

*Efficient Lifting of Symmetry Breaking Constraints for Complex Combinatorial Problems**

ALICE TARZARIOL and KONSTANTIN SCHEKOTIHIN

University of Klagenfurt, Klagenfurt, Austria

(e-mails: alice.tarzariol@aau.at, konstantin.schekotihin@aau.at)

MARTIN GEBSER

University of Klagenfurt, Klagenfurt, Austria Graz University of Technology, Graz, Austria

(e-mail: martin.gebser@aau.at)

MARK LAW

Imperial College London, London, UK

(e-mail: mark@ilasp.com)

submitted 16 May 2022; accepted 8 June 2022;

Abstract

Many industrial applications require finding solutions to challenging combinatorial problems. Efficient elimination of symmetric solution candidates is one of the key enablers for high-performance solving. However, existing model-based approaches for symmetry breaking are limited to problems for which a set of representative and easily solvable instances is available, which is often not the case in practical applications. This work extends the learning framework and implementation of a model-based approach for Answer Set Programming to overcome these limitations and address challenging problems, such as the Partner Units Problem. In particular, we incorporate a new conflict analysis algorithm in the Inductive Logic Programming system ILASP, redefine the learning task, and suggest a new example generation method to scale up the approach. The experiments conducted for different kinds of Partner Units Problem instances demonstrate the applicability of our approach and the computational benefits due to the first-order constraints learned.

KEYWORDS: answer set programming, inductive logic programming, symmetry breaking constraints

1 Introduction

Finding solutions to hard combinatorial problems is important for various applications, including configuration, scheduling, or planning. Modern declarative solving approaches allow programmers to easily encode various problems and then use domain-independent solvers to find solutions for given instances. The search performance depends highly on

* This work was partially funded by KWF project 28472, cms electronics GmbH, FunderMax GmbH, Hirsch Armbänder GmbH, incubed IT GmbH, Infineon Technologies Austria AG, Isovolta AG, Kostwein Holding GmbH, and Privatstiftung Kärntner Sparkasse. We thank the anonymous reviewers for their valuable suggestions and comments.

both the encoding quality and selected solver parameters. The latter issue can partially be solved using portfolio solvers that use machine learning to select the best possible parametrization of underlying solving algorithms (Hoos et al. 2014). However, writing an optimal encoding remains a challenge that requires an experienced programmer to clearly understand the problem up to details, such as possible structures present in instances, their invariants, and symmetries.

Automatic computation of symmetry breaking constraints (SBCs) can greatly simplify the programming task by extending a given encoding with constraints eliminating symmetric solution candidates, that is, a set of candidates where each one can be obtained from another by renaming constants (Margot 2007; Katebi et al. 2010; Walsh 2012). Existing approaches compute SBCs either for a first-order problem encoding or for only one instance of a given problem. The latter – *instance-specific* methods – compute SBCs online, that is, before or during each invocation of a solver (Puget 2005; Cohen et al. 2006; Drescher et al. 2011). As a result, the obtained SBCs are not transferable to other instances and their computation might significantly increase the total solving time since the problem of finding SBCs is intractable. The *model-based* methods are usually applied offline and aim at finding general SBCs breaking symmetries of a class of instances. However, these methods are either limited to local symmetries occurring due to definitions of variable domains (Devriendt et al. 2016) or require representative sets of instances (Mears et al. 2008; Tarzariol et al. 2021) to identify SBCs that highly likely eliminate symmetries for all instances of this class. Roughly, model-based approaches apply graph-based methods, for example, SAUCY (Darga et al. 2004), to find candidate symmetries of the given instances and then lift the obtained information to first-order SBCs.

The main challenge in the application of model-based approaches is that they must be able to access or generate instances that (i) comprise symmetries representative for the whole instance distribution and (ii) are simple enough to allow the implementation to compute all of their solutions. Consider a small example of the well-studied *Partner Unit Problem* (PUP) (Aschinger et al. 2011; Teppan et al. 2016), which is an abstract representation of configuration problems occurring in railway safety or building security systems. As shown in Figure 1a, the input of the problem is given by a set of units U and a bipartite graph $G = (S, Z, E)$, where S is a set of sensors, Z is a set of security/safety zones, and E is a relation between S and Z . The task is to find a partition of vertices $v \in S \cup Z$ into bags $u_i \in U$, presented in Figure 1b, such that the following requirements hold for each bag: (i) the bag contains at most $UCAP$ many sensors and $UCAP$ many zones and (ii) the bag has at most $IUCAP$ adjacent bags, where the bags u_1 and u_2 are adjacent whenever $v_i \in u_1$ and $v_j \in u_2$ for some $(v_i, v_j) \in E$. The given example shows the smallest instance representing a class of building security systems named *double* by Aschinger et al. (2011). Despite being the simplest instance, it has 145,368 solutions, 98.9% of which can be identified as symmetric (for instance, by renaming the units of a solution). Therefore, the enumeration of symmetries for PUP instances is problematic, even for the smallest and simplest ones.

The *model-based* method suggested by Tarzariol et al. (2021) applies an *instance-specific* tool for identifying symmetries of small but representative instances of a target distribution, and then generalizes respective examples by means of Inductive Logic Programming (ILP). To lift the symmetries, the method needs to investigate all solutions of the analyzed instances, making it inapplicable if no trivial satisfiable instances exist. In this work, we address such limitations and extend the framework's applicability to

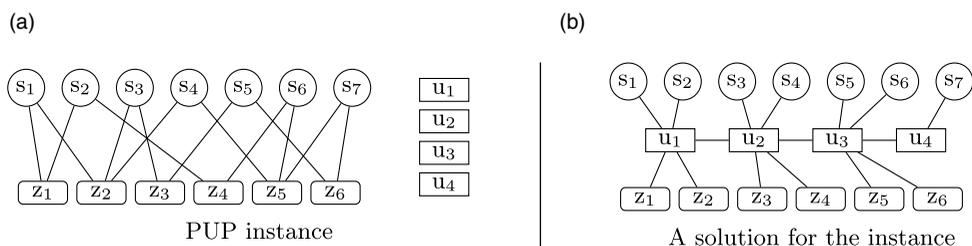


Fig. 1. Partner Unit Problem example with $UCAP = IUCAP = 2$.

combinatorial problems lacking trivial representative instances. In particular, our paper makes the following contributions:

- We propose a new definition of the ILP learning task and a corresponding implementation for the input generation, which allow the approach to scale with respect to the number of answer sets of the analyzed instances, thus, learning efficient first-order constraints.
- We provide a novel conflict analysis method for the learning system ILASP (Law *et al.* 2020) that significantly improves the efficiency of the constraint learning.
- We present an extensive experimental study conducted on three kinds of PUP benchmarks shows that the new method clearly outperforms the legacy approach in terms of learning and solving performance.

The paper is organized as follows: after a short introduction of preliminaries in Section 2, we introduce the revised learning framework and new conflict analysis method in Section 3. Results of our experimental study are presented in Section 4, followed by conclusions and future work in Section 5.

2 Background

In this section, we briefly introduce basic notions of Answer Set Programming, Inductive Logic Programming, and the work by Tarzariol *et al.* (2021), which uses ILP to learn first-order constraints from symmetries of ground ASP programs.

2.1 Answer Set Programming

Answer Set Programming (ASP) is a declarative programming paradigm based on non-monotonic reasoning and the stable model semantics (Gelfond and Lifschitz 1991). Over the past decades, ASP has attracted considerable interest thanks to its elegant syntax, expressiveness, and efficient system implementations, successfully adopted in numerous domains like, for example, configuration, robotics, or biomedical applications (Erdem *et al.* 2016; Falkner *et al.* 2018). We briefly present the syntax and semantics of ASP, and refer the reader to textbooks (Gebser *et al.* 2012; Lifschitz 2019) for more in-depth introductions.

Syntax. An ASP program P is a set of rules r of the form:

$$a_0 \leftarrow a_1, \dots, a_m, \text{not } a_{m+1}, \dots, \text{not } a_n,$$

where *not* stands for *default negation* and a_i , for $0 \leq i \leq n$, are atoms. An *atom* is an expression of the form $p(\bar{t})$, where p is a predicate, \bar{t} is a possibly empty vector of terms, and the predicate \perp (with an empty vector of terms) represents the constant *false*. Each *term* t in \bar{t} is either a variable or a constant. A *literal* l is an atom a_i (positive) or its negation *not* a_i (negative). The atom a_0 is the *head* of a rule r , denoted by $H(r) = a_0$, and the *body* of r includes the positive or negative, respectively, body atoms $B^+(r) = \{a_1, \dots, a_m\}$ and $B^-(r) = \{a_{m+1}, \dots, a_n\}$. A rule r is called a *fact* if $B^+(r) \cup B^-(r) = \emptyset$, and a *constraint* if $H(r) = \perp$.

Semantics. The semantics of an ASP program P is given in terms of its *ground instantiation* P_{grd} , which is obtained by replacing each rule $r \in P$ with its instances obtained by substituting the variables in r by constants occurring in P . Then, an *interpretation* \mathcal{I} is a set of (*true*) ground atoms occurring in P_{grd} that does not contain \perp . An interpretation \mathcal{I} *satisfies* a rule $r \in P_{grd}$ if $B^+(r) \subseteq \mathcal{I}$ and $B^-(r) \cap \mathcal{I} = \emptyset$ imply $H(r) \in \mathcal{I}$, and \mathcal{I} is a *model* of P if it satisfies all rules $r \in P_{grd}$. A model \mathcal{I} of P is *stable* if it is a subset-minimal model of the reduct $\{H(r) \leftarrow B^+(r) \mid r \in P_{grd}, B^-(r) \cap \mathcal{I} = \emptyset\}$, and we denote the set of all stable models, also called answer sets, of P by $AS(P)$.

2.2 Inductive Logic Programming

Inductive Logic Programming (ILP) is a form of machine learning whose goal is to learn a logic program that explains a set of observations in the context of some pre-existing knowledge (Cropper et al. 2020). Since its foundation, the majority of research in the field has addressed Prolog semantics although applications in other paradigms appeared in the last years. The most expressive ILP system for ASP is *Inductive Learning of Answer Set Programs* (ILASP), which can solve a variety of ILP tasks (Law et al. 2014, 2021).

Learning from Answer Sets. A *learning task* for ILASP is given by a triple $\langle B, E, H_M \rangle$, where an ASP program B defines the *background knowledge*, the set E comprises two disjoint subsets E^+ and E^- of *positive* and *negative examples*, and the *hypothesis space* H_M is defined by a language bias M , which limits the potentially learnable rules (Law et al. 2014). Each example $e \in E$ is a pair $\langle e_{pi}, C \rangle$ called *Context Dependent Partial Interpretation* (CDPI), where (i) e_{pi} is a *Partial Interpretation* (PI) defined as pair of sets of atoms $\langle T, F \rangle$, called *inclusions* (T) and *exclusions* (F), respectively, and (ii) C is an ASP program defining the *context* of PI e_{pi} . Given a (total) interpretation \mathcal{I} of a program P and a PI e_{pi} , we say that \mathcal{I} *extends* e_{pi} if $T \subseteq \mathcal{I}$ and $F \cap \mathcal{I} = \emptyset$. Given an ASP program P , an interpretation \mathcal{I} , and a CDPI $e = \langle e_{pi}, C \rangle$, we say that \mathcal{I} is an *accepting answer set* of e with respect to P if $\mathcal{I} \in AS(P \cup C)$ such that \mathcal{I} extends e_{pi} .

Each hypothesis $H \subseteq H_M$ learned by ILASP must respect the following criteria: (i) for each positive example $e \in E^+$, there is some accepting answer set of e with respect to $B \cup H$; and (ii) for any negative example $e \in E^-$, there is no accepting answer set of e with respect to $B \cup H$. If multiple hypotheses satisfy the conditions, the system returns one of those with the lowest cost. By default, the cost c_r of each rule $r \in H_M$ corresponds to its number of literals (Law et al. 2014); however, the user can define a custom scoring function for defining the rule costs. Law et al. (2018) extend the expressiveness of ILASP by allowing noisy examples. With this setting, if an example e is not covered, that is, there

is an accepting answer set for e if it is negative, or none if e is positive, the corresponding weight is counted as a penalty. If no dedicated weight is specified, the example's weight is infinite, thus forcing the system to cover the example. Therefore, the learning task becomes an optimization problem with two goals: minimize the cost of H and minimize the total penalties for the uncovered examples.

The *language bias* M for the ILASP learning task is specified using *mode declarations*. Constraint learning, that is, when the search space exclusively consists of rules r with $H(r) = \perp$, requires only mode declarations for the body of a rule: `#modeb(R,P,(E))`. In this definition, the optional element R is a positive integer, called *recall*, which sets the upper bound on the number of mode declaration applications in each rule. P is a ground atom whose arguments are placeholders of type `var(t)` for some constant term t . In the learned rules, the placeholders will be replaced by variables of type t . For each rule, there are at most V_{\max} variables and B_{\max} literals in the body, which are both equal to 3 by default. Finally, E is an optional modifier that restricts the hypothesis space further, limited in our paper to the `anti_reflexive` and `symmetric` options that both work with predicates of arity 2. When using the former, the atoms of the predicate P should be generated with two distinguished argument values, while rules generated with the latter take into account that the predicate P is symmetric.

In a constraint learning task, just as in other ILP applications (Cropper and Dumani 2020), the language bias must be defined manually for each ASP program P . A careful selection of the bias is essential since a too weak bias might not provide enough limitations for a learner to converge. In contrast, a too strong bias may exclude solutions from the search space, thus resulting in suboptimal learned constraints.

Conflict Driven Inductive Logic Programming (CDILP). Several ILASP releases have been developed in the last years, extending its learning expressiveness and applying more efficient search techniques (Law *et al.* 2020). Recently, Law (2022) introduced CDILP – a new search approach that overcomes the limitation of previous ILASP versions regarding scalability with respect to the number of examples and further aims at efficiently addressing tasks with noisy examples. The approach exploits a set CC of *coverage constraints*, each defined by a pair $\langle e, F \rangle$, where $e \in E$ is a CDPI and F is a propositional formula over identifiers for the rules in H_M . The formula is defined such that, for any $H \subseteq H_M$, if H does not respect F , then H does not cover e . CDILP interleaves the search for an optimal hypothesis H (for the current CC) with a “conflict analysis” phase. In this phase, ILASP identifies (at least) one example e not covered by H , which was not determined by the current CC ; then, it creates a new conflict for e and adds it to CC . If such an example does not exist, ILASP returns the current H as an optimal solution. Otherwise, the system repeats the procedure with the updated set of conflicts. Using the Python interface, PyLASP, one can apply different conflict analysis methods as long as they are proven to be *valid*, that is, a method must terminate and compute formulas F such that the current hypothesis H does not respect them. This requirement guarantees that the CDILP procedure terminates and returns an optimal hypothesis for the learning task.

There are currently three built-in methods for conflict analysis in ILASP, denoted by α , β , and γ , each of which determines a coverage constraint for an example e that is not covered by a given hypothesis H . In the most stringent case, γ , ILASP computes a coverage constraint that is satisfied by exactly those hypotheses that cover e . Identifying such a

comprehensive coverage formula has the advantage that any example will be analyzed in at most one iteration, so that ILASP with γ for conflict analysis usually requires a small number of iterations only. On the other hand, computing such a precise coverage formula is complex, meaning that an iteration can take long time. For this reason, α and β were introduced. Both methods yield smaller formulas that can be computed in less time. While this can lead to more iterations of CDILP, in some domains, the overall runtime benefits from significantly shorter iterations. However, our preliminary investigations showed that, for highly combinatorial PUP instances, even α and β struggle to compute coverage constraints in acceptable time. In this work, we thus introduce a new conflict analysis method that brings significant improvements over α , β , and γ on PUP instances.

2.3 Lifting SBCs for ASP

Tarzariol et al. (2021) presented an approach to lift ground SBCs for ASP programs using ILP. Their system takes four kinds of inputs: (i) an ASP program P modeling a combinatorial problem; (ii) two sets S and Gen of small satisfiable instances representative for a practical problem solved using P ; (iii) the hypothesis space H_M ; and (iv) the Active Background Knowledge ABK as an ASP program comprising auxiliary predicate definitions and constraints learned so far. The *generalization set* Gen contains instances used to generate positive examples that the set of learned constraints must preserve. As a result, we increase the likelihood for the learned constraints to generalize beyond the training examples. The instances of the *training set* S are passed to the instance-specific symmetry breaking system SBASS (Drescher et al. 2011) to identify the symmetries of each instance in S . The output of SBASS, Π , is a set of permutation group *generators* (also called *permutations*) subsuming groups of symmetric answer sets for the analyzed ground program. The framework by Tarzariol et al. (2021) uses this information to define the positive and negative examples for an ILP task and applies ILASP to solve it. The negative examples are associated with a weight, as the system aims at constraints that remove as many symmetric answer sets as possible but does not require eliminating all of them.

In their subsequent work, Tarzariol et al. (2022b) discuss four approaches for creating training examples with SBASS for ground programs P_i , obtained by grounding P with the instances $i \in S$. In particular, *enum* enumerates all answer sets of P_i and classifies each solution as a positive or negative example, according to the common lex-leader approach. That is, if an answer set \mathcal{I} can be mapped to a lexicographically smaller, symmetric answer set using the permutations Π , that is, if \mathcal{I} is dominated, it will produce a negative example. Otherwise, \mathcal{I} yields a positive example. In both cases, the inclusions are $\mathcal{I} \cap atoms(\Pi)$, where $atoms(\Pi)$ denotes the set of atoms occurring in Π , the exclusions are $atoms(\Pi) \setminus \mathcal{I}$, and the context is i . On the other hand, the *fullSBCs* approach exploits the CLINGO API to interleave the solving phase, which returns a candidate answer set \mathcal{I} , with the analysis of all its symmetric solutions. Thanks to the properties of permutation groups (Sakallah 2009), *fullSBCs* can determine all symmetric answer sets by repeatedly applying the permutations Π to \mathcal{I} until no new solutions can be obtained. This approach leads to a partition of the answer sets for an instance, where every partition cell consists of symmetric solutions. For each obtained cell, the system labels the smallest answer set

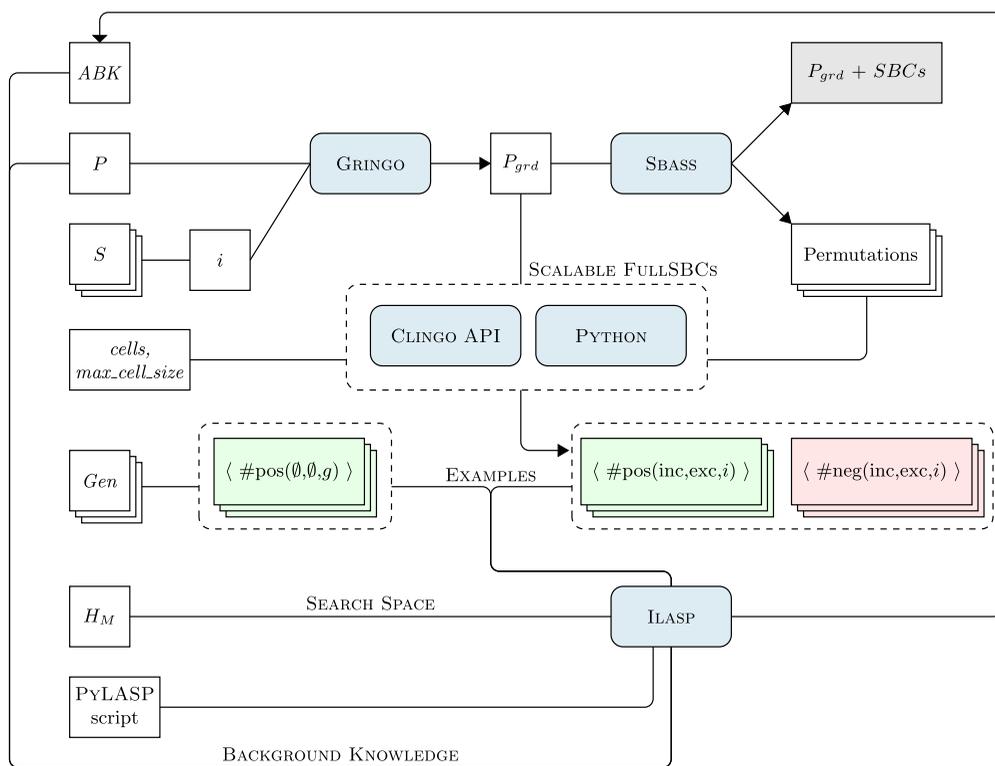


Fig. 2. Revised learning framework implementation.

as a positive example and all remaining ones as negative examples, thus achieving full rather than partial symmetry breaking, while the *enum* approach yields the latter only.

Tarzariol *et al.* (2022b) evaluate the performance of their methods on three versions of the pigeon-hole problem and the house-configuration problem (Friedrich *et al.* 2011). The *ABK* they use contains predicates emulating arithmetic built-ins, and the search space is split to apply the learning framework iteratively and thus increase the learning efficiency. Given that the considered problem instances are defined in terms of unary predicates and those in the training set S have a small number of solutions (from about a dozen up to a few hundred), the suggested formulation of the ILP task admits a fast learning of first-order constraints speeding up the solving of (unsatisfiable) instances. However, instances of complex application problems lack the presupposed characteristics, rendering the previously proposed approaches inapplicable and calling for a more scalable handling of training instances.

3 Method

This section presents an alternative version of the framework introduced by Tarzariol *et al.* (2021), extending its applicability. First, we propose a revised ILP learning task and procedures necessary to define inputs of this task for difficult combinatorial problems, as illustrated in Figure 2. Then, we describe a new conflict analysis method for the ILASP system enabling efficient constraint learning to handle this revised task.

3.1 Revised ILP task

Training Examples. Tarzariol et al. (2021) propose approaches to the example generation for an ILP learning task, which yield a number of examples proportional to the number of solutions for each problem instance in S . However, for difficult problems with many symmetries, even the simplest instances might yield a large number of answer sets. Their enumeration might thus take unacceptably long time. Moreover, even if the enumeration succeeds, the number of obtained examples is often too large to be handled by ILASP. To overcome these issues, we propose two scalable approaches to the generation of examples for each instance in S based on the *enum* and *fullSBCs* strategies by Tarzariol et al. (2021). The scalable *enum* strategy generates examples from a portion of solutions, which are sampled from at most n random answer sets. For each candidate solution, the lex-leader criterion is applied to determine whether another symmetric answer set dominates it. This approach does not guarantee a fixed ratio between positive and negative examples, and it might fail to identify symmetric answer sets, as inspecting single applications of permutation group generators does not achieve full symmetry breaking (Sakallah 2009). Nevertheless, the fixed number of considered answer sets provides means to limit the time required for the definition of a learning task. The scalable *fullSBCs* approach creates a set of examples configured by two parameters: (i) *cells* defines the number of cells of symmetric solutions to analyze and (ii) *max_cell_size* limits the maximal number of negative examples generated per cell. For each cell, the approach adds the first *max_cell_size* symmetric answer sets as negative examples. Next, it explores the whole cell of symmetric solutions and takes the smallest one as a positive example. As a result, this method generates a controlled number of positive and negative examples, regardless of how many solutions there may be for a given instance. In fact, at most *cells* many positive and *cells* \times *max_cell_size* many negative examples can be obtained in total.

Background Knowledge. Learning SBCs that improve the grounding and solving efficiency is crucial for difficult combinatorial problems. Therefore, the background knowledge *ABK* should provide necessary auxiliary predicates that allow an ILP system to incorporate such constraints in the search space. Previous formulations of *ABK* for learning SBCs have issues with expressing appropriate constraints since the provided predicates do not take the structure of problem instances into account. In this paper, we propose a new version of *ABK* comprising two new types of auxiliary predicates. The first type encodes local properties of nodes in an input graph, while the second enables a more efficient constraint representation. That is, for two different nodes N and M , the predicate `close(N,M)` holds if these nodes share a common neighbor. In case of bipartite graphs containing two types of nodes, A and B (standing for sensors and zones in case of PUP instances), we define two versions of this auxiliary predicate, distinguishing the two types by `closeA/2` and `closeB/2`. The second auxiliary predicate introduces an ordering on value assignments as follows: let $[1..max_x]$ and $[1..max_y]$ be the domains of two variables X and Y , and $p(X,Y)$ be a binary predicate that holds for at most one value $X \in [1..max_x]$ per $Y \in [1..max_y]$ in each answer set. Then, we define the following auxiliary predicate:

$$\begin{aligned} \text{pGEQ}(X,Y) & :- p(X,Y). \\ \text{pGEQ}(X,Y) & :- \text{pGEQ}(X+1,Y), 0 < X. \end{aligned}$$

If $\text{pGEQ}(X, Y)$ is true, we know that $\text{p}(X', Y)$ holds for some value X' equal to X or greater. A constraint may then contain $\text{pGEQ}(Y, Y)$ instead of the equivalent test $\text{p}(X, Y)$, $Y \leq X$, thus reducing the ground instantiation size. Moreover, this encoding can bring benefits for solving as well, as we obtain a more powerful propagation (Crawford and Baker 1994).

From a technical perspective, the framework presented by Tarzariol *et al.* (2021) runs SBASS on a ground program resulting from the union of P , an instance $i \in S$, and ABK . The inclusion of ABK caused no difference in the symmetries for their approach (without iterative constraint learning), given that the introduced auxiliary predicates do not affect atoms occurring in P . On the other hand, the predicate $\text{pGEQ}/2$ alters the identification of symmetries for atoms over the predicate $\text{p}/2$ contained in P . Hence, we do not necessitate ABK to contribute to a ground program passed to SBASS, as indicated in Figure 2.

Language Bias. To address difficult combinatorial problems, we suggest the following set of mode declarations to define the search space H_M :

```
#modeb(1,r(var(t),var(t))).
#modeb(1,close(var(t),var(t)),(symmetric,anti_reflexive)).
#modeb(2,pGEQ(var(t),var(t))).
#modeb(1,q(var(t),var(t))).
```

Assuming that the predicate $\text{r}/2$ specifies the graph provided by an instance, the mode declaration in the first line expresses that one such atom can occur per constraint. Then, for each $\text{close}/2$ or $\text{pGEQ}/2$ predicate in ABK , we include a respective mode declaration as in the second and third lines. Moreover, if atoms over another predicate $\text{q}/2$ occur in P , but not in ABK , we supply a mode declaration of the last kind, where considering binary predicates is sufficient for PUP instances. Let us notice that only one type t of variables is used for all mode declarations. In this way, there is no restriction on the variable replacements, and we may explore inherent yet hidden properties of the input labels. Furthermore, we introduce an alternative scoring function assigning the cost c_{id} to each rule $r_{id} \in H_M$. For every literal obtainable from the mode declarations, we overwrite its default cost 1 with a custom cost, except for literals over domain predicates like, for example, $\text{r}/2$. For literals over other predicates, in case the same variable is used for both of the contained arguments, the cost is 2, and 3 otherwise. For example, the cost of the constraint $:- \text{pGEQ}(V1, V1), \text{close}(V1, V2), \text{q}(V2, V3)$ would be equal to $1 + 1 + 1 = 3$ with the default scoring function but $2 + 1 + 3 = 6$ with our custom costs.

3.2 Conflict analysis

As discussed in Section 2.2, ILASP's Conflict Driven ILP (CDILP) approach requires a *conflict analysis* method that, given a hypothesis H and an example e that H does not cover, returns a coverage formula F . For ILASP to function correctly, this formula must (i) not be satisfied by H , and (ii) be satisfied by every hypothesis $H' \subseteq H_M$ covering e .

ILASP has several built-in methods for conflict analysis, some of which are described by Law (2022). However, for learning tasks with hypothesis spaces that consist of constraints only, these methods all behave equivalently when processing positive examples. For positive examples $e = \langle e_{pi}, C \rangle$ such that $B \cup C$ has many answer sets, the built-in methods can return extremely long coverage formulas that also take long time to compute. For this reason, we define a new conflict analysis method leading to shorter coverage constraints that can be computed much faster.

Definition 1

Let $T = \langle B, E, H_M \rangle$ be a learning task such that H_M consists of constraints, and let $e \in E^+$ be a positive example. For any $H \subseteq H_M$ that does not cover e , the *subsumption-based conflict analysis* method $sbca(e, H, T)$ returns the formula $\bigvee_{r \in H, r' \in H_M} \bigwedge_{r' \subseteq_{\theta} r} \neg r'_{id}$, where $r' \subseteq_{\theta} r$ denotes that r' subsumes r .

Example 1

Consider a scenario such that ILASP is run on a task T with the language bias given in the previous subsection. At some point in the execution, T may generate the hypothesis $H = \{:- \text{close}(V1, V2). \quad :- \text{not pGEQ}(V1, V1), \text{q}(V1, V1).\}$. Within the hypothesis space computed by ILASP,¹ the first rule is only subsumed by itself, and the second rule is subsumed by the following rules:

- | | |
|---|--|
| $:- \text{q}(V1, V1).$ | $:- \text{not pGEQ}(V1, V1), \text{q}(V1, V1).$ |
| $:- \text{q}(V1, V2).$ | $:- \text{not pGEQ}(V1, V1), \text{not pGEQ}(V1, V2), \text{q}(V1, V2).$ |
| $:- \text{pGEQ}(V1, V1).$ | $:- \text{not pGEQ}(V1, V1), \text{not pGEQ}(V2, V1), \text{q}(V1, V2).$ |
| $:- \text{pGEQ}(V1, V2).$ | $:- \text{not pGEQ}(V1, V1), \text{not pGEQ}(V2, V2), \text{q}(V1, V2).$ |
| $:- \text{not pGEQ}(V1, V1), \text{q}(V1, V2).$ | $:- \text{not pGEQ}(V1, V2), \text{not pGEQ}(V2, V1), \text{q}(V1, V2).$ |
| $:- \text{not pGEQ}(V1, V2), \text{q}(V1, V2).$ | $:- \text{not pGEQ}(V1, V2), \text{not pGEQ}(V2, V2), \text{q}(V1, V2).$ |
| $:- \text{not pGEQ}(V2, V1), \text{q}(V1, V2).$ | $:- \text{not pGEQ}(V2, V1), \text{not pGEQ}(V2, V2), \text{q}(V1, V2).$ |
| $:- \text{not pGEQ}(V2, V2), \text{q}(V1, V2).$ | |

Let r^1 be the first rule in H and r^2, \dots, r^{16} be the rules that subsume the second rule. In this case, for any positive example e that is not covered by H , we obtain the coverage constraint $sbca(e, H, T) = (\neg r^1_{id}) \vee ((\neg r^2_{id}) \wedge \dots \wedge (\neg r^{16}_{id}))$.

The following theorem shows that $sbca$ is a valid method for conflict analysis for positive examples, provided that the hypothesis space contains constraints only. This means that in our application domain, when using the $sbca$ method for some or all of the positive examples, ILASP is guaranteed to return an optimal solution for any learning task.

Theorem 1

Let $T = \langle B, E, H_M \rangle$ be a learning task such that H_M consists of constraints, and let $e \in E^+$ be a positive example. For any $H \subseteq H_M$ that does not cover e :

1. H does not satisfy $sbca(e, H, T)$.
2. Every $H' \subseteq H_M$ that covers e satisfies $sbca(e, H, T)$.

Proof

1. For each $r \in H$, $r \subseteq_{\theta} r$ implies that $\bigwedge_{r' \subseteq_{\theta} r} \neg r'_{id}$ is not satisfied by H . Thus, H cannot satisfy any of the disjuncts of $sbca(e, H, T)$, that is, it does not satisfy $sbca(e, H, T)$.
2. Assume for contradiction that some $H' \subseteq H_M$ covers $e = \langle e_{pi}, C \rangle$ but does not satisfy $sbca(e, H, T)$. Then, by the definition of covering e , there is some answer set $\mathcal{I} \in AS(B \cup C \cup H')$ that extends e_{pi} . As H' does not satisfy $sbca(e, H, T)$, H' must include constraints that subsume each of the constraints in H . This means that

¹ This space is smaller than the full hypothesis space as isomorphic rules are discarded. For instance, $:- \text{q}(V2, V2).$ is isomorphic to $:- \text{q}(V1, V1).$ and thus not considered by ILASP.

$AS(B \cup C \cup H') \subseteq AS(B \cup C \cup H)$. Hence, $\mathcal{I} \in AS(B \cup C \cup H)$ contradicts the condition that H does not cover e . □

The advantage of the *sbca* method for conflict analysis, over the methods that are already built-in to ILASP, is that this method does not need to compute answer sets of $B \cup C \cup H$ (for an uncovered example with context C). In other words, it does not need to consider the semantics of the current hypothesis H and can instead focus on purely syntactic properties. In combinatorial problem domains, where finding answer sets can be computationally intensive, our preliminary experiments showed that *sbca* is much faster than the existing conflict analysis methods of ILASP. On the other hand, the syntactic coverage constraints generated by *sbca* tend to be more specific than the semantic constraints computed by ILASP's built-in methods. That is, the formulas determined by *sbca* apply to fewer hypotheses and cut out fewer solution candidates, which in turn means that more (yet considerably faster) iterations of the CDILP procedure are required.

The *sbca* method is closely related to the ILP system POPPER (Cropper and Morel 2021), which also identifies syntactically determined constraints based on subsumption. Specifically, when POPPER encounters a hypothesis H entailing an atom that should not be entailed (a negative example for POPPER), all hypotheses subsuming H are discarded (as these would also entail the atom). Just like *sbca*, since POPPER uses syntactically determined constraints, it can compute these very quickly but may need many more iterations (compared to the semantic conflict analysis methods of ILASP). While both approaches are closely linked, POPPER learns under Prolog semantics. Therefore, POPPER cannot reason about logic programs with multiple answer sets and is inapplicable for learning ASP constraints in the combinatorial problem domains we address.

PyLASP Script. The *sbca* method is essential to compute coverage formulas for positive examples obtained from the generalization set Gen , including instances with a large number of answer sets, in acceptable time. However, for other examples, the existing conflict analysis methods of ILASP are better suited, as the obtained formulas are more informative. Hence, we extended the default PYLASP script of ILASP with means to specify examples requiring *sbca* usage by dedicated identifiers, associated with instances in Gen on which CLINGO takes more than 5 s for enumerating the answer sets.

4 Experiments

For testing our new method, we decided to use PUP configuration benchmarks (Teppan *et al.* 2016). We applied our framework to PUP instances supplied by Aschinger *et al.* (2011), studying the *double*, *doublev*, and *triple* instance collections with $UCAP = IUCAP = 2$. Instances of the same type represent buildings of similar topology with scaling parameters that follow a common distribution. Although the benchmark instances are synthetic, they represent a relevant configuration problem concerning safety and security issues in public buildings, like administration offices or museums. In addition, the scalable synthetic benchmarks are easy to generate and analyze. That is, the nodes corresponding to rooms are labeled in a specific order, and the topologies follow a clear pattern. The PUP instances and further details are provided by Tarzariol *et al.* (2022a).

In all three experiments, we learn first-order constraints using the same inputs to the ILP task, except for the two instance sets S and Gen : for S , we pick the smallest representative instance of the selected type, while Gen comprises the three smallest satisfiable instances other than the one in S . Components of the ILP task shared between the experiments with different kinds of instances include (i) the input program P as an ASP encoding of PUP that comprises no SBCs and originates from work by Dodaro et al. (2016), where it is referred to as ENC1; (ii) the background knowledge ABK defining the auxiliary predicates `close_sensors/2`, `close_zones/2`, `unit2zoneGEQ/2`, and `unit2sensorGEQ/2`; and (iii) the language bias M with mode declarations for the auxiliary predicates in ABK as well as the predicates `zone2sensor/2` and `partnerunits/2`, the latter taking the roles of `r/2` and `q/2` according to the scheme described in Section 3.1.

For each PUP instance type, we tested the two proposed approaches (*scalable enum* and *scalable fullSBCs*) combined with two versions of the scoring function: the default function of ILASP as well as the custom costs introduced in Section 3.1. With both example generation approaches, the sampling of answer sets was done using the `--seed <seed>` option of CLINGO (v5.4.0). We executed the experiments using 120 different random seeds to counterbalance the impact of randomness and get more reliable estimates of the relative learning and solving performance. For comparing the two approaches, we sampled the same number of examples configured by n for *scalable enum* and $cells \times max_cell_size$ for *scalable fullSBCs*. We limited the learning time to one hour, interrupting the process in case no positive or negative examples could be obtained from the analysis by SBASS.

In preliminary investigations, we tried the default PYLASP script of ILASP (v4.1.2) as well as the one we devised to apply the *sba* method for conflict analysis to positive examples from Gen . As expected, no run with the default script could be finished within one hour, and we thus focus below on results obtained with our new PYLASP script.

The learning efficiency for all three PUP instance types is such that *scalable fullSBCs* yields more successful ILP tasks, that is, tasks solved by ILASP within the time limit, than *scalable enum*.² Using the *scalable fullSBCs* strategy and 120 different random seeds, we were able to finish ILASP runs with 108 seeds for *double*, 88 for *doublev*, and 12 for *triple*. The application of *scalable enum* resulted in only 10, 18, or 6 successful runs, respectively, with some of the 120 seeds. Note that several *scalable enum* runs had to be canceled because the generated ILP tasks were partial, that is, without either positive or negative examples, so that trivial optimal hypotheses make learning with ILASP obsolete. As *scalable fullSBCs* is the by far more successful example generation strategy, we restrict the following considerations of solving performance to constraints learned with it.

Next, we compare the solving performance of CLINGO relative to constraints learned with the default scoring function of ILASP or the custom costs distinguishing domain predicates and variable recurrences as described in Section 3.1. To this end, we consider those of the 120 random seeds for which ILASP runs finished successfully with both of the scoring functions, and then provide the learned constraints as background knowledge to CLINGO along with a PUP instance. Each CLINGO run is limited to 300 s solving time, and Figure 3 displays box plots as well as average runtimes (Avg), standard deviation (Std), and number of timeouts with the total number of runs in parentheses (TO) for

² The results reflect the learning setting leading to the fastest runtime, using the alternative ordering *ord* but no *sat* mode (Tarzariol et al. 2022b). Detailed records are provided by Tarzariol et al. (2022a).

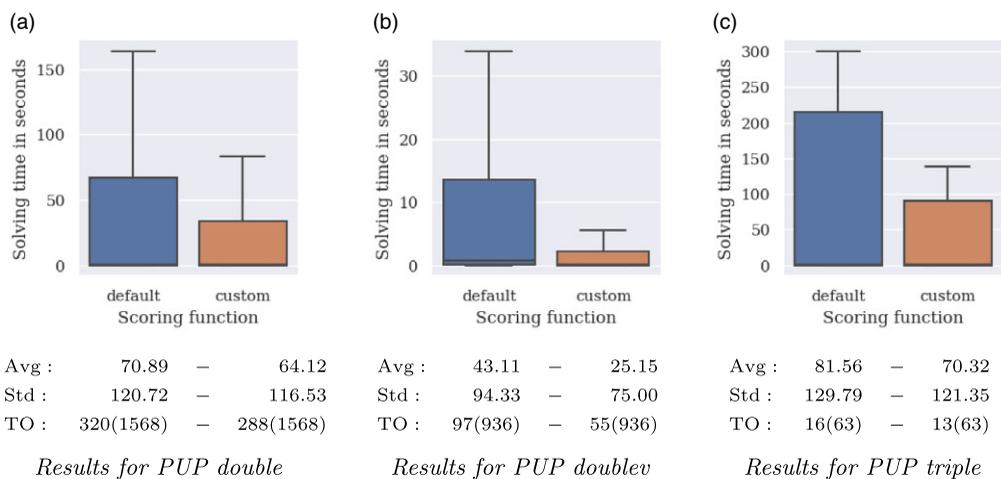


Fig. 3. Aggregated solving times for constraints learned with the two scoring functions.

the *double*, *doublev*, and *triple* instances. For all three PUP instance types and especially the hard instances whose runtime is above average, we observe significant speed-ups due to constraints learned by means of the custom scoring function. In fact, the custom costs give preference to first-order constraints whose ground instances apply and prune the search space more directly, thus benefiting the solving performance of CLINGO. We also applied the Wilcoxon Signed-Rank test (Wilcoxon 1992) for non-normally distributed runtimes, which confirms that the observed differences are statistically significant.

Finally, we contrast the best-performing learning setting, that is, the *scalable full-SBCs* strategy for example generation along with our custom scoring function for assigning costs to constraints, with originally proposed ASP encodings of PUP and instance-specific symmetry breaking. In detail, Tables 1–3 show runtime results for the following systems and encodings: (i) CLINGO on the plain ENC1 encoding; (ii) CLINGO on ENC1 augmented with the most efficient learned constraints among those aggregated in Figure 3 as *ABK*; (iii) CLINGO on the advanced ENC2 encoding (Dodaro et al. 2016), incorporating hand-crafted static symmetry breaking as well as an ordered representation (Crawford and Baker 1994) of assigned units similar to pGEQ/2 in Section 3.1; (iv) SBASS for permutation and ground SBCs computation on ENC1; and (v) CLASP π denoting the solving time of CLINGO on ENC1 augmented with ground SBCs by SBASS.³ PUP instances are named according to the scheme [un-]type-zones, where un indicates unsatisfiability due to including one unit less than required, type denotes the *double*, *doublev*, and *triple* collections by *dbl*, *dbl*v, or *tri*, and the number of zones provides a measure of size. Each run is limited to 600 s, and the TO entries mark unfinished runs.

Regarding different PUP encodings, ENC2 leads to more robust CLINGO performance than the simpler ENC1 encoding, even if ground SBCs from SBASS are included for CLASP π . That is, apart from a few shorter runs with ENC1 on satisfiable instances (*dbl*-30, *dbl*v-45, *tri*-18, and *tri*-21) in Tables 1–3, CLINGO scales better with ENC2, never times out on instances finished with ENC1 or possibly CLASP π , and is able to solve some instances (*dbl*-50,

³ In the online usage of instance-specific symmetry breaking, the runtimes of SBASS and CLASP π add up.

Table 1. *Runtimes for PUP double*

	ENC1	ABK	ENC2	SBASS	CLASP^π
dbl-10	0.02	0.01	0.01	0.04	0.02
un-dbl-10	505.91	0.01	0.16	0.03	TO
dbl-20	1.06	0.08	0.53	0.31	1.40
un-dbl-20	TO	0.34	TO	0.28	TO
dbl-30	1.70	0.46	2.21	3.12	0.91
un-dbl-30	TO	19.97	TO	3.19	TO
dbl-40	TO	5.50	11.50	14.58	482.17
un-dbl-40	TO	65.54	TO	9.52	TO
dbl-50	TO	54.89	542.09	57.91	TO
un-dbl-50	TO	61.27	TO	48.63	TO

Table 2. *Runtimes for PUP doublev*

	ENC1	ABK	ENC2	SBASS	CLASP^π
dblv-15	0.03	0.03	0.02	0.05	0.03
un-dblv-15	TO	0.04	0.98	0.04	68.30
dblv-30	0.24	0.53	0.19	0.72	0.84
un-dblv-30	TO	0.73	TO	0.59	TO
dblv-45	3.36	3.07	8.04	6.11	51.57
un-dblv-45	TO	12.32	TO	4.40	TO
dblv-60	50.18	4.81	2.93	31.74	57.24
un-dblv-60	TO	53.74	TO	22.34	TO
dblv-75	517.05	69.51	280.90	98.88	TO
un-dblv-75	TO	98.63	TO	79.33	TO

Table 3. *Runtimes for PUP triple*

	ENC1	ABK	ENC2	SBASS	CLASP^π
tri-12	0.02	0.03	0.01	0.04	0.02
un-tri-12	TO	0.13	3.38	0.03	TO
tri-15	0.03	0.04	0.03	0.08	0.24
un-tri-15	TO	73.88	393.70	0.05	TO
tri-18	0.08	1.62	0.62	0.16	1.05
un-tri-18	TO	72.03	TO	0.09	TO
tri-21	0.65	42.41	0.67	0.37	3.71
un-tri-21	TO	73.83	TO	0.19	TO
tri-24	7.20	78.18	2.80	0.62	12.57
un-tri-24	TO	269.38	TO	0.44	TO

un-tri-12, and un-tri-15) on which the latter two settings fail. Considering that SBASS produces ground SBCs within the time limit for all instances, the better performance with ENC2 suggests that its hand-crafted static symmetry breaking approach provides a more economic trade-off between the compactness and completeness of introduced SBCs. However, we checked that static symmetry breaking by unit labels counteracts SBASS to detect any instance-specific symmetries, so that ENC2 commits to value symmetries only.

Indeed, when comparing ENC2 to CLINGO on ENC1 with first-order constraints learned by ILASP as ABK, we observe further significant performance improvements thanks to our

approach, particularly on the unsatisfiable instances in Tables 1–3. That is, the learned *ABK* enables CLINGO to solve the considered PUP instances of three different types and efficiently prunes the search space, which must be fully explored in case of unsatisfiability. We checked that the learned constraints exploit the topology of instances to restrict the assignable units, particularly using the predicate `unit2sensorGEQ/2` for referring to the assignment of sensors. The specific sets of efficient first-order constraints learned for each of the three PUP instance types are provided in our repository (Tarzariol *et al.* 2022a).

5 Conclusions

This paper introduces an approach to learn first-order constraints for complex combinatorial problems, which cannot be successfully tackled by previous, less scalable methods (Tarzariol *et al.* 2022b). We devised and implemented two configurable strategies to generate the examples for an ILP task and extended the background knowledge by auxiliary predicates admitting a more compact representation and potentially stronger propagation of constraints on value assignments. Moreover, we introduced a custom scoring function taking the ground instances of first-order constraints into account, along with a new conflict analysis method for ILASP that enables a much faster handling of positive examples with many answer sets. The revised learning framework taking all proposed techniques together is able to learn efficient first-order constraints from non-trivial problem instances, as demonstrated on three kinds of PUP benchmarks – a challenging configuration problem. In the future, we hope to introduce automatic (re-)labeling schemes for constants appearing in instances to exploit common problem structure in a less input-specific way. Moreover, we aim at extending our learning framework further to enable the model-based analysis and breaking of symmetries for practically important optimization problems. Beyond graph-oriented applications, we plan to expand the scope of our constraint learning methods to further areas, such as scheduling domains with symmetries among tasks and resources, for example, several instances of products, machines, or workers with the same skills.

References

- ASCHINGER, M., DRESCHER, C., FRIEDRICH, G., GOTTLÖB, G., JEAVONS, P., RYABOKON, A. AND THORSTENSEN, E. 2011. Optimization methods for the partner units problem. In *International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming*. Lecture Notes in Computer Science, vol. 6697. Springer, 4–19.
- COHEN, D., JEAVONS, P., JEFFERSON, C., PETRIE, K. AND SMITH, B. 2006. Symmetry definitions for constraint satisfaction problems. *Constraints* 11, 2-3, 115–137.
- CRAWFORD, J. AND BAKER, A. 1994. Experimental results on the application of satisfiability algorithms to scheduling problems. In *AAAI Conference on Artificial Intelligence*. AAAI Press 1092–1097.
- CROPPER, A., DUMANČIĆ, S. AND MUGGLETON, S. 2020. Turning 30: New ideas in inductive logic programming. In *International Joint Conference on Artificial Intelligence*, 4833–4839. ijcai.org.
- CROPPER, A. AND DUMANI, S. 2020. Inductive logic programming at 30: A new introduction. URL: <https://arxiv.org/abs/2008.07912>.
- CROPPER, A. AND MOREL, R. 2021. Learning programs by learning from failures. *Machine Learning* 110, 4, 801–856.

- DARGA, P., KATEBI, H., LIFFITON, M., MARKOV, I. AND SAKALLAH, K. 2004. Saucy. URL: <http://vlsicad.eecs.umich.edu/BK/SAUCY/>. [Accessed on May 21, 2021].
- DEVRIENDT, J., BOGAERTS, B., BRUYNNOOGHE, M. AND DENECKER, M. 2016. On local domain symmetry for model expansion. *Theory and Practice of Logic Programming* 16, 5-6, 636–652.
- DODARO, C., GASTEIGER, P., LEONE, N., MUSITSCH, B., RICCA, F. AND SCHEKOTIHIN, K. 2016. Combining answer set programming and domain heuristics for solving hard industrial problems. *Theory and Practice of Logic Programming* 16, 5-6, 653–669.
- DRESCHER, C., TIFREA, O. AND WALSH, T. 2011. Symmetry-breaking answer set solving. *AI Communications* 24, 2, 177–194.
- ERDEM, E., GELFOND, M. AND LEONE, N. 2016. Applications of ASP. *AI Magazine* 37, 3, 53–68.
- FALKNER, A., FRIEDRICH, G., SCHEKOTIHIN, K., TAUPE, R. AND TEPPAN, E. 2018. Industrial applications of answer set programming. *Künstliche Intelligenz* 32, 2-3, 165–176.
- FRIEDRICH, G., RYABOKON, A., FALKNER, A., HASELBÖCK, A., SCHENNER, G. AND SCHREINER, H. (Re)configuration using answer set programming. In *IJCAI 2011 Workshop on Configuration 2011*, 17–24. CEUR-WS.org.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool Publishers.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9, 365–385.
- HOOS, H., LINDAUER, M. AND SCHAUB, T. 2014. claspfolio 2: Advances in algorithm selection for answer set programming. *Theory and Practice of Logic Programming* 14, 4-5, 569585.
- KATEBI, H., SAKALLAH, K. AND MARKOV, I. 2010. Symmetry and satisfiability: An update. In *International Conference on Theory and Applications of Satisfiability Testing*. Lecture Notes in Computer Science, vol. 6175. Springer, 113–127.
- LAW, M. 2022. Conflict-driven inductive logic programming. *Theory and Practice of Logic Programming, First View*.
- LAW, M., RUSSO, A. AND BRODA, K. 2014. Inductive learning of answer set programs. In *European Conference on Logics in Artificial Intelligence*. Lecture Notes in Computer Science, vol. 8761. Springer, 311–325.
- LAW, M., RUSSO, A. AND BRODA, K. 2018. Inductive learning of answer set programs from noisy examples. *Advances in Cognitive Systems* 7, 57–76.
- LAW, M., RUSSO, A. AND BRODA, K. 2020. The ILASP system for inductive learning of answer set programs. URL: <https://arxiv.org/abs/2005.00904>.
- LAW, M., RUSSO, A. AND BRODA, K. 2021. ILASP. URL: <https://www.ilasp.com>. [Accessed on May 21, 2021].
- LIFSCHITZ, V. 2019. *Answer Set Programming*. Springer.
- MARGOT, F. 2007. Symmetric ILP: Coloring and small integers. *Discrete Optimization* 4, 1, 40–62.
- MEARS, C., DE LA BANDA, M., WALLACE, M. AND DEMOEN, B. 2008. A novel approach for detecting symmetries in CSP models. In *International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming*. Lecture Notes in Computer Science, vol. 5015. Springer, 158–172.
- PUGET, J. 2005. Automatic detection of variable and value symmetries. In *International Conference on Principles and Practice of Constraint Programming*. Lecture Notes in Computer Science, vol. 3709. Springer, 475–489.
- SAKALLAH, K. 2009. Symmetry and satisfiability. In *Handbook of Satisfiability*. Frontiers in Artificial Intelligence and Applications, vol. 185. IOS Press, 289–338.
- TARZARIOL, A., GEBSER, M. AND SCHEKOTIHIN, K. 2021. Lifting symmetry breaking constraints with inductive logic programming. In *International Joint Conference on Artificial Intelligence*, 2062–2068. ijcai.org.

- TARZARIOL, A., GEBSEER, M. AND SCHEKOTIHIN, K. 2022a. ILP Symmetry Breaking. URL: https://github.com/prosyscience/Symmetry_Breaking_with_ILP/tree/pup/. [Accessed on January 28, 2022].
- TARZARIOL, A., GEBSEER, M. AND SCHEKOTIHIN, K. 2022b. Lifting symmetry breaking constraints with inductive logic programming. *Machine Learning, Special Issue on Learning and Reasoning*.
- TEPPAN, E., FRIEDRICH, G. AND GOTTLÖB, G. 2016. Tractability frontiers of the partner units configuration problem. *Journal of Computer and System Sciences* 82, 5, 739–755.
- WALSH, T. 2012. Symmetry breaking constraints: Recent results. In *AAAI Conference on Artificial Intelligence*. AAAI Press, 2192–2198.
- WILCOXON, F. 1992. Individual comparisons by ranking methods. In *Breakthroughs in Statistics*. Springer, 196–202.