

A Logical Framework for Depiction and Image Interpretation

Raymond Reiter*

*Department of Computer Science, University of Toronto,
Toronto, Ontario, Canada M5S 1A4*

Alan K. Mackworth*

*Department of Computer Science, University of British
Columbia, Vancouver, BC, Canada V6T 1W5*

ABSTRACT

We propose a logical framework for depiction and interpretation that formalizes image domain knowledge, scene domain knowledge and the depiction mapping between the image and scene domains. This framework requires three sets of axioms: image axioms, scene axioms and depiction axioms. An interpretation of an image is defined to be a logical model of these axioms.

The approach is illustrated by a case study, a reconstruction in first-order logic of a simplified map-understanding program, Mapsee. The reconstruction starts with a description of the map and a specification of general knowledge of maps, geographic objects and their depiction relationships. For the simple map world we show how the task level specification may be refined to a provably correct implementation by applying model-preserving transformations to the initial logical representation to produce a set of propositional formulas. The implementation may use known constraint satisfaction techniques to find the set of models of these propositional formulas. In addition, we sketch preliminary logical treatments for image queries, contingent scene knowledge, ambiguity in image description, occlusion, complex objects, preferred interpretations and image synthesis.

This approach provides a formal framework for analyzing and going beyond existing systems such as Mapsee, and for understanding the use of constraint satisfaction techniques. It can be used as a foundation for the specification, design and implementation of vision and graphics systems that are correct with respect to the task and algorithm levels.

1. Introduction

Computational vision requires, no less than any other area of artificial intelligence, representations of knowledge that are complete, correct, flexible and efficient. In pursuit of that goal researchers have exploited a wide variety of knowledge representation schemes including grammars, semantic nets, programs, logics, schemas, rules, constraints and neural nets. McCarthy and

* Fellow, Canadian Institute for Advanced Research

Hayes [16] proposed some adequacy criteria for knowledge representation schemes in general and used them to argue for a logical representation. Vision researchers have, by and large, ignored that suggestion. Clowes [5] and Huffman [8], for example, advocated a knowledge representation based on simple constraints in the scene domain, in a nonlogical framework. Mackworth [14] argued that any adequate representation scheme for visual knowledge should satisfy various criteria of descriptive and procedural adequacy. Here we can only briefly refer to some of them.

The relevant criteria of descriptive adequacy are *capacity*, *primitives*, *composition*, *specialization*, *subworlds*, *depiction* and *correctness*; the relevant criteria of procedural adequacy are *soundness*, *completeness*, *flexibility* and *efficiency*. If the generative power of the scheme is adequate then the *capacity* criterion is satisfied. A generative scheme must be based on *primitives* and it must provide rules for the *composition* of structured objects, whose descriptions can be refined through *specialization*. To satisfy the *subworlds* criterion the representation scheme must, minimally, maintain the distinction between knowledge of the image and knowledge of the scene; otherwise, elementary category errors, such as confusing a real scene edge with its depiction in the image, are bound to be made. Moreover, the scheme must carry information about the *depiction* relation itself: how objects in the scene domain appear in the image. The task specification must be precise to determine *correctness*. Only if the concept of an image interpretation is precisely defined can we determine if an implementation is *complete* and *sound*; that is, finds all and only the interpretations allowed by the general knowledge and a description of the particular image. A representation scheme achieves some measure of *flexibility* if it can exploit contingent knowledge or support both image interpretation and image generation. *Efficiency* can be evaluated by complexity analyses of the task and proposed algorithms using measures of time, space, number of processors or communication costs. Some measure of *completeness* or *soundness* may need to be sacrificed to *efficiency* through the use of approximation algorithms. This paper provides an adequate logical framework for depiction and interpretation, and demonstrates its application in a simple world.

Informally, to motivate the sceptical reader who asks "Why should I care?" we can only say that, to our knowledge, this is the first paper to provide a precise definition of the concept of an interpretation of an image. Furthermore, the point of much of the resulting logical manipulation is to show how the nonprocedural specification reduces to a constraint satisfaction problem (CSP). This is important for three reasons. First, there are well-understood algorithms for solving CSPs. Second, the CSP is logically equivalent to the original specification, so we have a correctness proof. Third, the transformation from specification to CSP explains and justifies the central role that CSPs play in model-based vision.

2. An Illustrative Specification: Mapsee's Sketch Maps

As an example of how one might logically specify the knowledge base for an image interpretation application, we focus on Mapsee, a long term research project at the University of British Columbia designed to interpret hand drawn sketch maps of geographical regions.

The Mapsee project is a series of experiments in visual knowledge representation [18]. Mapsee-1 (Mackworth [10]) used n -ary constraints in the scene domain and a network consistency constraint satisfaction algorithm. Mapsee-2 (Havens and Mackworth [7]) adopted a schema representation of knowledge that was simplified and enhanced in Mapsee-3 (Mulder [19]) and augmented with a hierarchical constraint satisfaction algorithm (Mackworth et al. [12]). These systems served as testbeds for new knowledge representation techniques and as useful artifacts in their own right, for example, acting as knowledge sources for the interpretation of satellite and aerial imagery, and as prototypes for more autonomous image understanding systems. However, since no precise definition of the notion of an interpretation has been provided and since much of the knowledge is procedurally encoded and distributed, it is not possible to determine if these programs are functioning correctly according to a formal specification of the task. One purpose of this paper is to provide a "logical reconstruction" of a fragment of the Mapsee project.

For expository purposes, we considerably simplify the kinds of image and geographic features which Mapsee deals with, as well as the kinds of knowledge it uses in image interpretation. As a further *caveat*, we emphasize that the following specification is appropriate for the world of sketch maps; other applications may require very different axioms and assumptions. Minimally, the principal feature which remains applicable is the definition below of just what one means by an interpretation of an image (namely, a model of the axioms). Having a formal definition allows system designers to address in their specifications and implementations the issues of descriptive and procedural adequacy introduced in Section 1.

The user of this proposed simple Mapsee would sketch the map input using a mouse or data tablet. The initial map description is a set of chains, where a chain is a list of consecutively drawn points connected by a line segment. From this description the system constructs an enriched description that explicitly determines topologically connected spatial regions and the various relations of the chains and regions to be described in Section 2.1. Initially, we assume a carefully drawn map (with no gaps at the intended chain junctions, for example) which ensures a unique image description but we show, in Section 7.3, that this assumption can be relaxed.

The task for this Mapsee is to compute the set of interpretations of the map as depicting a simple scene of roads, rivers, shorelines and areas of land and

water under various assumptions about what is permitted in the scene, to be described in Section 2.2, and how scene objects are depicted in images, to be described in Section 2.3

2.1. Specifying the image domain

We assume that there are just two kinds of image primitives—chains and regions—so that the taxonomy of image objects is given by Fig. 1, which pictorially represents the following first-order formulas¹.

$$(\forall x) \text{image-object}(x) \equiv \text{chain}(x) \vee \text{region}(x),$$

$$(\forall x) \neg(\text{chain}(x) \wedge \text{region}(x))$$

In addition there are the following relationships which may hold between image primitives:

- tee*(*c*,*c'*): chain *c* meets chain *c'* at a T-junction, as in Fig. 2(a)
- chi*(*c*,*c'*): chains *c* meets chain *c'* at a χ -junction, as in Fig. 2(b).
- bounds*(*c*,*r*): chain *c* bounds region *r*, as in Fig. 2(c).
- closed*(*c*): chain *c* is a simple closed figure, as in Fig. 2(d).
- interior*(*c*,*r*): an interior of closed chain *c* is region *r*, as in Fig. 2(e).
- exterior*(*c*,*r*): an exterior of closed chain *c* is region *r*, as in Fig. 2(f).

A given image will consist of finitely many chains and regions, together with finitely many instances of the above relations. Mapsee makes the following:

Closure Assumption (*Closed World Assumption* (Reiter [21]) *for the image domain*). All image domain predicates are completely known. This closure assumption is logically specified by *closure axioms* of the form:




Fig. 1 An image domain taxonomic hierarchy

¹ We denote image domain predicates using lower case characters and scene domain predicates in upper case

	Relation	Example
(a)	$tee(c,c')$	
(b)	$chi(c,c')$	
(c)	$bounds(c,r)$	
	$bounds(c,r)$	
(d)	$closed(c)$	
(e)	$interior(c,r)$	
(f)	$exterior(c,r)$	
	$exterior(c,r)$	

Fig 2 Relations in the image domain

$$(\forall x) chain(x) \equiv x = i_1 \vee \dots \vee x = i_m,$$

$$(\forall x) region(x) \equiv x = i_1 \vee \dots \vee x = i_n,$$

$$(\forall x, y) tee(x, y) \equiv (x = i_1 \wedge y = i'_1) \vee \dots \vee (x = i_k \wedge y = i'_k),$$

$$(\forall x, y) bounds(x, y) \equiv (x = i_1 \wedge y = i'_1) \vee \dots \vee (x = i_l \wedge y = i'_l),$$

etc.,

where the i and i' are all constants.

Example 2.1. Figure 3 shows a simple hand drawn sketch map with its chains and regions labeled by suitable constants. The closure axioms for this image




Fig 3 A sketch map

are:

$$(\forall x) \text{ chain}(x) \equiv x = c_1 \vee x = c_2 \vee x = c_3 \\ \vee x = c_4 \vee x = c_5 \vee x = c_6,$$

$$(\forall x) \text{ region}(x) \equiv x = r_1 \vee x = r_2 \vee x = r_3 \vee x = r_4,$$

$$(\forall x, y) \text{ tee}(x, y) \equiv (x = c_2 \wedge y = c_1) \vee (x = c_2 \wedge y = c_3) \\ \vee (x = c_4 \wedge y = c_5) \vee (x = c_3 \wedge y = c_5),$$

$$(\forall x, y) \text{ chu}(x, y) \equiv (x = c_3 \wedge y = c_4) \vee (x = c_4 \wedge y = c_3),$$

$$(\forall x, y) \text{ bounds}(x, y) \equiv (x = c_1 \wedge y = r_1) \vee (x = c_2 \wedge y = r_1) \\ \vee (x = c_3 \wedge y = r_1) \vee (x = c_3 \wedge y = r_2) \\ \vee (x = c_4 \wedge y = r_1) \vee (x = c_4 \wedge y = r_2) \\ \vee (x = c_5 \wedge y = r_1) \vee (x = c_5 \wedge y = r_2) \\ \vee (x = c_5 \wedge y = r_3) \vee (x = c_6 \wedge y = r_3) \\ \vee (x = c_6 \wedge y = r_4),$$

$$(\forall x) \text{ closed}(x) \equiv x = c_5 \vee x = c_6,$$

$$(\forall x, y) \text{ interior}(x, y) \equiv (x = c_6 \wedge y = r_4) \vee (x = c_5 \wedge y = r_3),$$

$$(\forall x, y) \text{ exterior}(x, y) \equiv (x = c_5 \wedge y = r_1) \\ \vee (x = c_5 \wedge y = r_2) \vee (x = c_6 \wedge y = r_3)$$

In addition to the closed world assumption for the image domain, Mapsee also makes the following assumption.

Unique Names Assumption (Reiter [22]). All image primitives (i.e. the chains and regions) are pairwise distinct. In other words if i and i' are different constants denoting image primitives, they denote different image primitives. Thus, the specification of the image domain includes the following *unique names axioms*:

$i \neq i'$ for all distinct constants i and i' mentioned in the closure axioms for *chain* and *region*.

Notice that we have been implicitly assuming suitable *type constraints* on the arguments of image predicates, e.g. that the first argument of *bounds* is a chain, and the second a region. We also want no constant mentioned in the closure axiom for *chain* to be mentioned in the closure axiom for *region*; otherwise, for any such constant i , both $chain(i)$ and $region(i)$ would hold, contradicting the taxonomic axiom $(\forall x) \neg(chain(x) \wedge region(x))$.

We make these two assumptions explicit by imposing the following simple requirements on the above closure axioms:

Coherence Requirements.

(C1) Each constant occurring in the closure axiom for *chain* is distinct from any occurring in the closure axiom for *region*.

(C2) All constants mentioned in the closure axioms for *tee* are mentioned in the closure axiom for *chain*. We impose this by the following image type constraint:

$$(\forall x, y) tee(x, y) \supset chain(x) \wedge chain(y) .$$

Similarly,

$$(\forall x, y) bounds(x, y) \supset chain(x) \wedge region(y) .$$

Similar axioms hold for the image predicates *chi*, *closed*, *interior* and *exterior*.

2.2. Specifying the scene domain

We assume that the taxonomy of scene objects is given by Fig. 4, which pictorially represents the following first-order formulas²:

$$\begin{aligned} (\forall x) SCENE-OBJECT(x) &\equiv LINEAR-SCENE-OBJECT(x) \vee AREA(x) , \\ (\forall x) \neg(LINEAR-SCENE-OBJECT(x) \wedge AREA(x)) , \end{aligned}$$

²Recall our convention that scene domain predicates are denoted by upper case characters, and image domain predicates by lower case




Fig 4 A scene domain taxonomic hierarchy

$$\begin{aligned}
 (\forall x) \text{ LINEAR-SCENE-OBJECT}(x) &\equiv \text{ROAD}(x) \vee \text{RIVER}(x) \vee \text{SHORE}(x) , \\
 (\forall x) \neg(\text{ROAD}(x) \wedge \text{RIVER}(x)) , \\
 (\forall x) \neg(\text{ROAD}(x) \wedge \text{SHORE}(x)) , \\
 (\forall x) \neg(\text{RIVER}(x) \wedge \text{SHORE}(x)) , \\
 (\forall x) \text{ AREA}(x) &\equiv \text{LAND}(x) \vee \text{WATER}(x) , \\
 (\forall x) \neg(\text{LAND}(x) \wedge \text{WATER}(x)) .
 \end{aligned}$$

In addition to this taxonomic information, we assume the following general facts about the real world of roads, rivers, shorelines, land and water

- (i) Rivers do not cross each other:

$$(\forall x, y) \text{ RIVER}(x) \wedge \text{RIVER}(y) \supset \neg \text{CROSS}(x, y)$$

- (ii) Shorelines form closed loops:

$$(\forall x) \text{ SHORE}(x) \supset \text{LOOP}(x)$$

- (iii) Rivers cannot form loops:

$$(\forall x) \text{ RIVER}(x) \supset \neg \text{LOOP}(x) .$$

(iv) The inside area of a shoreline is land iff its outside is water; its inside is water iff its outside is land:

$$\begin{aligned}
 (\forall x, y, z) \text{ SHORE}(x) \wedge \text{INSIDE}(x, y) \wedge \text{OUTSIDE}(x, z) \\
 \supset (\text{LAND}(y) \equiv \text{WATER}(z)) \wedge (\text{WATER}(y) \equiv \text{LAND}(z)) .
 \end{aligned}$$

(v) If a road or a river is beside an area then that area is land:

$$(\forall x, y) \text{ BESIDE}(x, y) \wedge (\text{ROAD}(x) \vee \text{RIVER}(x)) \supset \text{LAND}(y) .$$

(vi) Rivers flow into other rivers, or into shores:

$$\begin{aligned} (\forall x) \text{RIVER}(x) \supset (\exists y) \text{RIVER}(y) \wedge \text{JOINS}(x, y) \\ \vee (\exists z) \text{SHORE}(z) \wedge \text{JOINS}(x, z) . \end{aligned}$$

Finally, we require the following axioms which restrict the scene predicates to scene objects only:

Scene Predicate Type Constraint Axioms.

$$\begin{aligned} (\forall x, y) \text{CROSS}(x, y) \\ \supset \text{SCENE-OBJECT}(x) \wedge \text{SCENE-OBJECT}(y) , \\ (\forall x) \text{LOOP}(x) \supset \text{SCENE-OBJECT}(x) , \\ (\forall x, y) \text{INSIDE}(x, y) \\ \supset \text{SCENE-OBJECT}(x) \wedge \text{SCENE-OBJECT}(y) , \\ (\forall x, y) \text{OUTSIDE}(x, y) \\ \supset \text{SCENE-OBJECT}(x) \wedge \text{SCENE-OBJECT}(y) , \\ (\forall x, y) \text{BESIDE}(x, y) \\ \supset \text{SCENE-OBJECT}(x) \wedge \text{SCENE-OBJECT}(y) , \\ (\forall x, y) \text{JOINS}(x, y) \\ \supset \text{SCENE-OBJECT}(x) \wedge \text{SCENE-OBJECT}(y) . \end{aligned}$$

2.3. Specifying the image-scene domain mappings

In any given application, there will be relations which hold between the image and scene domains, for example, relations specifying how various three-dimensional objects project onto the two-dimensional image plane, or what kinds of scene objects are depicted by image objects. We refer to such relations as *mappings*, and represent them by a distinguished binary predicate $\Delta(i, s)$ meaning that image object i depicts scene object s .

In the case of Mapsee, the following assumptions are made:

(i) The world consists of image objects and scene objects, and these form a taxonomy:

$$\begin{aligned} (\forall x) \text{image-object}(x) \vee \text{SCENE-OBJECT}(x) , \\ (\forall x) \neg(\text{image-object}(x) \wedge \text{SCENE-OBJECT}(x)) . \end{aligned}$$

(ii) Every image object t depicts a unique scene object which we denote by $\sigma(t)$:

$$(\forall t) \text{image-object}(t) \supset \text{SCENE-OBJECT}(\sigma(t)) \\ \wedge \Delta(t, \sigma(t)) \wedge [(\forall s) \Delta(t, s) \supset s = \sigma(t)].$$

(iii) Every scene object is depicted by a unique image object:

$$(\forall s) \text{SCENE-OBJECT}(s) \supset (\exists! t) \text{image-object}(t) \wedge \Delta(t, s)$$

Assumptions (ii) and (iii) are very strong. For example, (ii) forces the conclusion that a noise patch in the image depicts something real in the scene, while (iii) precludes occluded objects in the scene. Clearly, there are settings where these assumptions are unwarranted, where some of (ii), (iii) and the other image, scene and mapping axioms require more complex representations. We gloss over this issue for now but return to it briefly in Section 7.4 where we sketch a logical treatment of occlusion.

(iv) Depiction holds only between image and scene objects:

$$(\forall t, s) \Delta(t, s) \supset \text{image-object}(t) \wedge \text{SCENE-OBJECT}(s)$$

(v) Taxonomic mappings:

– Regions in the image depict areas in the scene:

$$(\forall t, s) \Delta(t, s) \wedge \text{region}(t) \supset \text{AREA}(s)$$

– Chains in the image depict linear scene objects in the scene:

$$(\forall t, s) \Delta(t, s) \wedge \text{chain}(t) \supset \text{LINEAR-SCENE-OBJECT}(s).$$

(vi) Relational mappings:

– Tee relations in the image depict join relations in the scene, and vice versa:

$$(\forall i_1, t_2, s_1, s_2) \Delta(t_1, s_1) \wedge \Delta(t_2, s_2) \supset \text{tee}(t_1, t_2) \equiv \text{JOINS}(s_1, s_2).$$

– Similarly, for the other image relations (Fig. 2) and their corresponding scene relations:

$$(\forall i_1, t_2, s_1, s_2) \Delta(t_1, s_1) \wedge \Delta(t_2, s_2) \supset \text{chl}(i_1, t_2) \equiv \text{CROSS}(s_1, s_2), \\ (\forall i_1, i_2, s_1, s_2) \Delta(t_1, s_1) \wedge \Delta(t_2, s_2) \\ \supset \text{bounds}(i_1, t_2) \equiv \text{BESIDE}(s_1, s_2),$$

$$\begin{aligned}
& (\forall t, s) \Delta(t, s) \supset \text{closed}(i) \equiv \text{LOOP}(s), \\
& (\forall i_1, i_2, s_1, s_2) \Delta(i_1, s_1) \wedge \Delta(i_2, s_2) \\
& \quad \supset \text{interior}(i_1, i_2) \equiv \text{INSIDE}(s_1, s_2), \\
& (\forall i_1, i_2, s_1, s_2) \Delta(i_1, s_1) \wedge \Delta(i_2, s_2) \\
& \quad \supset \text{exterior}(i_1, i_2) \equiv \text{OUTSIDE}(s_1, s_2).
\end{aligned}$$

3. What Is an Interpretation?

In general, a logical specification of the relevant knowledge and underlying assumptions for an image understanding application will consist of:

- (i) *image axioms*: an axiomatization of the image domain,
- (ii) *scene axioms*: an axiomatization of the scene domain, and
- (iii) *mapping axioms*: an axiomatization of the mappings between the image and scene domains.

Sections 2.1, 2.2 and 2.3 provide an example of this tripartite specification, for the sketch map task.

With such an axiomatization in hand, we can provide a formal definition of an interpretation as follows:

Definition 3.1. An *interpretation* of an image is a model of the image, scene and mapping axioms.

We use the term “model” here in its strict logical sense (Mendelson [17]).³

At this point it is appropriate to say a few words about computational issues. Determining the models of an arbitrary set of first-order axioms is a wildly impractical task. To begin, it is undecidable in general whether such a set of formulas even has a model. Moreover, there may be infinitely many models. Is there anything special about vision which precludes these problems?

At this stage of our research we can only speculate. The most promising observation is that an image is finite. There are just finitely many primitive image objects and relations between objects. Provided the depiction relation allows for just finitely many scene objects corresponding to the image primitives, then all quantifiers will have finite range. As we shall see, this is the case for our sketch map domain. Whenever this is the case, quantified formulas reduce to propositional ones and image interpretations are all computable. It is unclear just how general this observation is. Very likely a variety of vision tasks

³The term “interpretation” has a logical meaning (Mendelson [17]) which differs from our use of the word. Since we are grounding high level vision in logic, there is a risk of terminological confusion. Since “interpretation” is so firmly entrenched in the computational vision literature, we choose to continue use of the term in this paper. We emphasize that its use does not refer to its logical meaning.

must be formalized before some general principles emerge regarding decidability issues.

4. Some Results Derivable from Mapsee's Axiomatization

Let MAP-AXIOMS be those axioms specified in Sections 2.1, 2.2 and 2.3 for our simplified Mapsee domain, namely the image axioms, the scene axioms, and the mapping axioms. In this section, we state various logical consequences of these axioms which will simplify the process of computing the interpretations for a hand drawn sketch map. We omit the proofs, which are contained in [24, Appendix A].

Notation. Whenever MAP-AXIOMS entails a closure formula of the form

$$(\forall x_1) \dots (\forall x_n) P(x_1, \dots, x_n) \\ \equiv (x_1 = t_1^1 \wedge \dots \wedge x_n = t_n^1) \vee \dots \vee (x_1 = t_1^m \wedge \dots \wedge x_n = t_n^m),$$

where the t_i^l are all terms, then $|P|$ denotes $\{(t_1^1, \dots, t_n^1), \dots, (t_1^m, \dots, t_n^m)\}$

Result 4.1 (Closure on image objects).

$$\text{MAP-AXIOMS} \\ \models (\forall x) \text{image-object}(x) \equiv \bigvee_{t \in |\text{chain}| \cup |\text{region}|} (x = t).$$

Result 4.2 (Closure for Δ)

$$\text{MAP-AXIOMS} \\ \models (\forall x, y) \Delta(x, y) \equiv \bigvee_{t \in |\text{image-object}|} (x = t \wedge y = \sigma(t))$$

Result 4.3 (Uniqueness of all objects). *If $I_m, I_n \in |\text{image-object}|$,*

- (1) $\text{MAP-AXIOMS} \models \sigma(I_m) \neq \sigma(I_n)$ *when $m \neq n$*
- (2) $\text{MAP-AXIOMS} \models \sigma(I_m) \neq I_n$,
- (3) $\text{MAP-AXIOMS} \models I_m \neq I_n$ *when $m \neq n$.*

Result 4.4 (Closure for scene objects).

$$\text{MAP-AXIOMS} \\ \models (\forall s) \text{SCENE-OBJECT}(s) \equiv \bigvee_{t \in |\text{image-object}|} (s = \sigma(t)).$$

Result 4.5 (Domain closure).

MAP-AXIOMS

$$\models (\forall x) \left[\bigvee_{i \in |\text{image-object}|} (x = i \vee x = \sigma(i)) \right].$$

Result 4.6 (Closure for linear scene objects and areas).

(1) MAP-AXIOMS

$$\models (\forall s) \text{LINEAR-SCENE-OBJECT}(s) \equiv \bigvee_{i \in |\text{chain}|} (s = \sigma(i)),$$

(2) MAP-AXIOMS

$$\models (\forall s) \text{AREA}(s) \equiv \bigvee_{i \in |\text{region}|} (s = \sigma(i)).$$

Result 4.7 (Closure for scene domain relations).

(1) MAP-AXIOMS

$$\models (\forall x, y) \text{JOINS}(x, y) \equiv \bigvee_{(i, i') \in |\text{tee}|} (x = \sigma(i) \wedge y = \sigma(i')),$$

(2) MAP-AXIOMS

$$\models (\forall x, y) \text{CROSS}(x, y) \equiv \bigvee_{(i, i') \in |\text{chl}|} (x = \sigma(i) \wedge y = \sigma(i')),$$

(3) MAP-AXIOMS

$$\models (\forall x, y) \text{BESIDE}(x, y) \equiv \bigvee_{(i, i') \in |\text{bounds}|} (x = \sigma(i) \wedge y = \sigma(i')),$$

(4) MAP-AXIOMS

$$\models (\forall x) \text{LOOP}(x) \equiv \bigvee_{i \in |\text{closed}|} (x = \sigma(i)),$$

(5) MAP-AXIOMS

$$\models (\forall x, y) \text{INSIDE}(x, y) \equiv \bigvee_{(i, i') \in |\text{interior}|} (x = \sigma(i) \wedge y = \sigma(i')),$$

(6) MAP-AXIOMS

$$\models (\forall x, y) \text{OUTSIDE}(x, y) \equiv \bigvee_{(i, i') \in |\text{exterior}|} (x = \sigma(i) \wedge y = \sigma(i')).$$

5. Simplifying MAP-AXIOMS

We now show how the results of the previous section allow us to systematically eliminate from consideration many of the axioms of MAP-AXIOMS. This in turn

will considerably simplify the task of determining all interpretations of an image, as we shall see in Section 6 below

Let SIMP-AXIOMS consist of the following groups of formulas:

(S1) The closure axioms for *tee*, *chi*, *bounds*, *closed*, *interior*, *exterior*, *chain* and *region* of Section 2.1, augmented by the closure formulas for *image-object*, Δ , SCENE-OBJECT, LINEAR-SCENE-OBJECT, AREA, JOINS, CROSS, BESIDE, LOOP, INSIDE and OUTSIDE, derived in the previous section.

(S2) Unique names formulas of Result 4.3, together with the domain closure formula of Result 4.5

(S3) (i) For $t \in |image-object|$,

$$\begin{aligned} &\neg ROAD(t) , \\ &\neg RIVER(t) , \\ &\neg SHORE(t) , \\ &\neg LAND(t) , \\ &\neg WATER(t) . \end{aligned}$$

(ii) For $s \in |AREA|$,

$$\begin{aligned} &\neg RIVER(s) , \\ &\neg ROAD(s) , \\ &\neg SHORE(s) , \\ &LAND(s) \vee WATER(s) , \\ &\neg LAND(s) \vee \neg WATER(s) . \end{aligned}$$

(iii) For $s \in |LINEAR-SCENE-OBJECT|$,

$$\begin{aligned} &\neg LAND(s) , \\ &\neg WATER(s) . \end{aligned}$$

(iv) For $s \in |LOOP|$,

$$\begin{aligned} &ROAD(s) \vee SHORE(s) , \\ &\neg ROAD(s) \vee \neg SHORE(s) , \\ &\neg RIVER(s) \end{aligned}$$

(v) For $s \in |LINEAR-SCENE-OBJECT| - |LOOP|$,

$$\begin{aligned} &ROAD(s) \vee RIVER(s) . \\ &\neg ROAD(s) \vee \neg RIVER(s) , \\ &\neg SHORE(s) . \end{aligned}$$

(S4) The following groups of formulas:

(i) For $(x, y) \in |\text{CROSS}|$,

$$\neg \text{RIVER}(x) \vee \neg \text{RIVER}(y).$$

(ii) For (x, y, z) such that $x \in |\text{LOOP}|$, $(x, y) \in |\text{INSIDE}|$, $(x, z) \in |\text{OUTSIDE}|$,

$$\text{SHORE}(x) \supset (\text{LAND}(y) \equiv \text{WATER}(z)).$$

(iii) For $(x, y) \in |\text{BESIDE}|$ and $x \notin |\text{LOOP}|$,

$$\text{LAND}(y).$$

For $(x, y) \in |\text{BESIDE}|$ and $x \in |\text{LOOP}|$,

$$\text{ROAD}(x) \supset \text{LAND}(y).$$

(iv) For $x \in |\text{LINEAR-SCENE-OBJECT}| - |\text{LOOP}|$,

$$\begin{aligned} \text{RIVER}(x) \supset & \left[\bigvee_{\{y | (x, y) \in |\text{JOINS}| \text{ and } y \notin |\text{LOOP}|\}} \text{RIVER}(y) \right] \\ & \vee \left[\bigvee_{\{z | (x, z) \in |\text{JOINS}| \text{ and } z \in |\text{LOOP}|\}} \text{SHORE}(z) \right]. \end{aligned}$$

Proposition 5.1. *MAP-AXIOMS and SIMP-AXIOMS are logically equivalent.*

Proof. See [24, Appendix B]. \square

In the next section we show how SIMP-AXIOMS may be used to compute interpretations of sketch maps.

6. Determining the Interpretations of a Map

It remains to compute the interpretations of a hand drawn sketch map, which means, by Definition 3.1, computing all models of MAP-AXIOMS, hence of SIMP-AXIOMS. All such models share the following properties:

(1) Suppose $|\text{image-object}| = \{i_1, \dots, i_n\}$. By the domain closure and unique names formulas of (S2), the universe of any such model consists of $2n$ pairwise unequal elements. If we denote the elements of this universe corresponding to $i_1, \dots, i_n, \sigma(i_1), \dots, \sigma(i_n)$ by themselves, then all models of SIMP-AXIOMS share the same universe $\{i_1, \dots, i_n, \sigma(i_1), \dots, \sigma(i_n)\}$ of pairwise unequal elements.

(2) The closure formulas of (S1) completely characterize their predicates. Accordingly each predicate with a closure axiom has the same extension in all models of SIMP-AXIOMS, and these extensions are known to us a priori. For example, $|BESIDE|$ is the extension of the predicate BESIDE common to all models of SIMP-AXIOMS.

The only predicates lacking closure formulas are ROAD, RIVER, SHORE, LAND and WATER. Thus, the models of SIMP-AXIOMS can differ from one another only in the extensions they assign to these predicates. It follows that the only formulas of SIMP-AXIOMS we need consider in computing these models are those of (S3) and (S4). Moreover, these are quantifier-free formulas, so the problem reduces to determining the set of all propositional models of a set of formulas of the propositional calculus. While this is in general an NP-hard problem, at least it is decidable and various algorithms are known (Bibel [1], Purdom [20], Mackworth [13]).

We illustrate the result of this calculation with the example sketch map of Fig. 3.

Example 6.1 (The map of Fig. 3). All models share the same universe $\{c_1, \dots, c_6, r_1, \dots, r_4, C_1, \dots, C_6, R_1, \dots, R_4\}$ where C_i and R_j denote $\sigma(c_i)$ and $\sigma(r_j)$ respectively. Table 1 summarizes the extensions common to all these models of the predicates with closure formulas in (S1).

It remains to determine all models of (S3) and (S4) which, for this example, are the following groups of formulas:

(S3) (i) For $t \in \{c_1, \dots, c_6, r_1, \dots, r_4\}$,

$\neg \text{ROAD}(t)$,
 $\neg \text{RIVER}(t)$,
 $\neg \text{SHORE}(t)$,
 $\neg \text{LAND}(t)$,
 $\neg \text{WATER}(t)$.

(ii) For $s \in \{R_1, \dots, R_4\}$,

$\neg \text{RIVER}(s)$,
 $\neg \text{ROAD}(s)$,
 $\neg \text{SHORE}(s)$,
 $\text{LAND}(s) \vee \text{WATER}(s)$,
 $\neg \text{LAND}(s) \vee \neg \text{WATER}(s)$.

(iii) For $s \in \{C_1, \dots, C_6\}$

$\neg \text{LAND}(s)$,
 $\neg \text{WATER}(s)$.

Table 1
The interpretations of predicates with closure formulas

Predicate	Extension ($ \text{Predicate} $)
<i>cham</i>	$c_1, c_2, c_3, c_4, c_5, c_6$
<i>region</i>	r_1, r_2, r_3, r_4
<i>tee</i>	$(c_2, c_1), (c_2, c_3), (c_4, c_5), (c_3, c_5)$
<i>chi</i>	$(c_3, c_4), (c_4, c_3)$
<i>bounds</i>	$(c_1, r_1), (c_2, r_1), (c_3, r_1), (c_3, r_2),$ $(c_4, r_1), (c_4, r_2), (c_5, r_1), (c_5, r_2), (c_5, r_3),$ $(c_6, r_3), (c_6, r_4)$
<i>closed</i>	c_5, c_6
<i>interior</i>	$(c_6, r_4), (c_5, r_3)$
<i>exterior</i>	$(c_5, r_1), (c_5, r_2), (c_6, r_3)$
<i>image-object</i>	$c_1, c_2, c_3, c_4, c_5, c_6, r_1, r_2, r_3, r_4$
Δ	$(c_1, C_1), (c_2, C_2), (c_3, C_3), (c_4, C_4), (c_5, C_5),$ $(c_6, C_6), (r_1, R_1), (r_2, R_2), (r_3, R_3), (r_4, R_4)$
LINEAR-SCENE-OBJECT	$C_1, C_2, C_3, C_4, C_5, C_6$
AREA	R_1, R_2, R_3, R_4
JOINS	$(C_2, C_1), (C_2, C_3), (C_4, C_5), (C_3, C_5)$
CROSS	$(C_3, C_4), (C_4, C_3)$
BESIDE	$(C_1, R_1), (C_2, R_1), (C_3, R_1), (C_3, R_2), (C_4, R_1),$ $(C_4, R_2), (C_5, R_1), (C_5, R_2), (C_5, R_3), (C_6, R_3),$ (C_6, R_4)
LOOP	C_5, C_6
INSIDE	$(C_6, R_4), (C_5, R_3)$
OUTSIDE	$(C_5, R_1), (C_5, R_2), (C_6, R_3)$
SCENE-OBJECT	$C_1, C_2, C_3, C_4, C_5, C_6, R_1, R_2, R_3, R_4$

(iv) For $s \in \{C_5, C_6\}$,

ROAD(s) \vee SHORE(s),
 \neg ROAD(s) \vee \neg SHORE(s),
 \neg RIVER(s).

(v) For $s \in \{C_1, \dots, C_4\}$,

ROAD(s) \vee RIVER(s),
 \neg ROAD(s) \vee \neg RIVER(s),
 \neg SHORE(s).

(S4) (i) \neg RIVER(C_3) \vee \neg RIVER(C_4).

(ii) SHORE(C_5) \supset LAND(R_3) \equiv WATER(R_1),

SHORE(C_5) \supset LAND(R_3) \equiv WATER(R_2),

SHORE(C_6) \supset LAND(R_4) \equiv WATER(R_3).

- (iii) $\text{LAND}(R_1)$,
 $\text{LAND}(R_2)$,
 $\text{ROAD}(C_5) \supset \text{LAND}(R_1)$,
 $\text{ROAD}(C_5) \supset \text{LAND}(R_2)$,
 $\text{ROAD}(C_4) \supset \text{LAND}(R_3)$,
 $\text{ROAD}(C_6) \supset \text{LAND}(R_3)$,
 $\text{ROAD}(C_6) \supset \text{LAND}(R_4)$,
- (iv) $\text{RIVER}(C_1) \supset \text{false}$,
 $\text{RIVER}(C_2) \supset \text{RIVER}(C_1) \vee \text{RIVER}(C_3)$,
 $\text{RIVER}(C_3) \supset \text{SHORE}(C_5)$,
 $\text{RIVER}(C_4) \supset \text{SHORE}(C_5)$

After a certain amount of simplification (which would require a propositional theorem prover in general) we obtain the following equivalent set of formulas.

The formulas of (S3)(i) above

For $s \in \{R_1, \dots, R_4\}$,

$\neg \text{RIVER}(s)$,

$\neg \text{ROAD}(s)$,

$\neg \text{SHORE}(s)$,

$\text{LAND}(R_1)$,

$\text{LAND}(R_2)$,

$\neg \text{WATER}(R_1)$,

$\neg \text{WATER}(R_2)$.

For $s \in \{R_3, R_4\}$,

$\text{LAND}(s) \vee \text{WATER}(s)$,

$\neg \text{LAND}(s) \vee \neg \text{WATER}(s)$.

The formulas of (S3)(iii) above.

The formulas of (S3)(iv) above.

For $s \in \{C_2, \dots, C_4\}$,

$\text{ROAD}(s) \vee \text{RIVER}(s)$,

$\neg \text{ROAD}(s) \vee \neg \text{RIVER}(s)$,

$\neg \text{SHORE}(s)$,

$\text{ROAD}(C_1)$,

$\neg \text{RIVER}(C_1)$,

$\neg \text{SHORE}(C_1)$.

The formulas of (S4)(i) above.

- SHORE(C_5) \supset WATER(R_3) ,
- SHORE(C_6) \supset LAND(R_4) \equiv WATER(R_3)
- ROAD(C_5) \supset LAND(R_3) ,
- ROAD(C_6) \supset LAND(R_3) ,
- ROAD(C_6) \supset LAND(R_4) .
- RIVER(C_2) \supset RIVER(C_3) ,
- RIVER(C_3) \supset SHORE(C_5) ,
- RIVER(C_4) \supset SHORE(C_5) .

It is a simple but tedious matter to determine all propositional models of these formulas; there are six of them, as summarized in Table 2. This means there are six possible interpretations of the original image.

The problem of determining all propositional models of these formulas can be formulated as a classical constraint satisfaction problem (CSP) [13] in two different ways. First, the problem of satisfiability of a propositional conjunctive normal form formula, SAT, is a CSP in which each atom is a variable with domain $\{true, false\}$ and each clause is a constraint on the values of the atoms in the clause. In an alternative formulation, there are ten variables $\{C_1, \dots, C_6, R_1, \dots, R_4\}$. For the variables $\{C_1, \dots, C_6\}$ the domain of possible values is $\{ROAD, RIVER, SHORE\}$; for the variables $\{R_1, \dots, R_4\}$ the domain of possible values is $\{WATER, LAND\}$. Each propositional formula corresponds to constraint (either unary, binary or ternary) on the sets of possible values allowed for the variables mentioned in the formula. Although, in general, CSPs are NP-hard there are several efficient approximation algorithms that may be useful. Network consistency approximation algorithms have been developed and used extensively in the Mapsee project [18].

In connection with implementing an image interpretation system, notice that the general form of SIMP-AXIOMS of Section 5, specifically the formula groups

Table 2
The six interpretations of the map of Fig 3

Extension	Predicate				
	ROAD	RIVER	SHORE	LAND	WATER
1	$C_1, C_2, C_3, C_4, C_5, C_6$			R_1, R_2, R_3, R_4	
2	C_1, C_2, C_3, C_4, C_5		C_6	R_1, R_2, R_3	R_4
3	C_1, C_2, C_3, C_4		C_5, C_6	R_1, R_2, R_4	R_3
4	C_1, C_2, C_3	C_4	C_5, C_6	R_1, R_2, R_4	R_3
5	C_1, C_2, C_4	C_3	C_5, C_6	R_1, R_2, R_4	R_3
6	C_1, C_4	C_2, C_3	C_5, C_6	R_1, R_2, R_4	R_3

(S3) and (S4), strongly suggests the use of a relational database system [15]. Predicates like CROSS, LOOP, etc. can be naturally viewed as relations, and |CROSS|, |LOOP|, etc. as their corresponding relational tables. For computations involving these tables, we can use the relational algebra which was designed specifically for the manipulation of such tables [15, Chapter 2]. For example by appealing to the join operator of the relational algebra, the formula group (S4)(ii) may be expressed as.

$$\text{For } (x, y, z) \in |LOOP| \bowtie_{1,1} (|INSIDE| \bowtie_{1,1} |OUTSIDE|), \\ \text{SHORE}(x) \supset \text{LAND}(y) \equiv \text{WATER}(z),$$

where $\bowtie_{i,j}$ indicates that the join is taken over the i th and j th columns of the first and second operands respectively of the join operator⁴.

By appealing to relational database systems in this way, computational vision can exploit the efficient storage, retrieval, and special purpose hardware of current and future database technologies. This can be especially important for vision applications since the relational tables obtained from complex images are likely to be quite large.

In connection with databases and vision, it is interesting to note that Bibel [2] proposes solving constraint satisfaction problems by means of the relational algebra. As we have just seen, SIMP-AXIOMS leads to a constraint satisfaction problem whose solution yields all interpretations of a sketch map. We therefore have the prospect of relational databases playing a major implementation role in high level vision.

7. Some Additional Features of this Framework for Depiction

We have emphasized that a logical foundation for high level vision provides a rigorous definition for the concept of an interpretation of an image. We have also demonstrated how logic can be used to refine a logical specification of an interpretation task to an algorithmic realization of this task. There are, however, other important advantages of a logical perspective. We sketch some of these here.

7.1. Incorporating contingent knowledge

In our axiomatization of hand drawn sketch maps, the scene axioms of Section 2.2 reflected general knowledge of the scene domain. These axioms were fixed in advance and, with the help of the other axioms, were refined to the groups of propositional formulas (S3) and (S4) of SIMP-AXIOMS. These formulas are used to determine all interpretations of a given image.

⁴The reader unfamiliar with the relational algebra can safely ignore this example. The important point is that the relational algebra provides operators for manipulating relational tables and that these have been implemented and optimized in current relational database systems.

It often happens, however, that *contingent* knowledge is available about a *particular* scene. Such knowledge is not universal to all scenes, nor can it be anticipated in advance. For example, we may know a priori something about the geographic region depicted by a particular sketch map, perhaps that the area contains a river with two tributaries, and it flows into a shore. This item of contingent knowledge is an additional constraint on the possible interpretations of the map, and must be exploited in computing these. The particular fact has the following logical representation:

$$\begin{aligned} & (\exists r, s) \text{ RIVER}(r) \wedge \text{SHORE}(s) \wedge \text{JOINS}(r, s) \\ & \wedge (\exists r_1, r_2) \text{ RIVER}(r_1) \wedge \text{RIVER}(r_2) \wedge r_1 \neq r_2 \\ & \wedge \text{JOINS}(r_1, r) \wedge \text{JOINS}(r_2, r) . \end{aligned}$$

Conceptually, to accommodate this new information, we need only add it to MAP-AXIOMS and find all models of the resulting formulas. Computationally, because MAP-AXIOMS must be refined to SIMP-AXIOMS, the contingent knowledge must similarly be refined. For the example at hand, it is straightforward to carry out this refinement using the methods referred to in Section 5. We obtain the formula

$$\bigvee_{(r, s) \in |\text{JOINS}|} [\text{RIVER}(r) \wedge \text{SHORE}(s) \wedge \bigvee_{\substack{\{(r_1, r_2) | r_1 \neq r_2 \text{ and } (r_1, r) \in |\text{JOINS}| \\ \text{and } (r_2, r) \in |\text{JOINS}|\}}} (\text{RIVER}(r_1) \wedge \text{RIVER}(r_2))] .$$

This can be added to SIMP-AXIOMS, and interpretations computed as before.

It is clear in general how contingent knowledge can be accommodated by a logical approach to high level vision, at least conceptually. One merely augments the axiomatization with the contingent facts. The interpretations of an image are the models of the enlarged axiom set. Computationally realizing this approach is another matter entirely. Such contingent scene knowledge must be transformed in exactly the same way as general scene knowledge as a first step in computing the interpretations, and these transformations must be algorithmically determined. For our simple sketch map world, specifying these transformations would be relatively straightforward, although we have not done so in this paper. For more general settings, the problem of automatically accommodating contingent knowledge remains a future research topic.

7.2. Querying an image

In many applications one is not concerned with finding some or all interpretations of an image. Rather, one is concerned with determining whether some

property of the scene is depicted in a given image. For example, in our map world, we might wish to know whether part of what the image depicts is a road leading to a shore. Formally, this query is

$$(\exists r,s) \text{ROAD}(r) \wedge \text{SHORE}(s) \wedge \text{JOINS}(r,s).$$

In general, a query Q can be any formula. If AXIOMS is a set of formulas formalizing the application under consideration, the query has answer “yes” provided it is true in all interpretations of the image, i.e. provided

$$\text{AXIOMS} \models Q.$$

Q has answer “no” provided it is false in all interpretations of the image, i.e. provided

$$\text{AXIOMS} \models \neg Q.$$

Otherwise, its answer is “possibly”, which is to say it is true in some, but not all interpretations of the image.

One approach to answering a query is to compute all interpretations of the image, then determine the truth values of Q in each such interpretation. The obvious problem with such an approach is that it is completely bottom up; the query does not participate in the computation of interpretations. If answering the query requires just a few image properties, or involves only a small local region of the image, we can hope to do better than a generate and test algorithm. The natural approach is to invoke a theorem prover, which attempts to derive one or both of Q and $\neg Q$ using AXIOMS as premises. Notice, however, that just as was the case for accommodating contingent knowledge, the axioms to be used for image interpretation will be some refined version of the original specification. In our map world, SIMP-AXIOMS is such a refinement of MAP-AXIOMS. The example query above would also have to be similarly refined to the equivalent

$$\bigvee_{(r,s) \in |\text{JOINS}|} \text{ROAD}(r) \wedge \text{SHORE}(s)$$

prior to a theorem proving computation with SIMP-AXIOMS as premises. Moreover, the theorem to be proved should be instrumental in guiding the search for its proof, so some mechanism will be required analogous to the set of support strategy in resolution theorem proving (Wos et al. [27]), or top-down derivations in PROLOG (Kowalski [9]). Since one can expect that this final theorem proving task will frequently be propositional, it is likely to appeal to constraint satisfaction techniques. In this case, we shall require mechanisms

whereby the theorem actively guides the search for solutions to a constraint satisfaction problem. Finally, when the answer to a query is “possibly”, we shall normally want to determine those image interpretations in which the query is true. All these issues remain totally unexplored in the vision setting.

7.3. Accommodating ambiguity in image descriptions

Ambiguity arises in vision in two fundamentally different ways. First, a well-specified image, for example the sketch map in Fig. 3, may have multiple scene interpretations. This scene ambiguity is reflected in the fact that the image, scene and mapping axioms may have multiple models (six in the case of Fig. 3). Second, the image itself may have multiple descriptions. Here we deal with this possibility.

The image axioms of Section 2 1 for our map world formalize the assumption that our information about the image is complete; the closure axioms state that we know all and only the instances of images relations like *tee* and *bounds*, while the unique names axioms provide complete information about the equality relation. This assumption of complete information is a gross simplification.

Consider Fig. 5 where the result of imperfect segmentation or careless drawing leaves open the possibility of a *tee* or a *chi* in the image. This setting can easily be represented by the image axiom

$$tee(c_1, c_2) \vee chi(c_1, c_2)$$

Of course, we now lose the closure axioms for *tee* and *chi*. This in turn leads to the loss of closure axioms for JOINS and CROSS which will have repercussions for the simplifications of MAP-AXIOMS derived in Section 5. Exploring the consequences of such ambiguities in the image description remains an open problem

Figure 6 illustrates a more interesting example of ambiguity in an image description, because it affects the treatment of the equality relation. The question is whether to treat chains c_1 and c_2 as a single continuous chain, in which case r_1 and r_2 must be distinct regions, or as two separate chains, in which case r_1 and r_2 are identical regions. We adopt the convention that $c_1 = c_2$ means that c_1 and c_2 define a single continuous chain, i.e. that this single chain




Fig 5 An image with two possible descriptions




Fig 6 A broken chain?

has two different *names*. Similarly with respect to the regions r_1 and r_2 . This setting can now be formalized as follows:

$$\begin{aligned}
 (\forall x) \text{ chain}(x) &\equiv x = c_1 \vee x = c_2 \vee x = c_3, \\
 (\forall x) \text{ region}(x) &\equiv x = r_1 \vee x = r_2, \\
 c_1 = c_2 &\equiv r_1 \neq r_2, \\
 c_1 \neq c_3, \quad c_2 \neq c_3, \quad c_i &\neq r_j
 \end{aligned}$$

Notice that closure axioms for *chain* and *region* are preserved. The unique names axioms of Section 2.1 are *not* preserved. Specifically, the image axioms do not contain the unique names axioms $c_1 \neq c_2$ and $r_1 \neq r_2$. In settings like this, where the full set of unique names axioms must be abandoned, an equality reasoner will be necessary for computing image interpretations. The consequences for vision of incomplete information about the equality relation remains an open problem.

It is precisely with respect to the specification of incomplete information that logic excels as a representation language. While the consequences of such incomplete axiomatizations may be far from obvious, there can be no question of just what it is about an image that is being formally specified. This is particularly important when the image description is ambiguous.

7.4. Occlusion

To this part our sketch map world admits only two-dimensional scenes. For example, MAP-AXIOMS precludes occlusion. This results from the mapping axiom (iii) in Section 2.3:

$$(\forall s) \text{ SCENE-OBJECT}(s) \supset (\exists ! i) \text{ image-object}(i) \wedge \Delta(i, s).$$

To see why, consider Fig. 7 which depicts a bridge passing over what might be a river or a road occluded by the bridge. If R denotes this occluded road or river, then $\Delta(c_1, R)$ and $\Delta(c_2, R)$; since $c_1 \neq c_2$ the uniqueness property of the above mapping axiom is violated.




Fig 7 An occluding bridge

To accomodate occlusions of this kind we must relax the above mapping to

$$(\forall s) \text{SCENE-OBJECT}(s) \supset (\exists i) \text{image-object}(i) \wedge \Delta(i, s) .$$

One consequence of this is that we lose the unique names formulas $\sigma(I_m) \neq \sigma(I_n)$ when $m \neq n$ for scene objects (see the proof of Result 4.3(i)). But as we are about to see, this price must be paid anyway in order to properly formalize occlusions of this kind.

Following the approach of the previous section, if a linear scene object is occluded so that its image contains two distinct chains c_1 and c_2 , we adopt the convention that $\sigma(c_1) = \sigma(c_2)$ means that the two chains depict one and the same scene object. Equivalently, $\sigma(c_1)$ and $\sigma(c_2)$ are two different names for the same scene object. With respect to Fig. 7, $\sigma(c_1) = \sigma(c_2)$ means that there is a single scene object depicted by the two chains c_1 and c_2 . We can formalize bridge occlusions by the following scene axiom:

$$(\forall b, s_1, s_2, l_1, l_2) \text{BRIDGE}(b, s_1, s_2) \wedge \text{JOINS}(l_1, s_1) \wedge \text{JOINS}(l_2, s_2) \\ \supset l_1 = l_2 \wedge (\text{ROAD}(l_1) \vee \text{RIVER}(l_1)) .$$

Here, $\text{BRIDGE}(b, s_1, s_2)$ means that b is a bridge with sides s_1 and s_2 .

Notice that this axiom forces us to abandon unique names for scene objects (Result 4.3(i)), much as the representation of ambiguous image descriptions of the previous section led to the rejection of some unique names for image objects. Notice also that the centrality of the equality relation for a proper treatment of occlusion is non unique to our analysis. Whenever Guzman's [6] SEE program uses the back-to-back T's heuristic to link two regions in an image of a polyhedral scene it is, in effect, declaring that those regions depict a single surface.

We do not presume to have solved the occlusion problem. Also there may well be other reasons for weakening the mapping axiom (iii) in Section 2.3. Many scene objects may not appear at all in the image because they are at the wrong scale, outside the frame of the map, inappropriate to the theme of the map or are totally occluded by, for example, a legend. The ramifications of abandoning unique names for image and scene objects requires exploration, as does the weakening of the mapping axiom (iii) in Section 2.3. What does emerge clearly is the centrality of the equality relation for reasoning about and representing occlusion.

7.5. Complex objects

In our treatment of sketch maps, we have considered only simple scene objects like roads and rivers, that is, objects with no component parts. Most vision settings involve complex objects consisting of aggregations of components which in turn may have components, etc. This observation has motivated the designers of several vision systems to incorporate *composition hierarchies* for the definition of complex objects [3, 7, 25]

We indicate how such complex objects may be treated in our logical setting. As an example, consider the concept of a river system which informally is a maximal collection of interconnecting rivers at least one of which flows into a shoreline. As in most treatments of composition in the vision literature, we appeal to a predicate $\text{PART-OF}(x, y)$ meaning that object x is a component of the more complex object y .

(1) Every river r is part of a unique river system which we denote by $\rho(r)$:

$$\begin{aligned} (\forall r) \text{ RIVER}(r) \\ \supset \text{RIVER-SYSTEM}(\rho(r)) \wedge \text{PART-OF}(r, \rho(r)) \\ \wedge (\forall y) \text{ RIVER-SYSTEM}(y) \wedge \text{PART-OF}(r, y) \supset y = \rho(r) \end{aligned}$$

(2) Properties of a river system.

$$\begin{aligned} (\forall x) \text{ RIVER-SYSTEM}(x) \\ \supset [(\forall y) \text{ PART-OF}(y, x) \supset \text{RIVER}(y)] \\ \wedge [(\forall r, p) \text{ RIVER}(r) \wedge \text{PART-OF}(p, x) \wedge \text{JOINS}(r, p) \\ \supset \text{PART-OF}(r, x)] \\ \wedge [(\exists s, z) \text{ SHORE}(s) \wedge \text{PART-OF}(z, x) \wedge \text{JOINS}(z, s)]. \end{aligned}$$

(3) Equality of river systems:

$$\begin{aligned} (\forall x, y) \text{ RIVER-SYSTEM}(x) \wedge \text{RIVER-SYSTEM}(y) \\ \supset x = y \equiv [(\forall p) \text{ PART-OF}(p, x) \equiv \text{PART-OF}(p, y)]. \end{aligned}$$

The introduction of complex objects into our sketch map world necessitates a number of minor changes to the axiomatization of Section 2. First, the scene domain taxonomy must be expanded to that of Fig. 8. Second, all references to the predicate SCENE-OBJECT in Section 2 and the subsequent analysis must now be replaced by the predicate $\text{SIMPLE-SCENE-OBJECT}$. In all other respects, the image interpretation process remains the same, with one exception. When the axioms for river systems are taken into account, there may be fewer interpretations of an image. This is not too surprising since adding axioms may eliminate

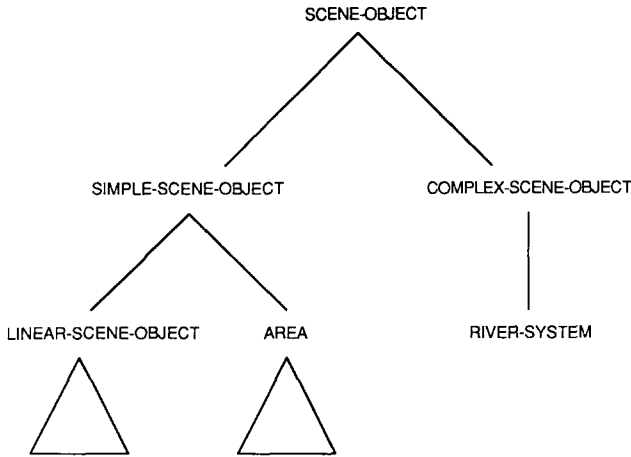


Fig 8 An expanded scene domain taxonomy




Fig 9 Two rivers?

models. For example, under MAP-AXIOMS the image of Fig. 9 has an interpretation in which RIVER(C_1) and RIVER(C_2), but the first two axioms above for river systems preclude this interpretation.⁵

In this section we have merely sketched how complex objects may be accommodated in a logical framework for depiction. The details of their logical representation, such as the axiomatization of RIVER-SYSTEM and PART-OF, remain to be worked out, as are algorithms for using such axioms in the interpretation process.

7.6. Characterizing preferred interpretations

On our account of high level vision, scene ambiguity is a purely logical property; multiple interpretations of an image arise from multiple models of the corresponding task axiomatization. The fact is, however, that frequently humans are unaware of all or even some of the ambiguities inherent in an image; certain interpretations are preferred over others.

⁵We omit the proof of this, although it is straightforward. The proof makes use of the taxonomy of Fig 8. It also requires unique names axioms of the form $\rho(x) \neq \sigma(y)$ i.e. that complex scene objects are different than simple scene objects

In this paper we have not addressed the important problem of characterizing preferred interpretations. At this level one can expect domain specific probabilistic information to be significant, as well as psychological data. It is possible that purely logical considerations will be relevant. For example, certain preferred interpretations may satisfy suitable extremal properties with respect to the space of all possible image interpretations. Such extremal properties arise in various formalizations of nonmonotonic reasoning [23]. In fact, since nonmonotonic reasoning is primarily concerned with plausible inferences, it is likely to play an important role in characterizing preferred (i.e. plausible) interpretations in vision.

Whatever considerations turn out to be relevant for characterizing preferred interpretations, we believe that a theory of high level vision must provide an account of all possible interpretations, not simply the psychologically preferred ones. In other words, it must provide a competence as well as a performance theory.

7.7. Graphics applications

Although we have concentrated on the task of interpreting images, the vision problem, the logic of depiction can equally well be applied to the task of generating images, the graphics problem [11]. One of the criteria of procedural adequacy is flexibility: the capacity of a knowledge representation scheme to support analysis and synthesis [14].

If we adopt the simple axioms of Section 2.3 then, based on the assumption that each scene object is depicted by a unique image object, we can postulate a function $\iota(s)$ satisfying the axiom:

$$\forall s \text{ SCENE-OBJECT}(s) \\ \supset \text{image-object}(\iota(s)) \wedge \Delta(\iota(s), s) \wedge [(\forall t)\Delta(t, s) \supset t = \iota(s)].$$

Coordinate frame transformations including metrical constraints on the scale, location and orientation of image and scene objects can be specified by the depiction relation $\Delta(t, s)$ or, equivalently, by the function $\iota(s)$.

To generate an image of a scene, one computes all models of the general image, scene and mapping axioms and the particular scene description. If the scene description is consistent (internally and with respect to the general axioms) and denotes a unique scene then it is well-specified in the sense that it is neither anomalous nor ambiguous. In that case there would be but one model of the axioms which would specify a unique image.

One of the advantages claimed for logic-based systems such as PROLOG is that there is often an element of "reversibility" in the definition of predicates: one can sometimes interchange the roles of input and output variables (Clocksin and Mellish [4]). However, in practice, one finds that PROLOG programs are

usually designed to exploit a particular direction of procedural interpretation. The analogy carries through to the logic of depiction. Just as we manipulated the axioms to support an efficient interpretation process, one would have to manipulate the axioms to support an efficient generation process. Although the knowledge base may have been optimized for a particular direction of use, these optimizations are model-preserving, which ensures that the same knowledge underlies image interpretation and generation. This guarantees, for example, that interpretation and generation are correct inverses of each other with the qualification, of course, that interpretation is, in general, a one-to-many mapping, and generation is many-to-one.

Using this approach, there are advantages for building user-computer interfaces. If an applications program is manipulating a database of objects a graphical display representing a view of those objects could be maintained by a separate system built on the principles outlined here. While the user actually interacts with the graphical description in the image domain both the user and the applications program can interpret the effects of each other's graphical actions in the scene domain.

Without changing the scene domain rules one can easily change the image formatting and object depiction rules. For example, if the applications program and the user are manipulating sets and set inclusion relationships then a scene configuration could be depicted as a conventional tree (as in Fig. 4) or the user may prefer to use Venn diagram conventions based on containment of closed regions (Wong [26]). The separation of the image, scene and mapping knowledge encourages the design of modular and correct graphics systems that go beyond device independence to image domain independence.

7.8. Beyond Mapsee

Many of the advantages of the logical framework discussed in Section 7 suggest that we can go beyond a reconstruction of some aspects of Mapsee. The extant Mapsee implementations cannot incorporate contingent knowledge, allow efficient responses to image queries, accommodate ambiguous image descriptions, deal sensibly with occlusion or generate maps; although, Mapsee-3 does deal well with complex objects. The framework presented here may prove to be a foundation for building better image-based systems.

8. Conclusion

We are far from having presented a completely adequate framework for depiction and image interpretation; however, we have outlined a formal treatment of a task level theory of model-based vision. General knowledge of the image domain, the scene domain and the depiction mapping can be expressed in first-order logic with equality. An interpretation of a particular image is a logical model of the general knowledge and a description of that

image. This perspective provides a purely logical account of scene ambiguity. It also provides a task level formulation of the interpretation problem. This specification is refined, through model-preserving transformations, to the equivalent problem of determining the satisfiability of a set of propositional formulas to which known constraint satisfaction algorithms can be applied.

This approach provides a framework for analyzing existing vision systems by a process of logical reconstruction. It also shows, for significant task domains, how to design and implement vision systems that are correct with respect to both the task and algorithm levels. The modular separation of the knowledge into three sets of axioms encourages portability and generality in the application of this framework for depiction to other domains. Consider, say, the task of interpreting diagrams of combinatorial logic circuits. Many of the image axioms will be unchanged. The generic classification of the axioms (namely, Taxonomy, Closure, Unique Names, Coherence and Type Constraints, Disjointness of Image and Scene Objects, Uniqueness of Depiction, Taxonomic Mapping and Relational Mapping) will survive. In any application, the foundation (namely, the definition of an interpretation as a model of an axiomatic formulation) will remain secure. In many applications, the methods used to transform to propositional form and the use of CSP techniques will, we hypothesize, still be appropriate. To that extent the framework is independent of the particular task and axiomatization exploited here as an example. Moreover, it apparently has applications in intelligent computer graphics as well.

We have also sketched logical approaches to the problems of contingent scene knowledge, image queries, ambiguity in image descriptions, occlusion and complex objects. These and many other issues of descriptive and procedural adequacy remain to be explored in depth.

ACKNOWLEDGMENT

We gratefully acknowledge several important discussions with Jan Mulder and Bill Havens and substantive critiques of earlier versions of the paper by Wolfgang Bibel, Ted Elcock, Paul Gilmore, Randy Goebel, Alex Kean, Bob Mercer and the referees. May Vink and Theresa Fong formatted the paper in their usual perfect ways. This work was supported by the Natural Sciences and Engineering Research Council of Canada and the Canadian Institute for Advanced Research.

REFERENCES

- 1 Bibel, W., On matrices with connections, *J ACM* **28** (4) (1981) 633-645
- 2 Bibel, W., Constraint satisfaction from a deductive viewpoint, *Artificial Intelligence* **35** (1988) 401-413
- 3 Brooks, R A., Symbolic reasoning among 3-D models and 2-D images, *Artificial Intelligence* **17** (1981) 285-348
- 4 Clocksin, W F. and Mellish, C S., *Programming in Prolog* (Springer, Berlin, 1981)
- 5 Clowes, M B., On seeing things, *Artificial Intelligence* **2** (1971) 79-116

- 6 Guzman, A , Decomposition of a visual scene into three-dimensional bodies, in *Proceedings AFIPS 1968 Fall Joint Computer Conference* (1968) 291–304
- 7 Havens, WS and Mackworth, A K , Representing knowledge of the visual world, *IEEE Comput* **16** (10) (1983) 90–96
- 8 Huffman, D A , Impossible objects as nonsense sentences, in B Meltzer and D Michie (Eds), *Machine Intelligence 6* (American Elsevier, New York, 1971), 295–323
- 9 Kowalski, R , *Logic for Problem Solving* (North-Holland, Amsterdam, 1979)
- 10 Mackworth, A K , On reading sketch maps, in *Proceedings IJCAI-77*, Cambridge, MA (1977) 598–606
- 11 Mackworth, A K , Recovering the meaning of diagrams and sketches, in *Proceedings Graphics Interface-83*, Edmonton, Alta (1983) 313–317
- 12 Mackworth, A K , Mulder, J A and Havens, WS , Hierarchical arc consistency Exploiting structured domains in constraint satisfaction problems, *Comput Intell* **1** (2) (1985) 71–79
- 13 Mackworth, A K , Constraint satisfaction, in S Shapiro (Ed), *Encyclopedia of Artificial Intelligence* (Wiley, New York, 1987) 205–211
- 14 Mackworth, A K , Adequacy criteria for visual knowledge representation, in Z Pylyshyn (Ed), *Computational Processes in Human Vision An Interdisciplinary Perspective* (Ablex, Norwood, NJ, 1988) 464–476
- 15 Maier, D , *The Theory of Relational Databases* (Computer Science Press, Rockville, MD, 1983)
- 16 McCarthy, J and Hayes, P. Some philosophical problems from the standpoint of artificial intelligence, in B Meltzer and D Michie (Eds), *Machine Intelligence 4* (American Elsevier, New York, 1969) 463–502
- 17 Mendelson, E , *Introduction to Mathematical Logic* (Van Nostrand, Princeton, NJ, 1964)
- 18 Mulder, J A , Mackworth, A K and Havens, WS , Knowledge structuring and constraint satisfaction The Mapsee approach, Tech Rept 1987CS-7, Division of Computer Science, Dalhousie University, Halifax, NS (1987), Tech Rept 87-21, Department of Computer Science, University of British Columbia, Vancouver, BC (1987)
- 19 Mulder, J A , Using discrimination graphs to represent visual knowledge, Ph D Thesis, Department of Computer Science, University of British Columbia, Vancouver, BC (1986)
- 20 Purdom, PW , Solving satisfiability with less searching, *IEEE Trans Pattern Anal Mach Intell* **6** (4) (1984) 510–513
- 21 Reiter, R , On closed-world data bases, in H Gallaire and J Minker (Eds). *Logic and Data Bases* (Plenum, New York, 1978) 55–76
- 22 Reiter, R , Equality and domain closure in first-order data bases, *J ACM* **27** (2) (1980) 235–249
- 23 Reiter, R , Nonmonotonic reasoning, *Ann Rev Comput Sci* **2** (1987) 147–187
- 24 Reiter, R and Mackworth, A K , The logic of depiction, Tech Rept RBCV-TR-87-18, Department of Computer Science, University of Toronto, Toronto, Ont (1987), Tech Rept 87-24, Department of Computer Science, University of British Columbia, Vancouver, BC (1987)
- 25 Tsotsos, J K , The role of knowledge organization in representation and interpretation of time-varying data The ALVEN system, *Comput Intell* **1** (1) (1985) 16–32
- 26 Wong, G , Depiction and domains in visual knowledge representation, M Sc Thesis, Department of Computer Science, University of British Columbia, Vancouver, BC (1986)
- 27 Wos, L , Robinson, G A and Carson, D F , Efficiency and completeness of the set of support strategy in theorem proving, *J ACM* **12** (1965) 536–541

Received August 1987, revised version received June 1988