Spatial Reasoning in Indeterminate Worlds

Janice Glasgow *
Department of Computing and Information Science
Queen's University, Kingston, Ontario
Canada K7L 3N6
janice@qucis.queensu.ca

Abstract

A possible worlds semantics for model-based spatial reasoning is presented. In this semantics, worlds are characterized by the alternative states that result from indeterminacy or partial knowledge. A world is represented as a set of symbolic arrays, where symbols in the array map to entities in the world and the relative locations of symbols correspond to the relative locations of entities. Deduction is carried out using a model-theoretic approach in which array representations are "inspected" using primitive array functions. Nonmonotonic reasoning using array representations is also discussed.

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. Just as mental models are pervasive to human problem solving, computational models for spatial reasoning provide a foundation for problem solving in AI.

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 world. Figure 1 illustrates a symbolic array representation for a map of Europe.[1] Symbols in the array correspond to the entities in the geographic domain and the relative locations of symbols in the array denote the relative directions among these entities. Each dimension in an array defines a linear order relation among entities in the domain. The order may correspond to relative location (e.g. left-of), geographic (e.g. north-of), temporal (e.g. before) or conceptual (e.g. taller-than) relations. In particular, we are concerned with transitive relations — i.e., relations r such that if $r(x, y)$ and $r(y, z)$ then $r(x, z)$. Topological relations, such as *touching, contained-in, bonded-to*, etc., can also be represented in an array. For the array representation of Europe in Figure 1, the *adjacent-to* relation in the array maps to the *borders-on* relation for the world.

The formalism presented in this paper borrows from previous research in the area of computational imagery (GP92; Gla93), which involves the study of AI knowledge representation and inferencing techniques that correspond to the representations and processes for mental imagery. In computational imagery, a mathematical theory of arrays provides a basis for representing and reasoning about visual (e.g. shape) and spatial (e.g. relative location) 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; whenever possible, the limitations of the human information processing system are overcome.

The research described in this paper extends work in computational imagery by presenting a formal semantics for spatial reasoning with array representations of worlds. The proposed formalism provides a foundation 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 representation. In the remainder of the paper we present a model-theoretic approach to spatial reasoning with array representations. An ongoing issue in AI is how to effectively update a knowledge base as new information is added or the world is transformed. The paper addresses this issue by demonstrating how nonmonotonic reasoning is achieved in the formalism.

*This research was supported by the Natural Science and Engineering Research Council (NSERC) of Canada and the Information Technology Research Corporation (ITRC) of Ontario.

[1] Note that adjacent cells in the array that contain identical symbols (e.g. Germany) are denoted by a single symbol with multiple indexes.

<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>Slovakai</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 1: Array representation of Europe

Deductive Reasoning with Spatial Models

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

left-of(a,b), *left-of(b,c)*, and
left-of(X,Y) $\wedge$ *left-of(Y,Z)* $\rightarrow$ *left-of(X,Z)*,

one can deduce *left-of(a,c)*. 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 (model) – $\boxed{a \mid b \mid c}$ – can be constructed in which *left-of* in the array corresponds to the *left-of* relation in the world being described. From this representation we can logically deduce the valid conclusion *left-of(a,c)* through the process of model inspection (also referred to as model checking).

Existing computational systems generally employ proof-theoretic deduction: 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 representation and the domain of interest. Conclusions are derived by applying functions that map 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 the spatial inferences that are made explicit by array inspection functions (GP92). In addition, the system is useful for reasoning about the indeterminate worlds resulting from uncertainty or incomplete information. This section describes how symbolic arrays can be used to represent worlds, including indeterminate worlds, consisting of entities and spatial relationships among the entities. A possible worlds semantics for deductive reasoning is also presented.

Array Representations

An *array representation* for a determinate world consists of a symbolic array, containing constant symbols corresponding to the entities in the world, and a set of array functions, which are used to determine the spatial relations among entities in the world. For example, a world described as:

The ball and the lamp are on the table and the lamp is to the right of the ball.

could be represented as the symbolic array

$$\mathcal{A} = \begin{array}{|c|c|}\hline \textit{ball} & \textit{lamp} \\ \hline \multicolumn{2}{|c|}{\textit{table}} \\ \hline \end{array},$$

where the symbols *lamp, ball* and *table* in array $\mathcal{A}$ are mapped to the corresponding entities in the world.

Truth of an atomic formula for an array representation is defined using primitive functions that "inspect" an array data structure. Assume that p is an n-ary spatial predicate symbol corresponding to a relation w_p in the world and $c_1, ..., c_n$ are constant symbols that denote entities in the world. Then the atomic formula $p(c_1, ..., c_n)$ is true for an array representation if and only if the function application $f_p(c_1, ..., c_n, \mathcal{A})$ evaluates to true, where f_p denotes the array function that inspects the array $\mathcal{A}$ to determine if the relation w_p holds for symbols $c_1, ..., c_n$. For example, in the previous array representation a function application *left-of(ball,lamp,$\mathcal{A}$)* would evaluate to *true*, whereas the expression *on-top(ball,lamp,$\mathcal{A}$)* would evaluate to *false*. A symbolic array $\mathcal{A}$ is said to represent a world w if and only if for all atomic formula $\phi = p(c_1, ..., c_n)$ in a specified language:

$$f_p(c_1, ..., c_n, \mathcal{A}) = \text{true} \text{ if and only if } (c_1, ..., c_n) \in w_p.$$

In such a case, we say that world w is *representable.*

An individual array representation models a determinate world in which all the spatial relations for the entities are explicitly specified or implied. However, a world may be indeterminate in the sense that its set of spatial relations is underspecified, resulting in ambiguity concerning the relative locations of certain entities. Indeterminacy generally implies the existence of alternative possible worlds, each of which is an extension of the indeterminate world (by adding more facts), and each of which is representable as an array. For example, a world described as – *The spoon is to the right of the fork and the knife is to the right of the fork* – suggests two consistent extensions, represented by the arrays:

fork	*spoon*	*knife*

and

fork	*knife*	*spoon*

.

An indeterminate world can be characterized by its complete (representable) extensions – i.e., those that have array representations. We say that world w' is an *extension* of world w, denoted $w \preceq w'$ if and only the two worlds consist of the same entities and all relations that hold in w also hold in w'. A world w is considered *possible* if there a representable world w' such that $w \preceq w'$. In general, we say an array $\mathcal{A}$ represents an indeterminate world w if and only if it represents a determinate extension of w.

Possible Worlds Semantics

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. A well-formed formula (wff) is necessarily true in a given world if it is true in all representable extensions of the world; a wff is possibly true if it is true in some representable extension of the world. Truth of a wff in a world w is defined recursively in terms of the truth of the atomic wffs and truth in the worlds that are extensions of w. In the following definition, a statement of the form $\models_w \phi$ denotes that the wff ϕ is true in world w.

Definition. Given a world w, we define truth of a wff ϕ in w as follows:

- If $\phi = p(c_1, ..., c_n)$ is an atomic wff then $\models_w \phi$ if and only if $(c_1, ..., c_n) \in w_p$, where w_p is the relation in w corresponding to predicate symbol p.
- $\models_w \phi \wedge \psi$ if and only if $\models_w \phi$ and $\models_w \psi$.
- $\models_w \phi \vee \psi$ if and only if $\models_w \phi$ or $\models_w \phi$.
- $\models_w \neg\phi$ if and only if not $\models_w \phi$.
- $\models_w \Box\phi$ if and only if $\models_{w'} \phi$ for all worlds w' such that $w \preceq w'$ and w' is representable.
- $\models_w \Diamond\phi$ if and only if $\models_{w'} \phi$ for some world w' such that $w \preceq w'$ and w' is representable.

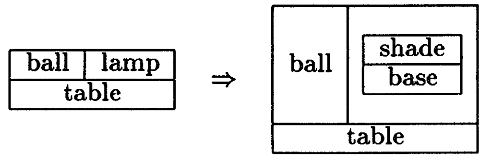

Figure 2: Embedded array representation

The proposed model theory assumes the principle of compositionality: the meaning of a wff in a possible world is determined totally by the meaning of its entities and their atomic relations. Possible worlds semantics is a version of model theory where truth of a wff in one world may depend on its truth in other possible worlds. Computationally, the possibility and necessity of truth for a wff in a world w can be determined by inspection of the array representations for the world.

Theorem: An atomic wff $\phi = p(c_1, ..., c_n)$ is necessarily true for a world w ($\models_w \Box\phi$) if and only if for all array representations $A_i =< S, \mathcal{A}_i, F >$ for w, $f_p(c_1, ..., c_n, \mathcal{A}_))$ = true for array function f_p.

Theorem: An atomic wff $\phi = p(c_1, ..., c_n)$ is possibly true for w ($\models_w \Diamond\phi$) if and only if for some array representation $A_i =< S, \mathcal{A}_i, F >$ for w, $f_p(c_1, ..., c_n, \mathcal{A}_))$ = true for array function f_p.

The proposed formalism for model-based reasoning was designed to capture and reason about the relevant spatial and structural qualities of a world. Although the examples presented are two-dimensional, the array theory on which the formalism is based (JG89) is not restricted – array functions have been developed for generating, transforming and inspecting arrays of any dimensionality. Ongoing research in this area involves several extensions to the formalism. One such extension involves the representation and inspection of structured worlds; 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. Reasoning at multiple levels of a parts hierarchy can be achieved using nested array representations where array symbols may define subarrays that correspond to the subworlds for the structured entities in the world. Figure 2 illustrates a representation where the symbol *lamp* denotes a subarray that represents the world corresponding to a structured entity.

The scheme is also being extended to model temporal worlds, where a temporal model is represented as a one-dimensional array consisting of discrete "snapshots" of worlds at progressive time steps. As well, inferences in the formalism need not be restricted to deductions. A model-based approach to analogical reasoning is also being developed using array representations.

Nonmonotonic Reasoning

Many current reasoning systems, such as first-order predicate logic, were designed for monotonic reasoning: if knowledge is added to a system then everything that was previously derivable is still derivable. Hoever, the domains that involve spatial reasoning often face the problems posed by uncertain or constantly changing knowledge where the property of monotonicity does not hold. A variety of representation schemes 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 in indeterminate worlds. Reiter's (Rei80) *default logic* allows inference rules of the form: *If A is provable and it is consistent to assume B then conclude C.* McDermott and Doyle (MD80) alternatively state defaults as sentences of the form: *If A holds and B is not disprovable then B.* Concepts such as "it is consistent to assume" and "is not disprovable" can be expressed and validated using the the concept of "possibility" in our formalism, i.e., B is not disprovable in w if B is true in some array representation for w. 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 represent the alternative interpretations arising from uncertainty.

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 for a given world. In the remainder of this section, we address the question of how such a knowledge base can be modified as the world is transformed by acquiring new knowledge or by modifying the existing spatial relations.

Knowledge Acquisition

In spatial reasoning systems, knowledge acquisition generally involves extending the spatial constraints for the world. Updating the array representations 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 the indeterminate world described by the atomic wffs *left-of(a,b)*, *left-of(a,c)*, *left-of(a,d)* and *left-of(b,d)*. This description suggests three representable extensions, corresponding to the following arrays:

a	b	c	d

$\mathcal{A}_1$

a	c	b	d

$\mathcal{A}_2$

a	b	d	c

$\mathcal{A}_3$

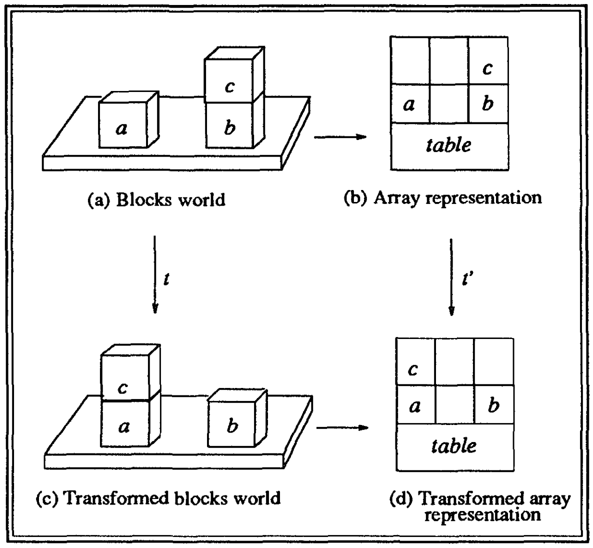

Figure 3: Representation of a blocks world

If the world is modified to include the spatial relation *left-of(c,d)*, then the array representation containing structure $\mathcal{A}_3$ would be eliminated from the knowledge base, since it is not consistent with the added spatial relationship.

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 such change is problematic in traditional reasoning systems, since it is necessary to consider the implications on the current 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 array functions to the representation. Thus, if t is a transformation that can be applied to a world w resulting in a world w', then we define a function t' such that if t' is applied to an array representation for w it would result in an array that represents w'.

To illustrate how the effects of transformations can be modeled, consider the blocks world in Figure 3(a) and its array depiction $\mathcal{A}$ in Figure 3(b). The blocks world resulting from the transformation t = *move block c to the top of block a* is illustrated in Figure 3(c). This change is modeled by applying a primitive array operation $t' = move\_rel$ to the parameter list $(c, a, \mathcal{A}, above)$. This function application results in an array that represents the transformed world, as depicted in Figure 3(d).

Array transformation functions may be complex and involve knowledge of the physical model for the entities in the domain. 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.

In summary, modifying a world by adding knowledge or by applying spatial transformations results in a new world, which can subsequently be used for reasoning about the validity of wffs. 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* (Rap71), 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.

Control Strategies for Model-based Reasoning

Model-based deduction can be carried out as a three step process: 1) a knowledge base array representations is constructed to represent the possible states of affairs (representable extensions) for the world; 2) transformations are performed on the representations in the knowledge base, corresponding to the transformations that occur in the world (this step is optional); and 3) conclusions are formed by applying inspection functions to the array representations. Alternative strategies can also be developed for model-based reasoning, depending on the form of the desired conclusion. Deductions that involve the possibility of a wff can be achieved by constructing a single model, corresponding to a possible world in which the premises and wff are true. Similarly, proving a wff invalid requires the construction of a single model in which the premises are true and the putative conclusion is false.

Model-based reasoning, as an alternative to theorem proving, has also been considered by Halpern and Vardi (1991). In this 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 (JL93). 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 reasoning, where an alternative model is generated if the current model becomes inconsistent.

Although the representation scheme was motivated by our understanding of cognitive processes, it was not intended to be model of cognition. The proposed computational approach to model-based reasoning 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 (JL93). In domains where the amount of indeterminacy is restricted, all possible array representations for a world can be generated and checked. Thus, no consistent interpretations are left unconsidered. The inferencing process for spatial representations also facilitates parallel implementations: multiple array representations can be constructed, transformed and inspected concurrently (GP92).

In conclusion, the proposed representation scheme for model-based reasoning provides an effective tool for performing spatial 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 or backtracking strategies can be developed.

Discussion

The concept of constructing knowledge representations that mirror the structure of the world is not unique to the proposed array representation. Hayes (Hay74) discusses *direct* representations in which there exist similarities between what is being represented and the medium of the representation. Sloman (Slo93) 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. Other hybrid approaches have been suggested for visual-spatial and model-based reasoning. Barwise and Etchemendy (BE92) proposed a system called *Hyperproof* which integrates diagrammatic reasoning with sentence-based logics. Myers and Konolige (MK92) 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.

Although interest and activity in spatial and diagrammatic reasoning is escalating, most of the research in this area is focussed on logic or analogical representations. What the array formalism offers is an intermediate representation that is less specific than a visual representation, yet less abstract than a logic representation. A characteristic of the array representation 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.

The array representation for spatial reasoning has measurable computational advantages over proof-theoretic logic systems. In particular, array models can be used to develop *vivid* knowledge bases. Levesque (Lev86) 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 corresponding 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.

Assuming that the array representations correctly model the world, the proposed knowledge representation scheme provides a complete and sound reasoning system that can perform under conditions of uncertainty or incomplete information. A model-theoretic formalism is used to make inferences about indeterminate worlds, using a three step process of constructing, transforming and inspecting array representations for the world. Thus, the process of generating syntactic proofs to derive spatial information is replaced by the process of model checking. The non-existence of a proof can be determined by finding an exception — i.e., an array representation in which the formula is refuted. The 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 proposed model-based approach to reasoning can be motivated and justified by human needs. Simon (Sim78) has proposed 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. Certainly 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. Although our scheme was motivated by human needs, it can overcome inherent limitations of the cognitive system.

In a recent debate, which was concerned with the advantages/disadvantages of descriptive versus depictive (model-based) representations, Levesque and Reiter (LR93) state that a reason to prefer descriptive (logic) representations is that they are "blessed with a *semantics*". Although logic-based representations can be advocated for their semantic clarity, we have shown that an intuitive semantics for model-based reasoning with array representations also exists.

References

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

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

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

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

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.

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

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

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

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.

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

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.

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.

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

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.

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