



Inductive Inference and Epistemic Modal Logic

Nina Gierasimczuk   

Department of Applied Mathematics and Computer Science,
Technical University of Denmark, Lyngby, Denmark

Abstract

This paper is concerned with a link between inductive inference and dynamic epistemic logic. The bridge was first introduced in [20, 21, 22]. We present a synthetic view on subsequent contributions: inductive truth-tracking properties of belief revision policies seen as belief upgrade methods; topological interpretation and characterisation of inductive inference; discussion of the adequacy of the topological semantics of modal logic for characterising inductive inference. We briefly present the topological Dynamic Logic for Learning Theory. Finally, we discuss several surprising results obtained in computational inductive inference that challenge the usual understanding of certainty, and of rational inquiry as consistent and conservative learning.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Models of learning

Keywords and phrases modal logic, dynamic epistemic logic, inductive inference, topological semantics, computational learning theory, finite identifiability, identifiability in the limit

Digital Object Identifier 10.4230/LIPIcs.CSL.2023.2

Category Invited Talk

1 Introduction

Modal logic is applied in a variety of domains. One particular direction, *epistemic logic*, interprets modalities as coding information states of an (artificial) agent. The dynamic extensions of epistemic logic, referred to with the umbrella term of Dynamic Epistemic Logic (DEL, see [7, 19, 8]), allow analysing the events of incorporating new information into the existing knowledge, for instance through communication. Such adoption of new information can be of course called “learning”, but inductive inference and computational learning theory have stronger demands: the process of learning is considered successful if over time it converges to a state of familiarity with the true underlying structure of the world. For instance we would say “Alice learned Polish”, if Alice managed to acquire sufficient familiarity with the grammar of the Polish language. Theories of inductive inference are concerned with this long-term horizon of learning, rather than with learning understood as ways of incorporating single pieces of information. In this survey paper we will discuss a connection between inductive inference and modal logic, by way of DEL. We hope to show that modal logic can contribute to our understanding of learning processes in general.

The link between dynamic epistemic logic and computational learning theory was first introduced in [20, 21, 22], where it was shown that exact learning in finite time (also known as finite identification, see [27, 29]) can be modelled in DEL, and that the elimination process of learning by erasing [26] can be seen as iterated upgrade of Dynamic Doxastic Logic. In the present paper we will outline the continuations developed in a number of contributions. In Section 3 we will discuss inductive truth-tracking properties of belief revision policies seen as belief upgrade methods [4, 6]. Section 4 will display a topological interpretation of inductive inference that allows characterising favourable conditions for learning in the limit in terms of general topology [5]. As modal logic can be given topological semantics, in Section 4.1 we will discuss a variety of modal logic characterisations of relevant kinds of topological spaces, to finally present Dynamic Logic for Learning Theory, which extends



© Nina Gierasimczuk;
licensed under Creative Commons License CC-BY 4.0
31st EACSL Annual Conference on Computer Science Logic (CSL 2023).

Editors: Bartek Klin and Elaine Pimentel; Article No. 2; pp. 2:1–2:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Subset Space Logics [16] with dynamic observation modalities and a learning operator [3]. Up until this point our set up abstracts away from computational constraints. In Section 5 we will discuss several surprising results obtained in computational inductive inference that challenge our understanding of knowledge as certainty, and of rational inquiry as requiring conjectures to be consistent (always accounting for all truthful information obtained so far) and conservative (never changing unless encountering information that is inconsistent with the current conjecture). In the final part of the paper we will mention some related directions of research.

1.1 Inductive Inference

Inductive inference is concerned with identifying a theory that underlies its partial presentation. Examples include isolating the right general scientific hypothesis on the basis of a limited number of empirical experiments, or finding a correct grammar that had produced finitely many observed sentences. This problem is difficult because each finite sample might be consistent with many different hypotheses of infinite generative capacities. *Computational* inductive inference poses effectivity constraints: hypotheses and learners should be computable objects. The field of *formal* learning theory (for an overview, see [33]) often abstracts away from the assumptions of computability. This means retaining some countability and enumerability assumptions, but giving up learner’s recursivity. As a consequence, within this more general setting, only the purely “structural” results of computational learning theory can be proven (for a formal account of this type of distinctions, see [15]). The first part of this paper, until Section 5, will focus on the structural aspects of learnable concepts, and the resulting notions of knowledge and belief. After that we will discuss some interesting consequences of computability.

We will start by discussing the basic setting of inductive inference and present some examples. Let \mathbb{N} be the set of all natural numbers (including 0), we call any $S \subseteq \mathbb{N}$ a language. In the general case, we will be interested in any countable class of languages together with its hypothesis space (a naming system), i.e., in the class $\mathcal{C} = (S_i)_{i \in I}$, the indices i will serve as names used by the learner to refer to a particular (possibly infinite) language. I is a countable index set, in most cases we will consider it to be \mathbb{N} . We will also sometimes take \mathcal{C} to be a finite set of (finite) languages.

► **Definition 1.** A text (positive presentation) τ of S is any infinite sequence enumerating all and only the elements from S (allowing repetitions). We will use the following notation: τ_n is the n -th element of τ ; $\tau \upharpoonright n$ is the sequence $(\tau_0, \tau_1, \dots, \tau_{n-1})$; $\text{set}(\tau)$ is the set of elements that occur in τ ; and if τ and σ be two sequences, then $\tau \wedge \sigma$ stands for their concatenation.

► **Definition 2.** A learner L is a function $L : \mathbb{N}^* \rightarrow \mathbb{N} \cup \{\uparrow\}$. The function is allowed to refrain from giving a natural number answer, in that case the output is marked by \uparrow . A learner L is (at most) once defined on $\mathcal{C} = (S_i)_{i \in \mathbb{N}}$ iff for any text τ for a language in \mathcal{C} and $n, k \in \mathbb{N}$ such that $n \neq k$ we have $L(\tau \upharpoonright n) = \uparrow$ or $L(\tau \upharpoonright k) = \uparrow$.

► **Definition 3** (Finite identifiability [24]). Let $\mathcal{C} = (S_i)_{i \in \mathbb{N}}$, and let L be a learner.

1. L finitely identifies S_i in \mathcal{C} on τ iff L is once defined on τ and the defined value is i .
2. L finitely identifies S_i in \mathcal{C} iff it finitely identifies S_i on every τ for S_i ;
3. L finitely identifies \mathcal{C} iff it finitely identifies every S_i in \mathcal{C} .

A class \mathcal{C} is finitely identifiable iff there is a learner L that finitely identifies \mathcal{C} .

► **Example 4.** We will now discuss two classes that are finitely identifiable, and one (very simple) class that is not. First consider the class containing all pairs of natural numbers, $\mathcal{C}_{pair} = \{\{n, k\} \mid n, k \in \mathbb{N}\}$. To see that \mathcal{C}_{pair} is finitely identifiable, consider the hypothesis space which uses a diagonal indexing of \mathcal{C}_{pair} . Let's take a learner that on the input of initial segments of τ outputs \uparrow as long as it sees only one number begin repeated, upon seeing the second number the learner outputs the corresponding index and, after that, switches back to the \uparrow answer forever. The learner finitely identifies every element of \mathcal{C}_{pair} because every pair is a subset of only one element of \mathcal{C}_{pair} .

Let us in turn consider a countable class of *infinite* languages. Let p_i be the i -th prime number (for $i \geq 1$), and $S_i = \{n \mid n \text{ is a multiple of } p_i\}$. $\mathcal{C}_{prime} = (S_i)_{i \in \mathbb{N} - \{0\}}$ is finitely identifiable. Our learner will refrain from a natural number answer as long as it receives non-primes. At the step when a prime is seen, it outputs the corresponding index, and repeats the “ \uparrow ”-answer from that point on. Since every prime number is an element of only one language in \mathcal{C}_{pair} , our learner's only conjecture will be the right one.

For contrast, let us look at a very simple class $\mathcal{C}_{simp} = \{S_0 = \{0\}, S_1 = \{0, 1\}\}$. It should be immediately evident this class is not finitely identifiable, but let us give a simple argument, as this example inspires the more complex kind of identifiability we will discuss shortly. Assume, for contradiction, that there is a learner that finitely identifies \mathcal{C}_{simp} . Then, by definition, on a finite initial segment of the text for S_0 the learner will output the index of S_0 , i.e., “0”, and it will be its only natural number answer. So, if after that it turns out that the text was in fact for S_1 , i.e., after that point the text will show 1, the learner will fail to identify S_1 on this text for S_1 , because it has already used its “one-shot” conjecture. This contradicts the assumption that the learner finitely identifies \mathcal{C}_{simp} .

The class \mathcal{C}_{simp} often inspires a (logical) objection: why is negative information disallowed? After all, if the learner is given an observation “not-1”, they will be able to settle the problem straight away. In this paper we restrict our attention to learning from positive information only. There are several reasons for this, in particular our desire to stay as close as possible to unsupervised learning. In any case, it is important to mention that learning from the so-called *informant*, i.e., from streams enumerating all positive and negative information consistent with the unknown language, is widely studied in inductive inference, recently even within the paradigm of finite identifiability (see [18, 32]).

Note that since in finite identifiability the learner is only allowed “one shot” at a correct conjecture, its hypotheses should be based on certainty, requiring the learner to have eliminated all other possibilities by the time the conjecture is made. Let us now turn to a more relaxed learner, which is allowed to change its mind, but is still required to *stabilize* to a correct answer after finitely many steps.

► **Definition 5** (Identifiability in the limit [24]). *Let $\mathcal{C} = (S_i)_{i \in \mathbb{N}}$, and let L be a learner.*

1. *L identifies S_i in \mathcal{C} in the limit on τ iff there is a $k \in \mathbb{N}$, s.t. for all $n \geq k$ $L(\tau \upharpoonright n) = i$;*
2. *L identifies S_i in \mathcal{C} in the limit iff it identifies S_i in the limit on every τ for S_i ;*
3. *L identifies \mathcal{C} in the limit iff it identifies in the limit every S_i in \mathcal{C} .*

A class \mathcal{C} is identifiable in the limit iff there is a learner L that identifies \mathcal{C} in the limit.

► **Example 6.** Let us start with a class that is identifiable in the limit. Consider $\mathcal{C}_{init} = \{S_n = \{0, \dots, n\} \mid n \in \mathbb{N}\}$, i.e., the class of finite initial segments of \mathbb{N} . The learner witnessing identifiability in the limit of this class, on the initial segment $\tau \upharpoonright n$ of a text τ for a language in \mathcal{C}_{init} , will output the largest n in $\text{set}(\tau \upharpoonright n)$. For every text for any language S_i in \mathcal{C}_{init} , i is the largest number that will ever be enumerated, so our learner will indeed stabilize on a correct language. For that to be possible, the learner must remain open-minded, i.e., the moment in which the right hypothesis is reached is not known to the learner.

Consider now a countable class of infinite sets containing complements of the sets in \mathcal{C}_{init} , i.e., the class $\mathcal{C}_{coinit} = \{S_n = \mathbb{N} - \{0, \dots, n\} \mid n \in \mathbb{N}\}$. This class is identified in the limit by a learner that always answers with the smallest natural number that was *not* eliminated so far. If we extend \mathcal{C}_{coinit} to include also the set of all natural numbers, i.e., we consider the class $\mathcal{C}_{coinit+} = \{S_0 = \mathbb{N}\} \cup \{S_{n+1} = \mathbb{N} - \{0, \dots, n\} \mid n \in \mathbb{N}\}$ we will retain identifiability in the limit: our learner should now answer “ $n + 1$ ” as long as the smallest number it hasn’t encountered so far is n , and upon receiving 0 the learner should change the conjecture to “0”, which stands for $S_0 = \mathbb{N}$.

Sometimes adding \mathbb{N} to an identifiable in the limit class can get us in trouble. Let’s return to the class \mathcal{C}_{init} and extend it with \mathbb{N} , i.e., consider $\mathcal{C}_{init+} = \{S_0 = \mathbb{N}\} \cup \{S_{n+1} = \{0, \dots, n\} \mid n \in \mathbb{N}\}$. This class is not identifiable in the limit. To see this we will construct an argument similar to the one for the class \mathcal{C}_{simp} in Example 4. Assume that \mathcal{C}_{init+} is identified in the limit by some learner L . Then L identifies in the limit $S_0 \in \mathcal{C}_{init+}$ on a text τ for S_0 . This means that after some finite initial segment $\tau \upharpoonright n$, L stabilizes to “0” (the index for \mathbb{N}). But $\tau \upharpoonright n$ is also an initial segment of a text τ' for S_n such that n is the largest number in $\text{set}(\tau \upharpoonright n)$. It follows that L does not identify in the limit S_n on the text τ' , which contradicts our assumption that L identifies \mathcal{C}_{init+} in the limit.

The concept of identifiability, its various kinds and characterisations, is fascinating and complex (for an introductory overview see, e.g., [30]). The above simple examples were selected because they will be instructive for us later on. For now there are two takeaways. Firstly, it’s important to realize that identifiability has a lot to do with inclusions among the sets in the class being learned. Secondly, the two types of identifiability differ with respect to their success conditions. Finite identifiability is “one-shot” learning, it results in (and requires) certainty about a correct conjecture. Identifiability in the limit does not – the learner is always allowed to change their conjecture, but (depending on the structure of the class) it can be successful in the sense that its conjecture will become “safe”, i.e., it will never change, because after some finite time it will not be contradicted by new (truthful) incoming information anymore.

1.2 Epistemic modal logic

In both kinds of identifiability reaching the correct hypothesis can be seen as arriving at a kind of knowledge. While finite identifiability results in *certainty*, identifiability in the limit carries *safe* or *reliable belief*. Such knowledge states are the topic of study of (multi-agent) epistemic and doxastic modal logic. This line of research interprets modal box \Box as a knowledge (or belief) operator, so that the intended meaning of the formula $\Box\varphi$ is “the agent knows that φ ” (or “the agent believes that φ ”).

► **Definition 7.** Let P be a countable set of propositions, with $p \in P$, and $\mathcal{A} = \{1, \dots, n\}$ is a set of agents, with $i \in \mathcal{A}$ as the name of some agent. The syntax of (epistemic) modal logic is given by:

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_i\varphi$$

Most commonly, epistemic languages are interpreted over Kripke models.

► **Definition 8.** A Kripke model M for n agents over P is a tuple $(W, v, (R_i)_{i \in \mathcal{A}})$, where: W is a non-empty set worlds; $v : P \rightarrow 2^W$ is a valuation; for each agent i , R_i is a binary accessibility relation on W . A pointed Kripke model is a pair (M, w) , of a Kripke model M and a designated world w from its domain W .

► **Definition 9.** We write $(M, w) \models \varphi$ to express that φ is true at w in M , and $(M, w) \not\models \varphi$ that φ is not true at w in M . The semantics of our language is defined in the following way:

$$\begin{aligned} (M, w) \models \top & \quad \text{always} \\ (M, w) \models p & \quad \text{iff } w \in V(p) \\ (M, w) \models \neg\varphi & \quad \text{iff } (M, w) \not\models \varphi \\ (M, w) \models \varphi \wedge \psi & \quad \text{iff } (M, w) \models \varphi \text{ and } (M, w) \models \psi \\ (M, w) \models \Box_i \varphi & \quad \text{iff for all } v \text{ with } (w, v) \in R_i, (M, v) \models \varphi \end{aligned}$$

The last clause of this definition complies to the intuitive understanding of what it means to know φ , i.e., φ being true in all worlds considered possible by (or “accessible to”) the agent. The properties of knowledge understood in this way will vary depending on the properties of the accessibility relations R (on the so-called frame conditions). The validities, when \Box_i is interpreted as knowledge, give additional insights into the properties of knowledge. For instance, the formula $\Box_i \varphi \rightarrow \Box_i \Box_i \varphi$ is read as “if the agent i knows that φ , then i knows that it knows φ ”, hence the name of this property: *positive introspection* of knowledge. This formula is valid in the class of models with transitive accessibility relation. We say that an axiom system Ax is a *logic of a class of models* \mathcal{M} iff Ax is sound and complete with respect to \mathcal{M} , see Table 1. Two systems are of special interest to us, the so-called $S5$ and $S4$. $S5$ (axioms $K+T+4+5$ in Table 1) is the logic of models with equivalence accessibility relations, while $S4$ ($K+T+4$), is the logic of locally connected, reflexive and transitive relations.

■ **Table 1** Validities with their epistemic interpretations and properties of R_i , for $i \in \mathcal{A}$.

Rules	
(MP)	if $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, then $\vdash \psi$
(Necc)	if $\vdash \varphi$, then $\vdash \Box_i \varphi$
Axioms	
(K)	$\Box_i(\varphi \rightarrow \psi) \rightarrow (\Box_i \varphi \rightarrow \Box_i \psi)$ (omniscience)
(T)	$\Box_i \varphi \rightarrow \varphi$ (truthfulness/reflexivity)
(D)	$\Box_i \varphi \rightarrow \neg \Box_i \neg \varphi$ (consistency/seriality)
(4)	$\Box_i \varphi \rightarrow \Box_i \Box_i \varphi$ (positive introspection/transitivity)
(5)	$\neg \Box_i \varphi \rightarrow \Box_i \neg \Box_i \varphi$ (negative introspection/Euclidean-ness)

Dynamic extension: Public Announcement Logic

Epistemic logic is made dynamic by adding action modalities. The first extension of this type was Public Announcement Logic (PAL, [31]), whose language is that of epistemic modal logic, enriched by the public announcement operator $[!\varphi]$.

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_i \varphi \mid [!\varphi]\varphi$$

The semantics of epistemic modal logic given in Definition 9 is accordingly extended in the following way: $(M, w) \models [!\varphi]\psi$ iff $(M, w) \models \varphi$ implies $(M|\varphi, w) \models \psi$, where $M|\varphi = (W', v', (R'_i)_{i \in \mathcal{A}})$ is defined as follows: $W' := \{w \in W \mid (M, w) \models \varphi\}$, $v' := v$ restricted to W' ; for each $i \in \mathcal{A}$, $R'_i := R_i \cap (W' \times W')$. Public announcement can be seen as incorporating new information in a radical way, by performing joint update – the accessibility relations of all agents are restricted to the set of worlds satisfying the incoming information φ , other worlds are simply disregarded.

Plausibility models and belief

Plausibility models (introduced in [9]) are a special kind of Kripke models in which the accessibility relations are *plausibility* relations. Plausibility satisfies the following conditions: it's reflexive, transitive, locally connected, and well-founded. It allows an interpretation of both certain knowledge (in the sense of $S5$), but also a weaker kind, called *safe belief* ($S4$), which is a belief that remains true under receiving true information.

We will use the symbol \preceq_i to denote plausibility of the agent i , with $v \preceq_i w$ standing for v is at least as plausible as w , which means that the worlds minimal according to \preceq_i will be seen as the most plausible ones for the agent i . The corresponding epistemic-doxastic language, called $K\Box$, includes two epistemic modalities, K_i interpreted as “hard information” possessed by the agent i , i.e., information that is true in all possible worlds in the connected component of agent i , while \Box_i stands for safe belief, i.e., the information that is true in all worlds that are at least as plausible for agent i .

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid \Box_i\varphi$$

This language is interpreted over a plausibility model, which is a tuple $M = (W, (\preceq_i)_{i \in \mathcal{A}}, v)$. We define the *connected component* of \preceq_i as follows: $cc_i(w) := \{v \in W \mid w(\preceq_i \cup \succeq_i)^*v\}$. The semantics for the propositional part is as usual, while the modal operators are defined in the following way:

$$\begin{aligned} (M, w) \models K_i\varphi & \quad \text{iff} \quad \text{for all } v \in cc_i(w) \text{ with } (M, v) \models \varphi \\ (M, w) \models \Box_i\varphi & \quad \text{iff} \quad \text{for all } v \text{ with } v \preceq_i w, (M, v) \models \varphi \end{aligned}$$

We can also define regular (possibly false) belief in φ , as φ being true in all most plausible worlds.

$$(M, w) \models B_i\varphi \quad \text{iff} \quad \text{for all } v \in \min_{\preceq_i}(W), (M, v) \models \varphi$$

This interpretation of belief is consistent with the standard approach of the belief revision theory (see, e.g., [36]). To know more about plausibility models and the logic $K\Box$ the reader is encouraged to consult adequate passages of [8].


2 Classes of languages as epistemic spaces

To express our inductive inference problems in the setting of epistemic modal logic we need to establish common ground between the two paradigms. In what follows we will consider the case of a single-agent learning, and so we will simplify the framework of epistemic logic accordingly, by omitting the agent indices $i \in \mathcal{A}$. First, we will express the classes of languages from Section 1.1 as single-agent *epistemic spaces*.


► **Definition 10.** An epistemic space is $\mathcal{S} = (X, \mathbb{U})$, where $\mathbb{U} \subseteq 2^X$ and both X and \mathbb{U} are at most countable.

Given a class of languages $\mathcal{C} = \{S_i \mid i \in I\}$, we set its corresponding epistemic space to be $\mathcal{S}_{\mathcal{C}} = (X_{\mathcal{C}}, \mathbb{U}_{\mathcal{C}})$, with $X_{\mathcal{C}} = \{x_i \mid i \in I\}$ and for each $u_n \in \mathbb{U}_{\mathcal{C}}$, $x_k \in u_n$ iff $n \in S_k$.


► **Example 11.** Recall the class \mathcal{C}_{simp} from Example 4, the corresponding epistemic space is $\mathcal{S} = (X, \mathbb{U})$, $X = \{x_0, x_1\}$, and $\mathbb{U} = \{u_0, u_1\}$ with $u_0 = \{x_0\}$, $u_1 = \{x_0, x_1\}$. The space is depicted below.



► **Example 12.** The epistemic space for $\mathcal{C}_{coinit+}$ from Example 6 is depicted below.



► **Example 13.** The epistemic space for \mathcal{C}_{init+} from Example 6 is depicted below.



Note that epistemic spaces can be seen as single-agent Kripke models with an equivalence accessibility relation – all worlds are considered equally plausible.

We will now reformulate the learning notions from Section 1.1 in terms of epistemic spaces.

► **Definition 14.** Let $S = (X, \mathbb{U})$ be an epistemic space. A data stream is an infinite sequence $\vec{U} = (u_0, u_1, \dots)$ from \mathbb{U} . A data sequence $\vec{U} \upharpoonright n$ is a finite initial segment of \vec{U} of length $n + 1$. A data stream \vec{U} is sound with respect to x if every element listed in \vec{U} contains x , and it is complete with respect to x if every element containing x is listed in \vec{U} .

► **Definition 15.** Let $S = (X, \mathbb{U})$ be an epistemic space and let σ be a data sequence. A learner L is a function that on a data sequence σ outputs a conjecture, a subset of X , $L(\sigma) \subseteq X$.

Thus, a learner outputs a subset of an epistemic space, which is more general than in the case of the classical learner in Section 1.1. This is in order to allow the learner to make more general conjectures.

► **Definition 16.** An epistemic space $S = (X, \mathbb{U})$ is finitely identified by a learner L if for every $x \in X$ and every data stream \vec{U} for x , L is once-defined and there is $k \in \mathbb{N}$ such that $L(\vec{U} \upharpoonright n) = \{x\}$. S is finitely identifiable if it is finitely identified by a learner L .

► **Definition 17.** An epistemic space $S = (X, \mathbb{U})$ is identified in the limit by a learner L if for every $x \in X$ and every data stream \vec{U} for x , there is $k \in \mathbb{N}$ such that $L(\vec{U} \upharpoonright n) = \{x\}$ for all $n \geq k$. S is identifiable in the limit if it is identified in the limit by a learner L .

3 Learning by belief revision

Learning often proceeds by an underlying implicit order of possible worlds. This brings to mind the plausibility relation described in Section 1.2. Let us consider the epistemic space corresponding to $\mathcal{C}_{\text{coin}^+}$ in Example 12. The successful learner, upon receiving information u , conjectures the *maximal*, according to the indices, hypothesis consistent with u . This can be interpreted as restricting the X to u , which leads the agent to “believe” everything that is true in the most plausible worlds in the restricted domain. This is how belief revision operators work in general, they adjust plausibility based on what is observed. In the contributions [22, 4, 6] it has been shown how belief revision methods can be seen as constructive strategies for converging to the right conjecture. An interesting question is how well various belief revision methods perform as learning engines. In order to capture this intuition we will now build up towards belief-revision based learners.

► **Definition 18.** *Let $\mathcal{S} = (X, \mathbb{U})$ be an epistemic space. A plausibility assignment PL is a map that assigns to \mathcal{S} some plausibility (total, reflexive, and transitive) relation \preceq on X , thus converting it into a plausibility space $PL(\mathcal{S}) = (X, \mathbb{U}, \preceq)$.*

Plausibility spaces are single-agent plausibility models from Section 1.2, and the language of $K\Box$ is interpreted on them in an analogous way. Note that the above definition allows the plausibility assignment to be non well-founded on X , i.e., there might be an infinite descending chain $s_0 \succ s_1 \succ s_2 \succ \dots$, where \prec is the strict plausibility relation, given by: $s \preceq t$ and $t \not\prec s$. This requires a modification of the definition of belief so that we cover the case when the minimum of \preceq does not exist. This is done in the following way: $(M, w) \models B_i \varphi$ iff there is a k such that for all $v \preceq_i k$, $(M, v) \models \varphi$.

Plausibility order will be revised upon receiving new information. The ways of revision vary, but they can be given a general definition.

► **Definition 19.** *A one-step revision method is a function R_1 that, for any plausibility space $\mathcal{B} = (X, \mathbb{U}, \preceq)$ and any $u \in \mathbb{U}$, outputs a new plausibility space $R_1(\mathcal{B}, u)$. A (iterated) belief revision method R is obtained by iterating a one-step revision method R_1 : $R(\mathcal{B}, \lambda) = \mathcal{B}$, for the empty data sequence λ , and for a data sequence σ and $u \in \mathbb{U}$, $R(\mathcal{B}, \sigma \wedge u) = R_1(R(\mathcal{B}, \sigma), u)$.*

► **Definition 20.** *Every belief revision method R , together with a plausibility assignment PL , generates a learning method L_R^{PL} , called a belief revision-based learning method, given by:*

$$L_R^{PL}(\mathcal{S}, \sigma) := \min R(PL(\mathcal{S}), \sigma).$$

With this apparatus at hand, we can ask the question: how good are belief revision methods as learning methods? Among many we will consider three widely studied belief-revision policies: conditioning, the minimal, and the lexicographic revision (a DEL account of those is given in [10]). Conditioning removes possible worlds inconsistent with the incoming information (just like public announcements do); minimal revision selects the best worlds satisfying the incoming information and makes them the most plausible; and lexicographic revision puts all the worlds satisfying the incoming information to be better than all the ones that do not (within the two clusters the order remains the same). The relevant question here is: given a belief revision method, can the belief-revision based learner identify in the limit everything that’s learnable by an arbitrary learning method? If it is so, we say that revision method generates a *universal* learning method. It turns out that under the condition of sound and complete streams conditioning and lexicographic revision are universal, while minimal revision *is not*. If we relax the condition of soundness on data streams (a stream is


fair, if it allows finitely many errors, and each error must be accounted for truthfully in the future), conditioning loses its power – once a possible world is removed on the basis of an incorrect piece of data, it cannot be revived, and so conditioning is no longer an universal learning method. Lexicographical revision remains universal on fair streams.

Another important insight from this analysis is that reaching the full learning power of belief revision requires non-standard plausibility relations, which is important because in belief revision theory plausibility is typically assumed to be well-founded. Under the assumption of well-foundedness none of the revision methods generates a universal learning method. The epistemic space of the class $\mathcal{C}_{\text{coin}it+}$ in Example 12 shows why – initially the learner must allow all possible worlds, and so must assume infinitely increasing plausibility, i.e., if $i < j$ then $x_j \prec x_i$. If a well-founded plausibility is initially assigned to this space, the learning process will inevitably overgeneralise while attempting to learn some $x \in X$.

4 Topological perspective on inductive inference

Whether or not an epistemic space is learnable has a lot to do with possible worlds being *separable* from others by the incoming data. General topology [25] not only puts separability properties as first class citizens, but it has also been proposed as a logic of observation [37]. In this perspective the points of the space represent possible worlds and basic open neighborhoods of a point code information that comes about from executing finitely many observations or measurements (for examples, see [3]).

Figure 1 shows several basic examples of how points can be separated by observations.



■ **Figure 1** Some topological separation properties.

The necessary condition for identifying individual worlds is T_0 : the space cannot contain two inseparable points, i.e., for any two worlds x and y , there must be an observation that separates them one way or the other, i.e., there must be an o such that $x \in o$ and $y \notin o$, or $y \in o$ and $x \notin o$. As a matter of fact, finite identifiability requires a stronger condition, in which every state has an observation that separates it from all other states in the space. Finding a condition for identifiability in the limit requires a more specific condition, concerning singleton worlds being *locally closed*.

► **Definition 21.** A topological space (X, \mathbb{O}) is T_d iff $\forall x \in X \exists U \in \mathbb{O} U \setminus \{x\} \in \mathbb{O}$.

In order to apply the tools of topology to epistemic spaces, we have to close them on finite intersections (corresponding to accumulating finitely many observations) and arbitrary unions. The latter addition does not disrupt our learning scenarios, since arbitrary unions do

not bring any extra information to the learning process. By considering a topology, we also allow our learners (in principle) to observe an \emptyset , i.e., contradictory datum – in the setting of learning from positive data, this will never happen.

► **Definition 22.** Let $\mathcal{S} = (X, \mathbb{U})$ be an epistemic space. A topology generated by \mathcal{S} is $\tau_{\mathcal{S}} = (X, \mathbb{O})$, where \mathbb{O} is defined in the following way:

1. if $u \in \mathbb{U}$, then $u \in \mathbb{O}$;
2. $\emptyset \in \mathbb{O}$.
3. if O is a finite subset of \mathbb{O} , then $\bigcap O \in \mathbb{O}$;
4. if O is an arbitrary subset of \mathbb{O} , then $\bigcup O \in \mathbb{O}$;
5. $X \in \mathbb{O}$.

We obtain the following characterisation identifiable in the limit epistemic spaces.

► **Theorem 23** ([5]). *An epistemic space \mathcal{S} is identifiable in the limit iff its generated topology $\tau_{\mathcal{S}}$ is T_d .*

Topological analysis of inductive inference has been picked up in similar settings: [34] propose the so-called *positive information topology*. [17] analyzes concept spaces (structures like our classes of languages in Section 1.1) and also obtains a T_d characterisation. That last contribution also contains interesting results involving topological and algebraic characterisations that could without a doubt be of use in logical approaches to learning.

4.1 Topological semantics of modal logic

Having established the topological characterisation of identifiable in the limit spaces, we can now ask the question of a corresponding modal logic. It's only natural we consider topological semantics. We will now interpret standard modal language on topological models via the so-called topo-semantics.

► **Definition 24.** A topological model (or a topo-model) $M = (X, \mathbb{O}, v)$ is a topology (X, \mathbb{O}) together with a valuation function $v : P \rightarrow 2^X$.


► **Definition 25** (Topo-semantics). Let $M = (X, \mathbb{O}, v)$ be a topological model and $x \in X$:

$$\begin{array}{ll}
 M, x \models \top & \text{always} \\
 M, x \models p & \text{iff } x \in v(p) \\
 M, x \models \neg\varphi & \text{iff not } M, x \models \varphi \\
 M, x \models \varphi \wedge \psi & \text{iff } M, x \models \varphi \text{ and } M, x \models \psi \\
 M, x \models \Box\varphi & \text{iff there is } U \in \tau(x \in U \text{ and for all } y \in U: M, y \models \varphi)
 \end{array}$$

Under this interpretation of modal box, we get, for instance, that $S4$ is the topo-logic of all topological spaces [28]. Unfortunately, it is also known that T_d spaces, the identifiability-adequate class of models, is not topo-definable. We can however change the way we view \Box , allowing belief to be false (see Figure 2). This agrees with the kind of belief occurring in learning in the limit.

► **Definition 26** (d -semantics). Given a topological model $M = (X, \mathbb{O}, v)$ and a state $x \in X$:

$$\begin{array}{ll}
 M, x \models_d \top & \text{always} \\
 M, x \models_d p & \text{iff } x \in v(p) \\
 & \dots \\
 M, x \models_d \Box\varphi & \text{iff } \exists U \in \tau(x \in U \ \& \ \forall y \in U - \{x\} \ M, y \models_d \varphi)
 \end{array}$$



■ **Figure 2** Relational (left), topo-semantic (center) and d-semantic (right) interpretations of \Box , with φ true in red areas.

Quite a lot is known about sound and complete d -axiomatisations (for an overview, see [11]): wKD45 is the d -logic of dense spaces; KD45 is the d -logic of DSO-spaces, where DSO stands for “derived sets are open”; GL is the d -logic of scattered spaces, where GL is the Grzegorzcyk axiom $\Box(\Box\varphi \rightarrow \varphi) \rightarrow \Box\varphi$; weak K4 is the d -logic of all topological spaces. Finally, we also have that K4 is the d -logic of all T_d -spaces. This is a technically valuable result, but given the relative poverty of our modal language it gives us very little insight into the possible *epistemic* interpretations of this modal language characterising spaces that are identifiable in the limit. To be of use a logic of learning theory should involve a semantic notion of a learner that drives beliefs, and also dynamic modalities to talk about receiving new information. Dynamic Logic for Learning Theory (DLLT, [3]), built upon Subset Space Logic [16] overcomes both challenges.

4.2 Dynamic Logic of Learning Theory

With Subset Space Logic (SSL, [16]) we are back to an epistemic-doxastic language with two modalities, like the one in Section 1.2.

► **Definition 27** (SSL Syntax). *Let P be a countable set of propositional symbols and $p \in P$.*

$$\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi \mid \Box\varphi$$

This logic is interpreted in topological models, in this context called *intersection models*.

► **Definition 28.** *An intersection model $M = (X, \mathbb{O}, v)$ is a topology (X, \mathbb{O}) together with a valuation function $v : P \rightarrow 2^X$.*

The meaning of formulas of SSL is defined in the following way.

► **Definition 29** (SSL Semantics). *Let $M = (X, \mathbb{O}, v)$, $U \in \mathbb{O}$, and $x \in U$.*

$$\begin{array}{ll} M, x, U \models p & \text{iff } x \in v(p) \\ M, x, U \models \neg\varphi & \text{iff } M, x, U \not\models \varphi \\ M, x, U \models \varphi \wedge \psi & \text{iff } M, x, U \models \varphi \text{ and } M, x, U \models \psi \\ M, x, U \models K\varphi & \text{iff } \forall y \in U \ M, y, U \models \varphi \\ M, x, U \models \Box\varphi & \text{iff } \forall O \in \mathbb{O} \text{ if } x \in O \subseteq U \text{ then } M, x, O \models \varphi \end{array}$$

In the above definition \Box (also called effort modality) accounts for *observational effort* – making the epistemic effort to obtain more information about a possible world. It denotes how shrinking the open neighborhood gives a more accurate approximation of the actual state of the world. Note that the formulas of SSL are evaluated at triples, the novel addition

being the third component, i.e., a particular open neighbourhood of the point at which the formula is evaluated. This can be seen as evaluating formulas at some stage of inquiry – a starting actual world, together with some truthful information already received.

Dynamic Logic of Learning Theory (DLLT, [3]) extends SSL in the following way.

► **Definition 30** (DLLT Syntax). *Let p and o be drawn from countable sets of propositional and observational symbols, P and O respectively.*

$$\varphi := p \mid o \mid L(\vec{o}) \mid \neg\varphi \mid \varphi \wedge \psi \mid K\varphi \mid \Box\varphi \mid [o]\varphi$$

Note that the language of DLLT allows talking directly about observations. The two remaining additions to the language are dynamic. The first one allows expressing Learner’s conjectures after being given some finite sequence of observations. The second is a more basic “learning” known already from PAL. The learning model has to accordingly involve a semantics notion of the learner.

► **Definition 31.** *The structure $M_L = (X, \mathbb{O}, \mathbb{L}, v)$ is a learning model consisting of a topology (X, \mathbb{O}) ; the learner $\mathbb{L} : \mathbb{O} \rightarrow 2^X$ is such that $\mathbb{L}(O) \subseteq O$, and if $O \neq \emptyset$ then $\mathbb{L}(O) \neq \emptyset$. Additionally: $\mathbb{L}(\vec{O}) := \mathbb{L}(\bigcap \text{set}(\vec{O}))$. Finally, $v : P \cup O \rightarrow 2^X$ is a valuation.*

► **Definition 32** (DLLT semantics). *Given a learning model M_L , $U \in \mathbb{O}$, and $x \in U$:*

$$\begin{array}{ll} M, x, U \models o & \text{iff } x \in v(o) \\ M, x, U \models L(o_1, \dots, o_n) & \text{iff } x \in \mathbb{L}(U, v(o_1), \dots, v(o_n)) \\ M, x, U \models [o]\varphi & \text{iff } x \in v(o) \text{ implies } M, x, U \cap v(o) \models \varphi \end{array}$$

Clauses for p , \neg , \wedge , K , and \Box are as in the semantics of SSL, see Definition 29.

The paper [3] provides a sound and complete axiomatisation of DLLT with respect to the learning models. The advantage of this logic is not only that it builds upon SSL, which with its effort modality \Box has a very natural epistemic interpretation, but also that it allows expressing the conditions of learnability discussed in the beginning of this paper. We have that: $M, x, U \models \Diamond Kp$ iff p is *finitely learnable* at x and $M, x, U \models p \wedge \Diamond \Box Bp$ iff p is *learnable in the limit* by \mathbb{L} at x (here B is an abbreviation $B^{\vec{o}}\varphi := K(L(\vec{o}) \rightarrow \varphi)$ and $B\varphi := B^\lambda\varphi$).

5 Knowledge in computable inductive inference

Computable inductive inference is concerned with learnability under the assumption of recursivity of the learner. It also often assumes that languages are, at least, recursively enumerable. Since the learner should be able to decide whether a given observation is consistent with a given hypothesis, indexed families of (non-empty) recursive languages are extensively studied as the background of learning. These are $\mathcal{C} = (S_i)_{i \in \mathbb{N}}$ for which a computable function $f : \mathbb{N} \times U \rightarrow \{0, 1\}$ exists that uniformly decides \mathcal{C} , i.e.:

$$f(i, w) = \begin{cases} 1 & \text{if } w \in S_i, \\ 0 & \text{if } w \notin S_i. \end{cases}$$

Even though the existence of such a computable function seems like a natural element of any agency, the logical approaches to knowledge representation, both belief revision and epistemic logic, do not consider this assumptions, or treat it as implicit. Belief revision operators are usually studied from the perspective of rationality postulates, and epistemic logic takes agents (learners) to be relations in a (dynamic) Kripke model, abstracting away from computability. The decision of whether or not $w \models p$ is taken to be an atomic step of computation which does not deserve much attention.

In this section we will assume the perspective of learning theory on certain aspects of knowledge that are usually taken for granted in epistemic logic and turn out to be restrictive when seen through the lenses of computable learning. We will start with finite identifiability and then move on to identifiability in the limit.

What are the conditions for reaching certainty? Intuitively, it seems that as soon as the agent receives information that is inconsistent with all the alternatives, she can be sure about the correct conjecture. Recall the concept of finite identifiability given in Definition 3. The necessary and sufficient condition for finite identifiability [29, 27] involves the *definite finite tell-tale set*, a more strict kind of tell-tale than the one involved in the characterisation of identifiability in the limit [2].

► **Definition 33** ([29, 27]). Let $\mathcal{C}=(S_i)_{i \in \mathbb{N}}$ be an indexed family of recursive languages. A set D_i is a definite finite tell-tale set (DFTT) for S_i in \mathcal{C} if

1. $D_i \subseteq S_i$,
2. D_i is finite, and
3. for any index j , if $D_i \subseteq S_j$ then $S_i = S_j$.

► **Theorem 34** ([29, 27]). An indexed family of recursive languages $\mathcal{C}=(S_i)_{i \in \mathbb{N}}$ is finitely identifiable from positive data iff there is an effective procedure $\mathcal{D} : \mathbb{N} \rightarrow \mathcal{P}^{<\omega}(\mathbb{N})$, given by $n \mapsto D_n$, that on input i produces a definite finite tell-tale of S_i in \mathcal{C} .

Let us again take a finitely identifiable class \mathcal{C} , and S_i in \mathcal{C} . Now, consider the collection \mathbb{D}_i of all DFTTs of S_i in \mathcal{C} .

► **Definition 35**. Let $\mathcal{C}=(S_i)_{i \in \mathbb{N}}$ be an indexed family of recursive sets. \mathcal{C} is finitely identifiable in the fastest way if and only if there is a learning function L such that, for each τ and for each $i \in \mathbb{N}$, $L(\tau \upharpoonright n) = i$ iff $\exists D_i^j \in \mathbb{D}_i$ ($D_i^j \subseteq \text{set}(\tau \upharpoonright n)$) and $\neg \exists D_i^k \in \mathbb{D}_i$ ($D_i^k \subseteq \text{set}(\tau \upharpoonright n - 1)$). We will call such L a fastest learning function.

The fastest learning thus our desired learner, as soon as a definite finite telltale arrives, certainty is reached and the correct conjecture is issued. It turns out that, even under the assumptions of uniform recursivity of the class and recursivity of the learner, things are not that simple.

► **Theorem 36** ([23]). There exists an indexed family of recursive sets $\mathcal{C}=(S_i)_{i \in \mathbb{N}}$ that is recursively finitely identifiable but is not recursively finitely identifiable in the fastest way.

The proof of this theorem relies on recursively inseparable sets [35]. The the witness class still possesses the well-separated topological structure, but a recursive fastest learner would be able decide that a minimal DFTT has been observed, and so it would provide a recursive separating set for something that is recursively inseparable (for details, see [23]).

Results similar in spirit are very common in identifiability in the limit. Two properties of learners deserve our special attention because of their close connection to epistemic logic and rationality postulates in belief revision. They are *conservativity* and *consistency*.

► **Definition 37**. A learning function L is conservative if, for each σ and x , $\text{set}(\sigma \wedge \langle x \rangle) \subseteq S_{L(\sigma)}$ implies $L(\sigma \wedge \langle x \rangle) = L(\sigma)$.

A property of this kind is a guiding principle of belief revision and dynamic epistemic logic. If the incoming information is already accounted for within the learner's conjecture, learning takes no effect. Under the assumptions of computability, conservativity turns out to be a restrictive property – there are uniformly recursive families that can be learned by a recursive learner, but not by a conservative recursive learner.

► **Theorem 38** ([39]). *There is a uniformly recursive family of languages \mathcal{C} , which is effectively identifiable in the limit, but which is not effectively identifiable by a conservative learner.*

Another crucial principle in learning is consistency, which in this case means consistency of the conjecture with the incoming sequence of information at every point of learning.

► **Definition 39.** *A learning function L is consistent if, for each σ , $\text{set}(\sigma) \subseteq S_{L(\sigma)}$.*

In AGM Belief Revision theory [1], this principle is called *success postulate* and is seen as one of the rationality postulate. Despite its very obvious character, for classes of recursive languages consistency turns out to be restrictive. For indexed families of recursive languages consistency is not limiting, but it turns out that sometimes issuing conjectures inconsistent with the information obtained so far can significantly speed up the learning process.

► **Theorem 40** ([38]). *There is a uniformly recursive class of languages that is polynomially effectively identifiable in the limit from informant but is not polynomially effectively consistently identifiable in the limit from informant.*

6 Conclusions

The standard notion of identifiability in the limit and its little cousin, finite identifiability, naturally correspond to safe belief and certain knowledge. Despite the clear intuitions, it is not easy to capture these analogies. Topological notion of epistemic effort and dynamic extensions of epistemic modal logic provide an adequate set of tools and an appropriate level of expressivity. The common ground between the theory of inductive inference and epistemic modal logic opens many opportunities for cross-fertilisation between the fields, e.g., understanding the limitations of computable learning in the context of belief revision, as outlined in Section 5. Another direction is looking into inductive inference of action models of dynamic epistemic logic [12, 13, 14]. Finally, (dynamic) epistemic modal logic is by default developed for multi-agent systems. It is interesting to investigate how that kind of multi-agency can be translated into the settings of computable inductive inference.

References

- 1 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.
- 2 Dana Angluin. Inductive inference of formal languages from positive data. *Information and Control*, 45(2):117–135, 1980.
- 3 Alexandru Baltag, Nina Gierasimczuk, Aybüke Özgün, Ana Lucia Vargas Sandoval, and Sonja Smets. A dynamic logic for learning theory. *Journal of Logical and Algebraic Methods in Programming*, 109:100485, 2019.
- 4 Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. Belief revision as a truth-tracking process. In K. Apt, editor, *TARK'11: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge, Groningen, The Netherlands, July 12-14, 2011*, pages 187–190. ACM, New York, NY, USA, 2011.
- 5 Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. On the solvability of inductive problems: A study in epistemic topology. In R. Ramanujam, editor, *Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015, Carnegie Mellon University, Pittsburgh, USA, June 4-6, 2015*, volume 215 of *EPTCS*, pages 81–98, 2015.
- 6 Alexandru Baltag, Nina Gierasimczuk, and Sonja Smets. Truth-tracking by belief revision. *Studia Logica*, 107(5):917–947, 2019.

- 7 Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In Itzhak Gilboa, editor, *TARK'98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 43–56, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.
- 8 Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.
- 9 Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *LOFT'08: Proceedings of 8th Conference on Logic and the Foundations of Game and Decision Theory*, number 3 in Texts in Logic and Games, pages 9–58. Amsterdam University Press, 2008.
- 10 Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 2:129–155, 2007.
- 11 Johan van Benthem and Guram Bezhanishvili. Modal logics of space. In Marco Aiello, Ian Pratt-Hartmann, and Johan Van Benthem, editors, *Handbook of Spatial Logics*, pages 217–298. Springer Netherlands, Dordrecht, 2007.
- 12 Thomas Bolander and Nina Gierasimczuk. Learning actions models: Qualitative approach. In Wiebe van der Hoek, Wesley H. Holliday, and Wen-Fang Wang, editors, *Logic, Rationality, and Interaction - 5th International Workshop, LORI 2015 Taipei, Taiwan, October 28-31, 2015, Proceedings*, volume 9394 of *Lecture Notes in Computer Science*, pages 40–52. Springer, 2015.
- 13 Thomas Bolander and Nina Gierasimczuk. Learning to act: qualitative learning of deterministic action models. *Journal of Logic and Computation*, 28(2):337–365, 2018.
- 14 Thomas Bolander, Nina Gierasimczuk, and Andrés Occhipinti Liberman. Learning to act and observe in partially observable domains. In Nick Bezhanishvili, Rosalie Iemhoff, and Fan Yang, editors, *Outstanding Contributions to Logic: Dick de Jongh on intuitionistic and provability logics*. Springer, 2023.
- 15 John Case and Timo Kötzing. Topological separations in inductive inference. *Theoretical Computer Science*, 620:33–45, 2016.
- 16 Andrew Dabrowski, Lawrence S. Moss, and Rohit Parikh. Topological reasoning and the logic of knowledge. *Annals of Pure and Applied Logic*, 78(1):73–110, 1996. Papers in honor of the Symposium on Logical Foundations of Computer Science, Logic at St. Petersburg.
- 17 Matthew de Brecht and Akihiro Yamamoto. Topological properties of concept spaces. *Information and Computation*, 208(4):327–340, 2010.
- 18 Dick de Jongh and Ana Lucia Vargas-Sandoval. Finite identification with positive and with complete data. In Alexandra Silva, Sam Staton, Peter Sutton, and Carla Umbach, editors, *Language, Logic, and Computation*, pages 42–63, Berlin, Heidelberg, 2019. Springer Berlin Heidelberg.
- 19 Hans van Ditmarsch, Wiebe Van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, The Netherlands, 2007.
- 20 Nina Gierasimczuk. Bridging learning theory and dynamic epistemic logic. *Synthese*, 169(2):371–384, 2009.
- 21 Nina Gierasimczuk. Learning by erasing in dynamic epistemic logic. In Adrian Horia Dediu, Armand Mihai Ionescu, and Carlos Martin-Vide, editors, *LATA'09: Proceedings of 3rd International Conference on Language and Automata Theory and Applications, Tarragona, Spain, April 2-8, 2009*, volume 5457 of *Lecture Notes in Computer Science*, pages 362–373. Springer, The Netherlands, 2009.
- 22 Nina Gierasimczuk. *Knowing One's Limits. Logical Analysis of Inductive Inference*. PhD thesis, Universiteit van Amsterdam, The Netherlands, 2010.
- 23 Nina Gierasimczuk and Dick de Jongh. On the complexity of conclusive update. *The Computer Journal*, 56(3):365–377, 2013.

- 24 E. Mark Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.
- 25 John L. Kelley. *General topology*. Springer-Verlag New York, 1975.
- 26 Steffen Lange, Rolf Wiehagen, and Thomas Zeugmann. Learning by erasing. In Setsuo Arikawa and Arun Sharma, editors, *ALT*, volume 1160 of *Lecture Notes in Computer Science*, pages 228–241. Springer, 1996.
- 27 Steffen Lange and Thomas Zeugmann. Types of monotonic language learning and their characterization. In *COLT'92: Proceedings of the 5th Annual ACM Conference on Computational Learning Theory, Pittsburgh, PA, USA, July 27-29, 1992*, pages 377–390. ACM, New York, NY, USA, 1992.
- 28 John C. C. McKinsey and Alfred Tarski. The algebra of topology. *Annals of Mathematics*, 1944.
- 29 Yasuhito Mukouchi. Characterization of finite identification. In Klaus Jantke, editor, *AII'92: Proceedings of the International Workshop on Analogical and Inductive Inference, Dagstuhl Castle, Germany, October 5-9, 1992*, volume 642 of *Lecture Notes in Computer Science*, pages 260–267. Springer, Berlin/Heidelberg, 1992.
- 30 Daniel Osherson, Michael Stob, and Scott Weinstein. *Systems that Learn*. M.I.T. Press, Cambridge, MA, USA, 1986.
- 31 Jan Plaza. Logics of public communications. In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
- 32 Ana Lucía Vargas Sandoval. *On the Path to the Truth: Logical and Computational Aspects of Learning*. PhD thesis, Universiteit van Amsterdam, The Netherlands, 2020.
- 33 Oliver Schulte. Formal Learning Theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2022 edition, 2022.
- 34 Olivier Schulte and Cory Juhl. Topology as epistemology. *Monist*, 79(1):141–147, 1996.
- 35 Raymond M. Smullyan. Undecidability and recursive inseparability. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 4(7–11):143–147, 1958.
- 36 Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In Brian Skyrms and William L. Harper, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134. Dordrecht: Kluwer, 1988.
- 37 Steven Vickers. *Topology Via Logic*. Cambridge University Press, 1996.
- 38 Rolf Wiehagen and Thomas Zeugmann. Ignoring data may be the only way to learn efficiently. *J. Exp. Theor. Artif. Intell.*, 6(1):131–144, 1994.
- 39 Thomas Zeugmann, Steffen Lange, and Shyam Kapur. Characterizations of monotonic and dual monotonic language learning. *Information and Computation*, 120(2):155–173, 1995.