1 Introduction: space, geometry, mind

This book is about the geometry of linguistic meaning. It outlines a theory the full sense of which depends on its geometrical formalism. This is a simple but hopefully consistent formalism that most people will recognise from high-school geometry. It is perhaps an approach that may strike readers as odd. Yet I think it is a perfectly natural way to approach the concerns that linguists, and especially cognitive linguists, have had for decades, in particular the relationship between space, cognition and language. If spatial concepts play a crucial role in linguistic meaning, then geometry ought to be a useful way to describe these meanings. What is more, the linguistic literature is full of hints that geometrical modelling is a natural way to go. Our textbooks and monographs are replete with arrows, axes, and words like *perspective*, *location*, *direction* and *distance*. Why multiply vague English expressions when we have a well-defined tried and tested mathematical formalism? Even more can be said in its favour. For geometry itself, and its standard notation, is motivated in a way that cognitive linguists are familiar with. It is, in its Euclidean form, based on human bodily experience. And it is Euclidean three-dimensional geometry that we shall use. True, multidimensional vector spaces are needed if we are trying to model meaning in an information-science or a connectionist framework (see for example Widdows 2004) but we are trying here to model meaning in a cognitive embodied framework, and our fundamental spatio-cognitive understanding is in fact three-dimensional, thus natural even when spatial cognitions are projected to non-spatial realms.

This is not a book that attempts mathematical proofs of the claims expressed in the diagrams, though it is not ruled out that a more rigorous mathematical demonstration would be worth attempting. Nor does the book use algebraic formulations, although in some cases the complex abstract relations expressed in the diagrams could be more simply displayed in that way. The guiding principle has been to rely on our visual intuitions about geometrical figures, always seeking to be as consistent and clear as possible. Of course, the concern is not just to have a neater, more encompassing, better motivated notation or diagramming system. The more one explores three-dimensional figures, the more it seems that this kind of diagram is able to capture some deep facts about words and constructions and perhaps also to resolve some puzzles or pseudo-puzzles.

1.1 Language and mind

We do not study language, or a language, only for the sake of knowing more about language or languages, but because language, any language, tells us something about some of the workings of the human mind, a view that is now common among cognitive linguists. This is not quite as straightforward as it might seem and needs a little clarification. The term ‘mind’ is deliberately chosen: studying language does not tell us about the brain, at least only in an extremely indirect manner. Investigating language tells us something about the sorts of things the human mind does or can do. I certainly do not mean that mind is independent of brain, but investigating language does not tell us directly about neurons and neuronal networks. And by investigating language, I do not mean investigating languages or investigating examples of the use of language in attested contexts.

Language is a tool for stimulating silent conceptualisation, and I do not mean only conscious conceptualisation. What sort of conceptualisations is language capable of stimulating? If we look into this question we are going to find, if we are lucky, only a small part of the conceptual activity that takes place in our minds. The term ‘conceptualisation’ is being used here in preference to ‘meaning’, since this term is often connected to the view that language elements (lexical items, constructions) have meaning in some independent language system in the mind. However, along with other cognitive linguists, I am adopting the view that there is no independent linguistic–semantic inventory, but rather that there is conceptualisation that can happen independently of language, but which language accesses. I also take the view that the particular structure of a language – its lexis and grammar – does not privilege or make more accessible or delimit some particular kinds of conceptualisation. That is, I think there is no clear evidence for the strong Whorf hypothesis – that, for example, a classifier language will make speakers more sensitive to, or restrict speakers to, certain kinds of conceptualisation. This is not to say that some empirical work does not give convincing evidence of certain limited effects (Lucy 1992, Levinson 1996, 2003), but these effects appear not to be extensive or not sufficient to make us think that linguistic meaning is something separate from a universal conceptualising ability.[1]

[1] It may be different if one’s focus is on discourse. It is simpler and faster to code those conceptual distinctions that one’s lexical or grammatical structure encodes but the generative properties of language mean that one is not limited to those structures in communicating conceptualisations of all sorts. If anything constrains them, it is cultural practice, which is also not an absolute constraint.

This view is of course different from the cognitive view inspired by Chomsky and his followers, according to which language is autonomous, modular, interfaced with a semantic component. I am not, however, throwing out the idea that the language system may be modular in some sense – but it is not sealed off or encapsulated (in the sense of Fodor 1983): rather, it is linked inextricably to non-linguistic cognitive abilities.

It is also different from denotational semantics, which is why I use the term 'semantics' throughout this book sparingly, usually to refer to the natural concepts associated with lexical items, that is, their 'semantic frame' (Fillmore 1982b). Since the 1980s cognitive scientists and linguists have realised that the meanings of words and constructions in languages are in some way built up on the basis of our embodied experience of physical space (Lakoff and Johnson 1980, 1999, Johnson 1987, Pinker 1997). Pinker sums up some of the conceptual elements that are important for the geometrical approach:

> Space and force pervade language. Many cognitive scientists (including me) have concluded from their research on language that a handful of concepts about places, paths, motions, agency, and causation underlie the literal or figurative meanings of tens of thousands of words and constructions, not only in English but in every other language that has been studied . . . These concepts and relations appear to be the vocabulary and syntax of mentalese, the language of thought . . . And the discovery that the elements of mentalese[2] are based on places and projectiles has implications for both where the language of thought came from and how we put it to use in modern times. (Pinker 1997: 355)

'Space' and 'force' are crucial to the approach developed in the present book, because they can be, and commonly are, formalised in geometric terms, as vectors. The nature of the linkage between spatial cognition and the structures of language is similar to metaphorical linkage: spatial, motor and visual systems have some sort of counterparts in the linguistic–communicative system. For example, humans can attend to detail or take in a gross gestalt, and linguistic constructions and lexical items also permit this kind of alternating focus. Depth perception enables us to judge one object as more distant from the self than another object opposed to another: similarly, linguistic structures can place information in the foreground or background that is, topicalise some information relative to other information (cf. Talmy 2000 [1983], Langacker 1995, Croft and Cruse 2004). Spatial expressions, and the spatial perceptual–cognitive systems, provide a source domain for linguistic expressions of time, in many, perhaps all, languages (Lakoff and Johnson 1980, 1999, Haspelmath 1997, Boroditsky 2000, Evans 2004, Evans and Chilton 2010).

[2] Or 'language of thought', Jerry Fodor's hypothesis that thinking is executed in a mental symbols system. It is not relevant to address the plausibility of that hypothesis here.

Such views as these are not uncommon amongst linguists. However, the apparatus that I am proposing is an unusual one in linguistics, though it is hinted at in many places, as will become clear by way of the references that I shall give as we move along. Moreover, the formal apparatus itself, as will be seen, is one that springs from natural foundations – namely, the fact (I am assuming it is a fact) that spatial location and orientation are experiences that organisms must cognise for their survival. We shall see, using the apparatus that I will outline, that these concepts of physical space can be used for other kinds of fundamental concept that languages use as their bedrock and about which they enable humans to communicate among one another. This idea is not new in cognitive linguistics (for example, Lakoff and Johnson 1980, 1999, Langacker 1987, 1991, Talmy 2000; cf. also 'localists', for example, Anderson 1971, Lyons 1977) or even outside it (cf. Jackendoff, for example, 1976, 2002a, Frawley 1992). What I'm suggesting is that there is a derived conceptual foundation that uses spatial cognition of various kinds, and that language in turn uses it as the most fundamental requirement for communication. Actually, this means that language is not simply a 'window' on the mind, but that investigating language and conceptualisation together, using the apparatus I describe below, will tell us something about both. It is perhaps not surprising that we need to view language through different kinds of apparatus and that indeed language may, after all, have components and modules that that require different instruments of investigation.

1.2 Formalisation

Linguistics is bedevilled by a proliferation of formalised notations and diagrams, the most systematic of which are bound up with propositional and predicate calculus and other elements of mathematical logic and reasoning. Others are ad hoc, intuitive and iconic. A few have origins in Euclidean geometry.[3]

1.2.1 Discourse Representation Theory

Possibly the most developed of discourse theories is the Discourse Representation Theory (DRT) of Kamp and Reyle (1993), a theory that has had some influence on the present approach. However, DRT does not claim to be cognitively motivated, although Kamp does occasionally maintain the cognitive *relevance* of DRT (see also Asher and Lascarides 2003: 376ff.). It is true

[3] This section is based on part of Chilton (2005).

that Kamp's project does resolve major problems, specifically anaphora, indefinite NP reference and definite NP reference, much debated in the logic-oriented tradition of twentieth-century philosophy. It is also true that in the more recent work, such as the SDRT (Segmented Discourse Representation Theory) of Asher and Lascarides (2003) and 'dynamic' versions of the theory (see for example Kamp *et al*. 2011), DRT goes still further in explaining various phenomena of discourse coherence and context dependence. However, the DRT apparatus itself and its newer incarnations may be constrained by their predicate calculus foundations: in any event, they do not appear to incorporate systematically the concepts of situatedness or embodiment. Discourse Representation Theory does not, for example, handle deixis in a naturalistic fashion, limiting deixis to objects in the non-linguistic context, similar to the way anaphoric referents are treated. For the essentially deictic concepts coded by linguistic tense, DRT also treats 'times' as referents more or less on a par with other discourse referents ('yesterday' has equivalent status to 'Peter' or 'the donkey'), which seems counter-intuitive. 'Yesterday' is the 'location', relativised to the speaker, at which, for instance, Peter beat his donkey, not a participant in that event. In the present theory, by contrast, it is assumed that temporal and spatial – and also modal – deixis are fundamental; consequently, the model integrates them into the representation of all utterances in discourse. It is the relativisation of 'yesterday' to another time point, that of the time of speaking, that is left out in DRT, yet this fact is precisely an intrinsic feature of discourse.

1.2.2 Mappings

Set theory and functions, alongside predicate calculus, have been fundamental to formal semantics and pragmatics. Within cognitive linguistics itself, set theory and functions have been used in a somewhat different way. Mappings across cognitively defined sets of one kind and another are systematically used in mental space theory (Fauconnier 1994, 1997), in blending theory (Fauconnier and Turner 2002), and more informally in conceptual metaphor theory (Lakoff and Johnson 1980, 1999). The strength of the concept of mapping lies in its potential mathematical clarity, which enables it to model the phenomena of mental spaces that were first precisely identified by the authors just mentioned. As in the case of predicate logic, however, there is no inherent connection between sets-and-mappings and cognitive or linguistic processes that are situated and embodied. Specifically, they do not incorporate deictic phenomena. The geometrical approach I am putting forward in this book owes much to the mental spaces approach, but recasts mappings across spaces as coordinate correspondences on three fundamental dimensions – three dimensions that we will spend some time explaining in the chapters to follow. This means that any point in the three-dimensional deictic space has three coordinates with respect to the three axes defining the space in which such a point occurs. Axis systems can be nested inside the base axis system, as will be seen, and points can be defined by their coordinates within a nested axis system. It is a logical consequence of the geometry that correspondences can be tracked across systems, relatively to the base system – or that what we might call in less abstract terms, the 'base world' or 'base reality' of the deictic centre, the speaker.

1.2.3 Diagrams

Diagrams are ubiquitous in linguistics, but in cognitive linguistics they arguably have a distinctive place in that particular approach to language. Talmy's detailed accounts, for instance, have this geometric quality (for example the account of path concepts in Talmy 2000 [1985]), as does Langacker's consistent use of pictorial diagrams (Langacker 1987, 1991, etc.). Langacker-diagrams, which are designed to capture a range of intuitions about linguistically encoded meanings (e.g. 'foregrounding', 'prominence', 'trajector', 'distance'), can probably be transposed into standard geometrical concepts. For instance, Langacker-diagrams generally encapsulate topological relations, directionality, relative distance and scalar magnitude. It is important that there is also a claim that the iconic diagrams capture *non*-linguistic perceptual or conceptual phenomena, with a particular emphasis on vision.

The most recent advocate of a more explicitly geometrical approach in cognitive grammar is Croft (2012), who develops 'geometric-cum-graph structures' diagrams for the purpose of describing lexical aspect (*Aktionsart*), the time-related event structure that is part and parcel of verb semantics, and also for the purpose of describing causal event structure of verbs such as *break*. These diagrams resemble familiar two-dimensional Cartesian coordinate diagrams. They use two clearly defined axes, a time axis and a q-axis, the latter providing a place for separate qualitative states in event structures (Croft 2012: 53–7). This system has the advantage over earlier theoretical models of making the time course of verbal aspect explicit and easily visualised, yielding a richer and more coherent description of the conceptual frames associated with the verbs of a language.

In his account of complex structure of causative verbs, Croft (2012: 212–17) says that he is using three-dimensional diagrams, the third dimension being the causal relation between two phases of a causal chain. However, the diagrams themselves do not show three dimensions defined by three axes, as one might expect. And a further detail appears: the placing not only of qualitative states q on the second axis but also participant entities associated with each *q*. In fact, the diagrams that model the causal structures of the sentence *Jack broke the vase* consist of two two-dimensional diagrams for two subevents: one for Jack making an impact, one for the vase's state of being broken. In addition 'the causal chain linking the individual participant subevents is represented in a third dimension'. It is not clear how this holds together geometrically – how the third dimension relates geometrically to the first two. It is also not clear how this third dimension, one reserved specially for causal relations, is to be defined. In Chapter 3, I attempt a different kind of geometrical modelling for causative verbs, bringing in another concept that comes for free with geometry, the concept of force vectors.

There are some similarities between Croft's geometric diagrams and the ones I develop in this book.[4] Both Croft's and mine use a time axis and an orthogonal axis relating to the structure of events. But there are also some significant differences. Whereas Croft's *q*-axis is specifically concerned with the stages in the unrolling of events (e.g. the successive states in an event such as the one denoted by the verb *melt*), my second axis provides coordinates for discourse entities, i.e. the participants in event structure. The most significant difference is that the abstract geometric space I work with is explicitly three-dimensional. As will be explained in Chapter 2, my geometric model has an essential third axis, which is epistemic.

1.3 Using geometry

Whereas Croft, and to a lesser extent Langacker, borrow elements of geometry to model meanings of linguistic expressions, the present book starts at a more fundamental level. To begin with, I attempt to apply geometry to the sorts of meaning that one would expect geometry to be well equipped to describe – spatial meanings. Then, building on the insight of many linguists that spatial meanings are somehow crucial for many kinds of more abstract meaning, I look into the potential of geometry for describing apparently non-spatial meanings, which may nonetheless be derived from, or linked to, spatial conceptualisation. This in turn means taking seriously certain elements of Euclidean (and Cartesian) geometry and developing out of them a descriptive model of language-based conceptual space. The elementary geometry of coordinate systems, vectors and transformations then becomes a heuristic for exploring linguistic–conceptual space as well as a descriptive model-building project.

The theoretical ideas that I develop in the present book draw on insights and methods found in the broad approaches outlined in the preceding sections,

[4] They were developed independently.

as well as many other meaning-based approaches to language. What I am putting forward is a research framework that investigates the applicability of coordinate vector geometry for the description and perhaps explanation of certain kinds of linguistic meaning. I am not suggesting that *all* linguistic meaning can be described or explained in this way, only that the range of lexical–constructional meanings that can be described in this way is wider than linguists may have thought.

There is nothing new about the use of coordinate geometry to account for spatial meanings, for example in prepositions and deictic expressions. Bühler, in writing of the 'deictic field', clearly has Cartesian coordinate geometry in mind. His description of what he calls 'the here-now-I system of subjective orientation' begins starkly:

> Let two perpendicularly intersecting lines on the paper suggest a coordinate system to us, 0 for the origin, the coordinate source . . . My claim is that if this arrangement is to represent the deictic field of human language, three deictic words must be placed where the 0 is, namely the deictic words *here*, *now* and *I*. (Bühler 1990 [1934])

Bühler's idea that deixis is one fundamental aspect of human language and has something to do with self, space and time leads naturally to geometric modelling. It is in some ways also the starting point for the Deictic Space Theory that I shall outline in this book. However, the space in question will be three-dimensional and the three axes will be defined conceptually in ways that are distinct from Bühler's ideas.[5] Bühler also spoke of 'displacement'[6] of the deictic centre away from the here and now of utterance (Bühler 1990 [1934]: 136–57). As will be seen in later chapters such deictic displacements can be modelled geometrically as transformations of coordinate systems. In more recent linguistic research, spatial expressions that relate objects to reference points other than self, particularly in spatial prepositions, have also been treated in explicitly geometric terms.

Highly abstract geometry and topology have been used to model language phenomena by Thom (1970) and Petitot (e.g. 1995). Gärdenfors (2000) argues for the geometrical representation of conceptualisation. Gallistel (1990) argues that low-dimensioned geometries are fundamental in the vertebrate nervous system. Of particular relevance is the empirical work of Paillard and Jeannerod amongst other neuroscientists, which shows that visuo-spatial

[5] The geometric approach may be in the background to the terminology used by Appolonius Dyscolus, as suggested by Dalimier (2001). Apollonius and the Stoics drew a broad distinction between content expressions and deictic expressions, as Bühler mentions.

[6] In Bühler (1990 [1934]) the translation is 'transposition', the translation used by Levinson (2003: 51). Bühler used *Versetzung* or *Verschiebung*, the latter term being associated with contemporary psychology. Bühler (1990) is an abridged and translated extract from his 1934 work.

processing and actions such as reaching and grasping depend on the brain's ability to work with egocentric coordinates systems whose origin is located at different parts of the body (shoulder, hand, retina, for example) and in different sensory modalities (Paillard 1991, Jeannerod 1997).[7] This work builds on other important empirical work (Ungerleider and Mishkin 1982, Goodale and Milner 1992, Milner and Goodale 1995) that indicates two pathways in visual processing, one relating to the identification of objects and the other locating objects in body-centred frames of reference. Such finding are highly suggestive both for investigating the linguistic phenomenon of deixis and for using a geometrical approach to linguistic description.

While the work mentioned so far focuses on egocentric coordinate systems, experimental findings of O'Keefe and Nadel (1978) point to the brain's ability to represent allocentric and absolute maps of an organism's spatial environment, as noted by Levinson (2003: 9–10). Levinson argues strongly for the need to recognise non-egocentric systems in the spatial semantics of certain languages. However, there need be no conflict between the two research emphases. It would appear likely that both egocentric (deictic) coordinate systems and non-egocentric systems are instantiated in neural structures and mental representations, including linguistic ones.

I shall not attempt to explore this body of scientific work further, nor attempt to apply any but the simplest geometrical concepts. Sufficient has, I hope, been said to justify the scientific appropriateness of a geometric approach in cognitive linguistics.

1.4 Space, situation and deixis

In the study of spatial semantics the most precise use of geometry, in particular of coordinate systems, has been Levinson (2003). As will be clear in Chapter 2, geometrical descriptions of linguistic expressions referring to physical space provide a foundation for the present book, though the aim is to move into an abstract conceptual geometry that can handle linguistic expressions and constructions that are not spatial in any obvious sense.[8]

This book is not concerned with spatial expressions as such – and it is important to emphasise this at the outset. But there is now a rich body of research into the precise nature of spatial meaning in language and the precise nature of the relationship between language and spatial cognition (see Evans and Chilton 2010 for a sample of recent theoretical and experimental work). In this work, while geometric aspects (or geometric descriptions) of spatial cognition are well established, geometry is not the whole story. In the case of

[7] Bühler was already aware of the shifting origo on the body: Bühler (1990 [1934]: 146–7).
[8] Parts of this section are based on Chilton (2005).

spatial prepositions, which have the most obvious links to spatial cognition, it is now widely recognised that for spatial prepositions a coordinate geometry alone is not adequate, though it is needed and is in some sense fundamental. Many linguists and psycholinguists have shown the crucial role played also by the shape, presence or absence of surface contact and the function of reference objects. What is loosely called 'topology' is also relevant (e.g. for prepositions like *in* which involve a concept of bounded space), although topology is indeed a kind of geometry (Herskovits 1986, Vandeloise 1991, Carlson-Radvansky *et al*. 1999, Coventry *et al*. 2001, Tyler and Evans 2003). It is also the case that when prepositions (e.g. *in front of*) locate objects relative to a reference point, it is not actually a point that is at issue, but a spatial region around the reference landmark. Such conceptualised regions are affected by the function and shape of the landmark itself (see for example Carlson 2010).

But coordinate systems (or reference frames) remain the foundation of the cognitive system underlying spatial orientation. The crucial point is that a rather simple geometrical formalism is an economical way, and indeed a natural way, of describing fundamental spatial meanings. On a more abstract level, what coordinate geometry enables us to do is to analyse 'point of view' or 'perspective', both in a physical–spatial and in various abstract sense. As is now recognised (see for example Tomasello 1999), one of the reasons why language is so complex and apparently redundant if we look at it from a truth-conditional position may well be that language 'is designed' to permit viewpoint alternation. Euclidean geometry is a well-understood way to analyse space the way humans experience it. This book is proposing that it is also a natural way to describe non-spatial linguistic meanings that are nonetheless in some way rooted in human spatial experience.

Once we introduce the geometry of coordinate systems it is a natural step to introduce vectors – mathematical objects conventionally visualised as arrows that have (i) direction and (ii) magnitude. In a coordinate system the position of a point can be given by the length and direction of a vector from the origin to the point. This approach can be used for explicating the denotation of certain prepositions, including *in front of*, by specifying a vector space in which all vectors have the same origin in some coordinate system (O'Keefe 1996, 2003, O'Keefe and Burgess 1996, Zwarts 1997). This space will be the 'search domain' within which an object can be said to be, for example, 'in front of John'.

Cognitive linguists often assert that linguistic meaning is situated and/or embodied (e.g. Lakoff and Johnson 1999, Croft and Cruse 2004, and, from a philosopher's perspective, Johnson 1987). To be situated means, for human language users, that our use of language will always assume and/or refer to the place and time of speaking, and will take a perspective on the surrounding