A new invariant for the AFP Horn learning algorithm

José L. Balcázar (UPC)

part of ongoing work (with unclear destiny)
with Lenny Pitt, Bardia Sadri (UIUC) and others

Learning via Queries

Active Learning: the learning process interacts with a fixed target concept

Concept modeled as a set of (positive) instances: inductive bias

Interaction through a Query Protocol

Most popular protocol:

Membership queries

Equivalence queries

Termination implies correctness

Propositional Logic Context

Instances are bit vectors of length n

Equivalently:

tuples in a relation of a binary database

propositional models

In the last view, each concept is then a propositional theory

Our Case

Inductive bias is: we assume that the target concept is described by a conjunction of m propositional Horn clauses

Known: Horn theories are exactly those closed under bitwise intersection

Upper bound (Angluin/Frazier/Pitt 1992): algorithm making $O(nm^2)$ queries

Lower bound (Arias/Khardon/Servedio 2003): certificate size $\Omega(nm)$

Interest

Outside visibility!

Khardon first-order extensions (ILP, MRDM)

Kivinen and Mannila (databases)

Kearns, Kautz, and Selman
(knowledge compilation)

...

Extremely sensitive algorithm that we should understand better

Context and Notation

H^* Horn formula on n variables

Clauses $\alpha \to v$, α a term, v a variable or $\Box$ (definite Horn clauses for today)

Equivalently, H^* is closed under bitwise intersection

Negative assignments denoted $x\ominus$ ($x \not\models H^*$)

Positive assignments denoted $z\oplus$ ($z \models H^*$)

The Angluin/Frazier/Pitt Algorithm

Maintain set P of positive examples seen so far

Maintain negative examples $x_1, \ldots, x_t$

(Each one chases a different clause of the target H^*)

From the x_i, construct a sensible hypothesis $H^\bullet$ consistent with P

Process the counterexample y using memberships

Processing the Counterexample (sort of)

$y\oplus$: add to P (*reduces* the hypothesis H by removing all clauses inconsistent with y)

$y\ominus$: compare to $x_1, \ldots, x_t$

$i := \min(\{j | (x_j \wedge y)\ominus, x_j \not\leq y\} \cup \{t+1\})$

(suggestion: imagine that $x_{t+1} = 111...1$)

$x_i := [x_i \wedge]y$ (*refines* a negative example)

Refinement: $x_i \not\leq y \iff (x_i \wedge y) < x_i$

Termination

We bound $t \leq m$, where m is the number of clauses in any arbitrary Horn axiomatization of H^*

Each x_i is refined at most n times

According to [AFP], the following suffices:

if $x_i \not\models (\alpha \to v)$ and $x_j \not\models (\alpha \to v)$ for $i \neq j$, then $H^* \not\models (\alpha \to v)$

Want: a stronger argument that explains what happens if we change the algorithm a bit

Towards a New Invariant

A preliminary invariant:

$$\forall i < j \; \exists z \oplus \; [x_i \wedge x_j \leq z < x_j]$$

$y\oplus$ does not affect it

x_i refined into $x_i \wedge y$ maintains it

y becoming x_{t+1} (one case):

$$x_i \wedge x_{t+1} = (x_i \wedge y)\oplus < y \text{ for all } i \leq t$$

About the Preliminary Invariant (I)

$$\forall i < j \ \exists z \oplus \ [x_i \wedge x_j \leq z < x_j]$$

Assume $x_i \not\models (\alpha \rightarrow v)$ and $x_j \not\models (\alpha \rightarrow v)$ for $i \neq j$

$x_i \wedge x_j \leq z$: thus $z \models \alpha$

$z \leq x_j$: thus $z \not\models v$, and $z \not\models (\alpha \rightarrow v)$

But $z \oplus$, so $H^* \not\models (\alpha \rightarrow v)$

About the Preliminary Invariant (II)

$$\forall i < j \ \exists z \oplus \ [x_i \wedge x_j \leq z < x_j]$$

x_j refined into $x'_j = x_j \wedge y$ maintains it:

$x_i \wedge x'_j = x_i \wedge x_j \wedge y =$

$$= (x_i \wedge y) \wedge (x_i \wedge x_j) \leq (x_i \wedge y) \wedge z \leq x_j \wedge y = x'_j$$

But $j > i$: assume i was not refined because $(x_i \wedge y) \oplus$

By closure of H^* under bitwise intersection, $z' = ((x_i \wedge y) \wedge z) \oplus$

Problem: maybe i was not refined because $x_i \leq y$

Intuition about the New Invariant

$H^\bullet$ has the form $\bigwedge\{x_i \to v\}$ including all clauses that are not disproved by some positive example

Consider $x^\bullet = \bigwedge\{v | H^\bullet \models (x \to v)\}$

Similarly $x^* = \bigwedge\{v | H^* \models (x \to v)\}$, $x^* \oplus$

The construction of $H^\bullet$ implies $x^* \leq x^\bullet$

If $y \models H^\bullet$ and $x_i \leq y$, then $x_i \leq x^* \leq x^\bullet \leq y$

For $i < j$, any y that refined x_j may have had two reasons for skipping x_i, and both are good but must be argued separately

The Invariant

$\forall i < j$

$$\exists z_{i,j} \oplus \exists x_{i,j} \ \exists y_{i,j}$$

$$x_j = x_{i,j} \wedge y_{i,j},$$

$$x_i^* \leq y_{i,j},$$

$$x_i \wedge x_{i,j} \leq z_{i,j} < x_{i,j}$$

Call that expression $I_{i,j}$

How to Refine Several Antecedents

Consider $I_{i,j}$ and $I_{j,i}$

One of them suffices for each pair

Keep a directed graph indicating $\neg I_{i,j}$

Use it to refine several x_i but ensuring that, for each j, either $I_{i,j}$ or $I_{j,i}$

(Details not yet worked out - unclear what it buys!)

Goal: an algorithm making $O(mn)$ queries