Tracking Regions using Conceptual Neighbourhoods

David Randell and Mark Witkowski
Dept. of Electrical and Electronic Engineering
Imperial College of Science, Technology and Medicine
Exhibition Road
London SW7 2BT
United Kingdom
{d.randell, m.witkowski}@imperial.ac.uk

Keywords: Spatial Reasoning, Conceptual Neighbourhood, Identity, Regions

Abstract

The use of *conceptual neighbourhood diagrams* (*CNDs*) is now fairly well established in AI literature. In AI applications, and in particular Qualitative Spatial Reasoning (*QSR*), the nodes of *CNDs* are typically populated by individual relations defined on regions or intervals, and the edges connecting adjacent nodes as continuous transformations between them. While *CNDs* have been used to measure the similarity between relations and their relata, region-identity is often assumed rather than explicitly determined by the representational theory used. In this paper we explore *QSR* theories and relation-based measures of similarity for region-identity in a dynamic setting. Ambiguity in potential identity mappings arising from using a weak qualitative similarity measure is reduced by combining conceptually related but distinct *QSR* theories, each supporting a *CND*. Results from a simulation program that implements the theory are discussed.

1 Introduction

Tracking of objects over time, is central to many machine vision applications. For example, this may be realised in the case of a robot identifying and then moving towards an object, or following it, or where we need to make sense of a sequence of interpreted images with the task of extracting an explanatory narrative. Either way, to do this we will require a sufficiently robust theory of the *continuity of object identity* that enables us to map objects to regions or clusters of spatial relationships or features extracted from a sequence of images. This needs to be sufficiently robust in order to make plausible matches in the presence of, for example, occlusion events, and changes of position or pose. One approach to this problem is to factor out a set of feature-rich properties predicated on individual bodies. Such properties might include low-level information about colour, texture and motion. Alternatively one might use high-level descriptions capturing spatial relation information between bodies (or at least their images). The former approach has been widely investigated; the latter is considered from both theoretical and practical standpoints here.

This paper reports on preliminary work where several relational-based Qualitative Spatial Reasoning (*QSR*) theories, each supporting a Conceptual Neighbourhood Diagram (*CND*), are used to determine the identity of tracked regions over time.

Our starting point is a sequence of temporally indexed images. Each image is assumed to be segmented into a set of unique named regions with spatial relations defined on the regions. Assuming relational informational only, our task is then to cross-correlate the regions in each image in terms of a 'best-fit' match so that the identity of the regions track through the sequence of images.

This task of reliably tracking symbolic objects over time is central to many AI applications and is made manifest in the well-known *symbol grounding* and *anchoring* problem. The main difference applied here is restricting the ontology to *spatial relational information* extracted from the domain. This stands in direct contrast to the more usual method of defining monadic region-identity properties on regions, and using these for recognition and tracking. The interest in relational information also forms part of a general investigation into the practical application of *QSR* formalisms for reasoning about real-world domains using robots supplied with vision sensors. See, for example, [Randell *et al*, 2001].

The utility and application *CNDs* is now fairly well established in AI literature and research. In their most abstract form, *CNDs* can be thought of as digraphs or as non-deterministic finite state automata, each with a set of nodes connected by paths representing transition functions. In AI applications, however, their role and interpretation has largely been for representing and reasoning about qualitative *meta-level* properties of theory fragments of space and time, with the nodes mapping to a set of relations and the paths connecting pairs of nodes to continuous transformations between them. Typically, the set of relations used to

populate the nodes of *CNDs* are defined to be Jointly Exhaustive and Pairwise Disjoint (*JEPD*), meaning that in any model used, at least one and at most one relation from the set is satisfiable [Cohn, 1997].

CNDs share many properties with (and in fact were predated by the use of) *Aspect Graphs (AGs)* [Koenderink and Van Doorn, 1979], which have proved popular in model-based vision research [e.g. Dickinson and Metaxsis, 1997; Herbin, 1998]. In contrast with *CNDs* the nodes of *AGs* represent individual prototypical aspects or viewpoints of a particular object type rather than relations between pairs of individuals, and the paths the direct transitions between them. Aspect Graphs have, for example, been used for object recognition and for encoding the similarity between views against stored prototypes [Cyr and Kimia, 2001]. Within *QSR* proper, *CNDs* have been successfully used in for encoding continuity in qualitative simulation programs [Cui, et al 1992], for encoding vagueness of the extent of regions [Lehmann and Cohn, 1994], and for querying spatial databases extracting models similar to the target one [Papadias and Delis, 1997]. However, most of these applications presuppose that the identity between compared relata is assumed, rather than supplied by the underlying theory.

The rest of this paper is structured as follows. In section 2 conceptual neighbourhood diagrams are discussed. In section 3 the *CND* structure is used to define several similarity distance functions mapping between individual relations and their relata, and between sets of these as a measure of the similarity between models. The distance functions are applied to several conceptually related but distinct spatial theories, and then combined to reduce ambiguity where model-isomorphism yields more than one potential identity assignment across models. Section 4 describes a simulation program that implements the theory. Section 5 presents some preliminary results using the program implementing the algorithm described in section 4, leading to a discussion of further work and conclusions drawn in section 6.

2 Conceptual neighbourhood diagrams

The term *conceptual neighbourhood* originates with Freksa's [Freksa, 1992a,b] analysis of the 13 interval relations defined by Allen's [Allen, 1981, 1983] temporal logic. Entries in cells of the logic's 13x13 *composition table*, were shown to exhibit clustering, forming a path connected subset when interpreted as nodes of a graph. These graphs were originally called "transition graphs" but are now known as conceptual neighbourhood diagrams (figure 1).

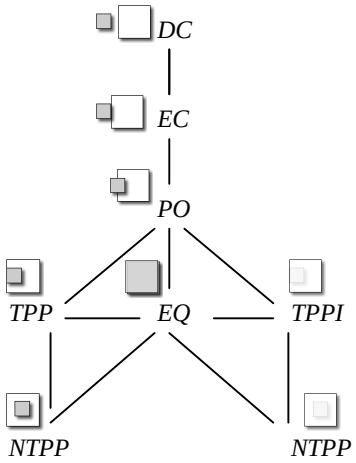

	DC	EC	PO	TPP	TPPI	EQ	NTPP	NTPPI
DC	DC EC PO TPP TPPI EQ NTPP NTPPI	DC EC PO TPP NTPP	DC EC PO TPP NTPP	DC EC PO TPP NTPP	DC	DC	DC EC PO TPP NTPP	DC
EC	DC EC PO TPPI NTPPI	DC EC PO TPP TPPI EQ	DC EC PO TPP NTPP	EC PO TPP NTPP	DC EC	EC	PO TPP NTPP	DC
PO	DC EC PO TPPI NTPPI	DC EC PO TPPI NTPPI	DC EC PO TPP TPPI EQ NTPP NTPPI	PO TPP NTPP	DC EC PO TPPI NTPPI	PO	PO TPP NTPP	DC EC PO TPPI NTPPI
TPP	DC	DC EC	DC EC PO TPP NTPP	TPP NTPP	DC EC PO TPP TPPI EQ	TPP	NTPP	DC EC PO TPPI NTPPI
TPPI	DC EC PO TPPI NTPPI	EC PO TPPI NTPPI	PO TPPI NTPPI	PO TPP TPPI EQ	TPPI NTPPI	TPPI	PO TPP NTPP	NTPPI
EQ	DC	EC	PO	TPP	TPPI	EQ	NTPP	NTPPI
NTPP	DC	DC	DC EC PO TPP NTPP	NTPP	DC EC PO TPP NTPP	NTPP	NTPP	DC EC PO TPP TPPI EQ NTPP NTPPI
NTPPI	DC EC PO TPPI NTPPI	PO TPPI NTPPI	PO TPPI NTPPI	PO TPPI NTPPI	NTPPI	NTPPI	PO TPP TPPI EQ NTPP NTPPI	NTPPI

Figure 1: *CND* and composition table for *RCC8*.

Conceptual neighbourhood diagrams as used in *QSR* are constructed in one of two common ways, either (i) from an axiomatic or logical theory that supplies the *JEPD* set of logical relations, or (ii) from descriptions assuming an algebraic model. Either way the *CND's* structure is generally stipulated, rather than derived from some underlying theory. For the purpose of this paper we will assume these *CNDs* are built from the former base, however this does not mean many of the properties discussed below fail to track through to other representational formalisms used.

With respect to the transformation functions used (and notion of continuity assumed) this depends on the nature of the domain and the interpretation of the axiomatised primitives. Many *CNDs* used in *QSR* model topological change but *CNDs* are by no means restricted to this, e.g. one can model continuity in terms of metrical or geometrical change too. For the purposes of this paper however, we shall illustrate the method using the well known spatial logic *RCC8* and a spatial analogue of Allen's interval logic of time, reducing the standard set of 13 relations to 3.

2.1 CNDs and Envisionment axioms

In the case of *RCC8* the structure of the *CND* is provided by a stipulated set of *envisionment axioms* [see e.g. Cui *et al*, 1992]. The stipulation is necessary since the axiomatised primitive connects relation *C/2* has a spatial interpretation only; it does not encode a theory of time, nor of change, nor of continuity over time, which will be required here. These axioms typically assume the regions in the projected state are non-null, i.e. they do not pass out of existence, as would be the case where two overlapping regions (with a non-empty product) separate; though as we will see below this restriction can be lifted.

3 Similarity and CNDs

Similarity is typically defined in terms of some metric that measures the *closeness* between specified elements, groups or variables of some model. In our case however we wish to compare and measure the distance between *sets* of spatial regions with respect to some background theory.

Let $x_1,\ldots,x_n$ be spatial region variables, and $\boldsymbol{X_1},\ldots,\boldsymbol{X_n}$ be sets of ground spatial *relations* defined on these. We additionally assume a set of related background theories: $\Phi_1,\ldots,\Phi_n$ each of which has a set of *JEPD* relations ($JEPD_\Phi$), and an associated *CND* (CND_Φ).

We first define the distance function: $dmin_\Phi(R_1(x_1,\ldots x_n),R_2(x_1',\ldots,x_n'))$, that maps R_1 and R_2 to two CND_Φ nodes and maps this to the minimum CND_Φ (node to node) distance:

$dmin_\Phi(R_1(x_1,\ldots x_n),R_2(x_1',\ldots,x_n')) = n;\ n \geq 0$
where:
$R_1(x_1,\ldots x_n) \in JEPD_\Phi$
$R_2(x_1',\ldots,x_n') \in JEPD_\Phi$
for every mapping: $x_1 \rightarrow x_1',\ldots,x_n \rightarrow x_n'$

Clearly, $dmin_\Phi(x,x)=0$, and $dmin_\Phi(x,y)>0$ iff $x \neq y$. From this we can see that each *CND* can now be presented as a similarity matrix that maps pairs of CND_Φ nodes to the minimal *CND* node-node path distance between them:

$dmin\Phi$	$R_1(x_1',\ldots,x_n')$	$R_2(x_1',\ldots,x_n')$	…	$R_n(x_1',\ldots,x_n')$
$R_1(x_1,\ldots,x_n)$	0	n_1	…	n_n
$R_2(x_1,\ldots,x_n)$	n_1	0	…	…
…	…	…	0	…
$R_n(x_1,\ldots,x_n)$	n_n	…	…	0

Table 1: Similarity matrix for CND_Φ

Similarly, we define the function: $Dmin_\Phi(\boldsymbol{X_1},\boldsymbol{X_2})$ that takes two sets of spatial relations $\boldsymbol{X_1}$ and $\boldsymbol{X_2}$ and returns the minimal CND_Φ distance between the sets.

$$Dmin_\Phi(\boldsymbol{X_1},\boldsymbol{X_2}) = \sum_i dmin_\Phi(R_1(x_1,\ldots x_n),R_2(x_1',\ldots,x_n')) = i;\ i \geq 0$$

where:
$R_1(x_1,\ldots x_n) \in \boldsymbol{X_1}$
$R_2(x_1',\ldots,x_n') \in \boldsymbol{X_2}$
for every mapping: $x_1 \rightarrow x_1',\ldots,x_n \rightarrow x_n'$

For example, assuming the spatial logic *RCC8* with its similarity matrix:

$dmin_{RCC8}$	DC	EC	PO	TPP	TPPI	EQ	NTPP	NTPPI
DC	0	1	2	3	3	3	4	4
EC	1	0	1	2	2	2	3	3
PO	2	1	0	1	1	1	2	2
TPP	3	2	1	0	2	1	1	2
TPPI	3	2	1	2	0	1	2	1
EQ	3	2	1	1	1	0	1	1
NTPP	4	3	2	1	2	1	0	2
NTPPI	4	3	2	2	1	1	2	0

Table 2: Similarity matrix for CND_{RCC8}

Then, for example:
$\boldsymbol{X_1} = \{dc(a,b),dc(b,a)\}$,
$\boldsymbol{X_2} = \{ntpp(a',b'),ntppi(b',a')\}$
$Dmin_{RCC8}(\boldsymbol{X_1},\boldsymbol{X_2}) = 8$
where: $\{a \rightarrow a', b \rightarrow b'\}$
and:
$\boldsymbol{X_3} = \{dc(a,b),po(b,c),dc(a,c),$
$dc(b,a),po(c,b),dc(b,a)\}$
$\boldsymbol{X_4} = \{ec(a',b'),ntpp(b',c'),dc(a',c'),$
$ec(b',a'),ntppi(c',b'),dc(c',a')\}$
$Dmin_{RCC8}(\boldsymbol{X_3},\boldsymbol{X_4}) = 6,$
where: $\{\{a \rightarrow a', b \rightarrow b', c \rightarrow c'\},\ldots\}$

The ellipsis $\{\{ x_1 \rightarrow x_1',\ldots,x_n \rightarrow x_n'\},\ldots\}$ indicated here shows that more than one set of mappings can satisfy some minimal *CND* distance.

Finally, we define two other functions: $Dmin_\Phi^{x \rightarrow x'}(\boldsymbol{X_1},\boldsymbol{X_2})$ and $n(Dmin_\Phi^{x \rightarrow x'}(\boldsymbol{X_1},\boldsymbol{X_2}))$ that respectively maps the minimised distance between sets to the set of assignments, and the number of these:

$$Dmin_\Phi^{x \rightarrow x'}(\boldsymbol{X_1},\boldsymbol{X_2}) = \bigcup_{x \rightarrow x'} Dmin_\Phi(\boldsymbol{X_1},\boldsymbol{X_2})$$

where:
$R_1(x_1,\ldots x_n) \in \boldsymbol{X_1}$
$R_2(x_1',\ldots,x_n') \in \boldsymbol{X_2}$
for every mapping: $x_1 \rightarrow x_1',\ldots,x_n \rightarrow x_n'$

By applying the function $Dmin_\Phi$ to the set of models being compared, we can define an ordering on the similarity measures returned for all possible mappings. In some cases a single model is returned, in other cases several models are generated. If several distinct models are returned we say the similarity measure used is *ambiguous*. This means that the theory encoded in CND_Φ is *underdetermined* with respect to the models being compared. If a single correspondence is sought, then either a richer representation, or a different

CND (and set of primitives) for the basis of our metric must be used. This is discussed in section 3.2 below.

3.1 Similarity and continuity

Each CND_{Φ} encodes a notion of *continuity* between its path-connected nodes; and this concept is ultimately reducible in terms of theory Φ's axiomatised primitives. Thus for example, CND_{RCC8} interprets continuity in terms of instantaneous topological change between states (as sets of relations); each state of which is reducible to the axiomatised primitive relation *C/2* defined on regions. It is therefore important to emphasise here that the distance functions used here, precisely encode this notion of similarity, and no more. The measure of similarity between models given by $Dmin_{\Phi}$ is none other than setting out the minimal number of transitions required, in order to achieve *model isomorphism*. This approach mirrors that already done for Aspect Graphs where the similarity between views to 2D shape metrics have been used to rate the similarity between unknown views of an object and stored prototypes [Cyr and Kimia, 2001]. Here, and in contrast, the logical aspect of the underlying theory and its primitives are factored out.

Two aspects of similarity are recognised here: (i) *relational (or structural) similarity* which concerns itself how various objects are connected but without regard to their shape and size; and (ii) *metric similarity* which is concerned with the similarity between corresponding objects themselves. We concentrate upon the former notion here, noting that the approach can be easily augmented with the latter. This is discussed briefly in section 5.

3.2 Isomorphism, ambiguity and granularity

We conjecture that *QSR* theories for which a distance function $dmin_{\Phi}$ can be set up may well be too weak for many applications. *RCC8*, for example, is quite a rudimentary measure of similarity. It is not even a metric since non-identical pairs of regions can have distance 0 - they just need to instantiate the same RCC relation. Thus two models that are *RCC8-isomorphic* will have distance 0, even though from a geometrical point of view, they look very different – see Figure 2.

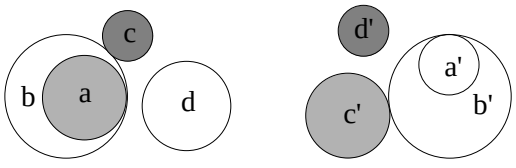

Figure 2: Geometrically dissimilar *RCC8-isomorphic* models. The two sets of regions: *{a,b,c,d}* and *{a',b',c',d'}* satisfy identical *RCC8* relations assuming the mappings $\{a \rightarrow a', \ldots, d \rightarrow d'\}$

However, this failure to distinguish between models, where such distinctions are sought, highlights a modelling problem. This problem ultimately stems from the primitives selected rather than from the measurement of similarity between models being proposed here. Indeed, for some applications this weak isomorphism may well be advantageous. Even assuming a weak target theory, the generality and modularity of the approach lends itself well to different techniques where this can be addressed.

Firstly, we can suitably enrich our chosen representational formalism (and thereby reduce, for example, geometric ambiguity) by coupling together and applying several theory related *CNDs*. This can be done either by extending a given theory, or by joining together two theories via a specified set of axioms that map between both theories. An example of the latter is the viewpoint dependent spatial logic *ROC20* (defined on bodies) that factors out and embeds the weaker spatial logic *RCC8* (defined on images). Here, for example, ambiguity stemming from the *RCC8* overlap relation when applied to the interpretation of images (as the geometric projection of bodies from some viewpoint) is strengthened by the use of stronger occlusion relation [Galton, 1994; Randell *et al*, 2001]. We give an example of this reduction of ambiguity in section 3.3 below, where a simplified spatial theory based on the left/right ordering of spatial regions is used to complement the assumed *RCC8* relations used in our model.

Secondly, another common QSR technique can be used that that maps a predefined $n \times m$ granularity matrix (or grid) to the extra-logical domain model. Relational properties of cell-cell adjacency of individual elements as well as properties of the cells themselves are then additionally taken into account, as when a cell is either empty, partially occupied or, completely occupied by some region in the domain model.

3.3 Reducing ambiguity by composing *CNDs*

Consider the two models depicted in Figure 3 immediately below. $\boldsymbol{X_1}$ maps to a set of ground spatial relations defined on the set of regions *{a,b,c}* and similarly with $\boldsymbol{X_2}$ mapping to the set of regions *{a',b',c'}*. The individual regions are uniquely labelled, and the spatial relations on the regions can now be determined. As given no identity map is made between paired regions across models, the models remain uncorrelated.

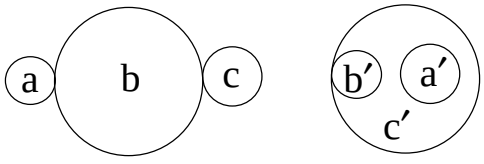

Figure 3: Uncorrelated models.

As an example, let us now assume we have a simple theory *Left* that gives relative left/right orientation information; of whether a region is either wholly to the left or right of another, and where in neither case we say the regions overlap with respect to their left-right most extents. We assume that the embedding space has a predefined left/right axis that is fixed with respect to some assumed viewpoint, as in the typical orientation of this page assumed by the reader when reading this text.

The axioms and definitions of *Left/2* are as follows:

(A1) $\forall x \neg Left(x,x)$
(A2) $\forall x \forall y \forall z [[Left(x,y) \& Left(y,z)] \rightarrow Left(x,z)]]$
(A3) $\forall x \forall y[Left(x,y) \rightarrow \forall z[P(z,x) \rightarrow Left(z,y)]]$
(D1) $Right(x,y) \equiv df.\ Left(y,x)$
(D2) $Overlaps(x,y) \equiv df.\ \neg Left(x,y) \& \neg Left(y,x)$

Axioms *(A1)* and *(A2)* define the primitive relation *Left/3* as irreflexive and transitive (and hence asymmetrical) while (*A3*) maps *Left/2* to the part-whole relation *P/2* of the spatial logic *RCC8*, i.e. if *x* is left of *y*, then any part of *x* is also left of *y*[1]. *(D1)* and *(D2)* provide definitions for when *x* is to the right (and the inverse of being to the left) of *y*, and where, with overlap, neither *x*, nor *y* is wholly left or right of the other. A casual inspection of the axioms and definitions will show that the set: *{Left(x,y),Overlaps(x,y),Right(x,y)}* is *JEPD*, and given this property, we assign these relations to a CND_{Left} similarity matrix:

$Dmin_{Left}$	*Left*	*Overlaps*	*Right*
Left	0	1	2
Overlaps	1	0	1
Right	2	1	0

Table 2: Similarity matrix for CND_{left}

We now stipulate the following sets of relations depicted in the models depicted in Figure 3, then compute product set of the minimal CND_{Left} distance between matched pairs of regions:

$\boldsymbol{X_1}$*= {Left(a,b),Left(b,c),Left(a,c),*
Right(b,a),Right(c,b),Right(c,a)}
$\boldsymbol{X_2}$ *= {Right(a',b'),Overlaps(b',c'),Overlaps(a',c'),*
Left(b',a'),Overlaps(c',b'),Overlaps(c',a')}

(a,b,c)	*L (a,b)*	*L (b,c)*	*L (a,c)*	*R (b,a)*	*R (c,b)*	*R (c,a)*	*n*
(a' ,b',c')	*R (a',b')*	*O (b',c')*	*O (a',c')*	*L (b',a)*	*O (c',b')*	*O (c',a')*	*8*
(a',c',b')	*R (a',c')*	*O (c',b')*	*O (a',b')*	*L (c',a')*	*O (c',b')*	*O (b',a')*	*8*
(b',a',c')	*R (b',a')*	*O (a',c')*	*O (b',c')*	*L (a',b')*	*O (c',a')*	*O (c',b')*	*4*
(b',c',a')	*R (b',c')*	*O (c',a')*	*O (b',a')*	*L (c',b')*	*O (a',c')*	*O (a',b')*	*8*
(c',a',b')	*R (c',a')*	*O (a',b')*	*O (c',b')*	*L (a',c')*	*O (b',a')*	*O (c',b')*	*4*
(c',b',a')	*R (c',b')*	*O (b',a')*	*O (c',a')*	*L (b',c')*	*O (a',b')*	*O (a',c')*	*4*

Table 3: *CND* similarity measures for *Left* assuming the model depicted in Figure 3. Note the shorthand *L* for Left, etc.

The top (header) row of Table 3 gives the set of relations for $\boldsymbol{X_1}$ and the cell entries immediately below for $\boldsymbol{X_2}$. The leftmost column gives the permutation for the set *{a' ,b' ,c' }*. The second row shows the assignment: *(a' ,b' ,c'): as Right(a' ,b'), Overlaps(b' ,c'), Overlaps(a' ,c')*, for which $n(Dmin_{Left}(\boldsymbol{X_i},\boldsymbol{X_j}))=8$, and where:

$\boldsymbol{X_i}$*= {Left(a,b),Left(b,c),Left(a,c),*
Right(b,a),Right(c,b),Right(c,a)}
$\boldsymbol{X_j}$*={Right(a',b'),Overlaps(b',c'),Overlaps(a',c'),*
Left(b',a'),Overlaps(c',b'),Overlaps(c',a')}

We can immediately see that: $n(Dmin_{Left}^{x \rightarrow x'}(\boldsymbol{X_1},\boldsymbol{X_2}))=3$; meaning three assignments have *Left-isomorphic* models and thus indistinguishable with respect to this reduced theory. In order to disambiguate the models, we may note that the explicit modelling of connectivity between the regions themselves which is captured by the logic *RCC8* is missing. So we apply the *RCC8* relation set to our model, and then compute our *CND* measures as before:

$\boldsymbol{X_1}$*= {Left(a,b),Left(b,c),Left(a,c),*
EC(a,b),EC(b,c),DC(a,c)}
$\boldsymbol{X_2}$ *= {Right(a',b'),Overlaps(b',c'),Overlaps(a',c'),*
DC(a',b'),TPP(b',c'),NTPP(a',c')}

(a,b,c)	*EC (a,b)*	*EC (b,c)*	*DC (a,c)*	*EC (b,a)*	*EC (c,b)*	*DC (c,a)*	*n*
(a',b',c')	*DC (a',b')*	*TPP (b',c')*	*NTPP (a',c')*	*DC (b',a')*	*TPPi (c',b')*	*NTPPi (c',a')*	*14*
(a',c',b')	*DC (a',c')*	*TPP (c',b')*	*NTPP (a',b')*	*DC (c',a')*	*TPPi (b',c')*	*NTPPi (b',a')*	*10*
(b',a',c')	*DC (b',a')*	*TPP (a',c')*	*NTPP (b',c')*	*DC (a',b')*	*TPPi (c',a')*	*NTPPi (c',b')*	*14*
(b',c',a')	*DC (b',c')*	*TPP (c',a')*	*NTPP (c',b')*	*DC (c',b')*	*TPPi (a',c')*	*NTPPi (b',c')*	*14*
(c',a',b')	*DC (c',a')*	*TPP (a',b')*	*NTPP (c',b')*	*DC (a',c')*	*TPPi (b',a')*	*NTPPi (b',c')*	*10*
(c',b',a')	*DC (c',b')*	*TPP (b',a')*	*NTPP (c',a')*	*DC (b',c')*	*TPPi (a',b')*	*NTPPi (a',c')*	*14*

Table 4: Similarity measures for *RCC8* assuming the model depicted in Figure 3.

[1] This is a very strong notion of *being left of* and is used here for simplicity only. We could have equally used a weaker more expressive version of this predicate, where *x* being left of *y* means some part of *x* is left of *y*, thus allowing for overlapping regions – see [Randell *et al*, 2001] for an example of this.

We first note that: $n(Dmin_{RCC8}^{x\to x'}(\mathbf{X_1},\mathbf{X_2}))=2$ meaning that, again, two assignments yield *(RCC8) isomorphic* models. But now by combining both theories, i.e by adding the values for each assigment for both matrices, we see that the target models are now distinguished with respect to the two theories *Left* and *RCC8*. The minimised value $n(Dmin_{Left}^{x\to x'}(\mathbf{X_1},\mathbf{X_2})) + n(Dmin_{RCC8}^{x\to x'}(\mathbf{X_1},\mathbf{X_2}))=14$ returns the singleton set $\{a\to b', b\to c', c\to a'\}$ corresponding to the assignment *(c' ,a' ,b')* and which is depicted in Figure 4 below:

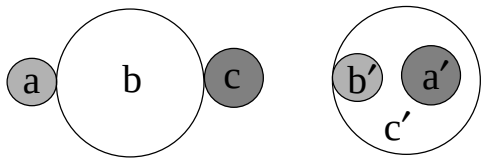

Figure 4: Correlated models assuming the similarity matrices for the theories *Left* and *RCC8*. Here similarly shaded regions indicate region identity with the mapping: $\{a\to b', b\to c', c\to a'\}$.

This example illustrates the general procedure, where *constraints* for model-isomorphism gets progressively stronger/weaker by the composition/reduction of functionally related similarity matrices.

4 Implementation

To support the theoretical analysis presented in this paper we have prepared a simulation program, implemented in C++ running under Windows, to investigate and evaluate the practical and pragmatic consequences of the issues raised in this paper, particularly the degree of ambiguity of region-identity for several distinct spatial theories. The program incorporates a GUI interface. Regions are represented by small bitmaps, which may be loaded into layers and subsequently moved (by the user directly or from a recorded file of movements) within a restricted operating area. These layers represent a notional depth (z-axis) to allow the calculation of the spatial relations in the presence of overlap, as is required, for instance, for the *LOS14* relation set. At present the program is limited to nine distinct regions, for reasons that will become apparent later.

Whenever any of the *n* regions move, the program recalculates the *RCC8*, *LOS14*, *Left/Right* and *Above/Below* relation sets and displays these in separate sub-windows. The Left/Right and Above/Below relations used are direct analogues of the temporal relations described by [Allen, 1994]. After computation, each relation set is presented as an *nxn* matrix of results in a sub-window. An example (for *LOS14*) is shown in figure 5. Note that the presentation is symmetric about the diagonal. The program also detects when a move results in a change of one or more relations. Not all regions movements cause any of the relations to change, and some will cause change to several relations simultaneously. This alone represents a useful tool for demonstrating the generation of different types of spatial relation.

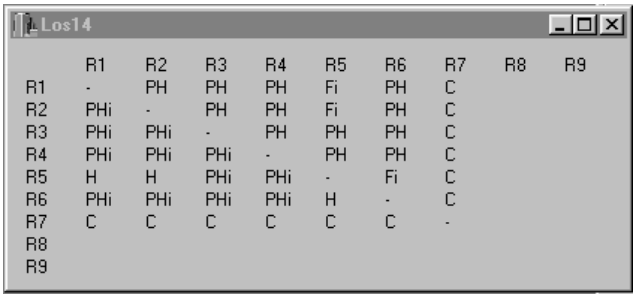

Figure 5: Sample *LOS14* relations matrix

Note that these are discrete representations of continuous descriptions. Adjacency (as, for instance, in the *RCC8* relations EC or TTP) is defined in terms of 8-connectivity between neighbouring pixels. Notions of discontinuity, overlap and containment follow from this. There are important issues relating to the use of such digitised regions, however the intended interpretation is clear here (refer to [Galton, 1999] or [Roy and Stell, 2002] for a more detailed discussion).

The-top level outline algorithm, implementing the process described in section 3.3, is as follows:

For any change of any relation, in any theory, between any pair of regions:
(i) for every pair of regions, record (and display) the appropriate $JEPD_\Phi$ relation (Φ = *RCC8*, *LOS14*, *Left/Right*, *Above/Below*), and compute the similarity measures as follows:
{
ii) Generate the departure set vector, ignoring duplication through argument order.
(iii) Initialise distance sums and counters.

iv) For each permutation of the ordered tuples of the relations:
{
(v) For each theory Φ:
{
(vi) look-up and sum the individual distances between elements in the departure set vector and the previous (saved) departure set vector ($n(dmin_\Phi)$).
(vii) If the new Φ distance sum is less than the existing Φ distance sum, record the new sum, reinitialise Φ counter. If equal, then increment the Φ counter only.
}
(viii) Total all individual Φ distance sums to form totalsum. If the new totalsum less than the existing, record the new and reinitialise the total

counter. If they are equal, then increment the total counter only.
}
(ix) Save the current departure set vector.
(x) Display (and record to file) the resulting values of the individual Φ distance sum and individual Φ counters, and the totalsum and total counter.
}

Step (i) only requires us to re-compute the similarity measures if some relation changes in any of the theories (Φ); for instance when any region moves, or changes size or shape, or is introduced or removed. Note that not all movements or changes necessarily lead to a change in any of the relation pairs. Step (ii) generates the elements in the departure set for *{r1...,rn}*. This step corresponds to the generation of the top row of table 4. Step (iv) generates each element of the target set in turn. The current program uses Dijkstra's algorithm for creating permutations in lexicographic order. This corresponds to each row in table 4. Steps (vi) and (vii) and then generate a new distance measure between the departure set vector and the target set vector for each model ($n(dmin_{\Phi})$). This is repeated for each of the (four) target theories (Step (v)).

The result of steps (v) to (vii) is to record the lowest distance sum for each of the theories (Φ distance sum), and the number of models giving rise to this lowest distance value (Φ counter). We will refer to the latter as an *ambiguity* measure (see section 5). Note than the individual distances between the current and previous departure set vectors are given within the program. The similarity matrix recording the relation-to-relation distances for *RCC8* is given in table 2. Those used for *LOS14* are derived from the CND given in Galton [1994] and those for *Left/Right* and *Above/Below* derived from the CND presented in [Freksa, 1992a] and also in [Papadias and Delis, 1977]. Further, it is easy to see that once a set of envisionment axioms are stipulated (or a digraph given) similarity matrices can be automatically computed, and models compared. This program modularity mirrors the general approach described in [Randell and Witkowski, 2002] and used to build composition tables for large axiomatic theories.

Step (vii) derives overall measures of conceptual distance and ambiguity across all the theories used (section 3). Step (ix) records the current departure set vector for use in the next computation and step (x) displays and records the results for the current step. A typical set of results obtained is shown in section 5.

4.1 A Note on Complexity and Performance

The number of computational steps inherent in the algorithm described in the previous section is clearly directly related to the number of regions, *n*, under consideration. Taking the "payload" of this algorithm to be the evaluation of the theory distance and ambiguity measures, we can see that step (iv) introduces a permutation ${}^nP_n$ or *n*! Step (v) introduces a repetition based on the number of theories, *m*, to be used. Step (vi) is proportional to the combination of the number of regions, ${}_nP^2$, noting the use of dyadic relations. This leads to a first, rough, approximation of ${}^nP_n$ x *m* x ${}_nP^2$ steps. The initial surprise, perhaps, is that the complexity of the underlying theory has no bearing on the timings; this complexity is reduced to a simple look-up operations provided by the similarity matrix (e.g. table 2).

Given the limiting computational overhead is determined by the permuted set of mappings, this quickly places an upper practical limit of the number of regions that can be considered. In practice, we note that on a relatively modest desktop computer (PIII, 933MHz, Windows 98), and despite the relatively heavy computational load imposed by the algorithm, up to seven regions (846720 steps according to the formulation above) may be processed with little perceptible delay. Addition of an eight region (~9M steps) introduces a small, but noticeable, delay and the introduction of the ninth (~105M steps) causes unacceptable (several seconds) delay in real time whenever the similarity must be computed. Sequences might still be processed in "batch" mode from a recorded file of movements; but without substantial algorithm redesign there is little realistic prospect of processing 10 or more regions with this approach in the foreseeable future.

Clearly we would like to maximise the number of regions that can be considered. Apart from minor improvements in code organisation, several steps might be considered in order to reduce the number of possible assignments to be made. For instance, the number of elements to be compared might be reduced. In the case of *RCC8*, for example, the function *sum(x,y)* can be recursively applied to pairs of regions of a finite model, to reduce the number of regions being compared until some specified upper limit is reached. The (reduced) models are compared, then we then use the theory itself to constrain the properties of the decomposed relation sets as they are recursively unpacked. For example, in the simple case where: $\boldsymbol{X_1}$= *{ntpp(a,b),ntpp(b,c)},* $\boldsymbol{X_2}$*={ntpp(x,y),ntpp(y,z)};* we can re-work the sets as: $\boldsymbol{X_1}'$*= {ntpp(sum(a,b),c)},* $\boldsymbol{X_2}'$*={ntpp(x,y)}.* We calculate $Dmin_{RCC8}(\boldsymbol{X_1}, \boldsymbol{X_2})=0$ where: *{sum(a,b)→x,c→y},* decompose the region *sum(a,b)*, then make the assignments. The formal correctness of this equivalence is provided by the *RCC8* theorem:

$$\forall x \forall y \forall z[[NTPP(sum(x,y),z)] \leftrightarrow [NTPP(x,z) \ \& \ NTPP(y,z)]]$$

5 Example Results

Figure 6 shows an example run using seven regions, and the associated ambiguity measures for each of the four theories

used, *RCC8*, *LOS14*, *Left/Right* and *Above/Below*, during an (arbitrary) sequence of movements of the regions. The images above the start, central and end points indicate the configuration of the regions at the indicated points.

At the start of the example each of the seven regions was disjoint. Each was then brought to the centre, some being hidden by the square region, some only overlapping. The central image indicates a stage just before this process was completed. Note that the central image also corresponds to the *LOS14* relation matrix shown in figure 5. The regions were then moved away again to result in the configuration shown in the right hand image. The shape, size (and colour) of the regions have no effect on the computation performed, but act as a useful aid to the user; only the relations as indicated are considered.

The overall higher ambiguity of the *RCC8* and, to a marginally lesser extent, the *LOS14* relations indicates that they represent a weaker description than those of the *Left/Right* and *Above/Below* relation sets. While all the regions are disconnected, the ambiguity of the *RCC8* and *LOS14* models is total (5040), as there is no possibility of resolving the ambiguity. The *RCC8* and *LOS14* measures are closely related, and *LOS14* reduces ambiguity only on the basis of relative depth to the viewpoint, but this effect is significant. The *Left/Right* and *Above/Below* measures are very similar, as expected, and largely independent of overlap. They are also highly effective at managing ambiguity. These Allen based relations are ineffective when identical objects are aligned to the orthogonal of the measure's axis. The total ambiguity measure indicates that, when taken together, this combination of similarity measures reduces the overall ambiguity to almost zero, note the single ambiguity arising at step 688.

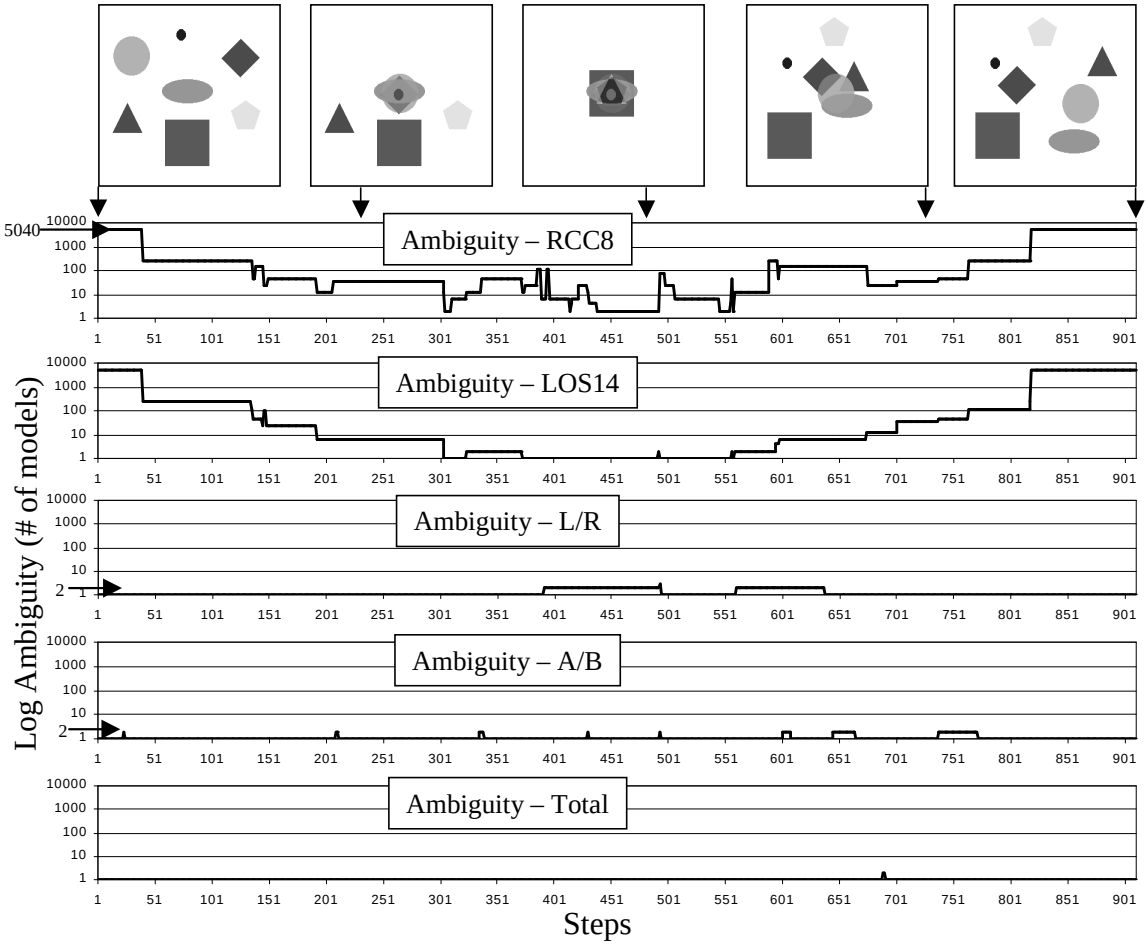

Figure 6: Sample run with seven regions showing the ambiguity measure for four theories.

Each of the five graphs in figure 6 record the ambiguity measure as computed during the sequence of movements. The seven regions represent a possible ambiguity of 5040 (7!) models. Note the use of logarithmic ambiguity scale.

6 Further work and suggested applications

In this section we discuss a few potential *QSR* applications for this approach. However, the technique is by no means restricted to this domain and can be applied to any theory that supports a *CND*, from which the *CND* distance is computed.

Where an extra-logical model is used to interpret sets of defined relations, any pair of models can be compared. This means we can either compare models where the identity of regions have been assumed and the the region labels have been pre-assigned (as in the case where we are comparing the *CND* distance between the sets: $\boldsymbol{X_1}$*= {dc(a,b)},* $\boldsymbol{X_2}$*={ntpp(a',b')},* assuming here:*{a→a',b→b'};* or where the mappings themselves are *not* pre-assigned and are to be computed by minimising the *CND* distance itself for all region combinations – for example where, $\boldsymbol{X_2}$ is an open-sentence: $\boldsymbol{X_1}$*= {dc(a,b)},* $\boldsymbol{X_2}$*={ntpp(x,y)}.* An example of this latter approach might be where we are using this technique to complement a region-matching algorithm used in a machine vision application, or for GIS data mining; where while we assume some correspondence exists between regions in a set of images (or set of geographical maps), an open question exists exactly what that correspondence is – see [Papadias and Delis, 1997]. As discussed above, this technique may also have an application in *symbol anchoring*; except here the distance function applied to the *CND* determines the best fit between objects in a registered set of images instead of matching properties of individual features.

One severe restriction assumed in this method, is that a one-to-one mapping is assumed between the sets of regions being compared. In many practical situations, however, this restriction rarely arises, as in a machine vision application where objects appear to appear and disappear, or fuse together or separate owing to the absence of noise, or from changes in resolution with variations in the observer-subject distance. This in part stems from the envisionment axioms used, where in the envisionment projection regions do not pass to null. However, each *CND* node can be augmented to allow for such transitions where, either a single region becomes null, or both do. An example of this may be where products of regions are being explicitly modelled, and where two overlapping regions separate (so the product becomes null), or separated, then overlap.

Two simple theories, *Left*, *RCC8* were to illustrate the method. Other natural extensions include encoding relational concepts such as being: *above/overlaps/below*, or being *smaller-than/same-size/larger-than*. For modelling geometrical solids, we could also consider *Aspect Graphs* themselves, define and decompose their node descriptions, and then model the *CND*s of these features. For example each aspect of a cube can be decomposed into sets of line junctions forming *acute/right/obtuse* angles that can in turn be represented as a *CND.*

Acknowledgements

Work described here has been supported by EPSRC project GR/S02891/01, "Spatial Reasoning and Perception in a Humanoid Robot". Special thanks to Murray Shanahan and to Antony Galton for constructive criticism on ideas worked into an earlier draft of this paper.

References

[Allen, 1981] Allen, J.F., An Interval-Based Representation of Temporal Knowledge, *Proc. IJCAI-81*, pp. 221-226.

[Allen, 1983] Allen, J.F., Maintaining Knowledge about Temporal Intervals, Comm. ACM 26(11), pp. 832-843.

[Cohn, 1997] Cohn, A.G., Qualitative spatial representation and reasoning techniques, *Proc. KI-97, Advances in Artificial Intelligence*, Springer LNAI 1303, pp. 1-30.

[Cui, *et al.*, 1992] Cui, Z, Cohn, A.G and Randell D.A., Qualitative simulation based on a logical formalism of space and time, *Proc. AAAI-92*, pp. 679-684.

[Cyr and Kimia, 2001], Cyr, C. M and Kimia, B. B. *3D Object Recognition Using Shape Similarity-Based Aspect Graph*, Int. Conf. on Computer Vision (ICCV'01), Vol. 1, pp. 254-261.

[Dickinson and Metaxas, 1997], Dickinson, S.J., Metaxas, D., Using Aspect Graphs to Control the Recovery and Tracking of deformable Models, In *International Journal of Pattern Recognition and Artificial Intelligence*, February 1997, pp 115-142.

[Freksa, 1992a] Freksa, C., Temporal Reasoning based on Semi-Intervals, *Artificial Intelligence*, Vol. 54, No. 1, pp. 199-227

[Freksa, 1992b] Freksa, C., Conceptual neighborhood and its role in temporal and spatial reasoning, In: Singh, M., Trave-Massuyes, L. (eds.) Decision Support Systems and Qualitative Reasoning, pp. 181-187, North-Holland, Amsterdam.

[Galton, 1994] Galton, A.P., Lines of Sight, *AISB Workshop on Spatial and Spatio-Temporal Reasoning*.

[Galton, 1999], Galton, A.P. (1999) The Mereotopology of Discrete Space, in *Proc. COSIT* 1999, pp. 251-266.

[Herbin, 1998] Herbin, S., Combining Geometric and Probabilistic Structure for Active Recognition of 3D Objects, in *European Conf. On Comp. Vision*, Berlin, Springer-Verlag LNCS 1407, pp. 748-764, .

[Koenderink and Van Doorn, 1979], Koenderink, J.J. and Van Doorn, A.J., The Internal representation of solid shape with respect to vision, *Biological Cybernetics*, Vol. 2, pp. 211-216.

[Lehmann and Cohn, 1994], The Egg/Yolk Relaibility Hierarchy: Semantic Data Integration using Sorts with Prototypes, Proc. Conf on Information Knowledge Management, ACM Press, pp. 272-279.

[Papadias and Delis, 1997], Papadia, D., and Delis, V., Relation-Based Similarity, in Proc. 5th ACM Workshop on GIS, Las Vegas, ACM Press.

[Randell, *et al.*, 1992] Randell, D. A., Cui, Z and Cohn A. G., A Spatial Logic Based on Regions and Connection, *Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning (KR-92)*, Morgan Kaufmann, San Mateo, pp. 165-176.

[Randell *et al*, 2001], Randell, D., Witkowski, M., and Shanahan, M., From Images to Bodies: Modelling Spatial Occlusion and Motion Parallax, *Proc 17th Int. Joint Conf. On Artificial Intelligence*, Morgan Kaufmann, Seattle, pp. 57-63.

[Randell and Witkowski, 2002], Building Large Composition Tables via Axiomatic Theories, in *Proc. KR 2002*, pp. 26-36.

[Roy and Stell, 2002] Roy, A.J. and Stell, J.G. (2002) A Qualitative Account of Discrete Space, GIScience 2002, pp. 276-290.