Self-organisation of conceptual spaces from quality dimensions

Paul Vogt

Abstract This chapter presents a discussion on how conceptual spaces can evolve from a set of quality dimensions, and how these spaces can become shared among a population of cognitive agents. An agent-based simulation of Steels' Talking Heads experiment is presented in which virtual agents construct novel concepts, as well as a shared, simplified language from scratch. Simulations demonstrate that the structure of a conceptual space (i.e. from what quality dimensions it is composed) can evolve in a population of communicating agents. It is argued that the underlying mechanisms involve the following factors: the environment of the agents, their embodiment and cognitive capacities, self-organisation, and cultural transmission.

1 Introduction

Conceptual spaces are constructed from quality dimensions (Gärdenfors, 2000), but how are quality dimensions selected to constitute a particular conceptual space? Is it the result of biological evolution? Or do the conceptual spaces emerge through ontogenetic development? And, if the latter, are they culturally determined and/or constrained through cognition, embodiment, or the ecological niche? I will argue that it is probably a combination of all these factors.

To answer these questions, let me start by briefly recapturing what quality dimensions are and how they constitute conceptual spaces. According to Gärdenfors (2000, p. 6) "the primary function of the quality dimensions is to represent various 'qualities' of objects ... [and] correspond to the different ways stimuli are judged to be similar or different". In visual perception, for instance, these qualities could be feature detectors such as hue, saturation and brightness to represent the conceptual space of colour, edge detectors that may combine to represent a shape, spatial de-

Paul Vogt
Tilburg centre for Cognition and Communication, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands, e-mail: p.a.vogt@uvt.nl

tectors combine to represent spatial locations, etc. It is beyond doubt that many (if not all) of these quality feature detectors are innate and have evolved biologically. It is also arguable that evolution has selected for particular configurations of feature detectors (or quality dimensions) that together form a particular conceptual space, such as the colour space. However, this does not necessarily hold for all possible conceptual spaces.

Take, for example, the space of spatial concepts. Languages have different ways of communicating spatial relations, based on three different frames of reference: relative to the target object (e.g. the box on the left of the tree), intrinsic to the target (e.g. I am in front of the box) or absolute (the box to the north). Often languages have a combination of two or three of these frames of reference, while other languages have only one of these (Majid, Bowerman, Kita, Haun, & Levinson, 2004). There is abundant evidence that the way people categorise and the way they communicate about spatial relations are tightly linked, so the people who speak using a particular frame of reference, also categorise the world that way (Majid et al., 2004; Haun, Rapold, Janzen, & Levinson, 2011). Such a Whorfian account (Whorf, 1956) not only holds for spatial concepts, but also for many other types of concepts (Bowerman & Levinson, 2001). Since the speakers of different languages categorise spatial relations so radically different, it is conceivable that their concepts are represented in conceptual spaces made up of different quality dimensions.

The spatial concepts that people use in their language community is clearly learnt, flexible, and depends to some extent on the physical environment (Haun et al., 2011). Tzeltal speakers, for instance, live in a hilly environment and communicate spatial relations in terms of being uphill or downhill. Now, imagine a Tzeltal speaker moving to the Netherlands where there are no hills. At first, he will have difficulty categorising the world in terms of spatial concepts. Nevertheless, he will be able to distinguish that objects are in different spatial locations. If he learns Dutch, he will learn that spatial relations are communicated in a relative frame of reference using left, right, front, etc. Its concepts would be represented in a different conceptual space constructed from a different set of quality dimensions. If he gets a child in the Netherlands and, as long as they will not leave the Netherlands, this child will learn to categorise the world the way Dutch people do in terms of left, right, etc., as well as the concepts of North, East, etc., but although the child may learn to speak Tzeltal, he will not be able to form the concepts of uphill and downhill. (Imagine Dutch speakers talking about things being uphill or downhill, while there are virtually no hills in the Netherlands.) To really being able to form such concepts, the child will need exposure to a hilly environment.

In this chapter, I will demonstrate how a group of virtual (i.e., simulated) robots can acquire various conceptual spaces from a given set of quality dimensions by developing a set of linguistic conventions from scratch through cultural evolution (I will take the biological evolution of conceptual spaces for granted). Before doing that, I will set the theoretical framework of these simulations in which I argue that the following factors are the driving force behind such a development: the environment (or ecological niche), embodiment (i.e. the physical properties of the agent), cognition (e.g., the way concepts are learnt and represented), self-organisation and cultural transmission. The robotic agents are necessarily abstracted away from human agents, so that the results are not directly generalisable to human cognition. The purpose of this chapter is therefore not to present a biologically plausible model of the formation of conceptual spaces from the basic primitives of quality dimensions, but to illustrate a number of likely mechanisms and properties that could explain how such a development could work.

2 The evolution of conceptual spaces

The theoretical framework is developed from an evolutionary linguistics point of view, because of the tight link between linguistic and conceptual structures (Majid et al., 2004; Bowerman & Levinson, 2001). In particular, the framework will be based on a hypothesised evolutionary transition from holistic protolanguages to more modern compositional languages (Wray, 1998). In order to do that, it is instrumental to define such languages:

Holistic languages are languages in which parts of expressions have no functional relation to any parts of their meanings. For instance, there is no part of the expression "bought the farm" that relates to any part of its meaning "died".

Compositional languages are languages in which parts of expressions do have a functional relation to parts of their meanings and the way they are combined. For instance, the part "John" in "John loves Mary" refers to a guy named John, likewise "loves" and "Mary" have their own distinctive meanings. In addition, this sentence has a different meaning in English when the word-order changes, as in "Mary loves John".

Based on these definitions, it is possible to conceive that if a particular meaning is associated with a holistic utterance, then this meaning could be represented in some N-dimensional conceptual space. However, when the same meaning would be associated with a compositional utterance, then parts of the utterance would be associated with individual concepts c_i, each represented within an n_i-dimensional conceptual space with $n_i \leq N$.

Alison Wray (1998) has argued that protolanguages were essentially holistic in nature and that from these initial stages language has gradually evolved into compositional languages. Although it has been argued that protolanguages were not holistic, but synthetic and instead consisted of multi-word utterances without a particular syntactic structure (Bickerton, 1984; Jackendoff, 2002), let us assume that Wray is correct. (Without justification, I believe that many of the underlying principles presented in this chapter would hold either way.) Then one could ask the question: what evolutionary mechanism(s) caused this transition? The nativist account would be that the population of language users have adapted biologically to learn and produce compositional languages (Pinker & Bloom, 1990). If this occurred through natural selection, this would require that individuals with a particular genetic mutation started using compositional language (at least to some extent), which made them evolutionary more advantageous, thus improving their chances of passing on this mutation, thus increasing the population of individuals using compositional language, etcetera. Although not impossible, biological evolution is a rather slow process that would take quite a number of generations before a mutation is spread among the entire population.

An alternative explanation takes the view that cultural evolution was the driving force behind the transition from holistic protolanguage to compositional language. In this viewpoint, put forward by Wray herself and soon adopted by Simon Kirby and colleagues (Brighton & Kirby, 2001; Kirby, 2001; Kirby, Smith, & Brighton, 2004; Kirby & Hurford, 2002; Kirby, Cornish, & Smith, 2008), the population of language users does not adapt to learn and use compositional languages, but the language adapts itself such that it can be learnt and produced by its users. This is an appealing explanation, not only because language change spreads faster across a population through cultural evolution, but also because when a genetic mutation yields a change in the language that other language users cannot deal with, the mutant language user does not conform to the other users, thus hampering effective communication.

The potential of this cultural evolutionary explanation for this transition has been demonstrated over and over again in computer simulations (Brighton & Kirby, 2001; Kirby, 2001; Kirby et al., 2004; Kirby & Hurford, 2002; Vogt, 2005a) and in psycholinguistic experiments (Garrod, Fay, Rogers, Walker, & Swoboda, 2010; Kalish, Griffiths, & Lewandowsky, 2007; Kirby et al., 2008). The typical approach in these simulations and experiments is based on iterated learning in which the language of one individual is passed on to a learner from a next generation, who in turn passes on the language to the next generation, and so forth. This thus creates a chain of generations of language users who each acquire the language from the previous generation.

The learners in this model are endowed with a learning mechanism that enables them to discover regular patterns in the input (both in speech and semantics) and when a regularity is discovered, a compositional representation can be constructed and used. This is especially useful when a language user wants to communicate a previously unseen meaning that is composed of several concepts for which the user knows words or utterances to express parts of the meaning, but not the whole meaning holistically. Kirby and colleagues have demonstrated that a transition from holistic languages to compositional language occurs when the language is transmitted through a bottleneck where the next generation needs to communicate about previously unseen meanings. The primary reason for this is that a bottleneck makes the transmission of holistic languages unstable, but not for compositional languages, as illustrated in Figure 1.

The abstractions and assumptions made in the iterated learning model, especially in the computational implementations, however, make it hard to generalise the results. For instance, it is typically assumed that each generation has only one individual and that only the individual from the older generation passes on language to the next generation, thus it rests entirely on vertical transmission. Consequently, the researcher has to impose the transmission bottleneck explicitly. In addition, in

Type	$G(n)$	Utterance	$G(n+1)$
Holistic	toma-[redsquare]	toma-[redsquare]	toma-[redsquare]
	tula-[greentriangle]	tula-[greentriangle]	tula-[greentriangle]
	bulo-[greensquare]	bulo-[greensquare]	bulo-[greensquare]
	rino-[redtriangle]		??-[redtriangle]
Compositional	toma-[redsquare]	to*ma*-[red*square*]	toma-[redsquare]
	bulo-[greentriangle]	**bu**lo-[greentriangle]	bulo-[greentriangle]
	buma-[greensquare]	**bu***ma*-[**green***square*]	buma-[greensquare]
	tolo-[redtriangle]		tolo-[redtriangle]

Fig. 1 This figure illustrates why holistic languages (upper part) are unstable when a population of generation $G(n+1)$ only observes three of the four utterances from generation $G(n)$'s language (i.e. word-meaning mappings). In this case, if generation $G(n+1)$ wishes to communicate about meaning [redtriangle], then this generation will have to create a new word. If the language were compositionally structured as in the bottom part of this figure, observing the aligning patterns from only three out of four utterances would allow the next generation to reconstruct the entire previous language. Hence transmitting a compositional language through a bottleneck is evolutionary more stable than transmitting holistic languages.

most computer simulations the semantics are predefined by the researchers, who thus ensure that there are clear decomposable semantic structures.

A more realistic model would assume a population containing many individuals from different generations, who can each pass on parts of the language to other individuals more akin to oblique and cultural transmission. This is important, because the dynamics of cultural evolution in vertical transmission – as in the iterated learning model – is quite different from the dynamics that can be observed in systems pertaining to oblique and horizontal transmission, which are more reminiscent of human cultural evolution (Cavalli-Sforza & Feldman, 1981). These systems allow for cultural traits, such as linguistic entities or memes, to evolve based on neo-Darwinian evolution in which variation, competition and self-organisation of traits play a crucial role (Boyd & Richerson, 2005; Croft, 2002; Mufwene, 2001). One advantage of a transmission system where the offspring can (try to) transmit knowledge to peers or to older generations while they are still learning, is that they will encounter new situations in which they may need to communicate about previously unseen items (cf. Fig. 1). In the iterated learning model such situations only occur after learning has stopped. The system of horizontal and oblique transmission thus provides learners with a natural implicit transmission bottleneck that triggers the emergence of compositionality (Vogt, 2005c).

A downside of predefining the agents' semantics – as is the case in most iterated learning models – is that this removes 1) the role that ontogenetic development of concepts can play in bootstrapping the emergence of compositionality (Vogt, 2006b), and 2) the individual variation in conceptualisation which is a crucial component of neo-Darwinian evolution. Moreover, enabling agents to develop categories/concepts from interacting (i.e. perceiving and acting in) the world, it becomes important to consider by what means the world is perceived and acted in. For example, a researcher should consider what sensors a robot may have. Are these cameras, touch sensors, a compass or a combination of these? And what type of information is filtered from these sensors? All these factors essentially define the agents' embodiment, which in turn defines what qualities the agent can perceive in the world, thus constricting the possible conceptual spaces that can be formed. Although agents with different physiological capacities can learn to communicate effectively – think of blind people, but also robots can do this (de Greeff & Belpaeme, 2011) – the question is to what extent they converge on internal conceptual representations. This question is even relevant for agents having the same bodies, but different experiences in the world.

It should be clear that agents form concepts that reflect the world they engage in – it is impossible for agents who only encounter a flat world to acquire concepts such as uphill or downhill. People who only live in a remote area of the Amazon and who have never visited or seen skyscrapers, will not be able to fully grasp the concept of a skyscraper. This does not only apply to basic concepts, but also to compositions of concepts and the structures thereof. For instance, consider the concept of a cup. A cup can hold many substances (coffee, tea, water, sugar, ...), have various colours (white, blue, orange, ...), shapes, textures, sizes, etc. When there is only one cup present in a particular context, such features are not so important, but when there are multiple cups around these features may become important. The way humans conceptualise the cup in these different situations is hard to tell, but looking at the ways humans refer to a particular cup in different situations suggest we structure our conceptual representation (Brennan & Clark, 1996; Koolen, Gatt, Goudbeek, & Krahmer, 2011). The way concepts are structured depends strongly on the objects' properties and the way we perceive them, which in turn tends to be reflected in the language. I would argue that this goes so far that much of the structure of our engagement in the world (and more particularly in our ecological niche) is reflected in the grammars of our language. Humans tend to manipulate some target in one way or another. This is how we universally behave in the world, and that is what is reflected in most languages spoken across the globe: Most (but not all) languages have linguistic structures in which sentences contain a subject, a verb and an object (Baker, 2003; Evans & Levinson, 2009). Hence, the way we interact with our environment (i.e. our situatedness) and consequently the structure of our environment, as well as our embodiment, influence the way we conceptualise the world. Culture and language are part of our environment and are thus not only manifestations of our conceptualisations, but also shape them.

In the remainder of this chapter, I discuss a model that tries to incorporate the fore-mentioned principles in a simulation in which a population of agents evolve a simple compositional language from scratch in two steps: first a holistic language is formed, second a transition towards a compositional language occurs (Vogt, 2005a, 2005c, 2007). This model combines some components of Kirby's iterated learning model (Kirby, 2001) – language learning and transmission over generations – with Luc Steels' language game model (Steels, 1997, 2003, 2012). This way, grammatical structures and – as part of this – conceptual spaces co-evolve through self-organisation driven by social interactions between agents and the cognitive learning mechanisms of these agents. As the agents are situated in a virtual environment where they are forced to communicate about the objects in the environment, the structure of the environment, as well as the agents' perceptual apparatus, constrain the conceptual structures of the emerging languages. The general principles of this system – especially with regards to the complex adaptive dynamics – are the same as in most of Steels' studies. However, where the formation of grammar in Steels' models relies on a complicated formalisation of cognitive grammars (Steels & Beule, 2006; Steels, 2012), the model presented here relies on a straightforward realisation of alignment-based learning (Zaanen, 2000) in combination with data-oriented parsing (Bod, Sima'an, & Scha, 2003).

3 Language games

The model simulates the Talking Heads experiment (Steels, Kaplan, McIntyre, & Van Looveren, 2002) in which a population of agents play a large number of guessing games –a variant of the language game– to develop a language that allows the population to communicate about their world. This world contains 120 coloured geometrical shapes (12 colours x 10 shapes) and the agents can only perceive the RGB values of the colour and one feature representing the shape. A guessing game is played by two agents: a speaker and a hearer. The aim of the game is for the hearer to guess what the speaker verbally refers to, and – where possible – each individual agent adapts its conceptual and linguistic representations such that the communication becomes more effective. The game consists roughly of the following steps: perception, conceptualisation, production, interpretation and adaptation.

These steps are explained in some detail in the remainder of this section, with a special focus on the emergence of conceptual spaces. It is beyond the scope of this chapter to present all details of the model, and the interested reader is referred to Vogt (2005a, 2005c).

3.1 Perception and conceptualisation

In each guessing game, a number of objects are randomly drawn from the world with a uniform distribution and are 'shown' to the agents as the context of the game. Suppose an agent sees the three objects on the top left of Figure 2: red square, yellow hexagon and purple circle. Using its perceptual apparatus, each object is transformed into a 4-dimensional vector representing the r, g and b values of the RGB colour space and a feature value representing the object's shape s. The red square is thus represented by vector $(1,0,0,1)$, the yellow hexagon by $(1,1,0,0.5)$ and the purple circle by $(1,0,1,0.57)$. These feature vectors represent the raw percepts of the objects.

Each feature of each percept is then categorised with a category from the relevant r, g, b and s quality dimensions. The categories divide each dimension in one or more segments and are represented by a prototypical value, as indicated by a

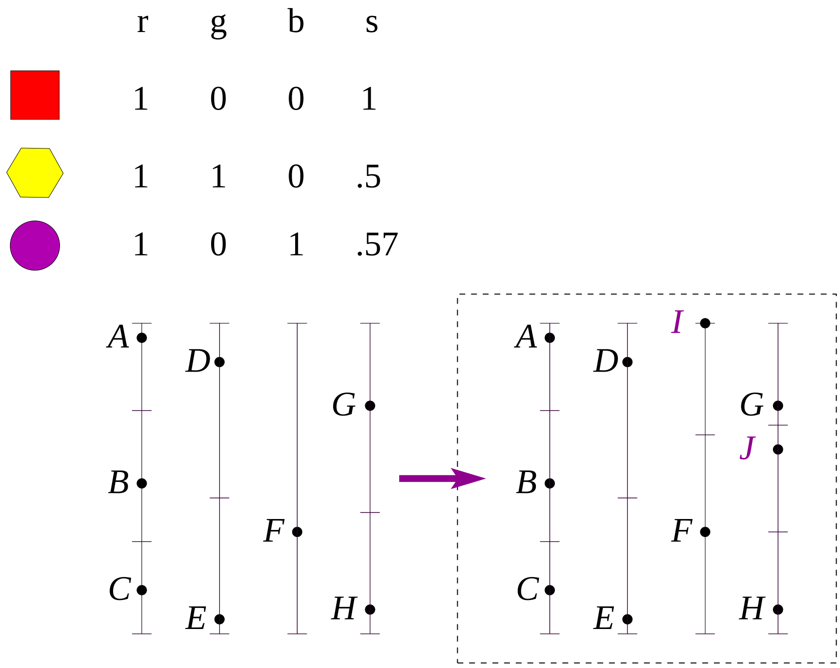

Fig. 2 This figure illustrates the conceptualisation and adaptation withing the discrimination game (see the text for details).

dot in Figure 2. The square would be thus categorised by the set $\{A,E,F,G\}$, the hexagon by $\{A,D,F,G\}$ and the circle by $\{A,E,F,G\}$. Such sets represent the objects' concepts as cubes in the 4-dimensional conceptual space. This is, probably, not a realistic representation of conceptual spaces, but it is a consequence of treating each quality dimension independently to facilitate their selection to be part of different conceptual spaces. More realistic implementations of conceptual spaces that would be applicable have been put forward in, e.g., Steels and Belpaeme (2005); Wellens, Loetzsch, and Steels (2008) and Vogt (2004).

In order to communicate effectively, the agents individually process discrimination games (Steels, 1996). The object of a discrimination game is to obtain a concept that represents an object such that it distinguishes this object from the other objects in the context. If the agent is to conceptualise the hexagon in contrast to the two other objects of Figure 2, the agent is successful as the concept $\{A,D,F,G\}$ is distinctive. If, however, the agent is to discriminate the square or the circle from the other two objects in this context, the agent fails as both objects have the same concept. If this occurs, the agent adapts its categories by adding the feature values of distinguishable dimensions as a prototypical exemplar to the appropriate quality dimensions. For instance, if the agent was trying to distinguish the circle from the other two objects, it would add the categories I and J to the original representation,

1	S $\rightarrow$ greensquare/(0,1,0,1)	0.2
2	S $\rightarrow$ A/rgb B/s	0.8
3	A $\rightarrow$ red/(1,0,0,?)	0.6
4	B $\rightarrow$ triangle/(?,?,?,0)	0.7

Fig. 3 This example grammar contains rules that rewrite a non-terminal into an expression-meaning pair (1, 3 and 4) or into a compositional rule that combines different non-terminals (2). Rule (2) is thus a rule that combines linguistic categories/conceptual spaces A/rgb and B/s (i.e., A relates to the RGB colour space and B to the shape space). For practical reasons concepts are presented as 4-dimensional vectors, where the first 3 dimensions relate to the RGB colour space (rgb) and the 4th relate to the shape feature (s); the question marks are wild-cards and indicate which quality dimension(s) is (are) not part of this conceptual space. Each rule has a rule score that indicates its effectiveness in past guessing games. Only sentences of one or two constituents are allowed in this grammar.

yielding the set of quality dimensions with categories as depicted in the dashed box of Figure 2.

Conceptual spaces in this model can be formed by taking one to four of these quality dimensions together, so there can be four 1-dimensional spaces, six 2-dimensional spaces, three 3-dimensional spaces and one 4-dimensional space. The concepts within each space can be used to represent the basic meanings in the agents' language. This way, conceptual spaces are constructed that could be interpreted as linguistic categories. Initially, the agents will only conceptualise percepts in the 4-dimensional space and associate such concepts with word-forms in a holistic manner. The purpose of this study is to demonstrate how agents can develop conceptual spaces of lower dimensions and use these coherently in language. To understand how this may be achieved it is important to understand how the agents represent, use and learn their language.

3.2 Production and interpretation

Once the agents have categorised the objects in the context, the speaker selects one object at random with a uniform distribution as the topic of the communication. This agent then searches its grammar for ways to produce an expression that conveys the topic's concept. The grammar (Figure 3) is an individual's competence and consists of simple rewrite rules that associate forms with concepts either holistically (e.g., rule 1) or compositionally (e.g., rule 2 combined with rules 3 and 4). The grammar may be redundant in that there may be rules that compete to produce or interpret an expression (cf. Batali, 2002; De Beule & Bergen, 2006; Steels, 2012). The speaker searches for those (compositions of) rules that match the topic's concept and if more than one are found, he selects the rule that has the highest rule score. If the speaker fails to produce an expression this way, a new form is invented as an arbitrary string and is associated with the topic's concept or – if a part of the concept matches some non-terminal rule – with the complement of this concept. For instance, if the speaker would want to produce an utterance expressing a red square $(1,0,0,1)$ and it knows a word for the colour red $(1,0,0,?)$ but not for square, then it invents a new word (e.g., ‘wateva’) to express square $(?,?,?,1)$ and adds this to its grammar.

In turn, the hearer tries to interpret the expression by searching its own grammar for (compositions of) rules that match both the expression and a concept relating to an object in the current context. If there is more than one such rule, the hearer selects the one with the highest score, thus guessing the object intended by the speaker. The hearer then ‘points’ to this object, and if this is the object intended by the speaker, the speaker acknowledges success; otherwise, the speaker points to the topic allowing the hearer to acquire the correct concept referring to the expression.

3.3 Adaptation

If the guessing game was successful, both the speaker and hearer increase the scores of the rules they used and lower the scores of those rules that compete with the used rules. If the game has failed, the scores of used rules are lowered and the hearer acquires the proper association between the heard expression and the topic’s concept. To this end, the hearer tries the following three steps until one step has succeeded:

1. If a part of the expression can be interpreted with a part of the topic’s concept, the rest of the expression is associated with the complement of the concept. For instance, if the hearer of the grammar shown in Figure 6 hears the expression “redcircle” referring to the concept (1,0,0,.5), the part “red”-(1,0,0,?) can be interpreted, so the hearer adds rule B$\rightarrow$circle/(?,?,?,.5) to its grammar.
2. If the above failed, the hearer searches its memory, where it stores all heard or produced expression-concept pairs, to see if there are instances that are partly similar to the expression-concept pair just heard. If some similarity can be found, the hearer will break-up the expression-concept pairs containing these similarities following certain heuristics, thus forming new compositional rules. Suppose, for instance, the hearer had previously heard the expression-concept pair “greensquare”-(0,1,0,1) and now hears “yellowsquare”-(1,1,0,1). The hearer can then break up these pairs based on the similarity “square”-(?,1,0,1), thus forming rules S$\rightarrow$C/r D/gbs, C$\rightarrow$green/(0,?,?,?), C$\rightarrow$yellow/(1,?,?,?) and D$\rightarrow$square-(?,1,0,1). Note that this is not the ideal break up, since it breaks apart the red component of the RGB colour space from the blue and green components and the shape feature (3). The next section shows that over time such mistakes diminish as a result of competition and selection.
3. If the above adaptations both fail, the heard expression-concept pair is incorporated holistically, leading to a new rule such as S$\rightarrow$yellowcircle/(1,1,0,.5).

At the end of these steps, the hearer performs a few post-processes to remove any multiple occurrences of rules and to update the grammar such that other parts of the internal language relates more consistently to the new knowledge. Full details of the model are found in (Vogt, 2005a; 2005b).

The three learning steps are the core cognitive mechanisms responsible for the co-evolution of linguistic structures and conceptual spaces. Basically, if there is no compositional structure yet in the rules of an agent, but there are regular patterns (i.e. similarities) in both forms and concepts, they are both split up. Yet, this does not necessarily mean these new rules will survive in the language. The way an agent breaks apart holistic expression-concept pairs depends on what the agent has acquired before, so it may make errors. However, later on in life the agent can recover from these errors when it hears new and different usages of parts of an expression. When that occurs, the agent adds new variants to its ‘pool’ of transmissible information units, which then compete for being used. Elements from these pools are selected based on their effectiveness in communication. If an element is used ineffectively, it is dampened and when it is used effectively it is reinforced, while competing ones are laterally inhibited. This competition yields a self-organising effect on the languages of the individual agents, but also brings about effectiveness at a global level, such that a globally shared language can evolve.

In the model, agents have four quality dimensions at their disposal and initially recruit them to form the conceptual space holistically. During development when the holistic expression-concept pairs are broken apart, the agents form new linguistic categories, each semantically relating to a conceptual space of lower dimensionality. The cognitive mechanism for breaking apart expression-concept pairs does not only require an alignment in expressions, but also in conceptual representations. This way a co-evolution of language and concept emerges that on the linguistic side is driven by cultural transmissions and on the conceptual side is facilitated and constrained by the environment (i.e. the objects in the world) and embodiment (i.e. the categorisation into quality dimensions). These processes are all mediated (i.e., facilitated and constrained) by the cognitive capacities of the agents.

4 Simulating the evolution of conceptual spaces

In order to illustrate the framework described in the first part of this chapter and to illustrate the conditions in which a compositional structure of conceptual spaces can emerge, two simulations were carried out. The first simulation, previously reported in Vogt (2006a), illustrates how the model evolves to a sub-optimal solution when there is no generational turnover, so where there is only horizontal transmission. The second simulation demonstrates that more optimal solutions emerge when there is a population flow such that the population contains multiple generations.

Before presenting the results, two measures need to be defined:

Communicative success measures the number of successful guessing games over a time window of 50 games.

Similarity measures the number of games in which both agents used the same syntactic structure over a time window of 50 games. A syntactic structure is considered similar if the words and the linguistic categories used are the same and in the same order. (A linguistic category is characterised by the dimensions that make up the conceptual space of a non-terminal node.)

Both measures are normalised to a value between 0 and 1. Communicative success informs us how successful the population becomes in communicating the referents. This measure, however, does not inform us how similar the internal languages are – the agents may well use different representations and nevertheless be successful in communication. Similarity informs us about the extent in which agents use the same grammatical constructions, thus to what extent they use the same conceptual spaces.

To show the evolution of conceptual spaces in more detail, I also present the relative frequencies of rule types used during successive periods of 10,000 guessing games. As the agents can break up the 4-dimensional conceptual space into two conceptual spaces of lower dimensions without having prior knowledge which dimensions should be separated, 15 different rule types (including the holistic type) can emerge. Only 5 rule types are inspected in this chapter (all other had very low frequencies):

I:	S→rgbs	holistic rule
II:	S→A/r B/gbs	red v. green, blue & shape
III:	S→B/gbs A/r	green, blue & shape v. red
IV:	S→C/rgb D/s	colour v. shape
V:	S→D/s C/rgb	shape v. colour

Rule type I concerns holistic rules in which word forms are associated with the 4-dimensional conceptual space. Rule types II and III are rules that combines the 1-dimensional conceptual space of the quality dimension that represents the red component of the RGB space with the 3-dimensional conceptual space containing the quality dimensions representing the green and blue RGB components, and the shape dimension. The difference between the two rule types is word-order. Rule types IV and V combines the 3-dimensional conceptual space that represents colour in the RGB space with the 1-dimensional shape space.

X	*Y*	*P*
r	gbs	0.297
g	rbs	0.200
b	rgs	0.256
rg	bs	0.117
rb	gs	0.144
gb	rs	0.117
rgb	s	0.075

Table 1 The probability *P* of finding in two different games a co-occurring structure in conceptual space *X* and not in *Y* in which case the 4-dimensional space may be segmented into these two spaces. These probabilities are based on the distribution of feature values that represent the different objects in the world. (This table is reproduced from Vogt, 2005b.)

One of the reasons for inspecting rule types II and III is that in this world, the probability of finding a regularity in the red component of the RGB space is substantially higher than finding any other regularity, such as those required to establish rules IV and V (Table 1). The probability of finding a regular pattern in the RGB space versus the shape space (cf. rule types IV and V) between two randomly selected objects is the chance that the two objects have the same colour (1/12) times the chance that the two objects have different shapes (9/10), which thus becomes $1/12 \cdot 9/10 = 0.075$. The probability of finding a regular pattern in the red component is much higher, because the 12 colours used in the simulation are highly regular in this dimension: 4 colours have value 0, 5 have value 1 and the others have unique values. Without showing the exact calculation, the average probability of finding a regularity in the red component of two randomly selected objects, while the values of the other dimensions differ, is 0.297. The probability of finding regularities in combinations of other dimensions (e.g. g-rbs, b-rgs, rg-bs, etc.) is somewhere in between (cf. Table 1). Although the rules for these combinations would occur more frequently by chance than rules of types IV and V, these are seldomly used by the agents, so their occurrences are not presented.

Despite the probability of finding a regularity in the red component is highest, rule types II and III which exploit this component are not efficient in terms of grammar size. This is because the complements of the red component in the RGB space are not very regular. In fact, the 12 colours have 9 different complements composed by the green and blue RGB components (three of which occur twice, both with red component values of 0 and 1). When combined with the 10 different shapes, the grammar to describe all 120 coloured shapes, would contain at least 96 rules: 5 to cover the red component, 90 to cover the gbs-space and one to describe word-order. In contrast, rules of type IV and V (i.e. those that combine colour with shape) only require a grammar of 23 rules: 12 to cover the rgb-space, 10 to cover the s-space and one to describe word-order. Thus, the two rule-types combining colour with shape are most optimal in terms of compressibility.

4.1 Horizontal transmission

The first simulation is the same as the one reported earlier in Vogt (2006a), but now discussed in the light of the framework set out earlier. This simulation involves a population of 50 agents from the same generation and is run for 1 million guessing games. In each game two agents are selected at random, one agent is arbitrarily assigned the role of speaker, and the other the role of hearer. The context size in each game was set to eight objects, randomly drawn from the world of 120 objects without replacement. Previous research has shown that there is little variation in the results when the simulations are replicated 10 times with different random seeds (Vogt, 2005a, 2005c). For the purpose of this chapter it is instructive to look at the results from one simulation run.

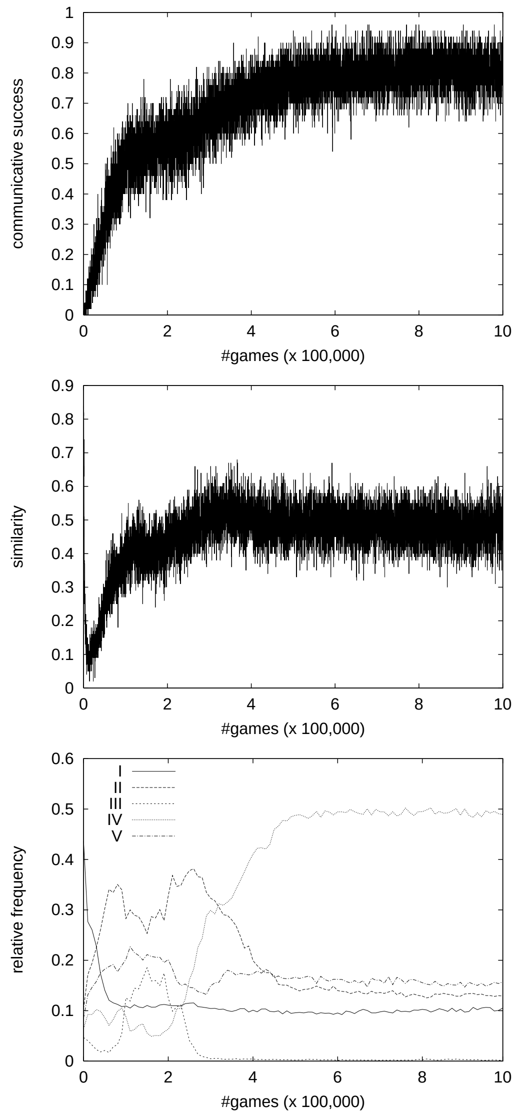

Fig. 4 The results of the first simulation. The graphs show communicative success (top), similarity (middle) and the competition diagram showing the evolution of rule types (bottom). These figures are reprinted with permission from Vogt (2006a).

Figure 4 shows the results of a typical simulation. The top graph shows that communicative success rapidly increases to a value near 0.5, after which it slowly increases to a value slightly above 0.8 and after around 500,000 guessing games, the system stabilised and more than 80% of the games are successful. Similarity (middle graph), however, increases to a value around 0.5, after which it stops increasing. So,

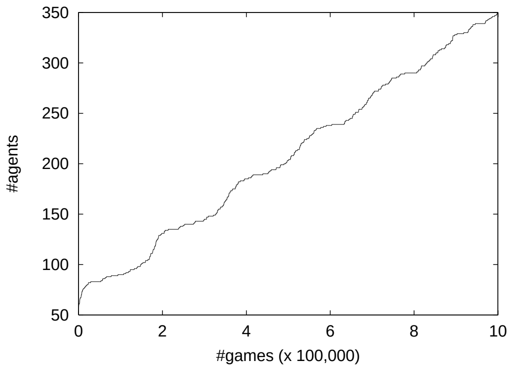

Fig. 5 The total number of agents that have entered the system over time.

in nearly half of the games, the agents use different internal grammar rules, even if they use the same utterances to refer to an object successfully. For example, some agents may use a holistic rule (type I), while others use may rule type II, III, IV or V.

The competition diagram (Fig. 4, bottom) shows the relative frequencies of the five rule types during this simulation. In the first 200,000 games, all rules types compete to be used. At the very early stages, the holistic rule (type I) occurs most frequently, but soon drops to a value near 0.2 after which it stabilises. So, in about 10% of the interactions, the agents use the 4-dimensional conceptual space to communicate objects. The other 90% are divided among all other rule types (including those not shown). After a bit more than 200,000 games, the frequency of rule type III drops to a value near 0, while rule types II and IV appear to compete for some more time until the system more or less stabilised after 500,000 games. From this time onward, the most frequently used rule type is number IV, followed by rule types V, II and I respectively.

So, although communicative success is high, similarity in the representation of the individual grammars (and consequently conceptual spaces) as used by the different agents has evolved into a sub-optimal system. About 70% of the rules used by the agents depend on conceptual spaces rgb and s (rule types IV and V), about 15% by conceptual spaces r and gbs (rule type II), and 10% by the 4-dimensional conceptual space (rule type I). This is sub-optimal, because the most efficient way of representing the grammars is by using rule types IV and/or V, since these require the least number of rules to capture the entire world.

4.2 Isotropic transmission

In the second simulation, the same model was used with the same parameter settings, but, instead of having one generation to simulate horizontal transmission, this sim-

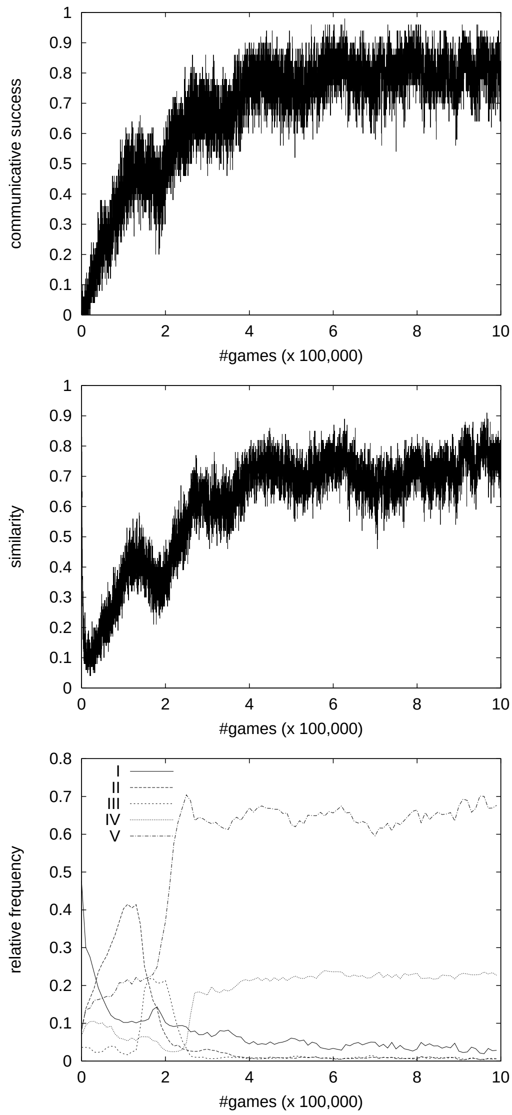

Fig. 6 The results of the second simulation. The graphs show communicative success (top left), similarity (bottom left) and the competition diagram showing the evolution of rule types.

ulation implements a more naturalistic population flow (cf. de Boer & Vogt, 1999; Steels & Kaplan, 1998). By allowing all agents speak to all other agents, this system implements *isotropic transmission* (Vogt, 2005c) that combines oblique, horizontal and upward forms of transmission. To implement a population flow, each agent was given an age measured in terms of the number of guessing games they played individually with a maximum set to 12,500 games. As before, the simulation starts with

50 agents, each initialised with an arbitrary age between 0 and 12,500 games. Each time an agent has played a game, this agent could die with a Gaussian probability distribution with the mean set to the maximum age and a standard deviation of 250. When one agent thus dies, a new agent is added to the population to keep the population size fixed at 50. New agents start with neither concepts nor grammar. Figure 5 shows the total number of agents that have entered the simulation over time.

Figure 6 shows the results of this simulation. The first thing that should be noted is that, in addition to the spikes, there are more fluctuations in the trends of the different graphs. These fluctuations coincide with increased influx of agents as shown in Figure 5. Apart from these fluctuations, it is apparent that communicative success rises to a similar level as in the previous simulation, but similarity rises to a substantially higher level and settles to fluctuate around 0.75. So, agents increasingly agree on using the same rule types. In particular, from about 300,000 games onward rule type V is most frequent (about 70%), followed by rule type IV (about 20%). The holistic rule type I continues to decrease from around 10% at 300,000 games to less than 5% at the end. The other rule types are only sparsely used.

These results demonstrate that when there is a generational turn-over, the language and conceptual spaces continue to evolve towards an optimal system where the grammar represents rules that combine colour with shape in slightly more than 90% of all cases, rather than stabilising in a sub-optimal system as in the previous simulation. So, the new agents rapidly learn the established language by acquiring and using the optimal rule types more effectively than the other rules.

5 Discussion

The simulations presented in the preceding section demonstrated how different conceptual spaces can emerge through cultural evolution. As argued in Section 2, the following factors are involved in this evolution: the environment, embodiment, cognition, self-organisation and cultural transmission. The remainder of this chapter will discuss how these factors contribute to the observed evolution in the simulations, starting with the first three factors, because they are highly interrelated.

The environment of the agents consists of objects that combined a given set of primary colours with a given set of basic shapes. As such, the most obvious way of expressing (and hence conceptualising) these objects is by colour and shape. The agents were embodied with feature detectors that represent the three dimensions of the RGB colour space and one detector that gives a value for each object, however, the agents had no way of telling which of these feature detectors belong to colour or shape and were treated independently. The quality dimensions these agents were endowed with constrained the way they could categorise the perceived objects and how these could be combined to form different conceptual spaces. The cognitive mechanisms were designed such that the agents could only acquire and use grammatical rules that either treated the semantics to be represented holistically or as a combination of two conceptual spaces with the restriction that all quality dimensions are used exactly once. As a result, the agents could construct a total of 15 different conceptual spaces to be used in 8 different combinations irrespective of word-order.

The way the colours and shapes were constructed to form the environment and the way agents could perceive these determined the distribution of focal values in each quality dimension. Since the way agents induce compositional rules from the observed input is based on discovering regular aligned patterns in two or more utterance-concept pairs (as outlined in Section 3.3), the probabilities of finding a regular pattern would drive the formation of grammatical rules as shown in Table 1. To some extent, this is observed in the simulations where rule types II and III occur frequently (at least in the beginning), but all other compositional rule types, except types IV and V, were hardly used. The explanation for this relates to an interaction between the environment, the cognitive learning mechanism and self-organisation.

The environment was constructed such that despite the probabilities of finding regular patterns in all combinations, except colour and shape, were higher, the combination of colour and shape would yield the most compact grammar to express the world. The utilisation of this property would not have happened without the feedback loop – the reinforcement of rule scores and the resulting self-organisation. When agents receive positive feedback, they increase the scores of rules that were used. In cases where agents have different ways of expressing an object by using different combinations of rules, they will select those rules that have the highest combined scores. Since the rules combining colour and shape could apply for all objects, these rule types are more likely to be reinforced and thus more likely to be re-applied. When these rules are more frequently re-applied by the speaker, this increases the chance that the hearer would discover a regular pattern in colour and shape. This positive feedback loop is a driving factor of self-organisation, similar to the way ant paths are formed (Prigogine & Strengers, 1984), and is considered one of the strongest factors for convergence in the language game paradigm (Steels, 1997). Although in the first simulation language evolved into system that incorporated rules combining colour and shape most frequently, a substantial amount of rules of types I and II remained. The constructions formed with these rules were so entrenched in the language that they were viable, also because the language had evolved into a stable system with no variation.

Variation, which is one of the crucial ingredients of (neo) Darwinian evolution (Boyd & Richerson, 2005; Darwin, 1968; Dawkins, 1976) and which is thought to be a driving factor of language change (Croft, 2000; Mufwene, 2001), occurs in the system through the speakers' invention of new words and through the acquisition of new constructs by hearers. In the first simulation, all variant constructions are created and spread among the population in the first, say 100,000 games or so, and after that the competition between the variants take over, which after approximately 500,000 games yield a stable system. The initial variation, subsequent competition and evolution to a stable system is characteristic of the language game model as is most clearly demonstrated in the naming game studied by Baronchelli, Felici, Caglioti, Loreto, and Steels (2006). When, as in the second simulation, newborn agents continue to enter the population and learn the language from scratch, the system no longer gets stuck in such a sub-optimal stable system.

The reason for a continued evolution is that these young agents create new variations in the pool of utterances. Often these new variations are errors or overextentions (Vogt, 2006b) that tend to be unlearnt during development, but sometimes these are new variations introduced by applying a compositional rule to previously unseen objects (a result of the implicit bottleneck, see Section 2 and Vogt, 2005c). Since the rules combining colour and shape tend to occur most frequently in the language (see competition diagrams of Figs. 4 and 6), it is most likely that these new variants reflect that structure. As a result, even more utterances that comply to these rules enter the language, increasing the chance for other agents to discover and use those regularities. This cultural transmission over generations thus strengthens the positive feedback loop that drives the self-organisation. Both language and concepts thus co-evolve to be learnt easier, as there are less rules to acquire (cf. Kirby & Hurford, 2002).

It is important to note that due to the –necessary– abstractions made in this model, it is hard to generalise the results from study to the way humans form conceptual spaces. The simulations are situated in a toy world, with homogeneous agents who can perceive the objects identically and without noise. In addition, the way concepts are constructed from independent quality dimensions is probably unrealistic. Moreover, the model assumes that during the course of human language evolution, protolanguages were essentially holistic and gradually evolved into compositional languages. This assumption is still very much under debate (Arbib & Bickerton, 2010). In spite of these abstractions, the model also contains a number of more realistic assumptions, such as a gradual generational turnover in the population, mechanisms that facilitate self-organisation, and general mechanisms for detecting regularities in the input. As a result, the present study illustrates plausible theoretical principles that may explain how conceptual spaces are shaped. Future modelling work should investigate the scalability of this model using a more realistic world (perhaps even the real world) and agents with more human-like like embodiment and cognition.

6 Conclusions

This chapter has investigated how conceptual spaces can emerge from quality dimensions based on the cultural evolution of compositional languages. The same principles have been demonstrated before in a series of studies where the population flow was implemented based on the iterated learning model in which the population always contained two generations (adults and children) and after a predetermined number of games, all adults die, children become adults and new children enter the population (see Vogt, 2007, for an overview). The differences between the present study and those previous studies concern the more gradual population flow and the focus on conceptual spaces.

The simulations have demonstrated that the evolution of conceptual spaces is driven by five crucial factors: environment, embodiment, cognition, self-organisation and cultural transmission. The emerging conceptual spaces reflect the structure of the environment. Its development within the agents is facilitated by the embodiment through its perceptual apparatus and the cognitive mechanisms. However, embodiment and cognition (and arguably the environment as well) are at the same time limiting factors. Would the agents have been able to perceive other qualities or to manipulate objects, then more complex languages could have evolved, provided the cognitive learning mechanisms would allow them to break apart the holistic utterances in more than two constituents.

The self-organisation results from the variation and competition in conceptual and linguistic structures, as well as the positive feedback loop driven by the learning mechanism. Cultural transmission across generations allows for additional variations to prevent the system entering a sub-optimal stable system and keep the evolution going. Gradually, the emerging language becomes easier to learn, which can catalyse cumulative cultural evolution (Boyd & Richerson, 2005; Vogt, 2006a). Due to the limitations that the model imposed on environment, embodiment and cognition, the linguistic structures and consequently the conceptual spaces evolved into a relatively stable state. However, if there was room for further development, more complex structures could have emerged.

Crucial to the design of this is the assumption that language and concepts coevolve. This is in line with the renewed appreciation of Whorf's linkage between language and thought (Bowerman & Levinson, 2001), and which may account for the cross-cultural differences in the ways languages express and conceptualise various aspects of the world, such as spatial relations (Haun et al., 2011; Majid et al., 2004). Although the present study did not focus on cultural differences in conceptualisation, the framework has the potential to explain these. To achieve this, future studies should incorporate more realistic scenarios based on data from different cultures, as for instance collected by Vogt and Mastin (2013).

Acknowledgements

The writing of this chapter was funded through a Vidi grant awarded by the Netherlands Organisation for Scientific Research (NWO, grant no. 276-70-018). I wish to thank Frank Zenker, Peter Gärdenfors and all participants of the Conceptual Spaces at Work symposium for their valuable contributions in discussing this research. Also, many thanks to Emiel Krahmer and an anonymous reviewer for their valuable comments on earlier versions of this manuscript.

References

Arbib, M. A., & Bickerton, D. (Eds.). (2010). *The emergence of protolanguage: Holophrasis vs compositionality*. Amsterdam: John Benjamins.

Baker, M. C. (2003). Linguistic differences and language design. *Trends in Cognitive Sciences*, *7*(8), 349–353.

Baronchelli, A., Felici, M., Caglioti, E., Loreto, V., & Steels, L. (2006). Sharp transition towards shared lexicon in multi-agent systems. *Journal of Statistical Mechanics*, *P06014*.

Batali, J. (2002). The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In T. Briscoe (Ed.), *Linguistic evolution through language acquisition: Formal and computational models* (chap. 5). Cambridge University Press. Available from http://cogsci.ucsd.edu/ batali/rephrase.ps.gz

Bickerton, D. (1984). The language bioprogram hypothesis. *Behavioral and Brain Sciences*, *7*, 173–212.

Bod, R., Sima'an, K., & Scha, R. (Eds.). (2003). *Data Oriented Parsing*. Stanford:CA, USA: Center for Study of Language and Information (CSLI) Publications.

Bowerman, M., & Levinson, S. C. (Eds.). (2001). *Language acquisition and conceptual development*. Cambridge: Cambrige University Press.

Boyd, R., & Richerson, P. (2005). *The origin and evolution of cultures*. Oxford: Oxford University Press.

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1482.

Brighton, H., & Kirby, S. (2001). The survival of the smallest: Stability conditions for the cultural evolution of compositional language. In J. Kelemen & P. Sosík (Eds.), *Proceeding of the 6th european conference on artificial life, ecal 2001*. Berlin Heidelberg: Springer-Verlag.

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton, NJ: Princeton University Press.

Croft, W. (2000). *Explaining language change: An evolutionary approach*. New York: Longman.

Croft, W. (2002). The Darwinization of linguistics. *Selection*, *3*, 75–91.

Darwin, C. (1968). *The origin of species*. Penguin Books.

Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.

de Boer, B., & Vogt, P. (1999). Emergence of speech sounds in changing populations. In D. Floreano, J.-D. Nicoud, & F. Mondada (Eds.), *Advances in artificial life: Proceedings of 5th european conference ecal'99.* Springer.

de Greeff, J., & Belpaeme, T. (2011). The development of shared meaning within different embodiments. In *Development and learning (icdl), 2011 ieee international conference on* (Vol. 2, pp. 1–6).

De Beule, J., & Bergen, B. K. (2006). On the emergence of compositionality. In A. Cangelosi, A. Smith, & K. Smith (Eds.), *The evolution of language:*

Proceedings of the 6th international conference on the evolution of language.

Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, *32*(05), 429–448.

Gärdenfors, P. (2000). *Conceptual spaces*. Bradford Books, MIT Press.

Garrod, S., Fay, N., Rogers, S., Walker, B., & Swoboda, N. (2010). Can iterated learning explain the emergence of graphical symbols? *Interaction Studies*, *11*(1), 33-50. Available from http://www.isrl.uiuc.edu/ amag/langev/paper/garrod10InteractionStudies.html

Haun, D., Rapold, C. J., Janzen, G., & Levinson, S. C. (2011). Plasticity of human spatial cognition: Spatial language and cognition covary across cultures. *Cognition*, *119*(1), 70–80.

Jackendoff, R. (2002). *Foundations of language*. Oxford University Press New York.

Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin and Review*, *14*(2), 288-294. Available from http://www.isrl.uiuc.edu/ amag/langev/paper/kalish07iteratedLearning.html

Kirby, S. (2001). Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, *5(2)*, 102-110.

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *PNAS*, *105*(31), 10681-10686.

Kirby, S., & Hurford, J. R. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (p. 121-148). London: Springer.

Kirby, S., Smith, K., & Brighton, H. (2004). From UG to universals: linguistic adaptation through iterated learning. *Studies in Language*, *28*(3), 587-607.

Koolen, R., Gatt, A., Goudbeek, M., & Krahmer, E. (2011). Factors causing overspecification in definite descriptions. *Journal of Pragmatics*, *43*(13), 3231–3250.

Majid, A., Bowerman, M., Kita, S., Haun, D., & Levinson, S. C. (2004). Can language restructure cognition? the case for space. *Trends in Cognitive Sciences*, *8*(3), 108–114.

Mufwene, S. S. (2001). *The ecology of language evolution*. Cambridge: Cambridge University Press.

Pinker, S., & Bloom, P. (1990). Natural language and natural selecion. *Behavioral and brain sciences*, *13*, 707-789.

Prigogine, I., & Strengers, I. (1984). *Order out of chaos*. New York: Bantam Books.

Steels, L. (1996). Perceptually grounded meaning creation. In M. Tokoro (Ed.), *Proceedings of the international conference on multi-agent systems.* Menlo Park Ca.: AAAI Press.

Steels, L. (1997). Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In J. Hurford, C. Knight, & M. Studdert-Kennedy (Eds.), *Approaches to the evolution of language.* Cambridge: Cambridge University Press.

Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Sciences*, *7*(7), 308-312.

Steels, L. (Ed.). (2012). *Experiments in cultural language evolution* (Vol. 3). John Benjamins.

Steels, L., & Belpaeme, T. (2005). Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, *28*, 469–529.

Steels, L., & Beule, J. D. (2006). Unify and merge in fluid construction grammar. In P. Vogt, Y. Sugita, E. Tuci, & C. Nehaniv (Eds.), *Symbol grounding and beyond: Proceedings of the third international workshop on the emergence and evolution of linguistic communication* (p. 197-223). Springer.

Steels, L., & Kaplan, F. (1998). Stochasticity as a source of innovation in language games. In *Proceedings of alive vi.*

Steels, L., Kaplan, F., McIntyre, A., & Van Looveren, J. (2002). Crucial factors in the origins of word-meaning. In A. Wray (Ed.), *The transition to language* (pp. 252–271). Oxford, UK: Oxford University Press.

Vogt, P. (2004). Minimum cost and the emergence of the Zipf-Mandelbrot law. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, & R. A. Watson (Eds.), *Artificial Life IX Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems* (pp. 214–219). The MIT Press.

Vogt, P. (2005a). The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, *167(1–2)*, 206–242.

Vogt, P. (2005b). Meaning development versus predefined meanings in language evolution models. In L. Kaelbling & A. Saffiotti (Eds.), *Proceedings of IJCAI-05* (pp. 1154–1159). IJCAI.

Vogt, P. (2005c). On the acquisition and evolution of compositional languages: Sparse input and the productive creativity of children. *Adaptive Behavior*, *13(4)*, 325–346.

Vogt, P. (2006a). Cumulative cultural evolution: Can we ever learn more? In S. Nolfi et al. (Eds.), *From animals to animats 9: Proceedings of the ninth international conference on simulation of adaptive behaviour.* Berlin: Springer Verlag.

Vogt, P. (2006b). Overextensions and the emergence of compositionality. In A. Cangelosi, A. Smith, & K. Smith (Eds.), *The evolution of language: Proceedings of the 6th international conference on the evolution of language.* World Scientific Press.

Vogt, P. (2007). Variation, competition and selection in the self-organisation of compositionality. In B. Wallace, A. Ross, J. B. Davies, & T. Anderson (Eds.), *The mind, the body and the world: Psychology after cognitivism?* (pp. 233–256). Exeter: Imprint Academic.

Vogt, P., & Mastin, J. D. (2013). Anchoring social symbol grounding in children's interactions. *Künstliche Intelligenz, 27*, 145–151.

Wellens, P., Loetzsch, M., & Steels, L. (2008). Flexible word meaning in embodied agents. *Connection Science*, *20*(2), 173-191.

Whorf, B. L. (1956). *Language, thought, and reality*. Cambridge Ma.: MIT Press.

Wray, A. (1998). Protolanguage as a holistic system for social interaction. *Language and Communication*, *18*, 47-67.

Zaanen, M. van. (2000). ABL: Alignment-based learning. In *Proceedings of the 18th international conference on computational linguistics (coling).*