Qualitative Spatial Representation Languages with Convexity *

Ian Pratt

1 Introduction

In recent years, there has been considerable interest within the AI community in *qualitative* descriptions of space. The idea is roughly this. Suppose we have a language in which we can say such things as "region a is convex" or "region b is a part of region c" and so on; then maybe such a language would enable us to characterize the spatial properties of everyday objects to an extent that suffices for—say—object recognition or route planning or commonsense mechanical reasoning. And maybe—so the thought goes—the use of merely qualitative descriptions might enable us to do all this while avoiding the computational complexity and error-sensitivity of numerical (coordinate) descriptions. Thus, the hope is that, by choosing an appropriate qualitative spatial description language, we might increase the effectiveness of an artificially intelligent agent operating in or reasoning about space.

However, such qualitative spatial representation languages are inevitably balanced on a semantic knife-edge: too little expressivity, and they are useless for the everyday tasks we want them for; too much, and they exhibit the over-precision which motivated qualitative representation languages in the first place. The aim of this paper is to demonstrate how sharp that knife-edge is, and thus to establish some limits on what such qualitative spatial description languages might be like. Specifically, we show that, once we can represent the property of convexity and the part-whole relation—modest assumptions by any standards—n-tuples of real polygons are completely determined by the sets of formulas they satisfy upto the fixing of three reference points. To be sure, we are not the first to express skepticism about the possibility of such languages, but this is, as far as we are aware, the first time such skepticism has been put on so firm a mathematical footing.

*With thanks to Dominik Schoop, Oliver Lemon, Tony Cohn and Ernest Davis for their comments. The author gratefully acknowledges the support of the Leverhulme Trust (Grant number F/120/AQ) and the European Commission's TMR Programme (contract number ERBFMBICT972035).

2 Background

Many systems of qualitative spatial reasoning have been proposed (for a recent survey, see [32]), and we certainly do not propose to review them all here. Of particular interest, however, is the language proposed by Randell, Cui and Cohn [25], whose primitives are the binary predicate symbol $C(x, y)$ and the function-symbol $\mathrm{conv}(x)$. The exact interpretation of this language is deliberately left open, partly, it seems, so as to allow for non-standard interpretations. But for definiteness, Randell, Cui and Cohn invite us to suppose that variables range over some unspecified set of well-behaved subsets of $\mathbb{R}^2$ (called *regions*), that $C(x, y)$ is satisfied by a pair of regions just in case their topological closures have some points in common, and that $\mathrm{conv}(x)$ is taken to denote the convex hull of its argument.

The original language envisaged by by Randell, Cui and Cohn incorporates the full power of first-order logic. Thus, despite the limited range of primitives, there is no restriction on the complexity of the formulas employing them. However, not all systems of qualitative spatial representaion allow such arbitrary logical combinations. In particular, Davis, Gotts and Cohn [12] have recently analysed a *constraint language* with two primitive binary predicates, $\mathrm{EC}(x, y)$ and $\mathrm{PP}(x, y)$, together with the unary predicate $\mathrm{conv}(x)$ and a stock of constant symbols $a, b, c, \ldots$. Here, the regions over which the names are interpreted are the regular closed sets of the Euclidean plane (i.e., those equal to the closure of their interior). The predicate $\mathrm{EC}(x, y)$ ("external contact") is satisfied by a pair of regions if and only if they have some boundary points but no interior points in common; $\mathrm{PP}(x, y)$ ("proper part") is satisfied by a pair of regions if and only if the first is a proper subset of the second; and $\mathrm{conv}(x)$ is satisfied by a region if and only if it is convex. In this context, a *constraint* is an atomic formula—e.g. $\mathrm{EC}(a, b)$ or $\mathrm{conv}(c)$—and assertions in the language, or *constraint networks*, as they are called, are simply conjunctions of constraints. Thus, this language lacks the full power of first-order logic, because there is no general quantification over regions. Nevertheless, Davis, Gotts and Cohn are able to prove a strong expressivity result. Specifically, they show that any two arrangements of bounded regions that are not related by an affine transformation can be distinguished by a constraint network.

In this paper, we study the first-order language whose signature consists of the two predicates $x \leq y$ and $\mathrm{conv}(x)$. We again take the variables of $\mathcal{L}$ to range over some set of regions in the plane, though we consider various choices for this set. The predicate $x \leq y$ is satisfied by a pair of regions if and only if the first is a subset of the second (the symbol $\leq$ is natural because we deal with Boolean algebras of regions); and $\mathrm{conv}(x)$ is satisfied by a region if and only if it is convex. We choose $\mathcal{L}$ because of its spartan set of primitives (smaller even than that studied by Davis, Gotts and Cohn) and great expressive power. We prove three main results. The first, similar to that of Davis, Gotts and Cohn, concerns the ability of $\mathcal{L}$-formulas to distinguish between collections of regions, given certain choices as to the set of regions over which the variables of $\mathcal{L}$ range. More precisely, we show that, by identifying regions with *polygons*, any two n-tuples of regions not related by an affine transformation satisfy different $\mathcal{L}$-formulas; moreover, by identifying regions with *rational polygons*, every n-tuple of regions satisfies a single $\mathcal{L}$-formula which determines it upto an affine transformation. Our second result concerns the ability of $\mathcal{L}$ formulas to define useful relations not specifically associated with affine geometry. More precisely, we show that the various *topological* relations are $\mathcal{L}$-definable, depending on which subsets of the plane we count as regions. Our third result concerns the ability of $\mathcal{L}$ to distinguish different choices for the set of regions over which quantification ranges. More precisely, we examine $\mathcal{L}$-interpretations based on the rational polygons, the polygons, the definable regular open sets, and the entire set of regular open sets; we show that all these interpretations make different sets of $\mathcal{L}$-sentences true.

More controversially, we claim that the language $\mathcal{L}$—and hence any first-order language in which its primitives are definable—does not deserve the label 'qualitative' at all. Coordinate position (upto the fixing of any three noncollinear reference points) has not disappeared from $\mathcal{L}$ at all: it has simply been made hopelessly and needlessly cumbersome. The point is a deeper one than the simiplicity of our proofs suggest, for it is intimately related to to fundamental results on the coordinatizability of axiomatically presented geometries. We argue that these results indicate that the some of the current interest in qualitative (understood: *nonquantitative*) spatial representation languages may be misplaced.

A note on accessibility. This paper uses various results from model theory and affine geometry. Since some readers may be unfamiliar with these fields, we have explained standard concepts and notation throughout; we have also been more explicit with some of the proofs than is usual in purely mathematical treatments.

3 Defining the ontology

Let $\mathcal{L}$ be the first-order language whose signature consists of the two predicates $x \leq y$ and $\text{conv}(x)$. We assume the variables of $\mathcal{L}$ range over some set A of subsets of $\mathbb{R}^2$. We define A in various ways below. Interpreting the primitives $\leq$ and conv as indicated above, we obtain, in the sense of model theory, an $\mathcal{L}$-structure $\mathfrak{A}$. $\mathfrak{A}$ determines, for any formula $\phi(\bar{x})$ with free variables $\bar{x} = x_1, \ldots x_n$, and any n-tuple of regions $\bar{a} = a_1, \ldots, a_n$ from A, whether $\bar{a}$ satisfies the property expressed by $\phi(\bar{x})$; if so, we write $\mathfrak{A} \models \phi[\bar{a}]$. (Note: the square brackets indicate that the $\bar{a}$ are *elements of the set* A, rather than *names in the language* $\mathcal{L}$. In fact, $\mathcal{L}$ has no names.) An $\mathcal{L}$-formula ϕ with no free variables is called an *$\mathcal{L}$-sentence*. Given an $\mathcal{L}$-structure $\mathfrak{A}$ and an $\mathcal{L}$-sentence ϕ, $\mathfrak{A}$ will determine a *truth value* for ϕ; if ϕ is true according to $\mathfrak{A}$, we write $\mathfrak{A} \models \phi$. We shall ask the following questions concerning the expressivity of $\mathcal{L}$. First: for given choices of A, to what extent is a given n-tuple $\bar{a}$ of A characterized by the $\mathcal{L}$-formulas it satisfies? Second: for given choices of A, what general properties (other than those expressed by the primitives) can be captured by $\mathcal{L}$-formulas?

Third: if we *vary* the choice of A, what difference does this variation make to the set of $\mathcal{L}$-sentences which turn out true? To answer these questions, we must examine the possible choices for A, and it is to this question we now turn.

It is fairly standard in treatments of qualitative spatial reasoning to confine attention to *regular* sets.

Definition 3.1 *Let X be a topological space and $x \subseteq X$. Then the set $\bigcup\{y \subseteq X | y \text{ open}, y \cap x = \emptyset\}$ is an open set in X called the* pseudocomplement *of x, written x'. We say that $x \subseteq X$ is* regular *if $x = x''$.*

It can easily be shown that a set if regular if and only if it is equal to the interior of its closure. Basically, we can think of regular sets in $\mathbb{R}^2$ as open sets with no internal cracks or point-holes. (Note that our use of *regular* coincides with the more usual term *regular open*. We could of course have used regular closed sets instead of regular open sets; nothing in the resulting account would have changed.) The following well-known theorem underlies the importance of the regular sets to qualitative spatial reasoning. We state it here without proof. (See, e.g. Koppelberg [18], pp. 26 and 60.)

Theorem 3.1 *Let X be a topological space. Then the set of regular sets in X forms a Boolean algebra* $\mathrm{RO}(X)$ *with top and bottom defined by $1 = X$ and $0 = \emptyset$, and Boolean operations defined by $x.y = x \cap y$, $x + y = (x \cup y)''$ and $-x = x'$.*

Thus the product of two regular sets is simply their intersection. The sum of two regular sets x and y is a little more complicated; *very* roughly, it is the union of $x \cup y$ with internal cracks filled in. Finally, the pseudocomplement, $-x$, of a regular set x is simply that part of the plane not occupied by x or its boundary. Let $\mathbb{R}^2$ denote the real plane with the usual Euclidean topology. Our domain of discourse will form a Boolean sub-algebra of $\mathrm{RO}(\mathbb{R}^2)$. Clearly, the Boolean functions $+$, $\cdot$ and $-$, as well as the constants 0 and 1, will all be $\mathcal{L}$-definable, and so we shall use these symbols in $\mathcal{L}$-formulas without further ado.

However, concentrating on regular sets by no means eliminates all strangely behaved spatial entities from the domain of quantification. For it is well-known that regular sets exist which could not possibly correspond to the space taken up on surfaces by ordinary objects and geographical abstractions. (See Pratt and Lemon [24] p. 232 for discussion of an example.) In order to rule out such pathalogical regular regions, and to simplify the technical discussion, we assume for most of this paper that all regions are regular *polygons*. We shall relax this assumption later.

Formally, any line in $\mathbb{R}^2$ cuts $\mathbb{R}^2$ into two residual domains, which we shall call *half-planes*. It is easy to see that these sets are regular, with each being the pseudocomplement of the other. Hence, we can speak about the sums, products and complements of half-planes in $\mathrm{RO}(\mathbb{R}^2)$.

Definition 3.2 *A* (real) polygon *is a Boolean combination of finitely many half-planes in $\mathbb{R}^2$.*

We denote the set of polygons by H, and will sometimes refer to it as the *polygonal domain.*

Of course, H is not the only well-behaved spatial domain we might choose. If a line is defined by an equation $ax + by + c = 0$, where a, b and c are rational numbers, we call it a *rational line*; and if a half-plane is bounded by a rational line, we call it a *rational half-plane.* Now we define:

Definition 3.3 *A* rational polygon *is a Boolean combination of finitely many rational half-planes in* $\mathbb{R}^2$.

We denote the set of rational polygons by G, and will sometimes refer to it as the *rational polygonal domain.*

Clearly, G and H are Boolean sub-algebras of $\mathrm{RO}(\mathbb{R}^2)$. We note that, in computer systems designed to manipulate plane spatial data, approximation by polygons is nearly universal. And since H—or perhaps, more modestly, G—is arguably the spatial ontology recognized by such systems, it follows that these polygonal ontologies are adequate to model planar arrangements in nearly all practical situations. (Remember, there is no limit to the number of straight edges in the boundaries of polygons.) We shall return later to the question of liberalizing this ontology; however, for the moment, we shall take either G or H to be the set of regions over which the variables of $\mathcal{L}$ range.

Having defined two suitable domains of discourse, it is straightforward to set up the $\mathcal{L}$-structures they give rise to.

Definition 3.4 *A nonempty set* $X \subseteq \mathbb{R}^2$ *is said to be* convex *if, for all* (ξ_1, ξ_2), $(\xi_1', \xi_2') \in X$ *and for all* $\alpha \in [0, 1]$, *we have* $(\alpha.\xi_1 + (1-\alpha).\xi_1', \alpha.\xi_2 + (1-\alpha).\xi_2') \in X$. *The empty set* $\emptyset$ *is taken to be nonconvex.*

Thus, $X \neq \emptyset$ is convex if the line segment connecting any two points in X lies within X.

Definition 3.5 *We define the* (real) polygonal model $\mathfrak{H}$ *to have the domain* H *and the following interpretations of the predicates in* $\mathcal{L}$*:*

1. $\leq^{\mathfrak{H}} = \{(a, b) \in H^2 | a \subseteq b\}$
2. $\mathrm{conv}^{\mathfrak{H}} = \{a \in H | \ a \text{ is convex}\}$

We define the rational polygonal model $\mathfrak{G}$ *exactly as for* $\mathfrak{H}$ *but with* $\mathfrak{H}$ *and* H *replaced throughout by* $\mathfrak{G}$ *and* G *respectively.*

Clearly, G is countable, whereas H is uncountable, so $\mathfrak{G}$ and $\mathfrak{H}$ are different structures; just how different we will discover anon. For the moment, however, many of our observations will apply to both $\mathfrak{G}$ and $\mathfrak{H}$; therefore we use $\mathfrak{F}$ to refer indeterminately to either. Note that we take the domain of $\mathfrak{F}$ to be F.

Having defined our domains of discourse, we can immediately establish some limitations on the expressive power of $\mathcal{L}$ over these domains.

Definition 3.6 *Let τ be a 1–1, onto mapping from $\mathbb{R}^2$ to itself. If τ maps lines in $\mathbb{R}^2$ to lines in $\mathbb{R}^2$ then we say that τ is a* collineation *(of the plane). If τ maps rational lines in $\mathbb{R}^2$ to rational lines in $\mathbb{R}^2$ then we say that τ is a rational* collineation *(of the plane).*

Collineations and rational collineations of the plane have a simple characterization.

Definition 3.7 *A mapping $\tau : \mathbb{R}^2 \to \mathbb{R}^2$ is an* affine *transformation if it is of the form*

$$\tau((\xi_1, \xi_2)) \quad = \quad (\xi_1, \xi_2)M + (\alpha, \beta),$$

where M is a nonsingular matrix and $(\alpha, \beta) \in \mathbb{R}^2$. If, in addition, α, β and the elements of M are all rational, we say that τ is rational affine.

Affine transformations are continuous with continuous inverses, whence:

Lemma 3.1 *Let τ be an affine transformation and a and b regular sets. Then $\tau(-a) = -\tau(a)$, $\tau(a \cdot b) = \tau(a) \cdot \tau(b)$ and $\tau(a + b) = \tau(a) + \tau(b)$.*

The following result is standard (see, e.g. Neumann *et al.* [20]).

Theorem 3.2 *Let τ be a 1–1 onto mapping from $\mathbb{R}^2$ to itself. Then τ is a collineation if and only if it is an affine transformation; moreover, τ is a rational collineation if and only if it is a rational affine transformation.*

Given any two triples p_1, p_2, p_3 and q_1, q_2, q_3 of noncollinear points in the plane, there is a (unique) affine transformation τ mapping p_i to q_i $(1 \leq i \leq 3)$. If the p_i and q_i have rational coordinates, then τ is rational. The importance of affine transformations for our purposes can be seen immediately:

Lemma 3.2 *Let τ be an affine transformation. Then τ induces an $\mathfrak{H}$-automorphism. If in addition τ is rational, then τ induces a $\mathfrak{G}$-automorphism.*

Proof: Clearly, affine transformations map half-planes to half-planes, and rational affine transformations map rational half-planes to rational half-planes. Hence, by lemma 3.1, any affine transformation maps H 1–1 onto itself, and any rational affine transformation maps G 1–1 onto itself. Moreover, by the definition of convexity, affine transformations must map convex elements of F to convex elements of F and nonconvex elements of F to nonconvex elements of F. □

We can now establish an upper bound on the expressivity of $\mathcal{L}$.

Definition 3.8 *Two n-tuples of regions $\overline{a}$ and $\overline{b}$ in H (or in G) are said to be* affine equivalent, *written $\overline{a} \sim \overline{b}$, if there is a (rational) affine transformation τ taking $\overline{a}$ to $\overline{b}$.*

Theorem 3.3 *Let $\bar{a}$ and $\bar{b}$ be n-tuples of F with $\bar{a} \sim \bar{b}$. Then $\bar{a}$ and $\bar{b}$ satisfy the same formulas in $\mathfrak{F}$.*

Proof: Immediate by lemma 3.2. □

Thus, $\mathcal{L}$-formulas cannot distinguish between affine-equivalent arrangements. However, as we shall show in the next section, theorem 3.3 has a converse: $\mathcal{L}$-formulas can always distinguish between non-affine-equivalent arrangements. As we shall see, this converse has serious consequence for the proposal to use $\mathcal{L}$ as a qualitative spatial representation language.

4 Expressivity in the polygonal models

Our next task is to establish the definability of various useful concepts in $\mathfrak{F}$. Most of these results are so obvious that we state them without proof. The formula

$$x > 0 \wedge -x > 0 \wedge \mathrm{conv}(x) \wedge \mathrm{conv}(-x) \quad (1)$$

is satisfied in $\mathfrak{F}$ by a if and only if a is a half-plane. We can informally identify such a region with the line bounding it, and we will sometimes refer a as a line. If, in addition, $a \in G$, then we refer to a as a rational line. (Hence, all rational lines are lines, but not vice versa.) To clarify proofs, we shall use letters l_1, l_2, m, m' etc. to denote half-planes when we want to think of them as lines. In addition, as a shorthand, we will often use the variables u_1, u_2, v, v' etc. as a shorthand for variables x, y, constrained to satisfy the formula (1).

Various relations involving lines are definable in $\mathfrak{F}$. Two lines l and l' are coincident (the same undirected line) if and only if l, l' satisfies the formula concident(u, u') given by

$$u = u' \vee u = -u' \ .$$

Likewise, l and l' are parallel if and only if l, l' satisfies the formula $\Pi(u, u')$ given by

$$\neg\mathrm{concident}(u, u') \wedge (u.u' = 0 \vee u. - u' = 0 \vee -u.u' = 0 \vee -u. - u' = 0).$$

Let the lines l_1, l_2 and l_3 be pairwise nonparallel and noncoincident. By inspection, l_1, l_2 and l_3 can divide the plane into either 6 or 7 residual domains, as shown in figure 1 a) and b). Since each residual domain corresponds to a nonzero product of the form $\pm l_1. \pm l_2. \pm l_3$ (remembering that the l_i are really half-planes), it is clear that we can write $\mathcal{L}$-formulas $\Gamma(u_1, u_2, u_3)$ and $\Delta(u_1, u_2, u_3)$ such that, for any lines l_1, l_2 and l_3, $\mathfrak{F} \models \Gamma(l_1, l_2, l_3)$ if and only if l_1, l_2 and l_3, all meet at a single point as in figure 1 a), and $\mathfrak{F} \models \Delta(l_1, l_2, l_3)$ if and only if l_1, l_2 and l_3 meet pairwise at three points as in figure 1 b).

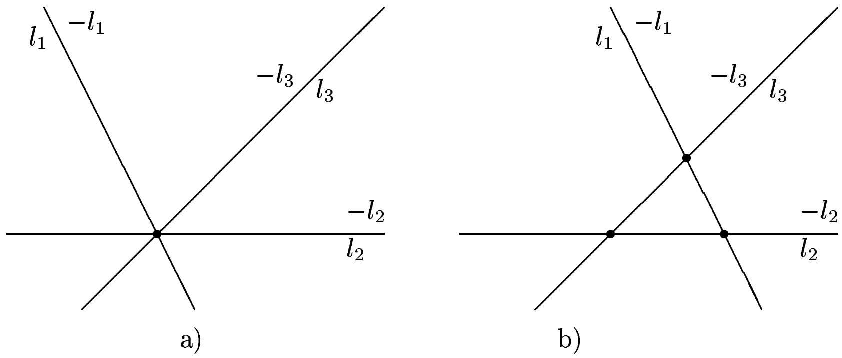

Figure 1: Intersection of three lines: a) $l_1.l_2.l_3 = -l_1. - l_2. - l_3 = 0$; b) only $-l_1. - l_2. - l_3 = 0$.

Lemma 4.1 *Let $\Delta_{\text{bound}}(u_1, u_2, u_3)$ be the formula*

$$\Delta(u_1, u_2, u_3) \wedge -u_1 \cdot -u_2 \cdot -u_3 = 0.$$

Then $a \in F$ satisfies

$$\exists u_1 \exists u_2 \exists u_3 (\Delta_{\text{bound}}(u_1, u_2, u_3) \wedge x = u_1 \cdot u_2 \cdot u_3)$$

if and only if a is a triangle. Hence, the property of being a bounded region is definable in $\mathcal{L}$.

Proof: If $\mathfrak{F} \models \Delta(l_1, l_2, l_3)$ then $l_1 \cdot l_2 \cdot l_3$ is the central triangular region (rather than one of the unbounded regions) in figure 1 b) if and only if the product $-l_1 \cdot -l_2 \cdot -l_3$ is zero. □

The following simple result contains the main idea in the following proofs.

Lemma 4.2 *Let the lines l_1, l_2, l_3 satisfy Δ, and let m_1, m_2, m_3 be parallel to l_1, l_2, l_3, respectively, intersecting at points O, P, Q, R, S as shown in figure 2. Then $\overline{OP} = \overline{PQ}$.*

Proof: Triangles OPR and SRP are congruent (2 angles and side); triangles SRP and PQS are congruent (2 angles and side). □

Lemma 4.3 *Let the lines l_1, l_2, l_3 satisfy Δ, and let n be any integer. Suppose that l_2 and l_3 intersect l_1 at (distinct) points O and P respectively (figure 3).*

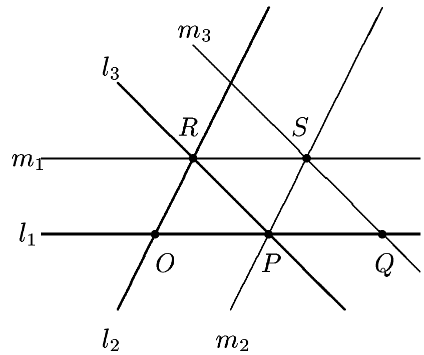

Figure 2: Copying a triangle formed by l_1, l_2, l_3. The labels "l_i" and "m_i" are written next to the lines to indicate that it does not matter which half-planes they denote.

Then there exists a formula $\phi_{l_1,l_2,l_3,n}(u_1, u_2, u_3, v)$ of $\mathcal{L}$ such that, for any line m intersecting l_1 at a point Q (where Q may coincide with either O or P) , $\mathfrak{F} \models \phi_{l_1,l_2,l_3,n}[l_1, l_2, l_3, m]$ if and only if $\overline{OQ} = n\overline{OP}$.

Proof: Immediate by repeating the construction of lemma 4.2 a number of times (in either direction along l_1) and using the formulas $\Gamma(u_1, u_2, u_3)$ and $\Pi(u_1, u_2)$. □

We can generalize this lemma to cover rational multipliers:

Lemma 4.4 *Lemma 4.3 holds when n is replaced by a rational number q.*

Proof: By two applications of lemma 4.3 and by use of the formula $\Gamma(u_1, u_2, u_3)$. □

We now have the main lemma on which all the other results depend.

Lemma 4.5 (Rational fixing lemma) *Let the rational lines l_1, l_2, l_3 satisfy Δ, and let m be any other rational line. Then there exists a formula $\phi_{l_1,l_2,l_3,m}(u_1, u_2, u_3, v)$ of $\mathcal{L}$ such that, for any line m', $\mathfrak{F} \models \phi_{l_1,l_2,l_3,m}[l_1, l_2, l_3, m']$ if and only if $m = m'$.*

Proof: If m is coincident with any of the l_i, then the result follows immediately. If m passes through the point of intersection of two of the l_i and is parallel to the third, then again the result follows immediately. Let us assume that neither of these conditions obtains.

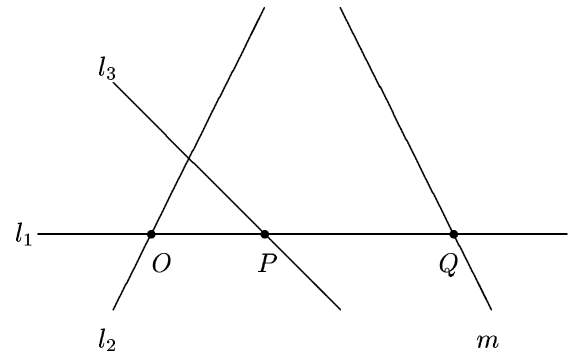

Figure 3: Fixing the intersection of l_1 and m.

By permuting the l_i if necessary, we can suppose that m intersects both l_1 and l_2 in two *distinct* points. (Otherwise, we would have one of the two special cases above.) Let l_2, l_3 and m intersect at l_1 at points O, P, and Q, respectively, as in figure 3. Since all lines are rational, $\overline{OP} = q\overline{OQ}$ for some rational number, so that by lemma 4.4 there is a formula fixing point Q given l_1, l_2 and l_3. By switching l_1 and l_2, we fix a second point on m and the result follows. □

We call $\phi_{l_1,l_2,l_3,m}(u_1, u_2, u_3, v)$ a *fixing formula* and say that $\phi_{l_1,l_2,l_3,m}$ *fixes m with reference to* l_1, l_2, l_3. Notice that if l_1^*, l_2^*, l_3^* satisfy Δ and m^* is a line such that $\mathfrak{G} \models \phi_{l_1,l_2,l_3,m}[l_1^*, l_2^*, l_3^*, m^*]$, then $\phi_{l_1,l_2,l_3,m}$ also fixes m^* with reference to l_1^*, l_2^*, l_3^*.

Definition 4.1 *A formula $\phi(\overline{x})$ is said to be* affine complete *in $\mathfrak{F}$ if any two n-tuples in F satisfying ϕ in $\mathfrak{F}$ are affine-equivalent. Similarly, a set of formulas $\Phi(\overline{x})$ is said to be* affine complete *in $\mathfrak{F}$ if any two n-tuples in F satisfying every formula in Φ in $\mathfrak{F}$ are affine-equivalent.*

Notation: If τ is an affine transformation, then we will denote the induced automorphism (of either $\mathfrak{G}$ or $\mathfrak{H}$) by $\tilde{\tau}$.

The following result establishes the descriptive power of the language $\mathcal{L}$ over the model $\mathfrak{G}$.

Theorem 4.1 *Every n-tuple in G satisfies an affine-complete formula in $\mathfrak{G}$.*

Proof: Let $\overline{a} = a_1, \ldots, a_n$ be an n-tuple in G. Choose an N-tuple of lines (half-planes) $\overline{l} = l_1, \ldots l_N$ such that every a_i is expressible as a Boolean combination of the l_j ($1 \leq i \leq n$, $1 \leq j \leq N$). Without loss of generality, assume $N \geq 3$ and that l_1, l_2, l_3 satisfy Δ_{bound} (as defined in lemma 4.1). Let $\psi_{\overline{l}}(u_1, \ldots, u_N)$

be the formula:

$$\Delta(u_1, u_2, u_3) \wedge \bigwedge_{4 \leq j \leq N} \phi_{l_1,l_2,l_3,l_j}(u_1, u_2, u_3, u_j) \ ,$$

where the ϕ_{l_1,l_2,l_3,l_j} are fixing-formulas as defined in lemma 4.5. Let $\bar{l}'$ be any N-tuple of rational lines. We show that $\mathfrak{G} \models \psi_{\bar{l}}[\bar{l}']$ only if $\bar{l} \sim \bar{l}'$. Suppose that $\mathfrak{G} \models \psi_{\bar{l}}[\bar{l}']$. Then for each j $(4 \leq j \leq N)$

$$\mathfrak{G} \models \phi_{l_1,l_2,l_3,l_j}[l'_1, l'_2, l'_3, l'_j] \ ,$$

and, in addition, l'_1, l'_2, l'_3 satisfy Δ_{bound}. It is a standard result of affine geometry that there is a (unique) rational affine transformation τ such that $\tilde{\tau}(l_j) = l'_j$ $(1 \leq j \leq 3)$. By lemma 3.2, we have, for each j $(4 \leq j \leq N)$,

$$\mathfrak{G} \models \phi_{l_1,l_2,l_3,l_j}[l'_1, l'_2, l'_3, \tilde{\tau}(l_j)] \ .$$

Since these are all fixing formulas, we have $\tau(l_j) = l'_j$, so that $\bar{l} \sim \bar{l}'$. Hence $\psi_{\bar{l}}(u_1, \ldots, u_N)$ is affine complete in $\mathfrak{G}$. Thus, $\bar{a}$ satisfies a formula $\phi(\bar{x})$ of the form

$$\exists u_1, \ldots, u_N(\psi_{\bar{l}}(u_1, \ldots, u_N) \wedge \bigwedge_{1 \leq i \leq n} x_i = t_i(u_1, \ldots, u_N))$$

in $\mathfrak{G}$, where the $t_i(u_1, \ldots, u_N)$ are Boolean combinations of the $u_1, \ldots, u_N$, and this formula is visibly affine complete in $\mathfrak{G}$. □

We now turn to the descriptive power of the language $\mathcal{L}$ over the model $\mathfrak{H}$. Let us say that l is a *bounding half-plane* of a region $a \in A$ if the boundary of l coincides with some finite line segment on the boundary of a. Then we have

Lemma 4.6 *If $a \in H$ then a is expressible as a Boolean combination of its bounding half-planes.*

Proof: Extend to infinity in both directions all the line segments in the boundary of a. Then a is obviously expressible as a sum in $\text{RO}(\mathbb{R}^2)$ of the residual domains of the resulting lines. □

Lemma 4.7 *Let $a \in H$ and let l be a half-plane. Then there is a $\mathcal{L}$-formula* $\text{bhp}(u, x)$ *satisfied by l, a if and only if l is a bounding half-plane of a.*

Proof: Let $\text{bhp}(u, x)$ be

$$\exists y(y{\cdot}x > 0 \wedge \text{conv}(y) \wedge ((y \leq u{+}x \wedge y{\cdot}u > 0 \wedge y{\cdot}u{\cdot}x = 0) \vee (y \leq -u{+}x \wedge y{\cdot}{-u} > 0 \wedge y{\cdot}u{\cdot}x = 0)))$$

□

The following result establishes the descriptive power of the language $\mathcal{L}$ over the model $\mathfrak{H}$.

Theorem 4.2 *Every n-tuple in H satisfies an affine-complete set of formulas in* $\mathfrak{H}$.

Proof: Let $\bar{a}$, $\bar{b}$ be n-tuples of H, such that $\bar{a} \not\sim \bar{b}$. We find a formula $\phi(\bar{x})$ separating $\bar{a}$ and $\bar{b}'$ in $\mathfrak{H}$.

Recall the abbreviation $\Delta_{\text{bound}}(u_1, u_2, u_3)$ introduced in lemma 4.1. We suppose first that $\bar{a}$ satisfies a formula $\exists \bar{u} \psi(\bar{u}, \bar{x})$ with $\psi(\bar{u}, \bar{x})$ of the form

$$\bigwedge_{1 \leq j \leq N} \bigvee_{1 \leq i \leq n} \text{bhp}(u_j, x_i) \wedge \bigwedge_{1 \leq i \leq n} x_i = t_i(\bar{u}) \wedge \Delta_{\text{bound}}(u_1, u_2, u_3),$$

where $\bar{u}$ is an N-tuple of variables and $t_i(\bar{u})$ $(1 \leq i \leq n)$ are Boolean expressions in the $\bar{u}$. Intuitively, this formula says that the bounding half-planes involved in the elements of $\bar{a}$ can be ordered such that the first three of them form a triangle in the sense of figure 1 b). We deal with the rare cases where $\bar{a}$ satisfies no such formula below. If $\bar{b}$ does not satisfy $\exists \bar{u} \psi(\bar{u}, \bar{x})$, we are done. So suppose otherwise, and let $\bar{l}$, $\bar{m}$ be N-tuples of half-planes in H such that

$$\begin{array}{rcl} \mathfrak{H} & \models & \psi[\bar{l}, \bar{a}] \\ \mathfrak{H} & \models & \psi[\bar{m}, \bar{b}]. \end{array}$$

Now let Π be the set of permutations π of $\{1, \ldots, N\}$ (where we write $\pi(\bar{m}) = m_{\pi(1)}, \ldots, m_{\pi(N)}$) such that

$$\mathfrak{H} \quad \models \quad \psi[\pi(\bar{m}), \bar{b}].$$

And for each $\pi \in \Pi$, let α_π be an affine transformation which maps l_1, l_2, l_3 to $m_{\pi(1)}, m_{\pi(2)}, m_{\pi(3)}$. We know that such α exists, because both l_1, l_2, l_3 and $m_{\pi(1)}, m_{\pi(2)}, m_{\pi(3)}$ form triangles.

Since $\bar{a} \not\sim \bar{b}$, and since $a_i = t_i[\bar{l}]$, $b_i = t_i[\pi(\bar{m})]$ for all i $(1 \leq i \leq n)$ and all $\pi \in \Pi$, we must have $\alpha(\bar{l}) \neq \pi(\bar{m})$, so that it is certainly possible to find a formula $\psi_\pi(\bar{u})$ not satisfied by $\pi(\bar{m})$ but satisfied by $\alpha(\bar{l})$ and hence by $\bar{l}$.

Now let $\phi(\bar{x})$ be the formula

$$\exists \bar{u} (\psi(\bar{u}, \bar{x}) \wedge \bigwedge_{\pi \in \Pi} \psi_\pi(\bar{u})).$$

Then it is clear that $\bar{a}$ satisfies $\phi(\bar{x})$. To see that $\bar{b}$ does not satisfy $\phi(\bar{x})$, suppose that $\bar{m}', \bar{b}$ satisfies $\psi(\bar{u}, \bar{x})$. Then obviously, $\bar{m}' = \pi(\bar{m})$ for some $\pi \in \Pi$, whence $\bar{m}'$ does not satisfy $\psi_\pi(\bar{u})$.

For the case where neither $\bar{a}$ nor $\bar{b}$ satisfies any formula $\exists \bar{u} \psi(\bar{u}, \bar{x})$, with $\psi(\bar{u}, \bar{x})$ as described above, we note that the bounding lines of the $\bar{a}$ (except possibly for one) must all be parallel (and similarly for $\bar{b}$). Then $\bar{a} \not\sim \bar{b}$ implies that the ratios of distances between these parallel lines are not all the same, in which case it is routine to find a distinguishing formula. □

5 Expressing topological properties in $\mathcal{L}$

So far, we have analysed $\mathcal{L}$'s expressivity in terms of the ability of its formulas to characterize individual n-tuples of F. However, that does not tell us very much about which general relations are definable by $\mathcal{L}$-formulas. One particularly salient class of relations are the topological relations such as the primitive $\mathrm{EC}(x, y)$ of Davies, Gotts and Cohn, and the many topological relations that can be defined in terms thereof. In this section, we show that various topological relations are indeed $\mathcal{L}$-definable over certain choices for A.

Within the structures $\mathfrak{F}$, it is very easy to define the topological relations $C(x, y)$ and $\mathrm{EC}(x, y)$, where, recall, $C(x, y)$ is satisfied by a pair of regions just in case their topological closures have a point in common, and $EC(x, y)$ is satisfied by a pair of (open) regions just in case they have no point in common, but their topological closures do.

Theorem 5.1 *The relations $C(x, y)$ $\mathrm{EC}(x, y)$ are $\mathcal{L}$-definable in $\mathfrak{F}$.*

Proof: If a' and b' are convex regions, then their closures have a point in common just in case it is impossible to separate a' and b' with two parallel (noncoincident) lines. If $a, b \in F$, then it is obvious that their closures have a point in common just in case there exist convex polygons $a' \leq a$, $b' \leq b$ whose closures have a point in common. □

We note that this definability result can certainly be generalized to domains of quantification richer than the polygons. However, it is doubtful whether it holds, for example, if the regions are taken to be all regular open sets.

Our next result has wider applicability. It is standard to take an open set to be *connected* just in case it is not the union of two nonempty, disjoint, open sets. A maximal connected subset of a set is called a *component* of that set. This leads us to define:

Definition 5.1 *Let A be a subset of* $\mathrm{RO}(\mathbb{R}^2)$. *We say that A is* closed under components *if, for all $a \in A$, if b is a component of a, then $b \in A$.*

Lemma 5.1 *The sets G, H and* $\mathrm{RO}(\mathbb{R}^2)$ *are all closed under components.*

Proof: Easy. □

For the next lemma, we introduce the abbreviation $x = x_1 \oplus x_2$ for the formula

$$x = x_1 + x_2 \wedge x_1 > 0 \wedge x_2 > 0 \wedge x_1 . x_2 = 0.$$

Lemma 5.2 *Let $a \in A$ with $G \subseteq A \subseteq \mathrm{RO}(\mathbb{R}^2)$, and let A be closed under components. Then a is connected if and only if $\mathfrak{A} \models \mathrm{con}[a]$ where $\mathrm{con}(x)$ is the formula:*

$$\forall x_1 \forall x_2 (x = x_1 \oplus x_2 \to \exists y (y \leq x \wedge y.x_1 > 0 \wedge y.x_2 > 0 \wedge \mathrm{conv}(y))).$$

Proof: Suppose that a is connected. Then if $a = a_1 \oplus a_2$, we have $a_1 \cup a_2 \neq a = a_1 + a_2$. So let p be a point in a but not in $a_1 \cup a_2$, and since $G \subseteq A$, let $b \in A$ be a small convex region such that $p \in b \leq a$. Certainly, then, $p \in \mathcal{F}(a_1) \cap \mathcal{F}(a_2)$, so that $b \cdot a_1$ and $b \cdot a_2$ are nonzero. Hence, a satisfies $\mathrm{con}(x)$.

Conversely, suppose a is not connected. Let $a_1 \in A$ be a component of a and $a_2 = a \cdot -a_1$ so that $a = a_1 \oplus a_2$. Moreover, suppose that b is such that $b \leq a$ with $b \cdot a_1$ and $b \cdot a_2$ nonzero. Then since a_1 is a component of a, b cannot be connected (for then $a_1 + b$ would be a connected subset of a strictly including a_1), so certainly b cannot be convex. Hence a does not satisfy $\mathrm{con}(x)$ in $\mathfrak{A}$. □

Thus, the property of being connected is $\mathcal{L}$-definable over a wide range of choices for the set A of regions.

We note in addition the following corollary concerning decidability. We call the set of sentences true in a structure $\mathfrak{A}$ the *theory* of $\mathfrak{A}$, denoted $\mathrm{Th}(\mathfrak{A})$. This set is said to be *decidable* if there is an effective computational procedure for determining whether any given sentence belongs to it—that is, whether any given sentence is true in $\mathfrak{A}$.

Corollary 1 *The sets of sentences $\mathrm{Th}(\mathfrak{G})$ and $\mathrm{Th}(\mathfrak{H})$ are both undecidable.*

Proof: Let $\mathcal{L}'$ be the first-order language whose primitives are $\mathrm{bounded}(x)$ and $\mathrm{con}(x)$, and let $\mathfrak{G}'$ be the $\mathcal{L}'$-structure with domain of quantification G and primitives interpreted in the obvious way; similarly for $\mathfrak{H}'$. It follows from Pratt and Lemon [24] (using lemma 4.14 and the proof of lemma 5.2) that $\mathrm{Th}(\mathfrak{G}') = \mathrm{Th}(\mathfrak{H}')$. Moreover, it is shown in Dornheim [13] that $\mathrm{Th}(\mathfrak{G}')$ is undecidable. By lemmas 4.1 and 5.2, we can regard $\mathrm{Th}(\mathfrak{G}')$ as, in effect, a subset of $\mathrm{Th}(\mathfrak{G})$. The result then follows. □

6 Distinguishing between models

So far, we have analysed $\mathcal{L}$'s expressivity in terms of the ability of its formulas to characterize individual n-tuples of F, and to express topological concepts. In this section, we show how the *sentences* of $\mathcal{L}$ (formulas with no free variables) distinguish between various domains of quantification—that is, between various choices as to what subsets of the plane count as bona fide regions. On the way, we prove a few easy corollaries of theorems 4.1 and 4.2 which are of general mathematical interest.

The first two of these corollaries rely on the notion of *atomicity* (see, e.g. Chang and Keisler [8], sec. 2.3).

Definition 6.1 *A formula $\phi(\bar{x})$ is said to be* complete *in a theory T if, for all formulae $\theta(\bar{x})$, exactly one of $T \models \phi \to \theta$ and $T \models \phi \to \neg\theta$ hold. A model $\mathfrak{A}$ is said to be* atomic *if any n-tuple $\bar{a}$ in A satisfies a formula $\phi(\bar{x})$ in $\mathfrak{A}$ such that $\phi(\bar{x})$ is complete in* $\mathrm{Th}(\mathfrak{A})$.

Intuitively, an atomic model is one in which every n-tuple of elements satisfies some $\mathcal{L}$-formula from which all of its $\mathcal{L}$-definable properties can be deduced.

Corollary 2 *The model $\mathfrak{G}$ is atomic.*

Proof: By lemma 3.2, every affine complete formula in $\mathfrak{G}$ is a complete formula in $\mathrm{Th}(\mathfrak{G})$. □

Atomicity is an important property of models, because atomic models count as being 'small'. If $\mathfrak{A}$ and $\mathfrak{B}$ are two $\mathcal{L}$-structures, we say that $\mathfrak{A}$ *can be elementarily embedded in* $\mathfrak{B}$ if there is a function $f : A \to B$ such that for all n-tuples $\bar{a}$ in A, and all $\mathcal{L}$-formulas,

$$\mathfrak{A} \models \phi[\bar{a}] \text{ if and only if } \mathfrak{B} \models \phi[f(\bar{a})] \ .$$

Intuitively, $\mathfrak{A}$ is a copy of a substructure of $\mathfrak{B}$ where any additional elements of $\mathfrak{B}$ make no difference to the $\mathcal{L}$-definable properties of the elements of this substructure. If $\mathfrak{A}$ and $\mathfrak{B}$ make the same *sentences* true, then we say that $\mathfrak{A}$ and $\mathfrak{B}$ are *elementarily equivalent* and write $\mathfrak{A} \equiv \mathfrak{B}$. Note that if $\mathfrak{A}$ can be elementarily embedded in $\mathfrak{B}$, then $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent; however, the converse entailment doe not in general hold. Thus the following result indicates the special status of $\mathfrak{G}$.

Corollary 3 *The model $\mathfrak{G}$ is* prime*: if $\mathfrak{A} \equiv \mathfrak{G}$ then $\mathfrak{G}$ can be elementarily embedded in $\mathfrak{A}$.*

Proof: It is a standard result of model theory (Chang and Keisler [8]: 2.3.4) that a model is prime if it is countable and atomic. □

In other words, any alternative spatial ontology making the same $\mathcal{L}$-sentences true as $\mathfrak{G}$ must contain a copy of $\mathfrak{G}$, together with a collection of elements making no difference to the properties satisfied by members of the copy of $\mathfrak{G}$. In that sense, no alternative ontology for $\mathrm{Th}(\mathfrak{G})$ is simpler than $\mathfrak{G}$ itself.

The following result also comes for free with the above analysis. An *automorphism* of an $\mathcal{L}$-structure $\mathfrak{A}$ is simply a 1–1 map from A to itself which preserves the interpretation of the primitive symbols of $\mathcal{L}$. It is often useful to know what the automorphisms of a structure are.

Corollary 4 *The group of automorphisms of $\mathfrak{G}$ is isomorphic to the rational affine group. The group of automorphisms of $\mathfrak{H}$ is isomorphic to the affine group.*

Proof: We show by a modification of the proof of theorem 4.1 that, if $f : \mathfrak{F} \to \mathfrak{F}$ is an automorphism, then there exists an unique (rational) affine transformation τ such that $f = \tilde{\tau}$; the corollary then follows from lemma 3.2, since we can take the mapping $\tau \mapsto \tilde{\tau}$ to be the isomorphism.

Consider the case $\mathfrak{F} = \mathfrak{G}$. Let f be an automorphism of $\mathfrak{G}$ and let l_1, l_2, l_3 satisfy $\Delta_{\text{bound}}(u_1, u_2, u_3)$. Then $f(l_1), f(l_2), f(l_3)$ satisfies $\Delta_{\text{bound}}(u_1, u_2, u_3)$.

Let τ be the unique rational affine transformation such that $\tilde{\tau}(l_j) = f(l_j)$ $(1 \leq j \leq 3)$. Let m be any half-plane in G. Since both f and $\tilde{\tau}$ are automorphisms, we have

$$\mathfrak{G} \models \phi_{l_1,l_2,l_3,m}[f(l_1), f(l_2), f(l_3), f(m)]$$
$$\mathfrak{G} \models \phi_{l_1,l_2,l_3,m}[f(l_1), f(l_2), f(l_3), \tilde{\tau}(m)] \ .$$

where $\phi_{l_1,l_2,l_3,m}$ is a fixing-formula. But then, given our above observations on fixing formulae, $f(m) = \tilde{\tau}(m)$. Hence $f = \tilde{\tau}$.

The case $\mathfrak{F} = \mathfrak{H}$ is analogous, except that we use a fixing set of formulas $\Phi_{l_1,l_2,l_3,m}$ in place of the single fixing formula $\phi_{l_1,l_2,l_3,m}$. □

The following theorem shows that $\mathcal{L}$ distinguishes between the models $\mathfrak{G}$ and $\mathfrak{H}$.

Theorem 6.1 *The $\mathcal{L}$-structures $\mathfrak{G}$ and $\mathfrak{H}$ are not elementarily equivalent.*

Proof: If the lines k, l, m form a triangle, let us denote the intersection points $m \cap l$ and $k \cap l$ by O and I, respectively. (See figure 4.) If j_b is any line intersecting l in a point B, it is easy to see that B lies on the same side of O as I if both k and j_b lie entirely outside the same quadrant created by l and m (marked with $\star$ in the diagram). Hence there is a formula $\text{pos}(u_b, u_k, u_l, u_m)$ satisfied by the lines j_b, k, l and m guaranteeing that k, l and m form a triangle with the points $l \cap j_b$ and $l \cap k$ lying on the same side of the point $l \cap m$ on l.

In figure 4, the lines j_a, j_b and j_c are shown intersecting l at the points A, B and C, respectively, and satisfy the conditions: $j_a || k$, $j_c \cap m = k \cap m = \{X\}$, $j_b || j_c$ and $j_b \cap m = j_a \cap m = \{Y\}$. By easy plane geometry, triangles OXI and OYA are similar, as indeed are triangles OXC and OYB. Hence:

$$\begin{aligned} \overline{OY}/\overline{OX} &= \overline{OA}/\overline{OI} \\ \overline{OY}/\overline{OX} &= \overline{OB}/\overline{OC}. \end{aligned}$$

Therefore,

$$\overline{OA}/\overline{OI} = \overline{OB}/\overline{OC}.$$

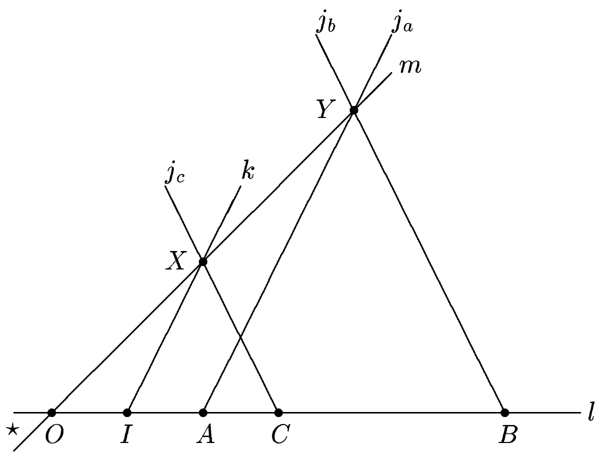

Figure 4: Multiplication of points on a line

whence

$$\frac{\overline{OA}}{\overline{OI}} \cdot \frac{\overline{OC}}{\overline{OI}} = \frac{\overline{OB}}{\overline{OI}}. \tag{2}$$

Intuitively, equation (2) says that, if A, B and C are taken to denote numbers on the line l, where O is the origin and $\overline{OI}$ the unit of measurement, then $B = A.C$. Since the conditions on j_a, j_b and j_c are certainly expressible in $\mathcal{L}$, there is a formula $\text{mult}(u_a, u_b, u_c, u_k, u_l, u_m)$ satisfied by j_a, j_b, j_c, k, l, m just in case k, l and m form a triangle, and j_a, j_b, j_c intersect l in points A, B, C satisfying equation 2.

Now let ϕ be the sentence

$$\forall u_b \forall u_k \forall u_l \forall u_m (\text{pos}(u_b, u_k, u_l, u_m) \rightarrow \exists u \text{mult}(u, u_b, u, u_k, u_l, u_m)).$$

Then it is obvious that $\mathfrak{H} \models \phi$. For let j_b, k, l, m satisfy $\text{pos}(u_b, u_k, u_l, u_m)$ in $\mathfrak{H}$, and let the points O, I, B be as in figure 4. Then we can find a suitable witness for u in ϕ by taking a line j intersecting l at a point A where $\overline{OA}/\overline{OI} = \sqrt{\overline{OB}/\overline{OI}}$. That $\mathfrak{G} \not\models \phi$ is equally obvious if we choose j_b intersecting l at a point B such that $\overline{OB}/\overline{OI}$ has an irrational square root. □

The following result is not surprising.

Theorem 6.2 *The model $\mathfrak{H}$ is not atomic.*

Proof: If $h \in H$, then, by theorem 4.2, let the set of formulas $\Phi_h(x)$ fix h upto affine equivalence. Suppose $\mathfrak{H}$ is atomic, and let $\phi_h(x)$ be a complete formula satisfied by x in $\mathfrak{H}$. Then $\mathrm{Th}(\mathfrak{H}) \models \phi_h \rightarrow \phi$ for every $\phi \in \Phi_h$, so ϕ_h fixes h upto affine equivalence. But this is a contradiction given that $\mathcal{L}$ is countable and that there are certainly uncountably many non-affine-equivalent elements of H. □

So far, we have focused exclusively on the polygonal models $\mathfrak{G}$ and $\mathfrak{H}$; and it is reasonable to insist that one's spatial ontology should be richer than this. For example, one could take as the set of regions the set $\mathrm{RO}(\mathbb{R}^2)$ of *all* regular open sets (i.e. not just polygons). Denote the resulting $\mathcal{L}$-structure by $\mathfrak{A}_{\mathrm{RO}}$. Then we have

Theorem 6.3 *$\mathfrak{A}_{\mathrm{RO}}$ is elementarily equivalent to neither $\mathfrak{G}$ nor $\mathfrak{H}$.*

Proof: Consider the sentence ϕ:

$$\exists x(x > 0 \wedge \forall y \neg(\mathrm{conv}(y) \wedge \mathrm{conv}(x.y) \wedge \mathrm{conv}(-x.y)))$$

(Recall that the empty set qualifies as being nonconvex.) That $\mathfrak{A}_{\mathrm{RO}} \models \phi$ is obvious by taking a circle as a witness for x. That $\mathfrak{F} \not\models \phi$ is equally obvious given that every nonempty element of F must have straight-line edges. □

Consider now the set of regular open *definable* sets, $\mathrm{ROD}(\mathbb{R}^2)$, where a subset of $\mathbb{R}^2$ is said to be definable just in case it is the set of values (ξ_1, ξ_2) which satisfy some formula $\psi(x, y)$ in the usual language of fields. Definable sets are regarded as generally well-behaved. So we might consider the $\mathcal{L}$-interpretation $\mathfrak{A}_{\mathrm{ROD}}$ where quantification rages over all regions in ROD. The proof of theorem 6.3 shows that $\mathfrak{A}_{\mathrm{ROD}}$ is elementarily equivalent to neither $\mathfrak{G}$ nor $\mathfrak{H}$, since open circular discs are certainly in ROD. However, we also have

Theorem 6.4 *$\mathfrak{A}_{\mathrm{RO}}$ and $\mathfrak{A}_{\mathrm{ROD}}$ are not elementarily equivalent.*

Proof: This proof uses results from Pratt and Lemon [24] on the first-order language whose primitives are $\leq$ and $\mathrm{con}(x)$, interpreted in the normal way. It follows from standard results on definable sets ([24] lemma 6.8) that the set ROD is closed under components in the sense of definition 5.1. Therefore, the formula given in lemma 5.2 defines the property of being connected in both $\mathfrak{A}_{\mathrm{RO}}$ and $\mathfrak{A}_{\mathrm{ROD}}$.

Thus, we can regard the following sentence as an $\mathcal{L}$-sentence, by replacing $\mathrm{con}(x)$ by its definition.

$$\forall x_1 \forall x_2 \forall x_3 (\Big(\bigwedge_{1 \leq i \leq 3} \mathrm{con}(x_i) \wedge \mathrm{con}(x_1 + x_2 + x_3)\Big) \rightarrow (\mathrm{con}(x_1 + x_2) \vee \mathrm{con}(x_1 + c_3))).$$

This sentence states that, if the sum of three connected regions is connected, then the first must be connected to at least one of the other two. It is observed in [24] that this sentence is true if quantification is restricted to either G, H or $\mathrm{ROD}(\mathbb{R}^2)$, but false if it is allowed to range over the whole of $\mathrm{RO}(\mathbb{R}^2)$. □

The main conclusion of this section is that $\mathcal{L}$ is able to distinguish between all the following choices for A in terms of the $\mathcal{L}$-sentences that these choices make true: G, H, $\mathrm{ROD}(\mathbb{R}^2)$ and $\mathrm{RO}(\mathbb{R}^2)$. We note in passing that, like $\mathrm{Th}(\mathfrak{G})$ and $\mathrm{Th}(\mathfrak{H})$, both $\mathrm{Th}(\mathfrak{A}_{\mathrm{ROD}})$ and $\mathrm{Th}(\mathfrak{A}_{\mathrm{RO}})$ are undecidable. But the details are routine and not worth rehearsing here.

7 Discussion and related work

Various logicians have sought to give deductive theories of space and space-time (Basri [3], Carnap [7], Goldblatt [15], Henkin, Suppes and Tarski [17]), many in terms of modal logics (Balbiani *et al.* [2], Rescher and Garson [26], Rescher and Urquhart [27], Segerberg [29], Shehtman [31], von Wright [33]). Most recent work on qualitative spatial reasoning has focussed on *mereotopological* languages, that is, those whose primitives are limited to mereological (part-whole) and topological relations (Whitehead [34], Clarke [9], [10], Biacino and Gerla [5], Gotts, Gooday and Cohn [16], Asher and Vieu [1], Borgo, Guarino, and Masolo [6], Roeper [28], Pratt and Lemon [24]).

Whether qualitative spatial representation languages have any practical uses is still unclear. For purely *topological* languages, which employ such primitives as $\mathrm{con}(x)$ and $C(x, y)$, the situation is at least somewhat promising. On the one hand, it is known that many of these languages are surprisingly expressive. Thus, results analogous to theorem 4.1 are available for these languages, but with homeomorphisms taking the place of affine transformations (see Papadimitriou, Suciu and Vianu [21], and, for a different approach Pratt and Schoop [23]). On the other hand, these purely topological languages do not permit the re-introduction of coordinate descriptions. In that respect, they manage to perform the semantic balancing act mentioned at the start of this paper: not so inexpressive as to be useless, yet inexpressive enough to remain distinctively qualitative.

The results established in this paper suggest that, as we move away from the a purely topological vocabulary, this balancing act becomes impossibly difficult. They might be summarized in the informal equation:

$$\text{Mereology} + \text{Convexity} = \text{Affine Geometry}.$$

Once we can represent the property of convexity and the part-whole relation—modest assumptions by any standards—n-tuples of real polygons are determined upto affine equivalence by the sets of formulas they satisfy, and n-tuples of rational polygons are so determined by a single formula. Such a language cannot be called *qualitative* in any meaningful sense. Complete characterizations of polygonal planar arrangements in terms of the properties expressible in $\mathcal{L}$ just *are* coordinate descriptions modulo the fixing of three noncollinear reference points.

Where, one is tempted to ask, have all the numbers come from? Answer: from the fact that affine geometries can be coordinatized by means of constructions describable in terms of the intersection and parallelism of lines, and hence in terms of the expressive resources of $\mathcal{L}$. Specifically, it transpires that any system of points and lines for which the usual axioms of affine geometry hold, together with the theorems of Desargues and Pappus, admits of a coordinatization over a suitable field. (See, e.g. Bennett [4].) And any spatial representation language which enables us to express basic coincidence relations involving points and lines will give us considerable access to this coordinatization. Many of the proofs developed in this paper are essentially a reconstruction of this coordinatization in the special cases of the affine geometries based on the real and rational planes. It is an interesting—and, as far as we are aware, unanswered—question to what extent the properties of the coordinatizing field can be fixed using first-order formulas ranging over various sets of regions in various affine geometries. However, aside from this technical issue, the prospects for the practical application of such languages appear bleak. Bearing in mind the robustness of the phenomenon of coordinatization, it is hard to see how any spatial description language, in which a small repertoire of apparently qualitative spatial primitives is combined with the full power of first-order logic, could possibly avoid the fate which befell $\mathcal{L}$. It is all the fault of logic and geometry: a first-order spatial description language must either lack the ability to express the part-whole relationship and the property of convexity, or must permit the re-introduction of numerical coordinates upto affine transformations. True, this observation does not eliminate the possibility that the affine primitives of $\mathcal{L}$, or any similar set, may form a convenient language in which to formulate simple problems, where the things we want to say often just happen to be expressible with simple, easy-to-reason-with formulas. But we see no grounds for optimism in this regard.

References

[1] Nicholas Asher and Laure Vieu. Toward a Geometry of Common Sense: a semantics and a complete axiomatization of mereotopology. In *International Joint Conference on Artificial Intelligence (IJCAI '95)*. AAAI Press, 1995.

[2] Philippe Balbiani, Luis Farinas del Cerro, Tinko Tinchev, and Dimiter Vakarelov. Geometrical Structures and Modal Logic. In Gabbay and Ohlbach, editors, *Practical Reasoning*, Lecture notes in Artificial Intelligence 1085, pages 43 –57. Springer-Verlag, Berlin, 1996. Formal Aspects of Practical Reasoning '96.

[3] Saul A. Basri. *A Deductive Theory of Space and Time.* Studies in Logic and the Foundations of Mathematics. North-Holland Publishing Co., Amsterdam, 1966.

[4] Brandon Bennett. Modal Logics for Qualitative Spatial Reasoning. *Bulletin of the IGPL*, 3, 1995.

[5] Loredana Biacino and Giangiacomo Gerla. Connection Structures. *Notre Dame Journal of Formal Logic*, 32(2):242 – 247, 1991.

[6] Stefano Borgo, Nicola Guarino, and Claudio Masolo. A Pointless Theory of Space Based on Strong Connection and Congruence. In L. C. Aiello, J. Doyle, and S. C. Shapiro, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96)*, pages 220 –229, San Francisco, CA., 1996. Morgan Kaufmann Publishers.

[7] Rudolf Carnap. *Introduction to Symbolic Logic and its Appplications.* Dover Publications Inc., New York, 1958.

[8] C. C. Chang and H. J. Keisler. *Model Theory.* North Holland, Amsterdam, 3rd edition, 1990.

[9] B. L. Clarke. A calculus of individuals based on 'connection'. *Notre Dame Journal of Formal Logic*, 23(3):204 – 218, 1981.

[10] B. L. Clarke. Individuals and Points. *Notre Dame Journal of Formal Logic*, 26(1):61 – 75, 1985.

[11] Ernest Davis. *Representations of Commonsense Knowledge.* Morgan Kaufmann, San Mateo, CA, 1990.

[12] Ernest Davis, Nicholas Mark Gotts, and Anthony G. Cohn. Constraint networks of topological relations and convexity. *Submitted*, 1999.

[13] Christopoh Dornheim. Undecidability of plane polygonal mereotopology. In *Proceedings of the 6th International Conference on Knowledge Representation*, pages 342–353, 1998.

[14] Gregory Dudek, Paul Freedman, and Souad Hadjres. Using local information in a non-local way for mapping graph-like worlds. In *13th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1639 – 1645. Morgan Kaufmann, 1993.

[15] Robert Goldblatt. *Mathematics of Modality.* CSLI lecture notes 43. CSLI, Stanford, 1993.

[16] Nicholas Gotts, John Gooday, and Anthony Cohn. A connection based approach to commonsense topological description and reasoning. *Monist*, 79(1):51–75, 1996.

[17] L. Henkin, P. Suppes, and A. Tarski, editors. *The Axiomatic Method, with Special Reference to Geometry and Physics.* North-Holland Publishing Co., Amsterdam, 1959.

[18] Sabine Koppelberg. *Handbook of Boolean algebras*, volume 1. North-Holland, 1989.

[19] B. Kuipers. Modeling Spatial Knowledge. *Cognitive Science*, 2(2):129 –153, 1978.

[20] Peter M. Neumann, Gabrielle A. Stoy, and Edward C. Thompson. *Groups and geometry.* Oxford University Press, Oxford, 1994.

[21] C. H. Papadimitriou, D. Suciu, and V. Vianu. Topological queries in spatial databases. In *Proceedings, PODS'96*, pages 81–92. ACM, 1996.

[22] I. Pratt and D. Schoop. A complete axiom system for polygonal mereotopology of the real plane. *Journal of Philosophical Logic*, 27(6), 1998. forthcoming.

[23] I. Pratt and D. Schoop. Expressivity in polygonal, plane mereotopology. Technical Report UMCS-98-5-1, University of Manchester, Department of Computer Science, United Kingdom, 1998.

[24] Ian Pratt and Oliver Lemon. Ontologies for plane, polygonal mereotopology. *Notre Dame Journal of Formal Logic*, 38(2):225–245, 1997.

[25] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic based on regions and connection. In B. Nebel, C. Rich, and W. Swartout, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR '92)*, pages 165–176, Los Altos, CA, 1992. Morgan Kaufmann Publishers.

[26] Nicholas Rescher and James Garson. Topological Logic. *Journal of Symbolic Logic*, 33:537–548, 1968.

[27] Nicholas Rescher and Alasdair Urquhart. *Temporal Logic*, volume 3 of *Library of Exact Philosophy.* Springer-Verlag, New York, 1971.

[28] Peter Roeper. Region-based topology. *Journal of Philosophical Logic*, 26:251–309, 1997.

[29] Krister Segerberg. Two dimensional modal logic. *Journal of Philosophical Logic*, 2:77–96, 1973.

[30] Murray Shanahan. Default Reasoning about Spatial Occupancy. *Artificial Intelligence*, 74:147 – 163, 1995.

[31] Valentin B. Shehtman. Modal Logics of Domains on the Real Plane. *Studia Logica*, 42:63–80, 1983.

[32] Oliviero Stock. *Spatial and Temporal Reasoning*. Kluwer, Dordrecht, 1997.

[33] Georg Henrik von Wright. A Modal Logic of Place. In E. Sosa, editor, *The Philosophy of Nicholas Rescher: Discussion and Replies*, pages 65–73. D. Reidel, Dordrecht, Holland, 1979.

[34] Alfred North Whitehead. *Process and Reality*. The MacMillan Company, New York, 1929.