A Semantics for Model-Based Spatial Reasoning

Janice Glasgow and Andrew Malton
Department of Computing and Information Science
Queen's University, Kingston, Ontario K7L 3N6
{janice,malton}@qucis.queensu.ca

Abstract

Model-based reasoning involves proving the truth of a proposition by computation in the semantic domain. In contrast, rule-based reasoning is proving truth by means of formal manipulation of formulas. A growing body of research in cognitive science suggests that human spatial reasoning is model-based, rather than rule-based.

The paper begins with a cognitive perspective of model-based reasoning. A semantic domain for spatial reasoning, based on a theory of symbolic arrays, is defined. A modal logic of spatial assertions for reasoning in indeterminate worlds is then presented, along with possible extensions that address structural hierarchy, temporal modalities, multiple views and analogy.

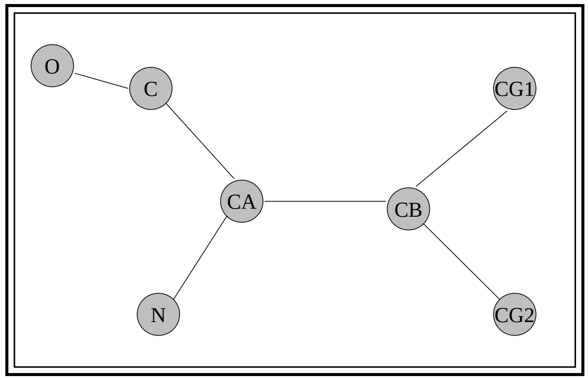

Figure 1: Spatial model of a molecular structure

1 Introduction

Psychologists have acknowledged that mental models are fundamental to human problem solving, particularly for their predictive and explanatory power in understanding human interactions with the environment and with others [66]. These models correspond to internalized representations that can be mentally inspected and transformed. Contemplate the planning problem of rearranging the furniture in your living room. One approach to solving this problem is to physically move the furniture about the room to evaluate the alternative arrangements. A less backwrenching approach is to mentally visualize and analyze the various possibilities. Mental models can also be applied metaphorically in problem solving; a heuristic cited by orators is to consider a speech as a voyage through a building where objects along the way act as cues to the next topic. Although some mental models may be specialized and require training to develop (e.g., models for reasoning about the physical behavior of complex mechanical devices), others are more accessible and correspond to everyday problem solving (e.g., a mental map for planning a route from your bedroom to the refrigerator).

Johnson-Laird [34] describes several types of mental models. The first, and most fundamental, is a *relational model,* which is a static frame consisting of tokens that represent entities in the world and a set of relations that define the physical relationships among entities. A *spatial model* is a relational model in which the relations of interest are spatial in nature; tokens are located within a symbolic, multi-dimensional space. For example, the graphical depiction in Figure 1 could be considered a spatial model of a molecular structure: each entity (atomic part) within the structure is represented as a symbolic token (node in the graph) and structural relations among entities (relative location and bonding) are represented using spatial dimensions and links. This model is neither complete nor totally accurate; knowledge about the bond lengths and angles, and the relative size of atoms is not captured. However, it does explicitly depict information that can be used for reasoning about molecular structure and interactions, while discarding irrelevant details.

Just as mental models are pervasive in human problem solving, computational models for spatial reasoning provide a foundation for problem solving in AI. One active area of research where models play an important role is qualitative physics [12]. Studies in this

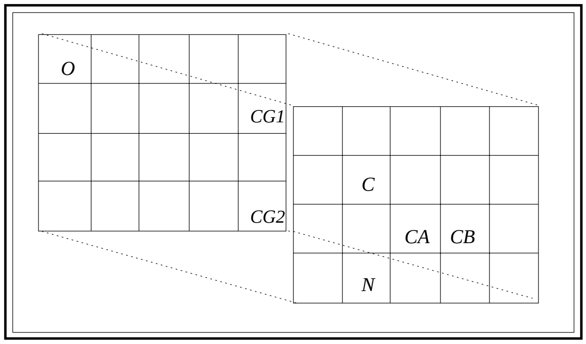

Figure 2: Array representation of molecular structure

problem domain focus on explaining the behavior of the physical world in qualitative terms, rather than by exact quantitative calculations. Mental models have also been studied in the context of mechanistic devices such as physical machines, electronic and hydraulic devices or reactors [9]. The dynamic nature of model-based reasoning makes it particularly suited for planning problems, since it can be used to simulate actions and events. The importance of spatial reasoning with models encompasses many other domains, including computer vision, robotic navigation and spatial databases.

This paper is concerned with the development of a computational methodology for spatial reasoning with models. A knowledge representation scheme is presented in which symbolic array data structures are used, in conjunction with imagery inspection and transformation operations, to reason about the spatial properties of a domain. Each dimension in an array defines a linear order relation among entities in the domain. The relative locations of symbols in the array may correspond to spatial (e.g., left-of), geographic (e.g., north-of), temporal (e.g., before) or conceptual (e.g., taller-than) relations in the world. In particular, we are concerned with transitive relations – that is, relations r such that if $r(x, y)$ and $r(y, z)$ then $r(x, z)$. For example, the spatial model in Figure 1 could alternatively be depicted as the symbolic array depicted in Figure 2. In this representation the spatial locations of symbols within the array correspond to the relative locations of atoms within a three-dimensional molecular structure. Topological relations, such as *bonded-to, contained-in, touching,* etc., can also be modelled in an array representation. In the above example, adjacency (within distance one in all dimensions) in the array denotes bonding between atoms in the molecular structure. Reasoning at varying levels of abstraction can be achieved in the scheme by defining array representations as recursive data structures; symbols in the array can themselves denote models at a more detailed level of abstraction.

The formalism presented in this paper has evolved from research in the area of computational imagery [24, 17, 18, 52], which involves the study of AI knowledge representation and inferencing techniques that correspond to the representations and processes for mental imagery. In the previously proposed scheme for computational imagery, a mathematical theory of arrays provides a basis for representing and reasoning about visual and spatial properties of entities in the world. Although results of cognitive studies offered initial motivation for the representations and functionality of the formalism, the ultimate concerns of research in computational imagery are expressive power, inferential adequacy and efficiency.

The paper presents a formal semantics for spatial reasoning with array representations (models) of possible worlds. This semantics offers a basis for deductive reasoning, where inferences are based on a semantic theory of relational deductions, rather than on a syntactic theory that depends on rules of inference. Incomplete or uncertain knowledge may result in worlds with multiple possible interpretations, where each consistent interpretation is represented by a unique array model.

To provide a cognitive perspective of the research area, the paper begins with a discussion of mental models and their role in spatial reasoning and problem solving. A representation scheme for spatial reasoning using symbolic arrays is presented in Section 3. Section 4 illustrates how array representations can be incorporated in a deductive reasoning scheme using a model-theoretic semantics. An ongoing issue in AI is how to update a knowledge base as new information is added or as a world is transformed. Section 5 considers this issue by demonstrating how nonmonotonicity can be addressed in the formalism. Extensions to the representation scheme – for reasoning with hierarchical models, for temporal and analogical reasoning, and for inspecting models from alternative perspectives – are introduced in Section 6. The paper concludes with a discussion of related work and a summary of the major contributions of the described research.

2 Mental Models

Results of experimental studies in cognitive psychology suggest that much of human problem solving is not achieved through rule-based reasoning, but rather through the manipulation of mental models. That is, humans often reason by constructing and transforming a class of representations that are structurally similar to the reality they depict. The primary purpose of this paper is to present a computational framework for reasoning about spatial models of the world. Although we do not suggest that the proposed representation scheme is necessarily a model of cognition, an understanding of the underlying principles and behavior of mental models is important to the development of AI systems for spatial reasoning and problem solving. In particular, it is useful to discriminate between mental models and other forms of mental representation.

This section describes the concept of mental models in terms of a set of underlying principles. It also discusses the role of mental models in human reasoning and relates models to representations for logic and mental imagery.

2.1 Principles of Mental Models

Although there does not appear to be an agreed on account for what constitutes a mental model, Johnson-Laird [34] has proposed some weak constraints on these representations:

- Mental models, and the machinery for constructing and interpreting them, are computable and finite.

- A description of a single state of affairs is represented by a single mental model. Indeterminacies are directly represented only if their use does not result in exponential growth in the number of models.

- The structure of a mental model is isomorphic to the structure of the state of affairs it represents; a model is constructed from tokens corresponding to the entities in the world.

Johnson-Laird asserts that models are akin to how people perceive the world, yet may be incomplete or simplified. Moreover, mental models are specific, and can be used to represent relations concerning space or time. Inferences are drawn, not through the application of formal rules, but through the construction and inspection of alternative models that are used to validate or refute a putative conclusion.

Problem solving with spatial models is often associated with the reasoning abilities of mental imagery. A large body of experimental data has been generated and theories proposed concerning the representations involved in imagery. These theories fall into three categories: 1) theories that suggest that image representations are analog or picture-like [37, 60], 2) theories that liken image representations to linguistic descriptions [54], and 3) those that suggest that there may exist multiple image representations, corresponding to different task demands [11]. Johnson-Laird proposes, as a resolution to the imagery debate, that there exist three kinds of representation involved in imagery: a propositional representation, a mental model and a visual image. What distinguishes a mental model from other forms of representation is the degree of specificity, which can be measured by the amount of information that is made explicit by the representation. Mental models are less specific than visual representations – they may discard features such as shape and size – yet they are more specific than propositional representations. A knowledge representation scheme for computational imagery has previously been proposed [24]. This scheme includes a semantic network representation for long-term memory, corresponding to Johnson-Laird's propositional representation, and two working-memory representations corresponding to the mental model and the visual representation. Computational imagery involves tools and techniques for visual-spatial reasoning, where image depictions are generated or recalled from long-term memory and then manipulated, transformed, scanned, associated with similar forms (constructing spatial analogies), and so forth.

2.2 Reasoning with Mental Models

One purpose of a mental model is to simulate and thus predict and/or plan for the behavior of a system. Humans are adept at reasoning about space, yet it is not well understood how this is accomplished. Forbus [12] suggests that it is not through logical theorem proving or through algebraic calculations, but through diagrammatic reasoning, that we achieve this competence. He states that the spatial structure of a diagram allows us to use our perceptual apparatus to inspect and interpret models in a way that is analogous to how we inspect and interpret entities in the world. He further conjectures that people can reason with less detailed representations than diagrams – representations that symbolically describe places and relationships among these places.

Theories of inference based on mental models have suggested that the processing of syllogisms can be achieved by the inspection of symbolic spatial models [31]. In such theories, a model is constructed in which tokens corresponding to entities in the world are mapped along an axis corresponding to comparative dimensions such as taller-than/shorter-than or older-than/younger-than. For example, the description "*John is taller than Mary and Mary is taller than Jane*" could be represented as an array

John	*Mary*	*Jane*

, where the *left-of* relation in the array corresponds to the transitive *taller-than* relation among entities in the world. Using this model, questions such as *Is John taller than Jane?* can be answered by applying inspection operations analogous to those used in visual inspection. Although it is possible to construct a logical description and rules of inference for syllogistic reasoning, experimental results suggest that mental models that incorporate array representations increase the efficiency and accuracy of problem solving involving transitive inferences [59, 38].

Several recent studies have focussed on the mental models resulting from textual descriptions of scenes and how these can be used to retrieve spatial information. Denis [10] presents experimental results that support a stage of text processing that represents a model of the world in terms of the spatial configuration among entities. In his study, subjects were presented with spatial descriptions such as:

> In the extreme north-west part there is a mountain. To the east, there is a forest. To the east of the forest, there is a lake. In the extreme south-west part, there is a meadow. To the east of the meadow, there is a cave. To the east of the cave, there is a desert.

which is representable as an array of the form:

mountain	*forest*	*lake*
meadow	*cave*	*desert*

.

Denis suggests that the cognitive advantage to such a model in linguistic text processing is that "it allows readers to make inferences without necessary recourse to computations based on formal logic".

Experiments carried out by Taylor and Tversky [68] presented subjects with both route and survey descriptions of spatial domains. Their findings included the fact that the subjects constructed mental models that were sufficiently abstract to allow inferences from alternative perspectives. They propose that the advantage of such a representation lies in its flexibility, since it supports exploration of a world from unique points of view as well as adaptation resulting from change in the environment. Other studies by Tversky [70, 71] provide evidence that spatial mental models might be distorted by an alignment with existing landmarks or frames of reference. For example, when college students were asked to choose a correct map of America from two possibilities, the majority chose the incorrect version which was altered so that South America appeared directly below North America. Related experiments show that people incorrectly believe that Reno is east of San Diego, based on the knowledge of the relative locations of their respective states [65].

Johnson-Laird [35] cites three fundamental differences between reasoning with mental models and reasoning with logical representations:

- Model-based reasoning is semantic: it relies on the construction and inspection of alternative models, where each model represents a unique state of affairs. Logic-based reasoning is generally syntactic: conclusions are formed by applying rules of inference to syntactic forms in order to derive new forms.

- In mental models, symbolic tokens correspond to individual entities; models do not contain variables. Much of logical reasoning is based on the instantiation of generalized terms containing variables.
- Whereas logical forms mirror the structure of discourse, mental models are structured to mirror the relations among entities.

Furthermore, Johnson-Laird suggests that the principles involved in mental models have serious advantages for computational reasoning. In particular, they allow for the integration of deductive and nonmonotonic reasoning: derivations occur by simple model checking (inspection of model representations) and updating of models can be achieved without the cost of undoing previously computed deductions resulting from default reasoning. In the following sections we support this claim by demonstrating how deduction and nonmonotinicity can be integrated in a computational approach to model-based spatial reasoning.

3 Representation of Spatial Models

This section presents a knowledge representation scheme for model-based reasoning in which, similar to the spatial component of computational imagery [24], symbolic arrays depict the entities and relations in a world. The scheme was developed using a theory of arrays. Array theory is the mathematics of rectangularly arranged, nested data objects [48]. An array consists of zero or more items held at positions along multiple axes, where rectangular arrangement is the concept of objects having spatial positions relative to one another in the collection. Similar to set theory, array theory is concerned with the concepts of aggregation, nesting and membership. An array can be considered as a multi-dimensional generalization of the list data structure used in Lisp [32]. The representation of spatial models involves a special class of arrays – those whose symbols and structure denote entities and their relative locations in the domain of interest. In order to specify spatial relations, a symbol may occupy one or more cells of an array. For example, the description – *The ball and the lamp are on the table; the lamp is to the right of the ball* – could be represented as the array:

ball	*lamp*
table	*table*

where the symbols *lamp, ball* and *table* are mapped to the corresponding entities in the world and the spatial relations of interest are *on-top-of* and *right-of*. For the purpose of this paper, we employ a convention for depicting arrays where adjacent identical symbols are represented in a single cell comprised of multiple locations in the array. For example, the above array would be depicted as:

<table><tr><td>ball</td><td>lamp</td></tr><tr><td colspan="2" align="center">table</td></tr></table>

.

Visual information such as shape, relative distance and relative size is often discarded in a model. However, if desired, distance and shape attributes can be preserved in the array representation. Figure 3(a) illustrates an island map similar to the one used by Kosslyn

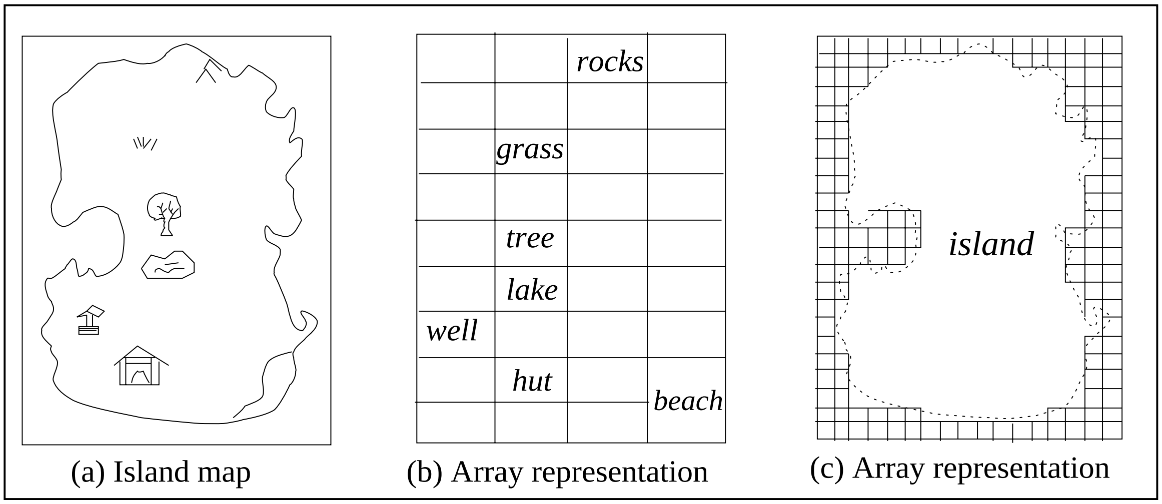

Figure 3: Representations of Kosslyn's island map

and colleagues [37] to study how humans store and inspect mental maps. Much of the information derivable through the visual inspection of the map image can also be inferred from the symbolic array representation in Figure 3(b). Geographic directions are determined in this representation by comparing the relative locations of entities, e.g. the hut is *south-of* the lake and *west-of* the beach. As well, it can be determined that the tree is *near* the lake and that the beach is *closer* to the hut than it is to the lake. Relative size and shape information can be preserved in a representation by increasing the granularity of the array. For example, the shape of the island map is computable from the array representation depicted in Figure 3(c).

A large collection of total, primitive functions, chosen to express fundamental properties of arrays, are described for array theory. These functions, which subsume most of the operations of APL and Lisp, have been implemented in the programming language Nial [33]. Array theory provides a high-level language that can be used for expressing and proving properties of spatial models. It is currently being employed to specify the primitive operations for constructing, transforming and inspecting array representations. Section 3.2 describes the primitive array functions for model-based reasoning that have been implemented in Nial.

In the remainder of this section, we define an approach to knowledge representation for spatial reasoning. The scheme consists of *array representations*, which model the entities and relations in the world, and a set of *primitive array functions* for generating, inspecting and transforming representations.

3.1 Array Representation

An array representation is constructed so that there is a correspondence between the structure of the symbolic array and the structure of the world being modeled. More precisely, a world is representable by an array if there exists a mapping between symbols in the array and entities in the world that preserves the relative location of entities. Array representations provide a basis for deductive reasoning in a spatial domain.

We define a *world* as a set of entities and a set of spatial relations that are defined on the entities. Our definition assumes a finite set P of predicate symbols that describe the relations of interest in the world.

Definition. A **world** w is defined as a pair $< S, R >$ such that:

- S is a finite set of symbols that denote the *entities* of interest in the world.
- R is a P-indexed set of *spatial relations* defined over the set of symbols S for the world. Each n-ary relation in R is defined in terms of a set of n-tuples containing entities in S. The notation w_p is used to denote the relation in R corresponding to predicate symbol $p \in P$.

Similar to a world, an array representation contains a set of spatially organized parts.

Definition. An **array representation** $\mathcal{A}$ is a multi-dimensional symbolic array. The set of symbols appearing in $\mathcal{A}$ is denoted $Sym(\mathcal{A})$. A symbol may occupy more than one location in $\mathcal{A}$, but each location contains at most one symbol.

The assumption that a location in an array contains at most one symbol corresponds to the fact that at most one entity can occupy a single location in space. This does not preclude, however, the concept of containment or the fact that a symbol may denote a complex entity consisting of subentities (see Section 6).

We specify a set $\mathcal{F}$ of predefined boolean array functions to be used to inspect an array representation. The set $\mathcal{F}$ is P-indexed in order to correspond to the relations in the world. An array representation is said to be a *model* for a world if all of the world's relations are representable by the corresponding array function in the set $\mathcal{F}$.

Definition.
Given a world $w =< S, R >$ and an array representation $\mathcal{A}$, an n-ary relation $w_p \in R$ is **represented** in $\mathcal{A}$ if and only if $S = Sym(\mathcal{A})$ and for all symbols $s_1, ..., s_n \in S$:

$$f_p(s_1, ..., s_n, \mathcal{A}) = true \text{ if and only if } (s_1, ..., s_n) \in w_p,$$

where $f_p \in \mathcal{F}$.

An array representation $\mathcal{A}$ is an **array model** for world w if and only if for every relation $w_p \in R$, w_p is represented in $\mathcal{A}$.

If there exists an array model for a world w then we say that w is **representable**.

The mapping between worlds and array models is not one-to-one. For a given world there exists an equivalence class of array representations that model the world. Two arrays belong

<table>
<tr><td></td><td></td><td></td><td></td><td>Norway</td><td>Sweden</td><td>Finland</td></tr>
<tr><td></td><td rowspan="3">Britain</td><td></td><td></td><td>Denmark</td><td></td><td></td></tr>
<tr><td>Ireland</td><td></td><td>Holland</td><td rowspan="2">Germany</td><td>Poland</td><td></td></tr>
<tr><td></td><td></td><td>Belgium</td><td>Czech Republic</td><td>Slovakia</td></tr>
<tr><td></td><td></td><td colspan="2" rowspan="2">France</td><td>Switzerland</td><td>Austria</td><td>Hungary</td></tr>
<tr><td></td><td></td><td rowspan="2">Italy</td><td colspan="2">? Yugoslavia ?</td></tr>
<tr><td>Portugal</td><td colspan="2">Spain</td><td></td><td></td><td>Greece</td></tr>
</table>

Figure 4: Array representation of Europe

to the same equivalence class if they represent the same spatial relations. For example, array representations $\boxed{\mathit{ball} \mid \mathit{box}}$ and $\boxed{\mathit{ball} \mid \quad \mid \mathit{box}}$ are equivalent with respect to the *left-of* relationship.

A representable world is *complete* in the sense that all spatial relationships among entities are made explicit by the relations in R, and can thus be represented by the array inspection functions in $\mathcal{F}$. To illustrate the concept of an array model, consider the representable world described by the set $S = \{$*Britain,Portugal,Spain,...*$\}$ of countries in Europe and their set R of corresponding geographical relations indexed by the set of predicate symbols $P = \{$*north-of, west-of, east-of, south-of and borders-on*$\}$. We define an array model $\mathcal{A}$ for w where:

- $\mathcal{A}$ is the array depicted in Figure 4.

- The array functions in $\mathcal{F}$ are defined to model the spatial relations in the world. For example, the relation $w_{west\text{-}of}$ is represented in $\mathcal{A}$ using the function $f_{west\text{-}of} \in \mathcal{F}$, which is defined so that an application of the form $f_{west\text{-}of}(s_1, s_2, \mathcal{A})$ returns the value true if and only if symbol s_1 occurs in a location that is to the left of the left-most occurrence of symbol s_2 in the array data structure $\mathcal{A}$. Similarly the relation $w_{borders\text{-}on}$ is represented in $\mathcal{A}$ using the function $f_{borders\text{-}on}$, such that $f_{borders\text{-}on}(s_1, s_2, \mathcal{A})$ evaluates to true just in the case where symbols s_1 and s_2 are situated in adjacent cells of array $\mathcal{A}$.

It is worth noting here that an array representation explicitly depicts the absence, as well as the presence, of entities at relative locations. In the above example, the symbolic array of Europe can be inspected to infer that there is no country that is directly east of Britain and west of Holland. This property is particularly valuable for problem domains such as planning and navigation.

3.2 Representation of Indeterminate Worlds

In the previous section we were concerned with array models for representable worlds, that is, those that completely specify the relationships among entities. It is also possible to model worlds that are indeterminate in the sense that the specified relations may imply alternative possible worlds (or array representations). In this section we define a more general notion of a model that encompasses both determinate and indeterminate worlds. Before doing so, it is important to distinguish between worlds that are possible and those that are not. As well, we introduce an extension relation between worlds.

A *possible world* is one that is extensible to a world that is representable. Consider the world $w =< S, R >$ such that $S = \{knife, spoon, fork\}$ and $R = \{w_{left\text{-}of}\}$. If $w_{left\text{-}of}$ is a transitive, asymmetric relation defined by the pairs $\{(spoon, fork), (fork, spoon)\}$, then w is not possible since we cannot construct an array representation that preserves the relations for w. On the other hand, if $w_{left\text{-}of} = \{(fork, knife), (knife, spoon)\}$, then w is possible, and in fact determinate, since there exists a unique state of affairs for w, represented as the array:

fork	*knife*	*spoon*

The third possibility is that the world is possible, but indeterminate. Assuming the relation $w_{left\text{-}of} = \{(fork, knife), (fork, spoon)\}$, there are two possible representable extensions of w, denoted as the arrays:

fork	*knife*	*spoon*

and

fork	*spoon*	*knife*

In general, an extension of a world is one in which new relationships are added; and a possible world is considered to be determinate if it has a unique representable extension. It is indeterminate if there exists at least two distinct extensions that are representable.

Definition:
A world $w =< S, R >$ is **extensible** to a world $w' =< S, R' >$, denoted $w \preceq w'$, if and only if for all $p \in P$, $w_p \subseteq w'_p$.

A world w is **possible** if there exists a world w' such that $w \preceq w'$ and w' is representable.

A world w is **determinate** if and only if it is possible and for all representable worlds w', w'' if $w \preceq w'$ and $w \preceq w''$ then $w' = w''$.

A world w is **indeterminate** if it is possible and not determinate, that is, there exists representable worlds w' and w'' such that $w \preceq w'$, $w \preceq w''$ and $w' \neq w''$.

We now extend the notion of array model to include indeterminate worlds. In general, a world will be modelled by a finite set of array representations corresponding to the representable extensions for the world. If w is not possible, then this set will be empty; if w is determinate then the set will contain a single representation. For indeterminate worlds the set will contain at least two array representations. A model is also characterized by the set $\mathcal{F}$ of functions that are used to inspect array functions in order to represent the relations in the world.

Definition:
A world w is **modelled** as a pair $< \mathcal{W}, \mathcal{F} >$ such that $\mathcal{W}$ is a finite set of array representations and $\mathcal{F}$ is a P-indexed set of boolean array functions. An array representation $\mathcal{A}$ is an element of $\mathcal{W}$ if and only if $\mathcal{A}$ is a unique representation (with respect to $\mathcal{W}$) for a representable extension w' of w.

In Section 4 we will develop a semantics for model-based deduction based on the model for a world. Truth in a model will be based on the inspection of the array representations for a world. In particular, a formula is possibly true in a model if there exists an array representation in $\mathcal{W}$ such that the formula holds. Similarly, a formula is necessarily true if it is true for all representations in $\mathcal{W}$. The semantics presented is a form of modal logic, in which the possible worlds for the model are equated with the possible array representations that depict the world.

The initial construction of array representations that mirrors the structure of a world is a domain specific task. There are, however, some fundamental principles that can be applied to determine the structure of the arrays and the mapping of array functions to spatial relations in the world.

- Each transitive spatial relation in the world corresponds to an array function that defines an ordering of symbols along a single dimension in the array representation.
- Transitive relations that correspond to orthogonal dimensions in the world correspond to array functions that operate on orthogonal dimensions in an array.
- Transitive relations that correspond to the same dimension in the world correspond to array functions that operate on the same dimension in the array.

In summary, if there exists a dependency between spatial relations in the world then the same dependency should exist in the array representation. A world is considered n-dimensional if n is the minimal dimensionality for a representation for a world.

3.3 Primitive Functions

Approaches to knowledge representation are distinguished by the operations that are performed to carry out inferencing. Model-based reasoning with array representations is achieved by applying functions that transform and inspect array data structures in ways that are analogous to the physical transformations and visual inspections that would be applied in the world being modeled.

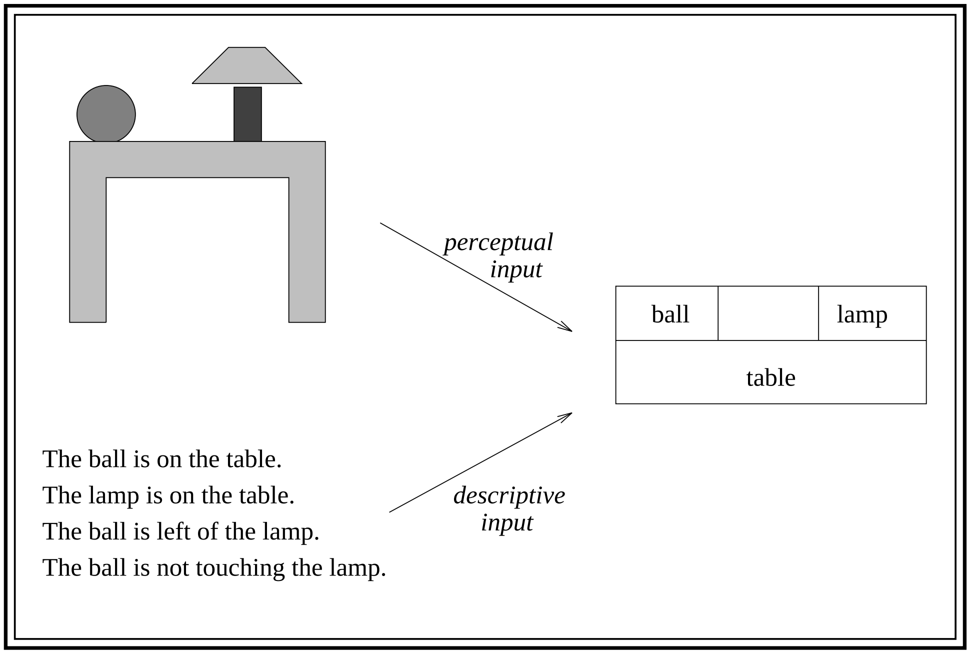

Figure 5: Initial construction of array representation

An array model for a world may initially be constructed through the interpretation of domain specific perceptual input or through the understanding of linguistic descriptions (see Figure 5). It is also possible to construct a model from existing data; algorithms have been developed to automatically transform the three-dimensional atomic coordinate information in molecular databases into three-dimensional array representations for the structures. The functions for model generation, however, are not concerned with how a model is initially generated, but with how an array representation can be constructed from model descriptions stored in long-term memory. In the current implementation of the representation scheme, these descriptions are represented as a collection of frames, where each frame in the knowledge base corresponds to description of a world, or of an entity in the world.[1] A frame may store an array representation $\mathcal{A}$ explicitly, or as a sparse array. For example, the array in Figure 5 could be described as the list of entity/locations: *(ball,(1,1)), (lamp,(1,3)), (table,(2,1),(2,2),(2,3))*. Non-spatial attributes of entities, such as the color of the ball, can also be stored in the frame knowledge base. The frames are organized in terms of conceptual and parts hierarchies, as illustrated in Figure 6.

Although some of the functions for model-based reasoning are domain specific, many are applicable across a variety of worlds. These functions may depend on a mapping between the predicate symbols for the world and the definitions in the theory. For example, the *left-of* primitive function would correspond to the $f_{west\text{-}of}$ function in the geographic example presented in Section 3.1. Table 1 describes the primitive functions for spatial reasoning that have been specified in array theory and implemented in the programming language Nial. Note that there are functions designed to a determine the possibility and necessity

[1] Although the current implementation for array models is frame based, it would also be possible to construct a logic-based approach to storing array representations[22].

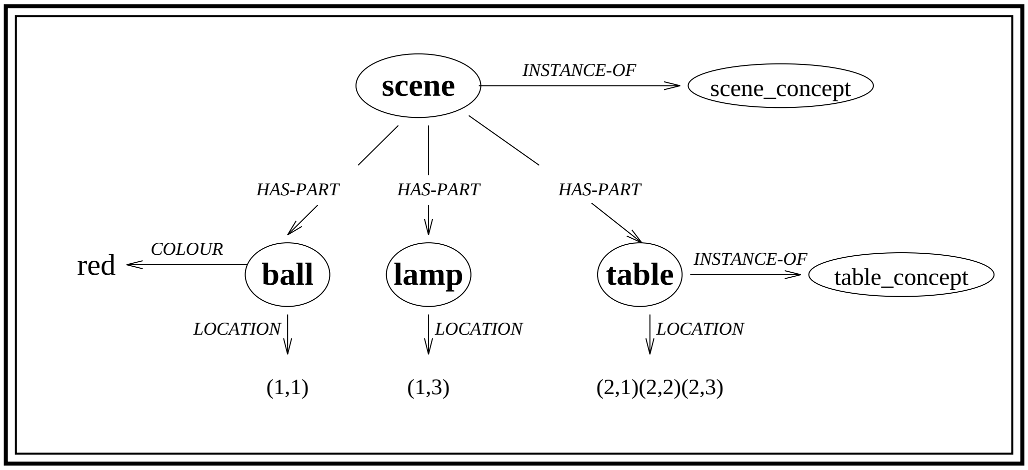

Figure 6: Frame hierarchy for long-term memory

of a relation holding given a set of array representations. These functions can be used for reasoning in indeterminate worlds (see Section 4).

The primitive array functions can be used to define more complex inspection and transformation operations for the language. For example, to determine whether one symbol is *north-west* of another, where *north-of* and *west-of* are predicates for the language, one could define the Nial function:

north-west-of is operation s_1 s_2 A $\{$*(inspect* s_1 s_2 A *north-of) and*
(inspect $s_1 s_2$ A *west-of)*$\}$

4 The Semantics of Spatial Deduction

Reasoning by deduction is the process of logically inferring a conclusion from a given set of premises. For example, from the premises:

left-of(fork,spoon),
left-of(spoon,knife), and
$\forall X \forall Y \forall Z$ *(left-of(X,Y)* $\wedge$ *left-of(Y,Z)* $\rightarrow$ *left-of(X,Z))*,

one can deduce *left-of(fork,knife)*. This form of reasoning, where conclusions are derived using the iterative application of syntactic inference rules, is referred to as *proof-theoretic*. Alternatively, the validity of an argument can be demonstrated using a *model-theoretic* approach. Given the above premises, an array representation $\mathcal{A}$ can be constructed for the world such that $\mathcal{A} = \boxed{fork \mid spoon \mid knife}$. Using this representation and the primitive inspection functions in $\mathcal{F}$, we can deduce the valid conclusion *left-of(fork,knife)* through the process of model inspection.

Most existing computational systems employ proof-theoretic deduction, where reasoning is carried out by applying rules that manipulate syntactic forms of expressions. The proposed system for spatial reasoning, however, relies on semantics, or the mapping between the representations and the domain of interest. Conclusions are derived by applying functions

FUNCTION	MAPPING	DESCRIPTION
Construction Functions:		
retrieve	$S \rightarrow F$	Retrieve frame representation for given array symbol.
reconstruct	$S \rightarrow A$	Retrieve frame representation and construct array from description.
compose	$A \times A \rightarrow A$	Compose two arrays into a single array that is consistent with other two.
Inspection Functions:		
left-of	$S \times S \times A \rightarrow B$	Determines if a symbol occurs to the left of the left-most occurrence of another symbol. (Corresponding functions have been defined for *right-of, above, below, in-front, behind* and *adjacent.*)
inspect	$S^n \times A \times P_n \rightarrow B$	Determines if two symbols are related by an n-ary predicate symbol $p \in P$.
possible	$S^n \times A^* \times P_n \rightarrow B$	Determines if a relation holds in at least one of the possible array representations.
necessary	$S^n \times A^* \times P_n \rightarrow B$	Determines if a relation holds in all of the array representations.
find	$S \times A \times P_2 \rightarrow S^*$	Returns all symbols that are related to a symbol by a binary predicate symbol $p \in P$.
equivalent	$A \times A \times P^* \rightarrow B$	Determines whether two arrays are equivalent with respect to a set of predicates.
Transformation Functions:		
put	$S \times L \times A \rightarrow A$	Places a symbol at specified locations in an array.
put_rel	$S \times S \times A \times P \rightarrow A$	Places a symbol relative to another symbol in an array.
move	$S \times L \times A \rightarrow A$	Moves a symbol to a new location set.
move_rel	$S \times S \times A \times P \rightarrow A$	Moves a symbol to a location relative to another symbol in an array.
move_forward	$S \times A \rightarrow A$	Moves a symbol one location forward in the direction specified by current orientation for symbol.
delete	$S \times A \rightarrow A$	Delete a symbol from an array.
orient	$S \times A \times O \rightarrow A$	Reorient the direction for a symbol.
rotate	$A \times O \rightarrow A$	Rotate an array based on the direction specified by an orientation vector.
focus	$S \times A \rightarrow A$	Replace a symbol with its array representation.
unfocus	$A \rightarrow A$	Undo the last focus transformation.
refocus	$A \rightarrow A$	Return to the original unnested array representation.

$S-$ array symbol; S^*- list of array symbols; $A-$ array; A^*- list of arrays; P_n- n-ary predicate symbol; $L-$ location set; $F-$ frame representation; $O-$ orientation vector; $B-$ boolean value.

Table 1: Primitive functions for model-based reasoning

that correspond to the relevant spatial relations in the world. Thus, reasoning with array representations can be thought of as a restricted form of model-theoretic deduction, one which is limited to inferences that are made explicit by inspection of array representations. The system can also be used to reason about the possible worlds resulting from uncertain or incomplete information. This section presents a semantics for deductive reasoning in such worlds.

4.1 Language and Semantics

A *language* $\mathcal{L}_w$ is introduced for a world w in order to formulate spatial propositions. $\mathcal{L}_w$ is a modal language employing predicate symbols P and constant symbols S. The compound formulas for the language are formed from the standard propositional operators, as well as modal operators, which will be used to reason about possibility and necessity in indeterminate worlds.

Definition. The **language** $\mathcal{L}_w$ for a world $w =< S, R >$ is defined as follows.

- Taking S as the set of *constant* symbols, and
- Taking P as the set of *predicate* symbols (Recall that the set of relations R in w, and the corresponding set $\mathcal{F}$ of array functions, are indexed by the set P.),
- The formulas in $\mathcal{L}_w$ consist of:
  - The *atomic formulas* $p(s_1, ..., s_n)$ where p is an n-ary predicate symbol in P and $s_1, ..., s_n$ are constant symbols in S.
  - The *compound formulas*, which are formed inductively from the atomic formulas by means of the unary operators $\neg$, $\Box$ and $\Diamond$ and the binary operators $\vee$ and $\wedge$.

Following, we present a possible worlds semantics for spatial reasoning based on a modal logic that accounts for the *necessity* and *possibility* of truth of a proposition. The notion of possible world is identified with an array representation that models an extensible, representable extension of a world. Truth of an atomic formula is determined by applying functions that inspect an array representation. A formula is necessarily true if it is true for all representations in the model. A formula is possibly true if there exists a representation for which it holds. In the following definition, a statement of the form $\mathcal{A} \models \phi$ denotes that the formula ϕ is true in array representation $\mathcal{A}$ for the model.

Definition. Let $\mathcal{L}_w$ be the language for a world w and let $< \mathcal{W}, \mathcal{F} >$ be a model for w. The truth of formulas in $\mathcal{L}_w$ in an array representation $\mathcal{A} \in \mathcal{W}$ is given as follows:

- For atomic formulas $p(s_1, ..., s_n) \in \mathcal{L}_w$,

  $\mathcal{A} \models p(s_1, ..., s_n)$ if and only if $f_p(s_1, ..., s_n, \mathcal{A}) = true$.

- For all compound formulas in $\mathcal{L}_w$:

  $\mathcal{A} \models \phi \wedge \psi$ if and only if $\mathcal{A} \models \phi$ and $\mathcal{A} \models \psi$.
  $\mathcal{A} \models \phi \vee \psi$ if and only if $\mathcal{A} \models \phi$ or $\mathcal{A} \models \phi$.
  $\mathcal{A} \models \neg\phi$ if and only if not $\mathcal{A} \models \phi$.
  $\mathcal{A} \models \Box\phi$ if and only if $\mathcal{A}' \models \phi$ for all $\mathcal{A}' \in \mathcal{W}$.
  $\mathcal{A} \models \Diamond\phi$ if and only if $\mathcal{A}' \models \phi$ for some $\mathcal{A}' \in \mathcal{W}$.

The modal logic belongs to the class $S5$ in the Lewis hierarchy, in fact, it extends this system since the truth of a modal formula is independent of the array representation. In particular, $\models \Box\phi$ if and only if $\mathcal{A} \models \Box\phi$ for any representation $\mathcal{A}$ in the model. Thus, questions of necessity (and possibility) refer to the whole model, not just to an individual representation. Consequently, any sequence of modalities is equivalent to the last one in the sequence. For example, $\Box\Box\Diamond\phi$ is equivalent to $\Diamond\phi$, and $\Diamond\Box\phi$ is equivalent to $\Box\phi$.

A proof system is *sound* with respect to a world if every formula that is valid is true for the world. Conversely, a system is *complete* with respect to a world if all true relations are provable. The properties of soundness and completeness are built into the definition of the model: an atomic formula is true in a representable world if and only if it holds for the array model of the world. In practical applications, we may wish to approximate a world rather than model it precisely, resulting in a system that could lead to incorrect conclusions. This does not imply that the inferencing process is unsound; rather, the knowledge used to represent the world is imprecise. Further discussions concerning the fidelity of reasoning with model-based representations have been presented by Goebel [25].

4.2 Strategies for Model-based Reasoning

Model-theoretic deductions can be achieved using a variety of control strategies. To prove the truth of a proposition, inference can be carried out as a three step process: 1) array representations are constructed to represent possible states of affairs in the world; 2) transformations are performed on the representations, corresponding to the transformations that occur in the world (this step is optional); and 3) conclusions are formed by applying inspection functions to the arrays. Alternative strategies can be developed for model-based reasoning, depending on the form of the desired conclusion. Deductions that involve the possibility of truth can be achieved by constructing a single array representation, corresponding to a possible world in which the premises and the given formula are all true. Similarly, proving a formula invalid requires the construction of a single array in which the premises are true and the putative conclusion is false.

Model-based reasoning, as an alternative to proof-theoretic deduction, has also been considered by others, including Halpern and Vardi [28]. In their work, an agent's knowledge is represented using a semantic model, where model checking is used to determine validity of a formula. For cases where the number of possible worlds grows exponentially, they suggest that heuristics could be used to focus attention on those worlds that are "most relevant" or "most likely".

Cognitive studies suggest that humans reason with a single model, even in situations that imply multiple states of affairs [34, 36]. If it is discovered that the current model does not correspond to the situation that is described then it is changed. A similar control strategy could be developed for a computational approach to model-based spatial reasoning, where a single array is used for reasoning and an alternative array is generated if the current one becomes inconsistent. Although our representation scheme was motivated by our understanding of cognitive processes, it can overcome some of the limitations of human information processing. Human errors occur in model-based deduction by failing to consider all possible interpretations compatible with a given set of facts [34]. In domains where the amount of indeterminacy is restricted, all possible array representations for a world can be generated, transformed and inspected. Thus, no consistent interpretations are left unconsidered. The array theory functions for spatial reasoning also facilitates parallel implementations [23]: multiple array representations can be transformed or inspected concurrently.

In summary, the proposed representation scheme for model-based reasoning provides an effective tool for performing inferences. Alternative control strategies can be constructed for carrying out deductions by generating, transforming and inspecting array representations. For cases where the number of array models is unmanageable, heuristic, parallel or backtracking strategies can be developed.

5 Nonmonotonic Reasoning

Proof systems such as first-order predicate logic were designed for monotonic reasoning: if new axioms are added to a system then everything that was previously derivable is still derivable. Many of the domains that involve spatial reasoning face the problems posed by uncertain or often changing knowledge where the property of monotonicity does not hold. A variety of logics have been developed in an attempt to accommodate nonmonotonic reasoning. These systems typically extend existing logics to include axioms and rules of inference that make it possible to reason with incomplete information. Reiter's default logic [56] allows inference rules of the form: *If X is provable and it is consistent to assume Y then conclude Z.* McDermott and Doyle [45] alternatively state defaults as sentences of the form: *If X holds and Y is not disprovable then Y.* Concepts such as "is consistent to assume" and "is not disprovable" can be related to the concept of "possibility" in our modal formalism. For example, a formula B is not disprovable in world w if B is true in some array representation in the model for w (that is, $\models \Diamond B$). Two issues that have to be addressed by nonmonotonic reasoning systems are:

- *How can inferences be made in the presence of incomplete knowledge?*
  In the previous section we presented a formalism for making inferences in the presence

of spatial indeterminacy. These inferences are achieved by constructing and inspecting symbolic arrays that model the alternative interpretations arising from uncertainty or incomplete information.

- *How is the knowledge base updated when new information is added?*
  A knowledge base for spatial reasoning can be defined as the set of array representations in the model for a world. In the remainder of this section, we address the question of how such a knowledge base is modified as the world is transformed by acquiring new knowledge or by modifying the existing spatial relations.

5.1 Knowledge Acquisition

In spatial reasoning systems, knowledge acquisition generally involves extending the spatial constraints for the world. Updating a model to accommodate such information is straightforward: the new world is modeled by eliminating from the knowledge base those representations that are inconsistent with the added information. Consider, for example, the indeterminate world described by the formulas *left-of(a,b)*, *left-of(a,c)*, *left-of(a,d)* and *left-of(b,d)*. This description implies three representable worlds, which can be modelled using the following arrays:

$$\underset{\mathcal{A}_1}{\begin{array}{|c|c|c|c|}\hline a & b & c & d \\ \hline\end{array}} \quad \underset{\mathcal{A}_2}{\begin{array}{|c|c|c|c|}\hline a & c & b & d \\ \hline\end{array}} \quad \underset{\mathcal{A}_3}{\begin{array}{|c|c|c|c|}\hline a & b & d & c \\ \hline\end{array}}$$

If the world is modified to include the spatial relationship *left-of(c,d)*, then the array representation corresponding to array $\mathcal{A}_3$ would be eliminated from the model, since it is not consistent with the added spatial relationship. Another way to view this is that the world that the array representation models is no longer an extension of the current world.

The above solution (and the current formalism) assumes that acquired knowledge does not extend the language for the world. Extending the theory to include adding information that involves new constant or predicate symbols is generally more complex. This situation could be addressed by reconstructing the model using the new language and including the updated information. It may also be possible to incrementally update a model by considering the current array representations for the world and modifying each of them individually. This can be achieved by extending the primitive *compose*[2] function to operate on indeterminate worlds. The composition of two array representations would result in a set (possibly empty) of array representations that corresponds to all of the consistent interpretations that are subsumed by the two arrays. Consider the array structure $\mathcal{A}_3$ above and a new spatial relationship *left-of(d,e)*. The function *compose* would be defined such that:

$$compose\left(\begin{array}{|c|c|c|c|}\hline a & b & d & c \\ \hline\end{array}, \begin{array}{|c|c|}\hline d & e \\ \hline\end{array} \right) = \begin{array}{|c|c|}\hline \begin{array}{|c|c|c|c|c|}\hline a & b & d & e & c \\ \hline\end{array} & \begin{array}{|c|c|c|c|c|}\hline a & b & d & c & e \\ \hline\end{array} \\ \hline\end{array}$$

[2]The complexity of an indeterminate version of the *compose* function would depend on the domain and the relations that are being preserved.

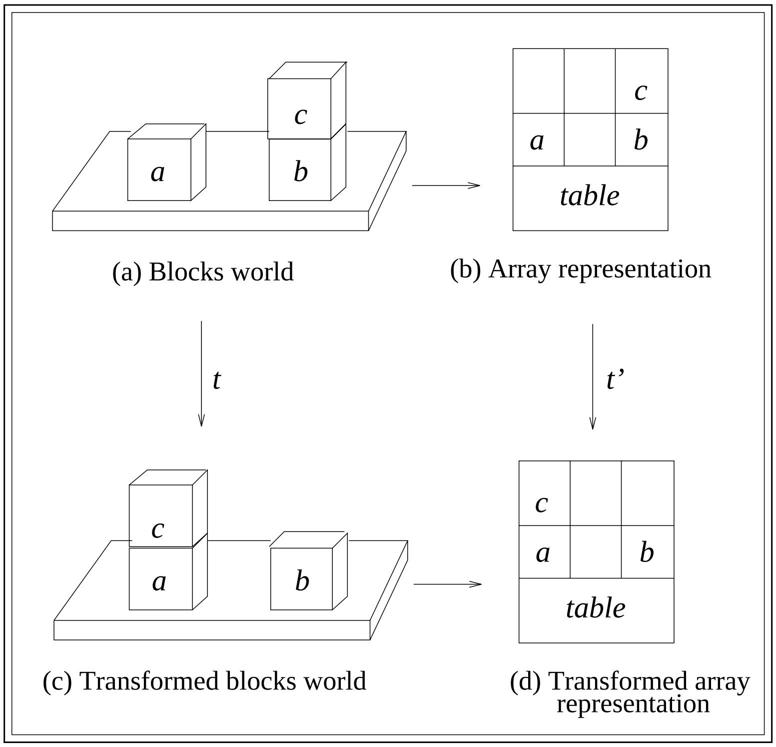

Figure 7: Representation of blocks world

5.2 Transforming a World

Spatial reasoning may involve applying transformations that result in changes to the relative locations of entities in the world. Reasoning in the presence of change is problematic in traditional reasoning systems, since it is necessary to consider the implications of change on the state of affairs. In the proposed scheme for model-based reasoning, however, the inferences arising from transformations on a world can be derived by applying analogous functions to the array representation: if t is a transformation that can be applied to a world w resulting in a world w', then we define an analogous array theory function t' such that if t' is applied to each of the array representations for w it would result in a set of arrays that represent the possible worlds for w'.

To illustrate how the effects of change can be modeled, consider the blocks world in Figure 7(a) and the corresponding array representation $\mathcal{A}$ in Figure 7(b). The blocks world resulting from the transformation t = *Move block c to the top of block a* is illustrated in Figure 7(c). This change is modeled in the array representation by applying the primitive operation $t' = move\_rel$ to the parameter list $(c, a, \mathcal{A}, above)$. The function application results in an array that represents the transformed world, as depicted in Figure 7(d). In the case where a transformation is applied to an indeterminate world, the array operation is applied to the array representation for each of the possible worlds. Array transformation functions may be complex and involve knowledge of the physical model for the entities in the world. For example, the transformation operation for *push* in the blocks world would have to take into account that if the block being pushed is supporting other blocks, then the locations of the supported blocks are also changed by the transformation.

Modifying spatial models by adding knowledge or by applying transformations results in a new model, which can subsequently be used for reasoning about validity. Note that it is not necessary to examine any previous deductions to determine whether relationships need to be deleted from the knowledge base, since the modified spatial relations are determined directly from inspection of the transformed array representations. Thus, the model-based approach to spatial reasoning addresses the *frame problem* [55], which is concerned with what relations are withdrawn or remain valid as change occurs in a world. This information is implicit in the transformation functions for the array. In the next section, we discuss how an ordered sequence of transformed representations can be used to model temporal reasoning in the theory.

6 Extensions to the Model

The formalism proposed for model-based reasoning was designed to capture and reason about the relevant spatial and structural qualities of a world. In the previous sections it was demonstrated that the formalism is applicable to indeterminate worlds and to nonmonotonic reasoning. In this section we introduce some potential extensions to the representation scheme. One possible extension is to consider worlds that are hierarchically structured and allow for reasoning at multiple levels of the decomposition hierarchy. This is achieved by assuming that an array model $\mathcal{A}$ for a world w is a recursive data structure, in which symbols in $Sym(\mathcal{A})$ may denote models of subworlds of w (i.e., worlds at a more detailed level of parts abstraction). Extensions to the array representation for temporal and analogical reasoning, and for viewing a model from internal perspectives are also considered. Our presentation of these extensions does not include a rigorous reformulation of the existing semantics; rather, it is meant to provide motivation for future research in the area.

6.1 Hierarchical Models

Results of cognitive studies suggest that mental models may be hierarchically organized and that reasoning takes place at varying levels of structural decomposition based on a *part-of* relation [46]. For example, when planning a route for a European vacation one might first consider a spatial model of the countries to be visited then later focus in on the details (regions or cities) of the individual countries. Similarly, when experts analyze the spatial structure of a protein molecule they generally begin by considering entities and relations at the level of secondary or tertiary structure, then subsequently attend to more detailed models containing entities such as residues or atoms. Our formalism allows for reasoning at multiple levels of parts decomposition through the use of nested array data structures: array symbols may denote subarrays that correspond to the subworlds for the structured entities in the world. Figure 8 illustrates a modified representation of the previously presented array for Europe where the symbol *Britain* has been replaced by a subarray that models the world corresponding to this symbol.

A hierarchical representation explicitly depicts relations among entities at multiple levels of the decomposition hierarchy, without having to specify inheritance laws. For example,

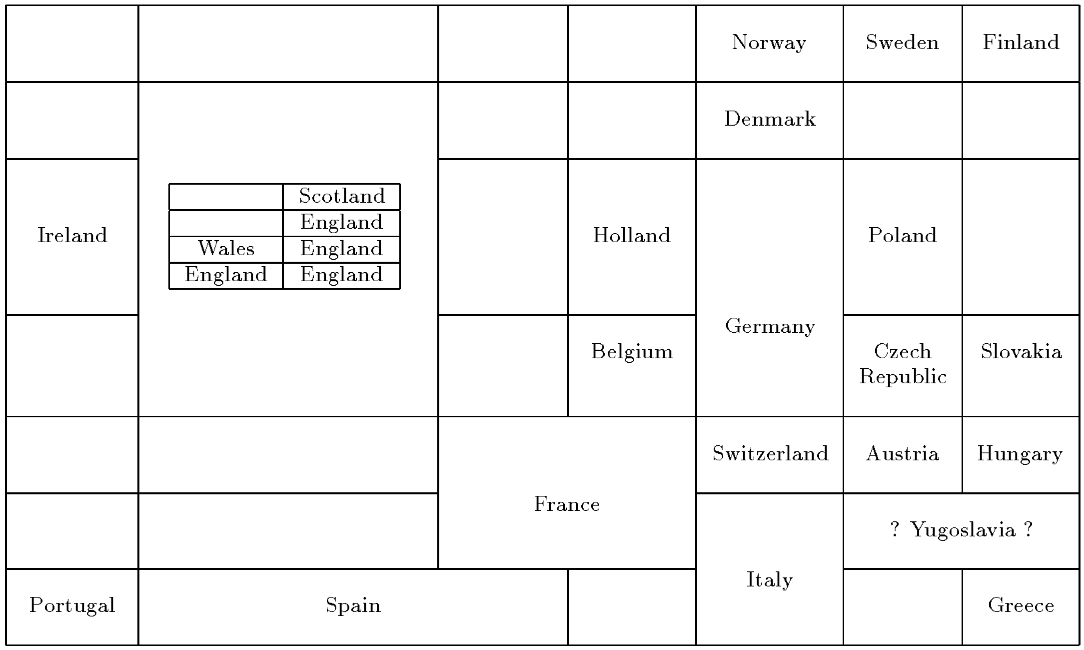

Figure 8: Embedded array representation of Europe

we can compute that England is *west-of* Holland in the embedded array representation of Figure 8 by extending the interpretation of the relation *west-of* to range over nested data structures. More precisely, a hierarchical world $w =< S, R >$ is a world in which the set S may denote structured entities. For all $s_i \in S$, if s_i denotes a structured entity, then there exists a world $w_{s_i} =< S_{s_i}, R_{s_i} >$ such that:

- $S_{s_i} \subseteq S$ denotes the set of subentities (parts) of the entity denoted by symbol s_i; and
- R_{s_i} is the set of spatial relations for the subworld denoted by s_i.

We assume that each subworld w_{s_i} has a model consisting of array representations. In a hierarchical world w, the set of spatial relations R would contain structural relations corresponding to the part-of hierarchy for the world, that is, *part-of*(s_1, s_2) is specified for all constant symbols s_1 and s_2 such that s_1 denotes a subentity of a structured entity denoted by s_2.

The primitive *focus* function (see Table 1) can be used to transform an array representation by replacing a constant symbol corresponding to a structured entity by the representation for the subworld corresponding to the symbol. Figure 9 depicts the successive arrays resulting from iteratively applying the *focus* function to an array model for North America.

The choice of a decomposition hierarchy for subworlds is dependent on the domain and the task demands. It is interesting to note that certain decompositions, combined with course-grained representations, lead to errors analogous to those displayed by humans (see Section 2.2). For example, the hierarchy displayed in Figure 9 could be used to deduce the false conclusion that the city Seattle, which is situated in the USA, is *south-of* the Canadian

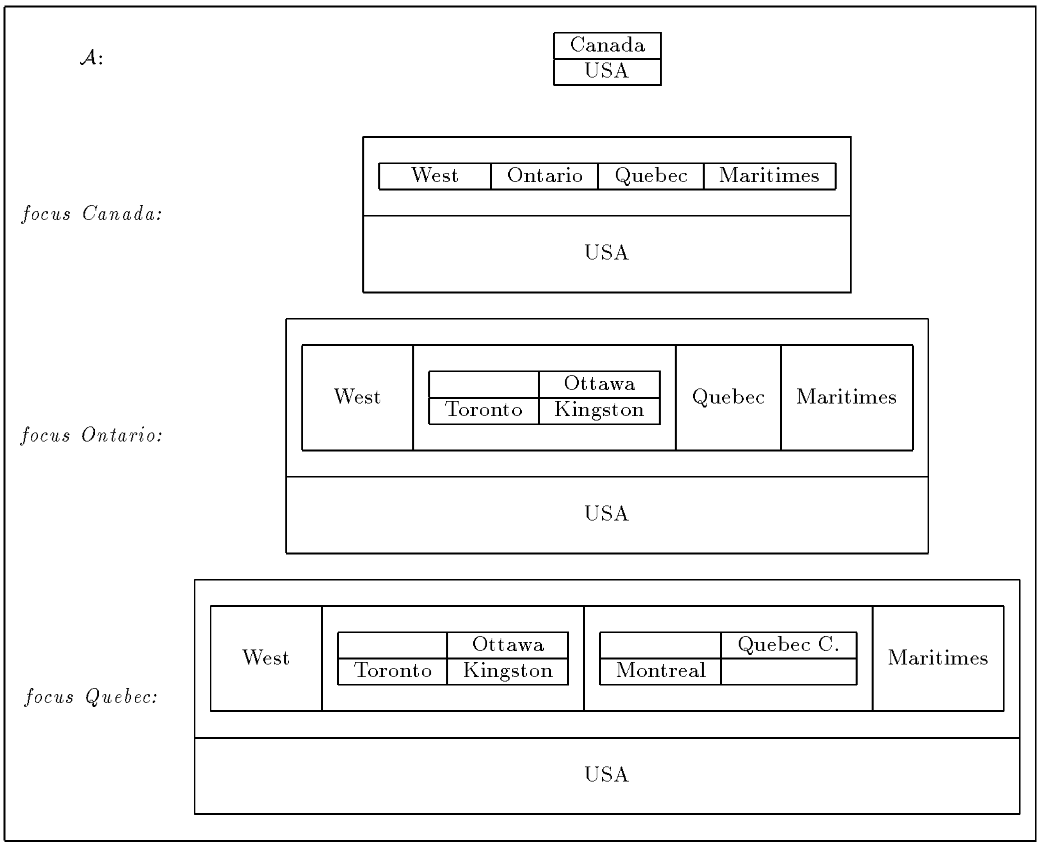

Figure 9: Successive applications of *focus* function

Vector	Orientation
(-1,0)	$\uparrow$
(-1,-1)	$\nwarrow$
(0,-1)	$\leftarrow$
(1,-1)	$\swarrow$
(1,0)	$\downarrow$
(1,1)	$\searrow$
(0,1)	$\rightarrow$
(-1,1)	$\nearrow$

Figure 10: Orientation vectors for a two-dimensional array

city Montreal. Such erroneous conclusions are avoidable by designing different decomposition hierarchies, or by using finer-grained array representations.

6.2 Multiple Views

Cognitive studies suggest that humans have the ability to inspect models from multiple perspectives, including alternative internal perspectives – sometimes referred to as reasoning in *egocentric space*. Different external views can be achieved in our scheme by applying rotational transformations to a multi-dimensional array.[3] For problem domains such as robot navigation, it is useful to reason about the world from an internal perspective, where inferences concerning relative locations depend on a current frame of reference. This can be achieved by extending the formalism to include functions for inspecting the array from a particular reference location and orientation. The orientation of a reference value in n-dimensional space can be specified as an n-vector. For two-dimensional arrays, the orientation vector takes on one of eight possible values, corresponding to reorientations by angles that are multiples of 45 degrees (see Figure 10). There are 27 possible orientation vectors for three-dimensional arrays.

Following is a description of a convention center used in experiments by Taylor and Tversky [67]:

> The entrance is on the east side of the building. As you enter, there is a water fountain on your left, and beyond it a bulletin board. As you walk down the corridor in front of you, you pass movie cameras on your right, then 35mm cameras. On your left is the office. As you reach the end of the corridor, the restrooms are directly ahead of you, side by side. Turn right and continue walking; the cafeteria will be on your left. Turn right again at the end of the corridor; the CD players will be just ahead on your left, and the televisions on your right. Farther up, the VCRs are on the right and stereo components on your left. Turn right at the end of the hallway; you will pass personal computers on your left hand and then find yourself back at the entrance.

[3] Only rotations that are multiples of 90 degrees can be applied to the axes of the array using the current implementation of the primitive *rotate* function.

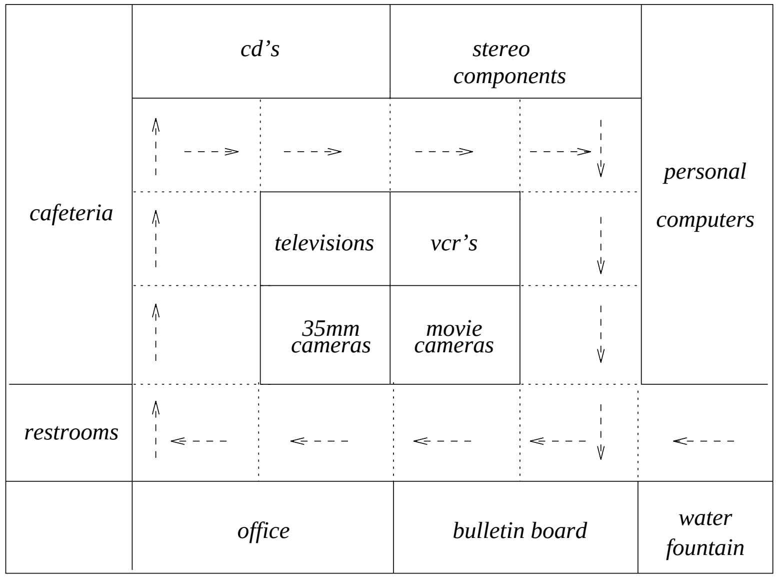

Figure 11: Array representation of convention center

The world described in this narrative is representable as the array depicted in Figure 11, where the dotted arrows denote the locations visited and the corresponding orientation vectors. Given such an array representation, inferences can be made with regards to the current location and perspective of the frame of reference. For example, a query of the form *"If you are facing entity s what is behind you?"* could be resolved by: 1) transforming the array by moving the reference point to an empty location adjacent to the symbol s; 2) updating the orientation vector to be directed towards the symbol s; and 3) applying the primitive *find* function with the parameter *behind* to determine the entities located behind (in the opposite direction from the orientation) the reference location. As will be discussed in the next section, movement within a world can also be modeled in a temporal framework.

6.3 Temporal Models

As stated earlier, temporal order relations can be modeled using a single dimension in the array representation. By varying the time dimension in such a representation, it would be possible to track the transformations that occur in the world and reason in terms of relations such as *before, during* and *after*. Johnson-Laird [34] defines a *temporal model* as a sequence of spatial models. In this spirit, temporal models can be represented in an extension to the proposed theory as one-dimensional arrays, where each element in the array denotes a discrete snapshot of a determinate world. Relative locations in the array would depict the temporal ordering of worlds: if array $\mathcal{A}_1$ is *left-of* array $\mathcal{A}_2$ then the world represented by $\mathcal{A}_1$ existed before the world represented by $\mathcal{A}_2$. Thus, array inspection operations can be performed to reason about $before/after$ relations among determinate worlds. Figure 12 illustrates three time steps resulting from transforming an array representation using primitive operations corresponding to *turn right* then *push desk*, where the symbol "$\rightarrow$" denotes the current

bed		closet	
bed			chair
	↑	desk	
	table	table	

Time = 1

bed		closet	
bed			chair
	→	desk	
	table	table	

Time = 2

bed		closet	
bed			chair
		→	desk
	table	table	

Time = 3

Figure 12: Temporal model for spatial reasoning

reference location and orientation in two-dimensional space.[4] Using this representation we can infer information such as: the table is behind you *before* the table is to your right. If the representation scheme for temporal reasoning stores the transformations that are applied at each time step, then the formalism could also be used to reason about *causality*. For example, the change in state reflected in moving from the world represented at Time=1 to the world represented at Time=2 was caused by the transformation *turn right*. Array theory operations could also be defined to compare two arrays to analyze the change that results from a transformation. Homogeneous approaches to representing time and space have been previously proposed by others, including Guesgen and Hertzberg [26] who extend Allen's temporal logic [1] with constraint satisfaction algorithms for spatiotemporal reasoning.

It has been suggested that the processing of transformations on mental models may occur in parallel. Ullman [73] has proposed the concept of *spatial parallelism*, which corresponds to the the same operations being applied concurrently to different locations in a representation, and *functional parallelism*, which occurs when different operations are applied simultaneously to the same location. These forms of parallelism can be achieved using the primitive functions of array theory, such as the second-order function *EACH*, which applies an operation to all of the arguments of an array. A function application *EACH* f $\mathcal{A}$ results in the operation f being applied (potentially in parallel) to each symbol in the array $\mathcal{A}$.

6.4 Analogical Reasoning

Model-based reasoning is not restricted to the application of deductive inferences; inductive and analogical reasoning can also be performed using array representations. *Analogical reasoning* involves the transfer of knowledge about a known world (often referred to as the *source* domain) to a world that is to be explained (referred to as the *target* domain). The computation of similarity is a central process of analogical reasoning; for spatial analogy this implies determining structural or spatial similarity. The spatial similarity between two array

[4] Note that this could also be modelled as a three-dimensional array where the third dimension denotes temporal ordering.

models can be measured in terms of the symbols and relations that they have in common. It can alternatively be measured in terms of the primitive transformations needed to bring the two models into equivalence.[5] Many types of transformations are possible, such as moving symbols, changing the orientation of a part, deleting parts, etc.[6] Given the three array representations:

	cup	
fork		*spoon*

$\mathcal{A}_1$

	fork
spoon	*cup*

$\mathcal{A}_2$

	cup
fork	*spoon*

$\mathcal{A}_3$

and the set of predefined array inspection functions $\mathcal{F}=\{f_{left\text{-}of}, f_{above}\}$, we could say that array $\mathcal{A}_1$ is more similar to $\mathcal{A}_3$ than it is to $\mathcal{A}_2$ since the application of a single transformation, $move\_rel(cup, spoon, \mathcal{A}, above)$ would result in an array equivalent to $\mathcal{A}_3$. No such simple transformation would bring $\mathcal{A}_1$ and $\mathcal{A}_2$ into equivalence.

Research in analogical reasoning often goes under the name of *case-based reasoning* [57]. A case-based reasoning system retrieves and adapts previous experiences in order to derive or criticize a solution to a new problem. An approach to molecular scene analysis that incorporates array representations for stored cases of determined scenes is currently being considered [20]. In this system, cases are retrieved and compared with a novel scene containing unidentified parts. Spatial analogies are of key importance in guiding the search towards a fully reconstructed model of a scene, since they can be used to anticipate the identity of the uninterpreted entities.

An *inductive inference* is one that proceeds from specific examples to generalizations. The I-MEM (for Image MEMory) system [8] is a spatial concept formation system that has been developed as a framework for efficient analogical classification of images using induction. The main premise of induction in the I-MEM system is that generalization is based on entity or relation deletion from a spatial representation. This approach to generalization can be defined for worlds and their models. Informally, a world w is a generalization of a world w' if and only if w is extensible to w', that is $w \preceq w'$. This notion of generalization can be implemented using models for w and w'. In particular, w is a generalization of w' if and only if all representations $\mathcal{A}'$ in the model for w', there exists an array representation $\mathcal{A}$ in the model for w such that $\mathcal{A}$ is equivalent to $\mathcal{A}'$.

The ability to determine spatial similarity and to carry out inductive and analogical reasoning is particularly critical for recognition, classification and learning tasks [74]. The I-MEM system has been applied to the discovery of three-dimensional spatial concepts in molecular databases [6] and to the classification of molecular motifs (reoccurring structures with similar spatial relations among entities) [7]. Thagard [69] has also considered analogy-based reasoning in the development of a system that uses symbolic array representations to help explain Dalton's atomic theory. Others have argued for the central role of imagery and models in scientific problem solving and have called for research on computational approaches to spatial analogy [47, 72].

[5] Two arrays are equivalent if they represent the same spatial relations.

[6] See Table 1 for a description of the currently implemented transformation functions.

7 Related Research

The work described in this paper extends and enhances research in computational imagery [24] in a number of ways:

- The spatial representation for imagery is related to cognitive studies of mental models.
- An array model is defined in terms of a mapping from from spatial functions in the array representation to spatial relations in the world.
- The concept of a model, consisting of a set of array representations, is defined for indeterminate worlds.
- A possible worlds semantics for model-based reasoning using array representations is presented.
- Nonmonotonic reasoning is discussed in relation to array representations.
- Potential extensions to the formalism – for reasoning with hierarchical models, for temporal and analogical reasoning, and for inspecting representations from multiple perspectives – are introduced.

The concept of constructing knowledge representations that mirror the structure of the world is not unique to the array models described in this paper. Hayes [29] discusses *direct* representations in which there exist similarities between what is being represented and the medium of the representation. Sloman [62] has also argued the pros and cons of analogical representations, and has concluded that a variety of representation formalisms – including those specialized for spatial reasoning – are important to AI problem solving [63]. Other hybrid approaches have been suggested for visual-spatial and model-based reasoning. Barwise and Etchemendy [2] have proposed a system called *Hyperproof* which integrates diagrammatic reasoning with sentence-based logics. Hyperproof uses both diagrams and logic notation to teach students how to reason logically. In subsequent work, Barwise and Etchemendy [3] present a formal semantics for reasoning with Hyperproof diagrams. Habel and colleagues [27] have developed a hybrid system consisting of a propositional and depictorial partonomy (organization of parts) for reasoning, where the depictorial partonomy reflects the hierarchy proposed in representations for visual processes. They suggest that the advantage of the depictorial representation in their system is that it facilitates an efficient attention-driven method for reasoning. Myers and Konolige [49] treat model-based manipulations as a form of inference within a classical logic system. More specifically, they store partially interpreted sensor data using an analogical representation that interacts with a general-purpose sentential language. A similar approach has been taken by Chandrasekaran and Narayanan [5], who have proposed an architecture where analogical representations derived from visual perception are used in combination with symbolic (propositional) representations. A technique for qualitative spatial reasoning, based on the directional orientation information made available through perceptual processes, has been presented by Freksa and Zimmermann [14]. In this work, orientations in two-dimensional space are defined by the relation between a vector and a point.

Visual-spatial reasoning techniques have also been considered in the context of specific application domains. Funt [15] represents and manipulates a visual analog of a world in order to predict potential instabilities and collisions in a physical domain. Several others have applied diagrams or analogical representations to qualitative physics problems [12, 16, 27, 51]. For the domain of route planning, Kuipers [40] has developed a program that determines a path between points by considering a hierarchical network of region representations. McDermott and Davis [44] describe a more general representation for route planning that stores the shapes and locations of entities in the world. Facts in this system are represented as propositions and spatial reasoning is carried out by special-purpose theorem proving modules. Other problem domains where diagrammatic reasoning has been applied include biology, architecture, geometry and theorem proving [50].

Research in geographical information systems and spatial databases has long been concerned with the issue of representing spatial knowledge. Samet [58] has proposed a method for storing geographic knowledge based on the recursive decomposition of space. In this work, the term *quadtree* is used to describe binary array data structures that iteratively subdivide regions into segments until blocks are obtained that consist entirely of 1s or entirely of 0s. These structures (and their three-dimensional counterpart, termed *octrees*) are efficiently stored and implemented as trees, where each node of the tree corresponds to a region in the decomposition hierarchy. Samet's data structures and algorithms for querying spatial data bases could potentially aid in the development of efficient implementations for the knowledge representation scheme proposed in Section 3. An alternative approach to reasoning in geographic systems has been described by Papadias and Sellis [53]. In their work, a symbolic two-dimensional array structure is used to preserve a set of spatial relations among geographic entities. Their approach is similar to a model for geographic information systems based on the array representation scheme proposed in this paper [19].

Spatial representations have also been considered in machine vision research. According to Biederman [4], the representation of objects can be constructed as a spatial organization of simple primitive volumes, called *geons*. The process of image analysis, as defined by Marr and colleagues [43], depends on a series of representations culminating in a three-dimensional model of the spatial relations among entities which makes explicit *what* is *where*. As in Marr's approach to computational vision, *molecular scene analysis* [13] is concerned with discovering what is present in the world and where it is spatially located. The act of determining the structure of a molecule is an interactive process consisting of a state space search of partially interpreted scenes, which can be represented and evaluated as three-dimensional symbolic array models [21].

Possible worlds semantics have been used in other areas of AI research to express uncertainty and deal with nonmonotonicity. Epistemic logic was proposed by Hintakka [30] to reason about the knowledge and belief of agents in the world: an agent *knows* a formula ϕ exactly if ϕ is true in all the worlds considered possible for the agent. A widely accepted semantics for modal logics is one introduced by Kripke [39]. In a Kripke model, truth is based on a binary accessibility relation between possible worlds. When applied to temporal logics, possible worlds are equated with points in time and accessibility is defined in terms of the linear order of time points. Thus, a formula is necessarily true for a possible world in a temporal logic if it is true in all future (accessible) possible worlds. In the semantics presented in this paper, all possible worlds are accessibly to one another, implying that the necessity and possibility of a formula is the same for all array representations in a model.

Although interest and activity in spatial and diagrammatic reasoning is escalating, research in this area has focussed on logical or analogical representations, corresponding to the propositional and visual representations proposed by Johnson-Laird [34]. What the array formalism offers is an intermediate representation which is less specific than a visual representation (it may discard visual details), yet less abstract than a logic representation (spatial information is made explicit). The choice of representation depends on the demands of a particular problem; array models are not appropriate for all spatial tasks. They are suitable, however, for reasoning in spatial domains where inferences depend on determining the relative location and topology of entities in the world. In the worst case, the complexity of such inferences are limited by the size of the array being inspected.[7]

8 Conclusions

A characteristic of the array formalism for model-based reasoning is that it brings relevant spatial properties to the forefront. The entities and spatial relations in the world are explicitly denoted as symbols and relations in a multi-dimensional array. This representation provides for a simplified model of the world – one that captures salient spatial features and suppresses unnecessary or irrelevant details. An advantage of the array representation lies in its succinct encoding and its provision for updating and change. It can also be distinguished from traditional logic representations by the fact that it imposes specificity on a representation, yet symbolic arrays are more abstract than the analogical representations that have been developed for diagrammatic reasoning.

The described knowledge representation scheme provides a complete and sound system that can perform under conditions involving uncertain or incomplete information. Model-theoretic reasoning is used to make inferences about determinate or indeterminate worlds, using a three-step process of generating, transforming and inspecting array representations for the world. Thus, the process of constructing syntactic proofs to derive spatial information is replaced by model checking. The non-existence of a proof can be determined by the existence of an array representation for the model in which the formula is refuted. Alternatively, the truth of a proposition can be verified by the inability to construct a representation that refutes the conclusion. The representation scheme provides a framework for integrating model-theoretic deduction with nonmonotonic reasoning in which representations are updated and reinterpreted as new information is acquired or as transformations are performed.

The array-based formalism for spatial reasoning has measurable computational advantages. In particular, they can be used to develop *vivid* knowledge bases. Levesque [41] defines a vivid knowledge base as one that is structured so that there is a one-to-one correspondence between the entities in the world and the symbols in the knowledge base, and for each simple relationship of interest in the world – in our case spatial relationships – there exists a corre-

[7]For a more detailed discussion of the computational advantages of symbolic array representations over traditional knowledge representation schemes, see [18].

sponding connection among symbols in the knowledge base. Levesque argues that the main advantage of vivid knowledge bases is that they provide for efficient worst case reasoning behavior, since calculating what is logically implicit generally reduces to retrieving what is explicit.

The model-based approach to reasoning can be motivated and justified by human needs. Simon [61] has suggested criteria for assessing and selecting representations based on information content and on ease of programming. These criteria are task dependent and partially rely on the ability of the programmer to represent the state of knowledge in the world and the transformations and inferences that may occur. Experimental results in cognitive psychology suggest that humans apply model-based reasoning for problem solving in a variety of domains. Consequently, a formalism that captures the representations and processes associated with model-based reasoning would facilitate the implementation of computational reasoning systems in such problem solving domains. In particular, such systems could incorporate the intuitions and heuristics that are applied to spatial models. Although the scheme can be motivated by human needs, it overcomes the inherent limitations of the cognitive system. In particular, control strategies can be developed so that no consistent interpretation for a world is overlooked.

In a recent debate concerning the advantages/disadvantages of descriptive versus depictive (model-based) representations, Levesque and Reiter [42] state that a reason to prefer descriptive (logic) representations is that they are "blessed with a *semantics*". Although proof-theoretic representations can be advocated for their semantic clarity, we have demonstrated in this paper that an intuitive semantics for model-based reasoning with array representations also exists. Furthermore, as pointed out by Stenning [64], the model-based deductions should be considered as inferences in a weak logical system that facilitates "easy but limited inferences". Although some strategies for spatial reasoning with the array models were discussed, a fully operational model-based reasoning system and its application in a variety of domains is a subject of ongoing research.

In summary, the array representation scheme provides an effective and efficient means for performing spatial reasoning. Extensions to the scheme – for temporal, inductive and analogical reasoning, for hierarchical worlds and for egocentric reasoning – are also proposed. Further research is required, however, to develop these extensions and to fully realize the applicability of the scheme to a variety of domains.

Acknowledgements: The research described in this paper was supported by grants from the Natural Science and Engineering Research Council (NSERC) of Canada and the Information Technology Research Corporation (ITRC) of Ontario.

References

[1] J.F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26:21–36, 1983.

[2] J. Barwise and J. Etchemendy. Hyperproof: Logical reasoning with diagrams. In *Working Notes of the AAAI Spring Symposium on Reasoning with Diagrammatic Representations*, 1992.

[3] J. Barwise and J. Ethemendy. Heterogeneous logic. In *Working notes of the 1993 IJCAI Workshop on Principles of Hybrid Representation and Reasoning*, 1993.

[4] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–147, 1987.

[5] B. Chandrasekaran and N.H. Narayanan. Towards a theory of commonsense visual reasoning. In K.F. Nori and C.E. Veni Madhavan, editors, *Foundations of Software Technology and Theoretical Science*, pages 388–409. Springer-Verlag: Berlin, 1990.

[6] D. Conklin, S. Fortier, and J.I. Glasgow. Knowledge discovery in molecular databases. *IEEE Transactions on Knowledge and Data Engineering*, pages 985–987, 1993. Special Issue on Learning and Discovery in Knowledge-Based Databases.

[7] D. Conklin, S. Fortier, and J.I. Glasgow. Representation for discovery of protein motifs. In D. Searls L. Hunter and J. Shavlik, editors, *Proceedings of the First International Conference on Intelligent Systems for Molecular Biology*. AAAI/MIT Press, Menlo Park CA., 1993.

[8] D. Conklin and J.I. Glasgow. Spatial analogy and subsumption. In Sleeman and Edwards, editors, *Machine Learning: Proceedings of the Ninth International Conference ML(92)*, pages 111–116. Morgan Kaufmann, 1992.

[9] J. de Kleer and J.S. Brown. Assumptions and ambiguities in mechanistic mental models. In A. Collins and E.E. Smith, editors, *Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence*, pages 270–287. Morgan Kaufmann, 1988.

[10] M. Denis. *Image and Cognition*. Harvester Wheatsheaf, London, 1991.

[11] M. Farah, K. Hammond, D. Levine, and R. Calvanio. Visual and spatial mental imagery: Dissociable systems of representation. *Cognitive Psychology*, 20:439–462, 1988.

[12] K.D. Forbus. Qualitative reasoning about space and motion. In D. Gentner and A.L. Stevens, editors, *Mental Models*, pages 53–74. Lawrence Erlbaum Assoc., Hillsdale, NJ, 1983.

[13] S. Fortier, I. Castleden, J. Glasgow, D. Conklin, C. Walmsley, L. Leherte, and F. Allen. Molecular scene analysis: The integration of direct methods and artificial intelligence strategies for solving protein crystal structures. *Acta Crystallographica*, D1, 1993.

[14] C. Freksa and K. Zimmermann. On the unilization of spatial structures for cognitively plausible and efficient reasoning. In *Proceedings of the IJCAI '93 Workshop on Spatial and Temporal Reasoning*, pages 61–66, 1993.

[15] B.V. Funt. Problem-solving with diagrammatic representations. *Artificial Intelligence*, 13:201 – 230, 1980.

[16] F. Gardin and B. Meltzer. Analogical representations of niave physics. *Artificial Intelligence*, 38:139–159, 1989.

[17] J.I. Glasgow. Imagery and AI – where do we go from here? *Computational Intelligence*, 9(4):424–435, 1993. Response to taking issue forum.

[18] J.I. Glasgow. The imagery debate revisited: A computational perspective. *Computational Intelligence*, 9(4):309–333, 1993. Taking issue paper.

[19] J.I. Glasgow. Representation of spatial models for geographic information systems. In N. Pissinou, editor, *Proceedings of the ACM Workshop on Advances in Geographic Information Systems*, pages 112–117, Arlington, VA, 1993.

[20] J.I. Glasgow, D. Conklin, and S. Fortier. Case-based reasoning for molecular scene analysis. In *Working Notes of the AAAI Spring Symposium on Case-Based Reasoning and Information Retrieval*, pages 53–62, Palo Alto, California, 1993.

[21] J.I. Glasgow, S. Fortier, and F.H. Allen. Molecular scene analysis: crystal structure determination through imagery. In L. Hunter, editor, *Artificial Intelligence and Molecular Biology*. AAAI Press, 1993.

[22] J.I. Glasgow, M.A. Jenkins, E. Blevis, and M. Feret. Logic programming with arrays. *IEEE Transactions on Data and Knowledge Engineering*, 3(3):307–319, 1991.

[23] J.I. Glasgow, M.A. Jenkins, C. McCrosky, and H. Meijer. Expressing parallel algorithms in Nial. *Parallel Computing*, 11(3):46–55, 1989.

[24] J.I. Glasgow and D. Papadias. Computational imagery. *Cognitive Science*, 16(3):355–394, 1992.

[25] R. Goebel. Spatial and visual reasoning: Human versus logic. *Computational Intelligence*, 9(4):406–412, 1993.

[26] H.W. Guesgen. A constraint-based approach to spatiotermporal reasoning. *Journal of Applied Intelligence*, 3:71–90, 1993.

[27] C. Habel, S. Pribbenow, and G. Simmons. Propositional and depictorial partonomies: A hybrid approach. In *Proceedings of IJCAI '93 Workshop on Principles of Hybrid Representation and Reasoning*, pages 62–71, 1993.

[28] J.H. Halpern and M.Y. Vardi. Model checking vs. theorem proving: a manifesto. In J.A. Allen, R. Fikes, and E. Sandewall, editors, *Principles of Knowledge Representation and Reasoning: Proceeding of the Second International Conference (KR '91)*, pages 325–334, 1991.

[29] P. Hayes. Some problems and non-problems in representation theory. In *Proceedings of AISB Summer Conference*, pages 63–79, University of Sussex, 1974.

[30] J. Hintikka. *Knowledge and Belief: An introduction to the logic of the two notions.* Cornell University Press, Ithaca, NY, 1962.

[31] J. Huttenlocher. Constructing spatial images: A strategy in reasoning. *Psychological Review*, pages 550 – 560, 1968.

[32] M.A. Jenkins and J.I. Glasgow. A logical basis for nested array data structures. *Programming Languages Journal*, 14(1):35 – 49, 1989.

[33] M.A. Jenkins, J.I. Glasgow, and C. McCrosky. Programming styles in Nial. *IEEE Software*, 86:46 – 55, 1986.

[34] P.N. Johnson-Laird. *Mental Models.* Harvard University Press: Cambridge, Mass., 1983.

[35] P.N. Johnson-Laird. *Human and Machine Thinking.* Lawrence Erlbaum Assoc: Hillsdale, NJ, 1993.

[36] P.N. Johnson-Laird and R.M.J. Byrne. *Deduction.* Lawrence Erlbaum Assoc: Hillsdale, NJ, 1991.

[37] S.M. Kosslyn. *Image and Mind.* Harvard University Press, 1980.

[38] S.M. Kosslyn. *Ghosts in the Mind's Machine: Creating and using images in the brain.* W.W. Norton: New York, 1983.

[39] S. Kripke. Semantical considerations of modal logic. *Acta Philosophica Fennica*, (16):83–94, 1963.

[40] B. Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–154, 1978.

[41] H.J. Levesque. Making believers out of computers. *Artificial Intelligence*, 30:81–108, 1986.

[42] H.J Levesque and R. Reiter. A counterexample is worth a thousand words. *Computational Intelligence*, 9(4):394–397, 1993. Commentary on taking issue paper.

[43] D. Marr. *Vision.* W.H. Freeman and Company: San Francisco, 1982.

[44] D.V. McDermott and E. Davis. Planning routes through uncertain territory. *Artificial Intelligence*, 22, 1984.

[45] D.V. McDermott and J. Doyle. Non-monotonic logic. *Artificial Intelligence*, 13:41–72, 1980.

[46] T.P. McNamara. Mental representation of spatial relations. *Cognitive Psychology*, 18:87 – 121, 1986.

[47] A.I. Miller. *Imagery in scientific thought.* MIT Press: Cambridge, Mass., 1986.

[48] T. More. Notes on the diagrams, logic and operations of array theory. In Bjorke and Franksen, editors, *Structures and Operations in Engineering and Management Systems.* Tapir Pub., Norway, 1981.

[49] K. Myers and K Konolige. Reasoning with analogical representations. In *Proceedings of the Conference on Principles of Knowledge Representations and Reasoning*, Los Altos, CA, 1992. Morgan Kaufmann.

[50] N. H. Narayanan, editor. *Working notes of the AAAI Spring Symposium on Reasoning with Diagrammatic Representations.* AAAI Press, Stanford University, Palo Alto, CA., 1992.

[51] N.H. Narayanan and B. Chandrasekaran. A computer model of diagrammatic reasoning. In *Working Notes of the AAAI Spring Symposium on Reasoning with Diagrammatic Representations*, pages 136–141, 1992.

[52] D. Papadias and J.I. Glasgow. A knowledge representation scheme for computational imagery. In *Proceedings of the Thirteenth Annual Meeting of the Cognitive Science Society*, pages 48–54, Chicago, 1991. Lawrence Erlbaum Associates: Hillsdale, NJ.

[53] D. Papadias and T. Sellis. The semantics of relations in 2d space using representative points: Spatial indices. In A.U. Frank and I. Campari, editors, *Spatial Information Theory: A theretical basis for GIS. Proceedings of the European Conference on Spatial Information Theory (COSIT '93)*, pages 234–247. Springer-Verlag Lecture Notes on Computer Science, volume 716, 1993.

[54] Z.W. Pylyshyn. The imagery debate: Analog media versus tacit knowledge. In N. Block, editor, *Imagery*, pages 151 – 206. MIT Press: Cambridge, Mass., 1981.

[55] B. Raphael. The frame problem in problem-solving systems. In Findler and Meltzer, editors, *Artificial Intelligence and Heuristic Programming*, pages 159 – 169. Edinburgh University Press, 1971.

[56] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13, 1980.

[57] C. Riesbeck and R. Schank. *Inside case-based reasoning.* Lawrence Erlbaum: Hillsdale, N.J,, 1989.

[58] H. Samet. *The Design and Analysis of Spatial Data Structures.* Addison-Wesley, 1989.

[59] P. Shaver, L. Pierson, and S. Lang. Converging evidence for the functional significance of imagery in problem solving. *Cognition*, 3:359–375, 1974.

[60] R.N. Shepard and L.A. Cooper. *Mental Images and their Transformations.* MIT Press: Cambridge, Mass., 1982.

[61] H.A. Simon. On the forms of mental representations. In W.C. Savage, editor, *Minnesota Studies in the Philosophy of Science, Vol. IX: Perception and Cognition: Issues in the Foundations of Psychology*. University of Minnesota Press: Minneapolis, 1978.

[62] A. Sloman. Afterthoughts on analogical representation. In *Proceedings of Theoretical Issues in Natural Language Processing*, pages 164–168, Cambridge, MA, 1975.

[63] A. Sloman. Varieties of fomalisms for knowledge representation. *Computational Intelligence*, 9(4):413–423, 1993.

[64] K. Stenning. Depictive versus descriptive representations: A distinction, but what is the difference? *Computational Intelligence*, 9(4):353–355, 1993. Commentary on taking issue paper.

[65] A. Stevens and P. Coupe. Distortions in judged spatial relations. *Cognitive Psychology*, 10:422–437, 1978.

[66] A. Stevens and D. Gentner, editors. *Mental Models*. Lawrence Erlbaum Associates: Hillsdale, NJ, 1983.

[67] H.A. Taylor and B. Tversky. Descriptions and depictions of environments. *Memory and Cognition*, 20(5):483–496, 1992.

[68] H.A. Taylor and B. Tversky. Spatial mental models derived from survey and route descriptions. *Journal of Memory and Language*, 31:261–292, 1992.

[69] P. Thagard and S. Hardy. Visual thinking in the development of Dalton's atomic theory. In *Proceedings of AI '92, the Ninth Canadian Conference on Artificial Intelligence*, pages 30–37, 1992.

[70] B. Tversky. Distortions in memory for maps. *Cognitive Psychology*, 13:407–433, 1981.

[71] B. Tversky. Distortions in memory for visual displays. In M.K. Kaiser and A. Grunwald, editors, *Pictorial communication in virtual and real environments*, pages 61–75. London: Taylor and Francis, 1991.

[72] R.D. Tweney. Five questions for computationalists. In J. Shrager and P. Langley, editors, *Computational Models of Scientific Discovery and Theory Formation*. Morgan Kaufmann, 1990.

[73] S. Ullman. Visual routines. In S. Pinker, editor, *Visual Cognition*, pages 97 – 159. MIT Press: Cambridge, Mass., 1984.

[74] S. Vosniadou and A. Ortony. Similarity and analogical reasoning: a synthesis. In *Similarity and Analogical Reasoning*, pages 1–17. Cambridge University Press, 1989.