A Mereological Approach to Representing Spatial Vagueness*

Nicholas Mark Gotts *Anthony G Cohn*
Division of Artificial Intelligence
School of Computer Studies
University of Leeds, Leeds LS2 9JT, UK
Telephone: +44-113-233-6806
Fax: +44-113-233-5468
Email: {ngotts,agc}@scs.leeds.ac.uk

Abstract: Spatial reasoning is crucial in many AI application domains, such as robotics, qualitative and naive physics, and some types of planning. Qualitative reasoning is often required; and in many of these cases, *uncertainty* or *imprecision* about the spatial extent of particular entities has to be represented and coped with. This paper develops an axiomatisation of a relation of 'crisping' (reducing imprecision or vagueness) between pairs of 'vague' spatial regions: those with indeterminate boundaries. This axiomatisation is then related to the previously developed 'egg-yolk' representation of vagueness, originally designed by (Lehmann and Cohn 1994) for database applications, then applied to expressing spatial vagueness by (Cohn and Gotts 1994).

Keywords: Qualitative spatial reasoning, knowledge representation, uncertainty, logic

* We gratefully acknowledge the support of the EPSRC under grant GR/H/78955 and also the CEC under the Basic Research Action MEDLAR 2, Project 6471. We also acknowledge useful discussions with the rest of the QSR group at Leeds and Eliseo Clementini.

1 Introduction: Vague Spatial Regions

> 'Traffic chaos enveloped central Amsterdam today, as the QR community gathered from all parts of the industrialised world.'

Where exactly are or were the limits of the traffic chaos, central Amsterdam, the industrialised world? They do not exist: some points or areas would be definitely within the regions these entities occupy, some definitely outside, but for many there is no right or wrong answer. How to represent and reason about such entities, and the 'vague' regions they occupy, is the topic of this paper.

Qualitative spatial reasoning is crucial in many AI application domains, e.g. robotics, qualitative and naive physics, and many kinds of planning. Often, there is *uncertainty* or *imprecision* about the spatial extent of particular entities (physical objects, regions defined by some property such as temperature, and/or socially defined regions such as those owned by persons or organizations). We concentrate here on the development of a representational formalism for such vague spatial regions; we consider the development of an adequate representation essential prior to detailed consideration of questions relating to reasoning and applications. We are sceptical about the merits of 'fuzzy' approaches to indeterminacy, believing that their use of real number indices of degrees of membership and truth are both counterintuitive, and logically problematic. We have no space to argue this controversial viewpoint here; see (Elkan 1994) and responses for arguments on both sides.

We need to say at least some of the same sorts of things about vague regions as about 'crisp' ones, with precise boundaries: that one contains another (southern England contains London, even if both are thought of as vague regions), that two overlap (the Sahara desert and West Africa), or that two are disjoint (the Sahara and Gobi deserts). In these cases, the two vague regions represent the space occupied by distinct entities, and we are interested in defining a vague area corresponding to the space occupied by either, by both, or by one but not the other. We may also want to say that one vague region is a 'crisper' version of another. For example, we might have an initial (vague) idea of the extent of a mineral deposit, then receive information reducing the imprecision in our knowledge. Here, the vagueness of the vague region is a matter of our ignorance: the entity concerned actually occupies a fairly well-defined region — though perhaps any entity's limits will be imprecise to some degree. In other cases, vagueness appears intrinsic: consider an informal geographical term like 'southern England'. The uncertainty about whether particular places (north of London but south of Birmingham) are included cannot be resolved definitively: it is a matter of interpretational context. A contrasting example is the region occupied by a cloud of gas from an industrial accident. Here we have two sources of intrinsic vagueness: the concentration of the gas is likely to fall off gradually as we move out of the cloud; and its extent will also vary over time, so any temporal vagueness (e.g., if we are asked about the cloud's extent 'around noon') will result in increased spatial vagueness.

In these cases of intrinsic vagueness, there is a degree of arbitrariness about any particular choice of an exact boundary, and often, none is required. But *if* we decide to define a more precise version (either completely precise, or less vague but still imprecise), our choice version is by no means *wholly* arbitrary: we can distinguish more and less 'reasonable' choices of more precise description.

Distinguishing ignorance-based from intrinsic vagueness is important, but many of the same problems of representation and reasoning arise for both. The distinction between purely spatial entities ('regions'), and what *occupies* such regions (physical objects in the broadest sense, or social phenomena such as areas of legal ownership and political influence), will also be significant. We are concerned mainly with regions here, but need to remember that it is what occupies a region that makes it interesting [1].

Since we present here an axiomatisation which we claim captures important aspects of the intuitive concept of a vague region (and of relations between vague regions, and between vague and crisp ones), we distinguish the intuitive terms 'crisp' and 'vague' from their formally defined equivalents, Crisp and NonCrisp. When we want to talk about regions which may be either Crisp or NonCrisp, we will refer to 'OCregions' (for 'Optionally Crisp regions'). We will use the informal term 'blurring' as the converse of 'crisping'.

Section 2 develops a set of axioms that expresses some of the things we intuitively want to hold about the relations between alternative, more or less vague, estimates of the spatial extent of some extended entity, when this is imprecisely known or defined. We express the relation between two such estimates, of which one is a refinement of the other, as $X \prec Y$, read: 'region X is a crisping of vague region Y'[2]. A parallel is drawn between the axiom-set developed, and the axiom-sets for mereology (theory of part-whole relations) discussed by (Simons 1987). Toward the end of section 2, a start is made on exploring the possible relations between vague regions representing the spatial extent of different entities, revealing considerable complexity. Section 3 introduces a previously developed formalism which can be used for this latter purpose (Cohn and Gotts 1994), linking it to the approach of section 2 and showing that it can serve to reduce this complexity; it also suggests why the parallels with mereology explored in section 2 exist. Section 4 discusses future work, including a possible application area.

2 Possible Axiom-Sets for the $\prec$ Relation

We assume first order predicate calculus with equality as a logical basis. What sort of properties do we want $\prec$ to have? First, $\prec$ should be asymmetric, irreflexive and transitive:

(A1)$\forall X, Y[X \prec Y \rightarrow \neg Y \prec X]$

(A2)$\forall X, Y, Z[[X \prec Y \wedge Y \prec Z] \rightarrow X \prec Z]$

(A1) ensures $\prec$ is asymmetric and hence irreflexive, (A2) that it is transitive. It will be convenient to add some definitions as part of the basic package:

(D1)$X \preceq Y \equiv_{def} X \prec Y \vee X = Y$

(D2)$X \succ Y \equiv_{def} Y \prec X$

(D3)$\mathsf{MA}(X,Y) \equiv_{def} \exists Z[Z \preceq X \wedge Z \preceq Y]$

(D4)$\mathsf{NMA}(X,Y) \equiv_{def} \neg\mathsf{MA}(X,Y)$

(D5)$\mathsf{Crisp}(X) \equiv_{def} \neg\exists Y[Y \prec X]$

(D6)$\mathsf{NonCrisp}(X) \equiv_{def} \exists Y[Y \prec X]$

(D7)$X \prec\!\!\prec Y \equiv_{def} X \preceq Y \wedge \mathsf{Crisp}(X)$.

$\preceq$ can be read: 'crisper than or equal to'. $\succ$ is the inverse of $\prec$, while MA and NMA stand for two regions being 'mutually approximate' or 'not mutually approximate': i.e, having, or not having, some common region which is $\preceq$ both. Whether Crisp OCregions actually exist can be determined by adding an additional axiom, or left open. $X \prec\!\!\prec Y$ can be read 'X is a completely crisp version of Y', or 'X is a complete crisping of Y'.

These basic axioms and definitions alone allow us to show that some intuitively correct properties of vague regions and crispings hold. For example, we can show that if X and Y are NMA, and Z is a crisping of X, it cannot be MA with Y:

[1] Here, we allow regions themselves to be vague; it would be possible to assign the vagueness of spatial extent only to physical objects and other 'region-occupiers', or to a function mapping them to sets of Crispregions instead.

[2] We will use upper-case italic letters for variables ranging over OCregions.

(T1)$\forall X, Y, Z$
$[[\mathsf{NMA}(X,Y) \wedge Z \prec X] \rightarrow \mathsf{NMA}(Z,Y)]]$.
If T1 is not true, then:
(1)$\exists X, Y, Z[\mathsf{NMA}(X,Y) \wedge Z \prec X \wedge \neg\mathsf{NMA}(Z,Y)]$
(2)$\exists X, Y, Z[\mathsf{NMA}(X,Y) \wedge Z \prec X \wedge \mathsf{MA}(Z,Y)]$
(1, D4)
(3)$\exists X, Y, Z[\mathsf{NMA}(X,Y) \wedge Z \prec X \wedge$
$\exists W[W \preceq Z \wedge W \preceq Y]]$ (2, D3)
(4)$\exists X, Y, Z, W[\mathsf{NMA}(X,Y) \wedge Z \prec X \wedge$
$W \preceq Z \wedge W \preceq Y]$(3)
(5)$\exists X, Y, Z, W[\mathsf{NMA}(X,Y) \wedge W \preceq X \wedge W \preceq Y]$
(1, 4, A2)
(6)$\exists X, Y[\mathsf{NMA}(X,Y) \wedge \mathsf{MA}(X,Y)]$ (5, D3)
(7)$\perp$ (6).

At least one further axiom seems necessary (A3): given one crisping of a NonCrisp region, there is an alternative, incompatible one; if this is not assured, it is not clear what could be meant by one region being a crisping of another.

(A3)$\forall X, Y[X \prec Y \rightarrow \exists Z[Z \prec Y \wedge \mathsf{NMA}(X,Z)]]$.
Beyond these three axioms, there are several independent 'dimensions' along which sets of axioms for vague regions can be extended. These can best be explored using a parallel between the relations of crisping and blurring on the one hand, and part/whole relations on the other. If we regard the crisping relation $\prec$ as analogous to a proper part relation, the kinds of distinction made between possible mereological systems by (Simons 1987) are very similar to those that arise in the case of crisping.

2.1 A 'Minimal Extensional Mereology' and its OCregion Counterpart

Early in his investigation of mereology, Simons (Simons 1987, pp.25-30) constructs a 'minimal extensional mereology', taking as primitive the proper part relation, which he symbolises $\ll$. The logical basis of the system is:

(SA0) Any axiom set sufficient for first-order predicate calculus with identity.

The first two axioms for $\ll$ are (using a different syntax from Simons):

(SA1) $\forall x, y[x \ll y \rightarrow \neg(y \ll x)]$

(SA2) $\forall x, y, z[[(x \ll y) \wedge (y \ll z)] \rightarrow x \ll z]$

These, like (A1) and (A2) above, simply assert that the system's basic relation is a strict partial ordering. Simons goes on to define part (symbolised '$<$'), in a way that parallels our definition of $\preceq$.

Simons' next step is to note that an individual cannot have a *single* proper part, just as we noted that a vague region cannot reasonably have a *single* crisping. He considers various axioms which ensure that if an individual has a proper part, it has at least two. After defining overlapping ('$\circ$') and disjointness (for which we use '$\|$', as Simons' symbol is unavailable), in ways directly corresponding to our definitions of MA and NMA, Simons chooses:

(SA3)$\forall x, y[x \ll y \rightarrow \exists z[z \ll y \wedge z \| x]]$.
This axiom he refers to as the *Weak Supplementation Principle* (WSP), and in asserting that any individual with a proper part has another that is disjoint with the first, it corresponds exactly to our axiom A3 above.

The axiom set SA0-3 still permits various models Simons regards as unsatisfactory, in which overlapping individuals do not have a unique product or intersection. One of these is shown as figure 1 (lines join 'parts' (below) to 'wholes' (above)). Such models are ruled out by adding:

(SA6)$\forall x, y[x \circ y \rightarrow \exists z \forall w$
$[w < z \equiv w < x \wedge w < y]]$,
which ensures the existence of such a unique product.

The corresponding axiom for $\prec$ would ensure the

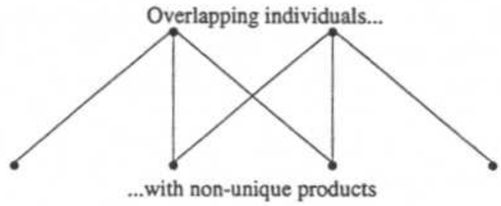

Figure 1: An Inadmissible Model of Part-Whole Relations

existence of a 'vaguest common crisping' (VCC), such that any other common crisping of the two is also a crisping of the VCC:

(A4)$\forall X, Y[\mathsf{MA}(X,Y) \rightarrow \exists Z \forall W$
$[W \preceq Z \equiv W \preceq X \wedge W \preceq Y]]$.
For the time being, we simply assert that we choose to add such an axiom to our set; later, we will indicate why.

2.2 Blurring, and the 'Complete Blur'

Simons notes that his 'minimal extensional mereology' of SA0-3 and SA6 is much weaker than the 'classical' mereological systems of (for example) (Tarski 1956), or (Leonard and Goodman 1940), even for small finite models, because the existence of sums or upper bounds is not guaranteed. (A mereological 'upper bound' (u.b.) of two individuals is an individual of which both are parts; a mereological 'sum' is an u.b. of which any part overlaps one of the two individuals summed.) He therefore considers how u.b.s and sums can be added. Correspondingly, we can consider how to ensure the existence of 'common blurrings' of pairs of OCregions. We will not follow Simons the whole way here, as we do not currently try to define *general* blurrings, over arbitrary sets of OCregions; only binary ones.

Simons suggests a considerable number of different axioms and combinations of axioms concerning the existence of mereological u.b.s, least u.b.s (an u.b. that is part of any other u.b. — if a sum exists it is always the least u.b., but a least u.b. need not be a sum) and sums, for pairs of individuals. It will be useful to describe his axioms, although we do not write them out:

(SA12) Guarantees the existence of an u.b. for any pair of individuals.

(SA7) Does the same, for pairs of overlapping individuals only.

(SA13) Guarantees the existence of a least u.b. for any pair of individuals.

(SA8) Does the same, for pairs of overlapping individuals only.

(SA14) Guarantees the existence of a sum for any pair of individuals.

(SA9) Does the same, for pairs of overlapping individuals only.

(SA10) Guarantees that a least u.b. exists for any pair with an u.b..

(SA11) Guarantees that a sum exists for any pair of individuals with an u.b..

(SA15) Guarantees that when a least u.b. exists, it is a sum.

(SA16) Guarantees a universal individual exists, which is an u.b. for any pair.

Clearly, a mereology does not need *all* of these: the unconditional existence guarantees (SA12, SA13 and SA14) subsume their conditional counterparts (respectively SA7; SA8 and SA10; SA9 and SA11), any guarantee of the existence of a sum makes the corresponding guarantee of a least u.b. or u.b. unnecessary, and if a least u.b. exists then so does an u.b.. SA14 in fact subsumes all of SA7-13 and SA15, and adding it to SA0-3 and SA6 creates a Boolean algebra minus a zero element. SA16, guaranteeing the existence of a universal region, is independent of all the others apart from SA7 and SA12, which it subsumes, but which are weaker. The strongest mereological system Simons proposes without the use of *general* least u.b.s or sums consists of SA0-3, SA6, SA14 and SA16.

The analogue of guaranteeing the existence of an u.b. is guaranteeing that of a 'common blurring' (CB) for a pair of OCregions: an OCregion to which both members of the pair are $\preceq$. A least u.b. would correspond to a 'crispest common blurring' (CCB) of a pair of OCregions, and a sum to an OCregion which is a crispest common blurring of which *any* crisping is MA to one of the pair (we could call it a 'blur sum' or BS). The analogue of the universal region of SA16 could be called the 'complete blur': a NonCrisp region which is a blurring of any other OCregion.

Which of these analogues are plausible? In particular, would we want to adopt analogues of SA14 and SA16? If not, what weaker substitutes should be adopted? There is not necessarily a 'right answer' here: different sets of axioms may be useful for different applications. Figure 2a shows why we probably would *not* want to adopt an analogue of SA14. In the left part of figure 2a, the inner and outer pairs of circles represent two NonCrisp regions: we might take the inner of each pair to represent an area definitely within the NonCrisp region, the outer an area definitely containing it. Imagine that these are, for example, two versions of the area enveloped by a flood received from different sources, each with some imprecision. What might their 'crispest common blurring' (CCB) look like? The right side of figure 2a shows one possibility: we retain the inner circle from the smaller of the pair, and the outer from the larger, and take these as representing inner and outer limits on a vaguer NonCrisp region which can be regarded as a blurring of both. However, the dashed circles suggest another possible crisping of this CCB, having *no* common crisping with either of the original pair. If this is admitted, then the CCB *cannot* be the pair's BS ('blur sum'): any crisping of the BS of two OCregions should be MA with one of the pair, just as any part of a mereological sum of a pair of individuals overlaps one of them.

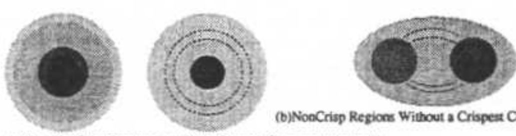

(b)NonCrisp Regions Without a Crispest Common Blurring?

(a) Two NonCrisp Regions and their 'Crispest Common Blurring'?

Figure 2: 'Crispest Common Blurrings'

If a CCB must exist for any pair of OCregions, but a BS may not, we need an analogue of SA13. We therefore adopt:

$$(A5)\forall X, Y[\exists Z[X \preceq Z \wedge Y \preceq Z \wedge \forall W[X \preceq W \wedge Y \preceq W \rightarrow Z \preceq W]]].$$

Depending on our interpretation of OCregions, we might even want to deny that a CCB must exist for any pair. Consider figure 2b: the lighter and darker circles represent the inner limits of two NonCrisp regions with a common outer limit (the containing oval). If we allow inner and/or outer limiting areas of a NonCrisp region to be discontinuous (multi-piece), then the CCB of these two NonCrisp regions can be arrived at by adding the two inner limiting regions to get the CCB's inner limiting region, and letting their common outer limit be the CCB's outer limit. However, if we do disallow discontinuous inner limiting regions, the inner limiting region of any u.b. of the two NonCrisp regions must include their inner limiting regions *and* a 'bridge' between them, such as those outlined by the upper, dashed pair of lines, or the lower, dotted ones. Clearly there is no unique such bridge and hence no CCB.

Similarly, whether we adopt an analogue of SA16 depends on the view we take of OCregions. One interpretation of a 'complete blur', a vague region which is a blurring of any other, is that *no* area is either definitely in such a region, or definitely out of it: for example, we might want to represent (our current state of knowledge concerning) the parts of the universe which have at any time been inhabited by extraterrestrial beings. If so, we should adopt:

(A6)$\exists X \forall Y[Y \preceq X]$.

If we only want to talk about NonCrisp regions with some limits on their blurring, we could adopt the negation of this:

(A6a)$\neg\exists X \forall Y[Y \preceq X]$[3].

To specify necessary and sufficient conditions for the identity of two OCregions, we need an axiom linking $\prec$ and $=$. We could choose:

(A7a)$X = Y \equiv \forall Z[X \prec Z \equiv Y \prec Z]$.

(X and Y are equal iff (X is a crisping of any Z iff Y is a crisping of that Z).) If we wanted to *define* equality in terms of $\prec$, rather than assuming it, we could make this a definition. Note that this axiom implies that there is at most one distinct OCregion with no blurrings (the complete blur, if it exists). If we axiomatised equality for OCregions instead as:

(A7b)$X = Y \equiv \forall Z[Z \prec X \equiv Z \prec Y]$,

this would require that there would be at most one OCregion with no crispings, that is, at most one Crisp region. In fact, there is an alternative we prefer to either of these – see (A7) in the next subsection.

2.3 Do Crisp Regions Exist?

Do Crisp OCregions exist? We have four alternative possibilities, according to whether A8, A8a, A8b, or none of the three is included in the set of axioms:

(A8)$\forall X[\exists Y[Y \preceq X \wedge \mathsf{Crisp}(Y)]]$.

(A8a)$\forall X[\exists Y[Y \prec X]]$

(A8b)$\exists X[\mathsf{Crisp}(X)] \wedge$
$\exists X \forall Y[Y \preceq X \rightarrow \exists Z[Z \prec Y]]$

The first of these asserts that all regions are Crisp, or have a complete crisping; the second that there are no Crisp regions. The third asserts that some regions are 'crispable', while others are not.

The existence or non-existence of Crisp regions is analogous to that of mereological 'atoms': individuals without proper parts. Here, bearing in mind the distinction made between spatial regions and the entities that occupy them, we will add A8 to our set of axioms: all NonCrisp OCregions are crispable. We could further assume that for every pair of OCregions such that one is a crisping of the other, there is a third which lies between them:

(A9)$\forall X, Y[X \prec Y \rightarrow \exists Z[X \prec Z \wedge Z \prec Y]]$.

This is a kind of 'denseness' axiom. Its mereological counterpart would assert that for any two regions of which one is a proper part of the other, there is a third which is a proper superpart of the first and a proper part of the second.

[3]Axioms belonging to our current axiom-set are numbered from (A1) to (A11). Possible alternatives are given an additional lower-case letter, e.g. (A6a).

In contrast to *regions*, we will allow that some physical and social entities are 'uncrispable'. In a fuller theory, linking spatial regions and entities with additional, non-spatial properties, we would need a function mapping things with spatial extent onto the regions they occupy. Some types of entities might be restricted to mapping onto NonCrisp OCregions; others, to Crisp OCregions. Furthermore, given a temporal theory, we could formally express the fact that for some kinds of physical entities, the regions would become crisper (or at least not less crisp) as our knowledge of their spatial extent grew, whilst for others, their crispness might not necessarily increase with time (e.g. the spatial extent of a river which varies over time).

Asserting that all NonCrisp OCregions are crispable gives us a way to deal with spatial relations between the entities associated with vague regions, by thinking in terms of the relations between members of their sets of complete crispings. As we noted earlier, we want to say that 'London' is a proper part of 'southern England', even though both are vague. We can now say that *any* complete crisping of the OCregion occupied by London is a proper part of *any* complete crisping of the OCregion occupied by southern England. We formally define the relation between the NonCrisp OCregions for London and southern England as follows:

(D8)$X \ll_{\mathsf{OC}} Y \equiv_{def}$
$\forall Z, W[[Z \prec\!\!\prec X \wedge W \prec\!\!\prec Y] \rightarrow Z \ll W]$.

Relations such as $\gg_{\mathsf{OC}}$ ($\gg$ is the inverse of $\ll$), $\circ_{\mathsf{OC}}$, $||_{\mathsf{OC}}$ can be defined analogously, and will have the same properties as their Crisp region counterparts: for example, $\ll_{\mathsf{OC}}$ will be asymmetric and transitive.

We could not define a relation $=_{\mathsf{OC}}$ in the same fashion, even if we had chosen to define $=$ in terms of $\ll$: for any two NonCrisp OCregions, even if they are identical, there must be complete crispings of the two that are *not* $=$. Instead, we adopt the following axiom :

(A7)$X = Y \equiv$
$\forall Z[Z \prec\!\!\prec X \rightarrow \exists W[W \prec\!\!\prec Y \wedge Z = W]] \wedge$
$\forall Z[Z \prec\!\!\prec Y \rightarrow \exists W[W \prec\!\!\prec X \wedge Z = W]]$

(for any complete crisping of either OCregion there is an equal complete crisping of the other). This also suggests the possibility of defining analogous weaker versions of the other mereological relations between OCregions e.g.:

(D9)$X \ll_{\mathsf{OC}_a} Y \equiv_{def}$
$\forall Z[Z \prec\!\!\prec X \rightarrow \exists W[W \prec\!\!\prec Y \wedge Z \ll W]] \wedge$
$\forall Z[Z \prec\!\!\prec Y \rightarrow \exists W[W \prec\!\!\prec X \wedge W \ll Z]]$.

This relation might hold where the stronger does not, for example between the OCregions for 'the wettest parts of Britain' and 'the wetter parts of Britain'. However we crisp the second, there will be a complete crisping of the first that is a proper part of it, and however we crisp the first, there will be a proper superpart complete crisping of the second, but *every*

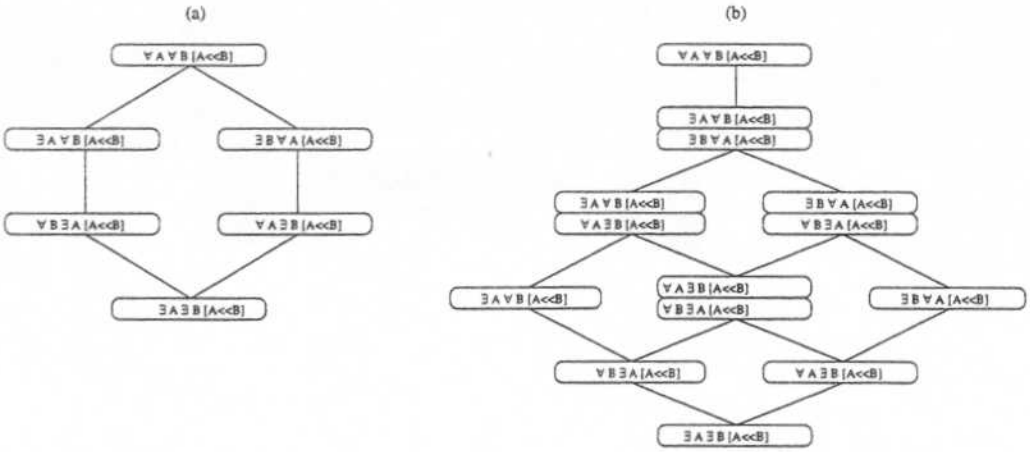

Figure 3: Possible assertions about $\ll$ between two sets of individuals

complete crisping of the first may not be a proper part (or even a part) of *every* complete crisping of the second. The relation $\ll_{\mathsf{OC}_a}$ and its inverse are still asymmetric and transitive. In fact, there are at least *nine* possible weakened versions of $\ll_{\mathsf{OC}}$, of which *eight* retain this property. Figure 3 shows the relationship between $\ll_{\mathsf{OC}}$ and the nine weaker possibilities: 3a illustrates the logical relationships between six assertions about pairs of *sets* of individuals, members of which may be related by $\ll$. '$\forall A \forall B[A \ll B]$' stands for the assertion that every member of one set (the As) is a proper part of every member of the other (the Bs); '$\exists A \forall B[A \ll B]$' means that there is some member of the As that is a proper part of every member of the Bs, and so on. Straight lines connect stronger assertions (above) to weaker (below). Figure 3b shows all the ten logically distinct conjunctions that can be produced from these six assertions. Each of the ten corresponds to a possible version of $\ll_{\mathsf{OC}}$, with the Xs and Ys being the sets of complete crispings of two OCregions: only the one at the bottom fails to retain transitivity.

3 The Egg-Yolk Theory

Given all these possibilities for generalizing mereological relationships from crisp to vague regions, an alternative approach is worth considering. (Lehmann and Cohn 1994) suggest an approach to spatial vagueness that involves using two (or more) concentric subregions, indicating degrees of 'membership' in a vague region. (In the simplest, two-subregion case, the inner is called the 'yolk', the outer the 'white', and the inner and outer subregions together the 'egg'.) Lehmann and Cohn first suggested the egg-yolk approach in the context of the problem of integrating heterogeneous databases, where the notions of 'regions' and 'spatial relations' are used metaphorically to represent sets of domain entities, and their relations. It was developed as an approach to expressing spatial vagueness itself in (Cohn and Gotts 1994). Here, we employ it as a means to understanding the alternative approach we develop.

The egg-yolk formalism as developed in (Lehmann and Cohn 1994) allows for just 5 'base relations' (DR, PO, PP, PPI and EQ) between any egg-egg or yolk-yolk pair, or any egg and the yolk belonging to another egg. (A yolk is always a PP of the corresponding egg.) EQ, PP, PPI and DR correspond to Simons' $=$, $\ll$, $\gg$, and $||$, except that they are defined in terms of a primitive $\mathsf{C}(x, y)$ (connection) (Randell, Cui and Cohn 1992), with which we need not be concerned. PO (partial overlap) is simply the relation that holds when none of the other four do: the five base relations (henceforth 'RCC-5' for historical reasons) form a jointly exhaustive and pairwise disjoint (JEPD) set: exactly one holds between any pair of (crisp) regions.

RCC theory asserts the existence of a universal region, Us (of which every other region is a PP), and provides quasi-Boolean functions (quasi-Boolean because no 'empty' or 'zero' region exists) on pairs of regions: $\mathsf{compl}(x)$ (the complement of a region in Us), and the region-sum, region-intersection, and region-difference of two regions ($\mathsf{sum}(x, y)$, $\mathsf{prod}(x, y)$ and $\mathsf{diff}(x, y)$). All these functions other than sum may return a NULL object instead of a region (e.g. prod will do so if the two argument regions are DR). The distinction between regions and NULL is dealt with using the sorted logic LLAMA (Cohn 1987).

The RCC-5 set produces 46 possible relations between a pair of egg-yolk pairs (see figure 4). In (Cohn and Gotts 1994) we argued that these 46 could be identified with the possible relations between complete crispings of two vague regions. Here, we take the OCregions of section 2 as our initial representation of vague regions, and show how the egg-yolk representation can be used to provide additional constraints for OCregion theory.

At first glance, there is one apparent problem with the egg-yolk approach: the most obvious interpreta-

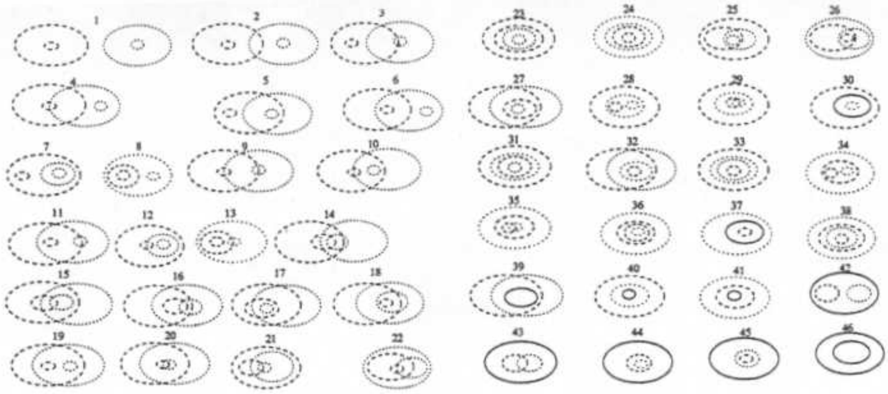

Figure 4: The 46 possible relationships between two egg-yolk pairs

tion is that it replaces the precise dichotomous division of space into 'in the region' and 'outside the region' of the basic RCC theory by an equally precise trichotomous division into 'yolk', 'white' and 'outside' — and this appears contrary to a key intuition about vagueness: that not only is there a 'doubtful' zone around the edges of a vague region, but that zone itself has no precise boundaries. So we want a way of using the egg-yolk formalism that is consistent with this.

We link the OCregions of section 2 (and the corresponding theory), with ordered pairs of RCC-5 regions, the first of the pair being a part, but not necessarily a proper part, of the second. If it *is* a PP, then the pair is an egg-yolk pair in the sense of (Lehmann and Cohn 1994), and the OCregion is NonCrisp. If not, the OCregion is Crisp. We now link the $\prec$ predicate of OCregion theory with the egg-yolk approach. We define a function ey to map an OCregion to an egg-yolk pair, and two functions eggof and yolkof, to map such egg-yolk pairs to the RCC-5 region comprising its egg and yolk respectively. We will normally write $\mathsf{ey}(X)$ as $\hat{X}$ for notational convenience. We have the following axiom for egg-yolk pairs:

(A10) $\forall X \mathsf{P}(\mathsf{yolkof}(\hat{X}), \mathsf{eggof}(\hat{X}))$

We then assert the following additional axiom concerning $\prec$:

$$\begin{aligned}&(\text{A11})\forall X, Y[X \prec Y \rightarrow\\ &\quad[[\mathsf{PP}(\mathsf{eggof}(\hat{X}), \mathsf{eggof}(\hat{Y}))\wedge\\ &\qquad\mathsf{P}(\mathsf{yolkof}(\hat{Y}), \mathsf{yolkof}(\hat{X}))]\vee\\ &\quad[\mathsf{P}(\mathsf{eggof}(\hat{X}), \mathsf{eggof}(Y))\wedge\\ &\qquad\mathsf{PP}(\mathsf{yolkof}(\hat{Y}), \mathsf{yolkof}(\hat{X}))]]].\end{aligned}$$

This axiom links $\prec$ to the predefined RCC-5 relations by an *implication*, not an *equivalence*: we do *not* specify that if the specified RCC relations hold between $\mathsf{eggof}(\hat{X})$, $\mathsf{yolkof}(\hat{X})$, $\mathsf{eggof}(\hat{Y})$ and $\mathsf{yolkof}(\hat{Y})$, the CR relation holds between X and Y, but these relations *must* hold for the $\prec$ relation to do so. We leave undefined what additional conditions, if any, must be met. This gives us the kind of indefiniteness in the *extent* of vagueness, or 'higher-order vagueness', that intuition demands. Consider the vague region 'beside my desk'. This can be regarded in OCregion theory as a NonCrisp region. There are some precisely defined regions, such as a cube 10*cm* on a side, 5*cm* from the right-hand end of my desk, and 50*cm* from the floor, that are undoubtedly contained within any reasonable complete crisping of this NonCrisp region. Others, such as a cube 50*m* on a side centred at the front, top right-hand corner of the desk, contain any such reasonable crisping. These two could correspond to the 'yolk' and 'egg' of an egg-yolk pair constituting the NonCrisp region 'beside my desk', forming a very conservative inner and outer boundary on its possible range of indefiniteness. However, some OCregions (Crisp and NonCrisp) lying between this pair would *not* make a reasonable crisping of this region: consider a volume including the 'yolk' of the pair, plus a layer one centimetre deep at the very top of the white. This meets all the conditions for a crisping of the specified OCregion, but is an absurd interpretation of 'beside my desk'. In general, we need not precisely specify the limits of acceptability. For specific applications, we could add further conditions on acceptable crispings (such as preserving particular topological features or relative proportions in different dimensions), and perhaps assert that (for that application) these conditions are sufficient.

Configuration 1 in figure 4, given the interpretation of OCregion region theory in terms of egg-yolk pairs of RCC-5 regions outlined here, clearly shows a pair of NonCrisp regions such that any pair of complete crispings of the two must be DR. Taking the left-hand egg-yolk pair as representing NonCrisp region X, and the righthand one NonCrisp region Y:

$$\begin{aligned}&\forall V, W[[V \mathrel{-\!\!\!\prec\!\!\prec} X \wedge W \mathrel{-\!\!\!\prec\!\!\prec} Y] \rightarrow\\ &\qquad\mathsf{DR}(\mathsf{eggof}(\hat{V}), \mathsf{eggof}(\hat{W}))]\end{aligned}$$

Similarly, configuration 2 represents a pair of NonCrisp regions such that, for any complete crisping of either, we can choose a complete crisping of the other that is DR from it, and there are also some complete crisping pairs of the two that are PO. (Cohn and Gotts 1994) shows how each of the 46 configurations can be distinguished in terms of the possible results of replacing one or both of the egg-yolk pairs with a single region-boundary lying within the white of the egg — a complete crisping of the vague region represented by the egg-yolk.

This way of interpreting OCregion theory explains why we found so many parallels with Simons' mereology. Under the egg-yolk interpretation, an OCregion amounts to a three-way division of Us into yolk, white, and non-egg. If we consider a set of all such divisions where no part of space is in the 'yolk' of one division and the 'non-egg' of another, we have a mereological system with all the possible precise boundaries as 'atoms'. Crisping expands yolk and/or non-egg at the expense of the white. One OCregion being a crisping of another is like one individual being a proper part of another because the white of the first is a proper part of the white of the second. We have a plausible candidate for the VCC of two MA OCregions: the VCC's yolk could be the sum of the yolks of its two blurrings, its egg the prod of the two blurrings' eggs (which, if the two are MA, must exist as a region). Similarly, the yolk of the CCB of *any* two OCregions might be defined as the prod of their yolks; its egg as the sum of their eggs.

The implications of these identifications remain to be explored. However, the egg-yolk model of the OCregion axioms does appear to provide a straightforward way to define OCregion extensions of the compl, sum, prod and diff functions defined within RCC. Using forms such as '$\langle \mathsf{eggof}(X), \mathsf{yolkof}(X)\rangle$'to represent the egg-yolk pairs of RCC-5 regions corresponding to NonCrisp OCregions, we extend the definitions of compl, sum, prod and diff as follows:

(D10) $\mathsf{compl}(\hat{X}) =_{def} \langle \mathsf{compl}(\mathsf{yolkof}(\hat{X})), \mathsf{compl}(\mathsf{eggof}(\hat{X}))\rangle$

(D11) $\mathsf{sum}(\hat{X}, \hat{Y}) =_{def} \langle \mathsf{sum}(\mathsf{eggof}(\hat{X}), \mathsf{eggof}(\hat{Y})), \mathsf{sum}(\mathsf{yolkof}(\hat{X}), \mathsf{yolkof}(\hat{Y}))\rangle$

(D12) $\mathsf{prod}(\hat{X}, \hat{Y}) =_{def} \langle \mathsf{prod}(\mathsf{eggof}(\hat{X}), \mathsf{eggof}(\hat{Y})), \mathsf{prod}(\mathsf{yolkof}(\hat{X}), \mathsf{yolkof}(\hat{Y}))\rangle$

(D13) $\mathsf{diff}(\hat{X}, \hat{Y})) =_{def} \mathsf{prod}(\hat{X}, \mathsf{compl}(\hat{Y}))$.

In the cases of (D12) and (D13), one or both of the components of the 'output' egg-yolk pair may be NULL (as indeed is also the case with (D10) when taking the complement of the 'complete blur' guaranteed by (A6)). These cases require further investigation to ensure their correct formal treatment.

Figure 5 shows the 46 possibilities assuming that RCC-5 calculus is used to relate eggs and yolks. We will briefly investigate here how egg-yolk theory can be used to explore the relations between vague regions expressed in figure 3. Table 1 shows the various sets of egg-yolk configurations which satisfy the upper five quantificational schema of figure 3a, but considering three of the other four mereological relations along with the proper part relation (we omit PPI as it is simply the dual of PP). Only PO differentiates all five quantificational schema. DR, PP (and PPI) collapse the distinction between ∃X∀Y and ∀Y∃X (and dually, between ∃Y∀X and ∀X∃Y). EQ only distinguishes two cases. Thus, in the egg-yolk interpretation there are not quite as many possible relations as figure 3 might suggest. Moreover, egg-yolk theory gives us a way to reason with vague regions using the existing mechanism of the RCC calculus.

4 Discussion: From Representation to Reasoning and Application

The paper has shown how two approaches to the representation of vague spatial regions, originally developed in competition with each other (although partly within the same research team), can be brought together, shedding new light on both and on the representational problems they were designed to solve. The next stage of our research is to investigate the potential of the two approaches, individually and in combination, to support *reasoning* about vague regions.

The mereologically-based approach outlined in section 2 has concentrated on capturing aspects of the relations between different, more or less vague versions of the same region. Here, we have shown that these different versions of a region can be regarded as forming a mereological structure, but most of the axioms adopted simply assert the existence of one vague region given the existence of others: more work is required on developing axioms which constrain the relations that can exist between different versions of a region in useful ways.

The other kind of relation we want a vague region formalism to support is that between different spatially extended entities, each occupying a vague region. For example, we want to be able to deduce that if London is a part of southern England, and southern England is disjoint from central Europe, then London is disjoint from central Europe — even if we regard all three as having indeterminate boundaries. For this kind of reasoning, we could make use of relations such as $\ll_{OC}$, defined toward the end of section 2; as noted there, $\ll_{OC}$ (and its inverse) are transitive. However, the link from the mereologically-based formalism to the egg-yolk approach made in section 3 makes available the resources developed within the larger RCC theory of spatial reasoning of which the latter forms a part (Randell et al. 1992). These include 'transitivity' or 'composition' tables of relations (Cui, Cohn and Randell 1993), which specify the set of possible relations between regions x and z, given those between the pairs x and y, and y and z. Further work is planned on this link between the two formalisms, and on applying their fusion to practical problems in qualitative spatial representation. The application of the theory to vague temporal intervals will also be investigated.

It is no accident that our motivating examples in this paper have been geographical: dealing with geographical regions with indeterminate boundaries is an important issue in current work on GIS (geographical information systems), as the recent meeting specif-

	∀X∀Y	∃X∀Y	∀Y∃X	∃Y∀X	∀X∃Y
PP	24	18,24,26,32, 33,37,38,45	18,24,26,32, 33,37,38,45	8,13,22,24,26,34, 35,36,37,38,41	8,13,22,24,26,34, 35,36,37,38,41
EQ	∅	∅	33,40,45,46	∅	36,41,44,46
DR	1	1,2,3,5,7	1,2,3,5,7	1,2,4,6,8	1,2,4,6,8
PO	14	3,5,7,9,10,11,12,14, 15,16,18,19,20,21,28,29,33	3,5,7,9,10,11,12,14,15, 16,18,19,20,21,28,29,32,33	4,6,8,9,10,11,13,14,15, 16,17,19,20,22,34,35,36	4,6,8,9,10,11,13,14,15,16, 17,19,20,22,27,34,35,36

Table 1: The egg-yolk configurations which satisfy various relations as defined by various quantificational schema.

ically on that issue (Burrough and Frank 1994) indicates. Digital mapping, satellite photography, and GIS have made greatly increased amounts of information available to geographers in recent years; the difficulty they face is to make effective use of this information. Appropriately organised summaries of what is known about spatial relationships of interest — and of what remains uncertain or ill-defined — are of central importance in making available information useful in hypothesis-building and testing, and it is here that we see work on the topology of regions with indeterminate boundaries as potentially of great value. Consider, for example, the questions surrounding global climate change, forest destruction and desertification, and similar actual or possible large-scale environmental problems. Information about the spatial properties and relations of many different kinds of entities and variables, from many different sources of varying reliability, is potentially relevant to unravelling the complex causal interactions involved.

To make the example more specific, what are the leading causes of the degradation and destruction of tropical forests? Logging, mining, road-building, cattle-ranching and cash crops for export, fuelwood gathering, increasing population, the movement of displaced people, government-promoted migration, patterns of land-ownership and the legal framework regulating it, measures purportedly aimed at conservation such as the declaration of national parks — all these and more have been blamed in one region or another (Hecht and Cockburn 1989, Brown and Pearce 1994). Complex interactions between factors differing from region to region are involved. To design and implement effective responses, a qualitative understanding of such interactions is necessary (though certainly not sufficient, as powerful economic and political interest groups are involved). For such an understanding to be developed, it is necessary to know how the spatial distributions of the putative causal factors relate to each other, and to the pattern of destruction in various parts of the world. Some of these factors have well-defined and readily determined spatial boundaries: land ownership and the laws regulating it, for example; but the majority do not.

In an initial search for plausible hypotheses, we would want to ensure that we do not miss potentially important possibilities, and therefore to make generous estimates of the maximum spatial extent of various possible causative factors. A large number of factors might then be shown as possibly overlapping the areas of rapid destruction. However, we would also wish to avoid erroneously concluding that a particular factor is essential, and would therefore wish to have safe minimum estimates of these factors' spatial extent as well. Given similar maximum and minimum estimates of the area of rapid forest destruction, we would then have a (probably large) set of qualitative hypothesis about the causally relevant factors. If we assume that all the important factors lie among those we are considering, then any Boolean combination of factors whose spatial extent could coincide with the areas of rapid destruction (under some crisping of the vague regions concerned) is a possible qualitative hypothesis.

It would then be useful to know the effect on this set of hypotheses of tightening the limits on the spatial extent of the areas of rapid destruction, and/or one or more of the putative causal factors, in order to help us decide what additional information would be most useful in distinguishing between them. To simplify matters, suppose we had just two possible factors in mind, and the areas of rapid destruction lay within the areas covered by factor A whatever estimates we used, but within the areas covered by factor B if and only if we used a narrow definition of the areas of rapid destruction, and a generous estimate of the areas affected by that factor. Additional information about the spatial extent of factor B's effect is more likely to be immediately useful in narrowing the range of hypotheses than information about that of factor A. To reach this conclusion, we must ourselves be employing some intuitive 'logic of vague regions'. Assigning essentially arbitrary numerical measures of 'degree of membership' to particular points, as in a 'fuzzy' approach, appears to us unlikely to be helpful. Clearly, however, the non-numerical approach we have begun to develop here needs considerably more work before its usefulness can be assessed.

References

Brown, K. and Pearce, D. (eds): 1994, *The Causes of Tropical Deforestation*, UCL Press.

Burrough, P. and Frank, A. M. (eds): 1994, *GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries*, Francis Taylor. To appear.

Cohn, A. G.: 1987, A more expressive formulation of many sorted logic, *Journal of Automated Reasoning* **3**, 113–200.

Cohn, A. G. and Gotts, N. M.: 1994, The 'egg-yolk' representation of regions with indeterminate boundaries, *in* P. Burrough and A. M. Frank (eds), *Proceedings, GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries,*, Francis Taylor. To appear.

Cui, Z., Cohn, A. G. and Randell, D. A.: 1993, Qualitative and topological relationships in spatial databases, *in* D. Abel and B. C. Ooi (eds), *Advances in Spatial Databases*, Vol. 692 of *Lecture Notes in Computer Science*, Springer Verlag, Berlin, pp. 293–315.

Elkan, C.: 1994, The paradoxical success of fuzzy logic, *IEEE Expert* **9**(4), 3–8. Followed by responses and a reply.

Hecht, S. and Cockburn, A.: 1989, *The Fate of the Forest: Developers, Destroyers and Defenders of the Amazon*, Verso.

Lehmann, F. and Cohn, A. G.: 1994, The EGG/YOLK reliability hierarchy: Semantic data integration using sorts with prototypes, *Proc. Conf. on Information Knowledge Management*, ACM Press, pp. 272–279.

Leonard, H. S. and Goodman, N.: 1940, The calculus of individuals and its uses, *Journal of Symbolic Logic* **5**, 45–55.

Randell, D. A., Cui, Z. and Cohn, A. G.: 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.

Simons, P.: 1987, *Parts: A Study In Ontology*, Clarendon Press, Oxford.

Tarski, A.: 1956, Foundations of the geometry of solids, *Logic, Semantics, Metamathematics*, Oxford Clarendon Press, chapter 2. trans. J.H. Woodger.