

Conflict Generalisation in ASP: Learning Correct and Effective Non-Ground Constraints

RICHARD TAUPE^{1,2}, ANTONIUS WEINZIERL³, and GERHARD FRIEDRICH²

¹ *Siemens AG Österreich,*

(*e-mail: richard.taupe@siemens.com*)

² *Alpen-Adria-Universität, Klagenfurt, Austria,*

(*e-mail: gerhard.friedrich@aau.at*)

³ *TU Wien (Vienna University of Technology), Austria,*

(*e-mail: antonius.weinzierl@kr.tuwien.ac.at*)

Abstract

Generalising and re-using knowledge learned while solving one problem instance has been neglected by state-of-the-art answer set solvers. We suggest a new approach that generalises learned nogoods for re-use to speed-up the solving of future problem instances. Our solution combines well-known ASP solving techniques with deductive logic-based machine learning. Solving performance can be improved by adding learned non-ground constraints to the original program. We demonstrate the effects of our method by means of realistic examples, showing that our approach requires low computational cost to learn constraints that yield significant performance benefits in our test cases. These benefits can be seen with ground-and-solve systems as well as lazy-grounding systems. However, ground-and-solve systems suffer from additional grounding overheads, induced by the additional constraints in some cases. By means of conflict minimization, non-minimal learned constraints can be reduced. This can result in significant reductions of grounding and solving efforts, as our experiments show.

KEYWORDS: Answer Set Programming, Deductive Learning, Non-Ground Nogood Learning

1 Introduction

Conflict-Driven Nogood Learning (CDNL) is a major success factor for high-performance state-of-the-art ASP systems (Gebser et al. 2012). When a conflict occurs, new propositional nogoods are learned that prevent the same conflict from re-occurring, which improves search performance.

Previous work has failed to address the question whether learned nogoods can be generalised and re-used to speed up the solving of different instances of the same problem. This paper aims to fill this gap. We present an extension of CDNL that learns non-ground constraints. The idea is that whole parts of the search tree can be pruned when these learned constraints are added to the original program. We presume the common distinction between an unvarying problem encoding and separate inputs consisting only of facts. A problem instance is specified by the problem encoding and a set of input facts. We aim to derive non-ground constraints from the problem encoding that are valid for all possible inputs and which can be employed to speed up solving new instances. In practice, instances are often not random but share some similarities that are reflected by nogoods learned during solving. This led us to assume the existence of non-ground nogoods capable of significantly speeding up solving of practical problem instances.

CDNL and Explanation-Based Learning (EBL) (Mitchell et al. 1986; DeJong and Mooney 1986; Hirsh 1987) are our starting point. EBL is a well-known logic-based machine learning technique which learns first-order rules that are entailed by the background knowledge (in our case, the problem encoding). We combine CDNL with EBL to learn non-ground nogoods while solving prior problem instances. Since the number of generalised nogoods can be overwhelming, choosing those that will actually pay off is particularly challenging. Our basic idea is to generalise those non-ground conflicts that occur most often, i.e., we generalise propositional nogoods learned from frequently violated nogoods. The underlying assumption is that nogoods learned from frequent conflicts will also be able to prevent many conflicts.

A realistic hardware configuration example and a graph colouring problem are used for demonstration and experimentation purposes. Results show that both lazy-grounding and ground-and-solve systems benefit from our approach. By adding learned constraints to the problem encodings, up to 64% more configuration instances can be solved and graph colouring instances can be solved much faster. Learning itself requires low computational resources.

We believe that we have developed an innovative tool supporting the design of efficient answer-set programs. This view is supported by encouraging experimental results. Our main contributions can be summarized as follows:

- We present a motivating practical example and describe how some redundant constraints can be derived from a given encoding. We also sketch how non-minimal nogoods can be reduced.
- Next, we introduce a novel approach combining CDNL and EBL to automatically learn correct non-ground nogoods while solving an answer-set program. A pseudo-code algorithm is given.
- We suggest techniques to make learned nogoods use only predicates from the input program, the means to choose from the wide range of possible generalisations, and methods to quickly conduct learning under limited resources.
- Finally, we experimentally compare the effects of various learned constraints on two domains. Results demonstrate the practical benefits of our approach.

2 Preliminaries

2.1 Answer Set Programming

An answer-set program P is a finite set of non-ground rules of the form

$$h \leftarrow b_1, \dots, b_m, \text{ not } b_{m+1}, \dots, \text{ not } b_n. \quad \langle I \rangle$$

where h and b_1, \dots, b_n are atoms. A classical atom is an expression $p(t_1, \dots, t_n)$ where p is an n -ary predicate and t_1, \dots, t_n are terms. A term is a variable, a constant, or a complex term involving arithmetics. A NaF-literal is either an atom a or its negation $\text{not } a$. Negation as failure (NaF) refers to the absence of information, i.e., an atom is assumed to be false as long as it is not proven to be true. If l is a literal, then \bar{l} denotes its complement (i.e., $\overline{\text{not } a} = a$ and $\bar{a} = \text{not } a$).

Given a rule r , $H(r) = \{h\}$ is called the *head* of r , and $B(r) = \{b_1, \dots, b_m, \text{ not } b_{m+1}, \dots, \text{ not } b_n\}$ is called the *body* of r . A rule r where the head is absent ($H(r) = \{\}$), e.g., $\leftarrow b$, is called *constraint*. A rule r where the body is absent ($B(r) = \{\}$), e.g., $h \leftarrow$, is called *fact*. Given an answer-set program P , the *universe* of P , denoted by U_P , is the set of constants appearing in P . By \mathcal{A} we denote the set of classical atoms constructible from predicates of P with constants in U_P . The set of all literals is denoted by $\mathcal{L} = \mathcal{A} \cup \{\text{not } a \mid a \in \mathcal{A}\}$. The facts function maps a program P to all ground atoms defined by facts in P , i.e., $\text{facts}(P) = \{H(r) \mid r \in P \text{ and } B(r) = \emptyset\}$.

The vars function maps any structure containing variables to the set of variables it contains, e.g., for a rule $r_1 : a(X) \leftarrow b(X, Y)$, $\text{vars}(r_1) = \{X, Y\}$. The set $\mathcal{V} \supset \text{vars}(P)$ includes all variables from a program and also variables that can additionally be used by a solver, it is usually infinite. A literal l or rule r is ground if $\text{vars}(l) = \emptyset$ or $\text{vars}(r) = \emptyset$, respectively. A program P is ground if all its rules $r \in P$ are. A substitution $\sigma : \mathcal{V} \rightarrow \mathcal{V} \cup U_P$ is a mapping from variables to variables or constants. A substitution is called *grounding* for a set of variables $V \subseteq \mathcal{V}$ if for every $v \in V$ there exists a constant $c \in U_P$ s.t. $\zeta(v) = c$. The function $v_A : \mathcal{A} \rightarrow \{\mathbf{T}, \mathbf{M}, \mathbf{F}, \mathbf{U}\}$ for a (partial) assignment A maps an atom to the truth value the atom is currently assigned in the given assignment (\mathbf{M} stands for must-be-true and \mathbf{U} for currently unassigned).

A nogood is a set of Boolean signed literals $\{\mathbf{T} a_1, \dots, \mathbf{T} a_n, \dots, \mathbf{F} a_{n+1}, \dots, \mathbf{F} a_m\}$ which means that an answer set may not satisfy all the literals in the nogood. A literal in a nogood differs from a NaF-literal in a rule semantically: In a nogood, negation as failure has no role. A nogood is violated by an assignment A if $v_A(a_i) \in \{\mathbf{T}, \mathbf{M}\}$ for all $i \in \{1, \dots, n\}$ and $v_A(a_j) = \mathbf{F}$ for all $j \in \{n + 1, \dots, m\}$. Sometimes, we will represent nogoods not by sets but by fixed-order multisets, because we need the possibility of duplicates and of fixed positions to be able to map literals between ground and non-ground nogoods. When this is the case, $\omega[i]$ denotes the literal at position i in nogood ω .

For ASP semantics, we refer to Faber et al. (2011).

Aggregates and choice rules are common extensions to the input language of ASP (Calimeri et al. 2020). Aggregate atoms are used to express arithmetic constraints on sets of atoms (e.g., cardinality constraints). A choice rule expresses that a subset of a set of atoms will be included in an answer set. This does not clash with the subset-minimality of answer sets because choice rules are translated to normal rules involving additional atoms whose predicates we call *choice predicates*. For simplicity, we only consider choice rules without bounds (which are the only ones currently supported by ALPHA). A choice rule is of the form $\langle I \rangle$, but the head h is not a classical atom but a choice atom of the form $\{C_1; \dots; C_n\}$, where for $n > 0$ and $0 \leq i \leq n$ each C_i is a *choice element*. A choice element has the form $a : l_1, \dots, l_k$, where a is a classical atom and l_1, \dots, l_k are literals for $k \geq 0$. If the body of a choice rule is satisfied, it means that for each C_i where the literals l_1, \dots, l_k are satisfied, the atom a may or may not be true in an answer set. A choice rule r can be translated to $2n$ rules, two for each choice element $a : l_1, \dots, l_k$:

$$a \leftarrow B(r), l_1, \dots, l_k, \text{not } \hat{a}. \quad \hat{a} \leftarrow B(r), l_1, \dots, l_k, \text{not } a.$$

For both aggregates and choice rules, see Calimeri et al. (2020) for details.

The goal of our method is to generate non-ground nogoods which can be re-used for solving further problem instances. More formally, let P^E be an ASP problem encoding. $InPred$ and $OutPred$ are sets of predicate symbols with defined arities. The input for P^E is specified by a set of ground facts InA containing only predicate symbols from $InPred$. The output of a problem instance $P^E \cup InA$ is the set of atoms of a stable model whose predicate symbols are in $OutPred$.

A conflict for an ASP program P is a constraint C s.t. $P \models_S C$ where \models_S refers to the skeptical ASP semantic, i.e., C must be true in all stable models of P .

Definition 2.1

A generalized conflict of $(P^E, InPred)$ is a constraint GC s.t. $P^E \cup InA \models_S GC$ for all finite sets of ground facts InA whose predicate symbols are in $InPred$.

Let \models_F be the implication relation based on standard FOL semantic. Since every stable model of an ASP program P is also a model for P if P is interpreted as an FOL program, it follows that if $P \models_F C$ then $P \models_S C$. It is well known that every stable-model of P is also a model of

$P \cup \text{comp}(P)$, i.e., the extension of P by the axioms of the Clark completion. Note that for problem instances $P^E \cup \text{InA}$ the Clark completion changes depending on InA . Since we are interested in conflicts which are independent of the input, we consider only the completion of the head literals which appear in P^E but not in InPred . We denote this restricted Clark completion for a program P and input predicates InPred as $\text{rComp}(P, \text{InPred})$. It holds that $\text{rComp}(P, \text{InPred}) \subseteq \text{comp}(P)$ and $P^E \cup \text{rComp}(P^E \cup \text{InA}, \text{InPred})$ is the same for all problem instances $P^E \cup \text{InA}$.

Corollary 2.1

Let a problem be defined by (P^E, InPred) and C is a conflict. If $P^E \cup \text{rComp}(P, \text{InPred}) \models_F C$ then C is a generalized conflict of (P^E, InPred) .

Proof sketch: Since $\text{rComp}(P, \text{InPred})$ depends only on P , it holds that if $P^E \cup \text{rComp}(P, \text{InPred}) \models_F C$ then $P^E \cup \text{rComp}(P, \text{InPred}) \cup \text{InA} \models_F C$ where InA is a finite set of ground facts whose predicate symbols are in InPred . Because every stable model of $P^E \cup \text{InA}$ is also a model for $P^E \cup \text{rComp}(P, \text{InPred}) \cup \text{InA}$ under FOL semantic and C is true in these models it follows that $P^E \cup \text{InA} \models_S C$. \square

Consequently, adding a generalized conflict GC to P^E does not change the set of outputs of (P^E, InPred) for any input.

2.2 Conflict-Driven Nogood Learning (CDNL)

Conflict-Driven Clause Learning (CDCL) (Silva et al. 2009; Lintao Zhang et al. 2001) is a SAT solving technique that extends DPLL by conflict analysis and enables the solver to learn new clauses and to do non-chronological backtracking (“backjumping”). While the original definitions of CDCL are based on clauses, i.e., disjunctions of literals, we use the notion of nogoods and speak of Conflict-Driven Nogood Learning (CDNL) (Gebser et al. 2012). The two variants are equivalent since the conjunction of classically negated nogoods can be directly transformed to an equivalent conjunction of clauses by applying De Morgan’s laws.

CDCL- and CDNL-based solvers usually characterize each atom $a \in \mathcal{A}$ by the following properties: its *truth value*, its *antecedent clause*, and its *decision level*, denoted respectively by $v(a) \in \{\mathbf{T}, \mathbf{F}, \mathbf{M}, \mathbf{U}\}$, $\alpha(a) \in 2^{\mathcal{L}} \cup \{\text{NIL}\}$, and $\delta(a) \in \{-1, 0, 1, \dots, |\mathcal{A}|\}$. The truth value of an atom is the value currently assigned to it. By the *antecedent clause* of an atom we mean the cause of its current truth value. The antecedent clause of an atom whose current truth value has been implied is the nogood that was unit at the time the implication happened and thus forced the atom to assume its current value. The value NIL is used instead if the current truth value of an atom results from a heuristic decision. The decision level of an atom denotes the depth of the decision tree at which it has been assigned, and -1 if the atom is unassigned. The decision level of a literal is the same as the literal’s atom’s decision level ($\delta(\text{not } a) = \delta(a)$), and the antecedent of a literal is the one of its atom ($\alpha(\text{not } a) = \alpha(a)$). An assignment’s decision level is the decision level of the atom most recently assigned (Silva et al. 2009; Lintao Zhang et al. 2001).

When a nogood is violated during solving (i.e., all literals in the nogood are satisfied), the conflict is analysed and a new nogood may be learned. Learning starts with the violated nogood and resolves one literal assigned at the current (most recent) decision level with its antecedent. This is repeated with the nogood resulting from resolution (the “resolvent”) as long as it is possible. Resolution stops when the only literal in the resolvent that has been assigned at the current decision level is a decision literal, i.e., its antecedent is NIL. The resulting nogood is then “learned”, i.e., added to the set of known nogoods, and the solver executes a backjump (Silva et al. 2009).

First UIP clause learning is a modification of this algorithm that already stops resolution when only one literal from the current decision level remains in the nogood, even if its antecedent is not NIL. This is correct because this literal is a unique implication point (UIP), a node in the implication graph that lies on every path from the last decision to the conflict. Each UIP could be used to learn a new nogood, but modern SAT solvers stop already at the first UIP, i.e., the one nearest to the conflict (Silva et al. 2009).

2.3 Explanation-Based Learning

Explanation-Based Learning (EBL), a.k.a. Explanation-Based Generalization (EBG), is a logic-based learning technique that, in contrast to Inductive Logic Programming (ILP), learns only general rules that are entailed by background knowledge alone (and not by background knowledge together with some new hypotheses). Therefore, in EBL nothing factually new is learned from the example. In EBL, two proof trees are created simultaneously (one for the concrete example and one with variables instead of constants). Then, new rules can be obtained from the non-ground proof tree. The *operationality criterion* is a restriction which predicates can be used to express the learned rules. The purpose of operationality is to use only predicates that are easy to solve. A trade-off between operationality and generality is usually an issue in EBL. By choosing adequate general rules, EBL makes a knowledge base more efficient for the kind of problems that one would reasonably anticipate (Mitchell et al. 1986; DeJong and Mooney 1986; Hirsh 1987; Mitchell 1997; Russell and Norvig 2010).

Van Harmelen and Bundy (1988) show that EBL/EBG and Partial Evaluation (PE) are equivalent to each other in the context of logic programming. PE is a program optimisation method that reformulates an input program in an equivalent but more efficient way. Weinzierl (2013) employs PE (and terms it “unfolding” of rules) in lazy-grounding answer set solving to learn new rules during solving. Learning is triggered by conflicts to prevent future occurrences of similar conflicts. Learned rules are constructed from constraints that are almost violated and in which atoms have been replaced by their definitions: The single literal in the constraint’s body that is not yet satisfied must not be satisfied. Learned rules are created and used online during solving.

3 A Motivating Example

The House Configuration Problem (HCP) serves as a motivating example. It is an abstraction of real-world hardware configuration problems and defined as follows:¹ Given an association of things to the persons they belong to, a domain of cabinet IDs, and a domain of room IDs, the task is to assign things to cabinets and cabinets to rooms such that there are at most five things in each cabinet and at most four cabinets in each room, and each room only contains things that belong to the same person. The problem encoding stays the same for all problem instances. Input facts of predicates `personTOthing/2`, `cabinetDomain/1`, and `roomDomain/1` vary from instance to instance. Each answer set specifies one valid configuration.

We are guided by the original definition of HCP by Friedrich et al. (2011) and limit ourselves to the configuration aspects of the problem, neglecting reconfiguration for presentation purposes. The complete encoding is given in Listing 1.²

¹ See <https://sites.google.com/view/aspcomp2019/problem-domains>, Friedrich et al. (2011), and Ryabokon (2015) for more complete descriptions of the problem.

² Compared to Ryabokon (2015), variables have been renamed for this paper and some simplifications have been made.

```

1 { cabinet(C) } :- cabinetDomain(C).
2 { room(R) } :- roomDomain(R).
3 room(R1) :- roomDomain(R1), roomDomain(R2), room(R2), R1 < R2.
4 cabinet(C1) :- cabinetDomain(C1), cabinetDomain(C2), cabinet(C2), C1 < C2.
5
6 { cabinetTOthing(C,T) } :- cabinetDomain(C), thing(T).
7 thingHasCabinet(T) :- cabinetTOthing(C,T).
8 :- thing(T), not thingHasCabinet(T).
9 :- thing(T), cabinetDomain(C1), cabinetTOthing(C1,T), cabinetDomain(C2),
   cabinetTOthing(C2,T), C1 < C2.
10 :- 6 <= #count { T : cabinetTOthing(C,T), thing(T) }, cabinet(C).
11
12 { roomTOcabinet(R,C) } :- roomDomain(R), cabinet(C).
13 cabinetHasRoom(C) :- roomTOcabinet(R,C).
14 :- cabinet(C), not cabinetHasRoom(C).
15 :- cabinet(C), roomDomain(R1), roomTOcabinet(R1,C), roomDomain(R2),
   roomTOcabinet(R2,C), R1 < R2.
16 :- 5 <= #count { C : roomTOcabinet(R,C), cabinetDomain(C) }, room(R).
17
18 personTOroom(P,R) :- personTOthing(P,T), cabinetTOthing(C,T), roomTOcabinet(R,C).
19 :- personTOroom(P1,R), personTOroom(P2,R), P1 < P2.
20
21 room(R) :- roomTOcabinet(R,C).
22 room(R) :- personTOroom(P,R).
23 cabinet(C) :- cabinetTOthing(C,T).
24 cabinet(C) :- roomTOcabinet(R,C).
25 thing(T) :- personTOthing(P,T).

```

Listing 1. An encoding for the House Configuration Problem

Ryabokon (2015) added the following redundant constraint to the problem:

$$\leftarrow \text{cabinetTOthing}(C, T1), \text{personTOthing}(P1, T1), \\ \text{cabinetTOthing}(C, T2), \text{personTOthing}(P2, T2), P1 < P2. \quad \langle 2 \rangle$$

However, this constraint can be learned automatically. Using ALPHA (Weinzierl 2017) to solve the smallest instance, the nogood that is violated most often and thus leads to the highest number of conflicts is $\{\mathbf{T} \text{ personTOroom}(P1, R), \mathbf{T} \text{ personTOroom}(P2, R), P1 < P2\}$ which corresponds to the constraint in Line 19 of Listing 1. This constraint forbids things of two different persons $P1$ and $P2$ from ending up in the same room R . Using our method, the following non-ground constraint can be learned at the first UIP of the implication graphs of all these conflicts:

$$\leftarrow \text{roomTOcabinet}(R, C), \text{cabinetTOthing}(C, T1), \text{personTOthing}(P1, T1), \\ \text{cabinetTOthing}(C, T2), \text{personTOthing}(P2, T2), P1 < P2. \quad \langle 3 \rangle$$

This constraint can also be easily generated by applying the rule in Line 18 to the constraint in Line 19 of Listing 1 and by factorizing different literals of predicate $\text{roomTOcabinet}/2$. The automatically learned constraint $\langle 3 \rangle$ is almost identical to the human-created one $\langle 2 \rangle$. However, it includes one additional literal $\text{roomTOcabinet}(R, C)$.

Learned non-ground constraints can be minimized by employing axioms of the Clark completion (Shepherdson 1984). The Clark completion of a logic program is computed by completing the head literals of the clauses of the program, including the completion of facts. Every stable model of a program P must also satisfy the Clark completion $\text{comp}(P)$. For example, the completion of $\text{cabinetHasRoom}(C)$ is $\exists R : \text{roomTOcabinet}(R, C) \leftarrow \text{cabinetHasRoom}(C)$. To minimize the learned nogood by resolution we substitute the existential quantifier by a Skolem function resulting in the rule $\text{roomTOcabinet}(r(C), C) \leftarrow \text{cabinetHasRoom}(C)$.

To minimize the conflict, we have to prove that constraint $\langle 2 \rangle$ is entailed. This is achieved by showing that the program in Listing 1 becomes unsatisfiable when the negation of constraint $\langle 2 \rangle$

is added to this program. Consequently, it is sufficient to show that the program becomes unsatisfiable when the facts $\text{cabinetTOthing}(c, t1)$, $\text{personTOthing}(p1, t1)$, $\text{cabinetTOthing}(c, t2)$, $\text{personTOthing}(p2, t2)$, $p1 < p2$ are added, where c , $t1$, $t2$, $p1$, and $p2$ are Skolem constants.

Every stable model of an ASP program is also a model of the program interpreted under the standard first-order logic (FOL) semantic. Consequently, if a program is unsatisfiable under FOL semantics, it is also unsatisfiable skeptical ASP semantics. From $\text{cabinetTOthing}(c, t1)$ and the rule in Line 23 we can deduce $\text{cabinet}(c)$. From the constraint in Line 14 of Listing 1, we can deduce $\text{cabinetHasRoom}(c)$ interpreting the sentences in Listing 1 as FOL clauses. Exploiting the rule provided by the Clark completion of $\text{cabinetHasRoom}(C)$ as shown above, we can deduce $\text{roomTOcabinet}(r(c), c)$. Finally, using all the deduced and given facts, constraint $\langle 3 \rangle$ is violated. Since this constraint is implied by the program and we have shown that it is violated if the negation of constraint $\langle 2 \rangle$ is added, constraint $\langle 2 \rangle$ is implied. As a result, we have shown that indeed the learned non-ground constraint can be reduced by dropping $\text{roomTOcabinet}(R, C)$.

4 Conflict Generalisation

We propose to combine CDNL with EBL to facilitate the learning of general (i.e., non-ground) constraints. This can be done online (during solving), or offline. In the latter case, learned constraints are computed and recorded. A selection can then be made, for example either automatically based on metrics or by a human, and useful constraints can be added to the original program to improve solving performance in the future. In this section, we describe how offline conflict generalisation can be implemented in an ASP solver and exemplify this by means of a prototypical implementation in the lazy-grounding system ALPHA (Weinzierl 2017). The constraints learned by this method can be used by any ASP system, not only by the system employed for the learning task, because they are stated in pure ASP-Core-2 (Calimeri et al. 2020), using only predicates from the original program.

Parallely conducting CDNL on the ground level and on the non-ground level when a conflict is encountered is the basic idea of combining CDNL with EBL to generalise learned nogoods. A key requirement for this is that for each ground nogood, a corresponding non-ground nogood is known that can be used for non-ground resolution. We will now describe how non-ground nogoods are maintained in ALPHA. This can be implemented similarly in any other ASP system.

In ALPHA, there are five types of nogoods:

- *static* nogoods that represent ground rules and are generated by the grounder;
- *support* nogoods that encode the situation that the body-representing atom of a rule must be true if the head of the rule is true;
- *learned* nogoods originating from CDNL;
- *justification* nogoods learned by justification analysis (Bogaerts and Weinzierl 2018);
- and *internal* nogoods containing solver-internal atoms.

Non-ground nogoods for *static* and *support* nogoods are produced by the grounder as described in the following paragraphs. Non-ground nogoods for *learned* nogoods are produced by the conflict generalisation procedure that is the main contribution of this paper. Non-ground nogoods are not maintained for *justification* nogoods (because these nogoods depend too heavily on the specific problem instance) and for *internal* nogoods (because these are irrelevant to CDNL).

We now describe how non-ground nogoods for static and support nogoods are produced by the grounder. Let r be a rule of the form $\langle I \rangle$ and σ be a substitution that is grounding for $\text{vars}(r)$. The

grounder produces a ground rule $r\sigma$ from r if no fact and no fixed-interpretation literal makes the body of $r\sigma$ false. Because of this, facts and fixed-interpretation literals are actually omitted from the ground rule produced by the grounder.³ For simplicity of presentation, however, we assume here that the grounder does not eliminate any atoms known to be true from generated nogoods. This ensures a one-to-one relationship between literals in ground nogoods and literals in non-ground nogoods. Additional effort is necessary to map from ground literals to non-ground literals if true literals are eliminated from ground nogoods, but this is purely a matter of implementation. Not eliminating any literals from non-ground nogoods, however, is crucial.

A body-representing atom $\beta(r, \sigma)$ is created for every ground rule $r\sigma$. Similarly, $\beta(r, \text{vars}(r))$ is a fictitious atom representing the body of a non-ground rule r . The latter atom contains a term that lists all the variables that occur in the rule r , so that they can affect unification when a body-representing literal is used for resolution. This will become crucial in Section 4.1.

The following static nogoods are produced by the grounder from a rule r and a grounding substitution σ (Leutgeb and Weinzierl 2017):

$$\{\mathbf{F} \beta(r, \sigma), \mathbf{T} b_1\sigma, \dots, \mathbf{T} b_m\sigma, \mathbf{F} b_{m+1}\sigma, \dots, \mathbf{F} b_n\sigma\} \quad \langle 4 \rangle$$

$$\{\mathbf{F} h\sigma, \mathbf{T} \beta(r, \sigma)\} \quad \langle 5 \rangle$$

$$\{\mathbf{T} \beta(r, \sigma), \mathbf{F} b_1\sigma\}, \dots, \{\mathbf{T} \beta(r, \sigma), \mathbf{F} b_m\sigma\} \quad \langle 6 \rangle$$

$$\{\mathbf{T} \beta(r, \sigma), \mathbf{T} b_{m+1}\sigma\}, \dots, \{\mathbf{T} \beta(r, \sigma), \mathbf{T} b_n\sigma\} \quad \langle 7 \rangle$$

$$\{\mathbf{T} h\sigma, \mathbf{F} \beta(r, \sigma)\} \quad \langle 8 \rangle$$

Nogood $\langle 8 \rangle$ is a so-called *support nogood* and is produced by ALPHA currently only when an atom occurs in the head of just a single rule. If r is a constraint, only one nogood is created, which consists of the whole body of $r\sigma$.

As a prerequisite for conflict generalisation, a solver must associate each ground nogood with a non-ground nogood. This non-ground nogood is obtained from a non-ground rule the same way as a ground nogood is obtained from a ground rule. This means that the non-ground nogoods are exactly the same ones as given above, except the substitution σ does not appear anywhere: The atom $\beta(r, \sigma)$ becomes $\beta(r, \text{vars}(r))$, and any other ground atom $a\sigma$ becomes just a .

4.1 Non-ground CDNL

In this section, we describe how we extend CDNL to learn non-ground nogoods. The non-ground nogoods learned that way can then just be used as constraints and be added to the original program. When learning only constraints and no other kinds of rules, \mathbf{F} in nogoods can just be replaced by negation as failure.

Our main conflict generalisation algorithm is shown in Algorithm 1. We represent nogoods by fixed-order multisets to be able to map between literals in ground and non-ground nogoods. The algorithm takes as input a violated ground nogood ω and the current assignment A . If the conflict occurred at decision level 0, UNSAT is returned in Line 2, otherwise two lists of learned nogoods are returned in Line 14. The first list contains ground nogoods and the second list contains non-ground nogoods. Each element of both lists corresponds to one UIP and both lists are ordered by distance from the conflict, i.e., the first element of both lists is a nogood learned at the first UIP and the last element of both lists is a nogood learned at the last UIP.

³ A more sophisticated grounder would even use the set of atoms derived by a stratified component of the program depending only on facts instead of $\text{facts}(P)$ to make generated rules more compact.

Algorithm 1 CDNL and conflict generalisation**Input:** ω : the violated antecedent, A : current assignment**Output:** a list of learned ground nogoods and a list of learned non-ground nogoods, or UNSAT

```

1: if  $\delta(A) = 0$  then
2:   return UNSAT
3: end if
4:  $LearnedNoGoods \leftarrow$  empty list
5:  $LearnedNonGroundNoGoods \leftarrow$  empty list
6:  $\Omega \leftarrow$  non-ground nogood associated with  $\omega$ 
7: while ( $l \leftarrow \text{FINDNEXTLITERALFORRESOLUTION}(\omega, A) \neq \text{NIL}$ ) do
8:    $\omega, \Omega \leftarrow \text{RESOLVE}(\omega, \Omega, l)$ 
9:   if  $\exists! l' \in \omega: \delta(\omega) = \delta(A)$  then
10:     append  $\omega$  to  $LearnedNoGoods$ 
11:     append  $\Omega$  to  $LearnedNonGroundNoGoods$ 
12:   end if
13: end while
14: return  $LearnedNoGoods, LearnedNonGroundNoGoods$ 

15: procedure  $\text{FINDNEXTLITERALFORRESOLUTION}(\omega, A)$ 
16:   return the literal  $l$  most recently assigned in  $A$  that appears in  $\omega$ 
        and for which it holds that  $\delta(l) = \delta(A)$  and  $\alpha(l) \neq \text{NIL}$ ,
        or NIL if no such literal exists
17: end procedure

18: procedure  $\text{RESOLVE}(\omega, \Omega, l)$ 
19:    $\omega' \leftarrow (\omega \setminus \{l\}) \cup (\alpha(l) \setminus \{\bar{l}\})$ 
20:    $L \leftarrow$  non-ground literal s.t.  $\omega[i] = l$  and  $\Omega[i] = L$  for some  $i$ 
21:    $\bar{L} \leftarrow$  non-ground literal s.t.  $\alpha(l)[j] = \bar{l}$  and  $\Omega[j] = \bar{L}$  for some  $j$ 
22:    $\Omega' \leftarrow$  the non-ground nogood associated with  $\alpha(l)$ , standardised apart from  $\Omega$ 
23:    $\sigma \leftarrow \text{unify}(L, \bar{L})$ 
24:    $\sigma \leftarrow \sigma \circ \text{UNIFYDUPLICATELITERALS}(\Omega\sigma, \Omega'\sigma, \omega, \alpha(l))$ 
25:    $\Omega' \leftarrow (\Omega \setminus \{L\})\sigma \cup (\Omega' \setminus \{\bar{L}\})\sigma$ 
26:   return  $\omega', \Omega'$ 
27: end procedure

28: procedure  $\text{UNIFYDUPLICATELITERALS}(\Omega, \Omega', \omega, \omega')$ 
29:    $\gamma \leftarrow \omega \cap \omega'$ 
30:    $\sigma \leftarrow$  empty unifier
31:   for  $l' \in \gamma$  do
32:      $L \leftarrow$  non-ground literal s.t.  $\omega[i] = l'$  and  $\Omega[i] = L$  for some  $i$ 
33:      $L' \leftarrow$  non-ground literal s.t.  $\omega'[j] = l'$  and  $\Omega'[j] = L'$  for some  $j$ 
34:      $\sigma \leftarrow \sigma \circ \text{unify}(L\sigma, L'\sigma)$ 
35:   end for
36:   return  $\sigma$ 
37: end procedure

```

As long as the current nogood contains a literal from the current decision level⁴ with non-NIL antecedent (Line 7), it is resolved with the antecedent of such a literal to produce a new current nogood ω (the resolvent, Line 8). If ω contains exactly one literal on the current decision level (Line 9), we have found a UIP and remember the current nogood (both ground and non-ground).

Resolution of ground nogoods is straightforward and well-known (Line 19): The resolvent is the union of two input nogoods (ω and the antecedent of l , $\alpha(l)$), minus l that occurs in the first input nogood and its complement \bar{l} that occurs in the second input nogood.

Following the resolution on the non-ground level is more complex, however, and several special cases have to be considered. First, the non-ground nogood corresponding to $\alpha(l)$ is identified and standardised apart from the current nogood (i.e., variables are renamed to avoid overlaps; Line 22).⁵ Then, two steps of unification are necessary: First, in Line 23, the two complementary resolution literals are unified s.t. variable occurrences are correctly updated in the resolvent. Then, an additional step is necessary only on the non-ground level: If the two ground input nogoods share some literals, duplicates are just removed during resolution, because a nogood is a *set* of literals. On the non-ground level, however, for each ground literal that occurs in both input nogoods, the corresponding non-ground literals must be unified before one of them can be removed (Line 24). This is a restricted form of factoring that is guided by ground resolution. Note that in the context of CDNL, it is not possible for two input nogoods to contain complementary literals apart from the resolution literal, because all antecedents must be unit to entail a literal. Finally, after applying the computed unifier to both non-ground nogoods, the resolution step is the same as on the ground level (Line 25).

The connection between our suggestion of non-ground CDNL and EBL becomes apparent when viewing the violated nogood as the training instance. The implication graph utilised by CDNL constitutes a proof that the nogood is violated by giving a derivation for each literal in the nogood. Each arc in the implication graph that originates from unit propagation can be seen as a definite Horn clause that derives the literal assigned by the propagation step. Operability in our setting is specified by demanding that nogoods are only learned at UIPs. The set of predicates allowed in learned nogoods is not restricted, because in ASP there are no predicates that are “easier to solve” than others. In lazy grounding, rules are even grounded only when (part of) their body is already satisfied (Taupe et al. 2019), which alleviates the overhead of additional rules. Currently, facts from the ASP program are not regarded by our approach because they are not included in nogoods generated by state-of-the-art answer-set solvers. This ensures that learned constraints depend only on the rules from the problem encoding.

Algorithm 1 is correct because it derives constraints that are implied when the input program P is interpreted as an FOL program, and because such constraints are also true in all stable models of P , as has been discussed in Section 2.1.

Replacing internal atoms in learned constraints. Learned constraints may contain literals of solver-internal predicates. Literals of two kinds of such predicates, namely body-representing and choice predicates, can easily be replaced by equivalent (sets of) literals of predicates from the problem encoding. Only then can resulting constraints be correctly added to the input program.

⁴ It would also be correct to choose a literal from a lower decision level for resolution, even though this is not done in CDNL. We will use such an extension of Algorithm 1 for parts of our experiments reported on in Section 5.

⁵ If no non-ground nogood is available for $\alpha(l)$, conflict generalisation can continue with another literal from the current decision level or just learn the current resolvent.

A body-representing atom such as $\beta(r, \text{vars}(r))$ represents the body of non-ground rule r . If a nogood contains a positive literal of such an atom, by definition this literal can just be replaced by the body of r . This is equivalent to resolving with nogood $\langle 4 \rangle$ above. If the head of r is the head of no other rule, $\beta(r, \text{vars}(r))$ can also be replaced by $H(r)$ by resolving with the support nogood $\langle 8 \rangle$. If a nogood contains a negative literal of a body-representing atom, it can be replaced by the negated head of r (which is equivalent to resolving with nogood $\langle 5 \rangle$ above) or by an arbitrary negated body literal of r (equivalent to resolving with any of the nogoods in $\langle 6 \rangle$ or $\langle 7 \rangle$).

For choice rules, internal literals of a different kind are created, as has been described in Section 2.1. Because of this, atoms \hat{a} may occur in learned nogoods. Due to the way choice rules are translated to normal rules, the literal $\mathbf{F} \hat{a}$ is equivalent to $\mathbf{T} a$ and can simply be replaced. Similarly, $\mathbf{T} \hat{a}$ is equivalent to $\mathbf{F} a$.

Choosing effective constraints. Since the number of learnt constraints might be overwhelming, strategies to choose effective ones are of vital importance. Our approach is simple but effective: We count how many conflicts could be avoided if the additional non-ground constraints were already included in the input program.

Each ground nogood that is violated during a run of CDNL belongs to a class of nogoods that share the same non-ground nogood. When running Algorithm 1 upon such a nogood violation, the lists of learned non-ground nogoods are associated with this class. When conflict generalisation terminates, the learned nogoods are printed together with the number of violations of their associated non-ground nogood, s.t. a (human) user can then select the most useful constraints. Since learned constraints can of course be equivalent to each other when variables are renamed, only one unique representation of each nogood is remembered. For example, the class of ground nogoods violated most often while ALPHA solves the HCP (Section 3) is identified by the non-ground nogood $\{\mathbf{T} \text{personTOroom}(P1, R), \mathbf{T} \text{personTOroom}(P2, R), P1 < P2\}$.

Every UIP provides the opportunity to learn a constraint. We focused on the first and last UIP in our experiments, because ground CDNL very successfully learns only from the first UIP, and because we expect constraints from the last UIP also to be useful since they contain the decision from the current decision level. Investigating the usefulness of other UIPs, and finding other quality criteria to discriminate among learned constraints, belong to future work.

If no limit is imposed on the conflict generalisation algorithm, it runs until the problem is solved (one or more answer sets are found, or unsatisfiability is proven). To increase efficiency, resource consumption may be limited by stopping after a certain time or number of conflicts and then collecting results as if the problem had been solved. In our experiments, conflict generalisation was very effective even if resources were heavily limited, because the class of nogoods violated most often emerged at the very beginning of the solving process.

4.2 Continuation of the motivating example

The House Configuration Problem (HCP) was presented as a motivating example in Section 3. We have already presented the constraint $\langle 3 \rangle$ that can automatically be learned at the first UIP.

When continuing along the implication graph until the last UIP, two different non-ground

nogoods can be learned: namely, $\Omega_{\min} \cup \{\beta(r_{r2c}, \text{vars}(r_{r2c}))\}$ and $\Omega_{\min} \cup B(r_{r2c})$, where:

$$\begin{aligned} \Omega_{\min} = & \{ \mathbf{T} \text{ cabinetTOthing}(C, T1), \mathbf{T} \text{ personTOthing}(P1, T1), \\ & \mathbf{T} \text{ cabinetTOthing}(C, T2), \mathbf{T} \text{ personTOthing}(P2, T2), P1 < P2 \} \\ r_{r2c}: & \text{roomTOcabinet}(R, C) \leftarrow \text{roomDomain}(R), \text{cabinet}(C), \text{not } \widehat{\text{roomTOcabinet}}(R, C). \end{aligned}$$

In both nogoods, internal atoms can be replaced by ordinary atoms as described above to yield the following unique non-ground nogood⁶ learned at the last UIP:

$$\begin{aligned} \leftarrow & \text{roomTOcabinet}(R, C), \underline{\text{roomDomain}}(R), \underline{\text{cabinet}}(C), \\ & \text{cabinetTOthing}(C, T1), \text{personTOthing}(P1, T1), \\ & \text{cabinetTOthing}(C, T2), \text{personTOthing}(P2, T2), P1 < P2. \end{aligned} \quad \langle 9 \rangle$$

This nogood contains additional literals of domain predicates (underlined for emphasis), which is the only difference to the nogood $\langle 3 \rangle$ learned at the first UIP.

Reduction of the non-minimal learned nogood as described in Section 3 has not yet been implemented in our system. An implementation based on first-order theorem proving is conceivable, in which undecidability could be avoided by imposing a bound on the number of constants, e.g., a maximal number of entities in a configuration.

5 Experimental Results

We conducted a set of experiments on encodings of the House Reconfiguration Problem (HRP) and on a Graph Colouring problem (Listing 2) to demonstrate the feasibility of our approach. HRP extends HCP from Section 3, which has disregarded reconfiguration. All encodings and instances used for our experiments are available on our website⁷ and in the supplementary material corresponding to this paper at the TPLP archives. The HRP encoding was closely based on the encoding by Ryabokon (2015), except that the redundant constraint $\langle 2 \rangle$ has been removed, as described in Section 3, and due to syntactic restrictions of ALPHA some aggregates had to be rewritten and optimization statements were not used at all. Note that for reasons of a fair comparison, all solvers used the same encodings in our experiments, even though solvers supporting aggregates could have profited from a more sophisticated encoding.

Graph Colouring problems are an abstraction to which many real-world problems can be mapped. We have designed 100 satisfiable and 100 unsatisfiable graph instances (“3CC” for “3-colourable chains”) containing repeated patterns that force some pairs of nodes to have the same colour. On 3CC, our approach was able to learn constraints that represent this pattern. This was only possible by adapting procedure FINDNEXTLITERALFORRESOLUTION in Algorithm 1 to also use literals from the next-lower decision level ($\delta(l) \geq \delta(A) - 1$) for resolution. The question how many decision levels to consider is an interesting topic for future work. From the constraints learned from the conflicts occurring most often, one from the first UIP and one from the last UIP were chosen manually for the experiments. Since these constraints represent only the conflict for one specific combination of the three colours, they were multiplied manually to cover all possible combinations.⁸ The constraints learned at the first UIP were then reduced manually, similarly

⁶ We use nogood notation and constraint notation interchangeably for learned nogoods.

⁷ <http://ainf.aau.at/dynacon>

⁸ Each colour is represented by a predicate in our encoding. Depending on the actual conflicts during a specific solver run, constraints representing other colour combinations may be produced automatically.

```

1 blue(N) :- node(N), not red(N), not green(N).
2 red(N) :- node(N), not blue(N), not green(N).
3 green(N) :- node(N), not red(N), not blue(N).
4 :- link(N1,N2), blue(N1), blue(N2).
5 :- link(N1,N2), red(N1), red(N2).
6 :- link(N1,N2), green(N1), green(N2).

```

Listing 2. An encoding for our Graph Colouring Problem

Table 1. Number of HRP instances solved by CLINGO (#) per encoding; grounding and solving times in seconds (1st, 2nd, 3rd quartile)

Encoding	#	grounding time			solving time		
		Q1	Q2	Q3	Q1	Q2	Q3
original	34	6.787	39.738	—	4.315	66.580	—
1st UIP	22	32.150	—	—	1.030	—	—
last UIP	22	34.833	—	—	0.940	—	—
reduced	38	7.850	52.026	244.256	1.070	9.270	90.575

as has been described in Section 3, yielding the following final constraint $\langle I0 \rangle$ for red (and two more constraints for the other colours):

$$\leftarrow \text{red}(N12), \text{not red}(N22), \\ \text{link}(N12, N11), \text{link}(N12, N21), \text{link}(N11, N21), \text{link}(N11, N22), \text{link}(N21, N22). \quad \langle I0 \rangle$$

The ASP solvers ALPHA⁹ v0.5.0 (Weinzierl 2017), DLV 2.0 (Alviano et al. 2017), and CLINGO 5.4.0 (Gebser et al. 2014) were used. ALPHA was configured to ground rules strictly lazily and constraints permissively, as recommended by Taupe et al. (2019). The JVM was called with parameters `-Xms1G -Xmx32G`. For each problem instance, solvers searched for 10 answer sets.¹⁰

Experiments were run on machines each with two Intel[®] Xeon[®] CPU E5-2650 v4 @ 2.20GHz with 12 cores each, 252 GB of memory, and Ubuntu 16.04.1 LTS Linux. Benchmarks were scheduled with ABC Benchmarking System (Redl 2016) and HTCondorTM.¹¹ PYRUNLIM¹² was used to measure time and memory consumption and to limit time consumption to 10 minutes per instance, memory to 40 GiB and swapping to 0. Care was taken to avoid interference between CPUs, e.g., by not running different benchmarks concurrently on the same machine.

To compare solving performance using encodings with and without learned constraints, cactus plots (Figs. 1 to 6) have been created in the usual way. The x axis gives the number of instances solved within real (i.e., wall-clock) time which is given on the y axis. Time is accumulated over all solved instances. Since we are investigating the effects of constraints on each solver

⁹ <https://github.com/alpha-asp/Alpha>

¹⁰ Obtaining more than one (maybe trivial) answer set is often desirable. The number 10 has been chosen arbitrarily.

¹¹ <https://github.com/cred1/abcbenchmarking>, <http://research.cs.wisc.edu/htcondor>

¹² <https://alviano.com/software/pyrunlim/>




Fig. 1. HRP results with ALPHA

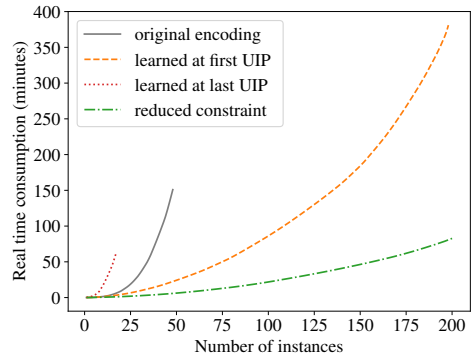


Fig. 2. 3CC results with ALPHA




Fig. 3. HRP results with DLV2




Fig. 4. 3CC results with DLV2




Fig. 5. HRP results with CLINGO




Fig. 6. 3CC results with CLINGO

and not comparing solvers against each other, maximum axis values vary between solvers. One curve has been drawn for each encoding: the original encoding, one encoding each with the additional constraint(s) learned automatically at the first/last UIP ($\langle 3 \rangle / \langle 9 \rangle$ for HRP), and one encoding in which only the reduced constraint(s) ($\langle 2 \rangle$ for HRP, $\langle 10 \rangle$ for 3CC) are added to the original encoding.

Learned constraints improved solving performance in many cases. For HRP, ALPHA benefited especially from the reduced constraint, but also from the constraint learned automatically at

the first UIP. DLV2 profited most from the first-UIP constraint, but the reduced one was not far behind. The last-UIP constraint affected DLV2's performance negatively, however. CLINGO profited only from the reduced constraint, while automatically learned constraints caused overall solving time to increase. Detailed analysis of grounding and solving times as reported by CLINGO (Table 1) shows that this was due to additional grounding effort induced by learned constraints. With automatically learned constraints, more than half of the instances could not be solved within 10 minutes, therefore no median times (second quartiles) can be shown. In almost all time-out cases, CLINGO did not manage to finish grounding and start solving. For 3CC, all solvers perform best with the reduced constraints, while the 1st-UIP constraints also improve performance.

Learning itself is cheap: On the first author's computer, the relevant HRP constraints are learned using the easiest instance in less than 6 seconds if search is stopped after 50 conflicts and in less than 14 seconds if search is carried out until the first answer set is found (our implementation has not yet been tuned for optimal performance). Performance for 3CC is similar. Computational complexity has not yet been analysed and could be addressed in future work.

The results show that our approach is able to improve ASP solving using both lazy-grounding and ground-and-solve systems. While effects vary between types of learned constraints and solver implementations, every system under investigation has profited from at least one learned constraint in both domains.

6 Conclusions

We have proposed an extension of CDNL that, while solving one problem instance, learns non-ground nogoods that can be used to speed up the solving of other problem instances. As far as we know, this is the first attempt to generalise and re-use knowledge learned during ASP solving. Experimental results showed compelling benefits of our approach: Both ground-and-solve systems and lazy-grounding systems performed significantly better on instances of a practical configuration problem and a graph colouring problem when using constraints learned by our method, solving more instances and/or solving the problems faster.

So far, we have been experimenting with encodings for two problem domains: a graph colouring problem, and an important part of many configuration domains (those where systems are composed hierarchically). However, it seems natural to assume that many domains feature redundant constraints that may not be obvious to a human modeller. We therefore see our approach mainly as a tool to support the design of efficient answer-set programs.

It remains to be clarified whether our approach could improve encodings used by the ASP competitions (Gebser et al. 2020), which generally are already heavily optimised, or whether it would prove more useful when applied to encodings devised by inexperienced modellers.

Acknowledgments. This work has been conducted in the scope of the research project *DynaCon* (FFG-PNr.: 861263), which is funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under the program "ICT of the Future" between 2017 and 2020,¹³ and in the scope of the research project Productive4.0, which is funded by EU-ECSEL under grant agreement no737459.

¹³ See <https://iktderzukunft.at/en/> for more information.

References

- ALVIANO, M., CALIMERI, F., DODARO, C., FUSCÀ, D., LEONE, N., PERRI, S., RICCA, F., VELTRI, P., AND ZANGARI, J. 2017. The ASP system DLV2. In *LPNMR*. LNCS, vol. 10377. Springer, 215–221.
- BOGAERTS, B. AND WEINZIERL, A. 2018. Exploiting justifications for lazy grounding of answer set programs. In *IJCAI*. ijcai.org, 1737–1745.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., MARATEA, M., RICCA, F., AND SCHAUB, T. 2020. ASP-Core-2 input language format. *Theory Pract. Log. Program.* 20, 2, 294–309.
- DEJONG, G. AND MOONEY, R. J. 1986. Explanation-based learning: An alternative view. *Mach. Learn.* 1, 2, 145–176.
- FABER, W., PFEIFER, G., AND LEONE, N. 2011. Semantics and complexity of recursive aggregates in answer set programming. *Artif. Intell.* 175, 1, 278–298.
- FRIEDRICH, G., RYBOKON, A., FALKNER, A. A., HASELBÖCK, A., SCHENNER, G., AND SCHREINER, H. 2011. (Re)configuration using answer set programming. In *Configuration Workshop*. CEUR-WS.org.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., AND SCHAUB, T. 2014. Clingo = ASP + control: Preliminary report. *CoRR abs/1405.3694*.
- GEBSER, M., KAUFMANN, B., AND SCHAUB, T. 2012. Conflict-driven answer set solving: From theory to practice. *Artif. Intell.* 187, 52–89.
- GEBSER, M., MARATEA, M., AND RICCA, F. 2020. The seventh answer set programming competition: Design and results. *Theory Pract. Log. Program.* 20, 2, 176–204.
- VAN HARMELEN, F. AND BUNDY, A. 1988. Explanation-based generalisation = partial evaluation. *Artif. Intell.* 36, 3, 401–412.
- HIRSH, H. 1987. Explanation-based generalization in a logic-programming environment. In *IJCAI*. Morgan Kaufmann, 221–227.
- LEUTGEB, L. AND WEINZIERL, A. 2017. Techniques for efficient lazy-grounding ASP solving. In *DECLARE*. LNCS, vol. 10997. Springer, 132–148.
- LINTAO ZHANG, MADIGAN, C. F., MOSKEWICZ, M. H., AND MALIK, S. 2001. Efficient conflict driven learning in a boolean satisfiability solver. In *ICCAD*. IEEE, 279–285.
- MITCHELL, T. M. 1997. *Machine learning, International Edition*. McGraw-Hill Series in Computer Science. McGraw-Hill.
- MITCHELL, T. M., KELLER, R. M., AND KEDAR-CABELLI, S. T. 1986. Explanation-based generalization: A unifying view. *Mach. Learn.* 1, 1, 47–80.
- REDL, C. 2016. Automated benchmarking of KR-systems. In *RCRA@AI*IA*. CEUR Workshop Proceedings, vol. 1745. CEUR-WS.org, 45–56.
- RUSSELL, S. J. AND NORVIG, P. 2010. *Artificial Intelligence – A Modern Approach, Third International Edition*. Pearson Education.
- RYBOKON, A. 2015. Knowledge-based (re)configuration of complex products and services. Ph.D. thesis, Alpen-Adria-Universität Klagenfurt.
- SHEPHERDSON, J. C. 1984. Negation as failure: A comparison of clark’s completed data base and reiter’s closed world assumption. *J. Log. Program.* 1, 1, 51–79.
- SILVA, J. P. M., LYNCE, I., AND MALIK, S. 2009. Conflict-driven clause learning SAT solvers. In *Handbook of Satisfiability*. IOS Press, 131–153.
- TAUPE, R., WEINZIERL, A., AND FRIEDRICH, G. 2019. Degrees of laziness in grounding – effects of lazy-grounding strategies on ASP solving. In *LPNMR*. LNCS, vol. 11481. Springer, 298–311.
- WEINZIERL, A. 2013. Learning non-ground rules for answer-set solving. In *2nd Workshop on Grounding and Transformations for Theories With Variables*. 25–37.
- WEINZIERL, A. 2017. Blending lazy-grounding and CDNL search for answer-set solving. In *LPNMR*. LNCS, vol. 10377. Springer, 191–204.