

From Spatial Cognition to Language

Boban Arsenijević

Similarities between aspects of spatial cognition and language are examined in the domains of type of computations (recursive, categorial), type of information used (descriptive and geometric), update procedures for the relevant context representations, and neuro-cognitive aspects (the role of the hippocampus). Striking similarities observed, and the fact that spatial cognitive capacities of all vertebrates are of approximately the same nature and complexity, narrow down the set of possible distinctive properties of the human cognition and the language faculty in the comparative cognition perspective. It is proposed that these properties are: (A) domain-general use of the otherwise similar computational capacities, (B) serialization of the computations of descriptive and geometric means of reference, and (C) increased importance of the update of mental representations by a group rather than just an individual.

Keywords: evolution of language, hippocampus, language faculty, spatial cognition

1. Introduction

This article compares two cognitive domains: spatial cognition and natural language. While the former is present in a quite sophisticated form at least in all vertebrates, and a number of other species, the latter is an exclusive characteristic of humans. From the aspect of evolution, this means that spatial cognition has been there in the animal world for a very long time, while language is a relatively new development. A look at some interesting similarities and differences between these two capacities may contribute to the theories of each of them. Yet, more can be revealed about a newer capacity by looking at an older one, than about an older capacity by looking at the one that is more recent. The interest of this article is set in this more informative direction: It aims at learning about language (its setting among other capacities, its origins, its structures) by comparing it to the spatial cognition capacity, more precisely to one of its

The present work is a result of my engagement at the project "The Origins of Truth and the Origins of the Sentence" funded by NWO (The Netherlands Organization for Scientific Research). I am grateful to two anonymous reviewers for their insightful comments. The article benefited a lot from my discussions with Wolfram Hinzen, Aritz Irurtzun, Marieke Schouwstra, and Hedde Zeijlstra. The mistakes are of course mine.



components: cognitive maps.

The article is organized as follows. In sections 2 and 3, I briefly outline the major aspects of spatial cognition and language, respectively, i.e. those that compare in the most insightful way from the perspective of the article. Section 4 points to the relevant similarities, and section 5 presents and discusses some differences. In section 6, I discuss the results of the comparison, especially concentrating on the possibility that language has evolved from the spatial cognition capacity, with the crucial step being an extension of the spatial computation to a domain-general use. Section 7 concludes.

2. Cognitive Maps

The field of spatial cognition presents one of the better-explored domains of the cognitive neuroscience. It has acquired a significant body of knowledge, which establishes quite precise links between the functional, representational and neurological aspects of the domain. Experimental work on a wide variety of species has resulted in a broadly accepted functional architecture of spatial cognition, and in precise linking of some of these functions to particular brain areas. One of the central fields of the theory of spatial cognition is concerned with cognitive maps — a component that is prominent in spatial cognition of all vertebrates, including humans (as opposed, for instance, to dead reckoning or the so-called compass mechanisms). I briefly present those among the core elements of the theory of cognitive maps that are of particular interest for this contribution.

Cognitive maps represent territories and involve two main types of information: the map of the territory, involving places, paths between them and their geometric configurations (spatial cues), and descriptions (object-specific cues) of each of the places involved, expressed in terms of a number of associated features (Vallortigara, Zanforlin & Pasti 1990, Vallortigara, Pagni & Sovrano 2004).¹ A third emerging type of information is the geometric information about a place, i.e. the set of relations of a place with other places and the set of geometrically relevant properties of a place (length, height, shape): This type involves geometric information, but presents part of the description of a particular place. Together, they form the representation of the spatial context.

For instance, consider the inside of a box with a rectangular base, painted white, with one red corner in which there is a small piece of meat. A rat, placed in this box, represents this territory as a spatial context. The context representation involves a map with a number of places: two long walls, two short walls, the floor, the ceiling, two corners with a long wall to the left and a short

¹ In this article, I group together the pure geometric cues and the landmark cues, since they both involve the component of a spatial structure, absent in pure descriptions. I am aware of the facts that imply that the two grouped types of cues are different and should be treated apart, but for the purposes of the contribution, their grouping does not have important consequences and is a handy simplification. In fact, it would even strengthen the point because (i) even in language there are global geometric cues (topic, focus, familiar) and the more local ones, relative to some prominent referents (e.g., proximity vs. distance in demonstratives and other elements) and (ii) introduction of one more category further stresses the categorical nature of spatial cognition, as a parallel with language.

one to the right, one of which with the additional properties of being red, emitting a smell of an edible thing and containing an object, and two corners with a long wall to the right and a short one to the left, distinguished by whether they have the 'red edible' corner to the left or to the right. As abundantly confirmed by experiments, ignoring the red color and the piece of meat, we get two pairs of indistinguishable walls and two pairs of indistinguishable corners.

The representation of a spatial context can be updated by new information acquired through a sensory input. It has been argued that this process goes via the match–mismatch procedure (Mizumori, Ragozzino & Cooper 2000), which can be briefly sketched as follows. At any point, the animal has a spatial context representation, constructed as a set of expectations for the territory it is located in. The sensory input is continuously matched with this set of expectations, leading to the preservation of the matched and correction of the non-matched expectations. The spatial context representation is thus subject to a constant update. In the given example, removing the animal from the box, moving the piece of meat to a neighboring corner and then bringing the animal back into the box would result in the update of its spatial context representation by specifying that the properties of being red and of having an edible object are now distributed over two corners, and the corner with the latter property now has a long wall to its right and a short one to its left, while the rest of the context stayed the same.

Finally, the spatial context representation serves as a background for different behavioral actions, such as movement, eating, drinking, removing obstacles, etc. These actions depend on motivational aspects, and most of them introduce the need for an update of the spatial context representation. This update may involve the integration of a path that the subject is moving on, the change of the subject's location in the map or the change of features of some place in the map (after the food is eaten, the place where it used to be loses the feature of containing an edible object).

Animals can compute complex structures from a spatial context representation, among which most prominently paths. Paths involve the source, direction and goal, but also possibly a number of places via which they reach the goal, and which may serve as intermediate cues for navigation while moving along the paths. During movement, the path needs to be regularly recomputed, updating the position of the subject on the path, and hence also on the map. Note that this involves computation of paths and places (and their features), and as such is a different mechanism from dead reckoning (vector-based computation of the position of a moving animal with respect to the starting position), although dead reckoning may be involved in the computation of paths.

When two places have the same description in terms of non-geometric features, their position relative to some other, unambiguously defined place may serve as the distinctive feature (as with the pair of corners without meat above). This implies that spatial computation involves hierarchical structures of the type '[THE_SHELTER [BETWEEN [THE_TREE AND THE_ROCK [ALONG [THE_WATER [BEHIND [THE_HILL]]]]]]]'. It takes a hierarchical structure to represent one place as specified by a description that involves another place (and its description).

One of the major roles in the computations producing the discussed

representations, and dealing with their update and use in other capacities, is played by the hippocampus, a brain area that can be identified in a broad range of animal species. There is evidence that this is where the spatial context representation and its update are handled (Nadel, Willner & Kurz 1985, Anagnostaras, Gale & Fanselow 2001). In addition, the hippocampus appears to have a role in the coordination between this representation and the peripheral modules: the sensory input, the motivational aspects, and the behavioral actions (Jakab & Leranth 1995, Markus *et al.* 1995, Wood *et al.* 2000). This means that among other things, the hippocampus is responsible for matching the sensory input with the spatial context representation, and for selecting parts of the spatial context representation to be matched with motivational aspects and patterns of behavioral actions. In other words, (i) it pairs the sensory input with a segment from the spatial context representation, where the latter can be seen as the interpretation of the former, and (ii) triggered by different motivational aspects, it matches segments of the spatial context representation with the adequate patterns of behavioral actions, usually realized by the motoric system.

3. Language

Natural language grammar is traditionally defined as the system that maps between the meaning and its physical carrier. The physical carrier is produced by the motoric system and perceived by the sensory system, while the meaning is taken to be some mental representation, directly or indirectly related to the real world. Arsenijević & Hinzen (2007) argue that there is no separate level of representation reserved for the meanings of linguistic expressions. Syntax *is* the specification of the compositional structure of the complex concept that we recognize as the intuition about the meaning of an expression. And syntax directly interfaces the discourse, and drives the integration of the expression, securing that the concepts associated to the terminal lexical units are integrated in the proper discourse domain and in the proper relation with respect to other such concepts in the expression. To sum up, semantics, as the intuition about the concepts derived by linguistic expressions, relates to two empirical domains: the syntactic structure that derives these concepts, and the effects of the integration of the expression into the discourse.

The use of language always involves a discourse: the representation of the aggregate body of information relevant for the current language use. This information includes contributions of the explicit linguistic expressions uttered so far in the communication situation, the immediately relevant presupposed material, and the directly relevant parts of the non-linguistic sensory input (the communication situation). The discourse consists of referents, their descriptions in terms of different predicates, and discourse functions (topic, focus). Discourse functions mark the position of a referent within the internal organization of the discourse — is it within the speaker's, the hearer's or some remote domain, is it part of the (recent) old information or not, is it a member of some relevant set, etc. The discourse also involves paths: A remote referent can be reached via the more proximal ones: an expression like *the dog of my friend's sister* would be used

if the speaker's friend in question is more topical in the discourse than the friend's sister, and the friend's sister is more topical than the friend's sister's dog, hence deriving the path among referents: *the friend – the sister – the dog*.

The discourse can be updated from the sensory input. If the discourse involves some relevant bit of information, which is denied, or corrected, by some expression uttered within this discourse, the discourse gets updated. Even if a bit of information became part of the discourse by the contribution of an earlier linguistic expression within the same communication situation, it can be denied or corrected (e.g., someone realizes that what he said a minute ago was wrong). This means that the discourse consists of expectations, which can be modified by a strong enough sensory input.

The discourse can trigger certain behavioral actions, falling into two important classes. The first is the production of a linguistic expression: an action effecting in yet another contribution to the discourse. This is a consequence of the fact that the discourse is usually shared, and those sharing it normally want to enrich it. A direct contribution to the discourse is achieved by producing a physical entity that presents a sensory input for the other persons sharing the discourse, thus making them update their discourse representations with the relevant material. The second class is simpler: Some updates in the discourse may introduce a direct instruction for the subject to take a certain behavioral action, assuming that a sufficient motivational support is provided (like the sentence *Leave me alone!*). These updates are referred to as speech acts.

In the domain of pure sensory input, linguistic expressions, serving as one of the possible sensory inputs leading to the discourse update, present (sets of parallel) linear strings (*the-dog-of-my-friend's-sister*). They have to be assigned hierarchical structures on their way to the discourse: [*the dog [of [my friend's] sister]*]. The hierarchical structures mediating the discourse update fall in the research domain of syntax. The choice of the hierarchical structure to be matched with a linear string is usually restricted by the relatively restricted expectations in the discourse. The proper one among them is in the default case uniquely determined by two types of information: the sequencing of units forming the linear string, their ordering, and their categorical and selectional properties memorized in the lexicon.²

It has been shown that in language comprehension, hippocampus plays a central role in the syntactic integration. Recording of Event-Related Potentials (ERP) shows that syntactically incorrect sentences elicit a negative deflection of 500–800 ms in this brain area (Meyer *et al.* 2005). This does not necessarily imply that hippocampus is directly involved in the processing of the syntactic structure (see Opitz & Friederici 2003 for arguments that the rule-based aspects of grammar are computed by other centers). It does mean, however, that the hippocampus plays a role in the discourse integration, and that this role is sensitive to whether the expression that is being integrated is assigned a valid syntactic structure or not. The special activity of the hippocampus may, for

² The phonetic string is in fact phonologically computed as not linear but hierarchical, with a relatively shallow hierarchical structure (compared to that of syntax, to which it imperfectly matches), but since this article concentrates on the structures with direct semantic effects, the issue is slightly simplified.

instance, be related to trying to use the information from the discourse to identify the most likely update — in the lack of precise information provided by syntax, but I refrain from going into speculations of this kind.

4. Parallels

This section presents some striking similarities between cognitive maps and the faculty of language, concerning their general architectures and the core structural properties of the computational mechanisms behind them.

Let us consider first the general architecture of the two capacities. Grammar, the core component of the faculty of language, has for more than a century been defined as a system that maps between physical objects (carriers of messages) and meanings (contents of messages). In the language use perspective, this renders three domains:

- (A) the mental representations corresponding to the meaning;
- (B) language production, as one aspect of the physical carrier pole;
- (C) language perception as its other aspect.

Grammar maps from the perceived linguistic material to meanings, and from meanings to the behavioral patterns engaged in the production of a message by speaking, signing or writing (or in other possible ways). Moreover, it has been argued that the notion of meaning should be dispensed with, in favor of a more precise model, in which it resolves into the lexical and syntactic material of a linguistic expression on the one hand, and the effects of the integration of the expression into the relevant discourse (Arsenijević & Hinzen 2007, Hagoort & van Berkum 2007). This gets us to the following picture: There is a computational module, grammar (plus lexicon), which drives the integration of the relevant type of sensory input into a special mental representation (the discourse) and triggers the adequate patterns of behavioral actions for the retrieved segments from the mental representation. The behavioral actions are sensitive to communicational and other motivations, and may be of two types: (i) language production and actions affecting the immediate context of communication and (ii) other behavioral actions coming as a response to the changes in the discourse. As mentioned above in section 3, the hippocampus has a central role in the coordination between the discourse and the sensory aspects of language, especially in the process of discourse integration of new material.

With a high degree of parallelism, the standard model of spatial cognition involves coordination between a mental representation (the spatial context representation), the relevant sensory input and the adequate behavioral patterns (Cheng & Newcombe 2005). The context is viewed as the continuously present background stimuli, but it also includes internal aspects such as plans, goals, motivation, types of behavioral activities involved (Markus *et al.* 1995). The sensory input that does not match the relevant segment of the spatial context representation gets integrated, leading to an update of the spatial context representation (Mizumori, Ragozzino & Cooper 2000). Patterns of behavioral

actions involve two types of motivation: the general curiosity about the spatial environment and the independent, non-spatial motivations such as hunger, thirst, fear, etc. (Voicu & Schmajuk 2001). They can be divided into behavioral actions with an immediate controlled effect on the spatial context (removing an obstacle, storing food in some place, changing the landscape by, e.g., digging) and those without such effects. As mentioned above in section 2, the coordination between the sensory input, the spatial context representation and the behavioral actions is shown to involve a significant role played by the hippocampus.³

The two mechanisms of update are characterized by one important difference: The update of a cognitive map by the sensory input is unconditional, while the update of the context representation by the linguistic input is not — the subject may as well decide to discard it if in conflict with some well-established old part of the representation (in effect similar like bees do it in the experiment described below on this page⁴). The choice that humans have in this respect can be attributed to their rich theory of mind. Humans can deal with a number of context representations, possibly embedded in one another. When talking to Bill, apart from her own immediate world-knowledge, Sue also deals with a representation of a context shared with Bill, and a representation of the context that she assigns only to Bill (including, e.g., points that they know they view differently). Any proposition contributed by Bill can be used to update any one of these representations, a combination of two, or all three of them. Even if Sue thinks that Bill is lying, she has to update the shared representation in some way. Hence, the necessity of update is common for the two capacities, but the linguistic update is characterized by the multiple possible target representations.

One of the phenomena most frequently identified as a characteristic property of language is the capacity of reference. By use of language, humans can talk about a particular object in reality, and assign it certain relevant properties, even when this object is out of the reach of any of their senses ('displaced reference'). A counterpart of this capacity can be identified in the domain of spatial cognition, in a number of different examples. For instance, the so-called homing species can compute paths leading to a particular location even when this location is far out of the reach of their (visual, olfactive, and other) senses. Examples like this allow for the possibility that the animal does not operate over a representation of the place of homing, but instinctively computes some complex paths that bring it to a territory where it can use other navigation mechanisms.⁵ However, an experiment with bees, referred to in Gallistel (in

³ The parallels outlined do not entail, of course, that the two capacities do not also have computational components specific for each of them, shared with other capacities, or engaged by spillover when the centers engaged are overloaded. In any case, it is beyond the goal of this contribution to make more concrete claims about the neurological aspects of the parallel. Also, the architectures as presented might reflect a more general architecture of any module involving sensory input, and not only between cognitive maps and language, but other similarities point to a tighter parallel between the two capacities discussed.

⁴ I thank an anonymous reviewer for pointing this out to me.

⁵ But even in this case, the difference assumed between the human awareness of a place and the animal's instinctive mechanical computation of certain geometrical and other properties is an empty one. Ultimately, our awareness of a place, or other referent, also has to be represented in terms of neuro-cognitive mechanisms.

press), proves that animals can indeed compute, using cognitive maps, information involving a particular remote place, about which they are currently receiving no sensory input. In the experiment, an artificial source of food, new in the territory, moved in three steps: from a flower field to the water, and to another flower field. Different groups of foragers visited each of the three places, and then went back to inform the community about the source of food. Their dance only had effect when they were informing about the first and the third step, i.e. when the source was in places already known as possible locations of sources of food. When informing about the source on the water, the bees ignored the dance. This shows not only that bees have representations of remote places, but also that they are able to evaluate information as true (or useful, trusted) or false (or useless). Similarly, the ability of animals to create new paths within the territory, such as yet unexplored shortcuts (Taylor, Naylor & Chechile 1999), implies that they have a representation of the place presenting the goal of the path. This all suggests that the core of computation of cognitive maps at least in some species involves a counterpart of reference: animals can compute the representation of a certain place even when it is absent from their immediate perceptive input. The fact that humans can perform more complex activities with respect to reference may be a consequence of the absence of restriction of the computation of reference to a small number of domains (e.g., space, social relations) and of the higher processing and memory human capacities of computation of the context and of referents. To the exclusion of domain generality and memory and processing capacities, the essential ingredients seem to be shared between a number of species.

The problem with reference is that there is no standard theory of it, and hence it is difficult to make a deeper comparison between linguistic reference in humans and that found in spatial cognition in a wider variety of species. At least at the descriptive level, however, some more substantial parallels can be made.

In language, referential expressions usually involve two ingredients. One is a specification of a discourse domain, and of a discourse function that the intended referent has in that domain, and it is achieved by different tools which include sentence typing, discourse function marking, and demonstrative pronouns. The other ingredient is a description, a complex concept built by a composition of a number of simpler ones. The description should be restrictive enough to reduce the number of possible candidates within the relevant discourse domain to one. The two ingredients act as the address and the name of the addressee: The discourse domain and the discourse function specify the city and the street, while the description singles out the addressee among the candidates with the same address. As an illustration, consider the passage in (1), and especially the underlined nominal expression.

- (1) Two girls went to the hairdresser. One of the girls was extremely tall. The hairdresser told the tall girl to come another day...

Its intonation, position within the sentence, and the definite article used suggest that the expression refers to an old, topical referent, part of the ongoing discourse of the preceding two sentences. The words *tall* and *girl* specify two properties, i.e.

two concepts; the syntax of the expression specifies that the two properties should be intersected, producing a more restrictive interpretation. This interpretation is the description of the referent — it is tall and girl at the same time (for sake of simplicity, I ignore the meanings of singularity and countability, also contributed by the expression). The discourse domain and the discourse function reduce the possible candidates to the hairdresser and the two girls. The description singles out the intended referent among the candidates, in this case the tall girl. Note that *girl* only would not be enough (hence the expression *the girl*, instead of the underlined *the tall girl*, would not be salient for the context).

In cognitive maps, as briefly outlined in section 2, two types of information are used: the geometric information about a place, specifying its location relative to different elements of the spatial context representation, and the descriptive information, specifying some properties that characterize the relevant place. Interestingly enough, it seems that animals with sophisticated spatial cognition, apart from humans, acquire the two types of information independently of each other, and do not combine them, but have to choose only one of them for every attempt to locate a particular place (and pick a pattern of a behavioral action); cf. Wang & Spelke (2002), Pearce *et al.* (2006). Apart from this aspect, which is discussed in section 5, there is a strong parallelism with language: Both cognitive maps and language use two types of information in locating a referent. One type locates the referent relative to the organization of the mental representation of the relevant context (the spatial context representation or the discourse representation), and the other involves features, i.e. non-geometric concepts, to specify a restriction that singles out the relevant referent from the set of suitable candidates. The size of the context representation, the possibility for it to involve domains currently inaccessible to the senses of the individual (including the abstract ones), and other mostly quantitative features of the particular system may lead to dramatic differences in their performance, but the ontological similarity remains a fact.

The parallel between cognitive maps and language related to the use of geometric and descriptive information goes even deeper. It appears that in familiar spatial environments (i.e. contexts), in a vast majority of (vertebrate) species, geometric cues are preferred to the descriptive ones (Gibson & Shettleworth 2003). This is explained by the fact that geometric cues are less likely to vary over time, compared to the visual, auditory, or other descriptive cues. For instance, a colorful flower may close during the day, or change its angle with the ground, but its location will stay the same. In new environments, however, descriptive cues are more important than the geometric ones. In fact, in a new spatial environment, geometric relations are still to be established by exploration — which is most naturally performed by using the more readily perceived descriptive cues.

In language, within an old discourse, the preferred way of locating its referents is through (repeated) personal and demonstrative pronouns, elements using geometric features of locality and proximity, rather than through repeated descriptions. Consider the example in (2), where pronominal elements are used in the reply in B, and descriptions in B'.

- (2) A: My friend Mary's math teacher wants her to be more active.
 B: Did he tell her what exactly he meant by that?
 B': # Did (your friend Mary's) math teacher tell (your) friend Mary what exactly (your friend Mary's) math teacher meant by telling your friend Mary that your friend Mary's math teacher wants your friend Mary to be more active?

It is a standard view in syntax and semantics that pronominal elements, including demonstratives, are directly related to the functional projection of the determiner, which specifies the features related to specificity and definiteness, both linked to the organization of the discourse (e.g., Kayne 2002). Moreover, demonstratives often involve the component of distance vs. proximity (*this* vs. *that*), which is in most cases related to the discourse organization and the abstract vicinity of the referent to the speaker or to the thematic background of the expression (Jayaseelan & Hariprasad 2001).⁶ This shows that language also prefers the use of geometric features in familiar discourses, i.e. for discourse-familiar referents. New discourses, new referents, and new choices among old groups of referents are better handled by the use of descriptions. This is why the subject of (2A), *my friend Mary's math teacher*, is used, as an expression involving a description: The referent has probably not appeared in the immediately preceding discourse, and therefore has not yet been assigned 'discourse-geometric' properties in the relevant discourse domain.

Another characteristic property of natural language grammar is that it crucially relies on a set of categories, which can embed in one another under certain structural restrictions, referred to by a number of different terms: subcategorization, selection, projection, etc. So, for instance, the category verb (VP in syntax) can embed immediately under the category tense (TP) — such as in $[_{TP} -ed [_{VP} walk]]$ — to give the interpretation of the location of the eventuality denoted by the verb in the past with respect to some reference time and/or with respect to the speech time (the tense operates over the meaning of the verb).⁷ However, TP cannot embed immediately under VP ($[_{VP} walk [_{TP} -ed]]$), to give, for instance, the interpretation of a past that has the property of walking, perhaps as opposed to running or flying (i.e. where the meaning of the verb operates over the tense). Even more importantly, TP cannot embed immediately under TP, nor can VP embed immediately under VP (two verbs can compose, but this is either a case of coordination or a more complex hierarchical structure than immediate embedding). All these embeddings become possible when mediated by other categories, and hence not immediate: The restrictions are local and target imme-

⁶ Imagine the following reply to (2A), in which the interlocutor switches to another world:

(i) B'': In a fairytale, he'd be her evil stepmother.

The dynamics of the use of pronouns and the features *distal* and *proximal* indicates that they are sensitive to a geometry independent of any particular world that belongs to the discourse, unless some particular worlds are involved in specifying the property that singles out the referent from the relevant set of alternatives.

⁷ The lexicalization of this structure as *walked* is an issue of morphology and phonology, it is not a source of controversies, and is not interesting for the current discussion.

diate embedding only.

Something very similar to this can be found in the domain of cognitive maps. As outlined in section 2, there are two major types of spatial objects: places and paths. These make two different categories. Moreover, it appears that paths are a more complex category, the definition of which may immediately involve places (paths must go from places, to places and/or via places, while places can be specified by descriptions that do not involve paths), such as, for example, [_{PATH} goal [_{PLACE} home]]. Places are defined based on the geometric properties of their position with respect to the territory and other places on it, especially landmarks, and based on the properties they have, such as color, smell, or shape (i.e. descriptive features). Paths are dynamically computed in every individual situation, because they directly depend on the current position of the subject, and the way in which it is changing. They are dynamic interpretations of geometric properties, actual in the respective situation. There are also other possible categories, such as landmarks (a subcategory of places) or geometrical structures, an issue I do not dwell on in this article because the two categories above suffice for the aimed argument: That categories, with the same type of embedding, and restrictions over the embedding are an important property of cognitive maps as well as of the language faculty.

The distinction between descriptive and geometrical features is another level of categorization. While animals seem to compute these two categories separately, each of them still embeds in the category of places: Places are determined by their descriptions and/or by their geometrical positions. Hence we get an even deeper hierarchical structure: [_{PATH} goal [_{PLACE} [_{DESCRIPTION} home]]]. Moreover, if the description involves an additional feature specified with respect to another place, a real recursive embedding takes place; for *the water near the home*, [_{PATH} goal [_{PLACE} [_{DESCRIPTION} water, near [_{PLACE} [_{DESCRIPTION} home]]]]], for example. Embedding of this kind is quite restricted: computation of cognitive maps probably can only handle structures with at most one round of embedding (one place described in term of one other place). Yet, this restriction may be imposed by the memory capacities, or by economy principles, rather than by the computational capacity, which is then genuinely recursive. If it is correct that the computation of cognitive maps involves structures with recursive embedding, this presents recursive computations as a much older development in the course of evolution than argued by Hauser, Chomsky & Fitch (2002).

It is possible to speculate in the direction of establishing parallels between places in cognitive maps and referents of nominal expressions in language on the one hand, and between paths in cognitive maps and eventualities in language. Both members of the former pair correspond to geometric points, and both members of the latter pair have linear structures. Moreover, both paths and eventualities include places and objects, respectively, as important defining elements in their structures. Without a neuro-cognitive, or at least a deeper cognitive support, this line of thinking remains in the domain of speculations.

In talking about the way language establishes reference, referents were discussed as located within discourses, but also within parts of discourses, discourse domains, which present the immediate thematic, temporal and spatial vicinity of the most prominent topical referents in a particular segment of the text. This im-

plies a hierarchical organization of the discourse, i.e. its division into domains, which are smaller and hence easier for retrieval and for locating referents in them. The very same has been argued to hold for cognitive maps: A territory is always divided into sub-territories, which may be further divided, in order to make the retrieval procedure faster and simpler (Schmajuk & Voicu 2006).

Finally, as presented in sections 2 and 3, in the computation of both cognitive maps and language, an important role is played by the hippocampus. In fact, some proposed descriptions of this role in the spatial domain are equally well applicable to its linguistic aspects. One of them defines the activity of the hippocampus as directly handling a coding of locations, events, behavioral strategies and their mutual relations into the context representation (Aggleton & Brown 1999). This could well serve as a description of the effects of discourse integration, if locations are taken to cover all referents, and the relevant context representation is taken to be the discourse. This touches on the important question of language-specific cognitive elements, supporting the view that most of the components of the language faculty are rather general and apply in other capacities as well. Such is the case with the procedures that integrate new material to some context representation, be it the spatial context, the discourse, or some other relevant representation, or with those which, influenced by motivational impulses, match segments of a context representation with behavioral strategies.

5. Differences

The major difference between natural language and spatial computation is that while in the latter only one individual integrates new information into the relevant mental representation, in the former the representation can be shared and updated by groups of individuals. In fact, the possibility that more than one individual shares the same discourse and participates in the process of its update is one of the properties of language that have influenced its current form the most, from the very existence of the phonology and phonetics (and hence the dual patterning of language), to a number of smaller differences at all levels at which spatial cognition and the language faculty can be compared, especially at those at which language involves an important role played by phonology.

Another interesting difference is briefly mentioned in section 4. While in language, descriptions and geometrical properties (i.e. those related to the organization of the discourse) appear as parts of one and the same unit of computation (a linguistic expression), in cognitive maps, these two types of information are separated and one unit of computation may consist of elements either only of the descriptive, or only of the geometrical nature. In fact, even in language, the two types of information are not really computed simultaneously. They appear strictly structurally segregated: The 'lower' structural domains are reserved for the descriptive content, while the 'higher' ones involve the discourse-related information (e.g., Rizzi 1997). This means that grammar performs serialized computations, where each unit of computation, a phase under the phase theory of syntactic computation (Chomsky 2001), consists of the

inner phase with the descriptive content, and the edge, whose contents are discourse-related (McNay 2006).

It is interesting that the acquisition of language in humans influences their spatial cognition. Some properties of the processing of cognitive maps, including the use of descriptive and geometric cues, undergo a drastic change around the age of six, which is also considered the critical period during which the individual rounds her acquisition of grammar (Hermer-Vazquez, Spelke & Katsnelson 1999).

A third difference, and the last one to be discussed in this section, concerns the domains of the two capacities. In the spatial computation, only spatial objects are categorized, all the other concepts falling in the category of 'the rest', i.e. of the material used only for descriptions. In natural language, referents are not restricted in any way; they do not even need counterparts in the real world: they can be abstract (as in: *The suspicion caused jealousy.*), non-existent (*A unicorn fell in love with Godzilla.*), or even impossible (*square circles on solid liquids*). This means that every concept can be used as a referent (jealousy, redness, distance) as well as (part of) a description. Referents are placed in an abstract space, the discourse, with its own organization that is only marginally influenced by the spatial relations in the real world. This makes the geometric properties of the discourse abstract and much more easily transformed than those of the (representation of) physical space. At the same time, it makes the process of updating and retrieving the context representation much more complex in language than in cognitive maps. One tool that language developed for this purpose is a richer set of categories. Instead of several categories that could be identified in cognitive maps, syntax disposes with several dozens of categories (at least according to the 'cartographic' approach to syntax, see, e.g., Cinque 1999).

The result, at the surface, where functional effects of the cognitive systems are observed, is a significant asymmetry. Spatial cognition, the domain in which animals show a high level of reference-based abilities, produces and deals with (retrieves, updates, combines) a set of SPATIAL contexts and sub-contexts stored in the long-term memory, amounting to the set of relevant territories in the life of an individual. Language, present only in humans, and based on their domain-general application of the recursive computational algorithms, produces and deals with a drastically larger set of discourses and discourse domains in the unified abstract macro-space of *all the available concepts and all their (possible) compositions*.⁸

A very important consequence of the domain-general application of recursive computation in humans is that they can not only embed (spatial) contexts in other (spatial) contexts, but also embed non-spatial contexts in non-spatial objects. In such a way, a powerful theory of mind can be derived: The description of each object potentially involves a context, or a set of contexts (her knowledge, views, beliefs). This may be the explanation for the universality and important role in grammar of the feature of animacy. Animacy marks objects that

⁸ The possibility that humans only differ from animals in their domain-general use of recursive computations is mentioned in Hauser, Chomsky & Fitch (2002), although it is not the claim they are directly arguing for.

can have their own context representations, i.e. 'points' in the discourse where another discourse can be embedded. And this altogether enables a multiple update, as the core of communication. While a unit of a spatial sensory input updates one spatial (sub-)context, a sentence in language may update a larger number of context representations, some of which can be embedded in the descriptions of referents of other contexts. Crucially, the context representation of one interlocutor contains representations of other interlocutors as objects, i.e. referents, and the description of these referents involve representations of their relevant context representations. Each of these embedded context representations is normally updated by each sentence uttered in the discourse, parallel to the update of the hierarchically highest context representation. On the surface, this looks like a *group update*: Each of the interlocutors represents a number of (sub-)contexts that have counterparts in all the other interlocutors, and all such sets of counterparts are updated in (nearly) the same way. Individuals of the group thus develop synchronized context-representations, which enable a synchronized functioning of the group. Apart from its cognitive and linguistic significance, this phenomenon plays an important role at the social level, which has probably been one of the ingredients in the selective pressure that pushed the evolution of language (Bickerton 1998).

6. Discussion

Similarities and differences between spatial cognition and language discussed in this article could be interpreted in three possible ways. One option is that the similarities observed are just a consequence of the methodological apparatus applied: Cognitive sciences deal with a set of general models, such as the division of different systems into computational and memory components. The fact that the same set of models can be fruitfully applied to different subjects of study does not guarantee that a deeper exploration would not uncover significant differences and require a modification of the models that would make them ontologically different from each other. This option presents a general danger for any theoretical work and hence will be ignored, leaving to the future research to prove it correct or wrong.

The second possibility is that the similarities are not more than that: (vague) similarities between two different systems. The weakest explanation would be that the similarities are accidental. A stronger one would be that they are a consequence of some general properties of cognition, i.e. of the neuronal systems in the brain, but that they still are disjoint systems. The strongest option under this interpretation is that the two systems share some components, for instance the computational module engaged in the retrieval and update, or the window to the long-term memory, and that this shared component is responsible for the shared properties between the two systems. This option agrees quite well with the neuro-cognitive data about the role of the hippocampus in both systems, as discussed in sections 2–4.

The third possibility is that the language faculty has evolved from the spatial cognition capacity. This is the strongest, and hence the theoretically most

interesting interpretation: It allows for the second possibility above as the description of the current relation between the two capacities, but it also hypothesizes on the origins of this relation. Therefore, but also because it is an attractive hypothesis, this interpretation receives a more extensive discussion in this section.

Different cognitive capacities have been suggested as the possible immediate origins of the complex computational patterns found in natural language, in the arithmetic capacity and in other sophisticated human cognitive capacities. Among them are the vocal production (Carstairs–McCarthy 1999), social cognition (Bickerton 1998), motorics (Jarvis 2007), and navigation (Bartlett & Kazakov 2005). In the remaining part of this section, I consider arguments in favor of spatial computation as a better candidate.

Virtually all animals, and even some plants, show some sort of sensitivity to aspects of space. Whenever this sensitivity is not a matter of a direct physical reaction, but requires the mediation of some biological process, it may be considered to involve computation. Hence, it is reasonable to think that spatial computation preceded any other kind of cognitive computation in animals. Moreover, it is a prominent possibility that other types of computation developed through the process of broadening, or shifting, the domain of application of the spatial computation, and of its gradual, or perhaps at times abrupt, sophistication. This is to say that all the types of computation that can be observed in animals today stem from the original purely spatial computation, which emerged very early in the animal evolution line.

Arguments in favor of this view are numerous. First, most other domains in which computation applies either can be seen as essentially spatial, or can be seen as metaphorically subjecting non-spatial data to spatial computation. Among the essentially spatial ones are the vision, the navigation and the motorics. Some others, like the cognition of time, planning, and language, involve such a high degree of spatial computations at each level, that they can easily be seen as originating from the spatial domain.

Apart from the similarities presented in section 4, there are many other spatial borrowings in the structures and computations involved in language. Even the metaphors used to talk about grammar are predominantly spatial. In phonology, an important role is played by the linearity of structures involved and by notions such as distance or adjacency, which are all essentially spatial. In syntax, again, there are syntactic trees, feature geometries, locality relations, movements, unifications, and so on. In semantics, operators have scope, variables get bound, predicates are bounded (e.g., with an upper bound), homogeneous, or scalar; even our intuitions about sets and quantification rely on spatial concepts. This not only illustrates the suitability of spatial relations in the theoretical modeling of grammar, but also suggests the possibility that the target of this modeling borrows a number of essentially spatial computational and structural patterns.

But more importantly, there are similar connections at the level of content. In lexical semantics, one can observe that for instance all prepositions, including the temporal ones, usually stem from words that had spatial meanings (see Lakoff & Johnson 1980 for the lexical semantic, but also for the general cognitive

status of spatial concepts). Other classes show similar, although usually less strong, relatedness to spatial meanings.

A second aspect in favor of the view that language has evolved from spatial cognition capacity comes from brain science. Jarvis (2007) reports about a series of experiments on different animal species with considerably complex computational capacity in the domain of vocal production. Their insights go in the direction of the generalization that brain centers engaged in vocal production in all the examined species are directly related to the brain centers engaged in motoric activities. The authors speculate that the former developed from the latter, by processes of specialization and adaptation to different tasks. Even the shapes and positions of the brain centers involved strongly suggest this conclusion: The brain center engaged in vocal production is either located within that engaged in motorics, or looks like its translated copy (i.e. it is located in the immediate vicinity and has approximately the same shape). It is very difficult to separate the centers engaged in motorics from those involved in spatial cognition. The entire motoric system has developed for functions directly related to space. Every activation of the motoric system has direct effects only in the spatio-temporal domain, and it is in space and time that they lead to the possible further effects, which achieve their actual function. Every possible function of an activity of the motoric system is a function from certain spatial relations. Jarvis reports about experiments designed to exclude the possibility that the activated centers are those engaged in navigation, and involved in the control of the motoric activities. However, navigation is only one specialized type of spatial computation and even a successful isolation of the navigation centers from the experiment does not mean the isolation of all aspects of spatial computation. In fact, conceptual considerations quite strongly suggest that no experiment can investigate the motoric cognition in full isolation from any aspects of spatial cognition, because the former does not exist without the latter. If Jarvis is right that at least some special types of computation, such as those of vocal production and learning, evolved from spatial cognition, by its extension into a particular non-spatial domain, then it is a prominent possibility that a change of the same type, but involving a larger number of domains, lead to the emergence of language. Consecutive development and adaptation of the newly emerged capacity lead to the language faculty as we have it today.

A third argument comes from an intriguing speculation by Krifka (2007), who argues that the subject–predicate, or topic–comment relation, which is central for the human language faculty, originates from the human property of handedness: the specialization of one hand for slow, heavy, rough tasks, and of the other for precise, quick, light tasks. In essence, this is an argument that a core property of language (but also of vision as a figure-ground distinction between the focused and the non-focused part) is argued to originate from an essentially spatially realized property of the cognition involved in motorics.

Pushing the hypothesis further, we may offer the following answers to some interesting questions of evolution of language. Language has evolved from spatial computation. The important changes that channeled this process are the following:

- (i) the extension of the spatial computation into non-spatial domains leading to a domain-general use of the computation;
- (ii) the serialization of the descriptive and geometrical domains, generalizing a sequence that specifies both the description and the geometric properties of a place, i.e. referent;
- (iii) the increasing functionality of a group update of the mental representations involved, mediated (or even pushed) by the development of phonological/ phonetic modules.⁹

Note that (i) and (ii) are well facilitated by the expansion of the number of categories. Principles of economy lead to the development of complex translation and (de-)compression procedures between segments from the discourse and phonological structures. In this view, syntax is to be divided into two systems: one, the 'conceptual syntax', determining the structure of (the concept specifying) the descriptive and geometric components of a discourse referent (close to the notion of conceptual semantic structures of Jackendoff 1999), and the other, the 'translation syntax', specifying the translation and (de-)compression rules between the structures generated by the 'conceptual syntax' and the corresponding phonological structures. Only one of the two, the 'conceptual syntax', is generative (engaged in producing and interpreting structures), while the other is only translational (interface computations). The former developed together with the development of spatial cognition capacity and its extension to other domains, while the latter is part of the development of language, and in particular of phonology.¹⁰

Out of the three important changes above, only the first one, the step of extending the spatial computation to a domain-general use, presents a qualitative change, which might have happened relatively abruptly, i.e. within a relatively short period of time, and a relatively small number of generations. Yet, it is equally possible that this change was gradual, originally involving an import of some pseudo-spatial concepts into the spatial domain, and then of the less spatial ones, until the full disappearance of domain boundaries for the application of spatial computation procedures. In any case, it may be a consequence of a fairly simple genetic change, or possibly just a cultural development: a series of breakthroughs of individuals incorporated into the culture and acquired by the entire community (due to its special organizational properties). The other two

⁹ Originally, a group update could have emerged when the situation in which a change in the immediate context was perceived by a group of individuals was utilized and became a part of the cultural load of a group, triggering some theory of mind effects in the individuals from the group. The next step is the emergence of behavioural strategies to trigger a group update in a controlled fashion, which became more and more systematic, and more and more phonological. The present view has nothing to say about whether this process was pushed by the group update or by some already existing system of vocal production and learning.

¹⁰ 'Conceptual syntax' is involved in the generation of every non-atomic concept. This is not to say that every time we use a complex concept, we generate it from scratch. More frequent complex concepts are stored in the memory, and can be used without generation. Sometimes they are also associated with phonological material, whether single morphemes or complex constructions.

changes are more likely to have been gradual, possibly driven by probabilistic changes and rounds of reanalysis. The serialization of the descriptive and geometric components might have kicked off as a product of the planning capacity, aimed to guarantee efficiency in navigation, which was generalized during a period of time, eventually becoming part of the computational procedure. The group update of the spatial context representation, and later discourse, is another phenomenon which exists in a number of animal species (e.g., the coordinated hunting strategies of some dolphin species or the food-caching jays discussed in Gallistel, in press), but as domain-specific.¹¹ Its extent in the behavior of humans differs from that in animals in a number of properties, such as for instance involving a complex and sophisticated intentionality.

7. Conclusion

The contribution pointed out some striking parallels between cognitive maps and the language faculty, from their architectures, to the role of categories, to reference, but also some interesting differences between the two capacities. The article concentrates on the possible explanations for the presented facts, paying a special attention to the possibility that language has evolved from spatial cognition by the switch of the genuinely spatial computation involved — to a domain-general use. Although the present view of the evolution of language is highly speculative, it presents a hypothesis that deserves serious consideration.

References

- Aggleton, John P. & Malcolm W. Brown. 1999. Episodic memory, amnesia, and the hippocampal-thalamic axis. *Behavioral & Brain Sciences* 22, 425-498.
- Anagnostaras, Stephan G., Greg D. Gale & Michael S. Fanselow. 2001. Hippocampus and contextual fear conditioning: Recent controversies and advances. *Hippocampus* 11, 8-17.
- Arsenijević, Boban & Wolfram Hinzen 2007. Phase derivations versus single unification spaces. Ms., University of Amsterdam & Durham University.
- Bartlett, Mark & Dimitar Kazakov. 2005. The origins of syntax: From navigation to language. *Connection Science* 17(3-4), 271-288.
- Bickerton, Derek. 1998. Catastrophic evolution: the case for a single step from protolanguage to full human language. In James R. Hurford, Michael Studdert-Kennedy & Chris Knight (eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases*, 341-358. Cambridge: Cambridge University Press.

¹¹ One possibility is that some general ability to conduct domain-general computations opened space for the switch of a number of originally domain-specific computational procedures to a domain-general use, including, but not limited to the application of spatial computations and the group update of the spatial context and discourse representations.

- Carstairs–McCarthy, Andrew 1999. *The Origins of Complex Language: An Inquiry into the Evolutionary Beginnings of Sentences, Syllables and Truth* (Oxford Linguistics). Oxford: Oxford University Press.
- Cinque, Guglielmo. 1999. *Adverbs and Functional Heads: A Cross-Linguistic Perspective* (Oxford Studies in Comparative Syntax). Oxford: Oxford University Press.
- Chomsky, Noam. 2001 Derivation by phase. In Michael Kenstowicz (ed.), *Ken Hale: A Life in Language* (Current Studies in Linguistics 36), 1-52. Cambridge, MA: MIT Press.
- Gallistel, Charles R. In press. The foundational abstractions. In Massimo Piattelli–Palmarini, Juan Uriagereka & Pello Salaburu (eds.), *Of Minds and Language: The Basque Country Encounter with Noam Chomsky*. New York: Oxford University Press.
- Gibson, Brett M. & Sara J. Shettleworth. 2003. Competition among spatial cues in a naturalistic food-carrying task. *Learning & Behavior* 31, 143-159.
- Hagoort, Peter & Jos J.A. van Berkum. 2007. Beyond the sentence given. *Philosophical Transactions of the Royal Society, B: Biological Sciences* 362, 801-811.
- Hauser, Marc D., Noam Chomsky & Tecumseh W. Fitch. 2002. The Language Faculty: What is it, who has it, and how did it evolve? *Science* 298, 1569-1579.
- Hermer–Vazquez, Linda, Elisabeth S. Spelke & Alla Katsnelson. 1999. Source of flexibility in human cognition: Dual task studies of space and language. *Cognitive Psychology* 39, 3-36.
- Hinzen, Wolfram & Boban Arsenijević 2007. Recursion as an epiphenomenon. Paper presented at the annual meeting of the *Deutsche Gesellschaft für Sprachwissenschaft (DGfS 2008) – Workshop on Foundations of Language Comparison: Human Universals as Constraints on Language Diversity*, Bamberg. [Universität Bamberg, 27-29 February 2008.]
- Jackendoff, Ray 1999. The representational structures of the language faculty and their interactions. In Colin M. Brown & Peter Hagoort (eds.), *The Neurocognition of Language*, 37-79. Oxford: Oxford University Press.
- Jakab, Robert L. & Csaba Leranthy. 1995. Septum. In George Paxinos (ed.), *The Rat Nervous System* (2nd edn.), 405-442. San Diego, CA: Academic Press.
- Jarvis, Erich D. 2007. A motor theory for vocal learning evolution. Paper presented at the symposium *Birdsong, Speech and Language: Converging Mechanisms*, Utrecht. [Universiteit Utrecht, 21 April 2007.]
- Jayaseelan, K.A. & Hariprasad, M. 2001. Deixis in pronouns and noun phrases. *Linguistic Analysis* 31(1-2), 132-149.
- Kayne, Richard S. 2002. Pronouns and their antecedents. In Samuel David Epstein & T. Daniel Seely (eds.), *Derivation and Explanation in the Minimalist Program* (Generative Syntax 6), 133-166. Malden, MA: Blackwell.
- Krifka, Manfred. 2007. Functional similarities between bimanual coordination and topic/comment structure. *Working Papers of the SFB 632: Interdisciplinary Studies on Information Structure* 8, 61-96. [To be reprinted in Regine Eckardt, Gerhard Jäger & Tonjes Veenstra (eds.), *Variation, Selection, Development: Probing the Evolutionary Model of Language Change* (Trends in

- Linguistics: Studies and Monographs 197). Berlin: Mouton de Gruyter.]
- Lakoff, George & Marc Johnson 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Markus, Etan J., Yu-Lin Qin, Brian Leonard, William E. Skaggs, Bruce L. McNaughton & Carol A. Barnes. 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience* 15, 7079-7094.
- McNay, Anna. 2006. Information structure at the phase level. *Online Proceedings of LingO – The First University of Oxford Postgraduate Conference in Linguistics*.
- Meyer, Patric, Axel Mecklinger, Thomas Grunwald, Jürgen Fell, Christian E. Elger & Angela D. Friederici. 2005. Language processing within the human medial temporal lobe, *Hippocampus* 15(4), 451-459.
- Mizumori, Sheri J.Y., Katharine E. Ragozzino & Brenton G. Cooper. 2000. Location and head direction representation in the dorsal striatum of rats. *Psychobiology* 28, 441-462.
- Nadel, Lynn, Jeffrey Willner & Elizabeth M. Kurz. 1985. Cognitive maps and environmental context. In Peter D. Balsam & Arthur Tomie (eds.), *Context and Learning*, 385-406. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schmajuk, Nestor & Horatiu Voicu 2006. Exploration and navigation using hierarchical cognitive maps. In Michael F. Brown & Robert G. Cook (eds.), *Animal Spatial Cognition: Comparative, Neural & Computational Approaches*. Published online in cooperation with Comparative Cognition Press of the Comparative Cognition Society (<http://www.comparativecognition.org>): <http://www.pigeon.psy.tufts.edu/asc/toc.htm>.
- Opitz Bertram & Angela D. Friederici. 2003. Interactions of the hippocampal system and the prefrontal cortex in learning language-like rules. *NeuroImage* 19, 1730-1737.
- Pearce, John M., Moira Graham, Marc A. Good, Peter M. Jones & Anthony McGregor. 2006. Potentiation, overshadowing, and blocking of spatial learning based on the shape of the environment. *Journal of Experimental Psychology: Animal Behavior Processes* 32, 201-214.
- Rizzi, Luigi. 1997. The fine structure of the left periphery. In Liliane Haegeman (ed.), *Elements of Grammar: Handbook of Generative Syntax* (Kluwer International Handbooks of Linguistics 1), 281-338. Dordrecht: Kluwer.
- Taylor, Holly A., Susan J. Naylor & Nicolas A. Chechile. 1999. Goal-directed influences on the representation of spatial perspective. *Memory and Cognition* 27(2), 309-319.
- Vallortigara, Giorgio, Piero Pagni & Valeria A. Sovrano. 2004. Separate geometric and non-geometric modules for spatial reorientation: Evidence from a lopsided animal brain. *Journal of Cognitive Neuroscience* 16, 390-400.
- Vallortigara, Giorgio, Mario Zanforlin & Giovana Pasti. 1990. Geometric modules in animals' spatial representations: A test with chicks (*Gallus gallus domesticus*). *Journal of Comparative Psychology* 104, 248-254.
- Voicu, Horatiu & Nestor A. Schmajuk. 2001. Three-dimensional cognitive mapping with a neural network. *Robotics and Autonomous Systems* 35, 21-35.
- Wang, Rrenxiao F. & Elisabeth S. Spelke, 2002. Human spatial representation: Insights from animals. *Trends in Cognitive Sciences* 6, 376-382.

Wood, Emma R., Paul A. Dudchenko, R. Jonathan Robitsek & Howard Eichenbaum. 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27, 623-633.

Boban Arsenijević
University of Amsterdam
Department of Theoretical Linguistics
Spuistraat 210
1012 VT Amsterdam
The Netherlands
b.arsenijevic@uva.nl