



10-2012

Toward a Memory Model for Autonomous Topological Mapping and Navigation: the Case of Binary Sensors and Discrete Actions

Dan P. Guralnik
University of Pennsylvania

Daniel E. Koditschek
University of Pennsylvania, kod@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ease_papers

 Part of the [Electrical and Computer Engineering Commons](#), and the [Systems Engineering Commons](#)

Recommended Citation

Dan P. Guralnik and Daniel E. Koditschek, "Toward a Memory Model for Autonomous Topological Mapping and Navigation: the Case of Binary Sensors and Discrete Actions", . October 2012.

BibTeX entry

```
INPROCEEDINGS{memory_Guralnik_Koditschek_Allerton_2012, author={Guralnik, D.P. and Koditschek, D.E.},
booktitle={Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on},
title={Toward a memory model for autonomous topological mapping and navigation: The case of binary sensors
and discrete actions}, year={2012}, pages={936-945}, keywords={SLAM (robots);learning (artificial
intelligence);mobile robots;navigation;path planning;autonomous goal directed planning;autonomous sensory
experience;autonomous topological mapping;binary sensor;discrete action;learning scheme;memory model;motion
planning;navigation;self-organizing database;symbolic representation;Computational
modeling;Databases;Navigation;Planning;Robot sensing systems;Trajectory}, doi={10.1109/
Allerton.2012.6483319}, }
```

Funded in part by the Air Force Office of Science Research under the MURI FA9550-10-1-0567 and in part by the National Science Foundation under CDI-II- 1028237.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/656
For more information, please contact repository@pobox.upenn.edu.

Toward a Memory Model for Autonomous Topological Mapping and Navigation: the Case of Binary Sensors and Discrete Actions

Abstract

We propose a self-organizing database for perceptual experience capable of supporting autonomous goal-directed planning. The main contributions are: (i) a formal demonstration that the database is complex enough in principle to represent the homotopy type of the sensed environment; (ii) some initial steps toward a formal demonstration that the database offers a computationally effective, contractible approximation suitable for motion planning that can be accumulated purely from autonomous sensory experience. The provable properties of an effectively trained database exploit certain notions of convexity that have been recently generalized for application to a symbolic (discrete) representation of subset nesting relations. We conclude by introducing a learning scheme that we conjecture (but cannot yet prove) will be capable of achieving the required training, assuming a rich enough exposure to the environment.

For more information: [Kod*Lab](#)

Disciplines

Electrical and Computer Engineering | Engineering | Systems Engineering

Comments

BibTeX entry

```
INPROCEEDINGS{memory_Guralnik_Koditschek_Allerton_2012, author={Guralnik, D.P. and Koditschek, D.E.}, booktitle={Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on}, title={Toward a memory model for autonomous topological mapping and navigation: The case of binary sensors and discrete actions}, year={2012}, pages={936-945}, keywords={SLAM (robots);learning (artificial intelligence);mobile robots;navigation;path planning;autonomous goal directed planning;autonomous sensory experience;autonomous topological mapping;binary sensor;discrete action;learning scheme;memory model;motion planning;navigation;self-organizing database;symbolic representation;Computational modeling;Databases;Navigation;Planning;Robot sensing systems;Trajectory}, doi={10.1109/Allerton.2012.6483319}, }
```

Funded in part by the Air Force Office of Science Research under the MURI FA9550-10-1-0567 and in part by the National Science Foundation under CDI-II- 1028237.

Toward a Memory Model for Autonomous Topological Mapping and Navigation: the Case of Binary Sensors and Discrete Actions

Dan P. Guralnik¹ and Daniel E. Koditschek²

Abstract—We propose a self-organizing database for perceptual experience capable of supporting autonomous goal-directed planning. The main contributions are: (i) a formal demonstration that the data-base is complex enough in principle to represent the homotopy type of the sensed environment; (ii) some initial steps toward a formal demonstration that the data base offers a computationally effective, contractible approximation suitable for motion planning that can be accumulated purely from autonomous sensory experience. The provable properties of an effectively trained data-base exploit certain notions of convexity that have been recently generalized for application to a symbolic (discrete) representation of subset nesting relations. We conclude by introducing a learning scheme that we conjecture (but cannot yet prove) will be capable of achieving the required training, assuming a rich enough exposure to the environment.

I. INTRODUCTION

A. Motivation

Mapping and navigation of an environment by an autonomous agent is a deep problem, long and well studied by neuroscientists [1], cognitive scientists [2], AI researchers [3], and, of course, the very earliest pioneers of robotics [4], continuing through the present day [5], [6]. We are interested in designing physical agents endowed with prescribed sensorimotor capabilities that can learn to negotiate useful traversals through an initially poorly known environment, E , in support of goals whose complexity grows in step with its increasing familiarity as manifest by the accumulation of sensorimotor experience.

It is generally accepted that contextual information regarding the state of the environment can be modelled as (at least) a topological space X [7], which, from the point of view of our agent, comes equipped with a *position map* $\text{pos} : X \rightarrow E$ fixing the relation between any given state of the environment and the unique position that might be occupied by the agent given that state [8] (since the agent interacts with the environment, we need to think of the agent's internal state, its position and posture as observables in X).

To qualify as autonomous it seems essential that an agent be able to: (1) maintain an evolving internal representation of the environment and its possible states; (2) identify the current state of the environment; and (3) plan its motion through the environment towards designated targets – or,

at least in support of some sensory designated goal condition. The traditional difficulty in building such autonomous systems arises from the apparently contradictory need for significant prior understanding by the designer of the specific geometric and physical constraints associated with the given task. This contradiction incurs a familiar tradeoff between robust "low level" vs. brittle "high level" approaches to autonomy. In particular, constraints of communication latency, storage space and power requirements tend to push "high level" solutions to any given instance of this problem toward complex and dedicated architectures that typically offer little in the way of formally guaranteed properties and result in systems that empirically operate properly only under numerous (and poorly characterized) assumptions that severely limit the practical applications of the design.

In this paper, we approach the traditional gap between robustness and prior knowledge by interpreting the problem of autonomy as requiring a self-organizing data base for perceptual experience capable of supporting goal-directed planning as well. More specifically, we presume that a machine with adequate mechanical and sensory capabilities, demands an "episodic memory" capable of collecting observations of the state space and inferring information about how its environment is organized in relation to its motor abilities¹. This memory mechanism must keep an efficient record of the causal relations between the actions it is capable of performing and the resulting state transitions; it must use this record to plan its trajectory through the environment; and it must be capable of revising this plan when faced with unexpected input. The main contribution is our effort to formalize (and aim for provable statements about) heretofore poorly defined universality properties that allow an agent to deal with arbitrary real-time input, thus obviating the need to plan in advance for all possible circumstances it may encounter in its environment.

What is a good representation of the state space X for a physically embodied agent? From the point of view of control theory, preference should be given to representations capable of recovering the homotopy type of X : many environments E have a piecewise smooth cellular structure (or are sufficiently well approximated by one for practical purposes), in which case one expects X to have a similar

¹Post doctoral fellow, Electrical & Systems Engineering, School of Engineering and Applied Sciences, University of Pennsylvania, guraldan at seas.upenn.edu

²Alfred Fitler Moore Professor, Electrical & Systems Engineering, School of Engineering and Applied Sciences, University of Pennsylvania, kod at seas.upenn.edu

¹Our construction is greatly impoverished relative to the hegemonic notion of "episodic memory" in psychology that requires: a sense of self; a sense of subjective time; and the ability to distinguish memories of the self at junctures of past and present subjective time [9]. Here, we are really only interested in a mechanism for accumulating episodic experience and organizing it into procedural knowledge for present or future use.

(though much more complicated) description.

As has long been understood within the field of motion planning [10], cellular structures are useful for spatial navigation, as they afford a formally straightforward and (sometimes) computationally tractable two stage approach: first solve the discrete navigation problem of planning a path in the dual graph of the cellular decomposition; then navigate through each of the cells in the discrete solution by applying an appropriate (depending on the problem) flow using local coordinates.

More generally, a flow with sufficiently tame attractor produces a strong deformation retraction of its domain onto its attractor, demonstrating that faithful representation of the homotopy type of X is a natural feature to require from a ‘brain’ for our agent if we expect this agent to exhibit complex and purposive patterns of interaction with its environment.

B. Our Contributions

In view of the above we operate on the assumption that, ultimately, our goal is to construct agents capable of maintaining an evolving, ever more precise, discrete topological model of the space X – say in the form of a finite cell complex – together with a set of computational tools allowing the agent to construct paths between cells (or sets of cells) of the model indexed by its available actions.

Our approach is based on [11], where it was proposed to base the construction of a memory structure on Sageev-Roller duality (see below, section III-B): it was pointed out that maintaining a record \mathbf{P} of the (Boolean) logical implications among the elements of a finite set of binary sensors on any space Z produces a contractible model $\text{Cube}(\mathbf{P})$ of Z , together with a tool for highly efficient (discrete) path planning within the model.

In just a few words, the main idea of [11] is that, given the ability to ask n fixed distinct binary questions about Z (using n fixed binary sensors indexed by \mathbf{P}), not all 2^n combinations of answers to these questions are realized by sensory equivalence classes of points in Z : given known implication relations among the sensors, some combinations of questions are *a priori* nonsensical and should never be considered by the agent; the complex $\text{Cube}(\mathbf{P})$ is *precisely* the result of discarding the corresponding vertices from the n -cube (viewed as a cubical complex), together with all adjacent faces. As luck would have it, complexes of this kind (or, rather, their 1-dimensional skeleta) turn out to possess a very strong convexity theory (see section III-C) enabling very simple path-planning algorithms, as we prove in section IV-B.

The work [11] does not address the practical problem of *how* to effectively obtain a record of implications among the available sensors through sampling. Neither does it provide any way to quantify the discrepancy between the model it proposes and the observed space Z . This work is our initial step in the direction of filling in these gaps.

Our first result is Theorem 3.1, providing sufficient conditions for obtaining a homotopy model of Z by excising a

certain sub-complex of $\text{Cube}(\mathbf{P})$. The contribution of such a result in the context of our class of applications is twofold:

- 1) it provides an explicit criterion for a binary sensorium to be rich enough to produce a homotopy model of the observed state space;
- 2) it states that learning the state space X (or environment E , etc.) up to homotopy using a sufficient binary sensorium \mathbf{P} is equivalent to learning to identify and avoid a particular sub-complex of $\text{Cube}(\mathbf{P})$.

Our second contribution is the proposal of an observation model and associated database allowing for the effective maintenance of a memory structure of the kind originally proposed in [11] and capable of executing basic motion planning under the assumption that the contractible world model is correct.

Our agents are equipped with binary transition sensors (section II), moving through an environment in discretely experienced steps. The database structure we propose is described in section IV-A. It allows the agent to effectively plan its motion in quadratic time in the number of sensors – see IV-B using a contractible model of $X \times X$ (constructed from a sampled record – see IV-C – of approximate implication relations among the sensors – see III-B.2). We prove that each update step of the database is achieved in quadratic time in the number of sensors (see IV-C). The database itself is of quadratic size in the number of sensors – a vast improvement over the generally exponentially sized space of sensory equivalence classes.

In this preliminary work, we are unable to resolve the inherent tension between computational effectiveness (practically efficient database management) and model fidelity (actually realizing the homotopy type in the evolving database). We emphasize that these two goals are, indeed, contradictory, since homotopy type is known both to be arbitrarily complex and to provide lower bounds on the complexity of membership problems (see [12], as well as [13] on the complexity of motion planning). Work in progress addresses the prospect for online introduction of new ‘‘introspective’’ sensors capable of representing ‘‘obstacles’’ (i.e., sensations that will never be experienced) and hence planning motions that respect them. Whether it is possible much less desirable to build computationally effective agents capable of actually reasoning about such topological obstacles is an interesting open question.

II. THE OBSERVATION MODEL

For the sake of simplicity we commit to a discrete time model: the lifetime of an agent is described by a *trajectory* $\varphi : \mathbb{T} \rightarrow X$, where \mathbb{T} is the set of non-negative integers. Any trajectory φ has a *differential* $d\varphi : \mathbb{T} \rightarrow X \times X$:

$$d\varphi(t) := \varphi(t) \times \varphi(t + 1), \quad (1)$$

We shall also restrict attention to binary sensors. By a *binary sensor on a space Z* we mean a $\{0, 1\}$ -valued function on Z . Equivalently, one could identify such a function $p \in 2^Z$ with the subset $p^{-1}(1) \subseteq Z$. In other words, a subset A of

Z is identified with its characteristic function χ_A , where $\chi_A(z) = 1$ iff $z \in A$.

Let us recall some standard notation for $\mathbf{2}^Z$:

- $\langle p | z \rangle := p(z)$, (evaluation)
- $\langle p + q | z \rangle := \langle p | z \rangle + \langle q | z \rangle$, (addition modulo 2)
- $\langle p^c | z \rangle := 1 + \langle p | z \rangle$, (complementation)
- $p \leq q \Leftrightarrow p^{-1}(1) \subseteq q^{-1}(1)$. (ordering)

Note that the addition operation above is addition modulo 2.

The set $\mathbf{2}^Z$ is also standardly referred to as a $|Z|$ -cube (see [14]) when viewed as the vertex set of a graph in which a pair of distinct vertices $p, q \in \mathbf{2}^Z$ are joined by an edge if and only if there is exactly one point $z \in Z$ in which $p(z) + q(z) = 1$. Thus, the combinatorial distance in this graph is precisely the Hamming distance.

The analogy to a cube deepens as we notice that $\mathbf{2}^Z$ may be viewed as the vertex-set $C^0(Z)$ of a cubical complex $C(Z)$ by setting $p, q \in \mathbf{2}^Z$ to belong in the same k -dimensional cube iff p and q are at Hamming distance at most k from each other. The lattice of subsets of the set $U = (p + q)^{-1}(1)$ is then in one-to-one correspondence with the set of vertices of the smallest cubical face of $C(Z)$ containing the vertices p and q .

In what follows, we will refer to both $C(Z)$ and its geometric realization as the Z -cube.

A. Binary transition sensors

A binary transition sensor is a binary sensor on the space $X \times X$. We let \mathbf{P} be a set of tags standing in one-to-one correspondence with the sensors available to our agent. It follows that an agent comes equipped with a *realization map* $r : \mathbf{P} \rightarrow \mathbf{2}^{X \times X}$, deterministically defining the physical capabilities of each individual sensor.

From the point of view of the agent, experiencing the transition $d\varphi(t)$ is tantamount to recording the fact that the following set of sensors was turned on (perhaps after already being on):

$$\text{Inp}^t = \{a \in \mathbf{P} \mid \langle r(a) | d\varphi(t) \rangle = 1\}. \quad (2)$$

It will be of technical convenience to assume henceforth that sensors come in complementary pairs: formally speaking, \mathbf{P} is equipped (once and for all) with a fixpoint-free involution $a \mapsto a^*$ (henceforth referred to as the *complementation operator* in \mathbf{P}), satisfying $r(a^*) = r(a)^c$ for all $a \in \mathbf{P}$. Also, for any subset $U \subset \mathbf{P}$ we shall use the symbol U^* to denote the image of U under the complementation operator.

Note that some sensors do not really record transitions: we say that $r \in \mathbf{2}^{X \times X}$ is a *state sensor*, if

$$\forall x, y \in X \quad \langle r | x \times y \rangle = \langle r | y \times x \rangle. \quad (3)$$

State sensors record properties of individual states, e.g: “the sky is green” and “there is a wall right in front of me”.

B. Actions

In the context of our problem every agent will be equipped with *actions*. For simplicity we assume X is deterministic in the sense that any action α taken by our agent has a well-defined outcome, which we denote by $\alpha x \in X$, for any given input state $x \in X$ (note that this restricts the agent to taking a single action at a time). We conclude that an action defines a binary transition sensor, providing us with a way to introduce actions into our observation model as follows:

- We designate a subset \mathbf{P}_{act} of \mathbf{P} as a subset of *action tags* and require \mathbf{P}_{act} to be closed under $a \mapsto a^*$, that is: \mathbf{P}_{act} is a sub-poc-set of \mathbf{P} ;
- For every $\alpha \in \mathbf{P}_{act}$, exactly one of $r(\alpha), r(\alpha^*)$ is a function on X when viewed as a relation in $X \times X$ (unless $|X| = 2$, but that is not an interesting case);
- At any time $t \in \mathbb{T}$, the decision to apply an action $\alpha \in \mathbf{P}_{act}$ is equivalent to setting

$$d\varphi = \varphi(t) \times (\alpha\varphi(t)). \quad (4)$$

In other words, taking an action results in determining the observation Inp^t which will take place during the transition $d\varphi(t)$. However, since the agent does not necessarily have precise knowledge of its current state, it still has a good chance of being surprised by the results of the chosen action if it happened to anticipate ‘landing’ in some state other than $\varphi(t+1) = \alpha\varphi(t)$. Thus, actions should be thought of as commands issued to the motors by the ‘brain’ and should not be mistaken for their consequences.

III. REPRESENTING A SPACE

We step away from our observation model for a short while. Once again let Z be an arbitrary topological space and suppose $r : \mathbf{P} \rightarrow \mathbf{2}^Z$ is a map satisfying the identity $r(a^*) = r(a)^c$.

A. Constructing a Homotopy Model from Binary Questions

Since the map $a \mapsto a^*$ defined on \mathbf{P} is fixpoint-free, \mathbf{P} has an even number of elements, $|\mathbf{P}| = 2d$. It is customary to refer to a subset $U \subset \mathbf{P}$ as a **-selection*, if, for any $a \in \mathbf{P}$ only one of the elements a, a^* belongs to U ; a *complete *-selection* is a *-selection of cardinality d .

We denote the set of all complete *-selections on \mathbf{P} by $S^0(\mathbf{P})$ to think of it as the set of vertices of a cubical complex $S(\mathbf{P})$, constructed by declaring $U, U' \in S^0(\mathbf{P})$ to belong to the same k -dimensional cube if and only if $|U \setminus U'| \leq k$. Note that this condition is symmetric, as $(U \setminus U')^* = U' \setminus U$. This also implies $S(\mathbf{P})$ is isomorphic to a d -dimensional cube, namely: if $V \in S(\mathbf{P})$ is picked arbitrarily, then it is easy to see that the mapping $U \mapsto \chi_U + \chi_V$ is an isomorphism of cubical complexes from $S(\mathbf{P})$ onto the V -cube $C(V)$.

We now turn to the study of realization maps $r : \mathbf{P} \rightarrow \mathbf{2}^Z$ where Z is an arbitrary topological space. Intuitively, the map r defines d binary questions about Z , and therefore partitions Z into at most 2^d observational equivalence classes [8] (call them r -classes for short). Formally, there is a map

$r^* : Z \rightarrow S(\mathbf{P})$ defined by $r^*(z) = \{a \in P \mid z \in r(a)\}$ and partitioning Z into point-preimages of r^* . Thus, the r -classes embed into $S^0(\mathbf{P})$, or, equivalently into the vertex set of a d -cube. Vertices of $S(\mathbf{P})$ appearing in the image of r^* will be called *consistent *-selections on \mathbf{P}* , or simply *r -consistent vertices*.

Let $\text{Dual}(\mathbf{P}, r)$ denote the sub-graph of $S^1(\mathbf{P})$ induced by the set of r -consistent vertices. Further, let $\text{Cube}(\mathbf{P}, r)$ denote the full sub-complex induced by this set in $S(\mathbf{P})$. In other words, $\text{Cube}(\mathbf{P}, r)$ is the smallest full sub-complex of $S(\mathbf{P})$ containing the r -consistent vertices. Under certain tameness assumptions on r , it is possible to prove the following:

Theorem 3.1 (Guralnik-Koditschek): Let Z be a topological space. Suppose $r : \mathbf{P} \rightarrow 2^Z$ be a realization and let \mathcal{C} denote the collection of cubes in $\text{Cube}(\mathbf{P}, r)$. For each $C \in \mathcal{C}$ let Z_C denote the union of all r -classes corresponding to the vertices of C . Suppose now that every $Z_C, C \in \mathcal{C}$ has an open neighbourhood $N_C \subset Z$ such that:

- 1) each N_C is contractible;
- 2) $\{Z_C\}_{C \in \mathcal{C}}$ and $\{N_C\}_{C \in \mathcal{C}}$ have isomorphic nerves².

Then $\text{Cube}(\mathbf{P}, r)$ is homotopy-equivalent to Z .

Proof: Let \mathcal{C} denote the collection of cubes in $\text{Cube}(\mathbf{P}, r)$. For any subset $A \subseteq \mathbf{P}$, define $\hat{r}(A) = \bigcap_{a \in A} r(a)$. The crucial technical observation is that the vertex set of any cube in \mathcal{C} has the form

$$V(A) = \{U \in S(\mathbf{P}) \mid A \subset U\} \quad (5)$$

for some suitably chosen $A \subset \mathbf{P}$. From here, the proof boils down to applying the nerve lemma twice, noticing that:

- 1) Let \mathcal{U} be the covering of $\text{Cube}(\mathbf{P}, r)$ with closed cubes. Observing that (realizations of) cubes are contractible subsets of $\text{Cube}(\mathbf{P}, r)$, we conclude that (the realization of) $\text{Cube}(\mathbf{P}, r)$ is homotopy equivalent to the realization of the nerve $N(\mathcal{U})$ of \mathcal{U} .
- 2) Now consider the covering \mathcal{V} of Z by sets of the form $\hat{r}(A)$, with A bounding a cube in $\text{Cube}(\mathbf{P}, r)$. Note that \mathcal{V} is in one-to-one correspondence with $\{Z_C\}_{C \in \mathcal{C}}$. The tameness assumption on the Z_C allows us to apply the nerve lemma again to conclude Z is homotopy equivalent to the realization of the nerve $N(\mathcal{V})$.

To finish the proof, it suffices to observe that \mathcal{U} and \mathcal{V} have isomorphic nerves: for $\hat{r}(A), \hat{r}(B) \in \mathcal{V}$ we have that $\hat{r}(A) \cap \hat{r}(B)$ is non-empty if and only if the cubes with 0-skeleta $V(A)$ and $V(B)$ intersect in a common face. ■

Loosely speaking, theorem 3.1 says that if the binary sensorium realized by r is rich enough to fulfill the requirements of the theorem, then $\text{Cube}(\mathbf{P}, r)$ is a precise homotopy model for Z . It is an interesting question, whether there exist verifiable conditions on the sensorium guaranteeing the hypothesis of this theorem. Furthermore, in the context of our application we must note that the map r (as well as the

²These conditions may be replaced by every Z_C itself being contractible, provided some topological tameness conditions on Z and the sets $r(a)$, $a \in \mathbf{P}$

space $Z = X \times X$ where the sensors are being realized) is unknown, largely rendering the above theorem useless unless we answer the following question:

Question 1: What information should an agent be extracting from its sensory readings in order to efficiently maintain increasingly useful approximations of $\text{Cube}(\mathbf{P}, r)$?

Partially answering this question is the job of the next section. Another question needing to be asked in the context of an agent navigating a physical environment is:

Question 2: What are some reasonable sufficient (perhaps geometric) conditions on a sensorium guaranteeing the assumptions of theorem 3.1 in a given context?

B. Reviewing Poc Sets and their Duals

The following notion, due to Martin Roller, is introduced in [15]: a partially ordered set (\mathbf{P}, \leq) with an order-reversing involution $a \mapsto a^*$ is called a *poc set*, if (1) it contains a minimum element denoted 0, and (2) for any element $a \in \mathbf{P}$ the relation $a \leq a^*$ implies $a = 0$ (the 0 element corresponds to the empty set when mapped into 2^Z under a realization in a space Z ; 0^* is realized as the whole of Z ; elements $a \in \mathbf{P}$ other than 0, 0^* are called *proper elements*).

A poc-set \mathbf{P} is said to be *discrete*, if it is discrete as a poset, that is: for every $a, b \in \mathbf{P}$, the set of elements $x \in \mathbf{P}$ satisfying $a < x < b$ is finite. A finite poc-set is, of course, discrete.

Example 1: 2^Z (for any set Z) is a poc set with respect to the inclusion ordering and the complementation operator defined by $p^* = p^c$.

If P, Q are poc sets, then a function $f : P \rightarrow Q$ is said to be a *poc morphism* if $f(0) = 0$, $f(a^*) = f(a)^*$ for all $a \in \mathbf{P}$ and $f(a) \leq f(b)$ holds whenever $a, b \in \mathbf{P}$ satisfy $a \leq b$. We say f is a *poc embedding* when the inequality $f(a) < f(b)$ holds in Q if and only if $a < b$ holds in \mathbf{P} . In particular, a poc embedding is injective.

1) *Coherent Selections:* It is easily verified that a pair of proper elements $a, b \in \mathbf{P}$, $a \notin \{b, b^*\}$, may only satisfy at most one of the following *nesting relations*:

$$a \leq b, a^* \leq b, a \leq b^*, a^* \leq b^*. \quad (6)$$

A set $A \subset \mathbf{P}$ where any pair of elements satisfies a nesting relation is said to be *nested*. If no two elements of A are nested, A is said to be *transverse*.

We say that a pair $a, b \in \mathbf{P}$ satisfying $a \leq b^*$ is *incoherent*. We say that a subset $U \subset \mathbf{P}$ is *coherent* if it contains no incoherent pair.

Returning to our specific application, suppose our agent, endowed with sensorium \mathbf{P} and its complementation operator, has (or thinks it has) prior knowledge of some implication relations among its sensors. Such knowledge corresponds to a poc-set structure on \mathbf{P} (after the trivial tags 0 and 0^* are added to the sensorium, again as a matter of technical convenience). If the agent's record of relations among its sensors is correct (though not necessarily complete), then the realization map $r : \mathbf{P} \rightarrow 2^Z$ is a poc morphism.

Extending the preceding definitions, we say that a subset $U \subset \mathbf{P}$ is r -consistent, if the intersection of $\{r(a)\}_{a \in U}$ is non-empty. Thus, every consistent subset of \mathbf{P} is coherent whenever r is a poc morphism (that is, $r(a) \subseteq r(b)$ whenever $a \leq b$ in \mathbf{P}). Loosely speaking, we would like to view our agent as working to maintain on its record an evolving poc-set structure \mathbf{P} with the aim to keep r as close to a poc morphism as possible while attempting to learn the complete set of implication relations among its sensors.

2) *Another Model:* Let $\text{Dual}(\mathbf{P})$ be the subgraph of $S(\mathbf{P})$ induced by the coherent vertices and let $\text{Cube}(\mathbf{P})$ be the cubical sub-complex of $S(\mathbf{P})$ induced by those vertices. In particular, $\text{Dual}(\mathbf{P}) = \text{Cube}(\mathbf{P})^1$.

When r is a poc morphism, we are guaranteed that $\text{Cube}(\mathbf{P})$ contains $\text{Cube}(\mathbf{P}, r)$, but has, potentially, much fewer vertices than the original cube $S(\mathbf{P})$. This fact may be understood as follows: *keeping a record of implication relations among sensors allows for a significant reduction of the amount of memory required for storing $\text{Cube}(\mathbf{P}, r)$.* To illustrate this idea, consider the two easily verified extremes:

- 1) \mathbf{P} is transverse iff $\text{Cube}(\mathbf{P}) = S(\mathbf{P})$.
- 2) \mathbf{P} is nested iff $\text{Cube}(\mathbf{P})$ is a tree.

We see that $|\text{Cube}(\mathbf{P})^0|$ varies all the way between $d+1$ and 2^d , where we recall that $2d = |\mathbf{P}|$ is the number of sensors in the sensorium.

Thus, seeking a finite graph/cell-complex representation for Z in $\text{Dual}(\mathbf{P})$ we gain a serious advantage over $\text{Cube}(\mathbf{P}, r)$: while in order to form (and maintain) $\text{Cube}(\mathbf{P}, r)$ as a sub-complex of $S(\mathbf{P})$ an agent will need to attempt searching through the exponentially-many vertices of the cube $S(\mathbf{P})$ (some of which may never be reached, as we know), holding *complete* information about $\text{Cube}(\mathbf{P})$ is tantamount to gradually obtaining a correct record of a ‘best fit’ poc-set structure on \mathbf{P} satisfying the requirement that the realization map r be a poc morphism.

We imagine a process where an agent is introduced into the environment at time $t = 0$ carrying a transverse poc-set structure \mathbf{P}^0 (i.e., no relations whatsoever except the trivial ones $0 < a$ and $a < 0^*$ hold for any $a \in \mathbf{P}$) which subsequently evolves (somehow) into a sequence of poc-set structures \mathbf{P}^t , each corresponding to the agent’s conjecture at time t regarding the inclusion relations within the family $\{r(a)\}_{a \in \mathbf{P}}$. At time $t = 0$ the agent ‘thinks’ the world is a cube, but, possibly, after 13 ticks of the clock the agent ‘notices’ (we will make that more precise later) some relations and rejects $\mathbf{P}^0 = \mathbf{P}^1 = \dots = \mathbf{P}^{12}$ in favor of a new structure \mathbf{P}^{13} that now contains a non-trivial relation. The result: a chunk of the original cube has just now been chopped off.

As time passes and the agent collects additional relations (perhaps discarding some on the way upon obtaining sufficient contradicting evidence), the optimist expects this excavation process to recover the complete list of containment relations among the sets $\{r(a) \cap I\}_{a \in \mathbf{P}}$, where I is the image

of the agent’s trajectory in Z . If this trajectory is ‘dense enough’ in Z , one could hope for a correct recovery of the containment relations in the image of the realization map.

What one *cannot* hope for, however, is for this process *alone* to ever recover $\text{Cube}(\mathbf{P}, r)$, unless Z was contractible to begin with. The following result (essentially due to Sageev [16], but proved by Roller [15] in this form) explains the reason why:

Theorem 3.2 (Sageev-Roller duality): For every finite poc-set \mathbf{P} , the complex $\text{Cube}(\mathbf{P})$ is a finite contractible non-positively curved cubical complex. Every finite contractible non-positively curved cubical complex arises in this way.

In other words, $\text{Cube}(\mathbf{P})$ plays the role of a “minimal contractible envelope” for all the possible $\text{Cube}(\mathbf{P}, r)$ satisfying the requirement that r is a poc morphism. At the same time, each of the $\text{Cube}(\mathbf{P}, r)$ could be arbitrarily complex homotopically, raising the question:

Question 3: What are we gaining by appealing to the Sageev-Roller construction?

First and foremost, it turns out we gain effectiveness of representation: maintaining an explicit representation of $\text{Cube}(\mathbf{P}, r)$ as a subset of $S(\mathbf{P})$ (that is, without keeping record of the possible poc set structure on \mathbf{P}) requires exponential space and time (in the number of sensors), while maintaining \mathbf{P} as a poc-set together with a sufficiently meaningful representation of $\text{Cube}(\mathbf{P})$ only takes quadratic space and time, as we prove in the next section.

Second, recalling that the realization map r is unknown to the agent, we strive to work with a world model with minimum possible redundancy (in terms of space). Being the smallest subset of $S(\mathbf{P})$ containing $\text{Cube}(\mathbf{P}, r)$ for all possible realizations r of a given poc set structure on \mathbf{P} – remember this requires r to be a poc-morphism – makes $\text{Cube}(\mathbf{P})$ a natural candidate to work with.

Finally, in stark contrast to other representation techniques known to us, where the model of the ambient space is constructed by agglomeration (“here is my current map, let us add this new chunk to it and glue those two other chunks together”), our approach is based on excavation of material from a contractible space (the complex $\text{Cube}(\mathbf{P})$) containing complete information about *all* possible connectivity relations among observation classes. As far as navigation and path-planning problems are concerned, the advantage of such an approach to representation lies in its potential for proposing unexplored paths between observation classes.

C. Convexity Theory in $\text{Dual}(\mathbf{P})$

Sageev-Roller duality is, in fact, a consequence of a very strong and restrictive convexity theory satisfied by graphs of the form $\text{Dual}(\mathbf{P})$ (here \mathbf{P} is an arbitrary finite poc set. We will now recall some facts from [15].

In any simple graph G , the interval $I(u, v)$ between two vertices u and v is the set of all vertices x in G satisfying $d(u, x) + d(x, v) = d(u, v)$, where d is the combinatorial

metric, defining the distance between two vertices is the minimum number of edges along a path joining them in G .

A graph G is said to be a *median graph*, if $I(u, v) \cap I(u, w) \cap I(v, w)$ contains precisely one vertex – denoted $\text{med}_G(u, v, w)$ – for any three vertices u, v, w in G . Figure 1 provides an example of a median graph.

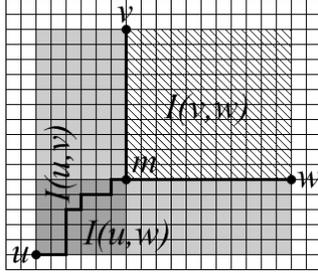


Fig. 1. The integer grid: example of computing a median in a median graph.

Theorem 3.3 ([15]): For any finite poc set \mathbf{P} , the graph $\text{Dual}(\mathbf{P})$ is a median graph. Conversely, every finite median graph is isomorphic to $\text{Dual}(\mathbf{P})$ for some poc set \mathbf{P} .

This theorem is, essentially, a reformulation of theorem 3.2, and it “lives” on the substrate of special properties of median graphs. Recall that a set $C \subset VG$ of vertices of a graph G is said to be *convex*, if $I(u, v) \subseteq C$ holds for any pair of vertices $u, v \in C$. A set of vertices $C \subset G$ is a *half-space*, if both C and $VG \setminus C$ are convex in G .

Median graphs turn out to have many half-spaces, as the following theorem and corollary show:

Theorem 3.4 (“Helly’s Theorem”, [15]): If C_1, \dots, C_m are all convex subspaces of a median graph G such that $C_i \cap C_j \neq \emptyset$ for all $1 \leq i, j \leq m$, then $\bigcap_{i=1}^m C_i \neq \emptyset$.

Corollary 3.5: Every convex subset of a median graph is the intersection of a family of half-spaces.

Lemma 3.6: Suppose C is a convex subset in a finite median graph G . Then for every vertex $v \in G$ there exists a unique vertex $v' \in C$ satisfying $d(v, v') \leq d(v, w)$ for every $w \in C$. This vertex is denoted by $\text{proj}_C v$. Furthermore, one has

$$d(v, \text{proj}_C v) = |\Delta(v, C)|, \quad (7)$$

where $\Delta(v, C)$ is the set of half-spaces containing v and disjoint from C . The vertex $\text{proj}_C v$ is called the closest point projection of v to C .

It turns out that the half-spaces in $\text{Dual}(\mathbf{P})$ are all of the form:

$$V(a) = \{A \in \text{Dual}(\mathbf{P}) \mid a \in A\} \quad (8)$$

Thus a convex subset $C \subset \text{Dual}(\mathbf{P})$ is necessarily of the form

$$V(B) = \bigcap_{b \in B} V(b) \quad (9)$$

for some suitably chosen subset $B \subseteq \mathbf{P}$ (actually, we have used this fact in the proof of theorem 3.1). For example, applying lemma 3.6 in this setting to $C = V(a)$ produces the following formula:

$$d(v, C) = |\{b \in v \mid b < a^*\}| \quad (10)$$

In particular, the projection $\text{proj}_C v$ has the formula

$$\text{proj}_C v = (v \setminus \{b \mid b < a^*\}) \cup \{b^* \mid b < a^*\}. \quad (11)$$

It is this particular property of $\text{Dual}(\mathbf{P})$ that is responsible for the low maintenance cost of our database structure (see IV-B).

D. Less Restrictive Language

It turns out that a poc set structure is a tricky thing to maintain dynamically. In our application, the agent is supposed to maintain an evolving poc-set structure $(\mathbf{P}^t)_{t \in \mathbb{T}}$ in an attempt to recover inclusions of the form $r(a) \subset r(b)$ ($a, b \in \mathbf{P}$) by accumulating sampled observations. Expecting such a recovery process to be precise would constitute naïve wishful thinking. Still, hoping for a statistical “ ϵ -approximation” of the set of such relations may stand a chance for success. This, however, requires a notion of negligibility to be introduced into the language of poc sets. We make the formal arrangements right here, and demonstrate their relevance to our application in the following section.

Definition 1: A partially ordered set (\mathbf{P}, \leq) with a fixpoint-free order-reversing involution and a minimum element (denoted 0) is called a *weak poc set*. An order-preserving and $*$ -equivariant map between weak poc sets is said to be a *weak poc morphism*.

What differentiates a weak poc set from a poc set is the possibility that sets

$$\hat{0} = \{a \in \mathbf{P} \mid a \leq a^*\}, \quad \hat{0}^* = \{a \in \mathbf{P} \mid a^* \leq a\} \quad (12)$$

may contain proper elements of \mathbf{P} .

Definition 2: Let $(\mathbf{P}, \leq, *, 0)$ be a weak poc set. An element of $\hat{0}$ is said to be *negligible*.

Suppose \mathbf{P} is a weak poc set. For each $a \in \mathbf{P}$ set $\hat{a} = \hat{0}$ if $a \in \hat{0}$, $\hat{a} = \hat{0}^*$ if $a \in \hat{0}^*$ and $\hat{a} = a$ otherwise. Denote $\hat{\mathbf{P}} = \{\hat{a} \mid a \in \mathbf{P}\}$ and set $\hat{a} \leq \hat{b}$ iff $a \leq b$ as well as $\hat{a}^* = \hat{a}^*$. The following lemma is an immediate consequence of the definitions. We omit the proof to save space:

Lemma 3.7: Suppose $(\mathbf{P}, \leq, *, 0)$ is a weak poc set. Then $(\hat{\mathbf{P}}, \leq, *, \hat{0})$ is a poc set with respect to the order and complementation defined above and the map $a \mapsto \hat{a}$ is a weak poc morphism. Moreover, if \mathbf{Q} is any poc set, and $f : \mathbf{P} \rightarrow \mathbf{Q}$ is any weak poc morphism, then there exists a unique poc morphism $\hat{f} : \hat{\mathbf{P}} \rightarrow \mathbf{Q}$ satisfying $\hat{f}(\hat{a}) = f(a)$ for all $a \in \mathbf{P}$.

The notions of a complete $*$ -selection, coherent subset and $\text{Cube}(\mathbf{P})$ remain the same – *verbatim!!* – for weak poc sets as they were defined for poc sets. For vertices of $\text{Cube}(\mathbf{P})$, we have:

Lemma 3.8: Let \mathbf{P} be a weak poc set. Suppose $U \subset \mathbf{P}$ is a coherent complete $*$ -selection. Then $\hat{0} \cap U = \emptyset$ and $\hat{0}^* \subset U$.

Proof: By definition of a complete $*$ -selection, it suffices to show $\hat{0} \cap U = \emptyset$. Suppose $a \leq a^*$ and $a \in U$.

For any $b \geq a$, if $b \notin U$ then $b^* \in U$ but then U contains the incoherent pair $\{a, b^*\}$ – a contradiction. Thus, for any coherent complete $*$ -selection we have that $a \in U$ and $b \geq a$ imply $b \in U$.

Back to our specific situation, we have $a \in U$ and $a \leq a^*$. But then $a^* \in U$ causes a contradiction. Hence no element of $\hat{0}$ is contained in U , as desired. ■

Thus every vertex $U \in \text{Cube}(\mathbf{P})$ gives rise to a vertex $\hat{U} = \{\hat{a} \mid a \in U\}$ in $\text{Cube}(\hat{\mathbf{P}})$ and this assignment is bijective. This provides the following easy corollary, whose proof we shall omit:

Corollary 3.9: Suppose \mathbf{P} is a weak poc set. Then the cubical map

$$\begin{cases} \text{Cube}(\mathbf{P}) & \rightarrow & \text{Cube}(\hat{\mathbf{P}}) \\ U & \mapsto & \hat{U} \end{cases} \quad (13)$$

is an isomorphism of cubical complexes.

In other words, a weak poc set and its “strict version” produce identical geometric models. However, a weak poc set structure in memory carries an additional benefit for our agent: it is conceivable that some sensations may seem negligible to the agent for a limited amount of time (due to their initial rarity), but become more meaningful as they become more frequent; other sensations may “fall out of favor”. We want the agent’s memory to have sufficient flexibility to allow it to treat the various sensations with a measure of respect commensurable with their relevance to the agent’s experience.

IV. ALL TOGETHER NOW: THE MEMORY STRUCTURE

It is time for us to provide a bird’s eye view description of a family of memory structures employing our observation model to construct ever-improving approximations of the ground truth order structure on \mathbf{P} , given a sufficiently dense trajectory $\varphi : \mathbb{T} \rightarrow X$: clearly, if the agent does not get to visit large chunks of E or if the agent dedicates unreasonable amounts of time to tiny portions of X , then its observations will become biased and could produce the wrong (weak) poc-set structure.

A. The data structure: Snapshots

The core idea behind our database structure is that in order to maintain a poc-set structure on the tag set \mathbf{P} one needs to maintain the following data structure:

Definition 3: A *snapshot* \mathbf{S} is a graph carrying the following structure:

- Vertices: to each $a \in \mathbf{P}$ we associate a *cell* σ_a , to serve as a vertex of \mathbf{S} . Each cell carries a binary *state variable* $\#a \in \{0, 1\}$.
The set $\#\mathbf{S} = \{a \in \mathbf{P} \mid \#a = 1\}$ is called the *state of the snapshot* \mathbf{S} .
- Edges: exactly one edge marked ab is present for each $a, b \in \mathbf{P}$ satisfying $a \notin \{b, b^*\}$.

Each edge $e = ab$ carries a non-negative integer counter denoted n_e .

Each edge e is either directed or undirected, subject to the requirements:

- 1) $e = 0a$ is always directed from σ_0 to σ_a ;
- 2) $e = ab$ is directed and points to σ_b if and only if b^*a^* is directed and points to a^* ;
- 3) if $e = ab$ is directed, then ab^* is not directed.

The restriction of \mathbf{S} to the subgraph induced by the cells $\sigma_a, \sigma_{a^*}, \sigma_b$ and σ_{b^*} will be denoted by $\mathbf{S}|_{ab}$ and generally referred to as a *square*.

The requirements regarding directions of edges are imposed to partially ensure that a snapshot represents a weak poc-set, as follows. For every snapshot \mathbf{S} , let us denote by $\overrightarrow{\mathbf{S}}$ the directed subgraph (containing all vertices and all directed edges) of the snapshot, discarding the weights. Define $a \leq b$ for $a, b \in \mathbf{P}$ whenever there is a directed vertex path in $\overrightarrow{\mathbf{S}}$ of the form $(\sigma_{a_0}, \dots, \sigma_{a_m})$ with $a_0 = a$ and $a_m = b$. It is clear that the map $a \mapsto a^*$ is order-reversing, and that (\leq) is a partial ordering on \mathbf{P} if and only if $\overrightarrow{\mathbf{S}}$ is acyclic (contains no *directed* cycles). We call (\leq) the *snapshot order*.

Definition 4: We say that a snapshot \mathbf{S} on a tag set \mathbf{P} is *acyclic*, if $\overrightarrow{\mathbf{S}}$ is acyclic. The weak poc-set structure induced on \mathbf{P} by the snapshot order will be denoted by $\text{Poc}(\mathbf{S})$.

Given what we have seen so far, the idea behind the notion of snapshot can be summarized by saying that we want an agent with trajectory φ to maintain a sequence $(\mathbf{S}^t)_{t \in \mathbb{T}}$ of snapshots satisfying the following (somewhat loosely stated) conditions:

- \mathbf{S}^t is acyclic for all t (then $\overrightarrow{\mathbf{S}}^t$ represents the conjecture of the agent regarding the correct contractible world model);
- \mathbf{S}^{t+1} is obtained from \mathbf{S}^t by an updating procedure whose only additional input is Inp^t ;
- $\#\mathbf{S}^t$ is coherent for all t , and represents the conjecture of the agent regarding the current state.

An agent could be initialized by setting \mathbf{S}^0 to be the *trivial snapshot*: the (acyclic) snapshot with $\#a = 0$ and $n_{ab} = 0$ for all $a, b \in \mathbf{P}$ and all edges unoriented, signifying that ‘the agent knows nothing’. Later snapshots (\mathbf{S}^t with $t > 0$) will represent the progressing state of the agent’s ‘knowledge of the world’.

The reason for the first two requirements is explained in the next paragraph, where applications of acyclic snapshots with coherent state are discussed. Note that the first requirement allows snapshots to represent *weak* poc-set structures – a necessary relaxation for inferring approximate inclusions between unknown sets through sampling – but will not allow the snapshot order to deteriorate into a meaningless structure where the order stops being anti-symmetric.

The second requirement relates to the overall goal of our construction: we want the updating operation on the space of relevant snapshots to give rise to a (topological) dynamical system; we expect that provable results on successful learning of an environment by an agent endowed with a memory

of the proposed architecture are inextricably linked with the structure of basins of attraction of this system. One such transition map on a preferred subspace of the space of acyclic snapshots with coherent state is introduced in IV-C.

The last requirement is debatable. A coherent state defines a convex subset $V(\sharp\mathbf{S}^t)$ of $\text{Cube}(\mathbf{P}^t)$ representing the conjectural region where the agent believes its last experienced transition has taken place. This also implies knowledge of the current state. Coherence of the state vector also allows for an easy interpretation of signal propagation in the snapshot in terms of the world model (as we demonstrate in the next paragraph). Yet these are but technical excuses to back a significant restriction of scope, and we intend to address this issue in forthcoming work.

B. Computing with Snapshots

Snapshots furnish us with a very efficient computational tool which can be used for planning purposes.

Definition 5 (Propagation): Let \mathbf{S} be an acyclic snapshot and let $a \in \mathbf{P}$. A new snapshot $a \cdot \mathbf{S}$ is obtained from \mathbf{S} as follows:

- Leave all edge markings unchanged;
- For all $a, b \in \mathbf{P}$, if $a \leq b$ in $\text{Poc}(\mathbf{S})$, then set $\sharp b := 1$;
- For all $a, b \in \mathbf{P}$, if $a \leq b^*$ in $\text{Poc}(\mathbf{S})$, then set $\sharp b := 0$.

Propagation can be implemented using the graph $\vec{\mathbf{S}}$: switch cell σ_a on into an excited state while sending a signal along edges emanating from it to turn on every cell dominated by σ_a in $\vec{\mathbf{S}}$. Additionally, since sensors are paired with their complements, any cell σ_b being switched on must immediately result in its twin σ_{b^*} turning off. Clearly, propagation is achieved in a time at most linear in the diameter of $\vec{\mathbf{S}}$ when physically implemented in this way.

We now consider some provable applications of the propagation computational mechanism. Suppose our agent froze in time, having a snapshot $\mathbf{S} = \mathbf{S}^t$, and needs to decide which action to take in order to solve a problem. Of course, the agent has no choice but to assume his contractible model of the world is correct – that is, that the realization map r is a weak poc-morphism of $\text{Poc}(\mathbf{S})$ into $\mathbf{2}^{X \times X}$.

Let us explore a couple of planning problems. The simplest planning problem in our context sounds like this: “given my current state and a sensation $a \in \mathbf{P}$, what should I do to reach a position where I can sense a ?”. The answer is given by the following lemma:

Lemma 4.1: Suppose \mathbf{S} is an acyclic snapshot with $\sharp\mathbf{S} \in \text{Dual}(\mathbf{S})$, and let $a \in \mathbf{P}$. Then $\sharp a \cdot \mathbf{S}$ is the image of the closest point projection of the vertex $\sharp\mathbf{S} \in \text{Dual}(\mathbf{S})$ to the half-space $V(a) \subset \text{Dual}(\mathbf{S})$.

Proof: This is an immediate application of equation 11 and the definition of propagation. ■

The lemma tells us that the result of propagating a through $\vec{\mathbf{S}}$ is combinatorially the closest vertex of $\text{Dual}(\mathbf{S})$, the contractible world model. By (10), the total number of cells turned on by the propagation process is greater than or equal to the distance of $\sharp\mathbf{S}$ to the half-space $V(a)$. Thus, choosing

the first action sensor to have switched on during the propagation procedure and updating \mathbf{S} with the consequences of that action is guaranteed to reduce the combinatorial distance of the current state to the half-space $V(a)$. At the same time, the computational complexity of finding the appropriate action is no more than linear in the number of sensors. These guarantees hold so long as $r : \text{Poc}(\mathbf{S}) \rightarrow \mathbf{2}^{X \times X}$ remains an order-preserving map, and the projection path from $\sharp\mathbf{S}$ to $V(a)$ avoids any r -inconsistent vertices.

If no such sensor was switched on, the agent will conclude there is no action it can take in order to start feeling a , and will need to set out on a different quest (interestingly, in the case that $r : \text{Poc}(\mathbf{S}) \rightarrow \mathbf{2}^{X \times X}$ is a poc-embedding – meaning that the contractible model is as precise as possible and has little chance of changing in the future – this situation truly implies there is nothing the agent can do at the moment in order to leave the half-space $V(a^*)$; it might as well just sit and wait for a miracle).

Finally, reaching a prescribed destination is equivalent to planning to reach a vertex $U \in \text{Dual}(\mathbf{S})$. If the required destination is incoherent, the agent refuses the task (as being an impossible task). If U is coherent, the following lemma proves that writing out $U = \{a_1, \dots, a_d\}$ and applying the preceding algorithm successively for any ordering of the a_i produces the required result.

Lemma 4.2: Suppose \mathbf{S} is an acyclic snapshot with coherent state, and let $U = \{a_1, \dots, a_m\}$ be a coherent subset of $\text{Poc}(\mathbf{S})$. Then $U \cdot \mathbf{S} := a_m \cdots (a_1 \cdot \mathbf{S})$ is a well-defined acyclic snapshot with coherent state.

Proof: It is easy to verify that, for any coherent pair $\{a, b\} \subseteq \mathbf{P}$ one has

$$\text{proj}_{V(a)} \circ \text{proj}_{V(b)} = \text{proj}_{V(a,b)}. \quad (14)$$

This implies that $\text{proj}_{V(a)}$ and $\text{proj}_{V(b)}$ commute, and the assertion of the lemma follows. ■

We observe that this procedure has at most quadratic complexity in the number of sensors (being the composition of a linear number of linear complexity algorithms).

C. Dynamics for Snapshots

Given a sensorium $(\mathbf{P}, *)$, we would like to restrict attention to a special collection of snapshots:

Definition 6: Let $\epsilon > 0$ and let \mathbf{P} be a finite set endowed with a fixpoint-free involution $a \mapsto a^*$. A *simple ϵ -Hebbian snapshot* on \mathbf{P} is a snapshot \mathbf{S} satisfying the following requirements:

- n_{ab} is a non-negative integer for every edge ab of \mathbf{S} ;
- There exists $t \in \mathbb{T}$ such that every edge ab in \mathbf{S} satisfies

$$n_{ab} + n_{ab^*} + n_{a^*b} + n_{a^*b^*} = t \quad (15)$$

- For every edge ab in \mathbf{S} , the edges of the square $\mathbf{S}|_{ab}$ are oriented according to figure 2 if and only if condition (\dagger) holds:

$$(\dagger) \quad n_{ab} < t\epsilon, \quad n_{ab^*}, n_{a^*b}, n_{a^*b^*} > 2t\epsilon$$

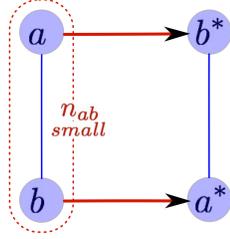


Fig. 2. determining edge orientations in a snapshot by restricting attention to $\mathbf{S}|_{ab}$. Note the symmetry of condition (\dagger) with respect to substituting a^* for a and/or b^* for b .

By a *simple Hebbian snapshot* we mean a simple ϵ -Hebbian snapshot for some suitable $\epsilon > 0$.

The reference to Hebb [17] in this definition is motivated by the way an ϵ -Hebbian snapshot \mathbf{S} decides on an implication relation $a \leq b^*$ (for $a, b \in \mathbf{P}$): we think of t as the ‘age’ of the snapshot, so that condition (\dagger) amounts to a statement about the observation “ a and b ” being decisively more rare (as quantified by the value ϵ) than any of the other three combined observations labeling the edges of the square $\mathbf{S}|_{ab}$. Thus, the (equivalent) statements “ a implies b^* ” and “ b implies a^* ” are put on record provided the corresponding synchronous observations have been made sufficiently many times over the agent’s lifetime.

Definition 7: Let $\epsilon > 0$ and suppose \mathbf{S} is a simple ϵ -Hebbian snapshot on \mathbf{P} . If $A \subset \mathbf{P}$ is a complete $*$ -selection, we define a new snapshot denoted $A * \mathbf{S}$ as follows:

- For every pair of proper elements $a, b \in A$, add 1 to the counter n_{ab} and recalculate $\vec{\mathbf{S}}$ according to condition (\dagger) to obtain the directed structure of $A * \mathbf{S}$;
- If A is coherent with respect to the new directed structure, update $\# \mathbf{S}$ to become $\# A * \mathbf{S} = A$;
- Otherwise, keep the previous state.

Observe that $A * \mathbf{S}$ can be obtained from \mathbf{S} in a time that is quadratic in $|\mathbf{P}|$ (the number of times we need to access a data element of the structure being commensurable with the number of edges in the complete graph on \mathbf{P}). This updating procedure leaves much to be desired. For the time being though, we will be satisfied with the limited result below, demonstrating the feasibility of our approach:

Proposition 4.3: Let $0 < \epsilon < 1$ and \mathbf{S}^0 be the *zero snapshot* on \mathbf{P} . Then for any sequence $(A^t)_{t=0}^\infty$ of complete $*$ -selections on \mathbf{P} , the rule $\mathbf{S}^{t+1} = A^t * \mathbf{S}^t$ produces a sequence $(\mathbf{S}^t)_{t=1}^\infty$ of acyclic simple ϵ -Hebbian snapshots.

Remark 1: throughout this proof, the time parameter will appear as a superscript in all variables, so that $\# a^t$ and n_{ab}^t will denote the state of cell a and the weight on the edge ab , respectively, given by the snapshot \mathbf{S}^t .

Proof: For any $B \subset \mathbf{P}$ and any $t \geq 1$ let n_B^t denote the number of indices $0 \leq j \leq t-1$ for which $B \subset A^j$. Note that $n_{\{a,b\}}^t = n_{ab}^t$, by construction. In particular, this

guarantees the identity

$$n_{ab}^t + n_{a^*b}^t + n_{ab^*}^t + n_{a^*b^*}^t = t \quad (16)$$

for every edge ab ($a, b \in \mathbf{P}$) and every $t \geq 1$, as desired. Note that no directed edges appear in \mathbf{S}^t before time $t = 4$, so that the first few snapshots ($\mathbf{S}^{0 \dots 3}$) are coherent regardless of the choice of input sets ($A^{0 \dots 4}$).

We now turn to proving $\vec{\mathbf{S}}^t$ is acyclic for all $t \geq 1$. It is straightforward to verify the following identity, holding for any $t \geq 1$ and any $a, b, c \in \mathbf{P}$ forming a $*$ -selection:

$$n_{ab^*}^t + n_{bc^*}^t + n_{ca^*}^t + n_{abc}^t + n_{a^*b^*c^*}^t = t \quad (17)$$

Let us suppress the superscript t until the moment it is required. Substituting x^* for each $x \in \{a, b, c\}$ we obtain

$$n_{a^*b} + n_{b^*c} + n_{c^*a} + n_{a^*b^*c^*} + n_{abc} = t \quad (18)$$

Subtracting (18) from (17) and denoting $\Delta(a, b) = n_{a^*b} - n_{ab^*}$ we obtain:

$$\Delta(a, b) + \Delta(b, c) + \Delta(c, a) = 0 \quad (19)$$

Observing $\Delta(a, b) = -\Delta(b, a)$ we finally arrive at

$$\Delta(c, a) = \Delta(c, b) + \Delta(b, a) \quad (20)$$

Moreover, using the identities $\Delta(a, a) = 0$ and $\Delta(a, a^*) = n_{a^*a} - n_{aa^*}$, it is straightforward to verify that this last identity holds even if the restriction that $\{a, b, c\}$ be a $*$ -selection is dropped. Finally, we also observe $\Delta(ab) > \epsilon t$ whenever $ab \in \vec{\mathbf{S}}^t$: by (\dagger^t) :

$$\Delta(ab) = n_{a^*b} - n_{ab^*} > 2\epsilon t - \epsilon t, \quad (21)$$

as $n_{ab^*}^t < \epsilon t$ while $n_{a^*b}^t > 2\epsilon t$.

Thus, if $(\sigma_{a_0}, \dots, \sigma_{a_m})$ is a directed vertex path in $\vec{\mathbf{S}}^t$, then

$$\Delta a_0 a_m = \Delta(a_0 a_1) + \dots + \Delta(a_{m-1} a_m) > m\epsilon t \quad (22)$$

Now, if $a_m = a_0$ we immediately arrive at a contradiction, as $\Delta(a_0, a_0) = 0$. We conclude that \mathbf{S} is an acyclic snapshot, as desired. \blacksquare

V. CONCLUSION

The result of the last section provides us with a soft guarantee of approximate (up to the arbitrary prescribed tolerance $0 < \epsilon < 1$) learning of the order structure in the sensorium with sufficient exposure to the environment, given observations are repeated at more or less regular intervals, that is – if the trajectory of the agent is uniformly distributed in the sense that

$$\lim_{t \rightarrow \infty} \frac{n_B^t}{t} = \Pr \left[\bigcap_{b \in B} r(b) \right] \quad (23)$$

holds for every $B \subset \mathbf{P}$. Due to the very nature of the problem of trying to infer an inclusion among (possibly infinite) sets from a finite number of samples, the notion of a poc set structure becomes too restrictive to fit our needs: while a sequence of ϵ -approximate inclusions cannot

end up producing a directed cycle (proposition 4.3), it is easily conceivable that the sets at the extremities of such a sequence could be complementary. However, by extending the Sageev-Roller construction of $\text{Cube}(\mathbf{P})$ to weak poc sets, we have shown that giving up the ban on negligible elements allows us to maintain a data structure capable of learning approximate inclusions without diminishing the descriptive power of the dual complexes as “world models” (see, e.g. corollary 3.9).

We would like to draw the reader’s attention to the almost complete separation of the learning process (updating of $\vec{\mathbf{S}}^t$) that we propose from the process of tracking the current state (updating of $\#S^t$) by the agent. This seems to be too simplistic, and we believe that deeper, more insightful update procedures of the state must exist, perhaps requiring a more sophisticated observation/action model. In particular, the necessity to maintain coherence of state at all times remains debatable: since one may *at best* hope to recover the correct poc set structure asymptotically, maintaining coherence every step of the way seems a redundant luxury costing the system extra computational effort. This thought motivates further study of the computational applications of propagation in the context of incoherent subsets of \mathbf{P} . Getting the model rid of the “stop the clock, let me think” assumption which currently allows our planning mechanism to run in between physical transitions is also a priority for future research.

Finally, the fact that our algorithm requires the agent to reject nonsensical observations without any further processing calls for a good dose of healthy skepticism regarding “asymptotical coherence”, yet the fact that, under our algorithm, sufficient repetition of seemingly incoherent observations forces the agent to acknowledge their validity seems to provide grounds for guardedly stating a conjecture:

Conjecture 5.1: With a given probability measure on $X \times X$, sufficient topological tameness conditions on the realization map r , $A^t = \text{Inp}^t$ for all $t \in \mathbb{T}$ and a fixed trajectory $\varphi : \mathbb{T} \rightarrow X$, if φ is uniformly distributed in the sense of (23), then $\#S^t$ remains coherent for $t \gg 0$.

If proven true, the meaning of this conjecture is that, given enough time and sufficient exposure, an agent employing our learning algorithm will be able to apply the motion planning algorithm from section IV-B meaningfully (with its success depending only on the topology of $\text{Cube}(\mathbf{P}^t, r)$).

ACKNOWLEDGEMENTS

This work was funded in part by the Air Force Office of Science Research under the MURI FA9550-10-1-0567 and in part by the National Science Foundation under CDI-II-1028237. The first author is grateful to Shai Revzen for some exceptionally interesting and invigorating discussions.

REFERENCES

[1] G. Buzsáki, *Rhythms of the Brain*. Oxford University Press, 2006.
 [2] C. R. Gallistel, “Animal cognition: The representation of space, time and number,” *Annual review of psychology*, vol. 40, no. 1, p. 155189, 1989.

[3] B. Kuipers, “An intellectual history of the spatial semantic hierarchy,” *Robotics and cognitive approaches to spatial mapping*, vol. 38, p. 243264, 2008.
 [4] H. P. Moravec, “The stanford cart and the cmu rover,” *Proceedings of the IEEE*, vol. 71, no. 7, p. 872884, 1983.
 [5] S. Thrun, W. Burgard, and D. Fox, “Probabilistic robotics,” 2006.
 [6] J. Leonard and H. Durrant-Whyte, “Mobile robot localization by tracking geometric beacons,” *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, p. 376382, Jun 1991.
 [7] S. Polyn, K. Norman, and M. Kahana, “A context maintenance and retrieval model of organizational processes in free recall,” *Psychological Review*, vol. 116, no. 1, p. 129156, 2009.
 [8] B. Donald and J. Jennings, *Sensor interpretation and task-directed planning using perceptual equivalence classes*, 1991, p. 190197. [Online]. Available: <http://ieeexplore.ieee.org/xpls/abs.all.jsp?arnumber=131977>
 [9] E. Tulving, “Episodic memory: From mind to brain,” *Annual Review of Psychology*, vol. 53, no. 1, p. 125, 2002.
 [10] J. T. Schwartz and M. Sharir, “On the piano movers problem: Ii. general techniques for computing topological properties of real algebraic manifolds,” *Advances in applied Mathematics*, vol. 4, no. 1, p. 298351, 1983.
 [11] D. P. Guralnik, “A formal approach to modeling the memory of a living organism.” [Online]. Available: [arXiv:1003.3821v1\[cs.AI\]](https://arxiv.org/abs/1003.3821v1)
 [12] A. C.-C. Yao, “Decision tree complexity and Betti numbers,” *J. Comput. System Sci.*, vol. 55, no. 1, part 1, pp. 36–43, 1997, 26th Annual ACM Symposium on the Theory of Computing (STOC '94) (Montreal, PQ, 1994). [Online]. Available: <http://dx.doi.org/10.1006/jcss.1997.1495>
 [13] M. Farber, “Topological complexity of motion planning,” *Discrete Comput. Geom.*, vol. 29, no. 2, pp. 211–221, 2003. [Online]. Available: <http://dx.doi.org/10.1007/s00454-002-0760-9>
 [14] D. B. West, *Introduction to Graph Theory*. Prentice-Hall, 2001.
 [15] M. Roller, “Poc sets, median algebras and group actions,” *University of Southampton, Faculty of Math. stud., preprint series*, 1998.
 [16] M. Sageev, “Ends of groups pairs and non-positively curved cube complexes,” *Proc. London Math. Soc.*, vol. 3, no. 71, pp. 586–617, 1995.
 [17] D. Hebb, *The Organization of Behavior: A Neuropsychological Theory*. Taylor & Francis, 2002. [Online]. Available: <http://books.google.com/books?id=gUtwMochA18C>