

Towards a general theory of topological maps

Emilio Remolina Benjamin Kuipers*
Computer Science Department
The University of Texas at Austin
Austin, TX 78712, USA
e-mail: {eremolin,kuipers}@cs.utexas.edu

April 1, 2002

Abstract

We present a general theory of topological maps whereby sensory input, topological and local metrical information are combined to define the topological maps explaining such information. Topological maps correspond to the minimal models of an axiomatic theory describing the relationships between the different sources of information explained by a map. We use a circumscriptive theory to specify the minimal models associated with this representation.

The theory here proposed is independent of the exploration strategy the agent follows when building a map. We provide an algorithm to calculate the models of the theory. This algorithm supports different exploration strategies and facilitates map disambiguation when perceptual aliasing arises.

1 Introduction

Topological maps are graph-like spatial representations. Nodes in such a graph often represent states in the agent's configuration space and edges represent system trajectories that take the agent from one state to another. A hierarchical structure can be accommodated on top of this "behavior graph", where nodes at one level of the hierarchy represent sets of nodes in lower levels. Despite their common use, there is no consensus about what topological maps are, or how they are built. The meanings of nodes and edges in a topological map varies according to the application as well as the algorithms used to build them. Richer structures than the graph-like description above are sometimes adopted as part of what a topological map is. Nevertheless, there are common elements to most of the topological map descriptions, namely, the use of sensory input descriptions in order to identify nodes, connectivity relations among nodes in the map, and local metrical information associated with edges in the map.

*This work has taken place in the Intelligent Robotics Lab at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Intelligent Robotics lab is supported in part by NSF grants IRI-9504138 and CDA 9617327, and by funding from Tivoli Corporation.

In this paper, we present a general theory of topological maps whereby sensory input, topological and local metrical information are combined to define the topological maps explaining such information. We take a declarative approach to define what topological maps are and how they are related to the information used to build them. We distinguish between the *causal graph*, which is a transition graph representation of regularities in action and sensory experience, and the *topological map*, which represents spatial properties of actions and of places and paths in the environment. We define topological maps as the minimal models of an axiomatic theory describing the relationships between the different sources of information explained by a map. We provide an algorithm to calculate the models of the theory. This algorithm supports different exploration strategies and facilitates map disambiguation when perceptual aliasing arises.

2 Related Work

Causal and topological maps have been mainly studied by cognitive theories of space and robotics. Cognitive theories of space are interested in the cognitive map, the human knowledge of large-scale space. Robotics is interested in representations of space that can be used (and learned) by an autonomous robot.

Computational theories of the cognitive map have been proposed by [Kuipers, 1978, Davis, 1983, McDermott and Davis, 1984, Leiser and Zilbershatz, 1989, Gopal *et al.*, 1989, O’Neill, 1991, Kortenkamp *et al.*, 1995]. These theories account for incomplete knowledge of space, use of multiple frames of reference, qualitative representation of metrical information, and connectivity relations among landmarks. The theories differ on how sensory information is represented, what a place is, and how the overall spatial knowledge is structured.

The use of topological maps in robotics varies according to the type of information used when building such map. [Rivest and Schapire, 1987, Dudek *et al.*, 1991, Dean *et al.*, 1993, Basye *et al.*, 1995] use the sequence of views and actions generated by the robot exploration to recover the minimum deterministic automaton consistent with such information. In these works, actions do not have any spatial properties associated with them. Metrical information associated with actions is considered by [Koenig and Simmons, 1996, Shatkay and Kaelbling, 1997], but there sensory information (views) is not used. The use of both sensory and metrical information is proposed by [Kuipers and Byun, 1988, Engelson and McDermott, 1992, Davis, 1983, Simmons and Koenig, 1995]. Among these works, [Kuipers and Byun, 1988, Davis, 1983] propose the use of multiple metrical frames of reference: the places in the topological map are not necessarily embedded in a single two dimensional Euclidean frame of reference, nor it is necessary to do so in order to create the topological map. [Kuipers *et al.*, 1993, Kuipers, 2000] propose the existence of topological objects (i.e. paths, regions) that can explain the agent’s experiences without relying on metrical information but rather qualitative spatial properties (i.e. travel, turnRight, turnLeft, turnAround) associated with actions.

The building of topological maps by physical robots has been described in most of the works above in addition to [Lee, 1996, Choset and Nagatani, 2001]. In these works

effort has been put on describing how the agent solves the problem of “perceptual aliasing” (i.e. different places that share the same view). Different exploration strategies as well as different discrimination procedures are proposed to solve this problem. The description of topological maps is usually closely tied to the algorithms and exploration strategy used by the agent. It is difficult then to know what topological maps are and how they are related to the agent’s experiences. In this paper we propose an axiomatic theory of topological maps. The task of building the map is then seen as an abduction task [Shanahan, 1996, Remolina and Kuipers, 1998] where the agent builds the map that best explains its observations.

Metrical maps are another spatial representation used in the robotics community [Elfes, 1987, Borenstein and Koren, 1991, Thrun *et al.*, 1998].¹ There, the location of objects in a two dimensional Euclidean space are used to explain the agent’s experiences. Topological maps as described in this paper can use metrical maps but they are confined to places, paths and local two dimensional frames of reference associated with regions. Our representation hypothesis is that an agent first builds a network of places on top of which metrical models are added, rather than to build first a single metrical map from which a network of places is derived.

This article is organized as follows: in Section 3 we define how the agent represents its experiences in the environment. Section 4 defines the causal map representation. The topological theory is presented in three parts: Section 5 introduces the main properties of paths and places. Section 6 adds boundary relations to this representation, and Section 7 defines the use of local metrical information. Section 8 presents our algorithms to build the topological maps associated with the agent’s experiences. Finally, we present our conclusions in Section 9.

3 The agent’s experiences in the environment

We assume that the continuous interaction of the agent and its environment is summarized by a discrete *view-action-view* sequence of the form

$$v_0, a_0, v_1, a_1, \dots, a_{n-1}, v_n \ . \quad (1)$$

A **view** represents a sensory description associated with an environment state. Only the name and not the internal structure of a view matters. The environment states where the views in sequence 1 were observed are called **distinctive states** (dstates). The same view can occur at different distinctive states (*perceptual aliasing*). It is possible for the agent to associate different distinctive state names with the same environment state. This is the case since the agent might not know at which of several environment states it is currently located. It is the purpose of the causal and topological theories (Sections 4 and 5) to deduce which of these dstates names refer to the same environment state. An **action** denotes a sequence of one or more control laws [Kuo, 1987] that take the agent

¹We refer the reader to Borenstein’s book [Borenstein *et al.*, 1996] (Chapter 8) for a review of the advantages and problems of using metrical maps as well as a description of several systems using this spatial representation.

from one dstate to the next.² The sequence 1 is transformed into a set of **schemas** of the form $\langle (v_i, ds_i), a_i, (v_{i+1}, ds_{i+1}) \rangle$, where ds_i is the dstate name associated with the environment state where view v_i is observed. A schema represents a particular action execution of the agent in the environment. An action execution is characterized in terms of the distinctive states the agent was at before and after the action was performed.

Example 1. Consider the environment in Figure 1. In order to go from distinctive state $ds1$ to distinctive state $ds2$, the agent executes the sequence of control laws $\langle get_into_corridor, follow_middle_line, localize \rangle$ where $get_into_corridor$ is a trajectory following control law that moves the agent from $ds1$ to a, $follow_middle_line$ is a trajectory following control law that takes the agent from a to b, and $localize$ is a hill-climbing control law that takes the agent from b to the distinctive state $ds2$. Environment states a and b are not distinctive states. At the distinctive state $ds2$ the agent is facing the wall ahead and it is equidistant from this wall and the intersection corners.




Figure 1: A sequence of control strategies $\langle get_into_corridor, follow_middle_line, localize \rangle$, takes the agent from distinctive state $ds1$ to distinctive state $ds2$. This continuous motion is represented by the schema $\langle (v1, ds1), a1, (v2, ds2) \rangle$, where $v1$ and $v2$ are the views at $ds1$ and $ds2$, and the action symbol $a1$ represents the sequence of control laws.

Distinctive state $ds3$ is at the same physical location as $ds2$ but with a different orientation. When the robot is at $ds3$, it is facing the open space (corridor) to the right of $ds2$. In order to go from distinctive state $ds2$ to distinctive state $ds3$, the agent executes the sequence of control strategies $\langle face_space_on_right, localize \rangle$. The schemas $\langle (v1, ds1), a1, (v2, ds2) \rangle$ and $\langle (v2, ds2), a2, (v3, ds3) \rangle$ are created, where $a1$ and $a2$ are action symbols representing the respective sequence of control laws.

{end of example}

4 Causal graphs

Schemas summarize the continuous interactions of the agent in the environment. This is done by storing the initial and final distinctive states (and their corresponding views) for any action execution. By considering only the views associated with the initial and final distinctive states of a schema, we define the *view graph* (Section 4.3.1, page 8),

²[Kuipers and Byun, 1988, Kuipers and Byun, 1991, Kuipers, 2000] illustrate how to create such a sequence. Distinctive states are the result of following *trajectory-following* and then *hill-climbing* control laws. The basin of attraction of the hill-climbing control laws absorbs accumulated error from each trajectory-following control law, along each action. If the basin is large enough, and error accumulates slowly enough, the action can be treated as absolutely *deterministic*.

which relates different views by actions linking them. By considering sequences of actions as well as views, the agent can further distinguish distinctive states. In Section 4.4 (Page 8) we define the predicate ceq which is the case for distinctive states that are not distinguishable by actions and views. We then define the *causal graph* whose nodes are classes of distinctive states (classes w.r.t ceq). This representation is akin to the view graph although it imposes further refinement in the set of environment states that are consistent with the agent experiences.

4.1 Ontology of the Causal theory

We use a first order sorted language in order to describe causal graphs. The sorts of such language include *views*, *actions*, *action types*, *action qualitative descriptions*, *distinctive states* and *schemas*. Next we present the predicate symbols and axioms associated with this ontology.

We use the predicate $\mathbf{View}(ds, v)$ to represent the fact that v is the *view* associated with *distinctive state* ds . We assume that a distinctive state has a unique view,^{3 4}

$$\exists!v \mathbf{View}(ds, v) . \quad (2)$$

However, we do **not** assume that views uniquely determine distinctive states (i.e. $\mathbf{View}(ds, v) \wedge \mathbf{View}(ds', v) \not\rightarrow ds = ds'$). This is the case since the sensory capabilities of an agent may not be sufficient to distinguish distinctive states.

An action has a unique type, either *travel* or *turn*, associated with it.⁵ These constant symbols define completely the sort of *action_types* (Axiom 3). The predicate $\mathbf{Action_type}(a, type)$ represents the fact that the type of action a is *type*. Formally,

$$turn \neq travel, \quad \forall atype \{ atype = turn \vee atype = travel \} , \quad (3)$$

$$\exists!type \mathbf{Action_type}(a, type) . \quad (4)$$

Turn actions have associated a unique qualitative description. The sort of qualitative descriptions is completely defined by the constant symbols *turnLeft*, *turnRight* and *turnAround* (Axioms 5 and 6). We use the predicate $\mathbf{Turn_desc}(a, desc)$ to indicate that $desc$ is the qualitative description of the *turn action* a . Formally,⁶

$$UNA[turnLeft, turnRight, turnAround] , \quad (5)$$

$$\forall desc \{ desc = turnLeft \vee desc = turnRight \vee desc = turnAround \} , \quad (6)$$

$$\mathbf{Turn_desc}(a, desc) \rightarrow \mathbf{Action_type}(a, turn) , \quad (7)$$

$$\mathbf{Action_type}(a, turn) \rightarrow \exists!desc \mathbf{Turn_desc}(a, desc) . \quad (8)$$

A **schema** represents a particular action execution of the agent in the environment. We use the following predicates to represent information associated with a schema: **action(s,a)** — *action* a is the action associated with *schema* s —, **context(s,ds)** — ds is

³Throughout this paper we assume that free variables in formulas are universally quantified.

⁴The formula $\exists!v P(v)$ means “*there exists a unique v s.t. P(v)*”. Formally, $\exists v \forall x [P(x) \equiv x = v]$.

⁵The type of an action will be important in the topological theory (Section 5). For completeness of the presentation we introduce this concept here.

⁶The notation $UNA[t_1, \dots, t_n]$ represents the uniqueness of names axioms for the grounded terms t_1, \dots, t_n . These axioms require that $t_i \neq t_j$ for $i \neq j$.

the starting *distinctive state* associated with the action execution represented by *schema* s —, and **result(s,ds)** — ds is the ending *distinctive state* associated with the action execution represented by *schema* s —. While we require a unique context and action associated with a schema, the result of a schema is optional (but unique if it exists):⁷

$$\exists!a \text{ action}(s, a), \exists!ds \text{ context}(s, ds), \text{result}(s, ds) \wedge \text{result}(s, ds') \rightarrow ds = ds' . \quad (9)$$

Most often we are interested in *complete* schemas: those for whom the resulting distinctive state exists. We use the predicate $CS(s, ds, a, ds')$ defined as

$$CS(s, ds, a, ds') \equiv_{def} \text{context}(s, ds) \wedge \text{action}(s, a) \wedge \text{result}(s, ds') \quad (10)$$

to express the fact that schema s represents an execution of action a which took the agent from *distinctive state* ds to distinctive state ds' .⁸

An action execution also has metrical information associated with it. This metrical information represents an estimate of, for example, the distance or the angle between the distinctive states associated with the action execution. We defer the study of metrical information associated with schemas until Section 7.

4.2 Schema notation

While schemas are explicit objects of our theory, it is convenient to leave them implicit. We introduce the following convenient notation:⁹

$$\begin{aligned} \langle ds, a, ds' \rangle &\equiv_{def} \exists s \ CS(s, ds, a, ds') \\ \langle v, a, v' \rangle &\equiv_{def} \exists s, ds, ds' \ \{CS(s, ds, a, ds') \wedge View(ds, v) \wedge View(ds', v')\} \\ \langle (v, ds), a, (v', ds') \rangle &\equiv_{def} \exists s \ \{CS(s, ds, a, ds') \wedge View(ds, v) \wedge View(ds', v')\} \\ \langle ds, type, ds' \rangle &\equiv_{def} \exists s, a \ \{CS(s, ds, a, ds') \wedge Action_type(a, type)\} \\ \langle ds, desc, ds' \rangle &\equiv_{def} \exists s, a \ \{CS(s, ds, a, ds') \wedge Turn_desc(a, desc)\} \end{aligned}$$

4.3 The E formulae.

The agent's experiences in the environment, E , are described in terms of CS , $View$, $Action_type$ and $Turn_desc$ formulae. Associated with E we have the sets $S(E)$, $DS(E)$, $V(E)$, $A(E)$ of schemas, distinctive states, views and action constant symbols occurring in E . We require all these symbols to be different (i.e. *uniqueness of names* assumption) and to completely define their corresponding sorts (*domain closure* assumption):

$$\begin{aligned} UNA[s_1, \dots, s_k], s_i \in S(E), \quad UNA[ds_1, \dots, ds_l], ds_i \in DS(E), \\ UNA[a_1, \dots, a_n], a_i \in A(E), \quad UNA[v_1, \dots, v_m], v_i \in V(E), \quad (11) \\ \forall s \bigvee_{s_i \in S(E)} s = s_i, \quad \forall ds \bigvee_{ds_i \in DS(E)} ds = ds_i, \\ \forall a \bigvee_{a_i \in A(E)} a = a_i, \quad \forall v \bigvee_{v_i \in V(E)} v = v_i. \end{aligned}$$

⁷Incomplete schemas allow the representation to account for common states of incomplete knowledge like “I could take you there, but I can't tell you how” [Kuijpers, 2000].

⁸CS stands for Causal Schema.

⁹Notice that we have “overloaded” the bracket notation depending on the type of its arguments.

The axioms above are not only required from a logical point of view, but make sense from the knowledge representation point of view. Domain closure axioms prevent models from including objects different from those experienced (named) by the agent. Each of the agent schemas represents a different experience and the agent names them with a different schema constant symbol. Different view symbols represent different sensory input. This is the case since the agent decides what view to associate with a sensory input. *Different distinctive state constant symbols might represent the same environment state.* Nevertheless, we assume that different distinctive state symbols are interpreted by different elements of the sort of distinctive states and we use the predicate *ceq* to indicate whether two distinctive states represent the same environment state.¹⁰

Finally, the type of actions as well as the qualitative description of turn actions have to be specified as part of the formulae E :

$$Action_type(a, type) \equiv \bigvee_{Action_type(a_i, type_i) \in E} [a = a_i \wedge type = type_i] \quad (12)$$

$$Turn_desc(a, desc) \equiv \bigvee_{Turn_desc(a_i, desc_i) \in E} [a = a_i \wedge desc = desc_i] \quad (13)$$

Definition 1.

Given a set E of *CS*, *View*, *Action_type* and *Turn_type* formulae,

$$COMPLETION(E)$$

denotes the union of E with Axioms 11 - 13. *{end of definition}*

Example 2. Consider the set of experiences E gathered by the agent while navigating the environment in Figure 2. The agent moves among intersections by performing action *ml*. The sensory input at the different intersections is very similar, and the agent associates the view $v+$ ¹¹ with the different distinctive states it found (i.e. a , b and c).




Figure 2: The agent moves among corridor intersections that have the same view $v+$. a , b and c are the distinctive states where this view is observed at.

The elements of E are as follows: $Action_type(ml, travel)$, $CS(s1, a, ml, b)$, $CS(s2, b, ml, c)$, $View(a, v+)$, $View(b, v+)$, and $View(c, v+)$.

The uniqueness of names axioms associated with E are $s1 \neq s2$ and $a \neq b \wedge a \neq c \wedge b \neq c$. The domain closure axioms associated with E are $\forall s \{s = s1 \vee s = s2\}$, $\forall ds \{ds = a \vee ds = b \vee ds = c\}$, $\forall a \{a = ml\}$ and $\forall v \{v = v+\}$.

¹⁰*ceq* stands for Causally Equal. See Section 4.4.

¹¹As with any other symbol name, the view name is arbitrary. The + in the view name is used to indicate that the view corresponds to a four corridor intersection. Later we use the symbol \square to indicate that the view corresponds to an end of corridor.

Finally, we also have the axioms $\forall a, desc \{Turn_desc(a, desc) \equiv false\}$ and $\forall a, type \{Action_type(a, type) \equiv [a = ml \wedge type = travel]\}$. *{end of example}*

4.3.1 The view graph

The *view graph* associated with a set of experiences E is the labeled graph $\langle Nodes, Edges, Labels \rangle$ such that:

- Nodes = $V(E)$, Labels = $A(E)$.
- Edges = $\{(v, a, v') : COMPLETION(E) \models \langle v, a, v' \rangle\}$.

When the same view occurs at different environment states, the view graph is not very informative. The agent has to use information other than the views alone in order to distinguish different environment states (see next section and Section 5). However, should the agent have enough sensory capabilities as to distinguish distinctive states by their views, then the view graph becomes a powerful spatial representation for reliable navigation. Work in [Schölkopf and Mallot, 1995, Franz *et al.*, 1998, Mallot and Gillner, 2000, Steck and Mallot, 2000] shows how the view graph is consistent with human navigation abilities.

4.4 The Causal theory

We use the predicate $ceq(ds, ds')$ to denote the fact that *distinctive states* ds and ds' are *causally* indistinguishable.¹² Informally, $ceq(ds, ds')$ is the case whenever distinctive states ds and ds' are indistinguishable by the actions and views in a given set of experiences E . The theory $CT(E)$ below defines the extent of the predicate ceq .

The causal theory associated with a set of experiences E , $CT(E)$, is the following nested abnormality theory (NATs) [Lifschitz, 1995] (see Appendix A, Page 43):

$$CT(E) = \tag{14}$$

$$COMPLETION(E),$$

$$Axioms 2 - 10,$$

$$\langle ds, a, ds' \rangle \wedge \langle ds, a, ds'' \rangle \rightarrow ds' = ds'', \tag{15}$$

$$CEQ_block = \tag{16}$$

$$\{ \text{max } ceq :$$

$$ceq(ds_1, ds_1),$$

$$ceq(ds_1, ds_2) \rightarrow ceq(ds_2, ds_1),$$

$$ceq(ds_1, ds_2) \wedge ceq(ds_2, ds_3) \rightarrow ceq(ds_1, ds_3),$$

$$ceq(ds_1, ds_2) \rightarrow View(ds_1, v) \equiv View(ds_2, v), \tag{17}$$

$$ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2) \tag{18}$$

$$\}$$

¹²In Section 5 we define when distinctive states are topologically indistinguishable.

Axiom 15 states our assumption that actions are deterministic. Axiom 17 states that indistinguishable distinctive states have the same view. Axiom 18 states that if distinctive states ds and ds' are indistinguishable, and action a is performed for both ds and ds' , then the resulting distinctive states must also be indistinguishable. Axioms 17 and 18 allow us to prove that if ds and ds' are two indistinguishable distinctive states, then any sequence of actions executed at ds and ds' will render the same sequence of views.

Lemma 1 *Let A denote a sequence of action symbols. Let $A(ds)$ denote the distinctive state symbol resulting from executing the sequence A starting at distinctive state ds , or \perp if A is not defined for ds .¹³ Then,*

$$ceq(ds_1, ds_2) \wedge A(ds_1) \neq \perp \wedge A(ds_2) \neq \perp \rightarrow View(A(ds_1), v) \equiv View(A(ds_2), v) .$$

There is a special case in which ceq is an equivalence relation without explicitly stating the axioms requiring so. This is the case when the result of every action at every distinctive state is known.

Definition 2. A set of experiences E is **complete** whenever

$$E \models \forall a, ds \exists ds' \langle ds, a, ds' \rangle .$$

{end of definition}

Theorem 1 *Let E be a complete set of experiences and let CEQ_block be defined as follows:*

$$\{ \text{max } ceq : \\ ceq(ds_1, ds_2) \rightarrow View(ds_1, v) \equiv View(ds_2, v), \\ ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2) \\ \}$$

Then, the predicate ceq is an equivalence relation.

Proof. See appendix B (page 46).

When a set of experiences is complete the predicate ceq captures the idea that two distinctive states are the same if they render the same views under any sequence of actions.

Theorem 2 *Let E be a complete set of experiences. Then,*

$$ceq(ds_1, ds_2) \equiv \forall A, v [View(A(ds_1), v) \equiv View(A(ds_2), v)] .$$

Proof. See appendix B (page 48).

Example 3. Consider the set of experiences E as in Example 2 (Page 7) (see Figure 3a). Since the same view is experienced at a , b and c , the extent of ceq is maximized by



Figure 3: (a) Distinctive states a , b and c cannot be causally distinguished. Topological information is needed in order to distinguish them. (see text) (b) a , b and c are distinguished given the new information $\langle c, \text{travel}, d \rangle$.

declaring $ceq = true$ (i.e. $\forall x, y ceq(x, y)$). Notice that axiom (18) is trivially satisfied since no action has been executed at c .

Though a , b and c were experienced at different environment states, they are declared causally indistinguishable. This happens because neither the actions nor the views in E provide enough information to distinguish them. By using topological information (i.e. the concepts of *path* and *place* in Section 5) we will be able to distinguish these distinctive states (see Example 5, page 18).

Suppose the agent continues exploring the environment and gets the new information $View(d, v \sqsupset)$, $CS(s3, c, ml, d)$, as suggested in Figure 3b. In virtue of lemma 1, it can be seen that $ceq(ds, ds') \equiv ds = ds'$, and consequently the agent concludes that all distinctive states refer to different environment states. *{end of example}*

Different models of $CT(E)$ generally arise when the set of experiences E is incomplete (i.e. the agent has not completely explore the environment) or when weak sensors determine the same view at different environment states.

Example 4. Consider the environment depicted in Figure 4. The agent visits the different distinctive states as suggested by their numbers in the figure. The same travel action ml is performed when traveling from a corner to the intersection (i.e. $\langle 1, ml, 2 \rangle$) and viceversa (e.g. $\langle 4, ml, 5 \rangle$). A turn around action is performed when reaching a corner (e.g. $\langle 3, \text{change_path_direction}, 4 \rangle, \langle 7, \text{change_path_direction}, 8 \rangle$, etc.). Assume that the different corners have the same views (i.e. $view(1) = view(4) = view(8)$, $view(3) = view(7) = view(11)$), and views associated with the other distinctive states are different.

Three models of $CT(E)$ can be associated with the exploration E of the T-environment:

1. Model 1: $ceq(8, 12), ceq(12, 8), ceq(x, x)$.¹⁴
2. Model 2: $ceq(1, 12), ceq(12, 1), ceq(x, x)$.
3. Model 3: $ceq(4, 12), ceq(12, 4), ceq(3, 11), ceq(11, 3), ceq(2, 10), ceq(10, 2), ceq(x, x)$.

In all the models above, $\neg ceq(1, 4)$, $\neg ceq(1, 8)$, $\neg ceq(4, 8)$. For instance, from $\langle 1, ml, 2 \rangle, \langle 4, ml, 5 \rangle$, and $view(2) \neq view(5)$ we conclude that $\neg ceq(1, 4)$. Although dstate 12 is at the same environment state as dstate 4, it is possible that $ceq(1, 12)$ or $ceq(8, 12)$. This is the case since no action has been performed at dstate 12.

¹³Given an action symbol A and distinctive state ds , $A(ds) = ds'$ if the schema $\langle ds, A, ds' \rangle$ has been observed, otherwise, $A(ds) = \perp$. Moreover, $A(\perp) = \perp$. The definition is then extended to action sequences in the standard way. Notice that $A(ds)$ is well-defined given our assumption that actions are deterministic (Axiom 15).

¹⁴The extent of ceq in model 1 is defined by $\{(8, 12), (12, 8)\} \cup \{(x, x) : x = 1, \dots, 12\}$.




Figure 4: The agent visits the different distinctive states in the order suggested by their numbers. The same view occurs at the different corners (i.e view(1)= view(4) = view (8)). Three different causal models can be associated with the agent exploration of this T-environment (see text).

Notice that the models of $CT(E)$ are maximal with respect to the set inclusion for ceq . The number of elements in the possible extents of ceq could vary, and consequently the number of different environment states represented by the models of $CT(E)$ will also vary. For instance, the three models above represent 11, 11 and 10 environment states respectively.

Finally, notice that all the models above are possible since at the causal level turn and travel actions do not bear any spatial meaning. When we consider topological information, only model 3 above will be possible (see Example 10, Page 21). *{end of example}*

4.5 The causal graph

The causal graph associated with a set of experiences E is the labeled graph $\langle Nodes, Edges, Labels \rangle$ such that:

- Nodes = $DS(E)/ceq$, Labels = $A(E)$,
- Edges = $\{([ds], a, [ds']) : COMPLETION(E) \models \langle ds, a, ds' \rangle \}$.

where $DS(E)/ceq$ denotes the set of equivalence classes of $DS(E)$ modulo ceq , and $[ds]$ denotes the equivalence class of ds given ceq .




Figure 5: (a)-(b). Causal graphs associated with the set of experiences in Figures 3a and 3b. (c) view graph associated with the set of experiences in Figure 3b. Notice that the causal and view graphs associated with the experiences in Figure 3a are isomorphic.

The problem of distinguishing environment states by outputs (views) and inputs (actions) has been studied in the framework of automata theory [Angluin, 1978, Gold, 1978, Rivest and Schapire, 1987, Basye *et al.*, 1995]. In this framework, the problem we address is the one of finding the minimum automaton (w.r.t. the number of states) consistent with a given set of input/output pairs. Without any particular assumptions about the environment or the agent’s perceptual abilities, the problem of finding this smallest automaton is NP-complete ([Angluin, 1978, Gold, 1978]).

5 Topological maps

Actions in the causal theory convey patterns of experience but not spatial configuration. Spatial configuration is considered by the topological theory where actions are categorized into two classes: *turns* and *travels*. Turns and travels are explained by a new ontology, that of *places* and *paths*. Turn actions leave the agent at the same place. Travel actions move the agent to a new place along a path.

Grouping places into *regions* allows an agent to reason efficiently about its spatial knowledge. Regions themselves can be grouped to form new regions forming a spatial abstraction hierarchy.¹⁵ In Section 6 (Page 26) we define *boundary regions* associated with paths. Informally, a path has associated three disjoint regions: the set of places in the path, the set of places to the left of the path, and the set of places to the right of the path. Boundary regions allow the agent to distinguish distinctive states, for two distinctive states can be considered different if they are in different boundary regions of the same path (see Example 16, Page 29).

Local metrical information derived during action execution is considered in the topological theory. For instance, the distances among places on a path or the angles among paths intersecting in a place can be accommodated in the topological map. We study the use of metrical information in Section 7 (Page 30).

5.1 Ontology of the Topological theory

The main purpose of the topological theory $TT(E)$ is to minimize the set of topological paths and topological places consistent with the given experiences E . The concepts of *path* and *place* are used to distinguish environment states that are not distinguishable by actions and views alone. We use the predicate $\mathbf{teq}(ds, ds')$ to indicate that distinctive states ds and ds' are topologically indistinguishable. This will be the case, when in addition to not being distinguishable by views and actions, ds and ds' are at the same place facing the same direction along the same path.

Within the sort of places, we distinguish between *topological places* and *regions*. A topological place is a set of distinctive states linked by turn actions. A region is a set of places. We use the predicates \mathbf{tpath} and $\mathbf{is_region}$ to identify these subsorts.

A path defines an order relation among places connected by travel with no turn actions. They play the role of streets in a city layout. Among paths, *topological paths* correspond to those paths whose places are topological places. We use the predicate

¹⁵In this article we do not consider this hierarchy.

tpath to identify these paths. A path connecting regions is called a **route**. A path has associated two directions, *pos* and *neg*. A path direction provides a frame of reference to establish the order in which places in a path are arranged or whether a place is to the right or left of the path (see boundary regions in Section 6). The sort of path directions is completely defined by *pos* and *neg*.¹⁶

The relations among distinctive states, places and paths are characterized in terms of the following predicates: **on**(**pa,p**) — *place p is on path pa* —, **order**(**pa,dir,p,q**) — *place p is before place q, when facing direction dir on path pa* —, **at**(**ds,p**) — *distinctive state ds is at place p* —, and **along**(**ds,pa,dir**) — *distinctive state ds is along path pa in direction dir* —.

We require the sorts of *places* and *paths* to be infinite. This is not to say that the topological map has infinite number of *places* or *paths*. Given a model of the theory, the topological map corresponds to the submodel obtained by restricting the different predicates to *topological places*, *regions*, *topological paths* and *routes*. Since *topological places* are identified with set of distinctive states and *topological paths* are identified with sequences of distinctive states, the topological map associated with a finite set of schemas (and so a finite set of distinctive states) has a finite number of *topological places* and *topological paths*. Our requirement of infinite *places* and *paths* allow us to compare any two models of the theory. Example 15 (Page 26) illustrates the use of this requirement.

5.2 The topological theory

The topological theory associated with E , $\mathbf{TT}(E)$, is the following nested abnormality theory (NATs) [Lifschitz, 1995] (see Appendix A, Page 43):¹⁷

$$\begin{aligned}
 \mathbf{TT}(E) = & \tag{19} \\
 & \text{there exist infinitely many places ,} \\
 & \text{there exist infinitely many paths ,} \\
 & \neg \exists p [tplace(p) \wedge is_region(p)] , \\
 & \neg \exists pa [tpath(pa) \wedge route(pa)] , \\
 & \mathbf{COMPLETION}(E) , \\
 & \mathbf{Axioms 2 - 10} , \\
 & \langle ds, a, ds' \rangle \wedge \langle ds, a, ds'' \rangle \rightarrow ds' = ds'' , \quad (\mathbf{Axiom 15}) \\
 & \mathbf{T_block} , \\
 & \mathbf{AT_block} .
 \end{aligned}$$

The block **T_block** defines the properties of the predicates \widehat{turn} , \widehat{travel} , and \vec{travel} . \widehat{turn} is the equivalence closure of the schemas $\langle \cdot, turn, \cdot \rangle$; \widehat{travel} and \vec{travel} are the equivalence and transitive closure of the schemas $\langle \cdot, travel, \cdot \rangle$ respectively.¹⁸

¹⁶For a direction *dir*, $-dir$ is defined such that $-pos = neg$ and $-neg = pos$.

¹⁷The condition that the sorts of places and paths are infinite is formalized by asserting the existence of a bijection between these sorts and the natural numbers.

¹⁸See Appendix D, Page 52.

The block **AT_block** is the heart of our theory. It defines how the agent groups distinctive states into *places*, and how *places* are ordered by *paths*. The purpose of this block is to define the extent of the predicates *tpath*, *tplace*, *at*, *along*, *order*, *on* and *teq*. The block has the associated circumscription policy¹⁹

$$\text{circ } tpath \succ tplace \text{ var } SSH\vec{pred}$$

where $SSH\vec{pred}$ stands for the tuple of predicates $\langle \mathbf{at}, \mathbf{along}, \mathbf{order}, \mathbf{on}, \mathbf{teq}, \mathbf{turn_eq}, \mathbf{travel_eq} \rangle$.²⁰ This circumscription policy states that a minimum set of topological paths is preferred to a minimum set of topological places. The block *AT_block* is defined as follows:²¹

$$AT_block = \tag{20}$$

$$\{ \text{max } teq : \tag{21}$$

$$teq(ds, ds) ,$$

$$teq(ds_1, ds_2) \rightarrow teq(ds_2, ds_1) ,$$

$$teq(ds_1, ds_2) \wedge teq(ds_2, ds_3) \rightarrow teq(ds_1, ds_3) ,$$

$$teq(ds_1, ds_2) \rightarrow View(ds_1, v) \equiv View(ds_2, v), \tag{22}$$

$$teq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow teq(ds'_1, ds'_2) , \tag{23}$$

$$teq(ds_1, ds_2) \rightarrow \forall p [at(ds_1, p) \equiv at(ds_2, p)] \wedge \tag{24}$$

$$\forall pa, dir [along(ds_1, pa, dir) \equiv along(ds_2, pa, dir)] ,$$

$$\langle ds, turn, ds' \rangle \rightarrow \neg teq(ds, ds') , \tag{25}$$

$$\langle ds, turnAround, ds' \rangle \wedge \langle ds, turnAround, ds'' \rangle \rightarrow teq(ds', ds'') , \tag{26}$$

$$\langle ds_1, turnAround, ds_2 \rangle \wedge \langle ds_2, turnAround, ds_3 \rangle \rightarrow teq(ds_1, ds_3) , \tag{27}$$

$$at(ds, p) \rightarrow tplace(p), \tag{28}$$

$$\exists! p at(ds, p), \tag{29}$$

$$turn_eq(ds_1, ds_2) \equiv \forall p [at(ds_1, p) \equiv at(ds_2, p)] , \tag{30}$$

$$\{ \text{min } turn_eq : \tag{31}$$

$$teq(ds_1, ds_2) \wedge teq(ds_3, ds_4) \wedge \widehat{turn}(ds_2, ds_3) \rightarrow turn_eq(ds_1, ds_4),$$

$$turn_eq(ds_1, ds_2) \wedge turn_eq(ds_2, ds_3) \rightarrow turn_eq(ds_1, ds_3)$$

$$\} \tag{32}$$

$$along(ds, pa, dir) \rightarrow tpath(pa), \tag{32}$$

$$\{ \text{min } along : \tag{33}$$

$$\langle ds, travel, ds' \rangle \rightarrow \exists pa, dir [along(ds, pa, dir) \wedge along(ds', pa, dir)] , \tag{34}$$

$$\langle ds, turnAround, ds' \rangle \rightarrow along(ds, pa, dir) \equiv along(ds', pa, -dir), \tag{35}$$

$$teq(ds_1, ds_2) \rightarrow along(ds_1, pa, dir) \equiv along(ds_2, pa, dir) \tag{36}$$

¹⁹The symbol \succ indicates prioritized circumscription (see Appendix A, Page 43).

²⁰The predicates *travel_eq* and *turn_eq* are “auxiliary” predicates used in our topological theory. Although they are completely defined in terms of *teq*, \widehat{turn} and \widehat{travel} , they need to vary in the circumscription policy.

²¹Figure 6, Page 16, summarizes the dependencies among the predicates defined by the *AT_block*.

$$\begin{aligned}
& \} \\
& \text{along}(ds, pa, dir) \wedge \text{along}(ds, pa1, dir1) \rightarrow pa = pa1 \wedge dir = dir1, \quad (37) \\
& \text{at}(ds_1, p) \wedge \text{at}(ds_2, p) \wedge \text{along}(ds_1, pa, dir) \wedge \text{along}(ds_2, pa, dir) \rightarrow \text{teq}(ds_1, ds_2) \quad (38)
\end{aligned}$$

$$\begin{aligned}
& [\langle ds, \text{turn_desc}, ds' \rangle \wedge \text{turn_desc} \neq \text{turnAround} \wedge \\
& \quad \text{along}(ds, pa, dir) \wedge \text{along}(ds', pa1, dir1)] \rightarrow pa \neq pa1, \quad (39)
\end{aligned}$$

$$\{ \text{min order} : \quad (40)$$

$$[\langle ds, \text{travel}, ds' \rangle \wedge \text{at}(ds, p) \wedge \text{at}(ds', q) \wedge \quad (41)$$

$$\text{along}(ds, pa, dir) \wedge \text{along}(ds', pa, dir)] \rightarrow \text{order}(pa, dir, p, q),$$

$$\text{order}(pa, pos, p, q) \equiv \text{order}(pa, neg, q, p), \quad (42)$$

$$\text{order}(pa, dir, p, q) \wedge \text{order}(pa, dir, q, r) \rightarrow \text{order}(pa, dir, p, r) \quad (43)$$

$$\}$$

$$\neg \text{order}(pa, dir, p, p), \quad (44)$$

$$\{ \text{min on} : \text{at}(ds, p) \wedge \text{along}(ds, pa, dir) \rightarrow \text{on}(pa, p) \} \quad (45)$$

$$\text{on}(pa, p) \wedge \text{on}(pa, q) \wedge \text{tpath}(pa) \rightarrow \quad (46)$$

$$\begin{aligned}
& \exists ds_1, dir_1, ds_2, dir_2 [\text{at}(ds_1, p) \wedge \text{along}(ds_1, pa, dir_1) \wedge \text{at}(ds_2, q) \wedge \\
& \quad \text{along}(ds_2, pa, dir_2) \wedge \text{travel_eq}(ds_1, ds_2)],
\end{aligned}$$

$$\{ \text{min travel_eq} : \quad (47)$$

$$\widehat{\text{travel}}(ds_1, ds_2) \rightarrow \text{travel_eq}(ds_1, ds_2),$$

$$\langle ds_1, \text{turnAround}, ds_2 \rangle \rightarrow \text{travel_eq}(ds_1, ds_2) \wedge \text{travel_eq}(ds_2, ds_1)$$

$$\text{teq}(ds_1, ds_2) \wedge \text{teq}(ds_3, ds_4) \wedge \text{travel_eq}(ds_2, ds_3) \rightarrow \text{travel_eq}(ds_1, ds_4),$$

$$\text{travel_eq}(ds_1, ds_2) \wedge \text{travel_eq}(ds_2, ds_3) \rightarrow \text{travel_eq}(ds_1, ds_3)$$

$$\}$$

$$\text{circ } \text{tpath} \succ \text{tplace } \text{var } \text{SSHpred} \quad (48)$$

$$\}$$

We discuss these axioms in turn.

Predicate *teq* is an equivalence relation. It stands for *topologically equal*. Whenever *teq*(ds_1, ds_2) is the case, we can consider ds_1 and ds_2 as denoting the same environment state: ds_1 and ds_2 cannot be distinguished by views and actions (Axioms 22 and 23), they are at the same place, and they are along the same paths (Axiom 24).

Axiom 25 states that a *turn* action takes the agent from one distinctive state to a different one. In particular we assume that a schema of the form $\langle ds, \text{Turn}, ds \rangle$ is not included in the agent's experiences. Axiom 26 states that there is a unique (modulo *teq*) distinctive state resulting from performing a turn around action. After two turn around actions the agent is back to the same dstate (Axiom 27). Turn around actions are special since they link distinctive states along the same path but in opposite directions (Axiom 35).

Axioms 29 and 30 state how the agent groups distinctive states into places. Ev-

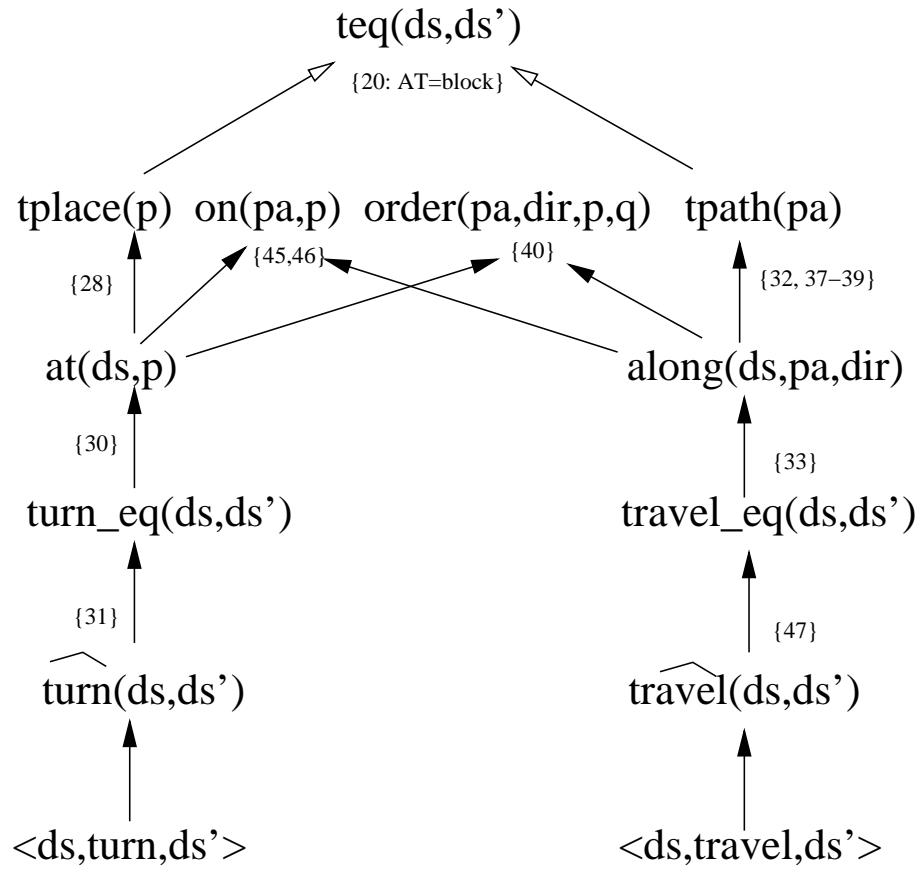


Figure 6: Dependency among predicates in $TT(E)$. Labels on the graph's arrows refer to the axioms relating the predicates pointed by the arrows.

Distinctive states related by turns modulo teq ($turn_eq$) must be *at* the same topological place ($tplace$). Distinctive states related by travels modulo teq ($travel_eq$) are *along* the same topological path ($tpath$). Knowing at which places and along which paths distinctive state are, determines what places are *on* what paths. The order of places on a path is derived from travels among distinctive states along a path.

Since the extents of $travel_eq$ and $turn_eq$ must be defined in order to determine places and paths, one has to know what distinctive states are teq . The arrows pointing to teq on the top of the diagram indicate that among the possible interpretations for teq , the preferred models of the theory select those that lead to a map where a minimum set of paths and places are needed to explain the schemas at the bottom of the diagram.

ery distinctive state is at a unique topological place (Axiom 29). Whenever the agent *turns*, it stays at the same topological place (Axiom 30). Distinctive states grouped into a topological place should be *turn* connected (modulo *teq*) (Axiom 30).²²

Travel actions among distinctive states are abstracted to topological paths connecting the places associated with such distinctive states. Travel axioms are explained in terms of the two related predicates, *along* and *order*. Both of these predicates are the minimum ones explaining travel actions and satisfying other properties included in their respective blocks 33 and 40.

Block 33 defines the predicate *along*. Whenever an agent *turns around*, it stays in the same path but facing the opposite path's direction (Axiom 35). Axiom 36 is a trivial consequence of the definition of *teq* but it has to be included in the block so that the interpretation of *along* has tuples other than the ones explicitly derived from schemas (see Example 7, Page 19).

There are further restrictions on the properties of *along*. For instance, a distinctive state is *along* at most one path (Axiom 37). Since Axiom 37 provides "negative" information about *along*, it does not need to be included in Block 33 (see Proposition 4 in [Lifschitz, 1994]). Axiom 37 prevents the existence of different paths that converge to the same distinctive state.

Turn actions other than *turnAround* change the path the initial and final distinctive states linked by the action are *along* (Axiom 39). This axiom allows the agent to conclude the existence of different paths once it turns right or left at a place (see Example 9, Page 20). This axiom prevents the existence of self-intersecting paths as illustrated in Figure 14 (Page 24).

Block 40 defines the predicate *order*. In addition to explaining travel actions, *order* defines an order among the places on a path satisfying the following two properties: i) the order of places in a given path direction is the inverse of the order of places in the other path direction (Axiom 42), and ii), the order of places in a path is transitive (Axiom 43).

There are further restrictions on the properties of *order*: i) the order of places in a path should be non-reflexive (Axiom 44), and ii) the agent has to have traveled among the places on the same path (Axiom 46). Since these requirements provide "negative" information about *order*, they do not need to be included in Block 40 (see Proposition 4 in [Lifschitz, 1994]). Notice that we rule out the existence of circular paths (Axiom 44). In Section 5.3 we will make this axiom a default.

Axiom 46 requires the agent to have traveled among the places on the same path. *travel_eq* defines when two distinctive states are linked by travel actions without turns (except for *turnAround* actions) (see Block 47). Example 8 illustrates how by using *travel_eq* the agent can minimize the set of topological paths.

Remark. We will be using the following properties of our theory. Axiom 37 in combination with Axioms 34, 41, and 44, imply that that whenever the agent has directly traveled between two distinctive states, the places associated with these distinctive states are different:

Corollary 1. $\vec{travel}(ds, ds') \rightarrow place(ds) \neq place(ds')$

where $\mathbf{place}(ds)$ denotes the unique topological place that distinctive state ds is

²²Block 31 states that the predicate *turn_eq* corresponds to the relation \widehat{turn} modulo *teq*.

at (Axiom 29). Moreover, consecutive travels among distinctive states occur along the same topological path.

Corollary 2.

$$\vec{travel}(ds, ds') \rightarrow \exists! pa, dir \ [order(pa, dir, place(ds), place(ds')) \wedge along(ds, pa, dir) \wedge along(ds, pa, dir)]$$

In order to prove that distinctive states ds_1 and ds_2 are at different topological places, one has to prove that $\neg turn_eq(ds_1, ds_2)$. The following theorem states a strong condition for when this is the case.

Theorem 3 *Let ds_1 be a distinctive state symbol such that*²³

$$\forall ds_2 \notin [ds_1]_{turn}, [ds_2]_{teq} \cap [ds_1]_{turn} = \emptyset .$$

Then, $\forall ds_2 \notin [ds_1]_{turn}, place(ds_2) \neq place(ds_1) .$

Proof. See Appendix C, Page50. \square

Finally, our circumscription policy 48 and the fact that the sorts of paths and places are infinite implies the following fact:²⁴

Theorem 4 *Any two models of the SSH topological theory have the same number of topological paths and the same number of topological places.*

Proof. See appendix C (page 51). \square

However, Theorem 4 does not mean that a unique map is necessarily associated with a set of schemas. As shown in Example 13 (Page 24) the SSH topological theory could have more than one non-isomorphic model.

{end of remark}

The next examples illustrate the interplay among the axioms in AT_block.

Example 5.




Figure 7: Distinctive states a , b and c cannot be distinguished at the causal level (see Example 3, Page 9). Using the concepts of *paths* and *places* these dstates are distinguished.

Consider the set of experiences E

$$\langle (a, v+), travel, (b, v+) \rangle, \langle (b, v+), travel, (c, v+) \rangle$$

²³Given an equivalence relation R , $[x]_R$ denotes the equivalence class of x according to R .

²⁴Recall that the interpretations for *tpath* and *tplace* are finite.

as in Example 3 (Page 9), Figure 7. From Corollary 1 we deduce that $place(a)$, $place(b)$ and $place(c)$ are all different places. From Corollary 2, the topological map associated with E has one topological path and three topological places. Distinctive states a and b can be distinguished though they are “causally indistinguishable” (i.e. $ceq(a, b) \wedge \neg teq(a, b)$).

{end of example}

Only distinctive states linked by turn actions can be grouped into a topological place (Axiom 30). Under incomplete information this constraint could imply the existence of more places than the ones needed in a map.

Example 6. Consider the set of experiences E indicated by the formulae

$$\langle a, travel, b \rangle, \langle b, turnAround, c \rangle, \langle c, travel, d \rangle ,$$

in addition to the views associated with the distinctive states. Moreover, assume that views uniquely distinguish the different distinctive states. The model for $TT(E)$ is presented in Figure 8c. The model has three places and one path. Not having a *turn* action relating d and a prevents the agent from grouping these distinctive states into the same place, as suggested in Figure 8b. Next we show why this is the case.




Figure 8: (a) The agent navigates a rectangle environment getting the experiences $\langle a, travel, b \rangle$, $\langle b, turnAround, c \rangle$, and $\langle c, travel, d \rangle$. The corresponding topological map has three places and one path (c) rather than two places and one path (b). Distinctive states a and d cannot be grouped into the same topological place since they are not linked by turn actions. Notice that the order of places in the path is not total. Should the agent turn around and experience the schema $\langle d, turnAround, a \rangle$, it will consider (b) as the topological map and disregard (c).

Since views uniquely distinguish distinctive states, then $teq(x, y) \equiv x = y$. From the definition of $turn_eq$ (Block 31), it follows then that $turn_eq = \widehat{turn}$. Since the only *turn* action mentioned in E is the one in schema $\langle b, turnAround, c \rangle$, we deduce that $\widehat{turn}(ds, ds') \equiv [ds = ds' \vee \{ds = b \wedge ds' = c\} \vee \{ds = c \wedge ds' = b\}]$. In particular, $\neg turn_eq(a, d)$. In virtue of Axiom 30 we cannot conclude that a and d are at the same topological place.

{end of example}

The next example shows the interplay between *teq* and *along* as well as the effect of maximizing *teq*.

Example 7.

Consider the set of schemas $\langle a, turnRight, b \rangle$, $\langle b, travel, c \rangle$, $\langle c, turnAround, d \rangle$, $\langle d, travel, e \rangle$, $\langle e, turnRight, a' \rangle$, $\langle a', turnRight, b' \rangle$ consistent with an agent going from one four-way intersection to another (Figure 9). Let's consider the models of these schemas. From our axioms, at least one path and three places must exist:




Figure 9: The agent moves back and forth from one intersection to the other. The second time the agent visits distinctive states a and b , it gives the names a' and b' . From the topological theory it follows that these names correspond to the previously visited a and b .

Places	Paths	Along	teq
$P = \{a, b\}$	Pa: b-c d-e	$\text{along}(b, Pa, \text{dir})$ $\text{along}(c, Pa, \text{dir})$	$\neg \text{teq}(a, b), \neg \text{teq}(c, d)$
$Q = \{c, d\}$		$\text{along}(d, Pa, -\text{dir})$ $\text{along}(e, Pa, -\text{dir})$	$\neg \text{teq}(e, a'), \neg \text{teq}(a', b')$
$R = \{e, a', b'\}$			

We know that $P \neq Q$ and $Q \neq R$. By having $\text{teq}(a, a')$, we can complete the model such that $P = R$. The maximization of teq will force the model to have $\text{teq}(b, b')$. By including Axiom 36 in the Block 33 we are allowed to have a model in which $\text{teq}(b, b')$ is the case. Notice that a travel action has not been performed at b' and so the schemas do not support a tuple of the form $\text{along}(b', \bullet, \bullet)$. *{end of example}*

Example 8. Consider the extension of the previous example where the schema $\langle b', \text{travel}, c' \rangle$ is obtained. Axiom 46 requires the agent to have traveled among places on the same path. As for places, we check this requirement “modulo” teq , since teq plays the role of equality in our theory. In this example, the agent concludes that $\text{teq}(c, c')$. Notice that $\neg \text{travel}(b, c')$ and $\text{travel_eq}(b, c')$ are the case. *{end of example}*

By requiring the agent to have traveled among the places on a same path (Axiom 46), different paths can be identified. The next example illustrates the case.

Example 9. Suppose the agent explores the environment depicted in Figure 10a




Figure 10: By requiring the agent to have traveled among the places on a same path (Axiom 46), different paths can be identified. (a) The agent visits the different distinctive states in the order $ds1, ds2, \dots, ds6$. (b) depicts the topological map associated with (a). Three paths instead of only two are required to explain the agent experiences (see text).

obtaining the following schemas:

$$\langle ds1, \text{travel}, ds2 \rangle \langle ds2, \text{turnRight}, ds3 \rangle \langle ds3, \text{travel}, ds4 \rangle \\ \langle ds4, \text{turnLeft}, ds5 \rangle \langle ds5, \text{travel}, ds6 \rangle$$

We assume that the agent associates different views with the different distinctive states in the example. Axiom 29 implies that there exist places A, B, C and D (see Figure 10b) such that

$$at(ds1, A), at(ds2, B), at(ds3, B), at(ds4, C), at(ds5, C), at(ds6, D) .$$

Moreover, Corollary 1 implies that $A \neq B$, $B \neq C$, $C \neq D$. Under our assumption that all distinctive states in the example have different views, it follows that $teq(ds_1, ds_2) \equiv ds_1 = ds_2$ and thus $\widehat{turn} = turn_{eq}$. Since $\neg\widehat{turn}(ds1, ds3)$, $\neg\widehat{turn}(ds1, ds5)$ and $\neg\widehat{turn}(ds2, ds6)$ are the case, A, B, C and D are all different. Axiom 34 implies that there exist paths $Pa, Pa1, Pa2$, and directions $dir, dir1, dir2$, such that:

$$order(Pa, dir, A, B), along(ds1, Pa, dir), along(ds2, Pa, dir),$$

$$order(Pa1, dir1, B, C), along(ds3, Pa1, dir1), along(ds4, Pa1, dir1),$$

$$order(Pa2, dir2, C, D), along(ds5, Pa2, dir2), along(ds6, Pa2, dir2) .$$

Schemas $\langle ds2, turnRight, ds3 \rangle$ and $\langle ds4, turnLeft, ds5 \rangle$, and Axiom 39 implies that $Pa \neq Pa1$, $Pa1 \neq Pa2$. Since $teq(ds_1, ds_2) \equiv ds_1 = ds_2$ and there is not $turnAround$ schemas in E , then $\widehat{travel} = travel_{eq}$. Consequently $\neg\widehat{travel}(ds1, ds4)$ and $\neg\widehat{travel}(ds1, ds5)$ are the case, and in virtue of Axiom 46 it follows that $Pa \neq Pa2$. *{end of example}*

Example 10. Consider the same T-environment exploration presented in Example 4 (Page 10) (see Figure 11). When using only causal information, three possible models are associated with the exploration. When using topological information, only one of these models is possible as illustrated next.




Figure 11: The agent visits the different distinctive states in the order suggested by their numbers. The same travel action ml is performed when traveling from a corner to the intersection (i.e. $\langle 1, ml, 2 \rangle$) and viceversa (i.e. $\langle 4, ml, 5 \rangle$). A turn around action is performed when reaching a corner (i.e. $\langle 3, turnAround, 4 \rangle, \langle 7, turnAround, 8 \rangle$, etc.). Assume that the different corners have the same views (i.e. $view(1) = view(4) = view(8)$, $view(3) = view(7) = view(11)$), and views associated with the other distinctive states are different. Three different causal models can be associated with the agent exploration of this T-environment but only one of them is consistent with topological information (see text).

The three causal models associated with T-environment are:

1. Model 1: $ceq(8, 12), ceq(12, 8), ceq(x, x)$.²⁵

²⁵The extent of ceq in model 1 is defined by $\{(8, 12), (12, 8)\} \cup \{(x, x) : x = 1, \dots, 12\}$.

2. Model 2: $ceq(1, 12), ceq(12, 1), ceq(x, x)$.
3. Model 3: $ceq(4, 12), ceq(12, 4), ceq(3, 11), ceq(11, 3), ceq(2, 10), ceq(10, 2), ceq(x, x)$.

We are to show that only model 3 above is consistent with topological information. For this we show the following three facts: (i) any model must have at least 2 tpaths and 5 tplaces,²⁶ (ii) there is a model with 2 tpaths and 5 tplaces (this is the intended model), (iii) a model of $\neg teq(2, 10)$ must have at least 6 tplaces. This last statement implies that models 1 and 2 above are not consistent with topological information.

From $\langle 1, travel, 2 \rangle$ and $\langle 2, travel, 3 \rangle$, Corollary 2 implies that there exist a path $Pa1$ and direction $dir1$ such that

$$along(1, Pa1, dir1), along(2, Pa1, dir1), along(3, Pa1, dir1) .$$

Moreover, Corollary 1 implies that

$$place(1) \neq place(2), place(2) \neq place(3), place(1) \neq place(3) .$$

From $\langle 3, turnAround, 4 \rangle, \langle 4, travel, 5 \rangle$, Axiom 35 and Corollary 2, it is the case that $along(4, Pa1, -dir1), along(5, Pa1, -dir1)$. Similarly, from $\langle 5, turnLeft, 6 \rangle, \langle 6, travel, 7 \rangle, \langle 7, turnAround, 8 \rangle, \langle 8, travel, 9 \rangle$ we conclude that there exist a path $Pa2$ and direction $dir2$ such that $Pa1 \neq Pa2$ (Axiom 39) and

$$place(5) \neq place(8), along(6, Pa2, dir2), along(7, Pa2, dir2), \\ along(8, Pa2, -dir2), along(9, Pa2, -dir2) .$$

From $\langle 9, turnRight, 10 \rangle, \langle 10, travel, 11 \rangle, \langle 11, turnAround, 12 \rangle$, there exist path $Pa3$ and direction $dir3$ such that $Pa2 \neq Pa3$ and

$$along(10, Pa3, dir3), along(11, Pa3, dir3), along(12, Pa3, -dir3) .$$

Theorem 3 allow us to conclude that $place(5) \notin \{place(1), place(2), place(3)\}$. The same argument shows that $place(8) \notin \{place(1), place(2), place(3), place(5)\}$. Consequently, a minimal model of the theory must have at least two tpaths and five tplaces.

Notice that in the intended model of the T-environment, $Pa1 = Pa3, dir1 = dir3, teq(2, 10), teq(3, 11)$ and $teq(4, 12)$. This model is indeed a model of $TT(E)$ since at least two topological paths and five topological places are needed to explain E , and consequently any model must have two topological paths and five topological places (Theorem 4).

If $\neg teq(2, 10)$ were the case, then Theorem 3 allows to conclude that $place(9) \notin \{place(1), place(2), place(3), place(5), place(8)\}$ and so the model will have at least six tplaces. Consequently $teq(2, 10)$ has to be the case in a minimal model of the theory. *{end of example}*

Example 11. Consider an extension of the previous example where we have the additional schemas $\langle 9, turnLeft, 5' \rangle, \langle 5', turnRight, 9 \rangle$. In this case, the intended model has *four* places and two paths. Notice that now the agent can conclude that $place(5) = place(2)$ by making $teq(5', 5)$ and so $turn_eq(5, 2)$. *{end of example}*

²⁶Since there is not a turn action between dstates $\{5, 6\}$ and dstates $\{2, 9, 10\}$, these dstates are not at the same topological place, as suggested by Figure 11.




Figure 12: The environment in (a) illustrates a case where different paths intersect at more than one place. Suppose the agent explores the environment by visiting the different distinctive states in the order $ds1, ds2, ds1, ds3, ds4, ds3, ds6, ds7, ds4, ds5, ds2$. (b) depicts the topological map associated with this environment.

The theory does not assume a “rectilinear” environment where paths intersect at most in one place. Consider the next example.

Example 12. Suppose the agent explores the environment depicted in Figure 12 obtaining the following schemas:

$$\begin{array}{lll}
 \langle ds1, turnAround, ds2 \rangle & \langle ds2, turnAround, ds1 \rangle & \langle ds1, travel, ds3 \rangle \\
 \langle ds3, turnRight, ds4 \rangle & \langle ds4, turnLeft, ds3 \rangle & \langle ds3, travel, ds6 \rangle \\
 \langle ds6, turnLeft, ds7 \rangle & \langle ds7, travel, ds4 \rangle & \\
 \langle ds4, turnRight, ds5 \rangle & \langle ds5, travel, ds2 \rangle &
 \end{array}$$

We assume that views uniquely distinguish the different distinctive states. From Corollary 1 there exist the different places A, B , and C suggested in the figure. In addition, Corollary 2 implies the existence of a path, Pa , and direction, say pos , such that $order(Pa, pos, A, B)$, $order(Pa, pos, B, C)$, $order(Pa, pos, A, C)$. Moreover, from schemas $\{\langle ds7, travel, ds4 \rangle, \langle ds5, travel, ds2 \rangle\}$ and Axiom 34, there exist paths $Pa1$, $Pa2$, and directions $dir1$, $dir2$, such that

$$\begin{array}{l}
 order(Pa1, dir1, C, B) \wedge along(ds7, Pa1, dir1) \wedge along(ds4, Pa1, dir1) , \\
 order(Pa2, dir2, B, A) \wedge along(ds5, Pa2, dir2) \wedge along(ds2, Pa2, dir2) .
 \end{array}$$

Since $along(ds6, Pa, pos)$, from Axiom 39 and schema $\langle ds6, turnLeft, ds7 \rangle$ we conclude that $Pa \neq Pa1$. Since we are minimizing paths, by setting $Pa2 = Pa$ and $dir2 = neg$, we obtain a minimal model for E . Notice that in this model, places B and C belong to two different paths, Pa and $Pa1$. *{end of example}*

There are some patterns of experience in which our theory is not applicable. In particular, Axiom 44 rules out circular paths and Axiom 37 rules out experiences where different paths merge into the same distinctive state. In Section 5.3 we extend the topological theory to deal with these type of paths.

There are patterns of experience that can be explained by more than one topological map. This happens because the experiences are incomplete, or the agent’s sensors are weak.²⁷

²⁷Since the positive and negative direction of a path are chosen arbitrarily (Axiom 34), there is not a unique minimal model for $TT(E)$. Given any model M of $TT(E)$ one could define another model M' of $TT(E)$ by choosing a path pa in M and reversing the roles of the directions pos and neg for pa . We will consider these “up to path direction isomorphic” models to be the same.




Figure 13: (a) The agent goes around the block visiting places A, B, \dots, F, C in the order suggested in the figure. Intersections B and C look alike to the agent. (b) and (c) represent two possible representations for the environment in (a). Topological information is not enough to decide whether the agent is back to B or C .

Example 13. Assume that the agent visits places A, B, C, D, E, F, C in the order suggested by Figure 13. Assume also that intersections look alike. In particular, places B and C look alike. Given this information, the agent is not able to decide whether it is back to B or C and consequently two minimal models can be associated with the set of experiences in this environment (Figures 13b,c).

Notice that if the agent accumulates more information, by turning at C and traveling to D , then it can deduce that the topology of the environment is the one in Figure 13b. This is the case since the views at C and D are different.²⁸ *{end of example}*

5.3 Coping with self intersecting paths

The topological theory presented in the previous section is adequate for representing environments where “complex” paths configurations do not occur. In particular, we assume that self-intersecting and convergent paths do not exist (see Figure 14). In this section we extend our theory to deal with these types of paths. Convergent paths are




Figure 14: (a) Self intersecting paths. (b) Convergent paths.

the standard counterexample for the axiom stating that distinctive states are along a unique path (Axiom 37). We replace Axiom 37 by the block

$$\begin{array}{l}
 \{ \text{min convergent_paths} : \\
 \quad [\text{along}(ds, pa, dir) \wedge \text{along}(ds, pa1, dir1) \\
 \quad \quad \wedge \neg [pa = pa1 \wedge dir = dir1]] \rightarrow \text{convergent_paths}(pa, pa1) \\
 \}
 \end{array}$$

²⁸Metrical information can be used to deduce the correct topology. See Example 17 (Page 35).

Self-intersecting paths are the standard counterexample for the axioms stating that turning changes the path (Axiom 39), at a place there is at most one distinctive state along a path direction (Axiom 38), and the order of places in a path is not reflexive (Axiom 44). We replace these axioms by the block

```
{min self_intersecting :
  order(pa, dir, p, p) → self_intersecting(pa),
  [(ds, turn_desc, ds') ∧ turn_desc ≠ TurnAround ∧ along(ds, pa, dir)
   ∧ along(ds', pa, dir1)] → self_intersecting(pa),
  [at(ds1, p) ∧ at(ds2, p) ∧ along(ds1, pa, dir) ∧ along(ds2, pa, dir)
   ∧ ¬teq(ds1, ds2)] → self_intersecting(pa)
}
```

While we have defined *convergent* and *self-intersecting* paths, we still need to state that by default these kind of paths do not exist. This is accomplished by giving priority to the minimization of these two predicates over any other predicate. The new circumscription policy associated with our theory becomes

circ $self_intersecting \succ convergent_paths \succ tpath \succ tplace$ **var** $SSH\vec{pred}$. (49)

The new theory is a *conservative* extension of our previous theory, since *any topological map with respect to our previous theory is a topological map according to the new theory*. In particular, the maps associated with examples 5 through 12 are still valid maps for the new theory. Next we study some cases we could not handle before.

Example 14. Suppose the agent has experienced the following schemas (see Figure 15):


$$\begin{array}{lll} \langle b, travel, d \rangle & \langle d, turnAround, c \rangle & \langle c, turnRight, e \rangle \\ \langle e, travel, a \rangle & \langle a, turnAround, b \rangle & \end{array}$$


Figure 15: Distinctive state a is along two different paths. These two paths are declared convergent paths in the model of our theory.

From Axiom 34 we know that exist paths Pa , $Pa1$ and directions dir , $dir1$ such that $along(b, Pa, dir)$, $along(d, Pa, dir)$, $along(e, Pa1, dir1)$ and $along(a, Pa1, dir1)$ are the case. Moreover, from Axiom 35 it follows that $along(b, Pa1, -dir1)$. We have two possible models for these schemas:

- **Model 1.** In this model $Pa \neq Pa1$. Consequently, $self_intersecting = false$ and $convergent_paths(Pa, Pa1)$ are the case.

- **Model 2.** In this model $Pa = Pa1$. Consequently, $self_intersecting(Pa)$ and $convergent_paths = false$ are the case.

We prefer model 1 over model 2 according to the circumscription policy 49. *{end of example}*

In order to compare any two possible maps, these maps must have a common sort of places and paths. Since a map can be arbitrarily large, no finite domain can be adequate and so we require the sorts of places and paths to be infinite (see block $TT(E)$, Page 13). Consider the following example (see also Theorem 4, Page 18).

Example 15. Consider the schema $\langle a, travel, b \rangle$ where a and b have the same view. The intended model has one topological path and two topological places. One expects that the path is not circular (self-intersecting), and so the existence of two places. However, without requiring the existence of enough places, the following model is also possible:

places = {A},	tplace = {A}	paths = {Pa},	tpaths = {Pa}
teq(a,b)		$self_intersecting(Pa)$	
at(a,A)	at(b,A)	along(a,Pa,pos),	along(b,Pa,pos)
		order(Pa,pos,A,A)	

In this model, $self_intersecting(Pa)$ must be the case, since the universe of places only has one place. Notice that when comparing two models according to the circumscription policy 49, the universe of $paths$ and $places$ in the models has to be the same. One can vary the interpretation of $tpath$, $tplace$, and so on, but **not** the universe of $paths$ and $places$. The model above is ruled out by requiring the universe of $places$ to have enough places. *{end of example}*

6 Boundary Regions

Topological paths play the role of *streets* in a city layout map. Streets are often used as a reference for specifying the location of a given place: a place will be either on the given street or in one of the “two sides” –left or right– of the street.

Mathematically, the concept of left and right of a topological path is related to the topological one of the interior and exterior of a curve. While not all curves have a well defined interior and exterior (for example, consider a spiral, or a fractal curve), closed not self-intersecting curves –*Jordan curves*– do have associated interior and exterior sets [Beardon, 1979].²⁹ Moreover, in order to go from the interior to the exterior (or vice versa) of the curve γ , one has to cross γ . Our analogy of topological paths and mathematical curves breaks down because in general the agent might be able to travel from one side of the path to the other without crossing the path. This can happen because of the agent’s inability to detect that it has crossed the path, or (more often) because paths are not long enough to divide the environment into two regions (for example, consider a dead-end street).

²⁹When the curve is removed, the plane is divided into two disjoint connected sets.

In order to determine boundary relations -the location of a place with respect to a path- we formally state the following heuristic. Suppose the agent is at an intersection on a given path, and it then turns right. If the agent now travels, any place it finds while traveling with no turns will be on the right of the starting path. While this heuristic draws the correct conclusion in a rectilinear environment, it may draw incorrect conclusions when paths are not straight. Consequently, we state our heuristic as a ‘defeasible’ rule so as not to conclude a boundary relation when inconsistent sources of information exist (see Figure 16).




Figure 16: Different environments illustrating how our default to determine boundary relations work. In (a) we conclude by default that place C is to the left of the path from A to B. In (b) we conclude nothing about the location of place D with respect to this path. In (c) we conclude that place C is to the left of the path from A to B. This is the case since there is no information to conclude otherwise.

TurnRight and *turnLeft* actions are used to define the relative orientation between paths at a given place (Section 6.1), relations that are then used to infer whether a place is on the left or the right of a given path (Section 6.2). The boundary relations inferred by an agent may not be complete: the agent does not necessarily know the location of each place with respect to each path. Nevertheless, the boundary relations inferred by the agent are useful to distinguish places otherwise not distinguishable by the topological maps as described so far (see Example 16, page 29).

6.1 Qualitative orientation of paths at a place

We extend the topological level in order to represent the relative orientation among paths that intersect at a given place. We use the predicates

totheLeftOf(*p*, *pa*, *dir*, *pa1*, *dir1*) , **totheRightOf**(*p*, *pa*, *dir*, *pa1*, *dir1*)

to represent the facts that (i) *p* is a *place* on both paths, *pa* and *pa1*, and (ii), when the agent is at *place p* facing on the direction *dir* of *pa*, after executing a turn left (right) action, the agent will be facing on the direction *dir1* of *pa1* (see Figure 17, Page 29).

The predicates *totheLeftOf* and *totheRightOf* are derived from the actions performed by the agent at a place:

$$\begin{aligned}
& \{ \text{min totheRightOf}, \text{min totheLeftOf} : & (50) \\
& \quad [\langle ds, \text{turnRight}, ds1 \rangle \wedge \text{at}(ds, p) \wedge \text{along}(ds, pa, dir) \wedge \text{along}(ds1, pa1, dir1)] \\
& \quad \quad \rightarrow \text{totheRightOf}(p, pa, dir, pa1, dir1), \\
& \quad [\langle ds, \text{turnLeft}, ds1 \rangle \wedge \text{at}(ds, p) \wedge \text{along}(ds, pa, dir) \wedge \text{along}(ds1, pa1, dir1)] \\
& \quad \quad \rightarrow \text{totheLeftOf}(p, pa, dir, pa1, dir1). \\
& \}
\end{aligned}$$

6.2 Left and Right of a path

A path has associated two regions: the places to the left of the path and the places to the right of the path. We use the predicates **leftOf**(**pa**, **dir**, **lr**) and **rightOf**(**pa**, **dir**, **rr**) to denote that *region lr (rr)* is the left (right) region of path *pa* with respect to the path's direction *dir*. The properties of these predicate are as follows:

$$\exists!lr \{ \text{leftOf}(pa, dir, lr) \}, \quad \exists!rr \{ \text{rightOf}(pa, dir, rr) \} \quad (51)$$

$$\text{leftOf}(pa, dir, r) \equiv \text{rightOf}(pa, -dir, r) \quad (52)$$

$$\{ \text{min is_region} : \text{LeftOf}(pa, dir, lr) \rightarrow \text{is_region}(lr) \} \quad (53)$$

$$\text{leftOf}(pa, dir, lr) \wedge \text{leftOf}(pa1, dir1, lr) \rightarrow pa = pa1 \quad (54)$$

Axiom 51 states the existence and uniqueness of a path's left/right regions. The domain of *leftOf* is restricted by Block 53 and Axiom 54. Since left/right regions of a path interchange when changing the path direction (Axiom 52), constraining the domain of *leftOf* imposes similar constraints on the domain of *rightOf*.

We use the predicate **in_region**(**p**,**r**) to indicate that *place p* is in *region r*. The domain of *in_region* is constrained by Axiom 55. The properties of *in_region* are defined in Block 56. A path has associated three disjoint set of places: the places on the path, and the places to the left/right of the path (Axioms 58 and 59).³⁰ Boundary relations are derived according to Axiom 60 and 61 (see Figure 17):

$$\text{in_region}(p, r) \rightarrow \text{is_region}(r), \quad (55)$$

$$\{ \text{min in_region} : \quad (56)$$

$$\{ \text{in_region} : \quad (57)$$

$$\text{on}(pa, p) \wedge \text{leftOf}(pa, dir, lr) \rightarrow \neg \text{in_region}(p, lr), \quad (58)$$

$$[\text{leftOf}(pa, dir, lr) \wedge \text{rightOf}(pa, dir, rr) \wedge \quad (59)$$

$$\text{in_region}(p, lr)] \rightarrow \neg \text{in_region}(p, rr),$$

$$[\text{totheRightOf}(p1, pa, dir, pa1, dir1) \wedge \text{order}(pa1, dir1, p1, p) \wedge \quad (60)$$

$$\text{rightOf}(pa, dir, rr) \wedge \neg \mathbf{Ab}(\mathbf{pa}, \mathbf{p})] \rightarrow \text{in_region}(p, rr),$$

$$[\text{toTheLeftOf}(p1, pa, dir, dir1, pa1) \wedge \text{order}(pa1, dir1, p1, p) \wedge \quad (61)$$

$$\text{leftOf}(pa, dir, lr) \wedge \neg \mathbf{Ab}(\mathbf{pa}, \mathbf{p})] \rightarrow \text{in_region}(p, lr)$$

$$\} \\ \}$$

³⁰The symmetry between *leftOf* and *rightOf* defined by Axiom 52 let us write our axioms in terms of only one of these predicates.




Figure 17: Path $Pa1$ is to the right of path Pa at place $p1$. Place p is after place $p1$ on path $pa1$. By default, we conclude that place p is to the right of path pa .

The outer block 56 states that normally boundary relations are false. This is the case since by default the agent does not know the location of a place with respect to a given path. The inner block 57 states under what conditions the agent can derive a boundary relation. For instance, according to Axiom 60, if at place $p1$ path $pa1$ is to the right of path pa , and place p is after $p1$ on path $pa1$, then normally it is the case that p is on the right of pa (see Figure 17).

6.3 Adding boundary relations to the topological map

While Block 56 defines the extent of the predicate in_region , it does not express our preference for models where in_region is maximal. We update the topological theory by including Axioms 50-61 inside the block **AT_block** (see Section 5.2, Page 14), and the new circumscription policy becomes

$$\mathbf{circ} \ tpath \succ \neg in_region \succ tplace \ \mathbf{var} \ newS\vec{S}Hpred$$

where $newS\vec{S}Hpred$ stands for the tuple of predicates

$$\langle at, along, order, on, teq, turn_eq, travel_eq, \\ \mathbf{totheRightOf}, \mathbf{totheLeftOf}, \mathbf{leftOf}, \mathbf{rightOf}, \mathbf{is_region} \\ \rangle.$$

The circumscription policy states that Axioms 60 and 61 should be used to draw conclusions even at the expense of having more places on the map. This is achieved by maximizing in_region over $tplace$ in the circumscription policy. By maximizing the extent of in_region at the expense of having possibly more places, boundary relations determine distinctions among environment states that could not be derived from the connectivity of places alone. The next example illustrates the case.

Example 16. Consider an agent visiting the different corners of a square room in the order suggested by Figure 18a. In addition, suppose the agent's sensory apparatus allows it to define $views$ by characterizing the direction of walls and open space. Accordingly, the agent experiences *four* different views, $v1-v4$, in this environment.

The agent's experiences, E , in this environment are:

$$View(ds1, v1), View(ds2, v2), View(ds3, v1), View(ds4, v2), View(ds5, v1), \\ \langle ds1, turnRight, ds2 \rangle, \langle ds2, travel, ds3 \rangle, \langle ds3, turnRight, ds4 \rangle, \langle ds4, travel, ds5 \rangle.$$




Figure 18: (a) Sequence of actions followed by an agent while navigating a square room. Starting at distinctive state $ds1$, distinctive states are visited in the order suggested by their number. (b) and (c) depict the resulting topological map without and using boundary regions, respectively.

Suppose that the agent does not use boundary regions when building the topological map. Then the minimal topological model associated with E has two paths³¹ and two places. In this model, $teq(ds1, ds5)$ is the case. The environment looks perfectly symmetric to the agent (Figure 18b).!!

Suppose now that the agent relies on boundary regions. Let P , Q , R , be the topological places associated with $ds1$, $ds3$ and $ds5$ respectively. From Axiom 34, let Pa , Pb , dir_a and dir_b be such that

$$\begin{aligned} &order(Pa, dir_a, P, Q), \text{ along}(ds2, Pa, dir_a), \text{ along}(ds3, Pa, dir_a), \\ &order(Pb, dir_b, Q, R), \text{ along}(ds4, Pb, dir_b), \text{ along}(ds5, Pb, dir_b), \end{aligned}$$

are the case. From Block 50 we conclude that $totheRightOf(Q, Pa, dir_a, Pb, dir_b)$. In the proposed model, the extent of in_region is maximized by declaring $Ab = false$ inside Block 57 and consequently from Axiom 60 we conclude $in_region(R, right(Pa, dir_a))$ where $right(Pa, dir_a)$ denotes the right region of Pa when facing dir_a .³² Moreover, from Block 56 we deduce $in_region(p, r) \equiv [p = R \wedge r = right(Pa, dir_a)]$. Finally, from Axiom 58 we conclude $P \neq R$ since $on(Pa, P)$ is the case. The resulting topological map is depicted in Figure 18c.

{end of example}

Boundary relations are in general not enough to distinguish different environment states. This is the case when the agent has weak sensors, the environment is symmetric, or the agent's experiences are incomplete. The use of local metrical information could help on those cases although metrical uncertainty could render this extra information useless. We discuss this issues in the next section.

7 Using local metrical information

Action executions have associated metrical information representing the observed magnitude of the action. For instance, after traveling the agent may have an estimate of the

³¹From $\langle ds3, turnRight, ds4 \rangle$ and Axiom 39 we can deduce that $Pa \neq Pb$ in Figure 18b.

³²Axiom 51 guarantees that there exists a unique such region.

distance between the “end places” of the travel action, and after turning, the agent may have an estimate of the angle turned.³³ This local metrical information is integrated into frames of reference associated with topological objects:

- Each path has associated a one dimensional frame of reference which assigns a position to each place in the path.
- Each place has associated a radial frame of reference which assigns a heading (angle) to each path the place belongs to.
- Regions or places might have associated two dimensional frames of reference which assign real valued tuples to certain places.³⁴

As *positions* and *headings* are derived from noisy data, there is uncertainty associated with their real values. Different representations for this uncertainty are possible: *intervals*, *probability distribution functions*, etc. As the agent repeatedly navigates among the same places and paths, new measure estimates are taken into account to update the uncertainty associated with positions and headings.³⁵ In our current work we use *intervals* to represent uncertainty in position and headings.³⁶ We use the predicate `action_execution(s, Int)` to state that the *interval Int* represents an estimate of the metrical information about the execution of the action associated with schema *s*. We use the notation $\langle \mathbf{ds}, (\mathbf{type\ Int}), \mathbf{ds}' \rangle$, where *type* is *travel* or *turn*, as an abbreviation for the formula

$$\exists s, a \{CS(s, ds, a, ds') \wedge action_type(a, type) \wedge action_execution(s, Int)\} .$$

How the estimates are to be interpreted depends on the type of action (turn or travel) the schema refers to. In the next sections we will describe how to do so.

7.1 One dimensional frames of reference

A path has associated a one dimensional frame of reference which assigns a location to each place on the path. This location is a real number,³⁷ representing the “distance” with respect to an arbitrary but fixed place on the path. The distance among places on a path are derived from estimates acquired when traveling among places on the path. These estimates have to be *consistent* so that positions can be associated with places. Next we formalize these ideas.

³³Notice that different kind of metrical estimates could be associated with a travel or turn action. For example, the agent could measure the arc length associated with a travel action. In addition, it could measure the minimum distance to an object on the left and the right sides at each point along the trajectory associated with a travel action [Kuipers, 2000].

³⁴Local analog maps [Elfes, 1987, Borenstein and Koren, 1991, Thrun, 1998] can also be associated with places [Kuipers, 2000].

³⁵For example, Kuipers and Byun [Kuipers and Byun, 1988, Kuipers and Byun, 1991] use a scalar Kalman-filter in order to update the distance between two places in a path.

³⁶However, the discussion in this section applies to other forms of representing uncertainty.

³⁷This real value represents a quantity whose magnitude is derived by the robot while navigating the environment. The units of this quantity can be *meters*, *feet*, or *number of wheel rotations*. Hereafter, we assume that all quantities are given in the same units.

The position of a place on a path is represented by the predicate **position1**(**path**, **place**, **position**). Positions along a path are unique and only assigned to places belonging to the path:

$$position1(pa, p, pos) \wedge position1(pa, p, pos') \rightarrow pos = pos' , \quad (62)$$

$$position1(pa, p, pos) \rightarrow on(pa, p) . \quad (63)$$

The distance between two places in a path is defined as the absolute value of the difference between their corresponding positions on the path. The predicate **path_distance**(**pa**, **p**, **q**, **d**) represents the fact that the distance between places p and q on path pa is d . The predicate *path_distance* is defined as follows:

$$path_distance(pa, p, q, d) \equiv \quad (64)$$

$$\exists pos_p, pos_q \{ position1(pa, p, pos_p) \wedge position1(pa, q, pos_q) \wedge d = |pos_p - pos_q| \} .$$

Estimates of the distance between places on a path are gathered while the agent navigates the environment. The predicate **path_distance**[~](**pa**, **p**, **q**, **I_d**) represents the fact that the closed interval I_d is an estimate of the distance between places p and q on path pa . Distance estimates are derived from experiences of the robot in the environment. Distance estimates are “*compounded*” to derive new estimates from known ones. Formally,

$$\{ \begin{aligned} &min\ path_distance^{\sim} : \\ &[\langle ds, (travel\ I_d), ds' \rangle \wedge at(ds, p) \wedge at(ds', q) \wedge along(ds, pa, dir) \wedge \\ &\quad along(ds', pa, dir)] \rightarrow path_distance^{\sim}(pa, p, q, I_d) , \\ &[order(pa, dir, p, q) \wedge order(pa, dir, q, r) \wedge path_distance^{\sim}(pa, p, q, I_{pq}) \wedge \\ &\quad path_distance^{\sim}(pa, q, r, I_{qr})] \rightarrow path_distance^{\sim}(pa, p, r, I_{pq} + I_{qr}) \\ &\} \end{aligned} \quad (65)$$

$$\quad (66)$$

where the addition of intervals is defined in the usual way: $[a, b] + [c, d] = [a + c, b + d]$. Finally, distance estimates are “*merged*” in order to have the “best” estimate associated with a distance. The predicate **path_distance**[⊗](**pa**, **p**, **r**, **I_d**) denotes the merging of distance estimates:³⁸

$$path_distance^{\otimes}(pa, p, r, I) \equiv_{def} I = \cap \{ I_{est} : path_distance^{\sim}(pa, p, q, I_{est}) \} . \quad (67)$$

The distance between places on a path must be *compatible* with all of its estimates. Formally,

$$path_distance^{\otimes}(pa, p, q, I_d) \rightarrow \exists d \in I_d\ path_distance(pa, p, q, d) . \quad (68)$$

When the agent has distance estimates available, **path_distance**[⊗](**pa**, **p**, **q**, **I_d**) is always the case for some interval I_d . If these estimates are consistent then $I_d \neq \emptyset$. Moreover, when the topological map is consistent with metrical information, it should be possible to assign locations to places on a path as specified by Axiom 64.

³⁸The operations of compounding and merging are standard operations in order to propagate uncertainty about the real value of a variable [Smith and Cheeseman, 1986]. They take different forms depending on how one represents uncertainty as well as on the dimensionality of the variables’ domains. In our case, these operations take the form of adding and intersecting intervals, respectively.

7.2 Radial frames of reference

Each place has a local frame of reference w.r.t. which path headings are associated. This information is represented by the predicate $\mathbf{radial}(\mathbf{p}, \mathbf{pa}, \mathbf{dir}, \mathbf{h})$ denoting the fact that *when the agent is located at place p, path pa could be followed in direction dir by facing the heading h w.r.t. the radial frame of reference local to p*. Headings take values in $[0, 2\pi)$. The formalization of radial frames of reference follows the same steps as for one dimensional frames of reference and will not be presented here.³⁹

7.3 Two dimensional frames of reference

A topological map does not explicitly represent the distance or direction between two arbitrary places. In order to do so, distances between places on a path as well as the angles between paths at a place must be combined. We use the predicate $\mathbf{location2}(\mathbf{p}, \mathbf{q}, \mathbf{l})$ to indicate that the location of place q with respect to the two dimensional frame of reference associated with place p is l (a real valued pair). We do not restrict what places are assigned locations with respect to a given two dimensional frame of reference.

When restricted to environments with “straight” paths, it is possible to state when a two dimensional frame of reference is *compatible* with the actual experiences of the robot. The next axioms state this requirement:⁴⁰

$$\begin{aligned} & location2(p, p1, l_{p1}) \wedge location2(p, p2, l_{p2}) \wedge path_distance^{\otimes}(pa, p1, p2, I_d) \\ & \rightarrow |l_{p1} - l_{p2}| \in I_d . \end{aligned} \quad (69)$$

$$\begin{aligned} & [location2(p, p1, l_{p1}) \wedge location2(p, p2, l_{p2}) \wedge location2(p, p3, l_{p3}) \wedge \\ & order(pa, dir, p1, p2) \wedge order(pa', dir', p2, p3) \wedge angle^{\otimes}(p2, pa, dir, pa', dir', I_{ang})] \\ & \rightarrow angle(-l_{p2}\vec{l}_{p1}, l_{p2}\vec{l}_{p3}) \in I_{ang} , \end{aligned} \quad (70)$$

where $angle(\vec{v}, \vec{w})$ denotes the angle in $[0, 2\pi)$ from vector \vec{v} to vector \vec{w} .

Axioms 69 and 70 assume that paths are straight. In order to deal with more general paths, one should include some parameters describing the shape of the path, or at least an estimate of the change in heading while traveling. For instance, in [Kuipers and Levitt, 1988] travel actions were represented as $\langle ds, (travel \ dist \ \Delta\theta), ds' \rangle$, where $dist$ corresponds to the distance between the places associated with ds and ds' , and $\Delta\theta$ corresponds to the change of orientation while traveling. However, there is not a statement of how this extra information is used or whether it suffices to describe appropriate metrical constraints for two dimensional frames of reference. While a more detailed account of the use of metrical information is desirable, including representing

³⁹Estimates of the angle between paths at a place are gathered from *turn* actions. Angle estimates are compounded and merged as we did for distances among places in a path. We use the predicates $\mathbf{angle}(\mathbf{p}, \mathbf{pa}, \mathbf{dir}, \mathbf{pa1}, \mathbf{dir1}, \mathbf{ang})$ -*ang* is the angle the agent will have to turn to face path *pa1* in direction *dir1* when it is at place *p* facing path *pa* in direction *dir*-, $\mathbf{angle}^{\approx}(\mathbf{p}, \mathbf{pa}, \mathbf{dir}, \mathbf{pa1}, \mathbf{dir1}, \mathbf{I}_{ang})$ - *I_{ang}* is an estimate of the angle at place *p* between path *pa* in direction *dir* and path *pa1* in direction *dir1*.

⁴⁰When curved paths are possible, the predicate *path_distance* represents distance *along the path*, not straight-line distance between end point. To handle curved paths, we have to separate those two concepts, or have estimates of both types of “distances”.

and reasoning about a path's shape, we have left this description outside the scope of this work.

7.4 Combining topological and metrical information

While radial and one dimensional frames of reference are associated with any place and path, respectively, there is not a general topological theory asserting when to create a two dimensional frame of reference, what places should be included in a such frame of reference, or how to assign place locations consistent with the estimates of distances and angles gathered by the agent.^{41 42}

In this section we formally state what it means for the topological map to be consistent with a given set of frames of reference. In order to do so, given distinctive states ds, ds_1, \dots, ds_n , we introduce the notation $\langle \mathbf{ds} : \mathbf{ds}_1, \dots, \mathbf{ds}_n \rangle$ to state that the places associated with the different ds_i have a location in the two dimensional frame of reference associated with ds 's place,

Definition 3. Let ds, ds_1, \dots, ds_n be a set of distinctive states. By definition,

$$\begin{aligned} \langle ds : ds_1, \dots, ds_n \rangle &\equiv_{def} & (71) \\ \exists p \left\{ at(ds, p) \wedge \bigwedge_{i=1}^n \exists p_i, l_i [at(ds_i, p_i) \wedge location2(p, p_i, l_i)] \right\} \end{aligned}$$

{end of definition}

By **2D_Frames** we denote the formula specifying any two dimensional frames of reference used by the agent. Without loss of generality, we require two dimensional frames of reference to be specified as in Definition 3. We require any model of the SSH to have only the two dimensional frames of reference specified in *2D_Frames*. In addition, the places belonging to a frame of reference should be only those explicitly stated in 71. These last two requirements can be stated as follows:

$$\{ min\ location2 : 2D_Frames \} \quad (72)$$

The topological theory includes local metrical information by adding Axioms 62 to 72 inside the block **AT_block** (see Page 14). The priority of predicates in the cir-

⁴¹Having a global frame of reference including all places in the map is usually inappropriate since the uncertainty associated with some places' locations in such a frame of reference may not allow the agent to draw useful conclusions. Instead, the agent can have multiple frames of reference as well as relations among the different frames of reference [McDermott and Davis, 1984, Kuipers, 2000]. As the agent explores the environment, new frames of reference are created when the current's location uncertainty with respect to the current frame of reference is larger than a given threshold [Moutarlier and Chatila, 1989, Engelson and McDermott, 1992].

⁴²The problem of assigning locations to places given some metrical constraints can be solved by borrowing methods from different fields. For example, estimation theory tells us how to estimate the true value of a given set of variables given noisy observations of the relations between those variables [Gelb, 1974, Smith and Cheeseman, 1986]. The robotics community has developed algorithms to solve a network of spatial relations [Durrant-Whyte, 1987, Durrant-Whyte, 1988a, Durrant-Whyte, 1988b, Moutarlier and Chatila, 1989]. Techniques from multidimensional scaling [Borg and Groenen, 1997] and nonlinear programming [Peressini *et al.*, 1988] can also be used.




Figure 19: (a) The robot goes around the block visiting distinctive states $ds1$ to $ds11$ in the order suggested by the figure. Distinctive state $ds11$ is observed at the same environment state as $ds4$. Assume distinctive states $ds1$ and $ds4$ look alike to the agent. (b) and (c) represent two possible topological maps for the environment in (a) (see Example 13, Page 24). The model in (c) can be discarded as it is not consistent with the available metrical information. (d) With $\pm 10^\circ$ noise associated with turn actions, the agent cannot use metrical information to discard the environment depicted in (c).

cumscription policy associated with *AT_block* remains the same. The predicates varied in the circumscription policy now include those predicates used to describe metrical information: *radial*, *position1*, *position2*, *path_distance*, *path_distance \approx* , *path_distance \otimes* , *angle*, *angle \approx* and *angle \otimes* .

The next examples illustrate how metrical information is used to disambiguate the topological map.

Example 17. Consider Example 13 (Page 24) where two topological maps are consistent with the agent’s experiences (see Figure 19). Suppose that “perfect” metrical information is available to the agent.

How does the agent figure out that it is back to $ds4$ rather than to $ds1$? As claimed in Example 13 (Page 24) both options $teq(ds4, ds11)$ and $teq(ds1, ds11)$ are topologically possible (Figures 19b,c). However, given the metrical information above, only the assumption $teq(ds4, ds11)$ is a consistent one. To deduce this fact, the agent includes the frame of reference $\langle ds4 : ds1, \dots, ds11 \rangle$ in E , which renders impossible $teq(ds4, ds11)$.

Should the metrical information have been less precise, the agent might not benefit from this extra metrical information. For example, suppose that instead of sharp 90° turn angles, there exists a $\pm 10^\circ$ uncertainty associated with the turn actions above (i.e. consider replacing $\langle ds1, (turn - 90^\circ), ds2 \rangle$ by $\langle ds1, (turn [-110^\circ, -80^\circ]), ds2 \rangle$).⁴³ In this case the agent cannot use metrical information to deduce that it is back to $ds4$ and it will have two topological maps consistent with its information.

{end of example}

The example above may suggest that metrical information is used to check whether an already built topological map is consistent with metrical information. However, by including Axioms 62-72 inside *AT_block*, metrical information is used while building the topological map. As the next example illustrates, this may imply that the agent identifies more places than it does when not using metrical information.

⁴³Whenever we use a number x instead of an interval, it is an abbreviation for $[x, x]$.

Example 18. Consider an agent visiting the different corners of a square room in the order suggested by Figure 20a. In addition, suppose the agent’s sensory apparatus allows it to define *views* by characterizing the direction of walls and open space so that all corners look alike to the agent (see Example 16, Page 29). Suppose the agent has access to perfect metrical information and uses it while building the metrical map.




Figure 20: (a) The agent visits distinctive states $ds1$ to $ds7$ by the order suggested in the figure. Suppose all corners look alike to the agent. In particular, $ds1$ and $ds7$ share the same view. (b) Topological map associated with (a) when metrical information is not available. (c) Topological map associated with (a) when metrical information is available. In this case, the places associated with $ds1$ and $ds7$ are different ($P \neq S$).

In order to decide whether the agent is back to $ds1$, the frame $\langle ds1 : ds1, \dots, ds7 \rangle$ is created. Given the available metrical estimates it is not possible to have $teq(ds1, ds7)$ while satisfying the metrical constraints. Consequently, the topological map will have *four* places instead of *three*, as illustrated in Figure 20c. *{end of example}*

While in the examples above all visited distinctive states were included in a two dimensional frame of reference, this is in general not the case. In the presence of metrical uncertainty, a global frame of reference may not provide useful information to determine whether two places are the same, or to estimate the distance between two arbitrary places.

8 Algorithms

In this section we present an algorithm for calculating the topological maps associated with a set of experience E . The models associated with the causal theory (Section 4) can be calculated as the answer sets [Gelfond and Lifschitz, 1991] of a logic program. This logic program is implemented in Smodels [Niemelä and Simons, 1997] as illustrated in [Remolina and Kuipers, 2001]. It is possible to calculate the topological maps by a similar logic program. However, the number of grounding rules associated with such a program turns out to be prohibitive for practical applications.

The algorithm for calculating topological maps (the models of $TT(E)$) is stated as a “*best first*” search. The states of the search correspond to partial models of $TT(E)$.⁴⁴ At each step of the search a schema $\langle ds, a, ds' \rangle$ has to be explained. Either the identity

⁴⁴A partial model of $TT(E)$ is a model of $TT(E')$, for some $E' \subseteq E$.

of ds' can be proved or a search branch is created for every previously known distinctive state ds'_i that cannot be proven to be different from ds' .⁴⁵ In the branch where $teq(ds'_i, ds')$ is the case, $\neg teq(ds'_j, ds')$, $i \neq j$ are also asserted. An additional branch is created where $\neg teq(ds', ds'_j)$ are asserted. This branch represents the possibility that ds' is indeed different from previously known dstates. The next state to explore is the one that is minimal according to the order associated with the circumscription policy for $TT(E)$. This search algorithm is described in Figures 21 and 22.⁴⁶

Find-Models (S)

```

{
;; S = s0, . . . , sn ; sequence of schemas such that result(si) = context(si+1)
queue = ∅ ;   models = ∅ ;
pmodel = create-new-pmodel(S);   insert(pmodel,queue) ;
while queue ≠ ∅ do
begin
  pmodel = get-next-pmodel(queue);
  s = get-next-schema(pmodel);
  Explain(pmodel,s) ;
  if (inconsistent(pmodel) ∨ has-extensions(pmodel)) then skip;
  else if total-model(pmodel) then insert(pmodel, models);
  else insert(pmodel,queue);
end
return models;
}

```

Figure 21: **Best first search algorithm used to calculate the models of TT(E)**. The queue contains *consistent* partial models (pmodels) to be expanded. At each step of the search, a minimal partial model is picked and the next schema from its list of associated schemas is explained. A pmodel has extensions when a branch has been created while explaining a schema. A pmodel is a *total-model* when it has no more schemas to explain. Figure 22 defines how a pmodel explains a schema and when extensions are created.

The three key steps in the search are (Figure 22): creating a set of possible candidates to branch (*possible-equal-dstates*), generating a set of extensions when needed (*create-possible-extensions*), and explaining a schema in a given partial model (*assert-schema*). Another important issue is to detect when a partial model becomes inconsistent. We use the predicate *inconsistent(pmodel)* to denote this fact and the rules

$$\begin{aligned}
x \stackrel{pmodel}{=} y \wedge x \stackrel{pmodel}{\neq} y &\rightarrow inconsistent(pmodel), \\
teq(x, y) \in pmodel \wedge \neg teq(x, y) \in pmodel &\rightarrow inconsistent(pmodel) .
\end{aligned}$$

In the next sections we will show how to rewrite the axioms in the topological theory so they can be fed to a theorem prover to deduce equality and inequality relations. We use

⁴⁵The identity of the schema's context (i.e. ds in $\langle ds, a, ds' \rangle$) is known at each step in the search.

⁴⁶A search state is implemented by a partial model, *pmodel*. Branches in the search are represented by creating *extensions* for the current search state (pmodel). That *pmodel'* is an extension of *pmodel* implies that *pmodel'* inherits from *pmodel* all known objects and facts. Partial models are described in Page 39.

```

Explain (pmodel, s)
{ ;; s is a schema  $\langle ds, a, ds' \rangle$ 
  candidates = {};
  if known-result(pmodel,s)
  then Assert-schema(pmodel,s);
  else begin
    candidates = possible-equal-dstates(pmodel,s);
    if candidates  $\neq \{\}$ 
    then create-possible-extensions(pmodel,s,candidates)
    else Assert-schema(pmodel,s)
  end
}

Known-result(pmodel, s)
{ ;; s is a schema  $\langle ds, a, ds' \rangle$ 
  ;; The notation  $obj \in pmodel$  indicates that object  $obj$  is
  ;; known in the partial model  $pmodel$ .
  return  $ds' \in pmodel \vee \exists ds^*, ds'^* \in pmodel [\langle ds^*, a, ds'^* \rangle \in pmodel \wedge teq(ds^*, ds)]$ ;
}

Assert-schema (pmodel, s)
{ ;; s is a schema  $\langle ds, a, ds' \rangle$ .  $ds$  is known in  $pmodel$ 
  assert  $s \in pmodel$ ;
  if  $\neg$  known-result(pmodel,s)
  then begin
    assert  $ds' \in pmodel$ ;
    Create places and paths needed to explain  $s$ .
  end
  else begin
    pick  $ds'^*$  s.t.  $\exists ds^* \in pmodel [teq(ds^*, ds) \wedge \langle ds^*, a, ds'^* \rangle \in pmodel]$ ;
    assert  $ds' \stackrel{pmodel}{=} ds'^*$  in pmodel;
  end
}

```

Figure 22: **Explaining a schema.** $known\text{-}result(pmodel, s = \langle ds, a, ds' \rangle)$ is the case when the equality class for ds' can be deduced in the partial model $pmodel$. $Possible\text{-}equal\text{-}dstates(cntx, s)$ returns $dstates$ known in $pmodel$, having the same view as ds' and that cannot be proven different from ds' in $pmodel$. For each $ds'' \in candidates$, $create\text{-}possible\text{-}extensions(pmodel, s, candidates)$ creates an extension of $pmodel$ where $teq(ds', ds'')$ is the case. If the identity of ds' can be established, then s is asserted in $pmodel$. This declares ds' to be known in $pmodel$ and creates the places and paths that explain s according to the axioms of the topological theory $TT(E)$.

the rule-based system Algernon [Crawford and Kuipers, 1991] as our theorem prover. In Section 8.2 (Page 40) we present an illustrative trace of the algorithm.

8.1 Partial Models

In addition to a list of schemas to explain, a partial model has associated a set of objects (i.e. distinctive states, schemas, places, paths) that are known in the model. The basic relation among pmodels is the one of *extensions*. That $pmodel'$ is an extension of $pmodel$ implies that all known objects and facts in $pmodel$ are known objects and facts in $pmodel'$ (i.e. $pmodel'$ inherits from $pmodel$ all known objects and facts). This inheritance property of extensions can be implemented in Algernon by rules like the next one:⁴⁷

$$at(ds, place, pmodel) \wedge extension(pmodel, pmodel1) \rightarrow at(ds, place, pmodel1)$$

Create candidates. $Possible\text{-}equal\text{-}dstates(pmodel, s = \langle ds, a, ds' \rangle)$ returns a list of states that are possible equal to ds' . These are dstates known in $pmodel$, having the same view as ds' and that cannot be proven different from ds' in $pmodel$. Given s , we filter out ds'' as equal to ds' using rules including:⁴⁸

$$\begin{aligned} s = \langle ds, turn, ds' \rangle \wedge at(ds, p) \wedge at(ds'', q) \wedge p \neq q &\rightarrow \neg teq(ds', ds'') & (73) \\ [s = \langle ds, travel, ds' \rangle \wedge along(ds, pa, dir) \wedge along(ds'', pa1, dir1) \wedge \\ \neg [pa = pa1 \wedge dir = dir1]] &\rightarrow \neg teq(ds', ds'') \\ [s = \langle ds, travel, ds' \rangle \wedge along(ds, pa, dir) \wedge at(ds, p) \wedge at(ds'', q) \wedge \\ order(pa, dir, q, p)] &\rightarrow \neg teq(ds', ds'') \end{aligned}$$

The rules above are derived from the axioms in our theory. For instance, rule 73 is derived from the fact that each distinctive state is at a unique place, and distinctive states that are related by turn actions are at the same place.

Equality relations among topological objects (i.e. dstates, places, paths) are proved using rules derived by rewriting topological axioms. These rules include:

$$view(ds_1, v_1) \wedge view(ds_2, v_2) \wedge v_1 \neq v_2 \rightarrow \neg teq(ds_1, ds_2) \quad (74)$$

$$\langle ds, turn, ds' \rangle \rightarrow \neg teq(ds, ds') \quad (75)$$

$$order(pa, dir, p, q) \rightarrow p \neq q \quad (76)$$

$$radial(p, ds_1, h_1) \wedge radial(p, ds_2, h_2) \wedge h_1 \neq h_2 \rightarrow ds_1 \neq ds_2 \quad (77)$$

$$position1(pa, dir, p_1, pos_1) \wedge position1(pa, dir, p_2, pos_2) \wedge pos_1 \neq pos_2 \rightarrow p_1 \neq p_2 \quad (78)$$

$$leftOf(pa, dir, p) \wedge on(pa, q) \rightarrow p \neq q \quad (79)$$

$$leftOf(pa, dir, p) \wedge on(pa_1, p) \rightarrow pa \neq pa_1 \quad (80)$$

$$at(ds, p) \wedge at(ds, q) \rightarrow p = q \quad (81)$$

$$along(ds, pa, dir) \wedge along(ds, pa_1, dir_1) \rightarrow pa = pa_1 \wedge dir = dir_1 \quad (82)$$

⁴⁷Our logic for partial models takes the basic ideas developed in the area of formally reasoning about contexts [McCarthy and Buvač, 1998].

⁴⁸Rules are *relativized* to a particular model. In the implementation, all the topological predicates have a last extra argument for a pmodel. For instance, instead of writing $at(ds, p)$ we write $at(ds, p, pmodel)$. $at(ds, p, pmodel)$ is the case when $at(ds, p)$ is true in the partial model $pmodel$ (i.e. $pmodel \models at(ds, p)$).

Rules 74 and 75 rely on the fact that dstates have a unique view and turn actions link different distinctive states (Axioms 22 and 39). Rule 76 uses the fact that paths are not circular in order to conclude that if p is before q then p and q must be different (Axiom 44). Rules 77 and 78 use radial and one dimensional frames of reference to conclude inequality of dstates and places, respectively (Axiom 62). Rules 79 and 80 use boundary relations in order to distinguish places and paths respectively (Axiom 58). Rules 81 and 82 state that each distinctive state is at a unique place, along a unique path direction (Axioms 29 and 37).

Assert schema. *Assert-schema*($pmodel, s$) creates the places and paths needed to explain s . Instead of asserting $s = \langle ds, a, ds' \rangle$ in $pmodel$, the algorithm asserts $s^* = \langle ds^*, a, ds'^* \rangle$ where ds^* and ds'^* are the representatives in $pmodel$ for the *teg* equivalence classes of ds and ds' . Asserting a schema in Algernon corresponds to creating the frame (object) representing the schema. Forward and backward chaining rules derived from the topological theory are then evaluated, and places and paths needed to explain s are created.

8.2 Trace example

We illustrate the topological map building algorithm with the environment of Figure 23a. Distinctive states are visited in the order suggested by the figure. Distinctive state 9 is at the same environment location as dstate 1. However, two topological map are possible: either the agent is back to dstate 1 or dstate 0 (this is Example 13, Page 24). After traveling from dstate 9 to dstate 10, only one topological map is possible (Figure 23b). Figure 24 illustrates the use of the topological rules to distinguish distinctive states that share the same view. Figure 25 shows when branches in the search are created and how they can be refuted as more information becomes available to the agent.⁴⁹




Figure 23: (a) Numbers identify the dstates created by the map building algorithm. Views associated with dstates are also shown. Dstates 1 and 9 are at the same environment location. (b) Places and dpaths created by the map building algorithm. Notice that $P1$ and $P5$ are two names for the same place.

⁴⁹In the implementation, *dpaths* represent ordered dstates linked by travel actions. Dpaths correspond to paths that only have one direction associated with them. Paths are created when the agent has traveled in both direction of a path. At that time, two dpaths are associated with the path, one for each path's direction.

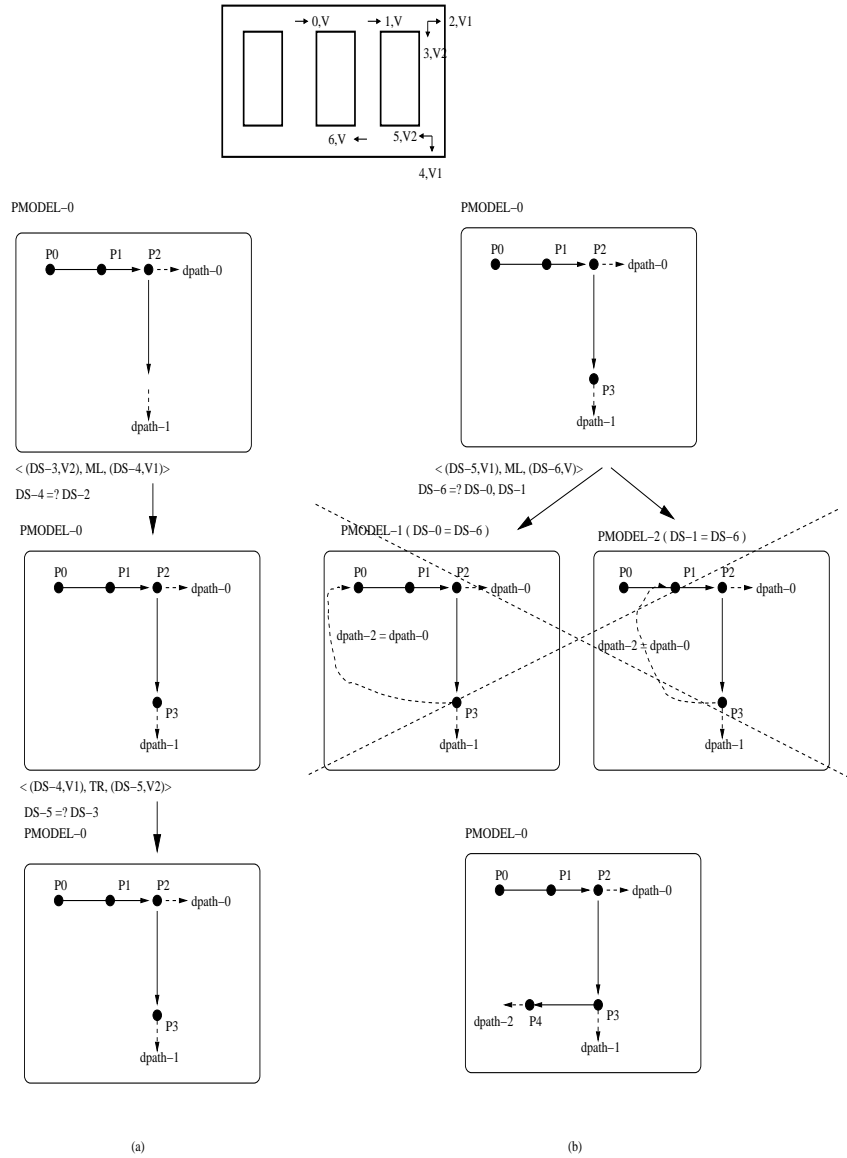


Figure 24: (a) When the agent reaches dstate 4, the same view $v1$ has been observed at dstate 2. Since dstate 2 is along $dpath-0$ and dstate 4 will be along $dpath-1$ (the agent just traveled from dstate 3 along $dpath-1$), dstate 4 and 2 are proven different. Dpaths 0 and 1 are different since there is a *turnRight* action relating them (Axiom 39). Place $P3$ is created to be the place dstate 4 is at (Axiom 29). Place $P3$ is proven different from place $P2$ since $P2$ is before $P3$ along $dpath-1$ (Corollary 1). Consequently, dstates 5 and 3 are proven different. (b) The agent travels to dstate 6 along $dpath-2$. Because $P3$ is to the right of $dpath-0$, $dpath-2$ cannot be the same as $dpath-0$. Consequently, dstate 6 is proven different from dstates 0 and 1.

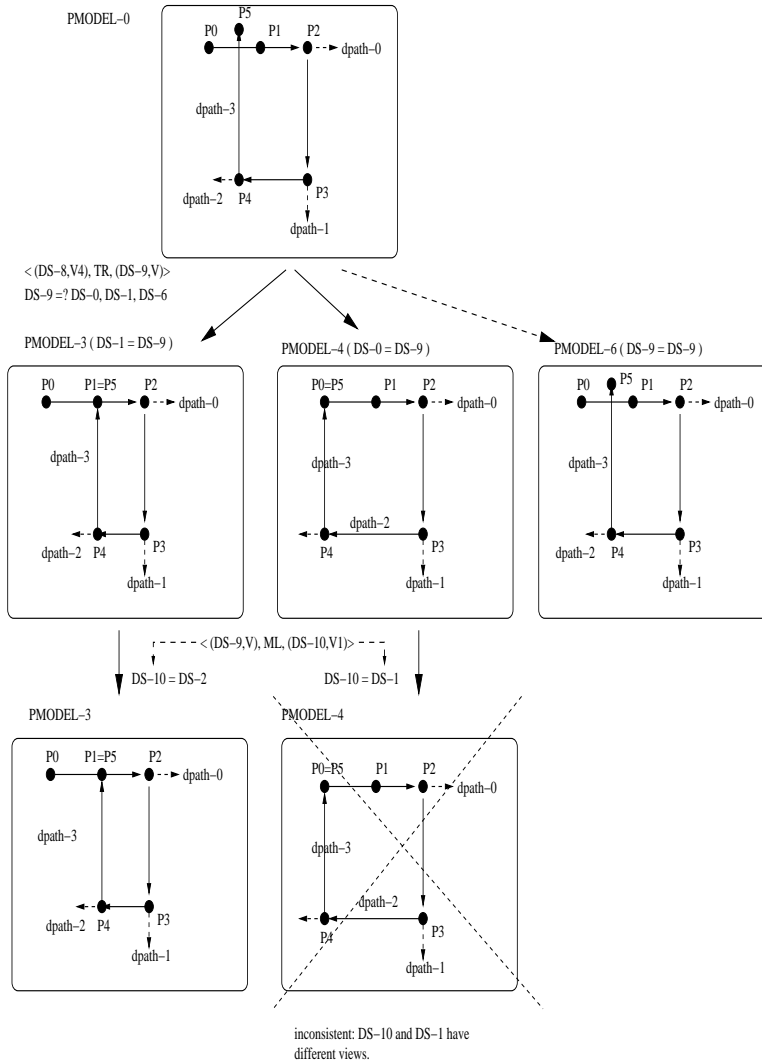


Figure 25: By the time the agent reaches dstate 8, six places (not five) are part of the map. Places P_5 and P_1 are not equal since the dstates of these places are not yet turn related (Axiom 30). Turning from dstate 8 to 9 leaves the agent with the three possibilities: (pmodel-3) dstates 9 and 1 are equal (and so $P_1 = P_5$, or (pmodel-4) dstates 9 and 0 are equal (and so $P_0 = P_1$, or (pmodel-6) dstate 9 is a new different dstate. That dstates 9 and 6 are different follows from the fact that places P_4 and P_5 are different. Pmodels 3 and 4 are minimal according to the topological theory circumscription policy. Pmodel-6 is not, but is left as a possible state in the search should new information render the other models inconsistent. The new schema $\langle 9, ML, 10 \rangle$ will render pmodel-4 inconsistent. Since actions are deterministic and dstates 0 and 9 are equal in this model, so should dstates 1 and 10. However, these dstates have different views so they cannot be equal. Pmodel-3 will then be the only map associated with the set of experiences.

9 Conclusions

What have we accomplished?. We have taken an informal description of the theory of topological maps and provided a formal account of the theory. In addition, we have extended the theory to handle perceptual aliasing, to describe environments with self intersecting and convergent paths, and to deal with local metrical information including uncertainty. The topological theory is independent of the agent's exploration strategy and of the algorithms used to build topological maps. We have taken the theory as a specification for a program able to keep track of different topological maps consistent with the agent's experiences in the environment. This program supports different exploration strategies as well as facilitates map disambiguation when the case arises.

A logical account of the causal, topological and local metrical theories was given using Nested Abnormality Theories. The minimality conditions embedded in the formalization define the preferred models associated with the theories. In Sections 4 through 7 we illustrated the main properties of the theories. In particular we showed how the minimal models associated with these theories are adequate models for the spatial knowledge an agent has about its environment. We also demonstrated how the causal, topological, and local metrical levels of the representation assume different spatial properties of the actions performed by the agent. This provides an increasingly refined ability to infer or refute equality relations (*ceq* and *teq*) among experienced environment states. By clarifying the ontology of causal and topological maps, and determining the dependency structure of the non-monotonic theory, we provide a solid foundation for general-purpose strategies for exploring unknown environments, or for disambiguating cases of perceptual aliasing.

Acknowledgments

A Nested Abnormality theories

In this appendix we define circumscription and nested abnormalities theories following [Lifschitz, 1994, Lifschitz, 1995]. The main idea of circumscription is to consider, instead of arbitrary models of an axiom set, only the models that satisfy a certain minimality condition (usually set inclusion).

Definition 4.[Circumscription] Let $A(P, Z_1, \dots, Z_m)$ be a sentence containing a predicate constant P and object, function and/or predicate constants Z_1, \dots, Z_m (and possibly other object, function and predicate constants). The *circumscription of P in A with varied Z_1, \dots, Z_m* is the sentence

$$A(P, Z_1, \dots, Z_m) \wedge \neg \exists p, z_1, \dots, z_m [A(p, z_1, \dots, z_m) \wedge p < P] \quad (83)$$

where $p < P$ denotes the formula

$$\forall x \{p(x) \rightarrow P(x)\} \wedge \exists x \{\neg p(x) \wedge P(x)\} .$$

We denote Formula 83 by *CIRC* $[A; P; Z]$. *{end of definition}*

Intuitively, the models of $CIRC [A; P; Z]$ are the models of A in which the extent of P cannot be smaller without losing the property A , even at the price of changing the interpretations of the constants Z .

It is often convenient to arrange different defaults by assigning priorities to them. Next we define two extensions to the basic definition of circumscription: parallel and prioritized circumscription.

Definition 5.[Parallel Circumscription] The *parallel circumscription*

$$CIRC [A; P^1, \dots, P^n; Z]$$

is the sentence $A(P, Z) \wedge \neg \exists p, z [A(p, z) \wedge p \prec P]$, where P stands for the tuple of predicates P^1, \dots, P^n and $p \prec P$ stands for the formula $\forall 1 \leq i \leq n \ p^i \leq P^i \wedge \exists 1 \leq i \leq n \ p^i < P^i$. {end of definition}

Definition 6.[Prioritized Circumscription] The *prioritized circumscription*

$$CIRC [A; P^1 \succ \dots \succ P^n; Z]$$

is the sentence $A(P, Z) \wedge \neg \exists p, z [A(p, z) \wedge p \prec P]$, where P stands for the tuple of predicates P^1, \dots, P^n and $p \prec P$ stands for the formula

$$\bigvee_{i=1}^n \left(\bigwedge_{j=1}^{i-1} (p^j = P^j) \wedge (p^i < P^i) \right) .$$

{end of definition}

The formula $p \prec P$ defines a *lexicographic* order among the predicates in p and P . Proposition 15 in [Lifschitz, 1994] shows that prioritized circumscription can be reduced to parallel circumscription as follows:

Theorem 5 *The circumscription $CIRC [A; P^1 \succ \dots \succ P^n; Z]$ is equivalent to*

$$\bigwedge_{i=1}^n CIRC [A; P^i; P^{i+1}, \dots, P^n, Z] .$$

Notation 1. $CIRC [A; P^1 \succ \dots \neg P_i \dots \succ P^n; Z]$ stands for the formula

$$CIRC [A \wedge not_P_i \equiv \neg P_i; P^1 \succ \dots not_P_i \dots \succ P^n; Z, P_i]$$

where not_P_i is a new constant predicate not occurring in A . {end of notation}

A.1 Nested Abnormality theories (NAT's)

Nested abnormality theories allows one to apply the circumscription operator to a subset of axioms, by structuring the knowledge base (the theory) into blocks. Each block can be viewed as a group of axioms that describes a certain collection of predicates and

functions, and the nesting of blocks reflects the dependence of these descriptions on each other.

Definition 7.[NAT's] Consider a second-order language L that does *not* include Ab among its symbols. For every natural number k , by L_k we denote the language obtained from L by adding Ab as a k -ary predicate constant. *Blocks* are defined recursively as follows: For any k and any list of function and/or predicate constants C_1, \dots, C_m of L , if each of A_1, \dots, A_n is a formula of L_k or a *block*, then $\{C_1, \dots, C_m : A_1, \dots, A_n\}$ is a *block*. The last expression reads: C_1, \dots, C_m are such that A_1, \dots, A_n . About C_1, \dots, C_m we say that they are *described* by this block.

The semantics of NAT's is characterized by a map φ that translates blocks into sentences of L . It is convenient to make φ defined also on formulas of the languages L_k . If A is such a formula, then $\varphi(A)$ stands for the universal closure of A . For blocks we define, recursively:

$$\varphi \{C_1, \dots, C_m : A_1, \dots, A_n\} = \exists ab \text{ CIRC } [\varphi A_1, \dots, \varphi A_n : ab : C_1, \dots, C_m] .$$

{end of definition}

Most often, it is desirable not to mention the predicate Ab at all. We will adopt the following notations:

- $\{C_1, \dots, C_m, \min P : A_1, \dots, A_n\}$ stands for

$$\{C_1, \dots, C_m, P : P(x) \rightarrow Ab(x), A_1, \dots, A_n\}$$

- $\{C_1, \dots, C_m, \max P : A_1, \dots, A_n\}$ stands for

$$\{C_1, \dots, C_m, P : \neg Ab(x) \rightarrow P(x), A_1, \dots, A_n\}$$

Definition 8. We extend the definition of *blocks* as follows: if A is a block, so is $\text{CIRC}[A; P^1 \succ \dots \succ P^n; Z]$. The semantics of NAT's is extended such that

$$\phi \text{CIRC}[A; P^1 \succ \dots \succ P^n; Z] = \text{CIRC}[\phi A; P^1 \succ \dots \succ P^n; Z] .$$

{end of definition}

As the next theorem shows, in some cases prioritized circumscription can be expressed using NAT's. In these cases however, the notation for prioritized circumscription is more compact than its equivalent NAT's. This motivates our previous definition.

Theorem 6 *Let A be a sentence such that Ab does not occur in A . Then,*

$$\text{CIRC}[A; P \succ Q; Z] = \{Z, \min Q : \{Z, Q, \min P : A\}\} .$$

B ceq properties

In this appendix we provide proofs for the different properties of the predicated *ceq* defined in Section 4 (Page 8).

Theorem 1. Let E be a complete set of experiences and let *CEQ_block* be defined as follows:

$$\{ \begin{array}{l} \text{max } \text{ceq} : \\ \text{ceq}(ds_1, ds_2) \rightarrow \text{View}(ds_1, v) \equiv \text{View}(ds_2, v), \\ \text{ceq}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow \text{ceq}(ds'_1, ds'_2) \\ \} \end{array}$$

Then the predicate *ceq* is an equivalence relation.

Proof. Let M_1 be a model for the axioms inside the *CEQ_block* as well as the other axioms of $CT(E)$. Let M_2 be a structure identical to M_1 except that

$$\text{ceq}^{M_2}(ds, ds') \equiv \text{ceq}^{M_1}(ds, ds') \vee ds = ds' .$$

We are to prove that M_2 is a model for the axioms inside the *CEQ_block* and consequently *CEQ_block* \models *ceq*(ds, ds).⁵⁰ Indeed,

- $M_2 \models \text{ceq}(ds, ds') \rightarrow \text{ceq}(ds', ds)$. In fact,

$$\begin{aligned} \text{ceq}^{M_2}(ds, ds') &\equiv \text{ceq}^{M_1}(ds, ds') \vee ds = ds' \\ &\rightarrow \text{ceq}^{M_1}(ds', ds) \vee ds' = ds \\ &\equiv \text{ceq}^{M_2}(ds', ds) \end{aligned}$$

- $M_2 \models \text{ceq}(ds, ds') \wedge \text{ceq}(ds', ds'') \rightarrow \text{ceq}(ds, ds'')$. In fact,

$$\begin{aligned} &\text{ceq}^{M_2}(ds, ds') \wedge \text{ceq}^{M_2}(ds', ds'') \\ &\equiv (\text{ceq}^{M_1}(ds, ds') \vee ds = ds') \wedge (\text{ceq}^{M_1}(ds', ds'') \vee ds' = ds'') \\ &\equiv (\text{ceq}^{M_1}(ds, ds') \wedge \text{ceq}^{M_1}(ds', ds'')) \vee (ds = ds' \wedge \text{ceq}^{M_1}(ds', ds'')) \vee \\ &\quad (\text{ceq}^{M_1}(ds, ds') \wedge ds' = ds'') \vee (ds = ds' \wedge ds' = ds'') \\ &\rightarrow \text{ceq}^{M_1}(ds, ds'') \vee (ds = ds' \wedge ds' = ds'') \\ &\equiv \text{ceq}^{M_2}(ds, ds'') \end{aligned}$$

- $M_2 \models \text{ceq}(ds, ds') \rightarrow \text{View}(ds, v) \equiv \text{View}(ds', v)$. In fact,

$$\begin{aligned} &\text{ceq}^{M_2}(ds, ds') \equiv \text{ceq}^{M_1}(ds, ds') \vee ds = ds' \\ &\rightarrow \forall v [\text{View}(ds, v) \equiv \text{View}(ds', v)] \vee ds = ds' \\ &\rightarrow \forall v [\text{View}(ds, v) \equiv \text{View}(ds', v)] \vee \forall v [\text{View}(ds, v) \equiv \text{View}(ds', v)] \\ &\equiv \text{View}(ds, v) \equiv \text{View}(ds', v) \end{aligned}$$

⁵⁰ M_2 satisfies the other axioms in $CT(E)$ since *ceq* does not occur in them.

- $M_2 \models ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2)$. In fact,

$$\begin{aligned}
& ceq^{M_2}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \\
& \equiv (ceq^{M_1}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle) \vee \\
& \quad (ds_1 = ds_2 \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle) \\
& \rightarrow ceq^{M_1}(ds'_1, ds'_2) \vee (\langle ds_1, a, ds'_1 \rangle \wedge \langle ds_1, a, ds'_2 \rangle) \\
& \stackrel{(15)}{\rightarrow} ceq^{M_1}(ds'_1, ds'_2) \vee ds'_1 = ds'_2 \\
& \equiv ceq^{M_2}(ds'_1, ds'_2)
\end{aligned}$$

Let's prove that $CEQ_block \models ceq(ds, ds') \rightarrow ceq(ds', ds)$. Let M_2 be a model identical to M_1 except that

$$ceq^{M_2}(ds, ds') = ceq^{M_1}(ds, ds') \vee ceq^{M_1}(ds', ds) .$$

By definition, ceq^{M_2} is symmetric. We need to prove that M_2 satisfy the axioms inside CEQ_block :

- $M_2 \models ceq(ds, ds') \rightarrow View(ds, v) \equiv View(ds', v)$. In fact,

$$\begin{aligned}
& ceq^{M_2}(ds, ds') \equiv ceq^{M_1}(ds, ds') \vee ceq^{M_1}(ds', ds) \\
& \rightarrow \forall v [View(ds, v) \equiv View(ds', v)] \vee \forall v [View(ds', v) \equiv View(ds, v)] \\
& \equiv View(ds, v) \equiv View(ds', v)
\end{aligned}$$

- $M_2 \models ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2)$. In fact,

$$\begin{aligned}
& ceq^{M_2}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \\
& \equiv [ceq^{M_1}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle] \vee \\
& \quad [ceq^{M_1}(ds_2, ds_1) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle] \\
& \rightarrow ceq^{M_1}(ds'_1, ds'_2) \vee ceq^{M_1}(ds'_2, ds'_1) \\
& \equiv ceq^{M_2}(ds'_1, ds'_2)
\end{aligned}$$

Finally, let's prove that $CEQ_block \models ceq(ds, ds') \wedge ceq(ds', ds'') \rightarrow ceq(ds, ds'')$. Let M_2 be a model identical to M_1 except that

$$ceq^{M_2} = transitive_closure(ceq^{M_1}) .$$

By definition, ceq^{M_2} is transitive. If ceq^{M_1} is reflexive and symmetric, so is ceq^{M_2} . We need to prove that M_2 satisfies the axioms inside CEQ_block :

- $M_2 \models ceq(ds, ds') \rightarrow View(ds, v) \equiv View(ds', v)$. In fact,

$$\begin{aligned}
& ceq^{M_2}(ds, ds') \\
& \equiv \exists ds^0, ds^1, \dots, ds^n [ds = ds^0, ds' = ds^n, ceq^{M_1}(ds^i, ds^{i+1}), 0 \leq i < n] \\
& \rightarrow \exists ds^0, ds^1, \dots, ds^n \\
& \quad [ds = ds^0, ds' = ds^n, View(ds^i, v) \equiv View(ds^{i+1}, v), 0 \leq i < n] \\
& \rightarrow \exists ds^0, ds^n [ds = ds^0, ds' = ds^n, View(ds^0, v) \equiv View(ds^n, v)] \\
& \equiv View(ds, v) \equiv View(ds', v)
\end{aligned}$$

- $M_2 \models ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2)$. In fact,

$$\begin{aligned}
 & ceq^{M_2}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \\
 & \equiv \exists ds^i (1 \leq i \leq n) [ds_1 = ds^1, ds_2 = ds^n, ceq^{M_1}(ds^i, ds^{i+1}), 1 \leq i < n] \\
 & \quad \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \\
 & \stackrel{hyp.}{\rightarrow} \exists ds^i \exists \langle ds^i, a, ds^{i'} \rangle \\
 & \quad [ds_1 = ds^1, ds_2 = ds^n, ds'_1 = ds^{1'}, ds'_2 = ds^{n'}, ceq^{M_1}(ds^i, ds^{i+1}), 1 \leq i < n] \\
 & \rightarrow \exists ds^{i'} [ds'_1 = ds^{1'}, ds'_2 = ds^{n'}, ceq^{M_1}(ds^{i'}, ds^{(i+1)'}), 1 \leq i < n] \\
 & \equiv ceq^{M_2}(ds'_1, ds'_2)
 \end{aligned}$$

□

When a set of experiences is complete the predicate ceq captures the idea that two distinctive states are the same if they render the same views under any sequence of actions. Assume that E is complete and let $A = a_1, \dots, a_n$ denote a sequence of actions. The term $A(ds)$ denotes the distinctive state resulting from executing A starting at ds . By definition, $A(ds) = ds$ if $n = 0$, $A(ds) = ds'$ such that $E \models \langle \langle a_1, \dots, a_{n-1} \rangle (ds), a_n, ds' \rangle$. Notice that the definition of $A(ds)$ makes sense since E is complete and actions are deterministic.

Theorem 2. Let E be a complete set of experiences. Then,

$$ceq(ds, ds') \equiv \forall A, v [View(A(ds), v) \equiv View(A(ds'), v)] .$$

Proof. Let M_1 be a model for the axioms inside the CEQ_block as well as the other axioms of $CT(E)$. Let M_2 be a model identical to M_1 except that

$$ceq^{M_2}(ds, ds') \equiv \forall A, v [View(A(ds), v) \equiv View(A(ds'), v)] .$$

By induction in the length of action sequences one can prove that $ceq^{M_1} \subseteq ceq^{M_2}$. Our proof is complete by showing that M_2 satisfies the axioms inside CEQ_block :

- $M_2 \models ceq(ds, ds') \rightarrow View(ds, v) \equiv View(ds', v)$. In fact, suppose $M_2 \models ceq(ds, ds')$ and consider the empty sequence of actions, $A = \{\}$, $A(ds) = ds$. Then

$$View(ds, V) \equiv View(A(ds), v) \equiv View(A(ds'), v) \equiv View(ds', v) .$$

- $M_2 \models ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2)$. In fact,

$$\begin{aligned}
 & ceq^{M_2}(ds'_1, ds'_2) \\
 & \equiv \forall A, v [View(A(ds'_1), v) \equiv View(A(ds'_2), v)] \\
 & \leftarrow \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \wedge \\
 & \quad \forall A, v [View(aA(ds_1), v) \equiv View(aA(ds_2), v)] \\
 & \leftarrow ceq^{M_2}(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle
 \end{aligned}$$

□

C teq properties

In this appendix we prove some properties of the SSH topological theory. Recall the SSH topological theory is defined as follows:

$$\begin{aligned}
TT(E) = & \\
& \text{there exist infinitely many places ,} \\
& \text{there exist infinitely many paths ,} \\
& \neg \exists p [tplace(p) \wedge is_region(p)] , \\
& \neg \exists pa [tpath(pa) \wedge route(pa)] , \\
& COMPLETION(E) , \\
& Axioms 2 - 10 , \\
& \langle ds, a, ds' \rangle \wedge \langle ds, a, ds'' \rangle \rightarrow ds' = ds'' , \quad (Axiom 15) \\
& T_block , \\
& AT_block = \tag{84} \\
& \{ \text{max teq :} \\
& \quad \Gamma \\
& \quad \mathbf{circ} \ tpath \succ tplace \ \mathbf{var} \ SSHpred \\
& \} \tag{85}
\end{aligned}$$

where Γ is the set of axioms defined in Page 14, and $SSHpred$ stands for the tuple of predicates $\langle \mathbf{at}, \mathbf{along}, \mathbf{order}, \mathbf{on}, \mathbf{teq}, \mathbf{turn_eq}, \mathbf{travel_eq} \rangle$.

Proposition 1 *Let M be a model of $TT(E)$. Then,*

- $M \models \forall pa, [tpath(pa) \equiv \exists ds, dir \ \mathbf{along}(ds, pa, dir)]$.
- $M \models \forall p, [tplace(p) \equiv \exists ds \ \mathbf{at}(ds, p)]$.

Proof.

$$\begin{aligned}
& CIRC[\Gamma; tpath \succ tplace; SSHpred] \\
& \equiv \{ \text{Proposition 15 in [Lifschitz, 1994]} \} \\
& CIRC[\Gamma; tpath; tplace, SSHpred] \wedge CIRC[\Gamma; tpath, tplace; SSHpred] \\
& \rightarrow \{ \text{def. of circumscription} \} \\
& CIRC[\Gamma; tpath]
\end{aligned}$$

Since $\Gamma = \Gamma'(tpath) \wedge [\mathbf{along}(ds, pa, dir) \rightarrow tpath(pa)]$ where $\Gamma'(tpath)$ is negative, then

$$\begin{aligned}
& CIRC[\Gamma; tpath] \\
& \equiv \\
& CIRC[\Gamma'(tpath) \wedge [\mathbf{along}(ds, pa, dir) \rightarrow tpath(pa)]; tpath] \\
& \equiv \{ \text{Proposition 4 in [Lifschitz, 1994]} \}
\end{aligned}$$

$$\begin{aligned}
& \Gamma'(tpath) \wedge CIRC[along(ds, pa, dir) \rightarrow tpath(pa); tpath] \\
& \rightarrow \{Proposition 1 \text{ in [Lifschitz, 1994]}\} \\
& [\exists ds, dir \text{ along}(ds, pa, dir)] \equiv tpath(pa)
\end{aligned}$$

Similarly, $\Gamma = \Gamma' \wedge [at(ds, p) \rightarrow tplace(p)]$ where $tpath$ does not occur in Γ' . Then,

$$\begin{aligned}
& CIRC[\Gamma; tpath \succ tplace; SSHpred] \\
& \rightarrow \{see \text{ above}\} \\
& CIRC[\Gamma; tpath, tplace; SSHpred] \\
& \rightarrow \{def. \text{ parallel circumscription}\} \\
& CIRC[\Gamma; tpath, tplace] \\
& \rightarrow \{def. \text{ parallel circumscription}\} \\
& CIRC[\Gamma' \wedge [at(ds, p) \rightarrow tplace(p)]; tplace] \\
& \equiv \{Propositions 1 \text{ and 4 in [Lifschitz, 1994]}\} \\
& \Gamma' \wedge [\exists ds, at(ds, p)] \equiv tplace(p)
\end{aligned}$$

□

Proposition 2 *The topological map associated with a finite set of experiences E has a finite number of topological paths and a finite number of topological places.*

Proof. Since a distinctive state is along at most one topological path (Axiom 37, Page 15), Proposition 1 implies that for any model M of $TT(E)$ there is an injection from $tpath^M$ into $distinctive-states^M$. Since $distinctive-states^M$ is finite so is $tpath^M$.

Similarly, since distinctive states are at a unique topological place (Axiom 29, Page 14), from Proposition 1 we conclude that the set of topological places in a model of $TT(E)$ is finite. □

Theorem 3. Let ds_1 be a distinctive state symbol such that

$$\forall ds_2 \notin [ds_1]_{turn}, [ds_2]_{teq} \cap [ds_1]_{turn} = \emptyset . \quad (86)$$

Then

$$\forall ds_2 \notin [ds_1]_{turn}, place(ds_2) \neq place(ds_1) .$$

Proof. The hypothesis of the theorem implies that

$$\forall ds_2 \notin [ds_1]_{turn}, \neg turn_eq(ds_2, ds_1) .$$

Indeed,

$$\begin{aligned}
turn_eq(ds_1, ds_2) & \equiv \exists b_0, \dots, b_n, b_{0'}, \dots, b_{n'} \text{ s.t.} \\
& \bullet b_0 = ds_2, b_{n'} = ds_1, \\
& \bullet teq(b_i, b_{i'}), i = 0, \dots, n \\
& \bullet \widehat{turn}(b_{i'}, b_{i+1}), i = 0, \dots, n-1 .
\end{aligned}$$

Let $1 \leq j \leq n$ such that $[\forall j \leq k \leq n, b_{k'} \in [ds_1]_{\widehat{turn}}]$ and $b_{(j-1)'} \notin [ds_1]_{\widehat{turn}}$. Notice that such a j exists since $ds_1 = b_{0'} \notin [ds_1]_{\widehat{turn}}$ and $ds_1 = b_{n'} \in [ds_1]_{\widehat{turn}}$. Consequently,

$$\begin{aligned}
& turn_eq(ds_1, ds_2) \\
& \rightarrow \\
& b_{j'} \in [ds_1]_{\widehat{turn}} \\
& \rightarrow \{teq(b_j, b_{j'})\} \\
& [b_j]_{teq} \cap [ds_1]_{\widehat{turn}} \neq \emptyset \\
& \rightarrow \{86\} \\
& b_j \in [ds_1]_{\widehat{turn}} \\
& \rightarrow \{\widehat{turn}(b_{(j-1)'}, b_j)\} \\
& b_{(j-1)'} \in [ds_1]_{\widehat{turn}} \\
& \rightarrow \\
& false
\end{aligned}$$

Thus $\neg turn_eq(ds_2, ds_1)$ should be the case. \square

Theorem 4. Any two models of the SSH topological theory have the same number of topological paths and the same number of topological places.

Proof. In order to prove that two models M_1 and M_2 of $TT(E)$ have the same number of topological paths (tpaths) and the same number of topological places (tplaces), it is enough to show that this is the case for models of the *AT_block* (Block 84). Suppose that $tpath^{M_1}$ has less elements than $tpath^{M_2}$, and so there exists an injection $\phi : tpath^{M_1} \rightarrow tpath^{M_2}$. One can extend ϕ to define an isomorphism from M_1 into M_2' , such that $M_2' \prec M_2$, where \prec is the order defined by the circumscription policy 85.⁵¹ In fact,

- Let $\phi : tplace^{M_1} \rightarrow tplace^{M_2}$ be an injection. Such an injection exists since $tplace^{M_1}$ is finite and $tplace^{M_2}$ is infinite.
- Let $\phi : S^{M_1} \rightarrow S^{M_2}$ be the identity over the sorts (S) of distinctive states, actions, views, schemas, path types and path directions. Recall we assumed a *Herbrand* interpretation for these sorts, where the corresponding universes are defined by the constant symbols in E .

The function ϕ above defines an isomorphic embedding from M_1 into M_2 in the standard way. In fact, $\phi(M_1) = M_2'$ is defined as follows:

- $tpath^{M_2'} = \phi(tpath^{M_1})$, $tplace^{M_2'} = \phi(tplace^{M_1})$.
- $teq^{M_2'} = \phi(teq^{M_1}) = \{teq(ds_1, ds_2) : M_1 \models teq(ds_1, ds_2)\} = teq^{M_1}$.
- $at^{M_2'} = \phi(at^{M_1}) = \{at(ds, \phi(p)) : M_1 \models at(ds, p)\}$.

⁵¹This proves that M_1 and M_2 have the same number of topological paths.

- $along^{M'_2} = \phi(along^{M_1}) = \{along(ds, \phi(pa), dir) : M_1 \models along(ds, pa, dir)\}$.
- $order^{M'_2} = \phi(order^{M_1}) = \{order(\phi(pa), dir, \phi(p), \phi(q)) : M_1 \models order(pa, dir, p, q)\}$.
- $on^{M'_2} = \phi(on^{M_1}) = \{on(\phi(pa), \phi(p)) : M_1 \models on(pa, p)\}$.
- $turn_eq^{M'_2} = \phi(turn_eq^{M_1}) = turn_eq^{M_1}$.
- $travel_eq^{M'_2} = \phi(travel_eq^{M_1}) = travel_eq^{M_1}$.

Notice that the language of Γ is defined by $\{tpath, tplace\} \cup SSHpred$. Thus $M_1 \models \Gamma$ implies $\phi(M_1) \models \Gamma$.⁵² Since $\phi(tpath^{M_1}) \subset tpath^{M_2}$, then $\phi(M_1) \prec M_2$, and so M_2 is not minimal, and is therefore not a model of $TT(E)$. It follows that M_1 and M_2 have the same number of topological paths.

Similar argument shows that M_1 and M_2 have the same number of topological places. If not, there would exist $\phi : tpath^{M_1} \rightarrow tpath^{M_2}$ a bijection and $\phi : tplace^{M_1} \rightarrow tplace^{M_2}$ an injection that allows us to apply the same argument as above. \square

D Theory axioms

This section presents some of the axioms not explicitly mentioned before.

D.1 T_block

The block **T_block** (see Page 13) defines the properties of the predicates \widehat{turn} , \widehat{travel} , and \vec{travel} . \widehat{turn} is the equivalence closure of the schemas $\langle \cdot, turn, \cdot \rangle$; \widehat{travel} and \vec{travel} are the equivalence and transitive closure of the schemas $\langle \cdot, travel, \cdot \rangle$ respectively.⁵³

$$\begin{aligned}
 T_block = & \tag{87} \\
 & \{ \min \widehat{turn}, \min \widehat{travel}, \min \vec{travel} : \\
 & \quad \langle ds, turn, ds' \rangle \rightarrow \widehat{turn}(ds, ds'), \\
 & \quad \langle ds, travel, ds' \rangle \rightarrow \widehat{travel}(ds, ds') \wedge \vec{travel}(ds, ds'), \\
 & \quad \widehat{turn}(ds, ds), \\
 & \quad \widehat{turn}(ds, ds') \rightarrow \widehat{turn}(ds', ds), \\
 & \quad \widehat{turn}(ds, ds') \wedge \widehat{turn}(ds', ds'') \rightarrow \widehat{turn}(ds, ds''),
 \end{aligned}$$

⁵²Notice that the circumscription policy varies all predicates in the language of Γ , and ϕ is the identity over all constant symbols in the theory, for otherwise, $\phi(M_1) \models \Gamma$ is not necessarily the case. In general the interpretations of an unary predicate (set) under a circumscriptive theory do not have the same number of elements. For example, consider the models of $CIRC[(P(0) \wedge P(1)) \vee P(2); P]$, where the interpretation of P could have one or two elements (this example is due to Vladimir Lifschitz).

⁵³A block of the form $\{C_1, \dots, C_n, \min P_1, \dots, \min P_k : A_1, \dots, A_m\}$ denotes the set of blocks $\{C_1, \dots, C_n, \min P_1 : A_1, \dots, A_m\}, \dots, \{C_1, \dots, C_n, \min P_k : A_1, \dots, A_m\}$.

$$\begin{aligned}
& \widehat{travel}(ds, ds), \\
& \widehat{travel}(ds, ds') \rightarrow \widehat{travel}(ds', ds), \\
& \widehat{travel}(ds, ds') \wedge \widehat{travel}(ds', dr) \rightarrow \widehat{travel}(ds, dr), \\
& \vec{travel}(ds, ds') \wedge \vec{travel}(ds', ds'') \rightarrow \vec{travel}(ds, ds'') \\
& \}
\end{aligned}$$

References

- [Angluin, 1978] Dana Angluin. On the complexity of minimum inference of regular sets. *Information and Control*, 39:337–350, 1978.
- [Basye *et al.*, 1995] K. Basye, T. Dean, and L.P. Kaelbling. Learning dynamics: System identification for perceptually challenged agents. *Artificial Intelligence*, 72(1):139–171, 1995.
- [Beardon, 1979] A. F. Beardon. *Complex Analysis*. John Wiley & Sons, New York, 1979.
- [Borenstein and Koren, 1991] J. Borenstein and Y. Koren. The vector field histogram - fast obstacle-avoidance for mobile robots. *IEEE Journal of Robotics and Automation*, 7(3):278–288, 1991.
- [Borenstein *et al.*, 1996] J. Borenstein, H.R. Everett, and L. Feng. *Navigating mobile robots: systems and techniques*. A K Peters, Wellesley, Massachusetts, 1996.
- [Borg and Groenen, 1997] Ingwer Borg and Patrick Groenen. *Modern multidimensional scaling: theory and applications*. Springer, New York, 1997.
- [Choset and Nagatani, 2001] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Trans. on Robotics and Automation*, 17(2):125–137, April 2001.
- [Crawford and Kuipers, 1991] J. Crawford and B. Kuipers. Algernon: a tractable system for knowledge representation. *SIGART Bulletin*, 2(3):35–44, 1991.
- [Davis, 1983] E. Davis. The MERCATOR representation of spatial knowledge. In A. Bundy, editor, *Proc. of the 8th IJCAI*, pages 295–301, Los Altos, 1983. Kaufmann Inc.
- [Dean *et al.*, 1993] T. Dean, K. Basye, and L. Kaelbling. Uncertainty in graph-based map learning. In Jonathan H. Connell and Sridhar Mahadevan, editors, *Robot Learning*, pages 171–192. Kluwer Academic Publishers, 1993.
- [Dudek *et al.*, 1991] Gregory Dudek, Jenkin Michael, Evangelos Miliotis, and David Wilkes. Robotic exploration as graph construction. *IEEE Trans. on Robotics and Automation*, 7(6):859–865, 1991.

- [Durrant-Whyte, 1987] H. F. Durrant-Whyte. Consistent integration and propagation of disparate sensor observations. *The International Journal of Robotics Research*, 6(3):3–24, 1987.
- [Durrant-Whyte, 1988a] H. F. Durrant-Whyte. *Integration, coordination and control of multisensor robot systems*. Kluwer Academic Publishers, Boston, 1988.
- [Durrant-Whyte, 1988b] H. F. Durrant-Whyte. Uncertain geometry in robotics. *IEEE Journal of Robotics and Automation*, 5(6):23–31, 1988.
- [Elfes, 1987] Alberto Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA3(3):249–265, 1987.
- [Engelson and McDermott, 1992] S.P. Engelson and D.V. McDermott. Error correction in mobile robot map learning. In *the IEEE International Conference on Robotics and Automation*, pages 2555–2560, 1992.
- [Franz *et al.*, 1998] M. Franz, B. Schölkopf, H. Mallot, and Bülthoff. Learning view graphs for robot navigation. *Autonomous robots*, 5:111–125, 1998.
- [Gelb, 1974] A. Gelb. *Applied Optimal Estimation*. MIT Press, Cambridge, MA, 1974.
- [Gelfond and Lifschitz, 1991] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.
- [Gold, 1978] E. Mark Gold. Complexity of automaton identification from given sets. *Information and Control*, 37:302–320, 1978.
- [Gopal *et al.*, 1989] S. Gopal, R. L. Klatzky, and T. R. Smith. Navigator: a psychologically based model of environmental learning through navigation. *Journal of Environmental Psychology*, 9:309–331, 1989.
- [Koenig and Simmons, 1996] S. Koenig and Reid Simmons. Passive distance learning for robot navigation. In *Proceedings of the Thirteenth International Conference on Machine Learning (ICML)*, pages 266 – 274, 1996.
- [Kortenkamp *et al.*, 1995] D. Kortenkamp, E. Chown, and S. Kaplan. Prototypes, locations, and associative networks (PLAN): towards a unified theory of cognitive mapping. *Cognitive Science*, 19:1–51, 1995.
- [Kuipers and Byun, 1988] B. Kuipers and Y. T. Byun. A robust qualitative method for spatial learning in unknown environments. In Morgan Kaufmann, editor, *AAAI-88*, 1988.
- [Kuipers and Byun, 1991] B. Kuipers and Y. T. Byun. A robot exploration and mapping strategy based on semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.
- [Kuipers and Levitt, 1988] Benjamin. Kuipers and T. Levitt. Navigation and mapping in large-scale space. *AI Magazine*, 9(2):25–43, 1988.

- [Kuipers *et al.*, 1993] B. Kuipers, R. Froom, W.Y. Lee, and D. Pierce. The semantic hierarchy in robot learning. In J. Connell and S. Mahadevan, editors, *Robot Learning*, pages 141–170. Kulwer Academic Publishers, 1993.
- [Kuipers, 1978] Benjamin Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978.
- [Kuipers, 2000] Benjamin Kuipers. The spatial semantic hierarchy. *Artificial Intelligence Journal*, 119:191–233, 2000.
- [Kuo, 1987] B. C. Kuo. *Automatic Control Systems*. Prentice-Hall, Inc., fifth edition, 1987.
- [Lee, 1996] W. Y. Lee. *Spatial Semantic Hierarchy for a Physical Mobile Robot*. PhD thesis, The University of Texas at Austin, 1996.
- [Leiser and Zilbershatz, 1989] David Leiser and Avishai Zilbershatz. THE TRAVELLER: a computational model of spatial network learning. *Environment and Behavior*, 21(4):435–463, 1989.
- [Lifschitz, 1994] Vladimir Lifschitz. Circumscription. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 297–352. Oxford University Press, 1994.
- [Lifschitz, 1995] Vladimir Lifschitz. Nested abnormality theories. *Artificial Intelligence*, (74):351–365, 1995.
- [Mallot and Gillner, 2000] H. A. Mallot and S. Gillner. Route navigating without place recognition: What is recognized in recognition-triggered responses? *Perception*, 29:43–55, 2000.
- [McCarthy and Buvač, 1998] John McCarthy and Sacha Buvač. Formalizing context (expanded notes). In A. Aliseda, R.J. van Glabbeek, and C. Westerståhl, editors, *Computing Natural Language*, volume 8L of CSLI Lecture Notes, pages 13–50. 1998.
- [McDermott and Davis, 1984] D. V. McDermott and E. Davis. Planning routes through uncertain territory. *Artificial Intelligence*, 22:107–156, 1984.
- [Moutarlier and Chatila, 1989] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modelling. In *5th International Symposium on Robotics Research*, pages 85–89, 1989.
- [Niemelä and Simons, 1997] I. Niemelä and P. Simons. Smodels - an implementation of the stable model and well-founded semantics for normal logic programs. In *4th International Conference on Logic Programming and Nonmonotonic Reasoning*, number 1265 in Lecture Notes in Computer Science, pages 420–429. Springer-Verlag, 1997.
- [O’Neill, 1991] Michael O’Neill. A biologically based model of spatial cognition and wayfinding. *Journal of Environmental Psychology*, 11:299–320, 1991.

- [Peressini *et al.*, 1988] Antony L. Peressini, Francis E. Sullivan, and K. Jerry Uhl. *The mathematics of nonlinear programming*. Springer-Verlag, New York, 1988.
- [Remolina and Kuipers, 1998] Emilio Remolina and Benjamin Kuipers. Towards a formalization of the Spatial Semantic Hierarchy. In *Fourth Symposium on Logical Formalizations of Commonsense Reasoning, London*, January 1998.
- [Remolina and Kuipers, 2001] Emilio Remolina and Benjamin Kuipers. A logical account of causal and topological maps. In *International Joint Conference in Artificial Intelligence (IJCAI-01)*, pages 5–11. AAAI press, Menlo Park, August 2001.
- [Rivest and Schapire, 1987] Ronald L. Rivest and Robert E. Schapire. A new approach to unsupervised learning in deterministic environments. In *Proceedings of the Fourth International Workshop on Machine Learning*, 1987.
- [Schölkopf and Mallot, 1995] B. Schölkopf and H. Mallot. View-based cognitive mapping and path planning. *Adaptive Behavior*, 3:311–348, 1995.
- [Shanahan, 1996] M. P. Shanahan. Noise and the common sense informatic situation for a mobile robot. In *AAAI-96*, pages 1098–1103, 1996.
- [Shatkay and Kaelbling, 1997] H. Shatkay and L. Kaelbling. Learning topological maps with weak local odometry information. In *IJCAI-97*, 1997.
- [Simmons and Koenig, 1995] R. Simmons and S. Koenig. Probabilistic Robot Navigation in Partially Observable Environments. In *IJCAI 95*, 1995.
- [Smith and Cheeseman, 1986] R. Smith and P. Cheeseman. On the representation of and estimation of spatial uncertainty. *The International Journal of Robotics Research*, 5:56–68, 1986.
- [Steck and Mallot, 2000] S. D. Steck and H. A. Mallot. The role of global and local landmarks in virtual environment navigation. *Presence*, 9(1):69–83, 2000.
- [Thrun *et al.*, 1998] S. Thrun, S. Gutmann, D. Fox, and B. Burgard, W. and Kuipers. Integrating topological and metric maps for mobile robot navigation: A statistical approach. In *AAAI-98*, pages 989–995, 1998.
- [Thrun, 1998] Sebastian Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.