

# *Iterative Learning of Answer Set Programs from Context Dependent Examples*

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 6 May 2016; revised 8 July 2016; accepted 22 August 2016*

---

## Abstract

In recent years, several frameworks and systems have been proposed that extend Inductive Logic Programming (ILP) to the Answer Set Programming (ASP) paradigm. In ILP, examples must all be explained by a hypothesis together with a given background knowledge. In existing systems, the background knowledge is the same for all examples; however, examples may be context-dependent. This means that some examples should be explained in the context of some information, whereas others should be explained in different contexts. In this paper, we capture this notion and present a context-dependent extension of the *Learning from Ordered Answer Sets* framework. In this extension, contexts can be used to further structure the background knowledge. We then propose a new iterative algorithm, ILASP2i, which exploits this feature to scale up the existing ILASP2 system to learning tasks with large numbers of examples. We demonstrate the gain in scalability by applying both algorithms to various learning tasks. Our results show that, compared to ILASP2, the newly proposed ILASP2i system can be two orders of magnitude faster and use two orders of magnitude less memory, whilst preserving the same average accuracy.

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

---

## 1 Introduction

Inductive Logic Programming (Muggleton 1991) (ILP) addresses the task of learning a logic program, called a *hypothesis*, that explains a set of examples using some background knowledge. Although ILP has traditionally addressed learning (monotonic) definite logic programs, recently, several new systems have been proposed for learning under the (non-monotonic) answer set semantics (e.g. (Ray 2009), (Corapi *et al.* 2012), (Athakravi *et al.* 2014), (Law *et al.* 2014) and (Law *et al.* 2015)). Among these, ILASP2 (Law *et al.* 2015) extended ILP to *learning from ordered answer sets* ( $ILP_{LOAS}$ ), a computational task that learns answer set programs containing normal rules, choice rules and both hard and weak constraints.

\* This research is partially funded by the EPSRC project EP/K033522/1 “Privacy Dynamics”.

Common to all ILP systems is the underlying assumption that hypotheses should cover the examples with respect to one fixed given background knowledge. But, in practice, some examples may be context-dependent – different examples may need to be covered using different background knowledges. For instance, within the problem domain of urban mobility, the task of learning journey preferences of people in a city may require a general background knowledge that describes the different modes of transport available to a user (walk, drive, etc.), and examples of which modes of transport users choose for particular journeys. In this case, the context of an example would be the attributes (e.g. the distance) of the journey. It is infeasible to assume that every possible journey could be encoded in the background knowledge – attributes, such as journey distances, may take too many possible values. But, encoding the attributes of observed journeys as contexts of the observations restricts the computation to those attribute values that are in the contexts.

In this paper, we present a generalisation of  $ILP_{LOAS}$ , called *context-dependent learning from ordered answer sets* ( $ILP_{LOAS}^{context}$ ), which uses *context-dependent* examples. We show that any  $ILP_{LOAS}^{context}$  task can be translated into an  $ILP_{LOAS}$  task, and can therefore be solved by ILASP2. Furthermore, to improve the scalability of ILASP2, we present a new iterative reformulation of this learning algorithm, called *ILASP2i*. This iterative approach differs from existing non-monotonic learning systems, which tend to be batch learners, meaning that they consider all examples at once. Non-monotonic systems cannot use a traditional cover loop (e.g., (Muggleton 1995)), as examples that were covered in previous iterations are not guaranteed to be covered in later iterations. However, ILASP2i iteratively computes a hypothesis by constructing a set of examples that are *relevant* to the search, without the need to consider all examples at once. Relevant examples are essentially counterexamples for the hypotheses found in previous iterations. This approach is a middle ground between batch learning and the cover loop: it avoids using the whole set of examples, but works in the non-monotonic case, as the relevant examples persist through the iterations. We show that ILASP2i performs significantly better than ILASP2 in solving learning from ordered answer set tasks with large numbers of examples, and better still when learning with context-dependent examples, as in each iteration it only considers the contexts of relevant examples, rather than the full set.

To demonstrate the increase in scalability we compare ILASP2i to ILASP2 on a variety of tasks from different problem domains. The results show that ILASP2i is up to 2 orders of magnitude faster and uses up to 2 orders of magnitude less memory than ILASP2. We have also applied both algorithms to the real-world problem domain of urban mobility, and explored in greater depth the task of learning a user's journey preferences from pairwise examples of which journeys are preferred to others. As we learn ASP, these user preferences can very naturally be represented as weak constraints, which give an ordering over the journeys. Our results show that ILASP2i achieves an accuracy of at least 85% with around 40 examples. We also show that, by further extending  $ILP_{LOAS}^{context}$  with ordering examples that express equal preferences, in addition to strict ordering, the accuracy can increase to 93%.

The rest of the paper is structured as follows. In Section 2 we review the relevant background. In Section 3 we present our new context-dependent learning from

ordered answer set task, and in Section 4 we introduce our new ILASP2i algorithm. In Section 5 we compare ILASP2i to ILASP2 on a range of different learning tasks and give a detailed evaluation of the accuracy of ILASP2i and compare its scalability with ILASP2 in the context of the journey planning problem. Finally, we conclude the paper with a discussion of related and future work.

## 2 Background

Let  $h, h_1, \dots, h_k, b_1, \dots, b_n$  be atoms and  $l$  and  $u$  be integers. The ASP programs we consider contain normal rules, of the form  $h :- b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n$ ; constraints, which are rules of the form  $:- b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n$ ; and choice rules, of the form  $l\{h_1, \dots, h_k\}u :- b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n$ . We refer to the part of the rule before the “:-” as the head, and the part after the “:-” as the body. The meaning of a rule is that if the body is true, then the head must be true. The empty head of a constraint means *false*, and constraints are used to rule out answer sets. The head of a choice rule is true if between  $l$  and  $u$  atoms from  $h_1, \dots, h_k$  are true. The solutions of an ASP program  $P$  form a subset of the Herbrand models of  $P$ , called the *answer sets* of  $P$  and denoted as  $AS(P)$ .

ASP also allows optimisation over the answer sets according to *weak constraints*, which are rules of the form  $:\sim b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n. [w@p, t_1, \dots, t_k]$  where  $b_1, \dots, b_n$  are atoms called (collectively) the *body* of the rule, and  $w, p, t_1 \dots t_k$  are all terms with  $w$  called the weight and  $p$  the priority level. We will refer to  $[w@p, t_1, \dots, t_k]$  as the *tail* of the weak constraint. A ground instance of a weak constraint  $W$  is obtained by replacing all variables in  $W$  (including those in the tail of  $W$ ) with ground terms. In this paper, it is assumed that all weights and levels of all ground instances of weak constraints are integers.

Given a program  $P$  and an interpretation  $I$  we can construct the set of tuples  $(w, p, t_1, \dots, t_k)$  such that there is a ground instance of a weak constraint in  $P$  whose body is satisfied by  $I$  and whose (ground) tail is  $[w@p, t_1, \dots, t_k]$ . At each level  $p$  the *score* of  $I$  is the sum of the weights of tuples with level  $p$ . An interpretation  $I_1$  *dominates* another interpretation  $I_2$  if there is a level  $p$  for which  $I_1$  has a lower score than  $I_2$ , and no level higher than  $p$  for which the scores of  $I_1$  and  $I_2$  are unequal. We write  $I_1 <_p I_2$  to denote that given the weak constraints in  $P$ ,  $I_1$  dominates  $I_2$ .

*Example 1*

Consider the set  $WS = \begin{cases} :\sim \text{mode}(L, \text{walk}), \text{crime\_rating}(L, R), R > 3. [1@3, L, R] \\ :\sim \text{mode}(L, \text{bus}). [1@2, L] \\ :\sim \text{mode}(L, \text{walk}), \text{distance}(L, D). [D@1, L, D] \end{cases}$

The first weak constraint in  $WS$ , at priority 3, means “minimise the number of legs in our journey in which we have to walk through an area with a crime rating higher than 3”. As this has the highest priority, answer sets are evaluated over this weak constraint first. The remaining weak constraints are considered only for those answer sets that have an equal number of legs where we have to walk through an area with such a crime rating. The second weak constraint means “minimise the number of buses we have to take” (at priority 2). Finally, the last weak constraint

means “minimise the distance walked”. Note that this is the case because for each leg where we have to walk, we pay the penalty of the distance of that leg (so the total penalty at level 1 is the sum of the distances of the walking legs).

We now briefly summarise the key properties of Learning from Ordered Answer Sets and ILASP2, which we extend in this paper to Context-dependent Learning from Ordered Answer Sets and ILASP2i. It makes use of two types of examples: partial interpretations and ordering examples. A partial interpretation  $e$  is a pair of sets of atoms  $\langle e^{inc}, e^{exc} \rangle$ . An answer set  $A$  extends  $e$  if  $e^{inc} \subseteq A$  and  $e^{exc} \cap A = \emptyset$ . An ordering example is a pair of partial interpretations. A program  $P$  bravely (resp. cautiously) respects an ordering example  $\langle e_1, e_2 \rangle$  if for at least one (resp. every) pair of answer sets  $\langle A_1, A_2 \rangle$  that extend  $e_1$  and  $e_2$ , it is the case that  $A_1 <_P A_2$ .

*Definition 1*

(Law *et al.* 2015) A Learning from Ordered Answer Sets ( $ILP_{LOAS}$ ) task  $T$  is a tuple  $\langle B, S_M, E \rangle$  where  $B$  is an ASP program, called the background knowledge,  $S_M$  is the set of rules allowed in hypotheses (the hypothesis space) and  $E$  is a tuple  $\langle E^+, E^-, O^b, O^c \rangle$ .  $E^+$  and  $E^-$  are finite sets of partial interpretations called, respectively, positive and negative examples.  $O^b$  and  $O^c$  are finite sets of ordering examples over  $E^+$  called, respectively, brave and cautious orderings. A hypothesis  $H$  is an inductive solution of  $T$  (written  $H \in ILP_{LOAS}(T)$ ) iff:  $H \subseteq S_M$ ;  $\forall e \in E^+, \exists A \in AS(B \cup H)$  st  $A$  extends  $e$ ;  $\forall e \in E^-, \nexists A \in AS(B \cup H)$  st  $A$  extends  $e$ ;  $\forall o \in O^b, B \cup H$  bravely respects  $o$ ; and,  $\forall o \in O^c, B \cup H$  cautiously respects  $o$ .

In (Law *et al.* 2015), we proposed a learning algorithm, called ILASP2, and proved that it is sound and complete with respect to  $ILP_{LOAS}$  tasks. We use the notation  $ILASP2(\langle B, S_M, E \rangle)$  to denote a function that uses ILASP2 to return an optimal (shortest in terms of number of literals) solution of the task  $\langle B, S_M, E \rangle$ . ILASP2 terminates for any task such that  $B \cup S_M$  grounds finitely (or equivalently,  $\forall H \subseteq S_M, B \cup H$  grounds finitely). We call any such task *well defined*.

### 3 Context-dependent Learning from Ordered Answer Sets

In this section, we present an extension to the  $ILP_{LOAS}$  framework called *Context-dependent Learning from Ordered Answer Sets* (written  $ILP_{LOAS}^{context}$ ). In this new learning framework, examples can be given with an extra background knowledge called the *context* of an example. The idea is that each context only applies to a particular example, giving more structure to the background knowledge.

*Definition 2*

A context-dependent partial interpretation (CDPI) is a pair  $\langle e, C \rangle$ , where  $e$  is a partial interpretation and  $C$  is an ASP program with no weak constraints, called a *context*. A context-dependent ordering example (CDOE)  $o$  is a pair of CDPIs,  $\langle \langle e_1, C_1 \rangle, \langle e_2, C_2 \rangle \rangle$ . A program  $P$  is said to bravely (resp. cautiously) respect  $o$  if for at least one (resp. every) pair  $\langle A_1, A_2 \rangle$  such that  $A_1 \in AS(P \cup C_1), A_2 \in AS(P \cup C_2), A_1$  extends  $e_1$  and  $A_2$  extends  $e_2$ , it is the case that  $A_1 <_P A_2$ .

*Example 2*

Consider the programs  $P = \{\text{coin}(1..2). 1\{\text{val}(C,h), \text{val}(C,t)\}1:- \text{coin}(C).\}, C_1 = \{\text{val}(1,V):- \text{val}(2,V).\}$  and  $C_2 = \{:- \text{val}(1,V), \text{val}(2,V).\}$ .  $AS(P \cup C_1) = \{\{\text{val}(1,h), \text{val}(2,h)\}, \{\text{val}(1,t), \text{val}(2,t)\}\}$  and  $AS(P \cup C_2) = \{\{\text{val}(1,h), \text{val}(2,t)\}, \{\text{val}(1,t), \text{val}(2,h)\}\}$ . Also consider the CDOE  $o = \langle\langle e_1, C_1 \rangle, \langle e_2, C_2 \rangle\rangle$ , where  $e_1 = e_2 = \langle\emptyset, \emptyset\rangle$ . Let  $W = \{:\sim \text{val}(C,t). [1@1,C]\}$ .  $P \cup W$  bravely respects  $o$  as  $\{\text{val}(1,h), \text{val}(2,h)\}$  is preferred to  $\{\text{val}(1,h), \text{val}(2,t)\}$ , but does not cautiously respect  $o$  as  $\{\text{val}(1,t), \text{val}(2,t)\}$  is not preferred to  $\{\text{val}(1,h), \text{val}(2,t)\}$ .

Examples with empty contexts are equivalent to examples in  $ILP_{LOAS}$ . Note that contexts do not contain weak constraints. The operator  $\prec_P$  defines the ordering over two answer sets based on the weak constraints in one program  $P$ . So, given a CDOE  $\langle\langle e_1, C_1 \rangle, \langle e_2, C_2 \rangle\rangle$ , in which  $C_1$  and  $C_2$  contain different weak constraints, it is not clear whether the ordering should be checked using the weak constraints in  $P$ ,  $P \cup C_1$ ,  $P \cup C_2$  or  $P \cup C_1 \cup C_2$ . We now present the  $ILP_{LOAS}^{context}$  framework.

*Definition 3*

A *Context-dependent Learning from Ordered Answer Sets* ( $ILP_{LOAS}^{context}$ ) task is a tuple  $T = \langle B, S_M, E \rangle$  where  $B$  is an ASP program called the background knowledge,  $S_M$  is the set of rules allowed in the hypotheses (the hypothesis space) and  $E$  is a tuple  $\langle E^+, E^-, O^b, O^c \rangle$  called the examples.  $E^+$  and  $E^-$  are finite sets of CDPIs called, respectively, positive and negative examples, and  $O^b$  and  $O^c$  are finite sets of CDOEs over  $E^+$  called, respectively, brave and cautious orderings. A hypothesis  $H$  is an inductive solution of  $T$  (written  $H \in ILP_{LOAS}^{context}(T)$ ) if and only if:

1.  $H \subseteq S_M$ ;
2.  $\forall \langle e, C \rangle \in E^+, \exists A \in AS(B \cup C \cup H)$  st  $A$  extends  $e$ ;
3.  $\forall \langle e, C \rangle \in E^-, \nexists A \in AS(B \cup C \cup H)$  st  $A$  extends  $e$ ;
4.  $\forall o \in O^b, B \cup H$  bravely respects  $o$ ; and finally,
5.  $\forall o \in O^c, B \cup H$  cautiously respects  $o$ .

In this paper we will say a hypothesis *covers* an example iff it satisfies the appropriate condition in (2)-(5); e.g. a brave CDOE is covered iff it is bravely respected.

*Example 3*

In general, it is not the case that an  $ILP_{LOAS}^{context}$  task can be translated into an  $ILP_{LOAS}$  task simply by moving all the contexts into the background knowledge ( $B \cup C_1 \cup \dots \cup C_n$  where  $C_1, \dots, C_n$  are the contexts of the examples). Consider, for instance, the  $ILP_{LOAS}^{context}$  task  $\langle B, S_M, \langle E^+, E^-, O^b, O^c \rangle \rangle$  defined as follows:

- $B = \emptyset. E^- = \emptyset. O^b = \emptyset. O^c = \emptyset$
- $S_M = \{\text{go\_out}:- \text{raining}. \text{go\_out}:- \text{not raining}.\}$
- $E^+ = \{\langle\{\text{go\_out}\}, \emptyset, \emptyset\rangle, \langle\emptyset, \{\text{go\_out}\}, \{\text{raining}\}\rangle\}$

This task has one solution:  $\text{go\_out}:- \text{not raining}$ . But, if we were to add all the contexts to the background knowledge, we would get a background knowledge containing the single fact  $\text{raining}$ . So, there would be no way of explaining both examples, as every hypothesis would, in this case, lead to a single answer set (either  $\{\text{raining}, \text{go\_out}\}$  or  $\{\text{raining}\}$ ), and therefore cover only one of the examples.

To capture, instead, the meaning of context-dependent examples accurately, we could augment the background knowledge with the choice rule  $0\{\text{raining}\}1$  and define the  $ILP_{LOAS}$  examples as the pairs  $\langle\{\text{go\_out}\},\{\text{raining}\}\rangle$  and  $\langle\{\text{raining}\},\{\text{go\_out}\}\rangle$ . In this way, answer sets of the inductive solution would exclude  $\text{go\_out}$  when raining (i.e., in the context of raining), and include  $\text{go\_out}$  otherwise, which is the correct meaning of the given context-dependent examples.

Definition 4 gives a general translation of  $ILP_{LOAS}^{context}$  to  $ILP_{LOAS}$ , which enables the use of ILASP2 to solve  $ILP_{LOAS}^{context}$  tasks. The translation assumes that each example  $ex$  has a unique (constant) identifier,  $ex_{id}$ , and that for any CDPI  $ex = \langle\langle e^{inc}, e^{exc}\rangle, C\rangle$ ,  $c(ex)$  is the partial interpretation  $\langle e^{inc} \cup \{\text{ctx}(ex_{id})\}, e^{exc}\rangle$ , where  $\text{ctx}$  is a new predicate. Also, for any program  $P$  and any atom  $a$ ,  $\text{append}(P, a)$  is the program constructed by appending  $a$  to the body of every rule in  $P$ .

**Definition 4**

For any  $ILP_{LOAS}^{context}$  task  $T = \langle B_1, S_M, \langle E_1^+, E_1^-, O_1^b, O_1^c \rangle \rangle$ ,  $\mathcal{T}_{LOAS}(T) = \langle B_2, S_M, \langle E_2^+, E_2^-, O_2^b, O_2^c \rangle \rangle$ , where the components of  $\mathcal{T}_{LOAS}(T)$  are as follows:

- $B_2 = B_1 \cup \{\text{append}(C, \text{ctx}(ex_{id})) \mid ex = \langle e, C \rangle \in E_1^+ \cup E_1^-\}$   
 $\cup \{1\{\text{ctx}(id_1), \dots, \text{ctx}(id_n)\}1. \mid \{id_1, \dots, id_n\} = \{ex_{id} \mid ex \in E_1^+ \cup E_1^-\}\}$
- $E_2^+ = \{c(ex) \mid ex \in E_1^+\}; E_2^- = \{c(ex) \mid ex \in E_1^-\}$
- $O_2^b = \{\langle c(ex_1), c(ex_2) \rangle \mid \langle ex_1, ex_2 \rangle \in O_1^b\}; O_2^c = \{\langle c(ex_1), c(ex_2) \rangle \mid \langle ex_1, ex_2 \rangle \in O_1^c\}$

We say that an  $ILP_{LOAS}^{context}$  task  $T$  is well defined if and only if  $\mathcal{T}_{LOAS}(T)$  is a well defined  $ILP_{LOAS}$  task. Before proving that this translation is correct, it is useful to introduce a lemma (which is proven in Appendix A). Given a program  $P$  and a set of contexts  $C_1, \dots, C_n$ , Lemma 1 gives a way of combining the alternative contexts into the same program. Each rule of each context  $C_i$ , is appended with a new atom  $a_i$ , unique to  $C_i$ , and a choice rule stating that exactly one of the new  $a_i$  atoms is true in each answer set. This means that the answer sets of  $P \cup C_i$ , for each  $C_i$ , are the answer sets of the combined program that contain  $a_i$  (with the extra atom  $a_i$ ).

**Lemma 1**

For any program  $P$  (consisting of normal rules, choice rules and constraints) and any set of pairs  $S = \{\langle C_1, a_1 \rangle, \dots, \langle C_n, a_n \rangle\}$  such that none of the atoms  $a_i$  appear in  $P$  (or in any of the  $C$ 's) and each  $a_i$  atom is unique:  $AS(P \cup \{1\{a_1, \dots, a_n\}1.\} \cup \{\text{append}(C_i, a_i) \mid \langle C_i, a_i \rangle \in S\}) = \{A \cup \{a_i\} \mid A \in AS(P \cup C_i), \langle C_i, a_i \rangle \in S\}$

**Theorem 1**

For any  $ILP_{LOAS}^{context}$  learning task  $T$ ,  $ILP_{LOAS}(\mathcal{T}_{LOAS}(T)) = ILP_{LOAS}^{context}(T)$ .

**Proof**

Let  $T = \langle B_1, S_M, \langle E_1^+, E_1^-, O_1^b, O_1^c \rangle \rangle$  and  $\mathcal{T}_{LOAS}(T) = \langle B_2, S_M, \langle E_2^+, E_2^-, O_2^b, O_2^c \rangle \rangle$ .  
 $H \in ILP_{LOAS}^{context}(T) \Leftrightarrow H \subseteq S_M; \forall \langle e, C \rangle \in E_1^+, \exists A \in AS(B_1 \cup C \cup H)$  st  $A$  extends  $e$ ;  $\forall \langle e, C \rangle \in E_1^-, \nexists A \in AS(B_1 \cup C \cup H)$  st  $A$  extends  $e$ ;  $\forall o \in O_1^b, B_1 \cup H$  bravely respects  $o$ ;  $\forall o \in O_1^c, B_1 \cup H$  cautiously respects  $o$   
 $\Leftrightarrow H \subseteq S_M; \forall ex \in E_1^+, \exists A \in AS(B_2 \cup H)$  st  $A$  extends  $c(ex)$ ;  $\forall ex \in E_1^-, \nexists A \in AS(B_2 \cup H)$  st  $A$  extends  $c(ex)$ ;  $\forall \langle ex_1, ex_2 \rangle \in O^b, B_2 \cup H$  bravely respects  $\langle c(ex_1), c(ex_2) \rangle$ ;  $\forall \langle ex_1, ex_2 \rangle \in O^c, B_2 \cup H$  cautiously respects  $\langle c(ex_1), c(ex_2) \rangle$  (by Lemma 1)

$\Leftrightarrow H \subseteq S_M; \forall e \in E_2^+, \exists A \in AS(B_2 \cup H)$  st  $A$  extends  $e$ ;  $\forall e \in E_2^-, \nexists A \in AS(B_2 \cup H)$   
 st  $A$  extends  $e$ ;  $\forall o \in O^b, B_2 \cup H$  bravely respects  $o$ ;  $\forall o \in O^c, B_2 \cup H$  cautiously  
 respects  $o$  □

Theorem 1 shows that, by using an automatic  $\mathcal{T}_{LOAS}$  translation, ILASP2 can be used to solve  $ILP_{LOAS}^{context}$  tasks. Although this means that any  $ILP_{LOAS}^{context}$  task can be translated to an  $ILP_{LOAS}$  task, context-dependent examples are useful for two reasons: firstly, they simplify the representation of some learning tasks; and secondly, the added structure gives more information about which parts of the background knowledge apply to particular examples. In Section 4 we present a new algorithm that is able to take advantage of this extra information.

### Theorem 2

The complexity of deciding whether an  $ILP_{LOAS}^{context}$  task is satisfiable is  $\Sigma_2^P$ -complete.

Theorem 2 (proven in Appendix A) implies that the complexity of deciding the satisfiability of an  $ILP_{LOAS}^{context}$  task is the same as for an  $ILP_{LOAS}$  task. Note that, similar to Theorem 2 in (Law *et al.* 2015), this result is for propositional tasks.

## 4 Iterative Algorithm: ILASP2i

In the previous section, we showed that our new  $ILP_{LOAS}^{context}$  task can be translated into  $ILP_{LOAS}$ , and therefore solved using the ILASP2 algorithm (Law *et al.* 2015). However, ILASP2 may suffer from scalability issues, due to the number of examples or the size and complexity of the grounding of the hypothesis space, when combined with the background knowledge. In this paper, we address the first scalability issue by introducing a new algorithm, ILASP2i, for solving (context-dependent) learning from ordered answer sets tasks. The algorithm *iteratively* computes a hypothesis by incrementally constructing a subset of the examples that are *relevant* to the search. These are essentially counterexamples for incorrect hypotheses. The idea of the algorithm is to incrementally build, during the computation, a set of relevant examples and, at each iterative step, to learn hypotheses with respect only to this set of relevant examples instead of the full set of given examples. Although we do not directly address the second issue of large and complicated hypothesis spaces, it is worth noting that by using the notion of context-dependent examples, the size of the background knowledge (and therefore the grounding of the hypothesis space) in a particular iteration of our algorithm may be much smaller. In fact, in Section 5 we show that the background knowledge of one learning task (learning the definition of a Hamiltonian graph) can be eliminated altogether by using contexts.

### Definition 5

Consider an  $ILP_{LOAS}^{context}$  learning task  $T = \langle B, S_M, \langle E^+, E^-, O^b, O^c \rangle \rangle$  and a hypothesis  $H \subseteq S_M$ . A (context-dependent) example  $ex$  is *relevant* to  $H$  given  $T$  if  $ex \in E^+ \cup E^- \cup O^b \cup O^c$  and  $B \cup H$  does not cover  $ex$ .

The intuition of ILASP2i (Algorithm 1) is that we start with an empty set of *relevant* examples and an empty hypothesis. At each step of the search we look for

an example which is relevant to our current hypothesis (i.e. an example that  $B \cup H$  does not cover). If no such example exists, then we return our current hypothesis as an optimal inductive solution; otherwise, we add the example to our relevant set of examples and use ILASP2 to compute a new hypothesis.

The notation  $\ll$ , in line 5 of algorithm 1, means to add the relevant example  $re$  to the correct set in *Relevant* (the first set if it is a positive example etc).

---

**Algorithm 1** ILASP2i

---

```

1: procedure ILASP2i ( $\langle B, S_M, E \rangle$ )
2:   Relevant =  $\langle \emptyset, \emptyset, \emptyset, \emptyset \rangle$ ;  $H = \emptyset$ ;
3:    $re = \text{findRelevantExample}(\langle B, S_M, E \rangle, H)$ ;
4:   while  $re \neq \text{nil}$  do
5:     Relevant  $\ll re$ ;
6:      $H = \text{ILASP2}(\mathcal{T}_{LOAS}(\langle B, S_M, \text{Relevant} \rangle))$ ;
7:     if ( $H == \text{nil}$ ) return UNSATISFIABLE;
8:     else  $re = \text{findRelevantExample}(\langle B, S_M, E \rangle, H)$ ;
9:   end while
10:  return  $H$ ;
11: end procedure

```

---

The function  $\text{findRelevantExample}(\langle B, S_M, E \rangle, H)$  returns a (context-dependent) example in  $E$  which is not covered by  $B \cup H$ , or *nil* if no such example exists. It works by encoding  $B \cup H$  and  $E$  into a meta program whose answer sets can be used to determine which examples in  $E$  are covered. This meta program contains a choice rule, which specifies that each answer set of the program tests the coverage of a single CDPI or CDOE example. For a positive or negative example  $ex = \langle e, C \rangle$ , if there is an answer set of the meta program corresponding to  $ex$  then there must be at least one answer set of  $B \cup C \cup H$  that extends  $e$ . This means that positive (resp. negative) examples are covered iff there is at least one (resp. no) answer set of the meta program that corresponds to  $ex$ . Similarly, CDOE's are encoded such that each brave (resp. cautious) ordering  $o$  is respected iff there is at least one (resp. no) answer set corresponding to  $o$ .  $\text{findRelevantExamples}$  uses the answer sets of the meta program to determine which examples are not covered. Details of the meta program are in Appendix B, including proof of its correctness.

It should be noted that in the worst case our set of relevant examples is equal to the entire set of examples. In this case, ILASP2i is slower than ILASP2. In real settings, however, as examples are not carefully constructed, there is likely to be overlap between examples, so the relevant set will be much smaller than the whole set. Theorem 3 shows that ILASP2i has the same condition for termination as ILASP2.

*Theorem 3*

ILASP2i terminates for any well defined  $ILP_{LOAS}^{context}$  task.

Note that although the algorithm is sound, it is complete only in the sense that it always returns an optimal solution if one exists (rather than returning the full set).



Table 1. The running times of ILASP2, ILASP2i and ILASP2i\_pt. TO stands for time out (6 hours) and OOM stands for out of memory.

Learning task	#examples				time/s			Memory/kB		
	$E^+$	$E^-$	$O^b$	$O^c$	2	2i_pt	2i	2	2i_pt	2i
Hamilton A	100	100	0	0	10.3	4.2	4.3	$9.7 \times 10^4$	$1.2 \times 10^4$	$1.2 \times 10^4$
Hamilton B	100	100	0	0	32.0	84.9	3.6	$3.6 \times 10^5$	$2.7 \times 10^5$	$1.4 \times 10^4$
Scheduling A	400	0	110	90	291.9	64.2	63.4	$2.7 \times 10^6$	$1.7 \times 10^5$	$1.7 \times 10^5$
Scheduling B	400	0	128	72	347.2	40.1	40.3	$5.2 \times 10^6$	$2.6 \times 10^5$	$2.6 \times 10^5$
Scheduling C	400	0	133	67	1141.8	123.6	124.2	$8.4 \times 10^6$	$4.9 \times 10^5$	$5.0 \times 10^5$
Agent A	200	0	0	0	444.5	56.7	39.1	$4.7 \times 10^6$	$3.7 \times 10^5$	$9.8 \times 10^4$
Agent B	50	0	0	0	TO	212.3	9.4	TO	$1.1 \times 10^6$	$1.8 \times 10^5$
Agent C	80	120	0	0	808.7	132.3	60.1	$2.9 \times 10^6$	$3.5 \times 10^5$	$8.4 \times 10^4$
Agent D	172	228	390	0	OOM	863.3	408.4	OOM	$2.4 \times 10^6$	$8.0 \times 10^5$

#### Theorem 4

ILASP2i is sound for any well defined  $ILP_{LOAS}^{context}$  task, and returns an optimal solution if one exists.

Note that in Algorithm 1 the translation of a context-dependent learning task is applied to the context-dependent task generated incrementally at each step of the iteration (see line 6) instead of pre-translating the full initial task. This has the advantage that the background knowledge of the translated task only contains the contexts for the relevant examples, rather than the full set. In Section 5 we compare the efficiency of ILASP2i on  $ILP_{LOAS}^{context}$  tasks that have been pre-translated with corresponding tasks that have not been pre-translated, and demonstrate that in the latter case ILASP2i can be up to one order of magnitude faster. We refer to the application of ILASP2i with an automatic pre-translation to  $ILP_{LOAS}$  as ILASP2i\_pt.

## 5 Evaluation

In this section, we demonstrate the improvement in performance of ILASP2i over ILASP2, both in terms of running time and memory usage. Although there are benchmarks for ASP solvers, as ILP systems for ASP are relatively new, and solve different computational tasks, there are no benchmarks for learning ASP programs. We therefore investigate new problems. To demonstrate the increased performance of ILASP2i over ILASP2, we chose tasks with large numbers of examples. We compare the algorithms in four problem settings, each including tasks requiring different components of the  $ILP_{LOAS}^{context}$  framework. We also investigate how the performance and accuracy vary with the number of examples, for the task of learning user journey preferences. All learning tasks were run with ILASP2, ILASP2i and ILASP2i\_pt<sup>1</sup>.

Our first problem setting is learning the definition of whether a graph is Hamiltonian or not (i.e. whether it contains a Hamilton cycle). Hamilton A is an  $ILP_{LOAS}$  (non context-dependent) task. The background knowledge  $B$  consists of the two choice rules  $1 \{ \text{node}(1), \text{node}(2), \text{node}(3), \text{node}(4) \} 4$  and  $0 \{ \text{edge}(N1, N2) \} 1 :-$

<sup>1</sup> For details of the tasks discussed in this section and how to download and run ILASP2, ILASP2i and ILASP2i\_pt, see <https://www.doc.ic.ac.uk/~m11909/ILASP>.

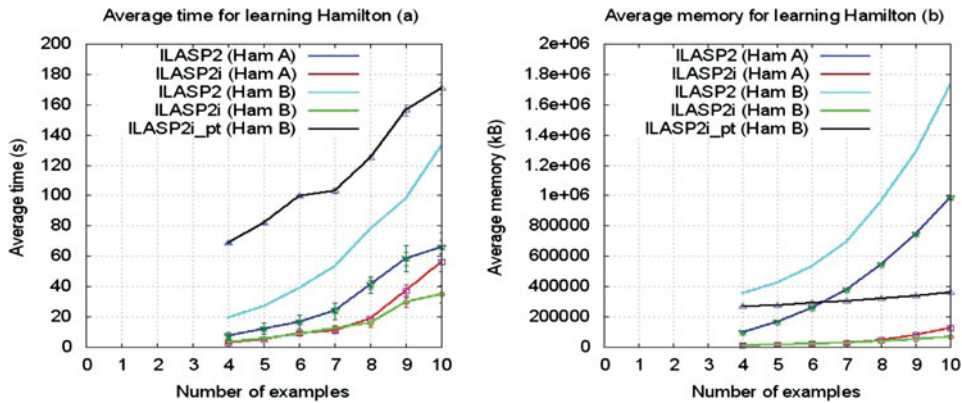


Fig. 1. (a) the average computation time and (b) the memory usage of ILASP2, ILASP2i and ILASP2i\_pt for Hamilton A and B.

node(N1),node(N2), meaning that the answer sets of  $B$  correspond to the graphs of size 1 to 4. Each example then corresponds to exactly one graph, by specifying which node and edge atoms should be true. Positive examples correspond to Hamiltonian graphs, and negative examples correspond to non-Hamiltonian graphs. Hamilton B is an  $ILP_{LOAS}^{context}$  encoding of the same problem. The background knowledge is empty, and each example has a context consisting of the node and edge atoms representing a single graph. ILASP2i performs significantly better than ILASP2 in both cases. Although ILASP2i is slightly faster at solving Hamilton B compared with Hamilton A, one interesting result is that ILASP2 and ILASP2i\_pt perform better on Hamilton A. This is because the non context-dependent encoding in Hamilton A is more efficient than the automatic translation (using definition 4) of Hamilton B.

To test how the size of the contexts affects the performance of the three algorithms, we reran the Hamilton A and B experiments with the maximum size of the graphs varying from 4 to 10. Each experiment was run 100 times with randomly generated sets of positive and negative examples (100 of each in each experiment). The results (figure 1) show that ILASP2i performs best in both cases - interestingly, on average, there is no difference between Hamilton A (non context-dependent) and Hamilton B (context-dependent) at first, but as the maximum graph size increases, the domain of the background knowledge in Hamilton A increases and so ILASP2i performs better on Hamilton B. Although ILASP2i\_pt is much slower on Hamilton B than Hamilton A, it uses significantly less memory on the former. As the performance of ILASP2i and ILASP2i\_pt is the same on any non context-dependent task, we do not show the results for ILASP2i\_pt on Hamilton A.

We also reconsider the problem of learning scheduling preferences, first presented in (Law *et al.* 2015). In this setting, the goal is to learn an academic's preferences about interview scheduling, encoded as weak constraints. Tasks A-C in this case are over examples with 3x3, 4x3 and 5x3 timetables, respectively. As this setting contains no contexts for the examples, the performance of ILASP2i and ILASP2i\_pt are relatively similar; however, for larger timetables both are over an order of magnitude faster and use over an order of magnitude less memory than ILASP2.

Interestingly, although ILASP2i does not directly attempt to scale up the size of possible problem domains (in this case, the dimensions of the timetables), this experiment demonstrates that ILASP2i does (indirectly) improve the performance on larger problem domains. One unexpected observation is that ILASP2i runs faster on task B than task A. This is caused by the algorithm choosing “better” relevant examples for task B, and therefore needing a smaller set of relevant examples. On average, the time for 4x3 timetables would be expected to be higher than the 3x3's.

Our third setting is taken from (Law *et al.* 2014) and is based on an agent learning the rules of how it is allowed to move within a grid. Agent A requires a hypothesis describing the concept of which moves are valid, given a history of where an agent has been. Agent B requires a similar hypothesis to be learned, but with the added complexity that an additional concept is required to be invented. While Agent A and Agent B are similar to scenarios 1 and 2 in (Law *et al.* 2014), the key difference is that different examples contain different histories of where the agent has been. These histories are encoded as contexts, whereas in (Law *et al.* 2014), one single history was encoded in the background knowledge. There are also many more examples in these experiments. In Agent C, the hypothesis from Agent A must be learned along with a constraint ruling out histories in which the agent visits a cell twice (not changing the definition of valid move). This requires negative examples to be given, in addition to positive examples. In Agent D, weak constraints must be learned to explain why some traces through the grid are preferred to others. This uses positive, negative and brave ordering examples. In each case, ILASP2i performs significantly better than ILASP2i\_pt, which performs significantly better than ILASP2 (ILASP2 times out in one experiment, and runs out of memory in another).

In our final setting, we investigate the problem of learning a user's preferences over alternative journeys, in order to demonstrate how the performance of the three algorithms varies with the number of examples. We also investigate how the accuracy of ILASP2i varies with the number of examples. In this scenario, a user makes requests to a journey planner to get from one location to another. The user then chooses a journey from the alternatives returned by the planner. A journey consists of one or more legs, in each of which the user uses a single mode of transport.

We used a simulation environment (Poxrucker *et al.* 2014) to generate realistic examples of journeys. In our experiment, we ran the simulator for one (simulated) day to generate a set of journey requests, along with the attributes of each possible journey. The attributes provided by the simulation data are: *mode*, which takes the value bus, car, walk or bicycle; *distance*, which takes an integer value between 1 and 20000; and *crime\_rating*. As the crime ratings were not readily available from the simulator, we used a randomly generated value between 1 and 5.

For our experiments, we assume that the user's preferences can be represented by a set of weak constraints based on the attributes of a leg. We constructed a set of possible weak constraints, each including at most 3 literals. Most of these literals capture the leg's attributes, e.g., *mode(L, bus)* or *crime\_rating(L, R)* (if the attribute's values range over integers this is represented by a variable, otherwise each possible value is used as a constant). For the crime rating (*crime\_rating(L, R)*), we also allow

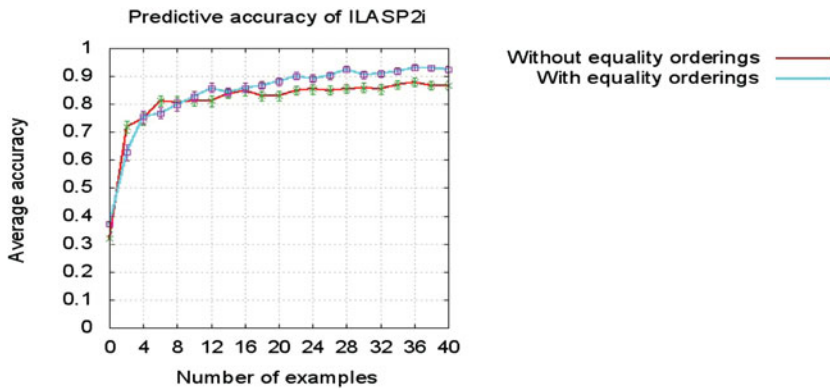


Fig. 2. average accuracy of ILASP2i

comparisons of the form  $R > c$  where  $c$  is an integer from 1 to 4. The weight of each weak constraint is a variable representing the distance of the leg in the rule, or 1 and the priority is 1, 2 or 3. One possible set of preferences is the set of weak constraints in Example 1.  $S_J$  denotes the set of possible weak constraints.

We now describe how to represent the journey preferences scenario in  $ILP_{LOAS}^{context}$ . We assume a journey is encoded as a set of attributes of the legs of the journey; for example the journey  $\{\text{distance}(\text{leg}(1), 2000), \text{distance}(\text{leg}(2), 100), \text{mode}(\text{leg}(1), \text{bus}), \text{mode}(\text{leg}(2), \text{walk})\}$  has two legs; in the first leg, the person must take a bus for 2000m and in the second, he/she must walk 100m. Given a set of such journeys  $J = \{j_1, \dots, j_n\}$  and a partial ordering  $O$  over  $J$ ,  $\mathcal{M}(J, O, S_J)$  is the  $ILP_{LOAS}^{context}$  task  $\langle \emptyset, S_J, E^+, \emptyset, O^b, \emptyset \rangle$ , where  $E^+ = \{\langle \langle \emptyset, \emptyset \rangle, j_i \rangle \mid j_i \in J\}$  and  $O^b = \{\langle \langle \langle \emptyset, \emptyset \rangle, j_1 \rangle, \langle \langle \emptyset, \emptyset \rangle, j_2 \rangle \rangle \mid \langle j_1, j_2 \rangle \in O\}$ . Each solution of  $\mathcal{M}(J, O, S_J)$  is a set of weak constraints representing preferences which explain the ordering of the journeys. Note that the positive examples are automatically satisfied as the (empty) background knowledge (combined with the context) already covers them. Also, as the background knowledge together with each context has exactly one answer set, the notions of brave and cautious orderings coincide; hence, we do not need cautious ordering examples for this task. Furthermore, since we are only learning weak constraints, and not hard constraints, the task also has no negative examples (a negative example would correspond to an invalid journey).

In each experiment we randomly selected 100 test hypotheses, each consisting of between 1 and 3 weak constraints from  $S_J$ . For each test hypothesis  $H_T$ , we then used the simulated journeys to generate a set of ordering examples  $\langle j_1, j_2 \rangle$  such that  $j_1$  was one of the optimal journeys, given  $H$ , and  $j_2$  was a non-optimal alternative to  $j_1$ . We then tested the algorithms on tasks with varying numbers of ordering examples by taking a random sample of the complete set of ordering examples.

The accuracy of ILASP2i for different numbers of examples is shown in Figure 2. The average accuracy converges to around 85% after roughly 20 examples. As we only gave examples of journeys such that one was preferred to the other, the hypotheses were often incorrect at predicting that two journeys were equal. We therefore introduced a new type of brave ordering example to ILASP2i, which

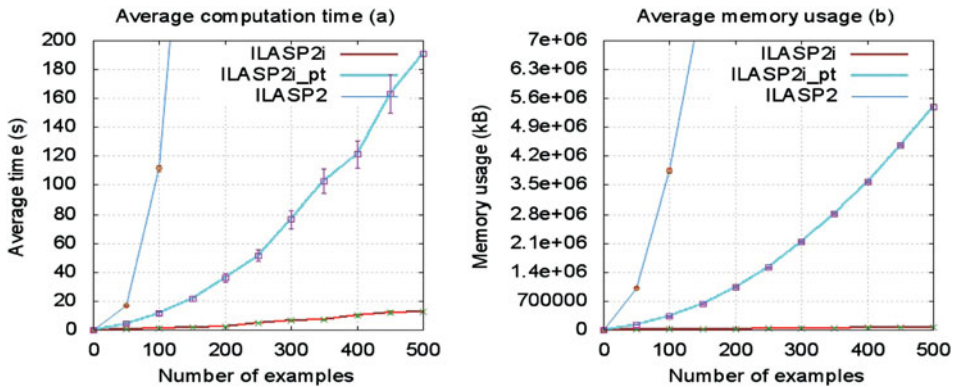


Fig. 3. (a) the average computation time and (b) the memory usage of ILASP2, ILASP2i and ILASP2i\_pt for learning journey preferences.

enables us to specify that two answer sets should be equally optimal. We ran the same experiment with half of the ordering examples as the new “equality” orderings. The average accuracy increased to around 93% after 40 examples. Note that as ILASP2 and ILASP2i return an arbitrary optimal solution of a task, their accuracy results, on average, are the same. We therefore only present the results for ILASP2i.

Figures 3(a) and (b) show the running times and memory usage (respectively) for up to 500 examples for ILASP2, ILASP2i and ILASP2i\_pt. For experiments with more than 200 examples, ILASP2 ran out of memory. By 200 examples, ILASP2i is already over 2 orders of magnitude faster and uses over 2 orders of magnitude less memory than ILASP2, showing a significant improvement in scalability. The fact that by 500 examples ILASP2i is an order of magnitude faster without the pre-translation shows that, in this problem domain, the context is a large factor in this improvement; however, ILASP2i\_pt’s significantly improved performance over ILASP2 shows that the iterative nature of ILASP2i is also a large factor.

## 6 Related Work

Most approaches to ILP address the learning of definite programs (Srinivasan 2001; Muggleton *et al.* 2014), usually aiming to learn Prolog programs. As the language features of Prolog and ASP are different (e.g. ASP lacks lists, Prolog lacks choice), a comparison is difficult. On the shared language of ASP and the fragment of Prolog learned by these systems (definite rules), a traditional ILP task can be represented with a single positive example (where the inclusions (resp. exclusions) of this example correspond to the positive (resp. negative) examples in the original task).

The idea of context-dependent example has similarities with the concept of *learning from interpretation transitions* (LFIT) (Inoue *et al.* 2014), where examples are pairs of set of atoms  $\langle I, J \rangle$  such that  $B \cup H$  must satisfy  $T_{B \cup H}(I) = J$  (where  $T_P(I)$  is the set of immediate consequences of  $I$  with respect to the program  $P$ ). LFIT technically learns under the supported model semantics and uses a far smaller

language than that supported by  $ILP_{LOAS}^{context}$  (not supporting choice rules or hard or weak constraints), but can be simply represented in  $ILP_{LOAS}^{context}$ . The head  $h$  of each rule in the background knowledge and hypothesis space should be replaced by  $j(h)$ , and each body literal  $b$ , by  $i(b)$ . Each example  $\langle I, J \rangle$  should then be mapped to a context-dependent positive example  $\langle \{j(a) \mid a \in J\}, \emptyset, \{i(a) \mid a \in I\} \rangle$ .

Other than our own frameworks, the two main ILP frameworks under the answer set semantics are *brave* and *cautious* induction (Sakama and Inoue 2009). As  $ILP_{LOAS}^{context}$  subsumes  $ILP_{LOAS}$ ,  $ILP_{LOAS}^{context}$  inherits the ability to perform both brave and cautious induction. ILASP2i is therefore more general than systems such as (Ray 2009; Corapi *et al.* 2012; Athakravi *et al.* 2014), which can only perform brave induction. In ILP, learners can be divided into batch learners (those which consider all examples simultaneously), such as (Ray 2009; Corapi *et al.* 2012; Athakravi *et al.* 2014; Law *et al.* 2014), and learners which consider each example in turn (using a cover loop), such as (Srinivasan 2001; Muggleton 1995; Ray *et al.* 2003). Under the answer set semantics, most learners are batch learners due to the non-monotonicity. In fact, it is worth noting that, in particular, although the HAIL (Ray *et al.* 2003) algorithm for learning definite clauses employs a cover loop, the later XHAIL algorithm is a batch learner as it learns non-monotonic programs (Ray 2009). One approach which did attempt to utilise a cover loop is (Sakama 2005). Their approach, however, was only sound for a small (monotonic) fragment of ASP if the task had multiple examples, as otherwise later examples could cause earlier examples to become uncovered.

The ILED system (Katzouris *et al.* 2015) extended the ideas behind XHAIL in order to allow incremental learning of event definitions. This system takes as input, multiple “windows” of examples and incrementally learns a hypothesis. As the approach is based on theory revision (at each step, revising the hypothesis from the previous step), ILED is not guaranteed to learn an optimal solution. In contrast, ILASP2i incrementally builds a set of relevant examples and learns a new hypothesis at each iteration. In so doing it is guaranteed to compute an optimal solution.

## 7 Conclusion

In this paper, we have presented an extension to our  $ILP_{LOAS}$  framework which allows examples to be given with extra background knowledge called the context of the example. We have shown that these contexts can be used to give structure to the background knowledge, showing which parts apply to which examples. We have also presented a new algorithm, ILASP2i, which makes use of this added structure to improve the efficiency over the previous ILASP2. In Section 5, we demonstrated that our new approach is considerably faster for tasks with large numbers of examples.

Unlike previous systems for learning under the answer set semantics, ILASP2i is not a batch learner and does not need to consider all examples at the same time, but instead iteratively builds a set of relevant examples. This combination of relevant examples and the added structure given by contexts means that ILASP2i can be up to 2 orders of magnitude better than ILASP2, both in terms of time and memory

usage. In future work, we intend to investigate how to improve the scalability of ILASP2i with larger hypothesis spaces and with noisy examples.

### Supplementary material

For supplementary material for this article, please visit <http://dx.doi.org/10.1017/S1471068416000351>

### References

- ATHAKRAVI, D., CORAPI, D., BRODA, K. AND RUSSO, A. 2014. Learning through hypothesis refinement using answer set programming. In *Inductive Logic Programming*. Springer, 31–46.
- CORAPI, D., RUSSO, A. AND LUPU, E. 2012. Inductive logic programming in answer set programming. In *Inductive Logic Programming*. Springer, 91–97.
- INOUE, K., RIBEIRO, T. AND SAKAMA, C. 2014. Learning from interpretation transition. *Machine Learning* 94, 1, 51–79.
- KATZOURIS, N., ARTIKIS, A. AND PALIOURAS, G. 2015. Incremental learning of event definitions with inductive logic programming. *Machine Learning* 100, 2-3, 555–585.
- LAW, M., RUSSO, A. AND BRODA, K. 2014. Inductive learning of answer set programs. In *Logics in Artificial Intelligence (JELIA 2014)*. LNAI, vol. 8761. Springer.
- LAW, M., RUSSO, A. AND BRODA, K. 2015. Learning weak constraints in answer set programming. *Theory and Practice of Logic Programming* 15, 4-5, 511–525.
- MUGGLETON, S. 1991. Inductive logic programming. *New Generation Computing* 8, 4, 295–318.
- MUGGLETON, S. 1995. Inverse entailment and Prolog. *New Generation Computing* 13, 3-4, 245–286.
- MUGGLETON, S. H., LIN, D., PAHLAVI, N. AND TAMADDONI-NEZHAD, A. 2014. Meta-interpretive learning: application to grammatical inference. *Machine Learning* 94, 1, 25–49.
- POXRUCKER, A., BAHLE, G. AND LUKOWICZ, P. 2014. Towards a real-world simulator for collaborative distributed learning in the scenario of urban mobility. In *Proceedings of the Eighth IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops*. IEEE Computer Society, 44–48.
- RAY, O. 2009. Nonmonotonic abductive inductive learning. *Journal of Applied Logic* 7, 3, 329–340.
- RAY, O., BRODA, K. AND RUSSO, A. 2003. Hybrid abductive inductive learning: A generalisation of Prolog. In *Inductive Logic Programming*. Springer, 311–328.
- SAKAMA, C. 2005. Induction from answer sets in nonmonotonic logic programs. *ACM Transactions on Computational Logic (TOCL)* 6, 2, 203–231.
- 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*.