



Lifting symmetry breaking constraints with inductive logic programming

Alice Tarzariol¹ · Martin Gebser^{1,2} · Konstantin Schekotihin¹

Received: 7 June 2021 / Revised: 29 September 2021 / Accepted: 7 February 2022 /
Published online: 19 April 2022
© The Author(s) 2022

Abstract

Efficient omission of symmetric solution candidates is essential for combinatorial problem-solving. Most of the existing approaches are instance-specific and focus on the automatic computation of Symmetry Breaking Constraints (SBCs) for each given problem instance. However, the application of such approaches to large-scale instances or advanced problem encodings might be problematic since the computed SBCs are propositional and, therefore, can neither be meaningfully interpreted nor transferred to other instances. As a result, a time-consuming recomputation of SBCs must be done before every invocation of a solver. To overcome these limitations, we introduce a new model-oriented approach for Answer Set Programming that lifts the SBCs of small problem instances into a set of interpretable first-order constraints using the Inductive Logic Programming paradigm. Experiments demonstrate the ability of our framework to learn general constraints from instance-specific SBCs for a collection of combinatorial problems. The obtained results indicate that our approach significantly outperforms a state-of-the-art instance-specific method as well as the direct application of a solver.

Keywords Answer set programming · Inductive logic programming · Symmetry breaking constraints

Editors: Nikos Katzouris, Alexander Artikis, Luc De Raedt, Artur d'Avila Garcez, Ute Schmid, Sebastijan Dumancic, Jay Pujara.

✉ Alice Tarzariol
alice.tarzariol@aau.at

Martin Gebser
martin.gebser@aau.at

Konstantin Schekotihin
konstantin.schekotihin@aau.at

¹ University of Klagenfurt, Klagenfurt, Austria

² University of Graz University of Technology, Graz, Austria

1 Introduction

Modern declarative programming paradigms allow for relatively simple modeling of various combinatorial problems. Nevertheless, solving these problems might become infeasible when the size of input instances and, correspondingly, the number of possible solution candidates start to grow (Dodaro et al., 2016). In many cases, these candidates are symmetric, i.e., one candidate can easily be obtained from another by renaming constants. In order to deal with large problem instances, the ability to encode *Symmetry Breaking Constraints* (SBCs) in a problem representation becomes an essential skill for programmers. However, the identification of symmetric solutions and the formulation of constraints that remove them is a tedious and time-consuming task. As a result, various tools emerged that avoid the computation of symmetric solutions by, for instance, automatically finding a set of SBCs using properties of permutation groups or applying specific search methods that detect and ignore symmetric states; see (Margot, 2010; Sakallah, 2009; Walsh, 2012) for an overview.

Existing approaches to SBC generation can be classified into *instance-specific* and *model-oriented* ones. The former methods identify symmetries for a particular instance at hand by computing and adding ground SBCs to the problem representation (Cohen et al., 2006; Drescher et al., 2011; Puget, 2005). Unfortunately, computational advantages do not carry forward to large-scale instances or advanced encodings, where *instance-specific* symmetry breaking often requires as much time as it takes to solve the original problem. Moreover, ground SBCs generated by *instance-specific* approaches are (i) not transferable, since the knowledge obtained is limited to a single instance; (ii) usually hard to interpret and comprehend; (iii) derived from permutation group generators, whose computation is itself a combinatorial problem; and (iv) often redundant and might result in a degradation of the solving performance.

In contrast, *model-oriented* approaches aim to find general SBCs that depend less on a particular instance. The method presented in Devriendt et al. (2016) uses local domain symmetries of a given first-order theory. SBCs are generated by identifying argument positions in atoms of a formula that comprise object variables defined over the same subset of a domain given in the input. As a result, the computation of lexicographical SBCs is very fast. However, the method considers each first-order formula separately and cannot reliably remove symmetric solutions, as it requires the analysis of several formulas at once. The method of Mears et al. (2008) computes SBCs by generating small instances of parametrized constraint programs, and then finds candidate symmetries using SAUCY (Codenotti et al., 2013; Darga et al., 2004) – a graph automorphism detection tool. Next, the algorithm removes all candidate symmetries that are valid only for some of the generated examples as well as those that cannot be proven to be parametrized symmetries using heuristic graph-based techniques. This approach can be seen as a simplified learning procedure that utilizes only negative examples represented by the generated SBCs.

In this work, we introduce a novel *model-oriented* method for Answer Set Programming (ASP) (Lifschitz, 2019) that aims to generalize the process of discarding redundant solution candidates for instances of a target domain using Inductive Logic Programming (ILP) (Cropper et al., 2020). The goal is to lift the SBCs of small problem instances and to obtain a set of interpretable first-order constraints. Such constraints cut the search space while preserving the satisfiability of a problem for the considered instance distribution, which improves the solving performance, especially in the case of unsatisfiability. The particular contributions of our work are:

- We suggest several methods to generate a training set comprising positive and negative examples, allowing an ILP system to learn first-order SBCs for the problem at hand.
- We define the components of an ILP learning task enabling the generation of lexicographical SBCs for ASP.
- We analyze the features and characteristics of the results obtained by our methods, as well as the effects of language bias decisions on several combinatorial problems.
- We present an approach that iteratively applies our method to revise constraints when new permutation group generators or more training instances become available.
- We conduct performance experiments on variants of the pigeon-hole problem as well as the house-configuration problem (Friedrich et al., 2011). The obtained results show that the SBCs generated are easy to interpret in most of the cases, and they result in significant performance improvements over a state-of-the-art *instance-specific* method as well as an ASP solver without SBCs.

This work extends the previous conference paper (Tarzariol et al., 2021) with two additional methods to obtain the positive and negative examples for an ILP task, i.e., an alternative atom ordering criterion and a full symmetry breaking approach, along with corresponding experimental results. Moreover, we provide much more detailed descriptions and analyses of experiments with the suggested language bias, the previous, and two new learning settings.

The structure of this paper is the following: a brief overview of the preliminaries is given in Sect. 2. Section 3 presents our approach, while Sect. 4 describes its implementation and specifies the components of an ILP learning task. In Sect. 5, we investigate observations from learning experiments conducted with our methods. Section 6 provides and discusses experimental results on the solving performance obtained for some combinatorial problems. Lastly, Sect. 7 concludes the paper and outlines directions for future work.

2 Background

This section introduces some basics and notations for ASP, symmetry breaking, and ILP.

2.1 Answer set programming

Answer Set Programming (ASP) (Lifschitz, 2019) 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, showing promising results in numerous domains like, e.g., industrial, robotics, and biomedical applications (Erdem et al., 2016; Falkner et al., 2018). We will briefly define the syntax and semantics of ASP, and refer the reader to (Gebser et al., 2012; Lifschitz, 2019) for more comprehensive 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 $p(\vec{t})$, where p is a predicate, \vec{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{i} is either a variable or a constant, and a *literal* l is an atom a_i (positive) or its negation $\text{not } a_i$ (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} , replacing each rule $r \in P$ with 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 $r \in P_{\text{grd}}$ 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 $r \in P_{\text{grd}}$. A model \mathcal{I} of P is *stable* if it is a subset-minimal model of the reduct $\{H(r) \leftarrow B^+(r) \mid r \in P_{\text{grd}}, B^-(r) \cap \mathcal{I} = \emptyset\}$, and we denote the set of all stable models of P by $AS(P)$.

2.2 Symmetry breaking

Most of the modern instance-specific approaches detect symmetries of a given object by representing it as an instance of the *graph automorphism problem*. This problem consists of finding all edge-preserving bijective mappings of a graph vertex set to itself. It is an attractive target of reduction since this problem can be solved efficiently for many families of graphs using methods from the *group theory*; see (Sakallah, 2009) for an overview.

A *group* is an abstract algebraic structure $\langle G, * \rangle$ where G is a non-empty set, closed under a binary associative operator $*$, such that G contains an identity element, and each element has a unique inverse. If the operator $*$ is implicit, the group is represented by G . A *subgroup* H of a group G is a group such that $H \subseteq G$. Given a set $X = \{x_1, \dots, x_n\}$ of n elements, a *permutation* of X is a bijection that rearranges its elements. The *symmetric group* S_X is defined by the set of all the $n!$ possible permutations of X , and its subgroups are called *permutation groups*. In *cycle notation*, we represent a permutation as a product of disjoint cycles, where each cycle $(x_1 \ x_2 \ x_3 \ \dots \ x_k)$ means that the element x_1 is mapped to x_2 , x_2 to x_3 , and so on, until x_k is mapped back to x_1 ; the elements mapped to themselves are not contained in the cycles.

Given a group G and a set X , the *group action* $\alpha : G \rightarrow S_X$ defines a permutation of the elements in X for each $g \in G$. Then, each permutation $\alpha(g)$ induces a partition on X , $\mathcal{P}^{\alpha(g)}(X)$, whose cells are called *orbits of X under $\alpha(g)$* . The cells of the finest partition on X that refines $\mathcal{P}^{\alpha(g)}(X)$ for each $g \in G$, $\mathcal{P}^G(X)$, are the *orbits of X under group G* .

Example 1 Given a set $\{a, b, c, d, e\}$ of atoms ordered lexicographically and a permutation $\pi = (a \ d \ e) (b \ c)$, let us consider a random interpretation $\{a, c\}$, respectively, the integer 00101. To find its symmetries with respect to π , we repeatedly apply the permutation to $\{a, c\}$ until no new interpretation is obtained, e.g., $a \mapsto d$ and $c \mapsto b$ yielding $\{b, d\}$, then, $d \mapsto e$ and $b \mapsto c$ yielding $\{c, e\}$, and so on. Thus, we get the orbit $\{\{a, c\}, \{b, d\}, \{c, e\}, \{a, b\}, \{c, d\}, \{b, e\}\}$, respectively, $\{00101, 01010, 10100, 00011, 01100, 10010\}$ with the integer representation.

Letting G be a permutation group for a set X of ground atoms, the orbits of the set of interpretations over X under G identify equivalence classes of the truth assignments for X . When taking truth assignments as binary integers determined by some total order of the elements in X , the *lex-leader* scheme consists of specifying a set of Symmetry Breaking

Constraints (SBCs) that may eliminate some interpretations but keep the smallest element in terms of the associated integer for each orbit. If the SBCs eliminate all symmetric assignments and preserve the smallest element of each orbit only, we obtain *full symmetry breaking*. However, dealing with up to $n!$ permutations of n variables explicitly would be an infeasible approach, and a more efficient alternative is to focus on a subset of G to refine $\mathcal{P}^G(X)$. If this leads to a partition finer than $\mathcal{P}^G(X)$, then we get *partial symmetry breaking*. In this case, the SBCs preserve the smallest element as a representative of each orbit but also other symmetric interpretations. Considering a set of *irredundant generators* K of G is an effective heuristic for partial symmetry breaking since such generators represent G compactly. A set $K \subset G$ of elements in a group $\langle G, * \rangle$ is a set of *generators* for G if every element of G can be expressed as a combination of finitely many elements of K under the group operation. K is *irredundant* if no proper subset of it is a set of generators for G .

Example 2 Let us consider the applications of the generator π of the orbit obtained in Example 1. For each interpretation in the orbit, we check whether π maps it into a smaller interpretation according to the integer representation. Thus, the interpretations $\{c, e\}$ and $\{b, e\}$ are eliminated since applying π yields $\{a, b\}$ or $\{a, c\}$, respectively, while the interpretations $\{a, c\}$, $\{b, d\}$, $\{a, b\}$, and $\{c, d\}$ are preserved. As three symmetric interpretations are obtained in addition to the smallest interpretation of the orbit, the representative $\{a, b\}$ (00011 in the integer representation), SBCs for direct applications of the generator π achieve partial symmetry breaking only.

Symmetry-Breaking Answer Set Solving (SBASS) (Drescher et al., 2011) detects and eliminates syntactic symmetries in ASP by adding ground SBCs to an input ground program P_{grd} . A symmetry of P_{grd} is given by a permutation π of ground atoms that keeps the program syntactically unchanged, i.e., P_{grd}^π has the same rules as P_{grd} , where P_{grd}^π is the set of rules obtained by applying π to the head and body literals of rules in P_{grd} . In the first step, SBASS transforms P_{grd} to a colored graph $\mathcal{G}_{P_{grd}}$ such that permutation groups of $\mathcal{G}_{P_{grd}}$ and their generators correspond one-to-one to those of P_{grd} . In the second step, it uses SAUCY (Codenotti et al., 2013; Darga et al., 2004) to find a set of group generators for $\mathcal{G}_{P_{grd}}$. Finally, for each found generator SBASS constructs a set of SBCs based on the lex-leader scheme and appends them to P_{grd} . Given the modified ground program, an ASP solver is provided means to avoid symmetric answer sets.

Example 3 To illustrate how SBASS works, let us consider the pigeon-hole problem, which is about checking whether p pigeons can be placed into h holes such that each hole contains at most one pigeon. An encoding in ASP of this problem is:

```
pigeon(X-1) :- pigeon(X), X > 1.
hole(X-1) :- hole(X), X > 1.
{p2h(P,H) : hole(H)} = 1 :- pigeon(P).
:- p2h(P1,H), p2h(P2,H), P1 != P2.
```

It takes as input the ground facts `pigeon(p)` and `hole(h)`. For example, solving the instance with $p = 3$ and $h = 3$ leads to six answer sets:

$$\begin{aligned}
AS_1 &= \{p2h(1, 1), p2h(2, 2), p2h(3, 3)\} = 100010001 \\
AS_2 &= \{p2h(1, 1), p2h(2, 3), p2h(3, 2)\} = 010100001 \\
AS_3 &= \{p2h(1, 2), p2h(2, 1), p2h(3, 3)\} = 100001010 \\
AS_4 &= \{p2h(1, 2), p2h(2, 3), p2h(3, 1)\} = 001100010 \\
AS_5 &= \{p2h(1, 3), p2h(2, 1), p2h(3, 2)\} = 010001100 \\
AS_6 &= \{p2h(1, 3), p2h(2, 2), p2h(3, 1)\} = 001010100
\end{aligned}$$

The binary integer given on the right is obtained with the following total order of atoms:

$$\begin{aligned}
p2h(3, 3) &> p2h(3, 2) > p2h(3, 1) > \\
p2h(2, 3) &> p2h(2, 2) > p2h(2, 1) > \\
p2h(1, 3) &> p2h(1, 2) > p2h(1, 1)
\end{aligned}$$

Applying SBASS to this pigeon-hole encoding, grounded with the previous input instance, produces the following set of generators:¹

$$\begin{aligned}
\pi_1 &= (p2h(3, 2) \ p2h(3, 3)) \ (p2h(2, 2) \ p2h(2, 3)) \ (p2h(1, 2) \ p2h(1, 3)) \\
\pi_2 &= (p2h(3, 1) \ p2h(3, 3)) \ (p2h(2, 1) \ p2h(2, 3)) \ (p2h(1, 1) \ p2h(1, 3)) \\
\pi_3 &= (p2h(2, 3) \ p2h(3, 3)) \ (p2h(2, 2) \ p2h(3, 2)) \ (p2h(2, 1) \ p2h(3, 1)) \\
\pi_4 &= (p2h(1, 1) \ p2h(3, 3)) \ (p2h(2, 1) \ p2h(2, 3)) \ (p2h(1, 3) \ p2h(3, 1)) \\
&\quad (p2h(1, 2) \ p2h(3, 2))
\end{aligned}$$

According to the lex-leader scheme, we should discard all answer sets but AS_6 , which is the only answer set such that applying either generator does not lead to a greater interpretation:

$$\begin{aligned}
\pi_1(AS_6) &= AS_4 \rightarrow AS_6 \leq AS_4 \text{ since } p2h(3, 1) = p2h(3, 1) \text{ and} \\
&\quad p2h(2, 3) > p2h(2, 2) \\
\pi_2(AS_6) &= AS_1 \rightarrow AS_6 \leq AS_1 \text{ since } p2h(3, 3) > p2h(3, 1) \\
\pi_3(AS_6) &= AS_5 \rightarrow AS_6 \leq AS_5 \text{ since } p2h(3, 2) > p2h(3, 1) \\
\pi_4(AS_6) &= AS_6 \rightarrow AS_6 \leq AS_6
\end{aligned}$$

Therefore, with this instance, we obtain full symmetry breaking by applying the lex-leader scheme for the irredundant generators returned by SBASS. However, symmetric solutions can be preserved for other inputs. E.g., with $p = 3$ and $h = 4$, two answer sets are preserved, while the generators describe a common cell with all answer sets and a single representative.

2.3 Inductive logic programming

Inductive Logic Programming (ILP) (Cropper et al., 2020) 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. Since its foundation, the majority of research in the field addresses Prolog semantics (Cropper and Muggleton, 2016; Muggleton, 1995; Srinivasan, 2004), even though applications in other logic paradigms appeared in the last years. The

¹ The generators are translated from SMOELS to symbolic representation, as described in Sect. 4.

most expressive ILP system for ASP is *Inductive Learning of Answer Set Programs* (ILASP) (Law et al., 2014, 2021), which can be used to solve a variety of ILP tasks.

In ILP, a learning task $\langle P, E, H_M \rangle$ is defined by three elements: a background knowledge P , a set of (positive and negative) examples E , and a hypothesis space H_M , which defines all the rules that can be learned. The learned hypothesis is a subset of the hypothesis space that satisfies a specified learning setting: for ILASP, the setting is *learning from answer sets* (Law et al., 2014). Before defining it, we introduce the terminology used by ILASP authors. A *Partial Interpretation* (PI) is a pair of sets of atoms, $e_{pi} = \langle T, F \rangle$, called *inclusions* (T) and *exclusions* (F), respectively. Given a (total) interpretation \mathcal{I} and a PI e_{pi} , we say that \mathcal{I} *extends* e_{pi} if $T \subseteq \mathcal{I}$ and $F \cap \mathcal{I} = \emptyset$. We can augment e_{pi} with an ASP program C to obtain a *Context Dependent Partial Interpretation* (CDPI) $\langle e_{pi}, C \rangle$. Given a program P , a CDPI $e = \langle e_{pi}, C \rangle$, and an interpretation \mathcal{I} , 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} .

A learning task for ILASP is given by an ASP program P as background knowledge, two sets of CDPIs, E^+ and E^- , as positive and negative examples, and the hypothesis space H_M defined by a language bias M , which limits the potentially learnable rules. The learned hypothesis $H \subseteq H_M$ must respect the following criteria: (i) for each positive example $e \in E^+$, there is some accepting answer set of e with respect to $P \cup H$; and (ii) for any negative example $e \in E^-$, there is no accepting answer set of e with respect to $P \cup H$. If multiple hypotheses satisfy the conditions, the system returns one of the shortest, i.e., with the minimum number of literals (Law et al., 2014). In Law et al. (2018), the authors extend the expressiveness of ILASP by allowing noisy examples. With this setting, if an example e is not covered (i.e., there is an accepting answer set for e if it is negative, and none, if it is positive), the corresponding weight is counted as a penalty. Therefore, the learning task becomes an optimization problem with two goals: minimize the size of H and minimize the total penalties for the uncovered examples.

Now, we will define the syntax of ILASP necessary for our work and refer the reader to the system's manual (Law et al., 2021) for further details. A CDPI is expressed as follows:

$$\#type(ID@W, \{Inc\}, \{Exc\}, \{C\}).$$

where `type` is either `pos` or `neg`, `ID` is a unique identifier for the example, `W` is a positive integer representing the example's weight (if not defined, the weight is infinite), `Inc` and `Exc` are two sets of atoms, and `C` is an ASP program. The language bias can be specified by *mode declarations*, which define the predicates that may appear in a rule, their argument types, and their frequency. Since in our work we aim to learn constraints, we restrict the search space just to rules r with $H(r) = \perp$. Hence, we only need to specify the mode declarations for the body of a rule, expressed by `#modeb(R, P, (E))` where `R` and `E` are optional and `P` is a ground atom whose arguments are placeholders of type `var(τ)` for some constant term τ . In the learned rules, the placeholders will be replaced by variables of type τ . The optional element `R` is a positive integer, called *recall*, which specifies the maximum number of times that the mode declaration can be used in each rule. Lastly, `E` is a condition that further restricts the hypothesis space. We limit our interest to the `anti_reflexive` option that works with predicates of arity 2. When using it, atoms of the predicate `P` should be generated with two distinguished argument values.

Choosing an appropriate language bias is still one of the major challenges for modern ILP systems. Whenever the bias does not provide enough limitations, the problem becomes intractable and ILASP might not be able to find useful constraints. In contrast, a too strong

bias may exclude solutions from the search space, thus resulting in suboptimal SBCs (Cropper & Dumančić, 2020).

3 Approach

We tackle combinatorial problems modeled in ASP such that the instances of a logic program P are generated by a discrete and often stationary stochastic process. Such situations occur, e.g., in industrial settings where the encoding of a manufacturing system is fixed and production orders vary. In this case, every problem instance can be seen as an outcome of the process. We assume that any instance (i) specifies the (*true*) atoms of unary domain predicates p_1, \dots, p_k in P , where c_i is the number of atoms that hold for each p_i ; and (*ii*) the satisfiability of the instance depends on the number of atoms for each domain predicate, but not on the values of their terms. Thus, without loss of generality, we consider the terms for each p_i to be consecutive integers from 1 to c_i .

Our method exploits *instance-specific* SBCs on a representative set of instances and utilizes them to generate examples for an ILP task. The learning method yields first-order constraints that remove symmetries in the analyzed problem instances as much as possible while preserving the instances' satisfiability. We consider the following two learning settings:

- *enum* is a cautious setting that preserves all answer sets that are not filtered out by the ground SBCs; and
- *sat* setting aims to learn tighter constraints which, however, preserve at least one answer set for each instance.

To compute the examples, our approach relies on small satisfiable instances (i.e., with a low value for each c_i), subdivided into two parts: S and Gen . Each instance $g \in Gen$

defines a positive example with empty inclusions and exclusions, and g as context. These examples, denoted by Ex_{Gen} , guarantee that the learned constraints generalize for the target distribution since they force the constraints to preserve some solution for each $g \in Gen$. The instances $i \in S$ are used to obtain positive and negative examples, representing answer sets of $P \cup i$ to be preserved or filtered out, respectively, by corresponding SBCs. We denote their union by Ex_S in Fig. 1, where positive examples represent whole answer sets in the *enum* setting, or like instances in Gen , consist of empty inclusions and exclusions along with the context i in *sat*.

An ILP task further requires background knowledge and a hypothesis space H_M . Both of them are defined by the user (for a possible instantiation, see Sect. 4.3). The *background knowledge* consists of a logic program P along with an *Active Background Knowledge*, denoted by ABK in Algorithm 1. We use ABK to simplify the management of auxiliary predicate definitions and constraints learned so far. The *hypothesis space* contains the mode declarations, and we assume it to be general enough to entail ground SBCs by learned first-order constraints. The remaining inputs of Algorithm 1 consist of the instances in Gen and S as well as the learning setting m . For each answer set \mathcal{I} of an instance $i \in S$ to be analyzed, the algorithm determines T and F by projecting \mathcal{I} to the atoms occurring in IG , denoted by $atoms(IG)$. Next, in line 10, the predicate $lexLead(\langle T, F \rangle, IG)$ evaluates to *true* if \mathcal{I} is dominated, i.e., \mathcal{I} can be mapped to a lexicographically smaller, symmetric answer set by means of some irredundant generator in IG . In this case, the negative

Algorithm 1: Framework to lift SBCs with ILP

```

input :  $P, ABK, H_M, Gen, S, m$ 
1  $Ex_{Gen} \leftarrow \emptyset;$ 
2  $Ex_S \leftarrow \emptyset;$ 
3 foreach  $g \in Gen$  do
4    $Ex_{Gen} \leftarrow Ex_{Gen} \cup \{pos(\emptyset, \emptyset, g)\};$ 
5 foreach  $i \in S$  do
6    $IG \leftarrow$  Set of irredundant generators for  $i;$ 
7   foreach  $\mathcal{I} \in AS(P \cup i \cup ABK)$  do
8      $T \leftarrow atoms(IG) \cap \mathcal{I};$ 
9      $F \leftarrow atoms(IG) \setminus \mathcal{I};$ 
10    if  $lexLead(\langle T, F \rangle, IG)$  then
11       $Ex_S \leftarrow Ex_S \cup \{neg(T, F, i)\};$ 
12    else if  $m = enum$  then
13       $Ex_S \leftarrow Ex_S \cup \{pos(\mathcal{I}, \emptyset, i)\};$ 
14    else
15       $Ex_S \leftarrow Ex_S \cup \{pos(\emptyset, \emptyset, i)\};$ 
16  $C \leftarrow$  Solve  $\langle P \cup ABK, Ex_{Gen} \cup Ex_S, H_M \rangle;$ 
17  $ABK \leftarrow ABK \cup C;$ 

```

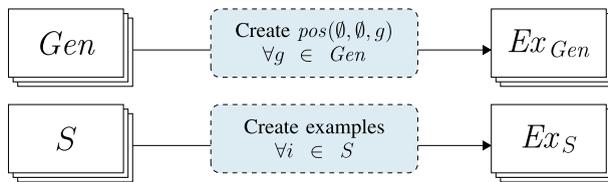


Fig. 1 ILP examples generation

example $neg(T, F, i)$ is added to Ex_S in order to eliminate \mathcal{I} , while $pos(\mathcal{I}, \emptyset, i)$ or $pos(\emptyset, \emptyset, i)$ is taken as the positive example otherwise, depending on whether the *enum* or *sat* setting is selected. Positive examples of the form $pos(\emptyset, \emptyset, g)$ are also gathered in Ex_{Gen} for instances $g \in Gen$, and solving the ILP task at line 16 gives new constraints C to extend ABK .

4 Implementation

The implementation of our framework relies on CLINGO (consisting of the grounding and solving components GRINGO and CLASP), SBASS and ILASP, and is available at Tarzariol et al. (2021). Figure 2 shows the pipeline to generate the examples for a given instance $i \in S$ (the for-loop at line 5 of Algorithm 1). First, the union of P , i , and ABK is grounded with GRINGO to get the ground program P_{grd} in SMODELS format. Then, the solver CLASP enumerates all its answer sets, obtaining $AS(P_{grd})$. Independently, SBASS is run on P_{grd} with the option `-show` to output a set of irredundant permutation group generators. This set contains the vertex permutations of $\mathcal{G}_{P_{grd}}$, expressed in cycle notation. We extract the cycles defined by vertices representing atoms of P_{grd} and transform them from SMODELS format back into their original symbolic representation (by a predicate and integer terms).

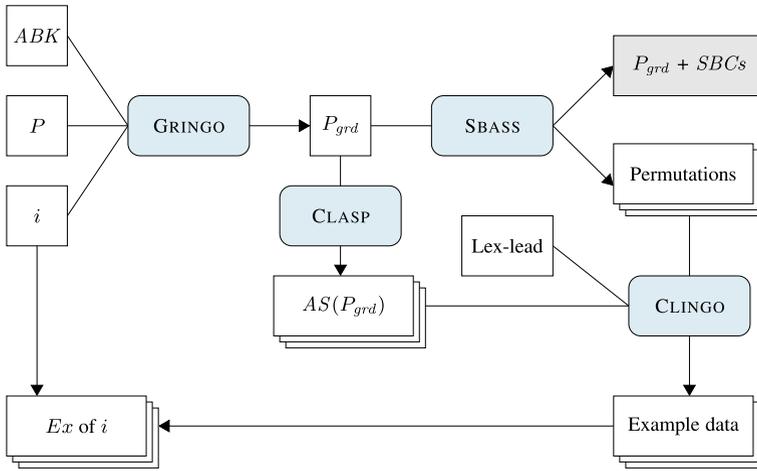


Fig. 2 Pipeline to compute examples from an instance i

Next, we identify the symmetric answer sets in $AS(P_{grd})$ by using an ASP encoding similar to the lex-leader predicate definition in Sakallah (2009) to evaluate SBCs. To this end, we implement an ordering criterion to compare atoms according to their signatures. Given two ground atoms $p_1(a_1, \dots, a_n)$ and $p_2(b_1, \dots, b_m)$, the first is considered smaller than the second if: (i) p_1 is lexicographically smaller than p_2 ; (ii) $p_1 = p_2$ and $n < m$; or (iii) $p_1 = p_2$, $n = m$, and there are constants $a_i < b_i$ such that $a_j = b_j$ for all $0 < j < i$. Our ASP encoding then checks whether an answer set $\mathcal{I} \in AS(P_{grd})$ is undominated by interpretations obtainable by applying the symbolic representation of some irredundant generator returned by SBASS to \mathcal{I} .

In case \mathcal{I} is dominated and thus must be eliminated as a symmetric answer set, we map it to a negative example with a unique identifier and a weight of 100. Due to the weights, ILASP returns a set of constraints even if some negative examples are not covered. Moreover, we use uniform weights so that all negative examples have the same relevance and as many as possible are to be eliminated. Lastly, answer sets that were not found to be dominated for any of the generators yield positive examples according to the selected setting—*enum* or *sat*. Such positive examples are unweighted so that the learned hypothesis must cover all of them.

4.1 Alternative atom ordering

Let us consider sets of n lexicographically ordered atoms that only differ in the values of the last terms in each atom. For two such sets $A = \{p(\bar{x}_1, a_1), \dots, p(\bar{x}_n, a_n)\}$ and $B = \{p(\bar{x}_1, b_1), \dots, p(\bar{x}_n, b_n)\}$ of atoms, where \bar{x}_i contains all terms but the last, the lex-leader scheme starts by checking the atoms with the greatest \bar{x}_i vectors until there are two constants $a_i \neq b_i$. Since various configuration problems yield answer sets of this kind, we devised an alternative atom ordering such that the lex-leader scheme starts from the smallest \bar{x}_i vectors when comparing two answer sets. To this end, an atom $p_1(a_1, \dots, a_n)$ is considered smaller than $p_2(b_1, \dots, b_m)$ if: (i) p_1 is lexicographically smaller than p_2 ; (ii) $p_1 = p_2$ and $n < m$; (iii) $p_1 = p_2$, $n = m$, and there are constants $a_i > b_i$ such that $i < n$ and $a_j = b_j$ for all $0 < j < i$; or (iv) $p_1 = p_2$, $n = m$, $a_i = b_i$ for all $0 < i < n$, and $a_n < b_n$.

This alternative ordering allows for more natural, undominated answer sets, as illustrated in the following example.

Example 4 Applying the alternative ordering criterion to the same input as described in Example 3, we get the following total order of atoms:

$$\begin{aligned} p2h(1, 3) &> p2h(1, 2) > p2h(1, 1) > \\ p2h(2, 3) &> p2h(2, 2) > p2h(2, 1) > \\ p2h(3, 3) &> p2h(3, 2) > p2h(3, 1) \end{aligned}$$

Thus, the integers associated with the answer sets become:

$$\begin{aligned} AS_1 &= \{p2h(1, 1), p2h(2, 2), p2h(3, 3)\} = 001010100 \\ AS_2 &= \{p2h(1, 1), p2h(2, 3), p2h(3, 2)\} = 001100010 \\ AS_3 &= \{p2h(1, 2), p2h(2, 1), p2h(3, 3)\} = 010001100 \\ AS_4 &= \{p2h(1, 2), p2h(2, 3), p2h(3, 1)\} = 010100001 \\ AS_5 &= \{p2h(1, 3), p2h(2, 1), p2h(3, 2)\} = 100001010 \\ AS_6 &= \{p2h(1, 3), p2h(2, 2), p2h(3, 1)\} = 100010001 \end{aligned}$$

Now the lex-leader scheme discards all but the answer set AS_1 , and three permutations map AS_6 to smaller answer sets:

$$\begin{aligned} \pi_1(AS_6) &= AS_4 \rightarrow AS_6 > AS_4 \text{ since } p2h(1, 3) > p2h(1, 2) \\ \pi_2(AS_6) &= AS_1 \rightarrow AS_6 > AS_1 \text{ since } p2h(1, 3) > p2h(1, 1) \\ \pi_3(AS_6) &= AS_5 \rightarrow AS_6 > AS_5 \text{ since } p2h(1, 3) = p2h(1, 3) \text{ and} \\ & \hspace{10em} p2h(2, 2) > p2h(2, 1) \\ \pi_4(AS_6) &= AS_6 \rightarrow AS_6 \leq AS_6 \end{aligned}$$

The answer set AS_1 contains atoms that are preserved by the general first-order constraint : $\neg p2h(P, H), P < H$, which removes all other symmetric solutions. Unlike that, when taking AS_6 as a representative solution, we have to distinguish particular cases for the assignment of the first and the last pigeon, resulting in longer and more specific constraints.

4.2 Exploiting generators for full symmetry breaking

When investigating irredundant generators to label an answer set as a positive or negative example according to the lex-leader scheme, there can be cases where the labeling achieves partial instead of full symmetry breaking. As illustrated in Example 2, this is because single applications of generators yield a subset of the orbit of an interpretation only. Thus, we implement an alternative setting to label the examples, named *fullSBCs*, which exploits generators to explore the whole orbit of symmetric interpretations for every answer set. For each of the obtained cells, we label the smallest answer set as a positive example and all the remaining ones as negative. This approach reduces the sensitivity of ILP tasks to particular irredundant generators returned by SBASS, allowing to achieve full symmetry breaking for any instance $i \in S$.

We implement this setting by means of the CLINGO API² to interleave the solving phase, which returns a candidate answer set, with the analysis of its orbit. Then, before continuing with the search for the next answer set, we prohibit the explored interpretations by feeding respective constraints to CLINGO. This setting allows for reducing the number of positive examples produced, and as we can configure it to sample a limited subset of all answer sets, it is also useful for dealing with underconstrained configuration problems that yield plenty of answer sets even for very small instances.

Example 5 To illustrate the *fullSBCs* setting, let us reconsider the pigeon-hole problem introduced in Example 3, where the instance with three pigeons and four holes leads to 24 solutions. Running sbass on this instance yields five generators, which identify a single cell since all the answer sets are symmetric. However, we only consider the first two generators in the following, allowing us to demonstrate the *fullSBCs* approach on an example with several, i.e., four, cells. The generators we inspect are:

$$\begin{aligned} \pi_1 &= (p2h(3, 2) \ p2h(3, 3)) \ (p2h(2, 2) \ p2h(2, 3)) \ (p2h(1, 2) \ p2h(1, 3)) \\ \pi_2 &= (p2h(3, 3) \ p2h(3, 4)) \ (p2h(2, 3) \ p2h(2, 4)) \ (p2h(1, 3) \ p2h(1, 4)) \end{aligned}$$

Let $AS_1 = \{p2h(1, 3), \ p2h(2, 2), \ p2h(3, 4)\}$ be the first answer set found. Then, before searching for other solutions, we repeatedly apply π_1 and π_2 to AS_1 to obtain the whole orbit of symmetric interpretations. The identified answer sets are:

$$\begin{aligned} AS_1 &= \{p2h(1, 3), \ p2h(2, 2), \ p2h(3, 4)\} = 010000101000 \\ \pi_1(AS_1) = AS_2 &= \{p2h(1, 2), \ p2h(2, 3), \ p2h(3, 4)\} = 001001001000 \\ \pi_2(AS_1) = AS_3 &= \{p2h(1, 4), \ p2h(2, 2), \ p2h(3, 3)\} = 100000100100 \\ \pi_1(AS_2) &= AS_1 \\ \pi_2(AS_2) = AS_4 &= \{p2h(1, 2), \ p2h(2, 4), \ p2h(3, 3)\} = 001010000100 \\ \pi_1(AS_3) = AS_5 &= \{p2h(1, 4), \ p2h(2, 3), \ p2h(3, 2)\} = 100001000010 \\ \pi_2(AS_3) &= AS_1 \\ \pi_1(AS_4) = AS_6 &= \{p2h(1, 3), \ p2h(2, 4), \ p2h(3, 2)\} = 010010000010 \\ \pi_2(AS_4) &= AS_2 \\ \pi_1(AS_5) &= AS_3 \\ \pi_2(AS_5) &= AS_6 \\ \pi_1(AS_6) &= AS_4 \\ \pi_2(AS_6) &= AS_5 \end{aligned}$$

Once we have computed all answer sets symmetric to AS_1 , we produce a positive example for the smallest answer set encountered, i.e., AS_2 , while the other five answer sets constitute negative examples. Now, we can proceed with the search for the next answer set, e.g., $AS_7 =$

² A complete reference documentation can be found at <https://potassco.org/clingo/python-api/current/>.

Algorithm 2: *fullSBCs* method to generate Ex_S for an instance i

```

input :  $P, i, ABK, IG$ 
1  $cnt \leftarrow \text{CLINGO.init}(P \cup i \cup ABK)$ ;
2 while  $\mathcal{I} \leftarrow cnt.get\_new\_solution()$  do
3    $min \leftarrow atoms(IG) \cap \mathcal{I}$ ;
4    $Q \leftarrow \text{new\_queue}()$ ;
5    $Q.push(min)$ ;
6    $seen \leftarrow \{min\}$ ;
7   while  $T \leftarrow Q.pop()$  do
8     if  $T > min$  then
9        $Ex_S \leftarrow Ex_S \cup \{neg(T, atoms(IG) \setminus T, i)\}$ ;
10    else if  $T < min$  then
11       $Ex_S \leftarrow Ex_S \cup \{neg(min, atoms(IG) \setminus min, i)\}$ ;
12       $min \leftarrow T$ ;
13    forall  $\pi \in IG$  do
14      if  $\pi(T) \notin seen$  then
15         $Q.push(\pi(T))$ ;
16         $seen \leftarrow seen \cup \{\pi(T)\}$ ;
17   $Ex_S \leftarrow Ex_S \cup \{pos(min, atoms(IG) \setminus min, i)\}$ ;
18   $cnt.ignore\_solutions(seen)$ ;

```

$\{p2h(1, 2), p2h(2, 1), p2h(3, 3)\}$, and repeat the application of generators to explore its cell, identifying another five symmetric solutions of which $\{p2h(1, 3), p2h(2, 1), p2h(3, 2)\}$ is the smallest. This process continues until all 24 answer sets, partitioned into four cells with a smallest representative for each, are explored.

Algorithm 2 outlines the *fullSBCs* approach, providing an alternative implementation of the for-loop at line 7 of Algorithm 1. In the first line, we create a search control object, *cnt*, using the CLINGO API. This object keeps track of already identified solutions and provides the *get_new_solution* method, which returns either a new answer set \mathcal{I} or *false* if all solutions have been exhausted. Similar to the previously presented approaches to example generation, we project the atoms of \mathcal{I} to $atoms(IG)$. The resulting interpretation *min* represents the smallest solution encountered so far in the current cell, and the set *seen* keeps track of already discovered interpretations belonging to the current cell. Starting with *min*, the queue Q collects the interpretations to which all irredundant generators will be applied to yield new symmetric interpretations. The while-loop at line 7 checks whether there is an interpretation, T , left to pop. Then, if T is greater than *min* (according to the applied atom ordering criterion), it constitutes a negative example, while a smaller T is taken as new smallest interpretation and the previous *min* instead becomes a negative example. Only after the cell has been completely explored, the interpretation *min* is eventually labeled as a positive example. Lastly, before querying *cnt* for the next solution, we eliminate answer sets subsumed by the explored interpretations in *seen* from the search space of *cnt*.

4.3 ILP learning task

After considering the example generation, we specify components of the ILP learning task suitable for the learning of constraints. The idea is to encode the predicates used by lex-leader symmetry breaking to order atoms and extract the maximal values for domain predicates. Since the mode declarations of ILASP (v4.0.0) do not support arithmetic built-ins such as $<$, we provide auxiliary predicates in *ABK* to simulate them. Presupposing the presence of unary domain predicates p_1, \dots, p_k with integers from 1 to c_i for each p_i , *ABK* defines the auxiliary predicates $\max p_i(c_i)$ for each p_i and $\text{lessThan}(t_1, t_2)$ for each pair of integers

$1 \leq t_1 < t_2 \leq \max\{c_i \mid i = 1, \dots, k\}$. These two predicates, based exclusively on syntactic properties of a considered problem, are minimal for overcoming limitations of *ILASP* to learn lex-leader SBCs. The selection of small yet representative instances for *S* and *Gen* depends on their hardness for learning. Regarding *S*, we pursued the strategy to empirically determine instances for which *SBASS* yields a manageable number of permutation group generators. As mentioned in Sect. 4.2, the irredundant generators alone sometimes achieve partial symmetry breaking, and we selected only instances without any or a small amount of “misclassified” answer sets. The instances in *Gen* are usually larger yet still solvable in a short running time to check that the learned constraints generalize.

The language bias of our learning task includes the mode declarations $\#modeb(2, p_i(\text{var}(t_i)))$ and $\#modeb(1, \max p_i(\text{var}(t_i)))$ for each domain predicate p_i , in which $\text{var}(t_i)$ is a placeholder indicating the domain for which each p_i holds. Moreover, for each (non-auxiliary) predicate P appearing in some of the generators computed for instances in *S*, we use $\#modeb(2, P)$, where the domains of variables in atoms of P are provided by a vector of the placeholders in $\{\text{var}(t_i) \mid i = 1, \dots, k\}$, depending on the role of P in the given program P . In addition, we include mode declarations $\#modeb(2, \text{lessThan}(\text{var}(t_i), \text{var}(t_j)))$ for all $i, j = 1, \dots, k$, with the option `anti_reflexive` in case $i = j$.

We decided to distinguish the variables’ types in the mode declarations in order to restrict the hypothesis space to rules such that a variable X of type t is included as an argument only in predicates defined over the same type t . To illustrate how this decision influences the search space of an ILP task, let us consider two extensions of the pigeon-hole problem introduced in Example 3, adding color and owner assignments. The pigeon-hole problem with colors associates a color with each pigeon and requires pigeons placed into neighboring holes to be of the same color. The version with colors and owners additionally assigns an owner to each pigeon and imposes the same constraint as with the colors for owners as well. For the pigeon-hole problem with colors, by using typed variables in the mode declarations, *ILASP* generates a search space of 1837 rules,³ while 9169 rules are obtained without distinguishing variables’ types. Regarding the extension to owners, this difference is even larger: 2895 rules using typed variables versus 21406 rules without distinguishing variables’ types.

To compare the learning performance of *ILASP*, we conducted several experiments on the pigeon-hole problem with colors and owners for a pool of instances⁴ and observed that applying our approach with typed variables in the mode declarations allows for learning constraints quicker than without distinguishing the types. When using the iterative approach described in Sect. 4.4, *ILASP* took on average less than two minutes to learn the shortest constraints related to holes, colors, and owners, and always finished in less than ten minutes. In opposite, a similar ILP task defined without distinction of variable types took on average thirty minutes, with cases where no hypothesis was found within an hour.

Reducing the hypothesis space has the potential drawback of learning less efficient rules since there can be situations where stronger constraints with fewer variables are excluded. For instance, a constraint like `:- pigeon(X), not p2h(X,X)` cannot be learned in the current setting, as the variable X is taken for a pigeon and a hole at the same time. However, we decided to use the restricted search space for our experiments in Sect. 6 because it leads to much better scalability of learning and constraints that still improve the solving performance.

³ For all our experiments, we used the default value (3 literals) for the *ILASP* parameter that defines the maximum number of literals that can occur together in the body of each rule of the hypothesis.

⁴ The collected data can be found at Tazariol et al. (2021). The experiments were run on an Intel® i7-3930K machine under Linux (Debian GNU/Linux 10), where each run of *ILASP4* was limited to 3600 s.

In fact, the ability to learn constraints in acceptable time is important for handling application scenarios better than with instance-specific symmetry breaking methods.

Example 6 To illustrate a feasible outcome of our ILP framework, let us inspect the constraints learned for the pigeon-hole problem and its instance with three pigeons and three holes, as also considered in Example 3. Applying the generators returned by SBASS to the six answer sets gives one positive and five negative examples, and the resulting ILP task is as follows:

```

%% Input encoding adapted for ILASP
pigeon(X-1) :- pigeon(X), X > 1.
hole(X-1) :- hole(X), X > 1.
0 {p2h(P,H)} 1 :- pigeon(P), hole(H).
:- p2h(P1,H), p2h(P2,H), P1 < P2.
:- p2h(P,H1), p2h(P,H2), H1 < H2.
:- pigeon(P), not p2h(P,_).

%% Active Background Knowledge
lessThan(X,Y) :- pigeon(X), pigeon(Y), X < Y.
lessThan(X,Y) :- hole(X), hole(Y), X < Y.
maxpigeon(X) :- pigeon(X), not pigeon(X+1).
maxhole(X) :- hole(X), not hole(X+1).

%% Negative examples
#neg(id1@100, {p2h(2,3), p2h(1,2), p2h(3,1)},
  {p2h(2,1), p2h(1,1), p2h(3,3), p2h(1,3), p2h(3,2), p2h(2,2)},
  {pigeon(3). hole(3).}).
#neg(id3@100, {p2h(2,1), p2h(3,2), p2h(1,3)},
  {p2h(1,1), p2h(3,3), p2h(3,1), p2h(2,2), p2h(2,3), p2h(1,2)},
  {pigeon(3). hole(3).}).
#neg(id4@100, {p2h(2,3), p2h(1,1), p2h(3,2)},
  {p2h(2,1), p2h(3,3), p2h(3,1), p2h(1,3), p2h(2,2), p2h(1,2)},
  {pigeon(3). hole(3).}).
#neg(id5@100, {p2h(2,1), p2h(3,3), p2h(1,2)},
  {p2h(1,1), p2h(3,1), p2h(1,3), p2h(3,2), p2h(2,3), p2h(2,2)},
  {pigeon(3). hole(3).}).
#neg(id6@100, {p2h(1,1), p2h(3,3), p2h(2,2)},
  {p2h(2,1), p2h(3,1), p2h(1,3), p2h(3,2), p2h(2,3), p2h(1,2)},
  {pigeon(3). hole(3).}).

%% Positive example produced with enum setting
#pos(id2, {p2h(3,1), p2h(2,2), p2h(1,3)}, {},
  {pigeon(3). hole(3).}).

%% Language bias
#modeb(2,p2h(var(pigeon),var(hole))).
#modeb(2,pigeon(var(pigeon))).
#modeb(2,hole(var(hole))).
#modeb(1,maxhole(var(hole))).
#modeb(1,maxpigeon(var(pigeon))).
#modeb(2,lessThan(var(hole),var(hole)),(anti_reflexive)).
#modeb(2,lessThan(var(pigeon),var(pigeon)),(anti_reflexive)).
#modeb(2,lessThan(var(hole),var(pigeon))).
#modeb(2 lessThan(var(pigeon) var(hole)))

```

Let us notice that the ASP input encoding in Example 3 has been adapted into an equivalent one above. Such a modification is necessary because the current version of `ILASP` does not support rules like `{p2h(P,H) : hole(H)} = 1 :- pigeon(P).` with the conditional operator “:” in the head. After running `ILASP`, the learned first-order constraints are:

```
:- p2h(X,Y), lessThan(Z,Y), maxpigeon(X).
% do not assign the pigeon with the max label to a hole
% other than the first one
:- p2h(X,Y), lessThan(X,Y), lessThan(Y,Z).
% for all but the last hole, do not assign a pigeon with
% a smaller label to the hole
```

4.4 Iterative learning

Inspired by the *lifelong learning* approach (Cropper et al. 2020), we apply our framework incrementally to a split learning task. This idea is especially useful if the ASP encoding presents several symmetries, where some of them are independent of the others. The iterative approach simplifies the learning task by exploiting the incremental applicability of ILP: first, it solves a subtask to identify a subset of symmetries, and before addressing the remaining ones, we integrate the constraints just learned into the background knowledge. To this end, we divide the hypothesis space for programs with three or more types of variables in the language bias. Then, in the first step, we provide a set S of instances to address their symmetries involving only two types of variables and define the search space with mode declarations restricted to the two types of variables considered. Next, we solve the ILP subtask and append the learned constraints to ABK . In the following steps, we repeat the procedure and analyze the same or different instances in S for symmetries going beyond those already handled by solving ILP subtasks with the mode declarations progressively extended to further types of variables.

To illustrate a concrete application scenario, reconsider the pigeon-hole problem with color and owner assignments, introduced in Sect. 4.3. For this problem, the search space is split into three incremental parts:

- the first is limited to predicates whose atoms exclusively include variables of the types `pigeon` and `hole`,
- the second part extends mode declarations by allowing atoms with variables of the type `color` too, and, finally,
- the third step includes variables of the type `owners`.

Initially, S contains instances with only one color and owner so that our framework produces examples entailing symmetries related exclusively to the pigeons’ placement. Next, we append the learned constraints to ABK and repeat the procedure by redefining S with instances with one owner but more than one color. Lastly, we analyze instances in S without restrictions on the numbers of colors and owners while considering the whole language bias. In this way, `ILASP` can learn new symmetries using predicates that involve more types of variables, as the language bias is progressively extended until it reaches the whole set of mode declarations.

By applying the iterative approach, the learning task can be decomposed into smaller and easier ILP subtasks. For an indication of the practical impact on the size of the search space(s), we note that `ILASP` generates 1040 rules for variables of the types `pigeon` and `hole` only, 1837 rules when variables of the type `color` are added, and 2895 rules with the full language bias for the pigeon-hole problem with colors and owners. That is, the search space for the ILP subtask in the last iteration includes the same rules as generated when addressing the full language bias in a single step, yet the background knowledge may already be extended by constraints reducing the number of (negative) examples still to investigate. We assessed the impact of the iterative approach in several experiments experiments (see footnote 4) and observed that it allows us to learn constraints much quicker than when tackling all symmetries in a single pass. By splitting the learning task, `ILASP` took on average less than two minutes to learn constraints related to pigeons and holes, then colors, and finally owner assignments. Unlike that, the ILP task that addresses the full language bias directly took on average more than thirty minutes to return the shortest hypothesis, where in some cases the search was not finished within one hour.

Splitting the learning task has the potential drawback that some of the symmetries can be lost in the process, as the updated `ABK` is considered in subsequent calls to `SBASS` for identifying remaining symmetries. However, for the combinatorial problems investigated in our experiments in Sect. 6, the results showed that even in case we learn constraints handling a subset of all problem symmetries, the solving performance benefits substantially.

5 Learning performance

We tested the different settings of our implementation over the two extensions of the pigeon-hole problem described in Sect. 4.3. For every setting, we used the same initial set of instances in `Gen`, auxiliary constraints in `ABK`, and mode declarations in the language bias (split to apply the iterative approach). For keeping the number of instances in `Gen` moderate, we hand-picked a few (satisfiable) instances to start from, applied our iterative learning approach, and then validated the learned constraints on other satisfiable instances as well. The instances for which learned constraints led to unsatisfiability were then also added to `Gen`, and we repeated the learning phase until all instances were found satisfiable together with the learned first-order SBCs.

In the following, we report representative results and conclusions drawn from the instances and records of learning experiments provided in our repository (Tarzariol et al., 2021).

5.1 Enum vs sat setting

The difference between the *enum* and *sat* setting lies in the positive examples generated for the instances in `S`: in the first setting, we explicitly list all undominated answer sets as positive examples, while the second produces just a general positive example with empty inclusions and exclusions. That is, the *sat* setting abstracts over undominated answer sets, as they are neither labeled as positive nor negative examples in the ILP task. In this case, `ILASP` aims at eliminating as many symmetric answer sets as possible while preserving the satisfiability of a given instance. This even means that the preserved answer sets, required in view of the general positive example, might belong to negative examples but are not covered by the learned constraints. In this way, we may, in general, learn alternative constraints that

preserve some specific pattern of solutions appearing in all satisfiable instances, regardless of symmetries, while representative solutions can be lost.

For example, in the first ILP iteration on the pigeon-hole problem with colors and owners, the instance with three pigeons and four holes (and only one color and owner) gives 24 answer sets, 22 of which are labeled as negative. In the *enum* setting, *ILASP* finds optimal constraints removing 12 negative examples and thus returns a hypothesis that applied to the same instance leaves 12 answer sets. In contrast, the *sat* setting enables learning of stronger constraints by *ILASP*, which preserve 2 answer sets only, both labeled as negative examples.

The complexity of the ILP task depends on the possibility of covering all negative examples. Since with *enum* we have tighter conditions on the candidate hypotheses, the search space is smaller than in the *sat* setting. Hence, the optimization problem regarding (weighted) negative examples addressed by *ILASP* takes in general longer for *sat*, but only if many negative examples cannot be covered even under relaxed conditions on the candidate hypotheses. On the other hand, if the language bias permits many hypotheses covering all or most of the negative examples, an ILP task is usually quickly solved with the *sat* setting. E.g., the instance with three pigeons, four holes, one color and owner took 72.8 seconds to be solved in the *enum* (eliminating 12 out of 22 symmetric answer sets) and just 27.7 seconds in the *sat* setting (eliminating 20 symmetric answer sets and the 2 unlabeled ones).

5.2 Alternative atom ordering

When answer sets for the combinatorial problem analyzed have the property described in Sect. 4.1, our alternative ordering criterion for the lex-leader scheme may distinguish the representative and symmetric solutions. Hence, ILP tasks can be solved with shorter constraints than for the default atom ordering. For instance, the setting illustrated in Example 6 yields the representative answer set $\{p2h(1, 1), p2h(2, 2), p2h(3, 3)\}$ instead of $\{p2h(1, 3), p2h(2, 2), p2h(3, 1)\}$. This allows *ILASP* to learn the short constraint $:- p2h(X, Y), lessThan(Y, X).$, expressing that no pigeon can be placed into a hole smaller than its label, which leaves just one answer set for instances with an equal number $p = h$ of pigeons and holes.

Given that the positive examples kept after checking direct applications of the irredundant generators returned by *SBASS* heavily depend on the computed generators, we found that often more positive examples are produced than for the default atom ordering. Namely, the generators preserve more symmetric solutions than the default ordering for the extensions of the pigeon-hole problem to colors as well as colors and owners. This leads to weaker (although shorter and easier to interpret) constraints, and better-suited ways of aligning generators with symbolic atom representations would be of interest.

5.3 Exploiting generators for full symmetry breaking

Section 4.2 describes an alternative implementation for labeling answer sets as positive or negative examples, called the *fullSBCs* setting. We can see the effects of always labeling the answer sets according to full SBCs on the same scenario as discussed in Sect. 5.1: instead of 22 negative and 2 positive examples generated with the *enum* setting, *fullSBCs* returns just one positive example, i.e., the representative of the single cell characterized by the generators of *SBASS*. As a consequence, instead of 72.8 seconds to return a hypothesis that produces 12 of the original 24 answer sets, with the ILP task defined based on *fullSBCs*, *ILASP* took 21.4 seconds to find a hypothesis that preserves only 4 answer sets.

Table 1 Runtime in seconds for pigeon-hole problem

	ENUM	SAT	ORD	FULL	BASE	SBASS	CLASP ^x
p50-h49	0.092	0.092	0.084	0.091	TO	52.829	1.900
p50-h50	0.080	0.080	0.073	0.082	0.116	53.408	1.417
p100-h99	0.635	0.715	0.640	0.715	TO	TO	–
p100-h100	0.707	0.627	0.574	0.622	1.047	TO	–
p200-h199	5.602	5.861	5.192	5.656	TO	TO	–
p200-h200	5.772	5.907	6.004	5.627	11.162	TO	–
p300-h299	21.606	21.166	19.162	21.963	TO	TO	–
p300-h300	20.824	20.629	19.477	20.631	38.723	TO	–
p400-h399	50.054	50.202	46.937	50.482	TO	TO	–
p400-h400	51.144	50.439	47.718	50.255	94.253	TO	–

Table 2 Runtime in seconds for pigeon-hole problem with colors

	ENUM	SAT	ORD	FULL	BASE	SBASS	CLASP ^x
c1-p12-h11	1.939	0.007	0.007	0.006	692.704	0.284	0.015
c1-p52-h52	0.147	0.100	0.101	0.100	0.145	63.455	1.083
c2-p12-h12	6.648	0.009	0.010	0.010	TO	0.092	TO
c2-p52-h53	0.559	0.298	0.421	0.293	TO	93.870	TO
c3-p12-h13	4.457	0.014	0.052	0.019	TO	0.228	TO
c3-p52-h54	4.974	0.545	1.072	0.628	1.674	449.324	TO
c4-p12-h14	4.195	0.022	3.026	0.075	TO	0.479	TO
c4-p52-h55	3.131	0.959	2.033	3.849	TO	550.065	TO
c5-p12-h15	5.673	0.035	138.861	0.251	TO	1.020	TO
c5-p52-h56	18.336	1.561	581.164	11.704	5.930	TO	–

We note that reducing the number of examples for an ILP task generated by some of our settings has a limited impact on ILASP, as its latest versions implement mechanisms to scale with respect to the number of examples (Law et al., 2016; Law, 2021). However, for instances with many answer sets, the *fullSBCs* approach can be helpful because equivalent answer sets need not be exhaustively computed by CLINGO.

6 Solving experiments

To evaluate our approach and the implementation design, we applied it to a series of combinatorial search problems. For each considered problem, we compared the running time of the original encoding, the version extended with our learned constraints, and the *instance-specific* approach of SBASS. The learned constraints depend on the instances used in *S* and *Gen* as well as how we apply the iterative learning approach. In the following, we report results for the constraints with good performance learned applying the definitions of Sect. 4.3.⁵ We ran our tests on an Intel® i7-3930K machine under Linux (Debian GNU/Linux 10), where each run was limited to 900 seconds and 20 GB of memory.

⁵ Detailed settings are provided at Tazariol et al. (2021).

Table 3 Runtime in seconds for pigeon-hole problem with colors and owners

	ENUM	SAT	ORD	FULL	BASE	SBASS	CLASP ^π
o1-c3-p12-h13	1.505	0.016	0.051	0.046	TO	0.204	TO
o1-c3-p52-h54	1.749	0.554	1.002	0.906	1.709	312.328	TO
o2-c3-p12-h13	1.081	0.017	0.070	0.045	TO	0.337	TO
o2-c3-p52-h54	1.017	0.705	1.561	1.231	TO	542.377	TO
o3-c1-p12-h13	1.511	0.016	0.053	0.035	TO	0.21	TO
o3-c1-p52-h54	1.726	0.551	0.996	0.903	1.698	313.062	TO
o4-c4-p12-h14	0.894	0.031	3.162	1.114	TO	1.102	TO
o4-c4-p52-h55	4.123	1.819	3.486	3.546	6.653	TO	–
o5-c5-p12-h15	0.684	0.061	97.453	21.272	TO	2.093	TO
o5-c5-p52-h56	13.984	3.045	573.523	200.467	11.622	TO	–

Table 4 Runtime in seconds for house-configuration problem

	ENUM	SAT	ORD	FULL	BASE	SBASS	CLASP ^π
p2-c6-t13	0.337	0.025	0.097	0.329	219.753	0.095	12.951
p2-c80-t160	4.928	5.001	5.166	5.024	6.583	TO	–
p3-c6-t13	0.342	0.031	0.093	0.424	254.065	0.242	73.041
p3-c80-t160	13.682	14.315	15.174	14.110	20.724	TO	–
p4-c6-t13	0.420	0.035	0.102	0.349	221.784	0.453	105.145
p4-c80-t160	27.437	28.866	28.145	27.299	40.121	TO	–
p5-c6-t13	0.418	0.042	0.101	0.397	236.961	0.890	405.461
p5-c80-t160	48.057	49.382	48.598	49.645	68.167	TO	–
p4-c7-t15	13.263	0.335	1.524	14.229	TO	0.729	TO
p15-c15-t30	5.655	6.777	2.700	2.525	4.155	TO	–

In Tables 1, 2, 3 and 4, the satisfiable instances are shown in even rows, while the odd rows contain unsatisfiable instances. The column BASE refers to CLINGO (v5.5.0) run on the original encoding, while ENUM, SAT, ORD, and FULL report results for the original encoding augmented with first-order constraints learned in the *enum*, *sat*, and *enum* with alternative atom ordering or *fullSBCs* setting, respectively. The time required by SBASS to compute ground SBCs is given in the corresponding column, and CLASP^π provides the solving time obtained with these ground SBCs. Runs that did not finish within the time limit of 900 seconds are indicated by TO entries.

We first tested the pigeon-hole problem, working without any division and iterative analysis of the language bias: the four learning settings led to similar performance constraints, although the ones obtained with the alternative ordering were shorter, as mentioned in Sect. 5.2. The running time comparison in Table 1 shows that all the settings of our approach bring about a similar speedup for solving satisfiable as well as unsatisfiable instances. In fact, the vast problem symmetries are cut by the learned first-order constraints. This is particularly important in case of unsatisfiability, where runs on the original encoding without additional constraints do not finish within the time limit. While SBASS also manages to handle the two smallest instances, the computation of permutation group

generators becomes too expensive when the instance size grows, in which case we cannot run CLASP^x with ground SBCs from SBASS .

Next, we tested the pigeon-hole problem adding color and owner assignments. For the pigeon-hole problem with color assignments, we divided the language bias into two parts: the first limiting to predicates whose atoms exclusively include variables of the types `pigeon` and `hole`, while the second part allows variables to be of the type `color` too. Likewise, the problem version with owners and colors required a third language bias extension to variables of the type `owner`. For both extensions of the pigeon-hole problem, the first-order constraints learned in *sat* turned out to be stronger than in the other settings. Nevertheless, all kinds of constraints helped to improve the search for solutions. Tables 2 and 3 show similar results: the constraints learned with the *sat* setting lead to the fastest running times for both satisfiable and unsatisfiable instances. The constraints learned with *enum* using the alternative ordering are shorter and easier to read than the other settings, but slightly less efficient since they break only a subset of all symmetries. In Table 2, the time took for identifying satisfiable and unsatisfiable instances is lower if we use the constraints learned with *fullSBCs* than those learned with *enum*; on the other hand, in Table 3, we observe the opposite behavior, especially for the last instances: for the pigeon-hole problem with colors and owners, we could have learned the same constraints in both settings because to obtain the constraints with *enum*, we used instances that identify full SBCs. However, we tested a different set of rules for *fullSBCs* since they were stricter than *enum*, concerning the pigeons' placement symmetries. Indeed, the first unsatisfiable instance with one color and owner was solved earlier by the constraint of *fullSBCs*. Lastly, for small unsatisfiable instances, the ground SBCs from SBASS lead to better performance than the constraints learned with the *enum* setting. However, as soon as the color (or owner) dimension grows, the runs of CLASP^x reach the timeout. This behavior is due to the redundancy of the ground SBCs, which slow down the search instead of facilitating it. For some of the satisfiable instances, finding a solution with the constraints learned in *enum* takes longer than with the original encoding alone. Nevertheless, the latter also has timeouts that do not occur with our learned first-order constraints.

To conclude, we applied the different settings of our approach to the house-configuration problem (Friedrich et al. 2011), which consists of assigning t things of p persons to c cabinets, where each cabinet has a capacity limit of two things that must belong to the same owner. Similarly to the pigeon-hole problem with color, we divided the language bias into two parts: the first limiting to predicates whose atoms exclusively include variables of the types `cabinet` and `thing`, while the second part allows variables to be of the type `person` too. The running times in Table 4 exhibit the same trend as observed on the previous problems that our first-order constraints help the search, especially those learned with the *sat* setting. For this problem, the constraints learned with the *enum* setting, the alternative ordering, and exploiting the full SBCs show similar performances. In some cases, the original encoding is quicker to solve satisfiable instances, although it takes considerably longer for unsatisfiable ones. On the other hand, SBASS brings a moderate speedup for unsatisfiable instances, but its performance suffers a lot when the problem size grows.

7 Conclusions

This paper introduces methods to lift the SBCs of combinatorial problem encodings in ASP for a target distribution of instances. Our framework addresses the limitations of common *instance-specific* approaches, like SBASS, since: (i) the knowledge is transferable, as learned constraints preserve the satisfiability for the considered instance distribution; (ii) the first-order constraints are easier to interpret than ground SBCs; (iii) the SBCs are computed offline, allowing for addressing large-scale instances, as shown in our experiments; and (iv) the learned constraints are non-redundant, avoiding performance degradation due to an excessive ground representation size. In the current implementation of our approach, ILASP learns shortest constraints that cover as many examples as possible, while there is no distinction regarding the solving performance of candidate hypotheses. Despite this, our experiments showed that the learned constraints significantly improve the solving performance on the analyzed problems. Moreover, the two example generation methods suggested in this work allowed for ILP tasks with (i) fewer positive examples and (ii) shorter learned constraints, in comparison to the two methods of our previous paper (Tarzariol et al., 2021). These results are due to the full symmetry breaking when enumerating all answer sets with the *fullSBCs* approach or an alternative atom ordering for the lex-leader scheme, respectively.

Nevertheless, there are still some limitations in the usability of our framework, which partially go back to the components used in our current implementation, i.e., SBASS, CLINGO, and ILASP. The SBASS tool does not support ASP programs with weak constraints (Calimeri et al., 2019), whose implementation is out of the scope of this work. However, extensions of instance-specific symmetry detection and model-oriented symmetry breaking to optimization problems are undoubtedly worthwhile. Optimization involves solving unsatisfiable subproblem(s) attempting (and failing) to improve an optimal answer set, where symmetry breaking is particularly crucial for performance. Concerning CLINGO, if a given encoding P leads to large ground instantiations, the addition of learned constraints does not reduce the size. Therefore, it would be desirable to directly incorporate the information about redundant answer sets into a modified encoding. For instance, for the pigeon-hole problem, this might prevent our method from even generating ground atoms representing the placement of a pigeon into some hole with a greater label. Lastly, ILASP does currently not scale well with respect to the size of the hypothesis space spanned by the language bias, which is a well-known issue tackled by next-generation ILP systems under development (Law et al., 2020, 2021).

At present, the successful application of our framework relies on the following characteristics of a combinatorial problem: (i) we can easily provide simple instances (i.e., the total number of solutions can be managed by our implementation) that entail the symmetries of the whole instance distribution; (ii) the object domains can be expressed in terms of unary predicates that hold for a range of consecutive integers; and (iii) the auxiliary predicate definitions suggested for *ABK* in Sect. 4.3 enable the learning of constraints that improve the solving performance. In particular, if it gets difficult to compute solutions for an instance in S to analyze, the formulation of an ILP task to learn constraints can become prohibitive.

In the future, we aim to investigate whether the learning of SBCs can be readily applied or further adapted to advanced industrial configuration problems, such as the *Partner Units Problem* (Dodaro et al., 2016), as well as complex combinatorial problems with specific instance distributions, like the labeling of *Graceful Graphs* (Petrie & Smith, 2003). For

such application scenarios, the language bias may be enriched, possibly extending the background knowledge with additional predicates characterizing the structure of instances. Moreover, for problem instances that yield a vast number of solutions, we can take advantage of the incremental implementation of the *fullSBCs* approach to limit the number of answer sets to consider as examples for an ILP task. Lastly, we intend to develop automatic mechanisms to select suitable instances for *S* and *Gen* from instance collections, support lifelong learning, and further optimize the grounding and solving efficiency of learned constraints.

Acknowledgements 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 helpful and constructive comments.

Funding Open access funding provided by University of Klagenfurt.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Calimeri, F., Faber, W., Gebser, M., Ianni, G., Kaminski, R., Krennwallner, T., Leone, N., Maratea, M., Ricca, F., & Schaub, T. (2019). ASP-Core-2 input language format. *Theory and Practice of Logic Programming*, 20(2), 294–309.
- Codenotti, P., Katebi, H., Sakallah, K., & Markov, I. (2013). Conflict analysis and branching heuristics in the search for graph automorphisms. In 25th IEEE International Conference on Tools with Artificial Intelligence, pp. 907–914. IEEE Computer Society.
- Cohen, D., Jeavons, P., Jefferson, C., Petrie, K., & Smith, B. (2006). Symmetry definitions for constraint satisfaction problems. *Constraints*, 11(2–3), 115–137.
- Cropper, A., Dumančić, S., & Muggleton, S. (2020). Turning 30: New ideas in inductive logic programming. In 29th International Joint Conference on Artificial Intelligence, pp. 4833–4839. ijcai.org
- Cropper, A., & Dumančić, S. (2020). Inductive logic programming at 30: A new introduction. <https://arxiv.org/abs/2008.07912>
- Cropper, A., & Muggleton, S. (2016). Metagol. <https://github.com/metagol/metagol>
- Darga, P., Katebi, H., Liffiton, M., Markov, I., & Sakallah, K. (2004). Saucy <http://vlsicad.eecs.umich.edu/BK/SAUCY/>
- Devriendt, J., Bogaerts, B., Bruynooghe, M., & 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., & 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., & Walsh, T. (2011). Symmetry-breaking answer set solving. *AI Communications*, 24(2), 177–194.
- Erdem, E., Gelfond, M., & Leone, N. (2016). Applications of ASP. *AI Magazine*, 37(3), 53–68.
- Falkner, A., Friedrich, G., Schekotihin, K., Taupe, R., & 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., & Schreiner, H. (2011). (Re)configuration using answer set programming. In: IJCAI 2011 Workshop on Configuration, pp. 17–24. CEUR-WS.org.
- Gebser, M., Kaminski, R., Kaufmann, B., & Schaub, T. (2012). *Answer set solving in practice*. Morgan and Claypool Publishers.

- Gelfond, M., & Lifschitz, V. (1991). Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9, 365–385.
- Law, M. (2021). Conflict-driven inductive logic programming. <https://arxiv.org/abs/2101.00058>
- Law, M., Russo, A., Bertino, E., Broda, K., & Lobo, J. (2020). FastLAS: Scalable inductive logic programming incorporating domain-specific optimisation criteria. In: 34th National Conference on Artificial Intelligence , pp. 2877–2885. AAAI Press.
- Law, M., Russo, A., & Broda, K. (2014). Inductive learning of answer set programs. In 14th European Conference on Logics in Artificial Intelligence , pp. 311–325. Springer.
- Law, M., Russo, A., & Broda, K. (2016). Iterative learning of answer set programs from context dependent examples. *Theory and Practice of Logic Programming*, 16(5–6), 834–848.
- Law, M., Russo, A., & Broda, K. (2018). Inductive learning of answer set programs from noisy examples. *Advances in Cognitive Systems*, 7, 57–76.
- Law, M., Russo, A., & Broda, K. (2021). Iasp. <http://www.ilasp.com>
- Law, M., Russo, A., Broda, K., & Bertino, E. (2021). Scalable non-observational predicate learning in ASP. In 30th International Joint Conference on Artificial Intelligence , pp. 1936–1943. ijcai.org.
- Lifschitz, V. (2019). *Answer set programming*. Springer.
- Margot, F. (2010). Symmetry in integer linear programming. In 50 Years of Integer Programming 1958–2008, pp. 647–686. Springer-Verlag.
- Mears, C., García de la Banda, M., Wallace, M., & Demoen, B. (2008). A novel approach for detecting symmetries in CSP models. In: 5th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems , pp. 158–172. Springer.
- Muggleton, S. (1995). Inverse entailment and Progol. *New Generation Computing*, 13(3–4), 245–286.
- Petrie, K., & Smith, B. (2003). Symmetry breaking in graceful graphs. In 9th International Conference on Principles and Practice of Constraint Programming , pp. 930–934. Springer.
- Puget, J. (2005). Automatic detection of variable and value symmetries. In 11th International Conference on Principles and Practice of Constraint Programming , pp. 475–489. Springer .
- Sakallah, K. (2009). Symmetry and satisfiability. In: Handbook of satisfiability, pp. 289–338. IOS Press.
- Srinivasan, A. (2004). The Aleph manual. <https://www.cs.ox.ac.uk/activities/programinduction/Aleph/>.
- Tarzariol, A., Gebser, M., & Schekotihin, K. (2021). ILP symmetry breaking. https://github.com/prosysscience/Symmetry_Breaking_with_ILP/tree/extended
- Tarzariol, A., Gebser, M., & Schekotihin, K. (2021). Lifting symmetry breaking constraints with inductive logic programming. In 30th International Joint Conference on Artificial Intelligence , pp. 2062–2068. ijcai.org.
- Walsh, T. (2012). Symmetry breaking constraints: Recent results. In 26th National Conference on Artificial Intelligence , pp. 2192–2198. AAAI Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.