Computational Learning in Dynamic Logics
Course Notes
Day 1

Nina Gierasimczuk and Caleb Schultz Kisby

@NASSLLI, June 2025

1 Notes on inductive inference

Let us consider a game of inductive inference (after Osherson et al., 1986). It starts by taking a class of sets $\mathcal{L} = \{\mathbb{N} - \{n\} \mid n \in \mathbb{N}\}$, where $\mathbb{N} = \mathbb{N}^+ \cup \{0\}$ is the set of all natural numbers. So, the following will be examples of sets in $\mathcal{L}$: $\{1, 2, 3, 4, 5, \ldots\}$, $\{0, 2, 3, 4, 5, \ldots\}$, $\{0, 1, 3, 4, 5, \ldots\}$, etc. I will now choose one of those sets secretly and your task will be to guess which one I have in mind. You will be guessing on the basis of elements of the set, which I will reveal to you one by one. Read the numbers in the sequence below and each time you read the new number guess one of the sets: $1, 3, 4, 2, 6, 7, 8, \ldots$. Let us stop at 8. Try to answer the following questions:

1. Are you confident about your current guess? What would make you change your guess?
2. What was your 'guessing rule'?
3. What winning condition would make this game interesting? How about: 'you win if at least one of your guesses is correct'?
4. What winning condition would make this game interesting? How about: 'you win if you make exactly one guess and that guess will be correct'?
5. And how about: 'you win if you succeed to make a right guess and never change your mind after that'? How many wrong guesses could you make under this condition?
6. Assume that I'll give you all and only truthful clues. What would be the guessing rule to win according to the last winning condition?
7. Add to $\mathcal{L}$ the set $\{0, 1, 2, 3, 4, 5, \ldots\}$. Is your guessing rule still good?
8. While keeping $\{0, 1, 2, 3, 4, 5, \ldots\}$ in, assume that I'll give you all and only truthful clues, and I'll guarantee they are ordered increasingly. Can you win the game?

Formal Learning Theory (FLT) is a mathematical framework to capture such learning effects. It goes back to the the 1960's; to Putnam (1965), Gold (1967), and Solomonoff (1964). This course is about a computational treatment of learning problems like the one in our motivating example. Implicitly, it is also about the problem of induction and related issues in epistemology and philosophy of science. This framework helps addressing the abstract problems of language learning and grammar inference, computable learning in artificial intelligence, but also issues in scientific inquiry and epistemology: fallible knowledge and reliable learning.

2 FLT: Frameworks Overview

The initial motivating game-example implicitly followed an underlying framework of learning. In FLT such framework is ofter referred to as *learning paradigm*, in which the following elements are specified:

1. Possible realities.
2. Hypotheses.
3. Information accessible to the learner.
4. Learner.
5. Success criterion.

2.1 Language Learning

The initial game is an example of *Language Learning*, which is also known as *Set Learning*, or *Numerical Paradigm*. Its specification is as follows:

1. Possible realities: sets of numbers.
2. Hypotheses: some names of sets.
3. Information accessible to the learner: sequences of numbers which are initial segments of infinite streams of elements of one of the sets.
4. Learner: a function that takes a sequence and outputs a hypothesis.
5. Success criterion: after finite number of outputs the answers stabilize on a correct answer.

Language learning is a subject of vast existing and on-going work of various levels of computational generality. Prominent examples of excellent overviews include those by Osherson et al. (1986); Martin and Osherson (1998); Zeugmann and Lange (1995); Lange et al. (2008).

2.2 Function Learning

Function learning is also known as *Learning of Functional Languages*, and can be captured by the following specification.

1. Possible realities: functions.
2. Hypotheses: names of functions.
3. Information accessible to the learner: sequences of pairs (argument, value).
4. Learner: function that takes a sequence and outputs a hypothesis.
5. Success criterion: after finite number of outputs the answers stabilize on a correct answer.

This setting has been applied, for example, to the philosophical problem of prediction, simplicity, and reliable belief revision by (Kelly, 1996). An overview of mathematical and computational results on the topic can be found in (Zeugmann and Zilles, 2008).

An Example: Eleusis Eleusis in an inductive inference card game. I ask: 'What is the rule behind this sequence of cards?'. Imagine I'm revealing the following sequence one-by-one and each time you see a new card you conjecture a rule:

A♠ Q♠ 3♠ A♠ Q♠ 4♡ . . .

First note that the problem of learning in Eleusis is of a different nature than that of the first example. Here, the conjecture might depend on the exact place of a card in the sequence, consider, e.g., the rule: 'Aces in even places'. Hence, for the formal treatment of the game the function learning paradigm would be more fitting.

Assume we have at our disposal unlimited amount of playing cards. How many different abstract scenarios? Quite many. Let us try to estimate this multitude. How many different playing cards do we have? How many different beginnings of infinite streams of length 1? How many different beginnings of length 2? Finally, how many different infinite sequences? Of course, infinitely many. Actually, uncountably many. To see why, assume, towards contradiction, that there countably many such sequences. Then they can be listed in the following way, each numbered a natural number.

1.	A♠	A♠	A♠	A♠	A♠	A♠	. . .		
2.	A♣	A♣	A♣	A♣	A♣	A♣	. . .		
3.	A♡	A♡	A♡	A♡	A♡	A♡	. . .		
4.	A♢	Q♠	3♠	8♡	2♡	5♠	. . .		
5.	A♠	Q♠	7♠	J♠	5♠	5♠	. . .		
. . .									
m.	A♣	A♡	A♣	A♡	A♣	A♢	. . .	m-th	
. . .									
									. . .

We will show that there must be an infinite sequence of cards that is not in this enumeration, effectively proving that such infinite sequences cannot be numbered with $\mathbb{N}$. We will construct the new, missing sequence, by making its nth element different than the nth element of nth sequence in the list. Note, that in the following enumeration some cards have been colored blue, just replace the blue element with any different card.

If we take the Eleusis problem to be that of exact prediction of the sequence, thinking of conjectures as exactly those infinite streams seems fitting. But then we will not be able to learn, since we will not be able to use the list of all conjectures in our background bookkeeping. This shows the importance of the hypothesis space in learning problems. Assume that in the game we require that the rule must be written in natural language on a piece of paper; or expressed by a natural language sentence; or in the extreme case, should be expressed in language but with a text no longer than 300 pages book. Finally, we could also require that the rule is encoded by a TM program. In all of those cases the descriptions are finite and there are countably many of them.

It must be that those rules cluster infinite streams in an adequate way. How many sequences comply to the rule: 'The sequence has solely A♠-cards.' How about: 'The sequence has solely ♠-cards' and 'The sequence has ♡-cards on even places'. Finally, note that the rules might be specified by the following prescription: 'The sequence is definable in first-order logic'.

Consider the following hypothesis spaces. As our learning setting requires, in each case one of the hypotheses is true, you are presented in a step-wise manner a stream consistent with this hypothesis. Each time you see a new datum you can change your conjecture. How much can you get to know and how quickly can you know it?

1. {(all cards are ♠), (all cards are ♢)}

2. $\{(\spadesuit$ at the 4-th position$), \neg(\spadesuit$ at the 4-th position$)\}$

3. $\{($exactly n cards are $\heartsuit) \mid n \in \mathbb{N}\}$

4. $\{($exactly n cards are $\heartsuit) \mid n \in \mathbb{N}\} \cup \{(\infty$ cards are $\heartsuit)\}$

As you probably noticed, case 3. is the one that is properly identifiable in the limit. Case 4. leads to trouble similar to those that the introduction of $\mathbb{N}$ led to in the initial example.

In Eleusis the rules are always expressible linguistically. What will happen if we require the hypotheses being expressible in first-order logic?

2.3 Model-theoretic Learning

Model-theoretic Learning (Martin and Osherson, 1997, 1998) is also known under the name of *First-Order Framework of Inquiry* and it constitutes a step towards a more logical treatment of the learning problem. In this it follows the line of thinking of the learning problem as a higher type of decidability, i.e., decidability in the limit (see Gold, 1965). The specification of the paradigm is given as follows:

1. Possible realities: models of a given signature.
2. Hypotheses: first order sentences.
3. Information accessible to the learner: sequences of atomic formulas and negations thereof.
4. Learner: function that takes a sequence and outputs a hypothesis.
5. Success criterion: after finite number of outputs the answers stabilize on a correct answer.

An Example: Learning about Orders Assume that our possible worlds are orders on $\mathbb{N}$, and that our language contains only the single binary predicate $<$ (and $=$). We will decide the following hypothesis space $\mathbb{P} = \{P_0, P_1\}$, where P_0 is a collection of strict total orders with a least point, and P_1 is a collection of strict total orders with a greatest point. The elements are given temporary names, via some variable assignment. Now you are given a sequence of clues concerning the order in question, for instance:

$$v_3{\neq}v_4,\ \neg v_0{<}v_0,\ v_1{<}v_9,\ v_{11}{=}v_{11},\ v_0{\neq}v_3, \ldots$$

The data consist of atomic propositions and their negations (e.g., of basic formulas). Your task is to solve the problem, i.e., to decide, possibly in the limit, $\mathbb{P}$ (the example comes forom Martin and Osherson, 1998).

2.4 Learning in Epistemic Spaces

Finally, the most recent setting for learning, gives a very general perspective on learning. It gives the advantage of thinking of learning in terms of belief-revision and possible world semantics.

1. Possible realities: possible worlds.
2. Hypotheses: sets of possible worlds.
3. Information accessible to the learner: sequences of propositions.

4. Learner: function that takes a sequence and outputs a proposition.
5. Success criterion: after finite number of outputs the answers stabilize on a proposition that is a singleton of the actual world.

For more details on this framework (see Gierasimczuk, 2010; Baltag et al., 2011, 2014b, 2019b; Gierasimczuk, 2023), the setting also leads to interesting general topological connection (Baltag et al., 2014a) and to a dynamic modal logic of learning theory (Baltag et al., 2019a).

2.5 Additional Notes on Paradigm Specification

Note that hypotheses are systematic descriptions of possible realities, they captured by what is sometimes called 'naming systems'. The hypotheses are finite descriptions of sets, e.g., Turing machines, grammars, natural numbers, logical formulas.

It is quite important to remember that in the interesting cases the data available at a given step presents only partial information about a possible reality. The character of data is determined by the setting, e.g., in language learning one might consider only positive or positive and negative information about a possible reality. In the basic settings data presented to the learner is arbitrary (conforming to the reality chosen by 'nature' in the beginning of the game), in some paradigms the learner can request particular information.

Identifiability in the limit is only one of many possible success criteria. Finite identifiability (see Mukouchi, 1992; Lange and Zeugmann, 1992; Gierasimczuk and de Jongh, 2013) requires that the learner in finite time arrives at complete *certainty* about the identity of the real world. Another condition is gradual identifiability—here after some time the learner starts giving answers which, according to some measure, keep getting closer and closer to the correct answer. While it is important to know about and appreciate the other possibilities, in this course we will focus on identifiability in the limit—we will require of the learner that after a finite time her answers will stabilize on a correct answer.

3 Epistemic Logic: validity arguments

Proposition 1. $\models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$

Proof. We need to show that the above formula is valid, i.e., is true in every possible-world model. In order to do that we take:

1. an arbitrary possible-worlds model M over n agents, $M = (S, \mathcal{K}_1, \ldots, \mathcal{K}_n, \pi)$,
2. an arbitrary $i \in \{1, \ldots, n\}$,
3. an arbitrary $s \in S$.

We need to show that $M, s \models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$.

There are two cases, the antecedent formula $K_i\varphi \wedge K_i(\varphi \rightarrow \psi)$ is either true or false in model M at the state s.

First consider that it is false, i.e., $M, s \not\models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi))$, in which case indeed $M, s \models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$ (by the semantics of $\rightarrow$).

Now consider that it is true, i.e., $M, s \models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi))$. Then $M, s \models K_i\varphi$ and $M, s \models K_i(\varphi \rightarrow \psi)$ (by the semantics of $\wedge$). Then, for all $t \in S$, such that $(s, t) \in \mathcal{K}_i$, $M, t \models \varphi$ and $M, t \models \varphi \rightarrow \psi$ (by the semantics of K). Therefore, for all $t \in S$, such that $(s, t) \in \mathcal{K}_i$, $M, t \models \psi$ (by the semantics of $\rightarrow$). So $M, s \models K_i\psi$ (by the semantics of K), and hence that $M, s \models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$. Since M, i, and s were chosen arbitrarily, we conclude that $\models (K_i\varphi \wedge K_i(\varphi \rightarrow \psi)) \rightarrow K_i\psi$. □

Proposition 2. *For all models M, if $M \models \varphi$, then $M \models K_i\varphi$.*

Proof. Let us take an arbitrary model M. To prove this proposition we need to only concern ourselves with models M such that $M \models \varphi$ (if $M \not\models \varphi$ the proposition is true, because it is an implication). Assume then that $M \models \varphi$, and so for all $s \in S$, $M, s \models \varphi$. In particular, for any fixed state $s \in S$, we get that $M, t \models \varphi$ at all $t \in S$, such that $(s, t) \in \mathcal{K}_i$. Hence $M, s \models K_i\varphi$, and, since s was chosen arbitrarily, $M \models K_i\varphi$. □

Proposition 3. $\models K_i\varphi \rightarrow \varphi$ *in the class of $S5$ models (i.e., models with equivalence accessibility relation).*

Proof. As in the proof of Proposition 1, we take an arbitrary model M, agent i and a state s in the model M. We assume that $M, s \models K_i\varphi$. Then $M, t \models \varphi$ for all t, such that $(s, t) \in \mathcal{K}_i$. Then, by the fact that $(s, s) \in \mathcal{K}_i$ (since $\mathcal{K}_i$ is **reflexive**), we obtain $M, s \models \varphi$. Since M and s were chosen arbitrarily, we conclude $\models K_i\varphi \rightarrow \varphi$. □

Proposition 4. $\models K_i\varphi \rightarrow K_iK_i\varphi$ *in the class of $S5$ models (i.e., models with equivalence accessibility relation).*

Proof. As before, we take an arbitrary M, i and s. Assume that $M, s \models K_i\varphi$. Consider any t such that $(s, t) \in \mathcal{K}_i$ and any u such that $(t, u) \in \mathcal{K}_i$. Since $\mathcal{K}_i$ is **transitive**, we have $(s, u) \in \mathcal{K}_i$. Since $M, s \models K_i\varphi$, we get $M, u \models \varphi$ (by the semantics of K). Thus, for all t such that $(s, t) \in \mathcal{K}_i$, we have $M, t \models K_i\varphi$ (by the semantics of K). Finally (again by the semantics of K) we obtain $M, s \models K_iK_i\varphi$, and so we conclude $\models K_i\varphi \rightarrow K_iK_i\varphi$. □

Proposition 5. $\models \neg K_i\varphi \rightarrow K_i\neg K_i\varphi$ *in the class of $S5$ models (i.e., models with equivalence accessibility relation).*

Proof. As before, we take an arbitrary M, i and s. Assume that $M, s \models \neg K_i\varphi$, then for some u, $M, u \models \neg\varphi$. Take any t, such that $(s, t) \in \mathcal{K}_i$. Since $\mathcal{K}_i$ is **symmetric** we have $(t, s) \in \mathcal{K}_i$ and since it is also **transitive** we get that $(t, u) \in \mathcal{K}_i$. Thus is follows that $M, t \models \neg K_i\varphi$ (by the semantics of K). Since this is true for all t, such that $(s, t) \in \mathcal{K}_i$ (again by the semantics of K) we obtain $M, s \models K_i\neg K_i\varphi$, and so we conclude $\models \neg K_i\varphi \to K_i\neg K_i\varphi$. □

4 Epistemic Logic: soundness and completeness

In this note we show that K_n is a sound and complete axiomatization with respect to $\mathcal{M}_n$ for the language $\mathcal{L}_n$ by first showing soundness (Theorem 4.1) and then completeness (Theorem 4.4).

4.1 Soundness of $\mathcal{K}_n$ wrt $\mathcal{M}_n$ for $\mathcal{L}_n$

Theorem 4.1 (Soundness)**.** *K_n is a sound axiomatization with respect to $\mathcal{M}_n$ for $\mathcal{L}_n$.*

Proof. To show that K_n is a sound axiomatization with respect to $\mathcal{M}_n$ for $\mathcal{L}_n$, we need to demonstrate that:

$$\text{for any formula } \varphi \in \mathcal{L}_n, \text{ if } \mathcal{K}_n \vdash \varphi, \text{ then } \mathcal{M}_n \models \varphi.$$

Let $\varphi \in \mathcal{L}_n$ and $\mathcal{K}_n \vdash \varphi$. That means that there exists a proof of φ in $\mathcal{K}_n$: a sequence of formulas $\varphi_0, \dots, \varphi_k$, with $\varphi_k = \varphi$, such that for all $\ell \in \{0, \dots, k\}$, φ_ℓ is either (a substitution of) an axiom of $\mathcal{K}_n$, or it is derived from previous steps of the proof by one of the rules of inference of $\mathcal{K}_n$. We need to show that then $\mathcal{M}_n \models \varphi$. We will prove that by induction on the length of the proof.

Base case: Let us assume that $k = 0$ (i.e., that the proof is of length 1), so $\varphi = \varphi_0$. Then it is only possible that φ is (a substitution of) an axiom of $\mathcal{K}_n$. There are two cases:

1. φ is (a substitution of) a propositional tautology. In this case φ is true under any propositional valuation, so it is true in all possible worlds of all possible-world models in $\mathcal{M}_n$, so φ is valid in $\mathcal{M}_n$.

2. φ is (a substitution of) the axiom $A2$:

$$(K_i\psi \wedge K_i(\psi \to \gamma)) \to K_i\gamma, \text{ where } i \in \{1, \dots, n\}.$$

   We need to show that the above formula is valid wrt $\mathcal{M}_n$, i.e., it is valid in every possible-world model, i.e., it is true in every possible world of every possible-world model. In order to do that we take an $M \in \mathcal{M}_n$, $M = (S, \mathcal{K}_1, \dots, \mathcal{K}_n, \pi)$, an arbitrary agent $i \in \{1, \dots, n\}$, and an arbitrary world $s \in S$.

   There are two cases: the antecedent formula $K_i\psi \wedge K_i(\psi \to \gamma)$ is either true or false in model M at the world s.

   First consider that it is false, i.e., $(M, s) \not\models (K_i\psi \wedge K_i(\psi \to \gamma))$. In that case indeed $(M, s) \models (K_i\psi \wedge K_i(\psi \to \gamma)) \to K_i\gamma$ (by the semantics of $\to$).

   Now consider that it is true, i.e., $(M, s) \models (K_i\psi \wedge K_i(\psi \to \gamma))$. Then $(M, s) \models K_i\psi$ and $(M, s) \models K_i(\psi \to \gamma)$ (by the semantics of $\wedge$). Then, for all $t \in S$, such that $(s, t) \in \mathcal{K}_i$, $(M, t) \models \psi$ and $(M, t) \models \psi \to \gamma$ (by the semantics of K). So, for all $t \in S$, such that $(s, t) \in \mathcal{K}_i$, $(M, t) \models \gamma$ (by the semantics of $\to$). That means that $(M, s) \models K_i\gamma$ (by the semantics of K), and hence that $(M, s) \models (K_i\psi \wedge K_i(\psi \to \gamma)) \to K_i\gamma$.

Since M, i, and s were chosen arbitrarily, we conclude that

$$\mathcal{M}_n \models (K_i\psi \wedge K_i(\psi \rightarrow \gamma)) \rightarrow K_i\gamma.$$

Induction step: Let us assume that $k > 0$, so $\varphi = \varphi_k$, in a proof $\varphi_0, \ldots, \varphi_k$ (of length $k+1$).

Induction hypothesis: Assume that for all $\ell < k$, $\mathcal{M}_n \models \varphi_\ell$.

We have to show that then φ $(= \varphi_k)$ is valid in $\mathcal{M}_n$. The arguments for the two cases where φ is (a substitution of) a propositional tautology or axiom $A2$ are identical to the ones in the base case. The two remaining cases are where φ was obtained by an application of one of the rules of $\mathcal{K}_n$ to the previous steps of the proof:

1. φ was obtained by applying $R1$ to some previous steps in the proof. That means that there is a $\gamma \in \mathcal{L}_n$ and $i, j < k$, such that $i \neq j$, $\varphi_i = \gamma$ and $\varphi_j = \gamma \rightarrow \varphi$. By the induction hypothesis, we know that then γ and $\gamma \rightarrow \varphi$ are valid in $\mathcal{M}_n$, i.e., true in all worlds in all possible-world models in $\mathcal{M}_n$. Then in all those worlds, also φ must be true, by the semantics of $\rightarrow$ in propositional logic. So, $\mathcal{M}_n \models \varphi$.

2. φ was obtained by applying $R2$ to some previous step in the proof. That means that there is a $\gamma \in \mathcal{L}_n$ and $j < k$, $\varphi_j = \gamma$ and $\varphi = K_i\gamma$, for some $i \in \{1, \ldots, n\}$. By the induction hypothesis we know that then γ is valid in $\mathcal{M}_n$, i.e., true in all worlds in all possible-world models in $\mathcal{M}_n$. Let us take an arbitrary model $M \in \mathcal{M}_n$ such that $M \models \gamma$, i.e., for all $s \in S$, $(M, s) \models \gamma$. In particular, for any fixed world $s \in S$, we get that $(M, t) \models \gamma$ at all $t \in S$, such that $(s, t) \in \mathcal{K}_i$. Hence $(M, s) \models K_i\gamma$, and, since s was chosen arbitrarily, $M \models K_i\gamma$. Because M was an arbitrary model in $\mathcal{M}_n$, we get that $\mathcal{M}_n \models \varphi$.

Thus, we have shown that for any formula $\varphi \in \mathcal{L}_n$, if $\mathcal{K}_n \vdash \varphi$, then $\mathcal{M}_n \models \varphi$. □

4.2 Completeness of $\mathcal{K}_n$ wrt $\mathcal{M}_n$ for $\mathcal{L}_n$

In order to show completeness we need to introduce several concepts and prove a lemma.

Definition 4.1. *Take an axiom system AX,*

1. *φ is AX-*consistent *if $\neg\varphi$ is not provable in AX.*
2. *A finite set $\{\varphi_1, \ldots, \varphi_k\}$ of formulas is AX-consistent if $\varphi_1 \wedge \ldots \wedge \varphi_k$ is AX-consistent.*
3. *An infinite set of formulas is AX-consistent if all of its finite subsets are AX-consistent.*

Definition 4.2. *A set F of formulas is a maximal AX-consistent set wrt a language $\mathcal{L}$ if:*

1. *it is AX-consistent, and*
2. *for all φ in $\mathcal{L}$ but not in F, the set $F \cup \{\varphi\}$ is not AX-consistent.*

Lemma 4.2 (Lindenbaum)**.** *Suppose the language $\mathcal{L}$ consists of a countable set of formulas and is closed wrt propositional connectives (so that if φ and ψ are in $\mathcal{L}$, then so are $\varphi \wedge \psi$ and $\neg\varphi$).*

In any axiom system AX that includes every instance of $A1$ and $R1$ for the language $\mathcal{L}$, every AX-consistent set $F \subseteq \mathcal{L}$ can be extended to a maximal AX-consistent set wrt $\mathcal{L}$.

Proof. Let F be an AX-consistent subset of formulas in $\mathcal{L}$. We will construct a sequence of AX-consistent sets:

$$F_0, F_1, F_2, \ldots$$

$\mathcal{L}$ is a countable language, we can enumerate its formulas:

$$\psi_1, \psi_2, \ldots$$

We define:

$$F_0 := F$$

$$F_{i+1} := \begin{cases} F_i \cup \{\psi_i\} & \text{if this set is } AX\text{-consistent} \\ F_i & \text{otherwise} \end{cases}$$

Each set in the sequence $F_0, F_1, \ldots$ is AX-consistent. F_i is a nondecreasing sequence of sets. We define F in the following way:

$$F := \bigcup_{i=0}^{\infty} F_i.$$

Note that each finite subset of F must be contained in F_j for some j, and thus must be AX-consistent (since F_j is AX-consistent). It follows that F itself is *AX-consistent.*

We claim that in fact F is a *maximal* AX-consistent set. To show this take any $\psi \in \mathcal{L}$ and $\psi \notin F$. Since ψ is a formula in $\mathcal{L}$, it must appear in our enumeration, say, as ψ_k. If $F_k \cup \{\psi_k\}$ were AX-consistent, then our construction would guarantee that $\psi_k \in F_{k+1}$, and hence that $\psi_k \in F$. Because $\psi_k = \psi \notin F$, it follows that $F_k \cup \{\psi\}$ is not AX-consistent. Hence $F \cup \{\psi\}$ is also not AX-consistent.

It follows that F is a maximal AX-consistent set.

□

Lemma 4.3. *If F is a maximal AX-consistent set, then it satisfies the following properties:*

1. *for every formula $\varphi \in \mathcal{L}$, exactly one of φ and $\neg\varphi$ is in F;*
2. *$\varphi \wedge \psi \in F$ iff $\varphi \in F$ and $\psi \in F$;*
3. *if φ and $\varphi \to \psi$ are both in F, then ψ is in F;*
4. *if φ is provable in AX, then $\varphi \in F$.*

Proof. (of 1.)

Let F be a maximal AX-consistent set, and let $\varphi \in \mathcal{L}$. We show that one of $F \cup \{\varphi\}$ and $F \cup \{\neg\varphi\}$ is AX-consistent. For assume to the contrary that neither of them is AX-consistent. It is not hard to see that $F \cup \{\varphi \vee \neg\varphi\}$ is then also not AX-consistent. So F is not AX-consistent, because $\varphi \vee \neg\varphi$ is a propositional tautology. This gives a contradiction.

If $F \cup \{\varphi\}$ is AX-consistent, then we must have $\varphi \in F$ since F is a maximal AX-consistent set. Similarly, if $F \cup \{\neg\varphi\}$ is AX-consistent then $\neg\varphi \in F$. Thus, one of φ or $\neg\varphi$ is in F.

Moreover, we cannot have both φ and $\neg\varphi$ in F, for otherwise F would not be AX-consistent. □

Theorem 4.4. *K_n is a complete axiomatization with respect to $\mathcal{M}_n$ for $\mathcal{L}_n$.*

Proof. We want to show that:

$$\text{for every formula } \varphi \in \mathcal{L}_n, \text{ if } \mathcal{M}_n \models \varphi, \text{ then } K_n \vdash \varphi.$$

It suffices to show that:

$$\text{every } K_n\text{-consistent formula in } \mathcal{L}_n \text{ is satisfiable with respect to } \mathcal{M}_n, \text{ (*)}$$

because if we knew that (*) is true we would get the theorem in the following way. Assume that $\mathcal{M}_n \models \varphi$. Assume for contraction that it is not the case that $\mathcal{K}_n \vdash \varphi$. Then it is also not the case that $\mathcal{K}_n \vdash \neg\neg\varphi$. This, by Definition 4.1, makes $\neg\varphi$ K_n-consistent. But then by (*) $\neg\varphi$ is satisfiable, so φ is not valid. We obtain contradiction.

So, indeed, we want to show (*). We will construct a special (so-called 'canonical') model $M^C \in \mathcal{M}_n$, whose worlds correspond to maximal K_n-consistent sets (denoted by V) in the following way.

Definition 4.3. *Given a set of formulas V, we define $V/K_i := \{\varphi \mid K_i\varphi \in V\}$.*

Example 1. *For example, if $V = \{K_1p, K_2K_1q, K_1K_3p \wedge q, K_1K_3q\}$, then $V/K_1 = \{p, K_3q\}$.*

Definition 4.4. *Let $M^C = (S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n)$, where:*

$$\begin{aligned} S &= \{s_V \mid V \textit{ is a maximal } K_n\textit{-consistent set}\} \\ \pi(s_V)(p) &= \begin{cases} 1 & \textit{if } p \in V \\ 0 & \textit{if } p \notin V \end{cases} \\ \mathcal{K}_i &= \{(s_V, s_W) \mid V/K_i \subseteq W\} \end{aligned}$$

Given that construction we want to show that:

$$(M^C, s_V) \models \varphi \text{ iff } \varphi \in V, \text{ (**)}$$

because if we knew that (**) is true, we would get (*) in the following way. By Lemma 4.2, if φ is K_n-consistent, then φ is contained in some maximal K_n-consistent set V. From (**) it follows that $(M^C, s_V) \models \varphi$, and so φ is satisfiable in M^C. Hence, φ is satisfiable wrt $\mathcal{M}_n$.

Now we show that: $(M^C, s_V) \models \varphi$ iff $\varphi \in V$ by induction on the structure of φ.

Base case: φ is a primitive proposition p. Then $(M^C, s_V) \models p$ iff $p \in V$, by π.

Inductive hypothesis: Assume that for $\gamma := \psi, \psi_1, \psi_2$ we have that $(M^C, S_V) \models \gamma$ iff $\gamma \in V$.

We have the following cases:

1. $\varphi := \neg\psi$. Then we have: $(M^C, s_V) \models \varphi$ iff $(M^C, s_V) \models \neg\psi$ iff it is not the case that $(M^C, s_V) \models \psi$ iff it is not the case that $\psi \in V$ iff $\neg\psi \in V$ iff $\varphi \in V$.

2. $\varphi := \psi_1 \wedge \psi_2$. The argument here is analogous to the one above.

3. $\varphi := K_i\psi$. We need to show that $(M^C, S_V) \models \varphi$ iff $\varphi \in V$, we will show the two directions separately.

   ($\leftarrow$) Assume that $\varphi \in V$. Then $\psi \in V/K_i$ and, by definition of $\mathcal{K}_i$, if $(s_V, s_W) \in \mathcal{K}_i$, then $\psi \in W$. Thus, using the induction hypothesis, $(M^C, s_W) \models \psi$ for all W such that $(s_V, s_W) \in \mathcal{K}_i$. By the definition of $\models$, it follows that $(M^C, s_V) \models K_i\psi$.

($\rightarrow$) Assume $(M^C, s_V) \models K_i\psi$. It follows that the set $(V/K_i) \cup \{\neg\psi\}$ is not K_n-consistent. So, there must be some finite subset, say $\{\varphi_1, \ldots, \varphi_k, \neg\psi\}$, which is not K_n-consistent. We have:

$K_n \vdash \varphi_1 \rightarrow \varphi_2 \rightarrow (\ldots \rightarrow (\varphi_k \rightarrow \psi) \ldots)$ (propositional logic)

$K_n \vdash K_i(\varphi_1 \rightarrow (\varphi_2 \rightarrow (\ldots \rightarrow (\varphi_k \rightarrow \psi) \ldots)))$ (R2)

$K_n \vdash K_i(\varphi_1 \rightarrow (\varphi_2 \rightarrow (\ldots \rightarrow (\varphi_k \rightarrow \psi) \ldots))) \rightarrow$

$K_i\varphi_1 \rightarrow (K_i\varphi_2 \rightarrow (\ldots \rightarrow (K_i\varphi_k \rightarrow K_i\psi) \ldots))$ (induction on k, $A2$, propositional logic)

$K_n \vdash K_i\varphi_1 \rightarrow K_i\varphi_2 \rightarrow (\ldots \rightarrow (K_i\varphi_k \rightarrow K_i\psi) \ldots)$ (R1)

$K_i\varphi_1 \rightarrow K_i\varphi_2 \rightarrow (\ldots \rightarrow (K_i\varphi_k \rightarrow K_i\psi) \ldots) \in V$ (Lemma 4.3)

Because $\varphi_1, \ldots, \varphi_k \in V/K_i$, we have $K_i\varphi_1, \ldots, K_i\varphi_k \in V$.

By Lemma 4.3, we have $K_i\psi \in V$.

□

References

Baltag, A., Gierasimczuk, N., Özgün, A., Vargas Sandoval, A. L., and Smets, S. (2019a). A dynamic logic for learning theory. *Journal of Logical and Algebraic Methods in Programming*, 109:100485.

Baltag, A., Gierasimczuk, N., and Smets, S. (2011). Belief revision as a truth-tracking process. In Apt, K., 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.

Baltag, A., Gierasimczuk, N., and Smets, S. (2014a). Epistemic topology. In progress.

Baltag, A., Gierasimczuk, N., and Smets, S. (2014b). Truth-tracking by belief-revision. to appear in Studia Logica 2017.

Baltag, A., Gierasimczuk, N., and Smets, S. (2019b). Truth-tracking by belief revision. *Studia Logica*, 107(5):917–947.

Gierasimczuk, N. (2010). *Knowing One's Limits. Logical Analysis of Inductive Inference.* PhD thesis, Universiteit van Amsterdam, The Netherlands.

Gierasimczuk, N. (2023). Inductive inference and epistemic modal logic (invited talk). In Klin, B. and Pimentel, E., editors, *31st EACSL Annual Conference on Computer Science Logic, CSL 2023, February 13-16, 2023, Warsaw, Poland*, volume 252 of *LIPIcs*, pages 2:1–2:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Gierasimczuk, N. and de Jongh, D. (2013). On the complexity of conclusive update. *The Computer Journal*, 56(3):365–377.

Gold, E. M. (1965). Limiting recursion. *Journal of Symbolic Logic*, 30(1):28–48.

Gold, E. M. (1967). Language identification in the limit. *Information and Control*, 10:447–474.

Kelly, K. T. (1996). *The Logic of Reliable Inquiry.* Oxford University Press, Oxford.

Lange, S. and Zeugmann, T. (1992). Types of monotonic language learning and their characterization. In *COLT*, pages 377–390.

Lange, S., Zeugmann, T., and Zilles, S. (2008). Learning indexed families of recursive languages from positive data: A survey. *Theor. Comput. Sci.*, 397(1-3):194–232.

Martin, E. and Osherson, D. (1997). Scientific discovery based on belief revision. *The Journal of Symbolic Logic*, 62(4):1352–1370.

Martin, E. and Osherson, D. (1998). *Elements of Scientific Inquiry.* MIT Press, Cambridge, MA, USA.

Mukouchi, Y. (1992). Characterization of finite identification. In Jantke, K., 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.

Osherson, D., Stob, M., and Weinstein, S. (1986). *Systems that Learn.* M.I.T. Press, Cambridge, MA, USA.

Putnam, H. (1965). Trial and error predicates and the solution to a problem of Mostowski. *The Journal of Symbolic Logic*, 30(1):49–57.

Solomonoff, R. (1964). A formal theory of inductive inference. Part I. *Information and Control*, 7(1):1–22.

Zeugmann, T. and Lange, S. (1995). A guided tour across the boundaries of learning recursive languages. In *GOSLER Final Report*, pages 190–258.

Zeugmann, T. and Zilles, S. (2008). Learning recursive functions: A survey. *Theor. Comput. Sci.*, 397(1-3):4–56.