



Towards a general theory of topological maps

Emilio Remolina^{a,*}, Benjamin Kuipers^{b,1}

^a *Stottler Henke Associates, Inc., 951 Mariners Island Blvd. #360, San Mateo, CA 94404, USA*

^b *Computer Science Department, The University of Texas at Austin, Austin, TX 78712, USA*

Received 25 March 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.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Topological maps; Causal maps; View graph; Spatial representations; Map building; Cognitive robotics; Spatial semantic hierarchy (SSH); Nested abnormality theories (NATs)

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

* Corresponding author.

E-mail addresses: eremolin@stottlerhenke.com (E. Remolina), kuipers@cs.utexas.edu (B. Kuipers).

¹ 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 the Texas Higher Education Coordinating Board, Advanced Technology Program (grant 003658-0656-2001), NSF grants IRI-9504138 and CDA 9617327, by funding from Tivoli Corporation, and by an IBM Faculty Research Award.

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.

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.

The major assumption underlying the topological approach to mapping is that there is a level of abstraction of the underlying environment at which actions are deterministic. In the Spatial Semantic Hierarchy [31], this is achieved by defining *distinctive states* and *actions* composed of trajectory-following and hill-climbing control laws such that actions are functionally deterministic when applied between distinctive states (see Fig. 1). There are two other assumptions that, when true, allow us to state the axiomatic theory in simpler terms. These are the assumptions that (a) a path does not intersect itself, and (b) a distinctive state corresponds to at most one path and one direction on that path. Section 5.3 describes the more elaborate default theory required to handle environments that violate these assumptions.

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 in [9,21,23,30,34,40,44]. 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. The sequence of views and actions generated by the robot exploration to recover the minimum deterministic automaton consistent with such information is used in [2,10,11,50]. In these works, actions do not have any spatial

properties associated with them. Metrical information associated with actions is considered in [22,53], but there sensory information (views) is not used. The use of both sensory and metrical information is proposed in [9,16,25,55]. Among these works, [9,25] 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. In [28,31] the existence of topological objects (i.e., paths, regions) is proposed 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.

In research on physical robots by Lee [33] and Choset and Nagatani [7], 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. The work by Choset and Nagatani [7] exploits the topology of the robot's free space to localize the robot on a partially constructed map. The map used in this work is the generalized Voronoi graph (GVG) which is a topological map that also encodes some metric information about the robot's environment. Our definition of topological maps includes but is not limited to GVGs. We propose an axiomatic theory of topological maps. The task of building the map is stated as an abduction task [47,52] where the agent's map correspond to the minimal models among those that explains its observations. Stating the minimality conditions as well as the ontology of the spatial representation is the content of this paper.

Metrical grid-based maps are another spatial representation used in the robotics community [4,15,59]. In these approaches 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.

The Spatial Semantic Hierarchy [31] assumes that an agent first builds a network of places and paths on top of which metrical models are added, rather than to build first a single metrical map from which a network of places and paths is derived. This assumption is motivated by research on human cognitive maps [37,46,54]. For the engineering tasks of robot exploration, mapping, and navigation, we believe that the “topology-first” approach is more efficient and robust. For example, Thrun et al. [59] propose a method for integrating topological and metrical paradigms to solve the concurrent mapping and localization problem studied in the mobile robotics community. The method has two phases. In the first phase, the topological mapping solves a global position alignment problem between potentially indistinguishable, significant places. The subsequent metric mapping phase produces a fine-grained metric map of the environment in high resolution. “This work illustrates that topological approaches indeed scale up to large and highly ambiguous environments. The environments tested here are difficult in that they possess large cycles, and in that local sensor information is insufficient to disambiguate locations” [59].

Finally, there are also feature-based spatial representations [58] where the map is a graph whose nodes represent observed features and whose edges represent geometric relationships between these features. Under these approaches the locations of geometric

features in the environment and the position of the vehicle is jointly estimated in a stochastic framework. Like grid-based methods, feature-based methods are subject to cumulative metrical error and the difficulty of properly closing large loops. A major benefit of topological maps is that the problem of correctly closing large loops is separated from the problem of metrical mapping of local environments. We refer the reader to Borenstein's book [5, Chapter 8] for a review of different approaches to map building.

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). Note that distinctive states represent not only location, but also the agent's orientation in the environment. 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 [32] that take the agent from one dstate to the next. For example, in [25,26,31] 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. Even with realistic levels of accuracy in the control laws, if the initial basin of attraction is large enough, and the hill-climbing control law is effective enough, the action become functionally deterministic (Fig. 1).

The sequence (1) is transformed into a set of **schemas** of the form $((v_i, ds_i), a_i, (v_{i+1}, ds_{i+1}))$, 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 Fig. 2. In order to go from distinctive state $ds1$ to distinctive state $ds2$, the agent executes the sequence of control laws




Fig. 1. Actions between distinctive states are functionally deterministic. The control laws making up an action have a basin of attraction surrounding the initial distinctive state ($ds1$). Any trajectory starting in that basin moves toward the fixed-point of the hill-climbing control law. Since any implementation has finite precision, the action terminates in a small region around the destination distinctive state ($ds2$). As long as the final region is small enough to be contained within the initial basin of attraction of every subsequent action departing from that state, then actions are functionally deterministic.




Fig. 2. A sequence of control strategies, $\langle \text{get_into_corridor}, \text{follow_middle_line}, \text{localize} \rangle$, 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.

$\langle \text{get_into_corridor}, \text{follow_middle_line}, \text{localize} \rangle$ where get_into_corridor is a trajectory following control law that moves the agent from $ds1$ to a , $\text{follow_middle_line}$ 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.

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 $\langle \text{face_space_on_right}, \text{localize} \rangle$. 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.

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 4.2.1), 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 4.3 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 **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,^{2,3}

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

However, we do **not** assume that views uniquely determine distinctive states (i.e., $\text{View}(ds, v) \wedge \text{View}(ds', v) \not\rightarrow ds = ds'$). 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.⁴ These constant symbols define completely the sort of *action_types* (Axiom 3). The predicate **Action_type**(*a*, *type*) represents the fact that the type of action *a* is *type*. Formally,

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

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

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 **Turn_desc**(*a*, *desc*) to indicate that *desc* is the qualitative description of the *turn action a*. Formally,⁵

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

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

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

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

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*),

² Throughout this paper we assume that free variables in formulas are universally quantified.

³ The formula $\exists!v P(v)$ means “there exists a unique *v* s.t. $P(v)$ ”. Formally, $\exists v \forall x [P(x) \equiv x = v]$.

⁴ The type of an action will be important in the topological theory (Section 5). For completeness of the presentation we introduce this concept here.

⁵ The notation $\text{UNA}[t_1, \dots, t_n]$ 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 $i \neq j$.

action a is the action associated with *schema* s ; **context**(s , ds), ds is 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):

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

Most often we are interested in *complete* schemas: those for whom the resulting distinctive state exists. Nevertheless, incomplete schemas allow the representation to account for common states of incomplete knowledge like “I could take you there, but I cannot tell you how” [31]. We use the (Causal Schema) predicate $CS(s, ds, a, ds')$ defined as

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

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' .

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.

While schemas are explicit objects of our theory, it is convenient to leave them implicit. We introduce the following convenient notation:⁶

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

4.2. 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):

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

⁶ Notice that we have “overloaded” the bracket notation depending on the type of its arguments.




Fig. 3. 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 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* (Causally Equal) to indicate whether two distinctive states represent the same environment state (Section 4.3).

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).

Example 2. Consider the set of experiences E gathered by the agent while navigating the environment in Fig. 3. 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 $v+$ ⁷ with the different distinctive states it found (i.e., a , b and c).

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 $\forall v \{v = v+\}$.

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

⁷ As 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.

4.2.1. The view graph

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

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

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 [17,38,51,57] shows how the view graph is consistent with human navigation abilities.

4.3. The causal theory

We use the predicate $\mathbf{ceq}(ds, ds')$ to denote the fact that *distinctive states* ds and ds' are *causally* indistinguishable. (In Section 5 we define when distinctive states are topologically indistinguishable.) Informally, $\mathbf{ceq}(ds, ds')$ 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 \mathbf{ceq} .

The causal theory associated with a set of experiences E , $CT(E)$, is the following nested abnormality theory (NATs) [36] (see Appendix A):

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

$\text{COMPLETION}(E),$

Axioms (2)–(10),

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

$$\text{CEQ_block} = \tag{16}$$

{ *max ceq*:

$\mathbf{ceq}(ds_1, ds_1),$

$\mathbf{ceq}(ds_1, ds_2) \rightarrow \mathbf{ceq}(ds_2, ds_1),$

$\mathbf{ceq}(ds_1, ds_2) \wedge \mathbf{ceq}(ds_2, ds_3) \rightarrow \mathbf{ceq}(ds_1, ds_3),$

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

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

}

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.

Given 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)).

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 . 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

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

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

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

Then, the predicate ceq is an equivalence relation.

Proof. See Appendix B. \square

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. \square

Example 3. Consider the set of experiences E as in Example 2 (see Fig. 4(a)). Since the same view is experienced at a , b and c , the extent of ceq is maximized by declaring




Fig. 4. (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 $\langle c, travel, d \rangle$.




Fig. 5. 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).

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

Although 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*, see Section 5) we will be able to distinguish these distinctive states (see Example 5).

Suppose the agent continues exploring the environment and gets the new information $View(d, v \sqcup)$, $CS(s3, c, ml, d)$, as suggested in Fig. 4(b). 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.

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

Example 4. Consider the environment depicted in Fig. 5. 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 vice versa (e.g., $\langle 4, ml, 5 \rangle$). A turn around action is performed when reaching a corner (e.g., $\langle 3, change_path_direction, 4 \rangle$, $\langle 7, change_path_direction, 8 \rangle$, 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:




Fig. 6. (a)–(b). Causal graphs associated with the set of experiences in Figs. 4(a) and 4(b). (c) View graph associated with the set of experiences in Fig. 4(b). Notice that the causal and view graphs associated with the experiences in Fig. 4(a) are isomorphic.

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

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 9 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).

4.4. 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 .

The problem of distinguishing environment states by outputs (views) and inputs (actions) has been studied in the framework of automata theory [1,2,20,50]. 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

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

assumptions about the environment or the agent’s perceptual abilities, the problem of finding this smallest automaton is NP-complete [1,20].

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. (In this article we do not consider this hierarchy.) In Section 6 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 17).

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.

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{tplace} 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 \mathbf{tpath} to identify these paths. A path connecting regions is called a **route**. A path has two directions, *pos* and *neg*, which can be thought of as referring to “upstream” and “downstream” in the order of places on the path. The path direction also serves as a frame of reference for specifying the boundary regions describing places to the left and right of the path (see Section 6). The sort of path directions is completely defined by *pos* and *neg*. For a direction dir , $-dir$ is defined such that $-pos = neg$ and $-neg = pos$.

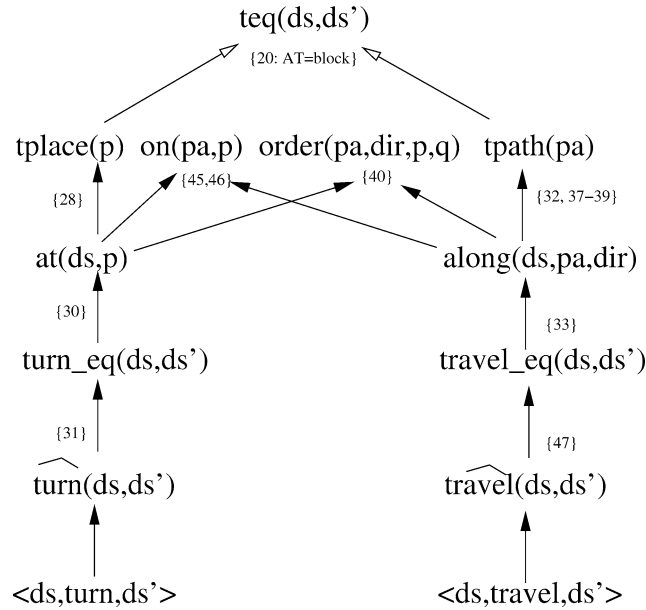


Fig. 7. Dependency among predicates in $TT(E)$. Labels on the graph's arrows refer to the axioms relating the predicates pointed by the arrows. Distinctive states related by turns modulo teq ($turn_eq$) must be *at* the same topological place ($tplace$). Distinctive states related by travels modulo teq ($travel_eq$) 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.

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 . Fig. 7 summarizes the dependencies among the above predicates. Section 5.2 formalizes these relationships.

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 countably 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 finite sets of distinctive states and *topological paths* are identified with finite 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*. We require infinite sorts of places and paths to avoid models being non-comparable due to a mismatch in the cardinalities of the sorts, as illustrated in Example 16.

5.2. The topological theory

The topological theory associated with E , $\mathbf{TT}(E)$, is the following nested abnormality theory (NATs) [36] (see Appendix A): (the condition that the sorts of places and paths are countably infinite is formalized by asserting the existence of a bijection between these sorts and the natural numbers)

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

The block $\mathbf{T_block}$ defines the properties of the predicates \widehat{turn} , \widehat{travel} , and \overrightarrow{travel} . \widehat{turn} is the equivalence closure of the schemas $\langle \cdot, turn, \cdot \rangle$; \widehat{travel} and \overrightarrow{travel} are the equivalence and transitive closure of the schemas $\langle \cdot, travel, \cdot \rangle$ respectively (Appendix D).

The block $\mathbf{AT_block}$ 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 policy

$$\mathbf{circ} \ tpath \succ tplace \ \mathbf{var} \ SSHpred$$

stating that a minimum set of topological paths is preferred to a minimum set of topological places. The symbol \succ indicates prioritized circumscription (see Appendix A). $SSHpred$ stands for the tuple of predicates $\langle \mathbf{at}, \mathbf{along}, \mathbf{order}, \mathbf{on}, \mathbf{teq}, turn_eq, travel_eq \rangle$. The 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. The block AT_block is defined as follows:

$$AT_block = \tag{20}$$

$$\begin{aligned}
 & \{ maxteq: \\
 & \quad teq(ds, ds), \\
 & \quad teq(ds_1, ds_2) \rightarrow teq(ds_2, ds_1), \\
 & \quad teq(ds_1, ds_2) \wedge teq(ds_2, ds_3) \rightarrow teq(ds_1, ds_3), \tag{21}
 \end{aligned}$$

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

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

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

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

$$\langle ds, turn, ds' \rangle \rightarrow \neg teq(ds, ds'), \quad (25)$$

$$\langle ds, turnAround, ds' \rangle \wedge \langle ds, turnAround, ds'' \rangle \rightarrow teq(ds', ds''), \quad (26)$$

$$\langle ds_1, turnAround, ds_2 \rangle \wedge \langle ds_2, turnAround, ds_3 \rangle \rightarrow teq(ds_1, ds_3), \quad (27)$$

$$at(ds, p) \rightarrow tplace(p), \quad (28)$$

$$\exists! p at(ds, p), \quad (29)$$

$$turn_eq(ds_1, ds_2) \equiv \forall p[at(ds_1, p) \equiv at(ds_2, p)], \quad (30)$$

$$\{min\ turn_eq: \quad (31)$$

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

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

}

$$along(ds, pa, dir) \rightarrow tpath(pa), \quad (32)$$

$$\{min\ along: \quad (33)$$

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

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

$$teq(ds_1, ds_2) \rightarrow along(ds_1, pa, dir) \equiv along(ds_2, pa, dir) \quad (36)$$

}

$$along(ds, pa, dir) \wedge along(ds, pa1, dir1) \rightarrow pa = pa1 \wedge dir = dir1, \quad (37)$$

$$at(ds_1, p) \wedge at(ds_2, p) \wedge along(ds_1, pa, dir) \wedge$$

$$along(ds_2, pa, dir) \rightarrow teq(ds_1, ds_2), \quad (38)$$

$$[\langle ds, turn_desc, ds' \rangle \wedge turn_desc \neq turnAround \wedge \quad (39)$$

$$along(ds, pa, dir) \wedge along(ds', pa1, dir1)] \rightarrow pa \neq pa1,$$

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

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

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

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

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

}

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

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

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

$$\exists ds_1, dir_1, ds_2, dir_2 [at(ds_1, p) \wedge along(ds_1, pa, dir_1) \wedge at(ds_2, q) \wedge \\ along(ds_2, pa, dir_2) \wedge travel_eq(ds_1, ds_2)],$$

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

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

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

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

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

}

$$\mathbf{circ}\ tpath \succ tplace\ \mathbf{var}\ SSH\vec{pred} \quad (48)$$

}

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 $\langle ds, Turn, ds \rangle$ 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. Every 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)). Block (31) states that the predicate *turn_eq* corresponds to the relation \widehat{turn} modulo *teq*.

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 Blocks (33) and (40), respectively.

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).

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 [35, Proposition 4]). Axiom (37) prevents the existence of different paths that converge to the same distinctive state (in Section 5.3 we will make this axiom a default). Finally, Axiom (38) states that there exist at most one distinctive state indicating a path's direction at a given place on the path.

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). This axiom prevents the existence of self-intersecting paths (Fig. 15).

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 [35, Proposition 4]). 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 whenever the agent has directly traveled between two distinctive states, the places associated with these distinctive states are different.

Corollary 1. $\overline{\text{travel}}(ds, ds') \rightarrow \text{place}(ds) \neq \text{place}(ds')$, where $\text{place}(ds)$ denotes the unique topological place that distinctive state ds is at (Axiom (29)). Moreover, consecutive travels among distinctive states occur along the same topological path.

Corollary 2.

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

In order to prove that distinctive states ds_1 and ds_2 are at different topological places, one has to prove that $\neg \text{turn_eq}(ds_1, ds_2)$. The following theorem states a strong condition for when this is the case. Given an equivalence relation R , $[x]_R$ denotes the equivalence class of x according to R .

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

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

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




Fig. 8. Distinctive states a , b and c cannot be distinguished at the causal level (see Example 3). Using the concepts of *paths* and *places* these dstates are distinguished.

Proof. See Appendix C. \square

Recall that the interpretations for *tplace* and *tplace* are finite. Our circumscription policy 48 and the fact that the sorts of paths and places are infinite implies the following fact:

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. \square

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

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

Example 5. Consider the set of experiences E

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

as in Example 3, Fig. 8. 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., $ceq(a, b) \wedge \neg teq(a, b)$).

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

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

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 Fig. 9(c). The model has three places and one path. Not having a *turn* action relating a and d prevents the agent from grouping these distinctive states into the same place, as suggested in Fig. 9(b). Next we show why this is the case.

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




Fig. 9. (a) The agent navigates a rectangle environment getting the experiences $\langle a, \text{travel}, b \rangle$, $\langle b, \text{turnAround}, c \rangle$, and $\langle c, \text{travel}, d \rangle$. 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 $\langle d, \text{turnAround}, a \rangle$, it will consider (b) as the topological map and disregard (c).




Fig. 10. 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 .

that $\widehat{\text{turn}}(ds, ds') \equiv [ds = ds' \vee \{ds = b \wedge ds' = c\} \vee \{ds = c \wedge ds' = b\}]$. In particular, $\neg \text{turn_eq}(a, d)$. In virtue of Axiom (30) we cannot conclude that a and d are at the same topological place.

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 $\langle a, \text{turnRight}, b \rangle$, $\langle b, \text{travel}, c \rangle$, $\langle c, \text{turnAround}, d \rangle$, $\langle d, \text{travel}, e \rangle$, $\langle e, \text{turnRight}, a' \rangle$, $\langle a', \text{turnRight}, b' \rangle$ consistent with an agent going from one four-way intersection to another (Fig. 10). Let us consider the models of these schemas. From our axioms, at least one path and three places must exist:

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

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 $\text{along}(b', \bullet, \bullet)$.




Fig. 11. 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 ds_1, ds_2, \dots, ds_6 ; (b) depicts the topological map associated with (a). Three paths instead of only two are required to explain the agent experiences (see text).

Example 8. Consider the extension of the previous example where the schema $\langle b', travel, c' \rangle$ 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 $teq(c, c')$. Notice that $\neg \widehat{travel}(b, c')$ and $travel_eq(b, c')$ are the case.

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 Fig. 11(a) obtaining the following schemas:

$$\langle ds_1, travel, ds_2 \rangle \langle ds_2, turnRight, ds_3 \rangle \langle ds_3, travel, ds_4 \rangle \\ \langle ds_4, turnLeft, ds_5 \rangle \langle ds_5, travel, ds_6 \rangle.$$

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 Fig. 11(b)) such that

$$at(ds_1, A), \quad at(ds_2, B), \quad at(ds_3, B), \quad at(ds_4, C), \quad at(ds_5, C), \quad at(ds_6, D).$$

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 $\neg \widehat{turn}(ds_1, ds_3)$, $\neg \widehat{turn}(ds_1, ds_5)$ and $\neg \widehat{turn}(ds_2, ds_6)$ are the case, A, B, C and D are all different. Axiom (34) implies that there exist paths Pa, Pa_1, Pa_2 , and directions dir, dir_1, dir_2 , such that:

$$order(Pa, dir, A, B), \quad along(ds_1, Pa, dir), \quad along(ds_2, Pa, dir), \\ order(Pa_1, dir_1, B, C), \quad along(ds_3, Pa_1, dir_1), \quad along(ds_4, Pa_1, dir_1), \\ order(Pa_2, dir_2, C, D), \quad along(ds_5, Pa_2, dir_2), \quad along(ds_6, Pa_2, dir_2).$$

Schemas $\langle ds_2, turnRight, ds_3 \rangle$ and $\langle ds_4, turnLeft, ds_5 \rangle$, and Axiom (39) implies that $Pa \neq Pa_1$, $Pa_1 \neq Pa_2$. 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 $\neg \widehat{travel}(ds_1, ds_4)$ and $\neg \widehat{travel}(ds_1, ds_5)$ are the case, and in virtue of Axiom (46) it follows that $Pa \neq Pa_2$.




Fig. 12. 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 vice versa (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).

Example 10. Consider the same T-environment exploration presented in Example 4 (see Fig. 12). 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.

The three causal models associated with T-environment are:

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

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

Moreover, Corollary 1 implies that

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

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

$$\begin{aligned} & \text{place}(5) \neq \text{place}(8), \quad \text{along}(6, Pa2, \text{dir}2), \quad \text{along}(7, Pa2, \text{dir}2), \\ & \text{along}(8, Pa2, -\text{dir}2), \quad \text{along}(9, Pa2, -\text{dir}2). \end{aligned}$$

From $\langle 9, \text{turnRight}, 10 \rangle$, $\langle 10, \text{travel}, 11 \rangle$, $\langle 11, \text{turnAround}, 12 \rangle$, there exist path $Pa3$ and direction $\text{dir}3$ such that $Pa2 \neq Pa3$ and

$$\text{along}(10, Pa3, \text{dir}3), \quad \text{along}(11, Pa3, \text{dir}3), \quad \text{along}(12, Pa3, -\text{dir}3).$$

Theorem 3 allow us to conclude that $\text{place}(5) \notin \{\text{place}(1), \text{place}(2), \text{place}(3)\}$. The same argument shows that $\text{place}(8) \notin \{\text{place}(1), \text{place}(2), \text{place}(3), \text{place}(5)\}$. 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, $Pa1 = Pa3$, $\text{dir}1 = \text{dir}3$, $\text{teq}(2, 10)$, $\text{teq}(3, 11)$ and $\text{teq}(4, 12)$. 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 $\neg \text{teq}(2, 10)$ were the case, then Theorem 3 allows to conclude that $\text{place}(9) \notin \{\text{place}(1), \text{place}(2), \text{place}(3), \text{place}(5), \text{place}(8)\}$ and so the model will have at least six tplaces. Consequently $\text{teq}(2, 10)$ has to be the case in a minimal model of the theory.

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

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 Fig. 13 obtaining the following schemas:

$$\begin{aligned} & \langle ds1, \text{turnAround}, ds2 \rangle \quad \langle ds2, \text{turnAround}, ds1 \rangle \quad \langle ds1, \text{travel}, ds3 \rangle \\ & \langle ds3, \text{turnRight}, ds4 \rangle \quad \langle ds4, \text{turnLeft}, ds3 \rangle \quad \langle ds3, \text{travel}, ds6 \rangle \\ & \langle ds6, \text{turnLeft}, ds7 \rangle \quad \langle ds7, \text{travel}, ds4 \rangle \\ & \langle ds4, \text{turnRight}, ds5 \rangle \quad \langle ds5, \text{travel}, ds2 \rangle. \end{aligned}$$

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 ad-




Fig. 13. 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.

dition, 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

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

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 $Pa1$.

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.

Since 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. However, no “up to path direction isomorphic” topological maps can explain the same pattern of experience. This happens because the experiences are incomplete, or the agent’s sensors are weak.

Example 13. Assume that the agent visits places A, B, C, D, E, F, C in the order suggested by Fig. 14. 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 (Figs. 14(b), (c)).

Metrical information can be used to deduce the correct topology (see Example 18). However, if the agent accumulates more information, by turning at C and traveling to D , then topological information suffices to deduce that the topology of the environment is the one in Fig. 14(b). This is the case since the views at C and D are different.




Fig. 14. (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 .



Fig. 15. (a) Self-intersecting paths. (b) Convergent paths.

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 (Fig. 15). In this section we extend our theory to deal with these types of paths. Converging paths are the standard counterexample for the axiom stating that distinctive states are along a unique path (Axiom (37)). We replace Axiom (37) by the block

```
{min convergent_paths:
  [along(ds, pa, dir) ∧ along(ds, pa1, dir1)
   ∧ ¬[pa = pa1 ∧ dir = dir1]] → convergent_paths(pa, pa1)
}
```

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

$$\mathbf{circ} \text{ self_intersecting } > \text{ convergent_paths } > \text{ tpath } > \text{ tplace } \mathbf{var} \text{ SSHpred.} \quad (49)$$

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 5–12 are still valid maps for the new theory. Next we study some cases we could not handle before.




Fig. 16. Distinctive state a is along two different paths. These two paths are declared convergent paths in the model of our theory.

Example 14. Suppose the agent has experienced the following schemas (Fig. 16):

$$\begin{aligned} \langle b, \text{travel}, d \rangle \quad \langle d, \text{turnAround}, c \rangle \quad \langle c, \text{turnRight}, e \rangle \\ \langle e, \text{travel}, a \rangle \quad \langle a, \text{turnAround}, b \rangle. \end{aligned}$$

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

- *Model 1.* In this model $Pa \neq Pa1$. Consequently, $\text{self_intersecting} = \text{false}$ and $\text{convergent_paths}(Pa, Pa1)$ are the case.
- *Model 2.* In this model $Pa = Pa1$. Consequently, $\text{self_intersecting}(Pa)$ and $\text{convergent_paths} = \text{false}$ are the case.

We prefer model 1 over model 2 according to the circumscription policy (49).

Example 15. Consider the set of experiences E

$$\langle (a, v+), \text{travel}, (b, v+) \rangle, \quad \langle (b, v+), \text{travel}, (c, v+) \rangle$$

as in Example 3. In the intended minimal model there are one path and three places. There are however other interpretations for the schemas. For example, the agent traveled from a to b along path Pa and then “changed” paths to travel from b back to a along path $Pa1$. In this “model” we have $\text{teq}(a, c)$, $Pa \neq Pa1$ and $\text{convergent_paths}(Pa, Pa1)$. The model has two paths and two places (less places than the intended model). By prioritizing paths over places we get rid of this model. The prioritization conveys the heuristic that “paths help to determine places”. In general if “concept X helps to determine concept Y ” then X has higher priority than Y in our circumscription policy.

Our requirement of infinite *places* and *paths* allow us to compare any two models of the theory (see Theorem’s 4 proof). This requirement also allow us to deal with unexpected models as illustrated in the next example.

Example 16. Consider the schema $\langle a, \text{travel}, b \rangle$ 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:

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

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 (infinite) places.

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: when the curve is removed, the plane is divided into two disjoint connected sets [3]. 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).

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 (Fig. 17).

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 17).




Fig. 17. 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.

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

$$\mathbf{totheLeftOf}(p, pa, dir, pa1, dir1), \quad \mathbf{totheRightOf}(p, pa, dir, pa1, dir1)$$

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 Fig. 18).

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

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

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 $\mathbf{leftOf}(pa, dir, lr)$ and $\mathbf{rightOf}(pa, dir, rr)$ 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:

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

$$\mathit{leftOf}(pa, dir, r) \equiv \mathit{rightOf}(pa, -dir, r), \tag{52}$$

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

$$\mathit{leftOf}(pa, dir, lr) \wedge \mathit{leftOf}(pa1, dir1, lr) \rightarrow pa = pa1. \tag{54}$$




Fig. 18. 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 .

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)). Boundary relations are derived according to Axioms (60) and (61) (see Fig. 18): (the symmetry between *leftOf* and *rightOf* defined by Axiom (52) let us write our axioms in terms of only one of these predicates)

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

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

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

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

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

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

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

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

$$[totheLeftOf(p1, pa, dir, dir1, pa1) \wedge order(pa1, dir1, p1, p) \wedge \quad (61)$$

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

}

}

Block (56) defines the extent of the predicate *in_region*. The outer preference minimizes *in_region*, so its positive instances only reflect actual observations. 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 Fig. 18). The predicate **Ab** inside Block (57) is the *auxiliary* “abnormality” predicate associated with a NAT block (Appendix A).

(See [36] for a similar formalization of the standard example: objects normally do not fly; birds normally do.) Some sufficient conditions for when Ab is the case can be derived from Block (57) as follows.

Let $leftOf'$ and $rightOf'$ denote the following abbreviations:

$$\begin{aligned} leftOf'(p, pa, dir) &\equiv \exists lr \{ leftOf(pa, dir, lr) \wedge in_region(p, lr) \}, \\ rightOf'(p, pa, dir) &\equiv \exists rr \{ rightOf(pa, dir, rr) \wedge in_region(p, rr) \}, \end{aligned}$$

which allow us to implicitly refer to the left and right regions associated with a path (these abbreviations make sense given Axiom (51)). Axioms inside Block (57) can be rewritten as follows:

$$\begin{aligned} on(pa, p) &\rightarrow \neg leftOf'(p, pa, dir) \wedge \neg rightOf'(p, pa, dir), \\ leftOf'(p, pa, dir) &\rightarrow \neg rightOf'(p, pa, dir), \\ [totheRightOf(p1, pa, dir, pa1, dir1) \wedge order(pa1, dir1, p1, p) \wedge \\ &\neg rightOf'(p, pa, dir)] \rightarrow Ab(pa, p), \\ [totheLeftOf(p1, pa, dir, dir1, pa1) \wedge order(pa1, dir1, p1, p) \wedge \\ &\neg leftOf'(p, pa, dir)] \rightarrow Ab(pa, p). \end{aligned}$$

Using this rewriting of Block (57), one can derive the following (among others) sufficient conditions to deduce Ab :

$$on(pa, p) \wedge [totheRightOf(p1, pa, dir, pa1, dir1) \wedge order(pa1, dir1, p1, p)] \rightarrow Ab(pa, p), \quad (62)$$

$$\begin{aligned} [totheRightOf(p1, pa, dir, pa1, dir1) \wedge order(pa1, dir1, p1, p) \wedge \\ totheLeftOf(p2, pa, dir, pa2, dir2) \wedge order(pa2, dir2, p2, p)] \\ \rightarrow Ab(pa, p), \end{aligned} \quad (63)$$

$$\begin{aligned} [totheRightOf(p1, pa, dir, pa1, dir1) \wedge order(pa1, dir1, p1, p) \wedge \\ totheRightOf(p2, pa, -dir, pa2, dir2) \wedge order(pa2, dir2, p2, p)] \\ \rightarrow Ab(pa, p). \end{aligned} \quad (64)$$

Conditions (62)–(64) show sufficient conditions for when Ab is the case, and consequently when the agent should not deduce boundary relations. (Condition (64) uses the symmetry between $leftOf'$ and $rightOf'$ defined by Axiom (52).) These conditions are in terms of predicates other than in_region , $leftOf$ and $rightOf$ whose extent is the purpose of Blocks (56) and (57).

6.3. Adding boundary relations to the topological map

We update the topological theory by including Axioms (50)–(61) inside the block **AT_block** (Section 5.2), and the new circumscription policy becomes

$$\mathbf{circ} \neg in_region \succ tpath \succ tplace \mathbf{var} new\overline{SSH}pred$$

where $new\overline{SSH}pred$ stands for the tuple of predicates

(*at, along, order, on, teq, turn_eq, travel_eq,*
totheRightOf, totheLeftOf, leftOf, rightOf, is_region
).

The circumscription policy states that Axioms (60) and (61) should be used to draw conclusions even at the expense of having more paths or more places on the map. This is achieved by maximizing *in_region* over *tpath* in the circumscription policy. This policy also prevents the theory from preferring pathological *tpaths* and *tplaces*. By maximizing the extent of *in_region* at the expense of having possibly more paths or 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 17. Consider an agent visiting the different corners of a square room in the order suggested by Fig. 19(a). 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),
(ds1, turnRight, ds2), (ds2, travel, ds3), (ds3, turnRight, ds4), (ds4, travel, ds5).

Suppose that the agent does not use boundary regions when building the topological map. From $\langle ds3, turnRight, ds4 \rangle$ and Axiom (39) we can deduce that $Pa \neq Pb$ in Fig. 19(b). Then the minimal topological model associated with *E* has two paths and two places. In this model, $teq(ds1, ds5)$ is the case. The environment looks perfectly symmetric to the agent (Fig. 19(b))!!

Suppose now that the agent relies on boundary regions. Let *P, Q, R*, be the topological places associated with *d1, d3* and *d5* respectively. From Axiom (34), let Pa, Pb, dir_a and dir_b be such that

$order(Pa, dir_a, P, Q), \quad along(ds2, Pa, dir_a), \quad along(ds3, Pa, dir_a),$
 $order(Pb, dir_b, Q, R), \quad along(ds4, Pb, dir_b), \quad along(ds5, Pb, dir_b),$




Fig. 19. (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.

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 (Axiom (60)) *in_region*($R, right(Pa, dir_a)$) where *right*(Pa, dir_a) denotes the right region of Pa when facing dir_a (Axiom (51)). 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 Fig. 19(c).

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 (see Example 19). 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 distance between the “end places” of the travel action, and after turning, the agent may have an estimate of the angle turned. 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 [31].

Action's executions 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. Local analog maps [4,15,60] can also be associated with places [31].

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. In order to propagate uncertainty about the real value of positions and headings we use the compound and merge operations [56]. These operations take different forms depending on how one represents uncertainty as well as on the dimensionality of the variables' domains. In our current work we use *intervals* to represent uncertainty in position and headings, and the compound and merge operations

correspond to add and intersect intervals, respectively. Nevertheless, the discussion in this section applies to other forms of representing uncertainty as long as the compound and merge operations are provided for that representation.

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 $\langle ds, (\mathbf{type\ Int}), ds' \rangle$, 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, representing the “distance” with respect to an arbitrary but fixed place on the path. This 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 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.

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:

$$position1(pa, p, pos) \wedge position1(pa, p, pos') \rightarrow pos = pos', \quad (65)$$

$$position1(pa, p, pos) \rightarrow on(pa, p). \quad (66)$$

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:

$$path_distance(pa, p, q, d) \equiv \quad (67)$$

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

Estimates of the distance between places on a path are gathered while the agent navigates the environment. The predicate **path_distance** \approx (**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,

$$\begin{aligned}
& \{ \text{minpath_distance}^{\approx}: \\
& \quad [\langle ds, (\text{travel } I_d), ds' \rangle \wedge \text{at}(ds, p) \wedge \text{at}(ds', q) \wedge \text{along}(ds, pa, dir) \wedge \quad (68) \\
& \quad \quad \text{along}(ds', pa, dir)] \rightarrow \text{path_distance}^{\approx}(pa, p, q, I_d), \\
& \quad [\text{order}(pa, dir, p, q) \wedge \text{order}(pa, dir, q, r) \wedge \text{path_distance}^{\approx}(pa, p, q, I_{pq}) \wedge \quad (69) \\
& \quad \quad \text{path_distance}^{\approx}(pa, q, r, I_{qr})] \rightarrow \text{path_distance}^{\approx}(pa, p, r, I_{pq} + I_{qr}) \\
& \}
\end{aligned}$$

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 $\text{path_distance}^{\otimes}(pa, p, r, I_d)$ denotes the merging of distance estimates:

$$\text{path_distance}^{\otimes}(pa, p, r, I) \equiv_{\text{def}} I = \bigcap \{ I_{est} : \text{path_distance}^{\approx}(pa, p, q, I_{est}) \}. \quad (70)$$

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

$$\text{path_distance}^{\otimes}(pa, p, q, I_d) \rightarrow \exists d \in I_d \text{ path_distance}(pa, p, q, d). \quad (71)$$

When the agent has distance estimates available, $\text{path_distance}^{\otimes}(pa, p, q, I_d)$ is always the case for some interval I_d . In a topological map $I_d \neq \emptyset$ (Axiom (71)) and it should be possible to assign locations to places on a path as specified by Axiom (67). The actual values of positions are not that important (there could be many ways to satisfy the metrical constraints). Their main use is to rule out possible interpretations of the theory where such positions do not exist given Axiom (71).

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 **radial**(**p**, **pa**, **dir**, **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. Estimates 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 **angle**(**p**, **pa**, **dir**, **pa1**, **dir1**, **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*; $\text{angle}^{\approx}(\mathbf{p}, \mathbf{pa}, \mathbf{dir}, \mathbf{pa1}, \mathbf{dir1}, I_{\text{ang}})$, I_{ang} is an estimate of the angle at place *p* between path *pa* in direction *dir* and path *pa1* in direction *dir1*.

7.3. Two dimensional frames of reference

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. Having 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 [31,40]. 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 [16,41].

The 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 [18,56]. The robotics community has developed algorithms to solve a network of spatial relations [12–14,41]. Techniques from multidimensional scaling [6] and nonlinear programming [45] can also be used.

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 **location2**(p, q, 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).

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:

$$\begin{aligned} & location2(p, p1, l_{p1}) \wedge location2(p, p2, l_{p2}) \wedge path_distance^{\otimes}(pa, p1, p2, I_d) \\ & \rightarrow |l_{p1} - l_{p2}| \in I_d, \end{aligned} \quad (72)$$

$$\begin{aligned} & [location2(p, p1, l_{p1}) \wedge location2(p, p2, l_{p2}) \wedge location2(p, p3, l_{p3}) \wedge \\ & order(pa, dir, p1, p2) \wedge order(pa', dir', p2, p3) \wedge \\ & angle^{\otimes}(p2, pa, dir, pa', dir', I_{ang})] \\ & \rightarrow angle(-l_{p2} \vec{l}_{p1}, l_{p2} \vec{l}_{p3}) \in I_{ang}, \end{aligned} \quad (73)$$

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

Axioms (72) and (73) 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 [27,42]. For instance, in [27] 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 and reasoning about a path's shape, we have left this description outside the scope of this work.

Using different metrical estimation approaches requires a reworking of the axioms in this section. In such case, the compound and merge operations (Axioms (68)–(70)) should be described differently. It is not difficult to define compound and merge operations for Gaussian representations of metrical uncertainty. More care will be required to update Axiom (71) which is used to refute inconsistent hypotheses, since no combination of Gaussians is logically inconsistent. A greater change will be needed in order to take into account the shape of paths when creating two dimensional frames of reference. Nevertheless, the presented axiomatization defines where in the theory the metrical information comes into place and suggests the type of axioms that need to be added.

7.4. Combining topological and metrical information

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,

$$\langle ds : ds_1, \dots, ds_n \rangle \equiv_{def} \quad (74)$$

$$\exists p \left\{ at(ds, p) \wedge \bigwedge_{i=1}^n \exists p_i, l_i [at(ds_i, p_i) \wedge location2(p, p_i, l_i)] \right\}.$$

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 (74). 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 (74). These last two requirements can be stated as follows:

$$\{min\ location2: 2D_Frames\}. \quad (75)$$

The topological theory includes local metrical information by adding Axioms (65) to (75) inside the block **AT_block** (Block (20)). The priority of predicates in the circumscription policy associated with AT_block remains the same. The predicates varied in the circumscription policy now include those predicates use to describe metrical information: *radial*, *position1*, *position2*, *path_distance*, *path_distance*[≈], *path_distance*[⊗], *angle*, *angle*[≈] and *angle*[⊗].

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

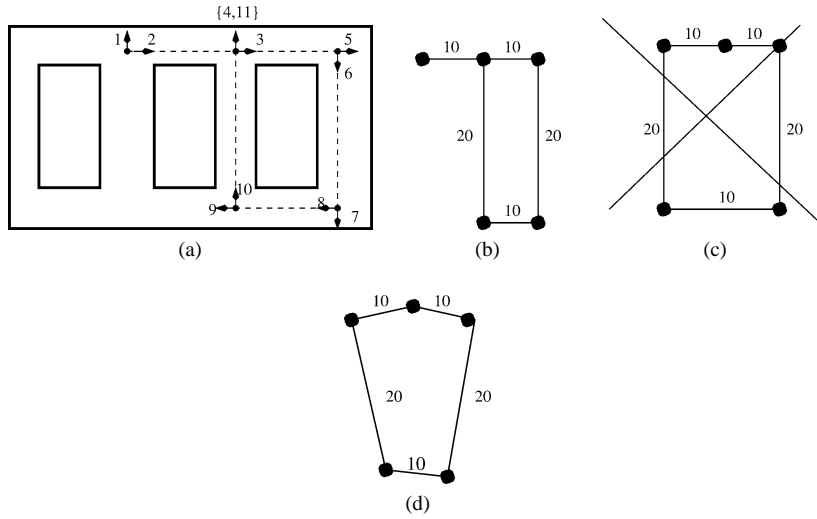


Fig. 20. (a) The robot goes around the block visiting distinctive states $ds1$ to $ds11$ in the order suggested by the figure. Distinctive state $ds11$ is observed at the same environment state as $ds4$. Assume distinctive states $ds1$ and $ds4$ look alike to the agent. (b) and (c) represent two possible topological maps for the environment in (a) (see Example 13). 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).

Example 18. Consider Example 13 where two topological maps are consistent with the agent’s experiences (see Fig. 20). 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 both options $teq(ds4, ds11)$ and $teq(ds1, ds11)$ are topologically possible (Figs. 20(b), (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 $\langle ds1, (turn [-110^\circ, -80^\circ], ds2) \rangle$).⁹ 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.

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 (65)–(75) 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.

⁹ Whenever we use a number x instead of an interval, it is an abbreviation for $[x, x]$.




Fig. 21. (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 ($P \neq S$).

Example 19. Consider an agent visiting the different corners of a square room in the order suggested by Fig. 21(a). 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 17). Suppose the agent has access to perfect metrical information and uses it while building the metrical map.

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 Fig. 21(c).

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 [19] of a logic program. This logic program is implemented in Smodels [43] as illustrated in [48]. 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. A search state is implemented by a partial model, $pmodel$. A partial model of $TT(E)$ is a model of $TT(E')$, for some $E' \subseteq E$ (Section 8.1). 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.

```

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;
}

```

Fig. 22. 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. Fig. 23 defines how a pmodel explains a schema and when extensions are created.

At each step of the search a schema $\langle ds, a, ds' \rangle$ has to be explained. Either the identity 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' . The identity of the schema's context (i.e., ds in $\langle ds, a, ds' \rangle$) is known at each step in the search.

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 Figs. 22 and 23.

The three key steps in the search are (Fig. 23): 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

$$x \stackrel{pmodel}{=} y \wedge x \stackrel{pmodel}{\neq} y \rightarrow inconsistent(pmodel),$$

$$teq(x, y) \in pmodel \wedge \neg teq(x, y) \in pmodel \rightarrow inconsistent(pmodel).$$

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 the rule-based system Algernon [8] as our theorem prover. In Section 8.2 we present an illustrative trace of the algorithm.

```

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
}

```

Fig. 23. 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 $TT(E)$.

8.1. Implementation

Our logic for partial models takes the basic ideas developed in the area of formal reasoning about contexts [39]. 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:

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

Create candidates. *Possible-equal-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:

$$s = \langle ds, turn, ds' \rangle \wedge at(ds, p) \wedge at(ds'', q) \wedge p \neq q \rightarrow \neg teq(ds', ds'') \quad (76)$$

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

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

The rules above are derived from the axioms in our theory. For instance, rule (76) 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. 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., $pmodel \models at(ds, p)$).

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

$$view(ds_1, v_1) \wedge view(ds_2, v_2) \wedge v_1 \neq v_2 \rightarrow \neg teq(ds_1, ds_2) \quad (77)$$

$$\langle ds, turn, ds' \rangle \rightarrow \neg teq(ds, ds'), \quad (78)$$

$$order(pa, dir, p, q) \rightarrow p \neq q, \quad (79)$$

$$radial(p, ds_1, h_1) \wedge radial(p, ds_2, h_2) \wedge h_1 \neq h_2 \rightarrow ds_1 \neq ds_2, \quad (80)$$

$$\begin{aligned} & position1(pa, dir, p_1, pos_1) \wedge \\ & position1(pa, dir, p_2, pos_2) \wedge pos_1 \neq pos_2 \rightarrow p_1 \neq p_2, \end{aligned} \quad (81)$$

$$leftOf(pa, dir, p) \wedge on(pa, q) \rightarrow p \neq q, \quad (82)$$

$$leftOf(pa, dir, p) \wedge on(pa_1, p) \rightarrow pa \neq pa_1, \quad (83)$$

$$at(ds, p) \wedge at(ds, q) \rightarrow p = q, \quad (84)$$

$$along(ds, pa, dir) \wedge along(ds, pa_1, dir_1) \rightarrow pa = pa_1 \wedge dir = dir_1. \quad (85)$$

Rules (77) and (78) rely on the fact that dstates have a unique view and turn actions link different distinctive states (Axioms (22) and (39)). Rule (79) 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 (80) and (81) use radial and one dimensional frames of reference to conclude inequality of dstates and places, respectively (Axiom (65)). Rules (82) and (83)

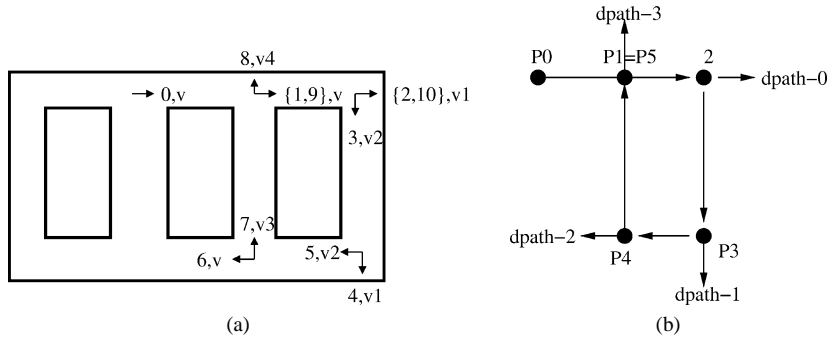


Fig. 24. (a) Numbers identify the dstates created by the map building algorithm. Views associated with dstates are also shown. Dstates 1 and 9 are at the same environment location. (b) Places and dpaths created by the map building algorithm. Notice that $P1$ and $P5$ are two names for the same place.

use boundary relations in order to distinguish places and paths respectively (Axiom (58)).¹⁰ Rules (84) and (85) state that each distinctive state is at a unique place, along a unique path direction (Axioms (29) and (37)).

Assert schema. $Assert\text{-}schema(pmodel, s)$ 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 teq 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 Fig. 24(a). 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 maps are possible: either the agent is back to dstate 1 or dstate 0 (this is Example 13). After traveling from dstate 9 to dstate 10, only one topological map is possible (Fig. 24(b)). Fig. 25 illustrates the use of the topological rules to distinguish distinctive states that share the same view. Fig. 26 shows when branches in the search are created and how they can be refuted as more information becomes available to the agent.¹¹

¹⁰ $leftOf(pa, dir, p)$ in the implementation is an abbreviation for Section 6.2's longer expression $leftOf(pa, dir, lr) \wedge in_region(p, lr)$.

¹¹ In 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.

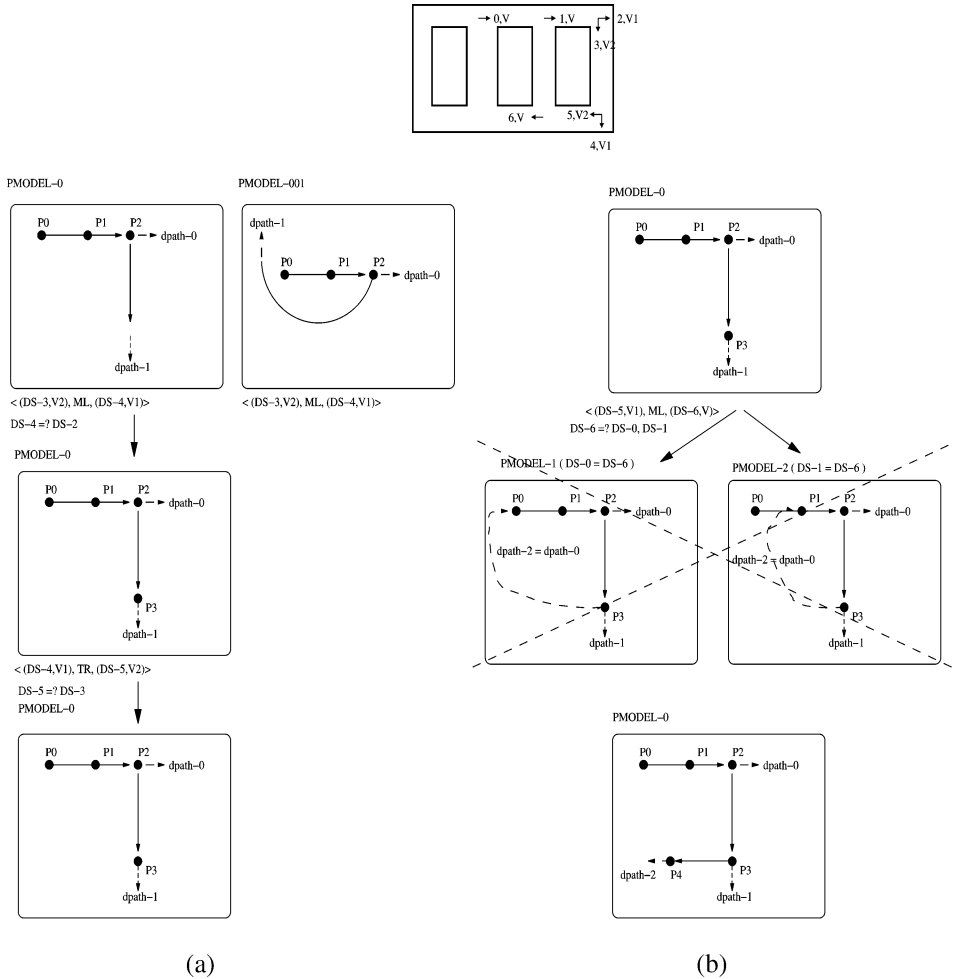


Fig. 25. (a) When the agent reaches dstate 4, the same view v_1 has been observed at dstate 2. Since dstate 2 is along $dpath-0$ and dstate 4 will be along $dpath-1$ (the agent just traveled from dstate 3 along $dpath-1$), dstate 4 and 2 are proven different. $Dpaths$ 0 and 1 are different since there is a *turnRight* action relating them (Axiom (39)). Place P_3 is created to be the place dstate 4 is at (Axiom (29)). Place P_3 is proven different from place P_2 since P_2 is before P_3 along $dpath-1$ (Corollary 1). Consequently, dstates 5 and 3 are proven different. There are however two possible models depending whether P_3 is to the right or not of $dpath-0$. Our boundary regions circumscription policy (Section 6.3) prefers PMODEL-0 in which P_3 is to the right of $dpath-0$ over PMODEL-001 in which no boundary relations exist. In this example, the search will never explore further the branch associated with PMODEL-001 because the branch associated with PMODEL-0 leads to a consistent map for the given experiences. (b) The agent travels to dstate 6 along $dpath-2$. Because P_3 is to the right of $dpath-0$, $dpath-2$ cannot be the same as $dpath-0$, which makes PMODEL-1 and PMODEL-2 inconsistent. The only remaining (and hence minimal) model is PMODEL-0, in which dstate 6 is different from dstates 0 and 1.

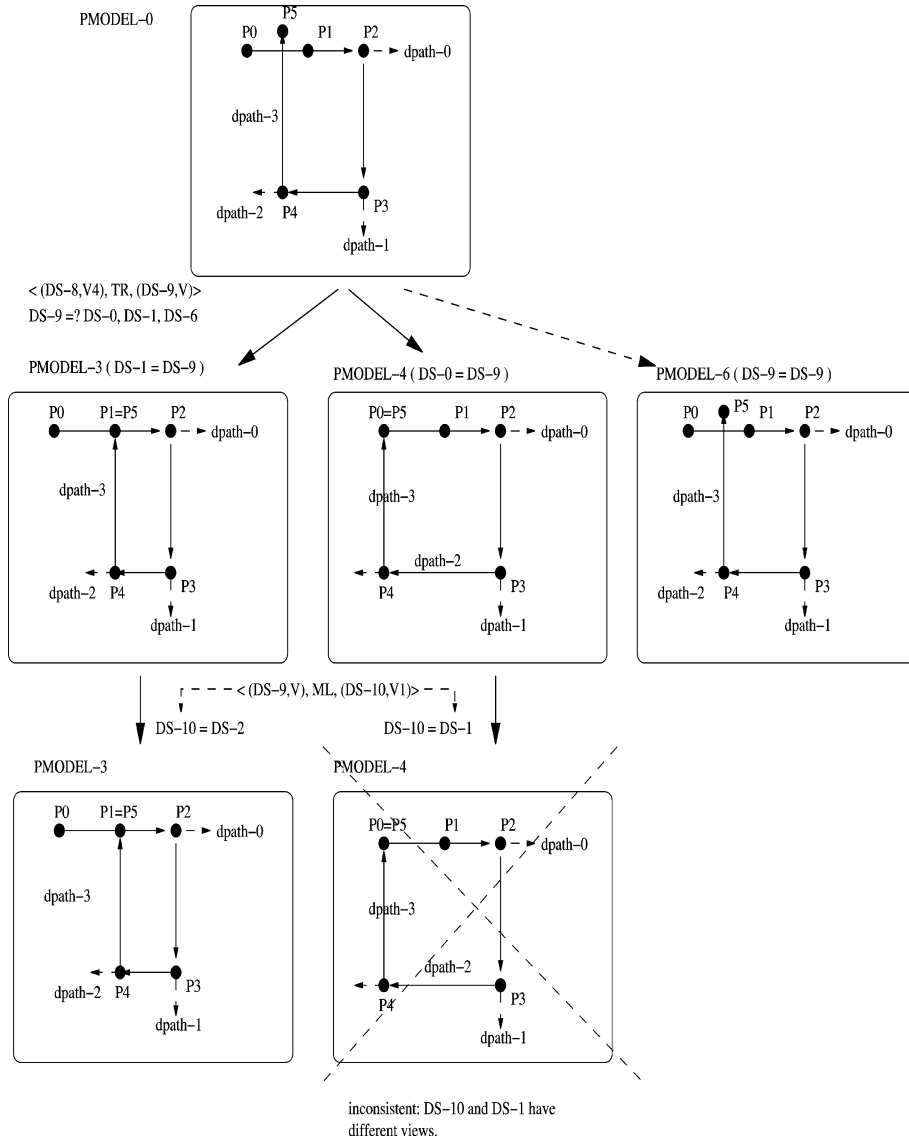


Fig. 26. By the time the agent reaches dstate 8, six places (not five) are part of the map. Places $P5$ and $P1$ are not equal since the dstates are these places are not yet turn related (Axiom (30)). Turning from dstate 8 to 9 leaves the agent with the three possibilities: (pmodel-3) dstates 9 and 1 are equal (and so $P1 = P5$, or (pmodel-4) dstates 9 and 0 are equal (and so $P0 = P1$), or (pmodel-6) dstate 9 is a new different dstate. That dstates 9 and 6 are different follows from the fact that places $P4$ and $P5$ 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 (9, ML , 10) 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 4–7 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.

The circumscription priority ordering embedded in the theory is a result of our research, as we experimented with various orders to determine which ordering defined models that corresponded to what is intuitively the “correct map” of the environment. Because we have no formal definition of what the correct map is, it is impossible to prove mathematically that the circumscription priority ordering is the correct one. Possibly future research can provide such a formal definition, but the difficulties arise from handling partial experience in the environment, or highly symmetrical environments with a great deal of perceptual aliasing.

How useful is this theory? This work defines topological maps independently of the algorithms used to create such maps. The theory is general in that it covers the major ideas in the field of spatial representation using topological maps. The theory is useful in that it specifies the minimal set of objects and relationships any topological map building implementation should have. Although our theory covers most of the known ideas about topological maps, it is not just a union of previous work in a common framework. The theory defines different spatial ontologies (causal, topological, metrical), illustrates what spatial knowledge is captured by each ontology, and then shows the relationships among these ontologies. The theory shows how the combined spatial knowledge associated with the different ontologies results in a different “map” from the one associated with each independent ontology.

The axiomatic theory has practical value. It has been used to build cognitive maps by both physical and simulated robots. [24,29,49] explicitly use the axiomatic theory described in this paper as well as the implemented algorithms in order to build topological

maps. [49] shows how a wheelchair robot builds the topological map of a building's floor. The major focus of this work was on testing the applicability and correctness of the axioms and algorithms here described.

In [24] topological maps are built as a mean to disambiguate distinctive states with the same view. The map provides an unambiguous assignment of distinctive states to views, which can then be used by the robot to “refine its views” so that it is possible to distinguish distinctive states from sensory information alone. “Lassie (the robot)... collected 240 images from 20 distinctive states. The topological map linking them contained seven places and four paths... By building the causal and topological map the robot is able to disambiguate all twenty distinctive states, even though there are only ten different views...” [24].

Finally, Fig. 9 in [29] describes experimental results where a simulated agent builds a topological map and learns boundary relations for grid-like environments. This work presents a computational hypothesis that describes how the “skeleton” of major paths emerges from the interaction of three factors: “(i) the topological map is represented as a bipartite graph of places and paths, where a path is a one-dimensional ordered set of places, (ii) a traveler incrementally accumulates topological relationships, including the relation of a place to a path serving as a dividing boundary separating two regions; and (iii) the wayfinding algorithm prefers paths rich in boundary relations so they are likely to acquire more boundary relations. This positive-feedback loop leads to an oligarchy of paths rich in boundary relations (i.e., the skeleton in the cognitive map)” [29].

Acknowledgements

We are grateful to Vladimir Lifschitz for his valuable feedback during this work. We also thank the anonymous referees for their helpful questions and suggestions.

Appendix A. Nested abnormality theories

In this appendix we define circumscription and nested abnormality theories following [35,36]. 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 A.1 (*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

$$A(P, Z_1, \dots, Z_m) \wedge \neg \exists p, z_1, \dots, z_m [A(p, z_1, \dots, z_m) \wedge p < P] \quad (\text{A.1})$$

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 (A.1) by $CIRC[A; P; Z]$.

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 A.2 (*Parallel circumscription*). The *parallel circumscription*

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

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

Definition A.3 (*Prioritized circumscription*). The *prioritized circumscription*

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

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

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

The formula $p < P$ defines a *lexicographic* order among the predicates in p and P . Proposition 15 in [35] shows that prioritized circumscription can be reduced to parallel circumscription as follows:

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

$$CIRC[A \wedge \text{not_}P_i \equiv \neg P_i; P^1 \succ \dots \text{not_}P_i \dots \succ P^n; Z, P_i]$$

where $\text{not_}P_i$ is a new constant predicate not occurring in A .

A.1. Nested abnormality theories (NATs)

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 A.4 (NATs). 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 NATs 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:

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

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 A.5. 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 NATs is extended such that

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

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 A.2. *Let A be a sentence such that Ab does not occur in A . Then,*

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

Appendix B. ceq properties

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

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

$$\begin{aligned} &\{\max \text{ceq}: \\ &\text{ceq}(ds_1, ds_2) \rightarrow \text{View}(ds_1, v) \equiv \text{View}(ds_2, v), \end{aligned}$$

$$\begin{aligned} & ceq(ds_1, ds_2) \wedge \langle ds_1, a, ds'_1 \rangle \wedge \langle ds_2, a, ds'_2 \rangle \rightarrow ceq(ds'_1, ds'_2) \\ & \} \end{aligned}$$

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

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

We are to prove that M_2 is a model for the axioms inside the CEQ_block and consequently $CEQ_block \models ceq(ds, ds)$.¹² Indeed,

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

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

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

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

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

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

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

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

¹² M_2 satisfies the other axioms in $CT(E)$ since ceq does not occur in them.

$$\begin{aligned} & \stackrel{(15)}{\rightarrow} ceq^{M_1}(ds'_1, ds'_2) \vee ds'_1 = ds'_2 \\ & \equiv ceq^{M_2}(ds'_1, ds'_2). \end{aligned}$$

Let us 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 us 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,

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

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 $E \models \langle a_1, \dots, a_{n-1} \rangle(ds), a_n, ds'$. 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,

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

Appendix C. teq properties

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

$$\begin{aligned}
TT(E) = & \\
& \text{there exist infinitely many places,} \\
& \text{there exist infinitely many paths,} \\
& \neg \exists p [tplace(p) \wedge is_region(p)], \\
& \neg \exists pa [tpath(pa) \wedge route(pa)], \\
& COMPLETION(E), \\
& \text{Axioms (2)–(10),} \\
& \langle ds, a, ds' \rangle \wedge \langle ds, a, ds'' \rangle \rightarrow ds' = ds'', \text{ (Axiom (15))} \\
& T_block, \\
& AT_block = \tag{C.1}
\end{aligned}$$

$$\begin{aligned}
& \{ maxteq: \\
& \quad \Gamma \\
& \quad \mathbf{circ} tpath \succ tplace \mathbf{var} SSH\vec{pred} \tag{C.2} \\
& \}
\end{aligned}$$

where Γ is the set of axioms defined on Block (20) (Section 5.2), and $SSH\vec{pred}$ stands for the tuple of predicates $\langle \mathbf{at}, \mathbf{along}, \mathbf{order}, \mathbf{on}, \mathbf{teq}, \mathbf{turn_eq}, \mathbf{travel_eq} \rangle$.

Proposition C.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.

$$\begin{aligned}
& CIRC[\Gamma; tpath \succ tplace; SSH\vec{pred}] \\
& \equiv \quad \{\text{Proposition 15 in [35]}\} \\
& CIRC[\Gamma; tpath; tplace, SSH\vec{pred}] \wedge CIRC[\Gamma; tpath, tplace; SSH\vec{pred}] \\
& \rightarrow \quad \{\text{def. of circumscription}\} \\
& CIRC[\Gamma; tpath].
\end{aligned}$$

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

$$\begin{aligned}
& CIRC[\Gamma; tpath] \\
& \equiv \\
& CIRC[\Gamma'(tpath) \wedge [along(ds, pa, dir) \rightarrow tpath(pa)]; tpath] \\
& \equiv \quad \{\text{Proposition 4 in [35]}\} \\
& \Gamma'(tpath) \wedge CIRC[along(ds, pa, dir) \rightarrow tpath(pa); tpath] \\
& \rightarrow \quad \{\text{Proposition 1 in [35]}\} \\
& [\exists ds, dir \text{ along}(ds, pa, dir)] \equiv tpath(pa).
\end{aligned}$$

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

$$\begin{aligned}
& CIRC[\Gamma; tpath \succ tplace; SSHpred] \\
& \rightarrow \quad \{\text{see above}\} \\
& CIRC[\Gamma; tpath, tplace; SSHpred] \\
& \rightarrow \quad \{\text{def. parallel circumscription}\} \\
& CIRC[\Gamma; tpath, tplace] \\
& \rightarrow \quad \{\text{def. parallel circumscription}\} \\
& CIRC[\Gamma' \wedge [at(ds, p) \rightarrow tplace(p)]; tplace] \\
& \equiv \quad \{\text{Propositions 1 and 4 in [35]}\} \\
& \Gamma' \wedge [\exists ds, at(ds, p)] \equiv tplace(p). \quad \square
\end{aligned}$$

Proposition C.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)), Proposition C.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)), from Proposition C.1 we conclude that the set of topological places in a model of $TT(E)$ is finite. \square

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

$$\forall ds_2 \notin [ds_1]_{turn}, [ds_2]_{teq} \cap [ds_1]_{turn} = \emptyset. \quad (C.3)$$

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,

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

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,

$$\begin{aligned}
&turn_eq(ds_1, ds_2) \\
&\rightarrow \\
&b_{j'} \in [ds_1]_{\widehat{turn}} \\
&\rightarrow \{teq(b_j, b_{j'})\} \\
&[b_j]_{teq} \cap [ds_1]_{\widehat{turn}} \neq \emptyset \\
&\rightarrow \{\text{C.3}\} \\
&b_j \in [ds_1]_{\widehat{turn}} \\
&\rightarrow \{\widehat{turn}(b_{(j-1)' }, b_j)\} \\
&b_{(j-1)' } \in [ds_1]_{\widehat{turn}} \\
&\rightarrow \\
&false.
\end{aligned}$$

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 (C.1)). 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 $M_2' < M_2$, where $<$ is the order defined by the circumscription policy C.2. This proves that M_1 and M_2 have the same number of topological paths. 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:

- $tpath^{M_2} = \phi(tpath^{M_1})$, $tplace^{M_2} = \phi(tplace^{M_1})$.
- $teq^{M_2} = \phi(teq^{M_1}) = \{teq(ds_1, ds_2) : M_1 \models teq(ds_1, ds_2)\} = teq^{M_1}$.
- $at^{M_2} = \phi(at^{M_1}) = \{at(ds, \phi(p)) : M_1 \models at(ds, p)\}$.
- $along^{M_2} = \phi(along^{M_1}) = \{along(ds, \phi(pa), dir) : M_1 \models along(ds, pa, dir)\}$.
- $order^{M_2} = \phi(order^{M_1}) = \{order(\phi(pa), dir, \phi(p), \phi(q)) : M_1 \models order(pa, dir, p, q)\}$.
- $on^{M_2} = \phi(on^{M_1}) = \{on(\phi(pa), \phi(p)) : M_1 \models on(pa, p)\}$.
- $turn_eq^{M_2} = \phi(turn_eq^{M_1}) = turn_eq^{M_1}$.
- $travel_eq^{M_2} = \phi(travel_eq^{M_1}) = travel_eq^{M_1}$.

Notice that the language of Γ is defined by $\{tpath, tplace\} \cup SSHpred$. Thus $M_1 \models \Gamma$ implies $\phi(M_1) \models \Gamma$. Notice 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).

Since $\phi(tpath^{M_1}) \subset tpath^{M_2}$, then $\phi(M_1) < M_2$, 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

Appendix D. Theory axioms

The block **T_block** inside Block (19) in Section 5.2 defines the properties of the predicates \widehat{turn} , \widehat{travel} , and \overline{travel} . \widehat{turn} is the equivalence closure of the schemas $\langle \cdot, turn, \cdot \rangle$; \widehat{travel} and \overline{travel} are the equivalence and transitive closure of the schemas $\langle \cdot, travel, \cdot \rangle$ respectively.¹³

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

¹³ A 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 $\{C_1, \dots, C_n, \min P_1 : A_1, \dots, A_m\}, \dots, \{C_1, \dots, C_n, \min P_k : A_1, \dots, A_m\}$.

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

References

- [1] D. Angluin, On the complexity of minimum inference of regular sets, *Inform. and Control* 39 (1978) 337–350.
- [2] K. Basye, T. Dean, L.P. Kaelbling, Learning dynamics: System identification for perceptually challenged agents, *Artificial Intelligence* 72 (1) (1995) 139–171.
- [3] A.F. Beardon, *Complex Analysis*, Wiley, New York, 1979.
- [4] J. Borenstein, Y. Koren, The vector field histogram—Fast obstacle-avoidance for mobile robots, *IEEE J. Robotics and Automation* 7 (3) (1991) 278–288.
- [5] J. Borenstein, H.R. Everett, L. Feng, *Navigating Mobile Robots: Systems and Techniques*, A.K. Peters, Wellesley, MA, 1996.
- [6] I. Borg, P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer, New York, 1997.
- [7] H. Choset, K. Nagatani, Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization, *IEEE Trans. on Robotics and Automation* 17 (2) (2001) 125–137.
- [8] J. Crawford, B. Kuipers, Algernon: A tractable system for knowledge representation, *SIGART Bull.* 2 (3) (1991) 35–44.
- [9] E. Davis, The MERCATOR representation of spatial knowledge, in: A. Bundy (Ed.), *Proc. IJCAI-83*, Karlsruhe, Germany, Morgan Kaufmann, Los Altos, CA, 1983, pp. 295–301.
- [10] T. Dean, K. Basye, L. Kaelbling, Uncertainty in graph-based map learning, in: J.H. Connell, S. Mahadevan (Eds.), *Robot Learning*, Kluwer Academic, Dordrecht, 1993, pp. 171–192.
- [11] G. Dudek, M. Jenkin, E. Milios, D. Wilkes, Robotic exploration as graph construction, *IEEE Trans. on Robotics and Automation* 7 (6) (1991) 859–865.
- [12] H.F. Durrant-Whyte, Consistent integration and propagation of disparate sensor observations, *Internat. J. Robotics Res.* 6 (3) (1987) 3–24.
- [13] H.F. Durrant-Whyte, *Integration, Coordination and Control of Multisensor Robot Systems*, Kluwer Academic, Boston, MA, 1988.
- [14] H.F. Durrant-Whyte, Uncertain geometry in robotics, *IEEE J. Robotics and Automation* 5 (6) (1988) 23–31.
- [15] A. Elfes, Sonar-based real-world mapping and navigation, *IEEE J. Robotics and Automation* 3 (3) (1987) 249–265.
- [16] S.P. Engelson, D.V. McDermott, Error correction in mobile robot map learning, in: *Proc. IEEE International Conference on Robotics and Automation*, Nice, France, 1992, pp. 2555–2560.
- [17] M. Franz, B. Schölkopf, H.A. Mallot, H.H. Bülthoff, Learning view graphs for robot navigation, *Autonomous Robots* 5 (1998) 111–125.
- [18] A. Gelb, *Applied Optimal Estimation*, MIT Press, Cambridge, MA, 1974.
- [19] M. Gelfond, V. Lifschitz, Classical negation in logic programs and disjunctive databases, *New Generation Computing* 9 (1991) 365–385.
- [20] E. Mark Gold, Complexity of automaton identification from given data, *Inform. and Control* 37 (1978) 302–320.
- [21] S. Gopal, R.L. Klatzky, T.R. Smith, Navigator: A psychologically based model of environmental learning through navigation, *J. Environmental Psychology* 9 (1989) 309–331.
- [22] S. Koenig, R. Simmons, Passive distance learning for robot navigation, in: *Proceedings of the Thirteenth International Conference on Machine Learning (ICML)*, Bari, Italy, 1996, pp. 266–274.

- [23] D. Kortenkamp, E. Chown, S. Kaplan, Prototypes, locations, and associative networks (PLAN): Towards a unified theory of cognitive mapping, *Cognitive Sci.* 19 (1995) 1–51.
- [24] B. Kuipers, P. Beeson, Bootstrap learning for place recognition, in: *Proc. AAAI-02*, Edmonton, AB, AAAI Press, 2002, pp. 174–180.
- [25] B. Kuipers, Y.T. Byun, A robust qualitative method for robot spatial learning, in: *Proc. AAAI-88*, St. Paul, MN, 1988, pp. 774–779.
- [26] B. Kuipers, Y.T. Byun, A robot exploration and mapping strategy based on semantic hierarchy of spatial representations, *Robotics Autonomous Syst.* 8 (1991) 47–63.
- [27] B. Kuipers, T. Levitt, Navigation and mapping in large-scale space, *AI Magazine* 9 (2) (1988) 25–43.
- [28] B. Kuipers, R. Froom, W.Y. Lee, D. Pierce, The semantic hierarchy in robot learning, in: J. Connell, S. Mahadevan (Eds.), *Robot Learning*, Kluwer Academic, Dordrecht, 1993, pp. 141–170.
- [29] B. Kuipers, D. Tecuci, B. Stankiewicz, The skeleton in the cognitive map: A computational and empirical exploration, *Environment and Behavior* 35 (1) (2003) 80–106.
- [30] B. Kuipers, Modeling spatial knowledge, *Cognitive Sci.* 2 (1978) 129–153.
- [31] B. Kuipers, The spatial semantic hierarchy, *Artificial Intelligence* 119 (2000) 191–233.
- [32] B.C. Kuo, *Automatic Control Systems*, Fifth Edition, Prentice-Hall, Englewood Cliffs, NJ, 1987.
- [33] W.Y. Lee, *Spatial semantic hierarchy for a physical mobile robot*, PhD Thesis, The University of Texas at Austin, 1996.
- [34] D. Leiser, A. Zilbershatz, THE TRAVELLER: A computational model of spatial network learning, *Environment and Behavior* 21 (4) (1989) 435–463.
- [35] V. Lifschitz, Circumscription, in: *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3, Oxford University Press, Oxford, 1994, pp. 297–352.
- [36] V. Lifschitz, Nested abnormality theories, *Artificial Intelligence* 74 (2) (1995) 351–365.
- [37] K. Lynch, *The Image of the City*, MIT Press, Cambridge, MA, 1960.
- [38] H.A. Mallot, S. Gillner, Route navigating without place recognition: What is recognized in recognition-triggered responses? *Perception* 29 (2000) 43–55.
- [39] J. McCarthy, S. Buvač, Formalizing context (expanded notes), in: A. Aliseda, R.J. van Glabbeek, C. Westerståhl (Eds.), *Computing Natural Language*, in: *CSLI Lecture Notes*, Vol. 8L, 1998, pp. 13–50.
- [40] D.V. McDermott, E. Davis, Planning routes through uncertain territory, *Artificial Intelligence* 22 (1984) 107–156.
- [41] P. Moutarlier, R. Chatila, Stochastic multisensory data fusion for mobile robot location and environment modelling, in: *Proc. 5th International Symposium on Robotics Research*, 1989, pp. 85–89.
- [42] A. Musto, K. Stein, K. Schill, A. Eisenkolb, W. Brauer, Qualitative motion representation in egocentric and allocentric frames of reference, in: *Proc. Fourth International Conference on Spatial Information Theory (COSIT'99)*, Springer, Berlin, 1999.
- [43] I. Niemelä, P. Simons, Smodels—An implementation of the stable model and well-founded semantics for normal logic programs, in: *Proc. 4th International Conference on Logic Programming and Nonmonotonic Reasoning*, in: *Lecture Notes in Computer Science*, Vol. 1265, Springer, Berlin, 1997, pp. 420–429.
- [44] M. O'Neill, A biologically based model of spatial cognition and wayfinding, *J. Environmental Psychology* 11 (1991) 299–320.
- [45] A.L. Peressini, F.E. Sullivan, K.J. Uhl, *The Mathematics of Nonlinear Programming*, Springer, New York, 1988.
- [46] J. Piaget, B. Inhelder, *The Child's Conception of Space*, Norton, New York, 1967, First published in French, 1948.
- [47] E. Remolina, B. Kuipers, Towards a formalization of the Spatial Semantic Hierarchy, in: *Proc. Fourth Symposium on Logical Formalizations of Commonsense Reasoning*, London, 1998.
- [48] E. Remolina, B. Kuipers, A logical account of causal and topological maps, in: *Proc. IJCAI-01*, Seattle, WA, AAAI Press, Menlo Park, CA, 2001, pp. 5–11.
- [49] E. Remolina, *A logical account of causal and topological maps*, PhD Thesis, The University of Texas at Austin, 2001.
- [50] R.L. Rivest, R.E. Schapire, A new approach to unsupervised learning in deterministic environments, in: *Proceedings of the Fourth International Workshop on Machine Learning*, 1987.
- [51] B. Schölkopf, H. Mallot, View-based cognitive mapping and path planning, *Adaptive Behavior* 3 (1995) 311–348.

- [52] M.P. Shanahan, Noise and the common sense informatic situation for a mobile robot, in: Proc. AAAI-96, Portland, OR, 1996, pp. 1098–1103.
- [53] H. Shatkay, L. Kaelbling, Learning topological maps with weak local odometry information, in: Proc. IJCAI-97, Nagoya, Japan, 1997.
- [54] A.W. Siegel, S. White, The development of spatial representations of large-scale environments, in: H. Reese (Ed.), *Advances in Child Development and Behavior*, Vol. 10, Academic Press, New York, 1975, pp. 9–55.
- [55] R. Simmons, S. Koenig, Probabilistic robot navigation in partially observable environments, in: Proc. IJCAI-95, Montreal, Quebec, 1995.
- [56] R. Smith, P. Cheeseman, On the representation of and estimation of spatial uncertainty, *Internat. J. Robotics Res.* 5 (1986) 56–68.
- [57] S.D. Steck, H.A. Mallot, The role of global and local landmarks in virtual environment navigation, *Presence* 9 (1) (2000) 69–83.
- [58] J. Tardos, J. Neira, P. Newman, J. Leonard, Robust mapping and localization in indoor environments using sonar data, *Internat. J. Robotics Res.* 21 (6) (2002) 311–330.
- [59] S. Thrun, S. Gutmann, D. Fox, W. Burgard, B. Kuipers, Integrating topological and metric maps for mobile robot navigation: A statistical approach, in: Proc. AAAI-98, Madison, WI, 1998, pp. 989–995.
- [60] S. Thrun, Learning metric-topological maps for indoor mobile robot navigation, *Artificial Intelligence* 99 (1) (1998) 21–71.