly due to their skill at score reading). However, the aesthetic experience of music evidently involves widespread areas of the brain.

Trainor & Schmidt (2003) pinpoint a basic four way split of emotional content of music involving the crossing of two dimensions, each with its distinct neurological distribution: that of approach/ withdrawal and that of intensity. So 'sad' is associated with slow, quiet, legato passages (far from metrical regularity), and 'happy' is associated with high-pitched, fast, staccato passages (close to metrical regularity). 'Joy' as an intense version of 'happy', as is 'fear' of 'sad' (op. cit.: 311). Just as emotional processes in general involve the right hemisphere more strongly than the left, so does music. In particular, the right ventro-medial part of the frontal lobes may be particularly active with more negative emotional content, the left correlate being more positive in valence/ affect (op. cit.: 317). The authors suggest that music may have evolved alongside emotion in order to further communication between children and caretakers. Later, when language begins to afford better control of emotional expression, it may 'go underground' but continue to be in touch with the deepest levels of emotions, allowing them now to be experienced without having to act on them – a vicarious arrangement, as it were (op. cit.: 318).

Music is not just 'about' emotions but directly evokes them – there is ample experimental evidence for 'embodied' autonomous responses, the same or overlapping with those of emotion in general. So music is not 'abstract$_1$' (disembodied) although dissonance (tension) and sadness don't necessarily evoke reactions of withdrawal in the listener, they can indeed be felt as beautiful from a somewhat detached perspective (cf. Trainor & Schmidt: 315 on the neural correlates of dissonance). The resolution of dissonance and its resolution has analogues in many other areas of cognition and bodily action. Another relevant aspect of music is the role of prediction in listening, a matter of 'top-down' projection from memory – we expect the next familiar phrase or can guess it if new but of a recognizable style. Use is made of this especially by classical composers (harmonic progression,

the alternation of phrases analogous to spoken dialogue, the framework of meter, etc.). Music is extended through time and so is our unfolding memory of musical pieces. When something novel or unexpected occurs specific brain activity is stimulated (and a concomitant emotional response ensues) as it is integrated into the holistic time-extended experience.

But is music abstract in any sense at all then? If it is not 'abstract$_1$' could it be 'abstract$_2$' ('simplified or generalized across instances')? Not exactly – the emotions it expresses are related to generalized types of response packed into patterns woven around them, but the structured patterns do not simply 'stand for them' like linguistic or visual symbols. Yet music can certainly be more or less 'detached' from emotion (as we saw in the case of abstract art) – think of the difference between the emotional bombardment of a Bruckner symphony compared with the gentle, vaguely melancholy irony of a Satie piano piece. Is not the latter somewhat more 'abstract$_1$' than the former? And is a tone poem (with at least some imposed referentiality) not somewhat less abstract than a classical sonata or fugue?

Polanyi has stressed the long-standing kinship of the abstract arts (music in particular) with mathematics and scientific theory in general (Polanyi 1958/1974: 193–5). This goes back at least as far as Pythagoras with his relating of musical harmony to the integer ratio of lengths of a struck cord. Also the search through the ages to formulate geometrical rules for harmonious pictorial and architectural structures illustrates this affinity. In common is the concept of aesthetic 'beauty' – of 'order'. According to Polanyi neither music nor mathematics 'denote' anything tangible, they are not 'about' anything, they simply "present their own striking sensuous presence". They are appreciated for the beauty of a set of complex relations embodied in them. Both articulate a vast range of rational relationships "for the mere pleasure of understanding them." This may be somewhat exaggerated as regards mathematics (which, as we shall see in 10.2, does contain its 'practical abstractions'), but his point as to the role 'intellectual passion' has played in both fields is well taken. We shall see in the final chapter how musical expression fits into the more general understanding of 'abstractness' that I shall settle upon.

CHAPTER 8

Neurological underpinnings

The world seen through hidden layers

8.1 Distributed networks and hierarchical structure

It is now time to look as far below the surface as we can get in our search for the roots of abstractness in the human brain. A good place to begin is with the notion of the now widely accepted 'parallel distributed processing' (PDP) view of the brain. This lies behind the connectionist models of recent decades. The basic idea in these computer simulations of the behaviour of neural assemblies in the brain is that there are many simple processing units ('neurons' or computer cells) all connected together in a complex matrix or network. Activation spreads among units according to the strengths of the connections between them, which may be adjusted according to experience (the amount of activation they receive from other units). This is parallel to the way neurons connect and enhance or inhibit each other's activation levels by 'Hebbian' learning ("neurons that fire together wire together"). Such systems are self-organizing, and can be said to learn – 'intelligent' behaviour (such as pattern or word recognition) gradually emerges through recurrent waves of activation that bring input units to match output units via 'hidden' units between them. These relate input to output patterns but do not themselves display a 'readable' encoding of either. Information within the system is thus widely distributed across units and the constant adjustment of their strengths occurs in a parallel fashion. Some of the earliest models of this kind are described in McClelland & Rumelhart (1988). Among the various algorithms employed are 'back-propagation', which works by successively reducing the error between the desired output of the network and the actual output at any stage.[1]

More recent versions include 'recurrent' connections from a hidden layer to a 'context' layer, which enables the system to learn to encode sequential information with the addition of 'context units' that feed the result of previous processing back into the internal representation. This provides the system with a kind of dynamic

1. It should be pointed out that there is no evidence (as yet) of an exact analog of back propagation in the human brain. However there is ample anatomical evidence of feed-forward as well as feed-back connectivity there.

memory (Elman et al. 1998: 119–20). Still more sophisticated simulations attempt to model such specific neural operations as the simultaneous object recognition and position tracking by the brain's dual visual streams, as in Mareschal et al. (1995). This is schematized in Figure 9 from Elman et al. (op. cit.: 154). Note the 'hidden' layer of cells between the input 'retina' and output from the 'response module'. This can be regarded as a series of commands to the child's motor responses for reaching towards the object. The system models the way in which an infant might take account of the details of an object when reaching for it and how this information can be coordinated with the infant's representation of the object's position. The network is 'recurrent', continually anticipating the future position of the (hypothetical) object based on the immediately preceding state of the system. Whether a reaching response was initiated or not depended on the identity of this object – whether it was pre-defined as very hot (too hot to handle) or cold. The network was required to learn to determine the object's identity wherever

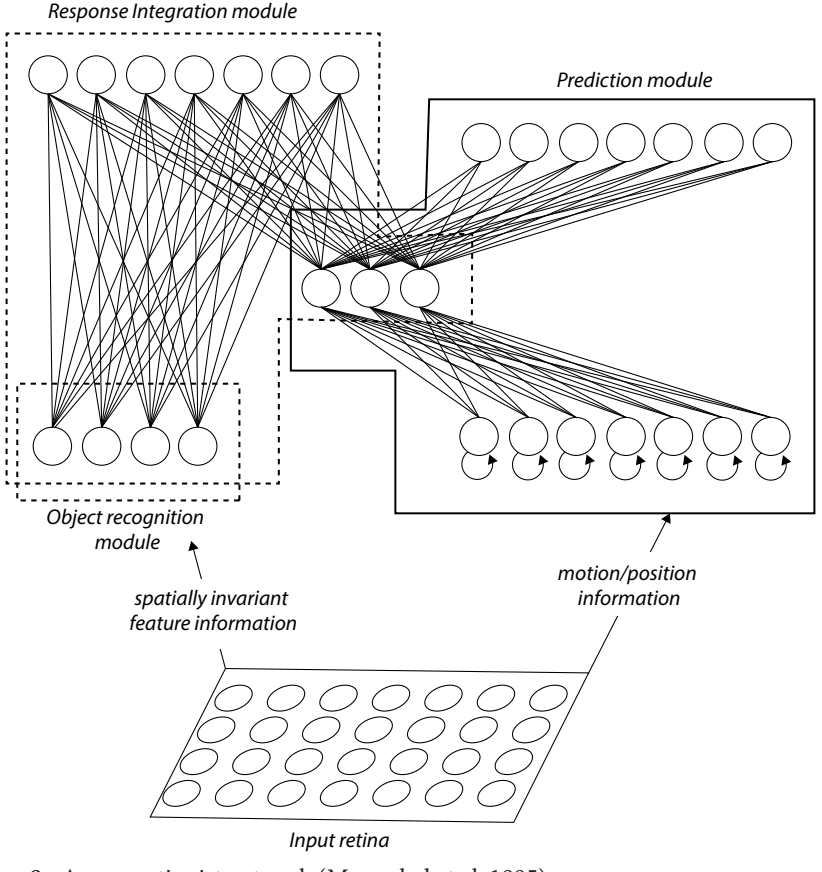

Figure 9. A connectionist network (Mareschal et al. 1995)

it occurred on the input 'retina' and to learn to predict the next position of the object on the retina. The two channels on the model correspond to the 'what' and the 'where' pathways in the brain's visual system. The 'where' channel kept track of the position of the object as it moved. The 'what' channel (the object recognition module in Figure 9) mapped the input to a set of complex cells that identified particular configurations of features in the input. (This allowed for up to four different features being activated.) The system used a learning algorithm to learn to recognize object invariances (wherever a coherent 'object' appeared in the input 'retina'). The results (the number of repeated passes or 'epochs' through the system needed to approach reliability) were impressive, but the model stopped short at the point of simulating actual motor responses.

Common to all such systems is the bypassing of the need for explicit 'rules' (e.g. the syntactic or phonological rules of generative linguistics) – these emerge as statistical generalizations across the input data. In other words by abstraction. These models are most successful at the phonological and word level. When applied to the extraction of meaning from linguistic input (typically pre-analysed into phonological units) it should be borne in mind that the semantic units employed in these systems – termed 'microfeatures' – do not necessarily correspond to anything overtly linguistic (i.e. correspond to words referring to them). They simply mediate between perceptual and linguistic units. Word meanings are 'attractors' or 'valley bottoms' in the landscape of semantic memory seen as one vast connectionist network. This is analogous to an associative or 'content-addressed access' memory system in recent computer design which compares input search data against stored data and returns a match. Whether such systems can ever be truly 'embodied' and in any sense understand the meaning of the distributionally extracted patterns they recognize is an open question. It may well be that at the higher levels of syntax and sentence-processing models operating with holistic symbols will be better suited and supplement lower level connectionist models, but it is nevertheless largely through lower level sensory processing that we 'see' the world around us. There is an intermediate 'functional' level between the two according to Jackendoff, namely his level of 'f-mind', of which only a small part emerges into consciousness. This is the level where he sees most linguistic processing takes place and crucially – unlike its neural 'computational' substrate – it is combinatorially organized (Jackendoff 2002: 21). I shall return to this in greater detail in 11.2. For more on networks and 'neuronal grammar' in the human brain see Pulvermüller (2002).

Continuing on our trajectory deeper into the organisation of the brain it should first be realized that the brain is far more complex than connectionist models might suggest. In particular the dynamics of the cortex depends on neural 'assemblies' rather than single neurons acting independently (cf. Hebb 1949). These

can be defined as 'synchronized oscillatory populations' (Singer 2013: 623), organised in a laminar fashion. Singer further describes the brain as a whole as a highly active pattern-generating system rather than just a stimulus-driven device, and more specifically as "a complex self-organised system with non-linear dynamics in which principles of distributed, parallel processing coexist with serial operations within highly interconnected networks" (op. cit.: 616–7).

Before going further into Singer's review of dynamic cortical systems (in the following section) let us recap what has been said already in 6.1 on dual stream models. According to Hickok & Poeppel's model of linguistic processing something similar to the dual streaming of visual perception is seen in the processing of linguistic input (Hickok & Poeppel 2007). The analysis of incoming verbal input splits after an early bilateral stage of acoustic 'spectrotemporal' analysis (close to primary auditory cortex in the superior temporal lobes) into a left hemisphere dorsal stream and bilateral ventral streams, the former mapping sound onto articulation/ speech production and the latter turning sound into meaning. They further hypothesize that there are parallel early stage routes in the mapping from acoustic input to phonological lexical representation according to time scale, one pathway sampling the input at a relatively fast ('gamma band') rate appropriate for segment-level information, the other (principally right hemisphere) at a slower ('theta band') rate for syllable level information (op. cit.: 398). The model posits a 'combinatorial network' specifically located in anterior regions of the left temporal lobe, communicating with Broca's region in the adjacent inferior temporal lobe (op. cit.: 395). It also presupposes – like nearly all models today – a widely distributed conceptual network (as discussed in 3.2).

Friederici's model of speech processing, introduced already in 2.2., makes explicit the time course from way-station to way-station, as indicated by well-known ERP signals (Friederici 2002). Prosodic modulation by the right hemisphere supplements linguistic expression with "what cannot be said in words alone" (like gesticulation). Friederici & Gierhan (2013) present a more detailed variant of the dual pathway model that involves at least two dorsal and two ventral streams. Two distinct areas that have been lumped together traditionally as 'Broca's area' (Brodman's areas 44 and 45) can be distinguished on this model according to their role in separate streams, and the separation of syntactic and semantic process into frontal and temporal regions respectively is shown to be over-simplified, with overlapping functions in both. The two dorsal routes support auditory-to-motor mapping and the processing of syntactically complex sentences respectively, and the ventral routes subserve on the one hand semantic processes and on the other basic syntactic processes (op. cit.: 252).

What of the hierarchical organization of the cortex? An influential approach is in terms of what Damasio calls 'convergence zones' (cf. Damasio & Damasio

1992: 65). Neural assemblies here code for dispositions such that modally specific areas of cortex simultaneously active in multi-modal perceptual events can re-activate the primary areas, thus recreating the multimodal experience. 'Convergence zones' are located in several distinct areas in the temporal, parietal and frontal lobes. Another way of looking at this is in terms of more locally restricted 'hubs'. Thus Patterson et al. (2007) present a 'distributed-plus-hub' model, where the hub is amodal, similar to Damasio's individual convergence zones for relating specific pairs of modal inputs but now located in one bilateral location in the extreme anterior temporal lobes. It allows for analogizing between modalities, recognizing categories across perceptual features (e.g. all fruit, despite their differences in colour and shape, etc.), and corresponds to a connectionist 'hidden layer' – being amodal and distributed it is reminiscent of a hologram. Binder & Desai (2011) propose – as mentioned in 5.3 – an 'embodied abstraction' approach combining modal and supramodal representations similar to Damasio's and contrasting with Patterson et al.'s unitary hub model. Also Pulvermüller opts for a multimodal (as opposed to abstract amodal) convergence zone model, with medial stations in long-distance cortico-cortical connections formed by correlation learning (Pulvermüller 2013: 459–460). Binder & Desai (op. cit.: 529) pinpoint temporal pole involvement in high-level representations of emotion and social concepts with an emotional content.

Whichever variant is nearest to the truth, it is clear that higher level control of complex organisms like the human body requires a high degree of abstraction and simplification in the control hierarchy, so that a simple conscious 'intention' at the highest level results in the rapid and unconscious cascading of complex lower level realizations of that intention. Working memory is assumed to be the highest level (conscious) control centre. Its temporal focus or 'duration' is of course short, but it must have access to long-term dispositions and plans stored in frontal cortex.

What is the role of the single neuron in all this? Whatever the specifics of that role are (and they can be multiple, a single neuron participating in numerous distinct networks), the individual neuron is only functionally 'meaningful' within larger neural assemblies that activate it to varying degrees (or not at all) depending on context. It is these that actively compare and abstract information in the process of building mental schemas – the "neural averaging of similar inputs" as Crick (1994: 185) puts it. These assemblies constitute the 'hidden layers' through which the external world as received from the sense organs is experienced. With the neural assembly we've probably reached the downward limits of the abstraction of meaningful patterns. The single neuron is just an input-output automaton. Meaning is emergent, it requires a context and someone or something for which it is meaningful. Only then can abstraction begin.

8.2 The binding problem and neural oscillation

This brings us to the question as to how separate paths of perceptual or linguistic analysis of a given complex input are reunified as holistic experiences – the so-called 'binding problem'. Recent research points more and more in the direction of synchronized network oscillation (involving the thalamus as well as the cortex) lying at the heart of what is going on here. Arnal et al. (2012) summarize this in their study of the oscillations that can be detected (by EEG measurements) in 'predictive timing' – 'when' information extractable from a signal distinct from predictive coding for 'what' information. The basic idea is that incoming multimodal input is 'phase-locked' with continuous on-going oscillations, aligning the stimulus with optimal peaks of the oscillations. These are carried by slow delta-theta waves involved in "setting temporal windows of sensory integration due to predictive mechanical entrainment" (a kind of underlying biological 'meter' relative to incoming stimuli). Other oscillatory frequency bands have different functions, thus the alpha band is probably to be associated with inhibitory mechanisms gating sensory information as a function of its cognitive relevance, and the beta band may be related to motor function and the timing of beats. Particularly significant for their study is gamma band activity, which increases as attention increases, whereas it decreases if a stimulus is expected (op. cit.: 392). Beta band activity is probably also crucial, associated with descending activation from higher to lower neural layers via feedback – it is traceable in prediction error generation. Gamma band activity is probably associated on the other hand with ascending (feedforward) activity – new neuronal assemblies are formed at this rhythm, whereas beta rhythm changes them into new patterns "via a rebound form inhibition" (op. cit.: 394–396). Basal oscillation rate increases in auditory cortex to allow for the speed of a speaker being attended to. 'Entrained' (synchronized) oscillations may become predictive by creating periodic temporal windows that higher-order regions can rely upon to read out encoded information (op. cit.: 391). Their model would seem to go a long way towards explaining the ERP responses in the models of Friederici (2002) and others.

As Hagoort (2005) has pointed out, the binding problem for language (at least at higher levels) is slower than for vision. Without drawing neurological oscillations into the picture, he analyses it in terms of 'unification'. Operations of this kind combine words into higher units, at the semantic (including referential) and phonological as well as syntactic processing levels. He proposes Brodmann's areas 47 and 45 as involved in semantic unification, areas 45 and 44 in syntactic unification, and areas 45 and (premotor) 6 in phonological unification, all in the left inferior frontal gyrus. His approach, by linking up lexical frames with identical roots and nodes and checking agreement (for number, gender, person,

etc.), is compatible with suggestions made by a number of prominent linguists. These include Chomsky's MERGE (as mentioned in 2.3) and Jackendoff's UNIFY (Jackendoff 2002: 164). Hagoort points out in support of his theory that the prefrontal cortex is good at holding on to information for processing temporally extended input 'on-line'.[2]

In his review of non-linear cortical dynamics Singer (2013) promotes the 'reservoir' (or 'echo-state' or 'liquid') model of the brain's computational activity. These recent kinds of model allow both for the temporal coordination of distributed processes and for the binding of local processes into coherent global states. They posit reciprocally coupled networks comprising "self-active non-linear units with random couplings that maintain their own dynamics and are engaged in active processing" (op. cit.: 619). Simple input is processed in multiple parallel sub-regions in a dynamic system displaying higher dimensionality. This 'high dimensionality' allows for the coexistence of a large number of potentially realisable states implemented by highly interconnected assemblies of (3-dimensional!) neural assemblies operating through time. Reverberations (or 'ripples') extending over a finite time before decaying facilitate memory for sequences. A unified output of the reservoir can in turn be recognized by a linear (or non-linear) readout 'classifier' assembly, optimally activated by a certain state of the 'liquid'. This allows for the storing of huge numbers of memories accessible at extremely high speed since many sub-states of extracted information can be simultaneously superimposed in parallel search and matching operations. Laminar structures such as neural layers are particularly suited to form map-like representations of concrete (spatial/ sensory) contents but also of abstract categories, and within the reservoir model they can be active in parallel at different frequency ranges (gamma vs. alpha and beta bands). Different kinds of synchrony are evoked by temporally structured stimuli and may function like a self-organizing 'glue' for neural assemblies. Inhibition by interneurons such as 'basket' cells probably have an important role here too, supporting coordination "by synchronising oscillating cell clusters at the appropriate phase … because cells embedded in oscillating circuits are exposed to precisely timed periodically repeating volleys of inhibition" (op. cit.: 618).

There is evidence for dynamic cortical architecture of this kind already detectable in the brain's 'resting state' activity (Singer 2013: 620). When exposed to simple stimuli, neurons form well-defined receptive fields resulting from the selective recombination of input connections and oscillations – an emergent property of recurrent networks. Singer's hypotheses relating to his model include these: "input

2. He also discusses more briefly distinct Memory (lexical-syntactic) and Control aspects of linguistic processing, the former involving left posterior superior temporal cortex, the latter dorsolateral prefrontal cortex.

signals cause selective stabilization of substrates that are often distinguished by enhanced coherence and reduced dimensionality; and reduced dimensionality and temporal coherence increases their saliency, promoting long-term stabilisation (in memory) by Hebbian modifications of synaptic connections" (op. cit.: 623).

8.3 Abstract words in the mental lexicon

Backing up a little from the deep neurological level let's return to the nature of abstract lexical items and their implementation in the cortex, taking up the thread from Barsalou & Wiemer-Hastings (2005), whose discussion of the complexity and variability of abstract words in the mental lexicon was touched upon in Chapter 2. Evidence has accumulated – starting with Kounios et al. (1994) – that abstract words are actually processed differently from concrete ones in the mental lexicon. Thus Kiehl et al. (1999) indicate with fMRI evidence a right hemisphere bias for abstract words, these specifically being implicated in the anterior part of the right superior temporal gyrus and the right inferior frontal gyrus, whereas concrete words are mostly processed in the left hemisphere – although semantic processing is likely to involve both hemispheres. This supports Paivio's dual code theory, in which abstract words are essentially coded in verbal representation only while concrete ones are coded both verbally and in perceptual imagery.

Pulvermüller (2013) proposes four distinct semantic mechanisms: referential semantics (between symbols and the objects and actions they refer to), combinatorial semantics (enabling the learning of symbolic meaning from context), emotional-affective semantics, and abstraction mechanisms for generalizing over a wide range of instances of semantic meaning (op. cit.: 458). These are seen as instantiated respectively in sensorimotor areas, limbic areas (and associated action areas), the periphery of the perisylvian region, and – significantly for abstract lexical items – multimodal convergence areas (op. cit.: 468). He stresses the importance of context for abstract meanings, and suggests that convergence zones (in particular those in inferior frontal cortex and superior temporal cortex) may support the learning of word combinations in context. Wary of purely 'embodied' models of the mental lexicon, he argues from a neurolinguistic point of view for complementary 'disembodied' and 'embodied' levels of meaning in interplay in the brain. He discusses evidence for abstract concepts activating frontal regions, areas known to coordinate multimodal information (op. cit.: 464–68). Abstract meaning is for him "generalized over variable concrete instantiations" – this is germane both to the Binder & Desai model sketched in Chapter 5 and to Barsalou et al.'s (1993) concept of 'transcendence', whereby knowledge of both individuals and types may become 'functionally detached' from the world model by being

instantiated in different spatio-temporal locations (where individuals or types are represented).

Abstract words typically contain significant emotional components and will thus display connections to cortical and limbic emotion centres (cf. Vigliocco et al. 2013). In Fortescue (2009: 54–7) the abstract$_1$ words 'liberty' and 'love' are analysed in terms of distributed networks with limbic 'tails'. Abstract emotion words in particular may be associated with prototypical actions. Pulvermüller suggests that such words may activate face and arm areas as well as the limbic system with its strong emotional content (Pulvermüller 2013: 464–66). Also idiomatic meanings of constructions may be more detached than their corresponding literal meanings, with stronger activation in multimodal convergence areas. Thus the processing of an expression like 'he cooked her goose' would appear to involve, simultaneously, distinct areas of the cortex for the 'disembodied' idiomatic meaning and for the concrete meanings of the individual words which constitute the expression – this goes counter to the view that idiomatic meanings are extracted as a subsequent step to literal processing.

Pulvermüller takes all word meanings as being instantiated in distributed networks, but 'abstract$_1$' (disembodied) meaning for him crucially involves the multimodal association areas (dorsolateral prefrontal, inferior parietal and anterior temporal regions of the cortex). Note the distinction between 'multimodal' and 'amodal': the former may be interpreted as referring to embodiment in concrete experiential situations involving sensory features of varying modality, unlike the latter, which is more overtly 'abstract$_1$'. Multimodal representations may nevertheless be 'abstract$_2$', i.e. 'schematized', simplified through juxtaposition and summation. It has long been known that the meanings of words referring to concrete sensory (including motor) experiences are indeed represented principally in or close to the corresponding sensorimotor regions of the cortex, but the location of abstract$_1$ meanings has remained less clear. There is however evidence of greater involvement of the inferior frontal gyrus and middle temporal gyrus in abstract concepts (and recall the evidence from Kiehl et al. above on right hemisphere involvement). Pulvermüller's suggestion is that what distinguishes abstract meanings in general is their weak, more diffuse and variable associations with sensorimotor experience, in a family resemblance pattern. (This supports Bersalou & Wiemer-Hastings' position.) These associations can be variously distributed, either (oversimplifying somewhat) in perceptually based cortex – e.g. in the case of 'beauty' – or in action-based cortex – e.g. in the case of 'free' or 'game' (op. cit.: 467). The boundary between 'abstract' and 'concrete' words is certainly not a clear-cut one.

8.4 Distributed context

Essential to the meaning of most abstract words and even many concrete ones is the context in which they are applied, and this is recognized in 'interactive' connectionist models that allow top-down influence – including contextually-based expectations – on processing from the outset. We have seen that the semantic neural representation of abstract words is highly distributed in the cortex, but this is likely to apply even more to the representation of contexts (whether static 'frames' or dynamic 'scripts' and 'scenarios'). If these are so diffuse how can they actually be involved in supporting the extraction of coherent meaning from sequences of words? There must be some way in which contexts, however diffuse and variable, are summated into discrete 'nodes' in the same way as the core meaning of individual words can be summated as abstract 'nodes' (correlating 'lexical concepts'). The way I suggested doing this in Fortescue (2009: 153–160) was in terms of 'mediatory columns' (neural assemblies) in right hemisphere associative cortex linking both concrete and (especially) abstract words with 'scenarios'. These are seen as correlating broadly distributed contents, overlapping with or containing others, but distinct enough to have determinate relationships with corresponding words stored in the left hemisphere. The model assumed no sharp break between episodic (experiential) and semantic memory – the latter is just 'averaged' over many occasions or else verbally anchored (by definition plus context). More specifically (and tentatively), I suggested a 'Contextual Symmetry' hypothesis, which posited right hemisphere 'macro-functional' scenarios symmetrically associated via cross-commissural links to corresponding left hemisphere phonological word representations and functioning as constraints on the semantics of the latter, thus ensuring harmony between the two hemispheres in engaging in everyday joint tasks that involve both perception and language. Recall from 5.2 the claim by Ferstl et al. (2005) that the right hemisphere is crucially involved in mental models.

Some of the specific 'scenarios' I analysed were 'SHOOTING GAME' (containing links to 'GUN', 'GAME ANIMAL', etc., and overlapping with other 'killing' scenarios such as 'MURDER' and 'ASSASSINATE'), 'TRUTH' and 'REALITY', 'CLAIM' (contained in the general scenario associated with 'say', i.e. sharing its semantic core), 'MUST', 'APPARENTLY', 'CATTLE' (within 'FARM ANIMALS', in turn within 'FARM'), and 'HOUSE', associated with scenarios of moving round the rooms of the house during daily activities. Also analysed was 'HOLIDAY', as described already in 4.4. I reproduce the diagram from Fortescue (2009: 56) in Figure 10. The format here represents the relationship between different levels of the mediatory column that associates the lexical concept concerned and the scenarios or frames representing its appropriate contexts of use. This needs a few words of explanation.

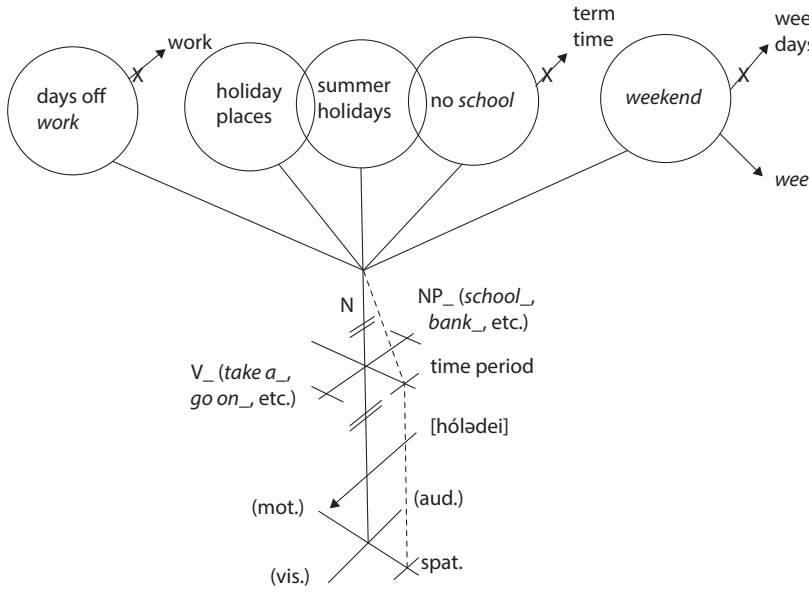

Figure 10. Contextual scenarios for the word 'holiday' (Fortescue 2009)

The lower half of the diagram schematizes the organization (in the left hemisphere) of the mediatory word column for [hóləʊdei], with input at mid-column from Wernicke's area already analysed phonetically. This reflects perceptual input from the thalamus to the 'granular' layer (level 4) of cortical columns.[3] Below it are the orthogonal axes for relevant sensory features (visual, auditory, motion and spatial respectively) – these have been left empty here since any sensory features needed for simulations of this 'abstract' word will be diffusely associated with the scenarios themselves (cf. Pulvermüller's suggestion in the preceding section). Further up in the column are links to the grammatically relevant features of the lexeme (a noun), indicated again on orthogonal axes, this time 'syntagmatic' as opposed to 'paradigmatic'. The former indicates the grammatical context the word is to be found in, respectively following a verb like 'take' or 'go on' or as the head of a noun phrase (e.g. 'school holiday') – I ignore the necessity of an indefinite article here. On the paradigmatic axis, the word belongs to the category 'time period'. At the top (before the branching of the 'balloons') are the long-distance links (via white matter fibres) to the overlapping contralateral areas where the approximately

3. See Mountcastle (1998: 165) as regards cortical columns. These are constituted by many 'mini-columns' of pyramidal cells bound together by dense, short-range horizontal connections. Each mini-column contains 80 to 300 neurons (Mountcastle op. cit.: 165). My 'mediatory columns' are cortical columns in associative 'convergence zones' that bind words to their distributed meanings, i.e. mediate between words and percepts and (more abstract) concepts.

five scenarios mentioned in 4.4. are taken to be located. The broken line connecting the paradigmatic feature 'time period' to the potential sensory dimension of space below reflects the construal of extension through time as analogous to the extension of physical objects in space. Positive links (much simplified) to further scenarios and frames (e.g. from 'weekend' to 'week') are indicated by arrows. Negative relations are indicated by a crossed out arrow, as between 'weekend' and 'week days' (which is itself included within the 'week' frame). How these scenarios can be distinguished one from the other, despite overlaps, will be returned to in 10.3. Essentially, irrelevant ones will be suppressed by context.

In Fortescue (2010a: 185–6) I analysed the word 'conquer' using the same format. This is a somewhat more complex case which involves two overlapping contralateral scenarios, one closely related to that for 'defeat' and constituted at its most abstract by a fight, a winner who beats his opponent, and a loser who acknowledges defeat. The other is more specific to the historical context in which the word is normally used, where there is a 'victor' (corresponding to the winner), a 'vanquished' (corresponding to the loser – either individuals or whole peoples/armies) who acknowledges defeat, and the conditions of surrender exacted by the former from the latter. The word is stylistically marked thanks to its association with a further frame (the broader context of 'historical narrative'). Specific problems arise with attempts to decompose the overall meaning of the word to semantic primitives, for instance with the feature 'acknowledges defeat': there are various ways in which defeat can be acknowledged depending on the type of 'fight' involved. In fact defeat does not always have to be acknowledged at all, just accepted. Yet this is hardly just a 'typicality' feature (in Jackendoff's sense) – it is crucial to the correct use of the word 'conquer'.

Some indirect supporting evidence for the Contextual Symmetry hypothesis is to be seen in Martin & Weisberg (2003). In an fMRI experiment they presented subjects with animated sequences of geometrical shapes that could be interpreted either as familiar social 'vignettes' (such as playing baseball, dancing, fishing, sharing, or scaring somebody) or as mechanical processes involving tools or machines. There was a clear right hemisphere preference for the former kind of presentation, in areas known to be active in the recognition of animate beings and human faces, including 'biological motion', but also in the amygdala, known to be involved in 'social' emotion. The authors speculated that such networks (corresponding to my 'macro-functional' scenarios) could have evolved from more basic image schemas for perceiving and knowing about animacy and agency.

A recent neurolinguistic approach to the role of context in comprehension is Skipper's (2015) NOLB model. This envisages multiple synchronous context

networks (or 'cnets') that cooperate and compete.[4] Such 'cnets' can be used to predict forthcoming sounds, words, phrases, etc. associated with the particular kind of context, all of which can constrain the interpretation of a speech event. Some contexts are exogenous – physical situations of language use, mouth movements, gestures, etc. – and some are endogenous – learned experiences/ memories, scripts, schemas, beliefs, etc. (op. cit.: 103–4). A 'cnet' network could in theory speed up processing by constraining language ambiguity (seen as widespread on all levels). Skipper claims that there is good evidence already for a plethora of contextual sub-networks (op. cit.: 118). The model is thus a generator of multiple level predictions, with analogs in the areas of vision, memory, etc., widely distributed in the brain. The model has not yet been elaborated to the point where it can describe how 'cnets' actually interact in detail, but hopefully that will be forthcoming. Such a model brings up again the point I raised in Chapter 1 as to how the 'brain as predictor' is related to the 'brain as abstractor'. Let me repeat here that abstraction and prediction are in fact two sides of the same coin since any prediction affecting future behaviour must perforce be based on abstraction from previous experience.

In 2.1 I presented a list of the principal kinds of 'context' relevant to interpreting a single sentence in English. I repeat these in Table 4, adding in parentheses my best guesses, given the neurolinguistic evidence that I have presented *en route*, as to the approximate cortical regions where these are localized. The first three are specific to the left hemisphere (in most speakers).

Table 4. Linguistic contexts and probable cortical locations

a)	phonological context in the recognition or production of successive phonemes (in the classical perisylvian speech centres)
b)	individual words – the semantic fields to which they belong (in the medial temporal lobe and angular gyrus, also premotor frontal cortex for action words)
c)	syntactic/ morphological constructions individual words are embedded in – including 'case' frames of individual lexemes (in Broca's area and inferior parietal lobes)
d)	textual context – including spoken discourse context involved in on-line mental model construction (in right hemisphere parietal and frontal cortex)
e)	the physical setting of the communicative situation (in the parietal lobes)
f)	speech act conditions – illocutionary intentions (in right hemisphere prefrontal cortex)

(*continued*)

4. It is complementary to the 'competition model' of Bates & MacWhinney (1982), which concerns production rather than comprehension.

5. This includes as regards attitude such things as irony, honesty, joking, confiding, comforting, and mocking, etc.

Table 4. (*continued*)

g)	facial expressions and gestures of interlocutor (in anterior right hemisphere, linked to the limbic system)
h)	attitude or intention towards interlocutor or subject – including perlocutionary intentions (in right hemisphere prefrontal cortex and the limbic system)[5]
i)	stylistic context or genre (in frontal cortex, right hemisphere)

This list can be extended with the following 'macro-functional' scenario types discussed in the previous section. Such entities are taken, as mentioned, to be mediated by nodes in right hemisphere convergence zones that bind them into coherent sensorimotor networks.

Table 5. Scenarios and probable cortical locations

j)	static 'frames' (in right hemisphere parietal cortex)
k)	dynamic action schemas (in right hemisphere premotor cortex)
l)	socio-cultural 'scenarios' (in right hemisphere prefrontal cortex)

Consider how these contexts are relevant to the reader's understanding of Alice's remark to the Cheshire Cat: *"I wish you wouldn't keep appearing and vanishing so suddenly: you make one quite dizzy."* Phonological context (a) is involved both in reading and hearing the individual words as they occur in sequence, aiding low-level predictions of what is to come. The semantic context in which they typically belong (b) is relevant to disambiguating words such as 'keep' which may either be an aspectual auxiliary or a full lexical word (here the former), or 'quite' (which can either be an intensifier – as here – or a 'dampener'). Context type (c), knowledge of the valency (argument structure) of individual words – especially verbs – is relevant to hearing words like 'wish' and 'make (me)', which evoke expectations of respectively a whole proposition or a predicate complement. The textual context (d) is the on-going story (the reader's mental model of it) – at this point Alice has already encountered the cat in the kitchen of the Duchess and also, shortly after she left, on the bough of a tree; an interesting exchange ensued during which the cat vanished and appeared again a couple of times; half expecting to see it again at any time she wandered on (heading for the March Hare's house) when the present incident occurred, the cat appearing suddenly on the bough of another tree. Context type (e) is the reader's own physical surroundings (this might be relevant to his or her recall of the story), and type (f) (speech act conditions) is relevant to our understanding of the exchange between Alice and the cat – she is expressing a wish, complying with the illocutionary conditions defining that act (wanting something to happen and knowing that the cat is in a position to comply, etc.). The facial expression of the cat (context type (g)) is to be a salient point in the

ensuing exchange – earlier on she had thought upon noticing its strange grin that "*It looked good-natured, but still it had very long claws and a great many teeth, so she felt it ought to be treated with respect*." This relates further to context type (h), Alice's general attitude towards the cat (a mixture of caution and curiosity), which continues throughout the present exchange. And stylistic context (i) – a mid-nineteenth century tale primarily for children – allows rapid understanding that 'one' (as in "makes one quite dizzy") is a natural way of referring to herself for a middle-class girl of the time.[6]

As for the three types of 'macro-contexts', there are both static frames and dynamic scenarios (or 'scripts') involved in the passage. The former lets us understand for example that the action takes place in a wood – but one constrained by being not too dense (so that something perched on a low bough is visible from the ground). There is a winding path leading through it – this is a schematic wood into which further compatible details like this can be filled in. A more dynamic context is that of Alice's walk through the woods, interrupted by a sight that causes her to halt. Finally, the conventional socio-cultural scenario of holding a dialogue with an interlocutor (all be it a strange one) is also schematic, filled in with some at least partially conventional content.

Let me end this chapter with a question: how could contexts like these be realized in a multidimensional 'reservoir' type of memory network of the kind that Singer (2013) describes? This still needs to be investigated in neurological detail, but presumably an answer would involve numerous parallel 'sub-states' or 'attractors' that respond to and constrain the overall meaning abstracted from the textual input. The output would be a holistic experiential simulation, a reconstructed sequence of self-coherent events – in effect a mental model built up with the aid of lexical scaffolding of the kind schematized by diagrams such as that in Figure 10. These may be understood as 'templates' for the elements of such simulations, only certain of the associated scenarios of which would be selected, others being suppressed by mutual contextual effects at a higher level of conformal 'blending' in the 'reservoir'.

6. One does not have to weigh up whether it might refer to quantitative 'one' (e.g. one 'quite dizzy' person or thing). The syntactic closure would in any case disambiguate this.

CHAPTER 9

Abstraction at work

It's child's play

9.1 The universal and the language-specific in language acquisition

The preceding chapter makes abstraction look like something exceedingly complex and non-linear (which it may be) but even young children do it with ease – better than adults in fact when it comes to learning a language. If one ignores the innate component in language acquisition, which most functional and cognitive linguists today regard as a matter of general dispositions rather than specific structural 'expectations', then the whole process of acquiring a language is one of abstracting it from the environment of speech that the child is exposed to in the first years of its life (I shall return to the usage-based perspective in 9.3). For a succinct, up-to-date account of the developmental sequences children go through, see Harley (2008: 103–152). De Villiers & de Villiers (1985) present a detailed account of the order of semantic and syntactic development in the acquisition of specifically English. This starts with babbling and the step-wise acquisition of phonemes, specific contrastive combinations of features permitted by the configuration of the speech organs when they are active. Vowels appear before consonants, stops before other consonants, labial consonants first, dentals last, and frequent sounds are mastered before infrequent ones. Already at this stage comprehension precedes production, so the child will typically have abstracted patterns from the speech heard around it that it can match with the sounds spontaneously produced in babbling – the former set of contrasting sounds that will be reinforced while others fall away.[1] On the way to complete mastery of the sound system of English, the child will produce various kinds of simplifications and approximations.

This gives way to a 'holophrastic' stage (single word utterances), lasting from about one to one and a half years of age, where the first words arise "out of situations where an exemplar of the category referred to by the word is present in view of parent and child" (Harley 2008: 127). To overcome the 'mapping problem' (what exactly words refer to in the environment) the child probably has some innate principles to help, e.g. the assumption that words/ names refer to whole objects

1. But see Harley (op. cit.: 135) for a caveat as to comprehension *always* preceding production.

not just their parts, as well as parallel attention on the part of the caretaker using a particular word or phrase. Children start using words at the semantic 'basic level' of categories (as discussed in 3.2) – superordinate words are more difficult to acquire, and learning abstract words like 'know', 'believe', etc. will come much later, requiring both exposure to their use and a fully developed 'theory of mind', i.e. the ability to recognize other people's intentions and desires. The role of imitation is limited and that of response to correction by adults hardly observed at all. However, the much simplified and exaggerated speech of caretakers to children ('Motherese') certainly has a facilitating role to play in the child's gradual communicative convergence with the adult world.

I should add a few further remarks about 'theory of mind' at this point. Nativists such as Fodor see it as innate, in the form of a 'belief-desire' module which should in theory be available from birth (Fodor 1992), but as Meltzoff argues, there must be a developmental component beyond whatever is innate here, with the ability to represent beliefs (as opposed to simple desires and physical intentions-in-action) requiring further development (Meltzoff 1995: 846–7). Most investigations of theory of mind in young children involve children already able to report on their experiences verbally. Typically they go through an early 'desire psychology' stage (recognizing desires in others), and only later reach a 'belief psychology' stage (involving representations of the intentions and beliefs of others). Aiming to ascertain whether pre-verbal children are already able to distinguish the actions of humans intentions from purely physical actions, Meltzoff devised a test in which young children aged 18 months were required to re-enact actions they observed of an adult attempting to fit different objects together but failing to carry them through to their conclusion. The children were able to infer the intentions and goals of the adult's behaviour and carry out the target action correctly, but they could not do this in a parallel test where actions analogous to the adult's were carried out by an inanimate object (two wooden prongs). In other words, completing the task presupposed recognizing that other human beings have intentions and feelings similar to their own. They know that doing purposeful things is an attribute of persons not things. Already at this young age more was required than just imitating physical actions (something new-borns can do), though the children concerned were still presumably well short of being able to ascribe adult-like belief states to others. As I suggested in 7.1., recognizing the emotional state of other children may be the pivot necessary for the initial development of theory of mind beyond the mere ability to imitate to recognizing desires – e.g. recognizing happy behaviour when a desire is fulfilled or frustration when it is not.

The order of acquisition of words is generally from semantically less to more complex, thus as regards dimension terms, 'big' and 'small' and 'tall' and 'short'

are learnt before 'thick' and 'thin' or 'deep' and 'shallow'. Verbs are more difficult to learn than nouns because of their reliance on complex linguistic contexts (op. cit.: 135). Tomasello (1992) describes how children profit maximally from their 'first verbs' (early verbs forming 'verb islands') to learn about constructions and to extend their knowledge of how verbs – the semantic and syntactic core of utterances – function in general. Bates & MacWhinney (1982) describe how children are able to assimilate abstract nouns on the basis of what they know already of concrete ones behaving in a certain way – by analogy and profiting from distributional criteria. They learn for example that all 'nouns' act in a similar way and are based on the semantic category of 'object' or 'thing', extendable to less perceptible concepts. A similar point is made by Feldman from within the Cognitive Linguistics paradigm – he sees metaphor as the means whereby children learn to understand and conceptualize the more abstract and cultural concepts needed in adult human society by linking them to basic subjective experience (Feldman 2008: 203). As is well known, children will define the meaning of abstract words in the adult language in terms of concrete action if asked, as in the example Vygotsky relates of a boy explaining that "Reasonable means when I am hot and don't stand in a draft" (Vygotsky 1962: 78).

At the end of the holophrastic stage children will have acquired about fifty words or so. This is followed by a two-word stage (at the age of approximately one and a half to two), in which pairs of words are conjoined (e.g. an agent noun and an action verb, either as a comment or as a request for an action). As this gradually progresses towards multi-word utterances there is a stage of 'telegraphic' speech (from approximately two to two and a half) in which bound morphemes and function words are still largely lacking. There is a basic pattern of development as inflectional and derivational morphemes are acquired, for instance, past tense endings. The child starts using fixed correct forms (learnt as wholes), followed by a stage with much regularisation of irregular pasts (revealing that the child is formulating rules), and finally a stage is reached approximating to correct adult usage (Harley 2008: 145ff.).

By the end of this stage children are capable of learning new words fast, either by relating their meaning to more basic words already learnt that share similar (but not identical) meanings, or from contextual bootstrapping, using context to 'fill in' likely meaning – much as adults do when coming across an unknown word when reading. This is an instance of the Gestalt principle of 'closure' applied to the comprehension of language rather than vision and may involve both 'syntactic bootstrapping', whereby children use their knowledge of how a word functions to help understand its meaning, and 'semantic bootstrapping', which assumes the child's ability to identify semantic roles like 'agent' which will generate expectations as

to the syntactic structures involved.[2] In fact by age 6 children will typically have acquired 8,000 root words at a rate of about five new words a day (Carey 1978: 264). This 'fast learning' (which seems at odds with the gradual transferral of episodic into long term memory mentioned in 5.3) has inspired connectionist models which appear to show how such learning is possible with simple algorithms and without assumptions of innateness (Elman et al. 1998: 156–8).

Carey distinguishes two kinds or stages of word learning. In 'fast mapping' new words are roughly placed within an incipient lexical field in a 'fragile' form based on the specific circumstances in which they are learnt (and the specific objects to which they are applied), whereas 'full mapping' – i.e. the consolidation of a word's full meaning within the lexical field – may take months if not years. She illustrates this with English children's acquisition of spatial adjectives, in which the understanding of dimensionality develops only gradually after the early grasp of 'polarity' (the comparative structure of 'big' vs. 'little'). This is facilitated by the child already understanding conceptually the standard size of typical items (Carey 1978: 279). Full mapping will extend the initial limitation of the adjective to one or two specific objects to a full adult-like assembly of object types to which it can apply (for instance 'high' applying to buildings but not to 'tall' people). She also discusses the acquisition of novel colour terms, and of causative verbs. Though the details vary, she emphasizes the importance of linguistic context in all cases, with two basic processes going hand in hand: a) discovering which semantic features of an already acquired conceptual system are relevant, and b) mapping these features onto words (adjusting the concepts to the words in the domain and v.v.). A typical context for fast mapping of the meaning of the new colour term 'chrome', for example, was the request by the adult to the child to "fetch the chrome tray, not the red one" (where there is a choice of just two objects and the child already understood 'red'). It would still be a long time before the child could place this new term (for a kind of olive green) correctly and productively within the conceptual domain of colour.

The semantics of children's words and the pragmatic intentions behind their multi-word utterances are not necessarily the same as those of adults. Typical of children's use of individual words is over-extension and under-extension compared to adult usage. There are various theoretical approaches to this phenomenon, including the 'prototype hypothesis', which assumes children learn early on to build their meanings around prototype instances – this meshes with the Cognitive Linguistics approach to prototype semantics generally – or the 'contrastive hypothesis', which

2. The latter is crucial to Pinker's (1984) 'linking roles' approach whereby the child is supposed to have innate knowledge about syntactic categories and their relationship to semantic ('thematic') roles.

sees the key features in learning the meaning of a word as those that differentiate it from related words (Harley op. cit.: 131–133). Consider for instance how a child might learn the word 'cat'. Perhaps he or she applies the word first to a household pet, a particular white cat, treating it as the prototype of its kind, perhaps overextending it to a small white dog or any fluffy white toy. Black and tabby cats might at first be excluded (and a Cheshire Cat most certainly would!). Less prototypical exemplars would be less likely to have the word applied to it but as more and more near-fits are encountered, the child's concept would tend to 'average them out'. On the other hand, there might be both a dog and a cat to tend with in the household, in which case the child would quickly learn that the two creatures displaying contrasting sets of features have different category names attached to them. Whatever combination of mechanisms is involved, the child will end up with a 'fuzzy' bundle of abstracted cat features that henceforth will be evoked by the word 'cat' much as they are for an adult.

The acquisition of negation (as mentioned, a highly abstract concept) is of particular interest. Pepperberg (1999: 80) summarizes research in this area (in the context of a discussion of how her African grey parrot Alex learnt to express absence of information) as follows. An early stage involving rejection – refusing objects or actions either gesturally or vocally – is followed by one in which the child begins to comment on absence or malfunction, then on objects dropped or that have disappeared, only then being able to express the (negative) veracity of a statement. There is some disagreement about the relative order of rejection and comment on absence. Thus Bloom (1970) found that children learning English first use 'no' or 'not' to indicate that an object is not in its usual place, then they extended this to cover the rejection of an offer or suggestion (as in "No car" said pushing away a toy car). And finally they extend the use of these words to situations in which they deny or contradict a positive assertion (e.g. "Not a truck" when pointing to a bicycle). The difference here may depend on the criteria used (e.g. syntactic combination in the case of Bloom's analysis). The distinction between these three stages (of successive abstraction) is at all events obscured by the fact that they all involve the same English forms.

As regards the pragmatic force behind young children's utterances, Halliday (1975) analyses this in terms of seven functional 'proto-speech acts', which do not necessarily develop in a direct line into adult ones. These are the instrumental (expressing a need); the regulatory (telling someone else what to do); the interactional (strengthening inter-personal bonds); the personal (expressing feelings); the heuristic (asking and learning about the environment); the imaginative (creating imaginary worlds, etc.); and the representational (the conveying of facts and information). It is from these roots that the adult interpersonal and ideational 'metafunctions' develop. As regards the transition to understanding abstract

words, Halliday (1993: 109) makes the point that there is a difference between how and when children start to master the general ($abstract_2$) vs. specific dimension and the abstract (i.e. $abstract_1$) vs. concrete dimension of meaning. The latter begins only around age 4 to 5, later than the former ability to handle words referring to classes of objects.[3]

So far we have been discussing rather universal processes, but cross-linguistic differences already start appearing at a surprisingly early stage, as illustrated by Bowerman (1996). Taking her starting point in Slobin's (1985) 'operating principles' – universal strategies such as paying attention to the beginnings and ends of words which are assumed to steer development rather than Chomskyan innate syntactic principles – Bowerman focuses not on what might be innate but on measurable effects on children's linguistic development stemming from the specific target language to which they are exposed. This appears to prompt the child to divide up his or her semantic 'space' along language-specific lines from an early age. There are also differences among children learning the *same* language. Thus it has long been recognized that children may follow either of two main pathways beyond early holophrasis as they develop towards adult language skills. Although all children go through a 'vocabulary explosion' roughly around eighteen months old, some children concentrate at first on elaborating names for things with a rapid increase in vocabulary but little production of whole clauses (and little increase in phonological fluency), while others are able to produce fluent but stereotyped clauses with adult-like intonation from an early stage but with little elaboration of vocabulary – Nelson (1973) calls these respectively the 'referential' and the 'expressive' style characterizing children around the age of 24 months.

It would seem that something analogous applies to adult languages: some reflect a more analytic 'nouny' pathway forward into complexity (languages like modern Indo-European ones) while others reflect a more 'verby' one (especially 'polysynthetic' ones). Symptomatic of the former is often a distinct class of 'adjective', whereas in the latter adjectives are usually subsumed under (stative) verbs. This closely parallels the distinction between dependent-marking and head-marking languages (those developing elaborate case and noun class systems, and those capitalizing on the synthesis of elaborate verbal predicates – cf. Nichols 1986). There may be some sociocultural factor behind the choice of routes here, for instance if greater value is given to fluent, stereotyped speech patterns in some cultures, as opposed to more weight given to the elaboration of nominal lexicon for naming objects and places in others. Polysynthetic languages (spoken mainly by hunter-gatherers) generally have a limited stock of lexical roots/stems, counterbalanced

3. His suggestion that this is concomitant with learning to read and write is obviously irrelevant for children in cultures lacking literacy.

by highly productive derivational resources (including noun incorporation). They also display morphosyntactic holophrasis, packing pronominal subjects and objects and much more into verbs to produce independent 'word-sentences'. (Recall what was said about them in 5.1 as regards the representation of events.) It is instructive to compare the ways in which children do in fact learn polysynthetic as opposed to analytic languages, although there is still to date relatively little documentation of the former as opposed to the extensive literature on the acquisition of European languages (English in particular). I turn to this in the following section.

9.2 Acquiring language amongst hunter-gatherers

As a child learns to master a polysynthetic language there is no inherent limit to the combinatorial variability of the morphological 'word-sentence', which can be quite syntax-like. In other words, 'holophrasis' will become a way of life, not just an infantile phase passed through – though it should be emphasized that these languages are by no means childlike, they are as sophisticated as any more analytic adult language (there are almost always alternative analytical clause types, for instance). It is thus to be expected that the course of development will be somewhat different in these languages, especially as regards the acquisition of lexicon, since the number of 'words' in them – though made up of a finite number of morphemes – is in principle infinite! In contrast, children learning more analytic, 'nouny' languages like English will reach a point where all the necessary lexical word forms in order to deploy the adult language in a socially satisfactory manner are acquired – i.e. further morphological development is 'capped' by the lexicon. Nominal lexicon may also reach that point for learners of polysynthetic languages, of course, but their ability to produce new verb forms will continue throughout life, limited purely by utility and individual verbal skill. These skills of interpretation, imitation and innovation may be passed on from generation to generation and nudge the language itself towards greater and greater complexity – 'more and more of the same'– until the need of simplification in order to communicate with neighbouring speech communities puts the brakes on.

But polysynthesis is not a unitary phenomenon (in fact the word 'type' applied to them is something of a misnomer, but a convenient one). There would appear to be three main sub-types and consequently three main historical 'routes' into polysynthesis: an 'affixing-only', a 'noun-incorporating' and a 'clause-combining' one (cf. Fortescue 2016a). These 'pure' sub-types are not mutually exclusive, however, and individual languages may combine traits of more than one of them. Studies have in fact been carried out documenting (roughly) how children acquire all three of these sub-types. Particularly well attested in this respect are the Inuit

languages of the Eskimo-Aleut family, representing the affixing-only path, as in Fortescue & Lennert Olsen (1992) and Allen (Forthcoming).

Perhaps the most egregious divergence in the developmental sequence of acquiring an Inuit language is the lack of a significant transition to a 'two-word' stage. Rather, Inuit children go through a 'two-morpheme' stage, with their utterances consisting (at least in West Greenlandic) of a stem plus an inflectional affix, combinations that can be shown to be productive, not just learnt as wholes. In a following step a number of other semantically important (and frequent) affixes are tentatively inserted between already familiar stems and inflections that act as 'scaffolding', and then a third can be added to such new familiar combinations (from the inside out, as it were) – compare Carey's 'fast mapping' in the previous section. This will eventually ramify into a productive and recursive system for building up a quasi-infinity of 'word-sentences'. Inuit children will end up mastering several hundred productive affixes, both inflectional and derivational, where English, for example, has just a handful. Syntax (as opposed to morphology) in these languages is rather simple, however, word order being rather free and determined by pragmatic concerns of highlighting, etc. (The inflections on both verbs and nouns unambiguously determine their relationship.)[4] Other areas in which learners of these languages show precocious development (where grammatical features of the language may 'pace' cognitive development) include sensitivity to aspectual distinctions (there is a large number of distinct aspectual affixes), and deictic reference. (All Eskimo languages distinguish a large number of demonstratives, including the vertical and interior/exterior dimensions as well as distance and visibility.) Conversely, the syntactic manipulation of word order for pragmatic ends is delayed, with frequent patterns just copied from the input utterances children hear around them.

Other studies, dealing with the acquisition of languages of the other two subtypes mentioned above, include Mohawk, a language of the noun-incorporating kind (Mithun 1989), and Murrinh-Patha, a northern Australian language of the clause-combining kind (Nordlinger et al. Forthcoming). The 'ways in' to these languages differ somewhat according to other factors than their historical sub-type – for example stressed syllables play an important role in acquiring Mohawk but not West Greenlandic. These salient positions within the word, regardless of the semantics, tend to be those first attended to and imitated and mastered in production. They form a useful starting point around which further developments can then accrue – unlike in languages like West Greenlandic which lacks contrastive

4. More typical for other polysynthetic languages is that nouns are not inflected and the relationship between them and the verbs 'sub-categorized' for them is looser, determined by word order or other kinds of morphological marking.

stress altogether. The existing studies do however all show a 'precocious' attention to morphology prior to syntax on the part of children learning them, as Bowerman's study would predict. They also show that children (given time) have no particular difficulty in abstracting patterns (phonological, morphological, syntactic and semantic) from input as complex as that which polysynthetic languages present the learner with. Unlike for adults who really have to struggle with them, for young children they are just puzzles to play with and solve – abstraction can be fun when it involves solving puzzles. Success for all children learning a first language has a further powerful affective component, the reward being the ability to interact fully with the surrounding world.

9.3 Learning and abstraction

We have already met Vygotsky's ideas on inner speech in both children and adults. As regards the general acquisition of language by children, he initiated an influential view of the process as socially embedded from the start. For him language is a social phenomenon acquired from the child's environment. He advocated the complex relation of language and thought (and as such was a precursor of modern relativism). The two have different ontogenetic sources which remain distinct in children up until the age of around three, when their external 'egocentric' monologues become internalized as inner speech. Before that, in the 'pre-intellectual' stage, words are not symbols but actual properties of things for them. Vygotsky saw language as crucial to various aspects of cognitive development. Thus he describes the acquisition of 'genuine concepts' as a gradual process of abstraction from vague heaps falling into 'family' groupings, blurred by successive superimpositions, to 'potential concepts' – the grouping of objects according to maximum similarity, such as all round as opposed to flat objects. Advanced 'complex thinking' in turn presupposes the assistance of 'potential concepts', and "only mastery of abstraction combined with advanced complex thinking enables the child to progress to the formation of genuine concepts". It is a matter of continual alternation from the particular to the general and from the general to the particular. The decisive role in the process is played by the word, which is used to direct the process of advanced concept formation. Finally, the child must learn to apply a concept grasped and formulated on the abstract level to new concrete situations that must be viewed in these abstract terms, an ability only mastered towards the end of adolescence (Vygotsky op. cit.: 76–80).

This approach contrasts with the more linear perspective of Piaget (1954), who saw language as arising out of preceding cognitive stages – those of successively 'sensorimotor', 'concrete', and 'formal' operations – and the social use of speech as

being a later development. For him thought determined language, which did not arise separately and only later become merged with thought, as Vygotsky had it. However this may be, the processes of 'accommodation' (to the environment) and 'assimilation' (within the child) that he introduced are germane to both perspectives – also to the dynamic sense of 'abstraction' that is crucial to the present volume. Piaget considered cognitive structures as arising from innate 'dispositions' while not themselves being innate.

Chomsky (1965: 32) went much further as regards the specific disposition for language, proposing a universal and innate 'language acquisition device' (LAD), needed to supplement the impoverished 'input from the stimulus' in order for language to be acquired. Later the term was replaced by the more abstract concept of 'Universal Grammar' (UG). Within his 'Principles and Parameters' framework such a device is supposed to generate a finite array of possible 'switch settings', out of which specific ones are chosen by the child according to the input he or she hears from the surrounding target language. Unfortunately there is little anatomical evidence for any such dedicated device (though plenty for the neural underpinnings necessary for language in specific areas of the cortex).

As to how the child accomplishes the feat of abstraction of a specific language from genetic 'dispositions' plus input from the social environment, there is an alternative to Chomsky's innate LAD (and Pinker's similarly innate 'linking rules'), one that has gained empirical support more and more in recent decades, namely the distributional approach, whereby information is seen as accessed directly from the environment with minimal assistance from innate knowledge. The child can be shown from an early age to abstract a great deal of information from its linguistic environment (much more than 'poverty of the stimulus' arguments suggested), namely by registering the distribution and frequency of input patterns, albeit with inherent constraints at different levels of analysis. The feasibility of this kind of probabilistic learning based on the statistical analysis of input is supported by connectionist models, especially when 'starting small', as discussed in Elman et al. (1998: 340–349). Many of the complex structures Chomsky claimed to be unlearnable owing to poverty of the stimulus can in fact be overcome by the child, who necessarily 'starts small', with a limited working memory capacity. This may be an advantage compared to late learners 'jumping right in', as expressed in the 'less is more' hypothesis of Newport (1990). Children learn simple sentence types before complex ones involving inversion, etc. that build upon them and presuppose working memory that has matured sufficiently to hold such complex entities. Certainly, as we have seen, when children acquire West Greenlandic they 'start small' and literally build up more complex morphological structures by trying out 'new' affixes (ones frequent in adult speech) in the middle of familiar word-sentence structures until these too form familiar contexts for further expansion.

Other cognitively oriented linguists have argued for similar views. Thus Bybee discusses evidence supporting Langacker's usage-based model (as in Langacker 1987, and in greater detail 1999: 91–145) in relation to child acquisition studies, observing how children utilize analogy to abstract repeated patterns from specific exemplars (Bybee 2010: 64–66). She points out that in learning more general, abstract construction schemas children do not necessarily lose all specific exemplars (some can still be accessed as single units). She cites, for instance, a study by Dąmbrowska & Lieven (2005) which indicates that the questions formulated by young children are heavily dependent on interrogative constructions they have already uttered. The authors suggest two basic operations that can be used to produce novel utterances: 'juxtaposition' (of two units) and 'superimposition' (by which a unit 'fills in' a schematic slot in another construction, expanding it). MacWhinney (2005) emphasizes the use children make of perspective in learning about language – understanding mutual perspective-taking is essential to successful communication and binds together direct experience, space/time deixis, plans, social roles and mental acts in language use (op. cit.: 198). For him, as for other functionalists and cognitivists who deny the strong involvement of innate knowledge, grammar is emergent. This is true 'emergent abstraction', becoming feasible as developmental conditions allow.[5]

A final aspect of language acquisition that has been used by nativists like Chomsky to argue for an innate language acquisition device is that of the critical period at which the conditions for learning a first language are optimal and beyond which learning becomes progressively more difficult. There are both nativist and 'maturational' explanations for this apparently 'pre-programmed' phenomenon. The latter type of explanation is premised on general neural and cognitive maturation having to reach a certain level of development before language learning can begin. Some doubts as to the rigid criticality of this period have recently arisen – thus Harley (2008: 78–79) prefers a weaker 'sensitive period' hypothesis. One way or another, children do end up acquiring probabilistic generalizations (if not completely general rules) about what they hear and reproduce, as seen in their overextensions (including past tense over-regularizations). In all the approaches characterized above abstraction from the environment is assumed. Even the innate LAD hypothesis requires input from adult speech outside the child's brain in order for specific parameters to be 'chosen', and that input has of course to be abstracted from the acoustic signal either in the step-by-step way sketched in 9.1 (analysis followed by synthesis, as with perception in general) or by a more purely distributional analysis.

5. An emergent property can be thought of in terms of the sudden transformation of water into ice or into steam according to precise conditions (here temperature). Successive atoms on the periodic table also possess emergent chemical properties by the addition of extra nucleons.

CHAPTER 10

The philosophical approach

10.1 Philosophers on ideas, universals and abstraction

Philosophers have always been concerned with the nature of abstraction – concepts, generalizations, simplifications, integrations, universals, essences, classes and so forth. A succinct account of the main landmarks through the history of western thought can be found in Russell (1959). The first to treat such entities explicitly was Plato, arguably the first rationalist – i.e. a thinker who proceeds essentially by deduction from first principles. His abstract universals covered both the form of objects (such as the idealized human form), mathematical forms, and abstract ideas such as 'beauty' and 'justice'.[1] His 'realism' (in the scholastic sense of the word, i.e. belief in universals as real 'things'), far from being ignored today, has been influential throughout the ages. This is especially so in mathematics and physics, through Galileo and Newton to Frege and Gödel in more recent times, in theories of number and class, mathematical values and axioms, etc., all of which can be seen as denizens of a 'third' realm, neither mental nor physical and yet, as Plato would have it, deeply involved in the world intelligible to the senses. According to Whitehead's famous acknowledgement of the breadth of Plato's vision, western philosophy can be seen as "a series of footnotes to Plato" (Whitehead 1978: 39).

By contrast, Aristotle as the first empiricist – i.e. a thinker who proceeds essentially by induction from instances – eschewed such ghostly entities, but nevertheless maintained certain metaphysical categories such as quantity, quality and substance applicable to all phenomena. He was a nominalist, holding that universals are only names (though he was to some degree also a 'realist' in maintaining the form/matter distinction). It took a very long time for philosophy to get over his conception of individual 'substance' and 'final causes', as adopted by Aquinas and integrated into official Christian theology. Duns Scotus (a precursor to Leibnitz) adopted Aristotle's form vs. substance (matter) distinction but gave it a twist, by seeing individuation solely in terms of form, i.e. of combinations of

1. Note that an 'idea' for Plato is an ideal type or pattern. In earlier usage, reflecting the word's origin in the verb *idein* 'see', it meant 'apperance, form, kind' (Buck 1988: 1213) , i.e. a mental 'picture'.

attributes – including position in space and time. Francis Bacon consequently laid the basis for empirical scientific investigation in inductive generalizations, natural hypotheses arising out of comparing things with the same property and subtracting all in which it is absent.

A very different perspective on form and universals was that of Descartes (whose rationalism strongly influenced Noam Chomsky). Doubting the reliability of input from the senses alone, he saw no clear way of distinguishing reality from dreaming – the only solution for him was to accept the total separation of the inner world of mind or soul from the outer one of matter. For him universals must be innate and come from the world of soul (and ultimately God). Philosophy on the continent continued in this rationalistic, essentially deterministic vein for some time, with elaborations in various different directions, first by holistic pantheist Spinoza (for whom understanding the causal chains linking everything into a totality allowed at least a kind of freedom). Leibnitz opposed Spinoza's 'whole' with a multiplicity of 'monads' (synchronized individual 'substances'). His rationalism was manifest not only in his invention of calculus (the precursor of symbolic logic) but also in his attempt to build a universal alphabet of human thought, an early precursor of semiotics.

Empiricist John Locke introduced a causal theory of perception, which assumes no innate knowledge, a young child's mind being a 'tabula rasa' in which all knowledge arises from experience. This presents to the mind two types of 'idea': simple ones are the immediate products of sensory stimulation, complex ones are constructed out of simple ones as the result of mental operations. Thus besides familiar types of object, unreal combinations such as 'unicorn' can also enter the mind, combining 'horse' and 'horn'. He divided simple ideas into primary qualities essential to all objects, like extension, shape and motion, as opposed to secondary qualities of objects due solely to their effect on our senses – colour, taste, etc. existing only in the mind (a view upheld by modern neuropsychology – we do indeed live in a 'virtual' world of sensory '*qualia*').[2] Hume extended the empiricist position into an extreme form of scepticism. For him there is no causality, only expectations based on induction – let alone a 'self' experiencing it. And induction (as opposed to deduction from axiomatic principles) is always guesswork/probabilistic prediction, so scientific laws are merely generalizations from inductive thinking. A still more extreme kind of scepticism – this time in stark contrast

2. Not to be confused with Pustejovsky's use of the term to refer to aspects of the meaning of lexical items – for him *qualia* may be 'constitutive' (the relation between an object and its constituent parts), 'formal' (that which distinguishes it within a larger domain), 'telic' (its purpose and function), and 'agentive' (factors involved in its origin) (Pustejovsky 1995). These are relevant to the 'micro-functional' features of my 'mediatory word columns' illustrated in 8.3.

to the materialism of the empiricists – is that of (Bishop) Berkeley, for whom *all* we can ever experience is 'ideas', not reality.

Kant synthesized the rationalist and empiricist views on 'ideas' (or 'metaphysical judgements') by positing his famous 'categories' such as substance, cause and effect, necessity, reality, totality, etc., which are restrictions on what can be abstracted from experience. They presuppose a spatial and temporal framework. Both categories and space/time are imposed by the human mind (as 'forms of intuition') – the noumenal 'thing in itself' is unknowable directly. His 'transcendental idealism' strongly influenced later phenomenologists and Gestalt psychologists. Hegel extended this to full-blown historical idealism – for him 'concepts' (i.e. ideas) are capable of being false (just as assertions or propositions can be), so scientific theories can always be false, and in any case they never tell the whole story. (Ultimate truth is only to be found in the 'Absolute' or universal mind.) His proposal of a historical dialectic was far-reaching: thesis leading to (contradictory) antithesis, leading in turn to synthesis, a new thesis, such that there is continual process, but only at the level of the absolute or ideal.

Bergson also focused on process, but in a way that united the idealistic and materialistic perspectives: he proposed a dichotomy between matter and the life force – the instinctual '*élan vital*', that is for ever coming into being. The life force is the urge to create novelty and to diversify, while matter is the entropic compulsion towards uniformity that dissipates energy and resists the flow of life. It is non-teleological, the font of free will, serving no absolute Aristotelian purpose. The intellect aligns with matter, classifying the seamless flow of life into discrete categories. Pragmatists like William James and John Dewey were by contrast not concerned with overarching theory. For them reasoning was a practical process of problem-solving, constructing hypotheses which could be weighed up and tested in action with alternatives. ("The true is that which works".) For another American pragmatist, Charles Sanders Peirce, the goal of knowledge was to counterbalance instability and doubt, to restore a kind of homeostasis in our behaviour. Besides his theory of the sign, which I have referred to already, he contributed the notion of 'abduction', which combines induction and deduction as (testable) hypothesis generation, the only form of reasoning that can produce novel ideas.

A surprisingly modern, perceptually based approach to abstraction is that of Alfred North Whitehead, whose dynamic 'philosophy of organism' was decidedly ahead of its time. Whitehead is a proponent of 'realism' (like his immediate predecessor Samuel Alexander with his 'emergent evolution'), but a modern kind of common-sense, scientific realism in which universals are not 'things' (as in scholastic 'realism') but potential 'forms of definiteness' in nature, not just 'names'. What is treated as mere 'attributes' in the empiricist tradition is seen by him as real internal relationships shaping the possible pathways dynamic processes can

take. His philosophical system is informed by his interest in psychology (through William James) as well his involvement with the rapidly advancing physical sciences, replacing the outworn concept of 'substance' with that of 'organism' and in place of time and space as absolutes treating the unitary 'event' (or 'actual occasion') as central. (Compare James' 'duration', the unit of psychological experience.) He further replaced Locke's split between primary and secondary qualities with another (that between 'causal efficacy' vs. 'perceptual immediacy'), which I shall return to in section 10.3. He emphasized the social and cultural context of ideas (especially in Whitehead 1947). In doing so he re-introduced subjective 'value' into an essentially scientific perspective – an attitude carried further in Polanyi's philosophy of 'personal knowledge'. His grand scheme can be said to combine Spinoza's holism with Leibnitz's monads, primed by Bergson's dynamism – though not his anti-intellectualism.[3] His 'eternal objects' (the 'concepts' of his system) owe much to Plato, but it is important to realize that these are explicitly to be taken as determined by the overall 'concrescence' of the universe-as-organism – they are 'pure potentials' (Whitehead 1978: 22). His warning against 'misplaced concreteness' in interpreting final realities beyond the 'actual occasion' could equally well be applied to 'eternal objects'. They can be equated with Locke's 'ideas' (subsequently abandoned by Hume). As he put it, aiming mainly at Cartesian dualism: "There is no going beyond actual entities to find something more real" (op. cit.: 27–28). His choice of terms was sometimes unfortunate and evoked more than a degree of scepticism in the philosophical climate of the time in England, which was much more hospitable to Russell's atomism and the burgeoning interest of professional philosophers in language and logic. Perhaps the most 'modern' aspect of Whitehead's system is his emergentism, according to which thought and the individual 'prehending subject' emerges from the world, not the reverse, as in Kant's system. For more general detail see Fortescue (2001a), a cognitive interpretation of Whitehead's philosophy with particular reference to language. In the following section I shall focus on just one aspect of his philosophy, his perspective on abstraction.

Following World War Two, interest in 'ideas' and 'concepts' waned, both in 'Anglo-Saxon' scientific materialism and in continental 'existentialism', directions which were diverging further and further from each other. In stark opposition to Whitehead's system-building holism, his erstwhile co-author of the "*Principia Mathematica*", Bertrand Russell, moved further in the direction of reductionist logical atomism – an exclusive focus on 'singulars' in the tradition of William of

3. His 'reality as process' approach can be said to go back to Heraclitus, for whom the universe was in permanent flux and strife, 'an everlasting fire', a viewpoint virtually lost to philosophy before Bergson and Whitehead (though ingredient in Hegel's idealism).

Occam, for whom concepts were all constructs of the mind, and before that of Democritus, for whom even the mind consisted of 'atoms' in constant flux. Russell had already developed his theory of types to elucidate the paradoxes that arise when logical levels are ignored (hence the absurdity to which statements about classes that include themselves can lead). He was sceptical of the surface form of sentences (a scepticism shared by Whitehead). Russell had a strong influence on the early Wittgenstein's 'picture theory' of meaning ("what we say reflects what there is"), while his theory of definite descriptions showed how it is logically possible to speak meaningfully of non-existing things, thus opening the way for ordinary language philosophy, one of the movements dominating post-war Britain.

From his later, quite different theoretical stance, Wittgenstein (1974/1953) insisted that meaning cannot be detached from the activities of language users: the role of context was now central. It is woven into his notion of 'language games', an analogy used by de Saussure (1916) in a purely structural context, and later formalized by Austin and then Searle as Speech Act theory. This was the beginning of 'ordinary language' philosophy ("how to do things with words"), in which concern for 'universals' was largely restricted to truth conditions for utterances. Strawson, however, representing this movement, utilized abstraction in a somewhat different way in characterizing the building up of more and more abstract sentence types from a 'basic combination' that consists of a concrete particular (expressed as a noun) plus a 'general concept' indicating some property of the particular (expressed as a predicate) (Strawson 1974). This can be extended to a higher level with non-particular subjects, and, by a further process of 'substantiation', to successively more abstract levels (approaching the actual capability of English) by inserting complex conceptual units as subject – e.g. nouns standing for a verbal action or an adjectival quality. Even an out-and-out materialist like Ryle, under his banner of "no ghost in the machine" (referring to the category mistake of Cartesian dualism), was concerned with such Russellian matters as the layering of abstractions, of classes of classes, etc. His claim that mental properties are simply dispositions to act sums up the dominant view of behaviourists and logical positivists of the period. A more recent approach to abstraction in the philosophy of language (Moltmann's) will be sketched in the following section.

Meanwhile on the continent, for Husserl and other phenomenologists following Brentano's concept of 'intentionality', the prevailing trend was inwards, continuing the direction of earlier idealism: conscious states refer to a content, which may be abstract (e.g. non-existing) or particular. Conscious state and object are two aspects of a single phenomenon, the intentional act. By stripping away all contingent aspects of conscious experience it was hoped that fundamental elements of mental processes could be reached. Empty speculation about the ultimate nature of objects was dismissed as just leading to Cartesian doubt. Even

Husserl's Kantian view on the exclusion of the external 'thing in itself' was rejected by Heidegger, whose exclusive focus on 'Being' (a massive abstraction!), led the way into existentialism.

Not all continental philosophy of the time was out of tune with the positivist strand of Anglo-Saxon philosophy of the time. Thus Frege's quantificational logic and philosophy of language was eagerly taken up by Russell and the early Wittgenstein, in particular his distinction between 'sense' and 'reference' (applicable also to mathematical equations), as was his analysis of sentences as couplings of mathematical functions and arguments, the meaning of individual terms requiring a sentential context. Also the structuralist approach to language of de Saussure was to prove widely influential. A language was for him an abstract synchronic system of signs – 'langue' – without reference to any particular historical implementation of the language. This relates directly to Chomsky's abstract notion of 'competence' (as opposed to concrete 'performance'). From a completely different perspective, mathematician Kurt Gödel saw abstract objects existing independently in a 'third realm' harking back to Plato. His famous theorem claimed that the truthful propositions generated by a mechanical system (e.g. the algorithms of Artificial Intelligence) can never be complete and can never emulate everything that a human mind can know (Gödel 1986).

In sum, all varieties of philosophical and scientific thought throughout the ages have by necessity involved abstractions of one sort or another. It is in the nature of the game. At this point a short excursion into the nature of mathematical abstractness – the purest form of abstraction of all – would seem appropriate. It will be seen that it too is amenable to a cognitive interpretation. Formal semantics straddles the divide.

10.2 Mathematical and formal semantic abstraction

One can take as starting point Whitehead and Russell's (1910–13) joint endeavour to reduce mathematics to logic (and numbers to classes) in the *Principia Mathematica*.[4] This was an attempt (only partially successful) to ground abstract number in the axioms and inferences of symbolic logic. The project carried forward the work of Leibnitz and Frege on mathematical logic and provided the basis for modern set theory, predicate logic and propositional logic. The inferences in *Principia Mathematica* include the following formal version of *modus ponens*:

4. Thus Russell defined a number as the class of all classes similar to a given class.

> "[If] '$\vdash . p$ ' and '$\vdash (p \supset q)$' have occurred, then '$\vdash . q$ ' will occur if it is desired to put it on record. The process of the inference cannot be reduced to symbols. Its sole record is the occurrence of '$\vdash . q$ ' [in other words, the symbols on the left disappear or can be erased]" (*PM* 1962: 9).

As for the axioms, the authors employed notation like the following, here for their 'Axiom of Reducibility' (*PM* 1962: xxxix):[5]

$$\varphi x \equiv_x \psi x .\supset. (x): f(\varphi \hat{z}) \equiv f(\psi \hat{z})$$

The authors provide the following 'reasonable meaning' of this formula "IF for all values of x the *truth-values* of the functions φ and ψ of x are [logically] equivalent, THEN the function f of a given $\varphi \hat{z}$ and f of $\psi \hat{z}$ are [logically] equivalent." As a layman I shall refrain from trying to comment on this 'reasonable meaning', which the authors described as 'obvious' (!). I am only interested in the *kind* of abstraction such formulae represent. They are clearly both abstract$_1$ and abstract$_2$. But what they are, Whitehead and Russell insisted, is essentially tools – usable mathematical techniques rather than just abstractions of philosophical interest. They argued that their axioms were defensible 'on inductive grounds' (1910, 2nd edition: 59):

> The reason for accepting an axiom, as for accepting any other proposition, is always largely inductive, namely that many propositions which are nearly indubitable can be deduced from it, and that no equally plausible way is known by which these propositions could be true if the axiom were false, and nothing which is probably false can be deduced from it. If the axiom is apparently self-evident, that only means, practically, that it is nearly indubitable; for things have been thought to be self-evident and have yet turned out to be false. And if the axiom itself is nearly indubitable, that merely adds to the inductive evidence derived from the fact that its consequences are nearly indubitable: it does not provide new evidence of a radically different kind. Infallibility is never attainable, and therefore some element of doubt should always attach to every axiom and to all its consequences. In formal logic, the element of doubt is less than in most sciences, but it is not absent, as appears from the fact that the paradoxes followed from premisses which were not previously known to require limitations.

The 'paradoxes' referred to are those that Russell unearthed in his theory of types. A concern with the nature of truth lies at the heart of the *Principia Mathematica*, but the authors recognized that the truth of an axiom from which all its further consequences derive is (in a sense) 'induced' from the 'indubitable' nature of the consequences themselves and the axiom itself may not be infallible (a point Gödel was to make much of). This is a curious kind of 'truth' and the possibility

5. This axiom was introduced as a means of overcoming the not completely satisfactory effects of Russell's theory of types.

of relating axiomatic formulae like that given above directly to 'reality' is seriously undermined by this caveat and led in fact to various objections to the whole endeavour (notably by the later Wittgenstein). Hence the importance the authors attached to the actual mathematical uses of their approach – and indeed the significance of their work for mathematics was lasting. We are justified then in characterizing such formulae as 'practical abstractions' arrived at by deduction from first principles that are not themselves infallible.

Whitehead and Russell's work can be compared with the attempt by Lakoff, starting from a completely different perspective (that of Cognitive Linguistics), to analyse the basis of mathematics not as something transcendent but as based on rational human activities (Lakoff 1987: 353–369). This he does in terms of now familiar 'image schemas', for example treating categories as containers, as illustrated in his schema-based explanation of *modus ponens* (op. cit.: 272). The different branches of mathematics, e.g. number theory, algebra and geometry, correspond to different kinds of human activities and cannot just be based on a unique set of axioms such as those characterizing set theory. Mathematics comprises an arsenal of different formal systems abstracted from such activities, each of which has its own 'logic'. This is perhaps not so different in spirit from *Principia Mathematica* after all. Compare for instance 'Force Dynamics' (as discussed in 3.1), which is compatible with Whitehead's non-linear, processual view of logic.[6]

For Langacker too conceptualization is a dynamic process at the very heart of Cognitive Linguistics. Specifically looking at some of the key terms of logical analysis, he uses the distinction between virtual and actual planes of meaning to explain quantifiers and quantifier scope in terms of 'fictivity' (Langacker 2005: 185). He analyses for example the semantics of 'each' (involving sequential 'fictive scanning') vs. 'every' (focusing on individuals within the set) and 'all' (which summates the set as a totality). These are different ways of generalizing across instances, all representable in similar simple schemas but 'profiled' differently. See also Sweetser (1990: 113–144) on the logic of conditionals within the Cognitive Linguistics framework.

The image schema approach leads us to reflect on algebraic formulae like the Mandelbrot set equation $z_{n+1} = z_n^2 + c$ given in 7.3. What is its relation to 'reality' if it does not belong in a purely Platonic world intelligible in terms of numbers but otherwise outside of common experience? Does the fact that its recursive application

6. In his later work he emphasized that the over-intellectualized approach to logic on the part of contemporary philosophers had obscured the main function of propositions in the nature of things. They are not primarily for belief but for 'feeling' at an unconscious physical level (Sherburne op. cit.: 63–4). Only in (some) higher 'intellectual prehensions' are propositions imbued with consciousness.

(filling in the variables with actual numbers corresponding to spatial coordinates) produces patterns that reflect the 'near chaotic' patterns of natural phenomena mean that it is just an invention of the human mind, an analogy that happens to fit such external patterns rather well? Or does it reflect something at the heart of Nature itself – its own non-bodily image schema, so to speak, quite independent of the human perspective? The 'laws of Nature' are all simple yet lead to and constrain an infinite variety of manifestations. This is where Whitehead's philosophy of organism and process comes into own, since it negates the Cartesian 'bifurcation of Nature'. Equations like the one above are exemplars of one type of 'eternal object', i.e. a potential constraint imposed by the structure of the universe as a whole on the direction in which reality can unfold. There is in principle no reason why such a pattern should not be realizable both in neural nets within individual brains and in the universe at large (an eminently 'realist' position). Universal constants like the speed of light or Planck's constant (the minimal quantum of action) are on the other hand certainly real 'forms of definiteness' at the macrocosmic and microcosmic scale of things, but they are only abstractions in the sense of theoretical constructs in a Popperian World 3 of cultural artefacts – these will be the principal subject of 11.3. As will be seen, this seemingly Platonic 'world' is non-existent apart from its 'ingression' in the material world (to use Whitehead's term).

Another type of 'eternal object' is the 'ideal' triangle, a quintessential Platonic form. It is by no means impossible to see how such an entity might be entertained by a finite brain. Barsalou (1999: 584–585) suggests how: what is needed is not a unique representation of a 'prototype' (which one exactly?) but a set of instructions to construct any shape fitting the definition of 'triangle'. Such a representation need not be holistic and is reducible to schematic components, different parts of a neurally representable triangle. Some of these would respond to the presence of lines independent of length, others to vertices between pairs of lines independent of the angle between them – "three qualitative detectors for lines coupled with three qualitative detectors for vertices that join them". This could define a generic triangle schema, some elements of which may be activated according to the task at hand while others are not.

Existing mathematical tools are still not able to tackle the biological activity of neurons directly, it would seem. The self-organizing, controlled chaos of the brain with its capability of coming up with novel activity patterns makes it different in kind from AI simulations. In the meantime Singer calls for a 'new' mathematics to handle the non-linear dynamics of neuronal processes – i.e. reflect them in schematizations as succinct as possible (Singer 2013: 624). No doubt many of the philosophical conundrums that remain regarding the relationship between mind and matter will fall away as more and more is understood of the hierarchical workings of the brain right down to the molecular level.

Closely related to mathematical abstractness is that of model-theoretic logic and formal semantics, as typified in 'Montague grammars'.[7] Such systems are not concerned with neural 'underpinnings' as opposed to a purely formal characterization of human language, but they offer further insights into the nature of abstractness within the general framework of the philosophy of language. Moltmann (2013) proposes a reformulation of the ontology of 'abstract objects' as expressed by natural language (specifically English), and attempts to answer the question: "In what way (if any) does natural language allow for reference to abstract objects?" Her answer is: "much more marginally than supposed". She argues that in linguistic semantics a wide range of 'abstract objects' (propositions, properties, etc.) is too often taken for granted, but that a great many of these are actually 'tropes' (op. cit.: 4). These are 'particularized properties', including such expressions as 'the beauty of the landscape', 'the heaviness of the stone', 'the number of planets' – they are 'abstract' only in the sense of attending to just one aspect of a particular and abstracting from all the others. But they are also 'concrete' in the sense that they can act as objects of perception and as *relata* of causal relations – they are 'causally efficacious', as in "The humidity of the air caused the softness of the wood" (op. cit.: 50).[8] This form of abstractness yields not only particularized properties but also degree-like, extent-like, number-like, and proposition-like objects "that may still be concrete entities".

Moltmann discusses five types of linguistic expressions apparently referring to abstract objects: NPs with nominalizations or functional nouns as heads (e.g. 'Socrates' wisdom' or 'the redness of the apple'); bare plurals and mass nouns and other kind terms (e.g. 'giraffes' or 'the desire to become rich'); expressions in subject or complement position that are not referential NPs (e.g. 'that' clauses); special quantifiers and pronouns (e.g. 'something', or 'everything'); and 'reifying terms' (e.g. 'the fact that it is raining', 'the possibility that it might rain', 'the property of being wise', 'the number two', and 'the colour green'). Only the latter, 'reifying terms' (referring to facts and states) are true abstractions in her view. They are non-referential and derivative, i.e. third order entities according to Lyons' distinction – it is these that are closest to the entities dealt with in the philosophy of mathematics

7. In Montague semantics the meaning of the whole proposition is a function of the meanings of its parts and their mode of syntactic combination. Montague was a mathematical logician who held the view that natural language is a formal language in the same way as predicate logic is. In his view the study of natural language belongs to mathematics, not to psychology (Montague 1970).

8. In the earlier history of philosophy they correspond to Aristotle's 'accidents' which 'instantiate' qualities as opposed to individuals 'having' a quality, as in 'the wisdom of Socrates' as opposed to 'Socrates is wise' (op. cit.: 46–7).

(op. cit.: 3). All the others are 'tropes'. Note that for Moltmann universals play a role in the ontology of natural language as pluralities of possible particulars "obtained by a form of abstraction from non-referential occurrences of expression". They can be conceived of as "pluralities of their instances", so words like 'mercy' or 'wisdom' do not (as usually assumed) refer to (single objects that are) universals, but rather stand for the pluralities of what would be the instances of the universal, namely tropes (op. cit.: 6).

This is not the place to enter into a detailed elaboration of this or any other variety of formal semantics, but just to illustrate the mathematical flavour of analyses in this area compare Moltmann's formula for the 'extended meaning' of the predicate 'exist' (Moltmann 2013: 22):

$$[\textit{exist}_{\text{ext}}] = \lambda i\, \lambda k\, [\exists d'(d' \, I_i \, k \,\&\, \text{exist}_i(d'))]$$

Notice the existential quantifier '∃' here, a symbol expressing the basic meaning 'there is', seemingly on a par of abstractness with mathematical symbols like '=' and '√'. But is it? None of these is really abstract_1 ('disembodied from specific sensory features') or abstract_2 ('simplified or generalized across instances') – they are functions or operators or relationships within formal systems and are not intended to symbolize or schematize 'things' in the physical world at all. In this they differ both from words and from such natural constraints on the physical world as the speed of light 'c' or Planck's constant 'h', or 'π' for that matter, which can certainly be considered (as argued above) to be universals.

It is worth considering the respective ontological status of a formal semantic formula such as that just given compared with that of a mathematical formula like $z_{n+1} = z_n^2 + c$, or a phrase structure rule such as S → NP+Aux+VP given in 2.3, or a 'law of nature' such as '$e = mc^2$'. One relevant distinction here is between descriptions and instructions: the formal semantic and the phrase structure formulae are purported descriptions of natural (human level) phenomena but can also be considered instructions of, on the one hand, how to interpret certain linguistic phenomena and on the other how to 'generate' an open-ended number of legitimate (and interpretable) syntactic structures – and as with descriptions in general, there are always alternative formulations. They are practical abstractions on a par with the axioms of *Principia Mathematica*. The Mandelbrot set formula is more clearly just an instruction (for how to generate certain fractally iterating patterns) – if its products happen to correspond to forms found in nature, that is of course interesting but the relationship is only approximate and the formula is not generative of natural phenomena as such (unless an autistic mathematical god is involved!). Einstein's famous formula is different again: it is neither descriptive nor instructional, it simply sets universal constraints on the way the universe actually functions. It is in a sense the macrocosmic 'context' of everything that is and, like

'h' and 'π', is independent of any human participation – except of course as regards the choice of symbols to represent the universal constraint.

In what respect are these 'concepts' all abstract? Is there some sense of abstraction they all share? And are we justified in calling them 'concepts' at all if some of them are independent of human minds entertaining them? Proposals as to how this and similar questions may be answered are to be found in the writings of Whitehead and Popper, on which the next two sections focus.

10.3 A Whiteheadian perspective on abstraction

Before getting to the role of abstraction in Whitehead's scheme of things, I must first introduce a few of his essential concepts. (This whole section can be skimmed over by readers who do not wish to tackle his admittedly somewhat daunting terminology head-on.) The basic unity of reality is the 'actual occasion' or 'actual entity', microcosmic drops of experience, units of process in continual 'concrescence' from multiple 'data' (input) into new forms of momentary unity forming data for ensuing occasions in a chain extending through time – the concrescence *is* the actual occasion. These units form interdependent aggregates that he terms 'nexūs', i.e. 'systematic schemes of relatedness' which, if organized socially, form 'societies', i.e. the macrocosmic objects we experience in everyday life, including as the most complex exemplar the human brain. The internal structure of the occasion is a succession of 'prehensions' leading to the 'satisfaction' of the concrescence as it perishes into an objective datum for successive occasions.[9] A prehension is the way an actual occasion relates and reacts to any other item in the universe, how it absorbs it into its own internal constitution. (Everything is to some degree related to everything else in Whitehead's philosophy of organism.) It consists of a subject (the actual occasion in which the prehension is a concrete element), the datum which is prehended, and its 'subjective form'. The 'subjective form' is how the subject of the prehension (the perceiver) feels the datum of the prehension.[10] It provides the '*qualia*' of Popper's subjective World 2, which we shall encounter in 11.3. It may include emotional valuation and intentionality but only in the case

9. Whitehead is ambiguous as to the 'duration' of the actual occasion – on the one hand it is the 'immediately perceived present' of the prehending subject (compare James' notion of the 'psychological duration' mentioned earlier), but he further states that the 'stages' of a concrescence are to be understood as logical rather than strictly temporal. A reasonable cognitive interpretation is to see these stages as temporally extended, though with a good deal of parallel processing.

10. In the case of a human subject this covers feelings and all manner of bodily reactions to perceptual and conceptual experience.

of actual occasions within a (human) brain can the subjective form involve the emergent property of consciousness. The term 'prehension', by the way, is adapted from Leibnitz's term 'apprehension' (Whitehead 1947: 300).

In its simplest 'physical' form an actual occasion is 'objectified' as datum to another (the subject of the prehension), and this could be at any level or scale of nexus in which the occasion is embedded, from elementary particle up to planets and stars. In Figure 11 from Sherburne (1981: 40) the successive phases of the concrescence are illustrated. (Note that the equivalent term 'feeling' is used here, a purely technical usage of this word equivalent to 'prehension'.) What is prehended in a 'conformal feeling' is a physical/ sensory datum, whereas in a conceptual feeling this datum is related to a corresponding 'eternal object' (b) – as I have already explained, 'eternal objects' are 'forms of potentiality' and under a cognitive interpretation can be understood as neurally instantiated concepts. (b') is a 'reversion' from the activated concept to other concepts through relevant associations (it is here novelty may enter the concrescence). The 'ingression' of 'eternal objects' in actual occasions can be understood as embodiment in the Cognitive Linguistics sense. Image schemas instantiating basic patterns of bodily experience are at all events compatible with this perspective, and Whitehead's 'concrescence', with its emphasis on novelty and process, resonates with cognitive blending and dynamic systems theory.

In a simple comparative feeling these eternal objects (or activated concepts) are summated (e.g. as the predicate of a 'proposition') and contrasted/ compared with the input 'datum' (e.g. as subject of the proposition), while in a complex

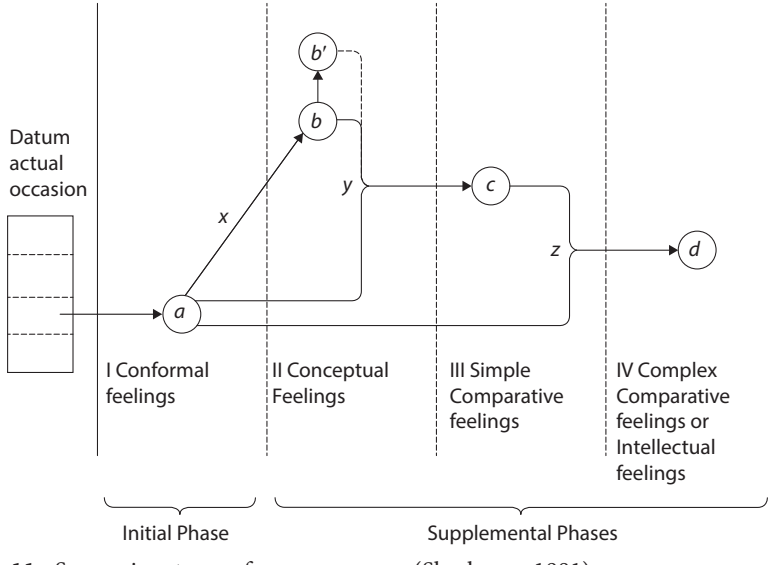

Figure 11. Successive stages of a concrescence (Sherburne 1981)

comparative (or 'intellectual') feeling a nexus of actual occasions (that in which the original datum is 'embedded') is contrasted with a proposition (c) the logical subject(s) of which is/are contained within the nexus. Whitehead's theory of prehensions is highly sophisticated and (in my opinion) his most valuable contribution to cognitive science since they can be directly related to brain states and processes (though his terminology may be an initial hindrance to this).

Not all occasions go through all these stages. The ones most relevant to our focus on abstraction lie at the higher comparative phases of the concrescence. Compare Figure 12 from Fortescue (2001a: 20) schematizing the three most important prehensions for present purposes, three species of 'propositional' prehension: the 'perceptive prehension' (a kind of simple comparative prehension), the more complex 'imaginitive prehension', and the most highly elaborated one, the 'intuitive judgment'. Here PS stands for the prehending subject of the prehension itself, LS is the logical subject of the prehended proposition, and EO is the relevant 'eternal object' or concept. The imaginative prehension integrates the prehension of a predicative pattern (an EO) with the 'indicative' prehension of a logical subject that does not actually display that pattern, in other words it 'imagines' the logical subject of the proposition in some other situation than that prehended in the immediate 'physical recognition' (it always involves 'reversions' and introduces novelty into a concrescence).[11] By contrast, in the 'perceptive prehension' there is no such integration since the logical subject prehended already displays the predicative pattern in question. There are also anticipatory prehensions linking present data to future potentiality. Finally there is the 'intuitive judgment' (Whitehead's spelling), the highest form of 'intellectual prehension', which integrates propositions with further prehensions of the nexus from which they (the propositions) derive, i.e. from the overall organic context in which an event or situation prehended is embedded. It is conscious and involves an extra layer of contrast whereby the full propositional potential of the logical subject is prehended, linking it to other known properties or associations that define its broader contextual nexus. The arrows indicate the prehensions themselves, successively more complex from type to type.

Other basic terms necessary to grasp in order to understand Whitehead's take on abstraction are 'causal efficacy' and 'presentational immediacy', his two distinct modes of perception. The former is the more primitive mode, present in all actual occasions and representing the inheritance of 'feeling' from past data – the feelings it transmits are vague and massive and arise in the initial conformal phase of a concrescence. It is perception as input from the sense organs filtered through and

11. A Whiteheadian predicate, comparable to Vygotsky's 'psychological predicate' discussed in 5.1, is a meaningful contrast within a system or nexus.

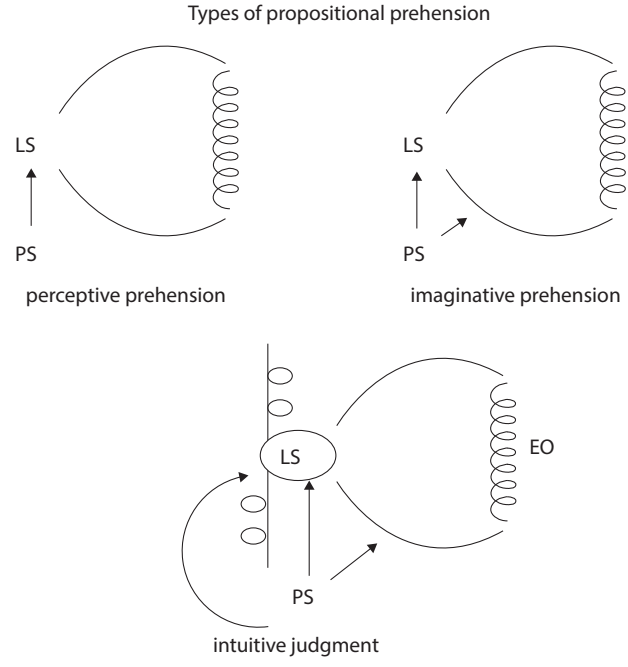

Figure 12. Types of propositional prehension (Fortescue 2001a)

interacting with experience, as accumulated in memory and imbued with emotional value. This is Whiteheadian 'embodiment'. Presentational immediacy, on the other hand, is the more sophisticated, clear and distinct consciousness of the extensive relations of the world, it is precise, spatially located, but cut off from the past and future. As a later stage of concrescence it seizes upon the vague but powerful input of causal efficacy and projects it into a contemporary region of space in sharp sensory detail. Both are involved in the mixed mode of 'symbolic reference', which is our ordinary mode of perceptual awareness (and not just a linguistic phenomenon). It provides a sense of bodily relevance to mere sensory patterns.[12]

But now back to our principal subject. Whitehead saw abstraction in terms of what he calls 'transmutation', a process whereby complex arrays of sensory input (prehended as 'physical feelings') are integrated and simplified in order to be experienced as unified 'conceptual feelings', as illustrated in Figure 13 (Sherburne

12. As Sherburne (op. cit.: 113) puts it: "The crux of symbolic reference is that the *sensa* that define regions in the contemporary nexus are not given in the mode of presentational immediacy, but are, rather, donated from the mode of causal efficacy and 'projected' onto the contemporary nexus revealed in presentational immediacy."

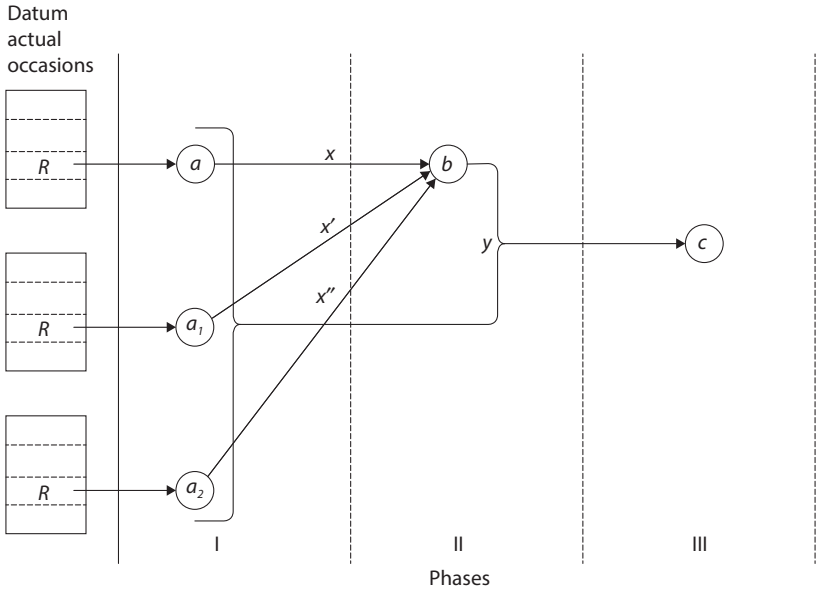

Figure 13. Whiteheadian 'transmutation' (Sherburne 1981)

op. cit.: 74).[13] Conceptual feelings constitute the basis of all higher cognition, and are arrived at by analogy (of datum to eternal object/ concept) – this involves the similarity between the various elements of the input 'datum' and the elimination of differences between them (Whitehead 1978: 251–254). As stated above, the 'datum' can be perceived either in the mode of 'causal efficacy' (vague subjective memories of bodily engagement with objects in the past) or 'presentational immediacy' (the distinct 'here and now' of sensory experience).[14] The continual transmutation from causal efficacy to presentational immediacy (in the process of symbolic reference) is the usual way in which we experience the perceptual world in terms of interconnected 'nexūs' (systematic schemes of relatedness), for example of individual objects, not just of successive actual occasions (op. cit.: 121). In fact transmutation presupposes such nexūs. This is the macrocosmos as we experience it as opposed to the microcosm of actual occasions which constitute it. As Sherburne puts it: "Transmutation is an operation whereby an aggregate of many actual occasions, forming a nexus, is prehended not as an aggregate, not as a many, but as a unity, as one macrocosmic entity" (Sherburne 1981: 73). This could be, for

13. Whitehead's usage is not to be confused with 'transmutation' in the more usual sense of transformation of one element into another (as in alchemy and radiochemistry).

14. These can perhaps be related respectively to the 'where/how' and the 'what' pathways of perception distinguished in contemporary neuroscience.

example, a particular tree persisting through time, with all the sensory qualities attributable to it. He further paraphrases Whitehead (op. cit.: 76): "Our usual way of consciously prehending the world is by these transmuted physical feelings. The intensity arising from the force of repetition makes this transmuted perception to be the prominent type of those feelings which in further integrations acquire consciousness as an element in their subjective form. It represents a simplification of physical feeling, effected in the course of integration."

Whitehead chooses as a simple example a perceptual datum of something red. It may contain various shades of red and different shapes and textures within the area which activates the concept of redness, for it need only do so approximately, by analogy with other experiences of objects exhibiting approximately the same wave length, eliminating any irrelevant differences. Of course this is a considerable over-simplification – see Varela et al. (1992: 157ff.) for a detailed account of how colour perception does not directly reflect wave-lengths of things in the outer world, but depends strongly on context and on the precise workings of a the distributed network system for maintaining 'colour constancy', etc. Colour is an emergent property of this network and can hardly be isolated from other aspects of visual processing.[15] The point is that this datum is experienced as exhibiting the '*qualia*' of redness by a higher level abstraction from simple 'physical feelings' – it is experienced via a 'transmuted feeling', and like all perceptual 'feelings' this has a 'subjective form'. On a neurological interpretation the process could be seen as one of abstracting a pattern from input that can be related (perhaps via a hidden layer of essentially arbitrary connections) to the neural assembly in the inferior temporal lobes that constitutes the concept 'redness' – in Whiteheadian terms the corresponding 'eternal object'. What is involved is the matching of input with previous experience of red objects 'transmuted' into a determinate but fluid memory trace capable of producing a unitary kind of response. This represents an essential step on the way to fully conscious perception, a gain in the power of abstraction that simple organisms are not capable of. The final step to full consciousness will according to Whitehead require the development of language and the kind of focus and distancing from the fuzziness of accumulated exemplars that words – e.g. 'red' – provide. For transmutation produces vagueness: the necessary grounds for the prehension of the relevant 'eternal object' is the recognition of only an approximate fit, one better at least than the potential fit with any competing candidate. This results in prototype effects, which is precisely what cognitive

15. Compare Davidoff's (1997) neurological model hypothesizing two separate modules, one for colour memory as such, the other for object memory containing colour information, both of which may provide information for colour naming in a given language.

linguists would want to see for linguistic as well as general conceptual categories (Fortescue 2001a: 69).

As regards 'projection', which I have argued for earlier as inseparable from abstraction, Whitehead uses the term to refer to the projection of sensory data in the mode of 'presentational immediacy' onto the vague sense of the outer world inherited from the immediate past in the mode of 'causal efficacy'. In other words it is again a matter of 'symbolic reference'. This is the reverse of the Humean input-output view of perception shared in our day by adherents of the digital computer model of the brain. (Hume made the sense of 'causal efficacy' an experience dependent on presentational immediacy rather than the reverse.) Whitehead sees our experience of body parts as primarily perceptions of them as 'reasons' for projected *sensa* (the eye is the 'reason' for the projected sight-*sensum*, etc. – Sherburne op. cit.: 124). The projection of the senses in 'presentational immediacy' is subject to error – it depends upon the state of the brain, which may hallucinate. Projection can be faulty, but this may be the price paid for the advantage of the increased focus of imagination that presentational immediacy provides.

10.4 'Transmutation' applied to the mental lexicon

Why have I gone out of the way to bring Whitehead into the picture at all? The answer lies in the theory of 'transmutation' sketched above, which, as we shall see in the final chapter, is the key to bringing together all the various senses of 'abstract' and 'abstraction' that we have seen so far. I need now to show how transmutation is involved in the way in which we focus the projection of otherwise 'fuzzy' concepts onto objects in the environment with the help of words. I shall use another diagram of the kind I introduced in 8.4., namely that for the word 'parrot' (Fortescue 2009: 15), given below in Figure 14. Though this is less 'fuzzy' than that for the hypernym 'bird', it is still applicable to variable instances around a 'prototypical'

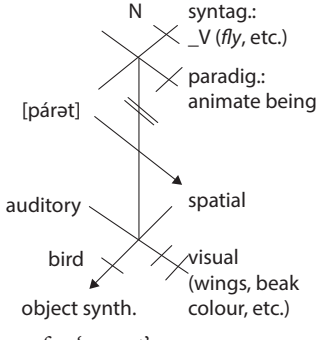

Figure 14. A mediatory column for 'parrot'

core (the flightless kakapo of New Zealand, for example, is decidedly peripheral to the category).

Here the lower half of the diagram indicates the sensory features of the prototypical parrot – the axes for an 'object' word are slightly different from those for 'event' or 'state' words, so besides such visual features as wings, beak of a certain shape, typical colours, etc., there is also indicated the hypernym 'bird' on an axis of 'object synthesis' (broadly lying along the temporal lobe and organized in a posterior-to-anterior direction). The phonetic shape (here simplified as [párət]) enters the column at mid position as before. The grammatically relevant features of the word are indicated further up the column – here the frequent syntagmatic combination with verbs like 'fly' and its paradigmatic status as indicating an animate being (this will constrain the types of sentences it can normally appear in).

This whole network will be activated in the process of 'transmutation' when there is a 'near enough' match between sensory input from a parrot in the observer's environment and the sensory features stored on the column in his or her long term memory, ready for the simulation of parrots. What transmutation does is to summate the various sensory 'affordances' presented by the object in the observer's field of vision and to reduce them to a single 'nexus' experienced as a unitary entity – an exemplar that can be matched to the generalized memory trace associated with the mediatory word column. It may even add some further detail to that trace, which is itself the 'summation' tuned over numerous instances (whether experienced live or from books or the visual media). Of course the word 'parrot' has many more diffuse associations than schematized here (the countries and types of environment they are likely to be found in, perhaps also experiences with individual parrots, etc.), but these are not immediately relevant to the activation and conventional use of the word 'parrot' – except in so far as the context is different from those in which parrots have previously been encountered.[16] As regards the general schematization of the process in Figure 13, this needs to be understood as the transmutation of individual roughly equivalent sensory data (analogous to 'pixels'), so similar parallel processes will be involved in all the channels involved in sensory perception, to be reunited as the holistic experience of an object – probably through synchronized oscillatory 'binding'.

The same process could be applied in theory to the analysis of linear speech input. Here a number of overlapping (presumably incremental) stages are involved. First, the initial data to which transmutation applies is the raw acoustic signal, consisting of modulating formants interspaced with short bursts of noise, 'clues' which have to be continuously transformed into strings of discrete phonemes for the next

16. This further information may be stored contralaterally, as suggested by the Contextual Symmetry hypothesis.

stage, the transmutation of phonemes into unitary morphemes, words and phrases. Again, as in the case of visual 'data', numerous parallel pathways of analysis can be envisaged at different levels of 'granularity'. The result will be the isolation of verbal 'objects' which can be matched to the contents of long-term memory, and meanings-in-context be evoked from their associations and combinations.

My contention that transmutation is not just relevant to our perception of individual 'objects' but also applies to the recognition of events or states can be illustrated with another mediatory column diagram, this time for 'climb', as in Figure 15 (Fortescue 2009: 104). Although Whitehead himself analysed event and state relations as the predicates of propositional prehensions, relating them directly to 'eternal objects', my cognitive interpretation sees the 'ingression' of eternal objects themselves in networks of neural columns. Not only kinds of object but also the kinds of properties they display and the kinds of relations between them are abstracted from experience. Predicates simply presuppose objects to which they may apply, as 'fly' presupposes a flyer, say a bird. The meanings of 'concrete' verbs of physical action are certainly no more abstract than that of many of the 'abstract' nouns we entertain in memory so there is no reason why they should be instantiated in a radically different way. My application of the process of transmutation to events should be understood as somewhat unorthodox. It should also be acknowledged that the neural processes whereby events or states are recognized as such are also more obscure to neuroscience than is the case with object recognition.

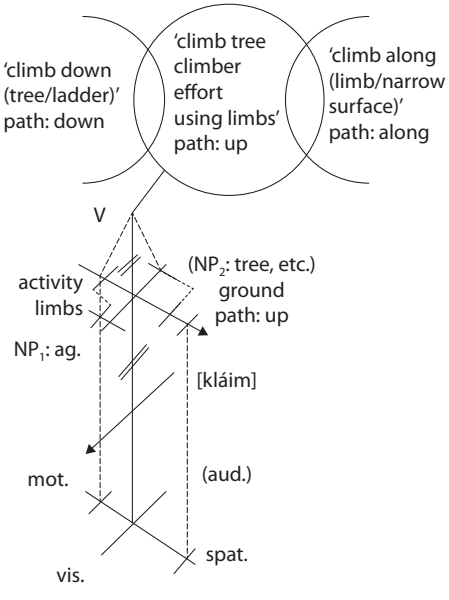

Figure 15. A mediatory column for 'climb'

In Figure 15 there is a single word ([kláim]) displaying a variety of sensorimotor interpretations while maintaining a common schematic core that may or may not represent a single 'image schema'. The latter would be indicated on the column by a specific congeries of features on the sensory axes at the base of the figure, but here I have left them empty (apart from the broken lines that indicate the projection upwards from general sensory features of motion and spatial 'up'-ness) in order to concentrate on the higher levels relevant to the use of the word. First, as usual, the word – a verb of motion – displays certain syntagmatic and paradigmatic features as indicated: it will be found in the syntactic context of an NP indicating the agent of the action and another indicating the ground of the motion (either a bare nominal like 'tree' or a prepositional phrase like 'up the ladder or 'along the ledge'), also optionally an adverbial indication of the path like 'up', which corresponds to the paradigmatic feature of that meaning. (In the prototypical case climbing is generally upwards.) Another prototypical paradigmatic feature is 'by the activity of the limbs'. Three distinct uses of the word are indicated in the scenario 'balloons' at the top: the most prototypical one involves climbing up a vertical 'ground' like a tree, with coordinated effort applied to successive leg and arm movements.[17] The two overlapping scenarios to the left and right concern the less typical notions of climbing down something (like a ladder) or along something (like a ledge). Note that many languages cannot apply their 'climb' verbs to these extended scenarios. Each of the three scenarios will have a particular 'feel' to them in internal simulations. The outcome of transmutation is something meaningful to the prehending subject, something that can be felt and simulated. There does not seem to be much point in calling what they have in common a generalized 'image schema' since they are each 'embodied' somewhat differently – just as the different meanings of 'open' discussed in 4.4. will be simulated differently according to the context or scenario in which it is used.

Each of these three scenarios can be recognized and experienced as unitary event types. The associated word 'climb' binds them all together, but on its own it does not distinguish the various extended scenarios or contexts to which it applies. What transmutation broadly understood can do in order to distinguish the scenarios themselves as unique despite their overlap (and still more diffuse associations in episodic memory) is to highlight only those features of the 'input' that match with the scenarios individually and to suppress the alternatives that do not maximally fulfil the match. (Recall that transmutation is an approximate process operating by 'best fit'.) Thus if the motion is upwards then the 'downwards'

17. I will not consider here purportedly 'metaphorical' extensions from a human climber to an inanimate one like an aeroplane. The case of birds climbing in the sky and insects climbing up stalks suggests that we are dealing here with under-specification rather than metaphor as such.

and 'along' scenarios will be suppressed as irrelevant. Transmutation, as I interpret it, is involved both in the isolation of perceptual/ conceptual event types and (by separate pathways) in the recognition of lexical words that may be associated with them. (Both involve physical, i.e. perceptual feelings.) And most important for present purposes is its role in recognizing context – for a 'transmuted feeling' can presumably be of any kind of nexus felt as a unity, not just 'objects' but also the broader perceptual nexus in which they are embedded. To fully understand the relationship between object and context (which may involve propositions at a higher level of abstraction) further 'intellectual' prehensions will be required.[18]

At the heart of transmutation is the relationship of 'contrast'. As we have seen, transmutation is an operation whereby an aggregate of many actual occasions, forming a nexus, is prehended not as an aggregate, not as a many, but as a unity, and the nexus, as one, is in contrast with (i.e. corresponds to) an 'eternal object'. "Whatever is a datum for a feeling has a unity as felt… this unity is a 'contrast' of entities. The third and fourth phases of concrescence are phases of 'integration and reintegration'. Feelings in these phases have complex groupings of earlier feelings in the concrescence as their data; this is to say that feelings in these phases have 'contrasts' as data." (Sherburne 1981: 51). In other words, concepts – contrasting with other concepts – are projected on the perceptual world in our everyday experience of it as made up of objects and relations. The 'integration and reintegration' can be understood as referring to the neurological 'binding problem' whereby multiple pathways of analysis of perceptual data are integrated at a subsequent stage.[19] But the concepts themselves, these contrasting 'forms of definiteness', are surely derived initially from the interplay between our biological inheritance and our perceptual experience of the world – the traffic is two-way, in fact three-way with the further involvement of language. Through transmutation even overlapping and diffuse conceptual networks may function as 'forms of definiteness' – as long as such activated networks contrast with others that are activated by somewhat different arrays of input 'data'. Without transmutation the world would be a blur, lacking all contrast between perceptible categories. I shall return again to the relationship between transmutation and abstraction in 11.5.

---

18. Recall the definition given earlier of the complex comparative (or 'intellectual') prehension in which a nexus of actual occasions (that in which the original datum is 'embedded') is contrasted with a proposition, the logical subject(s) of which is/are contained within the nexus.

19. Note that these later stages for Whitehead do not always involve high level 'intellectual' prehensions but may result in 'physical purposes' instead, in which – unlike in propositional prehensions – the eternal object concerned in evoking a feeling of adversion or aversion recedes "back into mere immanance", and the concrescence is 'satisfied' at the 'simple comparative' stage on Figure 11 (Sherburne op. cit.: 65).

10.5 A nudge from Wittgenstein: Similarity and context

Of cardinal importance for understanding abstraction is the notion of semantic 'similarity', since it is 'similar' things that are abstracted and generalized as categories, types or 'ideas'. Wittgenstein, considering the various meanings of the word 'game' (board-games, card games, ball games, children's games, etc.), famously analysed similarity in terms of 'family resemblance' (Wittgenstein 1974: 32). This boils down to the statement: "'This and similar things are called 'game'", as made when showing different games to the receiver of this ostensive 'definition'. It leads directly to the emphasis on prototypes and domains or contexts of use which characterizes Cognitive Linguistics today. Ordinary Language philosophers were also concerned with such matters, in particular with expressions of appearance vs. reality and with sensory and inferential evidence as expressed in everyday English. Thus Austin distinguished alternative expressions like "He looks/ appears/ seems guilty" and "The hill looks/ appears/ seems steep" in terms of what today would be called different kinds of evidentiality (Austin 1962: 36–7). He analyses 'look' as 'to have the (literal) look of', 'appear' as more or less the same but with an added nuance of 'under special circumstances' (which may suggest a broader reality behind mere appearance), and 'seem' as a matter of actual evidence (of any kind) bearing on whether the person really is guilty or the hill really is steep. Note that analysing all of these as reflecting the same underlying semantic 'primitive' would be difficult to justify.

Peirce, with his theory of the sign, introduced a three-fold distinction between icons, indexes and symbols. Indexes are signs that simply 'point to' a target by dint of association, like smoke to a fire, whereas symbols are arbitrary signs, such as words, 'standing for' a referent. All three are 'abstract' in one sense or another. Indexes and symbols are not relevant here, not being based on similarity, but icons are. Peirce sub-divided icons into images, diagrams, and metaphors. Many logograms and sign language signs are imagistic, as we saw in 2.5 (pictograms are better exemplars), but some involve schematic analogy rather than direct depiction, and these (like my 'mediatory column' diagrams) must qualify as Peircean diagrams. Like maps, these "represent the relations… of the parts of one thing by analogous relations in their own parts", i.e. involve a somewhat more abstract or indirect form of similitude (Peirce 1932: 2.277). We have already discussed metaphors in some detail. Philosophers like Davidson (1978) have further delved into the relation between simile and metaphor, and the basis of (non-metonymic) association and analogy has been investigated in considerable detail by Itkonen (2005), both within and outside of language.[20]

20. Itkonen makes the important distinction between analogy as structure and analogy as process, the latter a dynamic cognitive matter as seen in inferential processes such as (especially)

Many languages have grammaticalized systems of evidentiality, some (especially in the Amazon basin) with obligatory markers that indicate whether the information referred to was acquired by first hand witnessing (by sight or by hearing) as opposed to by inference or received through hearsay. In other words, the exact channel by which information was abstracted from an event or situation in order to be expressed in language has to be indicated in these languages (many of these are discussed in Aikhenvald 2004). In Fortescue (2010b) I discussed two languages with a variety of evidential morphemes (though none of these can be said to be grammatically obligatory). The languages are Japanese and West Greenlandic. The former has three suffix-like elements: *rashii* 'seeming, looking, acting like' (actually an adjective, primarily used of the here and now), *-soo* 'looking like' (a suffix, usually referring to the near future), and *yoo* 'seeming like (an independent nominal referring broadly to any of the senses except hearing). There is also a clitic *soo* in the sense of 'by hearsay'. West Greenlandic has suffixes *-(r)palaar* 'appear to be, act like a –' (also on verbal stems 'one can hear V'), *-(r)palug-* 'look, seem, act, sound like a –' (to an adjective or noun, also to verbal stems in the 'look like' sense), *-(r)pallag* 'act like a –' (which also has a 'hearsay' and a 'one can hear' sense with verbal stems), and *-gunar-* 'look like, seem to, evidently (usually visual but also from general evidence or logical inference)'. It also has a dedicated inflectional case of similitude: the so-called 'aequalis' case ending *-tut*, which, amongst its various functions, forms manner adverbials ('(do) in the manner of –').

Within the cognitive sciences the concept of similarity has been quite central, with perceptual categorization, prototype and metaphor theory figuring prominently (e.g. Taylor 1989 and Schroeder 2007). At the heart of this lies the question of what constitutes 'similarity' in relation to semantic extensions of more 'basic' meanings. Taylor's definition of metaphor as "perceived similarity across domains" (op. cit.: 139) is typical here. Hedged category assignment has been investigated in some detail by Lakoff (1972), and the perspectives of pragmatics and child language acquisition on the appearance/reality problem has been studied by Hansen & Markman (2004). As Gentner & Markman (1997: 48) put it, arguing for a continuum between analogy and similarity: "Analogy occurs when comparisons exhibit a high degree of relational similarity with very little attribute similarity. As the amount of attribute similarity increases, the comparison shifts towards literal similarity."

abduction. He points out parallels between language, vision and music, as well as logic, and argues against a strictly modular approach, since structures in one modality can be analogically reflected in others.

In Fortescue (2010b) I discussed the ways in which something can be considered 'like' something else and illustrated the way in which they are expressed in a variety of languages. They are shown in Table 6.

Table 6. Kinds of similarity

(a)	physically resembling by shape or other property
(b)	acting in the manner of, copying
(c)	of a similar dimension or degree (equivalent comparison)
(d)	(vaguely) reminiscent of, sort of a (hedge)
(e)	as if, seeming to (but not actually), pretending to, fake
(f)	much as on a previous occasion or at a different place
(g)	(exactly) as expected of such a person or thing
(h)	corresponding to, the equivalent (in value, function) of
(i)	according to, in the manner prescribed
(j)	likely to occur (epistemic)
(k)	by diagramic iconicity
(l)	of the same status (e.g. legal)

My conclusion was that there is a universal cognitive act of comparison which places pairs of *comparanda* on a continuum between SAME and DIFFERENT, with LIKE covering the middle ground, and that the different senses above depend on context (what kind of things are being compared). The relationship between 'comparison' and 'identification' is particularly close (or overlapping). They are nevertheless two distinct cognitive acts, namely comparing two things under one or more implicit conceptual types (what they have in common) and simply assigning a single token to an explicit type, i.e. categorization. The relationship of both to abstraction should be clear: the traits comprising the particular types involved must be abstracted from the tokens (for example by sensory inspection).

Bybee (2010: 57–75) specifically looks at the role similarity and analogy play in the introduction of novel items in known linguistic constructions during the development of children's language (as mentioned in Chapter 9), but she also considers the repertoire of adults, on the level of both morphology and syntax. This she does in the context of arguing for a usage-based, exemplar perspective in which productivity is premised on similarity to existing items, not on explicit 'rules'. As she puts it: "the probability and acceptability of a novel item is gradient and based on the extent of similarity to prior uses of the construction." In fact, she argues, analogy is in general probabilistic (op. cit.: 74). The creativity of language (which has from the start been the rallying cry of Chomsky and his followers) has

been exaggerated: the role of pre-packaged units ('prefabs') and analogical extensions from them is massive, both in child language (as discussed by Tomasello 1992) and in adults. Constructions, as opposed to completely general rules, lie on a continuum from local (and completely closed off) idioms to 'free' constructions with abstract slots fillable by any suitable lexical material. Their productivity is also gradient, with no clear-cut division between regular and irregular past tense morphological forms, for example (op. cit.: 73). So it is questionable whether children ever learn completely general as opposed to statistically probable 'rules'. Analogy also plays, as is well known, an important role in diachronic change, and as Bybee points out, it is also involved in the creation of new constructions out of old ones (op. cit.: 71). Frequency is relevant here too – high frequency forms are less likely to undergo analogical change. This is not a characteristic of general rules.

Let me conclude this chapter by returning for a moment to Wittgenstein and 'family resemblance'. There is one further lesson to be gleaned from his 'game' example, namely that abstraction expressed through language – albeit based on vague similarity and/or successive acts of comparison – may ultimately be indeterminate. It may not result in the singling out of one bounded 'concept'. The meaning of 'game' is unbounded, indefinitely extendable according to contemporary usage. So what do all uses of the word have in common? Do they need to have just one thing in common at all? The word is, after all, just a word. It may refer to different concepts, whether perceptually based or otherwise, variously grouped and combined according to context. Some of the relevant factors linking them are functional and emotional: the reason one plays games and how it feels to do so in terms of competitiveness, the recognition of rules, and the sheer enjoyment of partaking in an activity that is physically or mentally fun and has no serious consequences. Which particular features – and to what degree – are highlighted in different contexts may vary greatly, but the word evokes a certain common but vague affective feeling, the 'subjective form' with which we 'prehend' it. But this too will become less vague the more closely we constrain the context – the overall feeling of playing rugby is not going to be the same as that evoked by playing bridge, though some affective features will be the shared. This merely underlines the loose but inseparable bond between abstraction and wider context that I have stressed throughout this book.

CHAPTER 11

Abstraction all the way up

Its evolutionary purpose

11.1 Structural coupling with the environment: From amoeba to man

It is time then to look at that wider context and have a crack at answering the question as to the role abstraction has played in evolution – specifically the evolution of the human brain. This will lead to a more unified picture of what has up until now looked like a scattered phenomenon under numerous guises. Varela et al. (1992) present a useful starting point in what they call 'enactivism', a compromise between direct and indirect (representationalist) realism close to the Whiteheadian perspective. They describe the organism and its environment as 'structurally coupled' in embodied action, a relation that is constantly shifting and evolving. They define enaction in the following manner: "(1) perception consists in perceptually guided action and (2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided" (op. cit.: 173). 'Co-evolution' of organisms and environments results when they are coupled in this way. Like Whitehead the authors emphasize the emergence of higher cognitive functions by successive phases of abstraction – symbols are seen as higher level (approximate) summations of properties embedded in underlying distributive systems, e.g. neural nets (op. cit.: 101). Thus each cortical area higher in the brain's (partially parallel) semi-hierarchy extracts new features from those below and constructs new, more general features out of them. They agree with Gibson (1972) that the environment presents to the perceiving organism its own 'affordances', i.e. clues facilitating direct perception without the necessity of representations, though they insist that something more is needed – the perceived environment is not entirely independent, it is "enacted by histories of coupling", and perception is sensorimotor enactment. Their approach proceeds by specifying the sensorimotor patterns that enable action to be perceptually guided (op. cit.: 203–4).

Our understanding of the relationship between perceiver and environment has been furthered in the last two decades by the discovery of 'motor neurons' (Rizzolatti et al. 1996). These respond not only to actions performed by the perceiver but also to the same actions performed by others perceived as carrying them out, thus forming amongst other things a basis for empathy. The discovery

of these neurons has aroused considerable interest among cognitive linguists and psycholinguists since they cluster particularly densely in Broca's area (and the equivalent gesture areas of the great apes). This can be related to motor theories of speech whereby speech comprehension is directly coupled to speech production (see Harley 2008: 267–8), and also to the formation of embodied image schemas. Zwaan (2008: 16) discusses the implications of the distinction between broadly and narrowly tuned mirror neurons: the former, by abstracting away from specific features of an observed action, may represent the basis for fuzzy categories, approximate analogy, and metaphor itself.

McNeill (2012) specifically sees what he refers to as 'Meade's loop' – a thought-language-hand loop with mirror neurons in Broca's area – as forming the basis of metaphoricity (abstracted from action) and as facilitating reflexive thought. In terms of his proposed 'three ages' of acquisition, the learning of 'Meade's loop' occurs only in a final stage when gestural pantomime falls away at an age of around three or four (op. cit.: 168ff.). This stage involves learning to utilize motor neuron circuits not only in the recognition of the meaning of actions by others (an essential ingredient in 'theory of mind') but – by an extra inward 'twist' – also in recognizing the 'meaning' of one's own gestures to oneself. It is this, he claims in relation to speech that allows the full adult 'inhabitance' of language. These 'twisted' mirror neuron circuits bring one's own actional gestures into the service of symbolic imagery, thus generally enhancing the range of abstraction. They also determine the structure of Vygotskian 'psychological predicates', which for him equates to the dual speech-gesture 'growth point' – it consists of a predicate within a non-verbal context that reflects the speaker's intentions/goals. In other words, they 'orchestrate' inner thought as well as communicative expression (op. cit.: 68–69). Central to McNeil's theory is the ability of humans to combine language and perceptual imagery (including the kinaesthetic) in reflexive thought and memory. This presupposes the abstracting and relating of patterns in both language and imagery – also the representation of context types. Gazzola et al. (2006) have further presented evidence that mirror neurons can be activated by sounds as well as (visual) actions, supporting the notion of a specific audiovisual mirror system.

We are now ready to consider abstraction at successive evolutionary levels leading from its earliest manifestation up to ourselves. For it should be evident by now that I do not consider abstraction to be the sole prerogative of the human mind. Recall the three levels of rationality distinguished by Bermúdez as discussed in 6.2: Level 0 referred to simple instinctive behavioural dispositions (stimulus-response tropism or instinctive foraging patterns), and Level 1 referred to the specific behavioural choices of undertaking a particular course of action as opposed to another in a given situation. Finally Level 2, true conceptual thought or 'second order reflexion', requires language and what he calls 'intentional ascent', i.e. the

ability to hold thoughts in mind and think reflexively about them. This highest level requires the implementation of 'instrumental' belief-desire couplings, for instance knowledge of how to achieve a goal using the intermediary of tools. Such thought is compositional and infinitely modulatable according to circumstance and goal, but Bermúdez claims, as we saw, that other animals (mammals at least) can have thoughts of a compositional nature with determinate content already at Level 1. Also Barsalou (1999: 606–7) sees many animals as having perceptual symbol systems that permit them to simulate entities and events in their environment. He further speaks of ontogeny loosely recapitulating phylogeny in humans, so that children, already able to produce their own perceptual simulations (e.g. in imagination) learn to construct simulations productively from other people's utterances and to convey their own simulations once language begins to be acquired.

The story starts, as I have suggested, with animals possessing simple brains and thus capable of storing memories. Single-cell organisms reacting tropistically to light and food sources are clearly not acting rationally. The borderlines between Bermúdez's levels are not clear-cut (mainly because all creatures capable of higher level rationality *also* continue to use that of lower levels). However, it is reasonably clear that organisms displaying Level 0 rationality, such as birds in their foraging activities, do abstract information, patterns and expectations from their environment and store them 'for future use' as the basis for rudimentary 'decisions' when choices present themselves. Invertebrates and reptiles with simple brains may also qualify in so far as they can abstract maps of their territory and distinguish 'good' from 'bad' places on the basis of past experience. But unlike them, birds (or some of them) would seem capable at times of level 1 behaviour, momentarily suppressing purely instinctual responses to choose a particular course of action rather than another in a situation where either way seems a possibility – certainly mammals appear capable of this. The kind of abstraction this requires is not just of responses to specific 'good' vs. 'bad' situation types, but the ability to assess on the basis of past experience whether the current situation really does present sufficient grounds to choose the one course of action or the other (e.g. fight or flight). But this is not true Level 2 rationality, which is reflexive and requires self-awareness as possessed by adult humans. Where do the great apes belong on this scale? Though they lack what we as humans would unambiguously call 'language' they may well be capable of some degree of reflexive thought – recognizing themselves in mirrors as they can, for example, requires a fair degree of abstraction (not to mention mirror neurons!). As I have said, the borderlines are somewhat blurred here, but the ability that certain chimpanzees in captivity have to learn to manipulate visual symbols in an instrumental fashion surely counts as the dawning of Level 2 rationality.

The employment of symbols (rather than the acquisition of human-like grammar) is arguably what is criterial here – a view promoted by Deacon (1997). This is the point at which spatiotemporal image schemas are not only abstracted from interaction with one's environment and stored for future reference, but where they can be manipulated independently in novel combinations in imagination. What human language contributes over and above purely perceptual symbols is the ability to control this symbolic behaviour in a precise and open-ended manner, both in private thought and social interaction with others. It becomes possible to evoke from memory simulations of the past and based on them make predictions and plans about the future – and hold on to them. The greatly expanded memory capability of humans plus language provides the grounding of culture, held not just in one brain but in the inherited traditions of society.

Whitehead's analysis of the 'order of nature' in terms of a hierarchy of social nexūs – from the purely physical to the mental – is conveniently summarized in Sherburne (1966: 72–89). Within each level typical kinds of prehensions are taking place synchronically at their own scale and speed, every level of 'society' presupposing a higher one in which it is embedded, wheels within wheels – fractally, as it were. (Recall the definition of 'prehension' in 10.3: the way an 'actual occasion', the basic entity of reality, relates and reacts to any other item in the universe, how it absorbs it into its own internal constitution.) Whitehead traces the most basic prehensions present already within the lowest grade of inorganic nexus – that of oscillating atoms within rocks, for example, characterized by the massive 'objectification' of impinging environmental factors. So this clearly goes much deeper than the roots of perceptual/ conceptual abstraction that we are concerned with. However, when it comes to life forms, Whitehead can be interpreted as agreeing with Bermúdez that pre-linguistic animals are capable of the equivalent of level 1 rationality but not of level 2 (which presupposes the 'intuitive judgment' and language to support it).

He traces the emergence of life from non-life (extending the insensate physical responses of inorganic matter) to a stage where 'appetition' arose in societies (here = bodies) structured in such a way as to react to novel elements in the environment with internal novelty to match it. In Whitehead (1978: 103) he discusses the more subtle relationship between 'living societies' (e.g. simple cells made up of inorganic 'societies' that support them) and 'entirely living nexūs' (animal bodies, in which a superordinate 'regnant' nexus is in control and may make novel decisions). The latter is surely the prerequisite for Level 1 rationality. Higher organisms (approaching Level 2) add a further reflexive level to this cumulative process, developing the ability in humans to *think about* their diverse experiences. Evolution has generally gone from pure causal efficacy amongst the lowest forms of life (gross inheritance from occasion to occasion with little ingression of novelty) towards

integration – with sharper and sharper precision – of the senses in the mode of 'presentational immediacy', ultimately resulting in the capability for 'symbolic reference' in higher forms of life. Characteristic of the latter is a fluid interaction with novelty in the environment, something digital computer models, for example, lacking imagination and the possibility of choice among multiple reactions to the same input, cannot conceivably manage. Algorithms, however complex, are not enough. But nor is abstraction in isolation. Witness the following passage:

> Mankind is distinguished from animal life by its emphasis on abstractions. The degeneracy of mankind is distinguished from its uprise by the dominance of chill abstractions, divorced from aesthetic content…. A fortunate use of abstractions is of the essence of upward evolution. But there is no necessity of such good use.
>
> (Whitehead 1938/1968: 123)

Singer (2013), coming from a different direction altogether – the emergent dynamics of complex systems – also emphasizes that the brain is more than a stimulus-driven digital computer (op. cit.: 621–3). In arguing for a higher dimensional 'reservoir' organization of the brain he suggests a possible evolutionary advantage of systems in which oscillating nodes allow increasing dimensionality by introducing 'phase' as a variable. This could have resulted in a vastly increased memory capacity (as well as facilitating matching operations). Recall his definition of the brain in 8.1 as essentially a pattern-generating system. He points out that one of the prerequisites for such a system is already present in neural 'pacemakers' in invertebrates. When coupled with the environment these oscillating neurons generate complex, well-coordinated spatiotemporal activity patterns (op. cit.: 623).

11.2 Abstraction and consciousness

But where in the preceding does consciousness enter the picture, and what is its role in the abstraction of pattern from the environment? Jackendoff, while not concerning himself directly with the environment, has a unique perspective on consciousness as it relates to the different levels of functioning of the brain. I have already mentioned in 8.1 his notion of the 'f-mind'. In Jackendoff (1987) he developed his 'Intermediate Level' theory which lies behind this. As we shall see, this is in many ways compatible with Whitehead's understanding of consciousness, adumbrated in 10.3. The nub of this theory is that consciousness (awareness) does not occur at the highest level of thought but at an 'intermediate' level between the unconscious (and rapid) low-level functioning of computational modules and the equally inaccessible (but slower) functioning of thought as the complex understanding and manipulation of meaning. Awareness is borne by vehicles appropriate

to the individual modules of (at least) language, music and vision. In the case of language it is the phonological word – that is the level at which we are conscious of thinking (syntax and semantics, necessary for understanding, are largely hidden to awareness). As regards music it is the musical surface, and as regards vision Marr's '2 ½ D sketch' mentioned in 6.1.[1] These – the 'privileged' contents of short term memory – represent the nearest conscious approach we have to the brain's operations at the multiple levels of linguistic, musical and visual 'meaning'. He argues against an alternative 'all levels at once' theory of consciousness, and suggests that 'inner thought' is precisely experienced as a succession of phonological 'linguistic images' from which richer meaning can be projected, not directly in terms of manipulating conceptual structure (op. cit.: 289–91). The contents of 'richer' representations are simply *too* rich for short term memory to entertain.

The 'modules' Jackendoff refers to (and he suggests others) are much like those of Fodor described in 2.3, but are – unlike his – only partially 'encapsulated', interacting among themselves (but always at an implicit level). Each has its own internal hierarchical organization of sub-levels. There is also no 'slush' of central processing in the manner of Fodor, but rather a small number of relatively independent domains over which 'slow' processes are defined (op. cit.: 265–7). These are responsible for the integration of the output of computations into long term memory.

Jackendoff traces the emergence of this intermediate level of awareness through several stages on the way 'up' from the unconscious workings of 'the computational mind'. One essential ingredient is high-intensity 'attention' that selects the most 'salient' contents of on-going computation to provide the (partial) input to richer awareness. He summarizes his theory as follows (op. cit.: 298):

> The distinctions of form present in each modality of awareness are caused by/ supported by/ projected from a structure of intermediate level for that modality that is part of the matched set of short-term memory representations designated by the selection function and enriched by attentional processing. Specifically, linguistic awareness is caused by/ supported by/ projected from phonological structure; musical awareness from the musical surface; visual awareness from the 2 ½ D sketch.

Awareness understood in this manner does not directly reflect 'what is going on in the world' *or* in the mind – pure 'bottom-up' sensation in the manner of Hume is precluded. Unconscious 'understanding' unifies the separate modalities of awareness by 'parsing out experience across modalities', registering simultaneous manifestations of the same 'object' in different modalities (e.g. vision and touch). Awareness reflects an "amalgam of thought and the real world" (op. cit.: 300–1).

1. He argues that the full 3D image is analogous to the richer meaning of syntactic sentences.

The haptic sense of 'touch' is of particular relevance here, as being in some sense the most concrete of the senses. We ultimately rely on it for testing physical 'reality', as there is no physical distance between the object presented to the senses and the body of the 'prehender'. But this too may be the source of subtle illusion – touch too involves abstraction on the way between nerve endings and 'matching' with relevant memory traces and integration with input from the other senses (especially vision).

In the final formulation of his theory Jackendoff adds one more element, crossing the line decisively from the computational to the phenomenological: the affect or 'feel' of phenomenal entities (op. cit.: 304–310). In effect he distinguishes the 'form' from the 'feel' of such entities, giving as an example 'tip of the tongue' experiences when the feel but not the form is accessible to awareness. Affect is what distinguishes images ('from above') from percepts ('from below'), 'inner' from 'outer', 'ego-initiated' from 'non-ego-initiated', 'meaningful' from 'meaningless', 'familiar' from 'novel', 'effortful' from 'easy', and 'congruous' (or 'true') from 'incongruous', as well as 'affective' (positive or negative, 'nice' or 'yucky') from 'neutral', i.e. all manner of experiential *qualia*, all aspects of reality 'seen from within'. It cuts across forms, categories and even modalities. While central processes provide the content of thought and privileged 'linguistic images' provide the form, a special 'affect monitor' provides the 'feel' of entities held in short term awareness.

It is here that Jackendoff's thinking is closely parallel to Whitehead's, namely in the role of feeling in conscious awareness. For Whitehead this is the 'subjective form' adhering to all prehensions – including the higher level 'intellectual' propositional prehension relevant here. As we have seen, consciousness is for him an emergent property, the subjective form of the highest forms of prehension. In so far as all prehensions involving transmutation are abstractions, all such abstractions will have an integral element of subject form, its affect – in effect that which provides phenomenal experience with their experiential *qualia*. Transmuted physical feelings are, as discussed in 10.3, our usual way of consciously prehending the world in terms of objects and their properties. The question remains whether there are forms of abstraction *not* involving transmutation – I shall return to this in section 11.5.

Jackendoff does not explain how his 'affect monitor' might work, but he schematically represents in his Figure 15.2 the role it would play in introspection, parallel to 'ordinary' meaning abstraction processing.[2] For Whitehead such a separate 'monitor' would be superfluous since all prehensions intrinsically have a subjective form. It is at all events the addition of an affect component feeding an 'introspective processor' that Jackendoff's model of thought provides in order to

2. Possibly the amygdala might function that way – recall its role in emotional memory in 5.3.

produce what he calls 'conceptual structure$_2$'. Unenriched by such a component, the straightforward abstraction of 'conceptual structure$_1$' from auditory input (via 'translation' to syntactic and then conceptual representations) would result only in 'subliminal perception' on the basis of which it would be impossible to report the existence of the stimulus while at the same time (paradoxically) 'understanding' it (op. cit.: 315–317). Conceptual structure$_2$ is what can be verbally reported on, it is 'being aware of being aware of something', awareness as it were "looping back on itself" to produce full consciousness. Yet (and here is another surprising aspect of Jackendoff's model), this output as 'conceptual structure$_2$' (the 'stuff' of introspection) would actually be less rich than the 'privileged' contents of intermediate level short term memory, it would be fairly 'rudimentary' and 'pale' (op. cit.: 316).

The stream of consciousness is for him just evidence of deeper levels of thought taking place. The rules of grammar, like those of logic, are unavailable to direct awareness, though hypotheses about them can be couched in language which can of course be the object of awareness. One might say that the grounds of reality and truth can be intuited but not directly experienced, though we can abstract content from strings of words consciously held in short term memory. Language is not itself the 'medium of thought', it is in Jackendoff's eyes an *aid* to thought – it structures it, facilitates memory, and allows us to manipulate 'concepts about concepts' and to combine concepts in novel ways. Here, I think, we can equate 'concepts' with Whiteheadian 'eternal objects', which are also only accessible via specific kinds of prehension – and full consciousness requires language (read: phonological forms) in order to do so. But language also has its pitfalls such as making sharp conceptual distinctions where no such sharp boundaries exist 'out there'. The phonological forms of words are simply 'in registration' with images in other modalities that form part and parcel of their 'meaning'. In sum, consciousness for Jackendoff is not an especially high-level process in itself and not exactly what 'makes us human' (op. cit.: 325). It presupposes the unconscious computational mind of neural operations as its basis.

Whitehead, in a more reader-friendly context than "Process and Reality", has the following to say on abstraction (in a non-technical sense) and consciousness:

> The growth of consciousness is the uprise of abstractions. It is the growth of emphasis. The totality is characterized by a selection from its details. That selection claims attention, enjoyment, action, and purpose, all relative to itself… In the entertainment of abstractions there is always present a preservative instinct aiming at the renewal of connection, which is the reverse of abstraction. This reverse process, partly instinctive and partly conscious, is wisdom of that higher life made possible by abstraction… In the dim recesses behind consciousness there is the sense of realities behind abstraction. The sense of process is always present. There is the process of abstraction arising from the concrete totality of value

> experience, and the process points back to its origin. But consciousness, which is the supreme vividness of experience, does not rest content with the dumb sense of importance behind the veil. Its next procedure is to seek the essential connection within its own conscious area. This is the process of rationalization. This process is the recognition of essential connection within the apparent isolation of abstracted details. (Whitehead 1938/1968: 123–4)

Neither Jackendoff nor Whitehead have anything specific to say as to how consciousness might be instantiated in the human brain, but Damasio (2000) does. He distinguishes non-verbal 'core' consciousness as in the 'here and now' of unreflecting self-awareness in young infants (and indeed in adults much of the time). This may well be present in other animals than just humans. But the expansion into awareness of an 'autobiographical' past and future afforded by 'extended consciousness' is typical only of adult humans possessing language, although it may be present in simple form in some animals (op. cit.: 16). Extended consciousness does not first arise with language, however, though it is doubtless enhanced by it (op. cit.: 185). Consciousness begins for Damasio when brains acquire the power to 'tell stories' without words, using the non-verbal vocabulary of body signals (recall Barsalou's 'perceptual symbols'). The self emerges as a "feeling of a feeling" (op. cit.: 30–31). We have seen how Bermúdez' Level 1 rationality presupposes abstraction, and this corresponds to the kind of pre-verbal 'core consciousness' Damasio discusses. It is probably faster than our verbally enhanced extended consciousness, but still slow compared to entirely subconscious neural activity. He points at specific regions of the 'old brain' (all close to the midline) that are crucially involved in core consciousness, as opposed to the more broadly distributed areas involved in the higher 'autobiographical' memory of extended consciousness (op. cit.: 234–276).

We are led finally to address the concept of the 'self', among the highest abstractions we humans are capable of. Or is it? Is it abstract$_1$ or abstract$_2$? Or neither? In so far as we associate the concept with our conscious experiences as a 'thinking body' it can hardly be 'disembodied from sensory features', it is right there at our fingertips during every hour of our waking lives, it is nothing if not 'embodied'. And on the other hand it can hardly be said to be 'generalized across instances' – every individual experiences it as a unique entity. My 'self' is not your 'self'. Does it correspond to anything real at all? For Damasio the 'self' has a concrete neural basis – he makes some very specific suggestions as to the location of the parts of the brain responsible for what he calls the 'proto-self', the seat of core consciousness (Damasio 2000: 153–156). These include several brain-stem nuclei, the hypothalamus, the insular cortex, as well as the medial parietal cortices: "The ensemble of these cortices in the right hemisphere holds the most integrated representation of the current internal state of the organism at the level of the cerebral

hemispheres". Also for Whitehead the self is something quite concrete, namely the 'entirely living nexus' as described in 10.3.

For Dennett, however, a self is an abstraction "defined by the myriads of attributions and interpretations (including self-attributions and self-interpretations) that have composed the biography of the living body whose Centre of Narrative Gravity it is" (Dennett 1991: 426–7). This corresponds to Damasio's 'extended consciousness', not to 'core consciousness', and in so far as this is the mode of functioning of the 'autobiographical self' it is based on memory and imagination, and as I have argued throughout this book, these are indeed laid down by and evoked by processes of abstraction. Dennett's own 'Multiple Drafts' theory of consciousness denies the concrete existence of a unitary 'self'. For him all varieties of mental activity are "accomplished in the brain by parallel, multitrack processes of interpretation and elaboration of sensory inputs. Information entering the nervous system is under continuous 'editorial revision'" (Dennett op. cit.: 111). These parallel pathways include those that were described in connection with visual processing in 6.1. What Dennett's theory lacks here is any kind of final 'binding' to produce holistic visual experiences passed on to consciousness. His model keeps the parallel analyses as separate multiple 'narratives' that may or may not be transformed or revised and integrated into memory, perhaps overwriting or influencing other contents in the process (op. cit.: 134–5).

Jackendoff is wary of pronouncements about a 'self' which cannot ever be grasped in its entirety, but he too makes it clear that it is at all events not to be equated with some unitary 'executive control processor' (Jackendoff 2002: 22). Perhaps the most that one can say is that the 'self' is something 'real' although it cannot be directly experienced except by intuition via experience of the body and inner thought and feeling. As a cultural artefact 'clothed in words' it can most certainly be examined and discussed, as can such equally ineffable abstractions as 'God' or 'The Universe'. I shall return to 'abstract artefacts' in the following section.

11.3 Problem-solving and World 3 abstraction

We have seen many kinds of abstraction *en route* so far, but what do all they have in common, what is it that is abstracted and to what purpose? Abstraction is not just extraction of individual features from the environment, as in simple forms of perception, it can involve whole generalized schemas exemplified in particular situations, relationships between and within groups of people and institutions, intuitions of the intentions and feelings of others, the complex meaning of words and whole utterances in context, or the contexts themselves – social as well as spatiotemporal. If you are a linguist it could be a generalization underlying some

grammatical phenomenon that fits neatly into place within a more general theoretical framework; if you are a historian or a philosopher it could be some new addition to the stock of general ideas applying to a given age or to mankind as a whole; if an experimental scientist it could be the very laws of nature manifest in the results of your experiment. In all these cases abstraction is of information, information which in the broadest sense is useful. At the simple physical level this includes recognizing and abstracting essential nourishment from the environment or choosing a suitable mate, and at the highest cultural level it may serve the goal of gaining understanding and knowledge in a lasting form that supports its further manipulation and deployment. Whatever the level, the purpose of abstraction is not just to align input with output but to construct and update representations of the world in conformity with accumulating experience, in order to be able to project it back onto our immediate surroundings, guiding our further perceptions and actions – or into purely internal representations of possible, probable, or simply desired futures.

Karl Popper saw problem-solving as the source of evolution, the primal problem being survival (Popper 1973: 242). Abstraction clearly does serve evolution. The more – and the more efficiently – a creature can abstract information from the environment the better its chances of individual survival. The great apes as individuals are capable of limited forms of problem-solving, and even relatively primitive organisms may learn by trial and error, trying novel solutions, new expectations. But by 'problem-solving' Popper is referring not just to individuals but to Nature's own solutions to evolutionary quandaries. His is a rather special perspective on evolution, namely as 'evolution from within' or 'active Darwinism' (the organism actively trying to discover new ways of life in new ecological niches). He sees all living organisms (even ones as simple as the amoeba) as organized around a central control system (not necessarily a brain) genetically determining its 'goal'. This spearheads evolutionary development by being subject to mutations that can drastically change the organism's behaviour without necessarily changing the phenotype as such (Popper 1973: 281).[3] Some of the animal kingdom's solutions can be quite abstract – think of the social organisation of ants, or the patterns of communication among vervet monkeys warning their congeners of the proximity of specific kinds of enemy. Ever since Darwin replaced a 'designer' God with Nature itself as the self-organizing source of its own infinitely varied forms through the workings of long-term natural selection it has been clear that what a species abstracts from its environment is what helps its members survive. Thus, as a single example, the great apes and most monkeys developed red-sensitive cones in their

3. Note that Popper is – like Whitehead – a self-proclaimed realist as regards all three of his 'worlds' (1978: 151).

retinae 'in order to' distinguish ripe fruit (which cats can't), whereas they never developed cones sensitive to ultraviolet light for which they had no need – unlike birds, who utilize it, principally for sexual selection, without making themselves conspicuous to predators who can't see it (Parker 2005:149). For Popper the rationalist evolutionary processes can have a rationale without any overall plan (spirit or vital force) – they can never be predicted, but are emergent, and emergence cannot be rationally explained (a view close to Chomsky's on the origin of language).

A great leap forward in the ability to form and manipulate abstractions is of course that associated with the appearance of language in the human species – whether this was a relatively sudden occurrence due to a chance mutation, as Chomsky would have it,[4] or a more gradual affair, as Bickerton (1990) argues with his notion of a preceding stage of 'proto-language' (essentially "modern language minus syntax", a stage still present in for example pidgins). Givón (2002: 11–12) characterizes the development of modern language as one from pre-grammatical to grammatical mode, developing Bickerton's idea in a functionalist vein. Typical of this hypothetical pre-grammatical mode is simple syntax, the absence of morphology, the pragmatic use of word order, many pauses and errors, slow delivery, high mental effort, and high context dependence, whereas the 'grammatical mode' is more fluent, displays complex syntax, faster, automatic delivery and lower dependence on context. One might add that the acquisition of a fully modern language results in a far greater capability for abstraction than before – it distances us from the outer world of perceptual objects but at the same time allows us to select from the environment just those elements and relations that are relevant to our on-going interests and intentions and manipulate them without having to engage with them directly. It allows the imagination to flourish through verbally guided simulations that can be held for an appreciable time in powerful working memory circuits. It is only through language that we can formulate the highest abstractions of which we as cultural and imaginative beings are capable, notions such as 'the future', 'the world', 'time', 'the Good', etc. These things could no doubt be vaguely sensed by pre-linguistic man, but hardly held in sharp focus and discussed in the manner afforded by the wielding of linguistic symbols.

A slightly different elaboration of Bickerton's idea is that of Jackendoff (2002: 231–264), which involves a number of discrete steps, including importantly that at which hierarchical phrase structure developed out of 'protolanguage', as can be seen in Figure 16. This should be seen in the light of his view of Universal Grammar as a 'tool kit' of layered sub-components that can emerge incrementally

4. Thus he casts doubt on the role of natural selection in the appearance of language: 'It would be a serious error to suppose that all properties, or the interesting properties of the structures that have evolved [in language], can be 'explained' in terms of natural selection' (Chomsky 1975: 59).

(op. cit.: 263–4). It is relevant to add, however, that also tool-making and action planning are hierarchically organized and do not necessarily presuppose language (Higuchi et al. 2009). Recall Larsson's suggestion of their role in the evolution of language in 6.1. The evolution of language undoubtedly involved the coming together of multiple strands.

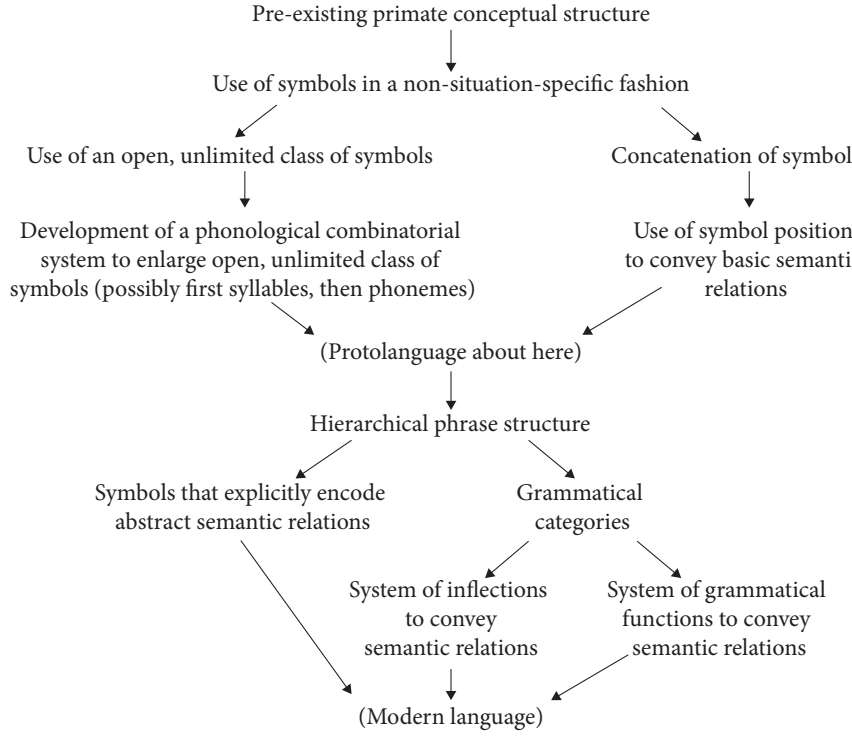

Figure 16. Incremental evolutionary steps (Jackendoff 2002)

The transition from "pre-existing primate conceptual structure" to "use of symbols in a non-situation-specific fashion" is also the starting point for Deacon's detailed account of the co-evolution of language and brain based on Peirce's theory of signs (Deacon 1997). Central to his thesis is his definition of mankind as the unique 'symbolic species'. The ability of the great apes to use symbols (as opposed to simple tools) is, he maintains, highly limited and not part of their social life in the wild. What they are naturally capable of is 'stimulus generalization' based on some rather abstract qualities of potential food items which they can distinguish. What the ability to label such categories in artificial situations with symbolic 'lexigrams' seems to foster is the ability to generalize to *new* items introduced – 'an incredible increase in learning efficiency' (Deacon op. cit.: 91–2). This is precisely what Pepperberg claims Alex, the African Grey parrot mentioned in 6.2, could do. In an

evolutionary context she emphasizes the importance of categorical classification for dealing with a changing environment – the ability to choose what to ignore as well as what to process. Both of these forms of abstraction are amply displayed by parrots and other intelligent non-human species (Pepperberg 1999: 324–6).

But we need to go back well beyond language and even symbol manipulation in our quest for the roots of abstraction. The dawning of consciousness, capable of abstracting new, more sharply focused forms of information from the environment beyond what stimulus-driven instincts could produce, must have emerged before full human language was achieved. It may arguably have emerged (in a rudimentary form) in other mammals too. We have seen Damasio's distinction between 'core' and 'extended' consciousness and how Bermúdez' Level 1 rationality corresponding to the former presupposes abstraction. Abstraction is hierarchically recursive through the course of evolution. Thus simple organisms may abstract information from their immediate environment guiding them to where good foraging lies. More advanced creatures can remember various good sources of food and generalize across such instances as representing a general sense of 'good food sources' – if one source is not available then another may be tried out. Primitive man must have been able to generalize still further to a notion of tribal 'territory' in which there are propitious foraging opportunities and which consequently needs to be defended. And advanced *homo sapiens*, possessing complex language, can generalize from experience of territory to the still more abstract notion of 'country' – not just a matter of one's immediate surroundings, and not just one's own. Finally, map-makers can produce faithful territorial maps, artefacts that last and can be transmitted from generation to generation. Something similar is seen ontogenetically in the way young children learn first about individual tokens of, say, cats, then about the category of all creatures displaying certain key cat traits (as constrained by adult usage), then finally about the hypernym category of 'animal' generalized across just those traits shared by all animate creatures. Once a child has learnt about numbers he or she is ready to learn how to manipulate them in basic arithmetic at school, then later how to substitute variables for numbers in algebra – and in adult life, if he or she becomes a mathematician, number theory will open vistas for still higher abstractions across abstractions.

To the question raised at the start of this section as to what all the forms of abstraction we have encountered so far have in common, one may refer to Popper's suggestion that problem-solving lies at the dynamic heart of evolution. Could abstractions themselves have emerged as problem-solving tools, hand in hand with evolution? To answer this we need to look at Popper's theory of the 'three worlds' (Popper 1978). 'World 1' is the world of material objects; 'World 2' that of subjective minds, of mental experience, thought and perception, and 'World

3' is that of the objective *products* of minds or living creatures. The latter is to be taken as the repository of human abstractions (like ideas of ethics, aesthetics and scientific theory) as well as more tangible cultural artefacts, all encoded and preserved in things of World 1 (Popper 1973: 118). World 3 should not be taken in a Platonic sense of complete detachment from the material world, precisely because its contents are embodied in World 1 objects (which is true of Whitehead's 'eternal objects' also), and this includes the formal systems of logic and linguistics which presuppose both human brains and (external) means of perpetuating them in written form. Although he has been accused of being a dualist, his subjective World 2 can also be understood as literally 'embodied' – in the human brain experienced 'from within' as it were (a common sense rather than a Platonic view). It is where '*qualia*' reside (cf. Whitehead's 'subjective form' defined in 10.3). It 'emerges as an evolutionary product of the world of organism' (Popper 1978: 166). What Popper and Whitehead have in common is the insistence on the emergence of novel properties – both link abstraction to emergence. This is 'abstraction by emergence' rather than 'abstraction by matching'. Abstractions of the more intangible World 3 'products' such as expressed ideas and theories (including Russell and Whitehead's 'practical abstractions') presuppose a 'prehending subject', and this is precisely what an individual World 2 mind is.

Primitive man came into a world dominated by abstractions – kinship relations, forms of social organisation, customs, beliefs, religion, language, etc., none of which were planned or intended according to Popper. The aim-structure of animals and men is not 'given' but develops with the help of feedback out of earlier aims, creating a whole new world of possibilities, one to a large extent autonomous and which we cannot further control or influence once emerged (Popper 1978: 117–8). Even mathematics, the most abstract product of all, is premised on a human creation – namely the sequence of natural numbers – but one creating its own autonomous domain in which man can discover further unsuspected relationships. This goes a long way towards accounting for Lakoff's objection to the widespread view of mathematics as belonging to a transcendent Platonic world mentioned in 10.1.

Dawkins' conception of the 'meme' (Dawkins 1976), like Popper's World 3 entities, requires a physical vehicle: it can only be transmitted through the senses – like the numerous ideas presented on the pages of the book that you have in your hands (many of which I have in turn abstracted from the written words of others). It is a unit of cultural transmission, premised on imitation and replication, in fact any cultural entity that an observer might consider a 'replicator', e.g. melodies, fashions and learned skills as well as pure 'ideas'. Memes generally replicate through humans, and because humans do not always copy memes perfectly they may change with time or combine with other memes to create new ones. Dawkins

likened the process by which memes survive and change through the evolution of culture to the natural selection of genes in biological evolution (op. cit.: 352).

Despite his emphasis (like Dawkins') on human cultural artefacts, Popper nevertheless characterized some pre-human problem-solving solutions as 'abstract', as we have seen. So what exactly is the relation between problem-solving and abstraction? It lies quite simply, I would suggest, in the fact that problem-solving of any kind requires the abstraction of information, of meaningful patterns, from the problematical situation requiring resolution before planning – or just allowing – the behavioural response that will overcome it can take place. This can be viewed from the point of view of the individual or of the species.

11.4 Negative entropy as the ultimate ground of abstraction

It might seem that we have reached the end of the journey with evolutionary problem-solving, but we can in fact delve still further into the roots of abstraction, beyond the purely human world. Physicist Murray Gell-Mann describes the advance of evolution as a combination of fundamental laws (including quantum uncertainty), macrocosmic chaos or chance ('frozen accidents'), and specific conditions. These function so as to produce 'effective complexity' in complex adaptive systems through a fruitful balance of randomness and regularity. Such systems find perceived regularities in the stream of incoming information and ignore the rest as random (Gell-Mann 1995: 368). At the fundamental physical level, quantum 'decoherence' of multiple probabilistic pathways results in the 'selection' of just one, as determined by the totality of the surrounding environment ('summing' and thereafter 'ignoring' its effects). To describe events measurable on the level of classical physics he employs the notion of 'coarse graining', a kind of abstraction across quantum level events. He further notes that systems displaying the highest level of complexity have a tendency to become still more complex, a development favoured by selective pressures (op. cit.: 369). This might seem to run counter to the Second Law of Thermodynamics which states that the degree of order in a system either remains the same or decreases but cannot increase ('entropy' or disorder can only increase). This can be restated in terms of (useful) information. But the law applies only to closed systems – the total amount of information in the surrounding universe may be maintained, just becoming dispersed and distributed more widely. Entropy, like 'effective complexity', depends on the coarse graining, the detail at which a system is described – in perfect detail it would not increase, just remain constant (op. cit.: 225–6). From this perspective evolution, adaptation and learning by complex adaptive systems are all, surprisingly, part of the winding down of the universe (op. cit.: 372).

To get a better handle on this we need to consider the converse of entropy, *negative* entropy, the development towards greater effective complexity, greater informational content, as we observe it operating in the development of complex life forms out of simpler ones. Schrödinger (1944) introduced this term in defining the difference between life and inert matter – he argued that living systems, not being closed, can defy the Second Law of Thermodynamics by organizing themselves so as to "feed on negative entropy", redistributing entropy (disorder) back to the environment, for example in the form of heat, and thus decreasing its (the environment's) own entropy. This could be conceptualized as a kind of Bergsonian non-linear, creative '*élan vital*' struggling against dull matter, only there is no struggle as such, no 'bifurcation of nature' if entropy ultimately remains constant. Perhaps a more useful way of envisaging the situation is to consider the interplay between the second and the first laws of thermodynamics, the latter stating that energy cannot be created or destroyed. As order is gradually redistributed in our present universe, nudged along by thermodynamic entropy and the other factors listed by Gell-Mann above, energy too becomes redistributed (though never lost). One logical way to maximize what energy is available in a sub-system within the totality would be to produce more and more informational order using less and less energy. Clark (2013: 6–7) discusses this in terms of 'the minimization of free energy' by the predictive brain. Singer sees the inner connectivity of the brain (the 'connectome') as constituting a compromise between randomness and regularity where complexity and 'dimensionality' are high. The resting state dynamics of cortical networks (i.e. with no additional energy requirements on processing input) suggests that they operate at close to a self-organized 'critical' state with maximal memory capacity (Singer 2013: 622).

The universe is not evolving towards some ultimate teleological state of being, but towards greater and greater anti-entropic discrimination, i.e. intelligence, in those of its sub-systems that permit it. This could be achieved by complex dynamic systems like living organisms that are capable of abstracting and utilizing finer and finer discriminations from the environment, integrating them efficiently into generalized schemas (existing or new). By this means such organisms would need to use less and less overall physical energy in analysing specific input from the environment as they increase at the same time in their power of survival-serving abstraction.

Following its massively energetic origins the expanding universe underwent a redistribution of energy into smaller and smaller imploding regions of protomatter (from galaxy to star to planet to individual organisms on their surface) needing successively less and less energy to maintain. With the appearance of life, requiring little – but constant – energy from the nearest star to thrive, the 'useful information' available to organisms, by being stored genetically, would have begun

to grow exponentially. The direction of entropy of the thermodynamically fated universe could be reversed in these physically diminutive anti-entropic 'eddies', at least long enough for advanced creatures with brains and a sophisticated system of memory storage to emerge and survive.

This can be seen as what drives the development of higher and higher forms of Whiteheadian 'prehension'. For the more abstract – the more generalized – the schemas abstracted, the more efficient the organism becomes in its actions and reactions to its environment. It allows more and more useful information to be packed into compact physical organs requiring relatively little energy to maintain – i.e. brains. The hierarchical organization of different levels of abstraction is more efficient still. Consider the pathways of control between recognizing a general situation type and responding to it with rapid, rational behaviour. Once such a schema has been internalized and its links to particular advantageous actions established, the organism need expend no excessive energy on planning a detailed course of action – this will now occur automatically, all the way down through the chain of control already established, from potentially conscious awareness of one's goal through routinized levels of sub-routines for fulfilling such goals, to the rapid unconscious activation of individual muscles. The physical response will itself be fine-tuned for maximally efficient use of what energy is available to the organism. This is true not only in the rationally thinking human brain, but in the instinctive responses of any living organism, right down to the lowliest amoeba. The drive at all levels is towards less and less energy to greater and greater effect.

Recall the suggestions in 8.1 as to how neural assemblies 'abstract' information through the sense channels – they actively compare incoming information in the process of forming and adjusting schemas by the "neural averaging of similar inputs", and these may in turn take their place in complex distributed networks within an overall multi-modal encyclopaedic memory system whose contents can be rapidly accessed by multi-level content-addressed searches. The activation of any part of a network potentially activates the rest without any further conscious effort, facilitating the manipulation of the concepts involved without having to reason step by step in order to get at the same information (obviously a more expensive process in terms of energy). Fine-grained knowledge is 'good for you' – and for the survival of your species. And to obtain more and more knowledge the brain no longer needs to get physically bigger once the crucial size and organization is reached that allows rapid 'symbolic reference' between all regions of the cortex and the tool *par excellence* to accomplish this is acquired, namely language. For language – words and propositions expressed by their combination – can 'stand for' complex realities in a much simpler, manageable form, shortcutting the need for detailed sensory simulations of objects and actions at every turn. The large-scale neural activity revealed by fMRI scans is probably misleading – the

computations going on at the neural level are probably on a much smaller spatial scale, deep within regions interconnected by white matter fibres rendered superfast and efficient by myelination. Such activity requires less energy than muscular activity of the body and certainly lies far below the reach of consciousness. This is possible because the brain functions both hierarchically and distributionally and is capable of projecting its fine discriminations to create further combinations and still finer internal tunings of already stored schemas. And the potentiality of the 'optimally' developed biological brain can be pushed even further by advances in computer technology, externally extending our calculating power enormously. The drive towards ever-increasing efficient miniaturization is not just the story of the recent history of computers but surely characterizes the development of human information packaging throughout the evolution of our species.

This phylogenetic development of the human brain to abstract and store information more and more efficiently is also reflected ontogenetically in the stepwise acquisition of language by children as we saw in Chapter 9. As Taylor puts it, children first learn to abstract generalized schemas – the basis of adult-like word meanings – after accumulating prototypes based on individual sensory tokens (Taylor 1989: 242). Thereafter they build on simpler word already acquired – along with their associated schemas and contexts – to accumulate more complex and abstract words. According to Pulvermüller & Schumann (1994: 698–9, 714) words learnt later (typically more abstract ones) tend to be concentrated (or anchored) in as yet not myelinated convergence zones, further away from the perisylvian 'core' language areas. They further see second language learning as exploiting these late-myelinating areas (op. cit.: 712–3). There is a 'vertical' as well as a horizontal aspect to this distribution: later acquisition exploits the more widely distributed 'A-system' of neural connections, which also myelinates later than local 'B-system connections', which may have more to do with phonology and syntactic behaviour than semantics (which is handled by the long-distance A-system). Be this as it may, adult semantic memory for individual word meanings is clearly more fine-grained than that of infants, reaching into more widespread regions of the cortex, including frontal regions where propositional and abstract knowledge appears to be concentrated. It is also more densely and efficiently packed – a decrease rather than an increase of entropy.

11.5 Conclusion: 'Abstract$_0$'

All that remains to do in order to conclude our search for the roots of abstraction is to reconsider the relationship between 'abstract$_1$' and 'abstract$_2$' (respectively 'disembodied from specific sensory features' and 'simplified or generalized across

instances'). Can we now characterize a more generalized kind of abstraction deployed by the human brain? We may in fact have seen the same basic phenomenon operating in varying guises and at multiple levels throughout this book – that is, an 'abstraction$_0$' that covers the senses of both 'abstract$_1$' and 'abstract$_2$'. Abstraction of the 'abstract$_1$' kind is a matter of degree on a continuum of distancing from sensory input, reflected as we have seen in 8.1 in the distinction between sensory-specific vs. amodal/ multimodal regions of the brain. 'Embodiment', I have argued, is also a matter of degree – the mental schemas the mind operates with can be at a second or even third remove from the senses. Abstraction of the schematic 'abstract$_2$' kind shares this characteristic with abstraction of the first kind then, only it relates more directly to the dimension of complexity by resulting in 'simplifications' of one degree or another. Recall Langacker's characterization of linguistic constructions as 'schemas' presented in 3.1: these are hierarchically organized at different levels of 'granularity' (i.e. of lexical specificity) and at the 'lowest' level may contain an open-ended number of lexically filled exemplars. 'Abstract$_1$' words may have simple or complex semantics but are (like all content words) in essence just units of symbolic 'shorthand' applicable across arrays of similar/analogous instances – their complexity resides in the complexity of their contexts of use. An 'abstract$_0$' item that is both 'abstract$_1$' and 'abstract$_2$' would be both simplified/ schematic and only indirectly derived from sensory input, however vividly experienced in terms of affect. It would lie on a different ontological plane from that of the products of direct perception, and could be instantiated, for example, in the dynamic behaviour of a neural assembly or network with no point-by-point relation to the input activating it. This might cover most of the more specific kinds of abstract meanings we have considered already. In fact it could also cover the problematic case of the 'self' met in 11.2 in so far as this ghostly entity is both 'only indirectly derived from sensory input' and 'schematic' or simplified in the sense of generalized across the whole gamut of our conscious experience of being embodied as an individual in the sensory world. So what is the nature of the 'abstraction$_0$' process that would produce such entities?

My contention that Whitehead's notion of 'transmutation' fits the bill rather well will come as no surprise. This, you will recall from 10.3, is defined as a process whereby complex arrays of sensory input, accessed through 'physical feelings', are integrated and simplified in order to be experienced as unitary objects with qualities accessed through corresponding 'conceptual feelings' (matching the input array to concepts). It is our usual way of consciously experiencing the world as consisting of objects and the relations among them. In my cognitive interpretation of Whitehead what these transmuted feelings prehend in general is 'forms of definiteness' manifest in the sensory input that correspond to/ resonate with neurally instantiated concepts (including the 'image schemas' of cognitive linguists). The

process of transmutation makes use of these to see the world both 'as it appears to be' in the here-and-now and in internalized mental models. In fact it is instrumental in the creation of the conceptual categories which the brain operates with in the first place, as infants learn about the world around them. They learn through pattern recognition, and this ability to abstract pattern is what the brain is there to do. The key property of transmutation over and above any more specific kind of abstraction is that it covers both the recognition and projection of concepts and schemas onto the world we see around us and the recognition and projection of the contexts in which those concepts and schemas occur (whether expressed in words or not), both in the external world and in memory. It abstracts from the environment enduring 'nexūs' or 'societies' defined in terms of clusters of shared properties or features. Furthermore, it is instrumental in originally forming and subsequently updating networks or assemblies in the brain capable of recognizing and simulating these properties – approximate similarity to which *is* the concept. In neural terms, 'context' can be seen as constraints on the applicability of concepts, in other words as the suppression of those of the multiple associations of a concept which do not mesh with a given situation. Transmutation must in sum be an innate capability of the neonate, perhaps the only one needed to interpret what the senses provide it with.

The above is after all just a common-sense realist answer to the age-old dispute between nominalists and realists – abstractions are 'real' in so far as they are embodied in matter. Concepts exist in our brains as potentials instantiated in neural networks for matching input with stored categories and relationships – and for forming mental models from their combinations. As potentials they must also be 'ingredient' in the workings of the material universe as a whole, since that is the source of the brain's reflection of them, the starting point for their internal transmutation. The simulations produced by words and mental models stored in the brain are just approximate analogues of external reality and the neural networks responsible for producing them do not necessarily bear any direct relationship with the sensory input from which they are abstracted (they are perhaps derived through 'hidden' layers). The human brain can moreover operate with words that 'stand for' potential simulations in suitable contexts without having to initiate expensive mental reconstructions every time they are used. We are perfectly used to consciously existing for extended stretches of time on a purely linguistic level of 'as if' reality – for example when we read a novel or a scientific report. We think nothing of it as long as that reality is internally coherent. Our awareness is focussed on the 'vehicles' of meaning (words on the page) while their meaning is 'elsewhere' (as Jackendoff would have it).

However, we must bear in mind that transmutation for Whitehead is essentially a matter of perception. Although it can be extended to the recognition of

events and states and not just physical objects (as I have argued in 10.3), is it also relevant to non-perceptual objects? These are supposed to be prehended directly by conceptual feelings of 'eternal objects', not by transmutation. Jackendoff too in his 'Intermediate Level' theory of consciousness is mostly concerned with perceptual entities. But he has also turned his attention to 'abstract objects' (Jackendoff 2002: 322–4). He suggests that conceptual structure may contain features that do not pertain to perception but which connect instead to the inferential system, and that these 'inferential' features could provide 'bridges' to other concepts. They could include features of an object's value, its uses to the observer, or the higher-level taxonomic category to which it belongs. Thus recognizing something as an 'artefact' does not simply involve recognizing certain perceptual features, rather it bridges to all sorts of inferences that can in turn potentially engage perceptual features. There are also concepts which have no perceptual features, only inferential ones – these are the meanings of extreme 'abstract$_1$' words like 'truth', 'belief' or 'obligation'. Such concepts are in a sense 'out there in the world' – Jackendoff shares Whitehead's realist stance, though he does not attempt to answer the question as to how exactly words connect with meanings 'out there'. This requires a closer investigation of 'context' and social meaning than he has provided us with.[5]

My own solution to this problem of the applicability of the process of 'transmutation' to 'abstract$_1$' meanings is much the same as Jackendoff's as regards access to 'abstract objects' – these can indeed be equated with (a sub-class of) Whiteheadian 'eternal objects' or 'forms of definiteness'. What transmutation does is to summate the various sensory 'affordances' presented by an object to the observer and to reduce them to a single nexus experienced as a unitary 'entity'. If such affordances adhering to objects and situations 'out there' can also be purely 'inferential' in Jackendoff's sense then transmutation can also be seen to be involved in understanding the nexus in which they are contained, however 'abstract' that might be. It might include for instance such categories as 'all one's relatives' or 'the truth of allegations about X in the media'. All that seems necessary is to accept that the recognition (and projection) of 'abstract objects' in 'the world out there' may involve a form of transmutation of objects defined not by perceptual features – or not by them alone – but by entities that can in principle only be grasped via inferential bridges. I shall return to these presently.

Allowing for this caveat then, all the senses of 'abstraction' that we have met underway would appear to be covered by this extended interpretation of 'transmutation'. These include (in approximate order of appearance) the abstraction of

5. He does however mention social concepts of kinship and dominance among monkeys as well as humans (Jackendoff 2002: 324).

linguistic theories from linguistic surface structures; abstract lexical items;[6] the (semi)abstractness of logograms and sign language hand shapes; the storage in memory of sensorimotor image schemas – including culturally relevant types of context, simplified episodic memory traces, and internally coherent mental models; visual abstraction from complex scenes; abstract emotion and its recognition in 'theory of mind'; the products of abstract art and other 'World 3' cultural artefacts; philosophical 'ideas' and 'universals'; the subject matter of mathematics; the abstraction of information relevant to survival by organisms in the course of evolution; and finally, the objects of consciousness.

The term covers both 'abstraction by matching' and 'abstraction by emergence', also the complex hierarchical levels of abstraction between say a book (abstracted and synthesized from an author's experiences) and a television version of it as observed by a viewer. Even the experience of music (i.e. of its World 3 products) can be understood in terms of the abstraction$_0$ of the relevant time-extended musical patterns, enriched by 'reversions' (associations), through the transmutation of sensory input. Important to note here is that the 'content' of music is from this perspective not emotion, as is commonly held. The emotion evoked lies in the Whiteheadian 'subjective form' (valuations of adversion vs. aversion, etc.) that is an integral part of any 'conceptual feeling'. There is plentiful evidence that complex brain chemistry involving specific neurotransmitters is initiated in response to music (cf. Panksepp & Bernatzky 2002). It is this that is experienced as the emotional 'content' – or rather accompaniment – of music. As with any emotion, this will evoke further bodily responses, as described in Chapter 7.

All these forms of abstraction fit into the picture then, except perhaps Gell-Mann's 'coarse graining'. This 'abstraction' of classical physical reality from a finer grained quantum level does not fit the definition any more than natural constants like the speed of light do, since it is not derived from sensory input – unless it should be taken as something theoreticians abstract, not Nature, in which case it certainly belongs along with other World 3 'abstract artefacts'![7]

But what of one further kind of abstraction, that typifying a broader social scale than the individual, like opinion polls, political elections, referendums, and statistics across human demographics and economics? The narrowing down of candidates in for example a presidential or party leadership race follows its own, man-made conventions, but essentially mirrors the 'Cohort' model of word rec-

6. These include items which are only abstract in the sense of 'under-specified', like the pronoun 'he' and words like 'climb' that have somewhat different meanings dependent on context.

7. Of course the standard 'Copenhagen interpretation' of quantum physics does crucially involve perception: probabilistic quantum wave functions are supposed to 'collapse' into the world of classical physics only when perceived by an observer via his or her experimental apparatus.

ognition from sensorimotor input (as described in 2.2). Only the input is not sensorimotor. Or is it? It depends on your perspective – there are several levels between the choices made by individual voters and the results declared by the vote-counters, and we have seen that abstraction can be hierarchically recursive and layered. Transmutation – i.e. 'abstraction$_0$' as I have defined it – presupposes a prehending subject, but perhaps such a 'demographic' abstraction as an election just reflects a longer and more complex succession of levels between the input and the result. Consider all the relevant sensory information the individual voter has absorbed from the media concerning the candidates – plus whatever bias past experience brings to bear on his or her choice on election day. The actual result at the higher social level, the generalized summation of individual votes across the community of voters, is itself a World 3 artefact, constrained and interpreted by cultural convention – it can be prehended by individual subjects via physical and conceptual prehensions of the usual kind.

The line needs to be drawn, however, at the *unintended* results of group activity such as the market forces described by Adam Smith (1982 [1776]) as guided by an 'invisible hand'. This presupposes a community of individuals acting according to their own egoistic goals and desires which have little to do with the overall outcome. Although those 'egoistical goals and desires' are premised on sensorimotor input to the individual's accruing experience on the basis of which he or she acts in a manner beneficial to themselves, they are, unlike in the case of voting, precisely *not* input to the overall result at the higher level, e.g. the equitable division of wealth. That is solely the outcome of 'natural' homeostatic processes at that higher level of nexus, entirely lacking in conscious intention at any level. It has more in common with the ant-hill, where the individual ant has no inkling of the highly organized society to which it is contributing. Of course all such things can nevertheless become World 3 theoretical 'artefacts' that display their own degree of abstraction over physical and social reality. Phenomena such as these are examples of abstractions that do not involve transmutation and therefore lack affect as their subjective *qualia*.

Another such phenomenon falling under the 'invisible hand' rubric is the purely linguistic one of 'grammaticalization'. Although this displays certain features typical of abstraction, in particular semantic 'bleaching', this is again independent of the intentions – but not necessarily the actions – of the individual perceiver/ speaker. The standard example given of this is the historical loss of content of the verb 'go' in English, starting from its full content sense of movement on foot, through a partial loss of the specificity of the movement in constructions like "He is going to visit Bill" (in the sense 'on his way to visit Bill'), and finally, in its completely 'bleached' form as a grammaticalized future tense marker, as in "He is going to like Bill" (Hopper & Traugott 1993: 88). Parallel but independent

pathways into the formation of future tense markers from verbs of motion are attested across a wide array of languages – it is pathways such as these that form the stuff of grammaticalization theory. Various proposals have been raised as to what does motivate such changes if not the individual speaker's intentions. We need not dwell on that here.

I have mentioned already a final kind of abstraction that could not be attained by transmutation understood in strictly Whiteheadian terms, namely the 'highest' abstractions of all, like 'the universe', 'time', 'the Good', etc. I have claimed previously that these can only be formulated by beings with language. That is not the same as saying they can be directly experienced through the intermediary of the words representing them. Transmutation requires a specific sensory array as input or at least corresponding inferential features. Prehension of the actual content of such a verbally clothed abstraction could only be by intuition or by inferential 'bridging' to concepts that *can* be perceptually entertained. (As Jackendoff would say of linguistic content in general: the grounds of reality and truth can be intuited but not directly experienced in a fully conscious fashion – it is too rich to be held in short term memory.) This kind of 'intuition' could occur with a variety of 'subjective forms' depending on the state and the 'autobiographical' background of the prehender – it could be semi-consciously or subliminally experienced and could be affect-laden or neutral. Paradoxically, the broadest possible abstractions could be endued with the highest level of affect for the prehender, but this does not render them 'embodied' in any meaningful sense. They lack *any* kind of grounding in the sensory world apart from affect and/or the results of secondary inferencing or association with symbolic exemplars. All I can say is that if this is abstraction$_0$ it is so at its most tenuous.

But enough of what may not count as abstraction$_0$, I need to sum up finally what lies at the heart of what does. I have stated that abstraction as a process requires a brain to do it and that evolution has through the millennia precisely pushed for the development of an organ specialized in this general function. Abstraction$_0$ is nothing less than the 'seemingly seamless' process of experiencing the world around you in terms of objects and relations between them, as defined by the concepts you have formed throughout your life by its application. The further projection of the patterns abstracted in memory on the immediate environment allows forward prediction (either conscious or unconscious, instantaneous or longer-term) which, although prone to error, ensures flexibility of response. Despite the obviously different levels of complexity and sophistication involved, there is no difference in essence between the abstraction of meaning from the language we hear around us and the general abstraction of information from the sensory environment with which we are structurally coupled – from the perspective of infants language presents just another layer of Gibsonian affordances that

the environment has to offer. At a deeper causal level beyond conscious perception and intentionality it is legitimate to see the ultimate roots of abstraction in negative entropy, whereby Nature constantly unfolds itself in creative novelty and is not just the passive 'victim' of the Second Law of Thermodynamics. The human brain is – as far as we know – the highest, most complex result of Nature's self-creating. Its continuing plasticity allows us to keep continually one step ahead of the game through flexible creativity, abstracting from the environment not just what it affords but what it *might* afford with a little rational intervention. Such intervention, like all the other operations of which the brain is capable, presupposes the fundamental one of abstraction. Here surely is the ultimate *raison-d'être* of the brain: as a dedicated abstraction engine capable of generating novel responses to novel situations rather than as a deterministic general purpose computer. It is capable of abstracting everything under the sun that might serve its purpose – including a hypothetical grin from a hypothetical Cheshire Cat.

References

Aikhenvald, A. (2004). *Evidentiality*. Oxford: Oxford University Press.

Allan, S. (Forthcoming). Polysynthesis in the acquisition of Inuktitut. In M. Fortescue, M. Mithun, & N. Evans (Eds.), *The Oxford Handbook of Polysynthesis*. Oxford: Oxford University Press.

Arnal, L. H. & Giraud, A.-L. (2012). Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences*, Vol. 16 (7), 390–398. doi: 10.1016/j.tics.2012.05.003

Austin, J. L. (1962). *Sense and Sensibilia*. Oxford: Clarendon Press.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation*, Vol. 8 (pp. 47–90). London: Academic Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioural and Brain Sciences*, 22, 577–660.

Barsalou, L. W., Yeh, W., Luka, B., Olseth, K., Mix, K., & Wu, L.-L. (1993). Concepts and meaning. In K. Beals, G. Cooke, D. Kathman, K. E. McCullough, S. Kita, & D. Testen (Eds.), *Chicago Linguistics Society 9: Papers from the Parasessions on Conceptual Representation*, vol. 2 (pp. 23–61). Chicago Linguistics Society.

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. A. Zwaan (Eds.), *Grounding Cognition* (pp. 129–163). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511499968.007

Bates, E. & MacWhinney, B. (1982). Functionalist approaches to grammar. In E. Wanner & L. R. Gleitmasn (Eds.), *Language Acquisition: the State of the Art* (pp. 173–218). Cambridge: Cambridge University Press.

Benedict, R. (1947). *The Chrysanthemum and the Sword*. London: Secker & Warburg.

Bermúdez, J. L. (2007). *Thinking Without Words*. Oxford: Oxford University Press.

Bickerton, D. (1990). *Language and Species*. Chicago: University of Chicago Press.

Binder, J. R. & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, Nov. 2011, vol. 15, no. 11, 527–536. doi: 10.1016/j.tics.2011.10.001

Bloom, L.M. (1970). *Language Development: Form and Function in Emerging Grammars*. Cambridge, MA: MIT Press.

Bohm, D. (1980). *Wholeness and the Implicate Order*. London/New York: Routledge [Ark paperback edition 1984].

Bohnemeyer, J., Enfield, N. J., Essegbey, J., & Kita, S. (2011). The macro-event property. In J. Bohnemeyer & E. Pederson (Eds.), *Event Representation in Language and Cognition* (pp. 43–67). Cambridge: Cambridge University Press.

Bourtchouladze, R. (2002). *Memories are Made of This*. New York: Columbia University Press.

Bowerman, M. (1974). Learning the structure of causative verbs: A study in the relationship of cognitive, semantic and syntactic development. *Papers & Reports on Child Language Development*. Stanford University, 8, 142–178.

Bowerman, M. (1996). The origins of children's spatial semantic categories, cognitive versus linguistic determinants. In J. J. Gumperz & S. C. Levinson (Eds.), *Rethinking Linguistic Relativity* (pp. 145–176). Cambridge: Cambridge University Press.

Briggs, J. (1970). *Never in Anger. Portrait of an Eskimo Family*. Cambridge, Mass.: Harvard University Press.

Brust, J. C. M. (2003). Music and the neurologist: a historical perspective. In I. Peretz & R. Zatorre (Eds.), *The Cognitive Neuroscience of Music* (pp. 181–191). Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198525202.003.0012

Buck, C. D. (1988) [1st edition 1949]. *A Dictionary of Selected Synonyms in the Principal Indo-European Languages*. Chicago/ London: University of Chicago Press.

Bugge, A. K., Lynge, K., Fuglesang-Damgaard, Ad., & Nielsen, F. (1980). *Dansk-Grønlandsk ordbog*. Ministeriet for Grønland: Kalâtdlit-nunânaKiterisitsissarfik.

Bybee, J. L. (2010). *Language, Usage and Cognition*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511750526

Carey, S. (1978). The child as word learner. In M. Halle, J. Bresnan, & A. Miller (Eds.), Linguistic Theory and Psychological Reality (pp. 264–293). Cambridge, Mass.: MIT Press.

Carroll, L. (1865). *Alice's Adventures in Wonderland*. London: Macmillan & Co.

Chafe, W. (1994). *Discourse, Consciousness, and Time*. Chicago: University of Chicago Press.

Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.

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

Chomsky, N. (1975). *Reflections on Language*. New York: Pantheon.

Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, N. (1995). *The Minimalist Program*. Cambridge, M.A.: MIT Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Brain and Behavioural Sciences* 36 (3), 1–73. doi: 10.1017/S0140525X12000477.

Crick, F. (1994). *The Astonishing Hypothesis*. London & New York: Simon & Schuster Ltd [Touchstone Books edition 1995].

Croft, W. (2002). *Radical Construction Grammar: Syntactic theory in Typological Perspective*. Oxford: Oxford University Press.

Croft, W & Cruse, D. A. (2004). *Cognitive Linguistics*. Cambridge: Cambridge University Press.

Dahl, Ö. (2004). *The Growth and Maintenance of Linguistic Complexity*. Amsterdam: Benjamins. doi: 10.1075/slcs.71

Damasio, A. (2000). *The Feeling of What Happens*. London: Vintage.

Damasio, A. & Damasio, H. (1992). Brain and language. *Scientific American*, September 1992, 63–71.

Dąmbrowska, E. & Lieven, E. (2005). Towards a lexically specific grammar of children's question constructions. *Cognitive Linguistics*, 16, 437–74.

D'Andrade, R. (1987). A folk model of the mind. In D. Holland, & N. Quine (Eds.), *Cultural Models in Language and Thought* (pp. 112–148). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511607660.006

Davidoff, J. (1997).The neuropsychology of color. In C. L. Hardin & L. Maffi (Eds.), *Color Categories in Thought and Language* (pp. 118–134). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511519819.005

Davidsen-Nielsen, N. (1996). Discourse particles in Danish. In E. Engberg-Pedersen, M. Fortescue, P. Harder, L. Heltoft, & L. Falster Jakobsen (Eds.), *Content, Expression and Structure. Studies in Danish Functional Grammar* (pp. 283–314). Amsterdam/Philadelphia: John Benjamins. doi: 10.1075/slcs.29.11dav

Davidson, D. (1978). What metaphors mean. In S. Sachs (Ed.), *On Metaphor* (pp. 29–44). Chicago: University of Chicago Press.

Davidson, D. (1980). *Essays on Actions and Events*. Oxford: Oxford University Press.

Davidson, M. (2002). *Studies in Southern Wakashan (Nootkan) Grammar*. PhD dissertation in Linguistics, State University of New York, Buffalo.

Dawkins, R. (1976). *The Selfish Gene*. Oxford: Oxford University Press.

Deacon, T, (1997). *The Symbolic Species. The co-evolution of language and the human brain*. London: Penguin Books.

de Bot, K. (2004). The multilingual lexicon. *International Journal of Multilingualism*, 1(1), 17–32. doi: 10.1080/14790710408668176

de Groot, A. M. B. (1992). Determinants of word translation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 1001–1018.

Dennett, D. C. (1991). *Consciousness Explained*. New York: Little, Brown & Company.

de Saussure, F. (1916). *Cours de linguistique générale*. [edited by C. Bally & A. Sechehaye]. Paris/ Lausanne: Payot.

de Villiers, J. G. & de Villiers, P. A. (1985). The acquisition of English. In D. I. Slobin (Ed.), *The Crosslinguistic Study of Language Acquisition, vol. 1: The Data* (pp. 27–139). Hillsdale NJ/ London: Lawrence Erlbaum Associates.

van Dijk, T. A. & Kintsch, W. (1983). *Strategies in Discourse Comprehension*. New York: Academic Press.

Dik, S. (1989). *The Theory of Functional Grammar, Part 1*. Dordrecht: Foris.

Dik, S. (1997). *The Theory of Functional Grammar, Part 2* (ed. by K. Hengeveld). Berlin/ New York: Mouton de Gruyter.

Dong, Y., Gui, S., & MacWhinney, B. (2005). Shared and separate meanings in the bilingual mental lexicon. *Bilingualism: Language and Cognition*, 8 (3), 221–238.

Dorais, L.-J. (2010). *The Language of the Inuit*. Montreal & Kingston: McGill-Queen's University Press.

Ekman, P. (1984). Expressions and nature of emotion. In K. Scherer & P. Ekman (Eds.) *Approaches to Emotion* (pp. 319–43). Hillsdale, N.J.: Erlbaum.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996) [1998 paperback edition]. *Rethinking Innateness, a Connectionist Perspective on Development*. Cambridge MA/London: MIT Press (a Bradford |Book).

Engberg-Pedersen, E. (1996). Iconic motivations in conflict: Language-specific structure and influence from the medium. In E. Engberg-Pedersen, M. Fortescue, P. Harder, L. Heltoft, & L. Falster Jakobsen (Eds.), *Content, Expression and Structure. Studies in Danish Functional Grammar* (pp. 39–64). Amsterdam/Philadelphia: John Benjamins. doi: 10.1075/slcs.29.04eng

Evans, V. (2009). *How Words Mean*. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780199234660.001.0001

Fauconnier, G. & Turner, M. (2002). *The Way We Think*. New York: Basic Books.

Feldman, J. A. (2008). *From Molecule to Metaphor*. Cambridge MA/ London: MIT Press (a Bradford Book).

Felleman, D. J. & Van Essen, D.C. (1991). Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex* 1(1), 1–47. doi: 10.1093/cercor/1.1.1

Ferstl, E. C., Rinck, M., & Yves von Cramon, D. (2005). Emotional and temporal aspects of situation model processing during text comprehension: an event-related fMRI study. *Journal of Cognitive Neuroscience*, 17:5, 724–739. doi: 10.1162/0898929053747658

Fillmore, C. J. & Atkins, B. T. (1992). Towards a frame-based lexicon: the semantics of RISK and its neighbours. In A. Lehrer & E. F. Kittay (Eds.), *Frames, Fields and Contrasts: New Essays in Semantic and Lexical Organization* (pp. 75–102). Hillsdale NJ: Lawrence Erlbaum Associates.

Fillmore, C. J., & Baker, C. (2000). Frame Net web site: http://www.icsi berkeley.edu/framenet.

Fodor, J. A. (1975). *The Modularity of Mind, an Essay on Faculty Psychology*. Cambridge, Mass. & London: MIT Press (A Bradford Book).

Fodor, J. A. (1992). A theory of the child's theory of mind. *Cognition*, 44, 283–296. doi: 10.1016/0010-0277(92)90004-2

Fodor, J. D., Fodor, J. A., & Garrett, M. (1975). The unreality of semantic representations. *Linguistic Inquiry*, 4, 515–31.

Foley, W. A. (1986).*The Papuan Languages of New Guinea*. Cambridge: Cambridge University Press.

Fortescue, M. (1988). *Eskimo Orientation Systems*. Copenhagen: Meddelelser om Grønland, Man and Society 11.

Fortescue, M. (1991). *Inuktun: An Introduction to the Language of Qaanaaq, Thule*. Institut for Eskimologis Skriftrække 15.

Fortescue, M. (2001a). *Pattern and Process: A Whiteheadian Perspective on Language*. Amsterdam & Philadelphia: John Benjamins. doi: 10.1075/hcp.6

Fortescue, M. (2001b). Thoughts about thought. *Cognitive Linguistics*, 12,1, 15–45. doi: 10.1515/cogl.12.1.15

Fortescue, M. (2007) How to catch a mental model by the tale. *Acta Linguistica Hafniensia*, 39, 125–152. doi: 10.1080/03740463.2007.10414603

Fortescue, M. (2009). *A Neural Network Model of Lexical Organisation*. London & New York: Continuum.

Fortescue, M. (2010a). A neural network approach to compositionality and co-compositionality. *The Mental Lexicon*, 5, 2 (special issue 'Words and their Meaning', Eds. M. Horne & M. Roll), 180–204. doi: 10.1075/ml.5.2.03for

Fortescue, M. (2010b). Similitude: a conceptual category. *Acta Linguistica Hafniensia*, 42(2), 117–142. doi: 10.1080/03740463.2010.521442

Fortescue, M. (2014). A network model of lexical organisation in the bilingual brain. In G. Rundblad, A. Tytus, O. Knapton, & C. Tang (Eds.), *Selected Papers from the 4th UK Cognitive Linguistics Conference*. London: UK Cognitive Linguistics Association. http://www.uk-cla.org.uk/proceedings/volume2.

Fortescue, M. (2016a). Polysynthesis: a diachronic and typological perspective. *In Oxford Research Encyclopedia (Linguistics)*. Oxford: Oxford University Press. doi: 10.1093/acrefore/9780199384655.013.152.

Fortescue, M. (2016b). The semantic domain of emotion in Eskimo and neighbouring languages. In P. Juvonen & M. Koptjevskaja-Tamm (Eds.), *The Lexical Typology of Semantic Shifts* (pp. 285–333). Berlin: Mouton de Gruyter. doi: 10.1515/9783110377675-010

Fortescue, M., Kaplan, L., & Jacobson, S. (2010). *Comparative Eskimo Dictionary with Aleut Cognates* (2nd, expanded edition). Fairbanks: Alaska Native Language Center Press.

Fortescue, M. & Lennert Olsen, L. (1992). The acquisition of West Greenlandic. In D. I. Slobin (Ed.) *The Crosslinguistic Study of Language Acquisition, vol. 3 (*pp. 221–308). Hillsdale NJ/London: Lawrence Erlbaum Associates.

Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Science*, 6, 2, 78–84. doi: 10.1016/S1364-6613(00)01839-8

Friederici, A. D. & Gierhan, S. M. E. (2013). The language network. *Current Opinion in Neurobiology*, 23, 250–254. doi: 10.1016/j.conb.2012.10.002

Garnham, A. & Oakhill, J. (1992). Discourse processing and text representation from a "Mental Models" perspective. *Language and Cognitive Processes*, 7(3/4), 193–204. doi: 10.1080/01690969208409384

Gazzola, V., Azih-Zadeh, L., & Keysers, C. (2006). Empathy and the somatotopic auditory mirror system in humans. *Current Biology*, 16, 18: 1824–1829. doi: 10.1016/j.cub.2006.07.072

Gell-Mann, M. (1995). *The Quark and the Jaguar*. London: Abacus [first edition 1994, London: Little, Brown and Company].

Gentner, D. & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52, 45–56. doi: 10.1037/0003-066X.52.1.45

Gernsbacher, M. A. (1990). *Language Comprehension as Structure Building*. Hillsdale, N.J.: Erlbaum.

Gibson, J. J. (1972). A direct theory of visual perception. In J. R. Royce & W. W. Rozeboom (Eds.), *The Psychology of Knowing*. New York: Gordon & Breach.

Givón, T. (1990). Verb serialization in Tok Pisin and Kalam: a comparative study of temporal packaging. In J. Verhaar (Ed.), *Melanasian Pidgin and Tok Pisin* (pp. 19–56). Amsterdam: Benjamins. doi: 10.1075/slcs.20.03giv

Givón, T. (2002). The visual information-processing system as an evolutionary precursor of human language. In T. Givón & B. F. Malle (Eds.) *The Evolution of Language out of Prelanguage*. Amsterdam/Philadelphia: John Benjamins. doi: 10.1075/tsl.53.03giv

Gleick, J. (1987). *Chaos*. London: Abacus (Little, Brown and Company).

Glenberg, A. M. (1997). What memory is for. *Behavioural and Brain Science*, 20, 1–53.

Globus, G. (1987). *Dream Life, Wake Life*. New York: State University of New York Press.

Gödel, K. (1986), *Collected Works I. Publications 1929–1936* (Eds. S. Feferman et al.), Oxford: Oxford University Press.

Goldberg, A. E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago/ London: University of Chicago Press.

Grady, J. E. (2005). Primary metaphors as inputs to conceptual integration. *Journal of Pragmatics*, 37–10: 1595–614.

Gumperz, J. J., & Levinson, S. C. (1996). *Rethinking Linguistic Relativity*. Cambridge: Cambridge University Press.

Hagoort, P. (2005). On Broca, brain, and binding: a new framework. *Trends in Cognitive Sciences*, vol. 9 (9), 416–423. doi: 10.1016/j.tics.2005.07.004

Halliday, M. A. K. (1975). *Learning How to Mean: Explorations in the Development of Language*. London: Edward Arnold. doi: 10.1016/B978-0-12-443701-2.50025-1

Halliday, M. A. K. (1985). *An Introduction to Functional Grammar*. London: Arnold.

Halliday, M. A. K. (1993). Towards a language-based theory of learning. *Linguistics and Education*, 5, 93–116. doi: 10.1016/0898-5898(93)90026-7

Halliday, M. A. K. & Hasan, R. (1976). *Cohesion in English*. London/ New York: Longman.

Hansen, M. B. & Markman, E. M. (2004). Appearance questions can be misleading: a discourse-based account of the appearance-reality problem. *Cognitive Psychology*, 50 (2005): 233–263. doi: 10.1016/j.cogpsych.2004.09.001

Harkins, J., & Wierzbicka, A. (Eds.) (2001). Emotions in Crosslinguistic Perspective, *Cognitive Linguistics Research*, 17. Berlin/New York: Mouton de Gruyter. doi: 10.1515/9783110880168

Harley, T.A. (2008). *The Psychology of Language* (3rd edition). Hove /New York: Psychology Press.

Hebb, D. (1949). *The Organisation of Behaviour: A Neuropsychological Theory*. New York: Wiley Science Editions.

Hickok, G. & Poeppel, D. (2007). The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8(5), 393–402. doi: 10.1038/nrn2113

Higuchi, S., Chaminade, T., Imamizu, H., & Kawato, M. (2009). Shared neural correlates for language and tool use in Broca's area. *NeuroReport* 20(15): 1376–1381. doi: 10.1097/WNR.0b013e3283315570

Hopper, P. J. & Traugott, E. C. (1993). *Grammaticalization*. Cambridge: Cambridge University Press.

Itkonen, E. (2005). *Analogy as Structure and Process*. Amsterdam: John Benjamins. doi: 10.1075/hcp.14

Jackendoff, R. (1985). *Semantics and Cognition*. Cambridge, MA: The MIT Press.

Jackendoff, R. (1987). *Consciousness and the Computational Mind*. Cambridge MA: MIT Press.

Jackendoff, R. (2002). *Foundations of Language, Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198270126.001.0001

Johnson-Laird, P. N. (1983). *Mental Models*. Cambridge: Cambridge University Press.

Johnson-Laird, P. N., & Oatley, K. (1992). Basic emotions, rationality and folk theory. *Cognition and Emotion*, 6: 201–223. doi: 10.1080/02699939208411069

Kari, J. (1990). *Ahtna Athabaskan Dictionary*. Fairbanks: Alaska Native Language Center Press, University of Alaska.

Kiehl, K. A., Liddle, P. F., Smith, A. M., Mendrek, A., Forster, B. B., & Hare, R. D. (1999). Neural pathways involved in the processing of concrete and abstract words. *Human Brain Mapping*, 7, 225–233. doi: 10.1002/(SICI)1097-0193(1999)7:4<225::AID-HBM1>3.0.CO;2-P

Kintsch, W. (1988). The use of knowledge in discourse processing: A construction-integration model. *Psychological Review*, 95, 163–182. doi: 10.1037/0033-295X.95.2.163

Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews, Neuroscience*, September 2001, vol. 2, 635–642. doi: 10.1038/35090055

Kounios, J. & Holcomb, P. J. (1994). Concreteness effects in semantic processing: ERP evidence supporting dual-coding theory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20, 804–823.

Labov, W. (1973). *Language in the Inner City*. Philadelphia: University of Pennsylvania Press.

Lakoff, G. (1972). Hedges: a study in meaning criteria and the logic of fuzzy concepts. *Proceedings of the Chicago Linguistics Society*, 8: 183–228.

Lakoff, G. (1987). *Women, Fire and Dangerous Things*. Chicago: Chicago University Press. doi: 10.7208/chicago/9780226471013.001.0001

Lakoff, G, & Johnson, M. (1980). *Metaphors We Live by*. Chicago: University of Chicago Press.

Langacker, R. (1987). *Foundations of Cognitive Grammar: Theoretical Prerequisites*, Vol. 1. Stanford: Stanford University Press.,

Langacker, R. (1998). Conceptualization, symbolization and grammar. In M. Tomasello (Ed.), *The New Psychology of Language* (pp. 1–39). Mahwah, New Jersey: Lawrence Erlbaum.

Langacker, R. (2000). *Grammar and Conceptualization*. Berlin & New York: Mouton de Gruyter.

Langacker, R. (2005). Dynamicity, fictivity, and scanning. The imaginative basis of logic and linguistic meaning. In D. Pecher & R. A. Zwaan (Eds.), *Grounding Cognition* (pp. 164–197). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511499968.008

Larsson, M. (2015). Tool-use-associated sound in the evolution of language. *Journal of Animal Cognition*, 18,5: 993–1005. doi: 10.1007/s10071-015-0885-x.

LeDoux, J. (1998). *The Emotional Brain*. London: Weidenfeld & Nicolson.

Lerdahl, F. & Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge MA: MIT Press.

Levelt, W. J. M. (1993). *Speaking*. Cambridge MA/ London: The MIT Press (A Bradford Book). [first edition 1989]

Levinson, S. (2003). *Space in Language and Cognition*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511613609

Lovick, O. (2012). Walking like a porcupine, talking like a raven: Figurative language in Upper Tanana Athabascan. In A. Idström & E. Piirainen (Eds.), *Endangered Metaphor* (pp.103–122). Amsterdam: John Benjamins. doi: 10.1075/clscc.2.05lov

Lucie-Smith, E. (1969). *Movements in Art since 1945*. London: Thames & Hudson.

Lyons, J. (1977). *Semantics*. Cambridge: Cambridge University Press.

MacWhinney, B. (2005). The emergence of grammar from perspective. In D. Pecher & R. A. Zwaan (Eds.) *Grounding Cognition* (pp. 198–223). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511499968.009

Mareschal, D., Plunkett, K., & Harris, P. (1995). Developing object permanence: A connectionist model. In J. D. Moore & J. F. Lehman (Eds.) *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 170–175). Makwah, NJ: Erlbaum.

Marr, D. (1983). *Vision*. San Francisco: W.H. Freeman.

Marslen-Wilson, W. D. (1989). *Lexical Representation and Process*. Cambridge MA: MIT Press.

Martin, A. & Weisberg, J. (2003). Neural foundations for understanding social and mechanical concepts. *Cognitive Neuropsychology*, 20 (3–6), 575–87. doi: 10.1080/02643290342000005

McClelland, J. L. & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18: 1–86. doi: 10.1016/0010-0285(86)90015-0

McClelland, J. L., & Rumelhart, D. E. (1988). *Explorations in Parallel Distributed Processing*. Cambridge MA: MIT Press.

McNeill, D. (2012). *How Language Began. Gesture and Speech in Human Evolution*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9781139108669

Meltzoff, A. (1995). Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5), 838–850. doi: 10.1037/0012-1649.31.5.838

Miller, G.A. & Johnson-Laird, P. N. (1976). *Language and Perception*. Cambridge: Cambridge University Press. doi: 10.4159/harvard.9780674421288

Mithun, M. (1989). The acquisition of polysynthesis. *Journal of Child Language*, 16, 285–312. doi: 10.1017/S0305000900010424

Moltmann, F. (2013). *Abstract Objects and the Semantics of Natural Language*. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780199608744.001.0001

Montague, R. (1970). Universal grammar. *Theoria*, 36, 373–398. doi: 10.1111/j.1755-2567.1970.tb00434.x

Mountcastle, V. B. (1998). *Perceptual Neuroscience. The Cerebral Cortex*. Cambridge MA & London: Harvard University Press.

Nelson, K. (1973). Structure and Strategy in Learning to Talk. *Monographs of the Society for Research in Child Development*, 38 (serial no. 149). doi: 10.2307/1165788

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science*, 14, 11–28. doi: 10.1207/s15516709cog1401_2

Nichols, J. (1986). Head-marking and dependent-marking grammar. *Language*, 62: 56–119. doi: 10.1353/lan.1986.0014

Nordlinger, R., Wigglesworth, G., Kelly, B., & Forshaw, B. (Forthcoming). The acquisition of Murrinh-Patha. In M. Fortescue, M. Mithun & N. Evans (Eds.) *The Oxford Handbook of Polysynthesis*. Oxford: OUP.

Paivio, A. (1986). *Mental Representations: a Dual Coding Approach*. New York/ Oxford: Oxford University Press.

Panksepp, J., & Bernatzky, G. (2002) Emotional sounds and the brain: the neuro-affective foundations of musical appreciation. Behavioural Processes, 60(2): 133–155. doi: 10.1016/S0376-6357(02)00080-3

Paradis, M. (2004). *A Neurolinguistic Theory of Bilingualism*. Amsterdam/ Philadelphia: John Benjamins. doi: 10.1075/sibil.18

Parker, A. (2005). *Seven Deadly Colours. The Genius of Nature's Palette*. London: The Free Press.

Pasamonik, C. (2012). "My heart falls out": Conceptualizations of body parts and emotion in Beaver Athabascan. In A. Idström & E. Piirainen (Eds.) *Endangered Metaphor* (pp. 77–102). Amsterdam: John Benjamins. doi: 10.1075/clscc.2.04pas

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do we know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12), 976–87. doi: 10.1038/nrn2277

Pawley, A. (2011). Event representation in serial verb constructions. In J. Bohnemeyer & E. Pederson (Eds.) *Event Representation in Language and Cognition* (pp. 13–42). Cambridge: Cambridge University Press.

Peirce, C. S. (1932). *Collected Papers of Charles Sanders Peirce*, vol. I–II. Cambridge, Mass.: Harvard University Press.

Pepperberg, I. M. (1999). *The Alex Studies*. Cambridge, Mass.: Harvard University Press. [paperback version 2002]

Petersen, R. (1976). Nogle træk i udviklingen af det grønlandske sprog efter kontakten med den danske kultur og det danske sprog. *Tidskrift Grønland*, 165–208.

Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic Books. doi: 10.1037/11168-000

Pinker, S. (1984). *Language Learnability and Language Development*. Cambridge, MA: MIT Press.

Pinker, S. (1994). *The Language Instinct*. New York: William Morrow & Company. doi: 10.1037/e412952005-009

Polanyi, M. 1958 [paperback edition 1974]. *Personal Knowledge. Towards a Post-critical Philosophy*. Chicago: University of Chicago Press.

Popper, K. (1973). *Objective Knowledge: an Evolutionary Approach*. Oxford: Oxford University Press. [corrected version, first published 1972]

Popper, K. (1978). Three Worlds. *The Tanner Lecture on Human Values*. Delivered by Karl Popper at the University of Michigan on April 7, 1978 (pp. 143–167).

Prinz, J. J. (2005). Passionate thoughts: The Emotional Embodiment of Moral Concepts. In D. Pecher & R. A. Zwaan (Eds.) *Grounding Cognition* (pp. 93–114). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511499968.005

Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, 9 (2), 129–54. doi: 10.1080/713752551

Pulvermüller, F. (2002). *The Neuroscience of Language*. Cambridge: Cambridge University Press.

Pulvermüller, F. (2013). How neurons make meaning: brain mechanisms for embodied and abstract-symbolic semantics. *Trends in Cognitive Sciences*, 17 (9), 458–470. doi: 10.1016/j.tics.2013.06.004

Pulvermüller, F. & Schumann, J. H. (1994). Neurobiological mechanisms of language acquisition. *Language Learning*, 44:4, 681–734. doi: 10.1111/j.1467-1770.1994.tb00635.x

Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge MA: MIT Press.

Reiser, M. (1993). *Memory in Mind and Brain: what Dream Imagery Reveals*. New Haven/ London; Yale University Press.

Rice, S. (1998). Giving and taking in Chipewyan: The semantics of THING-marking classificatory verbs. In J. Newman (Ed.) *The Linguistics of Giving* (pp. 97–134). Amsterdam: John Benjamins. doi: 10.1075/tsl.36.05ric

Rice, S. (2012). "Our language is very literal": Figurative expression in Dene Sųłiné [Athapaskan]. In A. Idström & E. Piirainen (Eds.) *Endangered Metaphor*, 21–76. Amsterdam: John Benjamins. doi: 10.1075/clscc.2.03ric

Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2), 131–41. doi: 10.1016/0926-6410(95)00038-0

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.) *Cognition and Categorization* (pp. 27–48). Hillsdale NJ: Lawrence Erlbaum Associates.

Russell, B. (1905). On Denoting. *Mind*, 14 (1905), 479–493. doi: 10.1093/mind/XIV.4.479

Russell, B. (1959). *Wisdom of the West*. London: Rathbone Books Limited.

Sacks, O. (1995). *An Anthropologist on Mars*. London: Picador.

Schroeder, T. (2007). A recipe for concept similarity. *Mind and Language*, 22, 1: 68–91. doi: 10.1111/j.1468-0017.2006.00300.x

Schrödinger, E. (1944). *What is Life? The Physical Aspect of the Living Cell*. Cambridge: Cambridge University Press.

Searle, J. R. (1969). *Speech Acts*. London/ New York: Cambridge University Press. doi: 10.1017/CBO9781139173438

Sherburne, D. W. (1981). *A Key to Whitehead's Process and Reality*. Chicago: University of Chicago Press. [originally published 1929 by Macmillan Publishing Co.]

Singer, W. (2013). Cortical dynamics revisited. *Trends in Cognitive Sciences*, Vol. 17 (12), 616–626. doi: 10.1016/j.tics.2013.09.006

Sinha, C., & Jensen de López, K. (2000). Language, culture and the embodiment of spatial cognition. *Cognitive Linguistics*, 11: 17–41.

Sinha, C., Da Silva Sinha, V., Zinken, J., & Sampaio, W. (2011). When time is not space: the social and linguistic construction of time intervals and temporal event relations in Amazonian culture. *Language and Cognition*, 3, 1: 137–169. doi: 10.1515/langcog.2011.006

Skipper, J. I. (2015). The NOLB model: a model of the natural organization of language and the brain. In R. M. Willems (Ed.) *Cognitive Neuroscience of Natural Language Usage* (pp. 101–134). Cambridge: Cambridge University Press. doi: 10.1017/CBO9781107323667.006

Slobin, D. I. (1985). Crosslinguistic evidence for the language-making capacity. In D. I. Slobin (Ed.) *The Crosslinguistic Study of Language Acquisition, vol. 2: Theoretical Issues* (pp. 1157–1256). Hillsdale NJ/ London: Lawrence Erlbaum Associates.

Slobin, D. I. (1996). From "thought and language" to "thinking for speaking". In J. Gumperz and S. Levinson (Eds.) *Rethinking Linguistic Relativity* (pp. 70–96). Cambridge: Cambridge University Press.

Smith, A. (1982). *The Wealth of Nations*. London: Penguin. [first published 1776]

Strawson, P. F. (1974). *Subject and Predicate in Logic and Grammar*. London: Methuen & Co.

Sweetser, E. (1990). *From Etymology to Pragmatics*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511620904

Talmy, L. (1996). Fictive motion in language and "ception". In P. Bloom, L. Nadel, & M. A. Peterson (Eds.) *Language and Space* (pp. 211–276). Cambridge, MA: MIT Press.

Talmy, L. (2000). *Towards a Cognitive Semantics*. Cambridge, MA: MIT Press.

Taylor, J. (1989). *Linguistic Categorization*. Oxford: Clarendon Press.

Thurston, W. R. (1988). How exotic languages build a lexicon: Esoterogeny in West New Britain. In R. Harlow & R. Hooper (Eds.) *VICAL 1: Papers from the Fifth International Conference on Austronesian Linguistics* (pp. 555–579). Aukland: Linguistic Society of New Zealand.

Tomasello, M. (1992). *First Verbs: A Case Study in early Grammatical Development*. New York: Cambridge University Press. doi: 10.1017/CBO9780511527678

Trainor, L. J., & Schmidt, L. A. (2003). In I. Peretz & R. Zatorre (Eds.) *The Cognitive Neuroscience of Music* (pp. 310–324). Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198525202.003.0020

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.) *Organization of Memory*, (pp. 381–403). New York: Academic Press.

Ullman, M.T. (2004). Contributions to memory circuits to language: The declarative/ procedural model. *Cognition*, 92, 231–270. doi: 10.1016/j.cognition.2003.10.008

Van Valin, Jr., R. D. (2001). *An Introduction to Syntax*. Cambridge: Cambridge University Press.

Varela, F. J., Thompson, E., & Rosch, E. (1992). *The Embodied Mind*. Cambridge, MA: MIT Press.

Vendler, Z. (1967). *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.

Vigliocco, G., Kousta, S.-T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2013). The neural representation of abstract words: the role of emotion. *Cerebral Cortex*. http://dx.doi.org/10.1093/cercor/bht025.

Vygotsky, L. S. (1962). *Thought and Language*. Cambridge MA: MIT Press. doi: 10.1037/11193-000

Whitehead, A. N. (1947). *Adventures of Ideas*. Cambridge: Cambridge University Press.

Whitehead, A. N. (1978). *Process and Reality* (corrected edition). New York: The Free Press (Macmillan) [original British edition 1929, Cambridge: Cambridge University Press]

Whitehead, A. N. (1938). *Modes of Thought*. New York: Macmillan [as Free Press paperback 1968]

Whitehead, A. N., & Russell, B. (1910–13). *Principia Mathematica*, vols. 1–3. Cambridge: Cambridge University Press. [1962 abbreviated edition, Cambridge: Cambridge University Press]

Wieger, L. (1965). *Chinese Characters*. New York: Paragon Book Reprint and Dover Publications. [first edition 1915]

Wierzbicka, A. (1992). The search for lexical primitives. In M. Pütz (Ed.), *Thirty Years of Linguistic Evolution* (pp. 215–42). Philadelphia/Amsterdam: John Benjamins. doi: 10.1075/z.61.20wie

Wittgenstein, L. (1974) [first edition 1953]. *Philosophical Investigations*. Oxford: Basil Blackwell.

Zelinsky-Wibbelt, C. (2011). Introduction. In C. Zelinsky-Wibbelt (Ed.) *Relations between Language and Memory* (pp. 1–23). Frankfurt am Main: Peter Lang.

Zwaan, R. A. (2008). Time in language, situation models, and mental simulations. *Language Learning*, 58, suppl. issue 1, 13–26. doi: 10.1111/j.1467-9922.2008.00458.x

Zwaan, R. A. & Madden, C. J. (2005). Embodied sentence comprehension. In D. Pecher & R. A. Zwaan (Eds.) *Grounding Cognition* (pp. 224–245). Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511499968.010

Name index

Subject index

The main thesis of this book is that abstraction, far from being confined to higher forms of cognition, language and logical reasoning, has actually been a major driving force throughout the evolution of creatures with brains. It is manifest in emotive as well as rational thought. Wending its way through the various facets of abstraction, the book attempts to clarify – and relate – the often confusing meanings of the word 'abstract' that one may encounter even within the same discipline. The unusual synoptic approach, which draws upon research in psychology, neural network theory, child language acquisition, philosophy and consciousness studies, as well as a variety of linguistic disciplines, cannot be compared directly to other books on the market that touch upon just one particular aspect of abstraction. It is aimed at a wide readership – anyone interested in the nature of abstraction and the cognitive processing and purpose behind it.

"Abstract, abstractness, abstraction: we know more or less what we mean by these terms. Or do we? Michael Fortescue, in this provocative and compelling book, analyzes types of abstraction and processes that involve them, in a wide-ranging survey of theories in linguistics, neuroscience, philosophy, and psychology; placing them in the context of cosmological and evolutionary considerations."

Chris Sinha, *Hunan University*

"The ability to abstract patterns from the particular is central to human thought and language. In this important and engaging book, Fortescue argues convincingly that abstraction is the hallmark of human cognition. Both compelling and insightful, the book brings together different approaches and findings from across the language and cognitive sciences. The result is a new and exciting perspective on the not-so-small matter of what makes us so smart."

Vyvyan Evans, *www.vyvevans.net*

ISBN 978 90 272 1361 7

John Benjamins Publishing Company