

Universal Memory Architectures for Autonomous Machines

Dan P. Guralnik[†], Daniel E. Koditschek[‡]

Abstract—We propose a self-organizing memory architecture for perceptual experience capable of supporting autonomous learning and goal-directed problem solving in the absence of any prior information about the agent’s environment. The architecture is simple enough to ensure (1) a quadratic bound (in the number of available sensors) on space requirements, and (2) a quadratic bound on the time-complexity of the update-execute cycle. At the same time, it is sufficiently complex to provide the agent with an internal representation which is (3) minimal among all representations of its class which account for every sensory equivalence class subject to the agent’s belief state; (4) capable, in principle, of recovering the homotopy type of the system’s state space; (5) learnable with arbitrary precision through a random application of the available actions.

The provable properties of an effectively trained memory structure exploit a duality between *weak poc sets* — a symbolic (discrete) representation of subset nesting relations — and *non-positively curved cubical complexes*, whose rich convexity theory underlies the planning cycle of the proposed architecture.

I. INTRODUCTION

A. Motivation

A major obstacle to autonomous systems synthesis is the absence of a capacious but efficient memory architecture. In humans, memory influences behaviour over a wide range of time scales, leading to the emergence of what seems to be a functional hierarchy of sub-systems [1]: from non-declarative vs. declarative through the split of declarative memory into semantic and episodic [2]; and on to theories of attention and recall [3]. This variety of scales is mirrored in the collection of problems addressed by the synthetic sciences: from learning dependable actions/motion primitives [4], [5]; through learning objects and their affordances [6], [7] to demonstration-driven task execution [8], [9]; through exploring and mapping an unknown environment [10], [11], [12], [13] and motion planning [14], [15], [16]; and on to general problem solving [17] using artificial general intelligence (AGI) architectures [18], [19], [20].

One idea stands out as common to all these approaches, beginning with the formal notion of a problem space introduced by Newell and Simon [17], [21]: the purpose of a memory architecture is to learn the transition structure of the state space \mathbf{X} of the system comprised of the agent and its environment

while processing the history of observations into a format facilitating improved future control.

It is often argued (e.g. [22], [23], [24]) that memory architectures for general agents should enjoy a high degree of domain- and task-independence. In general, however, clear definitions of notions such as ‘domain’ and ‘task’ are not readily forthcoming across the vast breadth of literatures discussing memory, agents and autonomy. Notions of ‘universal learners’ have been proposed [25] based on optimizing gain in estimators of predictive entropy, however there is evidence to suggest that the resulting level of generality may be insufficient for some tasks [26].

Absent broadly recognized formal foundations, we advance an architecture provably satisfying intuitive universality properties, including, most centrally: (1) interactions with the environment are encoded in the most generic, yet minimal, manner possible, while requiring no prior semantic information; and (2) learning obtains from direct binary sensory input, automatically developing appropriate contextual links between sensations of arbitrary modality. A key outcome is that the architecture encodes its observation history in a model space that supports the agent’s problem solving as a form of reactive motion planning whereby atomic computations provably correspond to nearest point projection in the reachable set.

B. Contribution

We consider a generic *discrete binary agent* (DBA): a machine sensing and interacting with its environment in discrete time, equipped with a finite collection Σ of Boolean-valued sensors, some of which serve as triggers for actions/behaviors (switched on and off at will).

Given an instance of a DBA interacting with an environment \mathbf{E} , it is natural to view the set Ξ of sensory equivalence classes of the associated transition system \mathbf{X} as a subset of the power set $\{0, 1\}^{\Sigma}$. It is generally accepted [27], [28] that a memory architecture must be capable of supporting internal representations rich enough to account for the diversity [29] of the transition system \mathbf{X} : Exact problem solving, when construed as abstract motion planning, requires an internal representation capable, eventually, of accounting for all the classes in Ξ and the transitions between them. Unfortunately, as expressed forcefully in [29] and as we review at length below, the task of obtaining an exact description of Ξ becomes intractable in the absence of strong simplifying assumptions about \mathbf{X} , as the number of sensors grows.

To circumvent this obstacle, rather than imposing any specific structure on \mathbf{X} , we propose to relax the requirement for precise reconstruction by introducing an approximation whose

[†]Post doctoral fellow, Electrical & Systems Engineering, School of Engineering and Applied Sciences, University of Pennsylvania, , 200 S 33rd Str., Moore Bldg. #203, Philadelphia, PA 19104, USA, guraldan at seas.upenn.edu

[‡]Alfred Fittler Moore Professor, Electrical & Systems Engineering, School of Engineering and Applied Sciences, University of Pennsylvania, 200 S 33rd Str., Moore Bldg. #203, Philadelphia, PA 19104, USA, kod at seas.upenn.edu

discrepancy from Ξ we characterize exactly and show to be the smallest possible in its (computationally effective) class of objects.

The new memory and control architecture we propose here consists of two layers:

- A data structure \mathbf{S} – called a *snapshot* – keeping track of the current state and summarizing observations in terms of collection of real-valued registers, of size quadratic in the number of sensors, summarizing the history of observations made by the agents.
- A reactive planner, built on a *weak poc set structure* \mathbf{P} ([30], [31] and defn. A.1) constituting a record of pairwise implications among the atomic sensations as observed by the agent; \mathbf{P} is computed from \mathbf{S} in each control cycle.

A crucial property of our architecture is that \mathbf{P} and \mathbf{M} are formally reconstructible from each other. The model space \mathbf{M} takes the form of a CAT(0) cubical complex, or *cubing*¹, whose 0-skeleton is contained in $\{0, 1\}^\Sigma$. As the snapshot \mathbf{S} is updated by incoming observations, the space \mathbf{M} , as encoded by \mathbf{P} , is transformed along with it. We can state our main contributions – albeit, necessarily, informally at this point – in terms of provable properties of the architecture and its model spaces:

- (i) **Universality of Representation.** \mathbf{M} is the minimal model guaranteed to represent all the sensory equivalence classes of *any* sensorium Σ satisfying the record \mathbf{P} (see A-E3). In particular, in the absence of information not already encoded in \mathbf{P} , it is impossible to distinguish the 0-skeleton of \mathbf{M} from the set of sensory equivalence classes, Ξ .
- (ii) **Topological Approximation.** As a topological space, \mathbf{M} is always contractible². Provided a sufficiently rich sensorium, the sub-complex $\mathbf{M}^\times \subset \mathbf{M}$ of faces all of whose vertices lie in Ξ inherits from \mathbf{M} the topology³ of the observed space \mathbf{X} (see appendix A-E3-A-E4).
- (iii) **Low-complexity, effective learning.** The proposed architecture requires quadratic space (in the number of sensors) for storage, and no more than quadratic time for updating. Furthermore, an agent picking actions at random learns an approximation of the resulting walk’s limiting distribution on \mathbf{X} (see II-E2).
- (iv) **Efficiency of Planning.** Planning the next action given a target sensation takes quadratic time in the number of sensors, while eliminating the need for searching in the model space. With sufficient parallel processing power, this bound may be reduced to a constant multiple of the height — the maximum length of a chain of implications of \mathbf{P} (see III-B).

To the best of our knowledge, this combination of provable properties has not previously appeared in the literature.

¹For a good introduction CAT(0) cubical complexes, see [32]. For a tutorial on cell complexes see [33], chapter 0 and appendix.

²The formal notion of being ‘hole-free’ — see [33], chapter 0.

³Up to homotopy equivalence — see definition in [33], chapter 0.

C. Overview and Related Literature

To establish the novelty of our contribution we now briefly review the copious literature bearing on these topics as arising from three distinct traditions: robotics; connectionist computation; and artificial general intelligence. After presenting our technical ideas we will explore at the end of the paper in a more discursive form their relation to and implications for the broader field.

1) *Relation to Mapping and Navigation:* Formulating navigation and mapping problems in terms of a point agent moving through a homotopically trivial ambient space while avoiding a collection of geometrically defined obstacle regions representing forbidden states is fundamental to motion planning [14], [15] and mapping [12], [34]. The ubiquity of obstacles in these settings introduces topological considerations whose primacy is well established in the algorithmic literature [34], [35], [36], [37], [38], [39], [40], governing the complexity of not only motion planning [41] but even set membership [42].

Our strategy is to reduce the general problem of memory storage and its use for motion planning in the underlying transition structure of a problem space (as sensed by a DBA) to the geometric problem of motion planning in the agent’s model space \mathbf{M} (playing the role traditionally assigned to Euclidean space). Generalizing the Euclidean setting, \mathbf{M} has a very strong convexity theory [31], [43] enabling low-cost greedy navigation.

The topological point of view has been shown to be well warranted in the discrete setting as well. As was demonstrated by Pratt [44], oriented topological structures (cubical complexes, in fact) may be used to encode the causal relations among actions and states in symbolic transition systems. Approaches generalizing Pratt’s have since been used to formulate very general models of reconfigurability and self-assembly [45], [46].

2) *Mechanisms for Learning and Planning:* Snapshots use an evolving estimate of pairwise intersections of sensor footprints to form a record of implications among the atomic sensations of the DBA. The necessity of such a record for planning goes back (at least) to [47], yet ideas about applying it as a way to encode context are fairly recent and specialized [4], [23], [48]. Our internal representation takes the additional step of applying this principle to *all* the sensations available to a DBA, including the control signals it uses to interact with its environment.

The resulting learning and control mechanisms may be realized in a highly simplified and idealized, yet highly plastic, network of neuron-like cells simulating the structure of \mathbf{P} [30]. This analogy with neural networks is not a coincidence: estimating *arbitrary* intersections from near-synchronous activation of sensors in a planar sensor field has been explored as a means for topological [39] as well as metric mapping by competitive attractor networks (RatSLAM [13], [24], [48]), as the study of the structure of stability properties vis-a-vis topology and plasticity in more general networks is just taking off [49], [50].

3) *Model Spaces:* The necessity to maintain high-dimensional representations of the state space \mathbf{X} poses a major challenge for current approaches to learning [51], [52],

[22] and general problem solving [53], [54]. The method closest to ours in its formalism seems to be that of [29] — and even lends itself to learning by a connectionist network [55] — but still requires an exponentially large representation for planning purposes. By contrast, in our case, the ability to translate action planning in \mathbf{M} into what is essentially a flow problem in a network constructed from the underlying sensorium obviates the need for maintaining \mathbf{M} in memory, allowing us to evade the curse of dimensionality. Nevertheless, we are still guaranteed a model space that is sufficiently rich to account for *all* sensory equivalence classes perceivable by the DBA [27].

The computational advantages of our approach come at a cost that is largely driven by topology, as expressed in (ii): \mathbf{M} necessarily has trivial topology⁴ [56], [31], [57], and our own result [30] establishes formal conditions on Σ under which the complex \mathbf{M}^\times reproduces the "topological shape"⁵ of \mathbf{X} (as discussed above), which may not be topologically trivial. The basic algorithm driving planning in our agents, however, achieves its efficiency by disregarding this mismatch. The introduction of auxiliary intrinsic motivation mechanisms [58], [59] as a means of steering the agent away from obstacle states in \mathbf{M} and towards desirable behaviours (not necessarily states!) seems to be a possible way out of this predicament, as well as towards a solution of the problem of closing the control loop. At this early stage, as a feasibility study for the overall approach, we only consider very simple excitation mechanisms causing the agent to choose actions with the desire to maximize *immediate* excitation gain, to the extent that may be sensed by Σ (and otherwise to choose random actions).

D. Organization of the Paper

Having already given proofs of the formal results underlying (i) and (ii) in our previous paper [30], we defer the technical discussion of poc sets to appendix A. This appendix is intended as an introductory overview of the theory of weak poc sets and the geometry of their dual spaces — our agents' model spaces — as well as a repository of proofs of technical results we require but could not find elsewhere in the literature.

Section II discusses (iii). We formally state the observation model for DBAs, describe snapshots and their learning mechanisms, and present our early numerical work illustrating the practical implications of the claims regarding learning.

Section III is dedicated to item (iv) in the list of contributions. Actions are introduced to the observation model, and control algorithms are defined and validated.

Finally, following an extended discussion of our results in relation to existing literature in section IV and the aforementioned appendix dealing with poc sets, a second appendix presents the proofs of technical results about snapshots.

II. SNAPSHOTS: FROM OBSERVATION SEQUENCES TO A MEMORY STRUCTURE

We begin with a formal statement of what we mean by a DBA and its observation model. We then proceed to construct

TABLE I
TABLE OF MATHEMATICAL SYMBOLS

	Topic/Notation	Ref.
DBA Model (general)		
\mathbf{E}	Environment (with points p, q, \dots)	Sec.II-A1
\mathbf{X}	State space of the experiment (with points x, y, \dots)	Sec.II-A1
pos	The position map $\mathbf{X} \rightarrow \mathbf{E}$	Sec.II-A1
\mathbb{T}	Time, the set of integers	Sec.II-A2
$ _t$	Reads as: "at time t "	Eqn.(1)
DBA model (sensing)		
Σ	Sensorium (elements are a, b, c, \dots), with involution $a \mapsto a^*$	Eqn.(6)
ρ	Realization map of the sensorium Σ	Eqn.(7)
$\langle a : x \rangle$	Evaluation, e.g. of $a \in \Sigma$ on $x \in \mathbf{X}$	Eqns.(4-5)
DBA computational model (at time t)		
$\mathbf{S} _t$	Agent's snapshot	Sec.III-B1
$\Gamma _t$	The derived poc graph, $\text{Dir}(\mathbf{S} _t)$	Sec.III-B1
$\mathbf{P} _t$	Derived (weak) poc set structure on Σ , $\text{Poc}(\mathbf{S} _t)$	Sec.II-B1
$\mathbf{M} _t$	The model space $\text{Cube}(\mathbf{P} _t)$	Sec.II-B3
$\mathbf{M}^\times _t$	The punctured model space $\text{Cube}^\times(\mathbf{P} _t)$	Def.III.8
$O _t$	Raw observation	Sec.II-B2
$S _t$	Recorded observation	Sec.II-B4
$A _t$	Decision (action) following the observation	
Contents/parameters of a snapshot \mathbf{S}		
\mathbf{K}_Σ	The complete graph on Σ with all aa^* edges removed	Def.II.2
$\#\mathbf{S}$	State of the snapshot	Def.II.3(a)
w_{ab}	Weight on the edge ab	Def.II.3(b)
τ_{ab}	Learning threshold for the pair $a, b \in \Sigma$	Def.II.3(c)
$\omega(ab)$	Orientation cocycle of \mathbf{S}	Prop.II.7
$\mu(ab)$	Dissimilarity measure of \mathbf{S}	App.B-C
Objects derived from a snapshot \mathbf{S}		
$\text{Dir}(\mathbf{S})$	Derived poc graph	Prop.II.8
$\text{Poc}(\mathbf{S})$	Derived weak poc set structure	Def.II.10
Weak poc sets and their duals		
P, Q, \dots	Poc sets (with and without indices)	Def.A.1
P°	The set dual of P , the 0-skeleton of $\text{Cube}(P)$	Def.A.8(b)
$\text{Dual}(P)$	Dual graph of P , the 1-skeleton of $\text{Cube}(P)$	Def.A.8(c)
$\text{Cube}(P)$	Dual cubing of the poc set P	Def.A.8(a)
$\text{Cube}^\times(P)$	The punctured dual $\text{Cube}(P, \rho)$ with respect to a realization ρ	Def.A.28
f°	The dual map $f^\circ : Q^\circ \rightarrow P^\circ$ of a poc morphism $f : P \rightarrow Q$	Defs.A.2,A.24

⁴Again, in the sense of \mathbf{M} being contractible.

⁵In the sense of *homotopy type* — [33], chapter 0.

snapshots, their updating mechanisms and the derived weak poc set structures, and conclude the section with results on the learning capabilities of snapshots. Table I reviews notation that will persist throughout the paper.

A. Observation Model for DBAs

1) *Environment and State*: We place an agent in an environment \mathbf{E} . The state space of the system will be denoted by \mathbf{X} , where we assume there is a map $\text{pos} : \mathbf{X} \rightarrow \mathbf{E}$ producing the location $\text{pos}(x)$ of the agent in \mathbf{E} , given the state $x \in \mathbf{X}$ of the system as a whole. As it turns out, no further mathematical structure on \mathbf{X} is required for the results that follow, hence, with a mind toward inviting the broadest range of applications, we impose none, much in the spirit of McCarthy and Hayes discussion of situation calculus [47].

2) *Time and Transitions*: We model time \mathbb{T} as the set of integers (the subjective time of the agent), with $t = 0$ corresponding to the initial time. The basic objects of study are then *trajectories*, or maps of the form

$$\varphi : \mathbb{T} \rightarrow \mathbf{X}, \quad \varphi = (\varphi|_t)_{t \in \mathbb{T}} \quad (1)$$

We define abstract transitions in \mathbf{X} as follows:

Definition II.1 (*n-transitions*). An element of the $(n+1)$ -fold Cartesian power \mathbf{X}^{n+1} will be referred to as an *n-transition*. For any trajectory $\varphi : \mathbb{T} \rightarrow \mathbf{X}$ and $n \geq 0$ we define the map

$$d^n \varphi : \begin{cases} \mathbb{T} & \rightarrow & \mathbf{X}^{n+1} \\ t & \mapsto & \varphi|_{t-n} \times \cdots \times \varphi|_t \end{cases} \quad (2)$$

We refer to 0-transitions as *states*, and to 1-transitions simply as *transitions*. \square

Any setting where \mathbf{E} , \mathbf{X} and the transition structure of the system are specified (though, possibly, in an implicit fashion), implies constraints on the set of achievable trajectories. We will refer to such settings as *experiments*, within the framework of which each allowed trajectory will be referred to as a *run of the experiment*, while observations produced by sensors during a run (below) will be called *experiences*.

3) *Discrete Binary Agents*: A *discrete binary agent* (DBA) is endowed with a collection of binary sensors indexed by a finite set Σ . We will assume that each $a \in \Sigma$ is assigned an *order* $n_a \geq 1$, and a realization $\rho(a) \subseteq \mathbf{X}^{n_a+1}$. We then say that a is a *n_a -sensor*, or a *sensor of degree n_a* . For example, a 0-sensor – or *state sensor* – responds to the system entering a certain subset of \mathbf{X} (a "macro-state"), while a 1-sensor responds to the system experiencing a transition of a particular kind.

Evaluation of sensors is best viewed in the context of trajectories: a n -sensor $a \in \Sigma$ is applied to a trajectory φ and assigned a value at time $t \in \mathbb{T}$ according to the rule

$$\langle a : \varphi \rangle|_t = 1 \stackrel{\text{def.}}{\iff} d^n \varphi|_t \in \rho(a) \quad (3)$$

Here $\langle a : \varphi \rangle|_t$ denotes the measurement provided by the sensor a at time t given the trajectory φ . To avoid a profusion of parentheses and subscripts we will generally use bracket

notation to denote the evaluation of Boolean- and scalar-valued functions:

$$\langle a : x \rangle := \mathbb{1}_{\rho(a)}(x) \text{ whenever } a \in \Sigma \quad (4)$$

$$\langle g : s \rangle := g(s) \text{ whenever } g : S \rightarrow [0, 1] \quad (5)$$

and so forth. The symbol $\mathbb{1}_A$ will always denote the indicator function of a set A with respect to the appropriate super-set.

We will assume that Σ comes endowed with a map $a \mapsto a^*$ satisfying the following for all $a \in \Sigma$:

$$a^{**} = a, \quad a^* \neq a, \quad (6)$$

as well as

$$\rho(a^*) = \rho(a)^c := \mathbf{X}^{n_a+1} \setminus \rho(a) \quad (7)$$

We also introduce the virtual sensors $\mathbf{0}, \mathbf{0}^* \in P$ evaluating to

$$\langle \mathbf{0} : \varphi \rangle|_t \equiv 0, \quad \langle \mathbf{0}^* : \varphi \rangle|_t \equiv 1 \quad (8)$$

on any trajectory φ and at any time $t \in \mathbb{T}$. For subsets $A \subseteq \Sigma$ we will always use the notation A^* to denote the set of all a^* , a ranging over A .

The database structure we will be using is designed to maintain an approximate record of the relations among sensors in Σ believed by the agent to hold true throughout time. This record at time $t \in \mathbb{T}$ is encoded in a weak poc set structure $\mathbf{P}|_t$ over Σ (definition A.1).

For two n -sensors $a, b \in \Sigma$ this requirement translates into $a \leq b$ in $\mathbf{P}|_t$ being treated (for planning purposes) as the inclusion $\rho(a) \subseteq \rho(b)$ in the space \mathbf{X}^{n+1} . Note how the equivalent containment $\rho(b^*) \subseteq \rho(a^*)$ is encoded by the contra-positive implication $b^* \leq a^*$, which, by the definition of a weak poc set, holds if and only if $a \leq b$ does.

When a, b have different orders we are forced to replace this requirement by a weaker one: at any time t , our agents will interpret $a \leq b$ in $\mathbf{P}|_t$ as

$$\langle a : \varphi \rangle|_{t'} \leq \langle b : \varphi \rangle|_{t'} \quad (9)$$

holding for all $t' \in \mathbb{T}$. In other words, our agents assume that relations among sensors do not change over time⁶.

For example, if $a, b \in \Sigma$ where a is a state sensor and b is a transition sensor, consider the statements:

$$(\dagger) \langle a : \varphi \rangle|_t \leq \langle b : \varphi \rangle|_t, \quad (\ddagger) \langle b : \varphi \rangle|_t \leq \langle a : \varphi \rangle|_t$$

treated as identities over both φ and t . We see that (\ddagger) states all transitions of type b must terminate in a state of type a , while (\dagger) means that *only* transitions of type b could produce state a . It is clear that both kinds of statement are essential for planning.

⁶This is not to say that our agents are not allowed to change their minds regarding which relations hold true and which do not: the purpose of keeping a dynamic record of relations is to eventually uncover the 'correct' relations.

B. The Model Spaces

1) *A Record of Implications:* Informally, a “record of implications in Σ ” is a partial ordering on Σ reflecting the standard interactions between Boolean complementation and Boolean implication. Formally, our DBA will maintain, at any time $t \geq 0$, a *weak poc set structure* $\mathbf{P}|_t$ on Σ consisting of a partial order relation \leq satisfying, for all $a, b \in \Sigma$:

- 1) $\mathbf{0} \leq a$;
- 2) $a \leq b \Rightarrow b^* \leq a^*$.

Note that $a^* \neq a$ by the construction of Σ and compare with the definition in appendix A.1.

2) *Observations as vertices of a cube:* From the agent’s viewpoint, the current state of the experiment at time t is completely characterized by the measurements $\langle a : \varphi \rangle|_t$, where φ is the agent’s trajectory. Equivalently, the state may be encoded in a subset $O|_t \subset \Sigma$ satisfying $|O|_t \cap \{a, a^*\}| = 1$ for all $a \in \Sigma$. Such subsets of Σ are called *complete *-selections*. An incomplete measurement of the state would then correspond to a subset $O \subset \Sigma$ satisfying $|O \cap \{a, a^*\}| \leq 1$ for all $a \in \Sigma$, called an (incomplete) **-selection* – see definition in appendix A.6, along with remarks on the notation to follow.

Thus, one thinks of the collection $S(\Sigma)^0$ of all complete **-selections* on Σ as enumerating the possible sensory equivalence classes in the sensed space. However, some of the elements in this collection are redundant given the record $\mathbf{P}|_t$: an implication $a \leq b$ means that no $O \in S(\Sigma)^0$ containing $\{a, b^*\}$ is expected by the agent to be witnessed by any observation (see fig.9). Formally, a set $O \subset \Sigma$ is *coherent* (definition A.7), if no pair of elements $a, b \in O$ satisfies $a \leq b^*$.

3) *The model spaces:* The model space $\mathbf{M}|_t$ corresponding to the record $\mathbf{P}|_t$ takes the form of a cubical complex — a topological space constructed from a collection of vertices (the 0-skeleton), a set of edges (the 1-skeleton), and successively higher dimensional connecting cells in the form of cubes [32]. We choose the vertex set of $\mathbf{M}|_t$ to coincide with the set of coherent **-selections* in $S(\Sigma)^0$. Edges are inserted to join any pair of vertices A, B satisfying $|A \setminus B| = 1$ (this condition turns out to be symmetric). The hop-distance on the resulting graph may be seen as a variant of the crude, ‘information motivated’, Hamming distance on $\{0, 1\}^\Sigma$. The 1-dimensional skeleton of $\mathbf{M}|_t$ is further enriched with higher dimensional cubes to yield the space $\text{Cube}(\mathbf{P}|_t)$, as described in appendix A (definition A.8) for the interested reader. While a fairly detailed knowledge of the geometry and topology of spaces obtained in this way is essential for following our formal arguments regarding the modeling capabilities of this class, much of it is unnecessary for this section’s account of how the agent obtains its representation of $\mathbf{M}|_t$, the record $\mathbf{P}|_t$.

4) *Maintaining a record of the current state:* Returning to the problem of representing the current state, observe that $\mathbf{P}|_t$ is expected to change as time progresses, possibly giving rise to observations $O|_t$ that are incoherent with respect to $\mathbf{P}|_t$, and therefore represent points ‘outside’ the model space. While the raw observation $O|_t$ *must* be applied to the agent’s data structure in hopes of improving $\mathbf{P}|_t$, the agent must resolve the contradiction within the framework of its current model,

replacing the incoherent complete observation $O|_t$ in its role as the record of the current state kept by the agent with a coherent but incomplete observation (104), $S|_t := \text{coh}(O|_t)$, satisfying certain naturality requirements — see appendix B-E2 for the complete technical discussion.

This means the agent resolves the contradiction at the price of introducing ambiguity into its record of the current state: instead of having a single vertex of $\mathbf{M}|_t$ representing the current state (“complete knowledge”), any vertex containing the set $S|_t$ may turn out to be the correct current state.

The complexity of coherent projection (lemma III.4) and its role in the agent’s reasoning processes, its interplay with the convexity theory of the model space $\mathbf{M}|_t$ and its interpretation as the basis for viewing our architecture as a connectionist model (albeit a very limited one) of cognition will all be discussed in section III-B.

C. Snapshots

In [30] we have introduced the rather loose notion of a *snapshot*, aiming to outline a class of database structures for dynamically maintaining weak poc-set structures from a sequence of observations made by an agent along a trajectory φ through \mathbf{X} . A rigorous treatment of this tool requires some careful definitions.

Definition II.2. Denote by \mathbf{K}_Σ the graph obtained from the complete graph over the vertex set Σ by removing all edges of the form aa^* , $a \in \Sigma$. Edges of \mathbf{K}_Σ will be referred to as *proper pairs* in Σ . We will abuse notation and write $ab \in \mathbf{K}_\Sigma$ for the edge $\{a, b\}$ of \mathbf{K}_Σ . \square

The graph \mathbf{K}_Σ is the scaffolding for snapshots:

Definition II.3 (Snapshot). A snapshot \mathbf{S} over Σ consists of the following:

- (a) **State.** Each vertex $a \in \Sigma$ of \mathbf{K}_Σ is assigned a binary state $\#_a \mathbf{S} \in \{0, 1\}$. The set

$$\#\mathbf{S} = \{a \in \Sigma \mid \#_a = 1\} \quad (10)$$

is called the *state vector of the snapshot* \mathbf{S} and is required to be a **-selection* on Σ (definition A.6).

- (b) **Edge weights.** Each edge $ab \in \mathbf{K}_\Sigma$ is assigned a non-negative real number denoted $w_{ab} = w_{ab}(\mathbf{S})$.
- (c) **Learning Thresholds.** Each edge $ab \in \mathbf{K}_\Sigma$ carries a non-negative real number $\tau_{ab} = \tau_{ab}(\mathbf{S})$ satisfying

$$\tau_{ab} = \tau_{a^*b} = \tau_{ab^*} = \tau_{a^*b^*} \leq \frac{1}{4}.$$

For every $ab \in \mathbf{K}_\Sigma$, the restriction of \mathbf{S} to the subgraph induced by the vertices a, a^*, b and b^* will be denoted by $\mathbf{S}|_{ab}$ and referred to as a *square* in \mathbf{S} . \square

The original motivation of [30] for the notion of a snapshot is twofold:

- 1) **Maintaining a consistent representation of the current state.** For this purpose we will generally assign the coherent projection of the current state measurement to be stored in $\#\mathbf{S}$.
- 2) **Learning implications in the sensorium.** To learn an estimate of the implication order on Σ inherited




Fig. 1. determining edge orientations in a snapshot by restricting attention to $\mathbf{S}|_{ab}$.

from its realization in \mathbf{X} it should suffice to maintain a system of weights w_{ab}^t on $\mathbf{S}|_t$ quantifying the relevance (e.g. frequency) of the event $a \wedge b$, allowing one to partially orient the snapshot according to the rule of thumb illustrated in figure 1.

The graphical representation derived from a snapshot in this manner does not automatically define a weak poc set, but is nearly there:

Definition II.4 (poc graph). A *poc graph* Γ over Σ is a subgraph of \mathbf{K}_Σ endowed with an orientation which satisfies, for every $ab \in \mathbf{K}_\Sigma$:

- If $ab \in \Gamma$ then $b^*a^* \in \Gamma$;
- If $ab \in \Gamma$ then $a^*b, ba^*, b^*a, ab^* \notin \Gamma$.

By abuse of notation, we use the symbol ab to mean the directed edge emanating from a and pointing to b (if it exists in Γ). \square

In order for a poc graph to represent a weak poc set structure on Σ one needs:

Lemma II.5 (derived poc set). *The transitive closure of the orientation relation on a poc graph Γ over Σ is a weak poc set structure on Σ if and only if Γ has no directed cycles.*

Proof. This follows directly from the discussion in example A-A2. \square

The rest of this section mainly deals with characterizing a large class of snapshots encoding acyclic directed poc graphs and with means of evolving snapshot representations of \mathbf{X} from trajectories. Given a trajectory φ of our agent through \mathbf{X} , the collection of *coincidence indicators*

$$c_{ab}^t := \langle a : \varphi \rangle|_t \cdot \langle b : \varphi \rangle|_t \quad (11)$$

may be used to evolve a sequence of snapshots $\mathbf{S}|_t$ representing, at any time $t > 0$, the cumulative influence of the agent's observations on its perception of implications in the sensorium.

D. Probabilistic Snapshots and Acyclicity

The following set-theoretic identities among the coincidence indicators are easily verified for all $a, b, c \in \Sigma$:

$$\begin{aligned} c_{aa^*}^t &= 0 \\ c_{ab}^t &= c_{ba}^t \\ c_{aa}^t &= c_{ab}^t + c_{ab^*}^t \\ c_{ab}^t + c_{a^*b}^t + c_{a^*b^*}^t + c_{ab^*}^t &= 1 \\ c_{ab^*}^t + c_{b^*c}^t + c_{ca^*}^t &= c_{a^*b}^t + c_{b^*c}^t + c_{c^*a}^t \end{aligned} \quad (12)$$

These identities motivate considering snapshots with weights obeying analogous constraints:

Definition II.6 (Probabilistic Snapshot). We say that a snapshot \mathbf{S} is *probabilistic*, if $\#\mathbf{S}$ is a coherent $*$ -selection and the edge weights satisfy the following:

- **Consistency constraint.** if $ab, ac \in \mathbf{K}_\Sigma$ then:

$$w_{ab} + w_{ab^*} = w_{ac} + w_{ac^*} \quad (13)$$

- **Normalization constraint.** for any $ab \in \mathbf{K}_\Sigma$:

$$w_{ab} + w_{a^*b} + w_{a^*b^*} + w_{ab^*} = 1 \quad (14)$$

- **Orientation constraint.** if $ab, bc, ac \in \mathbf{K}_\Sigma$ then:

$$w_{a^*b} + w_{b^*c} + w_{c^*a} = w_{ab^*} + w_{bc^*} + w_{ca^*} \quad (15)$$

We denote the set of all probabilistic snapshots over Σ by \mathcal{P}_Σ , or simply \mathcal{P} when there is no danger of confusion. \square

A fundamental observation regarding probabilistic snapshots is the following

Proposition II.7 (Acyclicity Lemma). *Suppose \mathbf{S} is a probabilistic snapshot over Σ and Γ is a poc graph satisfying the orientation cocycle condition:*

$$ab \in \Gamma \Rightarrow \omega(ab) := w_{a^*b} - w_{ab^*} > 0 \quad (16)$$

Then Γ contains no directed cycles.

Proof. See appendix B-A. \square

This proposition puts the vague notion from figure 1 on how to derive implications from a snapshot on a firm footing:

Proposition II.8 (Poc graphs from snapshots). *Let \mathbf{S} be a probabilistic snapshot. Construct a poc graph $\text{Dir}(\mathbf{S})$ by setting*

$$ab \in \text{Dir}(\mathbf{S}) \Leftrightarrow w_{ab^*} < \min \left\{ \tau_{ab}, w_{ab}, w_{a^*b}, w_{a^*b^*} \right\} \quad (17)$$

Then $\text{Dir}(\mathbf{S})$ is an acyclic poc graph. \square

Proof. The symmetries of τ_\bullet and w_\bullet immediately imply $ab \in \text{Dir}(\mathbf{S})$ iff $b^*a^* \in \text{Dir}(\mathbf{S})$. The strict inequality in (17) implies the second condition of a poc graph holds as well. Since the orientation cocycle is positive on every edge of $\text{Dir}(\mathbf{S})$ by definition, the acyclicity lemma applies. \square

The element of thresholding present in (17) may also be used as a part of the updating procedure of a probabilistic snapshot, without affecting the derived poc set:

Proposition II.9 (Snapshot Truncation). *Let \mathbf{S} be a probabilistic snapshot. Define a new snapshot $\lfloor \mathbf{S} \rfloor$ to have the same*

state as \mathbf{S} while for every $ab \in \mathbf{K}_\Sigma$ satisfying (17) the weights are updated as follows:

$$\begin{cases} w_{ab^*} & \mapsto 0 \\ w_{ab} & \mapsto w_{ab} + w_{ab^*} \\ w_{a^*b^*} & \mapsto w_{a^*b^*} + w_{ab^*} \\ w_{a^*b} & \mapsto w_{a^*b} - w_{ab^*} \end{cases} \quad (18)$$

Then $\text{Dir}(\mathbf{S}) = \text{Dir}(\lfloor \mathbf{S} \rfloor)$.

Proof. The proof amounts to a direct verification that the $\lfloor \mathbf{S} \rfloor$ is probabilistic, and that $ab \in \mathbf{K}_\Sigma$ satisfies (17) in $\lfloor \mathbf{S} \rfloor$ if and only if the same condition is satisfied by ab in \mathbf{S} . \square

Following lemma II.5 we may now safely define:

Definition II.10 (Derived Poc Set). Let \mathbf{S} be a probabilistic snapshot. Denote by $\text{Poc}(\mathbf{S})$ the weak poc set structure obtained by setting $a \leq b$ iff there exists a directed path in $\text{Dir}(\mathbf{S})$ from a to b . \square

We now proceed to introduce and study two possible snapshot constructions.

E. Empirical Snapshots

The empirical snapshot structure maintains an empirical approximation of the relative frequencies of observations of the form $a \wedge b$, $a, b \in \Sigma$. For any trajectory φ of the agent through \mathbf{X} we could try setting

$$w_{ab}^t = \sum_{k=1}^t c_{ab}^k, \quad (19)$$

with $\mathbf{S}|_t$ a trivial snapshot for all $t \leq 0$.

Definition II.11. A snapshot \mathbf{S} with $w_{\bullet} = 0$ is said to be *trivial* and denoted Null . \square

Properties (12) then imply that $\mathbf{S}|_t$ satisfies the consistency and cocycle constraints (defn. II.6) for all $t > 0$, and would satisfy the normalization constraint if we replace the weights w_{ab}^t by $\frac{1}{t}w_{ab}^t$ throughout.

1) *Construction and Properties:* The formal construction is as follows:

Definition II.12 (Empirical Snapshot). A snapshot \mathbf{S} over Σ is an *empirical snapshot* if the following conditions are satisfied:

- For all $ab \in \mathbf{K}_\Sigma$, $w_{ab} \in \mathbb{Z}_{\geq 0}$;
- For all $ab \in \mathbf{K}_\Sigma$, the expression

$$w_a := w_{ab} + w_{ab^*} \quad (20)$$

is independent of the choice of b , and vanishes only if $\#_a = 0$;

- The following expression does not depend on $a \in \Sigma$:

$$\text{Clk}(\mathbf{S}) := w_a + w_{a^*} \quad (21)$$

Denote the set of empirical snapshots over Σ by \mathcal{E}_Σ (or just \mathcal{E} when justified). \square

The evolution of an empirical snapshot under a sequence of observations is then defined through:

Definition II.13 (Empirical Update). Let \mathbf{S} be an empirical snapshot and let $O \subset \Sigma$ be complete $*$ -selection. The snapshot $O * \mathbf{S}$ is the empirical snapshot obtained from \mathbf{S} by setting

$$w_{ab}(O * \mathbf{S}) := w_{ab}(\mathbf{S}) + \langle \mathbf{1}_O : a \rangle \cdot \langle \mathbf{1}_O : b \rangle \quad (22)$$

for all $ab \in \mathbf{K}_\Sigma$. The state of $O * \mathbf{S}$ is set to $\text{coh}(O)$ (where, recall, this is the the coherent projection (104) computed with respect to the weak poc set structure derived from the new weights). \square

Definition II.14 (Evolution). We say that a snapshot \mathbf{T} over Σ is an evolution of a snapshot \mathbf{S} , either if $\mathbf{S} = \mathbf{T}$ or if there is a sequence $(O_k)_{k=1}^m$ of complete $*$ -selections in Σ such that $\mathbf{T} = O_m * \dots * O_1 * \mathbf{S}$. \square

Empirical snapshots are characterized by their ancestry:

Lemma II.15. A non-trivial snapshot \mathbf{S} over Σ is empirical if and only if it is an evolution of the trivial snapshot. \square

Proof. See appendix B-B. \square

Having characterized empirical snapshots as evolutions of the trivial one, we return to the observation that the weight $w_{\bullet}(\mathbf{S})/\text{Clk}(\mathbf{S})$ on \mathbf{K}_Σ — see (21) — defines a probabilistic snapshot. We may thus define $\text{Dir}(\mathbf{S})$ accordingly, by setting

$$ab \in \text{Dir}(\mathbf{S}) \Leftrightarrow w_{ab^*} < \min \left\{ \begin{array}{l} \tau_{ab} \cdot \text{Clk}(\mathbf{S}), \\ w_{ab}, w_{a^*b}, w_{a^*b^*} \end{array} \right\} \quad (23)$$

and conclude that:

Proposition II.16 (empirical implies acyclic). If \mathbf{S} is an empirical snapshot, then $\text{Dir}(\mathbf{S})$ as defined in (23) is an acyclic poc graph, and $\text{Poc}(\mathbf{S})$ as defined in defn. II.10 is a weak poc set structure on Σ . \square

We will henceforth refer to DBAs endowed with empirical snapshots and utilizing the empirical update as *empirical agents*.

2) *Performance of Empirical Agents:* In this paper we restrict attention to agents endowed with a fixed finite set of actions. An agent starting out at time $t = 0$ with a trivial snapshot $\mathbf{S}|_0$ has no knowledge of its environment, and is therefore assumed to engage in random exploration for some time, until actionable information becomes available. This motivates the question as to how well the memory structure of an empirical agent performs in this initial stage.

In the case where \mathbf{X} is finite and the agent's actions are deterministic it is easy to formulate this: Let Σ_{act} be the set of available actions, and consider the graph with vertex set \mathbf{X} , where a vertex x is joined to a vertex y labeled by an action $\alpha \in \Sigma_{act}$ if applying α at x results in y . Thus, \mathbf{X} becomes endowed with the structure of a Markov chain, where we draw actions uniformly at random in every state. Focus on the case when all the actions available to the agent are reversible in the sense that there is an edge from x to y if and only if there is an edge from y to x (loops are allowed as well). Then the corresponding Markov chain is a random walk and its stationary distribution over \mathbf{X} , denoted by π , is uniform [60] over each connected component of the resulting transition

system. Thus, each normalized weight w_{ab}^t/t is nothing but the empirical estimate of the joint probability, given by π , for $a, b \in \Sigma$ to fire synchronously.

Restricting to a reachability component, we may assume \mathbf{X} is connected. By abuse of notation, for $ab \in \mathbf{K}_\Sigma$ denote

$$\pi(ab) = \pi(\rho(a) \cap \rho(b)) \quad (24)$$

Let Dir^∞ denote the matrix⁷ with entries $\text{Dir}_{ab}^\infty = 1$ whenever $\rho(a)$ is contained in $\rho(b)$ up to a precision of τ_{ab} , that is:

$$\pi(\rho(ab^*)) < \min(\tau_{ab}, \pi(ab), \pi(a^*b^*), \pi(a^*b)) \quad (25)$$

and set $\text{Dir}_{ab}^\infty = 0$ otherwise. This matrix represents the true poc set structure to be learned by the agent, as determined by the fixed learning thresholds. Analogously we let Dir^t be the matrix with $\text{Dir}_{ab}^t = 1$ iff the directed edge $ab \in \mathbf{K}_\Sigma$ is contained in the derived poc graph of $\text{Dir}(\mathbf{S}|_t)$ (and 0 otherwise). A good measure of the agent's performance would be the behavior of the total error

$$\text{Err}(t) := \|\text{Dir}^t - \text{Dir}^\infty\|_1 \quad (26)$$

over time (the matrices viewed simply as vectors in ℓ^1 of the appropriate dimension). By Theorem 5.1 in [60], the agent's random walk converges to π at an exponential rate depending only on the transition structure of \mathbf{X} determined by the actions Σ_{act} . We conclude:

Proposition II.17. *Suppose a DBA performs a random walk on a connected \mathbf{X} , at each moment in time performing one of a fixed finite set of reversible deterministic actions. Then $\text{Err}(t)$ converges to zero at an exponential rate.*

With such strong performance guarantees for a broad class of empirical agents we are left to examine the variation in performance as a function of the geometry/topology of the environment (beyond the guarantees given by the preceding discussion) we have run simulations in the following settings:

- The agent performs a random walk along a path with 20 edges (example A-D1), choosing between one step forward and one step back uniformly at random for every $t \in \mathbb{T}$, learning the poc structure of a sensorium consisting of 20 'GPS' sensors, as described in example A-D1;
- The agent performs a random walk along a cycle with 20 edges, choosing between a clockwise and a counter-clockwise step uniformly at random for every $t \in \mathbb{T}$, and learning a sensorium consisting of 20 beacon sensors as described in example A-D3;
- The agent performs a random walk (up/down/left/right) on a square grid with 10 'GPS' sensors along each of the x - and y - axes;
- The agent performs a random walk along (forward/back) a path with 20 edges, but the sensors are chosen to have random activation fields (randomly chosen subsets of the set of vertices along the path); the sensor fields have been drawn anew prior to each separate run.

⁷Recall that $\text{Dir}(\mathbf{S})$ introduced in Prop. II.8 is a directed graph. The new notation is intended to connote a matrix representation of such a graph.




Fig. 2. Logarithmic plot of the mean number of incorrect edges in the derived poc graph of an empirical snapshot (20 sensors), for learning thresholds varying linearly between $\frac{1}{4}$ (cyan/light) and $\frac{1}{20^3}$ (blue/dark), averaged over 50 runs of random walks each.

The number of sensors is the same for each setting, and each agent carries out 50 runs of a length that is cubic in the number of sensors, starting at a random position with an empty snapshot. We have tested 10 different agents for each setting, corresponding to 10 different values of the learning threshold, spread linearly in the interval from $1/(20)^3$ to the maximal meaningful value of $1/4$ (where one should not expect much useful learning to occur).

The results are summarized in figure 2 plotting $\text{Err}(t)$, where we have replaced the matrix Dir^∞ as defined in (25) by the $\{0, 1\}$ -valued matrix

$$\text{Dir}_{ab}^\infty = 1 \Leftrightarrow \rho(a) \subseteq \rho(b) \quad (27)$$

to render the effect of choosing different values for the learning threshold more visible in the graph of $\text{Err}(t)$.

The resulting plots show significant, though subtle, differences in performance between the four settings, illustrating the similarities and differences in the weak poc set structures being approximated, most notably:

- The sharper initial decline in the mean deviation for (b) and (c) in comparison to (a) is expected due to the relative abundance of crossing in the former, as opposed to complete nesting (see definition A.22) in the latter.
- Performance in the random setting (d) seems to lag significantly behind performance in any of the structured settings.
- Performance in the completely nested setting (a) seems to provide exponentially fast learning no matter what; by contrast, the other settings seem to experience a transition between two modes, depending on how small the learning threshold τ is:
 - 1) For large τ , the deviation plateaus.
 - 2) For small enough τ , the deviation decreases to zero in finite time.

We expect the critical value of τ in any finite setting to be somewhere around the minimum probability of a state (under the stationary distribution of the random walk): in order for a relation $a < b$ to be put on record, it is necessary for the agent to have visited $\rho(a) \cap \rho(b^*)$ at a frequency below τ ; the smaller τ is, the fewer false relations will be recorded for posterity.

- Recalling that the poc set representing the ground truth in (c) is the direct sum (see A-D2) of two smaller copies of (a) having 10 sensors each, we see that the crossing relations between the x -axis sensors and their y -axis counterparts account for 800 of the 1600 entries (two 20×20 null sub-matrices) in the adjacency matrix of the derived graph. Thus, the two experiments are not that different: loosely speaking, the 10×10 square grid experiment projects onto a Cartesian product of two 10-path experiments where the random walk on the 10-path becomes a lazy walk with probability $\frac{1}{2}$ to stay put. In other words, the behavior of (c) may be inferred from the behavior of (a).
- Not so when comparing (a) and (c) with (b): note how the sub-critical values of the learning parameters in (a),(c) and (d) force the deviation plot to ‘plunge’ into the x -axis *versus* the horizontal asymptote behavior of (b). In view of theorem A.33), our guess is that the environment (the circle) not being contractible has something to do with this qualitative change in behavior, but this requires further investigation.

F. Discounted Snapshots

A notable weakness of empirical snapshots as a data structure is their potential high cost in space, due to the need for indefinitely maintaining integer-valued counters. In some sense, the *entire* history of the agent matters, and, in some sense, matters too much. This motivates the search for an alternative, more quantized, updating mechanism whose dependence on any given past observation weakens at a fixed rate.

1) Construction and Properties:

Definition II.18. (discounted update) Let $q \in [0, 1]$ and let \mathbf{S} be a probabilistic snapshot over Σ . For any complete $*$ -selection O on Σ we define the snapshot $O *_q \mathbf{S}$ to be the snapshot with weights determined by

$$w_{ab}(O *_q \mathbf{S}) := qw_{ab}(\mathbf{S}) + (1 - q) \langle \mathbf{1}_O : a \rangle \cdot \langle \mathbf{1}_O : b \rangle \quad (28)$$

The state of $O *_q \mathbf{S}$ is set to $\text{coh}(O)$, the reduction being computed with respect to the weak poc set structure derived from the new weights. Finally, define the q -discounted update of \mathbf{S} to be the snapshot $[O *_q \mathbf{S}]$ and we refer to q as the *decay parameter*. \square

A significant advantage of the discounted update is its applicability to arbitrary probabilistic snapshots:

Lemma II.19. *The q -discounted update of a probabilistic snapshot by a complete $*$ -selection is probabilistic.* \square

Proof. It is clear that the discounted update preserves probabilisticity. Proposition II.9 finishes the proof. \square

Consider the length of time (or the amount of evidence) it takes a discounted snapshot to acquire an implication, compared to the amount of evidence required for giving up an implication already on record.

Assuming a fixed value of the decay parameter q over a considerable length of time, a lower bound on the amount of time Δt required for w_{ab^*} to become small enough for a relation $a \leq b$ to be put on record is given by the situation when a long enough sequence of consecutive observations with $a \wedge b^*$ *not* occurring is made:

$$q^{\Delta t} < \tau_{ab} \Leftrightarrow \Delta t > \frac{\log_2 \tau_{ab}}{\log_2 q} \quad (29)$$

On the other hand, once the relation $a \leq b$ has been put on record, the number Δt of successive observations of $a \wedge b^*$ required for replacing this relation with $a \nabla b$ must satisfy:

$$\Delta t(1 - q) > \tau_{ab} \Leftrightarrow \Delta t > \frac{\tau_{ab}}{1 - q} \quad (30)$$

– this much is guaranteed by the truncation mechanism. Overall, it seems that choosing a value of q with $1 - q$ sufficiently small should produce meaningful learning: lower values of τ_{ab} make it both harder to learn and easier to unlearn a false relation, while maintaining a qualitative difference between the necessary requirements for either process.

Keeping q fixed over long periods of time places an emphasis on the values of the learning thresholds τ_{ab} . As these values do not have to be chosen uniformly over the snapshot, one might want to vary the values of the learning thresholds individually with the aim of altering the flexibility of the learning process in the corresponding square. This opens up a doorway to employing methods for varying the learning thresholds and the decay parameter in ways analogous to [26] and [22] as a means of improving the quality/dependability of the model space. The simulation results below emphasize the need for this kind of control, showing that a discounted agent is much more susceptible to changes in geometry and topology/combinatorics of the sensor fields than an empirical one.

2) *Performance Analysis:* Figure 3 compares the mean performance of time-discounted snapshot learning from a random walk in the four settings described earlier in II-E2, for the values of the decay parameter q given by $q = 1 - \frac{1}{2^{k+2}}$, $k = \{0, \dots, 9\}$.

One immediately notices, in comparison with the empirical case, that the dependence of the learning process on the discount parameter is not monotone: it would seem that a choice of $k = 5$ works best for all settings in terms of optimizing the eventual deviation, — though it is hard to say what ‘best’ would even mean for (d) — while a choice of $k = 4$ is more reasonable given the observed waiting time until meaningful learning occurs in the structured environments (a)-(c).

Similar observations to those made for the empirical case (figure 2) regarding the interplay between ‘learning modes’




Fig. 3. Mean number of incorrect edges in the derived poc graph of a discounted snapshot in 4 environments (20 sensors each) for varying values of the decay parameter, $q = 1 - \frac{1}{2^{k+2}}$, k from 0 (red/dark) to 9 (yellow/light), averaged over 50 runs of a random walk.

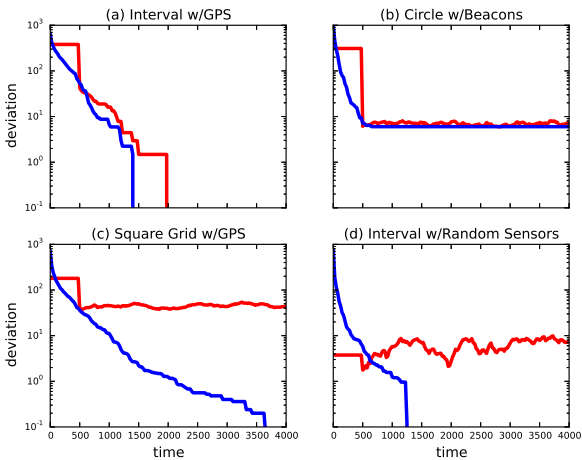


Fig. 4. A comparison of the mean number of incorrect edges in the derived poc graph as a function of time, for an empirical snapshot (blue) and a discounted one (red). Here $\tau = 1/20^3$ and $q = 1 - \frac{1}{2^6}$.

and geometry/topology can be made here as well, but are more subtle, as the comparison in figure 4 shows.

G. Further Adjustments to the Weak Poc Structure

The implication record constructed from a probabilistic snapshot in the preceding section does not recognize possible equivalences among sensations: if, for whatever reason, a relation of the form

$$w_{ab^*} = w_{a^*b} = 0 \quad (31)$$

takes place in a snapshot \mathbf{S} , it becomes reasonable to interpret it as the logical equivalence $a \Leftrightarrow b$, yet $\text{Dir}(\mathbf{S})$ will not register any relations in the square $\mathbf{S}|_{ab}$, barring an agent equipped with \mathbf{S} from utilizing the currently observed equivalence.

Thus, an adjustment of $\text{Dir}(\mathbf{S})$ is required if we are to allow our agents the advantage of reasoning about equivalences. The following extension of $\text{Dir}(\mathbf{S})$ turns out to serve our purposes for a restricted class of probabilistic snapshots:

Definition II.20. Let \mathbf{S} be a probabilistic snapshot. The poc graph $\text{Dir}(\mathbf{S})_0$ is defined to be the poc graph obtained from $\text{Dir}(\mathbf{S})$ by adding the directed edges ab, ba, a^*b^*, b^*a^* for each $ab \in \mathbf{K}_\Sigma$ satisfying (31). \square

It turns out that $\text{Dir}(\mathbf{S})_0$ gives rise to an adequate weak poc set structure and model space, provided \mathbf{S} satisfies the additional requirement:

Definition II.21. A snapshot \mathbf{S} is said to satisfy the *triangle inequality*, if

$$w_{a^*b} + w_{ab^*} + w_{b^*c} + w_{bc^*} \geq w_{a^*c} + w_{ac^*} \quad (32)$$

holds whenever $ab, bc, ac \in \mathbf{K}_\Sigma$. \square

A class of examples of special significance in this work is that of snapshots \mathbf{S} whose edge weights are derived from a measure μ on a space Z by pulling back along a realization $\rho: \Sigma \rightarrow Z$ as follows:

$$w_{ab} = \mu(\rho(a) \cap \rho(b)) \quad (33)$$

The triangle inequality for \mathbf{S} is then an immediate consequence of the well-known (e.g. [61], chapter 3) triangle inequality for measures:

$$\mu(A \Delta C) \leq \mu(A \Delta B) + \mu(B \Delta C), \quad (34)$$

where $A, B, C \subset Z$ are arbitrary measurable sets and $A \Delta B$ denotes the symmetric difference $(A \setminus B) \cup (B \setminus A)$.

The coincidence indicators c_{ab}^t of (11) are a special case of this example (where μ is an atomic measure), and so are empirical snapshots (as their weights are sums of coincidence indicators). Discounted snapshots fall into this class, too, as their weights are convex combinations of coincidence indicators.

Due to the technical nature of the interactions between $\text{Dir}(\mathbf{S})$ and the extension $\text{Dir}(\mathbf{S})_0$, we postpone the formal discussion of these interactions to appendix B-C. The bottom line, however, is that for any probabilistic snapshot structure satisfying the triangle inequality our agent may safely apply the control protocols of the next section to the *extended* poc graph derived from the agent's current snapshot to arrive at action choices while taking advantage of the perceived equivalences within the sensorium.

Although technically we are obliged to distinguish between $\text{Dir}(\mathbf{S})$ and $\text{Dir}(\mathbf{S})_0$, as well as between the weak poc set structures they correspond to, we will treat these objects as identical for the sake of simplifying the rest of the exposition.

III. CONTROL WITH SNAPSHOTS

This section introduces the basic control function of a snapshot. We begin with introducing a formalism designed to treat discrete actions as a sub-structure of the binary sensorium, and discuss the effect of this formalism on shaping the model space (III-A). We next turn to a discussion of the

snapshot $\mathbf{S}|_t$ as a highly efficient computational mechanism for coherent state-updating and for decision making based on the geometry of the model space $\mathbf{M}|_t$ (III-B).

At the technical level, this section requires an understanding of the convexity theory of cubings: The classical results are covered in appendix A-D, while our new technical results underlying the use of snapshots for greedy navigation in cubings are covered in appendix B-E.

Building on these results, section III-B2 introduces the mechanism of *signal propagation over a snapshot* which realizes the computation both of coherent projection and of closest point projections to prescribed convex subsets of the associated model space. This mechanism suggests a view of snapshot architecture as highly simplified connectionist architectures, and some related work in the literature is discussed.

At the heart of our proposed decision-making algorithm is an assumption that the sensorium is rich enough to detect direct causal relations between actions and other sensations. We provide a fairly broad formalization of this assumption in III-B3 (with an example in III-B4), and prove the ability of an agent to correctly ‘halucinate’ the immediate consequences of taking an action, provided sufficient exposure.

An algorithm using this tool to attempt greedy navigation over $\mathbf{M}|_t$ to a specified target state is proposed in III-B5, and some of its failure modes are discussed as a motivation for future research on judicious dynamical expansions of the sensorium which would allow the agent to overcome the navigational obstructions formed by states in $\mathbf{M}|_t$ having no witness in the situation space \mathbf{X} .

Finally, in III-C3, we explore the performance of some *excitation-driven DBAs*: agents endowed with an excitation level that changes depending on their distance from a target in the environment \mathbf{E} ; the agents are capable of sensing an increase or a decrease in excitation, and seek instant gratification in the sense of operating on the mandate to always pick an action guaranteeing an increase in excitation (or else act randomly). We compare the performance of such agents in the domains considered in II-E2 and II-F2; in these domains it is easy to guarantee arrival to the target given *a-priori* knowledge of the correct snapshot structure, but we are interested in the agent’s performance as they learn the problem “from scratch”.

A. Defining Actions

We will now restrict attention to DBAs with a sensorium Σ endowed only with state (degree 0) and transition (degree 1) sensors. As before, we denote the realization of a sensor $a \in \Sigma$ by $\rho(a)$, where $\rho(a) \subseteq X$ for a state sensor and $\rho(a) \subseteq X \times X$ for a transition sensor. Thus, state sensors and transition sensors may be viewed as Boolean and situational fluents over the situation space \mathbf{X} , which is sufficient for setting up a discussion of actions and competencies according to McCarthy and Hayes [47].

For our agents, we posit a set $\Sigma_{act} \subset \Sigma$ of transition sensors, each of which may be switched on and off *at will*, earning them the name of *actions*. To be precise, our requirements are:

- **Actions are binary.** We assume $\Sigma_{act} \cap \Sigma_{act}^* = \emptyset$, and we denote the poc subset $\Sigma_{act} \cup \Sigma_{act}^* \cup \{\mathbf{0}, \mathbf{0}^*\}$ by **Act**.
- **Every action has outcomes.** For any $\alpha \in \Sigma_{act}$ and $x \in \mathbf{X}$, the sets

$$\alpha(x) = \{y \in \mathbf{X} \mid x \times y \in \rho(\alpha)\} \quad (35)$$

are non-empty subsets of \mathbf{X} .

In this we depart slightly from the accepted notion of actions in the literature on transition systems of various flavors (e.g. [16],[62]), where actions are attached to states and the collection of actions available at each state may differ, depending on that state. Instead, we consider actions as nothing more than control signals, sent by the agent’s ‘mind’ to the agent’s ‘body’ in order to invoke (or not) one or more of a fixed set of available behaviors. It is the purpose of the ‘mind’ to identify whether or not a control signal produces meaningful outcomes as those outcomes are being sensed.

1) *Invoking Actions Synchronously:* Our sensor-centric approach to actions reflects the viewpoint that (1) an action $\alpha \in \Sigma_{act}$ taken at a state $x \in \mathbf{X}$ imposes a time-independent restriction on the set of states the system may enter in the following moment, and (2) the agent is capable — at least in principle — of observing its own decisions as they are being invoked. We must now discuss the precise extent to which these principles may or may not restrict our initial suggestion that the sensations in **Act** are controllable.

For example, consider the situation where the agent is not engaging in an action $\alpha \in \Sigma_{act}$ during a transition from state x to state y . This implies α^* is on during this transition, which restricts the possible values of y to $\mathbf{X} \setminus \alpha(x)$. Hence, not invoking any of the available $\alpha \in \Sigma_{act}$ must then restrict y to the intersection $\bigcap_{\alpha \in \Sigma_{act}} \mathbf{X} \setminus \alpha(x)$, the set of possible outcomes of the “no-action”.

More generally, allowing a number of actions to be taken at the same time (while not engaging in the rest) forces the following interpretation by our sensing model:

- 1) A *generalized action* by the agent is a complete *-selection A on **Act** (recall definition A.6);
- 2) The realization of a (generalized) action $A \in S(\mathbf{Act})^0$ is defined to be

$$\rho(A) = \bigcap_{\alpha \in A} \rho(\alpha), \quad (36)$$

or, equivalently, for every $x \in \mathbf{X}$, the set of possible outcomes of an action A equals

$$A(x) = \bigcap_{\alpha \in A} \alpha(x) \quad (37)$$

For this extended collection of actions one notices that the second requirement of an action — $A(x) \neq \emptyset$ for all $x \in \mathbf{X}$ — does not necessarily hold: for example, moving forward along a rail contradicts any motion in the opposite direction. We will say that a generalized action $A \in S(\mathbf{Act})^0$ is *admissible at* $x \in \mathbf{X}$ if $A(x) \neq \emptyset$, and that A is *admissible*, if it is admissible at x for all $x \in \mathbf{X}$.

Aside from setting natural bounds on the meaning of the initial statement that actions are available to the agent at will,

the notion of admissibility of a generalized action explains how to interpret the poc set structure induced on \mathbf{Act} from the realization ρ : if

$$\alpha < \beta \Leftrightarrow \rho(\alpha) \subset \rho(\beta) \quad (38)$$

happens to hold for $\mathbf{P}|_t$, then every generalized action admissible at a point $x \in \mathbf{X}$ defines a vertex of $\text{Cube}(\mathbf{Act})$ no matter the choice of $x \in \mathbf{X}$. Similarly, generalized actions not showing up as vertices $\text{Cube}(\mathbf{Act}|_t)$, where $\mathbf{Act}|_t$ denotes the restriction of the poc set structure $\mathbf{P}|_t$ to \mathbf{Act} , represent the agent's belief at time t regarding combinations of elementary actions it cannot achieve at that moment.

In the simple examples considered in this paper all agents will be endowed with a collection of mutually exclusive atomic actions. By this we mean that $\alpha < \beta^*$ holds for all $\alpha, \beta \in \Sigma_{act}$ ($\alpha \neq \beta$). Equivalently, only the "no-action", $\{\alpha^* | \alpha \in \Sigma_{act}\}$, and the "pure" actions $\{\alpha\} \cup \{\beta^* | \beta \in \Sigma_{act} \setminus \{\alpha\}\}$ are admissible, and the resulting cubing $\text{Cube}(\mathbf{Act})$ takes the form of a *starfish*: a tree with only one vertex of degree ≥ 2 given by the "no-action" and with a set of leaves in one-to-one correspondence with the set of "pure" actions (see example A-D1 and figure 11b).

2) *Observations*: The following set is the set of *observations* in \mathbf{P} (note that it is closed under the $*$ -operator):

$$\text{Obs} := (\Sigma \setminus \mathbf{Act}) \cup \{\mathbf{0}, \mathbf{0}^*\}, \quad (39)$$

and stands for the set of "passive" sensations, as opposed to actions. Sections II-C-II-G explain how a trajectory $\varphi|_t, t \geq 0$ may be used to form an evolving sequence of weak poc-set structures $(\mathbf{P}|_t)_{t \geq 0}$ over Σ , with each $\mathbf{P}|_t$ representing the agent's belief at time t regarding which implications among the sensors in Σ hold true throughout time. Two poc subsets of $\mathbf{P}|_t$ are formed by restricting its poc structure:

- $\mathbf{Act}|_t$ is the induced poc structure on \mathbf{Act} ;
- $\text{Obs}|_t$ is the induced poc structure on Obs .

We are interested in the interaction between these smaller poc sets and the full model space, $\text{Cube}(\mathbf{P}|_t)$. One has two surjections

$$\begin{aligned} \text{proj}_{act} &: \text{Cube}(\mathbf{P}|_t) \rightarrow \text{Cube}(\mathbf{Act}|_t) \\ \text{proj}_{obs} &: \text{Cube}(\mathbf{P}|_t) \rightarrow \text{Cube}(\text{Obs}|_t) \end{aligned} \quad (40)$$

defined, at the level of 0-skeleta, as follows: proj_{act} sends a coherent $*$ -selection A on $\mathbf{P}|_t$ to the $*$ -selection $A \cap \mathbf{Act}$, and similarly for proj_{obs} . Hence, at the level of 0-skeleta, there is a map:

$$\text{Cube}(\mathbf{P}|_t) \rightarrow \text{Cube}(\mathbf{Act}|_t) \times \text{Cube}(\text{Obs}|_t) \quad (41)$$

In fact, Sageev-Roller duality [31] implies a much more precise statement:

Proposition III.1. *The map above (41) is a median-preserving embedding of $\text{Cube}(\mathbf{P}|_t)$ in the Cartesian product $\text{Cube}(\mathbf{Act}|_t) \times \text{Cube}(\text{Obs}|_t)$.*

Proof. See proof in appendix B-D. \square

3) *Example: discrete path with motion*: To illustrate the description of the model space provided by proposition III.1, consider an agent moving in steps of unit length along a path of integer length $L > 1$. Formally, the environment is given by $\mathbf{E} = \{0, \dots, L\}$ and the agent has the actions defined by:

$$\begin{aligned} y \in \text{wait}(x) &\Leftrightarrow \text{pos}(y) = \text{pos}(x) \\ y \in \text{fwd}(x) &\Leftrightarrow \text{pos}(y) = \min\{L, \text{pos}(x) + 1\} \\ y \in \text{bck}(x) &\Leftrightarrow \text{pos}(y) = \max\{0, \text{pos}(x) - 1\} \end{aligned} \quad (42)$$

enabling motion from any vertex $k \in \mathbf{E}$ to the adjacent $k + 1$ and $k - 1$, when they exist. We also endow the agent with sensors $a_1, \dots, a_L \in \Sigma$ realized as:

$$\langle a_k : x \rangle = 1 \Leftrightarrow \text{pos}(x) < k \quad (43)$$

Up to symmetry, the only relations holding in the existing scheme are

$$a_1 < a_2 < \dots < a_L, \quad (44)$$

the 'starfish' relations for \mathbf{Act} :

$$\text{fwd} < \text{bck}^*, \quad \text{bck} < \text{wait}^*, \quad \text{wait} < \text{fwd}^*, \quad (45)$$

for the actions $\{\text{fwd}, \text{bck}, \text{wait}\}$, and the two relations

$$\text{fwd} < a_1^*, \quad \text{bck} < a_L, \quad (46)$$

indicating $\text{pos}(x) = 0$ may not be reached by applying fwd , while $\text{pos}(x) = L$ may not be reached by applying bck . No other relations hold universally. Let \mathbf{P} denote the poc set structure over Σ recording these relations.




Fig. 5. Model space for a DBA placed in a discrete path, depicted together with its projections to $\text{Cube}(\mathbf{Act})$ (right) and to $\text{Cube}(\text{Obs})$ (below). This is the case $L = 5$ of example III-A3.

$\text{Cube}(\mathbf{P})$ is the result of forming the Cartesian product of a 3-pronged starfish $\text{Cube}(\mathbf{Act})$ with the path of length L obtained⁸ as $\text{Cube}(\text{Obs})$, and then removing two squares as shown in figure 5, due to the relation in (46).

B. Reactive Planning

1) *Statement of the planning problem*: In this section we consider a DBA at time $t > 0$, equipped with a snapshot $\mathbf{S}|_t$ with a derived poc graph $\Gamma|_t = \text{Dir}(\mathbf{S}|_t)$ and associated weak poc set $\mathbf{P}|_t$ (but keep in mind the notational simplifications at the end of II-G). The agent's tasks at hand are:

⁸Compare with example A-D1 and figure 11.

- (T1) Predict the immediate outcome of any $\alpha \in \mathbf{Act}|_t$ (or, more generally, of any $A \in \text{Cube}(\mathbf{Act}|_t)$);
- (T2) Given a set $T \subset \Sigma$ of target sensations to be achieved jointly, decide on a (generalized) action $A|_t$ for the agent to invoke in the next transition.

Both tasks need to be achieved based on the agent's record of the current state, $S|_t = \#\mathbf{S}|_t$, which is a coherent (though not necessarily complete) *-selection on $\mathbf{P}|_t$ obtained from the complete *-selection $O|_t$ representing the agent's raw observation of the current state at time t , by coherent projection (104). We will keep all the above notation fixed through the rest of this section.

It is crucial that we interpret these tasks in terms of the model space $\mathbf{M}|_t = \text{Cube}(\mathbf{P}|_t)$. For any subset $B \subset \Sigma$, define the set:

$$\mathfrak{h}(B) := \{V \subset \Sigma \mid V \text{ is a vertex of } \mathbf{M}|_t \text{ containing } B\} \quad (47)$$

These are known to be *precisely* the convex subsets of the 1-skeleton of $\mathbf{M}|_t$ (see appendix A-D). Thus, the agent is assigned the problem of reaching the convex set $\mathfrak{h}(T)$ from a (possibly unknown) position in the convex set $\mathfrak{h}(S|_t)$.

2) *Signal Propagation over a Snapshot*: The purpose of $\Gamma|_t$ is to serve as an inference tool for the agent. Recall that $\Gamma|_t$ is formed from the weight structure of the snapshot $\mathbf{S}|_t$, which, in turn, is a result of updating the weights on $\mathbf{S}|_{t-1}$ with the *raw* observation $O|_t$. The last step of the update is the 'loading' of $\Gamma|_t$ with the current state $S|_t = \text{coh}(O|_t)$ of the agent.

Definition III.2. Let $B \subset \Sigma$. Denote by $[\Gamma|_t, B]$ the weighted graph obtained from $\Gamma|_t$ by attaching the Boolean weight $\langle \mathbb{1}_B : v \rangle$ to each vertex $v \in \Sigma$, and refer to it as $\Gamma|_t$ *being loaded with B*. \square

Definition III.3. A *propagation algorithm along $\Gamma|_t$* is any algorithm which, for any coherent load $B \subset \Sigma$ and any $T \subseteq \Sigma$ accepts $[\Gamma|_t, B]$ and T as input and produces as its output the loaded graph $[\Gamma|_t, R]$ where $a \in R$ if and only if:

- 1) there is a directed path in $\Gamma|_t$ from $B \cup T$ to a , or –
- 2) there is no directed path in $\Gamma|_t$ from a into T^* .

The set $R \subset \Sigma$ is said to be the *result of propagating the signal T along $[\Gamma|_t, B]$* . \square

Applying the convexity theory of the model space $\mathbf{M}|_t$ — specifically corollary B.16 and proposition B.4 — we find the following applications for propagation:

Lemma III.4 (Implementing the State Update). *For any propagation algorithm, propagating the signal $O|_t$ along $[\Gamma|_t, \emptyset]$ produces $S|_t = \text{coh}(O|_t)$.* \square

Lemma III.5 (Reasoning in Snapshots). *Let $T \subset \Sigma$ be any set. For any propagation algorithm, propagating the signal T along $[\Gamma|_t, S|_t]$ produces the projection in $\mathbf{M}|_t$ of the current state $\mathfrak{h}(S|_t)$ onto the reduced target $\mathfrak{h}(\text{coh}(T)) \subset \mathbf{M}|_t$.*

The first lemma explains how to implement the snapshot update, given a propagation algorithm:

Algorithm 1 A simple implementation of propagation of a signal T over a poc graph Γ loaded with B , based on depth-first search.

```

function MAIN( $\Gamma, B, T$ )      ▷ Propagating  $T$  over  $[\Gamma, B]$ 
  visited  $\leftarrow \emptyset$       ▷ A global variable
   $U \leftarrow \text{CLOSURE}(\Gamma, T)$ 
  return  $(B \cup U) \setminus U^*$ 
end function

function CLOSURE( $\Gamma, T$ )      ▷ Forward closure of  $T$  in  $\Gamma$ 
  for all  $a \in T$  do
    EXPLORE( $\Gamma, a$ )
  end for
  return visited
end function

procedure EXPLORE( $\Gamma, v$ )      ▷ Recursive step
  visited  $\leftarrow \text{visited} \cup \{v\}$ 
  for all  $w \in \text{CHILDREN}(\Gamma, v) \setminus \text{visited}$  do
    EXPLORE( $\Gamma, w$ )
  end for
end procedure

function CHILDREN( $\Gamma, v$ )      ▷ Children of  $v$  in  $\Gamma$ 
  return  $\{w \in \Sigma \mid vw \in \Gamma\}$ 
end function

```

- 1) use the raw observation $O|_t$ and $\mathbf{S}|_{t-1}$ to recalculate the edge weights for $\mathbf{S}|_t$;
- 2) compute the derived graph $\Gamma|_t$;
- 3) propagate the signal $O|_t$ over $[\Gamma|_t, \emptyset]$ to compute $S|_t = \text{coh}(O|_t)$.

The second lemma is the key tool for turning a propagation algorithm into the planning algorithms we discuss in the rest of this section.

In practice, one can implement propagation using a variant of depth-first search (DFS) on $\Gamma|_t$ [63], while maintaining an expanding record of vertices visited — see algorithm 1. This algorithm clearly has time complexity that is at most quadratic in the number of sensors, and we conclude:

Corollary III.6 (Quadratic Snapshot Maintenance). *Both the time and space complexity of updating the snapshot $\mathbf{S}|_{t-1}$ with an observation $O|_t$ to form $\mathbf{S}|_t$ are at most quadratic in $|\Sigma|$.* \square

A far more efficient implementation is possible provided sufficient parallel processing power, by realizing propagation directly on $\Gamma|_t$ in a distributed fashion, using corollary B.16: given $[\Gamma|_t, S|_t]$ and a target T one first follows all directed paths in $\Gamma|_t$ emanating from $(S|_t \cup T)$ loading the traversed vertices with 1, and then follows all *reverse* paths emanating from T^* and loads their vertices with zeros. Implementing this algorithm in practice is problematic for large $|\Sigma|$ in view of the high plasticity of the graph $\Gamma|_t$ and the potentially prohibitive requirement for the DBA to maintain up to $O(|\Sigma|)$ processes, all active at the same time. Despite its current impracticality, such an implementation seems evocative of the notion of neuronal networks. We discuss this tentative connection in section IV-B.

3) *Algorithm: Computing the consequences of an action:* Planning of any kind requires an ability to sense the context of an action. This ability may be imparted to the agent by introducing sensors of the form

$$\langle \alpha \wedge s : \varphi \rangle_t = \langle \alpha : \varphi \rangle_t \cdot \#_s \mathbf{S}|_{t-1} \quad (48)$$

where α is an action and $s \in \Sigma$ is any sensor. The idea behind constructing $\alpha \wedge s$ in this way is for snapshots to be able to detect implications of the form "invoking α when s is on leads to s' " simply as directed paths from $\alpha \wedge s$ to s' in the derived graph. From a formal point of view, allowing this kind of sensor requires the observability of the snapshot in \mathbf{X} , that is: in order for the values of sensors in Σ to be allowed as input to (possibly other) sensors in Σ , it is necessary by our formalism for \mathbf{X} to carry information regarding these values in its state.

The problem of constructing a judicious process of enriching the sensorium with an effective collection of introspective sensors is set aside for future research. Instead, in the present model and simulations used to illustrate these ideas for the purposes of this paper we have committed to a sensorium containing an over-abundance of such sensors:

- **Position Sensors.** We assume the environment \mathbf{E} is given as the union of a collection \mathcal{U} of its subsets, the agent being given a state sensor $\text{loc}[U]$ for each $U \in \mathcal{U}$ satisfying $\langle \text{loc}[U] : x \rangle = \langle \mathbb{1}_U : \text{pos}(x) \rangle$.
- **Actions.** A collection of actions (in the form of 1-sensors) is provided.
- **Contextualized actions.** For each $U \in \mathcal{U}$ and $\alpha \in \mathbf{Act}$ the agent is given the sensors $\alpha \wedge \text{loc}[U]$ and $\alpha \wedge \text{loc}[U]^*$.

Under these assumptions, the following result yields a mechanism allowing the agent to 'hallucinate' the broadest consequences of an action within the context of its current model space $\mathbf{M}|_t$:

Corollary III.7 (Computing the Consequences of an Action). *For any generalized action $A \subset \mathbf{Act}$, the result of applying A in the transition from time $t > 0$ to time $t + 1$ is the result of propagating the collection $\{\alpha \wedge \text{loc}[U]\}_{\text{loc}[U] \in \mathbf{S}|_t, \alpha \in A}$ along $[\Gamma|_t, \mathbf{S}|_t]$.* \square

Thus, propagation provides a provably correct and computationally efficient mechanism for predicting the immediate outcomes of an action, provided a sensorium of the above form and a snapshot faithfully recording the nesting relations among the sensors.

Combined with the results of section II demonstrating that learning the correct relations within a fixed sensorium with a high degree of fidelity is possible even for an agent performing a random walk, the last corollary suggests that effective (and efficient) planning and closed loop control — bundled together with life-long learning features — are entirely feasible for DBAs carrying a snapshot architecture. We discuss both problems in the following paragraphs.

4) *Example: discrete path with motion, revisited:* To illustrate the above, we continue example III-A3. Recalling $\mathbf{E} = \{0, \dots, L\}$ we see that the sensors a_k defined in (43)

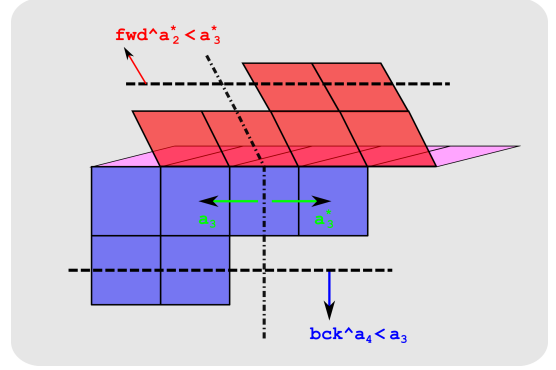


Fig. 6. Model space for an agent on a discrete path, with two added contextual action sensors.

may be rewritten as:

$$a_k = \text{loc}[U_k], \quad U_k = \{i \in \mathbf{E} \mid 0 \leq i < k\} \quad (49)$$

Thus, for example, adjoining the two sensors $\text{fwd} \wedge a_2^*$ and $\text{bck} \wedge a_4$ to Σ implies the relations

$$\text{fwd} \wedge a_2^* < a_3^*, \quad \text{bck} \wedge a_4 < a_3, \quad (50)$$

whose effect on $\text{Cube}(\mathbf{P})$, once they are learned by the agent, is shown in figure 6.

Further expanding Σ to include all the sensors

$$\begin{aligned} \text{fwd} \wedge a_k^*, \quad k = 1, \dots, L-1 \\ \text{bck} \wedge a_k, \quad k = 2, \dots, L \end{aligned} \quad (51)$$

turns $\text{Cube}(\mathbf{P})$ into the complex illustrated in figure 7. As shown in the figure, the order structure on \mathbf{P} encodes both large-scale geometry (the agent may use propagation to conclude "in order to reach $\mathfrak{h}(a_5^*)$, I need to reach $\mathfrak{h}(a_2^*)$ "), and the actions required to negotiate this geometry ("I know that $\text{fwd} \wedge a_1^*$ implies a_2^* , and I am currently in $\mathfrak{h}(a_1^*)$ ").




Fig. 7. Model space for an agent on a discrete path, enriched with a full complement of contextualized action sensors (51), and illustrating the geometry underlying planning by propagation III-B5.

5) *Algorithm: Greedy Reactive Planning (GRP) of Motion Towards a Specified Target:* The ability to compute the immediate consequences of any available action and the convexity theory of $\mathbf{M}|_t$ underlie the following greedy algorithm used to decide on an action to be taken for the purpose of achieving a *long-term* goal:

- 1) Given a set T of target sensations, propagating T over $[\Gamma|_t, S|_t]$ yields a list R of sensations characterizing the projection of the region $\mathfrak{h}(S|_t)$ representing current state in $\mathbf{M}|_t$ to the desired region $\mathfrak{h}(T)$.
- 2) Each of the elements of R may be considered as a subgoal, and a generalized action guaranteed to achieve as many of these subgoals as possible may be selected based on the corollary III.7; any ties are broken arbitrarily.
- 3) Once an action is invoked, the same target T is presented to the agent for an additional iteration of this procedure, until completion.

By lemma III.5, this algorithm is directly analogous to motion planning in the Euclidean plane in the absence of obstacles: the agent selects an action which, to the best of its knowledge, best approximates the greedy path towards the closest point of the indicated target. The next section will consider the kinds of problems arising in the presence of obstacles in the model space and some early numerical study undertaken to explore overcoming some of these problems.

C. Obstructions to Greedy Reactive Planning.

Where do obstructions to GRP in $\mathbf{M}|_t$ come from? Recall that every transition $x \times y \in X \times X$ capable of occurring in the given experiment determines a complete $*$ -selection A on Σ , by our observation model, via:

$$a \in A \iff \begin{cases} y \in \rho(a) & \text{if } a \text{ is a state sensor} \\ x \times y \in \rho(a) & \text{if } a \text{ is a transition sensor} \end{cases} \quad (52)$$

Thus, although $\mathbf{M}|_t$ does provide a universal model space for *any* realization of the weak poc set structure $\mathbf{P}|_t$, the agent is only capable of witnessing $*$ -selections of the above form, no matter the choice of action. This observation motivates the following definition:

Definition III.8 (Punctured Model Space). By the *punctured model space at time t* we mean the sub-complex $\mathbf{M}^\times|_t$ of $\mathbf{M}|_t$ induced⁹ by the set of vertices of $\mathbf{M}|_t$ of the form (52) (compare with the discussion in appendix A-E3).

Thus, in addition to the possibility that an agent will have a false implication on record (causing some sensory equivalence classes to be deemed incoherent until they are sufficiently sampled), it is also possible that $\mathbf{M}|_t$ contains obstacles to GRP in the form of vertices in $\mathbf{M}|_t \setminus \mathbf{M}^\times|_t \neq \emptyset$. In fact, the presence of obstacles of this kind is *guaranteed* by the main result of [30] — also reviewed in appendix A-E4, Theorem A.33 — at least in cases when \mathbf{E} does not have the homotopy

⁹Recall that a sub-complex L of a cell complex K is induced by a set of vertices $V \subset K^0$, if L contains every cell of K all of whose vertices belong to V .

type of a point and the covering \mathcal{U} of \mathbf{E} by location place fields satisfies the richness requirements placed on it by that theorem. We consider two of examples of this kind.

1) *Example: A Punctured Grid:* Consider the example of an agent navigating an $N \times N$ square grid (realized as a subset G_N of the integer grid) and equipped with a collection of position sensors identical to that of II-E2(d) and II-F2(d). Denoting $\text{pos}(x) := \xi \times \eta \in \mathbb{Z} \times \mathbb{Z}$ we have the sensors

$$\langle \mathbf{x}_i : x \rangle = 1 \iff \xi < i, \quad \langle \mathbf{y}_i : x \rangle = 1 \iff \eta < i, \quad (53)$$

for $i \in \{1, \dots, N-1\}$. This time, however, suppose that one interior vertex v_0 of the grid has been removed, so that $\mathbf{E} = G_N \setminus \{v_0\}$. As in the above simulations, we assume the agent is equipped with actions labelled `up`, `down`, `left` and `right` whose effect at each vertex is to move to the appropriate adjacent vertex of the integer grid if that vertex belongs to \mathbf{E} , and to remain in place otherwise. Suppose, for simplicity, that the snapshot structure for this agent is empirical.

For N sufficiently large, the statistical nature of the learning algorithm will cause the agent to learn the same weak poc set structure as in the case of v_0 being present: implications of the form

$$\text{up} \wedge \mathbf{y}_i^* < \mathbf{y}_{i+1}^*, \quad \text{up} \wedge \mathbf{x}_i < \mathbf{x}_i, \quad \mathbf{x}_i < \mathbf{x}_{i+1} \quad (54)$$

and their respective variations will be learned upon sufficient exposure, giving rise to the *same poc set structure as the one representing the complete grid G_N* . One could view this as a manifestation of the fact that our model spaces are always contractible (corollary A.31).

As a consequence, the agent is bound to attempt moving to the unavailable vertex v_0 at any time t when its position, p^t , is adjacent to v_0 and v_0 belongs to a shortest path in G_N joining p^t with the prescribed target. In such a situation, the agent is guaranteed to attempt motion in the direction of v_0 and fail (remain in place). Moreover, after sufficiently many such attempts the agent is bound to unlearn the implication responsible for this particular choice of action; this will have an overall negative impact on planning.

2) *Example: Agent on a Circular Rail:* A subtly different example is that of II-E2(b) and II-F2(b). Here, motion along a circular rail is modeled by setting \mathbf{E} to be the set of integers modulo N , with two available actions `fwd` and `bck` corresponding to the operations of adding and subtracting a unit, respectively (all arithmetic relating to the environment in this example is done modulo N). Position sensors have the form $\text{loc}[U_i]$ where $U_i = \{i-1, i, i+1\}$.

For simplicity consider a situation with N big and even, and assume the agent has complete knowledge of the correct poc set structure, which is the one generated by the relations (appendix A-A2):

$$\text{loc}[U_i] < \text{loc}[U_j]^* \iff \text{dist}(i, j) > 2, \quad (55)$$

where $\text{dist}(i, j)$ denotes the distance (modulo N) between the positions i and j , as well as:

$$\begin{aligned} \text{fwd} \wedge \text{loc}[U_i] &< \text{loc}[U_{i+1}], \\ \text{bck} \wedge \text{loc}[U_i] &< \text{loc}[U_{i-1}] \end{aligned} \quad (56)$$

for all $i \in \{0, \dots, N - 1\}$. Without loss of generality, the current state x of the system satisfies $\text{pos}(x) = 0$.

Let the specified target be $T = \{U_p\}$, where $p \in \{0, \dots, N - 1\}$ is sufficiently far from the origin (the current position of the agent) to accommodate a pair U_i, U_j such that:

- 1) $U_i \cup U_j$ does not intersect $U_p \cup U_0$;
- 2) $U_i \cup U_j$ separates p from 0 (on the circle).

Thus both the current state and the target region satisfy the constraints $\text{loc}[U_i]^*$ and $\text{loc}[U_j]^*$, which implies that *any* geodesic in the model space joining the current model state with the model target set passes through $\mathfrak{h}(\text{loc}[U_i]^*, \text{loc}[U_j]^*)$, yet it is clearly impossible to guarantee these constraints by *any* of the available actions.

It is, never the less, possible to extend this sensorium in a way that enables the effective learning of a target, as the numerical studies below demonstrate, by introducing into the environment a graded signal whose strength encodes a measure of distance to the target, while endowing the agent with sensors responding to the gradient of this signal.

3) *Closing the Loop: Excitation-Driven Navigation:* From the preceding examples it is clear that additional sensing — most probably involving information regarding transitions — is absolutely necessary for overcoming the geometric and topological obstructions to the GRP algorithm: while the GRP algorithm may be considered as providing a reasonable reference dynamics for reactive planning, one must consider possible means for *replanning in the face of failure*. We conjecture that the notion of a snapshot is sufficiently simple and agile for such purposes:

- **Control of learning thresholds.** At this stage of our research, no attempt is being made to control the learning thresholds τ_{ab} ; it seems plausible, however, that having a high level of "frustration" cause the lowering of a relevant threshold may be used as a tool for locating exceptions to poc relations.
- **Introduction of new sensors.** A principled mechanism is required for the introduction of *combinations of existing sensors*, such as Boolean functions thereof, or delayed conjunctions such as (48), to *become additional members of the sensorium*. In particular, such a mechanism must be capable of responding to exceptions, or failures of GRP, as we had already discussed above.

The need for self-adjustment in the sensorium opens the door to the introduction of auxiliary internal mechanisms of evaluating the position of the agent in the environment, or, more generally, the state of the system consisting of the pairing of the agent with the environment. For example, the settings described in the preceding paragraph suggests the introduction of an internal variable evaluating success (and failure) of invoking a planned action, while the need for closed-loop control suggests simple local control mechanisms based on an internally-defined 'navigation function' [15]. Many other ideas of internal behavior modulation ranging from varying notions of novelty, surprise and dependability [58], [64], [22], [4] all the way to a multivariate model of human neuro-modulation mechanisms [59] become relevant in this context.




Fig. 8. Mean deviation from target for empirical (blue) and a discounted (red) agents (20 sensors each), as a function of time in four different settings, averaging over 50 runs.

Excitation-Driven Agents. In the absence of tools for reactive replanning (our current situation), we have chosen to study a simplified notion of target, allowing us to close the control loop with a "motion" command generated with the aim to guarantee an immediate decrease in the value of an internal excitation signal.

The simplest instance of such a controller, applied to the navigation setting, seems to be the following. In addition to a sensorium of the form described above in III-B3 and the examples that followed, assume that the agent possesses a pair of sensors better and worse, responding to the decrease and increase, respectively, in a fixed measure of distance to a target point in the environment \mathbf{E} , over a single transition (think of this as a radically simplified sense of smell).

Starting out as a lazy random-walking agent¹⁰, the agent uses the algorithm of III-B5 at each step to obtain an action resulting with better (target specification $T = \{\text{better}\}$) as its first priority. In the case of failure to produce such an action, the agent attempts to guarantee worse*, periodically attempting a random action so as not to get stuck in place (upon having figured out that $\text{wait} < \text{worse}^*$). Figure 8 presents a comparison between the mean behaviors of four different such agents simulated in the same settings as those analyzed in section II.

It is important to stress that, following the discussion in III-A3 and the formal results of the preceding paragraphs, the guarantee of the agents in figure 8(a)-(c) finding their targets *given sufficient knowledge* of the correct poc structure on Σ is absolute. To see this, it suffices to verify for the true poc set structure on Σ that any position other than the target has associated with it a location sensor $a = \text{loc}[U]$ and an action α such that every state x with $\langle a : x \rangle = 1$ satisfies the requirement that $\alpha(x)$ is closer to the target than x is.

¹⁰An action `wait` is available now, to let the agent stay put when it has found the target.

The only remaining question is whether or not the control strategy we propose guarantees sufficient exposure for the agent to recover the poc relations necessary for it to capture the target. Figure 8 offers some evidence in favor of an affirmative answer, based on simulations.

IV. DISCUSSION

A. Topological Mapping and Planning

Mapping methods vary from probabilistic representations [10], [12], [65] to hybrid representations in which precisely mapped contractible local metric patches are integrated into a global map through ‘gluing instructions’ recorded in an annotated graph that is learned by the agent in the course of its travels ([35], [36], [37], [38]). The latter are known as *topological SLAM* (simultaneous localization and mapping) methods.

In the presence of significant obstacles in large scale environments, the lightweight ‘topological skelton’ of the environment recorded by a topological SLAM method provides valuable information on loop closure, which would otherwise be costly to obtain using global metric mapping techniques. This observation led to an extensive effort [66], [34] to produce a general notion of a *topological map*, based on the notion of *spacial semantic hierarchy* (SSH), that would allow for motion planning at varying scales. Other ways of leveraging planar topological descriptors to obtain simple and efficient data structures supporting localization and motion planning in simply connected planar domains have been explored as well, e.g. [67], [28] and [68].

In view of our stated need for guaranteed universality properties, the strong dependence of SLAM methods on Euclidean geometry made it necessary for us to adopt a significantly more abstract approach to localization arising from the point of view of sensor fields. A precursor of our approach is a family of SLAM algorithms, known as RatSLAM, employing a neural network to simulate the function of place/pose cells in rats, e.g. [13], [24], [48]. In a broad generalization of the neural computational engine underlying RatSLAM, [39] tests the hypothesis that place cells with sensor fields in *any* sufficiently dense configuration should make it possible [for a rat] not only to localize well, but also to accurately represent the topology of the environment by estimating the nerve¹¹ of the system of place fields from observations of near-synchrony in place-cell firing. It is shown through simulations that recovering the topological invariants of a connected planar arena, as well as some approximation of its geometry, is possible with a sufficiently dense network of convex place fields. Further evidence in support of this idea is the recently introduced [40] method for localization in an urban canyon, as well as the ELM architecture [23], using nerve-like information (nesting among convex polygons in the plane) as a means of encoding spatio-temporal context in an ‘episodic memory’ for a planar agent.

A significant drawback of nerve estimation is that, by definition, computing nerve of a covering requires expo-

ponential space in the cardinality of the covering¹². At the same time, restricting attention to pairwise intersections only, and focusing on those of them that are empty turns out to guarantee a universal model space for each combinatorial type of this ‘reduced nerve’, by Sageev-Roller duality [56], [31]. By construction, snapshots are estimators of this reduced nerve, turned into a computational tool for navigating the corresponding model space after converting all relations of the form ‘ $A \cap B$ negligible/unimportant’ into relations of the form ‘ A implies B ’.

The idea of leveraging nesting relations among geometric descriptors of events for the purpose of properly representing context is a well-recognized and widely used tool in the literature, for example: nesting of planar domains is used in [34], as well as the more recent [23], and a notion of nesting for actions is used in [4]. What is new about the snapshot architecture is its application of this principle to the *entire* sensorium, including the set of available actions.

Our numerical experiments with closed loop control suggest that a snapshot-driven agent with sufficient actuation and sufficiently rich sensorium is capable of learning a good approximation of the gradient field of a (discrete) Morse function despite been given no prior semantic information and starting out with random ‘motor babble’ for its navigation strategy.

In fact, the snapshot architecture is flexible enough to trivialize the task of introducing discrete variants of complex motivational systems [59] based on introspective sensing of signals quantifying (a) internally available resources (e.g. battery charge), (b) repulsion or fear of punishment, (c) attraction or anticipation of a reward signal (e.g., in the sense of navigation functions [15] or in the broad sense of Reinforcement Learning [51]), and even (d) frustration over the failure of a plan [69] and (e) measures of innovation [58], [64].

We expect systems such as (a)-(c) to contribute to an agent’s planning capabilities from the point of view of the variety of tasks they would enable. Even more significantly, we expect (d) and (e) to contribute to the agent’s ability to improve the quality of the topological representation encoded in its snapshot. Namely, (d) could be used to facilitate chunking by serving as a signal driving the creation of new conditional sensors detecting inconsistent states of the model space, while (e) could drive the learning of useful complex actions, as has already been proposed for many other architectures [25], [4], [22], [26], resulting in improving the connectivity of the model space. Endowing the snapshot architecture with these capabilities is the most immediate goal for follow-up research.

B. Connectionist Architectures

From the earliest days of the field, even extremely simple neural networks with a very small number of neurons have demonstrated the ability to perform complex learning and control tasks [70], [71], including complex symbolic structures such as context-sensitive grammars [72], [73], [74], and

¹¹See *nerve of a covering* in [33], section 3.3.

¹²As one must keep track of the intersections of all possible sub-collections of the covering

complex hierarchical schemata whose structural features vastly outnumber the physical components the actual network [75]. Even simple feed-forward networks have been shown to be capable of exercising ‘intelligent’ control through stigmergy [76], the depositing of tokens in the environment. Furthermore, the merging of reinforcement learning methods with the computational power of ‘deep learning’ networks [77], [78] has yielded architectures capable of matching and outperforming humans in complex tasks such as learning and playing video games [79] based only on the raw RGB output provided by the game console.

At the same time, the power of connectionist architectures comes with very limited formal understanding of how the internal representations they maintain encode symbolic representations in terms of problem spaces.

Though the connection demonstrated in this work falls short solving this problem (due to the over-simplification of the connectionist structure and due to realization of plasticity in the network by a non-neural controller), direct analogies with current studies of plasticity and neural coding give hope that a suitable generalization of the snapshot architecture could yield *both* strength of performance *and* provable guarantees of symbolic reasoning processes.

Of the classes of neural networks that are well understood, most relevant for our purposes is that of competitive attractor networks, whose strong stability properties [80] were applied to modeling the navigation mechanisms in rats, also inspiring the RatSLAM mapper [13].

Expanding on these results, [81], [49] sparked the discussion of the structure of the set of binary codewords corresponding to stable activity patterns of threshold-linear neural networks. This line of inquiry was picked up in [82], producing a combinatorial characterization of the possible codes in terms of the network’s organization; and in [50], initiating a rigorous study of the way codes vary as the synaptic weights are perturbed while subject to structural constraints, exposing interesting connections with topological invariants associated with these constraints.

The analogy with our work is straightforward. The process of obtaining the coherent projection of a binary observation by propagating it through the derived graph of a snapshot is completely analogous to the process of propagating a signal through a threshold-linear neural network and waiting for the network state to stabilize at a code word — especially when taking into account the excitatory nature of relations of the form $a \leq b$ and inhibitory role of relations of the form $a \leq b^*$ under propagation.

Chasing this analogy, it could be worthwhile investigating the degree to which the collection of codewords of a threshold-linear network conforms to the strict demands of median geometry (represented by coherent snapshot states), to establish a rigorous formal connection (if it exists) between the two architectures. A more general study of which neural learning methods [83] could be transferred into a snapshot architecture may, on one hand, expand the range of applications for snapshots, while, on the other hand, provide some existing architectures with a rigorous symbolic interpretation.

V. CONCLUSION

We introduce a new computationally efficient architecture intended to endow a generic discrete binary agent with the capacity to build over time an actionable model of its operations within a completely unknown dynamic environment, \mathbf{E} . The proposed architecture has a dual nature. On one hand, the agent maintains an evolving data structure, — the snapshot $\mathbf{S}|_t$ — of size quadratic in the number of sensors, encoding a planning mechanism based on propagation of excitation and inhibition signals through the highly plastic directed network $\text{Dir}(\mathbf{S}|_t)$, and is, thus, in a very crude sense, a connectionist learning and control architecture. On the other hand, the rather specific ordering properties of networks arising in this way (the derived ‘weak poc set structure’ $\mathbf{P}|_t$) also characterize any such network as encoding a system of ‘half-spaces’ in a geometric internal representation $\mathbf{M}|_t$ that is *just rich enough* to account for all sensory equivalence classes provided to the agent by its sensorium Σ . This duality affords the re-interpretation of snapshots as encoding a high-level symbolic representation of the problem space (i.e., the state space \mathbf{X} and the transition system induced on it by the agents interaction with its dynamic environment), through a mathematical formalism that rigorously supports symbolic planning with the efficiency and economy of a connectionist architecture.

Clearly, our current snapshot architectures (section II) lack certain key features found in existing AGI architectures. First and foremost among these is a mechanism for enriching the agent’s sensorium with sensors representing general Boolean predicates (or, even better, some limited LTL predicates), composed of the original atomic sensations. Of course, the problem lies not in proposing intuitively attractive approaches (there are many) but rather doing so in a principled, economical way that maintains the present combination of analytical and computational tractability. These ‘compound’ sensors are required for facilitating chunking and the learning of useful motor primitives. Another required feature is a capacity for symbolic abstraction (relating problem spaces via substitutions). While the duality theory of weak poc sets and their model spaces (appendix A-E) affords a rigorous formulation of enriched predicates and consequent symbolic abstraction, it is not yet clear how to engineer an enlarged snapshot-like architecture realizing such meta-extensions.

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.

REFERENCES

- [1] L. R. Squire, “Memory systems of the brain: a brief history and current perspective,” *Neurobiology of learning and memory*, vol. 82, no. 3, pp. 171–177, 2004.
- [2] E. Tulving, “Episodic memory: From mind to brain,” *Annual Review of Psychology*, vol. 53, no. 1, p. 125, 2002.
- [3] B. J. Baars, *A cognitive theory of consciousness*. Cambridge University Press, 1993.

- [4] J. Mugan and B. Kuipers, "Autonomous learning of high-level states and actions in continuous environments," *Autonomous Mental Development, IEEE Transactions on*, vol. 4, no. 1, pp. 70–86, 2012.
- [5] P. S. Thomas and A. G. Barto, "Motor primitive discovery," in *ICDL-EPIROB*, 2012, pp. 1–8.
- [6] D. Kraft, R. Detry, N. Pugeault, E. Baseski, F. Guerin, J. H. Piater, and N. Kruger, "Development of object and grasping knowledge by robot exploration," *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 4, pp. 368–383, 2010.
- [7] S. Ivaldi, S. Nguyen, N. Lyubova, A. Droniou, V. Padois, D. Filliat, P.-Y. Oudeyer, and O. Sigaud, "Object learning through active exploration," *Autonomous Mental Development, IEEE Transactions on*, vol. 6, no. 1, March 2014.
- [8] S. Tellex, P. Thaker, J. Joseph, and N. Roy, "Learning perceptually grounded word meanings from unaligned parallel data," *Machine Learning*, vol. 94, no. 2, pp. 151–167, 2014.
- [9] S. Chernova and A. L. Thomaz, "Robot learning from human teachers," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 8, no. 3, pp. 1–121, 2014.
- [10] 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.
- [11] B. Kuipers, "An intellectual history of the spatial semantic hierarchy," *Robotics and cognitive approaches to spatial mapping*, vol. 38, p. 243264, 2008.
- [12] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [13] M. J. Milford, G. F. Wyeth, and D. Prasser, "Ratslam: a hippocampal model for simultaneous localization and mapping," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 1. IEEE, 2004, pp. 403–408.
- [14] 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.
- [15] E. Rimon and D. E. Koditschek, "Exact robot navigation using artificial potential functions," *Robotics and Automation, IEEE Transactions on*, vol. 8, no. 5, pp. 501–518, 1992.
- [16] M. Erdmann, "On the topology of discrete strategies," *The International Journal of Robotics Research*, 2009.
- [17] A. Newell and H. A. Simon, *Human Problem Solving*. New Jersey: Prentice-Hall, Inc., 1972.
- [18] P. Langley, J. E. Laird, and S. Rogers, "Cognitive architectures: Research issues and challenges," *Cognitive Systems Research*, vol. 10, no. 2, pp. 141–160, 2009.
- [19] S. Franklin, T. Madl, S. D' Mello, and J. Snider, "Lida: A systems-level architecture for cognition, emotion, and learning," *Autonomous Mental Development, IEEE Transactions on*, vol. 6, no. 1, March 2014.
- [20] E. Nivel and K. R. Thórisson, "Towards a programming paradigm for control systems with high levels of existential autonomy," in *Artificial General Intelligence*. Springer, 2013, pp. 78–87.
- [21] A. Newell, "Reasoning, problem solving and decision processes : the problem space as a fundamental category," *Carnegie Mellon University Research Showcase CMU*, 1979. [Online]. Available: <http://repository.cmu.edu/cgi/viewcontent.cgi?article=2574&context=compisci>
- [22] L. C. Cobo, C. L. Isbell, and A. L. Thomaz, "Object focused q-learning for autonomous agents," in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 1061–1068.
- [23] D. Stachowicz and G. Kruijff, "Episodic-like memory for cognitive robots," *Autonomous Mental Development, IEEE Transactions on*, vol. 4, no. 1, pp. 1–16, 2012.
- [24] M. Milford and G. Wyeth, "Persistent navigation and mapping using a biologically inspired slam system," *The International Journal of Robotics Research*, vol. 29, no. 9, pp. 1131–1153, 2010.
- [25] J. Schmidhuber, "Ultimate cognition à la gödel," *Cognitive Computation*, vol. 1, no. 2, pp. 177–193, 2009.
- [26] G. Martius, R. Der, and N. Ay, "Information driven self-organization of complex robotic behaviors," *PloS one*, vol. 8, no. 5, p. e63400, 2013.
- [27] B. Donald and J. Jennings, "Sensor interpretation and task-directed planning using perceptual equivalence classes," in *Robotics and Automation, 1991. Proceedings., 1991 IEEE International Conference on*. IEEE, 1991, pp. 190–197.
- [28] B. Tovar, A. Yerzhova, J. O'Kane, and S. M. LaValle, "Information spaces for mobile robots," in *Robot Motion and Control, 2005. RoMoCo'05. Proceedings of the Fifth International Workshop on*. IEEE, 2005, pp. 11–20.
- [29] R. L. Rivest and R. E. Schapire, "Diversity-based inference of finite automata," *Journal of the ACM (JACM)*, vol. 41, no. 3, pp. 555–589, 1994.
- [30] D. Guralnik and D. Koditschek, "Toward a memory model for autonomous topological mapping and navigation: The case of binary sensors and discrete actions," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, 2012, pp. 936–945.
- [31] M. Roller, "Poc sets, median algebras and group actions," *University of Southampton, Faculty of Math. stud., preprint series*, 1998.
- [32] D. T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, ser. CBMS Regional Conference Series in Mathematics. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012, vol. 117.
- [33] A. Hatcher, *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [34] E. Remolina and B. Kuipers, "Towards a general theory of topological maps," *Artificial Intelligence*, vol. 152, no. 1, pp. 47–104, 2004.
- [35] N. Tomatis, I. Nourbakhsh, and R. Siegwart, "Hybrid simultaneous localization and map building: a natural integration of topological and metric," *Robotics and Autonomous systems*, vol. 44, no. 1, pp. 3–14, 2003.
- [36] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli, "Local metrical and global topological maps in the hybrid spatial semantic hierarchy," in *Robotics and Automation, 2004. Proceedings. ICRA04. 2004 IEEE International Conference on*, vol. 5. IEEE, 2004, pp. 4845–4851.
- [37] A. Ranganathan and F. Dellaert, "Online probabilistic topological mapping," *The International Journal of Robotics Research*, p. 0278364910393287, 2011.
- [38] S. Tully, G. Kantor, and H. Choset, "A unified bayesian framework for global localization and slam in hybrid metric/topological maps," *The International Journal of Robotics Research*, 2012.
- [39] C. Curto and V. Itskov, "Cell groups reveal structure of stimulus space," *PLoS Computational Biology*, vol. 4, no. 10, 2008.
- [40] J. Derenick, A. Speranzon, and R. Ghrist, "Homological sensing for mobile robot localization," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 572–579.
- [41] 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>
- [42] 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>
- [43] V. Chepoi, "Graphs of some CAT(0) complexes," *Adv. in Appl. Math.*, vol. 24, no. 2, pp. 125–179, 2000. [Online]. Available: <http://dx.doi.org/10.1006/aama.1999.0677>
- [44] V. Pratt, "Modeling concurrency with geometry," in *Proceedings of the 18th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM, 1991, pp. 311–322.
- [45] R. Ghrist and V. Peterson, "The geometry and topology of reconfiguration," *Advances in applied mathematics*, vol. 38, no. 3, pp. 302–323, 2007.
- [46] E. Klavins, R. Ghrist, and D. Lipsky, "A grammatical approach to self-organizing robotic systems," *IEEE Trans. Automat. Control*, vol. 51, no. 6, pp. 949–962, 2006. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2006.876950>
- [47] J. McCarthy and P. Hayes, *Some philosophical problems from the standpoint of artificial intelligence*. Stanford University USA, 1968.
- [48] M. J. Milford and A. Jacobson, "Brain-inspired sensor fusion for navigating robots," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 2906–2913.
- [49] R. H. Hahnloser, H. S. Seung, and J.-J. Slotine, "Permitted and forbidden sets in symmetric threshold-linear networks," *Neural Computation*, vol. 15, no. 3, pp. 621–638, 2003.
- [50] C. Curto, A. Degeratu, and V. Itskov, "Flexible memory networks," *Bulletin of mathematical biology*, vol. 74, no. 3, pp. 590–614, 2012.
- [51] A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 1998.
- [52] A. G. Barto and S. Mahadevan, "Recent advances in hierarchical reinforcement learning," *Discrete Event Dynamic Systems*, vol. 13, no. 4, pp. 341–379, 2003.
- [53] H. P. Helgason, "General attention mechanism for artificial intelligence systems," Ph.D. dissertation, Reykjavik University, 2013. [Online]. Available: <http://skemman.is/en/item/view/1946/16163>

- [54] B. R. Steunebrink, J. Koutník, K. R. Thórisson, E. Nivel, and J. Schmidhuber, “Resource-bounded machines are motivated to be effective, efficient, and curious,” in *Artificial General Intelligence*. Springer, 2013, pp. 119–129.
- [55] M. C. Mozer and J. Bachrach, “Slug: A connectionist architecture for inferring the structure of finite-state environments,” in *Connectionist Approaches to Language Learning*. Springer, 1991, pp. 35–56.
- [56] M. Sageev, “Ends of groups pairs and non-positively curved cube complexes,” *Proc. London Math. Soc.*, vol. 3, no. 71, pp. 586–617, 1995.
- [57] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, ser. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Berlin: Springer-Verlag, 1999, vol. 319.
- [58] A. Barto, M. Mirolli, and G. Baldassarre, “Novelty or surprise?” *Frontiers in psychology*, vol. 4, 2013.
- [59] B. Cox and J. L. Krichmar, “Neuromodulation as a robot controller,” *Robotics & Automation Magazine, IEEE*, vol. 16, no. 3, pp. 72–80, 2009.
- [60] L. Lovász, “Random walks on graphs: A survey,” *Combinatorics, Paul erdos is eighty*, vol. 2, no. 1, pp. 1–46, 1993.
- [61] M. M. Deza and M. Laurent, *Geometry of cuts and metrics*, ser. Algorithms and Combinatorics. Springer-Verlag, Berlin, 1997, vol. 15. [Online]. Available: <http://dx.doi.org/10.1007/978-3-642-04295-9>
- [62] R. S. Sutton, D. Precup, and S. Singh, “Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning,” *Artificial intelligence*, vol. 112, no. 1, pp. 181–211, 1999.
- [63] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, “Introduction to algorithms mit press,” *Cambridge, MA*, 2003.
- [64] J. Schmidhuber, “Formal theory of creativity, fun, and intrinsic motivation (1990–2010),” *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 3, pp. 230–247, 2010.
- [65] M. Montemerlo, S. Thrun, and B. Siciliano, *FastSLAM: A scalable method for the simultaneous localization and mapping problem in robotics*. Springer, 2007, vol. 27.
- [66] B. Kuipers, “The spatial semantic hierarchy,” *Artificial intelligence*, vol. 119, no. 1, pp. 191–233, 2000.
- [67] A. Yershova, B. Tovar, R. Ghrist, and S. M. LaValle, “Bitbots: Simple robots solving complex tasks,” in *PROCEEDINGS OF THE NATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE*, vol. 20, no. 3. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005, p. 1336.
- [68] E. Mair, M. Augustine, B. Jäger, A. Stelzer, C. Brand, D. Burschka, and M. Suppa, “A biologically inspired navigation concept based on the landmark-tree map for efficient long-distance robot navigation,” *Advanced Robotics*, vol. 28, no. 5, pp. 289–302, 2014.
- [69] S. Revzen, B. Ilhan, and D. Koditschek, “Dynamical trajectory replanning for uncertain environments,” in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 3476–3483.
- [70] R. S. Sutton and A. G. Barto, “An adaptive network that constructs and uses an internal model of its world,” *Cognition and Brain Theory*, vol. 4, no. 3, pp. 217–246, 1981.
- [71] A. G. Barto, R. S. Sutton, and C. W. Anderson, “Neuronlike adaptive elements that can solve difficult learning control problems,” *Systems, Man and Cybernetics, IEEE Transactions on*, no. 5, pp. 834–846, 1983.
- [72] J. L. Elman, “Finding structure in time,” *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.
- [73] D. Servan-Schreiber, A. Cleeremans, and J. L. McClelland, “Encoding sequential structure in simple recurrent networks,” DTIC Document, Tech. Rep., 1989.
- [74] J. L. Elman, “Distributed representations, simple recurrent networks, and grammatical structure,” *Machine learning*, vol. 7, no. 2-3, pp. 195–225, 1991.
- [75] M. Botvinick and D. C. Plaut, “Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action,” *Psychological review*, vol. 111, no. 2, p. 395, 2004.
- [76] J. R. Chung and Y. Choe, “Emergence of memory in reactive agents equipped with environmental markers,” *Autonomous Mental Development, IEEE Transactions on*, vol. 3, no. 3, pp. 257–271, 2011.
- [77] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [78] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, “Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations,” in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 609–616.
- [79] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” *CoRR*, vol. abs/1312.5602, 2013. [Online]. Available: <http://arxiv.org/abs/1312.5602>
- [80] M. A. Cohen and S. Grossberg, “Absolute stability of global pattern formation and parallel memory storage by competitive neural networks,” *Systems, Man and Cybernetics, IEEE Transactions on*, no. 5, pp. 815–826, 1983.
- [81] X. Xie, R. H. Hahnloser, and H. S. Seung, “Selectively grouping neurons in recurrent networks of lateral inhibition,” *Neural computation*, vol. 14, no. 11, pp. 2627–2646, 2002.
- [82] C. Curto, A. Degeratu, and V. Itskov, “Encoding binary neural codes in networks of threshold-linear neurons,” *Neural computation*, vol. 25, no. 11, pp. 2858–2903, 2013.
- [83] G. E. Hinton, “Connectionist learning procedures,” *Artificial intelligence*, vol. 40, no. 1, pp. 185–234, 1989.
- [84] J. R. Isbell, “Median algebra,” *Trans. Amer. Math. Soc.*, vol. 260, no. 2, pp. 319–362, 1980. [Online]. Available: <http://dx.doi.org/10.2307/1998007>
- [85] B. Nica, “Cubulating spaces with walls,” *Algebr. Geom. Topol.*, vol. 4, pp. 297–309 (electronic), 2004. [Online]. Available: <http://dx.doi.org/10.2140/agt.2004.4.297>
- [86] M. H. Stone, “The theory of representation for boolean algebras,” *Transactions of the American Mathematical Society*, vol. 40, no. 1, pp. pp. 37–111, 1936. [Online]. Available: <http://www.jstor.org/stable/1989664>
- [87] M. L. J. van de Vel, *Theory of convex structures*, ser. North-Holland Mathematical Library. Amsterdam: North-Holland Publishing Co., 1993, vol. 50.
- [88] M. Sholander, “Medians and betweenness,” *Proc. Amer. Math. Soc.*, vol. 5, pp. 801–807, 1954.
- [89] ———, “Medians, lattices, and trees,” *Proc. Amer. Math. Soc.*, vol. 5, pp. 808–812, 1954.
- [90] J. R. Isbell, “Median algebra,” *Trans. Amer. Math. Soc.*, vol. 260, no. 2, pp. 319–362, 1980. [Online]. Available: <http://dx.doi.org/10.2307/1998007>
- [91] H.-J. Bandelt and J. Hedlíková, “Median algebras,” *Discrete Math.*, vol. 45, no. 1, pp. 1–30, 1983. [Online]. Available: [http://dx.doi.org/10.1016/0012-365X\(83\)90173-5](http://dx.doi.org/10.1016/0012-365X(83)90173-5)
- [92] I. Chatterji, C. Druţu, and F. Haglund, “Kazhdan and Haagerup properties from the median viewpoint,” *Adv. Math.*, vol. 225, no. 2, pp. 882–921, 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.aim.2010.03.012>
- [93] D. Kozlov, *Combinatorial algebraic topology*, ser. Algorithms and Computation in Mathematics. Springer, Berlin, 2008, vol. 21. [Online]. Available: <http://dx.doi.org/10.1007/978-3-540-71962-5>

APPENDIX A A PRIMER ON SAGEEV-ROLLER DUALITY

The duality between poc sets and median algebras, going back to [84], was thoroughly studied by Martin Roller in [31], in a very successful attempt of pushing the envelope on a theory of actions of discrete groups on simply connected non-positively curved cubical complexes – henceforth referred to as *cubings* – pioneered by Michah Sageev in [56] and by Victor Chepoi [43], who characterized such complexes in terms of the convexity theory of their 1-dimensional skeleta.

This appendix provides a detailed review of the elements of this theory supporting the memory architecture proposed in this paper. This overview of the preliminary material is meant to extend the initial discussion provided in [30] as well as to illustrate it with examples, intended as bridges to our current application. In the end, the duality theory of poc sets will be called upon to provide the necessary formal guarantees that the proposed memory and control architectures actually do their job. We will mainly rely on [31] as a source of theoretical results, though sometimes it will be easier to use results from the elegant exposition in [85].

A. Basic Notions

We introduced the extension of the duality theory of poc sets to so-called *weak poc sets* in [30] out of the necessity to maintain poc sets as dynamical data structures.

Definition A.1 (Weak Poc Set). A partially-ordered set (P, \leq) endowed with an order-reversing fixpoint-free involution $a \mapsto a^*$ and having a minimum element $\mathbf{0} \in P$ is called a *weak poc set*. Note that $\mathbf{0}^*$ is a maximum for P . Thus, for all $a, b \in P$ one has:

- $0 \leq a = a^{**}$ and $a^* \neq a$;
- $a \leq b \Rightarrow b^* \leq a^*$.

An element $a \in P$ is said to be *negligible* if $a \leq a^*$, and *ubiquitous* if a^* is negligible. A *poc set* is a weak poc set in which $\mathbf{0}$ is the only negligible element. An element that is neither negligible nor ubiquitous is said to be *proper*. \square

Weak poc sets form a category:

Definition A.2 (Poc Morphism). A function $f : P \rightarrow Q$ between two weak poc sets is a *poc morphism*, if f is order-preserving, $*$ -equivariant and $f(\mathbf{0}) = \mathbf{0}$. The set of all poc morphisms as above will be denoted $\text{Hom}(P, Q)$. \square

1) *The Minimal Poc Set*: The set $\{0, 1\}$ with the relations $0 < 1$ and $1 = 0^*$ is a poc set, and it is denoted by $\mathbf{2}$. Clearly, there is only one poc morphism of $\mathbf{2}$ into any weak poc set P , but then there may be many poc morphisms of a weak poc set P onto $\mathbf{2}$.

2) *Generators and Relations*: A weak poc set $P = \langle S | R \rangle$ may be specified using a set S of generators and a set of relations R of the form $a < b$ or $a^* < b$ or $a < b^*$ for $a, b \in S^{13}$.

Formally, P is constructed as follows. Assume that the symbol $\mathbf{0}$ is not contained in S . First, set $S_{\pm} := (\{\mathbf{0}\} \sqcup S) \times \{+, -\}$ and define $(s, +)^* = (s, -)$ and $(s, -)^* = (s, +)$. Thus, S_{\pm} has a fix-point free involution $*$ defined on it. For simplicity, for each $s \in \{0\} \cup S$ we identify $(s, +)$ with s .

The relation set R is required to be a subset of $S_{\pm} \times S_{\pm}$. We then define an extension R_{poc} of R to be the intersection of all relations $W \subseteq S_{\pm} \times S_{\pm}$ that are reflexive, transitive and, in addition, satisfy the following:

- $(\mathbf{0}, a) \in W$ holds for all $a \in S_{\pm}$;
- For all $a, b \in S_{\pm}$, if $(a, b) \in W$ then $(b^*, a^*) \in W$.

We set P to be the quotient of S_{\pm} by the equivalence

$$x \sim y \Leftrightarrow (x, y) \in R_{poc} \wedge (y, x) \in R_{poc}$$

with the induced partial ordering

$$[x] \leq [y] \Leftrightarrow (x, y) \in R_{poc}.$$

For example, the notation

$$\langle a, b, c | a < c, b < c \rangle$$

stands for the poc set

$$P = \{\mathbf{0}, \mathbf{0}^*, a, b, c, a^*, b^*, c^*\}$$

having the order relations

$$\begin{aligned} \mathbf{0} < a < c < \mathbf{0}^*, & \quad \mathbf{0} < b < c < \mathbf{0}^* \\ \mathbf{0} < c^* < a^* < \mathbf{0}^*, & \quad \mathbf{0} < c^* < b^* < \mathbf{0}^* \end{aligned}$$

as well as the ones derived from these by transitivity. Thus, generators and relations provide a compact way of representing a (weak) poc set explicitly.

As another example, consider the poc sets

$$P = \langle a, b | a < b \rangle, \quad Q = \langle a, b | a^* < b \rangle$$

The partial assignment $f : P \rightarrow Q$ satisfying

$$f(a) = a^*, \quad f(b) = b$$

has one and only one extension to a poc morphism of P into Q .

3) *σ -Algebras as poc sets*: Let \mathcal{B} be a σ -algebra on a non-empty (possibly infinite) set \mathbf{X} . Then $(\mathcal{B}, \subseteq, F \mapsto \mathbf{X} \setminus F)$ is a poc set. In particular, the power set of \mathbf{X} , denoted $\mathbf{2}^{\mathbf{X}}$, obtains the structure of a poc set in this way. It is standard to identify $\mathbf{2}^{\mathbf{X}}$ with the space of functions $f : \mathbf{X} \rightarrow \mathbf{2}$: any such f will be identified with the subset $f^{-1}(1) \in \mathbf{2}^{\mathbf{X}}$. Note that the intersection and symmetric difference operators translate under this identification into pointwise multiplication and addition modulo 2, respectively. Recalling our notation (5) for the evaluation of functions, it will be convenient to extend it as follows:

$$\langle f : x \rangle = f(x), \quad \langle f^* : x \rangle = 1 + \langle f : x \rangle \quad (57)$$

The poc set structure on $\mathbf{2}^{\mathbf{X}}$ may then be written in functional form via

$$f \leq g \text{ in } \mathbf{2}^{\mathbf{X}} \Leftrightarrow \forall x \in \mathbf{X} \langle f : x \rangle \leq \langle g : x \rangle \text{ in } \mathbf{2} \Leftrightarrow fg = f, \quad (58)$$

where $f, g \in \mathbf{2}^{\mathbf{X}}$ are arbitrary elements.

Definition A.3 (realization). Let P be a weak poc set and let \mathbf{X} be a non-empty set. A *realization of P in \mathbf{X}* is a poc morphism $f : P \rightarrow \mathbf{2}^{\mathbf{X}}$. \square

A realization $f : P \rightarrow \mathbf{2}^{\mathbf{X}}$ provides a consistent way of regarding each $a \in P$ as a binary question over \mathbf{X} , so that the set of all $x \in \mathbf{X}$ with $\langle f(a) : x \rangle = 1$ is the set of all points where the question is answered affirmatively.

4) *Canonical Quotient*: Every weak poc set has a canonical poc set quotient \hat{P} obtained as the quotient of P by the equivalence relation

$$a \sim b \Leftrightarrow \begin{cases} a = b \text{ or} \\ a, b \text{ are both negligible, or} \\ a, b \text{ are both ubiquitous} \end{cases} \quad (59)$$

and endowed with the obvious ordering and involution.

Definition A.4. Let P be a weak poc set and let \hat{P} denote its canonical poc quotient. For every $a \in P$, we denote the equivalence class of a in \hat{P} with \hat{a} . The poc morphism $a \mapsto \hat{a}$ will be denoted by π_P . \square

The main characteristic of \hat{P} is the following elementary lemma:

¹³One may also use weak inequalities (\leq) to specify relations in R .

Lemma A.5. Let P be a weak poc set. Then any poc morphism $f : P \rightarrow Q$ of P into a poc set Q factors through π_P , that is: there exists one and only one poc morphism $\hat{f} : \hat{P} \rightarrow Q$ satisfying $f = \hat{f} \circ \pi_P$.

For example, seeking a realization of a weak poc set structure in any space \mathbf{X} is only possible after identifying all negligible elements with $\mathbf{0}$, because the only subset F of \mathbf{X} satisfying $F \subseteq \mathbf{X} \setminus F$ is $F = \emptyset$.

B. The Dual Graph of a Poc Set

The duality theory for poc sets is an extension of Stone Duality [86]. At the base of the construction are binary selections:

Definition A.6 (*-selections). Let P be a weak poc set. A subset $A \subset P$ is a **-selection on P* if no $a \in A$ has $a^* \in A$. A *-selection A on P is *complete* if for any $a \in P$ either $a \in A$ or $a^* \in A$. The set of all complete *-selections on P is denoted by $S(P)^0$. \square

The following is a metric on $S(P)^0$:

$$\Delta(A, B) = |A \setminus B| = |B \setminus A| = \frac{1}{2} |A \Delta B| \quad (60)$$

Indeed, fixing $A_0 \in S(P)^0$, an explicit isometry of $(S(P)^0, \Delta)$ onto $\mathbf{2}^{A_0}$ endowed with the Hamming distance is constructed by sending $A \in S(P)^0$ to the [indicator function of the] set $A_0 \setminus A$. Thus, $S(P)^0$ may be thought of as simply being the vertex set, or 0-skeleton, of the $\frac{|P|}{2}$ -dimensional standard unit cube, viewed as a combinatorial cubical complex, – we denote this complex by $S(P)$ – where a cubical face Q of $S(P)$ of dimension d corresponds to a maximal subset of $S(P)^0$ with diameter d as calculated in the metric $\Delta(\cdot, \cdot)$.

1) *Construction of the Dual Graph:* Some vertices of $S(P)$ cannot be witnessed in any realization of P :

Definition A.7 (Coherence). A pair of elements $a, b \in P$ is said to be *incoherent* if $a \leq b^*$. A subset A of a poc set P is said to be *coherent* if it contains no incoherent pair. \square

Definition A.8 (Dual graph, dual Cubing). Given a (finite) poc set P , one defines:

- (a) **The dual cubing of P** , denoted $\text{Cube}(P)$, is the sub-complex of $S(P)$ induced by the set of coherent vertices;
- (b) **The dual median algebra of P** , denoted P° , is the 0-skeleton of $\text{Cube}(P)$;
- (c) **The dual graph of P** , denoted $\text{Dual}(P)$, is the 1-skeleton of $\text{Cube}(P)$. \square

2) *A Simple Example:* To illustrate the definition, consider the poc set P whose Hasse diagram is given in figure 9(a). Given by generations and relations, P takes the form:

$$P = \langle a, b, c \mid a < c^*, b < c^* \rangle \quad (61)$$

A good way of thinking about a poc set is to pretend that it is realized in a space \mathbf{X} , so that our P is a collection of three binary questions (a, b and c) about \mathbf{X} , augmented with the complementary questions (a^*, b^* and c^*), together with an additional record of known implication relations between




Fig. 9. A simple poc set P on three generators (a) and the resulting cube complex (b), obtained by deleting all incoherent vertices from the cube $S(P)$ (c) – see example A-B2.

them ($a < c^*$ and $b < c^*$). In the absence of any implications on record, an observer endowed with P will model the space \mathbf{X} as the full 3-cube $S(P)$, where the proper elements of P correspond to the co-dimension one faces of the cube – fig. 9(c). On the other hand, knowledge of the above relations renders some of the vertices of $S(P)$ redundant, resulting in a reduction in the number of binary states necessary for modeling \mathbf{X} using the same three questions – fig. 9(b).

Remark A.9. We have chosen the term *coherent subset* over Roller's *filter-base* to better fit the context of our planning/sensing problem.

Remark A.10. The standard identification of $\mathbf{2}^P$ with the space of $\{0, 1\}$ -valued functions on P also puts the set P° of coherent vertices of $S(P)$ in one-to-one correspondence with the set $\text{Hom}(P, \mathbf{2})$ of poc morphisms of P onto the trivial poc set $\mathbf{2}$.

C. Poc Set Duals are Median Graphs

Graphs of the form $\text{Dual}(P)$ (for a weak poc set P) are completely characterized. We recall:

Definition A.11 (hop-distance, intervals). Let $G = (V, E)$ be a connected simple graph and let $u, v \in V$. The *hop distance* $d_G(u, v)$ is defined to be the minimum length of an edge-path in G joining u with v . The interval $I(u, v)$ is defined to be the set of all vertices $w \in V$ satisfying the equality $d_G(u, v) = d_G(u, w) + d_G(w, v)$. \square

A fundamental fact about the dual $\text{Dual}(P)$ of a poc set P is a quick corollary of the results in [85], section 4:

Lemma A.12. Let $G = \text{Dual}(P)$ for a finite poc set P . Then the metric Δ coincides with the hop metric on G .

An important and well-studied class of graphs is:

Definition A.13 (median graphs [43], [87]). A connected simple graph $G = (V, E)$ is said to be a *median graph*, if the set $I(u, v) \cap I(v, w) \cap I(u, w)$ contains exactly one vertex for each $u, v, w \in V$. This vertex is the *median* of the triple (u, v, w) and denoted by $\text{med}(u, v, w)$ – see figure 10. \square

Median graphs are a special subfamily of a family of ternary algebras, called *median algebras*, [88], [89], [90], [91]. Some modern generalizations and applications may be found in [92].

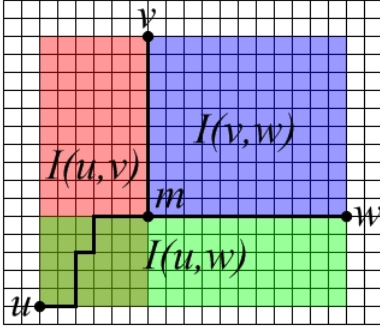


Fig. 10. Computing a median in the integer grid (see example A-D2 for a poc-set presentation).

Another way of stating the preceding lemma (again, see [85], section 4, where these results are derived in a much more general setting):

Theorem A.14. *The dual $G = \text{Dual}(P)$ of a finite poc set P is a finite median graph, with the median calculated according to the formula*

$$\text{med}(u, v, w) = (u \cap v) \cup (u \cap w) \cup (v \cap w) \quad (62)$$

for all $u, v, w \in P^\circ$.

In other words, the median of three coherent $*$ -selections is determined by a majority vote on the values of their observations.

D. Convexity

Median graphs have a very strong convexity theory. We recall:

Definition A.15 (Convexity). Let $G = (V, E)$ be a graph. A subset $K \subseteq V$ is said to be *convex* if $I(u, v) \subseteq K$ holds whenever $u, v \in K$. \square

Definition A.16 (Half-Spaces). Let $G = (V, E)$ be a graph. A subset $H \subseteq V$ is said to be a *half-space* of G , if both H and $V \setminus H$ are convex subsets of G . \square

For example, lemma A.12 can be used to quickly deduce

Lemma A.17. *Let P be a poc set. Then the half-spaces of $\text{Dual}(P)$ are precisely the subsets of P° of the form*

$$\mathfrak{h}(a) := \{u \in P^\circ \mid a \in u\} \quad (63)$$

where a ranges over P . In particular, subsets of P° of the form

$$\mathfrak{h}(K) := \{u \in P^\circ \mid K \subseteq u\} = \bigcap_{a \in K} \mathfrak{h}(a) \quad (64)$$

are convex in $\text{Dual}(P)$.

Remark A.18. Note also that $\mathfrak{h}(a^*) = P^\circ \setminus \mathfrak{h}(a)$ for all $a \in P$.

Much more can be said in general:

Theorem A.19 (Properties of Median Graphs, [31], section 2). *Let $G = (V, E)$ be a finite median graph. Then:*

- 1) *Every convex set is an intersection of halfspaces;*
- 2) *Any family of pairwise-intersecting convex sets has a common vertex (1-dimensional Helly property);*
- 3) *For any convex subset $K \subset V$, the subgraph of G induced by K is a median graph;*
- 4) *For any convex subset $K \subset V$ and any vertex $v \notin K$ there exists a unique vertex $\text{proj}_K v \in K$ at minimum hop distance from v .*
- 5) *For any convex subset $K \subset V$, the closest-point projection proj_{K^\bullet} is a median-preserving, distance non-increasing map of G onto the subgraph of G induced by K .*

Any graph $G = (V, E)$ generates a poc set $\text{Poc}(G)$: we let the underlying set of $\text{Poc}(G)$ be the set of all half-spaces of G , then we order it by inclusion and set $H^* = V \setminus H$ for the complementation operator. Of course, some graphs (e.g. any odd cycle) will generate the trivial poc set in this way, but not so for median graphs:

Theorem A.20 ([31], proposition 5.9). *Let G be a finite median graph. Then G is canonically isomorphic to $\text{Dual}(\text{Poc}(G))$ via the median-preserving map which sends each vertex v to the collection of halfspaces of G which contain v .*

An important conclusion (special case of proposition 6.11 in [31]) is:

Corollary A.21. *For any finite poc set P , the map $a \mapsto \mathfrak{h}(a)$ is a poc-isomorphism of P onto $\text{Poc}(\text{Dual}(P))$. In other words, the poc-set P may be reconstructed from its dual.*

From a practical standpoint, these two results offer an approach to understanding the geometry of $\text{Cube}(P)$ in terms of the order structure of P , which is the purpose of this and the following sections. As an aside, let us mention also that these results are best viewed together in categorical terms, as part of a restatement of the duality between the category of finite poc sets (with poc morphisms) and the category of finite median algebras (with median-preserving maps) — see appendix A-E below.

We must consider the possible relations (if any) among elements $a, b \in P$. Those are:

$$a \leq b, \quad a^* \leq b, \quad a^* \leq b^*, \quad a \leq b^* \quad (65)$$

It is easy to see that a pair of distinct *proper* elements will never satisfy two of the above conditions at the same time, as $\text{Cube}(P)$ provides us with a realization of P inside 2^{P° — after all, the last theorem tells us that:

$$a \leq b \Leftrightarrow \mathfrak{h}(a) \subseteq \mathfrak{h}(b) \quad (66)$$

Definition A.22 (Nesting and Transversality). Suppose a, b are proper elements of a weak poc set P . We say that they *cross* ($a \pitchfork b$), if none of (65) hold. Otherwise, we say they are *nested* ($a \parallel b$). A subset A of P is said to be *nested* if all its elements are pairwise nested, and *transverse* if its elements cross pairwise. \square




Fig. 11. Dual graphs for two arrangements of sensors along the real line (see A-D1): threshold sensors encoding a path (a), and beacon sensors encoding a ‘starfish’ (b).

Thus, the half-spaces of $\text{Dual}(P)$ are nothing more than the restriction to $\text{Dual}(P)$ of the half-spaces of the cube $S(P)^1$, with two of them nesting if and only if the corresponding elements of P are nested, that is, if and only if exactly one of the following holds:

$$\begin{aligned} \mathfrak{h}(a) \cap \mathfrak{h}(b) &= \emptyset, & \mathfrak{h}(a^*) \cap \mathfrak{h}(b) &= \emptyset, \\ \mathfrak{h}(a^*) \cap \mathfrak{h}(b^*) &= \emptyset, & \mathfrak{h}(a) \cap \mathfrak{h}(b^*) &= \emptyset \end{aligned} \quad (67)$$

We conclude that the more relations are on record in the order structure of P the fewer transverse sets there are to be found there. In other words, nesting relations are an obstruction to high-dimensional cubes in $\text{Cube}(P)$: each additional relation in P implies fewer faces of the original cube $S(P)$ survive the culling of incoherent vertices used for obtaining $\text{Cube}(P)$. At one extreme one finds $\text{Cube}(P) = S(P)$ when P itself (up to removing improper elements) is transverse. At the other extreme (exercise for the reader) $\text{Cube}(P)$ forms a tree if and only if P is nested.

1) *Example: the path of length N* : Consider an idealized point-robot situated on the interval $\mathbf{E} = [0, 1]$ and capable of moving about in this environment. Suppose the robot is endowed with binary sensors a_1, \dots, a_N , each responding to the robot’s position – denoted for now by x – according to the rule, say, that a_k turns on whenever $x < x_k := k/N$. It would be reasonable for us to wish for the robot to eventually be able to realize that a_k turning on implies a_{k+1} turning on, for all $k < n$. Forming the poc set

$$P = \langle a_1, \dots, a_N \mid a_k < a_{k+1}, k = 1, \dots, N-1 \rangle \quad (68)$$

it is easy to verify that $\text{Dual}(P)$ is the N -path – the path with $N+1$ vertices and N edges – whose vertices are all of the form

$$v_k = \{\mathbf{0}^*\} \cup \{a_j^*\}_{j>k} \cup \{a_i\}_{i \geq k}, \quad 0 \leq k \leq N \quad (69)$$

Please note that the choice of the points $x_k \in \mathbf{E}$ is immaterial – only their ordering should matter for the correctness of $\text{Cube}(P)$ as a discretized model of the ‘environment’ \mathbf{E} of our robot.

At the same time, imagine that the sensors a_k corresponded to ‘beacons’, with a_k turning on if and only if $|x - \frac{k}{N}| < \frac{1}{3N}$. Then a poc set description of the form

$$P = \langle a_1, \dots, a_N \mid a_k < a_j^*, 1 \leq k < j < N \rangle \quad (70)$$

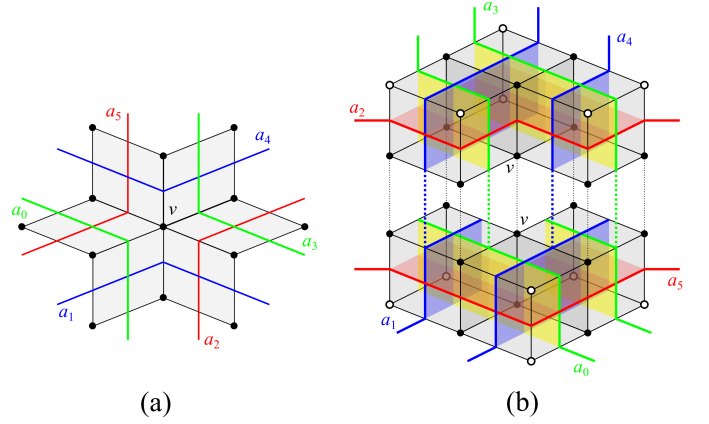


Fig. 12. Cubical models for example A-D3 with relations (a) $a_i < a_{i+x}^*$ where $x \in \{2, 3, 4\}$ and addition is modulo 6, and (b) only the relations $a_i < a_{i+3}^*$ are present. Black vertices are those coherent in for both poc set structures. Vertices painted white are coherent vertices for agent #2 that are incoherent for agent #1. The vertex v corresponds to the shared coherent $*$ -selection $\{a_0^*, \dots, a_5^*\}$.

would be more appropriate, indicating that the sensations a_k are mutually exclusive. The resulting dual would still have $N+1$ vertices and N edges, the vertices being:

$$\begin{aligned} v'_k &= \{\mathbf{0}^*, a_k\} \cup \{a_j^*\}_{j \neq k} \text{ for } 1 \leq k \leq N \\ v'_0 &= \{\mathbf{0}^*\} \cup \{a_1^*, \dots, a_N^*\} \end{aligned} \quad (71)$$

In both cases the dual graph is a tree (a *path* and a *starfish*), and it is hard to ignore the difference in the quality of its representation of the underlying space – see figure 11

2) *Example: direct sums of poc sets*: The easiest way to join two poc sets together is to form their direct sum:

Definition A.23. Let P and Q be poc sets. Their *direct sum* $P \vee Q$ is defined to be the quotient of their external disjoint union $P \sqcup Q$ by the identification $\mathbf{0}_P = \mathbf{0}_Q$ and $\mathbf{0}^*_P = \mathbf{0}^*_Q$, endowed with the following:

- $a < b$ in $P \vee Q$ iff $a, b \in P$ and $a < b$ or $a, b \in Q$ and $a < b$;
- $b = a^*$ iff both $a, b \in P$ and $b = a^*$ or $a, b \in Q$ and $b = a^*$.

(We abuse notation by identifying each element of $P \cup Q$ with the equivalence class in $P \vee Q$ of its natural representative in $P \sqcup Q$) \square

It is easy to verify, then, that

$$\text{Cube}(P \vee Q) \equiv \text{Cube}(P) \times \text{Cube}(Q) \quad (72)$$

where the isomorphism is that of cubical complexes. Intuitively, any proper elements $a \in P$ and $b \in Q$ satisfy $a \pitchfork b$, resulting in every cube in $\text{Cube}(P)$ and every cube in $\text{Cube}(Q)$ to form a product cube in $\text{Cube}(P \vee Q)$. For example, the grid in figure 10 may be thought of as the product of an N -path with an M -path (for the appropriate values of M and N) – hence the dual of the direct sum of two poc sets of the type discussed in the preceding example A-D1.

3) *Example: a cycle of length 6*: Imagine an agent – call it #1 – living on the unit circle $\mathbf{E} = \mathbb{S}^1$. We mark six vertices, spread uniformly along the circle, with the digits $\{0, \dots, 5\}$. Suppose that agent #1 is capable, for each position it occupies on \mathbf{E} , of asking any of the binary questions

- A_j : Am I positioned at arc length $< \frac{\pi}{3}$ from position j along \mathbf{E} ?

Agent #2 asks a slightly different set of questions:

- B_j : Am I positioned at arc length $< \frac{\pi}{2}$ from position j along \mathbf{E} ?

The questions available to either agent have sufficient resolution to pinpoint the agent's position wherever it is, but we claim that the collection $\{A_j\}_{j=0}^5$ is, in a sense, more efficient than $\{B_j\}_{j=0}^5$. Let $P = \{\mathbf{0}, \mathbf{0}^*\} \cup \{a_i, a_i^*\}_{i=0}^5$, where the a_i are symbols to represent the sensations corresponding to A_i for agent #1 and to B_i for agent #2. We compare the resulting embeddings $\rho_i : P \hookrightarrow \mathbf{2}^V$ defined by

$$\begin{aligned} \rho_1(a_j) &= A_j, \quad \rho_1(a_j^*) = V \setminus A_j, \\ \rho_2(a_j) &= B_j, \quad \rho_2(a_j^*) = V \setminus B_j, \end{aligned}$$

and with $\rho_i(\mathbf{0}) = \emptyset$ and $\rho_i(\mathbf{0}^*) = V$, of course. We observe that both representations of P in $\mathbf{2}^V$ form injective poc morphisms of P into $\mathbf{2}^V$ if P is given all relations of the form $a_i < a_{i+3}^*$ (addition modulo 6), yet only agent #1 can afford to also add the relations $a_i < a_{i+2}^*$ and $a_i < a_{i+4}^*$ to the record without losing the property of ρ_1 being a poc morphism. The difference between the duals is significant – see figure 12 – clearly showing the advantage of the compact and simple world map that agent #1 could deduce over the cumbersome monstrosity agent #2 must deal with. Note how the complex (a) in the figure may be obtained from (b) through deleting the vertices painted white – those are precisely the vertices of (b) forming incoherent families for the poc set structure represented in (a).

Two aspects of this example are noteworthy:

- 1) The less nested poc set of the two example poc set structures is capable of accommodating both agents, thus giving us the means for comparing them.
- 2) With none of the agents having direct access to the realization maps, they should be looking for efficient means of *evolving their maps* in an adaptive fashion so as to produce a good enough symbolic approximation of the ground truth.

E. Cubings and the Duality Theory of Weak Poc Sets

1) Sageev-Roller Duality from the Categorical Viewpoint:

In the finite case, the duality theory of poc sets has a very clean formulation in category-theoretical terms. For a quick review of the basic notions of Category Theory we refer the reader to [93] chapter 4, while here we will stick to the specific categories of interest:

- \mathbf{Poc}_f , the category of finite poc sets¹⁴, where each $P, Q \in \mathbf{Poc}_f$ have assigned to them the set $\text{Hom}(P, Q)$ of poc morphisms from P to Q ;

¹⁴One could work with the full category \mathbf{Poc} of all poc sets (rather than just the finite ones) but this introduces major complications that seem unnecessary given the current application. Similarly for the case of median graphs/algebras.

- \mathbf{Med}_f , the category of finite median graphs, where each $G, H \in \mathbf{Med}_f$ are assigned the set $\text{Hom}(G, H)$ of median-preserving maps from the vertex set of G to the vertex set of H (such maps are called *median morphisms*).

What connects the two categories is the assignment of the graph $\text{Dual}(P)$ to every poc set P . The important bit here is that this assignment is not confined to the level of objects, but, rather, extends over the level of maps as well in a natural way:

Definition A.24. Let $f : P \rightarrow Q$ be a morphism of weak poc sets. The dual map $f^\circ : Q^\circ \rightarrow P^\circ$ is defined to be the pullback map $f^\circ(A) = f^{-1}(A)$. \square

It is easy to verify that $f^\circ : Q^\circ \rightarrow P^\circ$ is a median-preserving map, that is:

$$f^\circ(\text{med}(u, v, w)) = \text{med}(f^\circ(u), f^\circ(v), f^\circ(w)) \quad (73)$$

where the medians are computed in the appropriate duals. Thus, a map $f \in \text{Hom}(P, Q)$ yields a map $f^\circ \in \text{Hom}(\text{Dual}(Q), \text{Dual}(P))$. Moreover, one easily checks that this is done in a way that respects composition, that is:

$$(g \circ f)^\circ = f^\circ \circ g^\circ \quad (74)$$

whenever the composition of the poc morphisms f, g is well-defined. This notion of map between categories is called a *functor*. The above constructions (of the dual graph and the dual map), together, are known as the *Sageev-Roller duality*.

Applying the results A.20 and A.21, we conclude that the above assignments form a *complete duality*, or *co-equivalence of categories*, between \mathbf{Poc}_f and \mathbf{Med}_f . That is, there are:

- **A correspondence between \mathbf{Poc}_f and \mathbf{Med}_f at the level of objects:** $P \mapsto \text{Dual}(P)$ is a one-to-one correspondence between the collection of finite poc sets and the collection of median graphs;
- **A correspondence between \mathbf{Poc}_f and \mathbf{Med}_f at the level of maps:** $f \mapsto f^\circ$ is a composition-reversing one-to-one correspondence between poc morphisms and median morphisms.

Thus, Sageev-Roller duality is a dictionary, translating order-theoretic statements about finite poc sets into graph-theoretic statements about finite median graphs and vice-versa. Loosely speaking, the aspects of Boolean Algebra covered by poc sets may be conveniently interpreted in terms of the convex geometry of median graphs, reasoned about within this framework, and the conclusions may then be translated back into the Boolean Algebra setting for the purpose of dealing with applications. We will now proceed to survey some of the contributions of this category-theoretic point of view to our application.

2) Extending Sageev-Roller Duality to Weak Poc Sets:

Recalling the fact that the category \mathbf{Poc}_f of proper finite poc sets is too restrictive for the purposes of our current application [30], the first order of business is to verify that Sageev-Roller duality extends to weak poc sets.

The first observation regarding dual maps is a consequence of the fact that no coherent subset of a weak poc set may contain a negligible element:

Lemma A.25. *Let P be a weak poc set and let $\pi : P \rightarrow \hat{P}$ denote the canonical projection. Then $p^\circ : \hat{P}^\circ \rightarrow P^\circ$ is a median isomorphism, making $\text{Cube}(P)$ and $\text{Cube}(\hat{P})$ isomorphic cubical complexes.*

Thus, weak poc sets are indistinguishable from poc sets, as far as dual graphs are concerned: applying Sageev-Roller duality one simply obtains

Corollary A.26. *For any weak poc set P , \hat{P} is naturally isomorphic to $\text{Poc}(\text{Dual}(P))$.* \square

At the same time, weak poc sets form a more flexible class of objects. In particular, weak poc set structures are easier to represent and evolve dynamically using snapshots II.

3) *Example: Realizations:* Suppose \mathbf{X} is the state space (possibly infinite) of some system, and Σ is a collection of binary sensors sensitive to the state of the system, such as in examples A-D1 and A-D3. Since the sensors are binary we may assume that each sensor $a \in \Sigma$ comes paired with a sensor $a^* \in \Sigma$ corresponding to the negation of a . In other words, the sensorium Σ comes equipped with a fixpoint-free involution $*$ and with a realization map $\rho : \Sigma \rightarrow \mathcal{B} \subseteq 2^{\mathbf{X}}$ satisfying $\rho(a^*) = \mathbf{X} \setminus \rho(a)$ for all $a \in \Sigma$, where \mathcal{B} is a prescribed σ -algebra of measurable events in \mathbf{X} . It also costs us nothing to assume there is a special sensor $\mathbf{0} \in \Sigma$ that is never on, that is: $\rho(\mathbf{0}) = \emptyset$.

Suppose now that, having spent some time observing state transitions in \mathbf{X} , we are able to write down some implication relations among the elements of Σ . These will be recorded in the form of a partial ordering (\leq). We would like to use our *a-priori* knowledge that $\rho(a^*) = \mathbf{X} \setminus \rho(a)$ and, naturally, we would like to believe that $a \leq b$ implies $\rho(a) \subseteq \rho(b)$, in which case ρ becomes a morphism of the weak poc set $\mathbf{P} = (\Sigma, \leq, *)$ into the poc set \mathcal{B} . Assuming this is correct, what can we say?

For any observed state $x \in \mathbf{X}$ the poc set \mathcal{B} supplies us with a coherent $*$ -selection $\mathcal{B}_x = \{A \in \mathcal{B} \mid x \in A\}$. The dual $\rho^\circ : \mathcal{B}^\circ \rightarrow \mathbf{P}^\circ$ then produces a vertex $v_\rho(x) := \rho^\circ(\mathcal{B}_x)$ in $\text{Cube}(\mathbf{P})$.

Definition A.27 (Consistent Vertices). Let ρ be a realization of a poc set P . Then the vertices of the form $v_\rho(x)$ as above are called *ρ -consistent vertices*.

For example, the vertex v in figures 12 (a) and (b) is inconsistent for either realization.

It is possible that $v_\rho : \mathbf{X} \rightarrow \mathbf{P}^\circ$ is not surjective, motivating the definition:

Definition A.28 (Punctured Dual). Let ρ be a realization of a poc set P . The *punctured dual* (with respect to ρ), denoted $\text{Cube}^\times(\mathbf{P})$, or $\text{Cube}(\mathbf{P}, \rho)$, is the cubical sub-complex of $\text{Cube}(\mathbf{P})$ induced by the set of ρ -consistent vertices of \mathbf{P}° which are contained in $\text{Cube}(\mathbf{P})$.

As a corollary of this discussion we obtain:

Corollary A.29. *$\text{Cube}(\mathbf{P})$ is the smallest cubical sub-complex of $S(P)$ with the property that, for any realization ρ of \mathbf{P} , if ρ is a poc morphism, then $\text{Cube}(\mathbf{P})$ contains all ρ -consistent vertices of $S(P)$.*

In other words, any realization that is also a poc morphism gives rise to a discrete representation of \mathbf{X} in $\text{Cube}(\mathbf{P})$ via the mapping $v_\rho : \mathbf{X} \rightarrow \mathbf{P}^\circ$. The benefit of maintaining a record of the order in Σ is our ability to discard the incoherent vertices of $S(P)$ (viewed as possible states of the observed system) without the risk of losing any information about \mathbf{X} , while possibly gaining some insight into the organization of \mathbf{X} , as stated in the introduction, contribution (i).

4) *Realizations, Cubings and Topology:* We recall a definition from [56]:

Definition A.30. A *cubing* is a simply connected, non-positively curved cubical complex. \square

We point the reader to [57] for a detailed account of non-positively curved metric spaces. For the purpose of this paper it will suffice to quote a corollary of the well-known Cartan-Hadamard theorem ([57], II.4.1):

Corollary A.31. *Cubings are contractible.*

We owe the following theorem in its full generality (finite and infinite cases) to the collective efforts of Michah Sageev [56], Martin Roller [31] and Victor Chepoi [43].

Theorem A.32. *The following are equivalent for a finite simple graph G :*

- 1) G is the 1-dimensional skeleton of a cubing;
- 2) G is a median graph;
- 3) G is isomorphic to $\text{Dual}(P)$ for some poc set P ;
- 4) G is the 1-dimensional skeleton of $\text{Cube}(P)$ for some poc set P .

Further developing the discussion in the preceding paragraph, we recall one of the main theorems of [30], applied to that setting:

Theorem A.33. *Let \mathbf{X} be a topological space and let \mathbf{P} be a poc set structure on a finite set Σ with realization $\rho : \mathbf{P} \rightarrow 2^{\mathbf{X}}$. Let \mathcal{C} denote the collection of cubes in the cubical complex $\text{Cube}^\times(\mathbf{P}) = \text{Cube}(\mathbf{P}, \rho)$. For each $C \in \mathcal{C}$ let $\mathbf{X}_C = (r^\circ)^{-1}(C)$ be the set of all points in \mathbf{X} witnessing C . If ρ is a poc morphism, and every \mathbf{X}_C , $C \in \mathcal{C}$ has an open neighbourhood $\mathcal{N}_C \subset \mathbf{X}$ such that:*

- 1) each \mathcal{N}_C is contractible;
 - 2) $\{\mathbf{X}_C\}_{C \in \mathcal{C}}$ and $\{\mathcal{N}_C\}_{C \in \mathcal{C}}$ have isomorphic nerves,
- then $\text{Cube}^\times(\mathbf{P})$ is homotopy-equivalent to \mathbf{X} .

To illustrate the theorem, consider figure 12 again: note how deleting the vertex v from either discrete model of the circle results in a space with the correct homotopy type (that of the circle). On the other hand, going back to example A-D1 and figure 11, the last theorem explains the qualitative difference in representations of the interval between the two provided models: while the ‘thresholds’ model satisfies the requirements of the theorem, the ‘beacons’ model possesses a vertex – the one marked V'_0 – whose set of witnesses is not connected.

This result could be interpreted as stating a condition on the *richness* of the sensorium (Σ, ρ) : the complex $\text{Cube}(\mathbf{P})$ provides an observer with a discretized contractible model of the state space \mathbf{X} of the observed system, while $\text{Cube}^\times(\mathbf{P})$ is a more realistic model of \mathbf{X} taking the standard form of a contractible space minus a set of obstacles.

5) *Example: A 'bad' Poc Morphism:* It is not true in general that the dual of a poc morphism $f : P \rightarrow Q$ extends to a morphism of graphs. For example, consider the situation

$$P = \langle a, b, c \mid a < b, b < c \rangle, \quad Q = \langle x, y \mid x < y \rangle \quad (75)$$

and $f : P \rightarrow Q$ is defined by $f(a) = f(b) = x$ and $f(c) = y$. The duals and dual map are illustrated in figure 13.




Fig. 13. The dual of a poc morphism is not necessarily a graph morphism (details in A-E5).

The absence of a canonical choice of extension for f° to a graph morphism of $\text{Dual}(Q)$ into $\text{Dual}(P)$ hints at a solution directly involving cubings: if one were to extend the range of f° to include the 2-dimensional cube shown in the figure, it would have been possible to extend f° to a cellular map taking the edge of $\text{Dual}(Q)$ crossed by x to an appropriately chosen diagonal of that cube. More generally, it is possible to extend f° to a continuous embedding of $\text{Cube}(Q)$ into $\text{Cube}(P)$ for any poc morphism $f : P \rightarrow Q$ by applying convexity properties of the canonical piecewise-Euclidean metrics on $\text{Cube}(P)$ and $\text{Cube}(Q)$ ([57], II.2.7). Thus, although median graphs are sufficient for describing the dual graphs of poc sets, describing the *dual morphisms* requires the higher dimensional geometry of cubings.

6) *Example: Degeneration:* Recall our promise to maintain weak poc set structures on a sensorium Σ *dynamically*, updating the ordering on Σ in real time. The duality theory of poc sets provides us with a hint as to how such maintenance should be done. The learning methods of section II are motivated by the analogy between the following observations and the ideas underlying Hebbian learning:

The kind of update we expect to see in an instance of learning is captured in the following simple example.

$$\begin{aligned} P_1 &= \langle a, b, c \mid a < c, b < c \rangle, \\ P_2 &= \langle a, b, c \mid a < b < c \rangle \end{aligned}$$

The two poc set structures have the same underlying set (denote it by P) and the identity map $f = \text{id}_P : P_1 \rightarrow P_2$ is a morphism, while the inverse map $g = \text{id}_P : P_2 \rightarrow P_1$ is not. Thinking of P_1 as representing an observer yet undecided regarding the nature of nesting (if any) of the pair $\{a, b\}$ and therefore maintaining $a \pitchfork b$ in P_1 , we see poc set




Fig. 14. The dual of a degeneration is an embedding of median graphs (details in A-E6).

P_2 as representing an observer with an identical set of beliefs except for the additional relation $a < b$. Figure 14 visualizes the dual map f° .

In general, if P_1 and P_2 are poc sets with the same underlying set P and $f = \text{id}_P : P_1 \rightarrow P_2$ is a poc morphism, then the dual f° has the following properties (see e.g. [31]):

Proposition A.34. *Suppose $f : P_1 \rightarrow P_2$ is a bijective poc morphism. Then:*

- 1) f° is injective ([31], proposition 7.8);
- 2) f° extends to an injective cellular embedding of $\text{Cube}(P_2)$ in $\text{Cube}(P_1)$;
- 3) The image of $\text{Cube}(P_2)$ under this embedding is a strong deformation retract of $\text{Cube}(P_1)$.

A more complicated instance of this setting is very nicely visualized in figure 12.

APPENDIX B PROOFS OF TECHNICAL RESULTS

A. Proof of proposition II.7

We first extend the weight function $ab \mapsto w_{ab}$ to a symmetric function of $\Sigma \times \Sigma$ by setting

$$w_{aa} = w_{ab} + w_{ab^*}, \quad w_{aa^*} = 0 \quad (76)$$

for any $a \in \Sigma$ and for any $b \in \Sigma$ with $ab \in \mathbf{K}_\Sigma$. The consistency constraint (13) implies that this extension is well-defined. This allows us to extend the definition of $\omega(\bullet)$ in (16) to the whole of $\Sigma \times \Sigma$ while satisfying the following identities

$$\omega(ab) + \omega(ba) = 0, \quad \omega(aa) = 0 \quad (77)$$

for all $a, b \in \Sigma$. Note that $\omega(\bullet)$, by definition, takes directed edges and loops for its input.

Further, the consistency identity allows us to write, for any $ab \in \mathbf{K}_\Sigma$:

$$\omega(aa^*) = w_{a^*a^*} - w_{aa} \quad (78)$$

$$= w_{a^*b} + w_{a^*b^*} - w_{ab} - w_{ab^*} \quad (79)$$

$$= \omega(ab) + \omega(ab^*) \quad (80)$$

We next observe that the cocycle constraint (15) may be rewritten in the form

$$\omega(ab) + \omega(bc) = \omega(ac) \quad (81)$$

Let us verify that this identity is, in fact, an identity over all $a, b, c \in \Sigma$. Due to the symmetries of $\omega(\bullet)$ in (77) it suffices to verify (81) only in the following cases:

- 1) The pair $ab \in \mathbf{K}_\Sigma$ and $c = a$: this case is taken care of by the anti-symmetry identity of $\omega(\bullet)$.
- 2) The pair $ab \in \mathbf{K}_\Sigma$ and $c = a^*$: this is precisely the statement of (80).
- 3) The pair $ac \in \mathbf{K}_\Sigma$ and $b, c \in \{a, a^*\}$: without loss of generality, either $b = c = a$ and (81) ends up claiming that $0 + 0 = 0$, or $b = a$ and $c = a^*$ – in which case the statement turns into the trivial identity $\omega(aa^*) = \omega(aa^*)$.

(cases 1-2 correspond to exactly two of the pairs being proper; case 3 accounts for all situations when none of the pairs is proper; having exactly one of the pairs proper is impossible).

Now suppose that $p = (a_0, \dots, a_m)$ is any directed vertex path in the given poc graph Γ . Then, applying the identity (81) repeatedly we obtain:

$$\omega(a_0 a_m) = \omega(a_0 a_1) + \dots + \omega(a_{m-1} a_m) \quad (82)$$

By the assumption on Γ , all the summands on the right hand side are positive. In particular, if p were a cycle with $a_m = a_0$ we would have obtained

$$0 = \omega(a_0 a_0) = \omega(a_0 a_m) > 0 \quad (83)$$

– a contradiction. We must therefore conclude that directed cycles in Γ are impossible, as desired.

B. Proof of lemma II.15

An evolution of the trivial snapshot is an empirical snapshot, for if $\mathbf{S} = \mathbf{S}|_t$ can be written in the form $\mathbf{S} = O_k * \dots * O_1 * \text{Null}$ then defining indicators

$$c_{ab}^k = \langle \mathbb{1}_{O_k} : a \rangle \cdot \langle \mathbb{1}_{O_k} : b \rangle \quad (84)$$

– compare with (11) – one would have the following identity holding for \mathbf{S} :

$$w_{ab} = \sum_{k=1}^t c_{ab}^k \in \mathbb{Z}_{\geq 0} \quad (85)$$

Since the functions c^k satisfy the consistency constraint, so does their sum $w_{..}$. The clock requirement is satisfied, too, since for any proper pair $\{a, b\}$ one has:

$$w_a + w_{a^*} = \sum_{k=1}^t \underbrace{(c_{ab} + c_{ab^*} + c_{a^*b} + c_{a^*b^*})}_{=1} = t \quad (86)$$

– compare with equation (12).

Conversely, due to the presence of an integer clock, it suffices to show that any empirical snapshot \mathbf{S} can be written in the form $\mathbf{S} = O * \mathbf{T}$ where $O = \# \cdot \mathbf{S}$ and T is either trivial or empirical. Let \mathbf{T} be the weighted graph obtained from \mathbf{S} by performing the following operations:

- 1) Subtract one unit from w_{ab} for every proper pair $a, b \in \Sigma$ satisfying $\{a, b\} \subseteq O$.
- 2) For any $a \in \Sigma$, set $\#_a$ for \mathbf{T} to 1 iff $w_a > w_{a^*}$ in \mathbf{T} .

The set $\#_{\mathbf{T}}$ is a complete $*$ -selection by construction, so it remains to verify the consistency and synchrony conditions for the new snapshot. For every proper pair $a, b \in \Sigma$, the fact that O is a $*$ -selection implies that all but one of the edge

counters in $\mathbf{T}|_{ab}$ coincide with their \mathbf{S} counterparts, while the exceptional one – let it be w_{ab} without loss of generality – is smaller than its counterpart in \mathbf{S} by one unit. Since the sum of edge counters in $\mathbf{S}|_{ab}$ is independent of the choice of square, we conclude the same is true for $\mathbf{T}|_{ab}$. To prove consistency we observe that $w_{ab} + w_{ab^*}$ suffering a decrease (of one unit) in the passage to \mathbf{T} implies $a \in O$ and hence $w_{ac} + w_{ac^*}$ must suffer a decrease as well for any $c \notin \{a, a^*\}$, since O is a complete $*$ -selection. Finally, with w_a being well-defined in \mathbf{T} for all $a \in \mathbf{P}$, the choice of $\#_{\mathbf{T}}$ guarantees that $\#_a = 1$ is only possible in \mathbf{T} if $w_a > 0$ in \mathbf{T} .

C. Equivalences in probabilistic Snapshots

Throughout this section, \mathbf{S} is a probabilistic snapshot satisfying the triangle inequality (32). The reason for the name is that the symmetric *dissimilarity measure* on $\Sigma \times \Sigma$ defined by

$$\mu(ab) := w_{a^*b} + w_{ab^*} \geq 0 \quad (87)$$

allows rewriting (32) in the form

$$\mu(ac) \leq \mu(ab) + \mu(bc) \quad (88)$$

for all $a, b, c \in \Sigma$.

Now let us turn to the purpose of this discussion, the analysis of equivalences in Σ that are observed in \mathbf{S} . Let $Eq(\mathbf{S})$ denote the undirected graph by vertex set Σ having edges ab and a^*b^* for every $ab \in \mathbf{K}_\Sigma$ satisfying (31). Let $\bar{\Sigma}$ denote the partition of Σ induced by the connected components of $Eq(\mathbf{S})$, and let $eq : \Sigma \rightarrow \bar{\Sigma}$ denote the map sending $a \in \Sigma$ to its block in $\bar{\Sigma}$. In other words, $eq(a)$ is nothing more than the list of all sensations $b \in \Sigma$ deemed equivalent to a , either directly or through a finite chain of equivalences.

Returning to the notation of appendix B-A and keeping in mind (76) we observe that:

$$a \equiv b \Rightarrow \omega(ab) = w_{a^*b} - w_{ab^*} = 0 \quad (89)$$

and since $\omega(\bullet)$ is additive (81) we conclude:

$$b \in eq(a) \Rightarrow \omega(ab) = 0 \quad (90)$$

In particular, since $\omega(ab) > 0$ whenever $ab \in \text{Dir}(\mathbf{S})$, we conclude:

Lemma B.1. *If $a, b \in \Sigma$ are connected by a directed path in $\text{Dir}(\mathbf{S})$ then $eq(a) \cap eq(b) = \emptyset$.*

Similarly, for the metric $\mu(\bullet)$ on Σ , we have that $\mu(ab) = 0$ if $a \equiv b$ by definition. The triangle inequality allows us to conclude:

Lemma B.2. *If $b \in eq(a)$ then $\mu(ab) = 0$.*

Proof. We have $b \in eq(a)$ iff there is a finite sequence of elements in Σ of the form

$$a = a_0 \equiv a_1 \equiv \dots \equiv a_n = b, \quad n \geq 0 \quad (91)$$

Applying the triangle inequality $n - 1$ times we obtain

$$\mu(ab) \leq \sum_{i=1}^n \mu(a_{i-1} a_i) = 0, \quad (92)$$

as desired. \square

In order to see that $\text{Dir}(\mathbf{S})$ actually defines a weak poc set structure on $\bar{\Sigma}$ we need the following:

Lemma B.3. *Given \mathbf{S} as above, for all $a \in \Sigma$ we have:*

- 1) $eq(a^*) = eq(a)^*$;
- 2) $eq(a^*) \neq eq(a)$.

(Recall the convention $A^* = \{x^* \mid x \in A\}$ for $A \subseteq \Sigma$)

Proof. For (1) it suffices to observe that $a \equiv b$ implies $a^* \equiv b^*$ by construction. Assertion (2) follows from observing that

$$eq(a^*) = eq(a) \Leftrightarrow eq(a^*) \cap eq(a) \neq \emptyset \Leftrightarrow a^* \in eq(a) \quad (93)$$

However, by the preceding lemma, $a^* \in eq(a)$ would imply $\mu(aa^*) = 0$ which, in turn, would contradict the obvious equality

$$\mu(aa^*) = w_{a^*a^*} + w_{aa} = 1 \quad (94)$$

finishing the proof. \square

The following proposition summarizes our progress thus far:

Proposition B.4. *Let \mathbf{S} be a probabilistic snapshot satisfying the triangle inequality. Then:*

- 1) *The operation $eq(a) \mapsto eq(a^*)$ defines a fixpoint-free involution on $\bar{\Sigma}$.*
- 2) *The directed graph Γ with vertex set $\bar{\Sigma}$ and with an edge pointing from $eq(a)$ to $eq(b)$ iff there exist $a' \in eq(a)$ and $b' \in eq(b)$ such that $a'b'$ is an edge of $\text{Dir}(\mathbf{S})$ is an acyclic poc graph.*
- 3) *Let $\mathbf{P} = \text{Poc}(\text{Dir}(\mathbf{S}))$ and $\bar{\mathbf{P}} = \text{Poc}(\Gamma)$. Then the quotient map $eq : \Sigma \rightarrow \bar{\Sigma}$ is a poc morphism of \mathbf{P} onto $\bar{\mathbf{P}}$ with the property that every fiber of eq is a transverse subset of \mathbf{P} .*
- 4) *For any subset $A \subset \Sigma$ one has*

$$A \uparrow = eq^{-1}(eq(A) \uparrow) \quad (95)$$

In particular, propagation over $\bar{\mathbf{P}}$ is equivalent to propagation over $\text{Dir}(\mathbf{S})_0$ (see defn. II.20). \square

D. Proof of proposition III.1

This proof requires the results of section A-E. The inclusions:

$$inc_{act} : \mathbf{Act}|_t \hookrightarrow \mathbf{P}|_t, \quad inc_{obs} : \mathbf{Obs}|_t \hookrightarrow \mathbf{P}|_t \quad (96)$$

satisfy $proj_{act} = inc_{act}^\circ$ and $proj_{obs} = inc_{obs}^\circ$, and since $\Sigma = \mathbf{Act} \cup \mathbf{Obs}$ and $\mathbf{Act} \cap \mathbf{Obs} = \{\mathbf{0}, \mathbf{0}^*\}$, we conclude that the identity map $id_{\mathbf{P}}$ of \mathbf{P} is, in fact, a surjective poc morphism

$$id_{\mathbf{P}} : \mathbf{Act}|_t \vee \mathbf{Obs}|_t \rightarrow \mathbf{P}|_t \quad (97)$$

(see A-D2 for a definition of $\mathbf{Act}|_t \vee \mathbf{Obs}|_t$). By proposition A.34, the dual of this map is a median-preserving embedding of cubical complexes:

$$id_{\mathbf{P}}^\circ : \text{Cube}(\mathbf{P}|_t) \hookrightarrow \text{Cube}(\mathbf{Act}|_t) \times \text{Cube}(\mathbf{Obs}|_t), \quad (98)$$

as required.

E. Local Structure of Duals and Greedy Navigation

In [30] we suggested exploring the link between the convexity theory of duals of weak poc sets and planning in DBAs, yet the formal results contained therein proved insufficient for supporting the planning algorithms proposed in this paper. This section fills in this gap.

Throughout this section we fix a finite weak poc set P and the median graph $\Gamma = \text{Dual}(P)$ (which is to say, Γ is an arbitrary finite median graph). We study the problem of computing the image of a non-empty convex subset $V(S)$ of Γ under the closest point projection of Γ to the convex subset $V(T)$.

1) *Gates:* We recall the following definitions and results from [31]:

Definition B.5 (Separator). Let $K, L \subseteq P^\circ$ be sets. The set

$$\text{sep}(K, L) = \{a \in P \mid K \subseteq V(a), L \subseteq V(a^*)\} \quad (99)$$

is called the separator of K and L . \square

The inequality $\Delta(u, v) \geq |\text{sep}(K, L)|$ follows immediately for all $u \in K$ and $v \in L$. This motivates:

Definition B.6 (Gate). Let $K, L \subseteq P^\circ$. A *gate* for K, L is a pair of points $u \in K, v \in L$ such that $\Delta(u, v) = |\text{sep}(K, L)|$. \square

The following result is well known in our setting:

Proposition B.7. *Let K, L be non-empty convex subsets of Γ and let $u \in K$ and $v \in L$. Then u, v form a gate for K, L if and only if $\text{proj}_K v = u$ and $\text{proj}_L u = v$. Moreover, any pair of non-empty convex subsets of Γ has a gate.*

We will apply this proposition without proof. An important consequence for us is the following:

Lemma B.8. *Suppose $K = \mathfrak{h}(S)$ and $S \subset P$ is coherent. Then, for any $a \in P$, if $K \subseteq \mathfrak{h}(a)$ then there exists $s \in S$ such that $s \leq a$.*

Proof. Let $u \in K$ and $v \in L := \mathfrak{h}(a^*)$ form a gate. Since $v \notin A$, there exists $s \in S$ such that $v \in \mathfrak{h}(s^*)$.

Suppose there were a $w \in B$ with $w \in \mathfrak{h}(s)$, and consider $m = \text{med}(u, v, w)$. Then $a \in v, w$ implies $a \in m$, but the inequality

$$\Delta(u, v) = \Delta(u, m) + \Delta(m, v) \geq \Delta(u, m) \quad (100)$$

implies $m = v$, since $v = \text{proj}_L u$. On the other hand, $s \in u, w$ implies $s \in m$ – a contradiction.

Thus, we have shown that $L = \mathfrak{h}(a^*)$ is contained in $\mathfrak{h}(s^*)$. Equivalently, $a^* \leq s^*$, which is the same as $s \leq a$. \square

The same kind of reasoning yields:

Lemma B.9. *Suppose K, L are non-empty convex subsets of $\text{Dual}(P)$. If $K \cap L \neq \emptyset$, then $\text{proj}_K L = \text{proj}_L K = K \cap L$.*

Proof. Clearly, if $v \in K \cap L$ then $\text{proj}_L(v) = v$, so $K \cap L \subset \text{proj}_L K$. For the reverse inclusion, suppose $v \in \text{proj}_L K$ and write $v = \text{proj}_L u, u \in K$. Pick any point $w \in K \cap L$. Setting

$m = \text{med}(w, v, u)$ we note that $m \in L$ (because $w, v \in L$) and

$$\Delta(u, v) = \Delta(u, m) + \Delta(m, v) \geq \Delta(u, m).$$

The uniqueness of projection forces $v = \text{proj}_L u$ to coincide with m . However, since $w, u \in K$ we also have $m \in K$, showing $v \in K \cap L$. \square

2) *The Coherent Projection:* We need to study a technical notion motivated by the necessity in correcting the observation of the current state as explained in section II-B. We recall the following standard notation for partially ordered sets:

$$a \uparrow = \{p \in P \mid a \leq p\}, \quad a \downarrow = \{q \in P \mid q \leq a\}, \quad (101)$$

and

$$A \uparrow = \bigcup_{a \in A} a \uparrow, \quad A \downarrow = \bigcup_{a \in A} a \downarrow. \quad (102)$$

Note that in a poc set P , one has the identities

$$A^* \uparrow = A \downarrow^*, \quad A^* \downarrow = A \uparrow^*. \quad (103)$$

For an arbitrary subset A of P we can define the following ‘correction’ of A :

Proposition B.10 (Coherent Projection). *Let P be a finite poc set and $A \subseteq P$ be any subset. Then the set*

$$\text{coh}(A) := A \uparrow \setminus A^* \downarrow = A \uparrow \setminus A \uparrow^* \quad (104)$$

is coherent, and satisfies the following properties:

- 1) if A is coherent then $\text{coh}(A) = A \uparrow$;
- 2) $\text{coh}(A) \uparrow = \text{coh}(A)$;
- 3) $\text{coh}(\text{coh}(A)) = \text{coh}(A)$;
- 4) if $A \in P^\circ$ then $A = \text{coh}(A) = A \uparrow$.

Proof. To show the coherence of $\text{coh}(A)$, let $b, c \in \text{coh}(A)$ with $b \leq c^*$; if $b \in A \uparrow$ then $c^* \in A \uparrow$ and $c \in A \uparrow^*$, contradicting $c \in \text{coh}(A)$.

For (1), A is coherent iff $A \uparrow$ and $A^* \downarrow$ are disjoint. Therefore, if A is coherent then $\text{coh}(A) = A \uparrow \setminus A^* \downarrow = A \uparrow$.

For (2), let $a \in \text{coh}(A)$ and $a \leq b$. Then $a \in A \uparrow$ implies $b \in A \uparrow$, and it suffices to verify $b \notin A \uparrow^*$. Indeed, were there $c \in A$ with $b^* \geq c$ then $a \leq b \leq c^*$ would have given $a \in A^* \downarrow$ in contradiction of $a \in \text{coh}(A)$. Thus, $b \in \text{coh}(A)$, as required.

For (3) since $\text{coh}(A)$ is coherent we have $\text{coh}(\text{coh}(A)) = \text{coh}(A) \uparrow$ by substituting $\text{coh}(A)$ instead of A in (1), and then we apply (2).

Finally, for (4), $A \in P^\circ$ means A is a coherent complete $*$ -selection, so $\text{coh}(A) = A \uparrow$ by (1) and it remains to show $A \uparrow = A$. Were there $b \in A \uparrow$ with $b \notin A$ we would have had $b^* \in A$, since A is a complete $*$ -selection. But then we would also have had $b, b^* \in \text{coh}(A)$, contradicting the coherence of $\text{coh}(A)$. \square

3) *Computing the Projection Maps:* For a vertex $u \in P^\circ$ and any subset $A \subset u$, one defines:

$$[u]_A := (u \setminus A) \cup A^* \quad (105)$$

Clearly, $[u]_A$ is a $*$ -selection. It is easily verified that $[u]_A$ is coherent if and only if there exists no pair $a \in A$ and $b \in u \setminus A$ satisfying $b < a$. This observation was first made in [56], leading to the following results in our setting:

Lemma B.11. *Let P be a finite weak poc set and let $u \in P^\circ$ be any vertex. Then the set $N(u)$ of vertices adjacent to u in $\Gamma = \text{Dual}(P)$ coincides with the set of all $[u]_a$, a ranging over the minset of u :*

$$\min(u) := \{a \in u \mid b < a \Rightarrow b \notin u\} \quad (106)$$

More generally, the cubes in $\text{Cube}(P)$ are characterized as follows:

Lemma B.12. *Let P be a finite weak poc set and $u \in P^\circ$ be a vertex. Then the cubes of $\text{Cube}(P)$ incident to u are in one-to-one correspondence with the transverse subsets of $\min(u)$.*

A particular application of these observations is an explicit construction of a geodesic path in Γ emanating from a given vertex u and terminating at its unique closest point projection $\text{proj}_{\mathfrak{h}(T)} u$:

Proposition B.13. *Let P be a finite weak poc set and suppose $u \in P^\circ$ is a vertex. Let T be a coherent subset of P . Then the following algorithm constructs a shortest path in Γ from u to $K = \mathfrak{h}(T)$:*

- 1) Find an element $b \in T \setminus u$; if no such element, stop and output u .
- 2) Find an element $c \leq b^*$ with $c \in \min(u)$;
- 3) Replace u by $[u]_c$ and go to the first step.

Proof. We have $u \in K$ iff $T \subset u$, which provides the stopping condition for the algorithm. Now, if $u \notin K$ and $b \in T \setminus u$ then for all $v \in K$ one has $v \in \mathfrak{h}(b)$ and $u \in \mathfrak{h}(b^*)$. Since $c \leq b^*$, we have $u \in \mathfrak{h}(c) \subseteq \mathfrak{h}(b^*)$, implying $v \in \mathfrak{h}(c^*)$ and $c \in u \setminus v$. As a result:

$$\Delta(v, [u]_c) = \Delta(v, u) - 1 \quad (107)$$

Having reduced $\Delta(u, v)$ by a unit for all $v \in K$, we have reduced $\Delta(u, K)$ by a unit as well. \square

Corollary B.14 (Projection of a Point). *Let P and T be as above. Then the closest point projection to $K = \mathfrak{h}(T)$ is given by the formula:*

$$\text{proj}_K u = (u \setminus T^* \downarrow) \cup T \uparrow = (u \cup T \uparrow) \setminus T^* \downarrow \quad (108)$$

Proof. The second equality follows from the DeMorgan rules and the fact that $T \uparrow \cap T^* \downarrow = \emptyset$ (since T is coherent).

Set $K = \mathfrak{h}(T)$ and proceed by induction on $\Delta(u, K)$. If $\Delta(u, K) = 0$, then $u \in K$ and therefore $T \subset u$. In addition, u is coherent and we conclude $T^* \downarrow \cap u = \emptyset$, leaving us with

$$u \setminus T^* \downarrow \cup T = u \cup T = u,$$

as desired. Now suppose $n := \Delta(u, K) > 0$. By the preceding proposition, there is $a \in T^* \downarrow \cap u$ such that $v := [u]_a \in P^\circ$, $\Delta(v, K) = n - 1$, and $\text{proj}_K u = \text{proj}_K v$. We thus have:

$$\text{proj}_K u = \text{proj}_K v = (v \setminus T^* \downarrow) \cup T \uparrow = (u \setminus T^* \downarrow) \cup T \uparrow,$$

the last equality being due to $a \in T^*$ and $a^* \in T$. Thus, the first identity has been proved. \square

4) Projecting a Convex Set to a Convex Set:

Proposition B.15. *Let K, L be non-empty convex subsets with $L = \mathfrak{h}(S)$ and $K = \mathfrak{h}(T)$. Then*

$$\begin{aligned} \text{proj}_K L &= \mathfrak{h}((S \uparrow \cup T \uparrow) \setminus T^* \downarrow) \\ &= \mathfrak{h}(T) \cap \mathfrak{h}(S \uparrow \setminus T \uparrow^*) \end{aligned} \quad (109)$$

Proof. Since T is coherent, $T \uparrow$ and $T^* \downarrow = T \uparrow^*$ are disjoint. This allows us to write:

$$\begin{aligned} \mathfrak{h}((S \uparrow \cup T \uparrow) \setminus T \uparrow^*) &= \mathfrak{h}(T \uparrow \cup (S \uparrow \setminus T \uparrow^*)) \\ &= \mathfrak{h}(T \uparrow) \cap \mathfrak{h}(S \uparrow \setminus T \uparrow^*) \end{aligned}$$

and the second equality in (109) follows from the identity $\mathfrak{h}(T) = \mathfrak{h}(T \uparrow)$. Denote $R = S \uparrow \setminus T \uparrow^*$ and $N = \mathfrak{h}(R)$.

For every $u \in L = \mathfrak{h}(S)$ we have $S \uparrow \subset u$, implying $\text{proj}_K u$ contains $T \uparrow \cup R$, by corollary B.14. Thus, $\text{proj}_K L \subset K \cap N$, as required.

For the converse, observe that the case $K \cap L \neq \emptyset$ was already dealt with in lemma B.9: if $K \cap L \neq \emptyset$, then

$$\text{proj}_K L = K \cap L = \mathfrak{h}(S \uparrow) \cap \mathfrak{h}(T \uparrow) = \mathfrak{h}(S \uparrow \cup T \uparrow)$$

In particular, $S \uparrow \cup T \uparrow$ is coherent, and hence does not intersect $T^* \uparrow$, and the formula (109) holds.

Thus we may henceforth assume $K \cap L = \emptyset$. Equivalently, $S \uparrow \cap T^* \downarrow \neq \emptyset$. In fact, by lemma B.8 we have $S \uparrow \cap T^* \downarrow = \text{sep}(A, B)$.

Starting with $v \in K \cap N$ we must show $v \in \text{proj}_K L$. Set $u = \text{proj}_L v$, $w = \text{proj}_K u$, and $m = \text{med}(u, v, w)$. Then $m \in K$ since $v, w \in K$. Since $K \cap L = \emptyset$, we have $\Delta(u, v) > 0$ and $\Delta(u, w) > 0$. Consider the point m : we have $m \in I(u, w)$ and $m \in K$; by the choice of w , m must equal w and therefore $w \in I(u, v)$. Thus, $w = \text{proj}_K u \in I(u, v)$ and $u = \text{proj}_L w$. By proposition B.7, the pair u, w is a gate for K, L and we have

$$u \setminus w = \text{sep}(L, K) = S \uparrow \cap T^* \downarrow.$$

Consider an element $a \in v \setminus u$. If $\mathfrak{h}(a) \cap L \neq \emptyset$, pick $u' \in \mathfrak{h}(a) \cap L$. Then $m = \text{med}(u, v, u')$ will satisfy $m \in \mathfrak{h}(a) \cap L$ as well as

$$\Delta(v, L) = \Delta(v, u) = \Delta(v, m) + \Delta(m, u).$$

Now, $\Delta(u, m) > 0$ since $u \in \mathfrak{h}(a^*)$ and a contradiction to $\text{uproj}_L v$ is obtained. Thus, $\mathfrak{h}(a) \cap L$ must be empty, which means $L \subseteq \mathfrak{h}(a^*)$. Applying lemma B.8 we obtain $a^* \in S \uparrow$.

Overall, we have shown that $v \setminus u \subseteq S \uparrow^*$. We will now verify that $v \setminus w = \emptyset$, finishing the proof. Indeed, were it not so, there would have been $h \in v \setminus w$. On one hand, $w \in I(u, v)$ implies $v \setminus w \subset v \setminus u$, and hence $h^* \in S \uparrow$. On the other hand, $h \notin w$ means $h^* \in w$ and therefore $h^* \notin \text{sep}(L, K) = S \uparrow \cap T \uparrow^*$, which forces $h^* \in R$. Since $R \subset v$

(by choice of v), we have $h^* \in v$, contradicting our choice of h . \square

We will need the following technical corollary for the purposes of propagation:

Corollary B.16. *Let $S, T \subset P$ be subsets and suppose S is coherent. Let $L = V(S)$ and $K = V(\text{coh}(T))$. Then:*

$$\text{proj}_K L = (S \uparrow \cup T \uparrow) \setminus T \uparrow^* = (S \uparrow \setminus T \uparrow^*) \cup \text{coh}(T) \quad (110)$$

Proof. Recall that $\text{coh}(T) = T \uparrow \setminus T \uparrow^*$, and set $J = T \uparrow \cap T \uparrow^*$, so that $T \uparrow = \text{coh}(T) + J$ and $T \uparrow^* = \text{coh}(T)^* + J$. Then,

$$\begin{aligned} (S \uparrow \cup T \uparrow) \setminus T \uparrow^* &= ((S \uparrow \cup \text{coh}(T) \cup J) \setminus \text{coh}(T)^*) \setminus J \\ &= (S \uparrow \cup \text{coh}(T)) \setminus \text{coh}(T)^* \end{aligned}$$

Since $\text{coh}(T) \uparrow = \text{coh}(T)$, the last expression equals $\text{proj}_K L$, by the preceding proposition. The proof of the second equality is similar. \square