ON THE DUALITY AND ON THE INTEGRATION OF PROPOSITIONAL AND SPATIAL REPRESENTATIONS

Christian Freksa and Thomas Barkowsky,
University of Hamburg, Department for Informatics and Cognitive Science Program

INTRODUCTION

A substantial amount of work in cognitive psychology concerns the question how many kinds of mental representations there are and how many representations are required, with answers ranging from zero (e.g. Brooks, 1991) to unlimited (e.g. Sloman, 1985). In the imagery debate, a controversy arose over the question whether mental representations have a propositional format (Pylyshyn, 1973), a pictorial format (Kosslyn and Pomerantz, 1977; Kosslyn, 1980), or whether they make use of both formats (Paivio, 1986; Habel *et al.*, 1995). Johnson-Laird's (1983) *mental models* provoke the question of their relationship to mental images.

External vs. internal representations

The present paper takes a computer science perspective at ways of representing knowledge: rather than interpreting empirical evidence for or against certain representation formats, we address the issue from an architect's point of view, discussing and designing knowledge representation and processing schemes for solving certain tasks (compare Sloman, 1994). Consequently, the article focuses on *external* representations of knowledge (for example on a sheet of paper) rather than on *internal* (i.e. mental) representations. However, external representations used by humans can be viewed as part of the cognitive system as a whole which in the sense of cognitive externalism (Peterson, 1996) comprises both, the internal representation and externalized components thereof. With this view on the internal / external distinction we have a means of experimenting with external representations and at the same time (non-invasively) exploring properties of internal mental representations. It will be

necessary to give a precise meaning to the notions relevant to this discussion, and we will therefore attempt to clarify some of these notions.

The study concerns visually mediated spatial knowledge. It is motivated by investigations on eye movements and visual perception (cf. Zangemeister *et al.*, 1996); by neuropsychological results on the representation and processing of visual information in vertebrates (cf. Davis, 1997); by attempts to exploit the spatial structure of diagrams for human-computer interfaces (e.g. Wahlster, 1994); and by attempts to exploit the spatial structure of diagrams for diagrammatic reasoning (cf. Schlieder, 1997; Barkowsky *et al.,* 1996; Glasgow *et al.*, 1995). The goal of this enterprise is to conduct the debate on spatial representations on a technical level of structures and processes allowing for an operational distinction between propositional and spatial use of knowledge. Applications of the study are seen in the design and development of systems for diagrammatic reasoning, in the representation of spatial knowledge for robots, and in the spatial representation of non-spatial knowledge for analogical reasoning.

Propositional vs. spatial formats: expressive power

When discussing the mutual merits of diagrams and textual representations for reasoning and problem solving with logicians, we received comments like "Before comparing diagrams with sets of propositions, first formalize your diagram in terms of well-defined formal expressions; then we will know what you are talking about." This statement expresses a presupposition (a) that propositional representations are at least as expressive as spatial representations and (b) that everything useful in a diagram can be translated into formal expressions without loss of utility. Furthermore the belief in the convertibility of diagrammatic into textual formats suggests that the transformation – when needed in a given situation – computationally come for free.

In this paper, we will challenge this view. Although we agree that *in principle* all facts and relations expressed in a diagram can be put into words and formal expressions, we will emphasize that diagrams and other images have a potential of being used in a way different from textual structures. This potential is not fully preserved in the translation of diagrams to text (cf. Myers and Konolige, 1995). The paper will elaborate on the issue of comparing images with textually expressed knowledge.

Before going into the comparison, however, we will introduce the notions *propositional* and *spatial* and we propose to use these notions in terms of knowledge processing criteria. More specifically, we present the two notions as two different *aspects* of representations (rather than two different *kinds* of representations, compare Rehkämper, 1995) and we illustrate, how these aspects may or may not be exploited in reasoning.

For the purposes of this paper, we will employ concrete examples from the visual domain. We will not answer the question whether and to what extent the notion *spatial* should be used outside a physical framework (for example for a computer representation). As is common in

the discussion of artificial intelligence algorithms, we will make a distinction between *domain knowledge* and *control knowledge*. This distinction, however, is only sensible with respect to a specific processing level within a complex system. Domain knowledge from one processing phase may be used as control knowledge in another processing phase; nevertheless it is helpful to conceptually distinguish the different *roles* knowledge can play on a given level. This distinction is related to *content* and *structure* of a representation system and includes its processes.

HOW CAN WE OPERATIONALIZE THE NOTIONS *PROPOSITIONAL* AND *SPATIAL?*

What types of instruments does a photographer, a painter, or a draftsperson need in order to go after his or her business, and what types of instruments does a reporter, an author, or a poet require for his or her work? Besides the idea or theme for their piece of work, the first group of professionals requires a notion of space and a spatial reference system or perspective for dealing with this space; the second group requires a notion of language and a system of phrases to express meanings by words. Both groups may use a spatial medium – typically a flat sheet of paper – and some sort of pen to register their thoughts. However, despite the similarity of the physical activities of the painter and the writer, the spatiality of the medium 'paper' plays a rather different role for the two. While the spatial arrangement of the pen marks on the paper is crucial for the painter and his product, it only plays a secondary role for the writer. The phrases used by the writer derive their primary significance from the meanings behind the word symbols, not from their shape or their position on the paper.

In the cognitive science literature, the term 'pictorial' representation is employed for the graphical use of the medium paper and the term 'propositional' representation for the linguistic use of the same medium. These absolute denotations suggest that representations be either of the one kind or of the other kind, particularly, when 'pictorial' and 'propositional' are used as antipodes. As we are concerned only with the spatial aspect of pictorial representations in this paper (and not with colors, for example), we speak of 'spatial' representations here. We speak of 'propositional' representations when we refer to the linguistic use of a spatial medium.

In the following discussion we will argue that 'propositional' and 'spatial' merely denote two different aspects of representations and that these aspects are not necessarily in conflict. To illustrate this point, let us consider the situation where graphical and linguistic aspects meet; they meet, for example, when character symbols are presented in certain fonts and when texts are formatted in certain layouts. Usually, we will consider a text presented in different fonts or layouts as the same text, thus implicitly referring to its propositional content. On the other hand, we will not consider two paintings the same, when the depicted entities are the same, but shape or location differs; thus, we implicitly take into account spatial content.

As long as representations employ only one mode – propositional or spatial – to convey their message, we may use these notions in an absolute and opposite sense. But there are representations that exploit both modes simultaneously. Examples are texts in which certain words or phrases are emphasized or otherwise distinguished by font or layout; poems, part of whose message is coded in the arrangement of words and phrases much like in pictures; geographic maps which would be useless if they contained no words or other symbols or if the symbols were arranged as in a regular text. These representations should neither be denoted as purely 'propositional' nor as purely 'spatial' representations nor do they employ multiple encoding of the same relations (cf. Paivio, 1986). Rather, we should speak of 'spatially represented aspects' and of 'propositionally represented aspects' within the same representation (cf. Palmer, 1978).

What, *where* and *how* in a visual setting

When we perceive some object in a visual scene (e.g. a picture, a text, or a geographic map), two types of knowledge are accessible: *what* the object is and *where* the object is located (cf. Farah *et al.*, 1988; Landau and Jackendoff, 1993). The *what*-aspect may be considered the essence of the scene (what is present in the scene? what is written in the text? what do we see in the picture?) while the *where*-aspect merely serves as the carrier of that essence. However, this way of regarding the relationship between a viewer and a scene may be nothing more than a coding convention: The canvas provides locations; the oil painting it carries is the piece of art.

It is also conceivable that scenes are depicted the other way around: the items in the scene serve as carriers, their location and arrangement provide the essence. Important is that we are not able to focus on both aspects simultaneously: either we focus on *what* is in the scene or *where* it is. Getting the full appreciation of *how* something is presented (including what and where) requires us to defocus from the individual aspects.

In any perception event, we may be interested (1) only in *what* we see, (2) only in *where* we see it, or (3) *what* we see *where*. For example, in bird watching (or in formal reasoning), the crucial question may be *what kind of bird is it* (or *what is the conclusion)?* We may ignore, *how* and *where* we found the bird or the conclusion. On the other hand, when we search for a certain object that cannot be mixed up with something else, like the car key or our glasses, we only are interested in *where* they are, the question of recognizing them will be ignored. However, in many situations we are concerned with the interaction of both aspects: To find out, what kind of bird is singing, we will search at appropriate locations. Finding a conclusion in a formal argument will be awfully difficult, when the facts and assumptions are randomly arranged. Finding keys or glasses is almost impossible, when we have no idea about their size and shape and about the possible actions which may have placed them from where they were

before to where they are now. In other words, *what* and *where* complement each other, but we may not always need both.

The explicit / implicit distinction

Specific knowledge about facts or relations, which can be directly read off a representation and used as is, is called *explicit* knowledge. Knowledge that is buried in facts and/or rules, which must be processed in order to be used, is called *implicit* knowledge. For example, when we have a rule stating that the square of a number is obtained by multiplying that number by itself, then we have *explicit* knowledge about what we mean by the square of a number (this knowledge is explicitly stated in the rule) and we have *implicit* knowledge about the value of the number 3 squared (as we can derive this value by applying the rule). When we apply the rule by multiplying 3 by itself, we make implicit knowledge explicit.

We can now raise the question if knowledge represented by a visual scene is explicit or implicit knowledge. Let us first consider the *what*-aspect of the knowledge: Suppose our scene is a photograph, i.e. a purely optical mapping from some environment to a film. Assuming 'true color' reproduction, at each location of the photograph knowledge about color is explicitly represented (it can be directly read off the photograph); other information (esp. what the objects in the scene are, what types of objects are depicted, etc.) must be derived through interpretation processes and is represented implicitly, at best.

Let us now consider the *where*-aspect. The location of the color spots on the photograph with respect to the colored objects is explicitly given. The location of the corresponding colors in the depicted scene, however, is implicitly represented and must be derived through an interpretation of the photograph. Consider now an intentionally constructed diagram of geometric entities instead of a photograph. Suppose we construct an equilateral triangle of a certain basis length using ruler and compass. Again, the locations of the lines on the diagram are given explicitly as constructed. But the angles of the triangle also are represented explicitly in the diagram, although they were not specified in the prescription of the construction; they can be read directly off the diagram. Thus, by depicting a triangle that was specified only by the length of its sides, knowledge about the angles of the triangle has been made explicit.

The knowledge about the angles in the triangle was specified implicitly through the specification of the length of the sides of the triangle together with the laws of plane geometry. How has this knowledge been made explicit, as we have not explicitly applied laws of plane geometry? The laws of plane geometry are implicitly preserved on a flat sheet of paper; the construction of the triangle under the constraints of these implicit laws effectively has made the angles of the triangle explicit.

An interesting aspect of this method of making implicit knowledge explicit is that the representation medium in effect prevents the graphical depiction of a triangle whose angles are underspecified or implicitly represented and forces the construction process to give the triangle

specific and explicit angles. Another interesting aspect is that the knowledge becomes explicit through a transformation from the *what* in the specification "equilateral triangle" to the *where* in the diagram. The transformation itself appears as a simple writing process rather than a typical interpretation process; the interpretation happens almost *by magic* through the structure of the representation medium.

Propositional and spatial channels of a representation

The propositional / spatial distinction used in this paper has a strong correspondence to the *what / where* distinction in visual perception. There are two main ways in which these two aspects can be simultaneously exploited in representations of 'real world situations' in which both, semantic and location information are significant: (1) the *what*-aspect of the represented situation can be represented propositionally and the *where*-aspect can be represented spatially; (2) the *what*-aspect of the represented situation can be (partially) coded spatially and spatial location can be coded propositionally. Thus, the propositional and the spatial aspects can be seen in a dual relationship to one another.

Geographic maps are an example for the first type of division of representational power; *what* is depicted in a map is coded propositionally by symbols whose meaning is specified in the map legend while the spatial locations of entities in the situation are coded through spatial locations on the map. The reference sections of city maps are examples for the second type of division of representational power: *what* is represented is spatially grouped either by similarity of name (alphabetically) or by similarity of function (e.g. hotels, motels, bed & breakfast; museums, movie theaters, opera houses; … grouped in sections) while their location is indicated propositionally through map coordinates.

When aspects other than semantic or pragmatic significance and spatial location are relevant, either the propositional or the spatial 'channel' or both can be used accordingly to convey those aspects.

With the 'channel' view of propositional and spatial information, two scenes (or images of scenes) with equal propositional content may be considered propositionally equivalent independent of the spatial arrangement of the propositions and two scenes with equal spatial arrangements can be considered spatially equivalent independent of the features represented in the scene.

Relation to syntax and semantics

In 'regular' propositional representations like natural or formal language expressions, spatial arrangement is used to structure propositional content. To a large extent, syntax serves to

separate meaningful legal expressions from illegal expressions; to a lesser extent syntax serves to provide an additional dimension for differentiating different meaningful information.

How about typical spatial representations? In principle, objects may be arbitrarily located, but the structure of the represented world determines which arrangements of objects and features make sense. In particular, an entity representing a central part of an object will be surrounded by entities belonging to the same object, where an 'object' can be a physical object or a connected group of entities, depending on the information conveyed by the spatial channel. Thus, freedom with respect to spatial arrangement may be offset by semantic constraints.

Coexistence of propositional and spatial aspects

In folk conversations about knowledge representation, we find statements like 'a natural language description *is propositional'* or 'a city map *is pictorial'* (or *'spatial'*) suggesting that *propositional* and *pictorial* (or *spatial*) are mutually exclusive absolute properties of the respective entities. The way in which the notions propositional and spatial have been introduced in this paper does not create a conflict between the two ways of representing knowledge; they can coexist peacefully. Thus, saying that 'a natural language description *is propositional'* may express that the relevant information represented in a natural language description is coded in a propositional way, i.e. the meaning of the description is conveyed through the semantics of the propositions rather than through the spatial arrangement of these propositions (of course we all know that for natural language this is true only as a rough approximation). Conversely, 'a city map *is spatial'* may mean that much of the knowledge contained in the map is coded spatially; this means that the knowledge contained in the map is conveyed to a large extent through the spatial arrangement (=syntax) of whatever is placed on the map.

In summary, there is a full range of possibilities to represent and convey knowledge about a domain: (1) Exclusively propositionally. This is particularly suitable, when the knowledge consists of a collection of eternally true facts and no two facts have to be considered more related to one another than others. (2) Exclusively spatially. This is suitable particularly, when the knowledge consists of a large number of interrelated relations and many of these relations can be simultaneously captured by the arrangement of a comparatively small number of objects; this is the case in photographs and maps, where an enormous number of spatial relations between objects in the domain is represented by the arrangement of a comparatively small number of objects. (3) A combination of (1) and (2): (3a) a natural language description of facts about the world in which the facts are ordered in the sequence of occurrence or detectability would be a spatially lean propositional representation; (3b) sheet music, where the relative locations of musical notes reflect both the occurrence of the corresponding sounds in time and their pitch would be a spatially rich propositional representation; (3c) a photograph containing an arrow annotated "that's me" or some other propositional symbols associated

would be a propositionally lean pictorial representation; while (3d) a comprehensive city map full of all sorts of symbols would be a propositionally rich spatial representation.

Interference of propositional and spatial modes of representation

We have shown that propositional and spatial representations are not *per se* in conflict. However, when we consider simultaneous representation of both aspects within one medium (i.e. 'single multimodal coding' of multiple aspects), conflicts may occur. The city map example shows that the 2-dimensional spatial background is used as a medium for carrying symbols. As long as the propositional content of the symbols is not coded by using the two spatial dimensions reflecting the spatial arrangements in the domain – for example by using color information – there is no conflict. However, when the representation of the propositional content of the symbols makes use of the same dimensions that carry the spatial information (e.g. shape and / or spatial extension), then we may get a conflict. For example, in small-scale road maps, width of the road symbol denotes type of road (rather than width of road) and thus, the propositional content of the map partially overwrites and falsifies its spatial content.

Do we also find the dual situation, i.e., do we have cases in which propositional content of a representation is modified by demands of the spatial representation? This is certainly the case. When we put prose into a poem, we may have to modify propositions of our original text to meet morphological and syntactic requirements of the lyric form of communication.

Nevertheless, it appears less typical or less likely that propositional content of a representation is overwritten by spatial content, for the following reason: When we think of propositions, we usually think of two-valued entities: propositions either hold, or they do not hold. Thus, overwriting a proposition would mean (a) to change its truth-value or (b) to make it meaningless as a proposition. Both options are none-or-all events.

Spatial knowledge usually is interpreted differently: a road depicted at roughly the correct location would still be considered this road rather than something completely different. Thus, the modification of spatial representations through the modification of propositional content typically may result in small gradual rather than radical changes; the spatial channel exhibits certain robustness (when interpreted appropriately). Thus, besides employing different channels, propositional and spatial knowledge are interpreted in different ways and modes.

Data structures and the role of interpretation processes

A static representation alone does not tell us whether it represents knowledge propositionally, spatially, or both. Which types of knowledge representation are involved depends on the processes used for accessing and interpreting the knowledge (cf. Palmer, 1978; Rumelhart and Norman, 1985; Furbach *et al.,* 1985). Processes that refer to explicitly given semantic

knowledge make use of propositional content while processes that exploit the spatial arrangement make use of the spatial content of the representation.

In visual inspection and interpretation processes, the propositional and spatial aspects work together to efficiently identify objects or 'understand' scenes. A priori expectations regarding objects to be seen focus attention to a location where this expectation is likely to be verified or refuted. Depending on the outcome of the verification procedure, a new expectation can be generated, etc. (Noton and Stark, 1971; Zangemeister *et al.*, 1996; Egner, 1994). In effect, the interaction of both aspects results in efficient recognition due to rapid restriction of the search space.

To computer scientists, this search principle is well known for data structures in internal computer representations. The main principle exploited in data structures is the arrangement of data according to a content-related ordering rule – although this arrangement is not necessarily spatial.

Dimensionality of propositional and spatial representations

Propositional and spatial representations complement each other in one more interesting way: the conceptual space of propositions is not restricted with respect to dimensionality while physical space is restricted to at most three dimensions which are strongly related to one another. As a consequence, propositional representations are suited to represent arbitrary sets of assertions, whether they can hold in the real world or not. This freedom has its prize: arbitrary inference processes may be carried out but they may be very complex and slow.

On the other hand, spatial representations are severely restricted. Only comparatively simple situations can be represented by the low dimensionality of physical space. The structural interdependence of the dimensions is so strong that the representation of few assertions *automatically* implies other facts (compare the geometry example in the beginning of this chapter). This is one of the strengths of this representational form: when we have a representation that complies with the constraints of the represented situation we may get many inferences rather inexpensively.

The easiest way to comply with the constraints of the represented situation is by having a very simple correspondence relation between the inherent structure of the represented world and the structure of the representing world (cf. Palmer, 1978). In the case of mapping a two-dimensional flat surface to another two-dimensional flat surface by an identity relation, for example, we can directly read off inferences corresponding to arbitrary operations in the represented world by applying identical operations in the representing world.

But this example also exhibits a problem with too simple correspondence relations: when we reconstruct the represented world instead of describing it, it is easy to 'reason' *in* the world, but it may be hard to reason *about* it. In other words, we can *simulate* concrete events, but we may

not be able to *abstract* to more general classes of events. This points to one of the major challenges to spatial representations.

How Can We Compare the Aspects *Propositional* and *Spatial?*

In the first part of this paper, we have shown two channels through which knowledge representations convey knowledge: one through which explicit knowledge is conveyed and one through which implicit knowledge is conveyed. As logicians we characterize implicit knowledge by making it explicit by means of propositions, e.g. by formally describing the knowledge structures. As cognitive beings, however, we are able to use implicit structures without making them explicit (cf. Karmiloff-Smith, 1996).

Humans can use implicit structures through implicit control. For example, we use the physical temporal sequence of events to structure our memory and to explain causal chains. We do not have to attach time-stamps or temporal relations to the representations of the events. Appropriate memory structures are able to preserve the sequence implicitly and by simply following this structure, temporal sequence underlying the causal chain is recovered. Using temporal structure in this way does not require understanding it. The great advantage of using the implicit structure of time is that no or much fewer alternatives must be considered and fewer decisions must be taken than by using explicit knowledge. For many purposes, this is all time is needed for.

Why do we have to consider fewer alternatives and have to take fewer decisions for implicit structures? Representing time explicitly and giving up or ignoring implicit structure means that we must reconstruct a given or potential sequence of events by comparing time stamps or by inspecting temporal relations and throw out those which are inconsistent with the natural ordering of temporal events. A process using physical time as a 'vehicle' to ride through a linear memory structure, on the other hand, does not have to analyze and throw out inconsistent orderings. Inconsistent temporal orderings do not exist in the real world. They only come into the picture through our abstractions of the world. By saying 'A is earlier than B', we introduce an idea allowing for the theoretical possibility that A could take place simultaneous with or later than B. By leaving things implicit, such possibilities do not come up and do not have to be dealt with.

This example shows two points. (1) Implicitly represented knowledge does not always have to be made explicit to be used; it can remain implicit when there are appropriate processes which make use of the knowledge inherent to the structure. (2) Explicit knowledge tends to slow down processing while implicit knowledge tends to speed up processing – provided appropriate processes are available.

The arguments apply to implicit *spatial* structure equally well. Consider the city map example again. We have two ways of finding a certain street on a map: propositionally, by matching symbolic labels on the map regardless of spatial location; and spatially, by using knowledge

about spatial relations between the street and other locations which are maintained in the map representation. A difference to the temporal example is that there may be more degrees of freedom and thus the structure-inherent restrictions may not be as severe as in the temporal case. For example, when we only know that the street is 'near the city center', we may have to search a two-dimensional vicinity rather than a linear structure; nevertheless, spatial nearness constrains the options dramatically compared to a pure propositional approach.

Limits of propositional and spatial knowledge

For the reasons stated above, pure propositional approaches are not used for dealing with real spatial problems. In practical life and in computer science we are always concerned with the question how to organize the propositions we are dealing with to use them efficiently. In other words, we create syntactic structures corresponding to spatial aspects. The partition of knowledge into propositional and structural aspects became particularly evident in the attempts of developing 'pure logic-programming languages', when it was realized that not using implicit structure like the sequence of propositions would make those languages unmanageable and useless. Thus, exploiting structural knowledge – or 'spatial knowledge' – in problem solving and in reasoning is not new. However, there is little work on the issue of comparing the propositional and spatial aspects of knowledge and on trading off the mutual benefits.

The limits of spatial knowledge are obvious: we cannot transcend the intrinsic constraints inherent in the representation structure. We cannot impose arbitrary new rules on using the knowledge as we can with purely propositional knowledge. The rules reside in the structure, and the structure is rather rigid.

Features of propositional and spatial knowledge

To get a handle at comparing the two aspects of knowledge, we will review some important features of propositional and spatial knowledge.

First of all, spatial knowledge represented in a spatial format can be seen as unprocessed knowledge. Specific processes are necessary to exhibit properties and relationships between the represented entities. In contrast, propositional knowledge can be seen as preprocessed or interpreted knowledge resulting from exhibition processes operating on raw information.

A closely related observation concerns the explicitness of different representational aspects. In a purely propositional representation format all knowledge relevant to the task at hand has to be worked out whereas in a spatial representation some (especially relational) knowledge comes for free. The focus therefore lies in designing processes that are capable of exhibiting implicitly contained pieces of knowledge when needed.

The features mentioned so far tend to the intuition that the spatial aspects of a given representation are much more constrained towards reality; i.e., in mimicking the crucial properties of the domain by a spatial representation we map the structural restrictions of the domain onto the representing structure. Propositional formats on the other hand allow for arbitrarily abstracting from the constraints relevant in the domain to be represented, as the structures are built by the propositions rather than supplied by the medium.

The last aspect mentioned here concerns the question whether a given form of representation is suitable for executing representational tasks, e.g. for simulating problem solving processes, or whether it is better suited for explanation tasks. Spatial forms of representation are more appropriate for efficiently performing a given task, as the constraints imposed on the representation structure by the spatiality can be efficiently tailored to the work to be performed. The propositional aspects of representation on the other hand tend to clarify features for explanation purposes because of their explicitness and their abstraction facilities.

Dimensions of the distinction

As already mentioned, knowledge representation is inseparably linked to knowledge processing. Thus, in comparing the propositional and the spatial aspects we have to consider both the static form of the representation and the processes that utilize the respective aspects. A comparison of the propositional and the spatial aspects of knowledge therefore must address a variety of issues relating (1) to the content and (2) to the use of the knowledge. The leading questions for the examination will be (a) what knowledge is expressed in the respective representations? (b) how can the two variants be created? and (c) how can the two variants be used?

To make use of the benefits of both aspects of representation, a representation ideally should exploit as much of both worlds as possible. Care must be taken, however, that both aspects do not interfere with each other when relying on the same representational resources, when a propositional and a spatial aspect need to occupy the same spatial position in the representation medium, e.g. in a map. Also, the question may arise whether a given piece of data has to be interpreted propositionally or spatially – this is a problem in map interpretation as well (see Barkowsky *et al.*, 1997). In the following, we will elaborate a bit further on the above questions. The question (a): what knowledge is expressed in the propositional and the spatial aspects (and how can they be compared to one another) occurs in a twofold manner, reflecting both possible directions of examination.

How can we determine propositional content of spatial structures and vice versa?

When determining propositional content of spatial structures, we can divide content represented in a spatial structure into explicit and implicit knowledge. Explicit knowledge (e.g. color of depicted entities) can easily be described propositionally by denoting the entities in question together with their colors. With regard to implicit knowledge, however, things are more difficult. For describing the pieces of knowledge contained in a spatial structure requires a formal description of the kinds of relationships that are intended to be read off the structure by appropriate processes. This kind of formal ontology of the knowledge structure at hand can be viewed as a model describing how to interpret the spatial structure, i.e., which are valid extraction operations on the structure. Usually, the opportunities to turn spatially implicit knowledge into explicit propositional knowledge are of overwhelming quantity; this makes this endeavor hardly feasible – except perhaps for very small worlds.

The dual task – converting propositional to spatial knowledge – turns out to be even harder. For expressing arbitrary information spatially is a highly demanding task (cf. the problems of visual communication, e.g. Bertin, 1981; Kosslyn, 1994). Mapping propositionally represented knowledge into spatial structures requires the identification of spatially representable regularities in the propositions. When we view the contents of propositions as explicit knowledge and recall that spatial representations carry much of their content implicitly, it becomes evident that turning propositions into spatial structures means converting explicit to implicit knowledge. This means that the inherent structure of the domain must be analyzed to integrate the individual pieces of knowledge contained in the given propositions. Thus, by transforming propositional representational aspects to spatial notations we achieve a significant knowledge gain about the inherent structure of the domain to be dealt with.

How can the two variants be created?

Having outlined the mutual transformations between propositional and spatial aspects of knowledge it becomes apparent that translating spatially represented knowledge into propositions does not meet the requirements needed for comparing them. Rather, the spatial representation of propositional knowledge seems to provide a much richer basis for examining the underlying knowledge structures.

So the central question will be what processes can do the mapping of specific spatial aspects and what are the structures they can use for the spatial representation. The task to be solved therefore comprises (a) the identification of the domain structure to be represented; (b) the design of appropriate processes to operate on this structure; and (c) the transfer of the targeted domain aspects to a representation structure suitable for the processes.

As we pointed out above, the task of representing propositional knowledge spatially is very difficult, in general. When designing representations from scratch spatial representation formats may be advantageous. Frequently, the structure of the domain merely has to be copied

by an appropriate representation structure; to build a propositional representation, in turn, requires the domain to be semantically interpreted before pieces of knowledge can be propositionally coded. This distinction is closely related to the simulation / explanation distinction discussed earlier.

How can the two variants be used?

In comparing the representational aspects *propositional* and *spatial* it turns out that (purely) propositional aspects are static in the sense that they are invariant with respect to the way they are used. Propositions are multi-purpose stand-alone representations whose semantics does not depend on context. It is up to the processes to connect them with each other and to establish global relationships.

Spatial representational aspects, in turn, are highly special-purpose in that they can be used only by specific processes that suit the structures they embody. So the design of a spatial representation structure always comprises the construction of both, the static structure and the processes that exhibit the structural properties. Preserving both, the close correspondence of the representation structure to the structure of the represented domain and the correspondence of processes facilitates the design task and increases the plausibility of the representation with respect to the world it stands for.

CONCLUSIONS

We have discussed the propositional and spatial representation of domain aspects from a knowledge structure and information processing perspective. We have shown that the two modes or "channels" of representation can be employed simultaneously within one representation medium and that both modes can complement one another in a variety of ways. This suggests that we can exploit the advantages of both modes when we find appropriate ways of combining them.

To compare the properties of the two representational modes we have proposed not only to express spatial knowledge in propositional form but conversely to express propositional knowledge in a spatial format. This is where we establish a close connection to mental models: mental models frequently are described as spatial arrangements suitable for visual or quasi-visual inspection where the spatial constraints support the problem solving process. Thus, mental models would be an instance of a representation that can make use of the transformation from the propositional to the spatial mode of representation.

If we view mental models as advantageous forms of *internal* representations it is plausible to assume that we will be able to take advantage of spatial modes of representations *externally*, as well. There is quite a bit of evidence that people do take advantage of spatial structures when

they transform propositionally presented information into sketches or maps to more easily understand what is going on.

To fully understand how these spatial structures can be exploited best, more research has to be done with respect to the processes operating on the representations. Here again, work on external representations has provided useful insights into these processes, both through analytical approaches (e.g. work on eye movements) and through synthetic approaches which we find in the 'active vision' research (cf. Mertsching, 1996). A better understanding is still needed with respect to representations and mechanisms employed in focussing and controlling attention and with respect to the granularity and structural organization of the knowledge involved in the attention focussing mechanisms.

Finally, we can draw a connection from spatial representation to discourse processing. As spatial representations of situations establish a coherent context, they can be instrumental in providing the reference frame for generating and interpreting discourse established in this context (cf. Tversky, 1996). In this way, the spatial medium may carry a considerable part of the communicative message, as the spatial structures do not have to be made explicit through the propositional channel of the discourse.

REFERENCES

Barkowsky, T., R. Röhrig and C. Freksa (1996). Operationalizing diagrammaticity. In: *ECAI-96 Workshop on Representations and Processes between Vision and Natural Language* (W. Maaß and P. McKevitt, chairs), pp. 27-34. Budapest, Hungary.

Barkowsky, T., C. Freksa, B. Berendt and S. Kelter (1997). Aspektkarten - Integriert räumlich-symbolische Repräsentationsstrukturen. In: *Perspektive in Sprache und Raum* (C. Umbach, M. Grabski and R. Hörnig, eds.), pp. 147-168. Deutscher Universitätsverlag, Wiesbaden.

Bertin, J. (1981). *Graphics and graphic information-processing*. de Gruyter, Berlin.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence,* **47**, 139-159.

Davis, J. (1997). *Mapping the mind.* Birch Lane Press, Secaucus.

Egner, S. (1994). *Generierung visueller Suchstrategien auf der Grundlage der Scanpath-Theorie*. Diplomarbeit, Fachbereich Informatik, Universität Hamburg.

Farah, M. J., K. M. Hammond, D. N. Levine and R. Calvanio (1988). Visual and spatial mental imagery: Dissociable systems of representation. *Cognitive Psychology,* **20**, 439-462.

Furbach, U., G. Dirlich and C. Freksa (1985). Towards a theory of knowledge representation systems. In: *Artificial Intelligence: Methodology, systems, applications* (W. Bibel and B. Petkoff, eds.), pp. 77-84. North-Holland, Amsterdam.

Glasgow, J., H. Narayanan and B. Chandrasekaran, eds. (1995). *Diagrammatic reasoning: Computational and cognitive perspectives*. MIT-Press, Cambridge, MA.

Habel, C., S. Pribbenow and G. Simmons (1995). Partonomies and depictions: a hybrid approach. In: *Diagrammatic reasoning: Computational and cognitive perspectives* (J. Glasgow, H. Narayanan and B. Chandrasekaran, eds.), pp. 627-653. MIT-Press, Cambridge, MA.

Johnson-Laird, P. N. (1983). *Mental models*. Harvard University Press, Cambridge, MA.

Karmiloff-Smith, A. (1996). Internal representations and external notations: a developmental perspective. In: *Forms of representation* (D. Peterson, ed.), pp. 141-151. Intellect Books, Exeter.

Kosslyn, S. M. (1980). *Image and mind*. Harvard University Press, Cambridge, MA.

Kosslyn, S. M. (1994). *Elements of graph design*. Freeman, New York.

Kosslyn, S. M. and J. P. Pomerantz (1977). Imagery, propositions and the form of internal representations. *Cognitive Science*, **9**, 52-76.

Landau, B. and R. Jackendoff (1993). "What" and "Where" in spatial language and spatial cognition. *Behavioral and Brain Sciences,* **16**, 217-265.

Mertsching, B., ed. (1996). *Aktives Sehen in technischen und biologischen Systemen*. Infix, Sankt Augustin.

Myers, K. and K. Konolige (1995). Reasoning with analogical representations. In: *Diagrammatic reasoning: Computational and cognitive perspectives* (J. Glasgow, H. Narayanan and B. Chandrasekaran, eds.), pp. 273-301. MIT-Press, Cambridge, MA.

Noton, D. and L. Stark (1971). Eye movements and visual perception. *Scientific American,* **224***, 34-43.

Paivio, A. (1986). *Mental representations - A dual coding approach*. Oxford University Press, Oxford.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In: *Cognition and categorization* (E. Rosch and B. B. Lloyd, eds.), pp. 259-303. Lawrence Erlbaum, Hillsdale, NJ.

Peterson, D., ed. (1996). *Forms of representation*. Intellect Books, Exeter.

Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: a critique of mental imagery. *Psychological Bulletin*, **80**, 1-24.

Rehkämper, K. (1995). Analoge Repräsentationen. In: *Bilder im Geiste: Zur kognitiven und erkenntnistheoretischen Funktion piktorieller Repräsentationen* (K. Sachs-Hombach, ed.), pp. 63-105. Rodopi, Amsterdam, Atlanta.

Rumelhart, D. E. and D. A. Norman (1985). Representation of knowledge. In: *Issues in cognitive modeling* (A. M. Aitkenhead and J. M. Slack, eds.), pp. 15-62. Lawrence Erlbaum, Hillsdale, NJ.

Schlieder, C. (1997). Räumliche Repräsentation im diagrammatischen Schließen. In: *Perspektive in Sprache und Raum* (C. Umbach, M. Grabski and R. Hörnig, eds.), pp. 127-145. Deutscher Universitäts-Verlag, Wiesbaden.

Sloman, A. (1985). Why we need many knowledge representation formalisms. In: *Research and development in expert systems* (M. Bramer, ed.), pp. 163-183. Cambridge University Press, New York.

Sloman, A. (1994). Explorations in design space. In: *Proceedings of the 11th Conference on Artificial Intelligence (ECAI'94)* (A. G. Cohn, ed.), pp. 578-582. Wiley, Chichester.

Tversky, B. (1996). Spatial perspective in descriptions. In: *Language and space* (P. Bloom, M. A. Peterson, L. Nadel and M. F. Garrett, eds.), pp. 463-491. MIT Press, Cambridge, MA.

Wahlster, W. (1994). Computational models of multimodal communication. In: *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI 94)* (A. G. Cohn, ed.), pp. 829-830. Wiley, Chichester.

Zangemeister, W. H., H. S. Stiehl and C. Freksa, eds. (1996). *Visual attention and cognition*. North-Holland, Amsterdam.

ON THE DUALITY AND ON THE INTEGRATION OF PROPOSITIONAL AND SPATIAL REPRESENTATIONS

Christian Freksa and Thomas Barkowsky,
University of Hamburg, Department for Informatics and Cognitive Science Program

running headline: ***Propositional and Spatial Representations***