09-05-05

Semantics, conceptual spaces and the meeting of minds

Massimo Warglien[1] and Peter Gärdenfors[2]

[1]Advanced School of Economics and Department of Business Economics, Ca' Foscari Università, Venezia, Italy
warglien@unive.it

[2]Lund University Cognitive Science, Kungshuset, S-22222 Lund, Sweden
Peter.Gardenfors@lucs.lu.se

Abstract: We present an account of semantics that is not construed as a mapping of language to the world, but mapping between individual meaning spaces. The meanings of linguistic entities are established via a "meeting of minds." The concepts in the minds of communicating individuals are modeled as convex regions in conceptual spaces. We outline a mathematical framework based on fixpoints in continuous mappings between conceptual spaces that can be used to model such a semantics. If concepts are convex, it will in general be possible for the interactors to agree on a joint meaning even if they start out from different representational spaces. Furthermore, we show by some examples that the approach helps explaining the semantic processes involved in the composition of expressions.

1. Introduction

Within traditional philosophy of language, a semantics is seen as a mapping between a language and the world (or several "possible worlds"). This view has severe problems. For one thing, it does not involve the users of the language. In particular, it does not tell us anything about how individual user can "grasp" the meanings determined by such a mapping (Harnad 1990, Gärdenfors 1997). Another tradition, cognitive semantics, brings in the language user by focusing on the relations between linguistic expressions and the user's mental representation of the meanings of the expressions, often in form of "image schemas." However, cognitive semantics has problems in explaining the social nature of semantics.

In this article, we propose a radically different view of semantics based on a "meeting of minds." According to this view, the meanings of expressions do not reside in the world or solely in the mental schemes of individual users, but they *emerge* from the communicative interactions between the language users. The fundamental role of human communication is indeed to affect the states of mind of others. A meeting of the minds means that the representations in the minds of the communicators will become sufficiently compatible.

As an example of how a meeting of minds can be achieved by communication, but without the aid of language, let us consider declarative pointing (Bates 1976, Brinck 2004, Gärdenfors and Warglien to appear). This act consists of one individual pointing to an object or spatial location and at the same time checking that the other individual (the "recipient") focuses his or her attention on the same object or location. The recipient in turn must check that the pointer notices that the recipient attends to the right entity. This attending to each others' attention is called "joint attention" (Tomasello 1999) and it is a good, but fallible, mechanism for checking that the minds of the interactors meet in focusing on the same entity. In passing, note that pointing, unlike language, is a continuous way of referring to the outer world (the direction of the finger is continuously variable).

As a matter of fact, achieving joint attention can be seen as reaching a fixpoint in communication. When my picture of what I point out to you agrees with my image of what you are attending to, my communicative intent is in equilibrium. Conversely, when what you attend to agrees with your image of what I want to point out to you, your understanding is in equilibrium (Gärdenfors and Warglien to appear).[1]

When the interactors are communicating about the external world, pointing is sufficient to make minds meet on a referent. However, when the interactors need to share referents in their inner mental spaces a more advanced tool is required. This is where language proves it mettle (Brinck and Gärdenfors 2003, Gärdenfors 2003, Gärdenfors and Osvath, to appear). In a sense, language is a tool for reaching joint attention in our inner worlds. As a matter of fact, Goldin-Meadow (2007, p. 741) goes beyond our metaphorical assertion and writes that in children "pointing gestures form the platform on which linguistic communication rests and thus lay the groundwork for later language learning."

We shall assume that our inner worlds can be modeled as spaces with topological and geometric structure. Here we will use conceptual spaces (Gärdenfors 2000) as the main modeling tools. The mental space that carries the meanings for a particular individual is partly determined from the individual's interaction with the world, partly from her interaction with others and partly from her interaction with herself (e.g. in the form of self-reflection). This view does not entail that different individuals mean the same thing by using an expression, only that their communication is sufficiently successful.

[1] Two features of this process should be noted: Firstly, joint attention requires a "theory of mind" that includes second order attention, in the sense that both communicative partners can attend to the attention of the other. Secondly, the two partners need not have the same image of each others' inner state: It is perfectly possible to reach joint attention without my picture of your attention being aligned with your attention. In other words, joint attention never requires leaving the subjective realms of the communicative partners.

As a comparison, consider the models of cognitive semantics (see e.g. Lakoff 1987, Langacker 1986, 1987, Croft and Cruse 2004, Evans 2006) where image schemas have been core carriers of meaning. An image schema is a conceptual structure that belongs to a particular individual. However, the mathematical structures of the image schemas are seldom spelled out.[2] To do this, it is natural to work with topological and geometric notions.

The image schemas of cognitive semantics are in general presented as structures that are *common* to all speakers of a language. However, in the socio-cognitive type of semantics we model in this paper, we do not assume that everybody has the same meaning space, but only that there exist well-behaved mappings between the meaning spaces of different individuals – "well-behaved" in the sense that the mappings have certain mathematical properties (to be specified below).

Our modeling takes a lot of inspiration from the communication games that have been studied by Lewis (1969, 1979), Stalnaker (1979) and others (e.g. Schelling (1960), Clark (1992), Skyrms (1998), Parikh (2000)). To this tradition we are adding some assumptions about the topological and geometric structure of the individual mental spaces that will allow us to specify more substantially how the semantics emerges and what properties it has. Linguistic acts are best seen as moves in such games. The players in a communication game have different payoff functions, but we also accept that they may have different individual meaning spaces. Indeed, we shall show that semantic equilibria can exist without assuming that the communicating individuals have the same mental spaces.

As long as communication is conceived as a process through which the mental state of an individual affects the mental state of another one, a "meeting of the minds" is a condition in which both individuals find themselves in compatible states of mind that do not require further processing. Just like bargainers shake hands after reaching an agreement on the terms

[2] For an early computational model, see Holmqvist (1993).

of a contract, speakers may reach a point in which both believe they have understood what they are talking about. Of course, they may actually mean different things, just like the terms of a contract might prove to be interpreted differently by the bargainers. But it is enough that, in a given moment and a given context, speakers may reach a point in which they feel there is a mutual understanding – no matter whether mutual agreement implies or not that they mean the same thing.

A very common mathematical way to define such kind of state would be to identify it as a fixpoint. A fixpoint x^* of a function $f(x)$ is a point in which the function maps x^* on itself ($f(x^*) = x^*$). But what kind of object is a function that reaches a fixpoint when minds agree? In linguistic communication, the most natural candidate for such a semantics is a function that maps language expressions on mental states, and vice-versa – a kind of interpretation function and its inverse expression function. So, in our framework minds meet when the interpretation function mapping states of mind on states of mind via language finds a resting point – a fixpoint.[3]

To provide a simple example of convergence to a fixpoint as a meeting of mind, let us return to the previous example about achieving joint attention via pointing, for example a child pointing out something to an adult. The individual mental spaces are in this case taken to be their visual fields (which may only partially overlap). The goal of the pointing is to make the adult react by looking at the desired point the visual field. The fixpoint is reached when the child sees that the adult's attention is directed at the correct point and the adult believes that her attention is directed to what is pointed at (see figure 1). More precisely, the fixpoint is characterized by four properties: (1) the attended object is on my line of gaze; (2) the attended object is on your line of gaze; (3) I see that your line of gaze has the right direction; and (4)

[3] An analogy is that communication is like a dog on a leash pulling a human towards something the dog has in mind. The dog will pull and its master will follow and until an equilibrium is reached. The place where they stop is literally a fixpoint.

you see that my line of gaze is in the right direction.[4] Thus our gazing lines intersect in correspondence of the object, and our representations of the other gaze are consistent with such effective gazes.

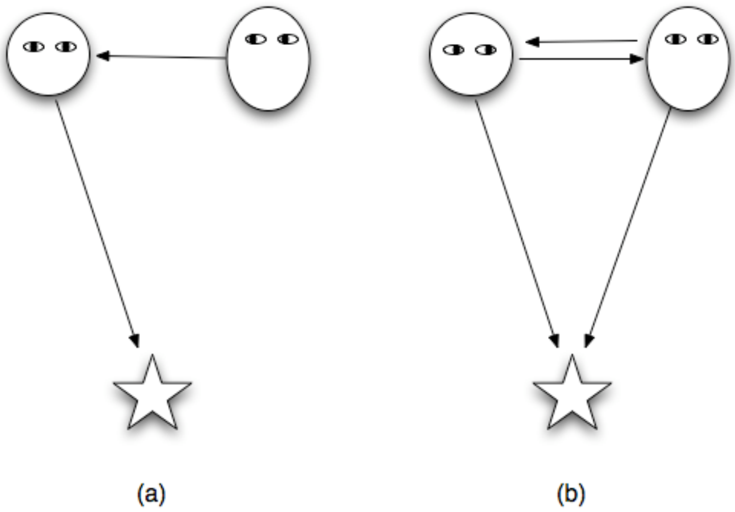

Figure 1: Joint attention as a fixpoint

Using fixpoints is, of course, not new to semantics. The semantics of programming languages often resort to fixpoints to define the "meaning" of a program: its meaning is where the program will stop (for a remarkable review, see Fitting 2002). In a different vein, Kripke's (1975) theory of truth is grounded on the notion of a fixpoint – in his case the fixpoints of a semantic evaluation function are at the focus of his interest. Fixpoints are also crucial in other fields, such as the study of semantic memory: content-addressable memories usually store information as a fixpoint of a memory update process (the canonical example being Hopfield neural net, see Hopfield 1982).

However, here we make a fairly different use of the fixpoint notion to define our "meeting of minds" semantics: We consider the fixpoints of an interactive, social process of meaning

[4] In fact what we observe is that the gaze of the other has the right angle. From that we calculate that the line of gaze has the right direction. The important fact to note is that this calculation involves a shift of perspective. For example when a child follows the gaze of somebody that focuses on a point outside its own visual field, it has to form an allocentric representation of the space. In general this can be modelled via a coordinate transformation together with an unlimited extension of the space.

construction and evaluation. From this point of view, our use of fixpoints resembles more the one made by game theorists to define equilibrium states of mutual compatibility of individual strategies. Furthermore, our fixpoints are of a topological nature, while those most used in computation and logics exploit properties of monotone functions on (partially) ordered structures.

Our argument is that some types of topological and geometric properties of mental representations afford meetings of minds, because they lend more naturally fixpoints to communication activities. Following Gärdenfors (2000), concepts will be represented as convex regions of mental spaces (see next section). Thus, we shift from the conventional emphasis on the way we share the same concepts to an emphasis on the way the “shapes” of our conceptual structures make it possible for us to find a point of convergence. A parallel with the pragmatics of conversation in the Gricean tradition comes to the mind: Just like maxims of conversation ensure that talk exchanges find a mutually accepted direction, we explore the complementary notion that the way we shape our concepts deeply affects the effectiveness of communication.

On this ground, we make an implicit selection argument: just like wheels are round because they make transportation efficient, we expect to identify the shapes of concepts that are selected to make communication smooth and memorizing efficient. As we will see, the convexity and compactness of concepts play central roles. In this way, constraints over the structure of concepts facilitate creation of coordinated meanings. In brief, our point is that communication works as long as it preserves the structure of concepts. This will later lead us to consider the role of continuous mappings in conveying similarity of meanings. An important point that we will elaborate in section 3.2, is that the preservation of similarities can be performed by a discrete system, to wit, the expressions of natural language. In this paper, we will focus on noun phrases and indexical expression, but our approach can be extended to other linguistic categories.[5]

2. The topology of conceptual spaces

It turns out that structural properties of conceptual representations that grant the existence of meetings of minds are to a large extent already familiar to cognitive semantics and in particular to the theory of conceptual spaces. These basic properties are the metric structure induced by similarity, the closed/bounded nature of concepts, convexity of conceptual representation, and the assumption that natural language, with all its resources, can "translate" (spatial) mental representations with reasonable approximation. In what follows, we will make more precise these notions and the role they play in a "meeting of minds" semantics theory.

Our first step is to assume, following Gärdenfors (2000), that conceptual spaces are construed out of primitive *quality dimensions* (often rooted in sensorial experience) and that similarity provides the basic metric structure to such spaces. The dimensions represent various "qualities" (color, shape, weight, size, position, etc) of objects in different domains.[6]

To be more precise, we recall that a metric space is a set of points with a measure of the degree of closeness (or distance) between such points. A metrizable topological space is a space whose topological structure is induced by some metric. Thus, our fundamental assumption is that conceptual spaces are metrizable, and that their specific metric structure is induced by a similarity relation. This leaves open the way to many different metric structures. While the nature of psychologically sound similarity measures is still highly controversial (and presumably differs between domains), numerous studies (Shepard 1987, Nosofsky 1988) suggest that it is a continuous function of Euclidean distance in the conceptual spaces.

[5] For some ideas on how verbs may be analysed, see Gärdenfors (to appear).

[6] It should be noted that conceptual spaces are used to represent relations between different concepts, which is more general than the use of space in representing single image schemas within cognitive semantics.

Consequently, we will assume, as a first approximation, that conceptual spaces can be modeled as *Euclidean spaces*. However, the general ideas may be extended to other metric structures (see e.g. Johannesson (2002)).

Following Gärdenfors (2000), we define *concepts* as *regions* of a conceptual space. Two properties of such regions are worth mentioning here. First, as long as given concepts are closed and bounded regions of Euclidean conceptual spaces, they acquire (by a corollary of the classical Bolzano-Weierstrass theorem) one more crucial topological property: *compactness*. One intuition underlying the compactness topological property is that it provides "enough" points that are near to a set – this proves to be a crucial property when fixpoints have to be defined. Furthermore, compactness allows approximating the whole space through a finite number of points, another property that will turn out to be fundamental in what follows.

Second, Gärdenfors (2000) proposed that concepts should be modeled as *convex regions* of a conceptual space. While convexity may seem a strong assumption, it is a remarkably regular property of many conceptual representations grounded in perception (e.g., color, taste, vowels) (Jäger 2007). The main argument for convexity in Gärdenfors (2000) was that it facilitates the *learnability* of concepts. Here we will argue that convexity is also crucial for assuring the *effectiveness of communication*. Although we will in this paper take individual concepts as given, clearly learnability and effectiveness of communication do interact in complementary ways in the process of acquiring individuals concepts.

There are interesting connections between analyzing concepts as convex regions and the *prototype theory* developed by Rosch and her collaborators (see, for example, Rosch 1975, 1978, Mervis and Rosch 1981, Lakoff 1987). When concepts are defined as convex regions of a conceptual space, prototype effects are indeed to be expected. In a convex region one can describe positions as being more or less central. In particular, in a Euclidean space one can calculate the centre of gravity of a region.

It is possible to argue in the converse direction too and show that if prototype theory is adopted, then the representation of concepts as convex regions is to be expected. Assume that some quality dimensions of a conceptual space are given, for example the dimensions of color space, and that we want to decompose it into a number of categories, for example color concepts.[7] If we start from a set of prototypes $p_1, ..., p_n$ of the concepts, for example the focal colors, then these should be the central points in the concepts they represent. The information about prototypes can be used to generate concepts by stipulating that any point p belongs to the same concept as the *closest* prototype p_i. It can be shown that this rule will generate a decomposition of the space – the so-called *Voronoi tessellation*. An illustration of the Voronoi tessellation is given in figure 2. The illustration is two-dimensional, but the tessellation can be extended to an arbitrary number of dimensions.

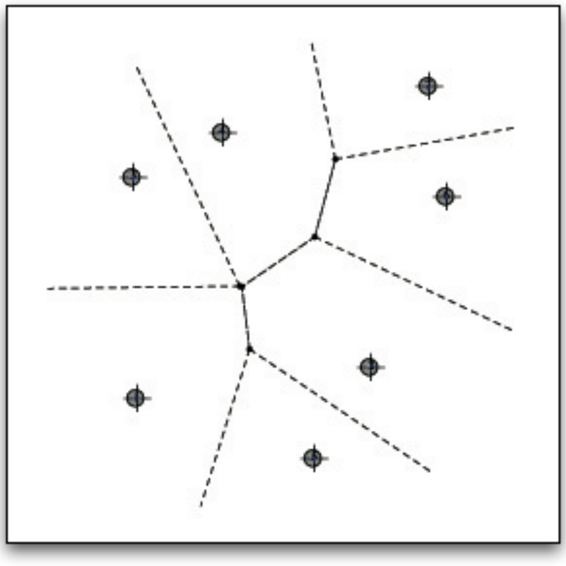

Figure 2. Voronoi tessellation of the plane into convex sets.

[7] Since borders of neighboring regions simultaneously belong to each such region, we call it a decomposition and not a partitioning. This preserves the compactness of the regions.

The basic assumption is that the most typical meaning of a word or a linguistic expression is the prototype of the convex region assigned to the word.[8] This mechanism is a very central principle in connecting the continuity of mental spaces and the discreteness of language.

A crucial property of the Voronoi tessellation of a conceptual space is that it always results in a decomposition of the space into *convex* regions (see Okabe, Boots and Sugihara 1992). In this way, the Voronoi tessellation provides a constructive geometric answer to how a similarity measure together with a set of prototypes determine a set of categories.

The Voronoi diagram has a dual, the so-called Delaunay triangulation, which will turn out to be useful in the sequel (Okabe, Boots and Sugihara 1992). The Delaunay triangulation is obtained by connecting two prototypes of cells that share a side by a line segment. Barring special cases, this procedure will result in a triangulation of the space (in the special cases, it can easily be extended to a triangulation). An important property of the triangulation is that contiguous prototypes will be connected. Furthermore, triangulations of convex sets play a special role in the approximation of continuous functions, as we will show later.

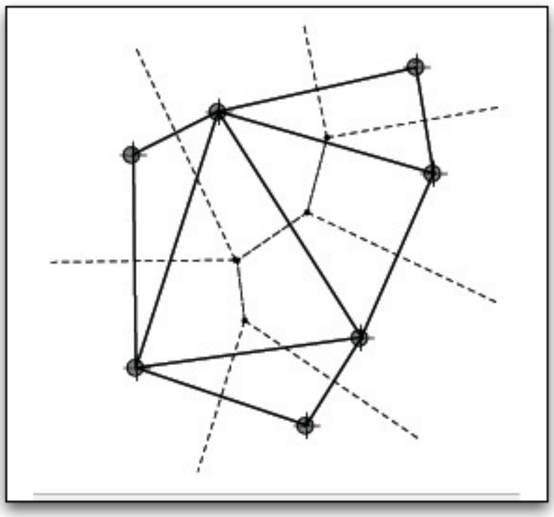

Figure 3. Delaunay triangulation (continuous lines) relative to a Voronoi diagram (dashed lines).

[8] In Gärdenfors (2000), it is argued that the region assigned to a linguistic expression may not be constant, but is in general dependent on the *context* of the use of the expression.

Euclidean metrics, compactness and convexity set the stage for our fixpoint argument. But before getting there, a last point must be made briefly. A basic tenet of cognitive semantics is that language can preserve the spatial structure of concepts. One way to express this is that language can preserve the neighborhood relations among points of conceptual spaces. In topology, a neighborhood preserving function is nothing but a continuous function. In other words, assuming that language can preserve neighborhood relations of conceptual spaces implies assuming that language can establish a continuous mapping between mental spaces of different individuals – and, as we shall see, a continuous mapping of the product space of individual mental spaces on itself. While this continuity assumption may seem extreme, it basically says that natural language must have enough plasticity to map neighborhoods of points in a conceptual space on neighborhoods of points in another conceptual space. Furthermore, we will show below that this assumption can be relaxed to assume that such continuous mappings can be suitably approximated.

3. Fixpoints

3.1 Existence

In our previous example concerning pointing, we assumed that the communicators shared more or less the same mental space – in these cases it was visual space. However, in the general setting, individuals will have different mental spaces. For simplicity let us assume that there are only two individuals with mental spaces C_1 and C_2, which we assume to be convex and compact. If I communicate with you, I alter your state of mind and your reaction will change my state of mind.[9] So communication can be described with the help of "semantic reaction functions" from (x_1,x_2) to (y_1,y_2) in the product space $C = C_1 \times C_2$. We assume that reaction functions are continuous, that is, small changes in the communication will result in small changes of the reaction. In section 3.3, we will provide an explicit example of this

[9] Note that at this stage, we are not modeling communication as such, only its effects on the mental spaces.

general mechanism. A fixpoint is now a point (x_1^*,x_2^*) where nobody changes his own state of mind.

As an elementary example, we can take the case of pointing and let C_1 and C_2 be the visual spaces of two individuals 1 and 2. In the simplest case, the reaction function starts from (x_1,x_2), where x_1 is the point to which 1 is pointing and x_2 is the current position of 1, which 2 is attending to see the pointing direction of 1. The semantic reaction function will map (x_1,x_2) to (x_1,x_1), which means that 2 is following 1's line of gaze to the point x_1. The resulting fixpoint is the one in which 1 and 2 attend to the same object.

Now all ingredients are there, and we can simply remind you of one of the most fundamental results of analysis, Brouwer's (1910) theorem: each continuous map of a convex, compact set on itself has at least one fixpoint. In the present context, the continuous map we are concerned is the semantic reaction function mapping the product space C onto C. In general, C can be the product of several individual spaces (and not just two). Figure 4a illustrates the fixpoint theorem for a function of a one-dimensional space in itself, and figure 4b shows the necessity of the continuity assumption.

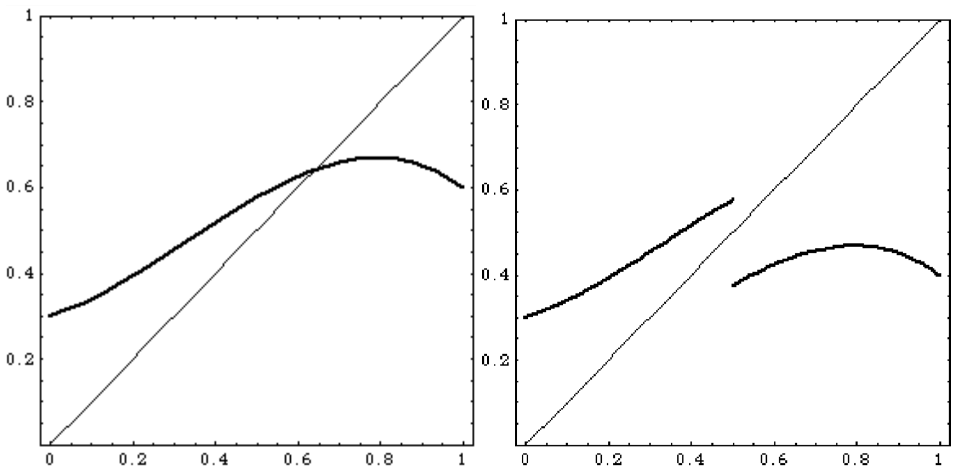

(a) (b)

Figure 4. (a) The fixpoint illustrated for a one-dimensional space

(b) Fixpoints may not exist if the function is not continuous.

This result basically tells us that, no matter what is the content of individual mental representations, provided that such representations are "well shaped" and that language is plastic enough to preserve the spatial structure of concepts, there will always be at least one point representing a "meeting of minds."

So far we have not mentioned the role of language in how fixpoints are reached. Since we are not telepathic, the mapping between individual conceptual spaces must be *mediated*. Language is the primary mediator (but also gestures and other visual tools can be used). Using language means that the speaker maps his mental space on some linguistic entities (from a language L) and that the hearer in turn maps these expressions on his mental space. Communicating linguistically between two individuals, with a product mental space $C_1 \times C_2$, is a composition of a function from C_1 to L and a function from L to C_2. This composition results in a modification of C_2, that is, a change of the hearer's mind. To put it simply, a linguistic message results in a change from (x_1,x_2) to (x_1,y_2) in the product space.[10] Similarly, when the hearer responds the mapping from C_1 via L to C_2 results in a move from (x_1,y_2) to (y_1,y_2) in the product space.

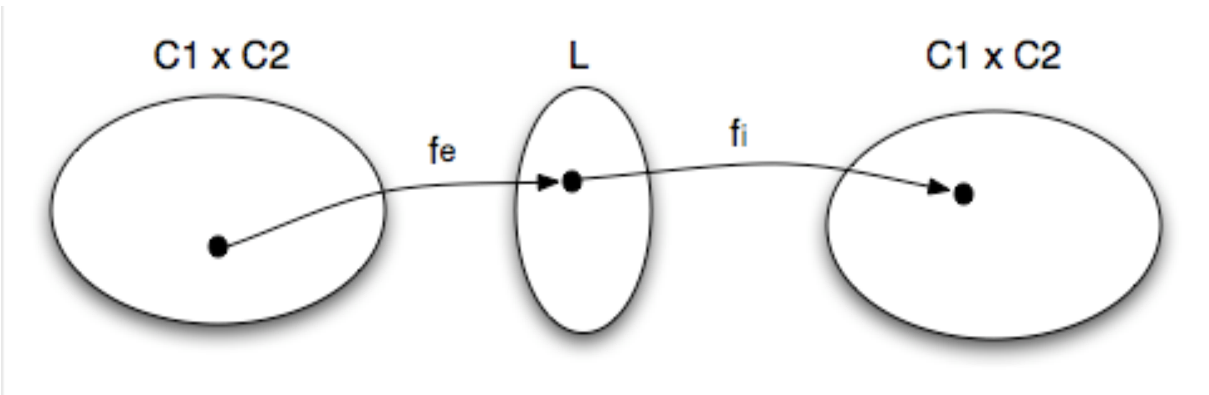

Figure 5: A semantic reaction function maps points in C1×C2 to C1×C2 via L (f_e is the expression function and f_i is the interpretation function).

[10] Actually x_1 is altered too, since the expectations of the speaker changes as well, in particular if individual 1 has a "theory of mind" and can predict the listeners reactions.

Brouwer's result depends on the reaction function being continuous. What is the meaning of continuity in the communication context? It is important to remember that the mental spaces of the communicators are based on similarity that provides them with the metric structure. Thus, in mental spaces "close to" means "similar to". Continuity has a well-known neighborhood-preserving property, which in our framework becomes similarity preserving. In brief, langue is presumed to preserve similarity in mental spaces.

The mapping from C_1 to C_2 need not go via a discrete language (verbal or signed), but other kinds of communicative means can be utilized, such as gestures, mimics and other visual tools. Using these means, we can actually *construct* continuous functions between the spaces. How the specific shape of such continuous function is construed will depend on a variety of pragmatic factors.

After this short and very informal mathematical detour, our central claim should become apparent: Whenever the facility to reach a meeting of minds matters, convex mental representations provide a background over which language can deploy most of its power. We see this as an indirect explanation of why concepts are in general convex. Please note that we are not claiming that convex representations are "faithful" representations of the world – we just claim that since they are effective one should find them quite widespread. In fact, our claim implies that one should expect to find convex representations even in cases in which they are biased representations of the world: Seeing a non-convex world with convex spectacles might be a peculiar bias arising from selective pressures towards effective communication.

3.2 Approximations of continuous functions

Brouwer's theorem shows the existence of fixpoints for any continuous function mapping a compact and convex space on itself. Cognitively, such a function is difficult to manage in terms of requirements on memory and communication systems. In particular, our claim that the semantic reaction function is continuous seems at odds with the discreteness of language resources.

To establish that there is no real conflict between the geometric nature of meanings and the discrete one of language (in particular of lexical resources), we will now briefly resort to a fundamental result of algebraic topology, showing that any continuous function between two Euclidean spaces can be approximated by a mapping between the vertices of some appropriate triangulation of the spaces. The result is a great economy in the cognitive resources needed to memorize and process such function. As we will show, the approximation may boil down to remembering and communicating about the prototypes of a Voronoi decomposition of the space. Such an approximation can thus serve as a bridge between the discreteness of language and the continuity of the (semantic) reaction function.

The basic idea can be stated quite simply, but requires some preliminary definitions.[11] Let X be a convex compact set in a Euclidean space C . A triangulation K of X is a decomposition of such set in a finite set of simplexes, where a simplex is the set of convex combinations of n independent points in some m-dimensional Euclidean space (basically, it is an m–dimensional generalization of a triangle). The triangulation further requires that two simplexes meet at (at maximum) one face or edge. The combinatorial structure thus generated constitutes what is called a geometric simplicial complex. While a geometric simplicial complex is not by itself a topological space, its points, topologized as a subspace of C, are a topological space, the polyhedron |K|. A *simplicial map* f: |K|→|L| between two polyhedra |K| and |L| is a function that maps vertices of |K| to vertices of |L|, and preserves simplexes - in other words, if a^0, a^1, ... a^n are the vertices a simplex in K, then $f(a^0)$, $f(a^1)$, ... $f(a^n)$ are vertices of a simplex in L (notice that $f(a^0)$, $f(a^1)$, ... $f(a^n)$ need not be all different points – repetition is allowed). Fig. 6a provides a simple illustration. Furthermore, it is required that if x is a convex combination of

[11] Here we only sketch the essential concepts. Any algebraic topology handbook will provide a more detailed treatment of the subject. We follow the presentation in Maunder (1980).

a^0, a^1, ... a^n, then f(x) must be a convex combination of f(a^0), f(a^1), ... f(a^n). It can be shown that f is continuous.

Clearly a simplicial map is a "simplex preserving" map. Yet, there is another very important property of simplicial maps. Given two convex, compact sets X and Y and a continuous function g:X→Y*,* there will always be a simplicial map *f* that approximates *g*, provided that X and Y can be triangulated at a sufficiently fine grain. By "*f* approximates *g*" it is meant that *f* is homotopic to *g*, or in plain words that *g* can be obtained from a continuous deformation of *f* This result is known as the *simplicial approximation theorem*. In fact, the theorem tells that any continuous map can be approximated by a piecewise linear map – fig. 6b provides an elementary example in which both X and Y are one-dimensional sets. An important fact is that the simplicial map *f* will preserve the fixpoint properties of *g*, and can be actually used to approximate the fixpoints of *g* (see once more fig. 6b for an illustration).

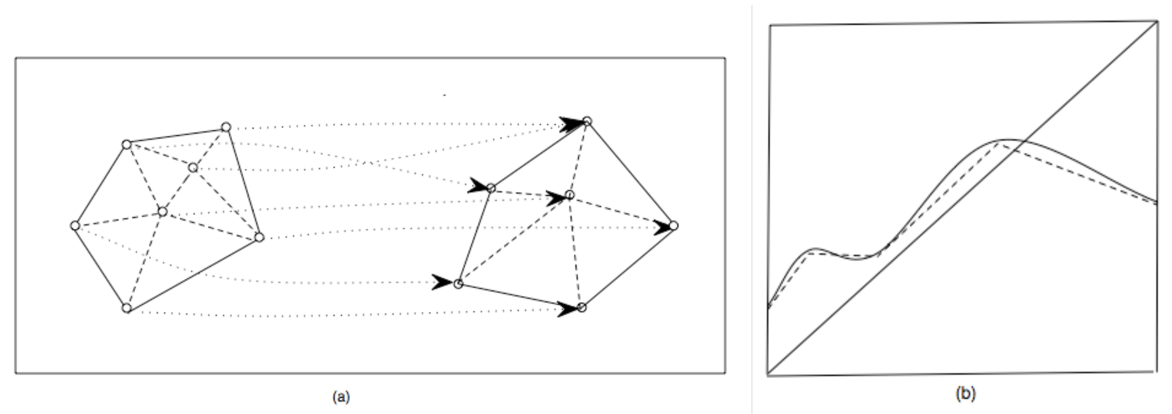

Figure 6: (a) mapping triangulation. (b) a piecewise linear approximation of a one-dimensional function.

As we have already seen, Voronoi tessellations provide a simple model of how categorization subdivides a conceptual space into convex sets. We have also seen that the Voronoi diagram has a dual, which is a set of triangles generated by joining contiguous prototypes, that is, the Delaunay triangulation. This suggests that prototypes generate a basic triangulation of conceptual spaces in which they play the role of simplicial vertices. A natural interpretation is that the prototypes can provide the building blocks of a simplicial approximation of a continuous map between mental spaces. The correspondence between prototypes and words (or other lexical elements) then explains how language can serve as a mediator between conceptual spaces, approximating a continuous function.

An important mathematical proviso concerning simplicial approximation is that, as seen above, it is not always possible to achieve such an approximation between two triangulated spaces, if it is not possible to further triangulate such spaces. In other words, the grain of the triangulation of the spaces may be insufficient to grant a simplicial approximation. Interestingly enough, human cognitive systems of categories have different levels of granularity, corresponding to different levels of prototypes (Rosch 1975, 1978). Thus moving across levels of categorization may ensure that finer triangulations can be constructed in order to achieve a simplicial approximation. Furthermore, just like categories can be refined "locally," it is not necessary to further triangulate all simplexes of a simplicial complex to achieve the required degree of decomposition. In fact, given a complex K, one can leave unchanged a subcomplex P and further triangulate the remaining subcomplex – generating what is called a subdivision relative to P.[12]

The upshot is that the mechanisms of language and linguistic categorization are sufficient to approximate continuity with economic means of discretization. It should be noted that the convexity of spaces plays two important roles here: It ensures triangulability and it allows reconstructing the behavior of approximated function as convex combination of the values of the approximated function in the correspondences of vertices (see figure 6).

3.3 Example: A coordination language game

Brouwer's theorem and the extension by simplicial approximations provide us with an existence result that guarantees that an appropriate meeting of minds can be found among a set of communicators that have convex and compact mental representations of meaning.

[12] See Maunder (1980), section 2.5.7.

However, the results do not in themselves say very much about the contents of the fixpoint or how it can be reached.

Jäger and van Rooij (2007) provide an example of how a meeting of minds functions. Their domain is the color space and the problem they approach is how a common meaning for color terms can develop in a communication game. In their example, there are only two players: s (signaler) and r (receiver). Jäger and van Rooij assume that the two players have a common conceptual space C for color. They define the space as a "continuous space" but from their following claims, it clearly must be a compact and convex space, such as a color circle or a color spindle. There is also a fixed and finite set M of n messages that the signaler can convey to the receiver. The color space C can also be interpreted as a state space from which Nature draws points according to some continuous distribution p. The signaler can choose a *decomposition* S of the space C in n subsets assigning to each color a unique message. The receiver can choose where to locate n points, corresponding to the meaning assigned to each of the n messages by the signaler.

The goal of the communication game is to maximize the average similarity between the intention of the signaler and the interpretation of the receiver. The communication game unfolds as follows: Nature chooses some point in the color space, according to some fixed probability distribution. The signaler s knows the choice of nature, but the receiver r does not. Then s is allowed to send one of the messages to r. The receiver r in turn picks a point in the color space. In the game, s and r maximize utility if they maximize the similarity between nature's choice and r's "interpretation." Here is it only assumed that the similarity is a monotonically decreasing function of the distance in the color space between nature's choice and r's choice.

A Nash equilibrium of the game is a pair (S*, R*), where S* is the sender's decomposition (in n subsets) of C and R* is the responder's n-tuple of points of C, such that both are a best response to each other. Jäger and van Rooij (2007) show how to compute the best response functions for each player. The central result of their paper can be restated by saying that if the color space is convex and compact and the probability and similarity functions are continuous, then there exists a Nash equilibrium, and it corresponds to a common Voronoi tessellation of the color space (which results in convex subsets).

They also show how, in a simplified evolutionary simulation of the game, discrete approximations of convex color regions can emerge as the evolutionary stable solutions of the game. Jäger and van Rooij's model is also interesting because it provides an illustration of how a discrete system of signs (there are only n signs in their communication game) can give rise to approximations of continuous functions mapping agents' mental representations on themselves. In their example, signs define an array of locations in the color space, and the "best response function" of s and r continuously maps configurations of such an array of points as responses to decompositions of C, and vice versa. In this language game, "language" has to be plastic enough to grant the continuity of the best response function, and the meaning space C must have enough topological structure to afford the existence of fixpoints. Language plasticity is given here by the possibility to continuously deform the decomposition S and the location of the points of R.

Notice that adding new signs would only involve *local* changes in the Voronoi tessellation. In other words, you do not have to revise all linguistic meanings each time you learn a new word.

4. Compositionality

4.1 Direct composition

A fundamental semantic property is that of *compositionality*. On a general level, compositionality directly emerges from our framework of space and functions. To give a trivial example, the meaning of "blue rectangle" is defined as the region which is the Cartesian *product* of the "blue" region of color space and the "rectangle" region of shape space (which we leave undefined here[13]). A noteworthy property is that the product of compact and convex sets is again a compact and convex set. Thus the structural properties of conceptual spaces are preserved under this basic semantic composition operator.

This product construction not only preserves topological properties but also the continuity of functions, in the sense that if functions f: A→X and g: B→Y are both continuous, then the product function h = (f,g): A×B→X×Y is continuous. Furthermore, the composition of continuous functions (g∘f) is again continuous. This allows us to concatenate functions preserving their basic properties.[14]

Conversely, the "blue rectangle" conceptual region can be decomposed into its generating regions "blue" and "rectangle" via *projection* (which in turn is a continuous function) from the product space to its component spaces. Once again, the compactness and convexity of "blue rectangle" region are preserved under projection.

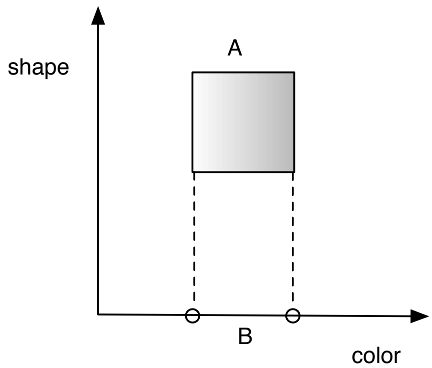

Figure 7: Projection on the color dimension.

Recursively, one can create composite concepts that preserve the basic topological properties of conceptual spaces, thus constructing ever richer concepts. On the other hand, there is a

[13] For an analysis, se Gärdenfors (2000), section 3.10.1.
[14] Although in a totally different sprit, Lewis (1970) uses compositionality of functions to analyse various linguistic categories.

lower bound to decomposition via projection. Either the projection is one-dimensional, where no further projection makes sense, or the projection is an integral set of dimensions, which we will call a *domain*.[15] For example, an object cannot be given a hue without also giving it brightness value. A domain is indecomposable in the sense that it cannot be reconstructed as the product of lower dimensional projections.

The proposed analysis of compositionality of meanings is not exactly the same as the classical Fregean notion. Traditionally, compositionality is defined as composing the meaning of words or expressions, while our analysis is generated from the composition of domains and functions. Since each domain is associated with a class of words, e.g. the class of color words, composing domains generates a composite conceptual space. From this the meanings of the composite expressions can be located as regions of the composite space.

This general presentation of compositionality implicitly assumes that the domains of the product constructions are independent. However, the reality of linguistic usage shows that that the spaces associated with composite expression are not totally independent, but some preprocessing must take place before they can be properly composed. As an example of some types of preprocessing, we will consider some cases of modifier-head composition.

4.2 Modifier – head composition

In the simplest cases, such as “blue rectangle”, where “blue” is the modifier and “rectangle” is the head, the two associated domains color and shape can be assumed to be independent. However, this is a rare case in actual language use. In general, our knowledge of the space associated with the head may affect our representation of the modifier. Thus white wine is not white, a large squirrel is not a large animal and a thick forest does not compare to thick hair.

[15] The notion of integral dimensions and their connection to domains are presented in Gärdenfors (2000), section 1.8.

In all those cases, there seems to be some preprocessing of the representation of the modifier space to adapt it to our knowledge of the head space.[16]

As an example of that some properties cannot be defined independently of other properties, consider ”tall.” This property is connected to the height dimension, but cannot be identified with a region in this dimension. To see the difficulty, note that a chihuahua is a dog, but a tall chihuahua is not a tall dog. This property presumes some *contrast class* given by some other property, since things are not tall in themselves but only in relation to some given class of things. Tallness itself is determined with the aid of the height dimension. For a given contrast class Y, say the class of dogs, the region H(Y) of possible heights of the objects in Y can be determined. An object can then be said to be a tall Y if it belongs to the “upper” part of the region H(Y).

For a contrast class such as skin color, one can map out the possible colors on the color spindle. This mapping will determine a subset of the full color space. Now, if the subset is *completed to a space with the same geometry* as the full color space, one obtains a picture that looks like figure 8.

In this smaller spindle, the color words are then used in the same way as in the full space, even if the hues of the color in the smaller space don't match the hues of the complete space. Thus, “white” is used about the lightest forms of skin, even though white skin is pinkish, “black” refers to the darkest form of skin, even though black skin is brown, etc.

[16] It is not the head, but rather the “modifier” that is modified before the composition of meanings takes place.

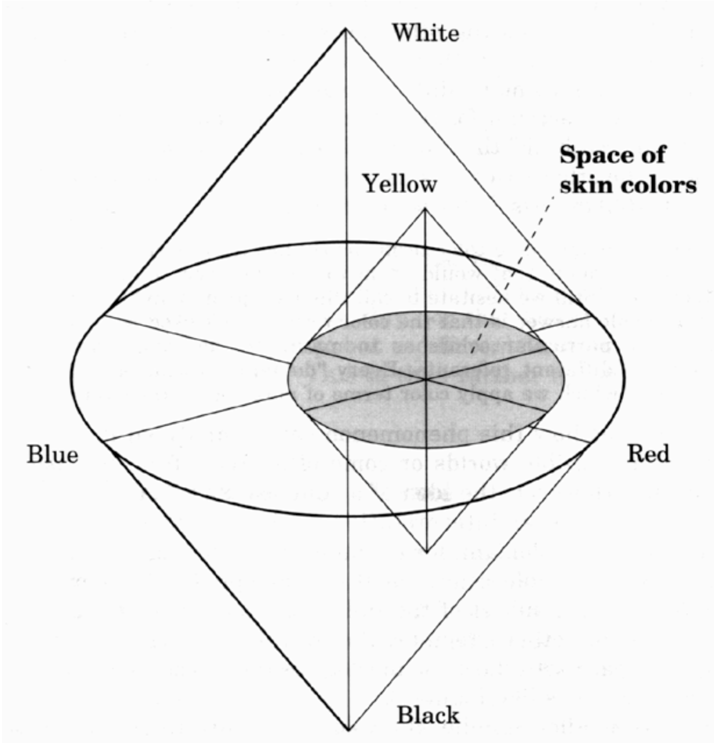

Figure 8: The subspace of skin colors embedded in the full color spindle (from Gärdenfors 2000, p. 121).

Once the head and modifier spaces are compact and convex regions of metric spaces, there always exists a way of rescaling the distances of the modifier space to fit with the constraints of the head space in a one-one correspondence. In this way all color words will be available to characterize the color of skins. The concept of *gauge* (also known as Minkovski functional) provides a natural bridge to model such contextual rescaling effect (Berge 1997). A gauge of a convex set (with an interior point 0 taken as the origo) is a generalized numeric function *j* defined as:

$j(x) = inf\{t: t>0, x \in tC\}$ if $x \in tC$ for at least one $t>0$ (where tC is C inflated by the factor t);

$j(x) = \infty$ otherwise

It is easily seen that if $x \in C$, $j(x) \leq 1$ (and, in particular $j(x) = 1$ if x is on the boundary of the set).

Consider two convex sets, C and D, both defined within a space X, with a common interior point 0 (taken arbitrarily as the origo). Let *j* and *k* be gauge functions for, respectively, C and D. One can define the following function σ: X→X:

$$y = \sigma(x) = \left(\frac{j(x)}{k(x)}\right)x$$

Such function (called *radial projection*) establishes a correspondence between each point of C and a corresponding point in D:

$$k(y) = k\left[\frac{j(x)}{k(x)}\right]x = \frac{j(x)}{k(x)}k(x) = j(x)$$

Figure 9 shows an example of correspondence between points on the boundaries (x_0 and y_0), and between points in the interior (x and y).

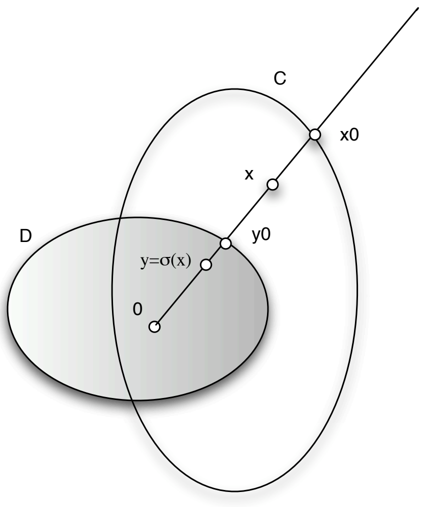

Figure 9: Radial projection.

It can be shown that radial projection establishes an homeomorphism between two convex sets – as long as two sets are convex and compact and have a common interior point, such homeomorphism always exists (Berge 1997, p. 167).

This construction allows us to formulate a general principle: If the region of the space representing the head contains a point that is shared with the space representing the modifier, this point can be taken as an origo of a transformation of the modifier space. This example can be generalized. The radial projection tells you how to import structure from other domains – and as long as concepts are convex and compact, such function always exists. Radial projection is a continuous function, so again all transformations preserve neighborhoods!

If the head and modifier only share some dimensions, the modifier (for example in "pet fish") or the head (for example in "stone lion") is projected onto the shared subspace and then expanded into a shared space by inverse projection. For example in *stone lion*, the representation of *stone* includes the property "non-living," while "living" is presumed by many of the domains of *lion*. These domains, like sound, habitat, behavior, etc, can thus not be assigned any region at all. By large, the only domain of *lion* that is compatible with *stone* is the shape domain. Consequently, the meaning of *stone lion* is an object made of stone that has the shape of a lion.

In Gärdenfors (2000, p. 122), the following general rule for the meaning of a composition of a head D and a modifier C was formulated: "The combination CD of two concepts C and D is determined by letting the regions for the domains of C, confined to the contrast class defined by D, replace the corresponding regions for D." We can now see that this principle follows as a consequence of the constructions presented in this section.

4.3 Metaphorical composition

Even if the head and the modifier do not share any dimension, we can still create a mapping between different domains by exploiting the convexity and compactness of the domains. Indeed an important implication of the existence of a radial projection is that any two convex compact spaces can be mapped via homeomorphism. This permits the creation of the *metaphor* effect, that is, the transfer of structure from one domain to another. Once this is made, we are back to the previous modifier action.

As a simple example, let us consider the expression "the peak of a career." The literal meaning of *peak* refers to a structure in physical space, namely the vertically highest point in a horizontally extended (large) object, typically a mountain. This structure thus presumes two spatial dimensions, one horizontal and one vertical (see figure 10a).

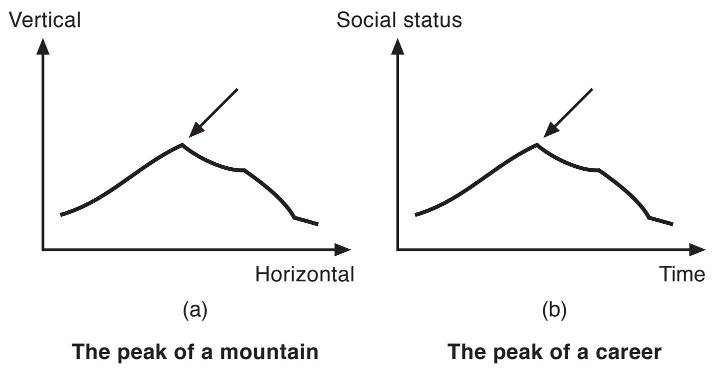

Figure 10: (a) A metaphorical reinterpretation of (b) the literal meaning of "peak" (from Gärdenfors 2000, p. 177).

A career is an abstract entity without location in space. So how can a career have a *peak*? What happens when we metaphorically talk about the peak of a career is that the same geometrical structure is applied to a two-dimensional space that consists of the *time* dimension (of the career), which is mapped on the horizontal spatial dimension, and a dimension of *social status* (or "level of accomplishment", e.g. in athletics) (see figure 10b). The latter dimension is normally conceived of as being vertical: we talk about somebody having a "higher" rank, "climbing" in the hierarchy, etc (see Lakoff and Johnson 1980).

A metaphor does not come alone – it is not only a comparison between two single concepts but it also involves an identification of the structure of two complete domains. Once a domain has been connected to another via metaphor, this connection may serve as a generator for new metaphors based on the same kind of relations (see also Lakoff and Johnson 1980, Tourangeau and Strenberg 1982 and Gärdenfors 2000, Section 5.4). In brief, metaphorical mappings involve whole systems of concepts.

To sum up, we have presented three different ways of composing a modifier and a head. The first one is just plain compositional product construction. The second presupposes an adaptation, in the form of a radial projection (via a gauge function) onto the head space, of the modifier space. The third one, which is involved in metaphor, requires in addition a homeomorphic mapping between spaces. The constructions presume that mental spaces are partitioned into domains. However, in all three cases, the topological properties of the spaces are preserved, which make them ready for further composition.

It should be noted the composition discussed in 4.1 is a special case of that in 4.2, which in turn is a special case of the compositions in 4.3. The compositions in 4.1 do not need to modify existing spaces. The compositions in 4.2 modify spaces that are naturally overlapping. Finally, the metaphors in 4.3 require establishing the homeomorphic correspondences between different spaces. As a consequence, we expect that the three levels of composition will require increasing cognitive processing.

5. How a meeting of minds is achieved in language games

Communication is for some reason. Until now we have postponed the issue of how motivation and the stakes of communication shape the semantic reaction function. Given what has been presented so far, it may seem that the semantics we propose is purely mentalistic. However, reality enters via the *payoffs* of communication. If meaning is not aligned with reality, then the communicators will suffer costs. Reality is what makes not all of our wishes come true, or as Philip K. Dick expresses it: “reality is that which, when you stop believing in it, doesn't go away.”

Reality often enters communication when we use *indexicals*. A paradigmatic case is pointing in combination with “this” or that.” As we mentioned already in the introduction, pointing is a way of coordinating our visual spaces, via a reference to the external world.

Schelling's (1960) type of coordination game has often been used to introduce simple forms of semantic equilibria (see also Lewis (1969)). In such games all equilibria have the same payoff, which implies that selection must be found by resorting to external factors such as conventions or perceptual salience. However, reality may enter communication games by making some equilibria superior in terms of payoffs. To make a Schelling style example, image that your have agreed on a meeting in a very large square. In this case there is an infinity of equilibria corresponding to the different location in the square. But it happens to be raining and there is only one spot in the square protected from rain. In this case, the obvious equilibrium to be selected is the one that lends the payoff being protected from the rain. So the payoff of reality can be used to select equilibria in communication games. In other cases, reality may void meetings of mind. Contracts offer a wealth of examples (Varzi and Warglien, to appear). A paradigmatic example concerns a contract (Sherwood vs. Walzer, 33 N.W. 919 – Michigan 1887) where the covenants had agreed for the sale of a non-pregnant cow, but the cow turned out to be pregnant. The court has considered the contract void, since "The thing sold and bought had in fact no existence. She was sold as a beef creature would be sold; she is in fact a breeding cow, and a valuable one".

In Jäger and van Rooij's (to appear) example, explicit payoffs for the success of communication were introduced and the payoffs generated the best response function that determined the fixpoint. More generally, considering the pragmatic factors that determine the payoff of communication leads naturally into communication games. In such a game, speech acts become moves that modify the conversational playground.

A typical move in a communication game that proposes to restrict the common ground is an *assertion* (Clark 1992, Warglien 2001). In accordance with this, Stalnaker (1979, p. 323), writes: "… the essential effect of an assertion is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed. This effect is avoided only if the assertion is rejectied." If your countermove is accepting the assertion, then we both contract our mental spaces (expnd our belief states). If, on the other hand, you reject the assertion, then another move will have to be attempted (for examples, see Clark (1992)).

A large portion of pragmatics has been concerned with language games – from Wittgenstein (1953) to Clark (1992) and further. Most of this literature assumes that any conversation will start from a strong set of common presuppositions (for example, sharing a set of possible worlds pace Stalnaker (1979) or conversational scores as in Lewis (1979). In contrast, the semantics presented here is based on *products* of mental spaces and not on intersection of spaces of possible world. This implies that we do not have to presuppose that the participants share mental spaces.

Another assumption of the traditional approach is that sentences affect the common ground through a reduction of the possibilities left open (Heim (1983), Clark (1992), Stalnaker (1979)). For example, Stalnaker assumes that the goal of a conversation is to reduce the common ground and converge to a smaller set of possible worlds. On our approach, we do not need the assumption of a strong shared representation. As long as the mental spaces of the participants satisfy our topological properties, the structure of these spaces together with contraction guarantees convergence. More precisely, if the product space of the mental spaces C is compact and the reaction function f: C→C is continuous and such that:

> If for any n, $f(C) \supseteq f^2(C) \supseteq \ldots \supseteq f^n(C)$ (repeated applications of f to its range generate a decreasing sequence) and for each x, $f(x) \neq \emptyset$, then there exists a compact subset K of the product space such that $f(K) = K$.[17]

[17] The proof immediately follows from theorem 8 in Berge (1997), p. 113.

Intuitively, if the communication game constantly contracts the mental spaces of the participants, convergence is unavoidable (given compactness). The typical conversational rules, as formulated by Grice (1975), are communicational institutions that help maintaining the contractive nature of communication and thus ensure convergence. In brief, compactness plus Grice's maxims result in convergence of communication. The condition that repeated applications of f to its range generate a decreasing sequence corresponds to the maxim of quantity (always be informative) and the condition that $f(x) \neq \emptyset$ corresponds to the maxim of quality (never say the false).

A second advantage of our model is that it allows us to consider a wide variety of communicative interaction types that correspond to different game types. For example, we want to distinguish between coordination and negotiation games. In coordination games, such as the Jäger and van Rooij color game, the participants have common interests, while in negotiation games the participants have an interest to reach a common agreement, but they have diverging interests on which agreement to reach.

In a semantic negotiation game, you have an interest to agree on a meaning, but you want to agree on different things. This is like in a contractual negotiation, where we have a surplus to divide. We may have a partial overlap in our representations of a certain concept, but in order to reach an agreement, for example in a contract, we need to negotiate a sufficiently common meaning. However, the process is obviously fallible. Once more there are many examples of contractual breaches that originate from different meanings associated with agreed contractual terms. In a famous court case, the contractors could not agree on the meaning of "chicken" that was the good to be delivered (FrigalimentImp. vs. B.N.S. Int'l Sales, 190 F. Supp 116 – S.D.N.Y. 1960).

If the mental spaces of the communicators are widely diverging, more radical communication methods must be applied in order to create a meeting of minds. This is when metaphors can be powerful tools. By applying a metaphor that exploits a domain that is shared by the communicators, the speaker can convey information about a domain that has no or only a vague correspondence in the spaces of the others. For instance, if you want to express an emotional experience that goes beyond the experiences of your interlocutors, a metaphorical description is often the only available resource.

6. Conclusion

This article presents a novel semantic theory based on the meeting of minds. It puts more structure into communication games by exploiting the topological and geometric structures of individual mental spaces, which are here modeled in terms of conceptual spaces. This approach emphasizes the shapes of representations instead of their contents. An advantage of our approach, in contrast to cognitive semantics, is that we do not assume that the communicators share mental spaces. This makes it possible for us to explain how people can misunderstand each other, but still avoid losing contact.

Putnam (1975) has criticized cognitive approaches to semantics, concluding that "meanings ain't in the head" and that reality thus must be a component in any reasonable semantic theory. To counter Putnam's argument, it can be said that our socio-cognitive framework shows that meanings are in the *heads* of language users when their minds meet.[18] Since our minds can meet over total fantasies, reality need not be a component for meaning to arise and for communication to succeed. On our approach, realism enters the picture when there are stakes to the communication, in the sense that the success depends on the outcome of some communication game.

Another advantage of our approach is that we can establish a connection between the discreteness of language resources and the continuity of thought. We have argued that language, in a broad sense including pragmatics and context, contains mechanisms to preserve

[18] See also Gärdenfors (2000), section 5.7.

similarity of meanings. We do this by showing that language, with its different resources, is an efficient way of approximating the continuity of semantic mappings. As a first study, we have focused on compositionality, but there are many other aspects of how linguistic structures can preserve similarity that should be investigated. We have even argued that semantics, in the form of a meeting of minds, can exist without language, for example in communicating by pointing. This is a great advantage of our approach, because it unites the different forms of communication and does not treat language as an exclusive carrier of meaning.

Finally, it should be emphasized that our presentation here is just a framework, not a model. But our framework lends itself naturally to modeling. We have only in very general term considered linguistic structures, in particular some forms of composition applied mostly to noun phrases. A next step would be to connect different linguistic categories to particular mathematical operations in our framework – something that might require more specific linguistic competence than ours. Some specific linguistic phenomena may also be modeled by weaker assumptions than we have used in this paper, for example, generalized convexity and purely ordinal topologies.

Acknowledgements

We thank Michael D. Cohen, Vince Crawford, Jans Erik Fenstad, Gerhard Jaeger, Paolo Pellizzari, Johan van Benthem, Joost Zwarts, and the participants at seminars at SALC (Università di Trento), and Mumbai for helpful discussions and comments.

References

Bates, E. (ed) (1976): *Language and Context. The Acquisition of Pragmatics*, Academic Press, New York.

Berge, C. (1997): *Topological Spaces*, Dover, Mineola (NY).

Brinck, I. (2004): "The pragmatics of imperative and declarative pointing," *Cognitive Science Quarterly 3*, 255-272.

Brinck, I., Gärdenfors, P. (2003): "Co-operation and communication in apes and humans," *Mind and Language 18*, 484-501.

Brouwer, L.E.J (1910): "Über ein eindeutige, stetige Transformation von Flächen in sich", *Mathematische Annalen 69*, 176-180.

Clark, H. (1992): *Arenas of language use*. University of Chicago Press, Chicago.

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

Evans, V. (2006): "Lexical concepts, cognitive models and meaning-construction", *Cognitive Linguistics 17*, 491-534.

Fitting, M. (2002): "Fixpoint semantics for logic programming: A survey", *Theoretical Computer Science 278*, 25-51

Gärdenfors, P (1997): "Does semantics need reality?", in *Does Representation Need Reality?*, Austrian Society of Cognitive Science Technical Report 97-01, Vienna, pp. 113-120.

Gärdenfors, P. (2000): *Conceptual Spaces: The Geometry of Thought*, MIT Press, Cambridge, MA.

Gärdenfors, P. (2003): *How Homo Became Sapiens: On the Evolution of Thinking*, Oxford University Press, Oxford.

Gärdenfors, P. (2007): "Representing actions and functional properties in conceptual spaces", in *Body, Language and Mind, Volume 1: Embodiment,* ed. by T. Ziemke, J. Zlatev and R. M. Frank, Mouton de Gruyter, Berlin, pp. 167-195.

Gärdenfors, P. and Osvath, M. (to appear): "The evolution of anticipatory cognition as a precursor to symbolic communication," to appear in *Evolution of Language: Biolinguistic Approaches*, Cambridge University Press, Cambridge.

Gärdenfors, P. and Warglien, M. (to appear): "The development of semantic space for pointing and verbal communication", to appear in *Conceptual spaces and the construal of spatial meaning. Empirical evidence from human communication,* Oxford University Press.

Goldin-Meadow, S. (2007): "Pointing sets the stage for learning language and creating language", *Child Development 78*, 741-745.

Grice, P. (1975): "Logic and conversation," in *Syntax and Semantics*, Vol. 3, Speech Acts, ed. by Peter Cole and Jerry L. Morgan. New York: Academic Press 1975, 41–58.

Harnad, S. (1990): "The symbol grounding problem", *Physica D. 42*, 335-46.

Heim, I. (1983): "On the projection problem for presuppositions", *Proc. of the West Coast Conference on Formal Linguistic*s, vol II, Stanford, CA, 114-125.

Hopfield, J.J. (1982): "Neural networks and physical systems with emergent collective computational abilities", *Proc. Nat. Acad. of Sci. 79*, 2554–2558.

Holmqvist, K. (1993): *Implementing Cognitive Semantics*, Lund University Cognitive Studies 17, Lund.

Jäger, G. and van Rooij, R. (2007): "Language structure: Psychological and social constraints", *Synthese 159*, 99-130.

Jäger, G. (2007): "The evolution of convex categories", *Linguistics and Philosophy 30*, 551-564.

Johannesson, M. (2002): *Geometric Models of Similarity*, Lund University Cognitive Studies 87, Lund.

Kripke, S. (1975): "Outline of a theory of truth", *Journal of Philosophy 72*, 690-716

Lakoff, G. and Johnson, M. (1980): *Metaphors We Live By*, The University of Chicago Press, Chicago, IL.

Lakoff, G. (1987): *Women, Fire, and Dangerous Things*. The University of Chicago Press, Chicago, IL.

Langacker, R. W. (1986): "An introduction to cognitive grammar," *Cognitive Science 10*, 1-40.

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

Lewis, D. (1969): *Convention*, Harvard University Press, Cambridge, MA.

Lewis, D. (1970): "General semantics", *Synthese 22*, 18-67.

Lewis D. (1979): "Scorekeeping in a language game", *Journal of Philosophical Logic 8*, 339-59.

Maunder C.R.F. (1980): *Algebraic Topology*, Cambridge University Press, Cambridge (UK).

Mervis, C., and Rosch, E. (1981): "Categorization of natural objects", *Ann. Rev. of Psychol.* 32, 89–115

Nosofsky, R. M. (1988a), "Similarity, frequency, and category representations," *Journal of Experimental Psychology: Learning, Memory and Cognition, 14*, 54–65.

Okabe, A., Boots, B., Sugihara, K. (1992*): Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, John Wiley & Sons, New York

Parikh, P. (2000): "Communication, meaning and interpretation", *Linguistics and Philosophy 23*, 185-212.

Putnam, H. (1975), “The meaning of 'meaning',” in Gunderson, K., ed., *Language, Mind and Knowledge*, University of Minnesota Press, Minneapolis, 131-193.

Rosch, E.: (1975) “Cognitive representations of semantic categories”, *Journal of Experimental Psychology: General 104*, 192–233.

Rosch, E. (1978): “Prototype classification and logical classification: The two systems”, in Scholnik, E. (ed.): *New Trends in Cognitive Representation: Challenges to Piaget’s Theory*, 73–86,Lawrence Erlbaum Associates, Hillsdale, NJ.

Schelling, T. (1960): *The Strategy of Conflict*, Harvard University Pres, Cambridge, MA.

Shepard, R. N. (1987), “Toward a universal law of generalization for psychological science,” *Science 237*, 1317–1323.

Skyrms, B. (1998): “Salience and symmetry-breaking in the evolution of convention”, *Law and Philosophy 17*, 411-418.

Stalnaker, R. (1979): “Assertion”. *Syntax and Semantics 9*, 315-332.

Tourangeau, R. and Sternberg, R. J. (1982), “Understanding and appreciating metaphors,” *Cognition 11*, 203-244.

Tomasello, M. (1999): *The Cultural Origins of Human Cognition*, Harvard Unversity Press, Cambridge, MA.

Varzi, A. and Warglien, M. (to appear): “Indeterminate contracts and semantic indeterminacy.”

Warglien, M (2001): “Playing conversation games”. Paper presented at the 2001 Wittgenstein Society Symposium, Kirchberg.

Wittgenstein, L. (1953): *Philosophical Investigations*, Blackwell’s, Oxford.