What makes a good teacher? A Computational Study

Nina Gierasimczuk

Institute for Logic, Language and Computation
Universiteit van Amsterdam

Logic Tea, November 30th 2009, Amsterdam

Outline

Outline

General Motivation: Learning

- Learning — getting to know and generalizing:
  - language learning (grammar inference),
  - scientific and medical (empirical) inquiry, etc.
- Worked out in formal learning theory and applied paradigms.
- Analyzing complexity is difficult.

General Motivation: Teaching

- Focus on single agent learning.
- Arbitrary environment.
- But complexity depends on data.
- Learning can be facilitated by clever teaching.
  - Is there a point to even start?
  - What information is relevant?
  - What is the most efficient way?

Outline

Motivation: illustration

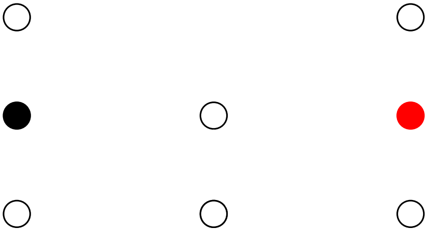

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: illustration

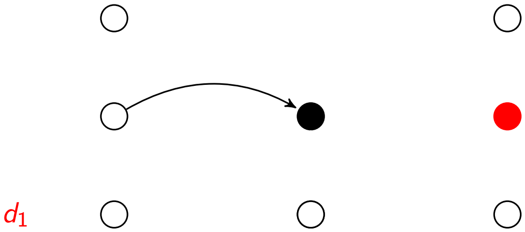

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: illustration

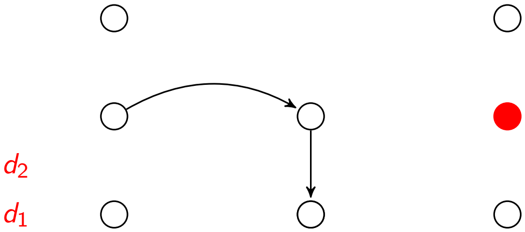

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: illustration

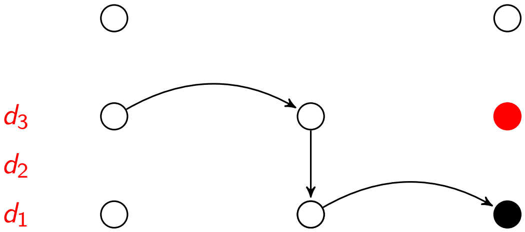

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: illustration

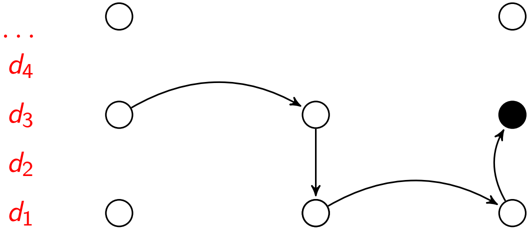

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: illustration

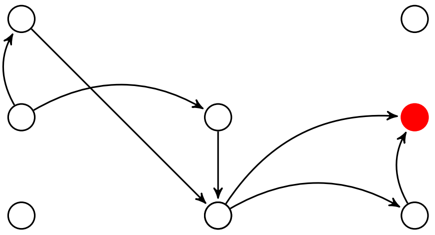

A class of hypotheses, with one distinguished goal and Learner generates a map of the possible mind changes of Learner.

Motivation: Learning/Teaching Game

- Learning situations as graphs.
- Learning as a game between Teacher and Learner played on a graph:
  - A step-by-step process in which Learner changes his state.
  - Successful if the goal is eventually reached.
  - Teacher's feedback rules out possible changes of mind.

Learning as a graph-game (1)

Learning Model	**The graph**
hypotheses	states
correct hypothesis	goal state
possibility of a mind change from hypothesis a to hypothesis b	edge from a to b
a mind change from hypothesis a to hypothesis b	transition from a to b
giving a counterexample that eliminates the possibility of a mind change from a to b	removing a transition between a and b

Sabotage?

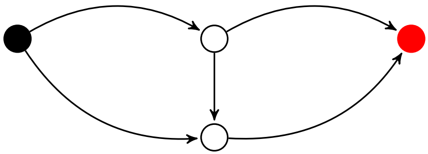

Sabotage?

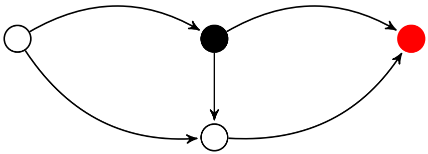

Sabotage?

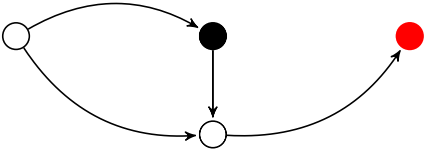

Sabotage?

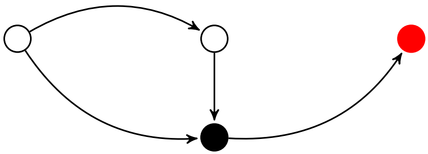

Sabotage?

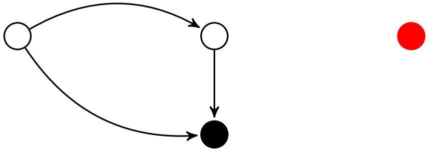

Sabotage Learning Game

Definition

A *Sabotage Learning Game* is a Game played between *Learner* and *Teacher* on a directed multi-graph with an initial vertex and a "goal" vertex.

Levels of cooperativeness

Levels of cooperativeness

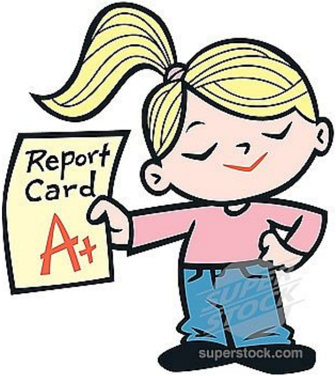

Levels of cooperativeness

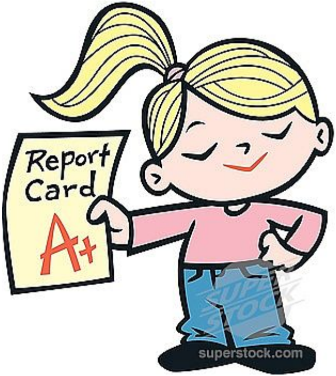

Various scenarios

Game	Winning Condition
SLG_{UE} unhelpful T., eager L.	Learner wins iff he reaches the goal state, Teacher wins otherwise.
SLG_{HU} helpful T., unwilling L.	Teacher wins iff Learner reaches the goal state, Learner wins otherwise.
SLG_{HE} helpful T., eager L.	Both players win iff Learner reaches the goal state, Both lose otherwise.

Sabotage Modal Logic

Sabotage Model

Take a finite Σ. The model $M = \langle W, (R_{a_i})_{a_i \in \Sigma}, Val \rangle$ is given by

$$W \neq \emptyset, \qquad R_{a_i} \subseteq W \times W, \qquad Val : \mathrm{PROP} \to \mathcal{P}(W)$$

Removal Operation

Let $M = \langle W, \{R_a \mid a \in \Sigma\}, Val \rangle$ be a Sabotage Model.

$$M^{a_i}_{(v,v')} := \langle W, R_{a_1}, \ldots R_{a_{i-1}}, R_{a_i} \setminus \{(v, v')\}, R_{a_{i+1}}, \ldots R_{a_n}, Val \rangle$$

Sabotage Modal Logic

Sabotage Modal Language

Language: $\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \Diamond_a \varphi \mid \diamondminus_a \varphi \qquad p \in \mathrm{PROP},\ a \in \Sigma$

Abbreviations: $\Diamond\varphi := \bigvee_{a\in\Sigma} \Diamond_a\varphi \qquad \diamondminus\varphi := \bigvee_{a\in\Sigma} \diamondminus_a\varphi$

Semantics

$$M, w \models \diamondminus_a \varphi \quad \text{iff} \quad \text{there is } (u, v) \in R_a \text{ s. t. } M^a_{(u,v)}, w \models \varphi$$

Theorem

Model checking of SML is PSPACE-complete (combined compl.).

Characterization Results

Theorem

Game	Existence of winning strategy	Winner
SLG_{UE} *unhelpful T., eager L.*	$\gamma_0^{UE} := goal,$ $\gamma_{n+1}^{UE} := goal \vee \Diamond \boxminus \gamma_n^{UE}$	*Learner*
SLG_{HU} *helpful T., unwilling L.*	$\gamma_0^{HU} := goal,$ $\gamma_{n+1}^{HU} := goal \vee (\Diamond \top \wedge (\Box \diamondminus \gamma_n^{HU}))$	*Teacher*
SLG_{HE} *helpful T., eager L.*	$\gamma_0^{HE} := goal,$ $\gamma_{n+1}^{HE} := goal \vee \Diamond \diamondminus \gamma_n^{HE}$	*Both*

Complexity of Sabotage-Type Learning

Theorem

Game	Winning Condition	Complexity
SLG_{UE}	*Learner wins iff he reaches the goal state, Teacher wins otherwise*	*PSPACE-complete*
SLG_{HU}	*Teacher wins iff Learner reaches the goal state, Learner wins otherwise.*	*PSPACE-complete*
SLG_{HE}	*Both players win iff Learner reaches the goal state. Both lose otherwise.*	*redNL-complete*

Local vs global moves

- Players' moves are of a different nature:
    - Learner moves by *local* transitions.
    - Teacher moves by *globally* removing an edge.
- Teacher only needs to act when necessary.

Strict vs non-strict alternation

Theorem

1. *Learner has a w.s. in SLG^*_{UE} iff he has a w.s. in SLG_{UE}.*
2. *Teacher has a w.s. in SLG^*_{HU} iff she has a w.s. in SLG_{HU}.*
3. *Teacher and Learner have a joint w.s. in SLG^*_{HE} iff they have a joint w.s. in SLG_{HE}.*

Corollary

- SLG^*: same complexity and characterization results as SLG.

High-level summary

Judging learnability/teachability with the assumption of

1. full cooperation — NL-complete (easy);
2. asymmetric obstruction — PSPACE-complete (difficult?).

Outline

Motivation

- High-level analysis, the cooperative case.
- How to make learning efficient?
- Explicit:
  - objects to be learned,
  - structure of information,
  - learning paradigm.

Framework

$\mathcal{L} = \{L_1, L_2, \ldots\}$ — an indexed family of languages.

Definition (Positive presentation)

By a positive presentation of L, ε, we mean an infinite sequence of elements from L such that it enumerates all and only the elements from L allowing repetitions.

Definition

- Let $\mathcal{L}$ be a class of languages, then $\mathcal{I}_{\mathcal{L}} = \{i \mid L_i \in \mathcal{L}\}$;
- $\varphi : (\bigcup \mathcal{L})^* \rightarrow \mathcal{I}_{\mathcal{L}}$ is a learning function.

Finite identifiability

Definition (Finite identification)

A learning function φ:

1. finitely identifies $L_i \in \mathcal{L}$ on ε iff, when inductively given ε, at some point φ outputs i, and stops;
2. finitely identifies $L_i \in \mathcal{L}$ iff it finitely identifies L_i on every ε for L_i;
3. finitely identifies $\mathcal{L}$ iff it finitely identifies every $L_i \in \mathcal{L}$;
4. a class $\mathcal{L}$ is finitely identifiable iff there is a learning function φ that finitely identifies $\mathcal{L}$.

Definition (Mukouchi 1992)

A set S_i is a definite finite tell-tale, DFTT, of $L_i \in \mathcal{L}$ if

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

Theorem (Mukouchi 1992)

A class $\mathcal{L}$ is finitely identifiable from positive data iff there is an effective procedure that enumerates all elements of a definite finite tell-tale S_i of L_i for any i.

Each set has a finite subset that distinguishes it from all other sets.

Definition

$\mathcal{L}$ — a class of languages, and $x \in \bigcup \mathcal{L}$. The eliminating power of x wrt $\mathcal{L}$ is determined by the function $El_{\mathcal{L}} : \bigcup \mathcal{L} \rightarrow \wp(\mathcal{I}_{\mathcal{L}})$, such that:

$$El_{\mathcal{L}}(x) = \{i | x \notin L_i\}.$$

Additionally, we will use $El_{\mathcal{L}}(X)$ for $\bigcup_{x \in X} El_{\mathcal{L}}(x)$.

$El_{\mathcal{L}}(x)$ gives indices of sets in $\mathcal{L}$ that are inconsistent with x.

Finite identifiability and eliminating power

Theorem

A class $\mathcal{L}$ is finitely identifiable from positive data iff there is an effective procedure that for any i enumerates all elements of a finite set $S_i \subseteq L_i$, such that

$$El_{\mathcal{L}}(S_i) = \{j | L_j \in \mathcal{L}\} - \{i\}.$$

Finite teachability is in PTIME

Theorem

Checking whether a finite class of finite sets is finitely identifiable is polynomial wrt the size of the class.

Proof.

1. Compute $El_{\mathcal{L}}(x)$.
2. Check whether for each $L_i \in \mathcal{L}$, $El_{\mathcal{L}}(L_i) = \mathcal{I}_{\mathcal{L}} - \{i\}$.

$\square$

Definition

$\mathcal{L}$ — a finitely identifiable class of languages, and $L_i \in \mathcal{L}$. A minimal DFTT of L_i in $\mathcal{L}$ is $S_i \subseteq L_i$, such that

1. S_i is a DFTT for L_i in $\mathcal{L}$, and
2. $\forall X \subset S_i \; El_{\mathcal{L}}(X) \neq \{j | j \in \mathcal{I}_{\mathcal{L}}\} - \{i\}$.

Example

$$\mathcal{L} = \{L_1 = \{1,3,4\}, L_2 = \{2,4,5\}, L_3 = \{1,3,5\}, L_4 = \{4,6\}\}$$

x	$El_{\mathcal{L}}(x)$
1	$\{2,4\}$
2	$\{1,3,4\}$
3	$\{2,4\}$
4	$\{3\}$
5	$\{1,4\}$
6	$\{1,2,3\}$

set	a minimal DFTT
$\{1,3,4\}$	$\{3,4\}$
$\{2,4,5\}$	$\{4,5\}$
$\{1,3,5\}$	$\{3,5\}$
$\{4,6\}$	$\{6\}$

Minimal finite teachability is in PTIME

Theorem

Let $\mathcal{L}$ be a finitely identifiable finite class of finite sets. Finding a minimal DFTT of $L_i \in \mathcal{L}$ can be done in polynomial time wrt to the size of the class.

Proof.

1. Set $X := L_i$.
2. Look for $x \in X$ such that $El(X - \{x\}) = \mathcal{I}_{\mathcal{L}} - \{i\}$.
   - If no such, then X is the desired reduct.
   - If there is, set $X := X - \{x\}$, and repeat the procedure.

□

Efficient teachability is NP-complete

Definition

$\mathcal{L}$ — a finitely identifiable class of languages, and $L_i \in \mathcal{L}$. A minimal DFTT of minimal size of L_i in $\mathcal{L}$ is a smallest $S_i \subseteq L_i$, such that

1. S_i is a DFTT for L_i in $\mathcal{L}$, and
2. $\forall X \subset S_i \; El_{\mathcal{L}}(X) \neq \{j | j \in \mathcal{I}_{\mathcal{L}}\} - \{i\}$.

Finding minimal DFTTs of minimal size

Perform a search through all the subsets of L_i starting from singletons, looking for the first $X_i \subseteq L_i$, such that $El_{\mathcal{L}}(X_i) = \mathcal{I}_{\mathcal{L}} - \{i\}$.

Example

$$\mathcal{L} = \{L_1 = \{1,3,4\}, L_2 = \{2,4,5\}, L_3 = \{1,3,5\}, L_4 = \{4,6\}\}.$$

x	$El_{\mathcal{L}}(x)$
1	$\{2,4\}$
2	$\{1,3,4\}$
3	$\{2,4\}$
4	$\{3\}$
5	$\{1,4\}$
6	$\{1,2,3\}$

set	min DFTTs of min size
$\{1,3,4\}$	$\{1,4\}$ or $\{3,4\}$
$\{2,4,5\}$	$\{2\}$
$\{1,3,5\}$	$\{1,5\}$ or $\{3,5\}$
$\{4,6\}$	$\{6\}$

NP-completeness

Min DFTTs of min size - a complete search through a large space.

Definition (MinDFTTMin Problem)

INSTANCE A finite class of finite sets $\mathcal{L}$, a set $L_i \in \mathcal{L}$, and a positive integer $k \leq |L_i|$.

QUESTION Is there a minimal DFTT $X_i \subseteq L_i$ of size $\leq k$?

Is $L_i \in \mathcal{L}$ finitely teachable from a sample of size k or smaller?

NP-completeness

Theorem

The MinDFTTMin *Problem is NP-complete.*

Definition (MinDFTTMin Problem)

Instance: A finite class of finite sets $\mathcal{L}$, a set $L_i \in \mathcal{L}$, and a positive integer $k \leq |L_i|$.

Question: Is there a minimal DFTT $X_i \subseteq L_i$ of size $\leq k$?

Definition (Minimal Cover Problem)

Instance: Collection P of subsets of a finite set F, positive integer $k \leq |P|$.

Question: Does P contain a cover for X of size k or less, i.e. a subset $P' \subseteq P$ with $|P'| \leq k$ such that every element of X belongs to at least one member of P'?

Low-level summary

- Judging finite teachability — PTIME;
- Teaching in a
  1. relevant way — PTIME (easy);
  2. the most efficient way — NP-complete (difficult?).

Outline

Conclusions and Further Work

Summary:

- Teachability can give a measure of difficulty for learning.
- Different approaches to the same epistemological problem: game theory, modal logic, formal learning theory, computational complexity.
- Higher resolution — higher complexity.

Further work:

- Specific graph structures for specific learning algorithms.
- Applicability to belief revision and belief merge.

References

N. Gierasimczuk, L. Kurzen and F. Velázquez-Quesada, Learning and Teaching as a Game: A Sabotage Approach, X. He, J. Horty, and E. Pacuit (Eds.): LORI 2009, LNAI 5834, pp. 119-132, 2009.

N. Gierasimczuk and D. de Jongh, On the Minimality of Definite Tell-tale Sets in Finite Identification, manuscript 2009.