



Topics in Cognitive Science 3 (2011) 778–795

Copyright © 2011 Cognitive Science Society, Inc. All rights reserved.

ISSN: 1756-8757 print / 1756-8765 online

DOI: 10.1111/j.1756-8765.2011.01151.x

Casimir: An Architecture for Mental Spatial Knowledge Processing

Holger Schultheis, Thomas Barkowsky

Department of Informatics, University of Bremen

Received 4 December 2009; received in revised form 8 October 2010; accepted 12 January 2011

Abstract

Mental spatial knowledge processing often uses spatio-analogical or quasipictorial representation structures such as spatial mental models or mental images. The cognitive architecture Casimir is designed to provide a framework for computationally modeling human spatial knowledge processing relying on these kinds of representation formats. In this article, we present an overview of Casimir and its components. We briefly describe the long-term memory component and the interaction with external diagrammatic representations. Particular emphasis is placed on Casimir's working memory and control mechanisms. Regarding working memory, we describe the conceptual foundations and the processing mechanisms employed in mental spatial reasoning. With respect to control, we explain how it is realized as a distributed, emergent facility within Casimir.

Keywords: Cognitive architectures; Mental imagery; Spatial mental models; Computational modeling; Spatial reasoning; Casimir

1. Introduction

As one aspect of spatial cognition, human spatial knowledge processing plays a central role in cognitive science, because the ability to process spatial information is essential for numerous kinds of cognitive tasks, be they immediately related to the physical environment or be they more abstract problems. Moreover, the forms of representation and the processing mechanisms for spatial knowledge in the human mind are also interesting from a *structural* point of view. Thus, mental spatial knowledge processing is an especially interesting field for computational cognitive modeling approaches.

In this contribution, we present our cognitive architecture Casimir that is dedicated to the explanation of cognitive processes in the context of human spatial knowledge processing.

Correspondence should be sent to Thomas Barkowsky, Department of Informatics, Universität Bremen, Enrique-Schmidt-Str. 5, 28359 Bremen, Germany. E-mail: barkowsky@informatik.uni-bremen.de

The purpose of the development of Casimir is to provide a conceptual framework that is capable of covering as many aspects of spatial mental knowledge processing as possible and to provide a modeling environment that enables the construction of specific models to explain detailed aspects in this cognitive domain. A key assumption underlying the development of Casimir is that mental spatial knowledge processing often uses spatio-analogical representation structures, that is, representations in which at least some of the spatial relations holding between parts of the representation are analogous to the spatial relations that hold between the entities denoted by those parts. Consequently, one main focus of Casimir's development is on devising such structures.

Besides Casimir, there are already a number of well-established cognitive architectures that are used in computational cognitive modeling, for example, ACT-R (Anderson et al., 2004), Soar (Newell, 1990), or EPIC (Kieras & Meyer, 1997). However, all these architectures mainly rely on the formal manipulation of abstract, that is, nonanalogical symbol structures. There are several attempts to extend these architectures such that they are capable of dealing with spatio-analogical representations, and there are also technical systems that employ visuo-spatial types of knowledge processing. However, all existing systems have more or less severe limitations regarding their capability of modeling phenomena of mental spatial knowledge processing. Against this background, this contribution presents a new architecture that is based on spatio-analogical processing capabilities and that aims at being suitable for describing all phenomena of human spatial knowledge processing. In the scope of presenting Casimir and its components, we will also highlight major differences to the already existing approaches.

Aiming at a comprehensive approach to modeling human spatial knowledge processing, Casimir integrates various types of knowledge involved in spatial reasoning tasks (such as topological relations, distance and orientation knowledge, and shape information) together with various types of processes (such as construction of spatial representations and reasoning with default assumptions). To adequately address as many aspects of human spatial knowledge processing as possible, Casimir is developed from an interdisciplinary perspective. Being technically realized in an artificial intelligence context, Casimir integrates research findings from cognitive psychology as well as from cognitive neuroscience. To evaluate the model on its intermediate developmental stages as well as to close gaps for which there are no results available in the literature, series of empirical studies are conducted.

In the following sections, we will give an overview of Casimir's main components and we will describe its essential processing capabilities. Given the space limitations in this contribution, however, we will not be able to go too much into detail with the individual components. A concluding section will refer to open issues and point to future developments.

2. Casimir's structure and components

A schematic overview of Casimir is provided in Fig. 1. As can be seen from the figure, the main parts of Casimir are the *long-term memory*, *working memory*, and a *diagram*

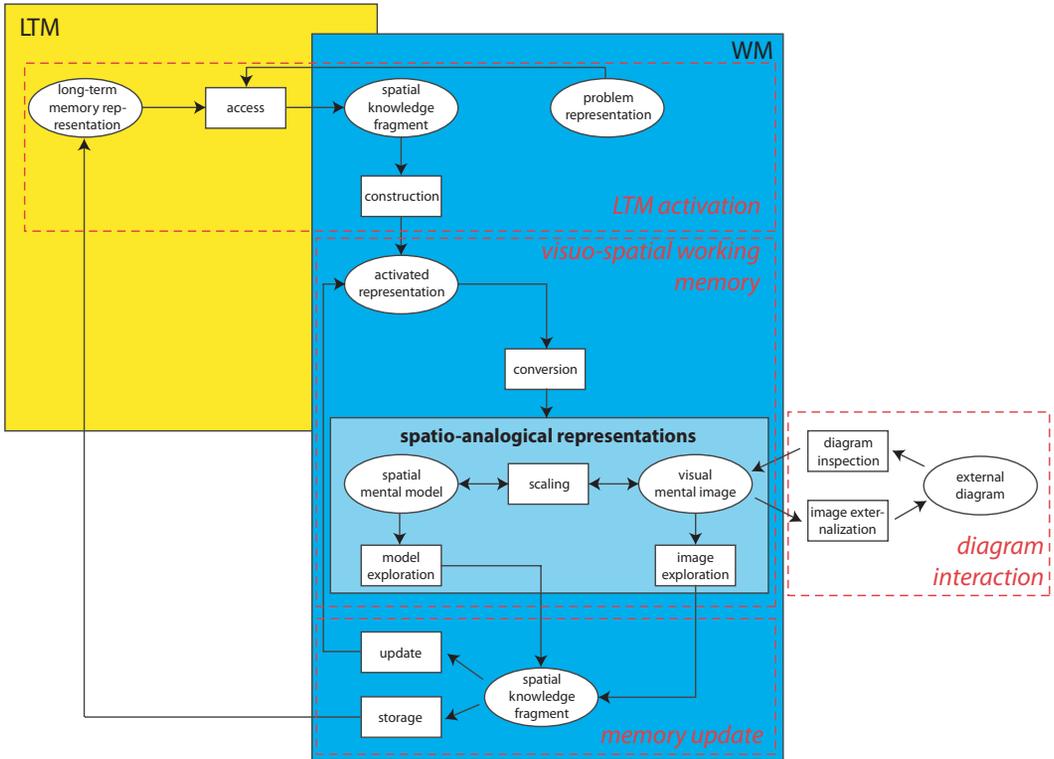


Fig. 1. Overview of the components of the computational cognitive architecture Casimir.

interaction component that realizes the interaction between internal working memory representations and external visualizations. All of these components comprise several representations and processes. The overall working of Casimir arises from the interaction between the different components as well as from the interaction between representations and processes within the components. As a result, proper *control* of the interplay of Casimir's parts is an important fourth component of the architecture. In the remainder of this section, all four components are described in turn.

2.1. Long-term memory

The representation of knowledge in Casimir's long-term memory bears some resemblance to semantic networks as proposed by, for instance, Collins and Loftus (1975). The basic building blocks of the representation are nodes and connections between nodes (Fig. 2A). Nodes represent categories of objects (e.g., city or street) or concrete objects (e.g., Paris or Champs-Élysées) as well as categories of relations (e.g., topological relation or distance relation) or concrete relations (e.g., overlaps or far) between objects. Employing nodes instead of connections to represent relations has the advantage that (a) relations of arbitrary arity (unary, binary, ternary, etc.) can be realized in the same way and (b)

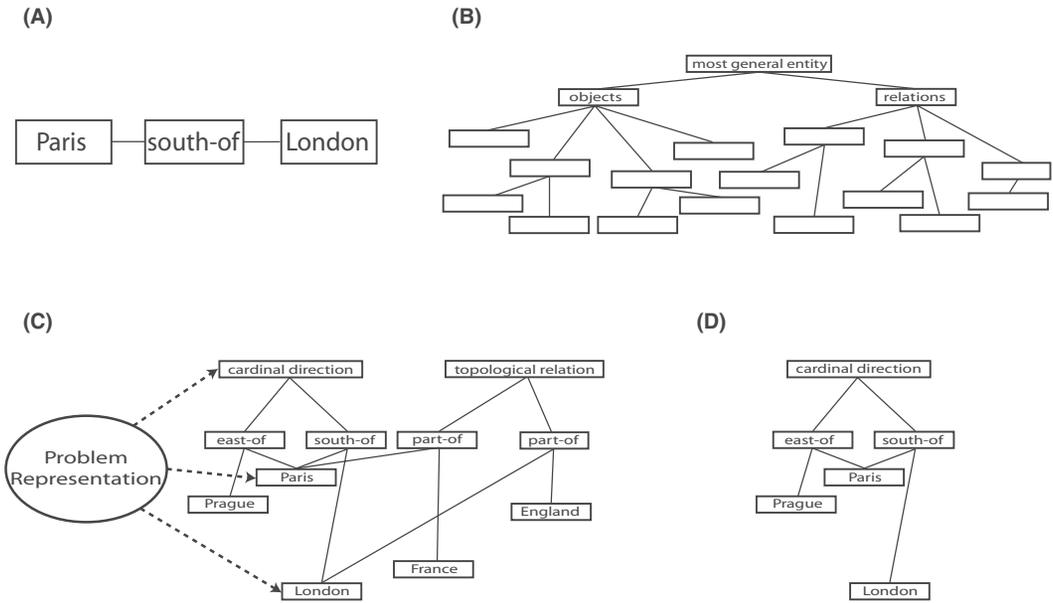


Fig. 2. Casimir’s long-term memory: (A) representation of the fact that Paris is south of London; (B) illustration of the tree-like ontology structure resulting from the use of subsumption relations; (C) on a retrieval request, the problem representation transfers initial activation (dashed arrows) to relevant nodes; (D) a subnet containing one or more knowledge fragments is the result of a retrieval request.

categories of relations can be represented explicitly. Connections signify associations between those entities represented by the nodes that are linked by the connections. Casimir’s long-term memory also contains information about how instances and subclasses of objects and relations are subsumed by (super)classes of objects and relations. Thus, enduring knowledge is represented as a subsumption hierarchy of (classes of) objects and relations that are associatively linked to each other (Fig. 2B).

Access to long-term memory is guided by the current problem representation and realized by spreading activation. The problem representation basically is a specification of the spatial knowledge required in the scope of the spatial task currently pursued and, thus, constitutes the reasoning goal. If, for instance, the current task requires the cardinal direction between Paris and London, the problem representation will comprise “Paris,” “London,” and “cardinal direction.” The entities that are part of the *problem representation* will lead to the initial activation of corresponding nodes in long-term memory if such nodes exist (Fig. 2C). By spreading activation, the initial activation is distributed across the node network according to the associations between the nodes.

The result of spreading activation is a pattern of activation over the nodes in long-term memory. Based on this pattern of activation, the retrieval result is determined as that subnet of the overall memory structure for which the sum of activations of all its nodes is the highest. The retrieval result is further constrained to be a proper subnet of the overall memory structure, that is, a set of nodes that are directly or indirectly associatively linked to each

other (Fig. 2D). This subnet contains one or more *spatial knowledge fragments* (e.g., that Paris is south of London) that are available for further processing in working memory.

Casimir's long-term memory has been shown to be able to account for a number of pertinent human memory effects (Schultheis, Lile, & Barkowsky, 2007b) such as, for instance, the fan effect (Anderson, 1974), effects arising from the hierarchical organization of memory (Collins & Quillian, 1969; Sharifian & Samani, 1997), and the Moses illusion (Erickson & Mattson, 1981; Park & Reder, 2004). In particular, the long-term memory component provides more parsimonious explanations of some of these effects than previously existing models (Schultheis, Barkowsky, & Bertel, 2006).

2.2. Working memory

Knowledge employed during spatial information processing may originate from memory or from current perceptions (e.g., a sketch of a certain spatial situation). In Casimir, knowledge available from these different sources is combined in working memory to establish a representational basis for further exploring and reasoning about the given spatial knowledge in the context of the current task. This section describes the representations and processes of Casimir's working memory. Since the treatment of working memory representations and processes in Casimir crucially hinges on a number of assumptions, we outline these assumptions first.

2.2.1. Basic assumptions

First, it is assumed that spatial information in human working memory is not represented in an abstract, propositional form but that it is rather dealt with in the form of *analogical representations* (e.g., Barsalou, 1999). The notion of analogical representations refers to the structural correspondence between the representations and the states of affairs they represent (Sloman, 1971, 1975), that is, at least some of the relations holding between parts of the representation are analogous to the relations that hold between the entities denoted by those parts. Prevailing examples of spatio-analogical representations in the human mind are *spatial mental models* (Mani & Johnson-Laird, 1982) and *visual mental images* (Shepard & Metzler, 1971) and, accordingly, these representations are important parts of Casimir's working memory.

Two examples of spatio-analogical representation structures, both representing the information "A is south of B" are shown in Fig. 3. The representation structure illustrated in Fig. 3A represents the entities A and B as nodes. A circular list is attached to each of these nodes and each cell of each list represents one cardinal direction. To represent which cardinal direction holds between two entities, the label of the corresponding node is stored in the corresponding cell of the list. Thus, to represent "A is south of B," the label A is stored in the "south" cell of B's list and the label B is stored in the "north" cell of A's list. The structure is analogical in the sense that the neighborhood relations between different cardinal directions is preserved: Those cardinal directions that are adjacent to each other (e.g., north and west) are represented by adjacent cells in the nodes' lists. Fig. 3B displays representation in an image-like format as may be assumed to underlie mental imagery. Spatial relations

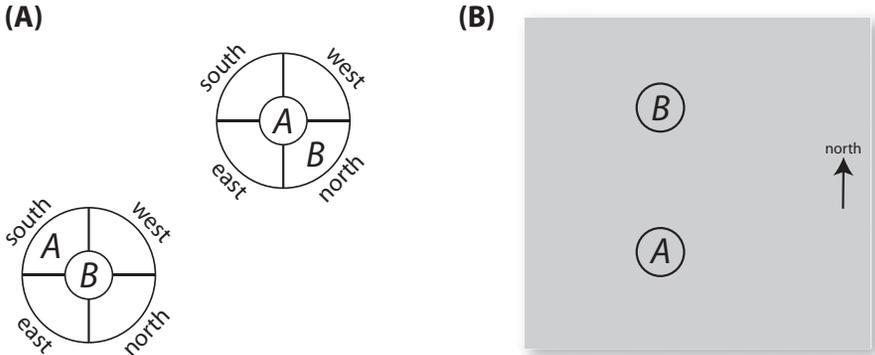


Fig. 3. Two analogical representation structures representing the information “A is south of B”: (A) a more abstract representation; (B) a more image-like representation. The image-like representation also represents not explicitly given information such as, for instance, distance information.

are represented by appropriately placing representations of the involved entities on an imaginary 2D plane (shown in light gray). As the representation in Fig. 3A, the image-like representation preserves the neighborhood structure of cardinal directions and, thus, is analogical.

Second, it is assumed that a strict distinction between spatial and visual aspects in mental spatial knowledge processing is unwarranted. Much research—at least implicitly—seems to assume that spatial information is represented and processed either in a nonmodal (such as spatial mental models) or a visual (such as mental images) format (e.g., Glasgow & Papadias, 1992; Kosslyn & Thompson, 2003; Kozhevnikov, Hegarty, & Mayer, 2002). However, as argued by Schultheis et al. (2007a), a strict distinction into spatial and visual formats does not adequately reflect cognitive reality. Rather than clearly belonging to only one of the two categories, some aspects in mental spatial knowledge processing may better be conceived of as being intermediate to these categories. Following this idea, working memory representations in Casimir are thought to lie along a continuum ranging from the spatial extreme such as simple nonmodal spatial mental models to the visual extreme such as mental images. As a result, an appropriate characterization of representations regarding spatial and visual aspects is a comparative one that allows identifying whether a representation is more spatial or more visual than others. Representations are deemed to be more visual with increasing (a) number of represented relations, (b) number of involved spatial knowledge types (e.g., distance, orientation, and topological knowledge), (c) specificity (i.e., to what extent the represented knowledge unambiguously specifies the location of the represented objects), and (d) exemplarity (i.e., to what extent concrete exemplars instead of prototypes are represented).

Third, it is assumed that the construction of working memory representations follows the principle of cognitive economy (Collins & Quillian, 1969): Representations are built up subject to the constraint that they remain as simple as possible given the current task requirements. Due to this property to adapt to the current task demands the representations are termed *scalable representation structures*.

Scalability shows itself in three ways: For one, the representations are scalable with respect to different knowledge types. Spatial information can convey something about the

orientation between objects (i.e., the direction from one object to the other), the distance between objects, or the topology of objects (i.e., whether and how objects are connected). In Casimir, these different types of knowledge are represented by separate structures and these structures are only built up when the current task involves knowledge of the corresponding type. If, for example, one first learns that Paris is north of Algiers in the scope of a certain task, representation of this fact will involve only that representation structure that exclusively represents orientation knowledge (in this case cardinal directions). Only if one further learns that Paris is in France, a representation structure for topological knowledge will be constructed to supplement the representation structure for orientation knowledge (see Fig. 4). As illustrated by this example, Casimir's working memory representation will scale to those representational facilities that are necessary to account for the presently employed knowledge types.

Furthermore, the representations are scalable in the sense that knowledge type-specific representations adapt to the current task demands. New entities are only added to a representation for knowledge type t when task-relevant information about type t relations between these entities becomes available. In addition, the granularity of type-specific representations, that is, the ability to distinguish between different relations, scales with task demands. Regarding knowledge about cardinal directions, for example, the complexity of the representation varies with the number of cardinal directions that are distinguished. If it is only known that Paris is north of Algiers the representation structure will employ a low granularity in distinguishing only the two cardinal directions north and south. If further processing results in knowledge involving additional cardinal directions (e.g., northeast and west), the representation will scale to be able to distinguish between all known cardinal directions.

A third way in which the representations are scalable is related to the assumption that representations lie along a spatial-visual continuum (see above). As representations scale by including additional knowledge types, additional relations, and by increasing granularity, they tend to move toward the visual extreme of the continuum. Put differently, representations are scalable in the sense that they change their location on the spatial-visual continuum. As a result, a more spatial representation may gradually change toward a full-blown mental image in the course of processing.

A number of behavioral, psychophysiological, and neuroscientific findings support the idea that spatial working memory representations scale with task demands, that is, are only

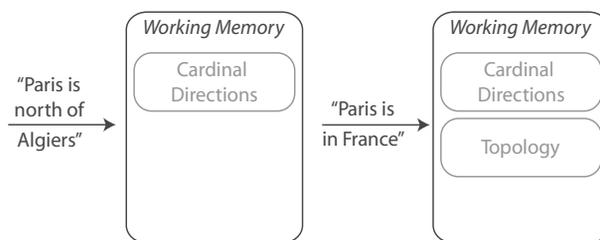


Fig. 4. Scalable representation structures: Knowledge-type specific representation structures are added on demand to working memory.

as complex as necessary for solving the current task. For instance, based on a literature review, Knauff (2009, p. 109) concludes that “A regular reasoning process [...] does not involve visual images but more abstract spatial representations—spatial mental models [...].” In a similar vein, the results reported by Sima, Lindner, Schultheis, and Barkowsky (2010) indicate that people employ more complex (image-like) representations only if the task explicitly requires them to do so while sticking to less complex representations otherwise.

Finally, the existence of preferences in mental spatial knowledge processing is assumed. Available information about a spatial situation not always unambiguously specifies all spatial relations between the involved entities. If, for example, you are told that a certain town *A* is east of another town *B* and that a third town *C* is northeast of *B*, the direction relation between *A* and *C* is undetermined. Depending on the distances between *A*, *B*, and *C*, town *A* may be east, southeast, south, or southwest of town *C* (see Fig. 5A, B, C, and D, respectively). When confronted with such indeterminate spatial information, humans seem to construct a single mental model, called *preferred mental model*, representing only one of the several spatial situations that are in accord with the provided spatial information. In particular, research has shown for a number of indeterminate situations and different knowledge types that the mental model constructed by different persons tends to be the same, that is, that there is a strong preference across people to represent one particular of the several possible situations (e.g., Jahn, Johnson-Laird, & Knauff, 2005; Knauff, Rauh, & Schlieder, 1995; Ragni, Tseden, & Knauff, 2007). Given this evidence, it is assumed that preferences in mental spatial knowledge processing constitute a general means for coping with indeterminacy of available spatial information and, thus, that Casimir’s working memory needs to instantiate such preferences.

Based on these four assumptions, spatial knowledge processing in Casimir is markedly different from existing architectural approaches. For example, though the cognitive architecture Soar has been extended by visual imagery facilities (Lathrop & Laird, 2007), this extension concentrates on *visual* imagery processes (i.e., on visual mental problem solving). Moreover, from a structural point of view, there are no representation structures that are specific for certain types of knowledge, but there is rather one digital image that is used as a universal structure. An alternative approach by Gunzelmann and Lyon (2007) presents a broad theoretical architecture for modeling human spatial information processing facilities based on ACT-R. Similar to Casimir, the intention is to be able to address—at least on a

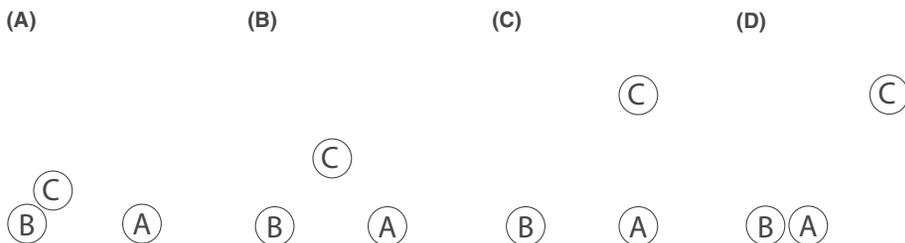


Fig. 5. Four possible representations of the situation “A is east of B” and “C is northeast of B.”

conceptual level—a wide range of phenomena of human spatial knowledge processing. Although this is a promising approach toward a comprehensive architecture for human spatial cognition, relying on ACT-R's buffer concept, the system does not feature specific spatio-analogical representation structures. In sum, there are a number of approaches addressing human spatial knowledge processing from a cognitive architecture-based modeling point of view. However, all these approaches fall short with respect to two central features we realize in Casimir: (a) they do not address the aspect of scalability with respect to working memory representations, that is, they either focus on visual or on spatial mental representations but they do not address the mental capability of using sparse spatial representations for simple spatial questions which may be gradually enriched up to the degree of complexity needed and which may be extended to fully fledged visual mental images and (b) they do not employ analogical representation structures that are specific for the type of visual/spatial knowledge they are used for.

As Shultheis, Bertel, and Barkowsky (2011) have shown, these features are crucial to realize representation structures that are parsimonious, flexible, and that mirror the reasoning preferences observed in human spatial reasoning.

2.2.2. Processing in working memory

As shown in Fig. 1, Casimir's working memory comprises a number of interacting processes and representations. This section describes these representations and processes as well as their interplay.

The basis for constructing spatio-analogical representations in Casimir is the *activated representation*. This representation combines spatial knowledge fragments from several sources to yield a set of fragments most relevant to the current task. One source of fragments is the results from long-term memory retrieval. Importantly, the activated representation may contain results of several different memory retrievals. As described above, each subnet retrieved from long-term memory contains one or more spatial knowledge fragments. On the one hand, not all of the retrieved knowledge fragments may be immediately relevant to the current task. On the other hand, further fragments may be required to appropriately solve the task. Against this background, one purpose of the activated representation is to accumulate the relevant fragments across several memory retrievals. A further source contributing to the activated representation are processes working on already existing spatio-analogical representations: Newly gained information is used to *update* the activated representation (Fig. 6C and D).

Conversion of the spatial knowledge fragments in the activated representation leads to the construction and extension of spatio-analogical representations such as *spatial mental models* and *visual mental images*. In line with the basic assumptions mentioned above, the conversion of the fragments occurs in a piecemeal fashion: Fragment by fragment is added to the spatio-analogical representations such that they scale to the current task demands (Fig. 6A and B). In particular, as part of this *scaling*, the representations may become more or less visual. Though, for illustrative purposes, Fig. 1 only displays representations at the endpoints of the spatial-visual continuum, representations intermediate to these extremes will commonly arise during processing.

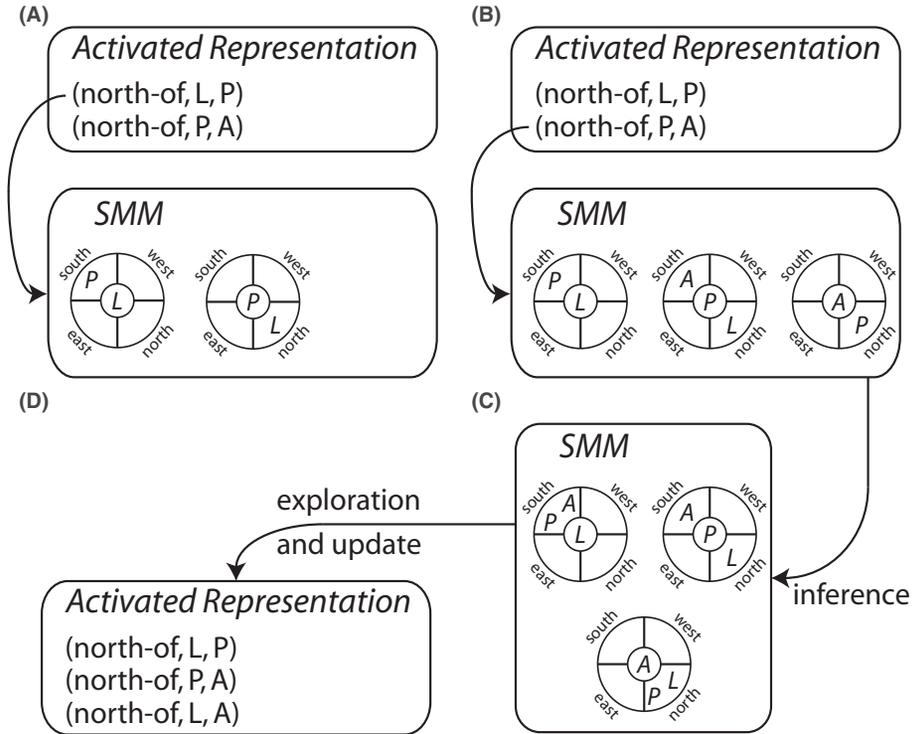


Fig. 6. Working memory processes: (A) and (B) piecemeal conversion from knowledge fragments to spatial mental model; (B) and (C) inference; (C) and (D) mental model exploration and update of activated representation. A, Algiers, L, London, P, Paris, SMM, Spatial Mental Model.

A further important part of processing in Casmir’s working memory is *exploration*. Exploration processes extract spatial information from spatio-analogical representations. This information may then be used to update the activated representation or may be stored more enduringly in long-term memory. As the description so far has focused on spatio-analogical representations that are built up by converting spatial knowledge fragments retrieved from long-term memory, it may not be immediately clear how new knowledge may be gained from exploration of the resulting representations. However, there are at least three ways in which new knowledge arises in the spatio-analogical structures.

First, new spatial information as perceived from the environment may become part of the working memory representation. For instance, a previously unknown spatial relation may become available through a diagram or sketch map drawn by another person. Perceiving this relation may lead to its integration into a spatio-analogical representation such as a mental image.

Second, reasoning about represented knowledge allows inferring new knowledge. If, for instance, you know that London is north of Paris and that Paris is north of Algiers, you may readily infer that London is north of Algiers (assuming you did not know that already, see

Fig. 6B and C). Such inferred knowledge becomes part of the existing representation and is thus available by exploration.

Third, already constructing a spatio-analogical representation may yield new knowledge. Remember that representations can vary with respect to their specificity, that is, with respect to how ambiguously the represented spatial situation is specified. Thus, the activated representation and the constructed/extended spatio-analogical representation may have differing specificity. If the specificity of the spatio-analogical representation is higher than the specificity of the activated representation, the conversion of knowledge fragments may yield, as a “by-product,” additional spatial information. For instance, consider the following—admittedly extreme—case: The next knowledge fragment of the activated representation that is to be converted is “Paris is south of London.” Due to previous demands of the task, the currently employed spatio-analogical representation is a mental image. Thus, the fragment has to be integrated into the existing mental image. Representing a direction relation between two entities in a mental image requires placing them accordingly in a quasi-pictorial two-dimensional plane. Such placement, however, inevitably also specifies a distance between the two entities. Consequently, after conversion of the direction fragment, information about the distance between London and Paris can be gained by exploring the mental image (see also Fig. 3B).

A first imagery component that realizes conversion, scaling, and exploration was implemented to generate mental images from more abstract spatial knowledge (Sima, 2010). Further development of the interplay between more visual and more spatial knowledge representations yielded a new computational theory of mental imagery that is able to explain several imagery phenomena within a more consistent framework than contemporary theories (Sima, 2011).

2.3. Diagram interaction

Performance in spatial tasks can often profit from employing external visualizations such as sketches, maps, or diagrams. One reason for the advantage provided by external visualizations is that their properties beneficially complement the properties of internal mental representations (Tversky, 2005). Accordingly, facilities that realize an interaction between internal representations and external visualizations are part of Casimir.

Two main processes are involved in the interaction. The *image externalization* process prepares selected parts of the internal representation such that they can be transferred to an external diagrammatic medium. The *diagram inspection* process, on the other hand, integrates results from visual perceptual processes into internal spatio-analogical representations.

For both of these processes, it is assumed that comparable basic attentional processes operate in visual perception and on internal spatio-analogical representations. Research on mental images, for example, indicates that inspection and construction of images involves an attentional focus that may scan over or zoom into the image (Kosslyn, 1994). Furthermore, more recent work supports the idea that such a resemblance extends beyond mental images to less visual spatial representations. Based on several experiments, Griffin and Nobre (2003) come to the conclusion that orienting attention to internal representations held

in working memory is similar to orienting attention to perceptual stimuli. Based on this similarity, the realization of image externalization and image inspection relies on a close coupling of attentional orienting to internal and external representations (Engel, Bertel, & Barkowsky, 2005). If, for instance, the attentional focus on the internal representation is currently on the entity representing Paris, the attentional focus on the corresponding external representation will tend to also be on the diagrammatic depiction of Paris and vice versa. Such attentional coupling facilitates the internal-external interaction, because if information is transferred from one representation to the other, attention in the receiving representation will tend to be close to or at the proper insertion point.

2.4. Control

Given the number of components, representations, and processes that are part of Casimir, coherent functioning of the overall architecture requires control. Thus, the realization of control is one of the crucial aspects of Casimir and this section details how control is achieved.

The implementation of control is guided by two main conceptions.

The first conception holds that control is distributed. Functionally, this means that there is no single and unitary component in the architecture that controls the interplay of the other components and their parts. The main reason to avoid a single controlling component is that such a component would constitute a homunculus that leads to infinite regress (Attneave, 1961): If all the responsibility for control is in the hands of a unitary component, the problem of how to realize control in the architecture has not been solved but only transformed to the problem of how to realize the control of the controlling component. In addition to such more theoretical considerations, neuroscientific evidence also suggests that control is distributed. A large number of brain imaging studies have revealed that not one but several different brain areas are involved in realizing control (e.g., Collette & van der Linden, 2002; Egner, Delano, & Hirsch, 2007).

The second conception holds that control is emergent. Viewing control as being emergent assumes that control arises from the interplay of components and parts that only support functions that are not (primarily) control functions (e.g., processing visual information, representing spatial relations, etc.). Note that this is different from the conception that control is distributed. Even if control is distributed, each of the involved components might support a certain well-circumscribed control function. Empirical support for the idea that control may be realized in an emergent way is provided by findings from studies of the functional neuroanatomy of the human brain. For example, Anderson (2007) has shown that activated brain areas and cognitive functions are related to each other by a many-to-many mapping: Each function activates many brain areas and each brain area is activated by many different functions. Thus, at least neuroanatomically, there is evidence that single (control) functions are realized by multiple parts of the human cognitive system. Against this background, the implementation of control in Casimir assumes distribution and emergence both for inter- and intracomponent control.

Fig. 1 illustrates the distribution and emergence of intercomponent control. As can be seen, none of the main components of Casimir, long-term memory, working memory, and the externalization component, is devoted to control all the other components. Furthermore, none of the components' parts primarily serves a control function. Rather, control of the components arises from processes that primarily deal with access to and construction of representations.

Consider, for example, the interaction between working memory and long-term memory. If, during processing, additional knowledge from long-term memory is needed, a request for this knowledge is sent to long-term memory in the form of a specification of the required knowledge (see above). As a result of the request, long-term memory will pass a network of spatial knowledge fragments to working memory. Thus, both modules exchange information that arises from processing within each component. Notably, the passed information is assembled without knowledge about the current state of the other component and, in particular, without the aim to control that state. Nevertheless, the passed information coordinates the interplay of long-term and working memory, because processing within the components partly depends on the received information. Regarding long-term memory, the nodes that are initially activated are determined by the received specification. Regarding working memory, certain processing steps can only occur once relevant knowledge has been retrieved from long-term memory. Thus, the two components control each other by exchanging information, although the exchanged information is not intended and constructed for controlling the other component.

To illustrate the realization of control within components, it is instrumental to consider the control of *spatial reference frames*. The organization of and access to a given spatial representation relies on a reference frame that provides a means for distinguishing parts of the represented space. As such, spatial reference frames are pervasive in any form of reasoning that involves spatial information. Importantly, numerous different reference frames can be distinguished and are, in principle, available for mental spatial knowledge processing. Levinson (2003), for example, distinguishes between absolute, relative, and intrinsic reference frames. Though the availability of multiple reference frames equips spatial knowledge processing with considerable flexibility, it also comes with a cost. In situations in which several mutually exclusive reference frames are available, the set of reference frames has to be controlled such that a single frame is selected for representation and processing.

The work of Levinson and colleagues (Levinson, 2003), for example, has shown that different reference frames are employed when people work on spatial transitive inference problems. In one experiment, participants first viewed two objects placed on a table (e.g., a circle and a square). After viewing the two objects, the participants were asked to turn 180° to face a second table. There were also two objects on the second table (e.g., a circle and a triangle). Importantly, exactly one of the two objects on the second table was identical to one of the objects on the first table. After viewing the objects on the second table, the participants had to turn 180° once more to face the first table again. This table now held only one object. This was the object that had previously been seen by the participants only on the first table (i.e., the square). The task of the participants was to place the object previously seen only on the second table (i.e., the triangle) such that after placement, the relation between

the two objects was consistent with what they had seen before (see Fig. 7 for an illustration of the task).

Participants can employ at least two different reference frames to solve this reasoning task: Both a relative (“The square is left of the circle”) and an absolute frame (“The square is south of the circle”) are possible. The result of the inference varies notably depending on which frame is selected for reasoning, and the observed solutions indicate that different participants use different reference frames (Levinson, 2003).

As spatial reasoning is a core issue addressed by Casimir, it includes mechanisms for realizing the control of reference frame selection. For instance, in the experimental task employed by Levinson (2003), reasoning will involve selecting either an absolute or a relative frame. If the absolute frame is chosen, a representation structure for cardinal directions (see Fig. 3A and 6) is employed for reasoning. If the relative frame is chosen, a representation structure for relative directions is employed. The mechanisms underlying selection are realized as a lateral inhibition-based competition between reference frames. Each reference frame is activated to a certain degree and the frames mutually influence their activation such that each frame excites itself and inhibits all other frames. In a situation where several reference frames are available and one has to be selected, each of the frames will be activated to a certain extent. The frames will then exchange activation until one of the frames is sufficiently more highly activated than all other frames. The most highly activated reference frame is selected for further processing. This realization of control of reference frame selection has been shown to mirror human reference frame selection in various spatial cognition tasks such as spatial term use (Schultheis, 2007b), imaginal perspective taking (Schultheis, 2007a), as well as mental image reinterpretation and spatial reasoning (Schultheis, 2009).

As in the case of intercomponent control, there is no entity that explicitly and primarily controls the selection of reference frames. Instead, control emerges from the interactive information (i.e., activation) exchange between entities of the architecture that do not (primarily) serve control functions. In this sense, also intracomponent control is distributed and emergent in Casimir.

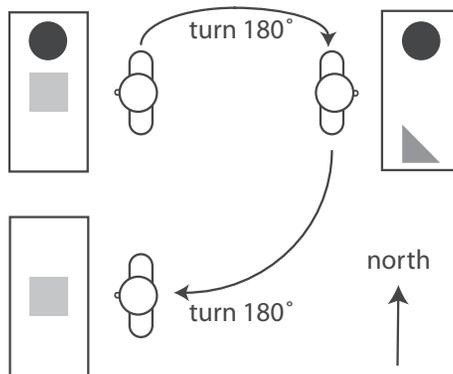


Fig. 7. Illustration of the reasoning task employed by Levinson (2003).

3. Conclusion and future work

In this contribution, we presented an overview of our computational cognitive architecture Casimir that is dedicated to the explanation of human spatial knowledge processing. Casimir provides a comprehensive conceptual framework that comprises crucial functional components involved in human spatial knowledge processing, that is, a long-term memory model that is capable of dealing with large bodies of domain-independent (spatial and non-spatial) knowledge, a working memory that is based on spatio-analogical representation structures (e.g., spatial mental models and visual mental images) which are instantiated in a flexible manner on demand, a diagram interaction component that is capable of externalizing mental images to model the mental interaction with external diagrammatic representations, and facilities for intra- and intercomponent control.

Besides being a comprehensive conceptual framework, Casimir is devised as a computational modeling environment that allows for the realization of specific models that explain dedicated phenomena of mental spatial knowledge processing. The predominant strength of Casimir—compared to existing architectures for cognitive modeling—is its specialization with respect to spatio-analogical representation structures together with the processes working on them. This emphasis on structural modeling aspects turns Casimir into an appropriate approach for explaining a wide range of mental phenomena in the area of spatial cognition. In this regard, Casimir exceeds the modeling capabilities of existing cognitive architectures.

Given the comprehensiveness of the approach pursued with Casimir, however, a number of important issues still await further investigation—in particular, regarding the working memory conception. The most challenging aspect is the structural integration of representations holding different types of spatial knowledge. There are structural dependencies between different types of knowledge (for instance, knowledge about the distance between two entities affects the orientation relations that are possible between them). These mutual dependencies between types of spatial relations are difficult to deal with from an artificial intelligence point of view since they are very much related to the analogical structure of a spatial medium (e.g., the earth's surface or a geographic map) and typically cannot be dealt with on a purely abstract basis.

Apart from this technical perspective, it is currently unclear whether and up to which degree humans are capable of dealing with such dependencies. From research results regarding cognitive maps, for instance, it is known that spatial representations in humans are often inconsistent, distorted, or even contradicting (e.g., Stevens & Coupe, 1978; Tversky, 1981; Tversky & Shiano, 1989). Thus, it may well be that certain dependencies between distance and topology (e.g., if *A* is far from *B*, *A* and *B* are disjoint) or between topology and direction (e.g., *A* inside *B* and *B* is south of *C*, then *A* is south of *C*) are neglected by humans. Further investigations are necessary to determine whether and up to which degree people are able to properly consider such dependencies between knowledge types. The fact that detecting and maintaining such dependencies represents a complex aspect of mental spatial reasoning may well be the reason that the proper consideration of all spatial relations involved (together with their dependencies) in a complex spatial setup requires building up

a fully specified mental image (Schultheis et al., 2007b). In any event, how humans deal with dependencies in integrating different types of spatial knowledge and how this can best be implemented in a cognitive architecture continues to be an important field for future research.

Another field for future investigation is related to the interaction with external representations while mentally dealing with spatial problems. Since cognitive off-loading (i.e., the externalization of spatial mental representations to a visuo-spatial representation medium, e.g., by drawing a sketch) and the visual availability of external spatial representations (e.g., a diagram or a map) may significantly influence the mental reasoning capabilities, this will have a crucial influence on the working memory characteristics, both quantitatively and qualitatively. Especially the extension of available resources for processes of spatial thinking gained by external representations and the more complex control operations necessary to interact with external visuo-spatial media can be expected to form a trade-off that needs to be adequately addressed in future versions of Casimir.

Acknowledgments

The work reported in this article was conducted in the scope of the project R1-[Image-Space] of the Collaborative Research Center SFB/TR 8 Spatial Cognition. Funding by the German Research Foundation (DFG) is gratefully acknowledged.

References

- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6, 451–474.
- Anderson, M. L. (2007). The massive redeployment hypothesis and the functional topography of the brain. *Philosophical Psychology*, 20, 143–174.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036–1060.
- Attneave, F. (1961). In defense of homunculi. In W. A. Rosenblith (Ed.), *Sensory communication: Contributions to the symposium on principles of sensory communication* (pp. 777–782). Cambridge, MA: MIT Press.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–609.
- Collette, F., & van der Linden, M. (2002). Brain imaging of the central executive component of working memory. *Neuroscience and Biobehavioral Reviews*, 26, 105–125.
- Collins, A., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review*, 82(6), 407–428.
- Collins, A., & Quillian, M. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Behavior*, 8, 240–248.
- Egner, T., Delano, M., & Hirsch, J. (2007). Separate conflict-specific cognitive control mechanisms in the human brain. *NeuroImage*, 35, 940–948.
- Engel, D., Bertel, S., & Barkowsky, T. (2005). Spatial principles in control of focus in reasoning with mental representations, images, and diagrams. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, & T. Barkowsky

- (Eds.), *Spatial Cognition IV – reasoning, action, interaction. International Conference Spatial Cognition 2004* (pp. 187–209). Berlin: Springer.
- Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, 20, 540–551.
- Glasgow, J., & Papadias, D. (1992). Computational imagery. *Cognitive Science*, 16, 355–394.
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, 15, 1176–1194.
- Gunzelmann, G., & Lyon, D. R. (2007). Mechanisms of human spatial competence. In T. Barkowsky, M. Knauff, G. Ligozat, & D. Montello (Eds.), *Spatial Cognition V: reasoning, action, interaction. Lecture Notes in Artificial Intelligence 4387* (pp. 288–307). Berlin: Springer.
- Jahn, G., Johnson-Laird, P. N., & Knauff, M. (2005). Reasoning about consistency with spatial mental models: Hidden and obvious indeterminacy in spatial descriptions. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, & T. Barkowsky (Eds.), *Spatial Cognition IV. reasoning, action, interaction* (pp. 165–180). Berlin: Springer.
- Kieras, D. E., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12, 391–438.
- Knauff, M. (2009). A neuro-cognitive theory of deductive relational reasoning with mental models and visual images. *Spatial Cognition & Computation*, 9, 109–137.
- Knauff, M., Rauh, R., & Schlieder, C. (1995). Preferred mental models in qualitative spatial reasoning: A cognitive assessment of Allen's calculus. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 200–205). Mahwah, NJ: Erlbaum.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Kosslyn, M. S., & Thompson, L. W. (2003). When is early visual cortex activated during visual mental imagery? *Psychological Bulletin*, 129, 723–746.
- Kozhevnikov, M., Hegarty, M., & Mayer, R. E. (2002). Revisiting the visualizer-verbalizer dimension: Evidence for two types of visualizers. *Cognition & Instruction*, 20, 47–78.
- Lathrop, S. D., & Laird, J. E. (2007). Towards incorporating visual imagery into a cognitive architecture. *Proceedings of the Eighth International Conference on Cognitive Modeling*. Ann Arbor, MI.
- Levinson, S. C. (2003). *Space in language and cognition*. Cambridge, UK: Cambridge University Press.
- Mani, K., & Johnson-Laird, P. N. (1982). The mental representation of spatial descriptions. *Memory & Cognition*, 10, 181–187.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Park, H., & Reder, L. M. (2004). Moses illusion: Implication for human cognition. In R. F. Pohl (Ed.), *Cognitive illusions* (pp. 275–291). Hove: Psychology Press.
- Ragni, M., Tsenden, B., & Knauff, M. (2007). Cross-cultural similarities in topological reasoning. In S. Winter, B. Kuipers, M. Duckham, & L. Kulik (Eds.), *Proceedings of 9th International Conference on Spatial Information Theory, COSIT 2007*, (pp. 32–46). New York: Springer.
- Schultheis, H. (2007a). A control perspective on imaginal perspective taking. In R. L. Lewis, T. A. Polk, & J. E. Laird (Eds.), *Proceedings of the 8th International Conference on Cognitive Modeling (ICCM), 2007, Ann Arbor, MI* (pp. 19–24).
- Schultheis, H. (2007b). A computational model of control mechanisms in spatial term use. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Meeting of the Cognitive Science Society* (pp. 1467–1472). Austin, TX: Cognitive Science Society.
- Schultheis, H. (2009). *Computational cognitive modeling of control in spatial cognition*. Lengerich, Germany: Pabst Science Publishers.
- Schultheis, H., Barkowsky, T., & Bertel, S. (2006). LTM^C — an improved long-term memory for cognitive architectures. In D. Fum, F. del Missier, & A. Stocco (Eds.), *Proceedings of the 7th International Conference on Cognitive Modeling (ICCM), 2006, Trieste* (274–279). Trieste: Edizioni Goliardiche.
- Schultheis, H., Bertel, S., & Barkowsky, T. (2011). *Modeling mental spatial reasoning about cardinal directions*. Manuscript in revision.

- Schultheis, H., Bertel, S., Barkowsky, T., & Seifert, I. (2007a). The spatial and the visual in mental spatial reasoning: An ill-posed distinction. In T. Barkowsky, M. Knauff, G. Ligozat, & D. R. Montello (Eds.), *Spatial Cognition V – reasoning, action, interaction* (pp. 191–209). Berlin: Springer.
- Schultheis, H., Lile, S., & Barkowsky, T. (2007b). Extending ACT-R's memory capabilities. In S. Vosniadou, D. Kayser, & A. Protopapas (Eds.), *Proceedings of EuroCogSci'07: The European Cognitive Science Conference* (pp. 758–763). Mahwah, NJ: LEA.
- Sharifian, F., & Samani, R. (1997). Hierarchical spreading of activation. In F. Sharifian (Ed.), *Proceedings of the Conference on Language, Cognition, and Interpretation* (pp. 1–10). Isfahan, Iran: IAU Press.
- Shepard, N. R., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*, 701–703.
- Sima, J. F. (2010). A computational account of complex mental construction. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 217–225). Philadelphia, PA: Drexel University.
- Sima, J. F. (2011). The nature of mental images – an integrative computational theory. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*.
- Sima, J. F., Lindner, M., Schultheis, H., & Barkowsky, T. (2010). Eye movements reflect reasoning with mental images but do not with mental models in orientation knowledge tasks. In C. Hölscher, T. Shipley, M. Olivetti-Belardinelli, J. Bateman, & N. Newcombe (Eds.), *Spatial Cognition VII* (pp. 248–261). Berlin: Springer.
- Sloman, A. (1971). Interactions between philosophy and artificial intelligence. *Artificial Intelligence*, *2*, 209–225.
- Sloman, A. (1975). Afterthoughts on analogical representations. In B. L. Nash-Webber & R. Schank (Eds.), *Proc. theoretical issues in natural language processing* (pp. 164–168). Cambridge, MA: Association for Computer Linguistics.
- Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, *10*, 422–437.
- Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, *13*, 407–433.
- Tversky, B. (2005). Functional significance of visuospatial representations. In P. Shah & A. Miyake (Eds.), *Handbook of higher-level visuospatial thinking* (pp. 1–34). Cambridge, England: Cambridge University Press.
- Tversky, B., & Shiano, D. J. (1989). Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of experimental psychology: General*, *118*(4), 387–398.