Learning through Belief Revision

Nina Gierasimczuk

DTU Compute
Technical University of Denmark

Knowledge Representation Conventicle
PRICAI 2019
Fiji, August 25th, 2019

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

Belief revision and learning

We investigate order-driven belief revision methods:

- in particular, some aspects of their long term behaviour:
- how good are they as learning methods?
- under what conditions are they reliable?

We then use the same apparatus for a more general purpose:

- topological characterisation of learnability and solvability;
- reaching out to obtaining an adequate logic of learnability.

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

Epistemic Spaces and Observables

Definition
An **epistemic space** is a pair $\mathbb{S} = (S, \mathcal{O})$ consisting of a countable state space S and a countable set of observables $\mathcal{O} \subseteq \mathcal{P}(S)$.

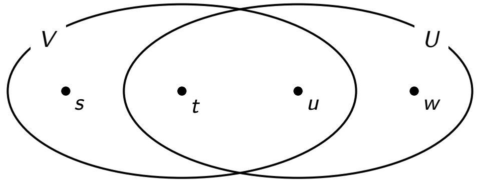

Learning: Streams of Observables

Definition
Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space.

- A **data stream** is an infinite sequence $\vec{O} = (O_0, O_1, \ldots)$ of data from $\mathcal{O}$.
- A **data sequence** is a finite initial segment of an $\vec{O}$;
  such a finite sequence of length $n-1$ is denoted by $\sigma = (\sigma_0, \ldots, \sigma_n)$.

Definition
Take $\mathbb{S} = (S, \mathcal{O})$ and $s \in S$. A data stream $\vec{O}$ is:

- **sound with respect to** s iff every element listed in $\vec{O}$ is true in s.
- **complete with respect to** s iff every observable true in s is listed in $\vec{O}$.

We assume that data streams are sound and complete.

Learning: Learners and Conjectures

Definition
Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space and let $\sigma_0, \ldots, \sigma_n \in \mathcal{O}$.
A **learner** is a function L that on the input of $\mathbb{S}$ and data sequence $(\sigma_0, \ldots, \sigma_n)$ outputs some set of worlds $L(\mathbb{S}, (\sigma_0, \ldots, \sigma_n)) \subseteq S$, called a **conjecture**.

Definition
$\mathbb{S} = (S, \mathcal{O})$ is **learnable by** L if

Learning: Learners and Conjectures

Definition
Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space and let $\sigma_0, \ldots, \sigma_n \in \mathcal{O}$.
A **learner** is a function L that on the input of $\mathbb{S}$ and data sequence $(\sigma_0, \ldots, \sigma_n)$ outputs some set of worlds $L(\mathbb{S}, (\sigma_0, \ldots, \sigma_n)) \subseteq S$, called a **conjecture**.

Definition
$\mathbb{S} = (S, \mathcal{O})$ is **learnable by** L if for every state $s \in S$

Learning: Learners and Conjectures

Definition
Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space and let $\sigma_0, \ldots, \sigma_n \in \mathcal{O}$.
A **learner** is a function L that on the input of $\mathbb{S}$ and data sequence $(\sigma_0, \ldots, \sigma_n)$ outputs some set of worlds $L(\mathbb{S}, (\sigma_0, \ldots, \sigma_n)) \subseteq S$, called a **conjecture**.

Definition
$\mathbb{S} = (S, \mathcal{O})$ is **learnable by** L if for every state $s \in S$
and for every sound and complete data stream $\vec{O}$ for s,

Learning: Learners and Conjectures

Definition
Let $\mathbb{S} = (S, \mathcal{O})$ be an epistemic space and let $\sigma_0, \ldots, \sigma_n \in \mathcal{O}$.
A **learner** is a function L that on the input of $\mathbb{S}$ and data sequence $(\sigma_0, \ldots, \sigma_n)$ outputs some set of worlds $L(\mathbb{S}, (\sigma_0, \ldots, \sigma_n)) \subseteq S$, called a **conjecture**.

Definition
$\mathbb{S} = (S, \mathcal{O})$ is **learnable by** L if for every state $s \in S$
and for every sound and complete data stream $\vec{O}$ for s,
there is $n \in \mathbb{N}$ such that:

$$L(\mathbb{S}, (O_0, \ldots, O_k)) = \{s\} \text{ for all } k \geq n.$$

An epistemic space $\mathbb{S}$ is **learnable** if it is learnable by a learner L.

Example of a Learnable Space

Let $\mathbb{S} = (S, \mathcal{O})$ such that $S = \{s_n \mid n \in \mathbb{N}\}$, $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k = \{s_i \mid 0 \leq i \leq k\}$. $\mathbb{S}$ is learnable.

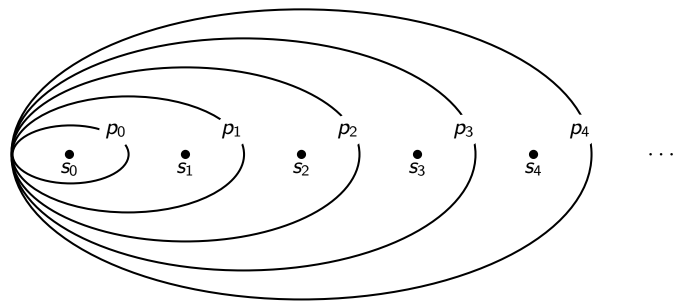

Example of a Learnable Space

Let $\mathbb{S} = (S, \mathcal{O})$ such that $S = \{s_n \mid n \in \mathbb{N}\}$, $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k = \{s_i \mid 0 \leq i \leq k\}$. $\mathbb{S}$ is learnable.

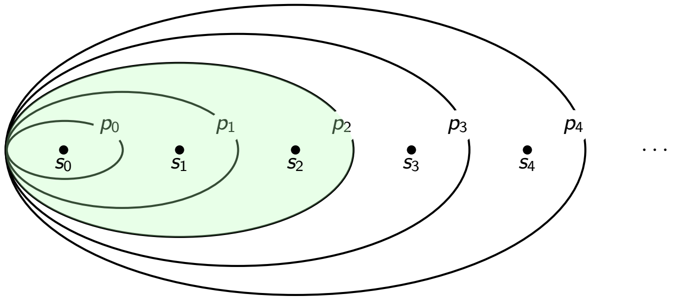

Example of a Learnable Space

Let $\mathbb{S} = (S, \mathcal{O})$ such that $S = \{s_n \mid n \in \mathbb{N}\}$, $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k = \{s_i \mid 0 \leq i \leq k\}$. $\mathbb{S}$ is learnable.

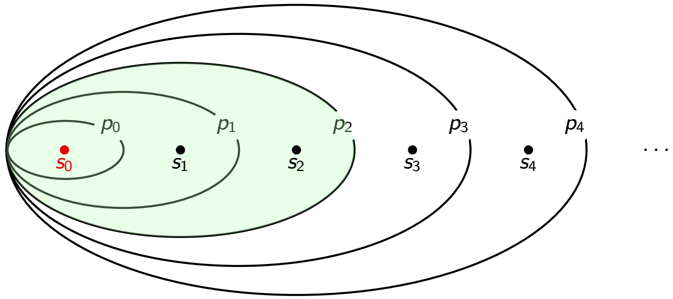

Example of a Learnable Space

Let $\mathbb{S} = (S, \mathcal{O})$ such that $S = \{s_n \mid n \in \mathbb{N}\}$, $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k = \{s_i \mid 0 \leq i \leq k\}$. $\mathbb{S}$ is learnable.

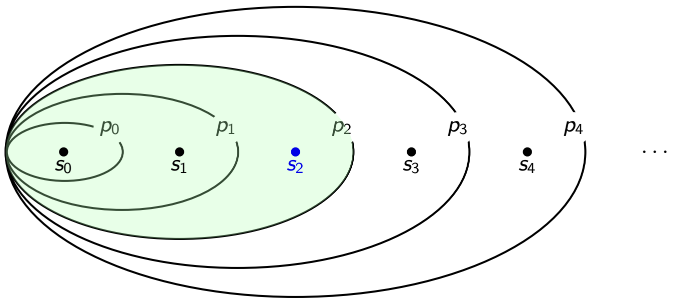

Example of a Non-Learnable Space

Consider $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \mathbb{N}\} \cup \{s_\infty\}$, and $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k := \{s_k, s_{k+1}, \ldots\} \cup \{s_\infty\}$. $\mathbb{S}$ is not learnable.

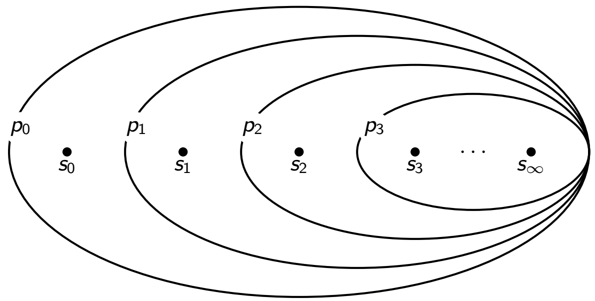

Example of a Non-Learnable Space

Consider $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \mathbb{N}\} \cup \{s_\infty\}$, and $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k := \{s_k, s_{k+1}, \ldots\} \cup \{s_\infty\}$. $\mathbb{S}$ is not learnable.

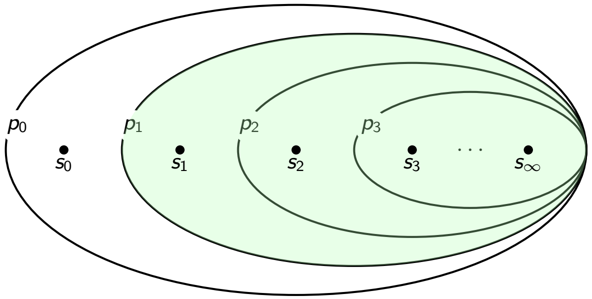

Example of a Non-Learnable Space

Consider $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \mathbb{N}\} \cup \{s_\infty\}$, and $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k := \{s_k, s_{k+1}, \ldots\} \cup \{s_\infty\}$. $\mathbb{S}$ is not learnable.

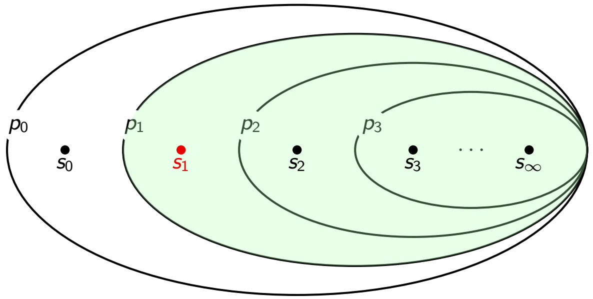

Example of a Non-Learnable Space

Consider $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \mathbb{N}\} \cup \{s_\infty\}$, and $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k := \{s_k, s_{k+1}, \ldots\} \cup \{s_\infty\}$. $\mathbb{S}$ is not learnable.

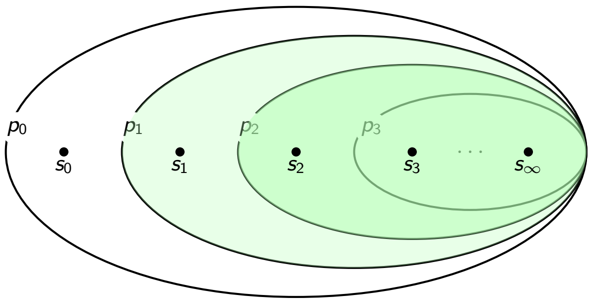

Example of a Non-Learnable Space

Consider $\mathbb{S} = (S, \mathcal{O})$, where $S := \{s_n \mid n \in \mathbb{N}\} \cup \{s_\infty\}$, and $\mathcal{O} = \{p_i \mid i \in \mathbb{N}\}$, and for any $k \in \mathbb{N}$, $p_k := \{s_k, s_{k+1}, \ldots\} \cup \{s_\infty\}$. $\mathbb{S}$ is not learnable.

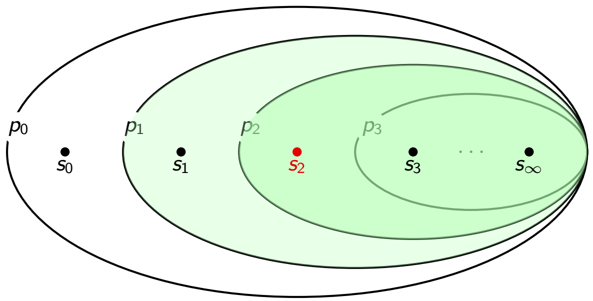

Questions, Answers, and Problems

Definition

A **question** $\mathcal{Q}$ is a partition of S, whose cells A_i are called **answers to** $\mathcal{Q}$. Given $s \in A \subseteq S$, $A \in \mathcal{Q}$ is called **the answer to** $\mathcal{Q}$ **at** s, denoted A_s.

Definition

$\mathcal{Q}'$ is a **refinement** of $\mathcal{Q}$ if all answers of $\mathcal{Q}$ is a disjoint union of answers of $\mathcal{Q}'$.

Definition

A **problem** $\mathbb{P}$ is a pair $(\mathbb{S}, \mathcal{Q})$ consisting of $\mathbb{S} = (S, \mathcal{O})$ and $\mathcal{Q}$ over S. $\mathbb{P}' = (\mathbb{S}, \mathcal{Q}')$ is a **refinement** of $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ if $\mathcal{Q}'$ is a refinement of $\mathcal{Q}$.

Illustration

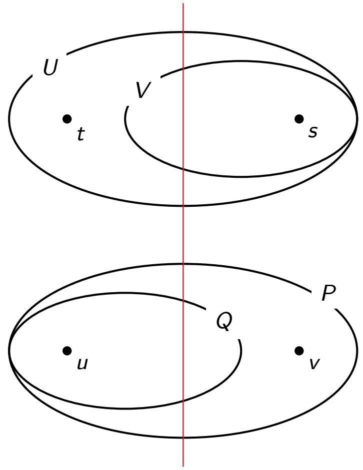

Illustration

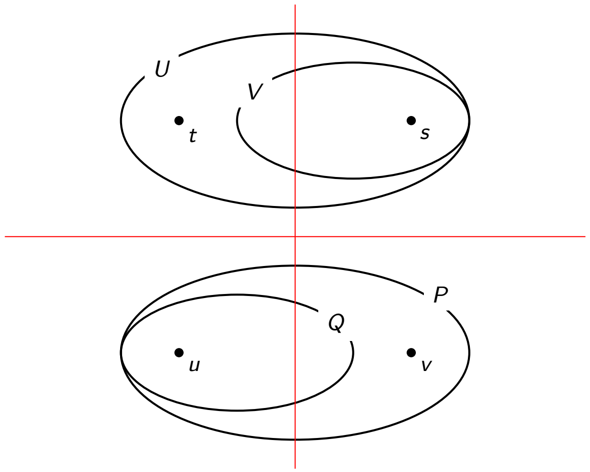

Solving in the Limit

Definition

A learning method L **solves a problem** $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ in the limit iff for every state $s \in S$ and every data stream $\vec{O}$ for s, there is an $n \in \mathbb{N}$ such that:

$$L(\mathbb{S}, (O_0, \dots, O_k)) \subseteq A_s \text{ for all } k \geq n.$$

A problem is **solvable in the limit** if there is a learner that solves it in the limit.

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

General Topology

Definition

A topology τ over a set S is a collection of subsets of S (open sets) s.t.:

1. $\emptyset \in \tau$,
2. $S \in \tau$,
3. for any $X \subseteq \tau$, $\bigcup X \in \tau$, and
4. for any finite $X \subseteq \tau$ we have $\bigcap X \in \tau$.

Separability by observations: Illustration

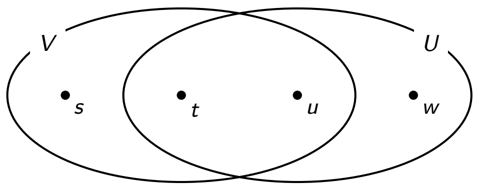

(A) t and u are not separable

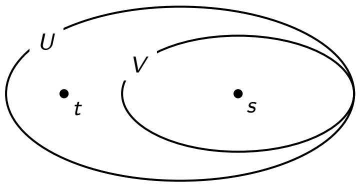

(B) weakly separated space, $T0$

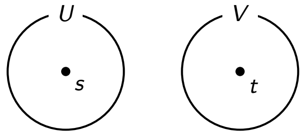

(C) strongly separated space, $T1$

Locally Closed and Constructible Sets

Definition

A topology τ is T_d iff for every $s \in S$ there is a $U \in \tau$ such that $U \setminus \{s\} \in \tau$, i.e., for every $s \in S$ there is a $U \in \tau$ such that $\{s\} = U \cap \overline{\{s\}}$.
T_d is a separation property between $T0$ and $T1$.

Definition

A set A is **locally closed** if $A = U \cap C$, where U is open and C is closed.

A set is **constructible** if it is a finite disjoint union of locally closet sets.

An ω-**constructible set** is a countable union of locally closed sets.

The Topology Associated with an Epistemic Space

Definition
The topology $\tau_{\mathbb{S}}$ associated with an epistemic space $\mathbb{S} = (S, \mathcal{O})$ is a collection of subsets of S of the following properties:

1. for any $O \in \mathcal{O}$ it is the case that $O \in \tau_{\mathbb{S}}$
2. $\emptyset \in \tau_{\mathbb{S}}$,
3. $S \in \tau_{\mathbb{S}}$,
4. for any $U \subseteq \tau_{\mathbb{S}}$, $\bigcup U \in \tau_{\mathbb{S}}$, and
5. for any $x, y \in \tau_{\mathbb{S}}$ we have $x \cap y \in \tau_{\mathbb{S}}$.

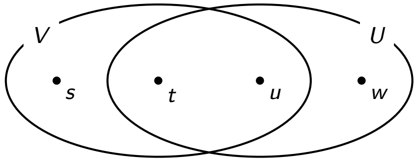

The Topology Associated with an Epistemic Space

Definition
The topology $\tau_{\mathbb{S}}$ associated with an epistemic space $\mathbb{S} = (S, \mathcal{O})$ is a collection of subsets of S of the following properties:

1. for any $O \in \mathcal{O}$ it is the case that $O \in \tau_{\mathbb{S}}$
2. $\emptyset \in \tau_{\mathbb{S}}$,
3. $S \in \tau_{\mathbb{S}}$,
4. for any $U \subseteq \tau_{\mathbb{S}}$, $\bigcup U \in \tau_{\mathbb{S}}$, and
5. for any $x, y \in \tau_{\mathbb{S}}$ we have $x \cap y \in \tau_{\mathbb{S}}$.

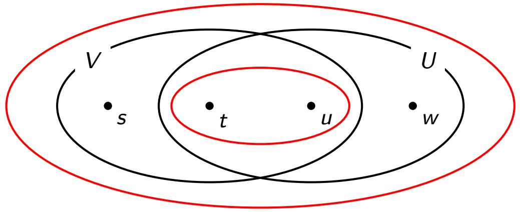

Characterization of Solvability in the Limit

Theorem
A problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ is solvable in the limit iff $\mathcal{Q}$ has a locally closed refinement.

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

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

Background

- Learning and belief revision go their separate ways,
- conjecture dynamics is a common theme.
- What are the principles of this dynamics?

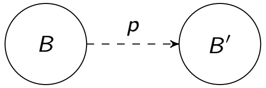

Background

- Learning and belief revision go their separate ways,
- conjecture dynamics is a common theme.
- What are the principles of this dynamics?

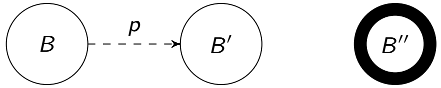

Truth-tracking!

Background

- Learning and belief revision go their separate ways,
- conjecture dynamics is a common theme.
- What are the principles of this dynamics?

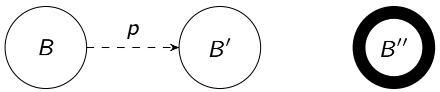

Truth-tracking!

Order-driven learning: Motivation

- Belief Revision: minimal states give beliefs.
- Computational Learning Theory: co-learning, learning by erasing.
- Philosophy of Science: Ockham's razor.

Plausibility Spaces

A **plausibility space**, $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, consists of an epistemic space $\mathbb{S} = (S, \mathcal{O})$ and a plausibility preorder $\preceq \subseteq S \times S$.

Plausibility Spaces

A **plausibility space**, $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$, consists of an epistemic space $\mathbb{S} = (S, \mathcal{O})$ and a plausibility preorder $\preceq \subseteq S \times S$.

Knowledge and Belief

$$\mathbb{B}_{\mathbb{S}} \models Kp \quad \text{iff} \quad S \subseteq p$$
$$\mathbb{B}_{\mathbb{S}} \models Bp \quad \text{iff} \quad \exists w \; \forall u \leq w \;\; u \in p$$

The latter, in well-founded spaces is equivalent to the standard definition:

$$\mathbb{B}_{\mathbb{S}} \models Bp \text{ iff } min_{\preceq} S \subseteq p.$$

Belief-Revision Methods

Definition

A **belief-revision method** is a function R that, for any plausibility space $\mathbb{B}_{\mathbb{S}} = (S, \mathcal{O}, \preceq)$ and any data sequence σ, outputs a new plausibility space:

$$R(\mathbb{B}_{\mathbb{S}}, \sigma) := (S^{\sigma}, \mathcal{O}, \preceq^{\sigma}).$$

DEL belief revision methods: Conditioning on p

- **Conditioning** eliminates all worlds of S that do not satisfy p.

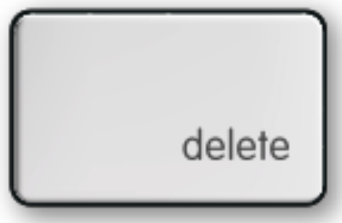

DEL belief revision methods: Upgrade with p

- **Lexicographic upgrade** rearranges the preorder by putting all worlds satisfying p to be more plausible than others.

DEL belief revision methods: Upgrade with p

- **Minimal upgrade** rearranges the preorder by making only the most plausible states satisfying p more plausible than all others, leaving the rest of the preorder the same.

The (Learning) Problem

Game Theory, Artificial Intelligence, Epistemic Logic, Belief Revision

equip agents with methods that allow them to change their beliefs and conjectures on the basis of new information.

How good are those strategies in the long run?
With respect to what criteria should they be evaluated?

Learning via Belief Revision

Definition
Every belief-revision method R, together with a prior plausibility $\preceq$, generates in a canonical way a learning method $L_R^{\preceq}$, called a **belief-revision-based learning method**, and given by:

$$L_R^{\preceq}(\mathbb{S}, \sigma) := \min_{\preceq} R((\mathbb{S}, \preceq), \sigma).$$

Definition
An epistemic space $\mathbb{S}$ is **identified in the limit by a belief-revision method** R if there exists a prior plausibility assignment $\preceq$ such that the generated learning method $L_R^{\preceq}$ identifies S in the limit.

Universality Results

Definition

A learning method L said to be **universal** if it can identify in the limit every epistemic space that is identifiable in the limit.

	Conditioning	Lexicographic	Minimal
Positive Streams	YES	YES	NO

Universality Results

Definition

A learning method L said to be **universal** if it can identify in the limit every epistemic space that is identifiable in the limit.

	Conditioning	Lexicographic	Minimal
Positive Streams	YES	YES	NO

Theorem

Minimal revision is not universal.

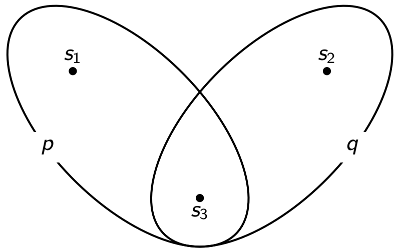

Universality Results

Definition

A learning method L said to be **universal** if it can identify in the limit every epistemic space that is identifiable in the limit.

	Conditioning	Lexicographic	Minimal
Positive Streams	YES	YES	NO

Theorem

There is no universal belief-revision method for well-founded preorder.

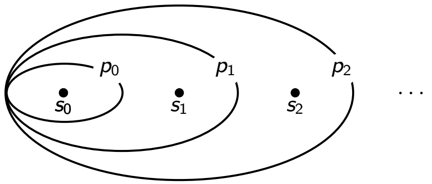

Detour: Direct Solvability by Conditioning

Definition

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

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

A problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ is **directly solvable by conditioning** if it is solvable by conditioning with respect to a total order $\trianglelefteq \subseteq \mathcal{Q} \times \mathcal{Q}$.

Ordering the States v. Ordering the Answers

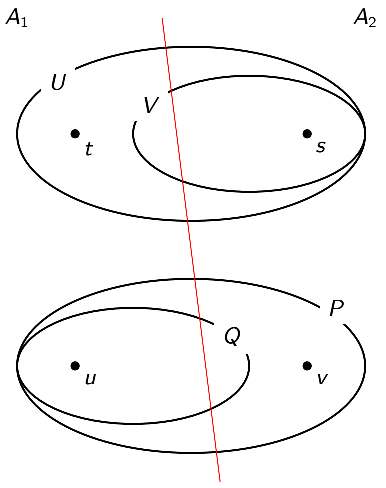

Linear Separation

Definition

A partition $\mathcal{Q}$ of a topological space (S, τ) is **linearly separated** if there exists some **total order** $\trianglelefteq$ on $\mathcal{Q}$ and a map $O : \mathcal{Q} \to \tau$, s.t. for all cells $A, B \in \mathcal{Q}$:

1. $A \subseteq O_A$;
2. if $B \triangleleft A$ then $O_A \cap B = \emptyset$ (where $\triangleleft$ is the corresponding **strict** order).

Theorem

A problem $\mathbb{P} = (\mathbb{S}, \mathcal{Q})$ is directly solvable by conditioning iff $\mathcal{Q}$ is linearly separated.

Universality: Building Stronger Refinements

Theorem

Conditioning is a universal problem solving method.

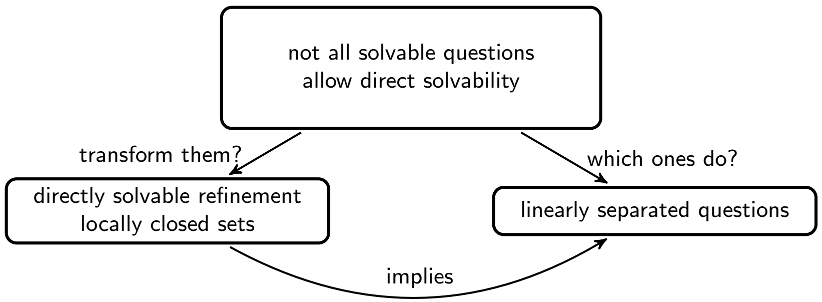

Can I observe $\neg p$? Negation-closed Observables

An epistemic space $\mathbb{B}_{\mathbb{S}}=(S, \mathcal{O})$ is **negation-closed** iff if $O \in \mathcal{O}$, then $S - O \in \mathcal{O}$.

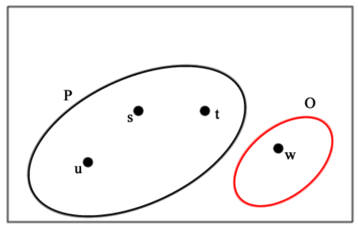

Can I observe $\neg p$? Negation-closed Observables

An epistemic space $\mathbb{B}_{\mathbb{S}}=(S,\mathcal{O})$ is **negation-closed** iff if $O\in\mathcal{O}$, then $S-O \in \mathcal{O}$.

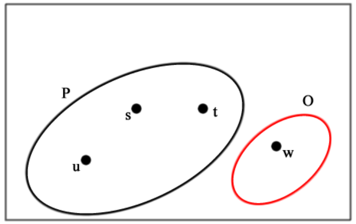

	Conditioning	Lexicographic	Minimal
Positive and Negative	YES	YES	NO
Fair Streams	NO	YES	NO

Outline

Introduction

Epistemic Spaces, Learnability and Solvability

Topo-characterizations of Learnability and Solvability

Constructive, Order-driven Learning: Belief Revision

Conclusion

Conclusion

- Iterated belief revision is given a DEL interpretation.
- Reliability of these policies as learning methods is evaluated.
- Learnability and solvability is characterised topologically.

THANK YOU!

THANK YOU!

- N. Gierasimczuk. Knowing One's Limits. Logical Analysis of Inductive Inference. PhD thesis, Universiteit van Amsterdam, The Netherlands, 2010.
- A. Baltag, N. Gierasimczuk, and S. Smets. Belief revision as a truth-tracking process. In K. Apt, editor, Proceedings of TARK'11, pages 187-190. ACM, 2011.
- A. Baltag, N. Gierasimczuk, and S. Smets. On the solvability of inductive problems: A study in epistemic topology. Proceedings of TARK'15, 2015.
- A. Baltag, N. Gierasimczuk, and S. Smets. Truth tracking by belief revision. Studia Logica 2018.
- A. Baltag, N. Gierasimczuk, A. Özgün, S. Smets, A. Vargas Sandoval. A Dynamic Logic for Learning Theory. Accepted to Journal of Logical and Algebraic Methods in Programming, 2019.