

Learning weak constraints in answer set programming

MARK LAW, ALESSANDRA RUSSO* and KRYSIA BRODA

Department of Computing, Imperial College London, SW7 2AZ
(e-mail: {mark.law09, a.russo, k.broda}@imperial.ac.uk)

submitted 29 April 2015; revised 3 July 2015; accepted 15 July 2015

Abstract

This paper contributes to the area of inductive logic programming by presenting a new learning framework that allows the learning of weak constraints in Answer Set Programming (ASP). The framework, called *Learning from Ordered Answer Sets*, generalises our previous work on learning ASP programs without weak constraints, by considering a new notion of examples as *ordered* pairs of partial answer sets that exemplify which answer sets of a learned hypothesis (together with a given background knowledge) are *preferred* to others. In this new learning task inductive solutions are searched within a hypothesis space of normal rules, choice rules, and hard and weak constraints. We propose a new algorithm, ILASP2, which is sound and complete with respect to our new learning framework. We investigate its applicability to learning preferences in an interview scheduling problem and also demonstrate that when restricted to the task of learning ASP programs without weak constraints, ILASP2 can be much more efficient than our previously proposed system.

KEYWORDS: Non-monotonic Inductive Logic Programming, Preference Learning, Answer Set Programming

1 Introduction

Preference Learning has received much attention over the last decade from within the machine learning community. A popular approach to preference learning is *learning to rank* (Fürnkranz and Hüllermeier 2003; Geisler *et al.* 2001), where the goal is to learn to rank any two objects given some examples of pairwise preferences (indicating that one object is preferred to another). Many of these approaches use traditional machine learning tools such as neural networks (Geisler *et al.* 2001).

On the other hand, the field of Inductive Logic Programming (ILP) (Muggleton 1991) has seen significant advances in recent years, not only with the development of systems, such as (Ray *et al.* 2004; Kimber *et al.* 2009; Corapi *et al.* 2010; Muggleton *et al.* 2012; Muggleton and Lin 2013), but also the proposals of new frameworks for learning (Otero 2001; Sakama and Inoue 2009; Law *et al.* 2014). In most approaches

* This research is partially funded by the 7th Framework EU-FET project 600792 “ALLOW Ensembles”, and the EPSRC project EP/K033522/1 “Privacy Dynamics”.

to ILP, a learning task consists of a background knowledge B and sets of positive and negative examples. The task is then to find a hypothesis that, together with B , covers all the positive examples but none of the negative examples. While in previous work ILP systems such as TILDE (Blockeel and De Raedt 1998) and Aleph (Srinivasan 2001) have been applied to preference learning (Dastani *et al.* 2001; Horváth 2012), this has addressed learning ratings, such as *good*, *poor* and *bad*, rather than rankings over the examples. Ratings are not expressive enough if we want to find an optimal solution as we may rate many objects as *good* when some are better than others. Answer Set Programming (ASP), on the other hand, allows the expression of preferences through *weak constraints*. In a usual application of ASP, one would write a logic program which has many answer sets, each corresponding to a solution of the problem. The program can also contain weak constraints (or optimisation statements) which impose an ordering on the answer sets. Modern ASP solvers, such as clingo (Gebser *et al.* 2011), can then find the optimal answer sets, which correspond to the optimal solutions of the problem. For instance, in a scheduling problem, we could define an ASP program, whose answer sets correspond to timetables, and weak constraints that represent preferences over these timetables (see (Banbara *et al.* 2013) for an example application of ASP in timetabling).

In this paper, we propose a new learning framework, *Learning from Ordered Answer Sets* (ILP_{LOAS}), that allows the learning of ASP programs with weak constraints. This framework extends the notion of learning from answer sets proposed in (Law *et al.* 2014), where ASP programs without weak constraints were learned using only positive and negative examples of partial answer sets. In our new learning task ILP_{LOAS} , additional examples are defined, as ordered pairs of partial answer sets, and the language bias captures a hypothesis space of ASP programs containing normal rules, choice rules and hard and weak constraints. A new algorithm is presented and proved to be sound and complete with respect to ILP_{LOAS} .

To demonstrate the applicability of our framework, we consider, as a running example, an interview timetabling problem and the task of learning, as weak constraints, academics' preferences for scheduling undergraduate interviews. An academic might be more comfortable interviewing for one course than another, might prefer not to have many interviews on the same day, or might hold both of these preferences but regard the former as more important. Given ordered pairs of partial timetables, our approach is able to learn these preferences as weak constraints.

The paper is structured as follows. Our new learning framework, ILP_{LOAS} is presented in Section 3. It extends the notion of Learning from Answer Sets (Law *et al.* 2014) to the new task of learning weak constraints. We discuss formal properties of the framework such as the complexity of deciding the existence of a solution. Our learning algorithm $ILASP2$ is described in Section 4, together with experimental results based on a scheduling example (Section 5). We also show that $ILASP2$ can have increased efficiency over our previous system when learning programs without weak constraints. Discussion on related and future work concludes the paper.

2 Background

In this section we introduce the concepts needed in the paper. Given any atoms $h, h_1, \dots, h_o, b_1, \dots, b_n, c_1, \dots, c_m$, $h \leftarrow b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$ is called a *normal rule*, with h as the *head* and $b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$ (collectively) as the *body* (“not” represents negation as failure); a rule $\leftarrow b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$ is a *hard constraint*; a *choice rule* is a rule $l\{h_1, \dots, h_o\}u \leftarrow b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$ (where l and u are integers) and its head is called an *aggregate*. A variable in a rule R is *safe* if it occurs in at least one positive literal in the body of R . A program P is assumed to be a finite set of normal rules, choice rules, and hard constraints. The Herbrand Base of P , denoted HB_P , is the set of variable free (ground) atoms that can be formed from predicates and constants in P . The subsets of HB_P are called the (Herbrand) interpretations of P . A ground aggregate $l\{h_1, \dots, h_o\}u$ is satisfied by an interpretation I iff $l \leq |I \cap \{h_1, \dots, h_o\}| \leq u$.

As we restrict our programs to sets of normal rules, (hard) constraints and choice rules, we can use the simplified definitions of the *reduct* for choice rules presented in (Law *et al.* 2015c). Given a program P and an Herbrand interpretation $I \subseteq HB_P$, the reduct P^I is constructed from the grounding of P in 4 steps: firstly, remove rules whose bodies contain the negation of an atom in I ; secondly, remove all negative literals from the remaining rules; thirdly, replace the head of any hard constraint, or any choice rule whose head is not satisfied by I with \perp (where $\perp \notin HB_P$); and finally, replace any remaining choice rule $\{h_1, \dots, h_m\} \leftarrow b_1, \dots, b_n$ with the set of rules $\{h_i \leftarrow b_1, \dots, b_n \mid h_i \in I \cap \{h_1, \dots, h_m\}\}$. Any $I \subseteq HB_P$ is an *answer set* of P if it is the minimal model of the reduct P^I . Throughout the paper we denote the set of answer sets of a program P with $AS(P)$.

Unlike hard constraints in ASP, *weak constraints* do not affect what is, or is not, an answer set of a program P . Hence the above definitions also apply to programs with weak constraints. Weak constraints create an ordering over $AS(P)$ specifying which answer sets are “better” than others. The set of *optimal* (best) answer sets of P is denoted as $AS^*(P)$. A *weak constraint* is of the form $\sim b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m \cdot [w@l, t_1, \dots, t_o]$ where $b_1, \dots, b_n, c_1, \dots, c_m$ are atoms, w and l are terms specifying the *weight* and the *level*, and t_1, \dots, t_o are terms. A weak constraint W is *safe* if every variable in W occurs in at least one positive literal in the body of W . At each *priority level* l , the aim is to discard any answer set which does not minimise the sum of the weights of the ground weak constraints (with level l) whose bodies are true. The higher levels are minimised first. Terms specify which ground weak constraints should be considered unique. For any program P and $A \in AS(P)$, $weak(P, A)$ is the set of tuples (w, l, t_1, \dots, t_o) for which there is some $\sim b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m \cdot [w@l, t_1, \dots, t_o]$ in the grounding of P such that A satisfies $b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$.

We now give the semantics for weak constraints (Calimeri *et al.* 2013). For each level l , $P_A^l = \sum_{(w, l, t_1, \dots, t_o) \in weak(P, A)} w$. For $A_1, A_2 \in AS(P)$, A_1 *dominates* A_2 (written $A_1 \succ_P A_2$) iff $\exists l$ such that $P_{A_1}^l < P_{A_2}^l$ and $\forall m > l, P_{A_1}^m = P_{A_2}^m$. An answer set $A \in AS(P)$ is *optimal* if it is not dominated by any $A_2 \in AS(P)$.

Example 1

Let P be the program consisting of `slot(m,1); slot(m,2); slot(t,1); slot(t,2);` and `0{assign(D,S)}1 ← slot(D,S`, which assigns 0 to 4 slots in a schedule (`slot(m,1)` represents slot 1 on Monday). Let W_1 , W_2 and W_3 be the weak constraints $:\sim \text{assign(D,S)}.[1@1]$, $:\sim \text{assign(D,S)}.[1@1,D]$ and $:\sim \text{assign(D,S)}.[1@1,D,S]$ respectively. Applying each weak constraint to P gives as its optimal answer set the one in which no slots are assigned. The remaining answer sets are ordered in the following way: W_1 considers all schedules in which slots have been assigned to be equally optimal, as there is only one unique set of terms t_1, \dots, t_n which is the empty set; W_2 minimises the number of days in which slots have been assigned, as there is one unique set of terms per day; and finally, W_3 minimises the number of assignments made, as each combination of day and slot has a unique set of terms.

In an ILP task, the hypothesis space is often characterised by mode declarations (Muggleton *et al.* 2012). A *mode bias* can be defined as a pair of sets of mode declarations $\langle M_h, M_b \rangle$, where M_h (resp. M_b) are the *head* (resp. *body*) declarations. Each mode declaration $m \in M_h$, or $m \in M_b$, is a literal whose abstracted arguments are either v or c . An atom a is *compatible* with a mode declaration m if replacing the instances of v in m by variables, and the instances of c by constants yields a . The search space is defined to be the set of rules of the form $h \leftarrow b_1, \dots, b_n, \text{not } c_1, \dots, \text{not } c_m$ where (i) h is empty, h is an atom compatible with some $m \in M_h$, or h is an aggregate $l\{h_1, \dots, h_k\}u$ such that $0 \leq l \leq u \leq k$ and $\forall i \in [1, k]$ h_i is compatible with some $m \in M_h$; (ii) $\forall i \in [1, n], \forall j \in [1, m]$ b_i and c_j are each compatible with at least one mode declaration in M_b ; and finally (iii) all variables in the rule are safe. We require the rules to be safe because ASP solvers such as *clingo* (Gebser *et al.* 2011) have this requirement. We denote the search space defined by a given mode bias (M_h, M_b) as $S_{LAS}(M_h, M_b)$.

In (Law *et al.* 2014), we presented a new learning task, *Learning from Answer Sets* (ILP_{LAS}) which used *partial interpretations* as examples. A partial interpretation e is a pair $\langle e^{inc}, e^{exc} \rangle$ of sets of ground atoms, called *inclusions* and *exclusions*. An answer set A is said to *extend* e if and only if $(e^{inc} \subseteq A) \wedge (e^{exc} \cap A = \emptyset)$. Given partial interpretations e_1 and e_2 , e_1 extends e_2 iff $e_2^{inc} \subseteq e_1^{inc}$ and $e_2^{exc} \subseteq e_1^{exc}$.

Definition 1

A *Learning from Answer Sets* task is a tuple $T = \langle B, S_{LAS}(M_h, M_b), E^+, E^- \rangle$ where B is the background knowledge, $S_{LAS}(M_h, M_b)$ is the search space defined by a bias $\langle M_h, M_b \rangle$, E^+ and E^- are sets of partial interpretations called the positive and negative examples. A hypothesis H is in $ILP_{LAS}(T)$, the set of all *inductive solutions* of T , if and only if $H \subseteq S_{LAS}(M_h, M_b)$; $\forall e^+ \in E^+ \exists A \in AS(B \cup H)$ such that A extends e^+ ; and finally, $\forall e^- \in E^- \nexists A \in AS(B \cup H)$ such that A extends e^- .

The task of ILP is usually to find *optimal* hypotheses, where optimality is often defined by the number of literals in a hypothesis. In ILP_{LAS} , aggregates are converted to disjunctions before the literals are counted, giving a higher “cost” for learning an aggregate; for example, $1\{p, q\}2$ is converted to $(p \wedge \text{not } q) \vee (q \wedge \text{not } p) \vee (p \wedge q)$ giving a length of 6. ILP_{LAS} aims at learning ASP programs consisting of normal

rules, choice rules and hard constraints. This paper extends this notion to a new learning task capable of learning weak constraints.

3 Learning from ordered answer sets

To learn weak constraints we extend the notion of mode bias with two new sets of mode declarations: M_o specifies what is allowed to appear in the body of a weak constraint; whereas M_w specifies what is allowed to appear as a weight. A positive integer l_{max} is also given to indicate the number of levels that can occur in H .

Definition 2

A mode bias with ordering is a tuple $M = \langle M_h, M_b, M_o, M_w, l_{max} \rangle$, where M_h and M_b are respectively head and body declarations, M_o is a set of mode declarations for body literals in weak constraints, M_w is a set of integers and l_{max} is a positive integer. The search space S_M is the set of rules R that satisfy one of the conditions:

- $R \in S_{LAS}(M_h, M_b)$.
- R is a safe weak constraint $:\sim b_1, \dots, b_i, \text{not } b_{i+1}, \dots, \text{not } b_j. [w@l, t_1, \dots, t_n]$ such that $\forall k \in [1, j] b_k$ is compatible with M_o ; t_1, \dots, t_n is the set of terms in b_1, \dots, b_j ; $w \in M_w, l \in [0, l_{max})$.

Note that even if we were to extend the learning task in Definition 1 with this new notion of mode bias, such a task would never have as its optimal solution a hypothesis which contains a weak constraint. This is because a Learning from Answer Sets task has only examples of what is, or is not, an answer set. Any solution containing a weak constraint W will have the same answer sets without W , and would be more optimal. We now define the notion of *ordering examples*.

Definition 3

An *ordering example* is a tuple $o = \langle e_1, e_2 \rangle$ where e_1 and e_2 are partial interpretations. An ASP program P *bravely respects* o iff $\exists A_1, A_2 \in AS(P)$ such that A_1 extends e_1 , A_2 extends e_2 and $A_1 \succ_P A_2$. P *cautiously respects* o iff $\forall A_1, A_2 \in AS(P)$ such that A_1 extends e_1 and A_2 extends e_2 , it is the case that $A_1 \succ_P A_2$.

Example 2

Consider the partial interpretations $e_1 = \langle \{\text{assign}(m, 1), \text{assign}(m, 2)\}, \{\text{assign}(t, 1), \text{assign}(t, 2)\} \rangle$ and $e_2 = \langle \{\text{assign}(m, 1), \text{assign}(t, 1)\}, \emptyset \rangle$. Let $o = \langle e_1, e_2 \rangle$ be an ordering example and recall P and W_1, \dots, W_3 from example 1. The only answer set of P that extends e_1 is m_1m_2 (where m_1m_2 denotes $\{\text{assign}(m, 1), \text{assign}(m, 2)\}$), whereas the answer sets that extend e_2 are $m_1t_1, m_1m_2t_1, m_1m_2t_1t_2$ and $m_1t_1t_2$. $P \cup W_1$ does not bravely or cautiously respect o as it gives to all these answer sets the same optimality; $P \cup W_2$ both bravely and cautiously respects o , as each pair of answer sets extending the partial interpretations is ordered correctly (i.e. answer sets extending e_1 have slots allocated in only one day whereas all the answer sets extending e_2 have slots assigned in two days). Finally, $P \cup W_3$ respects o bravely but not cautiously (the pair of answer sets m_1m_2 and m_1t_1 is such that $m_1m_2 \not\succeq_P m_1t_1$).

We can now define the notion of *Learning from Ordered Answer Sets (ILP_{LOAS})*.

Definition 4

A *Learning from Ordered Answer Sets* task is a tuple $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$ where B is an ASP program, called the background knowledge, S_M is the search space defined by a mode bias with ordering M , E^+ and E^- are sets of partial interpretations called, respectively, positive and negative examples, and O^b and O^c are sets of ordering examples over E^+ called brave and cautious orderings. A hypothesis $H \subseteq S_M$ is in $ILP_{LOAS}(T)$, the *inductive solutions* of T , if and only if:

1. Let M_h and M_b be as in M and H' be the subset of H with no weak constraints.
 $H' \in ILP_{LAS}(\langle B, S_{LAS}(M_h, M_b), E^+, E^- \rangle)$
2. $\forall o \in O^b$ $B \cup H$ bravely respects o
3. $\forall o \in O^c$ $B \cup H$ cautiously respects o

The notion of an optimal inductive solution of an ILP_{LOAS} task is the same as in ILP_{LAS} , with each weak constraint W counted as the number of literals in the body of W . Note that the orderings are only over E^+ rather than orderings over any arbitrary partial interpretations. We chose to make this restriction as we could not see a reason why a hypothesis would need to respect orderings which are not extended by any pair of answer sets of $B \cup H$. Note also that in the case where O^b and O^c are both empty, the task reduces to an ILP_{LAS} task.

Example 3

Consider the ILP_{LOAS} task T with the following background knowledge:

$$B = \left\{ \begin{array}{l} \text{slot}(m, 1..3). \text{slot}(t, 1..3). \text{slot}(w, 1..3). \\ \text{neq}(1, 2). \text{neq}(1, 3). \text{neq}(2, 1). \text{neq}(2, 3). \text{neq}(3, 1). \text{neq}(3, 2). \\ \text{neq}(m, t). \text{neq}(m, w). \text{neq}(t, m). \text{neq}(t, w). \text{neq}(w, m). \text{neq}(w, t). \\ \text{type}(m, 1, c1). \text{type}(m, 2, c2). \text{type}(m, 3, c2). \text{type}(t, 1, c2). \\ \text{type}(t, 2, c2). \text{type}(t, 3, c2). \text{type}(w, 1, c2). \text{type}(w, 2, c1). \text{type}(w, 3, c2). \\ 0\{\text{assign}(X, Y)\}1: \text{-slot}(X, Y). \end{array} \right\}$$

Using the notation from example 2, let T have the positive examples $e_1 = \langle \emptyset, m_2m_3t_1t_3w_1w_2 \rangle$, $e_2 = \langle m_1m_2, \emptyset \rangle$, $e_3 = \langle \emptyset, m_1t_2w_1w_2 \rangle$, $e_4 = \langle t_1t_2t_3, \emptyset \rangle$, $e_5 = \langle m_2m_3t_1t_2t_3w_1w_3, \emptyset \rangle$; $e_6 = \langle m_1w_1w_3, m_2m_3t_1t_2t_3w_2 \rangle$; two cautious orderings: $\langle e_1, e_2 \rangle$ and $\langle e_3, e_4 \rangle$; and one brave ordering: $\langle e_5, e_6 \rangle$. Consider S_M to be defined by the mode declarations: $M_h = M_b = \emptyset$; $M_o = \{\text{assign}(v, v), \text{neq}(v, v), \text{type}(v, c)\}$; $M_w = \{-1, 1\}$; and $l_{max} = 2$. Note that as each positive example is already covered by the background knowledge and there are no negative examples, it remains to find a set of weak constraints which meet conditions 2 and 3 of definition 4. One inductive solution H of T is $\{\sim \text{assign}(D, S1), \text{assign}(D, S2), \text{neq}(S1, S2). [1@1, D, S1, S2] \sim \text{assign}(D, S), \text{type}(D, S, c1). [1@2, D, S]\}$; this respects the first cautious ordering example because any timetable extending e_1 has at most one c1 course whereas e_2 has at least one, so e_1 is better or equal to e_2 on the highest priority weak constraint; even if they are equal, a timetable extending e_1 has at most one assignment per day and is, therefore, always better on the lower priority weak constraint. H also respects the other cautious ordering and the timetables $m_2m_3t_1t_2t_3w_1w_3$ and $m_1w_1w_3$ correspond to answer sets which demonstrate that the brave ordering is respected.

In fact, there is no shorter hypothesis which meets conditions 1 to 3 and so H is an optimal inductive solution; moreover, the other optimal solutions are equivalent hypotheses such as: $\{\sim \text{assign}(D, S1), \text{assign}(D, S2), \text{neq}(S1, S2). [1@1, D, S1, S2]; \sim \text{assign}(D, S), \text{not type}(D, S, c2). [1@2, D, S]\}$. These hypotheses represent the preferences described in the introduction. They express that the highest priority is to minimise the interviews for $c1$, and then to minimise the slots in any one day.

We now discuss some of the formal properties of ILP_{LOAS} . All learning tasks in the rest of this section are assumed to be propositional (B and S_M are both ground). The proofs for Theorems 1 to 3 can be found in (Law *et al.* 2015b). Theorems 1 and 2 state sufficient and necessary conditions for there to exist solutions for an ILP_{LOAS} task with an unrestricted search space (hypotheses can be any set of normal rules, choice rules and hard and weak constraints).

Theorem 1

Let T be the ILP_{LOAS} task $\langle B, E^+, E^-, O^b, O^c \rangle$. The following conditions (in conjunction) are sufficient for there to exist solutions of T : (i) $\forall e \in E^+$, there is at least one model of B which extends e ; (ii) $\forall e_1 \in E^+$, $\nexists e_2 \in (E^+ \cup E^-)$ such that e_1 extends e_2 ; (iii) there is no cyclic chain of ordering examples (in $O^b \cup O^c$) $\langle e_1, e_2 \rangle, \langle e_2, e_3 \rangle, \dots, \langle e_{n-1}, e_n \rangle, \langle e_n, e_1 \rangle$.

Theorem 2

Let T be the ILP_{LOAS} task $\langle B, E^+, E^-, O^b, O^c \rangle$. The following conditions are necessary for there to exist solutions of T : (i) $\forall e \in E^+$, there is at least one model of B which extends e ; (ii) $\forall e_1 \in E^+$, $\nexists e_2 \in E^-$ such that e_1 extends e_2 ; (iii) there is no cyclic chain of cautious orderings, $\langle e_1, e_2 \rangle, \langle e_2, e_3 \rangle, \dots, \langle e_{n-1}, e_n \rangle, \langle e_n, e_1 \rangle$.

Note that if we consider the usual setting where hypotheses come from a search space, the conditions in theorem 2 are still necessary, but the conditions in theorem 1 are no longer sufficient as, even if the conditions hold, the search space may be too restrictive. Theorem 3 states the complexity of deciding the existence of solutions for both ILP_{LAS} and ILP_{LOAS} tasks. The interesting property here is that deciding the existence of solutions for ILP_{LOAS} is in the same complexity class as ILP_{LAS} .

Theorem 3

Let T be any ILP_{LAS} or ILP_{LOAS} task. Deciding whether T has at least one inductive solution is NP^{NP} -complete.

4 Algorithm

We now describe our new algorithm, $ILASP2$, capable of computing inductive solutions of any ILP_{LOAS} task, and present its soundness and completeness results with respect to the notion of Learning from Ordered Answer Sets task given in Definition 4. We omit the proofs of the theorems in this paper, but they can be found in full in (Law *et al.* 2015b). For details of how to download and use our prototype implementation of the $ILASP2$ algorithm, see (Law *et al.* 2015a).

$ILASP2$ extends the concepts of *positive* and *violating* hypotheses, first introduced in our previous algorithm $ILASP$ (Law *et al.* 2014), to cater for the new notion of

ordering examples. A hypothesis is said to be *positive* if it covers all positive examples and bravely respects all the brave ordering examples. A positive hypothesis is defined to be *violating* if it covers at least one negative example or if it does not respect at least one of the cautious ordering examples. These two notions are formalised by Definitions 5 and 6.

Definition 5

Let $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$ be an ILP_{LOAS} task. Any $H \subseteq S_M$ is a *positive hypothesis* iff $\forall e \in E^+ \exists A \in AS(B \cup H)$ such that A extends e , and $\forall o \in O^b H \cup B$ bravely respects o . The set of positive hypotheses of T is denoted $\mathcal{P}(T)$.

Definition 6

A positive hypothesis H is a *violating hypothesis* of $T = \langle B, S_M, E^+, E^-, O^b, O^c \rangle$, written $H \in \mathcal{V}(T)$, iff at least one of the following cases is true:

- $\exists e^- \in E^-$ and $\exists A \in AS(B \cup H)$ such that A extends e^- . In this case we call A a violating interpretation of T and write $\langle H, A \rangle \in \mathcal{V}\mathcal{I}(T)$.
- $\exists A_1, A_2 \in AS(B \cup H)$ and $\exists \langle e_1, e_2 \rangle \in O^c$ such that A_1 extends e_1 , A_2 extends e_2 and $A_1 \not\prec_P A_2$ with respect to $B \cup H$. In this case, we call $\langle A_1, A_2 \rangle$ a *violating pair* of T and write $\langle H, \langle A_1, A_2 \rangle \rangle \in \mathcal{V}\mathcal{P}(T)$.

Example 4

Consider an ILP_{LOAS} task with B equal to P in Example 1 but with one additional fact `busy(1,1)`; positive examples $e_1^+ = \langle \{\text{assign}(t,1), \text{assign}(t,2)\}, \{\text{assign}(m,2)\} \rangle$ and $e_2^+ = \langle \{\text{assign}(m,2), \text{assign}(t,1)\}, \emptyset \rangle$; one negative example $e_1^- = \langle \{\text{assign}(m,1)\}, \emptyset \rangle$; and one cautious ordering $\langle e_1^+, e_2^+ \rangle$. S_M is unrestricted (hypotheses can be constructed from any predicate that appears in B and E). Three example hypotheses are given below. Note that when we describe answer sets we omit the facts in B .

$H_1 = \emptyset \in \mathcal{P}(T)$ as e_1^+ and e_2^+ are covered and O^b is empty; however, $H_1 \in \mathcal{V}(T)$ for two reasons: firstly it has a violating interpretation ($\{\text{assign}(m,1)\}$); secondly it has a violating pair ($\langle \{\text{assign}(t,1), \text{assign}(t,2)\}, \{\text{assign}(m,2), \text{assign}(t,1)\} \rangle$).

$H_2 = \{ \leftarrow \text{busy}(D,S), \text{assign}(D,S) \} \in \mathcal{P}(T)$. H_2 has no violating interpretations, but it has a violating pair ($\langle \{\text{assign}(t,1), \text{assign}(t,2)\}, \{\text{assign}(m,2), \text{assign}(t,1)\} \rangle$).

$H_3 = H_2 \cup \{ \sim \text{assign}(D,S).[1@1,D] \} \in \mathcal{P}(T)$. $H_3 \notin \mathcal{V}(T)$, as it has no violating interpretations and its weak constraints minimise the days assigned (so it cautiously respects the ordering example). It is, therefore, an inductive solution of the task.

One approach to computing the inductive solutions of an ILP_{LOAS} task would be to extend the original ILASP method with our new notions of positive and violating hypotheses. Given a positive integer n , ILASP worked by constructing an ASP meta representation of an ILP_{LAS} task T , called the *task program* T_{meta}^n , whose answer sets could be mapped back to the positive hypotheses of T of length n . T_{meta}^n could then be augmented with an extra constraint so that its answer sets corresponded exactly to the violating hypotheses of length n . ILASP first computed the violating hypotheses of length n , and then converted each of these to a constraint at the meta-level (ruling out that hypothesis). When T_{meta}^n was then augmented with these new constraints, its answer sets corresponded exactly to the positive hypotheses which were not computed the first time - the inductive solutions of length n .

The problem is that, in general, there can be many violating hypotheses which are shorter than the first inductive hypothesis and ILASP will compute all of them and add them into the task program as individual constraints. This scalability issue would be worsened if we were considering adding weak constraints to the search space. To overcome this, *ILASP2* adopts a different strategy: it eliminates *classes* of hypothesis, i.e. hypotheses that are violating for the same “reason”, namely they give rise to a particular violating interpretation or a particular violating pair of interpretations. The idea underlying the *ILASP2* algorithm is to make use of two sets VI and VP which accumulate, respectively, violating interpretations and violating pairs of interpretations that are constructed during the search. We start initially with two empty sets VI and VP and continually compute the set of optimal *remaining hypotheses* which do not violate any of the *reasons* in VI or VP . If a computed hypothesis gives rise to a new violating interpretation then this interpretation is added to VI , if it gives rise to a new violating pair of interpretations then this pair is added to VP . If no optimal remaining hypotheses are violating, then these hypotheses are the optimal inductive solutions of the task.

Definition 7

Let T be an ILP_{LOAS} task, VI and VP (resp.) be sets of violating interpretations and pairs of interpretations, and B be the background knowledge. Any $H \in \mathcal{P}(T)$ is a *remaining hypothesis* of T with respect to $VI \cup VP$ iff $VI \cap AS(B \cup H) = \emptyset$ and $\forall \langle I_1, I_2 \rangle \in VP$ if $I_1, I_2 \in AS(B \cup H)$ then $I_1 \succ_{B \cup H} I_2$. A remaining hypothesis H is a *remaining violating hypothesis* iff $\exists R$ such that $\langle H, R \rangle \in \mathcal{V}\mathcal{I}(T) \cup \mathcal{V}\mathcal{P}(T)$.

We use an ASP meta-level representation to solve our search for remaining hypotheses. As we rule out classes of hypothesis at the same time (rather than using one constraint per violating hypothesis), our meta-level representation is slightly more complex than that used in the original ILASP. Due to this complexity, we define this representation in the online appendix and give here the underlying intuition.

The intuition of our meta encoding is that for a given task T , we construct an ASP program T_{meta} whose answer sets can be mapped back to the positive hypotheses of T . Given an answer set A of T_{meta} we write $\mathcal{M}_{hyp}^{-1}(A)$ to denote the hypothesis represented by A . Each positive hypothesis may be represented by many answer sets of T_{meta} but if this hypothesis gives rise to a violating interpretation, then at least one of these answer sets will contain a special atom v_i . If the hypothesis gives rise to a violating pair of interpretations then at least one of the answer sets of T_{meta} representing the hypothesis will contain a special atom $v_p(t_1, t_2)$, where $\langle t_1, t_2 \rangle$ is a pair of identifiers corresponding to the cautious ordering example which is being violated. There is only one priority level in T_{meta} and the optimality of its answer sets is $2 * |H| + 1$ if the answer set does not contain the atom *violating* (*violating* is defined to be true if and only if v_i or at least one $v_p(t_1, t_2)$ is true) and $2 * |H|$ if it does. This means that for any hypothesis H , the answer sets corresponding to H that do contain *violating* are preferred to those which do not.

We can use T_{meta} to find optimal positive hypotheses of T . If these positive solutions are violating, then the optimal answer sets will contain *violating*. We can

then rule these hypotheses out. We can extract violating interpretations and violating pairs of interpretations from answer sets of T_{meta} , using the functions \mathcal{M}_{vi}^{-1} and \mathcal{M}_{vp}^{-1} respectively. Violating interpretations and violating pairs of interpretations are both called *violating reasons*. For any set of violating reasons $VR = VI \cup VP$, we then have a second meta encoding $VR_{meta}(T)$ which, when added to T_{meta} , rules out any hypotheses which are violating for a reason already in VR . This means that the answer sets of $T_{meta} \cup VR_{meta}(T)$ will represent the set of remaining hypotheses of T with respect to VR . These properties are guaranteed by Theorem 4.

Theorem 4

Given an ILP_{LOAS} task and a set of violating reasons VR , let AS be the set of optimal answer sets of $T_{meta} \cup VR_{meta}(T)$.

- If $\exists A \in AS$ such that *violating* $\in A$ then the set of optimal remaining violating hypotheses VH is non empty and is equal to the set $\{\mathcal{M}_{hyp}^{-1}(A) \mid A \in AS\}$.
- If no $A \in AS$ contains *violating*, then the set of optimal remaining hypotheses (none of which is violating) is equal to the set $\{\mathcal{M}_{hyp}^{-1}(A) \mid A \in AS\}$.

Algorithm 1 ILASP2

procedure ILASP2(T)

$VR = []$

$solution = solve(T_{meta} \cup VR_{meta}(T))$

while $solution \neq \text{nil}$ && $solution.optimality \% 2 == 0$ **do**

$A = solution.answer_set$

if $v_i \in A$ **then**

$VR += \mathcal{M}_{vi}^{-1}(A)$

else if $\exists t_1, t_2$ such that $v_p(t_1, t_2) \in A$ **then**

$VR += \mathcal{M}_{vp}^{-1}(A)$

end if

$solution = solve(T_{meta} \cup VR_{meta}(T))$

end while

return $\{\mathcal{M}_{hyp}^{-1}(A) \mid A \in AS^*(T_{meta} \cup VR_{meta}(T))\}$

end procedure

Algorithm 1 is the pseudo code of our algorithm ILASP2. It makes use of our meta encodings T_{meta} and $VR_{meta}(T)$. For any program P , $solve(P)$ is a function which, in the case that P is satisfiable, returns a pair consisting of an optimal answer set together with its optimality (as there is only one priority level in our meta encoding this is treated as an integer); if P is unsatisfiable then $solve(P)$ returns nil . While there are optimal remaining violating hypotheses, ILASP2 finds them and records the appropriate violating reasons. When there are no optimal remaining hypotheses which are violating then either the meta program will be unsatisfiable or the optimality of the optimal answer sets will be odd (as the optimality of any $A \in AS(T_{meta} \cup VR_{meta}(T))$ is $2 * |\mathcal{M}_{hyp}^{-1}(A)|$ if A contains *violating* and $2 * |\mathcal{M}_{hyp}^{-1}(A)| + 1$ if not), and so ILASP2 stops and returns the set of optimal remaining hypotheses.

Theorem 5 shows that ILASP2 is sound and complete with respect to the optimal inductive solutions of an ILP_{LOAS} task. This result relies on the termination of $ILASP2(T)$, which is guaranteed if $B \cup S_M$ grounds finitely.

Theorem 5

Let T be an ILP_{LOAS} task. If $ILASP2(T)$ terminates, then $ILASP2(T)$ returns the set of optimal inductive solutions of $ILP_{LOAS}(T)$.

5 Experiments

Although there are benchmarks for ASP solvers (Denecker *et al.* 2009), there are no benchmarks for learning ASP programs. In (Law *et al.* 2014) we discussed the example of learning an ASP program with no weak constraints, representing the rules of sudoku. Using the examples from the paper and a small search space with only 283 rules, the original ILASP algorithm takes 486.2s to solve the task. This is due to the scalability issues discussed in section 4 as there are 332437 violating hypotheses found before the first inductive solution. For the same task with ILASP2, there are only 9 violating reasons found before the first inductive solution, meaning that ILASP2 takes only 0.69s to solve the task.

As this is the first work on learning weak constraints, there are no existing benchmarks suitable for testing our approach of learning from ordered answer sets. We have, therefore, further investigated the interview scheduling example discussed throughout the paper. Our experiments, in particular, test whether ILP_{LOAS} can successfully learn weak constraints from examples of brave and cautious orderings. For the purpose of presentation, we assume our hypothesis space, S_M , to be defined by the mode declarations: $M_h = M_b = \emptyset$; $M_o = \{\text{assign}(v, v), \text{neq}(v, v), \text{type}(v, c)\}$; $M_w = \{-1, 1\}$; and finally, $l_{max} = 2$. We place several restrictions on the search space in order to remove equivalent rules. The size of S_M is 184 (our hypotheses can be any subset of these 184 rules, so even considering only hypotheses with up to 3 rules this gives over a million different hypotheses). The learning task uses background knowledge B from Example 3. As S_M only contains weak constraints, for any $H \subseteq S_M$, $AS(B \cup H) = AS(B)$. The learning tasks described in these experiments therefore correspond to learning to rank the answer sets of B .

For each experiment we randomly selected 100 hypotheses, each with between 1 to 3 weak constraints from S_M , omitting hypotheses that ranked all answer sets equally. The only atoms that vary in B are the `assign`'s. As there are 9 different slots, there are 2^9 answer sets of B (and many more partial interpretations which are extended by these answer sets). We say an example partial interpretation is *full* if it specifies the truth value of all 9 `assign` atoms, otherwise we describe the *fullness* as the percentage of the 9 atoms which are specified. In both experiments (for each of the 100 target hypotheses H_T), we generated ordered pairs of partial interpretations $o = \langle e_1, e_2 \rangle$ such that o was bravely respected. If o was also cautiously respected, then it was given as a cautious example (otherwise it was used as a brave example). In our first experiment we investigated the effect of varying the number of examples, and in the second we investigated the effects of varying the fullness of the examples.

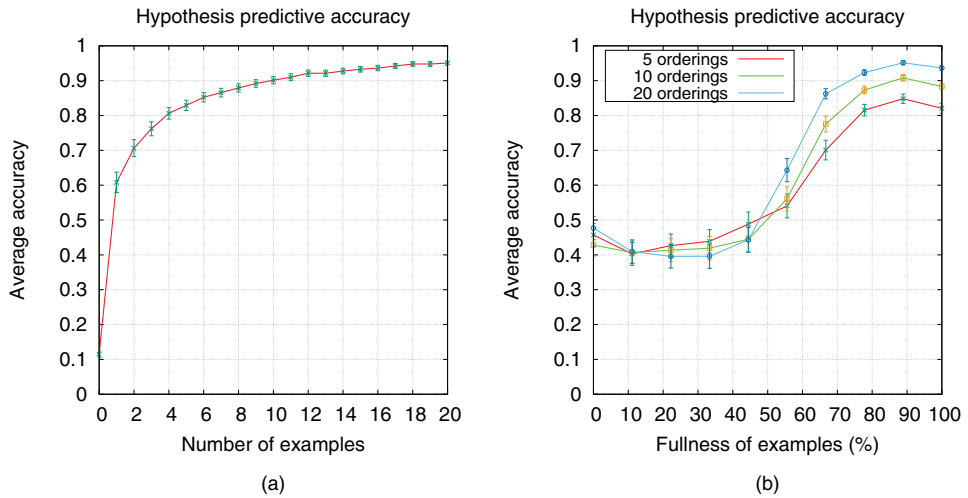


Fig. 1. Accuracy with varying (a) numbers of examples; (b) fullness of examples.

In both experiments, we tested our approach 20 times for each target hypothesis H_T . Each time, we used *ILASP2* to learn a hypothesis H_L which covered all examples. We then calculated the accuracy of H_L in predicting the pairwise ordering of answer sets in B (for each pair of answer sets $A_1, A_2 \in AS(B)$ we tested whether H_T and H_L agreed on the preference between them).

In our first experiment we investigated the effect of varying the number of examples from 0 to 20. The examples were of random fullness, each with between 5 to 9 assign atoms specified. Figure 1(a) shows the average predictive accuracy. Each point on the graph corresponds to 2000 learning tasks (100 target hypotheses with 20 different sets of examples). The error bars on the graph show the standard error. The results show that our method achieves 90% accuracy for this experiment with around 10 or more random examples.

For our second experiment we again tested our approach on 100 randomly generated hypotheses with 20 different sets of randomly generated examples. This time, however, we have kept the number of examples fixed at 5, 10 and 20 and varied the fullness of the examples. Results are shown in Figure 1(b). The graph shows that examples are only useful if they are more than 50% full. One interesting point to note is that the peak performance is with examples of around 90% fullness. This is because cautious ordering examples are actually more useful if they are less full (as there are more pairs which extend them); however, orderings are less likely to be cautiously respected when they are less full.

In our final experiment, we investigated the scalability of *ILASP2* by increasing both the number of days in our timetable and the number of examples. Figure 2 shows the average running time for *ILASP2* with 3, 4 and 5 day timetables (each with 3 slots) with up to 120 ordering examples. The learning tasks are targeted at learning the hypothesis from Example 3. We randomly generated ordering examples, as in the previous experiments with the slight difference that the fullness of the examples was unrestricted. As the hypothesis in these experiments does not use negative weights




Fig. 2. Average running time of ILASP2 with varying numbers of examples.

in either of the weak constraints, we also tested the average running time with a search space containing only positive weights. This means that S_M contained 92 weak constraints rather than the original 184. These experiments show that the time taken to solve an ILASP2 task is dependent not only on the number of examples, but also on the size of the domain and the size of S_M .

6 Related work

In (Law *et al.* 2014) we showed that any of the learning tasks in (Corapi *et al.* 2012; Ray 2009; Sakama and Inoue 2009; Otero 2001) could be expressed by ILP_{LAS} and computed by $ILASP$. As any ILP_{LAS} task can be (trivially) mapped into an ILP_{LOAS} (i.e. $O^b = \emptyset$ and $O^c = \emptyset$), ILP_{LOAS} inherits this property. None of the previous learning tasks (including ILP_{LAS}), however, can construct examples which incentivise the learning of a hypothesis containing a weak constraint. This is because they can only give examples of what should (or shouldn't) be an answer set of $B \cup H$. In addition, ILP_{LOAS} inherits the capability of ILP_{LAS} of supporting *predicate invention*, allowing new concepts to be invented whilst learning.

The ILASP2 algorithm is an extension of the original ILASP algorithm in (Law *et al.* 2014). It extends the concepts of positive and violating hypothesis to cover learning weak constraints (which was not possible in ILASP). For the simpler ILP_{LAS} tasks, ILASP2 is more efficient than ILASP. As discussed in section 4, the original ILASP algorithm has some scalability issues when there is a large number of violating hypotheses. We have shown in section 5 that by eliminating violating reasons rather than single violating hypotheses, ILASP2 can be much more efficient.

Also related to our work are existing approaches for *learning to rank*. These use non logic-based machine learning techniques (e.g. neural networks (Geisler *et al.* 2001)). Our approach shares the same advantages as any ILP approach versus a non logic-based machine learning technique: learned hypotheses are structured, human readable and can express relational concepts such as minimising the instances of particular combinations of predicates. Existing background knowledge can be taken into account to capture predefining concepts and the search can be steered towards particular types of hypotheses using a language bias. Furthermore, ILASP2 is also capable of learning preferences with weights and priorities, meaning that more structured and complex preferences can be learned.

An example of the use of an ILP system for learning constraints has been recently presented in (Lallouet *et al.* 2010) where timetabling constraints are learned from positive and negative examples. In this case the learned rules are hard constraints (e.g., enforcing that a teacher is not in two places at once). Examples of this kind are already computable by ILP_{LAS} , and so are also computable by ILP_{LOAS} .

7 Conclusion and future work

We have presented a new framework for ILP, Learning from Ordered Answer Sets, which extends previous ILP systems in that it is able to learn weak constraints and can be used to perform preference learning. The framework can represent partial examples under a brave and a cautious semantics. We have also put forward a new algorithm, ILASP2, that can solve any ILP_{LOAS} task for optimal solutions. This algorithm extends previous work for solving the simpler task ILP_{LAS} and resolves some of the scalability issues associated with the previous algorithm. Some scalability issues remain, especially when there is a particularly large hypothesis space and future work will focus on overcoming these. Current work also addresses extending the ILASP algorithm to support noisy examples.

References

- BANBARA, M., SOH, T., TAMURA, N., INOUE, K. AND SCHAUB, T. 2013. Answer set programming as a modeling language for course timetabling. *Theory and Practice of Logic Programming* 13, 4-5, 783–798.
- BLOCKEEL, H. AND DE RAEDT, L. 1998. Top-down induction of first-order logical decision trees. *Artificial intelligence* 101, 1, 285–297.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., RICCA, F. AND SCHAUB, T. 2013. ASP-Core-2 input language format. <https://www.mat.unical.it/aspcomp2013/files/ASP-CORE-2.0.pdf>.
- CORAPI, D., RUSSO, A. AND LUPU, E. 2010. Inductive logic programming as abductive search. In *ICLP (Technical Communications)*. 54–63.
- CORAPI, D., RUSSO, A. AND LUPU, E. 2012. Inductive logic programming in answer set programming. In *Inductive Logic Programming*. Springer, 91–97.
- DASTANI, M., JACOBS, N., JONKER, C. M. AND TREUR, J. 2001. Modeling user preferences and mediating agents in electronic commerce. In *Agent Mediated Electronic Commerce*. Springer, 163–193.
- DENECKER, M., VENNEKENS, J., BOND, S., GEBSER, M. AND TRUSZCZYŃSKI, M. 2009. The second answer set programming competition. In *Logic Programming and Nonmonotonic Reasoning*. Springer, 637–654.
- FÜRNKRANZ, J. AND HÜLLERMEIER, E. 2003. Pairwise preference learning and ranking. In *Machine Learning: ECML 2003*. Springer, 145–156.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T. AND SCHNEIDER, M. 2011. Potassco: The Potsdam answer set solving collection. *AI Communications* 24, 2, 107–124.
- GEISLER, B., HA, V. AND HADDAWY, P. 2001. Modeling user preferences via theory refinement. In *Proceedings of the 6th international conference on Intelligent user interfaces*. ACM, 87–90.

- HORVÁTH, T. 2012. A model of user preference learning for content-based recommender systems. *Computing and informatics* 28, 4, 453–481.
- KIMBER, T., BRODA, K. AND RUSSO, A. 2009. Induction on failure: learning connected horn theories. In *Logic Programming and Nonmonotonic Reasoning*. Springer, 169–181.
- LALLOUET, A., LOPEZ, M., MARTIN, L. AND VRAIN, C. 2010. On learning constraint problems. In *Tools with Artificial Intelligence (ICTAI), 2010 22nd IEEE International Conference on*. Vol. 1. IEEE, 45–52.
- LAW, M., RUSSO, A. AND BRODA, K. 2014. Inductive learning of answer set programs. In *Logics in Artificial Intelligence (JELIA 2014)*. Springer.
- LAW, M., RUSSO, A. AND BRODA, K. 2015a. The ILASP system for learning answer set programs. <https://www.doc.ic.ac.uk/~m11909/ILASP>.
- LAW, M., RUSSO, A. AND BRODA, K. 2015b. Proof of the soundness and completeness of ILASP2. https://www.doc.ic.ac.uk/~m11909/Proofs_for_ILASP2.pdf.
- LAW, M., RUSSO, A. AND BRODA, K. 2015c. Simplified reduct for choice rules in ASP. Tech. Rep. DTR2015-2, Imperial College of Science, Technology and Medicine, Department of Computing.
- MUGGLETON, S. 1991. Inductive logic programming. *New generation computing* 8, 4, 295–318.
- MUGGLETON, S., DE RAEDT, L., POOLE, D., BRATKO, I., FLACH, P., INOUE, K. AND SRINIVASAN, A. 2012. ILP turns 20. *Machine Learning* 86, 1, 3–23.
- MUGGLETON, S. AND LIN, D. 2013. Meta-interpretive learning of higher-order dyadic datalog: Predicate invention revisited. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 1551–1557.
- OTERO, R. P. 2001. Induction of stable models. In *Inductive Logic Programming*. Springer, 193–205.
- RAY, O. 2009. Nonmonotonic abductive inductive learning. *Journal of Applied Logic* 7, 3, 329–340.
- RAY, O., BRODA, K. AND RUSSO, A. 2004. A hybrid abductive inductive proof procedure. *Logic Journal of IGPL* 12, 5, 371–397.
- SAKAMA, C. AND INOUE, K. 2009. Brave induction: a logical framework for learning from incomplete information. *Machine Learning* 76, 1, 3–35.
- SRINIVASAN, A. 2001. The aleph manual. *Machine Learning at the Computing Laboratory, Oxford University*.