JS Gero, B Tversky and T Knight (eds), *Visual and Spatial Reasoning in Design III*, 2004 Key Centre of Design Computing and Cognition, University of Sydney, pp 163-182

DEVELOPING AN ONTOLOGY OF SPATIAL RELATIONS

JANE BRENNAN
University of Technology, Sydney, Australia

and

ERIC A MARTIN AND MIHYE KIM
University of New South Wales, Australia

Abstract. We propose a spatial ontology that brings together three aspects of spatial knowledge, namely connectivity, proximity and orientation. These aspects are rich enough to represent knowledge about physical space and each of them can be described in terms of a fixed subsumption hierarchy. The three subsumption hierarchies can then be combined into a relation hierarchy; the way the former are combined into the latter depends on the application domain. For an illustration, we examine how spatial knowledge is represented in natural languages, with an analysis of spatial prepositions in English as a particular case. We obtain a relation hierarchy $\mathscr{R}$ from the subsumption hierarchies using Formal Concept Analysis. We argue that $\mathscr{R}$ is a suitable ontology for the representation of physical space, as other natural languages would also result in the same relation hierarchy $\mathscr{R}$. It can also form part of computer-aided design systems to enable visual representations of verbal spatial descriptions. which might have been the result of discussion between designers.

1. Introduction

"Much of real world design takes place in domains with a spatial component" (Chandrasekaran 1999). Software tools that are to be developed to support the design process therefore need suitable representations of spatial knowledge. As Goel (1995) pointed out, sketches are a very important tool in the early stages of design, because they do not force designers to be committed to precise representations as, for example, CAD systems do. On the same account, qualitative representations are needed for spatial reasoning to support the early stages of the design process within a computing environment. This paper will examine spatial relations and show how generic qualitative representations can be derived.

Spatial relations can be subdivided into several classes and intensive research has been conducted to investigate their properties. Over the last few decades, three classes of spatial relations have emerged as natural candidates for representing spatial knowledge: connectivity, orientation and proximity. They have often been studied independently of each other. Only rarely have all three aspects been combined. An exception is the work of Kettani and Moulin (1999), who developed a spatial model to support navigation through natural language instruction. Their model is very specific and a more general model is still in need. In this paper, we propose an approach that combines connectivity, orientation and proximity together with their subsumption hierarchies. These hierarchies are discussed in more detail in Sections 2, 3 and 4. We show how Formal Concept Analysis (Wille 1982; Ganter and Wille 1999) can be used to define a relation hierarchy $\mathscr{R}$ representing an ontology for a particular domain, by locating each of the relevant aspects of spatial knowledge on some level of $\mathscr{R}$. The Formal Concept Analysis (FCA) is conducted on examples from the spatial knowledge domain. For the sake of simplicity, we will be analysing a set of English prepositions that describe spatial relations. The resulting relation hierarchy could therefore also be referred to as a "Meta-Lingua," since its definition is linguistically motivated.

However, as natural languages are often assumed to encode our internal representations of the real world, $\mathscr{R}$ offers a generic representation of spatial knowledge suitable for applications of different kinds. Note that this paper is not meant to answer the question whether universal concepts of the world actually exist in the reasoning agents' mind, but rather to explore the possibilities of a set of universal concepts to solve computational issues in spatial knowledge representation.

The research reported here stemmed from the observation that distinct natural languages can embrace very different concepts to describe the same spatial situation. For example, when comparing the use of *above* and *on* in English with the use of prepositions in the Mexican language Mixtec, a completely different structure is revealed (Regier 1996) as shown in Figure 1. Work employing universal primitives to represent these spatial situations using conceptual graphs (Sowa 2000) was reported in Brennan (1999). Spatial relations themselves were stipulated, while objects were described as concepts in terms of their orienting axes in 3-D space. The notion of shape, a very important point in design, can also be added to the object concept as needed. This representation is too weak to describe some spatial situations and contexts. However, the weakness of this representation can be overcome by refining the stipulated spatial relations. The relation hierarchy $\mathscr{R}$ provides a valuable tool to create these refinements.

We focus here on binary spatial relations, as these are the most common ones. We assume that spatial relations can be expressed in terms of the three relation classes of connectivity, proximity and orientation. We examine these classes and their associated subsumption hierarchies in the next three sections. In Sections 5 and 6, these hierarchies are incorporated into the relation hierarchy $\mathscr{R}$, representing a spatial ontology.

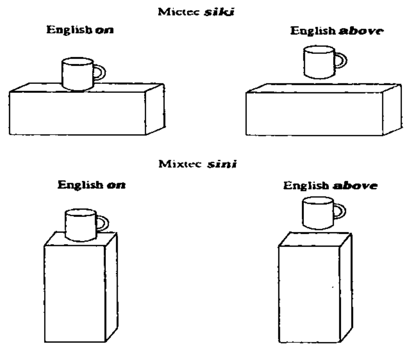

Figure 1. Disparate spatial concepts

2. Connectivity-Based Topology Relations

The Region Connection Calculus examines all possible connectivity relations between regions[1]. Randell et al. (1992) identified several connectivity-based topological relations such as discreteness or overlapping, and represented them in a subsumption hierarchy (or lattice). These connectivity notions are very intuitive and their naming is therefore adopted in this paper. However, it is important to note that our approach is not a mereo-topological but a set theoretical approach. While this does not interfere with the intuitive notions of connectivity such as overlapping or being part of, it does justify some modifications to the original RCC relations that we found useful. Some of the original RCC specialisations such as PP (i.e., proper part) and TPP (i.e., tangential proper part) are not included in our subsumption hierarchy of connectivity-based topology relations, in order to increase readability. In addition, we add the notion of pseudo-equality, which enables to represent intuitive notions such as objects being "next to each other". Recall that pseudo-equality allows for the distance between two points or sets (i.e., regions in this case) to be zero without the points being the same or the regions sharing any point. This has

[1] In RCC, regions are generally thought of as a representation of the space that is occupied by a spatial entity, i.e., an object that is abstracted for the purpose of conducting spatial reasoning.

been shown in Vakarelov et al. (2002) to be a very useful approach for region based theories of space. In the Region Connection Calculus (RCC) (Randell et al.1992), the external connectedness relation *EC* is stronger than the connectedness relation *C*. Two regions are connected if they share at least one point and externally connected if they share at least one closure point. In contrast, we ignore *EC* and define as a counterpart the relation $DC{=}_D$, which is stronger than $\neg C$. Intuitively, $DC{=}_D(x,y)$ means that x and y are "very close to each other," in the sense that they do not share any points (including closure points) but are close enough. The relation $=_D$ represents the fact that there are closure points from each of the sets that are pseudo-equal i.e., they are distinct points but the distance between them is zero. The relation $DC_{\neq D}$ has no counterpart in RCC, but it implies the traditional RCC notion of disconnectedness. Figure 2 shows all the relations used, and the resulting relation hierarchy.

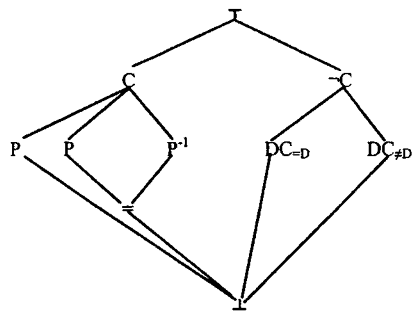

Figure 2. Lattice defining the relation hierarchy of connectivity-based topological relations where C(x,y) means x and y are connected; PO(x,y) means x and y have a nonempty intersection, and neither y is included in x nor x is included in y; P(x,y) means x is included in y; $P^{-1}(x,y)$ means x contains y; $DC_{=D}(x,y)$ means x and y are not connected but are pseudo-equal; $DC_{\neq D}$ means x and y are not connected and are not pseudo-equal

Though more specific than the RCC discreteness notion *DC* (due to its exclusion of pseudo-equality), the relation $DC_{\neq D}$ is not restrictive enough for many purposes such as the representation of natural language expressions of proximity, as it accounts for the large range of cases where the regions do not share a point or are not very close to each other. As discussed in detail in Brennan and Martin (2003), the notion of proximity needs to be carefully defined in order to account for different grades of disconnectedness.

3. Spatial Proximity Relations

Brennan and Martin (2003) presented a theory of nearness. It was assumed that objects are abstracted as points and positioned into a pseudo-metric space, with a pseudo-distance *D*. Points that are perceived as important by the cognitive agent are called sites. The agent's common-sense knowledge about an object abstracted as a site is reduced to a weight ω that codes various properties such as size, danger, or desirability of reaching the object. The weights are used to define the influence areas of these objects. For any site *p*, the influence area of *p* is denoted by *IA(p)*, and is computed from *D* and ω. The function *IA* is the basis for the formal definition of nearness, whose generic notion is assumed to satisfy axioms (*A1*) and (*A2*) below. Axiom (A1) is straightforward, just stating that every site is near itself. Linguistically, there might be cases where an object is not considered near itself, but formally this is a convenient assumption. From the case studies we have done, we could conclude that any two sites whose influence areas do not intersect have no nearness relation. Axiom (A2) expresses this property.

(A1) For all sites p, Near(p,p)
(A2) For all sites p,q, $IA(p) \cap IA(q) = \emptyset \rightarrow \neg Near(p,q)$

A "family" of nearness relations for specific distance and weight satisfying (*A1*) and (*A2*) were defined, Table 1, resulting in a relational tree shown in Figure 3, starting with s-near1 as the most general nearness notion of the "family." The various nearness relations are defined as follows, where the relations marked *s* are symmetric and the relations marked with *a* are asymmetric:

TABLE 1. Nearness Relations

s-near1(p,q)	$=_{Def}$	$IA(p) \cap IA(q) \neq \emptyset$
s-near2(p,q)	$=_{Def}$	*p belongs to IA(q)* or *q belongs to IA(p)*
a-near2(p,q)	$=_{Def}$	$(IA(p) \cap IA(q) \neq \emptyset)$ and $\omega(p) \leq \omega(q)$
a-near3(p,q)	$=_{Def}$	*p belongs to IA(q)*
s-near4(p,q)	$=_{Def}$	*p belongs to IA(q) and q belongs to IA(p)*
a-near4(p,q)	$=_{Def}$	*IA(p) is a subset of IA(q)*

Each of these six notions were shown to be useful in different contexts in a series of case studies. For example, within the context of small-scale spaces, i.e., spaces whose structure is within the sensory horizon of the agent, a magnetic field setting was examined which lead to the definitions of s-near1 and s-near4. This particular interpretation shows a space of three permanent two-bar magnets and a nail (i.e., an unmagnetised iron object). The scene is shown in Figure 4.

It is well known that magnets attract unmagnetised iron objects and attract or repel other magnets depending on their polarity. A bar magnet sets up a magnetic field in the space around it, and a second body responds to it (Sears et al. 2000). For example, if a nail happens to be within a magnetic field, it will be drawn towards the magnet. The second magnet does not have to be within the first magnet's field to either be drawn to or repelled from it, it is sufficient when the second magnet's field gets into contact with the first magnet's field.

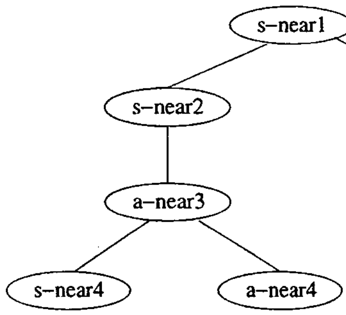

Figure 3. A "Family" of nearness relations for specific (D,ω)

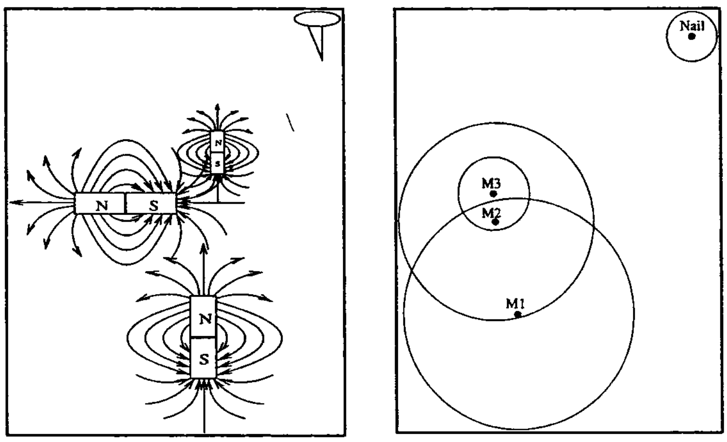

(a) Magnets and Nail (b) Abstracted Space

Figure 4. The influence of magnetic fields

In this interpretation, the term magnetic field describes the region within which the friction with the table top is not enough to stop an object (the magnets or the nail) from moving. This means that the influence areas depend on the frictional characteristics of the object being attracted or the frictional characteristics in addition to the magnetic properties of a magnet being attracted or repelled. In order to consider the spatial setting of the magnets, we assume that the observing agent is moving the magnets into the positions shown and then holding on to them. This allows their fields to intersect without them being physically moved at first encounter. The nail in the scene is not affected by any of the magnets. The nearness relations between these magnetic objects and even between unmagnetised iron objects and magnets are truly symmetric. This results in exactly one model of the scene, Table 2, where *T* stands for a *True* and *F* stands for a *False* nearness relation.

TABLE 2. Model of nearness in the scene in Figure 4

	M1	M2	M3	Nail
M1	T	T	T	F
M2	T	T	T	F
M3	T	T	T	F
Nail	F	F	F	T

The model clearly satisfies axiom (A1) with every object being near itself. Axiom (A2) is satisfied by the nail, whose influence area does not intersect with the influence area of any of the magnets, being not near any of the magnets. It can be observed that the magnets whose influence areas intersect are also near. The definition of s-near1 in Table 1 expresses this symmetric nearness notion. It is the most general nearness notion in our "family" of nearness relations for specific (D,ω). Its formulation was derived from the context of proximity spaces, in which two sets are near each other if they share at least one closure point. Adopting this to a universe containing distinctive points, i.e., sites and their associated areas of influence, nearness holds for two sites if their influence areas intersect.

This does not only make sense in the context of proximity spaces, but is also a reasonable approach for physical space. Worboys (2001) conducted studies in the domain of environmental spaces. His experimental results were analysed in the context of influence areas and used to validate the general nearness notion s-near1. It was found that s-near1 was satisfied in 99.56% of all empirical cases (230 out of 231 cases recorded by Worboys).

In addition to notion s-near1, it can also be noted that in the model shown in Table 2, two sites are always near whenever one of them belongs to the influence area of the other. For example, magnet M2 is in the influence area of magnet M1 and M1 in the influence of M2. The nearness notion s-near4 is a specialisation of s-near1 as defined in Table 1.

Asymmetric aspects of nearness arose from the examination of an environmental space setting, depicted in Figure 5(a). It is a small-scale space, because the scene can be observed from a single viewpoint by the cognitive agent. Moreover, since it is also interpreted in a natural language, linguistic restrictions are imposed on its possible models.

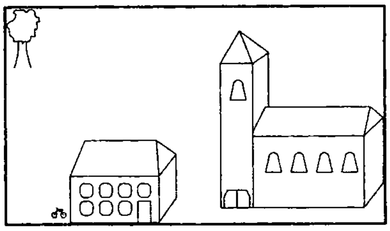

(a) Environmental Small-Scale Space

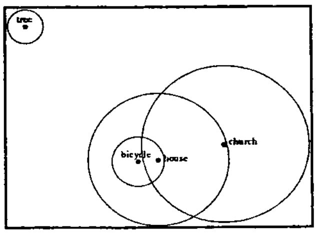

(b) Abstracted Space

Figure 5. Environmental small-scale space example

The scene contains a bicycle next to a house which is in the vicinity of a church. A tree is shown in the distance. According to Talmy (1983), natural language expressions representing spatial relations are commonly asymmetric. Such an asymmetry can occur when the objects differ greatly in the value of some common property such as size. While a small object might be near a large object, the large object is usually not correctly described as being near the small object, in natural language terms. For the scene depicted, the bicycle is definitely near the house, but not the other way around as can be seen in all four models shown in Tables 3 and 4.

TABLE 3. Models 1 and 2 of nearness in the environmental small-scale space of Figure 5

Model 1

	church	house	bicycle	tree
church	T	T	F	F
house	T	T	F	F
Bicycle	T	T	T	F
Tree	F	F	F	T

Model 2

	church	house	bicycle	tree
church	T	F	F	F
house	T	T	F	F
bicycle	T	T	T	F
tree	F	F	F	T

TABLE 4. Models 3 and 4 of nearness in the environmental small-scale space of Figure 5

Model 3

	church	house	bicycle	tree
church	T	T	F	F
house	T	T	F	F
bicycle	F	T	T	F
Tree	F	F	F	T

Model 4

	Church	house	bicycle	tree
church	T	F	F	F
house	T	T	F	F
bicycle	F	T	T	F
tree	F	F	F	T

The whole influence area of the bicycle is enclosed in the influence area of the house, but not conversely. This example and similar ones justify the introduction of the asymmetric nearness notion a-near4 defined in Table 1. The bicycle can also be considered near the church in certain contexts as shown in Models 1 and 2 in Table 3. For example, if the emphasis is on the bicycle being parked next to the house and near the church, and not, for example, near the station where it is usually parked. The bicycle is therefore considered to be near the church, while in this context, it would not be said that the church is near the bicycle. The asymmetric nearness notion a-near2 accounts for this kind of situations. The relationship between house and church is symmetric in some contexts, abstracted in Models 1 and 3, but asymmetric in other contexts, abstracted in Models 2 and 4 in Table 3 and 4 respectively.

In the FCA that will be performed in Section 5 on a set of spatial prepositions, proximity relations are only considered when regions are not connected. The reader should however keep in mind that proximity relations are also true when regions overlap. Although the original framework assumed a point-based universe for representing existing spatial entities, it is possible to adapt the theory to a region-based universe. The point-based universe is mainly chosen for simplicity and to derive more general properties of nearness thanks to a high level of abstraction. However, to investigate the properties of spatial knowledge, especially in the context of design, it is desirable to represent the original spatial entities as regions instead of points. These regions, as in RCC, represent the space that a physical object occupies. We assume that all the objects are two dimensional projections resulting in geometrical figures such as circle, triangles or rectangles. This is also important for consistency reasons, because proximity relations do provide a specification of the discrete relations discussed in Section 2, which assumed point-based regions.

In order to generalise the framework of Brennan and Martin (2003) to a region based approach, the notion of influence area needs to be redefined. We will adopt Kettani and Moulin's (1999) approach to influence areas, which assumes projections of buildings onto a map image resulting in geometric figures. The influence area is then calculated from the outer boundary of these figures using simple Euclidian geometry. While for practical applications, this approach will most likely be directly adopted, for the purpose of abstraction, we will be defining the influence areas of regions by considering these geometric figures as sets of points and generating their influence areas from their closure points and the stipulated notion of weight.

Definition 1 (Influence Area of Regions) *Let a region R be given. Let R' be the closure of R, and let $\omega(R)$ be the weight of R. The influence area of R is defined as the union of sets of points P where for each closure point r in*

R' and each point p in P the distance between p and r is at most equal to ω(R).

Note that the influence area of a region R contains R'. The influence areas of regions can be used to determine the degree of nearness between regions in the same way as for points. This means that regions can only be considered near if their influence areas intersect. Nearness notions for regions are denoted by adding a superscript R to the original notion's name when referring to objects abstracted into regions instead of points, in accordance with the fact that the notions for regions generalise the notions for points. The relation hierarchy of nearness relations, previously defined for a point-based universe and adapted to a region-based universe, is shown in Figure 6.

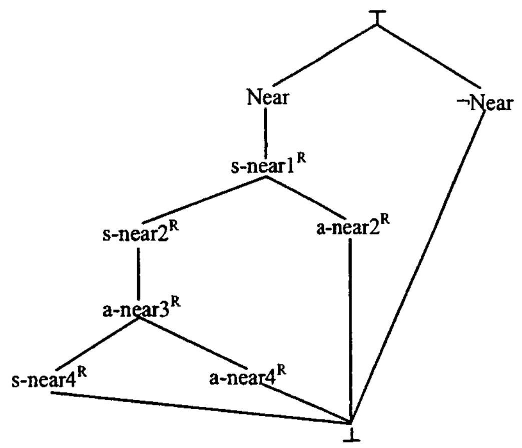

Figure 6. Lattice defining the relation hierarchy of proximity relations where Near(x,y) means x is near y and s-near1^R to a-near4^R denote different nearness notions between regions

4. Orientation Relations

Orientation or direction relations are also a very important aspect of spatial knowledge and can be used in conjunction with connectivity to describe the position of objects to each other in a qualitative way (Hernández 1994). Spatial descriptions of directions can be classified as either *relative* or *absolute*. The most studied relative reference system is no doubt the *left-right* system, but as we have previously seen, the Mexican language Mixtec does also use a relative reference system. However, the research in the field of spatial reasoning has mainly focused on the *Cardinal direction* system, for the obvious reason that many of its applications are within the field of Geographic Information Systems. What all frameworks do have in common is the alignment differentiation, with vertical and horizontal alignment differences born from the simple fact that gravity is omnipresent on this planet. In the spatial reasoning community, the vertical orientation relations have not been studied very deeply, as the focus is on geographic space which is generally restricted to two dimensional maps.

Our classification differs from the one that Frank (1998) proposed by categorising orientation relations into relative and absolute ones. The resulting relation hierarchy of orientation relations is shown in Figure 7. Note that only examples of the *Cardinal direction* system and the *left-right* system are shown for, respectively, absolute and relative reference systems. Other relations and systems are indicated by ellipsis and can easily be added to the conceptual structure as needed.

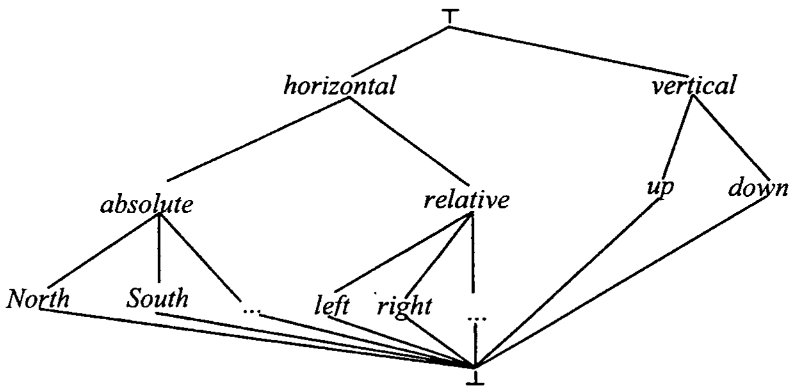

Figure 7. Lattice defining the relation hierarchy of orientation relations

We will now show with examples from natural spatial language how, by analysing their formal concepts, these three fixed subsumption hierarchies can be combined into an ontology for the domain of natural spatial language.

5. Formal Concept Analysis of Spatial Relations

The conceptual structure of spatial knowledge is assumed to be sufficiently represented by connectivity, orientation and proximity. But their associated hierarchies need to be merged in order to represent the "Meta-Lingua" we are striving for. It is therefore necessary to identify the correct links between these hierarchies in the final relation hierarchy $\mathscr{R}$. In order to achieve this, we analyse some examples of spatial relations, represented by spatial prepositions from English, in terms of the spatial primitives represented in the hierarchies of connectivity, orientation and proximity.

In the following we will analyse the English prepositions *on*, *in*, *above*, *under*, *in front of*, *behind*, *at*, *next to*, *near* and *off* as shown in Figure 8.[2]

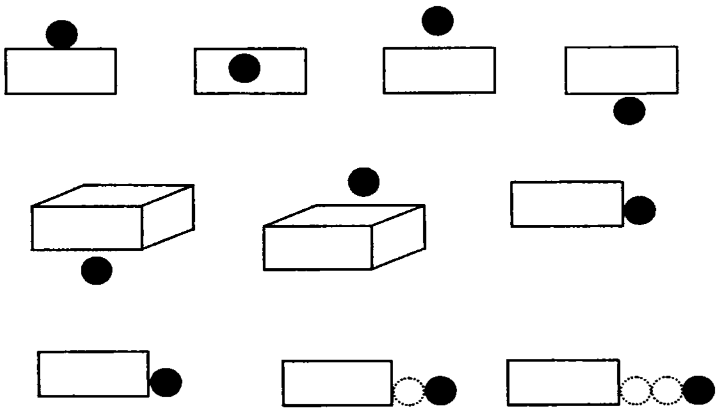

Figure 8. Illustration of the spatial prepositions analysed

The prepositions can be fully described in terms of connectivity, proximity and orientation, as shown in Table 5. For example, the preposition *in* can be adequately described by the connectivity-based topological relation *P*, meaning part of. None of the other sub-relations have any impact on this relation. The other relations are more complex and include orientation relations in their descriptions. For all relations that refer to objects that are not pseudo-equal, we need to consider proximity in addition to orientation. Note that s-near1^R is the most general notion of nearness, hence implies all the other ones. This fact will be accounted for in the Formal Concept Analysis.

We analyse the prepositions and their sub-relations, as stated in Table 5, using Formal Concept Analysis (Wille 1982; Ganter and Wille 1999), in order to derive a meaningful hierarchical structure for our spatial ontology. For the sake of an easy analysis and representation, the prepositions (i.e., objects) are represented by numbers as introduced in Table 5 and their sub-relations (i.e., attributes) are indexed alphabetically as in Table 6.

The context of the spatial relations represented by English spatial prepositions and the sub-relations they can be described with, is shown in Table 7. The rows represent the objects (spatial prepositions) and the

[2] This figure was inspired by Langenscheidt (1993).

columns represent the attributes (sub-relations with which the prepositions can be described). The symbol **x** signifies that the spatial prepositions in the corresponding row can be described (at least in part) with the sub-relation in the corresponding column.

TABLE 5. English prepositions described in terms of connectivity, proximity and orientation relations, i.e., objects

ID	Preposition	Connectivity	Proximity	Orientation
1	in	*P*		
2	on	$DC_{=D}$		*vertical, up*
3	above	$DC_{\neq D}$	$s\text{-}near1^R$	*vertical, up*
4	under	$DC_{\neq D}$	$s\text{-}near1^R$	*vertical, down*
5	in front of	$DC_{\neq D}$	$s\text{-}near1^R$	*horizontal, relative, front*
6	behind	$DC_{\neq D}$	$s\text{-}near1^R$	*horizontal, relative, back*
7	at	$DC_{=D}$		*horizontal*
8	next to	$DC_{\neq D}$	$s\text{-}near4^R$	*horizontal*
9	near	$DC_{\neq D}$	$s\text{-}near1^R$	*horizontal*
10	off	$DC_{\neq D}$	$\neg Near$	*horizontal*

TABLE 6. Alphabetically indexed sub-relations, i.e. attributes

ID	Relation	ID	Relation	ID	Relation	ID	Relation
a	*P*	**e**	*up*	**i**	*front*	**m**	$a\text{-}near2^R$
b	$DC_{=D}$	**f**	*down*	**j**	*back*	**n**	$a\text{-}near3^R$
c	$DC_{\neq D}$	**g**	*horizontal*	**k**	$s\text{-}near1^R$	**o**	$s\text{-}near4^R$
d	*vertical*	**h**	*relative*	**l**	$s\text{-}near2^R$	**p**	$a\text{-}near4^R$
						q	$\neg Near$

The context of the spatial relations represented by English spatial prepositions and the sub-relations they can be described with, is shown in Table 7. The rows represent the objects (spatial prepositions) and the columns represent the attributes (sub-relations with which the prepositions can be described). The symbol **x** signifies that the spatial prepositions in the corresponding row can be described (at least in part) with the sub-relation in the corresponding column.

Based on the formulae of Wille (1982), we determine all formal concepts of the context in Table 5. A formal concept is defined as a pair (X,Y) where X is the set of objects and Y is the set of attributes. The set X is called the *extent* and the set Y is called the *intent* of the concept (X,Y). Applying this to the context in Table 5, we can see that none of the attributes is applicable to all of the 10 objects as shown in the first line of Table 8. Object *1* is the extent of attribute *a* and the set of objects *2* and *7* is the extent of attribute *b*, as shown in lines 2 and 3 of Table 8. In order to find all possible extents, intersections are formed from the sets already formulated. For instance, the intersection of the extents of *a* and *b* is empty, as can be seen on line 4 in Table 8. Objects *2*, *3* and *4* form the extent of attribute *d*; the intersection of *d*'s extent with the extents of attributes occurring above *d* in the table, results in extents {*2*} and {*3,4*} shown on lines 7 and 8, respectively. Extents are unique, therefore if an extent already exists it is not added a second time. For this reason, attributes *l*, *m* and *n* are not added to the extent formulation, because they have exactly the same extent as attribute *k*. Using the extents i.e., sets of objects, we can then formulate the intents i.e., sets of attributes, from our context, as shown in the third column of Table 8. Due to their repetition, the attributes *k,l,m,n,p* are abstracted to one attribute and are denoted by *x* in the *Abstracted Intent* column of Table 8.

TABLE 7. Context of spatial relations: spatial prepositions and their attributes

Objects \ Attributes	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q
1 in	x																
2 on		x		x	x												
3 above			x	x	x						x	x	x	x	x	x	
4 under			x	x		x					x	x	x	x	x	x	
5 in front of			x				x	x	x		x	x	x	x	x	x	
6 behind			x				x	x		x	x	x	x	x	x	x	
7 at		x					x										
8 next to			x				x								x		
9 near			x				x				x	x	x	x	x	x	
10 off			x				x										x

From Table 8, we can now draw a concept lattice of the context in Table 7. This lattice is shown in Figure 9. For drawing the lattice, we use the abstracted intent. The structure of the concept lattice in Figure 9 outlines the possible arrangement of the sub-hierarchies, shown in Figures 2, 6 and 7, within the final relation hierarchy representing our ontology for concepts of physical space. We anticipate that this is possible, because the concept lattice of the FCA provides the implicit and explicit representations of the spatial data, to allow a meaningful and comprehensive interpretation of the information.

Our main interest lies in the sub-relations, with which the spatial prepositions can be described, and not with the prepositions themselves. In order to analyse the hierarchical structure of the sub-relations, a lattice, displaying the sub-relations only, is generated as shown in Figure 10. Note that more specialised levels in a lattice always inherit the sub-relations from more general levels. We have therefore only added the additional relations to each level or all relations for nodes that would otherwise have been empty, in order to make the lattice more readable.

TABLE 8. Formal concepts from the context in Table 7

Attribute	Extent	Intent	Abstracted Intent
	{1,2,3,4,5,6,7,8,9,10}	{ }	⊤
a	{1}	{a}	{a}
b	{2,7}	{b}	{b}
	{ }	{a,b,c,d,e,f,g,h,i,j,k,l,m,n,o,p,q}	⊥
c	{3,4,5,6,8,9,10}	{c}	{c}
d	{2,3,4}	{d}	{d}
	{2}	{b,d,e}	{b,d,e}
	{3,4}	{c,d,k,l,m,n,o,p}	{c,d,o,x}
e	{2,3}	{d,e}	{d,e}
	{3}	{c,d,e,k,l,m,n,o,p}	{c,d,e,o,x}
f	{4}	{c,d,f,k,l,m,n,o,p}	{c,d,f,o,x}
g	{5,6,7,8,9,10}	{g}	{g}
	{7}	{b,g}	{b,g}
	{5,6,8,9,10}	{c,g}	{c,g}
h	{5,6}	{c,g,h,k,l,m,n,o,p}	{c,g,h,o,x}
i	{5}	{c,g,h,i,k,l,m,n,o,p}	{c,g,h,i,o,x}
j	{6}	{c,g,h,j,k,l,m,n,o,p}	{c,g,h,j,o,x}
k	{3,4,5,6,9}	{c,k,l,m,n,o,p}	{c,o,x}
	{5,6,9}	{c,g,k,l,m,n,o,p}	{c,g,o,x}
o	{3,4,5,6,8,9}	{c,o}	{c,o}
	{5,6,8,9}	{c,g,o}	{c,g,o}
q	{10}	{c,g,q}	{c,g,q}

6. A Relation Hierarchy of Spatial Relations

As expected, the proximity relations are always a refinement[3] of the topological relation $DC_{\neq D}$. The final relation type hierarchy of spatial relations therefore has the relation hierarchy of proximity relations from Figure 6 as a refinement of $\mathrm{DC}_{\neq \mathrm{D}}$. Orientation is a refinement of both connectivity and proximity. These relations between the subsumption hierarchies of connectivity, proximity and orientation are shown in

[3] A relation *R* is a refinement of a relation *S* if *R is a subset of S*.

Figure 11. The resulting hierarchy is refined further by the relation hierarchy $\mathscr{R}$.

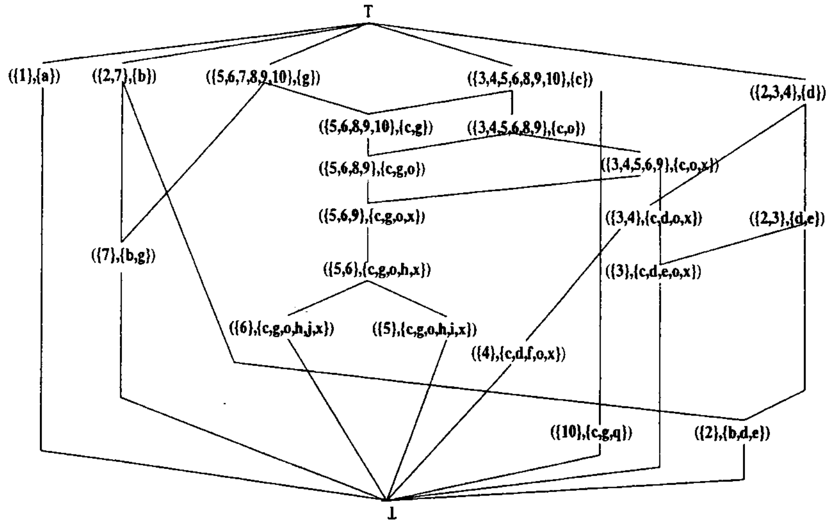

Figure 9. Concept lattice of the context in Table 7

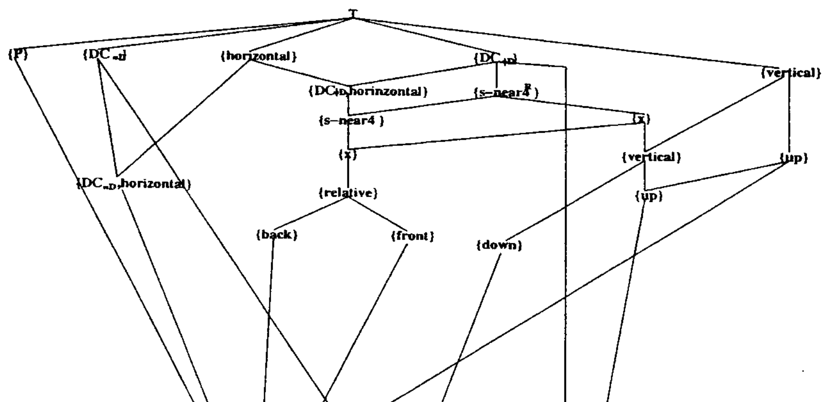

Figure 10. Sub-relation lattice where both occurrences of **x** stand for s-near1^R, s-near2^R, a-near2^R, a-near3^R and a-near4^R. More specialised levels in a lattice always inherit the sub-relations from more general levels. Therefore only the additional relations are added to each level or all relations for nodes that would otherwise have been empty.

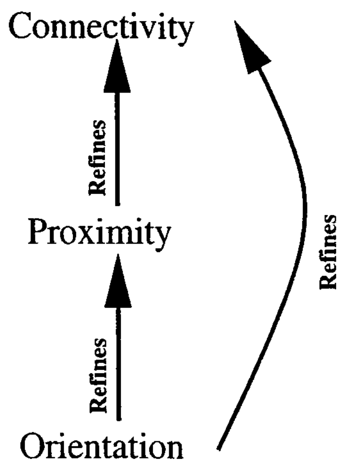

Figure 11. Relations between the subsumption hierarchies

Orientation only refines that part of the hierarchy for connectivity that is below $\neg C$, because for concepts such as the ones represented by the preposition *in*, orientation is not considered at all. This indicates that orientation relations should be located lower in the final hierarchy than connectivity relations. This is also a reasonable approach in the context of Gapp's (1994) potential field approach to representing orientation relations, where it is assumed that orientation is only considered if the reference object and the object to be localised are sufficiently close to each other. Therefore, orientation is not only a refinement of connectivity, but also of proximity.

The *Mixtec* prepositions *siki* and *šini* do not consider the topological relations, but only the orientation relations and the extent of the lower object. The extent of the objects has already been covered by the concept description in terms of the object's axes (Brennan 1999).

We therefore only need to add the orientation sub-hierarchy to the $\neg C$ branch without a need for the distinction between pseudo-equal and pseudo-unequal relations, which Mixtec does not account for. This way, the hierarchy can then cover spatial relations in general even covering for unfamiliar concepts such as the Mixtec examples. Topological and proximity relations can be subsumed in this branch of the relation hierarchy if needed. Figure 12 shows the relation hierarchy representing an ontology for concepts of physical space based on the sub-hierarchies previously discussed, with a hierarchical structure drawn from the Formal Concept Analysis of English spatial prepositions. Due to (drawing) space limitations, some of the sub-hierarchies are represented by symbols. For example, the relation denoted by *Orientation* can be subsumed by the relation sub-hierarchy for orientation relations shown in Figure 7.

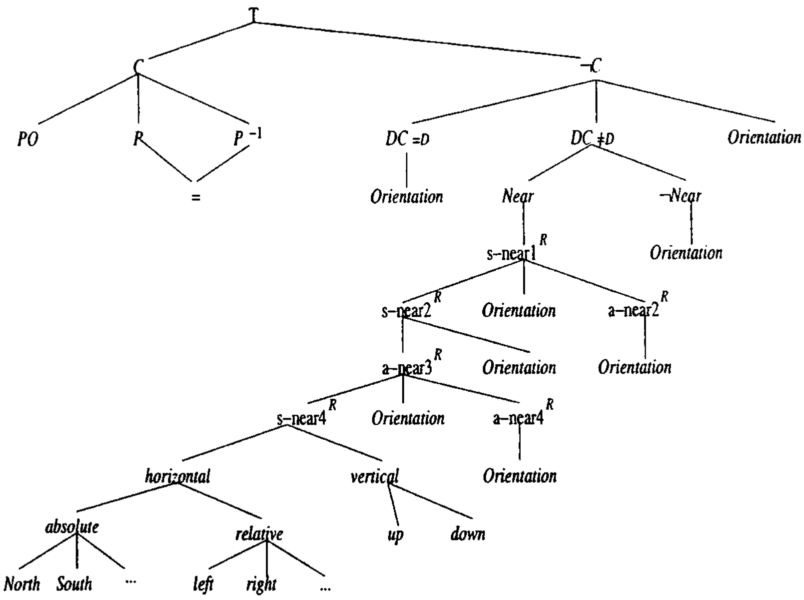

Figure 12. Relation hierarchy R for spatial relations

7. Example for the Use of the Relations Hierarchy $\mathscr{R}$ in Design

An example of the use of the relation hierarchy $\mathscr{R}$ in design could be a formal representation of discussions between the designer and a client. Once the client's requirements are mapped into some form of representation that corresponds to the entries in the hierarchy, tools could generate potential solutions. A more specific example could be as follows.

For the sake of simplicity, we will only consider simple geometric figures such as rectangle, triangle or circle, and assume design in the architectural domain. In the context of nearness notions, the influence areas of objects can here be defined in part by existing building regulations with the option of designer input i.e., the designer graphically placing objects in certain positions and assigning them to certain nearness notions as needed and for later usage. We assume that the designer requires two rectangles with A being smaller than B and states that he or she wants A to be on the left of B. If this is the only constraint given, we can now generate the possible layouts by identifying all the relations that represent a generalisation of the left-relation. The $DC_{=D}$ and all of the nearness relations should therefore be considered for the generation of possible layouts to be presented to the designer. A number of possible layouts are shown in Figure 13.

8. Conclusion and Outlook

This paper proposed a spatial ontology, bringing together three aspects of spatial knowledge: connectivity, proximity and orientation. We used FCA of English spatial prepositions together with the analysis of spatial concepts in other languages to define the hierarchical structure of the ontology.

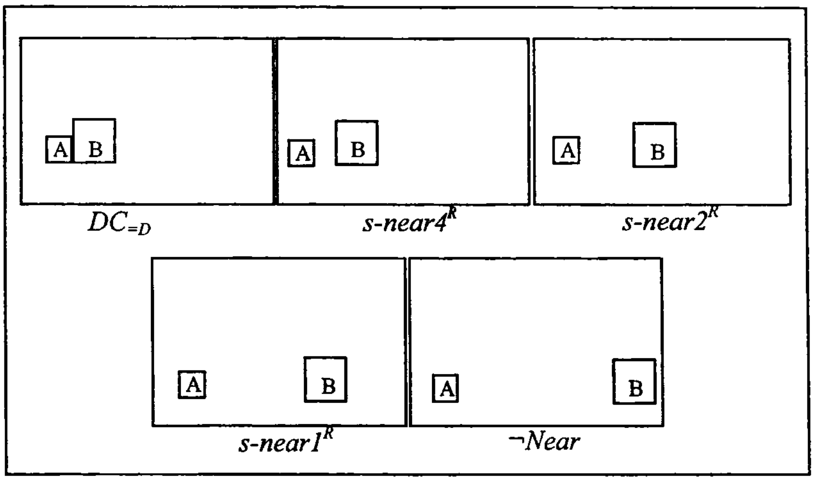

Figure 13. Possible layouts for rectangle A being to the left of rectangle B

The spatial ontology proposed in this paper provides a good starting point for a comprehensive representation of spatial knowledge that can be of interest not only to spatial reasoning, but also to the field of knowledge representation in general. The third spatial dimension (up-down) is often neglected, potentially causing problems in applications such as design or robotics. The representation of spatial knowledge we propose could help overcome this issue as it covers each dimension of the perceptual space.

Future work will focus on the formalisation of the ontology, as it offers very interesting prospects, not only from a spatial cognition point of view, but also in terms of a formal theory. Another possible and interesting direction that could be taken is the application of the ontology to knowledge representation systems or as part of computer-aided design systems that enable the visual representation of verbal spatial descriptions, which might have been the result of discussion between client and designer or a group of designers.

Acknowledgements

We would like to thank the reviewers of the paper and members of the *Creativity and Cognition Studios* at the University of Technology, Sydney for their valuable comments.

References

Brennan, J: 1999, Spatial universals as the human spatial notion, *in* WM Tepfenhart, W Cyre (eds), *Conceptual Structures: Standards and Practises*, Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, pp. 90-96.

Brennan, J and Martin, E: 2003, A theory of proximity relations, *UNSW Technical Report - TR0322.*

Chandrasekaran, B: 1999, Multimodal perceptual representations and design problem solving, invited paper, *in* JS Gero and B Tversky (eds), *Visual and Spatial Reasoning in Design,* Key Centre of Design Computing and Cognition, University of Sydney, pp. 3-14.

Frank, AU: 1998, *Formal Models for Cognition - Taxonomy of Spatial Location Description and Frames of Reference,* Lecture notes in computer science, 1404, Springer, Berlin.

Ganter, B and Wille, R: 1999, *Formal Concept Analysis - Mathematical Foundations,* Springer Verlag, Berlin.

Gapp, K-P: 1994, Basic meanings of spatial relations: Computation and evaluation in 3D space, *National Conference on Artificial Intelligence*, pp. 1393-1398.

Goel, V: 1995, *Sketches of Thought*, MIT Press, Cambridge, MA.

Gärdenfors, P and Williams, M-A: 2001, Reasoning about categories in computational spaces, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, pp. 385--392.

Hernández, D: 1994, *Qualitative Representation of Spatial Knowledge*, Number 804 in Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, Heidelberg.

Langenscheidt-Redaktion (ed): 1993, *Englische Präapositionen - Bedeutung und korrekter Gebrauch der englischen Verhältniswörter*, Langenscheidt.

Nicoletti, MC and Brennan, J: 2002, Learning spatial relations using an inductive logic programming system, *Computing and Informatics (formerly: Computers and Artificial Intelligence)* **21**(1):17-36.

Randell, DA, Cui, Z and Cohn, AG: 1992, A spatial logic based on regions and connection, *Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, pp. 165-176.

Regier, T: 1996, *The Human Semantic Potential - Spatial Language and Constrained Connectionism*, A Bradford Book, The MIT Press, Cambridge.

Sears, FW, Zemansky, MW, Young and Freedman: 2000, *University Physics with Modern Physics*, Addison-Wesley.

Sowa, JF: 2000, *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks Cole.

Talmy, L: 1983, *Spatial Orientation - Theory, Research, and Application*, How language structures space, Plenum Press, New York and London.

Vakarelov, D, Dimov, G, Düntsch, I and Bennett, B.: 2002, A proximity approach to some region-based theories of space, *Journal of Applied Non-Classical Logics*, **12**.

Wille, R: 1982, Restructuring lattice theory: An approach based on hierarchies of concepts, *in* I Rival (ed), *Ordered Sets*, Reidel, pp. 445-470.

Worboys, MF: 2001, Nearness relations in environmental space, *International Journal of Geographical Information Science* **15**(7):633-651.

Visual and Spatial Reasoning in Design III

Preprints of the
3rd International Conference on Visual and
Spatial Reasoning in Design
Massachusetts Institute of Technology, Cambridge, USA
22-23 July 2004

edited by
John S Gero
Barbara Tversky
Terry Knight