ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRY

DAN P. GURALNIK AND DANIEL E. KODITSCHEK

ABSTRACT. We propose a variant of iterated belief revision designed for settings with limited computational resources, such as mobile autonomous robots. The proposed memory architecture—called the *universal memory architecture* (UMA)—maintains an epistemic state in the form of a system of default rules similar to those studied by Pearl and by Goldszmidt and Pearl (systems Z and Z^+).

A duality between the category of UMA representations and the category of the corresponding model spaces, extending the Sageev-Roller duality between discrete poc sets and discrete median algebras provides a two-way dictionary from inference to geometry, leading to immense savings in computation, at a cost in the quality of representation that can be quantified in terms of topological invariants. Moreover, the same framework naturally enables comparisons between different model spaces, making it possible to analyze the deficiencies of one model space in comparison to others.

This paper develops the formalism underlying UMA, analyzes the complexity of maintenance and inference operations in UMA, and presents some learning guarantees for different UMA-based learners. Finally, we present simulation results to illustrate the viability of the approach, and close with a discussion of the strengths, weaknesses, and potential development of UMAbased learners.

1. INTRODUCTION

1.1. Motivation. Iterated belief revision (BR) deals with the problem of maintaining syntactic propositional knowledge representations that are sufficiently flexible to accommodate reasoning about a stream of incoming observations in the form of propositional formulae (over a finite alphabet of atomic propositions), while taking into account the possibility of any such observation being inconsistent with the current state of the knowledge representation. It is not unreasonable then to argue that BR operators should be used for maintaining well-reasoned internal representations for autonomous learning agents (see, e.g. [47]). However, one needs merely to observe the high computational costs associated with revision operators [31, 30] to conclude that such representations are too expensive to implement them in a mobile autonomous agent. Attempts at making the representations more palatable using prime forms [6, 33] have been made, but the fundamental complexity barriers remain [26].

We introduce a computationally cheap form of iterated propositional belief revision the universal memory architecture (UMA)—which harnesses the geometry of model spaces in place of the model-theoretic techniques characteristic of this field. The computational advantages come at the price of modifying the notion of an observation and restricting the syntactic form of the epistemic state maintained by the agent (understood in the broad sense of Darwiche and Pearl [11]) to a special type of default system in the sense of [39]. Most notably, observations are no longer allowed to take the form of arbitrary propositional formulae; rather, we restrict them to conjunctive monomials in the underlying propositional variables. Equivalently, an observation is a partial truth-value assignment to the agent's inputs. In addition, each observation is accompanied by a *value signal*—a quantity indicating a notion of the value of the experience to the agent at that time. ¹

These alterations to the classical setting of iterated BR are motivated by the prospect of implementing iterated BR on mobile robotic platforms in real time. While the Boolean component of the observation corresponds to the robot's raw sensory inputs, the value signal may correspond to an encoding of a task, or to feedback from a teacher. The limited form of the epistemic state maintained by an UMA instance reduces the space and time complexity costs of maintenance (applying the revision operator) and exploitation (e.g. inference) down to an absolute minimum, as we review next.

1.2. Contributions: Introduction and Analysis of UMAs. Motivated by the problem of realizing iterated belief revision and update in a bounded resources setting, we seek a class of lightweight general-purpose representations. From a learning perspective, ours is a problem of *learning from positive examples*: an observer of an unknown, unmodeled system \mathscr{S} experiences some process—a sequence of transitions—in that system through an array of Boolean sensors, and is required to reason about regularities in the observed sequence of experiences, constructing a formal theory of what is possible for that system.

We assume that observations occur in discrete time steps. An observation at time t will consist of (1) a complete truth-value assignment² Obs $|_t$ —the observation at time t—to a fixed set Σ —the sensorium—of Boolean queries of the agent's interactions with its environment; and of (2) a sample $\varphi|_t$ of a fixed value signal, φ .

Little needs to be assumed about the sensorium: for the purpose of this paper, we allow any query expressible as a Boolean function of the state history (finite or infinite) of the system \mathscr{S} (an appropriate formalism is developed in Section 2.1.1); it is also assumed that truth-value assignments $\mathsf{Obs}|_t$ are *consistent* in the sense that each agrees with the values of the available queries on the history that manifested at the corresponding time; finally, observations are assumed to be *time shift-invariant* in the sense that observing the same histories at different times must yield the same Boolean observation vector. The value signal, for now, is assumed to be *static*, in the sense that it factors through a function of the observation (more detail in Section 3.1). The architecture itself does not rely on any of these assumptions, but the learning guarantees we provide in this paper do.

An UMA representation integrates its accumulated experiences by repeatedly revising two structural components, based on the incoming observations: (a) a relation $G|_{t}$, called a *pointed complemented relation* (PCR), representing a system of implications, or *defaults*, which the agent believes to hold true among the queries in Σ ; and (b) a set $\operatorname{Curr}|_{t} \subset \Sigma$, representing the agent's belief regarding the current

¹The value signal should not be confused with the notion of reward, as used in Reinforcement Learning. One of our learning schemes (see Section 4.2) leads to a (partial) syntactic representation of the distribution from which observations are being drawn, and does not encode any preference of one state over another.

²Henceforth, the symbol $(|_{\star})$ appended to anything else should be read as "at time t".

state of the system. The machinery for maintaining these data structures will be referred to as a *snapshot*. Briefly, our results about UMA representations are as follows.

Universality of Representation. In our intended setting, the learner's sensors realize the formal sensorium Σ as a family of subsets of the space of histories, closed under complementation. The possible worlds actually witnessed by points of this space correspond to the learner's perceptual equivalence classes (in the sense of, e.g. [13, 42]). Intuitively, an element (a, b) of the PCR $G|_t$ should be seen as correct if no history falsifies the formula $a \to b$, and, more generally, if histories falsifying $a \to b$ are improbable, or insignificant according to the user's formal model of these notions.

It turns out that a PCR G supports a natural dual space, a set $\mathbf{M} = G^{\circ}$ of possible worlds canonically associated with the PCR. Recall that a possible world over Σ is a complete truth value assignment $\Sigma \to \{\bot, \top\}$. We prove that, given a PCR G over a set of literals Σ , its dual space \mathbf{M} has the following universality property (Proposition 2.22): \mathbf{M} is the smallest set of possible worlds over Σ which, for any realization ρ of Σ as a set of Boolean queries over a space \mathbf{X} not falsifying a relation listed in G, contains every model for ρ .

Returning to UMA learners, this means that the model space $\mathbf{M}|_{t}$ encoded by the PCR $G|_{t}$ is a minimal envelope for the true space of possible worlds, provided *just the information* that all the relations recorded in $G|_{t}$ are correct.

Computational Complexity. From a computational perspective, the maintenance costs of an UMA representation are roughly the same as those of maintaining a neural representation (=the cost of maintaining and using a matrix of weights), but with the added benefit of affording a formal understanding of the model space, its geometry, and its deficiencies. Here are some results, all of which are corollaries of the *geometric* properties of the class of model spaces defined by PCRs. Let N denote the cardinality of the sensorium Σ . Then:

- Maintaining an UMA snapshot structure requires $O(N^2)$ space;
- Update operations for learning the PCR structure require $O(N^2)$ time;
- Inference requires $O(N^2)$ time, reducible to O(N) on fully parallel hardware. ³

Multiple Learning Paradigms. The mathematical foundations for UMA provide sufficient flexibility to admit a variety of learning mechanisms and settings, spanning the range from probabilistic filtering, as proposed in [19], to a variation on [iterated] revision and update introduced in [11], while keeping maintenance costs down to the bare minimum (see preceding paragraph). Depending on the *snapshot type*, different learning scenarios and guarantees may be provided, while maintaining a uniform revision and update scheme at the symbolic level.

Flexibility of Representation. A central feature of the UMA architecture is that the duality theory of PCRs allows one to interpret maps between PCRs as maps between the associated model spaces and vice versa. This makes it possible to formally introduce—as well as operate with—notions of approximate equivalence, of redundancy and negligibility of queries. This also enables the study of the impact on

³We will remark that our current implementation is, in fact, an $O(N^3)$ implementation utilizing matrix multiplication on a GPU. This kind of implementation makes it possible to multiply fairly big matrices very quickly, improving on the performance of the naïve quadratic algorithm we provide later in this paper.

model space geometry of operations augmenting a sensorium with new queries (see, for example, Appendix A.2.4) or removing existing ones. In particular, this opens a way to formal (and, possibly, automated) cost/benfit analysis of such extension and pruning operations—a topic of ongoing research at the moment, which we will touch upon briefly in our final discussion of the results presented in this paper.

1.3. **Related Work.** Given the focus of this work on the representation of knowledge using defaults, we believe it is most tightly related to work in the field of propositional iterated belief revision. Early work in BR resulted in wide acceptance of the AGM framework [4, 3, 2] for maintaining a *belief set*—a deductively closed set of formulae representing the state of the observed system. Convenient, intuitive axioms for belief revision in the propositional setting, the KM axioms, were developed by Katsuno and Mendelzon in [25].

Pointing out some inadequacies of the KM axioms in the context of repeated application of revisions, Darwiche and Pearl (DP) argue in their seminal paper [11] that, to achieve the overarching goal of iterated revision, one must maintain a set of conditional statements—an *epistemic state*—which, upon revision by an incoming observation, always produces a belief set accommodating that observation (axiom $\mathbf{R} * \mathbf{1}$ of the DP system of axioms for iterated revision). Building on Spohn's framework of ordinal conditional functions [46] and its implications for ranked default systems [39, 17] and revision of the associated belief sets [18], they propose to view ranking functions as epistemic states (interchangeable with the associated system of ranked defaults), as they construct appropriate revision operators. Consequent work by many authors [24, 12, 27, 22, 34, 28]—much of it very new—considers different weaknesses and benefits of the DP axioms, relating to the effect of the order in which observations are made and the manner of mutual dependence they present, and resulting in a variety of iterated revision methods, as well as in some proposals to apply belief revision methods to the control of general agents [47] based on varying computational approaches to belief revision operators (e.g. [6, 33] on the use of prime forms for this purpose).

Clearly, the problems tackled by this field generalize the representation problem we posed at the beginning of Section 1.2, but one needs merely to observe the high computational costs associated with revision operators [31, 30] (or with computing normal forms and prime forms [26]) to reach the conclusion that the existing computational approaches cannot be considered viable candidates for a solution of the representation problem in any setting where computational resources are limited.

Aiming to reduce the computational burden on the learner, we shift attention from precise syntactic computation with arbitrary propositional formulae to imposing *radical* simplifying assumptions on the allowed model spaces. The postulated mode of interaction between the agent and its environment—specifically the fact that the agent is constrained to processing sequences of samples from the space **M** of realizable models (rather than arbitrary propositional formulae)—suggests constructing successive *upper approximations* $\mathbf{M}|_t \supset \mathbf{M}$ of **M**, belonging to a restricted class \mathscr{C} which satisfying the following intuitive properties:

- (1) Syntactic characterization of an element in \mathscr{C} is computationally inexpensive;
- (2) Each approximation is, in some sense, optimal/minimal among members of *C*, given its predecessor and the last observation;
- (3) Reasoning (e.g., forming a belief set) over a member of \mathscr{C} is cheap.

We present results on what is, in essence, the simplest possible class \mathscr{C} of model spaces satisfying these three requirements: the class of finite median algebras. This class of spaces is well studied, in several different guises, and in very disparate fields. These include: event structures in parallel computation [40]; median graphs in metric graph theory [8]; simply connected non-positively curved cubical complexes in formalizations of reconfiguration in robotic systems [16]; and the spectacular recent achievements in the topology of 3-dimensional manifolds by Agol [1] are much due to the notion of a cubulated group from Geometric Group Theory [50].

1.4. Structure of this Paper. In Section 2, we extend Sageev-Roller duality⁴, to obtain all finite median algebras as duals (model spaces) of PCRs, viewed as systems of defaults. Further, we explain how to reason over model spaces in this class by leveraging their geometry to avoid satisfiability checks, or any kind of explicit search in model space, for that matter. We then explain in Section 3 how, using UMA snapshot structures to perform a variant of iterated revision, where the model-theoretic outlook on the problem is replaced by its geometric counterpart arising by Sageev-Roller duality. We discuss the necessity of relaxing the DP axiom $\mathbf{R*1}$, and show there is a natural operator for computing a belief set, the *coherent projection*.

Section 4 presents two different classes of snapshot structures—mechanisms for learning PCR representations—one motivated by Goldszmidt and Pearl's interpretation of default reasoning as qualitative probabilistic reasoning [18], and the other based on statistical integration of the observed value signal. Finally, Section 5 presents two kinds of simulation studies:

- (1) First, in a range of settings with *a-priori* known (or readily computable) implications in the sensorium, we consider the deviation of the learned PCR from the ground truth as a function of the number of samples. This is done for both snapshot types, and under different exploration paradigms: sampling and diffusion.
- (2) Next, we consider settings closer to the heart of a roboticist. We implement agents with a reactive control paradigm based entirely on their internal UMA representations and conduct comparative simulation studies of their performance given different domains for exploration, and snapshot types.

We close with a discussion of our results and of avenues for additional research in Section 6.

2. Model Spaces for Systems of Approximate Implications.

In this section we construct a representation for finite median algebras (see above) that is sufficiently flexible to be maintained dynamically, and we explain how to reason over these representations. We review and apply existing results about the geometry of model spaces of this class of representations, leading to complexity bounds on maintenance and exploitation.

Section 2.1 formally introduces the basic formal notions required for discussing our representations. Section 2.2 constructs the model spaces as dual spaces of *pointed complemented relations* (PCRs) and discusses their universal properties.

 $^{{}^{4}}$ See [43] for a detailed development of that theory; chapters 6-7 of [50] for a brief intuitive review; and here, Appendix A for background material and examples developed specifically to support this paper.

Section 2.3 relates PCRs and their duals (the associated model spaces) to the earlier duality theory of poc sets that motivated our approach, showing that PCR duals are, in fact, poc set duals. Section 2.4 reviews known results about the geometry and topology of poc set duals. Finally, in Section 2.5 we discuss the connection between the geometry of PCR duals and algorithms enabling reasoning over PCRs.

2.1. **Pointed Complemented Relations (PCR)..** The nature of our application requires a generalization of the formal theory we are about to use, the Sageev-Roller duality theory of poc sets [43], prompting some changes in the language. We start with:

Definition 2.1 (pointed complemented set, PCS). A pointed complemented set is a set Σ endowed with a self-map $a \mapsto a^*$ satisfying $a^{**} = a$ and $a^* \neq a$ for all $a \in \Sigma$, and containing a distinguished element, denoted **0**. The element **0**^{*} will be denoted **1**. Whenever possible and safe, we will abuse notation and use the symbols **0**, **1**, * in different PCSs. For any $S \subset \Sigma$ we will denote by S^* the set of all $x^*, x \in S$. \Box

Definition 2.2 (PCS morphism). By a PCS morphism we mean a function $f : \Sigma_1 \to \Sigma_2$ between PCSs satisfying $f(\mathbf{0}) = \mathbf{0}$ and $f(a^*) = f(a)^*$ for all $a \in \Sigma_1$. The set of all PCS morphisms from Σ_1 to Σ_2 will be denoted by $\operatorname{Hom}_{PCS}(\Sigma_1, \Sigma_2)$. \Box

Example 2.3 (set families, power sets). Any collection $\mathscr{U} \subseteq \mathbf{2}^X$ of subsets of a fixed non-empty set \mathbf{X} satisfying (1) $\emptyset \in \mathscr{U}$, and (2) $A \in \mathscr{U} \Rightarrow \mathbf{X} \setminus A \in \mathscr{U}$. Then \mathscr{U} is a PCS with respect to the choices $\mathbf{0} := \emptyset$ and $A^* := \mathbf{X} \setminus A$.

The power set of a singleton is, up to isomorphism, the smallest PCS, which we denote by **2**, and identify with the set $\{\bot, \top\}$. Also, the power set $\mathbf{2}^X$ will be routinely identified with the set of all functions $X \to \{\bot, \top\}$.

Example 2.4 (PCS over an alphabet). Suppose \mathbb{A} is a finite collection of symbols, and think of them as atoms of the propositional calculus over \mathbb{A} . The extended collection of literals over \mathbb{A} ,

(1)
$$\Sigma(\mathbb{A}) := \{\bot, \top\} \cup \bigcup_{a \in \mathbb{A}} \{a, \neg a\}$$

may be thought of as a PCS when one declares $\mathbf{0} := \bot, \bot^* := \top, \top^* := \bot$ and $a^* := \neg a, (\neg a)^* := a$ for all $a \in \mathbb{A}$. Hereafter, \top and \bot stand for the truth values True and False, respectively.

The reason for considering PCSs is that *-selections "live on them":

Definition 2.5 (*-selection, the Hamming cube). Let Σ be a PCS. By a *-selection on Σ we mean a subset $S \subset \Sigma$ such that $S \cap S^* = \emptyset$. In addition, a *-selection Son Σ is complete, if $S \cup S^* = \Sigma$. The set of all *-selections $S \subset \Sigma$ with $\mathbf{1} \in S$ will be denoted by $\mathbf{S}(\Sigma)$, and referred to as the [combinatorial] Hamming cube on Σ . Its set of vertices, the complete *-selections in $\mathbf{S}(\Sigma)$, will be denoted by $\mathbb{H}(\Sigma)$. \Box

We now consider these notions in the context of our intended application.

2.1.1. Binary Sensing, Possible Worlds and Perceptual Classes. Suppose α is an observer of some system \mathscr{S} as it undergoes the transitions along a state trajectory $(p_t)_{t=-\infty}^{\infty}$, and suppose \mathbb{A} is a finite set of unique labels for the Boolean queries available to α —this observer's sensorium. We assume observations of \mathscr{S} by α begin at t = 0. It will not matter for our discussion whether the trajectory of \mathscr{S} in

any particular instance does indeed extend indefinitely into the past or future: if needed, one may set the value of p_t to be eventually constant (in either direction).

By a history of \mathscr{S} we mean a sequence of the form $\mathbf{x} := (x_s)_{s=-\infty}^0$, where x_s is a state of \mathscr{S} for all s, and x_0 represents the *current state* of the history \mathbf{x} ; x_{-1} represents the *preceding state*, and so on. Given a trajectory (p_t) of \mathscr{S} observed by α , at each time $t \ge 0$, the history that manifests at time t is given by $x_s := p_{t+s}$.

Henceforth, we let **X** denote the space of histories possible for the system \mathscr{S} given the initial history manifested at time t = 0 (as is the case in all physical systems, \mathscr{S} may have its own dynamics, disqualifying some histories from manifesting at any time t > 0, or making such events highly improbable). To say that α 's queries/sensors are time-shift invariant is to say that each query is represented by a fixed Boolean function of the manifested history. In other words, the sensorium is *defined* by a PCS morphism $\rho : \Sigma \to 2^X$, $\Sigma := \Sigma(\mathbb{A})$, with a sensor $s \in \Sigma(\mathbb{A})$ reporting \top on history $x \in \mathbf{X}$ if and only if $x \in \rho(s)$.

The mapping ρ induces a partition on X—its partition into perceptual classes as follows. Construct a map $\rho^* : \mathbf{X} \to \mathbb{H}(\mathbf{\Sigma})$ by setting $s \in \rho^*(x)$ if and only if $x \in \rho(s)$; each point is mapped to the set of queries (including complements) which evaluate to \top on that point. Two points $x, y \in \mathbf{X}$ are sensory-equivalent if $\rho^*(x) = \rho^*(y)$. The image $\mathbf{M}(\rho) := \mathrm{Im}(\rho^*)$ are the possible perceptual states of α in the system \mathscr{S} , given ρ and the system's initial history. We will also refer to a world/*-selection $u \in \mathbb{H}(\mathbf{\Sigma})$ as consistent, if, and only if $u \in \mathbf{M}(\rho)$, or, in other words, if and only if u is witnessed (through ρ) by a point of \mathbf{X} .

2.1.2. Concept Presentation of Perceptual States. Digging deeper into the formalism presented just now, observe that *-selections $S \subset \Sigma(\mathbb{A})$ are in one-to-one correspondence with vectors, as defined in concept learning [48]. Recall that a vector is an assignment $v : \mathbb{A} \to \{0, 1, \cdot\}$ of values standing for \bot , \top , and "undetermined", respectively, to the alphabet \mathbb{A} . A vector is total if it has no (•) values. The map $v \mapsto \sigma_v := v^{-1}(1) \cup (v^{-1}(0))^*$ is then a correspondence between vectors over \mathbb{A} and *-selections on the PCS $\Sigma(\mathbb{A})$, mapping the set of total vectors onto the set of complete *-selections. In more geometric terms, a complete *-selection—which corresponds to a complete conjunctive monomial (aka complete term) over \mathbb{A} —defines a vertex of the cube $[0, 1]^{\mathbb{A}}$, while a *-selection S with $|S| = |\mathbb{A}| - d$ corresponds to a d-dimensional face. We will refer to $[0, 1]^{\mathbb{A}}$ as the Hamming cube. The advantage of PCS terminology here is that *-selections on Σ enumerate the faces of the Hamming cube without us having to pick an origin for the cube.

Pushing the geometric viewpoint a bit further, we consider the notion of *concepts*. In [48], Valiant defines concepts as mappings F of the space of vectors to $\{\perp, \top\}$, satisfying the requirement that F(v) = 1 on a vector v if and only if F(w) = 1 for all total vectors w which agree with v on those $a \in \mathbb{A}$ where $v(a) \neq \bullet$. In other words, concepts correspond to collections K of *faces* of the Hamming cube, possibly of varying dimensions, satisfying the condition that a face F belongs to K if and only if every vertex of F lay in K. Such K are *precisely* the sub-complexes of the Hamming cube obtainable from it by vertex deletions.⁵

Now we return to the observer α and the system \mathscr{S} whose evolution it observes through the queries realized by $\rho: \Sigma(\mathbb{A}) \to 2^{\mathbf{X}}$, as discussed in the preceding section. Thinking of the space of perceptual classes $\mathbf{M}(\rho)$ as a concept gives rise to

⁵Similarly to case of graphs, the operation of deleting a vertex from a cubical complex requires the removal of all the adjoining faces.

a cubical sub-complex, say $\text{Cube}(\rho)$, of the Hamming cube, whose faces correspond to those *-selections on the PCS $\Sigma = \Sigma(\mathbb{A})$ that are witnessed (via ρ) by a point in **X**. Thus, precise reasoning and planning over $\mathbf{M}(\rho)$ depends on one's ability to efficiently capture/encode: (1) the notion of consistency produced by the map ρ ; (2) the topological properties (e.g. connectivity, contractibility) of $\text{Cube}(\rho)$; and (3) the geometric properties (e.g. shortest paths, curvature, isoperimetric inequalities) of $\text{Cube}(\rho)$. The class of approximating model spaces we propose to use as proxies for $\mathbf{M}(\rho)$ is a result of weakening this notion of consistency to the extreme, all the way to the notion of *coherence* discussed in the next section.

2.1.3. PCRs, Implications and Coherence.

Definition 2.6 (pointed complemented relation, PCR). Let Σ be a PCS. By a pointed complemented relation over Σ we mean a set $G \subseteq \Sigma \times \Sigma$ satisfying⁶ $\mathbf{0}a \in G$ and $ab \in G \Leftrightarrow b^*a^* \in G$ for all $a, b \in \Sigma$.

In the context of the representation problem, one should think of a PCR G over Σ as a record of Boolean implications believed to be valid over $\Sigma = \Sigma(\mathbb{A})$, conditioned on the particular space of histories being observed. In this respect, a PCR is a restricted form of the notion of a *system of defaults*, as discussed, e.g. in [18]. Some of these implications are specified directly ($ab \in G$ to be read as "it is believed that b follows from a"), while others are derived as their consequences, by transitive closure. Hence the following language:

Definition 2.7. Given a PCR G over a PCS Σ , for any $a, b \in \Sigma$, $S \subseteq \Sigma$, one defines the following:

- Write $a \leq_G b$ if ab lies in the reflexive and transitive closure of G;
- The *G*-equivalence class of $a \in \Sigma$, denoted $[a]_G$, is the equivalence class of a under the relation $a \sim b \Leftrightarrow a \leq_G b \land b \leq_G a$ on Σ ;
- The forward (backward) closure, $S \uparrow$ (resp. $S \downarrow$), of S with respect to G is the set of all $b \in \Sigma$ for which $a \leq_G b$ (respectively $b \leq_G a$) holds for some $a \in S$;
- Note that $S \subset S\uparrow$. One says that S is forward-closed if $S\uparrow = S$;
- Finally, we observe that $S^* \uparrow = S \downarrow *$ for all $S \subseteq \Sigma$.

We will often drop the subscripts G when no ambiguity can arise.

Definition 2.8 (PCR morphism). Let G_1, G_2 be PCRs over Σ_1, Σ_2 , respectively. A morphism of PCRS from G_1 to G_2 is a PCS morphism $f : \Sigma_1 \to \Sigma_2$, additionally satisfying $f(a) \leq f(b)$ in G_2 whenever $ab \in G_1$. The set of all morphisms from G_1 to G_2 will be denoted by $\operatorname{Hom}_{PCR}(G_1, G_2)$.

The primary example of a PCR for this work derives from the view of a power set as a PCS (Example 2.3):

Example 2.9 (Set Families as PCRs). Let $\mathbf{X} \neq \emptyset$ be a set. Then any collection \mathscr{U} of subsets of \mathbf{X} that is closed under complementation and satisfies $\emptyset \in \mathscr{U}$ gives rise to the PCR of all pairs (A, B) with $A \subseteq B$, $\mathbf{0} = \emptyset$ and $A^* = \mathbf{X} \setminus A$. In what follows, $\mathbf{2}^X$ will always be regarded as a PCR in this way, for any $\mathbf{X} \neq \emptyset$. \Box

Another 'canonical' example of a PCR to keep in mind is:

⁶To avoid a proliferation of parentheses, we write ab to denote the pair $(a, b) \in \Sigma \times \Sigma$.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRY9

Example 2.10 (Less classical PCRs). Let $\mathbf{X} \neq \emptyset$ be any set. Then $[0, 1]^{\mathbf{X}}$ may be endowed with the structure of a PCR by setting $\mathbf{0}(\mathbf{x}) := 0$, $\mathbf{x} \in \mathbf{X}$, and, for any $\psi, \varphi \in [0, 1]^{\mathbf{X}}$, setting $\psi^*(\mathbf{x}) := 1 - \psi(\mathbf{x})$, $\mathbf{x} \in \mathbf{X}$ and $(\varphi, \psi) \in G$ if and only if $\varphi(t) \leq \psi(\mathbf{x})$, $\mathbf{x} \in \mathbf{X}$.

Our notion of model for a PCR rests on the following weak form of consistency:

Definition 2.11. Let G be a PCR over Σ . A subset $S \subseteq \Sigma$ is said to be G-coherent, if no pair $a, b \in S$ satisfies $a \leq b^*$.

Note that a G-coherent set is always a *-selection on Σ . Furthermore:

(2) S is coherent $\Leftrightarrow S \cap S^* \downarrow = \emptyset \Leftrightarrow S \uparrow \cap S^* = \emptyset \Leftrightarrow S \uparrow \cap S^* \downarrow = \emptyset$,

so coherence is preserved by forward closure. Coherent, forward-closed sets may be thought of as the natural counterparts of the notion of a belief state in this setting. We now turn to studying the appropriate notion of model.

2.2. Model Spaces as Dual Spaces.

Definition 2.12 (duals). Let G be a PCR over Σ . The set G° of maximal G-coherent subsets of Σ is the *dual of* G. The set of all forward-closed G-coherent subsets will be denoted $\mathbf{C}(G)$.

A standard application of Zorn's lemma shows that any G-coherent subset of Σ is contained in an element of G° . Note also that $G^{\circ} \subseteq \mathbf{C}(G)$.

Example 2.13 (the orthogonal PCR and the Hamming cube). The simplest example of a dual space is one where the PCR in question is as small as possible. Let Σ be a PCS. The smallest PCR over Σ contains only pairs of the forms 0a and a1. We will denote this PCR by $\mathbb{O}(\Sigma)$ and refer to it as the *orthogonal PCR* over Σ . It is clear that $\mathbb{O}(\Sigma)^{\circ} = \mathbb{H}(\Sigma)$, the "Hamming cube" from Definition 2.5.

Example 2.14 ('bad' queries). The definitions given above do not preclude one from considering, for example, the PCR $G_1 = \{01, 10\}$. It is easy to see that $G_1^{\circ} = \{\emptyset\}$. At the same time, the smaller $G_2 = \{01\}$ has $G_2^{\circ} = \{1\}$. More generally, for any $a \in \Sigma$, having $a \leq a^*$ precludes a from belonging in any *G*-coherent set. In particular, if both $a \leq a^*$ and $a^* \leq a$ hold, then no *G*-coherent set is a complete selection on Σ .

Following the last example, two definitions are in order:

Definition 2.15. The *trivial* PCR, henceforth also denoted by **2**, is the PCR over $\Sigma = \{0, 1\}$ containing only **01**.

Definition 2.16 (negligible query, degenerate graph). Let G be a PCR over Σ . An element $a \in \Sigma$ is G-negligible, if $a \leq a^*$. Denote the set of negligible elements by N(G). We say that G is degenerate if Σ contains a negligible element whose complement is also negligible. Note that $N(G) \downarrow = N(G)$.

Proposition 2.17. For a PCR G over Σ , the following are equivalent:

- (1) G is non-degenerate;
- (2) Every element of G° is a complete selection on Σ ;
- (3) Some element of G° is a complete selection on Σ .

Proof. See Appendix B.1.

The impact of this result on our representation problem is twofold. First, it provides a clear and easily verifiable criterion for when the dual space of a PCR consists (only!) of possible worlds. Second, it introduces a new and consistent notion of a query of low import, not involving arbitrary choices such as thresholding.

Proposition 2.18. Let G be a non-degenerate PCR over the PCS Σ . Then the mapping χ : Hom_{PCR}(G, 2) \rightarrow G[°] defined by $\chi(f) = f^{-1}(1)$ is a bijection.

 \square

Proof. See Appendix B.2.

Remark 2.19. Note that the mapping χ is independent of the choice of G.

The last proposition explains the sense in which G° may be thought of as a *dual* space of G. As with other instances of duality, this is useful because it enables dual mappings:

Definition 2.20. Let $f : G_1 \to G_2$ be a PCR morphism. The *dual mapping* $f^{\circ} : G_2^{\circ} \to G_1^{\circ}$ is defined by $f^{\circ}(S) = f^{-1}(S)$. Alternatively, upon applying the identification in Proposition 2.18, for any $\varphi \in \operatorname{Hom}_{PCR}(G_2, 2)$, one has $f^{\circ}(\varphi) = \varphi \circ f$ to obtain an element of $\operatorname{Hom}_{PCR}(G_1, 2)$.

We remark that, since morphisms are composable (meaning that the composition $(f \circ g)(a) := f(g(a))$ of two morphisms is a morphism as well), so are their dual mappings, producing the identity $(f \circ g)^{\circ} = g^{\circ} \circ f^{\circ}$.

Example 2.21. Let G be a non-degenerate PCR over a PCS Σ . Then it is clear that the identity mapping $\iota : \mathbb{O}(\Sigma) \to G$ — that is: $\iota(a) = a$ for all $a \in \Sigma$ — is a morphism of PCRs. The dual mapping $\iota \circ : G^{\circ} \to \mathbb{H}(\Sigma)$ is then, clearly, an injection. This reflects the intuitive notion that the dual of any (non-degenerate) PCR may be "excavated" out of a standard Hamming cube by going over all G-incoherent pairs, one by one, and successively deleting any vertices of $\mathbb{H}(\Sigma)$ which contain the given pair.

We further specialize the example to our representation problem, considering the effect of fixing a PCR structure on a given PCS:

Proposition 2.22 (Universality of Representation). Let G be a non-degenerate PCR over Σ . Then, for any non-empty set **X** and every PCS morphism $\rho: \Sigma \to 2^{\mathbf{X}}$, the set $\mathbf{M}(\rho)$ of all complete *-selections witnessed (via ρ) by a point in **X** (in the sense of Section 2.1.1) is contained in G^o whenever ρ is a PCR morphism. Moreover, G^o is the smallest subset of $\mathbb{H}(\Sigma)$ having this property.

Proof. See Appendix B.3.

Thus, the dual G° of a non-degenerate G serves as a minimal model of the state space of the system (*agent* + *environment*), and remains valid under *any* change to this system for as long as ρ remains order-preserving. This is a form of robustness of the representation to changes in the coupling between the agent's sensory equipment and the environment: changes leaving the implication record invariant provide no reason for the agent to alter its reasoning.

2.3. Reducing PCR Representations. The universality of PCR duals motivates a deeper study of their properties, seeking a better understanding of the degree of redundancy in the description of G° by a PCR G. This is not a mere technical issue: while non-degeneracy guarantees the adequacy of our notion of an associated

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRM

"possible world", it is not obvious that it also provides for sufficient control over the quality of inference. The intended application—inferring *approximate* implications from partial observations—is well known to be problematic in the absence of simplifying assumptions (e.g. the ubiquitous restriction to directed acyclic graphs in the context of Bayesian networks). It is therefore crucial to clarify the precise formal sense in which a PCR may be viewed as encoding a "record of implications", which is the purpose of this section. A crucial notion in any such discussion is that of what it means for a query, as well as for the difference of two queries, to be negligible, because negligible but non-zero differences tend to accumulate in the transitive closure into material ones.

Looking more closely at the setting of the last proposition, notice that, for a fixed ρ , the assumption that ρ is a morphism translates into the following. The property $a \leq b \Rightarrow \rho(a) \subseteq \rho(b)$ for all $a, b \in \Sigma$ implies $\rho(a) = \emptyset$ for any $a \in N(G)$ (because \emptyset is the only negligible element of $\mathbf{2}^{\mathbf{X}}$); furthermore, $\rho(a) = \rho(b)$ must hold whenever a and b are G-equivalent (recall Definition 2.7). These identifications lead us to recall Roller's definition of a *poc set* from [43]:

Definition 2.23 (poc set). A poc set is a tuple $\mathbf{P} = (\Sigma, \leq, \mathbf{0}, *)$ where (Σ, \leq) is a partially ordered set with a minimum element $\mathbf{0}$, endowed with an order-reversing involution⁷ $a \mapsto a^*$ satisfying $\mathbf{0}^* \neq \mathbf{0}$ and $a \leq a^* \Rightarrow a = \mathbf{0}$ for all $a \in \Sigma$.

In other words, a poc set is a transitive and anti-symmetric PCR over Σ whose only negligible element is **0**.

Proposition 2.24 (canonical quotient). For any non-degenerate PCR G there exists a surjective PCR morphism $\pi_G : G \to \widehat{G}$ of G onto a poc set \widehat{G} such that any PCR morphism $f : G \to \mathbf{P}$ gives rise to one and only one PCR morphism $\widehat{f} : \widehat{G} \to \mathbf{P}$ satisfying $f = \widehat{f} \circ \pi$.

Proof. We defer the proof to Appendix B.4, but define the canonical quotient mapping here. We set:

(3)
$$\pi_G(a) := \begin{cases} [a]_G & \text{if } a \notin N(G) \cup N(G)^* \\ N(G) & \text{if } a \in N(G) \\ N(G)^* & \text{if } a \in N(G)^* \end{cases}$$

and let $\widehat{\Sigma} := \{\pi_G(a) | a \in \Sigma\}$, and setting $\pi_G(a) \leq \pi_G(b)$ to hold in \widehat{G} if and only if $ab \in G$. It remains to verify that (1) $\widehat{\Sigma}$ is a well-defined PCS; (2) \widehat{G} is a well-defined poc set structure over $\widehat{\Sigma}$; and (3) the assertions of the proposition hold. \Box

One should view this result as stating the precise conditions necessary for presenting a poc set in terms of a set of generators and a set of relations. However, the emphasis on what happens to morphisms leads to powerful realizations about dual spaces:

Corollary 2.25 (all duals are poc set duals). If G is a non-degenerate PCR then $\pi_G^\circ: \widehat{G}^\circ \to G^\circ$ is a bijection.

Proof. See Appendix B.5.

⁷That is, $a^{**} = a$ and $a \leq b \Rightarrow b^* \leq a^*$ for all $a, b \in \Sigma$.

Corollary 2.26 (naturality of canonical quotients). Let G, H be non-degenerate PCRs. Then, for every morphism $f : G \to H$ there exists one and only one morphism $\hat{f} : \hat{G} \to \hat{H}$ satisfying $\pi_H \circ f = \hat{f} \circ \pi_G$.

 \Box

Proof. See Appendix B.6.

 \Box

A particular consequence of the last corollary is that one also has $\pi_H^\circ \circ \hat{f}^\circ = f^\circ \circ \pi_G^\circ$. This means the dual maps of f and \hat{f} coincide up to the identifications between the pre- and post-projection duals. Thus, any results about poc set duals apply to duals of PCRs. In the next two sections we review these results, and then harness them in our construction of the *universal memory architecture* (UMA).

2.4. Convexity theory of PCR duals. To discuss the geometry of PCR duals, we need to endow PCRs with more structure. FROM THIS POINT ON, ALL PCRS WE CONSIDER WILL BE FINITE, WITH THE SOLE POSSIBLE EXCEPTION OF POWER SETS.

Definition 2.27 (Hamming metric). Let G be a PCR over Σ . The Hamming metric on G° is defined by $\Delta(u, w) = |\pi_G(u) \setminus \pi_G(w)|$, where $\pi_G : G \to \widehat{G}$ is the canonical quotient map. We define Dual(G) to be the simple⁸ graph with vertex set G° , and edges of the form $\{u, w\}$ for all $u, w \in G^{\circ}$ with $\Delta(u, w) = 1$.

In the case when G is already a poc set, two vertices $u, w \in \mathbf{P}^{\circ}$ form an edge if and only if $u \smallsetminus w$ is a singleton, that is: the perceptual classes represented by u and w differ by the truth value of a single query. The common edge they span in the Hamming cube $\mathbf{S}(\Sigma)$ corresponds to the *-selection $u \cap w$ in the concept presentation. In the general case (G not necessarily a poc set), since both u and w are coherent, each is the union of $N(G)^*$ with a number of G-equivalence classes $[a], a \notin N(G)$ (recall Definition 2.7 and Proposition 2.24). Thus u and w span an edge in G° if and only if $u \smallsetminus v = [a]$ for some $a \notin N(G) \cup N(G)^*$. Intuitively, we think of the different $b \in [a]$ as counting for a single Boolean query.

We briefly recall the graph-theoretic notion of convexity:

Definition 2.28 (convexity in graphs). Let $\Gamma = (V, E)$ be a graph and let $u, v \in V$. The *hop distance* $d_{\Gamma}(u, v)$ is defined to be the minimum length of an edge-path in Γ joining u with v. The *interval* I(u, v) is defined to be the set of all vertices $w \in V$ satisfying the equality $d_{\Gamma}(u, v) = d_{\Gamma}(u, w) + d_{\Gamma}(w, v)$. A set $C \subseteq V$ is said to be *convex* in Γ , if $I(u, v) \subseteq C$ holds for all $u, v \in C$. A set $H \subseteq V$ is a *half-space of* Γ , if both H and $V \smallsetminus H$ are convex sets in G. Finally, we denote by $\mathcal{H}(\Gamma)$ the poc set whose elements are the half-spaces of Γ (note that \emptyset is a half-space of Γ), ordered by inclusion, and with $H^* := V \smallsetminus H$.

We refer the reader to [38], section 4, for the (very elegant and much more general) proofs of the following two lemmas (stated there for poc sets, but valid for finite non-degenerate PCRs as well, due to Proposition 2.24 and its two corollaries):

Lemma 2.29. Let G be a finite non-degenerate PCR. Then the hop metric on Dual(G) coincides with the metric Δ .

⁸That is: loopless, unoriented, with no multiple edges.

							•	Г					Г						
							ľ												
_	_			_				⊢			_								_
	_			_				⊢		H	4	,	14	,)	_			_	-
		_		-			-	-	1	H	V	,	r i	9				_	-
+	Η	H	11	-	12)	-	⊢		-			-				-	-	-
+	1	ł	и	,	۲,	-						-	-	-	-		-		-
+	-	-	-	-			H	Ħ			-	-	⊢		-	-	H	•	
						Π											F	H	Þ
							7	1				١							
								1	l,	V	V	1							
								1	ĺ		1								
11	-																		
w	r .				1		1						1				1		

FIGURE 1. Computing a median in a rectangle Γ cut out of the integer grid (all vertices of the form $m \times n$, $m, n \in \mathbb{Z}$, with edges joining a vertex $m \times n$ to the vertices $(m \pm 1) \times n$ and $m \times (n \pm 1)$).

Lemma 2.30. Let G be a finite non-degenerate PCR. Then the half-spaces of Dual(G) are precisely the subsets of G° of the form⁹

(4)
$$\mathfrak{h}(a;G) := \{ u \in G^{\circ} | a \in u \} , \quad a \in G.$$

In particular, subsets of G° of the form

(5)
$$\mathfrak{h}(S;G) := \{ u \in G^{\circ} | S \subseteq u \} = \bigcap_{a \in S} \mathfrak{h}(a;G), \quad S \subset \Sigma$$

are convex in Dual(G), for any $S \subseteq G$.

Definition 2.31. To simplify notation, we will abuse it in the following ways:

Writing h(a), h(S) without specifying G will henceforth refer to the subsets of H(Σ), those are h(a; O(Σ)) and h(S; O(Σ)), respectively.

• When S is explicitly known, $S = \{a_1, \ldots, a_k\}$, we will write $\mathfrak{h}(a_1 \cdots a_k; G)$ instead of $\mathfrak{h}(S; G)$ when convenient.

As a side note, observe that $\mathfrak{h}(S;G) = \mathfrak{h}(S) \cap G^{\circ}$, where $\mathfrak{h}(S)$ coincides with the vertex set of a face of the hamming cube $\mathbb{H}(\Sigma)$. In particular, presenting any subset of G° as a concept is equivalent to decomposing it as a union of convex subsets of G° .

Median Graphs. The two preceding lemmas are results of Dual(G) being a *median* graph [8, 49]:

Definition 2.32. A connected simple graph $\Gamma = (V, E)$ is said to be a *me*dian 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 med(u, v, w) – see Figure 1. For median graphs $\Gamma_i = (V_i, E_i), i = 1, 2, a$ *median morphism* of Γ_1 to Γ_2 is a map $f: V_1 \to V_2$ which preserves medians: f(med(u, v, w)) = med(f(u), f(v), f(w)).

Median graphs are a special subfamily of *median algebras*, [44, 45, 23, 5]. Some modern generalizations and applications may be found in [7]. A central result in Sageev-Roller duality, specialized here to the finite case, and reformulated for non-degenerate PCRs is:

⁹Note that $\mathfrak{h}(a^*; G) = G^{\circ} \smallsetminus \mathfrak{h}(a)$ for all $a \in G$, by Proposition 2.17.

Theorem 2.33. The dual Dual(G) of a finite non-degenerate PCR G is a finite median graph, with the median calculated according to the formula:

(6)
$$\operatorname{med}(u, v, w) = (u \cap v) \cup (u \cap w) \cup (v \cap w), \quad u, v, w \in G^{\circ},$$

and with intervals in Dual(G) calculated according to the formula:

(7)
$$I(u,v) = \{ w \in G^{\circ} | u \cap v \subseteq w \} = \{ w \in G^{\circ} | w \subseteq u \cup v \} .$$

Conversely, if Γ is a finite median graph then Γ is naturally isomorphic to $\text{Dual}(\mathcal{H}(\Gamma))$ by sending every vertex v to the *-selection of all half-spaces of Γ which contain v.

This result is the consequence of a very strong convexity theory:

Theorem 2.34 (Properties of median graphs, [43], section 2). Let $\Gamma = (V, E)$ be a finite median graph. Then:

- (1) Any family of pairwise intersecting convex sets has a common vertex;
- (2) Every convex set is an intersection of halfspaces;
- (3) For any convex subset $K \subset V$, the subgraph of Γ induced by K is a median graph;
- (4) For any convex $K \subset V$ and any $v \in V \setminus K$ there is a unique vertex $\operatorname{proj}_{K}(v) \in K$ at minimum hop distance from v;
- (5) For any convex $K \subset V$, the nearest point projection $\operatorname{proj}_{K}(\bullet)$ is a median preserving, distance non-increasing retraction of Γ onto its subgraph induced by K.

Property (1) is often referred to as the Helly property.

The Helly property is, perhaps, the most notable of the results stated above. In
our setting of PCR duals, it may be interpreted as guaranteeing the satisfiability
of any family of conjunctive monomials over
$$\mathbb{A}$$
 in which every pair is separately
satisfiable.

Convex hulls. Given the central role of half-spaces in the convexity theory of median graphs, a notion of the set of half-spaces dual to a given set of vertices is useful:

Definition 2.35. For $K \subset \mathbb{H}$, its dual set of halfspaces, K^{\sharp} , is defined to be the set of all $a \in \Sigma$ with $K \subseteq \mathfrak{h}(a)$.

An immediate corollary of Theorem 2.34(2) is:

Corollary 2.36. Suppose G is a non-degenerate PCR, and $K \subseteq G^{\circ}$. Then K^{\sharp} is G-coherent and forward-closed, and the convex hull $\operatorname{hull}(K;G)$ of K in $\operatorname{Dual}(G)$ coincides with $\mathfrak{h}(K^{\sharp};G)$.

Thus, every convex subset of G° may be written as $\mathfrak{h}(S;G)$ for some $S \in \mathbf{C}(G)$. This representation is unique, by last assertion of the following lemma:

Lemma 2.37. Let G be a non-degenerate PCR over Σ . Then, for all $S, S_1, S_2 \subset \Sigma$:

- (1) $\mathfrak{h}(S;G) \neq \emptyset$ if and only if S is coherent;
- (2) For all $S_1, S_2 \subseteq \Sigma$ one has $\mathfrak{h}(S_1 \cup S_2; G) = \mathfrak{h}(S_1; G) \cap \mathfrak{h}(S_2; G);$
- (3) If $S_2 \subseteq S_1$ then $\mathfrak{h}(S_1; G) \subseteq \mathfrak{h}(S_2; G)$;
- (4) If S is coherent then $\mathfrak{h}(S;G) = \mathfrak{h}(S\uparrow;G)$;
- (5) If $S \in \mathbf{C}(G)$, then $\mathfrak{h}(S; G) = \mathfrak{h}(\min(S); G)$
- (6) For all $S_1, S_2 \in \mathbf{C}(G)$ one has $\mathfrak{h}(S_1; G) \subseteq \mathfrak{h}(S_2; G) \Leftrightarrow S_2 \subseteq S_1$.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRIS

Proof. See Appendix C.1.

Another important result helps bound the distance from the points of one convex set to another:

Lemma 2.38. Let $S, T \in \mathbf{C}(G)$ for a poc set G over Σ . Then $\Delta(u, \mathfrak{h}(T)) \leq |T \setminus S|$ for all $u \in \mathfrak{h}(S)$.

Proof. See Appendix C.4.2. \Box

This motivates the following definition for the general case:

Definition 2.39. Let G be a non-degenerate PCR over Σ and let $S, T \in \mathbf{C}(G)$. The *divergence* of S from T is defined to be $\text{Div}(S;T) := |T \setminus S|$.

Note how Div(S;T) seems independent of G; it is not, however, since it is only applied to upwards-closed coherent sets S,T. We will use this notion of divergence in Section 5.3, to drive the decision-making mechanism of the *binary UMA agents* briefly introduced there.

More details about the convexity theory of a median graph will be discussed in the appendices, as we go about proving our algorithmic results.

2.5. **Propagation: A Computational Workhorse.** We are now ready to present another central result of this paper: a low-complexity method for computing nearest point projections in Dual(G), which we call *propagation*. This method obviates the need for maintaining an explicit representation of each vertex of Dual(G) in memory, reducing space requirements for this architecture from $O(2^{|\Sigma|})$ in the worst case to $O(|\Sigma|^2)$. The time complexity is, at worst, $O(|\Sigma|^2)$, coming down to sub-linear on a fully parallel architecture, as will become evident below.

Definition 2.40 (coherent projection). Let G be a PCR over a finite PCS Σ . For any $T \subseteq \Sigma$, the set $\operatorname{coh}_G(T) := T \uparrow \ T^* \downarrow = T \uparrow \ (T \uparrow)^*$ is said to be the *G*-coherent projection of T.

Coherent projection itself plays an important role in obtaining an observer's belief state from its epistemic state (the learned PCR structure) and the latest observation (see Section 3.3).

The promised formula for computing projections works as follows.

Proposition 2.41. Let G be a PCR over a finite PCS Σ . Let $S, T \subset \Sigma$ and suppose S is G-coherent. Let $L = \mathfrak{h}(S; G)$ and $K = \mathfrak{h}(\mathsf{coh}_G(T); G)$. Then:

(8)
$$\operatorname{proj}_{K}(L) = (S \cup T) \uparrow \smallsetminus (T \uparrow)^{*} = (S \uparrow \smallsetminus T \uparrow^{*}) \cup \operatorname{coh}_{G}(T),$$

where $\operatorname{proj}_{K}(\bullet)$ is the nearest-point projection to K in $\operatorname{Dual}(G)$ defined in Theorem 2.34.

Proof. See Appendix C.4. \Box

This description of nearest point projection is easy to visualize as being computed by an algorithm propagating excitation among nodes of a directed graph:

Definition 2.42. Let G be a PCR over a finite PCS Σ . Let $S \subset \Sigma$. Denote by [G, S] the graph with vertex set Σ , edge set G and with Boolean weights $\lambda(a) = \mathbb{1}_S(a), a \in \Sigma$ attached to its vertices. We refer to it as G being loaded with S. \Box

Algorithm 1 Propagating a signal T over [G, S] using depth-first search.

function $PROPAGATE(G, S, T)$	
$\texttt{visited} \gets \varnothing$	
$T \leftarrow \text{CLOSURE}(G, T)$	
$S \leftarrow \text{CLOSURE}(G, S)$	
$\mathbf{return} \ (S \cup T) \smallsetminus T^*$	
end function	
function $CLOSURE(G, T)$	\triangleright Forward closure of T in G
for all $a \in T$ do	
EXPLORE(G, a)	
end for	
return visited	
end function	
procedure $EXPLORE(G, v)$	\triangleright Recursive step
$\texttt{visited} \gets \texttt{visited} \cup \{v\}$	
for all $w \in \text{CHILDREN}(G,v) \smallsetminus \texttt{visited do}$	
EXPLORE(G, w)	
end for	
end procedure	
function $CHILDREN(G, v)$	\triangleright Children of v in G
$\mathbf{return}\{w\in \mathbf{\Sigma} vw\in G\}$	
end function	

Definition 2.43. A propagation algorithm over G is any algorithm which, for any G-coherent load $S \subset \Sigma$ and any $T \subseteq \Sigma$ accepts [G, S] and T as input and produces as its output the loaded graph $[G, \mathsf{PROP}(T, S; G)]$, where

(9)
$$\mathsf{PROP}(T,S;G) := (S \cup T) \uparrow \backslash T^* \downarrow .$$

Note that coherent closure is obtainable via $\operatorname{coh}_G(T) = \operatorname{PROP}(T, \emptyset; G)$.

Envisioning G as describing a graph of 'cells' labeled by Σ and 'synapses' labeled by pairs $ab \in G$, the loaded graph [G, S] represents a state of the network indicating that the cells of S are in an excited state. A propagation algorithm should be seen as exciting, additionally, the cells of T and spreading this excitation along the directed connections while inhibiting a^* for each cell $a \in \Sigma$ encountered along the way. Realized on a modern day computer, this may be achieved in quadratic time in $|\Sigma|$. For example, propagation could be implemented using a variant of depth-first search (DFS) on $\Gamma|_t$, while maintaining an expanding record of vertices visited [9] see Algorithm 1. On a fully parallel machine allowing the 'cells' to compute their own excitation, the time complexity is clearly of the order of the longest directed vertex path in the network, which is sub-linear in $|\Sigma|$.

We now turn to a high-level description of the UMA architecture and its use of the results of this section.

3. UNIVERSAL MEMORY ARCHITECTURE (UMA): A HIGH-LEVEL VIEW.

In this section we provide a high-level description of the basic UMA functionalities: PCR update/revision and maintaining a belief state. ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRI7

3.1. **Observation Model.** Recall from Section 2.1.1 that an observer is given a set A of initial Boolean queries over the space of histories **X** of the observed system. The system of queries and their complements is modeled as a PCS morphism ρ : $\Sigma(\mathbb{A}) \to \mathbf{2}^{\mathbf{X}}$, which is unknown to the observer. The observer is presented with a sequence of observations $\mathbb{Obs}|_t \in \mathbb{H}(\Sigma(\mathbb{A}))$, and values $\varphi|_t \in \mathbb{R}_{\geq 0}$, $t \geq 0$, one per update cycle. One must distinguish between two settings:

Static signal.: The value signal $\varphi|_t$ only depends on the raw observation $Obs|_t$; **Dynamic signal.:** The value signal may produce $\varphi|_t \neq \varphi|_s$ while $Obs|_t = Obs|_s$.

While ultimately interested in covering the dynamic setting, we will only deal with the static setting in this paper. However, the setting being *static* by no means implies it is *unchanging*. We will see in Section 5 that instances of the static setting may, nevertheless, have rich and interesting dynamics. This will happen, in part, as a result of introducing *delayed queries*. By these we mean the following: if $\sharp: \mathbf{X} \to \mathbf{X}$ denotes the operation of truncating the last state from a given history, then, for any conjunction *a* of already available queries it is possible to introduce a new query of the form¹⁰ $\sharp a$, where $\sharp a$ reports its value according to the rule $x \in \rho(\sharp a) \Leftrightarrow \sharp x \in \rho(a), x \in \mathbf{X}$. Of course, implementing this operation requires that the UMA architecture retain the latest raw observation, but this seems like a small price to pay for increasing the range of application of the static setting.

The basic task of an UMA is to evolve a sequence $G|_t$, $t \ge 0$ of non-degenerate PCRs over $\Sigma = \Sigma(\mathbb{A})$ while aiming for the PCRs $G|_t$ to eventually satisfy the following:

- 'Completeness': $\rho: G|_t \to \mathbf{2}^{\mathbf{X}}$ is a PCR morphism, ensuring that every perceptual class is represented;
- 'Precision': $\mathbf{M}|_t := \text{Dual}(G|_t)$ is as close as possible to the true model space $\mathbf{M} := \rho^*(\mathbf{X})$.

These requirements should not be taken literally, however. For example, it stands to reason that in some contexts the observer could afford to misclassify a few perceptual classes of low import. We will see how—at least under some of the learning schemes we propose—these vague requirements become possible to state precisely in terms of PAC learning.

3.2. Maintaining a PCR presentation: Snapshot Structures. A rather restrictive notion of a *snapshot structure*—a method for learning a poc set structure from positive observations—was introduced by the authors in [19]. Here we merely review the main ideas to provide intuition, while deferring the formal constructions to Section 4.

Snapshot weights. Motivated loosely by Hebbian ideas about learning [21], we consider maintaining an evolving symmetric system of weights $\mathbf{w}_{\bullet}|_{t} = (\mathbf{w}_{ab}|_{t})_{a,b\in\mathbf{\Sigma}|_{t}}$,

with $\mathbf{w}_{ab}|_{t}$ quantifying in some prescribed way a notion of *cumulative degree of relevance* of the event $\mathfrak{h}(ab)$ to the observer, at time t.

In addition, rules to maintain $\mathbf{w}_{\bullet}|_{t}$ as time progresses, must be provided. First, a completion rule, to insert missing values into \mathbf{w}_{\bullet} when it undergoes an extension. Second, an update rule, computing $\mathbf{w}_{\bullet}|_{t+1}$ from $\mathbf{w}_{\bullet}|_{t}$ and the incoming observation.

¹⁰Here and on we abuse notation, applying the symbol \sharp to denote both a delayed query and the history truncation/shift operator. Which is which is clear from the context.

It is important for both rules to be as simple—and as local—as possible, so as not to sacrifice tractability. In our constructions, we constrain the update laws to ones where $\mathbf{w}_{ab}|_{t+1}$ depends only on $\mathbf{w}_{ab}|_{t}$, the value signal $\varphi|_{t+1}$, the truth value of the bit $Obs|_{t+1} \in \mathfrak{h}(ab)$ and possible global parameters (e.g. the system clock t). PCRs from snapshot weights. Inspired by the rough mechanism proposed in [19], we seek weight systems \mathbf{w}_{\bullet} for which the loosely specified rule—

(10)
$$ab \in G \iff \mathbf{w}_{ab^*}$$
 is negligible in comparison with $\mathbf{w}_{ab}, \mathbf{w}_{a^*b^*}$,

ranging over all $a, b \in \Sigma$ with $\{a, a^*\} \neq \{b, b^*\}$ is guaranteed to define a *non*degenerate PCR over Σ . The motivation for the rule is, of course, the fact that $\rho(a) \cap \rho(b^*) = \emptyset$ is equivalent to $\rho(a) \subseteq \rho(b)$, where ρ is the PCS morphism defining the semantics of the queries in Σ .

Finally, note how the properties of a PCR are guaranteed (to the extent that the rule is well-defined, of course), and non-degeneracy is the only remaining question. Of course, the precise notion of 'negligible' defined for the purpose of comparing weights is crucial, and is expected to greatly affect the quality and limitations of the emerging representations.

3.3. Maintaining a Belief State. Since, for each time t, we only get to observe states from $\mathbf{M}(\rho)$, we are facing the problem of having to learn negative statements—that is, the list of $G|_{t}$ -incoherent pairs—from the stream of positive examples $(\mathbf{Obs}|_{t}, \varphi|_{t})$. From what we have observed so far we must reason about what it is we might never encounter. Seeing that the implication record $G|_{t}$ is inherently uncertain, providing no guarantee at any time that the completeness requirement from Section 3.1 will be met, it is quite possible for the observation $\mathbf{Obs}|_{t+1}$ to land outside the model space $\mathbf{M}|_{t+1} = \mathbf{Dual}(G|_{t+1})$ despite its prior role in forming this model space, during the snapshot update. In fact, its value may be too low to trigger a revision of $G|_{t}$ into a $G|_{t+1}$ for which $\mathbf{Obs}|_{t+1}$ becomes coherent.

Contrary to the approach adopted by modern iterated revision schemes based on Darwiche and Pearl's [11], we do not insist on a revision forcing $Obs|_{t+1}$ into $\mathbf{M}|_{t+1}$. Instead, we apply $G = G|_{t+1}$ to the raw observation with aim to relax it, replacing it with a G-coherent and forward-closed set:

(11)
$$\operatorname{Curr}_{t+1} := \operatorname{coh}_G(\operatorname{Obs}_{t+1}) \subseteq \operatorname{Obs}_{t+1},$$

in the role of the current state of record, or the belief state. This way, UMA naturally resolves possible contradictions at the price of introducing ambiguity into its record of the current state: instead of marking a single vertex of $\mathbf{M}|_{t+1}$ as the current state, any vertex of the convex set $\mathfrak{h}(\mathbf{Curr}^{t+1})$ may turn out to be the correct current state from the observer's point of view.

The choice of the coherent projection for the purpose of forming the belief state is motivated by its geometric and categorical properties. In our class of model spaces it is a *canonical* method of producing coherent sets, as witnessed by the following two results:

Proposition 3.1 (Coherent Approximation). Let G be a PCR over Σ . Then, for any $A \in \mathbb{H}(\Sigma)$, if $B \in G^{\circ}$ realizes the Hamming distance

(12)
$$\Delta(A, G^{\circ}) := \min_{u \in G^{\circ}} \Delta(A, u)$$

-that is, if $\Delta(A, B) = \Delta(A, G^{\circ})$ -then we must have $B \in \mathfrak{h}(coh_G(A); G)$.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRIN

Proof. See Appendix C.2.

Thus, the operation $\operatorname{coh}_{\bullet}()$ yields the "best approximation" of $\operatorname{Obs}|_{t+1}$ by a convex subset of $\mathbf{M}|_{t+1}$, echoing the principle of minimal change as seen through Dalal's way [10] of quantifying the distance between theories. Moreover:

Proposition 3.2 (Coherent Projection). Let G be a PCR over Σ . Then the following hold for all $A \subseteq \Sigma$:

- (a) $\operatorname{coh}_G(A)$ is coherent and $\operatorname{coh}_G(A)\uparrow = \operatorname{coh}_G(A)$;
- (b) $\operatorname{coh}_G(\operatorname{coh}_G(A)) = \operatorname{coh}_G(A);$
- (c) $A \subseteq \operatorname{coh}_G(A)$ whenever A is G-coherent;
- (d) $\operatorname{coh}_G(A) = A$ if and only if A is G-coherent and $A \uparrow = A$.

In other words, as a self-map of 2^{Σ} , the operator $A \mapsto \operatorname{coh}_G(A)$ is an idempotent whose image coincides with $\mathbf{C}(G)$.

Proof. See Appendix C.3.

Note how properties (a) and (c) turn $\operatorname{coh}_G(\bullet)$ into a closure operator on the subspace of *G*-coherent sets with respect to inference (implication). At the same time, (b) and (d) characterize the set $\mathbf{C}(G)$ of all terms that are closed under inference.

Overall, Equation (11) provides an intriguingly natural way of maintaining an internal model and belief state with a built-in degree of resilience to observations that fail to make immediate sense to the agent given its epistemic state. Finally, the complexity of this computation is the complexity of propagation over $G|_{t+1}$, by Proposition 2.41 and the discussion following Definition 2.43.

4. LEARNING ALGORITHMS FOR UMAS: SNAPSHOT STRUCTURES.

4.1. Qualitative Snapshot Structures. The goal of this section is to construct a snapshot structure suitable for a scenario in which the learner's value signal is a ranking function in the sense of Pearl [39, 18] (which is a special form of Spohn's OCFs [46]), its values providing a qualitative notion of the degree of irrelevance of the current experience. Thus, an observation with $\varphi|_t = 0$ is considered desirable, while $\varphi|_t = 1, 2, \ldots$ renders an observation increasingly more irrelevant.

4.1.1. Rankings and 2-rankings. Throughout this section we let Σ be a PCS and let \mathbb{H} denote the Hamming cube $O(\Sigma)^{\circ}$. Also, let $\widehat{\mathbb{N}} := \{0, 1, 2, \dots, \infty\}$. We use the slight variation of the notion of a ranking from [39], which was introduced in [11]:

Definition 4.1. A ranking on \mathbb{H} is a function $\kappa : 2^{\mathbb{H}} \to \widehat{\mathbb{N}}$, satisfying:

- $\kappa(F) = \min_{\sigma \in F} \kappa(\sigma)$ for all $F \subset \mathbb{H}$;
- $\kappa(\sigma) < \infty$ for some $\sigma \in \mathbb{H}$;
- $\kappa(\emptyset) = \infty$.

Hereafter, we shall abuse notation, writing $\kappa(\sigma)$ to mean $\kappa(\{\sigma\})$ whenever $\sigma \in \mathbb{H}$. Note that the minimum value of a ranking κ is $\kappa(\mathbb{H})$.

Remark 4.2. Note that, since Σ is assumed to be finite, the first requirement may be replaced with the requirement that $\kappa(F_1 \cup F_2) = \min{\{\kappa(F_1), \kappa(F_2)\}}$ for all $F_1, F_2 \subset \mathbb{H}$.

The simplest examples of rankings seem to be:

Example 4.3 (point-mass ranking). Let $u \in \mathbb{H}$ and $r \in \widehat{\mathbb{N}}$, $r \neq \infty$. Then the following function $\delta_{u,r}$ is a ranking:

(13)
$$\delta_{u,r}(F) := \begin{cases} r & \text{if } u \in F \\ \infty & u \notin F. \end{cases}$$

Example 4.4 (pointwise minimum). If κ_1, κ_2 are rankings on $\widehat{\mathbb{N}}$, then the function $\kappa(F) := \min(\kappa_1(F), \kappa_2(F))$ is also a ranking. \square

Recall now the sets $\mathfrak{h}(S)$ from Lemma 2.30. They will help us study the interaction between rankings and concepts:

Definition 4.5. The concept representation of a ranking κ , is the function $\mathbf{w}_{S}^{\kappa} :=$ $\kappa(\mathfrak{h}(S))$, where S ranges over subsets of Σ . To simplify notation, we will often write $\mathbf{w}_{a_1a_2\cdots a_k}^{\kappa} := \mathbf{w}_S^{\kappa}$ whenever $S = \{a_1, \ldots, a_k\}$ is explicitly provided. **Remark 4.6.** Note that $w_S^{\kappa} = \infty$ if S is not a *-selection. Also, $w_{\emptyset}^{\kappa} = \kappa(\mathbb{H})$, the minimum value of κ .

Lemma 4.7 (triangle inequality). For any ranking κ on \mathbb{H} , the following holds $\mathbf{w}_{ac^*}^{\kappa} \geq \min\left\{\mathbf{w}_{ab^*}^{\kappa}, \mathbf{w}_{bc^*}^{\kappa}\right\} \text{ for all } a, b, c \in \mathbf{\Sigma}.$

Proof. See Appendix D.1.

We are interested in studying the interactions between rankings on \mathbb{H} and nondegenerate poc-graph structures on Σ . A weakened notion of ranking is required for this purpose.

Definition 4.8. A 2-ranking on Σ is a symmetric matrix $\mathbf{w}_{\bullet} = (\mathbf{w}_{ab})_{a,b\in\Sigma}$ with entries in $\widehat{\mathbb{N}}$, satisfying the following for all $a, b, c \in \Sigma$:

- $\begin{array}{ll} (1) & \mathbf{w}_{\mathbf{0}a} = \mathbf{w}_{aa^*} = \infty; \\ (2) & \min(\mathbf{w}_{aa}, \mathbf{w}_{a^*a^*}) = \min(\mathbf{w}_{bb}, \mathbf{w}_{b^*b^*}) < \infty; \end{array}$
- (3) $\mathbf{w}_{aa} = \min(\mathbf{w}_{ab}, \mathbf{w}_{ab^*});$
- $(4) \ \mathbf{w}_{ac^*} \geq \min{(\mathbf{w}_{ab^*},\mathbf{w}_{bc^*})}.$

We will say that a ranking κ agrees with \mathbf{w}_{\bullet} , if $\mathbf{w}_{ab}^{\kappa} = \mathbf{w}_{ab}$ for all $a, b \in \Sigma$. Also, we will abbreviate as follows: $\mathbf{w}_a := \mathbf{w}_{aa}$ and $\mathbf{w}_{\varnothing} := \min(\mathbf{w}_a, \mathbf{w}_{a^*})$, for all $a \in \Sigma$. Finally, note how $\mathbf{w}_a = \min\{\mathbf{w}_{a0^*}, \mathbf{w}_{a0}\}$ must hold, too, for all $a \in \Sigma$, by virtue of requirements 1. and 3.

Of course, the idea is to have a 2-ranking play the role of a snapshot weight, from which one needs to derive a non-degenerate PCR. In our learning setting, the best one could do is to derive from the samples of the value signal φ the 2ranking $(\mathbf{w}_{ab}^{\varphi})_{a,b\in\Sigma}$. The main question is, then, how much of the original φ could be recovered from this information. The following family of PCRs helps answer this question:

Proposition 4.9. Suppose \mathbf{w}_{\bullet} is a 2-ranking, and let $\mathbf{w}_{\sigma} \leq \delta \in \mathbb{N}$. Consider the PCRs on Σ defined by:

(14)
$$ab \in \mathbf{Res}(\mathbf{w}_{\bullet}; \delta) \iff \{a, a^*\} \cap \{b, b^*\} = \emptyset \text{ and } \mathbf{w}_{ab^*} > \delta$$

for $\delta < \infty$, and by:

(15)
$$ab \in \operatorname{Res}(\mathfrak{w}_{\bullet};\infty) \iff \{a,a^*\} \cap \{b,b^*\} = \emptyset, and \mathfrak{w}_{ab^*} = \infty.$$

Then $\operatorname{Res}(\mathbf{w}_{\bullet}; \delta)$ is a non-degenerate PCR for all $\delta \in \mathbb{N}$.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETR21

Proof. See Appendix D.2.
$$\Box$$

A surprising consequence of the non-degeneracy of these PCRs is the following corollary, leading to the conclusion that every 2-ranking has a ranking that agrees with it:

Corollary 4.10. Let \mathbf{w}_{\bullet} be a 2-ranking on Σ , and let $a, b \in \Sigma$. Set $r := \mathbf{w}_{ab}$. Then there exists a vertex $u \in \mathbb{H}$ such that the point mass ranking $\nu = \delta_{u,r}$ satisfies $\mathbf{w}_{pq}^{\nu} \geq \mathbf{w}_{pq}$ for all $p, q \in \Sigma$.

Proof. See Appendix D.3. \Box

Proposition 4.11. Let \mathbf{w}_{\bullet} be a symmetric \mathbb{N} -valued matrix. Then \mathbf{w}_{\bullet} is a 2-ranking if and only if there exists a ranking with which it agrees. Moreover, if \mathbf{w}_{\bullet} is a 2-ranking, then there exists one and only one ranking,

(16)
$$\widehat{\mathbf{w}}(u) := \max_{a, b \in u} \mathbf{w}_{ab}$$

that agrees with \mathbf{w}_{\bullet} and satisfies $\widehat{\mathbf{w}} \leq \kappa$ for every ranking κ that agrees with \mathbf{w}_{\bullet} .

Proof. See Appendix D.4. \Box

The upshot of the last proposition is that, henceforth, any 2-ranking may be treated as encoding a ranking. Formally:

Definition 4.12. Suppose \mathbf{w}_{\bullet} is a 2-ranking and κ is a ranking. The *completion* of \mathbf{w}_{\bullet} is the ranking $\widehat{\mathbf{w}}$ from the preceding proposition. The 2-restriction of κ is the 2-ranking, denoted $\kappa^{(2)}$, obtained from κ via the concept representation, that is: $\kappa_{ab}^{(2)} := \mathbf{w}_{ab}^{\kappa}$ for all $a, b \in \Sigma$. The 2-closure of κ is the ranking, denoted $\widehat{\kappa}$, obtained from κ as the completion of its 2-restriction. In particular one has $\kappa \geq \widehat{\kappa}$.

4.1.2. Derived PCRs and their duals. We now introduce the PCR used in the qualitative snapshot structure. As systems of defaults, these PCRs are strengthened (more restrictive) versions of the (ranked) default systems constructed by Goldzmidt and Pearl in [18], and they satisfy an analogous characterization.

Proposition 4.13. Suppose w_{\bullet} is a 2-ranking. For $0 \le \delta < \infty$, let its derived PCR be defined by:

$$(17) \qquad ab \in \mathbf{Der}(\mathbf{w}_{\bullet}; \delta) \Leftrightarrow \left\{ \begin{array}{l} \{a, a^*\} \cap \{b, b^*\} = \varnothing, \ and \\ \mathbf{w}_{ab^*} = \infty \ or \ \mathbf{w}_{ab^*} > \delta + \max(\mathbf{w}_{ab}, \mathbf{w}_{a^*b^*}), \end{array} \right.$$

and for $\delta = \infty$ let it be defined by:

(18)
$$ab \in \mathbf{Der}(\mathbf{w}_{\bullet}; \infty) \Leftrightarrow \begin{cases} \{a, a^*\} \cap \{b, b^*\} = \varnothing, and \\ \mathbf{w}_{ab^*} = \infty and \mathbf{w}_{ab}, \mathbf{w}_{a^*b^*} < \infty \end{cases}$$

Then, $\mathbf{Der}(\mathbf{w}_{\bullet}; \delta)$ is a non-degenerate PCR for all $\delta \in \widehat{\mathbb{N}}$.

Proof. Let $G = \mathbf{Der}(\mathbf{w}_{\bullet}; \delta)$ and $R = \mathbf{Res}(\mathbf{w}_{\bullet}; \delta)$. Once again, the basic properties of a PCR are baked into the definition of G. Furthermore, observe that $ab \in G$ implies $ab \in R$ (though not the other way around). In particular, we have $N(G) \subseteq N(R)$ and it follows that $N(G) \cap N(G)^* \subseteq N(R) \cap N(R)^* = \emptyset$, as required. \Box

Definition 4.14. Proposition 4.11 and Definition 4.12 make it possible for us to abuse notation and talk about the residual and derived PCRs of a ranking by setting $\operatorname{Res}(\kappa; \delta) := \operatorname{Res}(\kappa^{(2)}; \delta)$, and $\operatorname{Der}(\kappa; \delta) := \operatorname{Der}(\kappa^{(2)}; \delta)$, dropping all mention of δ when $\delta = 0$, as before. Of course, κ may be replaced with its 2-closure $\hat{\kappa}$ throughout .

We proceed to study properties of derived PCRs and their duals, to verify their utility to our representation problem. Specifically, we are interested in the geometry of level sets, as we try to answer the question: how well does the 2-restriction of a ranking κ capture the set of global minimum points of κ (the most meaningful states according to κ)?

Definition 4.15. Given an integer $\epsilon \geq 0$ and a 2-ranking w_{\bullet} , denote:

(19)
$$\mathbb{M}(\mathbb{w}_{\bullet};\epsilon) := \left\{ a \in \mathbf{\Sigma} \, \middle| \, \mathbb{w}_a < \mathbb{w}_{a^*} - \epsilon \right\}$$

The set $M(w_{\bullet}) := M(w_{\bullet}; 0)$ will be referred to as the *minset of* w_{\bullet} . By virtue of Proposition 4.11, this notion extends to rankings as follows:

(20)
$$\mathbb{M}(\kappa;\epsilon) := \left\{ a \in \mathbf{\Sigma} \, \middle| \, \mathbf{w}_a^{\kappa} < \mathbf{w}_{a^*}^{\kappa} - \epsilon \right\} \,,$$

with $M(\kappa) := M(\kappa; 0)$ being the minset of κ .

It is clear that a global minimum point of a ranking κ must contain $M(\kappa)$. Hence, $\mathfrak{h}(M(\kappa))$ contains all global minima of κ , but what does this have to do with the derived PCR and its dual? The main result is as follows:

Proposition 4.16. Let κ be a ranking on \mathbb{H} and set $G = \mathbf{Der}(\kappa)$ and $M = \mathbb{M}(\kappa)$. Let F and \widehat{F} be the sets of global minima of κ and $\widehat{\kappa}$, respectively. Then $F \subseteq \widehat{F} \subseteq G^{\circ}$ and $\widehat{F} = \mathfrak{h}(M; G)$. Moreover, \widehat{F} is the convex hull of F in $\mathbf{Dual}(G)$

Proof. See Appendix D.5.

Upon inspection, the details of the proof generate the impression that $\hat{\kappa}$ is, for lack of a better word, a form of convex smoothing of κ , the last proposition showing how the collection of possibly disparate minimum points of κ coalesces into a convex plateau of minimum points of $\hat{\kappa}$ in the dual space of the derived PCR.

4.1.3. A Snapshot Structure to Learn a Ranking. We return to our learning problem. Suppose φ is a fixed ranking on \mathbb{H} , and we are given a sequence of samples $\varphi|_t = \varphi(\mathsf{Obs}|_t)$, where $\mathsf{Obs}|_t \in \mathbb{H}$ are the observations made by our agent. We will assume $\varphi|_t < \infty$ for all t, reserving $\varphi(u) = \infty$ for the impossible observations.

We must define the *weight update* taking place in response to an incoming observation; and the *weight extension* in response to a query being added to the sensorium.

Weight update (static case). For our snapshot structure, we propose the following update rule for the snapshot weights:

(21)
$$\mathbf{w}\Big|_{0} = \delta^{(2)}_{\mathbf{0}\mathsf{bs}}\Big|_{0}, \varphi(\mathbf{0}\mathsf{bs}}\Big|_{0}), \quad \mathbf{w}\Big|_{t+1} := \min\left\{\mathbf{w}\Big|_{t}, \delta^{(2)}_{\mathbf{0}\mathsf{bs}}\Big|_{t+1}, \varphi(\mathbf{0}\mathsf{bs}}\Big|_{t+1})\right\}.$$

By Example 4.4, ${\tt w}_{\bullet}\big|_t$ is a 2-weight for every $t\geq 0,$ giving rise to a non-degenerate PCR in the form of

(22)
$$G|_{\star} := \mathbf{Der}(\mathbf{w}_{\bullet}|_{\star}) \; .$$

Since the sequence of weights is pointwise non-increasing, its convergence is guaranteed. Moreover, exposure to (at most) $N = \binom{|\Sigma|}{2} - \frac{|\Sigma|}{2}$ observations covering all pairs $\{a, b\}$ with $\{a, a^*\} \cap \{b, b^*\} = \emptyset$, sampling a minimum rank world in $\mathfrak{h}(ab)$ for each pair $\{a, b\}$ at least once, will result in \mathbf{w}_{\bullet} coinciding with $\varphi^{(2)}$. This motivates the question "How much less exposure is required for delivering the same result on average, in, say, an appropriately formulated PAC setting?", and emphasizes the good fit of ranking-based snapshot structures to settings featuring a teacher.

4.2. Statistical Integrators of a Real-Valued Signal. The original suggestion of [19] for maintaining a system of weights in the role of a snapshot structure was based on the idea that $|\mathbf{w}_{ab}|_t$ should be the empirical estimate at time t of the probability of the event $a \wedge b$, so that $ab \in G|_t$ could be put on record if and only if $|\mathbf{w}_{ab*}|_t < \min(|\mathbf{w}_{ab}|_t, |\mathbf{w}_{a^*b^*}|_t, |\mathbf{w}_{a^*b}|_t, |\tau_{ab}|_t)$, where $\tau_{ab}|_t$ is a fixed threshold. That is, the implication $a \to b$ is put on record whenever the event $a \wedge \neg b$ has sufficiently low empirical probability. We have since found out that the improved formalization provided by Propositions 2.17 and 2.24 enables the use of a far more general weight update scheme that is capable of incorporating a value signal into the learner's reasoning while also taking into account the observed frequency of events.

4.2.1. Real-valued 2-weights. Once again, the learner is presented with a sequence of observations $u_t \in \mathbb{H}$, accompanied by the signal $\varphi|_t = \varphi(u_t)$. This time we require that the value signal $\varphi|_t$ presented to the agent at time t is a real number greater than or equal to 1, where a higher value of φ indicates a more meaningful state of the observed system.

Definition 4.17. A real-valued 2-weight on a PCS Σ is a symmetric, real-valued function $\mathbf{w}_{\bullet} = (\mathbf{w}_{ab})_{a,b\in\Sigma}$ on $\Sigma \times \Sigma$, satisfying the following requirements for all $a, b, c \in \Sigma$:

- (1) $\mathbf{w}_{ab} \ge 0$, $\mathbf{w}_{0a} = 0$ and $\mathbf{w}_{aa^*} = 0$;
- (2) $\mathbf{w}_{\varnothing} := \mathbf{w}_a + \mathbf{w}_{a^*} = \mathbf{w}_b + \mathbf{w}_{b^*};$

 $(3) \hspace{0.1cm} \mathtt{w}_{a}:=\mathtt{w}_{aa}=\mathtt{w}_{ab}+\mathtt{w}_{ab^{*}};$

- $(4) \hspace{0.1in} \mathtt{w}_{ab^{*}} + \mathtt{w}_{bc^{*}} + \mathtt{w}_{ca^{*}} = \mathtt{w}_{a^{*}b} + \mathtt{w}_{b^{*}c} + \mathtt{w}_{c^{*}a};$
- $(5) \ \mathtt{w}_{ac^*} + \mathtt{w}_{a^*c} \leq \mathtt{w}_{ab^*} + \mathtt{w}_{a^*b} + \mathtt{w}_{bc^*} + \mathtt{w}_{b^*c}.$

When $\mathbf{w}_{ab} = 0$ for all $a, b \in \Sigma$, we say \mathbf{w}_{\bullet} is trivial.

The following example provides motivation for the definition:

Example 4.18. Suppose $(\mathbf{X}, \mathscr{B}, \mu)$ is a measure space and $\varphi : \mathbf{X} \to \mathbb{R}$ is a nonnegative function in $L_1(\mu)$. Suppose $\rho : \mathbf{\Sigma} \to \mathscr{B}$ is a PCS morphism, when \mathscr{B} is viewed as a sub-PCS of $\mathbf{2}^{\mathbf{X}}$ (recall Example 2.3). Then $\mathbf{w}_{ab} := \int_{\rho(a) \cap \rho(b)} \varphi d\mu$ is a real-valued 2-weight. Indeed, since the integral of a non-negative function is non-negative, the requirements 1.-5. become corollaries of various set-theoretic identities applied to $A = \rho(a), B = \rho(b)$ and $C = \rho(c)$, respectively:

- (1) $\varnothing \cap A = \varnothing, A \cap (\mathbf{X} \setminus A) = \varnothing.$
- (2) $\mathbf{X} = A \cup (\mathbf{X} \smallsetminus A) = B \cup (\mathbf{X} \smallsetminus B),$
- $(3) A = (A \cap B) \cup (A \setminus B),$
- (4) $(A \smallsetminus B) \cup (B \smallsetminus C) \cup (C \smallsetminus A) = (B \smallsetminus A) \cup (C \smallsetminus B) \cup (A \smallsetminus C)$ (see Figure 2),
- (5) $A \vartriangle C \subseteq (A \bigtriangleup B) \cup (B \bigtriangleup C),$

where $A \bigtriangleup B := (A \smallsetminus B) \cup (B \smallsetminus A)$, for short.



FIGURE 2. The set-theoretic identity underlying requirement 4. of a real-valued 2-weight (Definition 4.17), as explained in Example 4.18.

Example 4.19 (point mass weight). Similarly to the qualitative setting, the simplest example of a weight of this form is given by a point-mass measure on \mathbb{H} :

(23)
$$\mathbf{w}_{ab} = r \cdot \delta_u(\mathfrak{h}(ab)), \quad \delta_u(F) := \begin{cases} 1 & \text{if } u \in F \\ 0 & \text{if } u \notin F, \end{cases}$$

where $F \subset \mathbb{H}$ (Compare with Example 4.3).

4.2.2. Derived PCRs and their duals. The resulting notion of a derived PCR requires a system τ_{\bullet} of threshold values, denoted $\tau_{ab}|_{t} \in (0, 1)$, $a, b \in \Sigma$, satisfying the identities

(24)
$$\tau_{ab}\Big|_{t} = \tau_{ba}\Big|_{t} = \tau_{a^{*}b}\Big|_{t}$$

for all $a, b \in \Sigma$ and $t \ge 0$. This makes it possible to construct a non-degenerate PCR as follows:

Proposition 4.20. For any choice of threshold values τ_{\bullet} satisfying Equation (24), if w_{\bullet} is non-trivial, then

(25)
$$ab \in \mathbf{Der}(\mathbf{w}_{\bullet}, \tau_{\bullet})|_{t} \stackrel{def.}{\iff} \begin{cases} \mathbf{w}_{ab^{*}}|_{t} < \min(\tau_{ab} \cdot \mathbf{w}_{\varnothing}, \mathbf{w}_{ab}, \mathbf{w}_{a^{*}b^{*}}, \mathbf{w}_{a^{*}b}) \\ or \\ \mathbf{w}_{ab^{*}} = \mathbf{w}_{a^{*}b} = 0. \end{cases}$$

defines a non-degenerate PCR.

Proof. See Appendix E.1.

Let $G = \mathbf{Der}(\mathbf{w}_{\bullet})$ for a real-valued 2-weight \mathbf{w}_{\bullet} . A notion analogous to that of a minset may be considered in the real-valued setting, taking into account the reversal of the value hierarchy (now, bigger values of φ are considered the most significant):

(26)
$$\mathbf{M}(\mathbf{w}_{\bullet}) := \{ a \in \mathbf{\Sigma} \, | \mathbf{w}_a > \mathbf{w}_{a^*} \} ,$$

The argument that $M = M(\mathbf{w})$ is *G*-coherent and forward-closed, for any choice of the thresholds τ_{\bullet} , is the same as the one given for minsets in the qualitative setting (Lemma D.1 in Appendix D.5), upon reversing the relevant inequalities. This time around, however, the non-empty convex subset $F = \mathfrak{h}(M; G)$ of $\mathtt{Dual}(G)$ does not directly relate to extreme points of the value signal φ in \mathbb{H} , but, rather, to a notion of *center of mass* of G° with respect to \mathbf{w}_{\bullet} , seen as a representation of the distribution of the value signal over \mathbb{H} .

4.2.3. Snapshot update. Similarly to the qualitative setting, in the real-valued setting we will also be assembling our estimate of the [integrals of the] observed value signal from point-masses, this time replacing minimization with linear combinations. The update rule for a *discounted integrator snapshot* takes the form:

(27)
$$\begin{cases} \mathbf{w}_{ab}\big|_{o} := \varphi\big|_{o} \cdot \delta_{u_{0}}(\mathfrak{h}(ab)), \\ \mathbf{w}_{ab}\big|_{t+1} := q\big|_{t} \cdot \mathbf{w}_{ab}\big|_{t} + (1-q\big|_{t}) \cdot \varphi\big|_{t+1} \cdot \delta_{u_{t+1}}(\mathfrak{h}(ab)), \end{cases}$$

where the $q|_t \in (0,1]$ are the discount coefficients, $t \ge 1$. The fact that $w_{\bullet}|_{t+1}$ is a convex combination of real-valued 2-weights ensures that $w_{\bullet}|_{t+1}$ is a real-valued 2-weight as well.

Types of update. We studied two variants of the discounted integrator snapshot:

(1) Empirical Snapshot. In this case, one sets $q|_t := \frac{t+1}{t+2}$, resulting in

(28)
$$\mathbf{w}_{ab}\big|_{t} = \frac{1}{t} \sum_{s=0}^{t} \varphi\big|_{s} \cdot \delta_{u(s)}(\mathfrak{h}(ab)) \,,$$

which is the empirical estimate for the integral of φ over $\rho(a) \cap \rho(b)$. For this snapshot type, we used fixed thresholds $\tau_{ab} = \tau$.

(2) Fixed Discount Snapshot. Here one sets $q|_t := q$, a constant, playing the role of a rate at which information acquired about the signal 'fades' unless continually reinforced by incoming observations:

(29)
$$\mathbf{w}_{ab}\big|_{t} = q^{t}\varphi\big|_{0} \cdot \delta_{u(0)}(\mathfrak{h}(ab)) + (1-q)\sum_{s=1}^{t} q^{t-s}\varphi\big|_{s} \cdot \delta_{u(s)}(\mathfrak{h}(ab)) \,.$$

The eventual purpose of using an update of this form is to accommodate settings where φ has multiple peaks, as well as, possibly, the dynamic setting, provided the value signal changes sufficiently slowly.

PAC learning guarantees. The notion of probably approximately correct (PAC) learning introduced by Valiant [48] is one framework within which the quality of UMAs based on real-valued snapshots could be discussed. The assumptions of this setting are that the observations $(u_t)_{t\geq 0}$ are i.i.d. samples of a fixed distribution on \mathbb{H} , in which case, for any fixed pair $a, b \in \Sigma$ with $\{a, a^*\} \cap \{b, b^*\} = \emptyset$, one could think of the sequence of input values $X_{ab}|_t := \varphi|_t \cdot \delta_{u_t}(\mathfrak{h}(ab))$ as a sequence of i.i.d. samples of a random variable $X_{ab} \in [0, A]$, where A is an upper bound on the value signal φ . Equation Equation (27) then lets us think of $\mathbf{w}_{ab}|_t$ as random variables $Y_{ab}|_t$ constructed according to $Y_{ab}|_{t+1} = q|_t Y_{ab}|_t + (1 - q|_t) X_{ab}|_{t+1}$. Applying induction one immediately verifies that $\mathbb{E}[Y_{ab}|_t] = \mathbb{E}[X_{ab}]$ for all $t \geq 0$. It thus becomes reasonable to ask how many samples are required in order to bring the probability that $|Y_{ab}|_t - \mathbb{E}[X_{ab}]| > \varepsilon$ below a specified threshold. Valiant [48] had long ago observed that Chernoff bounds are a powerful tool for answering such questions. Computing Chernoff bounds for our setting yields:

Proposition 4.21 (PAC learning in empirical snapshots). Given $\delta > 0$, the empirical snapshot learning mechanism attains a precision of δ on all weights, with probability $1 - \delta$ from a number of *i.i.d* randomized samples that is at most linear in $\frac{1}{\delta}$, at a rate depending only on the value signal.

Proof. See Appendix E.2.

Our simulation results indicate that similar guarantees could be expected for the discounted setting, but the standard Chernoff-inspired approaches for leveraging the independence of the observations do not seem to work. Since discounted snapshot learning makes it easier for the representation to recover from false implications, it is important to ascertain whether or not a result of the form Proposition 4.21 could be proved, and if not—in what circumstances it might fail.

Other learning scenarios. The PAC learning guarantees of the preceding paragraph are predicated on the assumption that the sequence of observations is statistically independent. This assumption becomes unreasonable for an observer of a system whose state evolves continuously over time, subject to some internal dynamics, in which case it is often unlikely that contiguous observations will be uncorrelated.

A fairly general model of such settings is provided by Markov chains [42], where the underlying Markov process models the (uncertain) dynamics of the observed system. In our setting, one regards $\rho^*(\mathbf{X})$ —the set of observable possible worlds in $\mathbb{H}(\mathbf{\Sigma})$ (Section 2.1.1)—as the set of states of a fixed (albeit unknown) Markov process. Then, by the ergodic theorem for Markov chains [15], one has:

Proposition 4.22. Suppose the sequence of observations $u_t \in \mathbb{H}$ is sampled from an a-periodic, irreducible, positive-recurrent Markov chain with limiting distribution π . Then the empirical snapshot weights $\mathbf{w}_{ab}|_t$ learned from the constant value signal $\varphi|_t = 1$ converge to the marginals $\int_{\rho(a)\cap\rho(b)} d\pi$, for all $a, b \in \Sigma$.

In particular, any thresholded implications derived from the real-valued 2-weight $\mathbf{w}_{ab} := \int_{\rho(a) \cap \rho(b)} d\pi$ will be recovered in this process.

Finally, it follows from the decomposition theorem for Markov chains [15] that the ergodicity assumption in the above proposition does not impose undue restrictions on our model, as we only expect an agent to learn implications from recurring observations anyway. We also note that the special case of lazy random walks guarantees an exponential rate of convergence to the limiting distribution in many interesting cases (see Theorem 5.1 of [32] and Theorem 9 of [42]).

5. SIMULATIONS.

We present two kinds of simulation studies. Section 5.2 illustrates the preceding results about learning with different snapshot types in a sample of 'toy' settings. Section 5.3 explains how to construct simple UMA-based binary agents, whose performance is considered in Section 5.4.

5.1. Simulation settings. Each setting considered in Section 5.2 consists of an observer/agent \mathscr{A} situated in a discrete environment, **E**. For simplicity, the queries assigned to \mathscr{A} are functions of the agent's current position in the environment, which we denote by $pos(t) \in \mathbf{E}$. Let $[N] := \{0, \ldots, N\}$. The environments and sensory endowments we consider are:

- Discretized interval with GPS. Here $\mathbf{E} = [N]$, and \mathscr{A} has queries $\mathbb{A} = \{a_1, \ldots, a_N\}$, with a_i holding true at time t iff pos(t) < i;
- Discretized circle with beacons. Now set $\mathbf{E} = [N-1]$ with a_i (i = 0, ..., N-1) holding true iff pos(t) is close enough to *i*, modulo N;
- Discretized interval with random position sensors. $\mathbf{E} = [N]$ again, and $\mathbb{A} = \{a_1, \ldots, a_N\}$, with a_i true at time t iff $pos(t) \in A_i$, where $A_i \subseteq \mathbf{E}$ are chosen uniformly at random ahead of each simulation run.

We consider different value signals, all set to be functions of the position, depending on snapshot type:

• Qualitative Snapshots. Two natural choices of the signal are considered,

(30)
$$\varphi_d(p) := \begin{cases} 0 & \text{if } p = T, \\ 1 & \text{if } p \neq T, \end{cases} \text{ and } \varphi_s(p) := \texttt{dist}(p, T) ,$$

where T should be regarded as a "target" position of high significance.

• **Real-valued Snapshots.** To parallel the "sharp peak"/"dull peak" signal variants from the qualitative setting, we pick:

(31)
$$\varphi_d(p) := 1 + \operatorname{diam}(\mathbf{E}) - \operatorname{dist}(p, T)$$
, and $\varphi_s(p) := \varphi_d(p)^4$,

respectively. For discounted snapshots, the discount coefficients were picked to be q = 0.999. Learning thresholds are constant, where relevant, and are chosen to equal $\frac{1}{2N}$ to ensure correct learning of implications among the initial sensors by the real-valued snapshots.

5.2. Simulation results for observers. To assess the speed and quality of PCR learning, we track the error-rate of the learned PCR representation—the fraction of correctly learned PCR implications—over time.

5.2.1. Repeated i.i.d. sampling (PAC-style setting). Figure 3 compares logarithmic plots of two mean error rates over 100 observation sequences generated by repeated i.i.d. uniform sampling of positions from the environment, for the settings described in Section 5.1 for N = 20:

- (1) **Solid lines.** The mean fraction of incorrect implications in the learned PCR relative to the expected PCR for the given learner in each setting, as a function of time;
- (2) **Dashed lines.** The mean fraction of incorrect implications in the transitive closure of the learned PCR relative to the poc set of actual implications among the provided sensors, as a function of time;
- (3) **Shaded regions** depict the mean±standard deviations for the corresponding quantities.

The first most notable feature of the figures—beyond confirming (and, in fact, exceeding) the theoretical results—is the complete agreement of the curves for all six learners on the interval (left column). Since the sensors in this case are nested, the poc set of true implications coincides with the derived PCR induced by the expected weights and is recovered quickly and completely.

Next, on the circle we begin to see the difference between the quality of the learned PCR and the quality of the inferred system of implications as compared to the real ones. This deterioration in quality was to be expected, as transitive closure enables the deduction of implications from chains of *approximate* implications recorded in the PCR. Observe that the discrepancy is bigger for the sharp peak settings, in which a very small degree of significance is assigned to positions farther away from the target. This difference is most notable in the qualitative learners: while completely absent in the dull peak setting, it is very visible in the sharp peak setting. We account for these differences, among other things, in the detailed analysis of the true PCR provided in Appendix F.

A similar discrepancy is visible, but less pronounced in the third column, though we must keep in mind that, in this column, each run was executed with a different



FIGURE 3. Evolution over time of the log mean error rates (curves in bold, logarithmic scale), plus/minus standard deviation (shading), over 100 runs, of PCR representations acquired by i.i.d. uniform random sampling, for different snapshot types: empirical (red), discounted (green), qualitative (blue). Columns, left to right: interval w/GPS sensors; circle w/beacon sensors; and interval w/random sensors. Rows, top to bottom: φ_1 ("dull peak" value signal), φ_2 ("sharp peak" value signal), as in Equations (30) and (31). See Section 5.2.1 for more details.

random collection of sensors. The differences are less pronounced than on the circle because in the sensorium we have chosen for the circle there is very little nesting, while in a random sensorium, the probability of nesting is non-negligible. Nesting relations $\rho(a) \subset \rho(b)$ in the sensorium forces $\mathbf{w}_{ab^*}|_t = 0$ in real-valued snapshots, and $\mathbf{w}_{ab^*}|_t = \infty$ in qualitative snapshots at all times t, guaranteeing that a < b will be learned with sufficient exposure. The rotation-invariant sensorium we chose for the circle has very little nesting, and hence much more room for error if the provided value signal happens to discount too many positions as being insignificant. Deeper differences arise as a result of the circle's non-trivial homotopy type, which we discuss in Section 5.4.2 and further in Section 6.

Finally, let us remark that we do not yet have a good explanation for the good behavior of the discounted learners. We were unable to prove any concentration inequalities for the discounted weight update to parallel the ones obtained for the empirical one. Moreover, the quality of learning appears to be very sensitive to the choice of discount parameter. In fact, it was this difficulty with appropriately selecting and controlling the discount parameter that motivated the construction of qualitative learners in the first place. 5.2.2. Lazy random walk (learning from "motor babble"). For a robotic system, a more realistic mode of sampling from the environment is "motor babble": a random walk on **E** generated by repeated i.i.d. sampling from the space of available actions/decisions. In this mode, each instance of the agent \mathscr{A} is constrained to a small set of available actions, depending on **E**:

- Discretized interval. The allowed actions are a single step to the right $(rt: p \mapsto \min\{N, p+1\})$, a step to the left $(lt: p \mapsto \max\{p-1, 0\})$, or to remain in place;
- Discretized circle. Similarly, on the circle rt: p → p + 1(mod N), or lt: p → p - 1(mod N) or to do nothing at all.

Figure 4 shows the evolution of the error rates we had considered earlier in Section 5.2.1, in the new sampling mode.



FIGURE 4. Evolution over time of the mean error rates (curves in bold, logarithmic scale), plus/minus standard deviation (shading), over 100 runs, of PCR representations acquired by a lazy walk over the environment, for different snapshot types: empirical (red), discounted (green), qualitative (blue). Columns, left to right: interval w/GPS sensors; circle w/beacon sensors; and interval w/random sensors. Rows, top to bottom: φ_1 ("dull peak" value signal), φ_2 ("sharp peak" value signal). See Section 5.2.2 for details.

This set of plots provides a good illustration of the robustness of UMA learning especially with qualitative snapshots—where the quality of learning improves over time (though now at a much slower pace, due to the change in the sampling process), as the observer gains more exposure to the observed system.

5.2.3. Learning the target set over time. We compare how UMA learners of different snapshot types develop their notion of the target set, $M(w_{\bullet})$, over time. For this

purpose, Figure 5 shows this evolution for a single run from a separate batch of lazy random walk observations in a smaller environment (N = 10), over a shorter period of time (500 cycles). The features observed in this plot are, however, typical of the runs we generated for Figure 4. "Downgrading" the experiment to a smaller environment enabled faster learning, and hence plotting the run at a lower resolution, without requiring the reader to magnify the plot attempting to discern its significant features.



FIGURE 5. Evolution of an UMA observer's notion of target set (vertical bars drawn in grey) towards the true target (pink bar), in the interval w/GPS setting, N = 10. The agent's position (red dot) evolves as a symmetric random walk. See Section 5.2.3 for details.

Observe the eventual precision and efficiency of the qualitative reasoners, compared to the drift (away from the target) clearly noticeable for the real-valued learners. Also note some initial delay in learning the target (in comparison with other types) in the discounted learners: the value of q places a bound on how quickly an implication may be learned.

Both these observations are typical of all the batches we have observed. This suggests the qualitative UMA learners as the best bet for upgrading UMAs to perform learning in the dynamic setting. This also suggests that the real-valued learners could benefit from more careful shaping of the value signal, with significantly sharper peaks, as well as from lower values of the discount parameter (for discounted learners), if learning on shorter time scales is important.

5.3. Binary UMA agents. Postponing a more general formal definition of a binary UMA agent to another paper, let us describe just the simple sub-class of these agents considered here.

Actions as agents. Given the environment **E** and the associated set of queries A as described above in Section 5.1, we regard each of the actions α available to \mathscr{A} as an individual agent \mathscr{A}_{α} , in charge of making the decision whether to act (α) or not to act (α^*). Any conflicts between decisions made by different \mathscr{A}_{α} are, at this stage of development, arbitrated by hard-wiring (see example in Section 5.4.1 below).

Extended query set. For \mathscr{A}_{α} to be capable of considering the consequences of its decisions, we have to extend \mathbb{A} so as to enable reasoning about the past. Specifically, each \mathscr{A}_{α} is assigned a value signal φ_{α} , and an initial set of queries $\Sigma_{\alpha} := \Sigma(\mathbb{A} \cup \sharp \mathbb{A})$, where (\sharp) is the *delay operator*: the query $\sharp q$ holds true at time t + 1 if and only if q held true at time t.

UMA representation conditional on action. The BUA \mathscr{A}_{α} maintains *two* snapshots, $\mathbf{w}^{\beta}_{\bullet}, \beta \in \{\alpha, \alpha^*\}$. The 2-weight $\mathbf{w}^{\beta}_{\bullet}$ is updated precisely in those transitions in which \mathscr{A}_{α} acted according to β . Thus, at any time $t, \mathbf{w}^{\beta}_{\bullet}|_{t}$ may be used to infer implications conditioned on β taking place, by computing a derived graph, $G^{\beta}|_{t}$.

Prediction. Given the current state $\operatorname{Curr}^{\beta}|_{t}$ at time t as represented by the β snapshot, $\beta \in \{\alpha, \alpha^*\}$, the prediction for time (t+1) given β , $\operatorname{Pred}^{\beta}|_{t+1}$, is defined to be the coherent projection of $\sharp\operatorname{Curr}^{\beta}|_{t}$ with respect to $G^{\beta}|_{t}$. This is the collection of sensations which \mathscr{A}_{α} can prove will occur if β is chosen to take place, provided, of course, $G^{\beta}|_{t}$ persists into the (t+1)-st cycle.

Decision. At the same time, each of the agent's two snapshots has a notion of where it is that the agent *should be*: the subset $\mathfrak{h}(\mathfrak{M}(\mathfrak{w}^{\beta}_{\bullet}); G^{\beta}|_{t})$. A simple way for \mathscr{A}_{α} to make a choice of $\beta \in \{\alpha, \alpha^{*}\}$ is to pick the value of β for which $\mathsf{Div}(\mathsf{Pred}^{\beta}|_{t+1}; \mathfrak{M}(\mathfrak{w}^{\beta}_{\bullet})|_{t})$ is smaller, and to flip an even coin in the case of a tie (recall Definition 2.39).

5.4. Simulation results for agents.

5.4.1. Sniffy: locating a stationary target using "place field" sensors. Consider an agent \mathscr{A} in one of the two fixed settings described above in Section 5.1, with two actions rt and lt, as defined in Section 5.2.2, implemented as BUAs according to Section 5.3 with the value signals given in Equations (30) and (31). To minimize interference between rt and lt, we impose a hard-wired arbitration mechanism: if lt and rt decide to act at the same time, a Bernoulli $-\frac{1}{2}$ random trial decides which one of them to suppress.

At the beginning of each simulation run, Sniffy (our pet agent \mathscr{A}) and its target are placed in random positions in **E**, denoted pos(0) and *T*, respectively. The agent then experiences a training period during which every decision by every BUA is overridden by a random one, resulting in a lazy random walk. Once the training period is over, the BUAs are given control authority, with Sniffy acting according to their decisions.

Finally, following the indications of Section 5.2.3, we have chosen to replace the discount parameter of q = 0.999 with $q = 1 - \frac{1}{N+1}$, to enable a faster response by the discounted learners.

Figure 6 reports the results of our simulations. Each plot shows the mean, plus/minus standard deviation, over 100 distinct runs, of the distance of the agent to its target as a function of time, in each setting.

Appendix A.2.4 discusses the representations expected to arise in the case of the interval in some detail, explaining Sniffy's success in that environment, shown in the figure. However, we also notice a deterioration of the results as Sniffy is moved from



FIGURE 6. BUA implementations of Sniffy (Section 5.4.1) learning to find a target, with different snapshot types. Each graph shows the evolution of Sniffy's mean distance to the target, plus/minus std. deviation, over 100 runs.

the interval to the circle. This is due to subtle interactions between the propagation mechanism generating the BUAs' predictions (which drives decision-making), and the non-trivial homotopy type of the circle, which forces inconsistent states into all the model spaces involved (the latter, we recall, are always contractible). This discrepancy between the topology of UMA model spaces and the spaces they come to model provides the main motivation for our future project of studying the control of situated agents by *networks* of BUAs, where the deliberation among agents is meant to generate an emergent joint representation of reactive behavior patterns with the competence to overcome topological constraints and obstacles (more in Section 6).

5.4.2. What did Sniffy learn on the circle? All the graphs in Figure 6 indicate a significant change of behavior at the end of training. It therefore seems sensible to attempt splitting the set of runs in each setting into those finishing closer to the target than to its antipodal point on the circle, as shown in Figure 7. What emerges is that all the learned representations experience difficulties dealing with the situation loosely characterized as "Sniffy approaches the point on the circle antipodal to the target". Note that the "dull peak" qualitative learners emerge as the most apt, both in terms of efficiency and in terms of separation between the desirable and undesirable modes of behavior. In this setting, the target clearly emerges as an attracting point except for a small neighbourhood of its antipode, which seems to play the role of an unstable equilibrium. This is reminiscent of gradient descent over the function $f(x) = \operatorname{dist}(x, T)$ on the unit circle, viewed as a differentiable manifold: the target T is a robust attractive equilibrium, complemented by an unstable equilibrium that is forced by the non-trivial homotopy type of the circle.



ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETR33

FIGURE 7. Splitting the population of Sniffy's runs on the circle according to their termination points: red paths terminate inside $\rho(a_T)$; blue paths terminate outside of $\rho(a_T)$.

Since qualitative snapshots enable direct computation of the eventual values of the snapshot weights, it becomes possible to obtain explicit insights into the behavior learned by Sniffy in this setting. We refer the reader to Appendix F for a detailed discussion proving the preceding claims.

6. DISCUSSION.

Motivated by the goal of implementing well-reasoned general learning on mobile robots, this paper introduces algorithms implementing a simplified version of iterated belief revision and update that is consistent with budgetary constraints on storage space and computational complexity, collectively named "universal memory architectures" (UMAs). We establish and study the mathematical language necessary for the analysis of UMA instances, and show how the standard model-theoretic approach to belief revision gets naturally replaced by the study of the geometry of convex sets in the model spaces represented by UMAs.

By construction, UMA representations are systems of default rules that are closed under counter-positives. We show that such representations may be learned both by means of sampling and statistical integration of a real-valued signal (empirical and discounted snapshots, Section 4.2), as well as by means of aggregating samples of a ranking function on the space of possible worlds, in the sense of Spohn [46] and Pearl [39] (qualitative snapshots Section 4.1). In the latter case, we are able to guarantee the correct encoding of the convex hull, in the learned geometry, of the set of minimum rank worlds, provided sufficient exposure. Finally, we show the potential of UMA representations for the motivating application by considering its behavior in a pair of simple learning settings simulating a standard task formulation from Robotics: localize a target in the presence of (highly impoverished) sensing in a global frame (Section 5.4.1).

The need for expanding the set of queries ('self-enrichment'). It is important to state clearly the limitations of UMA learners in the form presented in this paper. From a practical perspective, attempting to learn a PCR structure for a fixed sensorium will yield *no learning at all* in the case of an arbitrary and/or 'unstructured' binary sensorium such as the pixel grid of a B/W video camera, where no two pixels are

a-priori correlated. Consider an even simpler example: the situation of $a, b, c \in \Sigma$ satisfying $\rho(a) \cap \rho(b) \subseteq \rho(c)$, $\rho(a) \notin \rho(c)$ and $\rho(b) \notin \rho(c)$ cannot be encoded by a PCR unless the query set Σ explicitly contains an element whose realization is $\rho(a) \cap \rho(b)$. Finally, it is clear that PCRs are not geared for studying temporal interactions unless explicitly outfitted with appropriate queries (as in the example of BUAs in Section 5.3).

Accepting the above as the price of the radical reduction in computational costs achieved by UMA-based learning (as compared to unrestricted iterated belief revision), a natural avenue for increasing the descriptive power of an UMA representation is to allow the set of queries Σ to expand (by adding 'meaningful' queries) and contract (by coalescing related queries, or deleting uninformative ones) over time, in a controlled fashion, at a known and minimal cost in computational resources. The fixed sensorium Σ should be replaced with a sequence $\Sigma|_t$, as the map $\rho: \Sigma \to 2^X$ is replaced with a sequence $\rho|_t: \Sigma|_t \to 2^X$. Still, the advantage of UMA representations over others is in their efficiency at encoding a model space and reasoning about it in terms of its convex subspaces. This motivates the search for an enrichment method that meets the lower complexity bound for representing the observed system.

Looking for such a method, one must be mindful that the expansion steps cannot be arbitrary, as it is necessary for each map $\rho|_{t+1} : \Sigma|_{t+1} \to 2^X$ to be uniquely determined by its predecessor $\rho|_t$ and the limited information that was available to the UMA at time t. This suggests two natural elementary expansion operations, which also happen to interact well with our detailed understanding of the geometry of duals:

- **Append a conjunction.:** Adding a query of the form $q = a_1 \land \dots \land a_k$ for some $a_1, \dots, a_k \in \Sigma|_t$, to form $\Sigma|_{t+1} = \Sigma|_t \cup \{q, q^*\}$, forces the extension of ρ via $\rho(q) = \rho(a_1) \cap \dots \cap \rho(a_k)$.
- Append a delayed sensor.: Let $\sharp : \mathbf{X} \to \mathbf{X}$ denote the operation of truncating the last state from a given history; Then it is possible to introduce a query of the form $q = \sharp a$ for $a \in \Sigma$, where $\sharp a$ reports the value of a preceding the current one, or, in other words: $x \in \rho(\sharp a) \Leftrightarrow \sharp x \in \rho(a)$.

Observing $\sharp(a \wedge b) = \sharp a \wedge \sharp b$ for all $a, b \in \Sigma$, we conclude that any composition of the above extension operations determines a unique extension of the original ρ . Hence, an UMA endowed with these enrichment operations is capable, in principle, of eventually representing very rich theories of the observed system, both in terms of Boolean relations among the original sensors and in terms of temporal properties provided we are willing to accept the cost in resources. Clearly, the burden is on us to decide when an extension is in order; for what purpose; and how to prevent the population of added sensors from exploding to a prohibitive size.

In the presence of delayed queries, the situation lends itself to the formation of a prediction operator, extending the simplistic one constructed in Section 5.3. This makes it possible to formulate learning objectives concerning the quality of prediction. Our ongoing work exploring analogies with perceptron learning [37] is directed towards studying the problem of optimizing prediction through gradual extension of the sensorium using the operations just formulated.

Agents. The stated motivation for this project was that of producing computationally efficient agents whose reasoning is grounded in a suitably relaxed—though still formally reasoned—form of iterated BR. At the same time, the model spaces encoded by UMAs are uniquely suited for *reactive control*: the selection of a control instruction in direct response to a localized (in time, as well as in space) perception of the task. At all times that the goal set is represented by a coherent selection on Σ (that is, the goal set is non-empty and convex in the relevant model space), propagation may be used to produce the nearest point projection paths from the current state to the goal set, within the model space, helping determine the appropriate actions as those provably propelling the agent roughly along one of these paths, using the mechanism described in Section 5.3.

A pertinent question for our current research is whether or not it is possible to employ self-enrichment procedures (see preceding paragraph) to guarantee—at least for some classes of problems—the emergence of a representation with the property that an agent's predictions from time t never fall outside the perceived current state at time (t + 1), for all t large enough. If, and when, that becomes possible, one will have to conclude that any planning failure is due to an obstacle in the relevant UMA model space(s) originating from an attempt to navigate into an impossible perceptual class. This would open the door to methods for efficient representation of such classes, as well as the leveraging of such representations for correcting the simplistic control scheme of navigation along geodesics.

Improving representation using multiple agents. The possible presence of obstacles focuses our attention on another important deficiency of UMA representations. While the concept representations they encode are *always* contractible when regarded as cubical complexes (see Appendix A.3 for more details), the concept representations corresponding to the ground truth will, more often than not, possess cavities/holes, serving the role of obstacles to navigation along geodesics in the UMA model space, and driving up the complexity of continuous planning [14]. An example of this phenomenon is already encountered in our simulations of target localization on the circle, in Section 5.4.2, and investigated in detail in Appendix F.

A possible solution to this problem might lie with the accumulation of a flexible collection of specialized agents, each with its own sensors and its own value signal; each correctly representing some aspects of the 'physical' agent's tasks, while having to rely on others in regions of its model space where its predictions fails. Fairly detailed descriptions of communities of this form have been proposed as possible models of human cognition by Minsky [35, 36], and studying the dynamics of such communities, charged with governing a situated agent, poses many interesting challenges.

In this context it is important to note that very recent results [41], demonstrating smooth(!) reactive switching between different control alternatives (behaviors/actions) using value-based motivational dynamics, provide a basis for speculation that (1) such methods may be applicable to our setting, too; and (2) formal understanding of the dynamics of the putative Minskian "societies" of UMAs just mentioned may be well within our reach.

Developing this approach will require the study of multi-agent systems incorporating means for the formation of "BUA coalitions", for lack of a better term, to be recruited for action under appropriate circumstances. This is also where we expect the mathematical theory behind UMAs to prove most useful. Its categorical underpinnings (the fact that model spaces arise as dual spaces; see Section 2.2) provide a rigorous framework for comparing different models of the same system, and for studying the interaction between different perceptual components of a single model (see Appendices A.2.4 and F where we carry out such detailed analysis).

ACKNOWLEDGEMENTS

This research was developed in part with funding from Air Force Research Lab (AFRL) grant FA865015D1845 (subcontract 669737-1), and in part with funding from the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Lab (AFRL) under agreement number FA8650-18-2-7840. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. The authors are grateful to Siqi Huang, a Penn CGGT Master's graduate, for his relentless work developing a hardware-accelerated implementation of the UMA architecture, making the simulations in this study possible. We also thank Kostas Karydis for helping proof-read some of the initial material on ranking-based UMAs during his last months as a post-doctoral fellow at Penn's GRASP lab.

References

- Ian Agol, Daniel Groves, and Jason Manning. The virtual Haken conjecture. Doc. Math, 18:1045–1087, 2013.
- [2] Carlos E Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic*, 50(2):510–530, 1985.
- [3] Carlos E Alchourrón and David Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48(1):14–37, 1982.
- [4] Carlos E Alchourrón and David Makinson. On the logic of theory change: Safe contraction. Studia logica, 44(4):405-422, 1985.
- [5] Hans-J. Bandelt and Jarmila Hedlíková. Median algebras. Discrete Math., 45(1):1–30, 1983.
- [6] Craig Boutilier. A unified model of qualitative belief change: A dynamical systems perspective. Artificial Intelligence, 98(1):281–316, 1998.
- [7] Indira Chatterji, Cornelia Druţu, and Frédéric Haglund. Kazhdan and Haagerup properties from the median viewpoint. Adv. Math., 225(2):882–921, 2010.
- [8] Victor Chepoi. Graphs of some CAT(0) complexes. Adv. in Appl. Math., 24(2):125–179, 2000.
- [9] TH Cormen, CE Leiserson, RL Rivest, and C Stein. Introduction to Algorithms. Cambridge, MA,, 2003.
- [10] Mukesh Dalal. Investigations into a theory of knowledge base revision: preliminary report. In Proceedings of the Seventh National Conference on Artificial Intelligence, volume 2, pages 475–479. Citeseer, 1988.
- [11] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. Artificial intelligence, 89(1-2):1–29, 1997.
- [12] James Delgrande and Yi Jin. Parallel belief revision: Revising by sets of formulas. Artificial Intelligence, 176(1):2223-2245, 2012.
- [13] Bruce Donald and James Jennings. Sensor interpretation and task-directed planning using perceptual equivalence classes. In *Robotics and Automation*, 1991. Proceedings., 1991 IEEE International Conference on, pages 190–197. IEEE, 1991.
- [14] Michael Farber. Topological complexity of motion planning. *Discrete Comput. Geom.*, 29(2):211–221, 2003.
- [15] William Feller. An introduction to the theory of probability and its applications, vol. 1. Wiley, 1967.
- [16] Robert Ghrist and Valerie Peterson. The geometry and topology of reconfiguration. Advances in applied mathematics, 38(3):302–323, 2007.
- [17] Moisés Goldszmidt and Judea Pearl. System-Z+: a formalism for reasoning with variablestrength defaults. In Proceedings of the ninth National conference on Artificial intelligence-Volume 1, pages 399–404. AAAI Press, 1991.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETR37

- [18] Moisés Goldszmidt and Judea Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. Artificial Intelligence, 84(1):57–112, 1996.
- [19] D.P. Guralnik and D.E. 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*, pages 936– 945, 2012.
- [20] Allen Hatcher. Algebraic topology. Cambridge University Press, Cambridge, 2002.
- [21] D.O. Hebb. The Organization of Behavior: A Neuropsychological Theory. Taylor & Francis, 2002.
- [22] Kern-Isberner, Gabriele and Brewka, Gerhard. Strong syntax splitting for iterated belief revision. In Proceedings of the 26th International Joint Conference on Artificial Intelligence, pages 1131–1137. AAAI Press, 2017.
- [23] John R. Isbell. Median algebra. Trans. Amer. Math. Soc., 260(2):319-362, 1980.
- [24] Yi Jin and Michael Thielscher. Iterated belief revision, revised. Artificial Intelligence, 171(1):1–18, 2007.
- [25] Hirofumi Katsuno and Alberto O Mendelzon. Propositional knowledge base revision and minimal change. Artificial Intelligence, 52(3):263–294, 1991.
- [26] Alex Kean and George Tsiknis. An incremental method for generating prime implicants/implicates. Journal of Symbolic Computation, 9(2):185–206, 1990.
- [27] Gabriele Kern Isberner and Daniela Huvermann. Multiple Iterated Belief Revision Without Independence. In *FLAIRS Conference*, pages 570–575, 2015.
- [28] Sébastien Konieczny and Ramón Pino Pérez. Putting More Dynamics in Revision with Memory. In International Workshop on Logic, Rationality and Interaction, pages 614–627. Springer, 2017.
- [29] Dmitry Kozlov. Combinatorial algebraic topology, volume 21 of Algorithms and Computation in Mathematics. Springer, Berlin, 2008.
- [30] Paolo Liberatore. The complexity of iterated belief revision. In International Conference on Database Theory, pages 276–290. Springer, 1997.
- [31] Paolo Liberatore and Marco Schaerf. Belief revision and update: Complexity of model checking. Journal of Computer and System Sciences, 62(1):43-72, 2001.
- [32] László Lovász. Random walks on graphs: A survey. Combinatorics, Paul erdos is eighty, 2(1):1–46, 1993.
- [33] Jerusa Marchi, Guilherme Bittencourt, and Laurent Perrussel. Prime forms and minimal change in propositional belief bases. Annals of Mathematics and Artificial Intelligence, 59(1):1–45, 2010.
- [34] Hua Meng, Hui Kou, and Sanjiang Li. Belief revision with general epistemic states. In AAAI, pages 1553–1559, 2015.
- [35] Marvin Minsky. K-Lines: A theory of Memory. Cognitive science, 4(2):117–133, 1980.
- [36] Marvin Minsky. The Society of Mind. Simon & Schuster, 1988.
- [37] Marvin Minsky, Seymour A Papert, and Léon Bottou. Perceptrons: An introduction to computational geometry. MIT press, 2017.
- [38] Bogdan Nica. Cubulating spaces with walls. Algebr. Geom. Topol., 4:297–309 (electronic), 2004.
- [39] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, pages 121–135. Morgan Kaufmann Publishers Inc., 1990.
- [40] Vaughn Pratt. Modeling concurrency with geometry. In Proceedings of the 18th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, pages 311–322. ACM, 1991.
- [41] Paul Reverdy and Daniel E Koditschek. A dynamical system for prioritizing and coordinating motivations. SIAM Journal on Applied Dynamical Systems, 17(2):1683–1715, 2018.
- [42] Ronald L Rivest and Robert E Schapire. Diversity-based inference of finite automata. Journal of the ACM (JACM), 41(3):555–589, 1994.
- [43] M.A. Roller. Poc Sets, Median Algebras and Group Actions. University of Southampton, Faculty of Math. stud., preprint series, 1998.
- [44] Marlow Sholander. Medians and betweenness. Proc. Amer. Math. Soc., 5:801–807, 1954.
- [45] Marlow Sholander. Medians, lattices, and trees. Proc. Amer. Math. Soc., 5:808-812, 1954.

- [46] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In Causation in decision, belief change, and statistics, pages 105–134. Springer, 1988.
- [47] Son Thanh To, Tran Cao Son, and Enrico Pontelli. A generic approach to planning in the presence of incomplete information: Theory and implementation. Artificial Intelligence, 227:1–51, 2015.
- [48] Leslie G Valiant. A theory of the learnable. Communications of the ACM, 27(11):1134–1142, 1984.
- [49] M. L. J. van de Vel. Theory of convex structures, volume 50 of North-Holland Mathematical Library. North-Holland Publishing Co., Amsterdam, 1993.
- [50] Daniel T. Wise. From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry, volume 117 of CBMS Regional Conference Series in Mathematics. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012.

APPENDIX A. APPENDIX: THE DUALITY THEORY OF FINITE POC SETS.

The purpose of this appendix is to review known results about the geometry of duals of finite poc sets, while illustrating them with simple examples which emphasize our application. An additional goal is to provide a sufficient technical background for proofs of new results in the appendices that follow.

The concept presentation of the dual of a poc set leads to more intuitive understanding of the geometry of poc set duals. Recall from Section 2.1.2 that the concept representation of a subset $V \subset \mathbb{H}$ of vertices of the Hamming cube over a PCS Σ encodes the set of (cubical) faces of the Hamming cube obtained by deleting all faces containing at least one vertex of $\mathbb{H} \setminus V$. The resulting structure is a (rather special) *cubical complex*¹¹. One way in which such cubical complexes are special is that they are completely determined by their *1-dimensional skeleton*—their collections of vertices and edges.

The resulting freedom to consider a higher dimensional "enveloping structure" for $Dual(\mathbf{P})$ when \mathbf{P} is a poc set over Σ turns out to be useful in many ways, some of which we intend to explore in this section.

Definition A.1 (Dual Cubing). Let **P** be a poc set structure over a finite PCS Σ . The *dual cubing* Cube(**P**) is the cubical complex obtained as the concept representation of the subset $\mathbf{P}^{\circ} \subset \mathbb{H}(\Sigma)$.

In the very least, the ability to refer to Cube(P) will make it easier to visualize the graph Dual(P), exposing its higher dimensional structure and bringing order to what otherwise would have been a chaos of edges (e.g. Figure 11). The notion of a dual cubing also makes it easier to understand cartesian products of dual graphs (Appendix A.2.2 below). Finally, we will use the dual cubing to explain some fundamental properties and limitations of PCR presentations (Appendix A.2.3 below) relating to their universality (Proposition 2.22).

A.1. Nesting, Transversality and Cubes. Fix a poc set P over a finite PCS Σ . The purpose of this section is to present the known characterizations of the cubes arising in Cube(P). Some additional standard terminology will be needed. The following are from [43], Section 1.4:

Definition A.2 (proper elements, proper pairs). Let Σ be a PCS. A *proper element* of Σ is any element $a \in \Sigma$ such that $a \notin \{0, 0^*\}$. A pair $\{a, b\}$ of proper elements in Σ is said to be proper, if $b \notin \{a, a^*\}$.

 $^{^{11}\}mathrm{See}$ [29], Chapter 2, for a very brief introduction to polyhedral (in particular, cubical) complexes.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETR39



FIGURE 8. Half-spaces of \mathbf{P}° and the hyperplanes of $\text{Cube}(\mathbf{P})$ that generate them (see discussion following Lemma A.4). Introducing the nesting relation $a < c^*$ (top,a) produces the dual cubing (top,b), whose hyperplanes arise as the restrictions of the canonical hyperplanes of the Hamming cube (top,c). Additional relations (bottom,a) further reduce the dual and its hyperplanes (bottom,b), as fewer faces of the dual each hyperplane of the Hamming cube (bottom,c).

Definition A.3 (nesting, transversality). Let **P** be a poc set. For any proper $a, b \in \mathbf{P}$ at most one of the following holds:

(32)
$$a \ge b, \quad a \ge b^*, \quad a < b, \quad a < b^*.$$

If any one of the above relations holds, we will say that a and b are nested. Otherwise, we say that a and b are transverse. Furthermore, for any $A \subset \mathbf{P}$, we say that A is nested (transverse), if every two elements of A are nested (resp. transverse). \Box

Recalling that a poc set is, first and foremost, a partially ordered set, for any subset $S \subset \mathbf{P}$ it makes sense to consider

(33)
$$\min(S) := \{a \in S \mid (\forall b \in S) (b \le a \to b = a)\}.$$

Since **P** is finite, $\min(S)$ is non-empty whenever S is. The following is Proposition 10.1 of [43], restricted to the finite case and parsed into more elementary language:

Lemma A.4 (when a vertex meets a cube). Let \mathbf{P} be a finite poc set and let $v \in \mathbf{P}^{\circ}$. Let Q be a d-dimensional cube of $\mathsf{Cube}(\mathbf{P})$. Then $v \in Q$ if and only if $\min(v)$ contains a transverse subset T of \mathbf{P} with the property that every vertex $u \in Q$ is of the form $u = [v]_s := (v \setminus S) \cup S^*$ for some $S \subseteq T$. \Box

In particular:

- every edge (1-cube) containing v is spanned by v and a vertex of the form $[v]_a := (v \setminus \{a\}) \cup \{a^*\}$ for some $a \in \min(v)$;
- every square (2-cube) containing v is spanned by v, $[v]_a$, $[v]_b$ and $[v]_{ab}$ for some transverse pair $\{a, b\} \subseteq \min(v)$.

These properties give rise to a new understanding of how the half-spaces $\{\mathfrak{h}(a; \mathbf{P})\}_{a \in \mathbf{P}}$ in Dual(**P**) interact with the geometry of Cube(**P**):one could think of the splitting of \mathbf{P}° in the form $\mathfrak{h}(a; \mathbf{P}) \cup \mathfrak{h}(a^*; \mathbf{P})$ as the result of cutting the cubing Cube(**P**) along the *hyperplane* arising as the union of perpendicular bisectors of edges of the form $\{v, [v]\}$ —see Figure 8.

In addition, the last lemma plays a crucial role in deducing some fundamental properties of $Cube(\mathbf{P})$ (Proposition 10.2 of [43]):

Theorem A.5. Let **P** be a finite poc set. Then $Cube(\mathbf{P})$ is contractible.¹²

Moreover, the lemma implies that $Cube(\mathbf{P})$ is non-positively curved (see [50], Section 2.1). This produces a characterization of complexes of the form $Cube(\mathbf{P})$ (Theorem 10.3 of [43]):

Theorem A.6 (characterization of cubings). A cubical complex arises as the dual of a finite poc set if and only if it is contractible and non-positively curved. \Box

All the above apply in far more general settings than the finite one: the interested reader should consult [43].

A.2. Examples of Duals. To improve the reader's intuition regarding dual graphs of poc sets, as well as to illustrate one of the example simulations (Section 5.4.1), we consider a sequence of examples in light of the results of Section 2.4.

A.2.1. Example: a bead on a string. Suppose the system being observed consists of a bead strung on a tight piece of string. The observed state of the system is modeled by the interval [0, 1] in the obvious way, so the space of histories **X** is the set of sequences $x = (x_n)_{n=-\infty}^0$, where $pos(x) := x_0$ corresponds to the current position of the bead given x, x_{-1} is the previous position of the bead, and so on. Let us set $\mathbf{\Sigma} = \{\mathbf{0}, \mathbf{0}^*, a_1, a_1^*, \dots, a_L, a_L^*\}$ with two different poc set structures, \mathbf{P} and \mathbf{Q} , defined by the relations $a_k < a_{k+1}, 1 \le k < L$ in \mathbf{P} and $a_i < a_j^*$, $1 \le i < j \le L$ in \mathbf{Q} . These may be regarded as PCR representations of two different sensoria constructed as follows. Let $p_1 < \dots < p_L$ in (0,1) be points that are pairwise at least ϵ apart, $0 < \epsilon < \frac{1}{2(L+1)}$. Then \mathbf{P} may be realized by setting $x \in \rho(a_k) \Leftrightarrow pos(x) < p_k$ ("threshold sensors"), while \mathbf{Q} may be realized, for example, by $x \in \rho(a_k) \Leftrightarrow dist(pos(x), p_k) < \epsilon$ ("beacon sensors").

The vertices of Dual(**P**) have the form $V_k = \{\mathbf{0}^*\} \cup \{a_j^*\}_{j>k} \cup \{a_i\}_{i\geq k}, 0 \leq k \leq L$, with an edge joining V_k to V_{k+1} for all k < L (recall that edges in Dual(**P**) are edges of the Hamming cube $\mathbb{H} = \mathbb{H}(\Sigma)$). The graph Dual(**Q**) has a different collection of vertices, dictated by the fact that all pairs $\{a_i, a_j\}$ with $i \neq j$ are incoherent: there is a 'special' vertex $V'_0 = \{\mathbf{0}^*, a_1^*, \dots, a_L^*\}$ and a collection of 'generic' ones, $V'_k = \{\mathbf{0}^*, a_k\} \cup \{a_j^*\}_{j\neq k}$; all the $V'_k, k > 0$, are adjacent to V'_0 , and no other pair of vertices are adjacent. Figure 9 shows Dual(**P**) (left), which is an *L*-path, and Dual(**Q**) (right), which we will refer to in the future as a *starfish*. Note how, of the two model spaces, Dual(**P**) seems to provide the better discretization of [0, 1]. Note that both duals are trees. This is a manifestation of the well-known fact that Dual(**P**) is a tree if and only if **P** is nested (that is, any two elements of **P** are nested).

 $^{^{12}}$ Contractibility of a topological space is a fundamental notion in Topology, formalizing the idea of a "space with no holes". See [20], Chapter 0 for a quick and very intuitive introduction.



FIGURE 9. The dual graphs Dual(P) and Dual(Q) for the two arrangements of sensors along the real line described in Appendix A.2.1: 'threshold' sensors encoding a path (left, satisfying **P**), and 'beacon' sensors encoding a starfish (right, satisfying **Q**), respectively.

A.2.2. *Example: Cartesian products of duals.* The easiest way to join two poc sets together is to form their direct sum:

Definition A.7. Let **P** and **Q** be discrete poc sets. Their *direct sum* $\mathbf{P} \vee \mathbf{Q}$ is defined to be the quotient of their external disjoint union $P \sqcup Q$ by the identification $\mathbf{0}_{\mathbf{P}} = \mathbf{0}_{\mathbf{Q}}$ and $\mathbf{0}^*_{\mathbf{P}} = \mathbf{0}^*_{\mathbf{Q}}$, endowed with the following:

- $a \leq b$ in $\mathbf{P} \vee \mathbf{Q} \Leftrightarrow (\{a, b\} \subseteq \mathbf{P} \text{ and } a \leq b$ in $\mathbf{P})$ or $(\{a, b\} \subseteq \mathbf{Q} \text{ and } a \leq b$ in $\mathbf{Q})$;
- $a = b^*$ in $\mathbf{P} \vee \mathbf{Q} \Leftrightarrow (\{a, b\} \subseteq \mathbf{P} \text{ and } a = b^*$ in $\mathbf{P})$ or $(\{a, b\} \subseteq \mathbf{Q} \text{ and } a = b^*$ in $\mathbf{Q})$.

We abuse notation by identifying each element of $\mathbf{P} \cup \mathbf{Q}$ with the equivalence class in $\mathbf{P} \vee \mathbf{Q}$ of its natural representative in $\mathbf{P} \sqcup \mathbf{Q}$.

Consider the two inclusion maps, $p: \mathbf{P} \hookrightarrow \mathbf{P} \lor \mathbf{Q}$ and $q: \mathbf{Q} \hookrightarrow \mathbf{P} \lor \mathbf{Q}$, each of which is an injective poc morphism. The dual maps p° and q° give rise to the median morphism $\mu : (\mathbf{P} \lor \mathbf{Q})^{\circ} \to \mathbf{P}^{\circ} \times \mathbf{Q}^{\circ}$ defined by $\mu(w) = (p^{\circ}(w), q^{\circ}(w))$, where $p^{\circ}(w) = w \cap \mathbf{P}$ and $q^{\circ}(w) = w \cap \mathbf{Q}$, by definition. Since every proper pair $a, b \in \mathbf{P} \lor \mathbf{Q}$ with $a \in \mathbf{P}$ and $b \in \mathbf{Q}$ satisfies $a \pitchfork b$, it follows that $u \cup v$ is coherent for any $u \in \mathbf{P}^{\circ}$ and $v \in \mathbf{Q}^{\circ}$, and we conclude that μ is bijective.

Finally, recall that an edge in $\text{Dual}(\mathbf{P} \vee \mathbf{Q})$ joining $w = \mu(u, v)$ with $w' = \mu(u', v')$ occurs iff $|w \setminus w'| = 1$. Since the intersection of \mathbf{P} with \mathbf{Q} in $\mathbf{P} \vee \mathbf{Q}$ is trivial, in terms of u, u', w, w' we obtain:

$$w \smallsetminus w'| = |u \smallsetminus u'| + |v \smallsetminus v'| ,$$

so that w, w' span an edge if and only if exactly one of the pairs $\{u, u'\}$ or $\{v, v'\}$ spans an edge. Thus, μ is a median isomorphism of the dual graphs and we have:

Corollary A.8. Let **P**, **Q** be discrete poc sets. Then the mapping

(34)
$$\mu: \begin{cases} \operatorname{Dual}(\mathbf{P} \lor \mathbf{Q}) \to \operatorname{Dual}(\mathbf{P}) \times \operatorname{Dual}(\mathbf{Q}) \\ w \mapsto (w \cap \mathbf{P}, w \cap \mathbf{Q}) \end{cases}$$

is a median-preserving graph isomorphism.

For an alternative argument, note that for any $u \in \mathbf{P}^{\circ}$ and $v \in \mathbf{Q}^{\circ}$, if $S \subset \min(u)$ and $T \subset \min(v)$ are transverse sets, then $S \cup T \subset \min(u \cup v)$ and is a transverse set



FIGURE 10. Illustrating the case L = 4 in the example of Appendix A.2.3: images of the realization maps (left), and the corresponding perceptual classes in the model space Cube(P) dual to a common PCR representation (right).

in $\mathbf{P} \lor \mathbf{Q}$. Therefore, by Lemma A.4, every cube in $\mathsf{Cube}(\mathbf{P}) \times \mathsf{Cube}(\mathbf{Q})$ corresponds to a unique cube in $\mathsf{Cube}(\mathbf{P} \lor \mathbf{Q})$. Thus μ from the corollary is much more than an isomorphism of graphs: it extends to an isomorphism of cubical complexes from $\mathsf{Cube}(\mathbf{P} \lor \mathbf{Q})$ onto $\mathsf{Cube}(\mathbf{P}) \times \mathsf{Cube}(\mathbf{Q})$.

A.2.3. Example: representing a circle. Similarly to the example of a bead on a straigh piece of string (Appendix A.2.1), one could consider a bead on a circular bracelet, replacing the interval [0,1] with the unit circle $\mathbb{S}^1 \subset \mathbb{C}$ in the complex plane. This time, let $p_0, \ldots, p_{L-1} \in \mathbb{S}^1$ be a cyclically ordered collection of marker points, say, $p_k := \exp(\frac{2\pi i \mathbf{i}}{L})$.

We will compare several different representations over the PCSs:

(35)
$$\Sigma = \Sigma(L) = \{\mathbf{0}, \mathbf{0}^*\} \cup \{a_0, a_0^* \dots, a_{L-1}, a_{L-1}^*\}, \quad L \in \mathbb{N}.$$

We regard Σ as a sensorium whose realization $\rho = \rho(L, \epsilon)$ is defined by setting $x \in \rho(a_k)$ for a history x if and only if the currect state x_0 lies in the open circular arc segment of \mathbb{S}^1 centered at p_k and having radius ϵ . Depending on the choice of L and ϵ , different PCRs (and duals) may arise. Specifically, We consider the examples with L = 4 and $\epsilon = \frac{\pi}{4}, \frac{\pi}{3}$; with L = 6 and $\epsilon = \frac{\pi}{3}, \frac{\pi}{2}$, to illustrate possible differences and shared qualities.

Jack Sparrow's compass, L = 4. Rather than keep track of the indices modulo 4 in this example, let us identify it with a day-to-day object: a compass. We denote

(36)
$$\mathbf{n} := a_0, \quad \mathbf{s} := a_2, \quad \mathbf{w} := a_1, \quad \mathbf{e} := a_3.$$

Figure 10(left) depicts the subsets of \mathbb{S}^1 which determine $\rho(a_i)$, i = 0, 1, 2, 3, for the realizations in the cases $\epsilon = \frac{\pi}{3}$ (A) and $\epsilon = \frac{\pi}{4}$ (B). Thinking of \mathbb{S}^1 as the space of all possible positions of a compass needle—the needle of *this* compass points in the direction of your heart's greatest desire and that may not be a visit to the magnetic north pole—one should think of, e.g., $N_\alpha := \rho_\alpha(\mathbf{n})$, $\alpha \in \{A, B\}$, as the set of positions of the needle with which observer α associates an affirmative answer to the question "Is the needle pointing North?". The difference between the two examples is that $\rho_B(\mathbf{n}), \rho_B(\mathbf{s}), \rho_B(\mathbf{w}), \rho_B(\mathbf{e})$ are pairwise disjoint, while ρ_A realizes the major directions so that only opposites are disjoint. Let G denote the poc set structure on Σ with relations of the form¹³

$$(37) n < s^*, w < e^*.$$

Then ρ is a poc morphism for either choice of ϵ , and the right hand side of Figure 10 illustrates the perceptual classes of ρ (yellow highlighting) together with the edges they induce in the ambient structure, Dual(G). Note how case (A) produces an embedded cycle sub-graph in Dual(G)—a coarse but topologically faithful reconstruction of S^1 , which is homotopically non-trivial—while case (B) produces a tree, a space homotopically equivalent to a point.

While providing an illustration for Proposition Proposition 2.22, this example also highlights the necessity in discussing what properties of the realization map ρ could guarantee a degree of fidelity of the observer's reconstruction of the observed space (the space of histories \mathbf{X} ? the 'environment' \mathbb{S}^1 ?) as, say, the sub-graph of Dual(G) induced by the perceptual classes.

Higher dimensions, L = 6. Figure 11 compares the dual graphs/cubings of two poc set representations, each optimal for its corresponding choice of the value of ϵ . The case $\epsilon = \frac{\pi}{3}$ again has the property that non-consecutive $\rho(a_i)$ are disjoint, implying that ρ is a poc isomorphism of the poc set structure

(38)
$$a_i \pitchfork a_j \Leftrightarrow |i-j| \le 1, \quad a_i < a_j^* \Leftrightarrow |i-j| > 1,$$

onto its image in $2^{\mathbb{S}^1}$. Denote this poc set structure on Σ by \mathbf{P}_1 . The case $\epsilon = \frac{\pi}{2}$ has fewer nesting relations among the $\rho(a_j)$, because every three consecutive sets of this form have a point in common. Formally, ρ is a poc isomorphism of the poc set structure described by:

(39)
$$a_i \pitchfork a_j \Leftrightarrow |i-j| \le 2, \quad a_i < a_j^* \Leftrightarrow |i-j| > 2$$

onto its image. Denote this poc set structure on Σ by \mathbf{P}_2 . Note that the identity map id: $\mathbf{P}_2 \to \mathbf{P}_1$ is a poc morphism, while its inverse is not: a poc morphism f is allowed to map a transverse pair to a nested one, but not the other way around. The dual of this map embeds $Dual(P_1)$ in $Dual(P_2)$. This embedding can be seen clearly in Figure 11(right).

A.2.4. Example: moving bead on an interval. Returning to the 'thresholds' example of Appendix A.2.1, we would like to consider it from the point of view of the agents described in Section 5.4.1.

The interval [0,1] from Appendix A.2.1 will now be replaced with [0,L], L a positive integer, for convenience. Once again we are given position sensors $a_1, \ldots, a_L \in \Sigma$ —more sensors will be added to Σ in a moment—with realizations $x \in \rho(a_j) \Leftrightarrow pos(x) < j$, where we recall that $x = (x_t)_{t=-\infty}^0, x_t \in [0, L]$ is our current notion of a history, and $pos(x) := x_0$ is the current position of the bead on [0, L].

This time we are interested in reasoning about the possible motion of the bead along the interval, so we introduce delayed sensors $\sharp a_i$ into Σ alongside the original position sensors. Formally, the delay operator acts on histories via $(\sharp x)_t := x_{t-1}$, and acts on sensors via $x \in \rho(\sharp a) \Leftrightarrow \sharp x \in \rho(a)$, that is: at any time, the current value of $\sharp a$ coincides with value of a in the previous cycle.

 $^{^{13}}$ We regard the indices in this example and any arithmetic operations on them as being defined modulo L.



FIGURE 11. Duals of 'optimal' poc set representations for two distinct realizations of the example of Appendix A.2.3, in the case L = 6: $Cube(\mathbf{P}_1)$ is 2-dimensional (left), while $Cube(\mathbf{P}_2)$ is 3-dimensional (center). Vertices painted black represent consistent models, while vertices painted white represent coherent, but inconsistent models. The diagram to the right shows what is left of $Cube(\mathbf{P}_2)$ when the inconsistent models are deleted from the complex, leaving a sub-complex homotopyequivalent to the circle \mathbb{S}^1 .

The bead is endowed with two actuators. One, named rt, whose action at time t pushes the bead one unit to the right along the interval. The only exception is the position L: if rt is applied there, its contribution to the motion of the bead will be nil. Similarly, an actuator named lt pushes the bead one unit toward the endpoint 0 of [0, L], with no effect when the bead is already there. Finally, turning on both actuators at the same time results in no motion of the bead in either direction.

Two agents, also named **rt** and **lt**, are each in charge of deciding, respectively, whether to act (turn on their assigned actuator for the duration of one time interval), or not. Each agent $\alpha \in \{\texttt{rt}, \texttt{lt}\}$ maintains two PCR representations: G^{α} is updated conditioned on the agent having acted, and G^{α^*} is updated conditioned on α resting.

Here we will consider the poc set representations we would like each agent to learn, as we attempt to draw their dual cubings. For this purpose, we analyze nesting relations in the sensorium

(40)
$$\Sigma_0 := \{\mathbf{0}, \mathbf{0}^*\} \cup \{a_j, a_j^*\}_{j=1}^L \cup \{\sharp a_j, \sharp a_j^*\}_{j=1}^L.$$

In the absence of any additional assumptions, the following relations are consistent with the selected realization (and are, therefore, desirable as part of any learned poc set structure on Σ_0): $a_1 < a_2 < \cdots < a_L$ encodes the geometry of the interval, and further implies also $\sharp a_1 < \sharp a_2 < \cdots < \sharp a_L$; not knowing anything about the actions taken by the actuators one may only be certain of the relations (\ddagger) $\sharp a_j < a_{j+1}$, $a_j < \sharp a_{j+1}$, for j < L. Denote the resulting poc set structure on Σ_0 by \mathbf{P}_0 .

Leveraging our understanding of cartesian products (Appendix A.2.2), we set \mathbf{Q} to be the sub-poc set of \mathbf{P}_0 restricted to just the position sensors a_j , while $\sharp \mathbf{Q}$ will be the sub-poc set of \mathbf{P}_0 over the delayed position sensors $\sharp a_j$. Then the identity mapping $\mathbf{id} : \mathbf{Q} \lor \sharp \mathbf{Q} \to \mathbf{P}_0$ is a poc morphism, whose dual map is a median-preserving embedding of cubical complexes, of $\operatorname{Cube}(\mathbf{P}_0)$ in the square $(L+1) \times (L+1)$ grid arising as $\operatorname{Cube}(\mathbf{Q}) \times \operatorname{Cube}(\sharp \mathbf{Q})$. We conclude that $\operatorname{Cube}(\mathbf{P}_0)$



FIGURE 12. Discrete motion of a bead along the interval [0, L], illustrating Appendix A.2.4: $\text{Cube}(\mathbf{P}_0)$ (red highlighting) shown as a subcubing of the 2-dimensional grid $\text{Cube}(\mathbf{Q}) \times \text{Cube}(\sharp \mathbf{Q})$ (left); the cubing representing the correct implications for rt and for \mathtt{lt}^* shown as a subcomplex of $\text{Cube}(\mathbf{P}_0)$ (yellow highlighting, right), while the cubing for \mathtt{lt} and for \mathtt{rt}^* is its diagonal mirror image.

is the cubical complex shown in Figure 12(left), by applying the relations (\ddagger) to erase redundant squares from the grid.

Now suppose that all the correct relations have been (somehow) learned and represented in the collection of PCRs G^{rt} , $G^{\text{rt}*}$, G^{lt} and $G^{\text{lt}*}$. Because the synchronous application of **rt** and **lt** yields no motion, there is no way to discriminate between G^{rt} and $G^{\text{lt}*}$, as well as between G^{lt} and $G^{\text{rt}*}$. However, the former two poc set presentations will have obtained the relations $a_j < \sharp a_j$ in addition to those of \mathbf{P}_0 , causing the dual cubing to grow even smaller, as highlighted on the right-hand side of Figure 12 (of course, a symmetric situation arises for the other pair of indistinguishable representations).

To end this section we note that, by hard-wiring the actuators to never execute both rt and lt at the same time, it is possible to disambiguate the representations. In this regime, the (optimally learned) representations G^{rt^*} and G^{lt^*} will remain the same, while the PCRs G^{rt} and G^{lt} will each experience a collapse to a nontrivial canonical quotient: the PCR G^{rt} witnesses $\sharp a_j$ if and only if it witnesses a_{j+1} (the diagonal vertices in Figure 12 are inconsistent given rt is active). The situation is symmetric (but not identical) for G^{lt} , and we obtain four distinct "world views" for each of the observers.

A.3. Homotopy type of the observed space. The phenomenon witnessed by the examples of Appendix A.2.3 is very general, and brings to bear on the capabilities and limitations of knowledge representation using PCRs.

For a fixed PCS Σ , a fixed space **X** and PCS morphism $\rho : \Sigma \to 2^{\mathbf{X}}$, recall (Section 2.1.1) the subset $\mathbf{M}(\rho)$ of the Hamming cube $\mathbb{H} = \mathbb{H}(\Sigma)$ consisting of those models $u \in \mathbb{H}$ for which $\bigcap_{a \in u} \rho(a)$ is non-empty—the set of possible worlds with respect to ρ . Let $\mathsf{Cube}(\rho)$ denote the cubical complex corresponding to the

concept presentation of $\mathbf{M}(\rho)$ —the set of cubical faces of \mathbb{H} all of whose vertices lie in $\mathbf{M}(\rho)$.

The authors proved in [19] that, for sufficiently tame topological spaces **X** and PCS morphisms $\rho : \Sigma \to 2^{\mathbf{X}}$, the following holds:

Theorem A.9 (Recovery of Homotopy Type). Suppose that, for every cube $C \in Cube(\rho)$, the set

(41)
$$Z_C := \bigcup_{u \in C} \bigcap_{a \in u} \rho(a)$$

is contractible. Then $Cube(\rho)$ is homotopy equivalent to **X**.

In other words, if the collection of queries available to the observer is sufficiently rich that obviously contractible subspaces of $Cube(\rho)$ (cubes) are witnessed by contractible subspaces of **X**, then $Cube(\rho)$ has, in the formal sense provided by algebraic topology, the same shape as the observed space **X**.

In particular, under the condition of the theorem, if \mathbf{P} is a poc set structure on Σ and ρ is a poc morphism, then the universality of representation by PCRs (Proposition 2.22) implies that $\text{Cube}(\rho) \subseteq \text{Cube}(\mathbf{P})$, making $\text{Cube}(\mathbf{P})$ into a minimal contractible model space for \mathbf{P} housing a homotopy model of the observed space, and the discrepancy between the two is precisely the set of unobservable perceptual classes.

To illustrate the theorem, let us return to the examples of the preceding paragraph to observe that none of the phenomena we have encountered there had happened by accident. For any $L \ge 2$, a choice of $0 < \epsilon \le \frac{\pi}{L}$ leads to $\text{Cube}(\mathbf{P})$ being a tree (a 'starfish') containing the vertex $v = \{a_j^*\}_{j=0}^{L-1}$. Since the set of points in S^1 witnessing this vertex is disconnected (see Figure 10), the hypothesis of the last theorem fails, making it possible for $\text{Cube}(\rho)$ to be contractible, which is exactly what happened for L = 4, $\epsilon = \frac{\pi}{4}$. At the same time, any choice of $\frac{\pi}{L} < \epsilon \le \frac{\pi}{2}$ results in the hypothesis of the theorem being fulfilled, which is why, in the three other cases considered here, $\text{Cube}(\rho)$ is homotopy-equivalent to the circle.

Finally, we would like to emphasize that—similarly to the examples considered above—neither the tameness assumptions on **X** and ρ nor the hypothesis of the last theorem are excessive in standard Robotics settings. First, since the sensor values are often functions of merely the last few visited states, the realization map $\rho: \Sigma \to 2^{\mathbf{X}}$ will often factor, up to sufficient approximation, through $2^{E \times \cdots \times E}$ where E is the configuration space of the robotic system (similarly to the role played by the circle and the interval in all the preceding examples). Second, E is often a manifold, possibly with corners, or a cellular complex; in the absence of chaotic behavior, and provided sufficient sensing, it becomes possible to construct a sufficiently fine mesh of sensor values for "chopping up" the reduced history space $E \times \cdots \times E$ into small contractible regions as required by our theorem.

APPENDIX B. APPENDIX: BASIC RESULTS ABOUT PCRS.

B.1. **Proof of Proposition 2.17.** Suppose G is non-degenerate. Take any $S \in G^{\circ}$ and any $a \in \Sigma$. One of the following holds:

- $S \cup \{a\}$ is coherent, and hence $a \in S$ (by the maximality property of S) and $a \leq a^*$;
- $S \cup \{a^*\}$ is coherent, in particular $a^* \in S$ and $a^* \not\leq a$;

46

• Neither of the above.

In the third case there are two possibilities. Either S is empty, in which case the statement is that neither $\{a\}$ nor $\{a^*\}$ are coherent, which means that both $a \leq a^*$ and $a^* \leq a$ hold, and putting a inside $N(G) \cap N(G)^*$ — a contradiction; or there exist $b, c \in S$ such that $a \leq b^*$ and $a^* \leq c^*$. But then $c \leq b^*$ — a contradiction to S being coherent. Thus we are left with $a \in S$ or $a^* \in S$ for each $a \in \Sigma$, as desired.

The second assertion trivially implies the third, and the third implying the first follows from the remark preceding Definition 2.16. $\hfill \Box$

B.2. **Proof of Proposition 2.18.** It is clear that χ is injective. Any $f \in \text{Hom}_{gc}(G, \mathbf{2})$ is a function of Σ to $\mathbf{2}$, a two-point set, and is therefore characterized by the (possibly empty) set of points on which it obtains the value $\mathbf{1}$.

Now let us verify that χ is well-defined, that is: that the set $S = f^{-1}(\mathbf{1})$ is a maximal coherent subset of Σ with respect to G. Indeed, were $a, b \in S$ such that $a \leq b^*$, this would force $\mathbf{1} = f(a) \leq f(b)^* = \mathbf{0}$ — a contradiction.

Finally, we prove the surjectivity of χ . Given a maximal coherent set S, Proposition 2.17 implies S is a selection on Σ . This means that the function $f: \Sigma \to 2$ defined by $f(a) = \mathbf{1} \Leftrightarrow a \in S$ satisfies the identity $f(a^*) = f(a)^*$. We claim that f is a PCR morphism. Since $\chi(f) = S$, proving this claim will finish the proof of the current proposition.

Suppose f is not a morphism. Then there is $ab \in G$ satisfying $f(a) \not\leq_G f(b)$. In the current setting this is tantamount to $f(b) = \mathbf{0}$ and $f(a) = \mathbf{1}$, or, equivalently, $f(b^*) = f(a) = \mathbf{1}$. In turn, this means $a, b^* \in S$. However, S is forward-closed (as is any maximal coherent set), so $a \in S$ and $ab \in G$ imply $b \in S$. With $b^* \in S$ we obtain a contradiction.

B.3. **Proof of Proposition 2.22.** The proof extends a standard argument from Sageev-Roller duality theory. Given **X** and ρ , pick any point $x \in \mathbf{X}$. By definition, $\xi = \rho^*(x)$ belongs in G° if and only if no $a, b \in \xi$ satisfy $a \leq b^*$ in G. Since ρ is order-preserving, having $a \leq b^*$ for $a, b \in \xi$ would imply $\rho(a) \cap \rho(b) = \emptyset$ while $x \in \tilde{\rho}(a) \cap \tilde{\rho}(b)$ at the same time—contradiction. Thus, $\xi \in G^\circ$ for all choices of $x \in \mathbf{X}$, proving the first assertion of the proposition. To verify the second one, consider the choice of $\mathbf{X} = G^\circ$ with $\rho : \mathbf{\Sigma} \to \mathbf{2}^{\mathbf{X}}$ given by $\rho(a) = \{U \in G^\circ | a \in U\}$. It is easily verified that ρ is a morphism and that $\rho^* : \mathbf{X} \to G^\circ$ is the identity map (and hence surjective), finishing the proof.

B.4. Proof of Proposition 2.24. Let G be a fixed non-degenerate PCR over Σ . For every $a \in \Sigma$, recall $[a]_G = a \uparrow \cap a \downarrow$, and recall the definition of $\pi = \pi_G$:

(42)
$$\pi_G(a) := \begin{cases} [a]_G & \text{if } a \notin N(G) \cup N(G) \\ N(G) & \text{if } a \in N(G) \\ N(G)^* & \text{if } a \in N(G)^* \end{cases}$$

Here are a few natural observations:

- For any $a \in \Sigma$, $[a]_{G^*} = [a^*]_G$, where for any set $S \subseteq \Sigma$ we remember that $S^* := \{a^* | a \in S\}.$
- Since N(G) is backwards-closed, N(G) is a union of strong components of G: indeed, if $a \in N(G)$ then every $b \in [a]_G$ satisfies $b \leq a$, which implies $b \in N(G)$; hence $[a]_G \subseteq N(G)$.
- Analogously for $N(G)^*$, since it is forward-closed.

This allows for the construction of a new PCR \widehat{G} over the PCS $\widehat{\Sigma} := \{\pi(a) | a \in \Sigma\}$, by setting $\widehat{G} = \{\pi_G(a)\pi_G(b) | a \leq_G b\}$. We claim that \widehat{G} induces on $\widehat{\Sigma}$ the structure of a poc set. For this it will suffice to show that \widehat{G} is a non-degenerate PCR and a partial order.

First we show that $\widehat{\Sigma}$ is a PCS. The identity $\pi(a^*) = \pi(a)^*$ yields $\pi(a)^{**} = \pi(a^{**}) = \pi(a)$ for all $a \in \Sigma$. Suppose some $a \in \Sigma$ satisfied $\pi(a)^* = \pi(a)$. Since G is non-degenerate, this means $a \notin N(G)$ and $a \notin N(G)^*$. But then $\pi(a) = [a]_G$, at the same time, and the equality $[a^*]_G = [a]_G$ implies both $a \leq a^*$ and $a^* \leq a$ —contradicting non-degeneracy. We conclude that $\pi(a)^* \neq \pi(a)$ for all $a \in \Sigma$, and, since π is surjective, $\widehat{\Sigma}$ is a PCS.

It is clear now that \widehat{G} is a PCR, by construction. Suppose now that $A \in \widehat{\Sigma}$ lay in $N(\widehat{G})$. Then $A \leq A^*$, and writing $A = [a]_G$, $a \in \Sigma$ we obtain $a \in N(G)$, showing that A = N(G). Thus, $N(\widehat{G})$ is trivial, as desired.

 $\widehat{\mathbf{G}}$ is partially ordered by general considerations, so to conclude that $\widehat{\mathbf{G}}$ is a poc set, it remains to verify that N(G) is its minimum. Now, $\mathbf{0} \in N(G)$ and $\widehat{\mathbf{0}} = \pi(\mathbf{0}) = N(G) \in \widehat{\mathbf{\Sigma}}$ imply that the edge $\widehat{\mathbf{0}}\pi(a) \in \widehat{G}$ for all $a \in \mathbf{\Sigma}$. Since π is surjective, $\widehat{\mathbf{0}}$ is the minimum element of $\widehat{\mathbf{\Sigma}}$ with respect to the new partial order.

Finally, let **P** be any poc set, and let $f: G \to \mathbf{P}$ be any PCR morphism. Then f is constant on $\pi(a)$ for all $a \in \Sigma$, which defines the injective set map $\hat{f}: \Gamma \to \mathbf{P}$ via $\hat{f}([a]) = f(a)$. This map is a PCR morphism of complemented graphs by construction, and is, therefore, a morphism of \hat{G} into **P**. If $f': \hat{G} \to \mathbf{P}$ is any poc morphism satisfying $f = f' \circ \pi$, then for any $a \in \Sigma$ we have $f'(\pi(a)) = f(a) = \hat{f}(\pi(a))$. Since π is surjective, f' coincides with \hat{f} .

B.5. **Proof of Corollary 2.25.** Let G be a PCR over a PCS Σ , and let $\pi_G : G \to \widehat{G}$ be the canonical quotient map. We apply Proposition 2.24 with $\mathbf{P} = \mathbf{2}$.

For any morphism $\varphi : G \to \mathbf{2}$ there exists one and only one morphism $\widehat{\varphi} : \widehat{G} \to \mathbf{2}$ satisfying $\varphi = \widehat{\varphi} \circ \pi_G$. Now, thinking of $\widehat{\varphi}$ as an element of \widehat{G}° , we may write, by Proposition 2.18 and Definition 2.20, $\pi_G^\circ(\widehat{\varphi}) = \widehat{\varphi} \circ \pi_G = \varphi$.

On the other hand, for any $\psi: \widehat{G} \to \mathbf{2}$, we may write $\widehat{\pi}_{G}^{\circ}(\overline{\psi}) = \widehat{\psi} \circ \widehat{\pi}_{G} = \psi$, the last inequality following from the uniqueness assertion of Proposition 2.24, applied to the morphism $\psi \circ \pi_{G}$.

We conclude that the map $\operatorname{Hom}_{cg}(G, \mathbf{2}) \to \operatorname{Hom}_{cg}(\widehat{G}, \mathbf{2})$ defined by $\varphi \mapsto \widehat{\varphi}$ is an inverse of $\pi_G^{\circ} : \operatorname{Hom}_{cg}(\widehat{G}, \mathbf{2}) \to \operatorname{Hom}_{cg}(G, \mathbf{2})$, as desired. \Box

B.6. **Proof of Corollary 2.26.** We apply Proposition 2.24 again, to the PCR G, the poc set $\mathbf{P} = \widehat{H}$ and the morphism $g := \pi_H \circ f : G \to \widehat{H}$, to conclude there exists one and only one morphism $\widehat{g} : \widehat{G} \to \widehat{H}$ satisfying $g = \widehat{g} \circ \pi_G$. Substituting $g = \pi_H \circ f$ we see that \widehat{g} is the required morphism.

APPENDIX C. APPENDIX: CONVEXITY THEORY OF PCR DUALS.

The purpose of this section is to provide a self-contained account of our results regarding coherent projection and the use of propagation for the computation of nearest-point projections in poc set duals. Throughout this section, G will be a fixed non-degenerate PCR over a finite PCS Σ . Moreover, without loss of generality (through replacing G with its canonical poc quotient \hat{G}), we may assume G is a

poc set. Since G is fixed, we will simplify notation by writing (\leq) instead of (\leq_G) , 'coherent' instead of 'G-coherent', and so on, throughout this section.

C.1. **Proof of Lemma 2.37:** Property (1) is a restatement of the fact that, if G is non-degenerate, then every coherent subset of Σ is contained in a coherent complete *-selection. Items (2,3) are straightforward from the definition.

For item (4), observe that, if $u \in \mathfrak{h}(a; G)$ and $a \leq b$, then $b \in u$ as well: indeed, if $b \notin u$, then $b^* \in u$ and $a \leq (b^*)^*$ means u is incoherent. We conclude that $a \leq b$ implies $\mathfrak{h}(a; G) \subseteq \mathfrak{h}(b; G)$, from which (4) readily follows.

To prove item (5), we observe that $S = S \uparrow \text{ implies } S = \min(S) \uparrow$, and then we apply (4).

Finally, one direction of (6) amounts to (3). To prove the converse, suppose $S_1, S_2 \in \mathbf{C}(G)$ are such that $(\dagger) \ \mathfrak{h}(S_1; G) \subseteq \mathfrak{h}(S_2; G)$, and suppose there exists $s \in S_2 \setminus S_1$. If there is a $u \in \mathfrak{h}(S_1; G)$ such that $s \in \min(u)$, then its neighbor $v := [u]_s$ in $\mathtt{Dual}(G)$ (recall Lemma A.4) contains S_1 but not S_2 , contradicting (\dagger) . We are left to prove that $\mathfrak{h}(S_1; G)$ must contain such a u.

Indeed, pick any $v \in \mathfrak{h}(S_1; G)$ such that the number N(v) of $a \in v$ satisfying a < s is smallest possible. Since $s \in S_2$, we have $s \in v$, by (\dagger). If $s \notin \min(v)$ (otherwise we are done), then N(v) > 0 and we may find an $a \in \min(v)$ with a < s. Consider the vertex $[v]_a$: since $S_1 = S_1 \uparrow$, we conclude that $a \notin S_1$, so $S_1 \subset [v]_a$ and $s \in [v]_a$, with $N([v]_a) = N(v) - 1$. This contradicts the choice of v, and we are done.

C.2. Proof of Proposition 3.1: Suppose $B \in G^{\circ}$ is such that $\Delta(A, B) \leq \Delta(A, u)$ for all $u \in G^{\circ}$. We must show that $\operatorname{coh}(A) \subseteq B$.

Suppose $a \in \operatorname{coh}(A) \setminus B$. Then $a^* \in B$ and there is an element $b \in \min(B)$ with $b \leq a^*$. Note that $u := [B]_{b}$ is then also an element of G° , by Lemma A.4.

Now, if $b \in A$, then $a \in A \uparrow *$, contradicting $a \in \operatorname{coh}(A) = A \uparrow \backslash A \uparrow *$. Therefore, $b^* \in A$ (since A is a complete *-selection), but then u satisfies:

(43)
$$\Delta(A, u) = |u \setminus A| = |B \setminus A| - 1 < \Delta(A, B) \le \Delta(A, u)$$

—a contradiction again. We conclude that $coh(A) \subseteq B$, as desired.

C.3. **Proof of Proposition 3.2:** Recall that $A \subseteq A\uparrow$, $A\uparrow\uparrow = A\uparrow$ and $A^* \downarrow = A\uparrow^*$ for all $A \subseteq \Sigma$. We check that $\operatorname{coh}(A)$ is coherent for all A. For suppose that $b, c \in \operatorname{coh}(A)$ satisfy $b \leq c^*$. Then $b \in A\uparrow$ implies $c^* \in A\uparrow\uparrow = A\uparrow$, and therefore $c \in A\uparrow^*$. But then c cannot lie in $\operatorname{coh}(A)$.

Next, we verify that $\operatorname{coh}(A)$ is forward-closed. It suffices to verify $\operatorname{coh}(A) \uparrow \subseteq$ $\operatorname{coh}(A)$. By definition we have $\operatorname{coh}(A) \subseteq A \uparrow$, hence $\operatorname{coh}(A) \uparrow \subseteq A \uparrow \uparrow = A \uparrow$, and it remains to check that no $b \in \operatorname{coh}(A) \uparrow$ belongs to $A \uparrow *$; were there such a b, there would have been $a \in \operatorname{coh}(A)$, $c \in A$ with $a \leq b$ and $c \leq b^*$, implying $a \leq c^*$ — a contradiction to $a \notin A^* \downarrow = A \uparrow *$. This proves (a).

Now let us calculate: $\operatorname{coh}(\operatorname{coh}(A)) = \operatorname{coh}(A) \uparrow \operatorname{\smallsetminus} \operatorname{coh}(A) \uparrow * = \operatorname{coh}(A) \setminus \operatorname{coh}(A)^* = \operatorname{coh}(A)$, the last equality due to $\operatorname{coh}(A)$ being coherent. At the same time, if A itself is coherent then $\operatorname{coh}(A) = A \uparrow \supseteq A$. Moreover, this shows $\operatorname{coh}(A) = A$ whenever $A \in \mathbf{C}(G)$. Finally, if $A = \operatorname{coh}(A)$ then $A \in \mathbf{C}(G)$ because $\operatorname{coh}(A)$ always is.

C.4. **Proof of Proposition 2.41:** The proof of the projection formula will require additional notions and results from [43], which we now recall.

C.4.1. Separators and Gates.

Definition C.1. For any $K, L \subseteq G^{\circ}$, the set

(44)
$$\operatorname{sep}(K,L) := \{a \in G | K \subseteq \mathfrak{h}(a;G)\}, \ L \subseteq \mathfrak{h}(a^*;G)\}$$

is called the separator of K and L in G° .

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

Definition C.2. Let $K, L \subseteq G^{\circ}$. A gate for K, L is a pair of points $u \in K, v \in L$ such that $\Delta(u, v) = |sep(K, L)|$.

The following result is well known in our setting:

Proposition C.3. Let K, L be non-empty convex subsets of a median graph and let $u \in K$ and $v \in L$. Then u, v form a gate for K, L if and only if $\operatorname{proj}_K v = u$ and $\operatorname{proj}_L u = v$. Moreover, the pair K, L has a gate.

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

Lemma C.4. Suppose $S \subset \Sigma$ is coherent, and $K = \mathfrak{h}(S;G)$. Then, for any $a \in \Sigma$, if $K \subseteq \mathfrak{h}(a;G)$ then there exists $s \in S$ such that $s \leq a$.

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

Suppose there existed a $w \in L$ with $w \in \mathfrak{h}(s; G)$, and consider $m = \operatorname{med}(u, v, w)$. Then, $a^* \in v, w$ implies $a^* \in m$, but the inequality

(45)
$$\Delta(u,v) = \Delta(u,m) + \Delta(m,v) \ge \Delta(u,m)$$

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

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

Lemma C.5. Suppose K, L are non-empty convex subsets of Dual(G). 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 $\operatorname{proj}_L(v) = v$, so $K \cap L \subset \operatorname{proj}_L K$. For the reverse inclusion, suppose $v \in \operatorname{proj}_L K$ and write $v = \operatorname{proj}_L u$, $u \in K$. Pick any point $w \in K \cap L$. Setting $m = \operatorname{med}(w, v, u)$ we note that $m \in L$ (because $w, v \in L$) and

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

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

We are now ready for the proof of one more lemma.

C.4.2. Proof of Lemma 2.38. Since $T \in \mathbf{C}(G)$ and u is a complete *-selection, we have $\operatorname{sep}(\mathfrak{h}(T), u) = u^* \cap T = T \setminus u$. Since $S \subseteq u$, we have $T \setminus u \subseteq T \setminus S$. Overall, this yields $\Delta(u, \mathfrak{h}(T)) = |\operatorname{sep}(u, \mathfrak{h}(T))| \leq |T \setminus S|$, as required. \Box

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRA

C.4.3. Computing Nearest Point Projection Maps. We now offer an explicit construction of a geodesic path in Dual(G) emanating from a given vertex u and terminating at its unique nearest point in a specified convex target set:

Proposition C.6. Suppose $u \in G^{\circ}$ is a vertex. Let $T \subseteq \Sigma$ be a coherent subset. Then the following algorithm constructs a shortest path in Dual(G) from u to $K = \mathfrak{h}(T;G)$:

- (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 $[u]_{a}$ and return to the first step.

Proof. We have $u \in K$ if and only if $T \subset u$, 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 $v \in \mathfrak{h}(b; G)$ and $u \in \mathfrak{h}(b^*; G)$. Since $c \leq b^*$, we have $u \in \mathfrak{h}(c; G) \subseteq \mathfrak{h}(b^*; G)$, implying $v \in \mathfrak{h}(c^*; G)$ and $c \in u \setminus v$. As a result:

(46)
$$\Delta(v, [u]_{\circ}) = \Delta(v, u) - 1$$

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

Corollary C.7 (Projection of a Point). Let G and T be as above. Then the closest point projection to $K = \mathfrak{h}(T; G)$ is given by the formula:

(47)
$$\operatorname{proj}_{K} u = (u \smallsetminus T^{*} \downarrow) \cup T \uparrow = (u \cup T \uparrow) \smallsetminus T^{*} \downarrow .$$

Proof. Note that the second equality follows from the DeMorgan rules and the fact that $T \uparrow \cap T^* \downarrow = \emptyset$, since T is coherent. We now prove the first equality.

Set $K = \mathfrak{h}(T; G)$ 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

(48)
$$u \smallsetminus T^* \downarrow \cup T = u \cup T = u,$$

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

(49)
$$\operatorname{proj}_{K} u = \operatorname{proj}_{K} v = (v \smallsetminus T^{*} \downarrow) \cup T \uparrow = (u \smallsetminus T^{*} \downarrow) \cup T \uparrow,$$

the last equality being due to $a \in T^*$ and $a^* \in T$.

C.4.4. Projecting a Convex Set to a Convex Set.

Proposition C.8. Let K, L be non-empty convex subsets of Dual(G) with $L = \mathfrak{h}(S;G)$ and $K = \mathfrak{h}(T;G)$. Then:

(50)
$$proj_{K}L = \mathfrak{h}((S\uparrow \cup T\uparrow) \smallsetminus T^{*}\downarrow;G)$$
$$= \mathfrak{h}(T;G) \cap \mathfrak{h}(S\uparrow \smallsetminus T\uparrow^{*};G) .$$

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

$$\begin{split} \mathfrak{h}((S \uparrow \cup T \uparrow) \smallsetminus T \uparrow {}^*;G) &= \mathfrak{h}(T \uparrow \cup (S \uparrow \smallsetminus T \uparrow {}^*);G) \\ &= \mathfrak{h}(T \uparrow;G) \cap \mathfrak{h}(S \uparrow \smallsetminus T \uparrow {}^*;G) \end{split}$$

and the second equality in Equation (50) follows from the identity $\mathfrak{h}(T;G) = \mathfrak{h}(T\uparrow ;G)$. Denote $R = S\uparrow \backslash T\uparrow *$ and $N = \mathfrak{h}(R;G)$.

For every $u \in L = \mathfrak{h}(S; G)$ we have $S \uparrow \subset u$, implying $\operatorname{proj}_{K} u$ contains $T \uparrow \cup R$, by Corollary C.7. Thus, $\operatorname{proj}_{K} L \subset K \cap N$, as required.

For the converse, observe that the case $K \cap L \neq \emptyset$ was already dealt with in Lemma C.5: if $K \cap L \neq \emptyset$, then

$$\operatorname{proj}_{K} L = K \cap L = \mathfrak{h}(S\uparrow;G) \cap \mathfrak{h}(T\uparrow;G) = \mathfrak{h}(S\uparrow \cup T\uparrow;G)$$

In particular, $S \uparrow \cup T \uparrow$ is coherent, and hence does not intersect $T^* \uparrow$, and the formula Equation (50) holds.

Thus we may henceforth assume $K \cap L = \emptyset$. Equivalently, $S \uparrow \cap T^* \downarrow \neq \emptyset$. In fact, by Lemma C.4 we have $S \uparrow \cap T^* \downarrow = \operatorname{sep}(A, B)$.

Starting with $v \in K \cap N$ we must show $v \in \operatorname{proj}_K L$. Set $u = \operatorname{proj}_L v, w = \operatorname{proj}_K u$, and $m = \operatorname{med}(u, v, w)$. 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, $w = \operatorname{proj}_K u \in I(u, v)$ and $u = \operatorname{proj}_L w$. By Proposition C.3, the pair u, w is a gate for K, L and we have

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

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

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

Now, $\Delta(u, m) > 0$ since $u \in \mathfrak{h}(a^*; G)$ and a contradiction to $u \operatorname{proj}_L v$ is obtained. Thus, $\mathfrak{h}(a; G) \cap L$ must be empty, which means $L \subseteq \mathfrak{h}(a^*; G)$. Applying Lemma C.4 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 $h^* \notin \operatorname{sep}(L, K) = S \uparrow \cap T \uparrow *$, which forces $h^* \in R$. Since $R \subset v$ (by choice of v), we have $h^* \in v$, contradicting our choice of h.

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

Corollary C.9. Let $S, T \subset P$ be subsets and suppose S is coherent. Let $L = \mathfrak{h}(S;G)$ and $K = \mathfrak{h}(\mathfrak{coh}_T(;)G)$. Then:

(51)
$$\operatorname{proj}_{K}L = (S \uparrow \cup T \uparrow) \smallsetminus T \uparrow^{*} = (S \uparrow \smallsetminus T \uparrow^{*}) \cup \operatorname{coh}(T).$$

Proof. Recall that $\operatorname{coh}(T) = T \uparrow \backslash T \uparrow *$, and set $J = T \uparrow \cap T \uparrow *$, so that $T \uparrow = \operatorname{coh}(T) + J$ and $T \uparrow * = \operatorname{coh}_T(*) + J$. Then,

$$\begin{split} (S \uparrow \cup T \uparrow) \smallsetminus T \uparrow^* &= ((S \uparrow \cup \operatorname{coh}(T) \cup J) \smallsetminus \operatorname{coh}(T)^*) \smallsetminus J \\ &= (S \uparrow \cup \operatorname{coh}(T)) \smallsetminus \operatorname{coh}(T)^* \end{split}$$

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

APPENDIX D. APPENDIX: QUALITATIVE SNAPSHOTS (PROOFS)

D.1. **Proof of Lemma 4.7.** For all $a, b, c \in \Sigma$ one has $\mathfrak{h}(ac^*) = \mathfrak{h}(abc^*) \cup \mathfrak{h}(ab^*c^*) \subseteq \mathfrak{h}(bc^*) \cup \mathfrak{h}(ab^*)$. Thus, either the minimum of κ over $\mathfrak{h}(ac^*)$ is attained at a point of $\mathfrak{h}(bc^*)$ or it is attained at a point of $\mathfrak{h}(ab^*)$ (or both). Therefore one has $\mathfrak{w}_{ac^*}^{\kappa} \geq \mathfrak{w}_{ab^*}^{\kappa}$ or $\mathfrak{w}_{ac^*}^{\kappa} \geq \mathfrak{w}_{bc^*}^{\kappa}$, as required.

ITERATED BELIEF REVISION UNDER RESOURCE CONSTRAINTS: LOGIC AS GEOMETRES

D.2. **Proof of Proposition 4.9.** Denote $G = \operatorname{Res}(\mathtt{w}, \delta)$ for the rest of this proof. First, we need to show that $ab \in G$ implies $b^*a^* \in G$. This is baked into the definition, as $\mathtt{w}_{b^*a^{**}} = \mathtt{w}_{ab^*}$. Also, $\mathbf{0}a \in G$ is satisfied because $\mathtt{w}_{\mathbf{0}a} = \infty$. Finally, applying Lemma 4.7 we conclude that, for all $a, b \in \Sigma$ one has $a \leq_G b \Rightarrow \mathtt{w}_{ab^*} > \delta$ when $\delta < \infty$, and $a \leq_G b \Rightarrow \mathtt{w}_{ab^*} = \infty$ when $\delta = \infty$. In particular, were $a \in N(G) \cap N(G)^*$, then $a \leq_G a^*$ would have implied $\mathtt{w}_a = \mathtt{w}_{aa^{**}} > \delta$ (or equals ∞ if $\delta = \infty$), while $a^* \leq_G a$ would have given $\mathtt{w}_{a^*} = \mathtt{w}_{a^*a^*} > \delta$ (or ∞ , respectively)). But that would have meant $\mathtt{w}_{\varnothing} \in (\delta, \infty]$ — a contradiction.

D.3. **Proof of Corollary 4.10.** With $r = \mathbf{w}_{ab} \geq \mathbf{w}_{\varnothing}$, we consider the PCR $G = \operatorname{\mathbf{Res}}(\mathbf{w}_{\bullet}; r)$, for which we have $p \leq_G q$ if and only if $\mathbf{w}_{pq^*} > r$. By Proposition 4.9, G is non-degenerate, hence there exists $u \in G^{\circ} \subseteq \mathbb{H}$. We set $\nu = \delta_{\nu,r}$. For any $p, q \in u$, since $p \not\leq q^*$, we must have $\mathbf{w}_{pq} \leq r = \mathbf{w}_{pq}^{\nu} = \mathbf{w}_{ab}$. At the same time, if $\{p,q\} \not\subseteq u$, then $\mathbf{w}_{pq} \leq \infty = \mathbf{w}_{pq}^{\kappa}$ again. Thus u is the desired vertex of \mathbb{H} . \Box

D.4. **Proof of Proposition 4.11.** Sufficiency follows from Lemma 4.7 and the observations following Definitions 4.1 and 4.5. Now, suppose \mathbf{w}_{\bullet} is a 2-ranking, and consider the set \mathscr{K} of all rankings κ satisfying $\mathbf{w}_{ab}^{\kappa} \geq \mathbf{w}_{ab}$ for all $a, b \in \Sigma$. By Example 4.4, the family \mathscr{K} is closed under taking pointwise minima. Since Σ is finite, \mathscr{K} must have a minimum element.

Let $\widehat{\mathbf{w}}$ be given by Equation (16). To prove that it coincides with the minimum of \mathscr{K} it suffices to verify that (a) $\widehat{\mathbf{w}} \leq \kappa$ for all $\kappa \in \mathscr{K}$, and that (b) $\widehat{\mathbf{w}}$ agrees with \mathbf{w}_{\bullet} .

Fix $\kappa \in \mathscr{K}$. Then, for any $u \in \mathbb{H}$ and $a, b \in u$ we have $\mathbf{w}_{ab}^{\kappa} \leq \kappa(u)$ because κ is a ranking, and $\mathbf{w}_{ab} \leq \mathbf{w}_{ab}^{\kappa}$ by the particular choice of κ , proving (a). Finally, to prove (b), it suffices to verify that, for every $a, b \in \Sigma$, there exists a ranking $\nu \in \mathscr{K}$ with $\mathbf{w}_{ab}^{\nu} = \mathbf{w}_{ab}$. Indeed, Corollary 4.10 provides just such a ranking ν , setting $r = \mathbf{w}_{ab}$, which finishes the proof.

D.5. **Proof of Proposition 4.16.** The first immediate observation regarding minsets is the following observation: An immediate result is this:

Lemma D.1. Let $\delta, \epsilon \geq 0$, and let \mathbf{w}_{\bullet} be a 2-ranking. Let $G = \mathbf{Der}(\mathbf{w}_{\bullet}; \delta)$ and $M = \mathbb{M}(\mathbf{w}_{\bullet}; \epsilon)$. Then $M \in \mathbf{C}(G)$ and $\mathfrak{h}(M; G) \neq \emptyset$.

Proof. Let $G = \mathbf{Der}(\mathbf{w}_{\bullet}, \delta)$ and $M = \mathbf{M}(\mathbf{w}_{\bullet}, \epsilon)$ for some $\delta, \epsilon \geq 0$. Clearly, M is a *-selection. Now suppose that $a, b \in \Sigma$ satisfy $ab \in G$. If $a \in M$, then $\mathbf{w}_a < \mathbf{w}_{a^*} - \epsilon$ and $\mathbf{w}_{ab^*} > \mathbf{w}_{ab}, \mathbf{w}_{a^*b^*}$. But $\mathbf{w}_{b^*} = \min\{\mathbf{w}_{ab^*}, \mathbf{w}_{a^*b^*}\}$ then implies $\mathbf{w}_{b^*} = \mathbf{w}_{a^*b^*}$, and hence also $\mathbf{w}_{a^*} \leq \mathbf{w}_{a^*b^*} = \mathbf{w}_{b^*}$. Similarly, $\mathbf{w}_a = \min\{\mathbf{w}_{ab^*}, \mathbf{w}_{ab}\} = \mathbf{w}_{ab} \geq \mathbf{w}_{b}$, and we have $\mathbf{w}_b \leq \mathbf{w}_a < \mathbf{w}_{a^*} - \epsilon \leq \mathbf{w}_{b^*} - \epsilon$, proving $b \in M$, and we conclude that M is a forward-closed *-selection. In particular, it is G-coherent.¹⁴

To prove Proposition 4.16, we need to analyze the relationship between level sets of the rankings κ and $\hat{\kappa}$. We have the following lemma:

Lemma D.2. Let κ be a ranking in \mathbb{H} and fix a value $r \in \widehat{\mathbb{N}}$, $\mathbf{w}_{\emptyset}^{\kappa} \leq r < \infty$. For the sub-level sets $F = [\kappa \leq r]$ and $\widehat{F} = [\widehat{\kappa} \leq r]$ one has: (a) $F \subseteq \widehat{F}$, (b) $F^{\sharp} = \widehat{F}^{\sharp}$, and (c) $\widehat{F} \subseteq \mathfrak{h}(F^{\sharp})$.

¹⁴For, suppose $a, b \in M$ and we had $a \leq_G b^*$; then we would have also had $b^* \in M$ because M is forward-closed, contradicting M being a *-selection.

Proof. Since $\hat{\kappa} \leq \kappa$, we have $F \subseteq \hat{F}$, which, in turn, implies $\hat{F}^{\sharp} \subseteq F^{\sharp}$. Conversely, if $a \in F^{\sharp}$, then $F \subseteq \mathfrak{h}(a)$ implies $\kappa(u) > r$ for all $u \in \mathfrak{h}(a^*)$; equivalently, $\kappa(\mathfrak{h}(a^*)) > r$. Since this information carries over to the 2-restriction of κ , we conclude that $\hat{\kappa}(\mathfrak{h}(a^*)) > r$ as well, which means $a \in \hat{F}^{\sharp}$, verifying (b). Assertion (c) follows directly from (b) via $\hat{F} \subseteq \mathfrak{h}(\hat{F}^{\sharp}) = \mathfrak{h}(F^{\sharp})$.

One can say more about the lowest level sets of κ :

Lemma D.3. In the notation of Lemma D.2, let $G = \mathbf{Der}(\kappa; \delta)$ with $\delta \in \widehat{\mathbb{N}}$. Then:

- (a) If $r \leq \delta + \mathbf{w}_{\varnothing}^{\kappa} \in \widehat{\mathbb{N}}$, then $\widehat{F} \subseteq G^{\circ}$;
- (b) If $\delta = 0$, then $\mathfrak{h}(F^{\sharp}) \cap G^{\circ} \subset \widehat{F}$.

Proof. To prove (a), consider an arbitrary $u \in \widehat{F}$. Since u is a complete *-selection it suffices to show that it is forward-closed with respect to G. Take any $x \in u$ and $y \in \Sigma$. If $xy \in G$, then:

$$\mathbf{w}_{xy^*}^{\kappa} > \delta + \max\{\mathbf{w}_{xy}^{\kappa}, \mathbf{w}_{x^*y^*}^{\kappa}\} \ge \delta + \mathbf{w}_{\varnothing}^{\kappa}.$$

However, if $y \notin u$, then we have $\{x, y^*\} \subseteq u$. Since $\hat{\kappa}(u) \leq r$, we must have $\mathbf{w}_{xy^*}^{\kappa} \leq r$, by Equation (16). This, however, contradicts our assumption regarding r.

To verify (b), suppose $u \in \mathfrak{h}(F^{\sharp}) \cap G^{\circ}$, but $\widehat{\kappa}(u) > r$. Applying Equation (16) again, we conclude there is a pair $x, y \in u$ with $\mathbf{w}_{xy}^{\kappa} > r$. In particular, no element of F is contained in $\mathfrak{h}(xy)$, which leads to the following two complementary cases:

- Both $\mathfrak{h}(xy^*)$ and $\mathfrak{h}(x^*y)$ contain elements of F. Then we have $\mathfrak{w}_{xy}^{\kappa} > r \ge \mathfrak{w}_{xy^*}^{\kappa}, \mathfrak{w}_{x^*y}^{\kappa}$, which means $xy^* \in G$ and contradicts $u \in G^{\circ}$.
- Either $F \cap \mathfrak{h}(x) = \emptyset$ or $F \cap \mathfrak{h}(y) = \emptyset$. In other words, either $x^* \in F^{\sharp}$ or $y^* \in F^{\sharp}$. Since $F^{\sharp} \subseteq u$, we conclude that one of x^*, y^* must lie in u. This contradicts the fact that $x, y \in u$, since u is a complete *-selection.

Thus, either case yields a contradiction, finishing the proof. \Box

We are finally ready to prove Proposition 4.16.

Setting $r = \kappa(\mathbb{H})$ —the minimum value of κ —we apply Lemma D.3(a) to conclude that $\widehat{F} \subseteq G^{\circ}$. From Lemma D.3(b) and Lemma D.2(c) we obtain that $\mathfrak{h}(F^{\sharp};G) =$ $\mathfrak{h}(F^{\sharp}) \cap G^{\circ}$ coincides with \widehat{F} . Since $F \subseteq \widehat{F}$ by Lemma D.2(a) we conclude that $F \subset G^{\circ}$ and we may apply Corollary 2.36 to deduce that \widehat{F} is the convex hull of F in Dual(G). Finally, consider the set $M = \mathfrak{M}(\kappa)$. For any $a \in \Sigma$, at least one of $\kappa(\mathfrak{h}(a)), \kappa(\mathfrak{h}(a^*))$ equals $\kappa(\mathbb{H})$. Therefore, we have $a \in M$ if and only if $\kappa(\mathfrak{h}(a)) = \kappa(\mathbb{H})$ and $\kappa(\mathfrak{h}(a^*)) > \kappa(\mathbb{H})$, if and only if $a \in F^{\sharp}$ for our current choice of $F = [\kappa \leq \kappa(\mathbb{H})]$. This finishes the proof. \Box

APPENDIX E. APPENDIX: REAL-VALUED SNAPSHOTS (PROOFS).

E.1. Proof of Proposition 4.20. We must show that the derived PCR $G = \text{Der}(\mathbf{w}_{\bullet})$ of a real-valued 2-weight \mathbf{w}_{\bullet} satisfying the requirements 1.-5. of Definition 4.17 is non-degenerate. Recall Equation (25), defining G given an assignment of thresholds τ_{\bullet} :

 $ab \in G \quad \Leftrightarrow \quad \mathbf{w}_{ab^*} \Big|_t < \min(\tau_{ab} \cdot \mathbf{w}_{\varnothing}, \mathbf{w}_{ab}, \mathbf{w}_{a^*b^*}, \mathbf{w}_{a^*b}) \text{ or } \mathbf{w}_{ab^*} = \mathbf{w}_{a^*b} = 0 \,,$

Define functions $\omega, \partial : \Sigma \times \Sigma \to \mathbb{R}$ via $\omega(ab) := \mathbf{w}_{a^*b} - \mathbf{w}_{ab^*}$ and $\partial(ab) := \mathbf{w}_{a^*b} + \mathbf{w}_{ab^*}$. From properties 1. and 4. of the 2-weight \mathbf{w}_{\bullet} , one has the identities $\omega(ab) = -\omega(ba)$, $\omega(aa) = 0$ and $\omega(ac) = \omega(ab) + \omega(bc)$. From properties 2. and 5. of the 2-weight w. one also obtains the identities $\partial(ab) = \partial(ba) > 0$, $\partial(aa) = 0$, $\partial(ac) < \partial(ab) + \partial(bc)$, and $\partial(aa^*) = \mathbf{w}_{\emptyset}$.

We are ready to prove the proposition. Suppose $a \leq_G a^* \leq_G a$ for some $a \in \Sigma$, and find a sequence $a_0, \ldots, a_m, \ldots, a_n$ with $a_0 = a, a_m = a^*, a_n = a$ and $a_{k-1}a_k \in G$ for $k = 1, \ldots, n$. We then must have $\omega(aa) = \sum_{k=1}^n \omega(a_{k-1}a_k) \ge 0$, with equality if and only if $\omega(a_{k-1}a_k) = 0$ for all $k = 1, \ldots, n$. By the definition of G, this implies $\mathbf{w}_{a_{k-1}^*a_k} = \mathbf{w}_{a_{k-1}a_k^*} = 0$ for all $k = 1, \dots, n$.

But then we also have $\mathbf{w}_{\emptyset} = \partial(aa^*) \leq \sum_{k=1}^m \partial(a_{k-1}a_k) = 0$, which is only possible when \mathbf{w}_{\bullet} is trivial—a contradiction.

E.2. Proof of Proposition 4.21. The proof of this proposition follows a standard scheme, widely attributed to Chernoff, and is only included here for the sake of completeness.

Recall that the Kullback-Leibler divergence of a $\mathbf{Ber}(q)$ random variable relative to a **Ber**(p) random variable is given by¹⁵:

(52)
$$\mathbf{D}_{\kappa L}(q||p) := q \log \frac{q}{p} + (1-q) \log \frac{1-q}{1-p}.$$

We require the following standard lemma:

Lemma E.1 (KL-divergence bound). Let $p, q \in (0, 1)$, and consider the function $f(\zeta) := e^{-q\zeta} \left(1 - p + pe^{\zeta}\right)$ over the interval $(0, \infty)$. Then

$$\inf_{\zeta>0} f(\zeta) = \begin{cases} \mathbf{e}^{-\mathbf{D}_{KL}\left(q \| p\right)} & \text{if } q > p \\ 1 & \text{otherwise,} \end{cases} \qquad \mathbf{D}_{KL}\left(q \| p\right) := q \log \frac{q}{p} + (1-q) \log \frac{1-q}{1-p} \end{cases}$$

Proof. Differentiating f one obtains:

$$\begin{split} f'(\zeta) &= \mathrm{e}^{-q\zeta} \left(-q(1-p) + p(1-q) \mathrm{e}^{\zeta} \right) \,, \\ f''(\zeta) &= \mathrm{e}^{-q\zeta} \left(q^2(1-p) + p(1-q)^2 \mathrm{e}^{\zeta} \right) > 0 \,. \end{split}$$

The function f has only one critical point:

$$f'(\zeta_0) = 0 \Leftrightarrow p(1-q)\mathbf{e}^{\zeta_0} = q(1-p) \Leftrightarrow \zeta_0 = \log \frac{q(1-p)}{p(1-q)},$$

and the value of f at ζ_0 is the claimed value, $f(\zeta_0) = e^{-\mathbf{D}_{KL}(q||p)}$.

Finally,
$$\zeta_0 > 0$$
 if and only if $q(1-p) > p(1-q)$, which is tantamount to $q > p$.

The setting for learning snapshot weights described in Proposition 4.21 simplifies to the following. Suppose $X \in [0, A]$ is a non-constant random variable, and let
$$\label{eq:alpha} \begin{split} \alpha &:= \frac{\mathbb{E}[\mathbf{X}]}{A}. \text{ We posit a sequence } \mathbf{X}\big|_t, \, t \geq 0 \text{ of i.i.d. random variables } \mathbf{X}\big|_t \sim \mathbf{X}. \\ \text{ We have the following probability bounds:} \end{split}$$

Lemma E.2. Let t be a non-negative integer and let $\delta > 0$. Then, for Y := $\frac{1}{t+1}\sum_{s=0}^{t} \mathbf{X}\Big|_{s}$ one has:

(53)
$$\mathbf{Pr}(|\mathbf{Y} - \mathbb{E}[\mathbf{X}]| \ge \delta) \le \mathbf{e}^{-(t+1)\mathbf{D}_{KL}\left(\beta \| \alpha\right)} + \mathbf{e}^{-(t+1)\mathbf{D}_{KL}\left(1 - \gamma \| 1 - \alpha\right)},$$

where $\beta = \alpha + \frac{\delta}{A}$ and $\gamma := \alpha - \frac{\delta}{A}.$

¹⁵We will always mean the natural base, e, of the logarithm when using the notation $log(\bullet)$.

To prove Proposition 4.21, observe that the first bound—the standard Chernoff bound—guarantees exponentially fast convergence in probability of the empirical snapshot weights $\mathbf{w}_{ab}|_t$ to the mean value of the signal over the domain $\rho(a) \cap \rho(b)$ when we take $\mathbf{X}|_t := \varphi|_t \cdot \delta_{u_t}(\mathfrak{h}(ab))$.

We are left to verify the Chernoff bound. Proof of Lemma E.2: Recall $\alpha := \frac{\mathbb{E}[X]}{A}$, and observe that $0 < \alpha < 1$ because X is non-constant. We proceed in the standard way to obtain a bound for the empirical estimate of the sample mean. For every *fixed value* of t, and recalling $\beta = \alpha + \frac{\delta}{A}$ one has:

$$P = \mathbf{Pr}(\mathbf{Y} \ge \mathbb{E}[\mathbf{X}] + \delta) = \mathbf{Pr}\left(\frac{\mathbf{Y}}{A} \ge \alpha + \frac{\delta}{A}\right) = \mathbf{Pr}\left(\sum_{s=0}^{t} \frac{\mathbf{X}|_{s}}{A} \ge \beta(t+1)\right)$$
$$= \mathbf{Pr}\left(\exp\left(\zeta \sum_{s=0}^{t} \frac{\mathbf{X}|_{s}}{A}\right) \ge \mathbf{e}^{\zeta\beta(t+1)}\right) \le \mathbf{e}^{-\zeta\beta(t+1)}\mathbb{E}\left[\exp\left(\zeta \sum_{s=0}^{t} \frac{\mathbf{X}|_{s}}{A}\right)\right]$$

by Markov's inequality; also, since the $X|_s$ are independent we have:

$$= \mathbf{e}^{-\zeta\beta(t+1)} \prod_{s=0}^{t} \mathbb{E}\left[\mathbf{e}^{\zeta \frac{\mathbf{X}_{s}}{A}}\right] \le \mathbf{e}^{-\zeta\beta(t+1)} \prod_{s=0}^{t} \mathbb{E}\left[1 - \frac{\mathbf{X}_{s}}{A} + \frac{\mathbf{X}_{s}}{A} \mathbf{e}^{\zeta}\right]$$

using the inequality (†) $e^{x\lambda} \leq 1 - x + xe^{\lambda}$ for $x \in [0, 1], \lambda > 0$. Finally, this yields: $P \leq e^{-\zeta\beta(t+1)} (1 - \alpha + \alpha e^{\zeta})^{t+1}$.

Using Lemma E.1 to minimize the right hand side over $\zeta > 0$, we obtain, for every fixed t:

$$P \le \exp\left(-(t+1)\mathbf{D}_{KL}\left(\beta \|\alpha\right)\right) \,,$$

as claimed. Now replace X with A - X, Y with A - Y, and recall $\gamma := \alpha - \frac{\delta}{A}$. We obtain:

$$\mathbf{Pr}(\mathbf{Y} \le \mathbb{E}[\mathbf{X}] - \delta) = \mathbf{Pr}(A - \mathbf{Y} \ge A - \mathbb{E}[\mathbf{X}] + \delta) = \mathbf{Pr}(A - \mathbf{Y} \ge \mathbb{E}[A - \mathbf{X}] + \delta)$$

$$\le \exp\left(-(t+1)\mathbf{D}_{KL}\left(1 - \alpha + \frac{\delta}{A} \| 1 - \alpha\right)\right)$$

$$= \exp\left(-(t+1)\mathbf{D}_{KL}\left(1 - \gamma \| 1 - \alpha\right)\right),$$

finishing the proof.

APPENDIX F. APPENDIX: DEBUGGING SNIFFY ON THE CIRCLE.

The purpose of this appendix is to explain in some detail the reasons for Sniffy's behavior on the circle in its qualitative BUA incarnation, as discussed in Section 5.4.2. We proceed in a manner similar to the discussion of an agent on the interval from Appendix A.2.4.

F.1. Sensors and Relations. Let a_k denote the sensor centered at $k \in \mathbf{E} = \{0, \ldots, 19\}$, reporting 1 at time t if and only if $dist(k, pos(t)) \leq 4$, where pos(t) is the position occupied by Sniffy at time t. It will be convenient to think of \mathbf{E} as a copy of the additive group ZZ_{20} , keeping in mind its action on subsets $S \subseteq \mathbf{E}$ given by $k + S := \{i - k \mid i \in S\}$.

It is then easy to verify that $a_k < a^*_{k+\{9,10,11\}}$ are the only relations among the a_k . Consequently, the analogous relations $\sharp a_k < \sharp a^*_{k+\{9,10,11\}}$ must also hold for all

k (see Appendix A.2.3 for details). Motion is described by the conditional relations rt : $\sharp a_k < a_{k+1}$ and \mathtt{lt} : $\sharp a_k < a_{k-1}$, leading to the unconditional implications $\sharp a_k < a_{k+10}$ (compare with the case of the interval discussed in Appendix A.2.4).

Finally, observing that the entire setting is rotation-invariant, without loss of generality we may assume for the rest of this section that Sniffy's target is located at position T = 0. Setting $M := \{a_0, a^*_{10}\} \cup \{a_{\pm 1}, \ldots, a_{\pm 4}\} \cup \{a_{\pm 5}, \ldots, a_{\pm 9}\}^*$, the *eventual* target sets (minsets) determined by the individual snapshots are:

(54)
$$\begin{aligned} M(\mathbf{rt}^*) &= M(\mathbf{lt}^*) &:= M \cup \sharp M, \\ M(\mathbf{rt}) &:= M \cup \sharp (1+M), \\ M(\mathbf{lt}) &:= M \cup \sharp (-1+M), \end{aligned}$$

where, due to the hard-wired arbitration enforcing $\neg(\texttt{rt} \land \texttt{lt})$ at all times, the derived PCR $G^{\texttt{rt}}$ of the rt snapshot identifies each a_k with $\sharp a_{k-1}$; and, similarly, $G^{\texttt{lt}}$ identifies each a_k with $\sharp a_{k+1}$.

Finally, note that given \mathbf{rt}^* it is impossible to witness $a_{k+9} \wedge \sharp a_k$ or $a_{k+10} \wedge \sharp a_k$ (while $a_{k+11} = a_{k-9}$ is still possible in conjunction with $\sharp a_k$, if \mathbf{lt} is active). We conclude that the relations $\sharp a_k < a^*_{k+9}, a^*_{k+10}$ hold in $G^{\mathbf{rt}^*}$ for every k, as do their analogous counterparts in $G^{\mathbf{lt}^*}$.

F.2. The "dull peak" value signal. The weights recorded on any snapshot in this case are $\{0, 1, \infty\}$, implying that (1) any implications appearing in Sniffy's four snapshots at any time are a subset of the implications listed in the preceding section; and (2) any raw observation generated by Sniffy is coherent for any of its snapshots.

Let $k \in \mathbf{E}$. Each snapshot forms its prediction for the next state by propagating $\sharp(-k+M)$, and giving rise to:

(55)
$$\begin{array}{rcl} \Pr ed^{lt} &=& \sharp (-k+M) \cup (1-k+M) \\ \Pr ed^{lt^*} &=& \sharp (-k+M) \cup \operatorname{coh}_{lt^*} ((-k+M) \cup (1-k+M)) \\ &=& \sharp (-k+M) \cup ((-k+M) \smallsetminus \{a^*_{k-5}, a_{k+4}\}) \end{array}$$

In order to compute the divergences from the targets, we first note that, for $k, \ell \in \mathbf{E}$:

(56)
$$|(k+M) \smallsetminus (\ell+M)| = |M \smallsetminus (\ell-k+M)| = 2\min\{9, \mathtt{dist}(k,\ell)\}$$

For the lt snapshot this results in:

Div(Pred^{it}; M(lt)) = M(lt) \ Pred^{it}
=
$$|\sharp(-1+M) \setminus \sharp(-k+M)| + |M \setminus (1-k+M)|$$

= $|(-1+M) \setminus (-k+M)| + |M \setminus (1-k+M)|$
= $4\min\{9, dist(1,k)\}$.

By symmetry, we conclude that the rt snapshot has:

(58)
$$\operatorname{Div}(\operatorname{Pred}^{\operatorname{rt}}; M(\operatorname{rt})) = 4\min\{9, \operatorname{dist}(-1, k)\}$$

Now, for the lt^{*} snapshot we have:

$$\begin{aligned} \mathtt{Div}(\mathtt{Pred}^{\mathtt{lt}^*}; M(\mathtt{lt}^*)) &= |M \smallsetminus (-k+M)| + |M \smallsetminus ((-k+M) \smallsetminus \{a^*_{k-5}, a_{k+4}\})| \\ &= 4\min\{9, \mathtt{dist}(0, k)\} + \delta(k), \end{aligned}$$

where $\delta(k) = 2$ for $-9 \le k = 0$, $\delta(k) = 0$ for $0 < k \le 9$, $\delta(k) = 1$ for k = 10. By symmetry:

(59)
$$\operatorname{Div}(\operatorname{Pred}^{\operatorname{rt}}; M(\operatorname{rt})) = 4\min\{9, \operatorname{dist}(0, k)\} + \delta(-k).$$

Armed with these formulae, we go over the possibilities and conclude:

- The BUA lt is active if and only if $k \in T + \{1, \dots, 11\} \pmod{20}$;
- The BUA rt is active if and only if $k \in T \{1, \dots, 11\} \pmod{20}$.

In particular, $k \in T \pm \{0, ..., 8\}$ is a basin of attraction for the target position, T, while $k \in T + \{9, 10, 11\}$ is a region where both lt and rt seek to be active, triggering the hard-wired arbitration mechanism.

F.3. The "sharp peak" value signal. Despite the value signal being more informative than that of its $\{0, 1\}$ -valued "dull-peak" counterpart, a "sharp peak" qualitative agent's performance on the target-finding task is clearly worse (Figure 7).

In a nutshell, the reason for this deficiency is that the limiting PCR satisfies the additional relations $a_{\pm 9} < a_{\pm 8} < \ldots < a_{\pm 1}$, which we verified by hand. As a result, properties (1) and (2) stated in Appendix F.2 for the "dull peak" setting will not hold in this one. In fact, just these extra relations (there may be others) suffice for the current state representation of any point other than k = 0, 10 in any of Sniffy's snapshots to degenerate (through coherent projection) into a less and less complete *-selection as Sniffy's physical distance from the target (in the environment) increases, while the quality of prediction deteriorating accordingly.

Specifically, any $k \in B := 10 \pm \{0, 1, 2, 3\}$ yields a raw observation containing a_{10}, a^*_0 and both of a_9 and a_{11} . On one hand, a_9, \ldots, a_1 and a_{11}, \ldots, a_{19} are directed paths in each of Sniffy's PCRs. On the other hand, though, so are a_k, a^*_{k+10} . Thus, (1) the coherent projection of the raw observation generated by k is merely $\{a_{10}, a^*_0\}$, showing that Sniffy is unable to distinguish among the points of B; and (2) if k is moved closer to the target, fewer conflicts of the above form will affect coherent projection.

In total, the observations above suffice for explaining the most visible differences (see Figure 7) between the behaviors of the two variants of Sniffy in the qualitative learning regime.

ELECTRICAL & SYSTEMS ENGINEERING, SCHOOL OF ENGINEERING & APPLIED SCIENCES, UNI-VERSITY OF PENNSYLVANIA, PENN ENGINEERING RESEARCH & COLLABORATION HUB (PERCH), 3401 GRAYS FERRY AVE., PENNOVATION CENTER, BUILDING 6176, 3rd FLOOR, PHILADELPHIA, PA 19146

E-mail address: guraldan@seas.upenn.edu

ELECTRICAL & SYSTEMS ENGINEERING, SCHOOL OF ENGINEERING & APPLIED SCIENCES, UNI-VERSITY OF PENNSYLVANIA, PENN ENGINEERING RESEARCH & COLLABORATION HUB (PERCH), 3401 GRAYS FERRY AVE., PENNOVATION CENTER, BUILDING 6176, 3rd Floor, Philadelphia, PA 19146

E-mail address: kod@seas.upenn.edu