

On the Solvability of Inductive Problems: A Study in Epistemic Topology

Alexandru Baltag

Nina Gierasimczuk

Sonja Smets

Institute for Logic, Language and Computation
University of Amsterdam, The Netherlands

A.Baltag@uva.nl

Nina.Gierasimczuk@gmail.com

S.J.L.Smets@uva.nl

We investigate the issues of inductive problem-solving and learning by doxastic agents. We provide topological characterizations of solvability and learnability, and we use them to prove that AGM-style belief revision is “universal”, i.e., that every solvable problem is solvable by AGM conditioning.

1 Introduction

When in the course of observations it becomes necessary for agents to arrive at a generalization, they should declare, along with their conjecture, the extent of their certainty. The problem of induction seems formidable if a standard of absolute certainty is imposed on the learner. Indeed, as is well-known in Philosophy of Science, the so-called problem of empirical underdetermination (i.e., the fact that typically the data are compatible with more than one hypothesis) rules out any chance of obtaining infallible knowledge in empirical research. But apart from the conclusions based on absolute certainty (cf. [13, 10, 15]), learners can produce hypotheses based on *beliefs*. It is thus strange that Formal Learning Theory and Belief Revision Theory developed completely independently from each other, and that they have generally maintained their distance ever since.

However, there does exist a line of research that combines belief revision with learning-theoretic notions, line pursued by Kelly [21, 26], Kelly, Schulte and Hendricks [19], Martin and Osherson [28] and ourselves [13, 3, 4, 14]. In this paper we continue this research program, using topological characterizations and methods.

An *inductive problem* consists of a state space, a family of “potential observations”, and a “question” (i.e., a partition of the state space). These observations provide data for learning. The problem is *solvable* if there exists a learner that, after observing “enough” pieces of data, eventually stabilizes on the correct answer. A special case of solvability is *learnability in the limit*, corresponding to the solvability of the “ultimate” question: ‘What is the actual state of the world?’. This notion matches the usual learning-theoretic concept of *identifiability in the limit* [31, 16, 29].

The aim of the paper is twofold. First, we give topological characterizations of the notions of solvability (and learnability), in terms of topological separation principles. Intuitively, the ability to reliably learn the true answer to a question, is related to the possibility to “separate” answers by observations. The second goal is to use these topological results to look at the “solving power” of well-behaved doxastic agents, such as the ones whose beliefs satisfy the usual *KD45* postulates of doxastic logic, as well as the standard AGM postulates of rational belief-revision [1]. We look at a particularly simple and canonical type of doxastic agent, who forms beliefs by *AGM conditioning*.

Our main result is that AGM conditioning is *universal for problem-solving*, i.e., that every solvable problem can be solved by AGM conditioning. This means that (contrary to some prior claims), AGM

belief-revision postulates are not an obstacle to problem-solving. As a special case, it follows that AGM conditioning is also “universal for learning” (every learnable space can be learned by conditioning).¹

The close connections between Epistemology and General Topology have already been noticed long ago [32, 20]. Based on these connections, Kevin Kelly started a far-reaching program [20, 22] meant to import ideas and techniques from both Formal Learning Theory and Topology into mainstream Epistemology, and show their relevance to the induction problem in Philosophy of Science. A further connection is the one with Ockham’s Razor, that would

(...) guarantee that always choosing the simplest theory compatible with experience and hanging on to it while it remains the simplest is both necessary and sufficient for efficiency of inquiry. [22]

Simplicity has been claimed to have topological characteristics—the simplicity order should in some way follow the structure imposed on the uncertainty range by possible tests and observations. It has also been linked with the notion of minimal mind change, where the learning agent keeps the conjecture changes to a minimum [20, 30].

Taken together, our results can be seen as a vindication both of the general topological program in Inductive Epistemology [20, 22] and of the AGM Belief Revision Theory [1]. On the first front, our general topological characterizations of learning-theoretic concepts seem to confirm Kelly’s long-standing claim that Inductive Epistemology can be seen mathematically as a branch of General Topology. On the second front, our universality result seems to vindicate Belief Revision Theory as a canonical form of learning.²

2 Epistemic Spaces and Inductive Problems

Definition 1. An epistemic space is a pair $\mathbb{S} = (S, \mathcal{O})$ consisting of a state space S and a countable (or finite) set of observable properties (“data”) $\mathcal{O} \subseteq \mathcal{P}(S)$. We denote by $\mathcal{O}_s := \{O \in \mathcal{O} \mid s \in O\}$ the set of all observable properties (holding) at a given state s .

One can think of the states in S as “possible worlds”, in the tradition of Kripke and Lewis. The sets $O \in \mathcal{O}$ represent properties of the world that are in principle observable: if true, such a property will eventually be observed (although there is no upper bound on the time needed to come to observe it).

To keep things simple, we assume that at each step of the learning process only one property is observed. As for the countability of the set \mathcal{O} , it is natural to think of observables as properties which can be expressed by means of a language or numerical coding system, generated from a grammar with a finite vocabulary. Any such family \mathcal{O} will be (at most) countable.

We denote by \mathcal{O}^\cap the family of all finite intersections of observations from \mathcal{O} , and by \mathcal{O}^* the family of all finite sequences of observations. Such a finite sequence $\sigma = (O_0, O_1, \dots, O_i) \in \mathcal{O}^*$ is called a *data sequence*, and its i -th component is denoted by $\sigma_i := O_i$. It is easy to see that both \mathcal{O}^\cap and \mathcal{O}^* are countable.

A *data stream* is a countable sequence $\vec{O} = (O_0, O_1, \dots) \in \mathcal{O}^\omega$ of data from \mathcal{O} (here, ω is the set of natural numbers, so \mathcal{O}^ω is the set of all maps assigning an observable property to every natural number). We use the following notation: \vec{O}_n is the n -th element in \vec{O} ; $\vec{O}[n]$ is the initial segment of \vec{O} of length n ,

¹This special case is a topological translation of one of our previous results [3, 4]. However, the result about problem-solving universality is not only new and much more general, but also much harder to prove, involving new topological notions and results.

²And in the same time (if we adopt a “simplicity” interpretation of the prior), this last result can be seen as a vindication of Ockham’s razor (in line with Kevin Kelly’s program).




Figure 1: A problem \mathbb{P} (left-hand side) and its refinement \mathbb{P}' (right-hand side), see Example 3

(O_0, \dots, O_{n-1}) ; $\text{set}(\vec{O}) := \{O \mid O \text{ is an element of } \vec{O}\}$ is the set of all data in \vec{O} ; $*$ is the concatenation operator on strings.

The intuition is that at stage n of a data stream, the agent observes the information in O_n . A data stream captures a possible future history of observations in its entirety, while a data sequence captures only a finite part of such a history.

Given a state $s \in S$, a *data stream for s* is a stream $\vec{O} \in \mathcal{O}^\omega$ such that $\mathcal{O}_s = \{O \in \mathcal{O} \mid \bigcap_{i=0}^n O_i \subseteq O \text{ for some } n \in \omega\}$. Such a stream is “sound” (every data in \vec{O} is true at s) and “complete” (every true data is entailed by some finite set of observations in \vec{O}).

Example 1. Let our epistemic space $\mathbb{S} = (S, \mathcal{O})$ be the real numbers, with observable properties given by open intervals with rational endpoints: $S := \mathbb{R}$, $\mathcal{O} := \{(a, b) \mid a, b \in \mathbb{Q}, a \leq b\}$, where $(a, b) := \{x \in \mathbb{R} \mid a < x < b\}$. For instance, observables may represent measurements of a physical quantity (such as a position along a one-dimensional line) that takes real numbers as its possible values. In such case, for any state $x \in \mathbb{R}$ and any two sequences $a_n, b_n \in \mathbb{Q}$ of rational numbers, such that $a_n \leq x \leq b_n$ and both sequences converge to x , the sequence $(a_0, b_0), \dots, (a_n, b_n), \dots$ is a (sound and complete) data stream for x .

Other examples include standard n -dimensional Euclidean spaces, e.g., $S = \mathbb{R}^3$ with \mathcal{O} consisting of all open balls with rational radius and center.

Definition 2. An inductive problem is a pair $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ consisting of an epistemic space $\mathbb{S} = (S, \mathcal{O})$ together with a “question” \mathcal{Q} , i.e., a partition³ of S . The cells A_i of the partition \mathcal{Q} are called answers. Given $s \in S$, the unique $A \in \mathcal{Q}$ with $s \in A$ is called the answer to \mathcal{Q} at s , and denoted A_s . We say that a problem $\mathbb{P}' = (\mathbb{S}, \mathcal{Q}')$ is a refinement of another problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ (or that the corresponding question \mathcal{Q}' is a refinement of the question \mathcal{Q}) if every answer of \mathcal{Q} is a disjoint union of answers of \mathcal{Q}' .

The most refined question concerns the identity of the real world.

Example 2. The learning question on a space S is $\mathcal{Q} = \{\{s\} \mid s \in S\}$ and corresponds to ‘What is the actual state?’.

Example 3. Let $\mathbb{S} = (S, \mathcal{O})$, where $S = \{s, t, u, v\}$, $\mathcal{O} = \{U, V, P, Q\}$, with $U = \{s, t\}$, $V = \{s\}$, $P = \{u, v\}$, $Q = \{u\}$. Take the problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$, given by the question $\mathcal{Q} = \{\{t, u\}, \{s, v\}\}$ depicted on the left-hand side of Figure 1. This can obviously be refined to obtain the problem $\mathbb{P}' = (\mathbb{S}, \mathcal{Q}')$ given by the learning question $\mathcal{Q}' = \{\{s\}, \{t\}, \{u\}, \{v\}\}$ for this space, as depicted on the right-hand side of Figure 1.

³This means that $\bigcup_{i \in I} A_i = S$, and $A_i \cap A_j = \emptyset$ for all $i \neq j$.

3 Learning and Problem-Solving

Definition 3. Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space and let $\sigma_0, \dots, \sigma_n \in \mathcal{O}$. An agent (also called a “learner”, or a “learning method”) is a map \mathcal{L} that associates to any epistemic space \mathbb{S} and any data sequence $(\sigma_0, \dots, \sigma_n)$ some family $\mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n) \subseteq \mathcal{P}(S)$ of subsets of S , satisfying a “consistency” condition: $\emptyset \notin \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$ whenever $\bigcap_{i=0}^n \sigma_i \neq \emptyset$.

Intuitively, after observing the data sequence $\vec{\sigma} = (\sigma_0, \dots, \sigma_n)$, we can say that agent \mathcal{L} believes a proposition P after observing the data sequence $\vec{\sigma} = (\sigma_0, \dots, \sigma_n)$, and write $B_{\mathcal{L}}^{\vec{\sigma}}P$ iff $P \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$. We can also interpret this as a *conditional belief*, rather than as revised belief, the agent believes every $P \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$ conditional on $\sigma_0, \dots, \sigma_n$. However, in the end we are of course interested in the actual revised beliefs after observing the data, so the assumption in this case is that conditional beliefs guide the agent’s revision strategy: they “pre-encode” future belief revisions, to use a term coined by Johan van Benthem [6]. The above consistency simply means that each of the agent’s beliefs is consistent whenever the observed data are consistent.

A *doxastic agent* is one whose set $\mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$ of beliefs forms a (*proper*) *filter* on S when observing consistent data; in other words, her beliefs are (consistent when possible, and also) *inference-closed* (i.e., if $P \subseteq Q$ and $P \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$, then $Q \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$) and *conjunctive* (i.e., if $P, Q \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$ then $(P \cap Q) \in \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$). Hence, for any doxastic agent \mathcal{L} and every consistent data sequence $\vec{\sigma}$, the belief operator $B_{\mathcal{L}}^{\vec{\sigma}}$ (as defined above) satisfy the usual *KD45* axioms of doxastic logic.

A *standard agent* is a doxastic agent \mathcal{L} whose beliefs form a *principal filter*, i.e., all her beliefs are entailed by one “strongest belief”; formally, a doxastic agent \mathcal{L} is standard iff for every data sequence $\vec{\sigma}$ over any epistemic space \mathbb{S} there exists some set $L_{\mathbb{S}}(\vec{\sigma})$, such that

$$\mathcal{L}_{\mathbb{S}}(\vec{\sigma}) = \{P \subseteq S \mid L_{\mathbb{S}}(\vec{\sigma}) \subseteq P\}.$$

It is easy to see that in this case, we must have $L_{\mathbb{S}}(\vec{\sigma}) = \bigcap \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$. Indeed, we can equivalently define a doxastic agent \mathcal{L} to be standard iff $\bigcap \mathcal{L}_{\mathbb{S}}(\vec{\sigma}) \in \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$ holds for all data sequences $\vec{\sigma}$. *Standard agents are globally consistent* whenever possible: $\bigcap \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n) \neq \emptyset$ whenever $\bigcap_{i=0}^n \sigma_i \neq \emptyset$.

Traditional learning methods in Formal Learning Theory correspond to our standard agents, and they are typically identified with the map L (given by $L_{\mathbb{S}}(\sigma_0, \dots, \sigma_n) := \bigcap \mathcal{L}_{\mathbb{S}}(\sigma_0, \dots, \sigma_n)$). From now on we follow this tradition, and refer to standard agents using the map L . But in general we do *not* restrict ourselves to standard agents.

An *AGM agent* is an agent \mathcal{L}^{\leq} who forms beliefs by *AGM conditioning*, i.e., it comes endowed with a map that associates any epistemic space \mathbb{S} some total preorder⁴ $\leq_{\mathbb{S}}$ on S , called “*prior*” *plausibility relation*; and whose beliefs after observing any data sequence $\vec{\sigma} = (\sigma_0, \dots, \sigma_n)$ are given by

$$\mathcal{L}_{\mathbb{S}}^{\leq}(\vec{\sigma}) := \{P \subseteq S \mid \exists s \in \bigcap_{i=0}^n \sigma_i \forall t \in \bigcap_{i=0}^n \sigma_i (t \leq s \Rightarrow t \in P)\}.$$

Intuitively, $t \leq s$ means that t is *at least as plausible* as s (according to our agent). So, an AGM agent believes P conditional on a data sequence $\vec{\sigma}$ iff P is true in all the states (consistent with the data) that are “plausible enough”.

⁴A total preorder on S is a binary relation \leq on S that is reflexive, transitive, and connected (i.e., for all $s, t \in S$, we have either $s \leq t$ or $t \leq s$).

It is easy to see that *every AGM agent is a doxastic agent*: $\mathcal{L}_{\mathbb{S}}^{\leq}(\vec{\sigma})$ is a proper filter whenever $\bigcap_{i=0}^n \sigma_i \neq \emptyset$; hence, the beliefs of an AGM agent satisfy the usual *KD45* axioms of doxastic logic (when learning any consistent data sequence).

Moreover, it is well-known that in fact, *the beliefs of AGM agents satisfy all the so-called AGM axioms from Belief Revision Theory* [1]: if, for any data sequence $\vec{\sigma} = (\sigma_0, \dots, \sigma_n)$, we set $T = \mathcal{L}(\sigma_0, \dots, \sigma_n)$, and for any new observation $\phi \in \mathcal{O}$ we set $T * \phi = \mathcal{L}(\sigma_0, \dots, \sigma_n, \phi)$, then the resulting revision operator $*$ satisfies all the AGM postulates. In fact, for any AGM agent \mathcal{L} , if we interpret the operator $B_{\mathcal{L}}^{\vec{\sigma}}$ (as defined above) as representing a conditional belief $B^{\sigma_0 \wedge \dots \wedge \sigma_n}$, then the sound and complete logic of these conditional belief operators is the so-called Conditional Doxastic Logic [8, 5] (which is itself just a repackaging of the AGM postulates in the language of conditional logic).

Observation 1. *Given a total preorder \leq on S and a subset $A \subseteq S$, set*

$$\text{Min}_{\leq}(A) := \{s \in A \mid s \leq t \text{ for all } t \in A\}$$

for the set of \leq -minimal states in A . Let $\vec{\sigma} = (\sigma_0, \dots, \sigma_n)$ be any data sequence such that $\text{Min}_{\leq}(\bigcap_{i=0}^n \sigma_i) \neq \emptyset$. Then $\mathcal{L}_{\mathbb{S}}^{\leq}(\vec{\sigma})$ is the principal filter generated by $\text{Min}_{\leq}(\bigcap_{i=0}^n \sigma_i)$, i.e., we have

$$\mathcal{L}_{\mathbb{S}}^{\leq}(\sigma_0, \dots, \sigma_n) := \{P \subseteq S \mid \text{Min}_{\leq}(\bigcap_{i=0}^n \sigma_i) \subseteq P\}.$$

In general though, the filter $\mathcal{L}_{\mathbb{S}}^{\leq}(\vec{\sigma})$ is not principal. So AGM agents are *not* necessarily standard agents. But there is an important case when they are standard: whenever the preorder $\leq_{\mathbb{S}}$ is well-founded in every space \mathbb{S} (i.e., there are no infinite chains $s_0 > s_1 > s_2 \dots$ of more and more plausible states). It is easy to see that the map L associated to a standard AGM agent is given by the set of \leq -minimal states consistent with the data:

$$L_{\mathbb{S}}^{\leq}(\sigma_0, \dots, \sigma_n) := \text{Min}_{\leq}(\bigcap_{i=0}^n \sigma_i).$$

Intuitively, this means that a standard AGM agent believes a proposition P iff P is true in all the “most plausible” states consistent with the data.

The original semantics of AGM belief was given using only standard AGM agents. But this semantics was in fact borrowed by Grove [18] from Lewis’ semantics for conditionals [27], which did *not* assume well-foundedness.⁵

Definition 4. *Let \mathbb{S} be an epistemic space. An agent \mathcal{L} verifies a proposition $A \subseteq S$ in the limit if, for every state $s \in S$ and every data stream \vec{O} for s , we have $s \in A$ iff there exists some $k \in \omega$ such that $A \in \mathcal{L}_{\mathbb{S}}(\vec{O}[n])$ for all $n \geq k$. For standard agents, this means that $L_{\mathbb{S}}(\vec{O}[n]) \subseteq A$ for all $n \geq k$. A set $A \subseteq S$ is verifiable in the limit if there exists some agent that verifies A in the limit.⁶*

An agent \mathcal{L} falsifies a proposition $A \subseteq S$ in the limit if, for every state $s \in S$ and every data stream for \vec{O} for s , we have $s \notin A$ iff there exists some $k \in \omega$ such that $A^c \in \mathcal{L}(\mathbb{S}, \vec{O}[n]) \subseteq A^c$ for all $n \geq k$ (here, as in the rest of this paper, $X^c := S \setminus X$ stands for the complement of X). For a standard agent, this means $L(\mathbb{S}, \vec{O}[n]) \subseteq A^c$ for all $n \geq k$,

A proposition $A \subseteq S$ is falsifiable in the limit if there exists some agent that falsifies A in the limit.

A proposition $A \subseteq S$ is decidable in the limit if it is both verifiable and falsifiable in the limit.

⁵Indeed, Lewis’ definition of conditionals has a similar shape to our above definition of (conditional) beliefs for non-standard AGM agents.

⁶For a discussion of the relationship between verifiability and learnability see, e.g., [20, 12].

An agent \mathcal{L} solves a problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ if, for every state $s \in S$ and every data stream \vec{O} for s , there exists some $k \in \omega$ such that $A_s \in \mathcal{L}_{\mathbb{S}}(\vec{O}[n])$ for all $n \geq k$ (recall that A_s is true answer to \mathcal{Q} at s). For a standard agent, this means that $L_{\mathbb{S}}(\vec{O}[n]) \subseteq A_s$ for all $n \geq k$. A problem is solvable (in the limit) if there exists some agent that solves it.

An epistemic space $\mathbb{S} = (S, \mathcal{O})$ is learnable (by an agent \mathcal{L}) if the (problem given by the) learning question $\mathcal{Q}_{\mathbb{S}} = \{\{s\} \mid s \in S\}$ is solvable (by \mathcal{L}).

All the above notions have a standard counterpart, e.g., A is standardly verifiable if there exist some standard agent that verifies it; \mathbb{P} is standardly solvable if it can be solved by some standard agent, etc.

Note that standard learnability is essentially the same as Gold's *identifiability in the limit* [29, 17].

Examples and Counterexamples: An example of *non-learnable* space $\mathbb{S} = (S, \mathcal{O})$ is obtained by taking four abstract states $S = \{s, t, u, w\}$ and two observable properties $\mathcal{O} = \{V, U\}$, with $V = \{s, t, u\}$ and $U = \{t, u, w\}$, as depicted in Figure 2. Since states s and t satisfy the same observable properties, no learning method will ever distinguish them.




Figure 2: A non-learnable space

But even spaces in which no two states satisfy the same observations can still be non-learnable, e.g., *all the n -dimensional Euclidean spaces from Example 1 are not learnable* (though, as we will see, many questions are solvable and many subsets are decidable over these spaces). Another example of *non-learnable* space is given in Figure 3: formally, $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \omega\} \cup \{s_\infty\}$, and $\mathcal{O} = \{O_i \mid i \in \omega\}$, and for any $i \in \omega$, $O_i := \{s_i, s_{i+1}, \dots\} \cup \{s_\infty\}$.




Figure 3: Another non-learnable space

In contrast, an example of *learnable* space is in Figure 4: formally, $S = \{s_n \mid n \in \omega\}$ consists of countably many distinct states, with $\mathcal{O} = \{O_n \mid n \in \omega\}$, where $O_n = \{s_0, s_1, s_2, \dots, s_n\}$. A standard agent that can learn this space in the limit is given by setting $L(\sigma_1, \dots, \sigma_n)$ to be the *maximum number* (in the natural order) in $\bigcap_{i=0}^n \sigma_i$, whenever there is such a maximum number, and setting $L(\sigma_1, \dots, \sigma_n) := \bigcap_{i=0}^n \sigma_i$ otherwise.

Proposition 1. Let \mathbb{S} be an epistemic space, $A \subseteq S$ a proposition and $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ an inductive problem. Then we have the following:

- A is verifiable (falsifiable, decidable) in the limit iff it is standardly verifiable (falsifiable, decidable) in the limit.
- \mathbb{P} is solvable iff it is standardly solvable.




Figure 4: A learnable space

- \mathbb{S} is learnable iff it is standardly learnable.

Proof. Let $A \subseteq S$ be a set that is verifiable (falsifiable, decidable) by an agent \mathcal{L} on an epistemic space \mathbb{S} . We construct a standard agent that does the same thing, by setting, for every data sequence $\vec{\sigma} \in \mathcal{O}^*$: $L_{\mathbb{S}}(\vec{\sigma}) := A$ if $A \in \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$, $L_{\mathbb{S}}(\vec{\sigma}) := A^c$ if $A \notin \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$ but $A^c \in \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$, and $L_{\mathbb{S}}(\vec{\sigma}) := S$ otherwise. Also, on any *other* space $\mathbb{S}' = (S', \mathcal{O}')$, we set by default $L_{\mathbb{S}'}(\vec{\sigma}') := S'$.

Similarly, let $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ be a problem that is solvable by \mathcal{L} . Let \leq be some arbitrary well-order of the set \mathcal{Q} . (Such a well-order exists, by the Well-Ordering Theorem.) We construct a standard agent who also solves \mathbb{P} , by setting $L_{\mathbb{S}}(\vec{\sigma}) := A$ if A is the first answer in \mathcal{Q} (according to \leq) such that $A \in \mathcal{L}_{\mathbb{S}}(\vec{\sigma})$ holds; and $L_{\mathbb{S}}(\vec{\sigma}) := S$ if no such answer exists. (As before, we can extend our agent to any other space $\mathbb{S}' = (S', \mathcal{O}')$, by setting $L_{\mathbb{S}'}(\vec{\sigma}') := S'$.)

By applying this to the learning problem $\mathcal{Q} = \{\{s\} \mid s \in S\}$, we obtain the similar result for learnability. \square

In conclusion, everything that can be learned by any agent can also be learned by some standard agent. However, this is no longer true when we restrict to more canonical types of agents (such as AGM agents).

Proposition 2. *There exist spaces that are learnable, but not learnable by standard AGM agents. Hence, there exist solvable problems that are not solvable by standard AGM agents.*

Proof. Consider a counterexample from [13, 3, 4]. Take the epistemic model from Figure 4. This space is learnable, and thus learnable by AGM conditioning, but it is not learnable by standard conditioning. Indeed, this space is learnable by conditioning *only* with respect to the following *non-wellfounded* prior: $s_0 > s_1 > \dots > s_n > s_{n+1} > \dots$ \square

4 The Observational Topology

In this section, we assume familiarity with the following notions: *topology* τ (identified with its family of *open* subsets) over a set S of points, *topological space* (S, τ) , *open* sets, *closed* sets, *interior* $\text{Int}(X)$ and *closure* \bar{X} of a set X , (open) *neighborhood* of a point s , *base* of a topology and *local base* (of neighborhoods) at a point. We use letters U, U' , etc., for open sets in τ , and letters C, C' , etc., for closed sets.

A space is said to be *second-countable* if its topology has a countable base. Given a topological space (S, τ) , the *specialization preorder* $\sqsubseteq \subseteq S \times S$ is defined in the following way: for any $s, t \in S$, we set

$$s \sqsubseteq t \text{ iff } \forall U \in \tau (s \in U \Rightarrow t \in U).$$

Separation Principles. In this paper we use four key topological separation notions. The first is the well-known separation axiom $T0$, which will be satisfied by all the topologies that arise in our setting. The second is the separation axiom TD . This condition (together with countability) will be shown to characterize *learnable* spaces. The next two notions are analogues of TD separation for *questions*. Instead of asking for open sets that separate points (states), these conditions require the existence of open sets that separate *answers* (to the same question). The concept of *locally closed questions* is a first analogue of TD , and it will be shown to characterize in some sense *solvable* problems. Finally, the notion of *linearly separated questions* is a stronger analogue of TD for questions, which characterizes a stronger type of solvability, what we will call *direct solvability by (AGM) conditioning*.

Definition 5. A topological space (S, τ) satisfies the separation axiom $T0$ if the specialization preorder is actually a partial order, i.e., it is antisymmetric: $s \sqsubseteq t \sqsubseteq s$ implies $s = t$. Equivalently, if $s \neq t$, then there exists some “separating” open U , such that either $s \in U, t \notin U$, or $s \notin U, t \in U$.

The space (S, τ) satisfies the separation axiom TD iff for every point $s \in S$, there is an open $U_x \ni x$ such that $y \not\sqsubseteq x$ for all $y \in O_x \setminus \{x\}$. Equivalently: for every $s \in S$ there is an open $U \in \tau$ such that $\{s\} = U \cap \{s\}$.

Essentially, $T0$ says that every two points $s \neq t$ can be separated (by an open U) one way or another (i.e., either $s \in U, t \notin U$, or $s \notin U, t \in U$), while TD essentially says that every point s can be separated (by an open neighborhood) from all the points $t \neq s$ that are inseparable from s .⁷

Definition 6. Given a topological space (S, τ) , a set $A \subseteq S$ is locally closed if it is the intersection $A = U \cap C$ of an open set U with a closed set C . Equivalently, if it is of the form $A = U \cap \bar{A}$ for some open U .

A set is ω -constructible if it is a countable union of locally closed sets.

A question \mathcal{Q} (partition of S) is locally closed if all its answers are locally closed. A problem \mathbb{P} is locally closed if its associated question is locally closed.

Essentially, locally closed questions are partitions with the property that every “answer” (i.e., partition cell) A can be separated (by an open neighborhood) from all the non- A -states that are inseparable from A .⁸

Definition 7. A question \mathcal{Q} is linearly separated if there exists some total order \triangleleft on the answers in \mathcal{Q} , such that $A \cap \overline{\bigcup_{B \triangleleft A} B} = \emptyset$. In other words, every answer A can be separated (by some open $U_A \supseteq A$) from the union of all the previous answers: $U_A \cap B = \emptyset$ for all $B \triangleleft A$.

Essentially, a linearly separated question is one whose answers can be totally ordered by a “plausibility” (or “simplicity”) order, in such a way that every answer A can be separated (by an open neighborhood $U_A \supseteq A$) from all answers that are “more plausible” (or “simpler”) than A .

Definition 8. The observational topology $\tau_{\mathbb{S}}$ associated with an epistemic space $\mathbb{S} = (S, \mathcal{O})$ is the topology generated by \mathcal{O} (i.e., the smallest collection of subsets of S , that includes $\mathcal{O} \cup \{\emptyset, S\}$ and is closed under finite intersections and arbitrary unions).

From now on, we will always implicitly consider our epistemic spaces \mathbb{S} to also be topological spaces $(S, \tau_{\mathbb{S}})$, endowed with their observational topology $\tau_{\mathbb{S}}$. Every topological property possessed by the associated topological space will thus be also attributed to the epistemic space.

⁷A point y is “inseparable” from x if every open neighborhood of y contains x , i.e., y and x are in the topological refinement order $y \sqsubseteq x$.

⁸Here, a state t is said to be “inseparable” from a set A if there is no open neighborhood $U \ni t$ that is disjoint from A .

Observation 2. Every epistemic space is $T0$ and second-countable. A (sound and complete) data stream for s is the same as a local neighborhood base at s .

Proposition 3. Every ω -constructible set can be written as a disjoint countable union of locally closed sets.

Proof. In order to prove this, we first recall some standard topological notions and results: A set is called *constructible* if it is a finite disjoint union of locally closed sets. Obviously, all locally closed sets are constructible. It is known that constructible sets form a Boolean algebra, i.e., the family of constructible sets is closed under complementation, finite unions, and finite intersections.

Suppose $A = \bigcup_{i \in \omega} A_i$, where all A_i are locally closed. Then we can rewrite A as a disjoint union $A = \bigcup_{i \in \omega} B_i$, where we have set $B_i = A_i \setminus (\bigcup_{k < i} A_k) = A_i \cap \bigcap_{k < i} A_k^c$, for every i . Since B_i 's are generated from locally closed sets using complementation and finite intersections, they must be constructible. Hence, each B_i can be written as disjoint finite unions of locally closed sets $B_i = \bigcup_{1 \leq j \leq i} B_{ij}$. Hence, we can write $A = \bigcup_{i \in \omega} \bigcup_{1 \leq j \leq i} B_{ij}$ as a disjoint countable union of locally closed sets. \square

Definition 9. A pseudo-stratification is a finite or ω -long sequence of locally closed sets $\langle A_i \mid i < \lambda \rangle$ (where $\lambda \in \omega \cup \{\omega\}$), which form a partition of S satisfying the following condition:

$$\text{if } j < i \text{ then either } A_i \cap \overline{A_j} = \emptyset \text{ or } A_i \subseteq \overline{A_j}.$$

Proposition 4. Every countable locally closed question can be refined to a pseudo-stratification.

Proof. Suppose $\Pi = \{A_i \mid i \in \omega\}$ is a countable locally closed question (partition of S). We first show the following:

Claim. There exists a family $\{(\Pi_i, <_i) \mid i \in \omega\}$, satisfying

- (1) each Π_i is a finite partition of A_i into locally closed sets;
- (2) each $<_i$ is a total order on Π_i ;
- (3) if $j < i$, $E \in \Pi_j$, $B \in \Pi_i$, then either $B \subseteq \overline{E}$ or $B \subseteq \overline{E}^c$;
- (4) if $B, E \in \Pi_i$, $E <_i B$, then $B \subseteq \overline{E}^c$.

Proof of Claim: We construct $(\Pi_n, <_n)$ by recursion: for $n = 0$, set $\Pi_0 := \{A_0\}$, with $<_0$ trivial. For the step $n + 1$: assume given $\{(\Pi_i, <_i) \mid i \leq n\}$ satisfying the above four conditions (for $i \leq n$). We set

$$\Pi_{n+1} := \{B_f \mid f : \bigcup_{i=1}^n \Pi_i \rightarrow \{0, 1\}\},$$

where for each function $f : \bigcup_{i=1}^n \Pi_i \rightarrow \{0, 1\}$ we have set

$$B_f := A_{n+1} \cap \bigcap \{\overline{E} \mid E \in f^{-1}(0)\} \cap \bigcap \{\overline{E}^c \mid E \in f^{-1}(1)\}.$$

It is obvious that the B_f 's are locally closed (given that A_{n+1} is locally closed) and that they form a partition of A_{n+1} . So condition (1) is satisfied.

It is also easy to check condition (2) for $i = n + 1$: let $j < n + 1$, $E \in \Pi_j$ and $B_f \in \Pi_{n+1}$. Then we have either $f(E) = 0$, in which case $B_f \subseteq \overline{E}$ (by construction of B_f), or else $f(E) = 1$, in which case $B_f \subseteq \overline{E}^c$.

To construct the order $<_{n+1}$, observe first that there is a natural total order $<^{(n)}$ on the disjoint union $\bigcup_{i=1}^n \Pi_i$, namely the one obtained by concatenating the orders $<_0, <_1, \dots, <_n$. More precisely, if, for

every $B \in \bigcup_{i=1}^n \Pi_i$, we set $i(B)$ to be the unique index $i \leq n$ such that $B \in \Pi_i$, then the order $<^{(n)}$ is given by setting: $B <^{(n)} E$ iff either $i(B) < i(E)$, or else $i(B) = i(E)$ and $B <_{i(B)} E$.

Now, the order $<_{n+1}$ on B_f 's is given by the lexicographic order induced by $<^{(n)}$ on the functions f (thought as "words" written with the letters 0 and 1). More precisely, we set:

$$B_f <_{n+1} B_g$$

iff there exists some set $E \in \bigcup_{i=1}^n \Pi_i$ such that

$$\left(\forall E' <^{(n)} E \ f(E') = g(E'), \text{ but } f(E) < g(E) \right),$$

where $<$ is the usual order $0 < 1$ on $\{0, 1\}$. Clearly, $<_{n+1}$ is a total order on Π_{n+1} , so condition (2) is satisfied.

Finally, we check condition (4) for $n+1$, let $B_f, B_g \in \Pi_{n+1}$ such that $B_f <_{n+1} B_g$. By definition of the order $<_{n+1}$, this means that there exists some $E \in \bigcup_{i=1}^n \Pi_i$ such that for all $E' <^{(n)} E$ we have $f(E') = g(E')$ but $f(E) < g(E)$, i.e., $f(E) = 0$ and $g(E) = 1$. By the construction of B_f 's, $f(E) = 0$ implies that $B_f \subseteq \overline{E}$, from which we get $\overline{B_f} \subseteq \overline{E}$, and thus $\overline{E}^c \subseteq \overline{B_f}^c$. Similarly, $g(E) = 1$ implies that $B_g \subseteq \overline{E}^c$. So we have $B_g \subseteq \overline{E}^c \subseteq \overline{B_f}^c$, and thus by transitivity of inclusion we get $B_g \subseteq \overline{B_f}^c$. This completes the proof of our Claim.

Given now the above Claim, we can prove our Lemma by taking as our refined partition

$$\Pi' := \bigcup_{i \in \omega} \Pi_i.$$

Clearly, Π' is a refinement of Π consisting of locally closed sets. We now define a well-order $<'$ on Π' as the concatenation of all the \leq_i 's.⁹ Obviously, $<'$ is a total order of type $\leq \omega$ on Π' , so we get finite or ω -long sequence that enumerates Π' . The above properties (3) and (4) ensure that this is a pseudo-stratification. \square

Lemma 1. *Given a pseudo-stratification $\langle A_i \mid i < \lambda \rangle$ (of length $\lambda \leq \omega$), there exists a λ -long sequence of open sets $\langle U_i \mid i < \lambda \rangle$, satisfying:*

- (1) $U_i \cap \overline{A_i} = A_i$;
- (2) if $j < i$ and $U_i \cap A_j \neq \emptyset$, then $A_i \subseteq \overline{A_j}$.

Proof. We know that each A_i is locally closed, so there exists some open set $U^{A_i} \in \tau$ such that $U^{A_i} \cap \overline{A_i} = A_i$. Now, for all $i \in \omega$ set

$$U_i := U^{A_i} \cap \bigcap \{ \overline{A_j}^c \mid j < i, A_i \subseteq \overline{A_j}^c \}.$$

Let us first check that the sequence $\langle U_i \mid i < \lambda \rangle$ satisfies condition (1):

$$\begin{aligned} U_i \cap \overline{A_i} &= (U^{A_i} \cap \bigcap \{ \overline{A_j}^c \mid j < i, A_i \subseteq \overline{A_j}^c \}) \cap \overline{A_i} \\ &= (U_i \cap \overline{A_i}) \cap \bigcap \{ \overline{A_j}^c \mid j < i, A_i \subseteq \overline{A_j}^c \} \\ &= A_i \cap \bigcap \{ \overline{A_j}^c \mid j < i, A_i \subseteq \overline{A_j}^c \} = A_i \end{aligned}$$

⁹Once again, one can specify this more precisely by first defining $i: \Pi' \rightarrow \omega$ by choosing $i(B)$ to be the unique index i such that $B \in \Pi_i$, and finally defining: $B <' E$ iff either $i(B) < i(E)$, or else $i(B) = i(E)$ and $B <_{i(B)} E$.

Second, let us check condition (2): Suppose that we have $j < i$ and $U_i \cap A_j \neq \emptyset$, but $A_i \not\subseteq \overline{A_j}$. Since $(A_i)_{i < \lambda}$ is a pseudo-stratified sequence, from $j < i$ and $A_i \not\subseteq \overline{A_j}$ we can derive $A_i \subseteq \overline{A_j}^c$. By the construction of U_i , this implies that $U_i \subseteq \overline{A_j}^c$, and hence that $U_i \cap A_j \subseteq \overline{A_j}^c \cap A_j \subseteq \overline{A_j}^c \cap \overline{A_j} = \emptyset$, which contradicts the assumption that $U_i \cap A_j \neq \emptyset$. \square

Lemma 2. *Every pseudo-stratification is linearly separated.*

Proof. Let $\Pi = \{A_i \mid i < \lambda\}$ be a pseudo-stratification (with $\lambda \leq \omega$), and let $\langle U_i \mid i < \lambda \rangle$ be a sequence satisfying the conditions of Lemma 1. It is clear that, in order to prove our intended result, it is enough to construct a total order \trianglelefteq on the set $\{i \in \omega \mid i < \lambda\} = \lambda \subseteq \omega$, such that

$$U_i \cap A_j \neq \emptyset \Rightarrow i \trianglelefteq j.$$

For this, we first define a reflexive relation R on λ , by setting

$$iRj \iff U_i \cap A_j \neq \emptyset.$$

Claim: There are no non-trivial cycles

$$i_1 R \cdots i_n R i_1 \quad (\text{with distinct } i_k\text{'s}).$$

Proof of Claim: Let $i_1 R \cdots i_n R i_1$ be a non-trivial cycle of *minimal length* $n \geq 2$. There are two cases:

Case 1: $n = 2$, i.e., $i_1 R i_2 R i_1$ with $i_2 \neq i_1$. We must have either $i_1 < i_2$ or $i_2 < i_1$. Without loss of generality, we can assume $i_1 < i_2$ (otherwise, just swap i_1 and i_2 , and use the cycle $i_2 R i_1 R i_2$). From $i_2 R i_1$, we get $U_{i_2} \cap A_{i_1} \neq \emptyset$. This together with $i_1 < i_2$, gives us $A_{i_2} \subseteq \overline{A_{i_1}}$ (by condition (2) from Lemma 2), and hence $U_{i_1} \cap A_{i_2} \subseteq U_{i_1} \cap \overline{A_{i_1}} = A_{i_1}$. From this, we get that $U_{i_1} \cap A_{i_2} = (U_{i_1} \cap A_{i_2}) \cap A_{i_2} \subseteq A_{i_1} \cap A_{i_2} = \emptyset$ (since $i_1 \neq i_2$, so A_{i_1} and A_{i_2} are different answers, hence disjoint), so we conclude that $U_{i_1} \cap A_{i_2} = \emptyset$. But on the other hand, from $i_1 R i_2$ we get $U_{i_1} \cap A_{i_2} \neq \emptyset$. Contradiction.

Case 2: $n > 2$. Since all the i_k 's are distinct, there must exist a (unique) smallest index in the cycle. Without loss of generality (since otherwise we can rearrange the indices, permuting the cycle), we can assume that i_3 is the smallest index. (Note that, since $n > 2$, there must be at least three distinct successive indices i_1, i_2, i_3 .) So $i_3 < i_1$ and $i_3 < i_2$. From $i_2 R i_3$ we get $U_{i_2} \cap A_{i_3} \neq \emptyset$. Since $i_3 < i_2$, it follows that $A_{i_2} \subseteq \overline{A_{i_3}}$ (by Lemma 2). But on the other hand, $i_1 R i_2$ gives us $U_{i_1} \cap A_{i_2} \neq \emptyset$. We hence obtain $U_{i_1} \cap \overline{A_{i_3}} \neq \emptyset$. This, together with $i_3 < i_1$, gives us $A_{i_1} \subseteq \overline{A_{i_3}}$ (again by Lemma 2). From this, we derive $A_{i_1} \subseteq U_{i_1} \cap \overline{A_{i_3}}$ (since $A_i \subseteq U_i$ for all i). Let now $s \in A_{i_1}$ be any state satisfying the answer $A_{i_1} \subseteq U_{i_1} \cap \overline{A_{i_3}}$. So we have $s \in U_{i_1}$ and $s \in \overline{A_{i_3}}$, which together imply that $U_{i_1} \cap A_{i_3} \neq \emptyset$ (since $s \in \overline{A_{i_3}}$ implies that every open neighborhood of s intersects A_{i_3}). Hence, we have $i_1 R i_3$, which means we can shorten the cycle by eliminating i_2 , we obtain contradiction.

Given the above Claim, it follows that the transitive closure R^* is a partial order on λ (which obviously includes R). By the Order Extension Principle, we can extend R^* to a total order \trianglelefteq on λ , which still includes R . \square

5 Topological Characterization of Solvability

Definition 10. Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space, L be a standard agent, $A \subseteq S$, and $s \in A$. An A -locking sequence for s (with respect to L) is a data sequence $\sigma = (O_1, \dots, O_k)$, such that:

- (1) σ is sound for s , i.e., $s \in \bigcap_{1 \leq i \leq k} O_i$;
(2) if δ is any data sequence sound for s , then $L(\mathbb{S}, \sigma * \delta) \subseteq A$.

For a given data sequence σ , we denote by L_A^σ the set of all states in A having σ as an A -locking sequence, i.e.,

$$L_A^\sigma := \{s \in A \mid \sigma \text{ is an } A\text{-locking sequence for } s \text{ wrt } L\}.$$

Lemma 3. *If A is verifiable in the limit by a standard agent L , then $\bigcup_{\sigma \in \mathcal{O}^*} L_A^\sigma = A$.*

Proof. Suppose not. Let A be verifiable in the limit, but such that $A \neq \bigcup_{\sigma \in \mathcal{O}^*} L_A^\sigma$. Since all $L_A^\sigma \subseteq A$, this means that $A \not\subseteq \bigcup_{\sigma \in \mathcal{O}^*} L_A^\sigma$, i.e., there exists some state $s \in A$ for which there is no A -locking sequence. This means that every data sequence σ that is sound for s can be extended to a sequence δ that is also sound for s and has $L(\delta) \not\subseteq A$.

Let now \vec{O} be a (sound and complete) data stream for s . We construct a new infinite data stream \vec{V} , by defining increasingly longer initial segments δ_k of \vec{O} , in countably many stages: we first set $V_0 = O_0$, thus obtaining an initial segment $\delta_0 = (O_0) = (V_0)$; at the $k+1$ -th stage, given some initial segment $\delta_k = (V_0, V_1, \dots, V_{n_k})$ (of some length n_k), we built our next initial segment by taking any extension δ_{k+1} of the sequence $\sigma_k = (V_0, \dots, V_{n_k}, O_{n+1})$ that is sound for s and has $L(\sigma_k) \not\subseteq A$. The resulting infinite stream \vec{V} is a (sound and complete) stream for s (the completeness of \vec{V} with respect to s follows the fact that this stream includes all the elements of \vec{O}), but which contains arbitrarily long initial segments σ_k with $L(\sigma_k) \not\subseteq A$. Since $s \in A$, this contradicts the assumption that A is verifiable in the limit. \square

Lemma 4. *If $A \subseteq S$ is verifiable in the limit by a standard agent L , then for every data sequence $\sigma = (O_1, \dots, O_k)$, the set L_A^σ is locally closed.*

Proof. Let $O := \bigcap_{i=1}^k O_i$ be the intersection of all the observations in σ . We will show that

$$O \cap \overline{L_A^\sigma} = L_A^\sigma,$$

from which the desired conclusion follows.

(\supseteq) If $s \in L_A^\sigma$, then σ is an A -locking sequence for s , hence σ is sound for s , and thus $s \in \bigcap_{i=1}^k O_i = O$.

(\subseteq) Suppose that $s \in O \cap \overline{L_A^\sigma}$. We prove two claims:

Claim 1: For every data sequence δ that is sound for s and extends σ , we have $L_{\mathbb{S}}(\delta) \subseteq A$.

Proof of Claim 1: Let $\delta = (\delta_1, \dots, \delta_n)$ be a data sequence that is sound for s (i.e., $s \in \delta_i$ for all $i = 1, \dots, n$) and extends σ , i.e., $n \geq k$ and $U_i = O_i$ for all $i \leq k$). Hence, $\bigcap_{i=1}^n \delta_i$ is an open neighborhood of s , and $s \in \overline{L_A^\sigma}$, so there must exist some $t \in \bigcap_{i=1}^n \delta_i$ such that $t \in L_A^\sigma$. Hence, $t \in A$ and σ is an A -locking sequence for t . But δ extends σ and is sound for t , so (by the definition of σ being an A -locking sequence for t), we have that $L(\delta) \subseteq A$, which concludes the proof of Claim 1.

Claim 2: We have $s \in A$.

Proof of Claim 2: Let \vec{V} be a stream for s that extends σ (such a stream must exist, since σ is sound for s : just take any stream for s and prefix it with σ). Then, for every $n \geq k$, the sequence $\delta_n = (V_1, \dots, V_n)$ is sound for s and extends σ . Hence, by the above Claim, we must have that $L_{\mathbb{S}}(V_1, \dots, V_n) \subseteq A$ for all $n \geq k$. But we assumed that A is verifiable in the limit, so we must have $s \in A$, which concludes the proof of Claim 2.

From Claims 1 and 2 together, we conclude that σ is an A -locking sequence for $s \in A$, hence $s \in L_A^\sigma$. \square

Theorem 1. *Given an epistemic space (S, \mathcal{O}) , a set $A \subseteq S$ is verifiable in the limit iff it is ω -constructible.*

Proof. (\Leftarrow) Assume $A = \bigcup_n (U_n \cap C_n)$ is a countable disjoint union of (mutually disjoint) locally closed sets $U_n \cap C_n$ (with U_n open and C_n closed). We define a standard agent L for A on finite data sequences $\delta = (O_1, \dots, O_k)$, by setting $L(\mathbb{S}, \delta) = A^c$, if we have $\bigcap_j O_j \not\subseteq U_n$ for all $n \in \omega$; $L(\mathbb{S}, \delta) = A^c$ (where A^c is the complement of A), if $\bigcap_j O_j \subseteq C_n^c$ holds for the first index $n \in \omega$ such that $\bigcap_j O_j \subseteq U_n$; and $L(\mathbb{S}, \delta) = A$ otherwise. Then it is easy to see that L verifies A in the limit.

(\Rightarrow) Suppose that A is verifiable in the limit. By Proposition 1, it is then verifiable by a standard agent L . By Lemma 1, A is the union of all sets L_A^σ for all finite data sequences σ . But there are only countably many such sequences, so this is a countable union. Moreover, by Lemma 2, each L_A^σ is locally closed. Hence A is a countable union of locally closed sets, i.e., an ω -constructible set. \square

Corollary 1. *A is decidable in the limit iff both A and A^c are ω -constructible.*

Proof. Follows trivially from the above results. \square

Theorem 2. *Let $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ be an inductive problem on an epistemic space \mathbb{S} . The following are equivalent:*

- (1) \mathbb{P} is solvable (in the limit);
- (2) the associated question \mathcal{Q} is an (at most) countable family of ω -constructible answers;
- (3) \mathcal{Q} has an (at most) countable locally closed refinement.

Proof. (1) \Rightarrow (2) : Let \mathbb{P} be a solvable problem. By Proposition 1, there exists some standard agent that solves it. Let L be such a standard agent that solves \mathbb{P} .

Claim: Every answer $A \in \mathcal{Q}$ is verifiable in the limit.

Proof of Claim: Let $A \in \mathcal{Q}$ be an answer. We construct a standard agent L^A that verifies it, by setting $L_{\mathbb{S}}^A(\sigma) := A$ iff $L_{\mathbb{S}}(\sigma) \subseteq A$, and $L_{\mathbb{S}}^A(\sigma) := A^c$ otherwise. It is easy to see that L^A verifies A .

Using the Claim and Lemma 3, we obtain that, for each answer $A \in \mathcal{Q}$, there exists some data sequence $\sigma \in \mathcal{O}^*$ such that $L_{\mathbb{S}}(\sigma) \subseteq A$. But \mathcal{O}^* is countable, so there can be only countably many answers in \mathcal{Q} .

By the claim above, Lemma 3 and Lemma 4, we obtain that every answer $A \in \mathcal{Q}$ is a countable union of locally closed sets, hence it is ω -constructible.

(2) \Rightarrow (3) : By (2), \mathcal{Q} is (at most) countable, say $\mathcal{Q} = \{A_i \mid i \in \omega\}$, and also each answer $A_i \in \mathcal{Q}$ is ω -constructible, hence it can be written as a countable disjoint union of locally closed sets $A_i = \bigcup_{k \in \omega} A_i^k$ (where all A_i^k 's locally closed and mutually disjoint). Then the question $\{A_i^k \mid i \in \omega, k \in \omega\}$ is a refinement of \mathcal{Q} , which is countable and locally closed.

(3) \Rightarrow (1) : Let $\mathcal{Q}' = \{B_i \mid i \in \omega\}$ be a countable closed refinement of \mathcal{Q} . By Corollary 1, every answer $B \in \mathcal{Q}'$ is decidable, and so by Proposition 1, we can choose for each $B_i \in \mathcal{Q}'$ some standard agent L_i that decides B_i . We define now a new standard agent L , by:

$$L_{\mathbb{S}}(\sigma) := \bigcup \{B_i \mid i \in \omega \text{ such that } L_i(\sigma) \subseteq B_i\}.$$

It is easy to see that this agent L solves \mathcal{Q}' , and since \mathcal{Q}' is a refinement of \mathcal{Q} , L also solves \mathcal{Q} . \square

Corollary 2. *An epistemic space $\mathbb{S} = (S, \mathcal{O})$ is learnable in the limit iff it is countable and satisfies the TD separation axiom.*

Proof. Apply Theorem 2 to the learning question $\{\{s\} \mid s \in S\}$, noticing that the fact that all its answers are ω -constructible is equivalent to all singletons being locally closed, which is just another formulation of the TD axiom. \square

6 Universality of Conditioning

Our aim in this section is to show that *AGM* conditioning is “universal”: every solvable problem can be solved by some *AGM* agent. First, we introduce an auxiliary notion, that of a problem being *directly solvable by AGM conditioning*.

Given a question \mathcal{Q} on an epistemic space (S, \mathcal{O}) , any total order $\trianglelefteq \subseteq \mathcal{Q} \times \mathcal{Q}$ on (the answers of) the question \mathcal{Q} induces in a canonical way a total preorder $\leq \subseteq S \times S$, obtained by:

$$s \leq t \text{ iff } A_s \trianglelefteq A_t$$

(where A_s is the unique answer $A_s \in \mathcal{Q}$ such that $s \in A_s$).

Definition 11. A problem $\mathbb{P} = (S, \mathcal{Q})$ is directly solvable by conditioning if it is solvable by *AGM conditioning with respect to (a prior \leq that is canonically induced, as explained above, by) a total order $\trianglelefteq \subseteq \mathcal{Q} \times \mathcal{Q}$ on (the answers of) the question \mathcal{Q} .*

Direct solvability by conditioning essentially means that the problem can be solved by a conditioning agent *who does not attempt to refine the original question*: she forms beliefs only about the answers to the given question, and is thus indifferent between states satisfying the same answer. Direct solvability by conditioning is thus a very stringent condition, and unsurprisingly this form of conditioning is *not* universal.

Proposition 5. (K. Genin, personal communication) *Not every solvable problem is directly solvable by conditioning.*

Proof. Let \mathbb{P} be the problem in Example 3, depicted on the left-hand side of Figure 1. It is easy to see that this problem cannot be directly solvable by conditioning. Indeed, if $\{t, u\} \triangleleft \{s, v\}$ then v is not learnable by \triangleleft -conditioning; if $\{s, v\} < \{t, u\}$ then t is not learnable by \triangleleft -conditioning; while if $\{t, u\}$ and $\{s, v\}$ are equally plausible, then neither t nor v are learnable.

But \mathbb{P} can be refined to a directly solvable problem, namely the “learning question” \mathbb{P}' (depicted on the right-hand side of Figure 1), which *can* be directly solvable (e.g. if we set $\{t\} \triangleleft \{s\} \triangleleft \{v\} \triangleleft \{u\}$). As a consequence, \mathbb{P} can itself be solved by (non-direct) conditioning (with respect to the order $t < s < v < u$). \square

This counterexample suggests a way to prove our intended universality result: it is enough to show that every solvable problem has a refinement that is directly solvable by conditioning. To do this, we first need a structural characterization of direct solvability.

Lemma 5. (Topological Characterization of Direct Solvability by Conditioning) *A problem $\mathbb{P} = (S, \mathcal{Q})$ is directly solvable by conditioning iff \mathcal{Q} is linearly separated.*

Proof. Left-to-right implication: Suppose that \mathbb{P} is directly solvable by conditioning with respect to (a prior \leq that is canonically induced by) a total order $\trianglelefteq \subseteq \mathcal{Q} \times \mathcal{Q}$. Then, for every $s \in S$ choose some sound and complete data stream $\vec{O}^s = (O_n^s)_{n \in \omega}$ for s (with $O_n^s \in \mathcal{O} \subseteq \tau_{\mathbb{S}}$). Direct solvability by conditioning implies then that there exists some N_s such that $\text{Min}_{\leq} (O_1^s, \dots, O_{N_s}^s) \subseteq A_s$. Set $U_s := \bigcap_{i=1}^{N_s} O_i^s \in \tau_{\mathbb{S}}$, so that we have $s \in U_s$ and $\text{Min}_{\leq} U_s \subseteq A_s$. Then set $U_A := \bigcup_{s \in A} U_s \in \tau_{\mathbb{S}}$ for every answer $A \in \mathcal{Q}$. We claim that U_A “separates” A from the union of all the answers $B \triangleleft A$ (as linear separation demands): indeed, by the construction of U_A , it is obvious that (1) $A \subseteq U_A$, and also that $\text{Min}_{\leq} U_A \subseteq A$. By unfolding the last clause in terms of \trianglelefteq , we obtain that: $A \trianglelefteq B$ holds for all $B \in \mathcal{Q}$ such that $U_A \cap B \neq \emptyset$. Since \trianglelefteq is a total order on

\mathcal{Q} , this is equivalent to: (2) $U_A \cap B = \emptyset$ for all $B \triangleleft A$. By (1) and (2) together, we obtain that \mathcal{Q} is linearly separated.

Right-to-left implication: Suppose \mathcal{Q} is linearly separated. Let \trianglelefteq be a total order on \mathcal{Q} that linearly separates it. This means that, for every answer $A \in \mathcal{Q}$, there exists some open set $U_A \in \tau_{\mathcal{S}}$ such that $A \subseteq U_A$ and $U_A \cap B = \emptyset$ for all $B \triangleleft A$. For each $s \in S$, we set $U_s := U_{A_s}$ (where A_s is the unique answer $A_s \in \mathcal{Q}$ with $s \in A_s$).

Let \leq be the total preorder on S canonically induced by the order $\trianglelefteq \subseteq \mathcal{Q} \times \mathcal{Q}$ (by $s \leq t$ iff $A_s \triangleleft A_t$). We show now that \mathbb{P} is directly solvable by conditioning with respect to \leq . For this, let $s \in S$ be any state, and $\vec{O} = (O_n)_{n \in \omega}$ be a sound and complete stream for s . Completeness of the stream implies that there must exist some $N \in \omega$ such that $\bigcap_{i=1}^N O_i \subseteq U_s$.

To conclude our proof, it is enough to show the following

Claim: For every $n \geq N$, we have

$$s \in \text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right) \subseteq A_s.$$

First, let us see why this Claim is enough to give us direct solvability by conditioning. The fact that $s \in \text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right)$ implies that $\text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right) \neq \emptyset$, for all $n \geq N$. A previous observation tells us that, when applied to such data streams, the AGM agent \mathcal{L}^{\leq} produces a ‘‘principal filter’’, given by

$$\mathcal{L}^{\leq}(O_1, \dots, O_n) = \{P \subseteq S \mid \text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right) \subseteq P\}.$$

By the Claim above we have $\text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right) \subseteq A_s$, and hence we obtain $A_s \in \mathcal{L}^{\leq}(O_1, \dots, O_n)$, for all $n \geq N$.

Proof of Claim: Let $n \geq N$. To prove the Claim, it is enough to show the following two implications (for all states t):

- (1) $t \in \bigcap_{i=1}^n O_i \Rightarrow s \leq t$;
- (2) $t \in \text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right) \Rightarrow A_t = A_s$.

To show (1), let $t \in \bigcap_{i=1}^n O_i$. Then $t \in U_s$ (since $\bigcap_{i=1}^n O_i \subseteq \bigcap_{i=1}^N O_i \subseteq U_s$), so $U_s \cap A_t \neq \emptyset$. Hence (by linear separation) we must have $A_s \triangleleft A_t$, i.e., $s \leq t$.

To show (2), let $t \in \text{Min}_{\leq} \left(\bigcap_{i=1}^n O_i \right)$. This implies that $t \leq s$ (since $s \in \bigcap_{i=1}^n O_i$). But by (1), we also have $s \leq t$, and hence $s \leq t \leq s$. This means that $A_s \triangleleft A_t \triangleleft A_s$. But \trianglelefteq is a total order on \mathcal{Q} , so it follows that $A_t = A_s$. \square

Theorem 3. *AGM conditioning is a universal problem-solving method, i.e., every solvable problem is solvable by some AGM agent.*

Proof. Let \mathbb{P} be a solvable problem. From Theorem 2, Proposition 3 and Lemma 2, it follows that \mathbb{P} has a linearly separated refinement \mathbb{P}' . By Lemma 5, that refinement is (directly) solvable by an AGM agent L^{\leq} . It is obvious (from the definition of solvability) that any doxastic agent which solves the more refined problem \mathbb{P}' solves also the original problem \mathbb{P} . \square

Corollary 3. *AGM conditioning is a universal learning method, i.e., every learnable space is learnable by some AGM agent.*

Proof. Apply the previous result to the finest question $\mathcal{Q} := \{\{s\} \mid s \in S\}$. \square

In contrast, recall that the counterexample in Proposition 2 showed that *standard* AGM agents have a very limited problem-solving power. Standard conditioning is not a universal learning method (while general AGM conditioning is universal). This means that *allowing prior plausibility orders that are non-wellfounded is essential for achieving universality of conditioning*. Beliefs generated in this way may occasionally fail to be globally consistent. (Indeed, note that in the counterexample from Proposition 2, the beliefs of the non-standard AGM agent who learns the space are initially globally inconsistent. In conclusion, *occasional global inconsistencies are the unavoidable price for the universality of AGM conditioning*.)

7 Conclusions and Connections to Other Work

The general topological setting for problem-solving assumed here is a variation of the one championed by Kelly in various talks [23] and in unpublished work [24, 25], though until recently we did not realize this close similarity. Our topological characterizations of verifiable, falsifiable and decidable properties are generalizations of results by Kelly [20], who proved characterizations for the special case of Baire spaces.¹⁰ Our result on learning-universality (Corollary 3) is also a generalization of analogue results by Kelly [21, 26], and Kelly, Schulte and Hendricks [19]. But our generalization to arbitrary spaces is highly non-trivial, requiring the use of the *TD* characterization. In contrast, the Baire space satisfies the much stronger separation axiom *T1*, which trivializes the specialization order, and so the proof of learning-universality is much easier in this special case: *any* total ω -like ordering of the space can be used for conditioning. Nevertheless, in a sense, this result is just a topological re-packaging of one of our own previous results [13, 3, 4].

While writing this paper, we learned that our *TD* characterization of learnability (Corollary 2) was independently re-proven by Konstantin Genin ([11], unpublished manuscript), soon after we announced its proof. This characterization is actually a topological translation of a classical characterization of identifiability in the limit [2], and in fact it also follows from a result by de Brecht and Yamamoto [9], who prove it for so-called “concept spaces”.

Our key new results are far-reaching and highly non-trivial: the topological characterization of solvability (Theorem 2), and the universality of AGM condition for problem-solving (Theorem 3). They required the introduction of new topological concepts (e.g., pseudo-stratifications and linearly separated partitions), and some non-trivial proofs of new topological results.

Philosophically, the importance of these results is that, on the one hand they fully vindicate the general topological program in Inductive Epistemology started by Kelly and others [20, 30], and on the other hand they reassert the power and applicability of the AGM Belief Revision Theory against its critics. To this conclusion, we need to add an important proviso: our results show that, in order to achieve problem-solving universality, AGM agents need to (a) be “creative”, by *going beyond the original problem* (i.e., finding a more refined problem that can be solved directly, and forming prior beliefs about the answer to this more refined question), and (b) admit *non-standard priors*, which occasionally will lead to *beliefs that are globally inconsistent* (although still locally consistent). Such occasional global inconsistencies can give rise to a type of “infinite Lottery Paradox”. But this is the price that AGM agents *have to pay* in

¹⁰In unpublished work [25] the authors claim a characterization of solvability in a general setting. Their characterization is slightly “looser” than ours, and can be easily obtained from ours. Our tighter characterization is the one needed for proving universality.

order to be able to solve every solvable question. Whether or not this is a price that is worth paying is a different, more vague and more “ideological” question, although a very interesting one. But this question lies beyond the scope of this paper.

8 Acknowledgments

We thank Johan van Benthem, Nick Bezhanishvili, Konstantin Genin, Thomas Icard and Kevin Kelly for their useful feedback on issues related to this paper. Johan helped us place belief-based learning within the larger context of long-term doxastic protocols [7], and beyond this he gave us his continuous support and encouragement for our work on this line of inquiry. Nick pointed to us the connections between our work and the notions of TD-space and locally closed set. Konstantin pointed to us the connections to the notion of stratification and gave the counterexample proving Proposition 5. His critical feedback on our early drafts was really essential for clarifying our thoughts and cleaning up our proofs, and so it’s fair to say that this paper in its current form owes a lot to Konstantin Genin. Thomas Icard’s comments on a previous draft and our friendly interactions with him on related topics during our Stanford visits are very much appreciated. Finally, Kevin Kelly’s work forms of course the basis and the inspiration for ours. Our frequent discussions with him in recent years influenced the development of our own perspective on the topic. He also gave us excellent reference tips concerning the history of the connections between topology and formal epistemology, as well as concerning his more recent work on related issues.

Nina Gierasimczuk’s work on this paper was funded by an Innovational Research Incentives Scheme Veni grant 275-20-043, Netherlands Organisation for Scientific Research (NWO). Sonja Smets was funded in part by an Innovational Research Incentives Scheme Vidi grant from NWO, and by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no. 283963.

References

- [1] Carlos Alchourrón, Peter Gärdenfors & David Makinson (1985): *On the Logic of Theory Change: Partial Meet Contraction and Revision Functions*. *Journal of Symbolic Logic* 50, pp. 510–530, doi:10.2307/2274239.
- [2] Dana Angluin (1980): *Inductive inference of formal languages from positive data*. *Information and Control* 45(2), pp. 117–135, doi:10.1016/S0019-9958(80)90285-5.
- [3] Alexandru Baltag, Nina Gierasimczuk & Sonja Smets (2011): *Belief revision as a truth-tracking process*. In K. Apt, editor: *Proceedings of TARK’11*, ACM, pp. 187–190, doi:10.1145/2000378.2000400.
- [4] Alexandru Baltag, Nina Gierasimczuk & Sonja Smets (2015): *Truth Tracking by Belief Revision*. Technical Report, ILLC Report PP-2014-20. To appear in *Studia Logica*.
- [5] Alexandru Baltag & Sonja Smets (2008): *A Qualitative Theory of Dynamic Interactive Belief Revision*. In G. Bonanno, W. van der Hoek & M. Wooldridge, editors: *Proc. of LOFT’7, Texts in Logic and Games* 3, Amsterdam University Press, pp. 9–58.
- [6] Johan van Benthem (2011): *Logical Dynamics of Information and Interaction*. Cambridge University Press, doi:10.1017/CBO9780511974533.
- [7] Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi & Eric Pacuit (2009): *Merging Frameworks for Interaction*. *Journal of Philosophical Logic* 38(5), pp. 491–526, doi:10.1007/s10992-008-9099-x.
- [8] Oliver Board (2004): *Dynamic interactive epistemology*. *Games and Economic Behavior* 49(1), pp. 49–80, doi:10.1016/j.geb.2003.10.006.

- [9] Matthew de Brecht & Akihiro Yamamoto (2010): *Topological properties of concept spaces*. *Information and Computation* 208(4), pp. 327–340, doi:10.1016/j.ic.2009.08.001.
- [10] Cédric Dégremont & Nina Gierasimczuk (2011): *Finite identification from the viewpoint of epistemic update*. *Information and Computation* 209(3), pp. 383–396, doi:10.1016/j.ic.2010.08.002.
- [11] Konstantin Genin (2015): *Linearizing a Countable TD Space*. Unpublished manuscript.
- [12] Nina Gierasimczuk (2009): *Identification through Inductive Verification*. In: *Proceedings of TBILLC'07, LNCS 5422*, Springer, pp. 193–205, doi:10.1007/978-3-642-00665-4_16.
- [13] Nina Gierasimczuk (2010): *Knowing One's Limits. Logical Analysis of Inductive Inference*. Ph.D. thesis, Universiteit van Amsterdam, The Netherlands.
- [14] Nina Gierasimczuk, Vincent F. Hendricks & Dick de Jongh (2014): *Logic and Learning*. In Alexandru Baltag & Sonja Smets, editors: *Johan van Benthem on Logic and Information Dynamics, Outstanding Contributions to Logic 5*, Springer International Publishing, pp. 267–288, doi:10.1007/978-3-319-06025-5_10.
- [15] Nina Gierasimczuk & Dick de Jongh (2013): *On the Complexity of Conclusive Update*. *The Computer Journal* 56(3), pp. 365–377, doi:10.1093/comjnl/bxs059.
- [16] E. Mark Gold (1965): *Limiting Recursion*. *JSL* 30(1), pp. 28–48, doi:10.2307/2270580.
- [17] E. Mark Gold (1967): *Language Identification in the Limit*. *Information and Control* 10, pp. 447–474, doi:10.1016/S0019-9958(67)91165-5.
- [18] Adam Grove (1988): *Two modellings for theory change*. *Journal of Philosophical Logic* 17, pp. 157–170, doi:10.1007/BF00247909.
- [19] Kevin T. Kelly, Oliver Schulte & Vincent Hendricks (1997): *Reliable Belief Revision*. In Maria Luisa Dalla Chiara, Kees Doets, Daniele Mundici & Johan van Benthem, editors: *Logic and Scientific Methods, Synthese Library 259*, Springer Netherlands, pp. 383–398, doi:10.1007/978-94-017-0487-8_20.
- [20] Kevin T. Kelly (1996): *The Logic of Reliable Inquiry*. Oxford University Press, Oxford.
- [21] Kevin T. Kelly (1998): *The Learning Power of Belief Revision*. In: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '98*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 111–124. Available at <http://dl.acm.org/citation.cfm?id=645876.671884>.
- [22] Kevin T. Kelly (2008): *Ockham's Razor, Truth, and Information*. In Pieter Adriaans & Johan van Benthem, editors: *Handbook of the Philosophy of Information*, Elsevier, pp. 321–359, doi:10.1016/B978-0-444-51726-5.50014-5.
- [23] Kevin T. Kelly (2011): *An erotetic theory of empirical simplicity and its connection with truth*. Unpublished manuscript.
- [24] Kevin T. Kelly (2015): *Notes on a General Topological Paradigm*. Unpublished manuscript.
- [25] Kevin T. Kelly & Hanti Lin (2011): *A simple theory of theoretical simplicity*. Unpublished manuscript.
- [26] Kevin T. Kelly (1999): *Iterated Belief Revision, Reliability, and Inductive Amnesia*. *Erkenntnis* 50(1), pp. 7–53, doi:10.1023/A:1005444112348.
- [27] David Lewis (1969): *Convention*. Cambridge: Harvard University Press.
- [28] Eric Martin & Daniel Osherson (1997): *Scientific Discovery Based on Belief Revision*. *Journal of Symbolic Logic* 62(4), pp. 1352–1370, doi:10.2307/2275647.
- [29] Hilary Putnam (1965): *Trial and Error Predicates and the Solution to a Problem of Mostowski*. *Journal of Symbolic Logic* 30(1), pp. 49–57, doi:10.2307/2270581.
- [30] Olivier Schulte & Cory Juhl (1996): *Topology as Epistemology*. *Monist* 79(1), pp. 141–147, doi:10.5840/monist19967916.
- [31] Ray J. Solomonoff (1964): *A formal theory of inductive inference*. *Information and Control Part I*, 7(1):1–22, 1964. *Part II*, 7(2):224–254, 1964, doi:10.1016/S0019-9958(64)90223-2.
- [32] Steven Vickers (1996): *Topology Via Logic*. Cambridge University Press.