

TECHNICAL REPORT

Modeling Variations of First-Order Horn Abduction
in Answer Set Programming

Peter Schüller*

Institut für Logic and Computation, Knowledge-Based Systems Group

Technische Universität Wien, Austria

Computer Engineering Department, Faculty of Engineering

Marmara University, Turkey

schueller.p@gmail.com

February 1, 2018

Abstract

We study abduction in First Order Horn logic theories where all atoms can be abduced and we are looking for preferred solutions with respect to three objective functions: cardinality minimality, coherence, and weighted abduction. We represent this reasoning problem in Answer Set Programming (ASP), in order to obtain a flexible framework for experimenting with global constraints and objective functions, and to test the boundaries of what is possible with ASP. Realizing this problem in ASP is challenging as it requires value invention and equivalence between certain constants, because the Unique Names Assumption does not hold in general. To permit reasoning in cyclic theories, we formally describe fine-grained variations of limiting Skolemization. We identify term equivalence as a main instantiation bottleneck, and improve the efficiency of our approach with on-demand constraints that were used to eliminate the same bottleneck in state-of-the-art solvers. We evaluate our approach experimentally on the ACCEL benchmark for plan recognition in Natural Language Understanding. Our encodings are publicly available, modular, and our approach is more efficient than state-of-the-art solvers on the ACCEL benchmark.

1 Introduction

Abduction [Pei55] is reasoning to the best explanation, which is an important topic in diverse areas such as diagnosis, planning, and natural language understanding (NLU).

We here focus on a variant of abduction, used in NLU, where the primary concern is to find an explanation of a given input (sentence) with respect to an objective function. Knowledge is expressed in First Order (FO) Horn logic axioms. For example ‘a father of somebody is male’ can be expressed as follows, where capital letters are variables which are universally quantified unless explicitly indicated otherwise.

$$inst(X, male) \Leftarrow fatherof(X, Y).$$

Abduction aims to find a set of explanatory atoms that make a set of goal atoms true with respect to a background theory (i.e., a set of axioms). If $inst(tom, male)$ is part of a goal then abduction can explain this goal atom with the atom $fatherof(tom, mary)$ where $mary$ is another person of interest.

*This work is a significant extension of [Sch15]; major additions are preference relations COH and WA, revised encodings, increase performance, on-demand constraints, and flexible value invention. This work has been supported by Scientific and Technological Research Council of Turkey (TUBITAK) Grant 114E777. This document is a preprint of [Sch16] with minor formatting corrections.

Using abduction, we can interpret whole natural language texts, for example ‘Mary lost her father. She is depressed.’ can be interpreted using knowledge about losing a person, death of a person, and being depressed, such that we obtain an abductive explanation that represents ‘Mary’s father died, and this is the reason for her depression’.

Abductive reasoning in FO Horn logic yields an infinite space of potential inferences, because backward reasoning over axioms can produce existentially quantified variables, which can introduce new terms (value invention). For example the above axiom is transformed as follows.

$$inst(X, male) \Rightarrow \exists Y : fatherof(X, Y).$$

To achieve decidability, we need to limit value invention, which leads to a challenging trade-off: the more we limit value invention, the more (potentially optimal) solutions we lose.

A second challenge in FO logic is, that terms (input and invented) can be equivalent to other terms. Equivalent terms make distinct atoms equivalent, which is used in an inference called *factoring*. In the above example, we can say that $is(mary, depressed)$ is factored with $is(she, depressed)$ under the assumption that the equivalence $mary = she$ holds.

A crucial issue when using abduction for NLU is the choice of an appropriate *preference* among possible abductive explanations. *Cardinality minimality* of the set of abduced atoms is a frequently used preference, however in NLU two other preferences have turned out to be more effective: *coherence* [NM90, Ng92], and *weighted abduction* [Sti89, HSME93, SM11], which are based on a proof graph induced by back-chaining and unification operations.

Several tools for realizing abduction with such preferences exist: *Phillip*, based on Integer Linear Programming (ILP) [YII⁺15] and its precursor *Henry-n700* [II13] as well as an approach based on Markov Logic [BHD⁺11]. The problem of termination is solved in [BHD⁺11] by instantiating existential terms only with terms from the input (no value invention), while [II13, YII⁺15] solves the issue by inventing a new term only if no previously invented term is present in the head of the axiom.

Unfortunately, using only input terms eliminates many valid solutions in NLU applications, for example when processing a text about a son and a grandfather, we would be unable to perform reasoning about the father (because it does not exist as a constant). The alternative approach of blocking value invention if an invented value is involved in the rule improves the situation, however it is an ad hoc solution and its implications on solution quality have not been formally or experimentally analyzed.

In addition to decidability issues, global consistency constraints are necessary to yield practically meaningful abductive explanations, however existing solvers *Henry-n700* and *Phillip* realize each possible form of a global constraint (e.g., unique slot values) in a separate checking procedure and make it difficult to experiment with additional constraints.

Towards overcoming some of these problems, we realize abduction in the declarative formalism of Answer Set Programming (ASP) [Lif08] which allows modular modeling of combinatorial problems based on clearly defined formal semantics. Primary motivation for this work is to obtain a more flexible framework where variations of Skolemization, objective functions and global constraints on abduction can be studied easily and where novel preferences can be studied, such as [Sch14] that can comfortably be represented in ASP but not in other solvers. Our secondary motivation is, to use ASP for realizing a task that it is not typically used for, and to study how far we can go in this task.

Realizing FO Horn abduction in ASP poses further challenges for the following reasons:(i) ASP semantics is based on the Unique Names Assumption (UNA), which means that distinct terms cannot be equivalent (e.g., $mary = she$ is not expressible in a built-in feature of ASP); moreover (ii) ASP rules have no built-in support for existential variables in rule heads, which is necessary for value invention during back-chaining as shown above. In particular, Skolemization using function symbols, i.e., replacing $\exists Y : fatherof(X, Y)$ by $fatherof(X, sk(X))$ where sk is a new function symbol, does not guarantee a finite instantiation.

We tackle the above challenges and present an ASP framework for solving FO Horn abduction problems for objective functions weighted abduction, coherence, and cardinality minimality. We describe insights on the structure of the problem as well as insights on the efficiency of straightforward versus more involved ASP formulations. Our formulation allows a fine-grained configuration of Skolemization for tackling cyclic background theories, moreover it permits the usage of global constraints of any form that is expressible in ASP. Experiments show, that our framework is faster than the state-of-the-art solver *Phillip* [YII⁺15] on the ACCEL benchmark [NM92] for plan recognition in NLU.

In detail, we make the following contributions.

- We provide a novel uniform formalization of abduction with preference relations weighted abduction, coherence, and cardinality minimality (Section 2).
- We present an ASP encoding that realizes back-chaining in ASP by deterministically representing an abductive proof graph and guessing which parts of that graph to use. For value invention we use uninterpreted function terms, and we explicitly represent an equivalence relation between terms to model unification (Section 3.2).
- We describe canonicalization operations on proof graphs, show that they do not eliminate optimal solutions, and use these transformations for encoding factoring efficiently (Section 3.3).
- We present an alternative ASP encoding which does not represent a proof graph, instead it generates abduced atoms, defines truth using axioms, and tests if goals are reproduced (Section 3.4).
- We give ASP encodings for realizing the objective functions weighted abduction, coherence, and cardinality minimality (Section 3.5).
- We study an alternative method for value invention by replacing uninterpreted function terms with external computations. This provides us with a fine-grained control over Skolemization, which is more flexible than state-of-the-art solutions for achieving decidability. We formalize this extension using the HEX formalism, and show termination guarantees for arbitrary (i.e., including cyclic) knowledge bases (Section 4.1).
- We apply a technique used to increase performance of the Henry-n700 solver [II13] for weighted abduction to our encodings by introducing on-demand constraints for transitivity of the equivalence relation and for ensuring acyclicity of the proof graph. We formalize this using HEX and describe an algorithm for computing optimal models in the presence of on-demand constraints using the Python API of Clingo [GKKS14] (Section 4.2).
- We perform computational experiments on the ACCEL benchmark [NM92] where we measure and discuss resource consumption in terms of space, time, and solver-internal statistics, as well as solution quality in terms of the objective function (Section 5). For experimental evaluation we use the Python API of Clingo, Gringo [GKKS11] with either Clasp [GKK⁺15] or Wasp [ADLR15], and the Phillip [YII⁺15] solver for weighted abduction which is based on C++ and Integer Linear Programming (ILP).

We discuss related work in Section 6 and conclude in Section 7.

Appendices provide additional information: all ASP encodings in their complete version; verbose listing of rewriting and answer set of the running example; proofs for correctness of encodings; and algorithms for realizing on-demand constraints.

Our framework, including experimental instances and algorithms, is available online.¹

2 Preliminaries

We give a brief introduction of abduction in general and variations of First Order Horn abduction as used in Natural Language Processing, describe the ACCEL benchmark which contains instances of such reasoning problems, and give preliminaries of ASP and HEX.

In the following, in logical expressions and ASP rules we write variables starting with capital letters and constants starting with small letters.

2.1 Abduction and Preferences on Abductive Explanations

Abduction, originally described in [Pei55], can be defined logically as follows: given a set B of background knowledge axioms and an observation O , find a set H of hypothesis atoms such that B and H are consistent and reproduce the observation, i.e., $B \cup H \not\models \perp$ and $B \cup H \models O$. In this work we formalize

¹<https://bitbucket.org/knowlp/asp-fo-abduction>

axioms and observations in First Order logic: the observation O (also called ‘goal’) is an existentially quantified conjunction of atoms

$$\exists V_1, \dots, V_k : o_1(V_1, \dots, V_k) \wedge \dots \wedge o_m(V_1, \dots, V_k) \quad (1)$$

and an axiom in B is a Horn clause of form

$$q(Y_1, \dots, Y_m) \Leftarrow p_1(X_1^1, \dots, X_{k_1}^1) \wedge \dots \wedge p_r(X_1^r, \dots, X_{k_r}^r). \quad (2)$$

where $\mathcal{X} = \bigcup_{1 \leq i \leq r} \bigcup_{1 \leq j \leq k_r} X_j^i$ is the set of variables in the body, $\mathcal{Y} = \bigcup_{1 \leq i \leq m} Y_i$ is the set of variables in the head, $\mathcal{Y} \subseteq \mathcal{X}$ and we implicitly quantify universally over \mathcal{X} . In the variant of abduction we consider here, the set H of hypotheses can contain any predicate from the theory B and the goal O , hence existence of a solution is trivial. A subset S of constants from B is declared as ‘sort names’ that cannot be equivalent with other constants. Given B , O , and S , we call the tuple (B, O, S) an *abduction instance*. Unless otherwise indicated, we assume that B is acyclic.

Example 1 (Running Example). Consider the following text

‘Mary lost her father. She is depressed.’

which can be encoded as the following observation, to be explained by abduction.

$$name(m, mary) \wedge lost(m, f) \wedge fatherof(f, m) \wedge inst(s, female) \wedge is(s, depressed) \quad (3)$$

Given the set of axioms

$$inst(X, male) \Leftarrow fatherof(X, Y) \quad (4)$$

$$inst(X, female) \Leftarrow name(X, mary) \quad (5)$$

$$importantfor(Y, X) \Leftarrow fatherof(Y, X) \quad (6)$$

$$inst(X, person) \Leftarrow inst(X, male) \quad (7)$$

$$is(X, depressed) \Leftarrow inst(X, pessimist) \quad (8)$$

$$is(X, depressed) \Leftarrow is(Y, dead) \wedge importantfor(Y, X) \quad (9)$$

$$lost(X, Y) \Leftarrow is(Y, dead) \wedge importantfor(Y, X) \wedge inst(Y, person) \quad (10)$$

and sort names

$$person \quad male \quad female \quad dead \quad depressed \quad (11)$$

we can use abduction to conclude the following: (a) loss of a person here should be interpreted as death, (b) ‘she’ refers to Mary, and (c) her depression is because of her father’s death because her father was important for her.

We obtain these because we can explain (3) by the following abductive explanation which contains atoms and equivalences.

$$name(m, mary) \quad fatherof(f, m) \quad is(f, dead) \quad m = s \quad (12)$$

The first two atoms directly explain goal atoms. We can explain the remaining goal atoms using inference from rules and factoring (which represents unification in a certain reasoning direction).

$$inst(f, male) \quad [\text{inferred via (4) using (12)}] \quad (13)$$

$$inst(m, female) \quad [\text{inferred via (5) using (12)}] \quad (14)$$

$$inst(s, female) \quad [\text{goal, factored from (14) using (12)}]$$

$$importantfor(f, m) \quad [\text{inferred via (6) using (12)}] \quad (15)$$

$$inst(f, person) \quad [\text{inferred via (7) using (13)}] \quad (16)$$

$$is(m, depressed) \quad [\text{inferred via (9) using (12) and (15)}] \quad (17)$$

$$is(s, depressed) \quad [\text{goal, factored from (17) using (12)}]$$

$$lost(m, f) \quad [\text{goal, inferred via (10) using (12), (15), and (16)}] \quad (18)$$

Note that there are additional possible inferences but they are not necessary to explain the goal atoms. Moreover, there are several other abductive explanations, for example to abduce all goal atoms, or to abduce $inst(m, pessimist)$ and $lost(m, f)$ instead of abducing $is(f, dead)$. \square

Preferred explanations. In the presence of multiple possible explanations, we are naturally interested in obtaining a set of *preferred* explanations. In this work we consider three preference formulations: (CARD) cardinality-minimality of abduced atoms, (COH) ‘coherence’ as described in [NM92] and in slight variation in more detail in [NM90], and (WA) ‘weighted abduction’ as initially formulated in [HSME93].

COH is based on connectedness between observations and explanations, while WA finds a trade-off between least-specific and most-specific abduction, depending on the explanatory power of more specific atoms. Both objective functions are based on an inference procedure that induces a proof graph by means of backchaining and unification.

Towards formalizing these preference functions, we next formalize such an inference procedure. The following definitions are based on [Sti89, IOIH14] but additionally define an explicit proof graph corresponding to the inferences leading to a hypothesis. Here we consider only single-head axioms.

Definition 1. A *hypothesis* is a conjunction of atoms or equivalences between terms. Given an abduction instance $A = (B, O, S)$ the set $\hat{\mathcal{H}}(A)$ of all hypotheses of A is the largest set containing hypotheses obtained by extending $\mathcal{H} = \{O\}$ using back-chaining and unification.

Back-chaining: given an atom P which is part of a hypothesis ($P \in H, H \in \mathcal{H}$), such that P unifies with the head $Q = q(Y_1, \dots, Y_m)$ of an axiom (2) with substitution θ , back-chaining adds to \mathcal{H} a new hypothesis $H \wedge P'_1 \wedge \dots \wedge P'_r$, where P'_i is the substituted version $\theta(p_i(X_1^i, \dots, X_{k_i}^i))$ of the i -th body atom of axiom (2).

Unification: given hypothesis $H \in \mathcal{H}$ with distinct atoms $P, Q \in H$ that unify under substitution θ such that all $X \mapsto Y \in \theta$ obey $X, Y \notin S$, unification adds hypothesis $H \wedge \bigwedge \{X = Y \mid X \mapsto Y \in \theta\}$ to \mathcal{H} .

Note that $\hat{\mathcal{H}}(A)$ is potentially infinite. We sometimes leave A implicit.

Example 2 (continued). Three hypotheses for the abduction instance in Example 1 are (a) the original goal G as shown in (3), which intuitively means that we do not justify any atom in the goal by inference, instead we abduce all atoms in the goal; (b) the hypothesis $G \wedge is(f, dead) \wedge m = s \wedge inst(f, male) \wedge inst(m, female) \wedge importantfor(f, m) \wedge inst(f, person) \wedge is(m, depressed)$ which corresponds to (13)–(18) and includes the abductive explanation (12); and (c) the hypothesis $G \wedge name(s, mary) \wedge m = s \wedge inst(f, person) \wedge inst(f, male) \wedge fatherof(f, n_2) \wedge m = n_2 \wedge is(n_1, dead) \wedge f = n_1 \wedge is(f, dead) \wedge importantfor(f, m) \wedge importantfor(n_1, m)$ which applies Skolemization during back-chaining and represents a variation of explanation (12). Details about this hypothesis can be found in Example 3 and in Figure 1. \square

A hypothesis $H \in \hat{\mathcal{H}}$ does not contain any information about how it was generated. For the purpose of describing the cost function formally, we define proof graphs G wrt. hypotheses H .

Definition 2. Given an abduction instance $A = (B, O, S)$, a *proof graph* G wrt. a hypothesis $H \in \hat{\mathcal{H}}(A)$ is an acyclic directed graph consisting of nodes $N(G) = \{P \in H \mid P \text{ is not an equality}\}$ and edges $E(G)$ are recursively defined by the inference operations used to generate atoms $P \in H$: (a) back-chaining of P induces an edge from all body atoms P'_i to P , and (b) unification of P with Q induces either an edge from P to Q or an edge from Q to P .

We denote by $A(G) = \{a \in N(G) \mid \nexists b : (b, a) \in E(G)\}$ the set of nodes that are neither back-chained nor unified. Note that the term ‘factoring’ is used to denote unification with direction, this is discussed in detail in [Sti89]. There can be multiple proof graphs with respect to a single hypothesis, and these graphs differ only by factoring directions. Figure 1 depicts a proof graph which is discussed in Example 3.

Intuitively, an edge in the proof graph shows how an atom is justified: by inference over an axiom (backchaining) or by equivalence with another atom (factoring).

Equipped with these definitions, we next formalize the objective functions of interest.

Definition 3. Given a proof graph G wrt. a hypothesis H of an abduction instance (B, O, S) ,

- CARD = $|A(G)|$,
- COH = $|\{(a, b) \mid a, b \in O, a < b, \text{ and } \nexists n \in N(G) \text{ such that from } n \text{ we can reach both } a \text{ and } b \text{ in } G\}|$ where the relation $<$ is an arbitrary fixed total order over O (e.g., lexicographic order),
- WA = $\sum_{a \in A(G)} \min cost(a)$, where $cost : N(G) \rightarrow 2^{\mathbb{R}}$ labels each atom in the graph with a set of cost values.

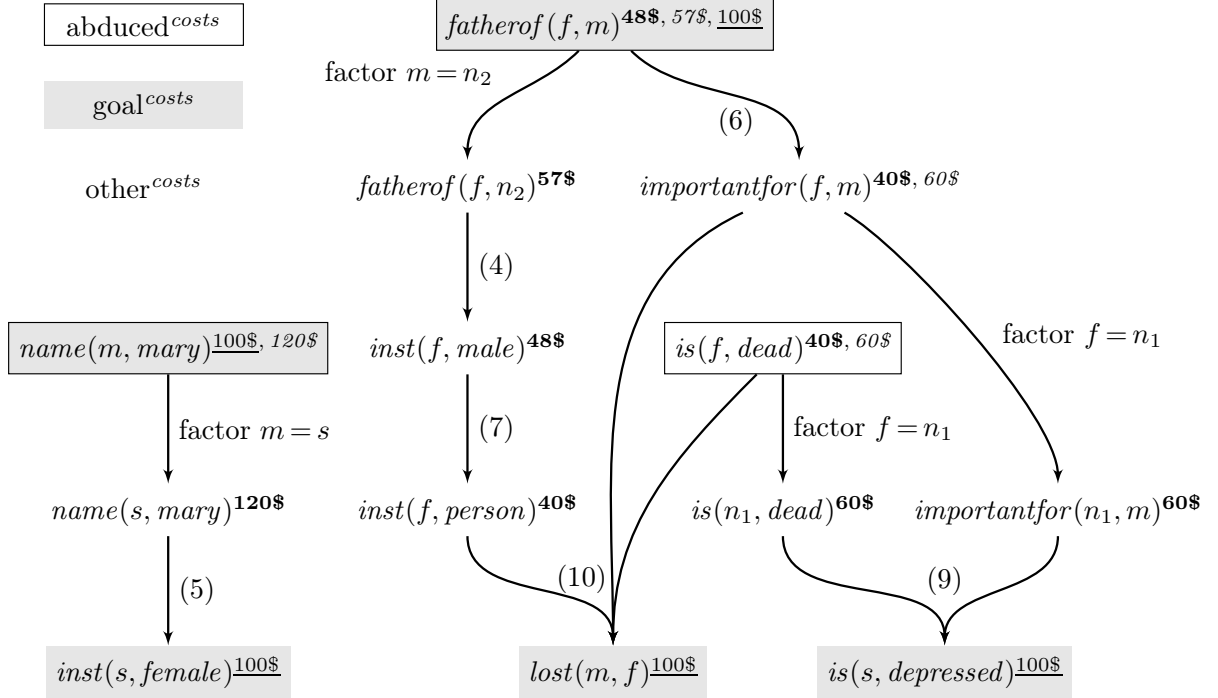


Figure 1: Abductive proof graph of Example 1 including costs of nodes for WA. Edges annotated with (X) are inference edges induced by back-chaining over axiom (X). We underline initial goal costs, use italic font for factoring costs, and bold font for back-chaining costs.

For the definition of *cost* in WA we require that each axioms of form (2) has weights w_1, \dots, w_r corresponding to its body atoms, and initial costs $ic(o)$ for each observation $o \in O$.

Then *cost* is initialized with \emptyset for each node and recursively defined as follows:

- goal nodes $o \in O$ obtain cost $cost(o) = cost(o) \cup \{ic(o)\}$;
- if P was back-chained with an axiom with body atoms P'_1, \dots, P'_r and $c = \min cost(P)$, then c is added to each body atom P'_i after adjusting it using the respective cost multiplier w_i , formally $cost(P'_i) = cost(P'_i) \cup \{c \cdot w_i\}$ for $1 \leq i \leq r$;
- if P was unified with Q such that there is a factoring edge $(Q, P) \in G$ from Q to P , then we add the smallest cost at P to Q : $cost(Q) = cost(Q) \cup \{\min cost(P)\}$.

Note that the formalization in [IOIH14, above (1)] assigns unification cost to the equality, but does not use that cost in case of multiple unifications, hence we do not use such a formalization. Moreover, deleting cost values from the hypothesis with higher cost in a unification (as shown in [IOIH14, Fig. 1, ‘Output’ vs ‘Backward-chaining’]) contradicts the formalization as a cost ‘function’ that maps hypotheses to costs. Therefore our formalization defines *cost* to map from atoms in a hypothesis to multiple ‘potential’ costs of that hypothesis. Note that due to acyclicity of the graph, no cost in *cost* recursively depends on itself, and that back-chaining can create hypothesis atoms containing a part of the goal, therefore goal nodes can have a cost lower than that goal’s initial cost (e.g., $fatherof(f, m)$ in Figure 1).

The formalization in this section was done to provide a basis for showing correctness of canonicalization operations on proof graphs and for showing correctness of ASP encodings. To the best of our knowledge, no precise formal description of proof graphs and associated costs of COH and WA exists in the literature, therefore we here attempt to formally capture the existing descriptions [NM92, HSME93, II13].

Example 3 (continued). The proof graph of Example 1 is depicted in Figure 1 where we also show the set of costs of each node using objective WA. The total cost of this graph is $100\$ + 48\$ + 40\$ = 188\$$. Objective CARD has cost 3 because we abduce 3 atoms, and objective COH has cost 6: let goal node set $A = \{inst(s, female), name(m, mary)\}$ and $B = \{fatherof(f, m), lost(m, f), is(s, depressed)\}$, then nodes within A and within B are reachable from some node, however pairs $\{(a, b) \mid a \in A, b \in B\}$ of nodes are not reachable from any node, and each of these $|A| \cdot |B| = 6$ pairs incurs cost 1. \square

Note that in this work we consider only hypotheses and proof graphs where an atom is either justified by a single inference, or by a single factoring, or not at all (then it is abduced).

Computational Complexity. Existence of an abductive explanation is trivial, because we can abduce any atom and the goal is the trivial explanation. However, finding the optimal abductive explanation with respect to an objective function is not trivial. With unlimited value invention and cyclic theories the problem of finding the optimal explanation is undecidable, as we cannot bound the size of the proof graph or the number of additionally required constants for finding the optimal solution.

To the best of our knowledge, the computational complexity of deciding whether an abductive explanation is optimal wrt. one of the objective functions CARD, COH, and WA, has not been formally studied so far, although for CARD related results exist. Section 6 discusses related complexity results.

2.2 ACCEL Benchmark

The ACCEL benchmark² [NM92, Ng92] contains a knowledge base with 190 axioms of form (2), defines a set of sort names that observe the UNA, and contains 50 instances (i.e., goals) with between 5 and 26 atoms in a goal (12.6 atoms on average). Axioms contain a single head and bodies vary between 1 and 11 atoms (2.2 on average). ACCEL axioms contain no weights and goal atoms contain no initial costs. For experiments with WA we follow existing practice (cf. [III13]) and set initial costs to $ic(o) = 100\$$ for all $o \in O$ and for each axiom we set weights to sum up to 1.2, i.e., we set $w_i = 1.2/r$, $1 \leq i \leq r$.

In addition to axioms, goals, and sort names, ACCEL contains constraints that forbid certain combinations of atoms to become abduced at the same time (assumption nogoods) and constraints that enforce functionality for certain predicate symbols (unique slot axioms). We next give two examples.

Example 4. An example for an assumption nogood is, that we are not allowed to abduce an event G to be part of a ‘go’ event S , and at the same time abduce that a person P is the ‘goer’ of G .

$$\bar{\Delta} S, G, P : \{go_step(S, G), goer(G, P)\} \in H \text{ for all } H \in \hat{\mathcal{H}} \quad (19)$$

An example for a unique slot axiom is, that the ‘goer’ of an event must be unique.

$$\bar{\Delta} G, P_1, P_2 : P_1 \neq P_2 \wedge \{goer(G, P_1), goer(G, P_2)\} \in H \text{ for all } H \in \hat{\mathcal{H}} \quad (20)$$

□

2.3 Answer Set Programming

We assume familiarity with ASP [GL88, Lif08, EIK09, GKKS12] and give only brief preliminaries of HEX programs [EFI⁺16] which extend the ASP-Core-2 standard [CFG⁺12]. We will use programs with (uninterpreted) function symbols, aggregates, choices, and weak constraints.

Syntax. Let \mathcal{C} , \mathcal{X} , and \mathcal{G} be mutually disjoint sets of *constants*, *variables*, and *external predicate names*, which we denote with first letter in lower case, upper case, and starting with ‘&’, respectively. Constant names serve as constant terms, predicate names, and names for uninterpreted functions. The set of *terms* \mathcal{T} is recursively defined, it is the smallest set containing $\mathbb{N} \cup \mathcal{C} \cup \mathcal{X}$ as well as uninterpreted function terms of form $f(t_1, \dots, t_n)$ where $f \in \mathcal{C}$ and $t_1, \dots, t_n \in \mathcal{T}$. An *ordinary atom* is of the form $p(t_1, \dots, t_n)$, where $p \in \mathcal{C}$, $t_1, \dots, t_n \in \mathcal{T}$, and $n \geq 0$ is the *arity* of the atom. An *aggregate atom* is of the form $X = \#agg\{e_1; \dots; e_k\}$ with variable $X \in \mathcal{X}$, aggregation function $\#agg \in \{\#min, \#max\}$, $k \geq 1$ and each aggregate element e_i , $1 \leq i \leq k$, is of the form $t:a$ or t with $t \in \mathcal{T}$ and a an atom. An *external atom* is of the form $\&f[y_1, \dots, y_n](x_1, \dots, x_m)$, where $y_1, \dots, y_n, x_1, \dots, x_m \in \mathcal{T}$ are two lists of terms (called *input* and *output* lists, resp.), and $\&f \in \mathcal{G}$ is an external predicate name. An external atom provides a way for deciding the truth value of an output tuple depending on the input tuple and a given interpretation. A term or atom is *ground* if it contains no sub-terms that are variables.

A *rule* r is of the form $\alpha_1 \vee \dots \vee \alpha_k \leftarrow \beta_1, \dots, \beta_n, \mathbf{not} \beta_{n+1}, \dots, \mathbf{not} \beta_m$ where $m, k \geq 0$, α_i , $0 \leq i \leq k$ is an ordinary atom and β_j , $0 \leq j \leq m$ is an atom, and we let $H(r) = \{\alpha_1, \dots, \alpha_k\}$ and $B(r) = \{\beta_1, \dots, \beta_n, \mathbf{not} \beta_{n+1}, \dots, \mathbf{not} \beta_m\}$. A *program* is a finite set P of rules. A rule r is a *constraint*, if $k = 0$ and $m \neq 0$, and a *fact* if $m = 0$.

²Available at <ftp://ftp.cs.utexas.edu/pub/mooney/accel> .

A *weak constraint* is of form $\sim \beta_1, \dots, \beta_n, \mathbf{not} \beta_{n+1}, \dots, \mathbf{not} \beta_m. [w@1, t_1, \dots, t_k]$ where all β_j are atoms, and all t_i are terms such that each variable in some t_i is contained in some β_j (note that 1 in $w@1$ shows the ‘level’ which we do not use).

Semantics. Semantics of a HEX program P are defined using its Herbrand Base HB_P and its ground instantiation $grnd(P)$. An aggregate literal in the body of a rule accumulates truth values from a set of atoms, e.g., $C = \#min\{4; 2 : p(2)\}$ is true wrt. an interpretation $I \subseteq HB_P$ iff $p(2) \in I$ and $C = 2$ or $p(2) \notin I$ and $C = 4$. Using the usual notion of satisfying a rule given an interpretation, the FLP-reduct [FPL11] fP^I reduces a program P using an answer set candidate I : $fP^I = \{r \in grnd(P) \mid I \models B(r)\}$. I is an answer set of P ($I \in AS(P)$) iff I is a minimal model of fP^I . Weak constraints define that an answer set I has cost equivalent to the term w for each distinct tuple t_1, \dots, t_k of constraints that have a satisfied body wrt. I . Answer sets of the lowest cost are preferred.

Safety and Splitting. Programs must obey *syntactic safety restrictions* (see [CFG⁺12]) to ensure a finite instantiation. In presence of loops over external atoms, HEX programs additionally must obey restrictions to ensure finite instantiation. A *splitting set* [LT94] of a program P is any set U of literals such that, for every rule $r \in P$, if $H(r) \cap U \neq \emptyset$ then $B(r) \subseteq U$. The set of rules $r \in P$ such that $B(r) \subseteq U$ is called the bottom $b_U(P)$ of P relative to U . Given splitting set U of program P , $I \in AS(P)$ iff $I = X \cup Y$ where $X \cap Y = \emptyset$, $X \in AS(b_U(P))$, and $Y \in AS(e_U(P \setminus b_U(P), X))$ where $e_U(Q, J)$ partially evaluates Q wrt. J . Splitting sets were lifted to HEX in [EIST06, EFI⁺16].

Syntactic Sugar. Anonymous variables of form ‘_’ are replaced by new variable symbols. Choice constructions can occur instead of rule heads, they generate a set of candidate solutions if the rule body is satisfied; e.g., $1 \leq \{p(a); p(b); p(c)\} \leq 2$ in the rule head generates all solution candidates where at least 1 and at most 2 atoms of the set are true. The bounds can be omitted. In choices, the colon symbol ‘:’ can be used to generate a list of elements (similar as in aggregates), for example $\{p(X) : q(X), \mathbf{not} r(X)\}$ encodes a guess over all $p(X)$ such that $q(X)$ is true and $r(X)$ is not true.

3 ASP Encodings

We next describe ASP encodings for modeling abduction with partial UNA and value invention. All encodings consist of a deterministic part that instantiates all atoms that can potentially used to build a hypothesis, i.e., to justify the goal. Some encodings also explicitly represent inferences leading to these atoms. All encodings guess an equivalence relation over all terms in these ‘potentially interesting’ atoms, to handle term equivalence. In BWD encodings, the actually used proof graph is nondeterministically guessed, while encoding FWD-A performs a guess of abduced atoms and checks if the goals become true given these atoms. Variations of BWD encodings perform factoring in different ways. We next give common aspects of all encodings, then give each encoding in detail, and finally provide a summary of differences between encodings (Table 1).

A detailed example containing a rewriting of our running example and an answer set corresponding to Figure 1 is given in the appendix.

We represent an atom of the form $p(a, b)$ as a term $c(p, a, b)$, which allows us to quantify over predicates using ASP variables (e.g., $c(P, X, Y)$). We represent each atom in a goal (1) as a fact

$$goal(c(o_1, v_1, \dots, v_k)). \quad (21)$$

where $v_i \notin S$ are constants corresponding to existentially quantified variables V_i .

We mark each sort name $s \in S$ using a fact

$$sortname(s). \quad (22)$$

Example 5 (continued). Goal and sort names of our running example are represented as

$$\begin{aligned} &goal(c(name, m, mary)). \quad goal(c(lost, m, f)). \quad goal(c(fatherof, f, m)). \\ &goal(c(inst, s, female)). \quad goal(c(is, s, depressed)). \quad sortname(person). \\ &sortname(male). \quad sortname(female). \quad sortname(dead). \quad sortname(depressed). \end{aligned}$$

□

3.1 Rules common to all encodings

Goals are potential interesting facts, i.e., potential nodes of the proof graph:

$$pot(X) \leftarrow goal(X). \quad (23)$$

Potential interesting facts provide potentially relevant terms in the Herbrand Universe (HU).

$$hu(X) \leftarrow pot(c(_, X, _)). \quad (24)$$

$$hu(X) \leftarrow pot(c(_, _, X)). \quad (25)$$

Note that (24) and (25) assume, that all atoms in B and O have arity 2. This assumption is made only in these two rules, which can be generalized easily to arbitrary arities.

We call terms in HU that are not sort names ‘User HU’, represent them in predicate uhu , and guess a relation eq among pairs of these terms.

$$uhu(X) \leftarrow hu(X), \mathbf{not} \text{ sortname}(X). \quad (26)$$

$$\{ eq(A, B) : uhu(A), uhu(B), A \neq B \} \leftarrow. \quad (27)$$

Relation eq holds symmetric on HU, and it is a reflexive, symmetric, and transitive (equivalence) relation.

$$eq(A, A) \leftarrow hu(A). \quad (28)$$

$$\leftarrow eq(A, B), \mathbf{not} eq(B, A). \quad (29)$$

$$\leftarrow eq(A, B), eq(B, C), A \neq B, B \neq C, A \neq C, \mathbf{not} eq(A, C). \quad (30)$$

Note, that we will later create instantiations of constraint (30) in a lazy manner (on-demand), therefore we use a constraint to ensure transitivity, and not a rule.

3.2 BWD: Defining Back-chaining Proof Graph, Guess Active Parts

We next encode the maximal back-chaining proof graph from Definition 2 in ASP by: (i) deterministically defining the maximum possible potential proof graph by back-chaining from the goal and creating new constants when required; (ii) guessing which parts of the proof graph are used, i.e., which atoms are back-chained over which axioms, and which bodies of axioms must therefore be justified; (iii) factor atoms with other atoms and mark the remaining atoms as abduced.

Potential Proof Graph. Building the potential proof graph is realized by rewriting each axiom of form (2) into a deterministic definition of potential inferences from the axiom’s head atom, and defining which body atoms become part of the hypothesis due to such an inference.

We first give this rewriting as an example and then formally.

Example 6. Axiom (9) from our running example is translated into the following rules.

$$mayInferVia(r_1, c(is, X, depressed), l(Y)) \leftarrow pot(c(is, X, depressed)), Y = s(r_1, \text{“Y”}, X). \quad (31)$$

$$inferenceNeeds(c(is, X, depressed), r_1, c(importantfor, Y, X)) \leftarrow mayInferVia(r_1, c(is, X, depressed), l(Y)). \quad (32)$$

$$inferenceNeeds(c(is, X, depressed), r_1, c(is, Y, dead)) \leftarrow mayInferVia(r_1, c(is, X, depressed), l(Y)). \quad (33)$$

Here r_1 is a unique identifier for axiom (9). Rule (31) defines all possible substitutions of the axiom, including value invention via Skolemization which is represented in the last argument of $mayInferVia$ in the term $l(\dots)$. Rules (32) and (33) define which body atoms become part of the hypothesis because of back-chaining. Note that Skolemization is here realized with an uninterpreted function term $s(r_1, \text{“Y”}, X)$ that takes as arguments the unique axiom identifier r_1 , the name of the skolemized existential variable “Y” (to skolemize several variables in one axiom independently), and all variables (here only X) in the body of the axiom. \square

For each body atom that can be added to the hypothesis by an inference, $inferenceNeeds(Head, Rule, Body)$ is defined. To allow back-chaining from $Body$, we define $Body$ as potentially interesting.

$$pot(P) \leftarrow inferenceNeeds(_, _, P) \quad (34)$$

Axioms rewritten as in (31)–(33) together with (23) and (34) form a deterministic ASP program, that, given a set of goal atoms $goal(A)$, defines the union of all possible proof graphs.

This graph is finite under the assumption that the knowledge base is acyclic, i.e., that no circular inferences are possible over all axioms in the knowledge base. (For reasons of presentation we will maintain this assumption while presenting the basic encodings and eliminate the assumption in Section 4.1.)

For readability we gave the rewriting in Example 6. Formally, an axiom of form (2) is rewritten into

$$\begin{aligned} mayInferVia(a, c(q, Y_1, \dots, Y_m), l(Z_1, \dots, Z_v)) \leftarrow \\ pot(c(q, Y_1, \dots, Y_m)), Z_1 = s(a, 1, Y_1, \dots, Y_m), \dots, Z_v = s(a, v, Y_1, \dots, Y_m) \\ inferenceNeeds(c(q, Y_1, \dots, Y_m), a, c(p_i, X_1^i, \dots, X_{k_i}^i)) \leftarrow \\ mayInferVia(a, c(q, Y_1, \dots, Y_m), l(Z_1, \dots, Z_v)) \quad \text{for } i \in \{1, \dots, r\} \end{aligned} \quad (35)$$

where a is a unique identifier for that particular axiom, $Z_1, \dots, Z_v = \mathcal{X} \setminus \mathcal{Y}$ is the set of variables occurring in the body but not in the head, and the second argument of the uninterpreted function $s(\cdot)$ is a unique identifier for each skolemized variable in this axiom.

Lemma 1. *Given an abduction instance $A = (B, O, S)$, let $P_{bos}(A) = P_b \cup P_o \cup P_s$ where P_b is the rewriting of each axiom in B as (35), P_o the rewriting of O as (21), and P_s the rewriting of S as (22). Let $P_{bpt}(A) = P_{bos}(A) \cup \{(23), (34)\}$. Then $P_{bpt}(A)$ has a single answer set I that represents the union of all proof graphs of all hypotheses of A that are generated by back-chaining according to Def. 2, with nodes $\{P \mid pot(P) \in I\}$ and edges $\{(Q, P) \mid mayInferVia(R, P, L) \in I, inferenceNeeds(P, R, Q) \in I\}$.*

Representing a Hypothesis. Based on the potential proof graph defined above, we now formulate in ASP the problem of guessing a hypothesis, i.e., selecting a connected part of the potential proof graph as a solution to the abduction problem. Nodes of the proof graph are represented as $true(\cdot)$.

If an atom P is a goal, it is true.

$$true(P) \leftarrow goal(P). \quad (36)$$

If an atom P is true, it is back-chained ($infer(\cdot)$), factored, or abduced (the latter two represented as fai).

$$1 \leq \{ infer(P) ; fai(P) \} \leq 1 \leftarrow true(P). \quad (37)$$

Each atom P that is marked as inferred in the proof graph, has to be back-chained via exactly one axiom R .

$$1 \leq \{ inferVia(R, P) : mayInferVia(R, P, _) \} \leq 1 \leftarrow infer(P). \quad (38)$$

If back-chaining an atom P would add body atom Q to the hypothesis, then we define Q as true.

$$true(Q) \leftarrow inferVia(R, P), inferenceNeeds(P, R, Q). \quad (39)$$

This encoding guesses the back-chaining part of a particular proof graph and hypothesis.

Proposition 2. *Given an abduction instance $A = (B, O, S)$, let P_{gp} consist of rules (36)–(39). Then answer sets $AS(P_{bpt}(A) \cup P_{gp})$ correspond 1-1 with proof graphs G induced by hypotheses $H \in \hat{\mathcal{H}}(A)$ via back-chaining: edges $E(G)$ are represented as $inferVia(\cdot, \cdot)$, nodes $N(G)$ as $true(\cdot)$, back-chained atoms as $infer(\cdot)$, and other atoms in $fai(\cdot)$.*

3.3 Factoring

So far we merely encoded the back-chaining part of proof graphs. It remains to deal with unification, which allows us to identify those atoms in a hypothesis that incur a cost in CARD and WA because they must be abducted. For that we use rules (24)–(30) which guess an equivalence relation $eq(\cdot, \cdot)$ over the Herbrand Universe such that constants that are not sort names can be equivalent with other constants.

Lemma 3. *Given an abduction instance $A = (B, O, S)$, let P_{eq} consist of rules (24)–(30). Then $AS(P_{bpt}(A) \cup P_{eq})$ contains one answer set for each equivalence relation on the HU of $I \in AS(P_{bpt}(A))$ represented in predicate eq such that sort names are singleton equivalence classes.*

Atoms that are not back-chained are represented as $fai(P)$. These must either be unified or abducted.

3.3.1 BWD-G: Guess Factoring

This method guesses whether an atom is factored or abducted and represents for factored atoms with which other atom they have been unified. As the deterministically defined proof graph does not contain factoring between inference steps, we require factoring with inferred atoms to obtain all possible proof graphs. (We discuss and relax this restriction in Section 3.3.3.)

For an atom in H that is not inferred, we guess if it is factored or abducted.

$$1 \leq \{ factor(P) ; abduce(P) \} \leq 1 \leftarrow fai(P). \quad (40)$$

If a factored atom $A_1 = c(P, S_1, O_1)$ unifies with an inferred atom $A_2 = c(P, S_2, O_2)$ that is not below A_1 in the proof graph, then represent that A_1 is factored via A_2 .

$$\begin{aligned} factorVia(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow factor(c(P, S_1, O_1)), infer(c(P, S_2, O_2)), \\ \mathbf{not\ below}(c(P, S_2, O_2), c(P, S_1, O_1)), eq(S_1, S_2), eq(O_1, O_2). \end{aligned} \quad (41)$$

We define $below(A_1, A_2)$ as a partial order over atoms such that A_1 is below A_2 whenever inference of A_1 requires A_2 , and whenever A_1 is factored via A_2 . Intuitively, ‘below’ can be read as ‘closer to goal nodes’.

$$below(P, Q) \leftarrow inferVia(R, P), inferenceNeeds(P, R, Q). \quad (42)$$

$$below(A, C) \leftarrow below(A, B), below(B, C). \quad (43)$$

$$below(P, Q) \leftarrow factorVia(P, Q). \quad (44)$$

Note that without $\mathbf{not\ below}(\dots)$ in (41), we would obtain cyclic proof graphs where an atom justifies itself. This would affect all objective functions we study, therefore we need to eliminate such cases.

If a factored atom unifies with an abducted atom, represent that this is the case.

$$\begin{aligned} factorVia(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow \\ factor(c(P, S_1, O_1)), abduce(c(P, S_2, O_2)), eq(S_1, S_2), eq(O_1, O_2). \end{aligned} \quad (45)$$

Finally, we require that all factored atoms are unified with another atom.

$$factorOk(P) \leftarrow factorVia(P, _). \quad (46)$$

$$\leftarrow factor(P), \mathbf{not\ factorOk}(P). \quad (47)$$

We do not prove correctness of BWD-G, as encoding BWD-A is similar and has better performance.

3.3.2 BWD-AI: Abduced/Inferred Cluster Factoring

As an alternative to guessing which atoms are factored and which are abducted, we next define deterministically, that every atom that can be factored with an inferred atom must be factored with that atom, and that all remaining sets X of atoms that unified wrt. eq are factored with the (lexicographically) smallest atom in that equivalence class X of atoms, in the following called ‘cluster’.

To that end, instead of (41) we use the following rule.

$$\begin{aligned} & \text{factorViaI}(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow \text{fai}(c(P, S_1, O_1)), \text{infer}(c(P, S_2, O_2)), \\ & \quad \mathbf{not} \text{ below}(c(P, S_2, O_2), c(P, S_1, O_1)), \text{eq}(S_1, S_2), \text{eq}(O_1, O_2). \end{aligned} \quad (48)$$

We represent atoms that are factored via inferred atoms using predicate *factorI* and we represent what remains to be factored or abduced in *fa*. Moreover, we define that *factorViaI* entails *factorVia*.

$$\text{factorI}(P) \leftarrow \text{factorViaI}(P, _). \quad (49)$$

$$\text{fa}(P) \leftarrow \text{fai}(P), \mathbf{not} \text{ factorI}(P). \quad (50)$$

$$\text{factorVia}(A, B) \leftarrow \text{factorViaI}(A, B). \quad (51)$$

Next we deal with these remaining atoms: we define a partial order over atoms that unify under equivalence *eq* and factor these clusters with the (lexicographically) smallest element as follows.

$$\begin{aligned} & \text{factorCluster}(c(P, S_2, O_2), c(P, S_1, O_1)) \leftarrow \text{fa}(c(P, S_1, O_1)), \text{fa}(c(P, S_2, O_2)), \text{eq}(S_1, S_2), \\ & \quad c(P, S_1, O_1) < c(P, S_2, O_2), \text{eq}(O_1, O_2). \end{aligned} \quad (52)$$

$$\text{factorClusterAbove}(A) \leftarrow \text{factorCluster}(A, _). \quad (53)$$

$$\begin{aligned} & \text{factorVia}(A, B) \leftarrow \text{factorCluster}(A, B), \\ & \quad \mathbf{not} \text{ factorClusterAbove}(B). \end{aligned} \quad (54)$$

Note that (52) defines the partial order, (53) represents elements that are not the smallest in the cluster, and (54) maps the partial order into *factorVia* using the smallest element as the atom that all others are unified with. Finally, we define *below* using rules (42)–(44), and we define that every hypothesis atom that could not be factored is abduced.

$$\text{factor}(P) \leftarrow \text{factorVia}(P, _). \quad (55)$$

$$\text{abduce}(P) \leftarrow \text{fa}(P), \mathbf{not} \text{ factor}(P). \quad (56)$$

Note that this encoding represents a restricted set of solutions compared to the previous one, however because of the symmetry of unification, by fixing the direction of factoring we merely canonicalize solutions and cannot lose optimal solutions.

Example 7 (continued). To illustrate, that the factoring method of BWD-AI does not eliminate optimal solutions, consider the arc between the abduced atom *fatherof*(*f*, *m*) and the factored atom *fatherof*(*f*, *n*₂) in Figure 1: if we reverse this arc, then the former becomes factored and the latter abduced. The number of abduced atoms stays the same, therefore CARD is not affected; reachability stays the same so COH is not affected; and costs propagate the other direction: *fatherof*(*f*, *m*)^{48\$, 100\$} obtains cost via (6) and initial goal cost; *fatherof*(*f*, *n*₂)^{48\$, 57\$} obtains cost from factoring and via (4), so WA remains unchanged.

Similarly, the factoring edge between *name*(*m*, *mary*) and *name*(*s*, *mary*) could be reversed: then *name*(*m*, *mary*)^{100\$} obtains only initial goal cost and we would abduce *name*(*s*, *mary*)^{100\$, 120\$} and minimum cost of WA remains 100\$ for these nodes, moreover the number of abduced atoms (CARD) and reachability (COH) stays the same.

For reversing factoring arcs between two non-abduced atoms, consider reversing the arc between *importantfor*(*f*, *m*) and *importantfor*(*n*₁, *m*) in Figure 1: reversing that arc makes the former atom factored (with costs 40\$) and the latter abduced (with costs 40\$ and 60\$), and back-chaining using (6) can be done from *importantfor*(*n*₁, *m*) instead, which yields the abduced atom *fatherof*(*n*₁, *m*)^{48\$, 57\$, 100\$} instead of *fatherof*(*f*, *m*). Note that *fatherof*(*f*, *n*₂) can still be factored with that new abduced atom as *f* = *n*₁. □

As we are solving an optimization problem, dealing with a subset of solutions that has been canonicalized (by enforcing an arbitrary order on factoring) can be an advantage for efficiency, as certain symmetric solutions are automatically excluded.

We do not prove correctness of this factoring variant, as the following variant has better performance.

3.3.3 BWD-A: Abduced Cluster Factoring

Finally, we apply an even stronger canonicalization to the proof graph: we assume factoring only happens with abduced atoms. We first show, that every proof graph, that contains factoring with an inferred atom, can be transformed into a proof graph where inferences between factored atom and abduced atoms is duplicated and all factoring is done with abduced atoms.

Example 8 (continued). The proof graph in Figure 1 can be transformed into a graph where factoring happens only with abduced atoms: instead of factoring atom $importantfor(n_1, m)^{60\$}$ with atom $importantfor(f, m)^{40\$, 60\$}$, we can back-chain from the former over axiom (6) which yields the atom $fatherof(n_1, m)^{72\$}$ in the graph. This atom can now be factored with $fatherof(f, m)$ at the top, which obtains the set $\{48\$, 72\$, 100\}$ of costs and therefore keeps the same minimum cost. \square

Importantly, the metrics we consider do not change when we perform this canonicalization.

Proposition 4. *Given a proof graph G of a hypothesis $H \in \hat{\mathcal{H}}(A)$ of an abduction instance A , there is a proof graph G' and hypothesis H' of same cost wrt. CARD, COH, WA where factoring is only performed with atoms in $A(G')$.*

Proof. We show how to push factoring edges closer to abduced atoms without changing the objective function value. As the graph is acyclic, this operation can be continued until we only unify with abduced atoms.

Given an atom $P \in H$ and an edge $(P, Q) \in E(G)$ of factoring Q with P using substitution θ , and $P \notin A(G)$. Then either (i) P is factored with Q' , i.e., $(Q', P) \in E(G)$, or (ii) P is back-chained on axiom r with k body atoms, i.e., $(P'_i, P) \in E(G)$ for $1 \leq i \leq k$. In case (i) we can factor Q with Q' instead of with P , which pushes factoring one edge closer to abduced atoms. This does not affect CARD as $A(G) = A(G')$, reachability stays the same so COH is not affected, and the minimal cost of P and Q is propagated to Q' as before the change, so WA is not affected. In case (ii) we can back-chain from Q with axiom r , creating edges $(\theta^{-1}(P'_i), Q)$, adding nodes $\theta^{-1}(P'_i)$ to G' if they do not exist (implicitly they are already contained in H due to equivalences) and adding factoring edges from $(P'_i, \theta^{-1}(P'_i))$ to G' . This reduces the number of inference edges between factored and abduced atoms in the graph by 1. Similar as before, reachability and number of abduced atoms stays constant. For WA, cost might increase for $\theta^{-1}(P'_i)$ but stays the same for Q' . Therefore, we do not lose optimal solutions by restricting factoring to abduced atoms. \square

By a similar argument, the order of factoring in such a proof graph does not matter, so we can also canonicalize factoring direction.

Proposition 5. *Given a proof graph G of a hypothesis H where factoring is only performed with atoms in $A(G)$, an abduced atom $P \in A(G)$, and atoms Q_1, \dots, Q_k that are factored with P . Then we can swap an arbitrary Q_i , $1 \leq i \leq k$, with P , factor all other Q_j , $j \neq i$ with Q_i , factor P with Q_i , and all objective functions stay the same.*

Proof. As all Q_1, \dots, Q_k, P unify, we can arbitrarily choose one of them as representative and factor all others with it. This does not increase the number of abduced atoms in CARD, this does not affect reachability in COH, and costs of all atoms are propagated to the chosen representative, and the minimum cost in WA stays the same. \square

To realize this canonicalization, we use rules (52)–(56) and add the following rule, such that every atom that is not back-chained is factored with abduced atoms if possible, otherwise abduced.

$$fa(P) \leftarrow fai(P). \quad (57)$$

Note that this encoding does not require any guesses to determine what is factored and what is abduced, moreover there is no need for the definition of *below*.

Proposition 6. *Given an abduction instance $A = (B, O, S)$, let $P_{\text{BWD-A}}(A) = P_{\text{bpt}}(A) \cup P_{\text{gp}} \cup P_{\text{eq}} \cup P_c$ where $P_c = \{(52)–(56), (57)\}$, then answer sets $AS(P_{\text{BWD-A}}(A))$ of $P_{\text{BWD-A}}(A)$ are in 1-1 correspondence with proof graphs G and hypotheses $H \in \hat{\mathcal{H}}(A)$ where factoring is performed only with $A(G)$ and only with lexicographically smaller atoms.*

3.4 FWD-A: Guess Abduced Atoms, Forward Inference, Check Goal

The previous encodings are based on explicitly representing proof graphs. We next describe an encoding that is more in the spirit of the generate-define-test paradigm of Answer Set Programming [Lif02]: we again propagate potentially interesting truth values, however we do not keep track of the inferences. We guess which of these potentially interesting truth values is abduced or factored, and use another rewriting of the axioms to reproduce their semantics, i.e., we define that the head of an axiom is true if all its bodies are true. Finally we check that all goals become true. For example axiom (9) is translated into

$$\text{infer}(c(is, X, depressed)) \leftarrow \text{true}(c(\text{importantfor}, Y, X)), \text{true}(c(is, Y, dead)). \quad (58)$$

$$\text{pot}(c(\text{importantfor}, Y, X)) \leftarrow \text{pot}(c(is, X, depressed)), Y = s(r_1, y, X). \quad (59)$$

$$\text{pot}(c(is, Y, dead)) \leftarrow \text{pot}(c(is, X, depressed)), Y = s(r_1, y, X). \quad (60)$$

where r_1 again is a unique identifier for axiom (9).

We guess potential atoms as factored or abduced, define truth from factoring, abducing, and inference, and require that goals are true.

$$\{ \text{fai}(X) : \text{pot}(X) \} \leftarrow . \quad (61)$$

$$\text{true}(X) \leftarrow \text{fai}(X). \quad (62)$$

$$\text{true}(X) \leftarrow \text{infer}(X). \quad (63)$$

$$\leftarrow \text{goal}(A), \text{not true}(A). \quad (64)$$

This realizes abduction in the classical generate-define-test way. The only thing missing is factoring to determine which atoms actually need to be abduced. For that we add the following rule to define which atoms are factored or abduced.

$$\text{fa}(X) \leftarrow \text{fai}(X), \text{not infer}(X). \quad (65)$$

We complete the encoding by cluster factoring rules (52)–(56) and common rules (23)–(30).

Because we do not have an explicit representation of the proof tree, only the CARD metric is applicable. Moreover, we cannot factor with inferred atoms as there is no way to rule out circular inference, hence we only study the most restricted factoring variant.

For readability we gave the rewriting as an example. Formally, an axiom of form (2) is rewritten into

$$\begin{aligned} \text{infer}(c(q, Y_1, \dots, Y_m)) &\leftarrow \text{true}(c(p_1, X_1^1, \dots, X_{k_1}^1)), \dots, \text{true}(c(p_r, X_1^r, \dots, X_{k_r}^r)). \\ \text{pot}(c(p_i, X_1^i, \dots, X_{k_i}^i)) &\leftarrow Z_1 = s(a, 1, Y_1, \dots, Y_m), \dots, Z_v = s(a, v, Y_1, \dots, Y_m), \\ &\text{pot}(c(q, Y_1, \dots, Y_m)). \quad \text{for } i \in \{1, \dots, r\} \end{aligned}$$

where a is a unique identifier for that particular axiom and $Z_1, \dots, Z_v = \mathcal{X} \setminus \mathcal{Y}$. Note that this means that the resulting rules will all be safe.

We do not prove correctness of this encoding as it is only applicable to CARD and does not have good performance compared with other encodings.

3.5 Encodings for Preferred Solutions

We next describe program modules that realize objective functions when we add them to the previously given encodings.

Cardinality Minimality. For realizing objective CARD we use the following weak constraint.

$$\leftarrow \text{abduce}(P). \quad [1@1, P] \quad (66)$$

Coherence Metric. For COH we represent which nodes are reachable from which goal node.

$$\text{reach}(P, P) \leftarrow \text{goal}(P). \quad (67)$$

$$\text{reach}(Q, \text{From}) \leftarrow \text{reach}(P, \text{From}), \text{inferVia}(R, P), \text{inferenceNeeds}(P, R, Q). \quad (68)$$

$$\text{reach}(Q, \text{From}) \leftarrow \text{reach}(P, \text{From}), \text{factorVia}(P, Q). \quad (69)$$

Table 1: Comparison of key aspects of ASP encodings BWD-G, BWD-AI, BWD-A, and FWD-A.

	BWD-G	BWD-AI	BWD-A	FWD-A
Objective functions	supports CARD, COH, and WA			supports only CARD
Atom justification (reasoning method)	backward inference from goals via guessed inferences in proof graph			forward inference from abduced atoms
Proof graph (deterministic)	represent potentially true atoms and inferences			represent potentially true atoms
Nondeterministic guess (atom justification)	inferred vs. abduced vs. factored	inferred vs. abduced or factored		false vs. abduced or factored
Factoring canonicalization	none	factor first with inferred, then with abduced atoms	deterministic factoring with abduced atoms	
Factoring acyclicity	check required (predicate <i>below</i>)		check not required (always acyclic)	

We represent pairs of distinct goal atoms that have a common reachable atom, and we create a weak constraint that incurs a cost corresponding to pairs of goal atoms without a common reachable atom.

$$reachFromBoth(P, Q) \leftarrow goal(P), goal(Q), P < Q, reach(N, P), reach(N, Q). \quad (70)$$

$$\Leftarrow goal(P), goal(Q), P < Q, \mathbf{not} \ reachFromBoth(P, Q). \quad [1@1, P, Q] \quad (71)$$

Weighted Abduction. For realizing WA we represent potential costs as an integers. We seed cost with \$100 for goal atom assumption cost.

$$pcost(P, 100) \leftarrow goal(P). \quad (72)$$

As common practice for applying WA to ACCEL, we realize axiom costs such that body cost factors sum up to 1.2. For that we require for each axiom R a fact $numberOfBodyAtoms(R, N)$ to be defined such that N is the number of body atoms of R . We also require a minimum cost of 1 which prevents spurious very deep proof trees from being optimal due to cost 0 at abduced atoms.

$$pcost(Q, Mc) \leftarrow inferVia(R, P), inferenceNeeds(P, R, Q), \quad (73)$$

$$Mc = \#max \{ (C * 6/5)/N ; 1 \}, pcost(P, C), numberOfBodyAtoms(R, N).$$

These computations are handled during instantiation. They can be generalized to assumption weights that are individually given for each axiom as facts, without causing a change the rest of the encoding.

We propagate cost across factoring edges, define cost to be the minimum cost found at all abduced atoms, and minimize that cost using a weak constraint.

$$pcost(Q, C) \leftarrow factorVia(P, Q), pcost(P, C). \quad (74)$$

$$cost(P, C) \leftarrow abduce(P), C = \#min \{ Ic : pcost(P, Ic) \}. \quad (75)$$

$$\Leftarrow cost(P, C). \quad [C@1, P] \quad (76)$$

Proposition 7. Let $P_{CARD} = \{(66)\}$, $P_{COH} = \{(67)-(71)\}$, and $P_{WA} = \{(72)-(76)\}$. Then the cost of answer sets $I \in AS(P_{BWD-A}(A) \cup P_{CARD})$, $I \in AS(P_{BWD-A}(A) \cup P_{COH})$, and $I \in AS(P_{BWD-A}(A) \cup P_{WA})$ is the objective function $CARD(G)$, $COH(G)$, and $WA(G)$, respectively, of the proof graph G represented in I , where for WA costs are rounded to integers and at least 1.

3.6 Summary

Table 1 gives an overview of our encodings. Encodings in the BWD family deterministically represent the maximal potential proof graph and all its inferences, while FWD-A represents only atoms in the hypothesis. Justification of atoms in the proof graph is ensured from goals to abducibles (backward) in the BWD encodings, and from abduced atoms to goals (forward) in FWD-A. The type of justification of

goals and atoms in the hypothesis is nondeterministically guessed as one of three classes by BWD-Gand one of two classes by BWD-AI and BWD-A. FWD-A guesses truth of abduced or factored atoms (two classes). Factoring is canonicalized to various extent, and acyclicity of factoring is implicitly ensured in BWD-A and FWD-A but is encoded explicitly in other encodings.

4 Extensions

The ASP encodings given so far are realized in pure ASP-Core-2 syntax and do not require additional features specific to particular solver tools. However, these encodings have two drawbacks: they can represent only acyclic theories, and they have performance issues related to the size of instantiation of the transitivity constraint for *eq*. We next formalize two extensions of these encodings and describe their computational realization.

In Section 4.1 we introduce Flexible Value Invention for fine-grained control of Skolemization, which makes our encodings applicable to cyclic theories. In Section 4.2 we show how to replace certain constraints in our encodings with lazy variants to reduce grounding size and potentially improve evaluation performance. Both extensions are described formally in the HEX formalism. Section 4.3 discusses how we realized these extensions using the Python library of Clingo.

4.1 Flexible Value Invention for Cyclic Knowledge Bases

The encodings in Section 3 assume that the knowledge base is acyclic, which ensures finite proof trees and a finite instantiation of our ASP encodings in the presence of Skolemization with Uninterpreted Function terms as done in our axiom rewritings.

Example 9. As an example of a cyclic knowledge base consider the following two axioms

$$p(A, b) \Leftarrow q(A, C), t(C, b). \quad (r_1)$$

$$t(D, b) \Leftarrow p(D, b). \quad (r_2)$$

where a goal of $p(a, b)$ yields the following infinite backward chaining instantiation of axioms in the proof tree

$$\begin{aligned} p(a, b) &\Leftarrow q(a, s(r_1, \text{“}C\text{”}, a)), t(s(r_1, \text{“}C\text{”}, a), b). && \text{(via } r_1) \\ t(s(r_1, \text{“}C\text{”}, a), b) &\Leftarrow p(s(r_1, \text{“}C\text{”}, a), b). && \text{(via } r_2) \\ p(s(r_1, \text{“}C\text{”}, a), b) &\Leftarrow q(s(r_1, \text{“}C\text{”}, a), s(r_1, \text{“}C\text{”}, s(r_1, \text{“}C\text{”}, a))), \\ &\quad t(s(r_1, \text{“}C\text{”}, s(r_1, \text{“}C\text{”}, a)), b). && \text{(via } r_1) \\ t(s(r_1, \text{“}C\text{”}, s(r_1, \text{“}C\text{”}, a)), b) &\Leftarrow p(s(r_1, \text{“}C\text{”}, s(r_1, \text{“}C\text{”}, a)), b). && \text{(via } r_2) \\ &\vdots \end{aligned}$$

where C is first skolemized as $s(r_1, \text{“}C\text{”}, a)$ (see (31)) but then used again to back-chain over the first axiom which leads to another Skolemization. This leads to undecidability as we cannot know when we have generated ‘enough’ distinct variables to find the optimal solution. \square

The ACCEL benchmark is described as being acyclic [NM92] however it contains one cyclic axiom and this contains a comment that suggests that the respective axiom has been added after publication of [NM92]. To evaluate ACCEL, or any cyclic theory with our encodings, we therefore need to exclude axioms to break cycles, or infinite instantiations will occur. However, in knowledge representation, knowledge is sometimes naturally expressed in cyclic axioms, and we would like to handle such knowledge bases. In particular the cyclic axioms in ACCEL are required to obtain correct solutions for some instances, so we do not want to dismiss such axioms.

We next use external atoms instead of uninterpreted function symbols to achieve a *Flexible Value Invention* where we can control when to block value invention. By blocking certain value inventions, we ensure a finite instantiation of our encoding which thereby allows computation of optimal solutions. If we do *not* use cyclic axioms and limit value invention, we obtain a subset of all acyclic proof graphs and a sound approximation of the optimal solution. If we use cyclic axioms, they extend the proof graph in

ways that are impossible with acyclic axioms. The variations of Flexible Value Invention (shown in the following) permit usage of cyclic axioms and allow for controlling the trade-off between instantiation size and distance from the optimal solution.

For Flexible Value Invention, we first outsource value invention into an external atom $\&skolem$. Formally, instead of skolemizing variables $Z_v \in \mathcal{X} \setminus \mathcal{Y}$ in the rewriting (35) as

$$Z_v = s(a, v, Y_1, \dots, Y_m) \quad (77)$$

where a is the axiom identifier, v is the variable index, and Y_1, \dots, Y_m are head variables, we use

$$\&skolem[a, v, Y_1, \dots, Y_m](Z_v). \quad (78)$$

Example 10. Instead of (31) in the BWD encodings (Example 6), we rewrite into the rule

$$\begin{aligned} \text{mayInferVia}(r_1, c(is, X, depressed), l(Y)) \leftarrow \\ \text{pot}(c(is, X, depressed)), \&skolem[r_1, \text{“Y”}, X](Y) \end{aligned} \quad (79)$$

and instead of (59) and (60) in the FWD-A encoding we rewrite the body elements of the axiom into

$$\begin{aligned} \text{pot}(c(\text{importantfor}, Y, X)) \leftarrow \text{pot}(c(is, X, depressed)), \&skolem[r_1, \text{“Y”}, X](Y). \\ \text{pot}(c(is, Y, dead)) \leftarrow \text{pot}(c(is, X, depressed)), \&skolem[r_1, \text{“Y”}, X](Y). \end{aligned}$$

□

This way we outsource Skolemization, i.e., building a new unique constant term Z_v from terms a, v, Y_1, \dots, Y_m — or the decision not to build such a term — to an external computation. We next realize several Skolemization methods that limit value invention in different ways.

The original Skolemization with uninterpreted functions can be emulated by defining $f_{\&skolem}$ as

$$f_{\&sk^\infty}(I, R, V, Y_1, \dots, Y_m, Z) = 1 \text{ iff } Z = s(R, V, Y_1, \dots, Y_m).$$

This shows that we can still express the original Skolemization (without guaranteeing decidability).

Example 11 (continued). The external atom $\&sk^\infty[r_1, \text{“Y”}, X](Y)$ is true iff $Y = s(r_1, \text{“Y”}, X)$. This means that instantiating (31), which contains an uninterpreted function, and instantiating (79), which contains an external computation, will create the same ground Skolem terms. □

A simple way for ensuring termination is the following function.

$$f_{\&sk^{P^1}}(I, R, V, Y_1, \dots, Y_m, Z) = 1 \text{ iff } \begin{cases} Z = s(R, V, Y_1, \dots, Y_m) \\ \text{and no } Y_i, 1 \leq i \leq m, \text{ is of the form } s(\cdot, \cdot, \dots). \end{cases} \quad (80)$$

This prevents value invention if any of the terms Y_1, \dots, Y_m is an invented value, which is a very restrictive criterion: it blocks all value invention where at least one parent is an invented value.

Example 12 (continued). The external atom $\&sk^{P^1}[r_1, \text{“Y”}, X](Y)$ is true if $Y = s(r_1, \text{“Y”}, X)$ and X is not a term of form $s(\cdot, \cdot, \cdot)$. Instantiating rule (79) with $X = m$ then yields a single external atom $\&sk^{P^1}[r_1, \text{“Y”}, m](s(r_1, \text{“Y”}, m))$ which evaluates to true. Therefore, the rule head is instantiated as $\text{mayInferVia}(r_1, c(is, m, depressed), l(s(r_1, \text{“Y”}, m)))$. Assume that we have an additional axiom which contains $is(X, depressed)$ in the head and $is(X, dead)$ or $\text{importantfor}(X, Z)$ in the body. Such an axiom allows cyclic back-chaining over $is(X, depressed)$, which yields another instantiation of (79) with body literals $\text{pot}(c(is, s(r_1, \text{“Y”}, m), depressed))$ and $\&sk^{P^1}[r_1, \text{“Y”}, s(r_1, \text{“Y”}, m)](Y)$. In this case the external atom will not be true for any ground term Y : it blocks Skolemization as the input term is already a Skolem term. Without this blocking (e.g., with $\&sk^\infty$) we would obtain an infinite instantiation. □

We can extend $f_{\&sk^{P^1}}$ to block value invention only if some grandparent is an invented value.

$$f_{\&sk^{P^2}}(I, R, V, Y_1, \dots, Y_m, Z) = 1 \text{ iff } \begin{cases} Z = s(R, V, Y_1, \dots, Y_m) \\ \text{and no } Y_i, 1 \leq i \leq m, \text{ is of the form } s(\cdot, \cdot, U_1, \dots, U_k) \\ \text{with some } U_j, 1 \leq j \leq k, \text{ of the form } s(\cdot, \cdot, \dots). \end{cases} \quad (81)$$

This can be further generalized to $f_{\&sk^{P^i}}$, where $i \in \{1, 2, \dots\}$, indicates that in the $i-1$ -th nesting level of terms Y_i , $1 \leq i \leq m$, terms must not be invented values. P^1 corresponds to the method used in Henry-n700 for achieving termination (Naoya Inoue 2015, personal communication).

External oracle functions $f_{\&sk^{P^i}}$, $1 \leq i$, guarantee finite instantiation of cyclic abduction problems.

Proposition 8. *Given a cyclic or acyclic abduction instance $A = (B, O, S)$, let $P_{\text{BWD-A}}^{P^i}(A)$ be the program $P_{\text{BWD-A}}(A)$ after replacing all body atoms of form (77) by body atoms of form (78) where $\&skolem = \&sk^{P^i}$, then $\text{grnd}(P_{\text{BWD-A}}^{P^i}(A))$ is finite for $i \in \{1, 2, \dots\}$.*

Proof. $P_{\text{BWD-A}}^{P^i}(A)$ is finite and contains only safe rules, therefore the only source of infinite instantiation can be the instantiation of terms of unlimited depth. Except for the first rule in the axiom rewriting (35), all rules have heads with term nesting level equal or lower than in the body, hence the only rule that can generate terms of unlimited depth is the first in (35). In that rule, term $l(\dots)$ is created, but it is only used in the other rewritten rules in (35) where only the arguments of $l(\dots)$ are used, and in (38), where this term is discarded, therefore $l(\dots)$ cannot be infinitely nested. The only source of terms of infinite depth is the external atom of form $\&sk^{P^i}[a, i, Y_1, \dots, Y_m](Z_i)$ and it causes infinite instantiation only if $f_{\&sk^{P^i}}(I, R, V, Y_1, \dots, Y_m, Z)$ is true for infinitely many distinct terms Z . $f_{\&sk^{P^i}}$ is 1 only if no input Y_i , $1 \leq i \leq m$, has a subterm at nesting level $i-1$ that is of form $s(\dots)$, and R, V can be a finite number of constants from $P_{\text{BWD-A}}^{P^1}(A)$. Moreover there is a finite number of terms that can be built from constants in $P_{\text{BWD-A}}^{P^1}(A)$ with function symbol s , and not having a subterm of form $s(\dots)$ below nesting level $i-1$. Hence there is a finite number of tuples R, V, Y_1, \dots, Y_m for which $f_{\&sk^{P^i}}$ evaluates to true. As Z depends on R, V, Y_1, \dots, Y_m , $f_{\&sk^{P^i}}$ evaluates to true for a finite number of tuples R, V, Y_1, \dots, Y_m, Z and instantiation is finite with respect to a given program $P_{\text{BWD-A}}^{P^i}(A)$. \square

Limiting value invention this way is not the only way: nontermination of value invention always involves a certain rule and a certain existential variable of that rule being instantiated over and over again in a cycle. Therefore, we next formulate an external Skolemization oracle that blocks Skolemization only if a child term was generated for the same rule and variable.

$$f_{\&sk^{G^1}}(I, R, V, Y_1, \dots, Y_m, Z) = 1 \text{ iff } \begin{cases} Z = s(R, V, Y_1, \dots, Y_m) \text{ and} \\ \text{no } Y_i, 1 \leq i \leq m, \text{ has a sub-term of form } s(R, V, \dots). \end{cases} \quad (82)$$

This function also ensures finite instantiation.

Proposition 9. *Given a cyclic or acyclic abduction instance $A = (B, O, S)$, let $P_{\text{BWD-A}}^{G^1}(A)$ be the program $P_{\text{BWD-A}}(A)$ after replacing all body atoms of form (77) by body atoms of form (78) where $\&skolem = \&sk^{G^1}$, then $\text{grnd}(P_{\text{BWD-A}}^{G^1}(A))$ is finite for $i \in \{1, 2, \dots\}$.*

Proof. As in the proof of Proposition 8, the only reason for infinite instantiation can be the external atom. Assume towards a contradiction, that the instantiation is infinite. Then $f_{\&sk^{G^1}}(I, R, V, Y_1, \dots, Y_m, Z)$ must be true for an infinite number of terms Z . As we have a finite number of constants in $P_{\text{BWD-A}}^{G^1}(A)$ and Z is instantiated using this set of constants and function symbols of form $s(\dots)$ with finite arity, for an infinite number of terms we require infinite nesting depth of terms of form $s(\dots)$. As the set of possible tuples (R, V) used as inputs of $f_{\&sk^{G^1}}$ is finite, we must repeat (R, V) in some subterms of Z to reach an infinite amount of them. However $f_{\&sk^{G^1}}$ is false for such terms, contradiction. \square

As before, we can further generalize to $f_{\&sk^{G^i}}$ where $i \in \{1, 2, \dots\}$ indicates that a Skolem term may contain at most i layers of sub-terms from the same rule and variable.

4.2 On-Demand Constraints

In the set of rules common to all encodings (Section 3.1), we represented transitivity of the equivalence relation eq as a constraint (30) instead of using the more commonly used rule

$$eq(A, C) \leftarrow eq(A, B), eq(B, C), A \neq B, B \neq C, A \neq C.$$

The formulation as a constraint allows us to eliminate (30) from the encoding and lazily create only those instances of (30) that are violated during the search for the optimal solution. Such a lazy instantiation

is easy for constraints but not supported for rules in current solvers, as adding a new rule changes the solver representation (usually the Clark completion) of all rules with the same head in the program.

Formally we represent lazy constraints in HEX (cf. [EFI⁺16, 2.3.1]) by replacing (30) with a constraint

$$\leftarrow \mathbf{not} \ \&transitive[eq](). \quad (83)$$

where the external computation oracle is defined as follows:

$$f_{\&transitive}(I, p) = 1 \text{ iff the relation } \{(A, B) \mid p(A, B) \in I\} \text{ is transitive.} \quad (84)$$

Moreover, if we find a transitivity violation, i.e., a new triple (A, B, C) such that $\{p(A, B), p(B, C)\} \subseteq I$ and $p(A, C) \notin I$, we add a new nogood into the solver that prevents $p(A, B) \wedge p(B, C) \wedge \neg p(A, C)$ for future answer set candidates.

Similarly, we can ensure that *below* is a partial order, by replacing (43) with a guess of the relevant part of the extension of predicate *below*³ as follows

$$\{ \textit{below}(Q, P) \} \leftarrow \textit{fai}(P), \textit{infer}(Q).$$

and require acyclicity of *below* using a constraint of form (83) with external atom $\&acyclic[\textit{below}]()$, an oracle function that is true iff relation $\{(A, B) \mid p(A, B) \in I\}$ is acyclic, and nogood generation for all basic cycles in that relation.

4.3 Implementation

We formulated Flexible Value Invention and On-Demand Constraints using HEX as a formal framework. In preliminary experiments we identified a performance problem in the dlhex solver that was not possible to fix easily, therefore we decided to realize the extensions using the Python libraries of Gringo and Clasp [GKKS14]. This posed additional challenges that we discuss next.

Flexible Skolemization. Flexible Skolemization can be realized purely during instantiation by replacing external atoms of the form $\&skolem[a, i, Y_1, \dots, Y_m](Z_i)$ by the expression $Z_i = @skolem(a, i, Y_1, \dots, Y_m)$ and implementing a Python function `skolem` that generates constants according to the various semantic definitions for limited value invention that are described in Section 4.1.

Note that we can only handle this kind of external atoms in grounding because the value of the oracle function $f_{\&skolem}$ does not depend on the interpretation I .

On-demand Constraints. Different from Flexible Skolemization, on-demand constraints are handled during solving. For that, Clasp provides an interface for registering a callback function which receives answer set candidates. In that callback we can add nogoods to the solver. However, answer set enumeration modes of the current Python API of Clasp do not work well together with this idea: either we enumerate models of increasing quality, or we enumerate models without using the objective function. In the former case an on-demand constraint can invalidate the first model of optimal quality, which causes no further answer set candidates to be found, (there are no better ones). The latter case is a blind search for better solutions which is prohibitively slow.

To realize on-demand constraints with reasonable efficiency, we created algorithm FINDOPTIMAL-MODEL which first finds an optimistic bound for the objective function and then backtracks to worse bounds using on-demand constraints. This algorithm is of interest only until the Clasp API supports changing enumeration mode or objective bound from within the callback, hence we show details only in the appendix.

Global Constraints of ACCEL. A final implementation aspect is the realization of global constraints of the ACCEL benchmark.

³I.e., the part used in (41) and (48).

Assumption nogoods and unique-slot constraints can be represented uniformly for all encodings in ASP constraints. We here show the encoding strategy by means of examples. Assumption nogoods as exemplified in (19) are encoded in ASP as

$$\leftarrow \text{abduce}(c(\text{go_step}, S, G_1)), \text{abduce}(c(\text{goer}, G_2, P)), \text{eq}(G_1, G_2)$$

where we take into account term equivalence.

Unique slot axioms as exemplified in (20) are encoded in ASP as

$$\leftarrow \text{true}(c(\text{goer}, G_1, P_1)), \text{true}(c(\text{goer}, G_2, P_2)), \text{eq}(G_1, G_2), P_1 < P_2, \mathbf{not} \text{eq}(P_1, P_2)$$

where we take into account term equivalence both for the entity that must be the same to violate the constraint (G_1, G_2) , and for the entity that is enforced to be unique, i.e., must not be the same to violate the constraint (P_1, P_2) . Note that condition $P_1 < P_2$ achieves symmetry breaking during instantiation.

5 Experimental Evaluation

We evaluated the above encodings, on-demand constraints, and flexible value invention using the ACCEL benchmark described in Section 2.2. The encodings and instances we used in experiments are available online.⁴ The benchmarks were performed on a computer with 48 GB RAM and two Intel E5-2630 CPUs (total 16 cores) using Ubuntu 14.04. As solvers we used the Python API of Clingo⁵ 4.5.4⁶ [GKKS14] to implement on-demand constraints and flexible Skolemization as described in Section 4.3, and we tested pure ASP encodings also with Gringo⁵ 4.5.4⁶ [GKKS11] as grounder and both solvers Clasp⁵ 3.1.4 [GKK⁺15] and Wasp⁷ version f9d436 [ADLR15]. We also make experiments to compare with state-of-the-art approaches Henry-n700⁸ version 4b0d900 [III13] and its successor Phillip⁹ version 5612b13 [YII⁺15]. Each run was limited to 5 minutes and 5 GB RAM, HTCondor was used as a job scheduling system, each run was repeated 5 times and no more than 8 jobs were running simultaneously. For Clasp we used the setting `--configuration=crafty` which turned out to be superior to all other preset configurations of Clasp. For Wasp we used the default configuration (core-based OLL algorithm) which performs equal or better compared with other settings (note in particular, that configurations `basic`, `mgd`, and `opt` for option `--weakconstraints-algorithm` perform clearly worse than the default).

In the following tables, columns Opt, To, and Mo give the number of instances for which an optimal solution was found, the timeout was reached, and the memory limit was exceeded, respectively. These numbers are summed over instances and averaged over runs. Columns T and M show time (seconds) and memory (MB) requirement, averaged over instances and runs.

The remaining columns show detailed diagnostics of the solver and objective function, where provided by the respective tool. T_{grd} and T_{slv} give grounding and solving time, respectively, as reported by running first Gringo and then Clasp or Wasp. In experiments with the Python implementation, T_{grd} includes solver preprocessing which cannot be separated from grounding in Clasp API, and T_{slv} contains pure search time. Further metrics are only available if an optimal solution could be found and proved as optimal: Obj shows the objective function, $|Odc|$ the number of on-demand constraints, $|Sk|$ the number of created Skolem constants, $|Chc|/|Conf|$ the number of choices and conflicts encountered, and $|Ru|$ the number of rules in the program. To permit a meaningful comparison of these values, we average them over the 17 easiest instances. We only show averages if all runs found the optimum solution, otherwise we write ‘*’. For large numbers, K and M abbreviate 10^3 and 10^6 .

Preliminary encodings in [Sch15] were able to represent acyclic theories and therefore only suitable for objective CARD. Equivalence was represented by a relation between terms and representative terms in the respective equivalence class (BACKCH) or not at all (SIMPL). We will not make numerical comparisons as the performance of BACKCH and SIMPL is significantly worse than the performance of encodings in this work in all cases (in particular memory is exhausted in more than 50% of instances with Clasp).

⁴<https://bitbucket.org/knowlp/asp-fo-abduction>

⁵<http://potassco.sourceforge.net/>

⁶Including a patch that will be contained in future versions and eliminates a bug with aggregates.

⁷<https://github.com/alviano/wasp/>

⁸<https://github.com/naoya-i/henry-n700>

⁹<https://github.com/kazeto/phillip>

Table 2: Experiments with *Clasp* (C) and *Wasp* (W) on pure ASP encodings, and the Phillip system (P).

Encoding	Solver	Opt #	To #	Mo #	T sec	M MB	Obj 1	T_{grd} sec	T_{slv} sec	$ Chc $ #	$ Conf $ #	
CARD	BWD-G	C	15	28	7	223	1962	*	*	*	*	*
	BWD-G	W	36	0	14	68	2440	9.3	5.5	10.6	10K	758K
	BWD-AI	C	35	9	6	125	1870	9.3	5.3	25.9	74K	544
	BWD-AI	W	36	0	14	64	2435	9.3	5.9	9.5	10K	143K
	BWD-A	C	43	1	6	99	1778	9.3	5.0	22.5	14K	189
	BWD-A	W	45	0	5	58	1689	9.3	5.2	6.0	11K	66
	FWD-A	C	35	10	5	129	1765	9.3	5.2	22.1	128K	7K
	FWD-A	W	44	0	6	58	1697	9.3	5.6	6.0	14K	7K
COH	BWD-G	C	37	5	8	122	1980	26.8	5.7	29.7	121K	2K
	BWD-G	W	34	2	14	81	2613	26.8	6.0	13.0	20K	1M
	BWD-AI	C	39	5	6	113	1867	26.8	6.2	26.3	82K	1K
	BWD-AI	W	35	0	15	64	2492	26.8	5.4	9.2	19K	90K
	BWD-A	C	42	3	5	102	1760	26.8	5.1	22.5	76K	334
	BWD-A	W	44	0	6	58	1696	26.8	5.3	5.9	8K	2K
WA	BWD-G	C	6	36	8	262	1977	*	*	*	*	*
	BWD-G	W	34	2	15	120	2761	482.6	5.3	53.0	953K	16M
	BWD-AI	C	31	13	6	156	1888	482.6	5.3	36.5	134K	8K
	BWD-AI	W	35	0	15	83	2530	482.6	5.4	18.5	393K	1M
	BWD-A	C	40	4	6	111	1774	482.6	5.1	24.1	138K	3K
	BWD-A	W	44	0	6	66	1703	482.6	5.1	8.1	159K	29K
	-	P	0	50	0	300	230	*	*	*	*	*

Basic ASP Encodings. Table 2 shows experimental results for encodings BWD-G, BWD-AI, BWD-A, and FWD-A for objectives CARD, COH, and WA using Gringo for grounding and *Clasp* (C) or *Wasp* (W) for solving. For WA we also compare with Phillip (P). BWD-G performs worst with respect to all metrics, it performs significantly worse than BWD-AI with *Clasp*, while it performs just a bit below BWD-AI with *Wasp*. BWD-A performs best with respect to all metrics, and for CARD, *Wasp* performs nearly the same with the FWD-A encoding.

Comparing *Clasp* and *Wasp* shows that *Wasp* is able to find the optimal solution faster than *Clasp* except in cases where *Wasp* exceeds the 5GB memory limit, which happens mainly for the encodings containing a higher amount of guesses (BWD-G, BWD-AI). Another difference is the number of choices and conflicts: *Wasp* generates fewer choices for CARD and COH, but more choices for WA, moreover *Wasp* often generates more conflicts. These differences on the same encoding can be explained by different ASP optimization algorithms: for *Clasp* the used (default) configuration BB is based on adding constraints to a relaxation of the instance, while for *Wasp* the used (default) configuration OLL is based on unsatisfiable cores (for a discussion of optimization approaches, see [ADMSR15]).

Due to its unfavorable performance, we omit results for BWD-G in the following.

The Henry-n700 solver realizes WA and on the ACCEL benchmark results of around 10 sec per instance have been reported [III13]. Phillip is the successor of Henry-n700 and adds heuristics for a partial instantiation of the most relevant portions of the proof graph [YII⁺15]. We experimented with Henry-n700 and Phillip on the original ACCEL knowledge base. Unfortunately we were not able to reproduce the results reported for Henry-n700: as shown in the table, all 50 instances timed out, moreover an ILP solution was found only for 4 instances and the costs of these solutions were more than 25% above the optimal result. As we had problems with the Gurobi license, we used the open-source ‘lpsolve’ ILP backend in experiments, however Gurobi cannot improve the situation much because most instances timed out during instantiation. From the authors of Henry-n700 and Phillip we obtained a transformed version of ACCEL that was used to produce their published results; unfortunately that rewriting is incompatible with the current version of Henry-n700 as well as Phillip, however we noticed that the rewriting makes some simplifying assumptions on ACCEL: it interprets constants in goals as sort names, i.e., distinct entities in the goal can never denote the same entity (which makes, e.g., coreference resolution impossible); moreover sort names are compiled into predicates which creates many distinct predicates of lower

Table 3: Experimental Comparison of using uninterpreted function symbols versus Python for Skolemization.

Encoding	Uninterpreted Function Skolemization								Python Skolemization						
	Opt #	To #	Mo #	T sec	M MB	T_{grd} sec	T_{slv} sec	Opt #	To #	Mo #	T sec	M MB	T_{grd} sec	T_{slv} sec	
CARD	BWD-AI	36	5	9	114	1987	29.2	1.1	37	5	8	109	1989	29.4	1.1
	BWD-A	43	0	7	89	1890	28.3	0.1	43	0	7	88	1891	28.3	0.1
	FWD-A	34	9	7	120	1892	26.0	1.0	39	4	7	100	1893	25.9	0.8
COH	BWD-AI	38	3	9	111	1998	31.3	3.7	39	2	9	102	1997	31.1	2.0
	BWD-A	42	1	7	94	1895	26.8	0.4	41	2	7	93	1896	27.4	0.6
WA	BWD-AI	29	12	9	151	2006	30.5	12.6	33	8	9	136	2004	28.9	8.6
	BWD-A	42	1	7	96	1899	26.2	1.2	42	1	7	95	1901	26.2	1.1

arity and makes instantiation easier. Also assumption constraints are not realized. For reasoning with the complete set of rules in ACCEL, our approach is significantly faster than the state-of-the-art solver Phillip.

Skolemization. Table 3 compares the encodings and objectives from Table 2 with their counterparts using Python Skolemization. For a fair comparison, we here do not compare with pure Gringo+Clasp, but we use our algorithm in Python in both cases (even for ASP encodings that do not use Python Skolemization): this explains differences to Table 2 and why we do not experiment with Wasp here. For these and previous experiments we use only acyclic axioms and do not limit Skolemization, so the only difference between the program is the shape of constant symbols: nested uninterpreted function terms of form $s(\cdot, \cdot, \cdot)$ versus constants of form p_i , $i \in \mathbf{N}$, generated with Python. Although Python Skolemization provides higher flexibility than uninterpreted function terms, there is no noticeable effect on efficiency of the 17 easiest instances, and across all 50 instances, Python Skolemization has a positive effect on efficiency of the BWD-AI and FWD-A encodings, while performance of the most efficient encoding BWD-A is unchanged.

It is not apparent why some encodings improve performance with Python Skolemization. Structurally, the instantiated program entering the solver is exactly the same in both Skolemization methods, except for Skolem constants instead of uninterpreted function terms in the symbol table. However, we additionally observed, that the order of rules created by Gringo changes between both Skolemization methods. From this we conclude that the order of rules matters, and that there is potential for optimization in solvers, moreover this suggests that Clasp is sensitive to the order of rules it consumes before solving. (We conjecture, that efficiency is not affected by the form of strings in the symbol table.)

We conclude that in addition to the flexibility of Python Skolemization we can gain efficiency.

On-Demand Constraints. Table 4 shows experimental results for comparing two methods for ensuring acyclicity of the relation *below* and transitivity of the relation *eq*. The rows with (R) use encodings as given in Section 3 while those with (O) use on-demand constraints as described in Section 4.2. For a fair comparison, all runs were performed with algorithms based on the Clasp Python API (Section 4.3).

We observe that on-demand constraints significantly reduce instantiation time and memory usage in all encodings and all objective functions. We can see the difference between instantiating the full encodings or encodings without rules (30) and (43) in column T_{grd} . We observe that instantiating transitivity and acyclicity dominates the instantiation time, and that back-chaining in ASP is fast. Therefore we did not perform experiments for creating the proof graph outside ASP (this would correspond to the architecture of Phillip, where the proof graph is created in C++ and solved in ILP).

Interestingly, for the COH objective, encoding BWD-AI outperforms BWD-A with on-demand constraints (although by a small amount, and although the instantiation time of the easiest 17 instances of BWD-A is smaller than the one of BWD-AI).

As reported in the ILP-based solvers Henry-n700 and Phillip, on-demand constraints turn out to be important for managing bigger instances in this reasoning problem: we observe increased performance and a significant reduction in memory usage.

Table 4: Managing acyclicity of *below* and transitivity of *eq* using rules (R) versus on-demand constraints (O).

	Encoding	Method	Opt #	To #	Mo #	T sec	M MB	$ Odc $ #	T_{grd} sec	T_{slw} sec	$ Chc $ #	$ Conf $ #	$ Ru $ #
CARD	BWD-AI	R	36	5	9	114	1987	0	29.2	1.1	72K	553	3M
	BWD-AI	O	49	1	0	28	196	19	1.0	0.2	42K	5K	129K
	BWD-A	R	43	0	7	89	1890	0	28.3	0.1	14K	185	3M
	BWD-A	O	50	0	0	3	79	19	0.4	0.0	18K	3K	34K
	FWD-A	R	34	9	7	120	1892	0	26.0	1.0	131K	7K	3M
	FWD-A	O	39	11	0	88	87	213	0.6	0.6	165K	14K	39K
COH	BWD-AI	R	38	3	9	111	1998	0	31.3	3.7	99K	2K	3M
	BWD-AI	O	48	2	0	21	244	1185	1.2	1.1	226K	47K	252K
	BWD-A	R	42	1	7	94	1895	0	26.8	0.4	80K	287	3M
	BWD-A	O	47	3	0	27	134	1863	0.5	0.2	147K	30K	66K
WA	BWD-AI	R	29	12	9	151	2006	0	30.5	12.6	113K	8K	3M
	BWD-AI	O	37	13	0	102	286	2663	1.4	2.4	211K	37K	257K
	BWD-A	R	42	1	7	96	1899	0	26.2	1.2	127K	3K	3M
	BWD-A	O	49	1	0	42	146	5256	0.5	0.8	253K	23K	66K

Table 5: Comparison of Skolemization limits with BWD-A encoding and objective functions COH and WA.

	Obj. Limit	Opt #	To #	Mo #	T sec	M MB	Obj 1	$ Sk $ #	T_{grd} sec	T_{slw} sec	$ Chc $ #	$ Conf $ #	$ Ru $ #
COH	P^1	50	0	0	2	37	33.9	23	0.4	0.0	4K	48	42K
	P^2	45	2	3	68	1313	26.8	121	23.3	0.5	83K	630	3M
	∞	41	2	7	93	1896	26.8	128	27.4	0.6	90K	948	3M
	G^1	33	1	16	127	2845	26.8	163	41.4	1.2	169K	917	6M
	G^2	25	0	25	157	3635	26.8	199	76.8	2.2	286K	2K	12M
WA	P^1	50	0	0	2	38	595.3	23	0.4	0.0	5K	111	44K
	P^2	43	4	3	79	1318	488.2	121	21.5	1.5	156K	2K	3M
	∞	42	1	7	95	1901	482.6	128	26.2	1.1	129K	2K	3M
	G^1	33	1	16	130	2853	448.5	163	41.7	2.5	243K	2K	6M
	G^2	24	1	25	159	3642	419.6	199	77.4	5.0	479K	3K	12M

Table 6: Experiments with Clasp (C) and Wasp (W) about realizing global constraints versus omitting them.

Encoding	Solver	Global Constraints in ASP							Global Constraints Omitted							
		Opt #	To #	Mo #	T sec	M MB	Obj 1	$ Conf $ #	Opt #	To #	Mo #	T sec	M MB	Obj 1	$ Conf $ #	
CARD	BWD-AI	C	35	9	6	125	1870	9.8	815	35	9	6	130	1870	9.8	1K
	BWD-AI	W	36	0	14	64	2435	9.8	194K	36	0	14	64	2442	9.8	188K
	BWD-A	C	43	1	6	99	1778	9.8	248	42	3	5	104	1755	9.8	560
	BWD-A	W	45	0	5	58	1689	9.8	73	45	0	5	59	1688	9.8	52
	FWD-A	C	35	10	5	129	1765	9.8	11K	34	11	5	133	1765	9.8	26K
	FWD-A	W	44	0	6	58	1697	9.8	2K	44	0	6	57	1695	9.8	2K
COH	BWD-AI	C	39	5	6	113	1867	29.4	666	40	4	6	111	1870	29.4	614
	BWD-AI	W	35	0	15	64	2492	29.4	77K	35	0	15	64	2489	29.4	124K
	BWD-A	C	42	3	5	102	1760	29.4	461	42	2	6	103	1783	29.4	369
	BWD-A	W	44	0	6	58	1696	29.4	1K	44	0	6	59	1699	29.4	563
WA	BWD-AI	C	31	13	6	156	1888	487.0	14K	29	14	7	162	1903	487.0	15K
	BWD-AI	W	35	0	15	83	2530	487.0	2M	35	1	14	82	2524	487.0	1M
	BWD-A	C	40	4	6	111	1774	487.0	4K	39	6	5	117	1757	487.0	6K
	BWD-A	W	44	0	6	66	1703	487.0	25K	44	0	6	68	1701	487.0	22K

Cyclic Theories and Limited Skolemization. Table 5 shows the results of experiments with cyclic theories and limited value invention as defined in Section 4.1. For encoding BWD-A and all objective functions, we evaluate the acyclic theory with unlimited $f_{\&sk\infty}$ Skolemization, and the theory including the cyclic axiom with parent- and rule-based Skolemization limits $f_{\&sk\alpha}$ with $\alpha \in \{P^1, P^2, G^1, G^2\}$.

We show only COH and WA with BWD-A, because CARD and other encodings show the same trends.

We observe that time (T) and memory (M) usage, number of generated Skolem constants ($|Sk|$), grounding time (T_{grad}), and size of the instantiation ($|Ru|$), are ordered $P^1 < P^2 < \infty < G^1 < G^2$ for both objective functions. Solving time (T_{slv}), choices ($|Chc|$), and conflicts ($|Conf|$), are also nearly always ordered like that.

Regarding the objective function, the ∞ method does not use the (single) cyclic axiom in ACCEL, while the other methods use that axiom. P^1 and P^2 permit a limited amount of value invention based on invented values, and allow fewer inferences than ∞ on the acyclic axioms, which results in a higher cost of the optimal solution. G_1 and G_2 block value invention only when it reaches the same rule (not other rules), therefore they allow a superset of the inferences of ∞ . This explains, why Obj of ∞ is above the one of P_2 and below the one of G_1 (recall that we display Obj only for the 17 instances where all runs found the optimum).

Regarding efficiency, one or two generations of invented values across rules (G^1, G^2) drastically increase memory exhaustion, while strictly limiting value invention (P^1) makes the problem easy to solve and impairs solution quality.

Global Constraints. Table 6 shows a comparison between realizing unique-slot constraints and assumption constraints in ASP constraints versus not considering these constraints.

Global constraints have no significant effect on efficiency.

We see no effect on the objective function when removing global constraints. This seems counter-intuitive: more constraints should intuitively increase cost of the optimal solution. It turns out that these cases are rare: we can observe an increased cost of optimal solutions if we limit Skolemization using method P^1 . We conclude that global constraints in the ACCEL benchmark have a small impact on solution quality.

Other Experiments. For combining on-demand constraints with optimization, we also investigated alternatives to algorithm FINDOPTIMALMODEL (see Section 4.3 and the appendix) where we used Clasp assumptions and Clasp externals for deactivating certain optimization criteria during some parts of the search. These alternatives perform significantly worse, moreover they are involved and require rewriting weak constraints into normal rules, therefore we decided to omit further details about these experiments.

Realizing global constraints of ACCEL also in an on-demand manner did not yield significantly different results from using the pure ASP versions, therefore we omit these results from the presentation.

We experimented with projecting answer sets to the atoms that are relevant for the objective function, which yielded a significant reduction in log file size (because we print the solution) but no significant reduction in time or memory.

6 Related Work

The idea of abduction goes back to Peirce [Pei55] and was later formalized in logic.

Abductive Logic Programming (ALP) is an extension of logic programs with abduction and integrity constraints. Kakas et al. [KKT92] discuss ALP and applications, in particular they relate Answer Set Programming and abduction. Fung et al. describe the IFF proof procedure [FK97] which is a FOL rewriting that is sound and complete for performing abduction in a fragment of ALP with only classical negation and specific safety constraints. Denecker et al. [DdS98] describe SLDNFA-resolution which is an extension of SLDNF resolution for performing abduction in ALP in the presence of negation as failure. They describe a way to ‘avoid Skolemization by variable renaming’ which is exactly what we found to increase performance in flexible Skolemization (recall that we create numbered constants p_i instead of structured terms $s(\dots)$ in Python). Kakas et al. describe the \mathcal{A} -System for evaluating ALP using an algorithm that interleaves instantiation of variables and constraint solving [KVD01]. The CIFF framework [MTS⁺09] is conceptually similar to the \mathcal{A} -System but it allows a more relaxed use of negation. The SCIFF framework [ACG⁺08] relaxes some restrictions of CIFF and provides facilities for modeling agent interactions. In [GLR⁺15], SCIFF was used to realize semantics of Datalog[±] [CGL09] which natively supports existentials in rule heads (i.e., value invention) as opposed to ASP. The focus of SCIFF is on finding abductive explanations, while our focus is to find preferred abductive explanations according to objective functions. Realizing objective functions requires modifying the SCIFF engine (Evelina Lamma 2015, personal communication) therefore we did not perform experiments comparing SCIFF with our encodings.

Implementations of ALP, have in common that they are based on evaluation strategies similar to Prolog [MTS⁺09]. In [MTS⁺09], CIFF is compared with ASP on the example of n-queens and the authors emphasize that CIFF has more power due to its partial non-ground evaluation. However, they use a non-optimized n-queens encoding for that comparison, and optimized n-queens encodings for Clingo[GKKS12] are known to yield orders of magnitude better performance than naive encodings, hence partial non-ground evaluation is not necessarily a guarantee for better performance. Different from CIFF and earlier ALP implementations, our approach instantiates one Boolean variable for each node in the potential proof graph and then searches for the best solution, while methods that create nodes on demand (such as CIFF) can completely eliminate certain nodes from instantiation, while instantiating other nodes multiple times.

The AAA (ATMS-based Abduction Algorithm) reasoner [Ng92, NM92] combines Prolog resolution with ATMS-based caching for realizing abduction. For ACCEL, AAA realizes CARD and COH metrics and enforces assumption-nogoods and unique-slot constraints in dedicated (imperative) procedures.

The Henry-n700 reasoner [II13] realizes WA by creating an ILP instance with C++ using back-chaining and then finding optimal solutions for the ILP instance. The newest version of Henry-n700 is called Phillip [YII⁺15]; this solver adds heuristics that partially instantiate the proof tree according to relatedness between predicates (although without formal proof of the correctness or worst-case approximation error). This two-step approach is similar to our approach in ASP: our encodings cause the ASP grounder to perform back-chaining in the knowledge base, and after instantiation the solver searches for optimal solutions satisfying all rules and constraints. A big performance improvement for Henry-n700 was the usage of Cutting Plane Inference [II13]. We mimic the approach with on-demand constraints in ASP, and we can observe similar improvements in instantiation size and time, however solve time increases by a larger amount for many instances, hence this approach is not sufficient for achieving a corresponding performance boost in ASP. Within the ASP community, on-demand constraints are related to methods of lazy instantiation of ASP programs, as done in the solvers ASPeRiX [LBSG15], OMiGA [DtEF⁺12], Galliwasp [MG13], and recently also in the IDP system [DDSB15]. These systems apply lazy instantiation to all rules and constraints in the program, whereas we make only certain problematic constraints lazy.

Probabilistic abduction was realized in Markov Logic [RD06] in the Alchemy system [KSR⁺10] al-

though without value invention [BHD⁺11, SM11], i.e., existential variables in rule heads are naively instantiated with all ground terms in the program. A corresponding ASP encoding for the non-probabilistic case for CARD exists [Sch15, SIMPL], however it shows prohibitively bad performance.

The termination proofs we do are related to the notion of *Liberal Safety* in HEX programs [EFKR13], however Liberal Safety requires either specific acyclicity conditions (which are absent in our encodings), or conditions on finiteness of the domain of certain attributes of the external atom (that our Skolemization atoms do not fulfill). Hence we had to prove termination without using Liberal Safety.

In the area of Automated Theorem Proving, algorithms search for finite models (or theorems, unsatisfiability proofs) in full first order logic without enforcing UNA and including native support for Skolemization (cf. [Sut09]). These algorithms focus on finding a feasible solution and do not contain support for preferences (optimization criteria). However, the main emphasis of our abduction problems is to find solutions with optimal cost (recall that our problems always have the trivial solution to abduce all input atoms). To tackle our abduction problem with such theorem provers, it would be necessary to transform the optimization problem into a decision problem and perform a search over the optimization criterion, calling the prover several times. Related to theorem proving, a hypertableaux algorithm for coreference resolution is described in [BK00]. This algorithm is inspired by weighted abduction, however it does not use preferences and relies solely on inconsistency for eliminating undesired solutions.

Computational Complexity. The complexity of abduction in *propositional* theories in the presence of the CARD objective has been analyzed in [BATJ91, EG95], and in [EGL97], the propositional case of abduction in logic programs is studied and extended to function-free logic programming abduction (Sec. 6), under the restriction that only constants from observations and knowledge base (there called ‘manifestations’ and ‘program’) are used and that the UNA holds for all terms. However, in our variant of abduction the optimal solution may use a set (of unspecified size) of constants that are not present in the input and there is potential equality among certain input constants and constants originating in value invention. Hence, existing results can be seen as lower bounds for hardness but do not directly carry over to our scenario.

In an acyclic theory, our reasoning problem is related to non-recursive negation-free Datalog theories and non-recursive logic programming with equality, which has been studied (although not with respect to abductive reasoning) in [DEGV01].

Creating the largest possible proof graph for a given goal and knowledge base can be done by reversing axioms and evaluating them with the goal; the complexity of this problem (definite not range-restricted logic program without function symbols) was shown to be PSPACE-complete [VV98, Thm. 4.1].

7 Conclusion

We have created a flexible and publicly available framework for realizing variations of cost-based FO Horn abduction represented in the declarative reasoning framework of Answer Set Programming [Lif08] that allows us (i) to modularly encode additional objective functions based on the abductive proof graph, and (ii) to add global constraints of arbitrary complexity and size. Our encodings use a modular translation of axioms into ASP rules, i.e., each axiom can be translated independent from other axioms. As preference relations we realized cardinality-minimality, coherence [NM92], and weighted abduction [HSME93, Sti89]. We evaluated our framework on the ACCEL benchmark [NM92] and found that we have significantly higher performance than state-of-the-art solver Phillip [YII⁺15] which is the successor system of Henry-n700[II13]. In our experiments, Wasp [ADLR15] solves instances faster than Clasp [GKK⁺15], however Wasp uses more memory for programs with a high degree of nondeterminism.

For realizing value invention we experimented with uninterpreted functions and with external computations providing new values to the program. Performing Skolemization with external computations provides fine-grained control for deciding when to instantiate a term and when to refuse further instantiation. This allows us to ensure and formally prove decidability when performing abduction in cyclic FO Horn knowledge bases. Fortunately, this flexibility does not impair computational efficiency.

An important topic in this research is encoding the proof graph. Usually, in ASP we are not interested in the order of inferences made in the program or in the dependencies or equivalences between atoms — those are handled transparently in the solver. However, for modeling preference functions COH and WA, which are defined on proof graphs, we must explicitly represent a proof graph, including back-chaining

and unification, in our ASP encoding. We also experiment with an alternative encoding (FWD-A) for representing objective CARD: this encoding performs abduction without representing a proof graph, it is based on forward inference, has good performance with the *Wasp* solver.

The BWD-A encoding performs best, intuitively because it makes the most strict canonicalization operations on the graph by requiring factoring to happen only with abduced atoms. Proof cost is not affected by this canonicalization, however proof graphs will contain duplicate inferences for atoms that are equivalent due to term equivalence: these atoms could be first factored and inferred later. Yet this seemingly wasteful proof graph does not diminish performance in ASP, because we anyway need to instantiate the whole potential proof graph, so all inferences trees are available even if we do permit factoring anywhere in the tree. Postponing factoring to abduced atoms has the effect that we need to handle term equivalence only for these atoms, which is an advantage for efficiency.

As ASP provides no native support for term equivalence, we encode equivalence and unification explicitly. We guess an equivalence relation and check its reflexivity, symmetry, transitivity with constraints. Explicit representation of transitivity of equivalence has been shown to be a major performance issue in *Henry-n700*, as it causes instantiation of a cubic amount of rules [II13]. This performance issue also becomes apparent in our ASP representation, and we apply the solution from *Henry-n700* to our encodings by using on-demand constraints, which we describe formally in the HEX formalism [EFI+16]. Realizing on-demand constraints in presence of optimization is nontrivial in current solvers, and we describe an algorithm based on the Python API of *Clingo*[GKKS14]. On-demand constraints significantly reduce memory usage by partially instantiating transitivity constraints of the term equivalence relation (note that the potential proof graph is still fully instantiated). This is consistent with results reported for *Henry-n700* and *Phillip* for weighted abduction [II13] with *ILP* as a solver backend and not surprising as ASP and *ILP* are related methods for solving combinatorial problems [LJN12].

Future work. The major motivation for this work was to obtain a more flexible framework where variations of objective functions and constraints on abduction can be studied. In the future we intend to perform research in this direction. Moreover we want to apply our encodings to other datasets like the one derived from the Recognizing Textual Entailment (RTE) [DGM06] challenge.

Among the encodings we experiment with, the most obvious and straightforward encoding (BWD-G) has the worst performance, and small encoding changes as well as bigger changes, that realize symmetry breaking based on a theoretical analysis of the problem, are required for achieve acceptable performance (BWD-A). Interestingly, the *Wasp* solver is able to compensate for the more nondeterministic representation: it performs similar on BWD-G and BWD-AI encodings, opposed to *Clasp* which performs significantly worse on BWD-G. We conclude that automatic program optimization and supporting tools for diagnosing performance issues are open problems in ASP and fruitful topics for future work.

Acknowledgements

We thank Naoya Inoue, Evelina Lamma, Christoph Redl, and the anonymous reviewers and RCRA workshop participants for constructive feedback about this work. We thank Mario Alviano, Carmine Dodaro, Roland Kaminski, and Benjamin Kaufmann for support regarding *Clasp*, *Gringo*, and *Wasp*.

References

- [ACG+08] Marco Alberti, Federico Chesani, Marco Gavanelli, Evelina Lamma, Paola Mello, and Paolo Torroni. Verifiable agent interaction in abductive logic programming: The SCIFF framework. *ACM Transactions on Computational Logic*, 9(4):Article No. 29, 2008.
- [ADLR15] Mario Alviano, Carmine Dodaro, Nicola Leone, and Francesco Ricca. Advances in WASP. In *International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR)*, pages 40–54, 2015.
- [ADMSR15] Mario Alviano, Carmine Dodaro, Joao Marques-Silva, and Francesco Ricca. Optimum stable model search: algorithms and implementation. *Journal of Logic and Computation*, 2015.

- [BATJ91] Tom Bylander, Dean Allemang, Michael C Tanner, and John R Josephson. The computational complexity of abduction. *Artificial Intelligence*, 49(1-3):25–60, 1991.
- [BHD⁺11] James Blythe, Jerry R Hobbs, Pedro Domingos, Rohit J Kate, and Raymond J Mooney. Implementing Weighted Abduction in Markov Logic. In *International Conference on Computational Semantics (IWCS)*, pages 55–64, 2011.
- [BK00] Peter Baumgartner and Michael Kühn. Abducing Coreference by Model Construction. *Journal of Language and Computation*, 1(2):175–190, 2000.
- [CFG⁺12] Francesco Calimeri, Wolfgang Faber, Martin Gebser, Giovambattista Ianni, Roland Kaminiski, Thomas Krennwallner, Nicola Leone, Francesco Ricca, and Torsten Schaub. ASP-Core-2 Input language format. Technical report, ASP Standardization Working Group, 2012.
- [CGL09] Andrea Cali, Georg Gottlob, and Thomas Lukasiewicz. Datalog+/-: A Unified Approach to Ontologies and Integrity Constraints. In *International Conference on Database Theory*, pages 14–30, 2009.
- [DdS98] Marc Denecker and Danny de Schreye. SLDNFA: An abductive procedure for abductive logic programs. *The Journal of Logic Programming*, 34(2):111–167, 1998.
- [DDSB15] Broes De Cat, Marc Denecker, Peter Stuckey, and Maurice Bruynooghe. Lazy Model Expansion: Interleaving Grounding with Search. *Journal of Artificial Intelligence Research*, 52:235–286, 2015.
- [DEGV01] Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3):374–425, 2001.
- [DGM06] Ido Dagan, Oren Glickman, and Bernardo Magnini. The PASCAL Recognising Textual Entailment Challenge. In *Machine Learning Challenges*, pages 177–190. Springer, 2006.
- [DtEF⁺12] Minh Dao-tran, Thomas Eiter, Michael Fink, Gerald Weidinger, and Antonius Weinzierl. OMiGA: An Open Minded Grounding On-The-Fly Answer Set Solver. In *Logics in Artificial Intelligence (JELIA)*, pages 480–483, 2012.
- [EFI⁺16] Thomas Eiter, Michael Fink, Giovambattista Ianni, Thomas Krennwallner, Christoph Redl, and Peter Schüller. A model building framework for Answer Set Programming with external computations. *Theory and Practice of Logic Programming*, 16(04):418–464, 2016.
- [EFKR13] Thomas Eiter, Michael Fink, Thomas Krennwallner, and Christoph Redl. Liberal Safety for Answer Set Programs with External Sources. In *AAAI Conference on Artificial Intelligence*, pages 267–275, 2013.
- [EG95] Thomas Eiter and Georg Gottlob. The Complexity of Logic-Based Abduction. *Journal of the ACM*, 42(1):3–42, 1995.
- [EGL97] Thomas Eiter, Georg Gottlob, and Nicola Leone. Abduction from logic programs: Semantics and complexity. *Theoretical Computer Science*, 189(1-2):129–177, 1997.
- [EIK09] Thomas Eiter, Giovambattista Ianni, and Thomas Krennwallner. Answer Set Programming: A Primer. In *Reasoning Web Summer School*, Lecture Notes in Computer Science, pages 40–110. Springer, 2009.
- [EIST06] Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits. Effective Integration of Declarative Rules with External Evaluations for Semantic-Web Reasoning. In *European Semantic Web Conference (ESWC)*, pages 273–287, 2006.
- [FK97] Tze Ho Fung and Robert Kowalski. The IFF proof procedure for abductive logic programming. *The Journal of Logic Programming*, 33(2):151–165, 1997.
- [FPL11] Wolfgang Faber, Gerald Pfeifer, and Nicola Leone. Semantics and complexity of recursive aggregates in answer set programming. *Artificial Intelligence*, 175(1):278–298, 2011.

- [GKK⁺15] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, Javier Romero, and Torsten Schaub. Progress in clasp Series 3. In *International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR)*, pages 368–383, 2015.
- [GKKS11] Martin Gebser, Roland Kaminski, Arne König, and Torsten Schaub. Advances in gringo series 3. In *International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR)*, pages 345–351, 2011.
- [GKKS12] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. *Answer Set Solving in Practice*. Morgan Claypool, 2012.
- [GKKS14] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Clingo = ASP + Control: Extended Report. Technical report, University of Potsdam, 2014.
- [GL88] Michael Gelfond and Vladimir Lifschitz. The Stable Model Semantics for Logic Programming. In *International Conference and Symposium on Logic Programming (ICLP/SLP)*, pages 1070–1080, 1988.
- [GLR⁺15] Marco Gavanelli, Evelina Lamma, Fabrizio Riguzzi, Elena Bellodi, Riccardo Zese, and Giuseppe Cota. An abductive Framework for Datalog+/- Ontologies. In *International Conference on Logic Programming (ICLP)*, *Technical Communications*, number 1433. CEUR-WS.org, 2015.
- [HSME93] Jerry R Hobbs, Mark Stickel, Paul Martin, and Douglas Edwards. Interpretation as abduction. *Artificial Intelligence*, 63(1-2):69–142, 1993.
- [II13] Naoya Inoue and Kentaro Inui. ILP-based Inference for Cost-based Abduction on First-order Predicate Logic. *Journal of Natural Language Processing*, 20(5):629–656, 2013.
- [IOIH14] Naoya Inoue, Ekaterina Ovchinnikova, Kentaro Inui, and Jerry Hobbs. Weighted Abduction for Discourse Processing Based on Integer Linear Programming. In *Plan, Activity, and Intent Recognition*, pages 33–55. Elsevier, 2014.
- [KKT92] A C Kakas, R A Kowalski, and F Toni. Abductive Logic Programming. *Journal of Logic and Computation*, 2(6):719–770, 1992.
- [KSR⁺10] S Kok, M Sumner, M Richardson, P Singla, H Poon, Lowd D, J Wang, A Nath, and P Domingos. The Alchemy system for statistical relational AI. Technical report, Department of Computer Science and Engineering, University of Washington, 2010.
- [KVD01] Antonis C. Kakas, Bert Van Nuffelen, and Marc Denecker. A-System: Problem solving through abduction. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 591–596, 2001.
- [LBSG15] Claire Lefèvre, Christopher Béatrix, Igor Stéphan, and Laurent Garcia. Asperix, a first order forward chaining approach for answer set computing. *Theory and Practice of Logic Programming*, 2015. To appear, arXiv:1503.07717.
- [Lif02] Vladimir Lifschitz. Answer set programming and plan generation. *Artificial Intelligence*, 138(1-2):39–54, 2002.
- [Lif08] Vladimir Lifschitz. What Is Answer Set Programming? In *AAAI Conference on Artificial Intelligence*, pages 1594–1597, 2008.
- [LJN12] Guohua Liu, Tomi Janhunen, and I Niemelä. Answer set programming via mixed integer programming. In *Principles of Knowledge Representation and Reasoning (KR)*, pages 32–42, 2012.
- [LT94] Vladimir Lifschitz and Hudson Turner. Splitting a Logic Program. *International Conference on Logic Programming (ICLP)*, 19(1):23–37, 1994.

- [MG13] Kyle Marple and Gopal Gupta. Galliwasp: A goal-directed answer set solver. In *Logic-Based Program Synthesis and Transformation (LoPSTR)*, pages 122–136, 2013.
- [MTS⁺09] Paolo Mancarella, Giacomo Terreni, Fariba Sadri, Francesca Toni, and Ulle Endriss. The CIFF Proof Procedure for Abductive Logic Programming with Constraints: Theory, Implementation and Experiments. *Theory and Practice of Logic Programming*, 9(6):691–750, 2009.
- [Ng92] Hwee Tou Ng. *A General Abductive System with Applications to Plan Recognition and Diagnosis*. Phd thesis, University of Texas at Austin, 1992.
- [NM90] Hwee Tou Ng and Raymond J Mooney. On the Role of Coherence in Abductive Explanation. In *National Conference on Artificial Intelligence*, pages 337–342, 1990.
- [NM92] Hwee Tou Ng and Raymond J Mooney. Abductive Plan Recognition and Diagnosis: A Comprehensive Empirical Evaluation. In *Knowledge Representation and Reasoning (KR)*, pages 499–508, 1992.
- [Pei55] C S Peirce. Abduction and Induction. In *Philosophical Writings of Peirce*, chapter 11, pages 150–156. Dover Publications, 1955.
- [RD06] Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine Learning*, 62(1-2):107–136, jan 2006.
- [Sch14] Peter Schüller. Tackling Winograd Schemas by Formalizing Relevance Theory in Knowledge Graphs. In *International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 358–367. AAAI Press, 2014.
- [Sch15] Peter Schüller. Modeling Abduction over Acyclic First-Order Logic Horn Theories in Answer Set Programming: Preliminary Experiments. In *International Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA)*, volume 1451, pages 76–90. CEUR-WS.org, 2015.
- [Sch16] Peter Schüller. Modeling Variations of First-Order Horn Abduction in Answer Set Programming. *Fundamenta Informaticae*, 149(1–2):159–207, 2016.
- [SM11] Parag Singla and Raymond J Mooney. Abductive Markov Logic for Plan Recognition. In *AAAI Conference on Artificial Intelligence*, pages 1069–1075, 2011.
- [Sti89] Mark Stickel. Rationale and methods for abductive reasoning in natural-language interpretation. In *Natural Language and Logic*, pages 233–252, 1989.
- [Sut09] Geoff Sutcliffe. The TPTP problem library and associated infrastructure: the FOF and CNF parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.
- [VV98] S. Vorobyov and A. Voronkov. Complexity of nonrecursive logic programs with complex values. *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, page 253, 1998.
- [YII⁺15] Kazeto Yamamoto, Naoya Inoue, Kentaro Inui, Yuki Arase, and Jun’ichi Tsujii. Boosting the Efficiency of First-Order Abductive Reasoning Using Pre-estimated Relatedness between Predicates. *International Journal of Machine Learning and Computing*, 5(2):114–120, 2015.

8 Appendix: Complete Encodings

We give full encodings in the following.

Complete Encoding BWD. Each axiom of form (2) is rewritten into the following set of ASP rules:

$$\begin{aligned}
& \text{mayInferVia}(a, c(q, Y_1, \dots, Y_m), l(Z_1, \dots, Z_v)) \leftarrow \\
& \quad \text{pot}(c(q, Y_1, \dots, Y_m)), Z_1 = s(a, 1, Y_1, \dots, Y_m), \dots, Z_v = s(a, v, Y_1, \dots, Y_m) \\
& \text{inferenceNeeds}(c(q, Y_1, \dots, Y_m), a, c(p_i, X_1^i, \dots, X_{k_i}^i)) \leftarrow \\
& \quad \text{mayInferVia}(a, c(q, Y_1, \dots, Y_m), l(Z_1, \dots, Z_v)) \quad \text{for } i \in \{1, \dots, r\}
\end{aligned}$$

where a is a unique identifier for that particular axiom and $Z_1, \dots, Z_v = \mathcal{X} \setminus \mathcal{Y}$.

The encoding contains the following rules.

$$\begin{aligned}
& \text{pot}(X) \leftarrow \text{goal}(X). \\
& \text{pot}(P) \leftarrow \text{inferenceNeeds}(_, _, P) \\
& \text{true}(P) \leftarrow \text{goal}(P). \\
& 1 \leq \{ \text{infer}(P) ; \text{fai}(P) \} \leq 1 \leftarrow \text{true}(P). \\
& 1 \leq \{ \text{inferVia}(R, P) : \text{mayInferVia}(R, P, _) \} \leq 1 \leftarrow \text{infer}(P). \\
& \text{true}(Q) \leftarrow \text{inferVia}(R, P), \text{inferenceNeeds}(P, R, Q). \\
& \text{hu}(X) \leftarrow \text{pot}(c(_, X, _)). \\
& \text{hu}(X) \leftarrow \text{pot}(c(_, _, X)). \\
& \text{uhu}(X) \leftarrow \text{hu}(X), \text{not } \text{sortname}(X). \\
& \{ \text{eq}(A, B) : \text{uhu}(A), \text{uhu}(B), A \neq B \} \leftarrow . \\
& \text{eq}(A, A) \leftarrow \text{hu}(A). \\
& \leftarrow \text{eq}(A, B), \text{not } \text{eq}(B, A). \\
& \leftarrow \text{eq}(A, B), \text{eq}(B, C), A \neq B, B \neq C, A \neq C, \text{not } \text{eq}(A, C).
\end{aligned}$$

Complete Factoring Encoding BWD-G.

$$\begin{aligned}
& \text{below}(P, Q) \leftarrow \text{inferVia}(R, P), \text{inferenceNeeds}(P, R, Q). \\
& \text{below}(P, Q) \leftarrow \text{factorVia}(P, Q). \\
& \text{below}(A, C) \leftarrow \text{below}(A, B), \text{below}(B, C). \\
& 1 \leq \{ \text{factor}(P) ; \text{abduce}(P) \} \leq 1 \leftarrow \text{fai}(P). \\
& \text{factorVia}(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow \text{factor}(c(P, S_1, O_1)), \text{infer}(c(P, S_2, O_2)), \text{eq}(S_1, S_2), \\
& \quad \text{eq}(O_1, O_2), \text{not } \text{below}(c(P, S_2, O_2), c(P, S_1, O_1)). \\
& \text{factorVia}(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow \text{factor}(c(P, S_1, O_1)), \text{abduce}(c(P, S_2, O_2)), \\
& \quad \text{eq}(S_1, S_2), \text{eq}(O_1, O_2). \\
& \text{factorOk}(P) \leftarrow \text{factorVia}(P, _). \\
& \leftarrow \text{factor}(P), \text{not } \text{factorOk}(P).
\end{aligned}$$

Complete Factoring Encoding BWD-AI.

$$\begin{aligned}
& \text{below}(P, Q) \leftarrow \text{inferVia}(R, P), \text{inferenceNeeds}(P, R, Q). \\
& \text{below}(P, Q) \leftarrow \text{factorVia}(P, Q). \\
& \text{below}(A, C) \leftarrow \text{below}(A, B), \text{below}(B, C). \\
& \text{factorViaI}(c(P, S_1, O_1), c(P, S_2, O_2)) \leftarrow \text{fai}(c(P, S_1, O_1)), \text{infer}(c(P, S_2, O_2)), \text{eq}(S_1, S_2), \\
& \quad \text{eq}(O_1, O_2), \text{not } \text{below}(c(P, S_2, O_2), c(P, S_1, O_1)). \\
& \text{factorI}(P) \leftarrow \text{factorViaI}(P, _).
\end{aligned}$$

$$\begin{aligned}
& fa(P) \leftarrow fai(P), \mathbf{not} \text{ factorI}(P). \\
& \text{factorVia}(A, B) \leftarrow \text{factorViaI}(A, B). \\
& \text{factorCluster}(c(P, S_2, O_2), c(P, S_1, O_1)) \leftarrow fa(c(P, S_1, O_1)), fa(c(P, S_2, O_2)), eq(S_1, S_2), \\
& \quad eq(O_1, O_2), c(P, S_1, O_1) < c(P, S_2, O_2). \\
& \text{factorClusterAbove}(A) \leftarrow \text{factorCluster}(A, _). \\
& \text{factorVia}(A, B) \leftarrow \text{factorCluster}(A, B), \\
& \quad \mathbf{not} \text{ factorClusterAbove}(B). \\
& \text{factor}(P) \leftarrow \text{factorVia}(P, _). \\
& \text{abduce}(P) \leftarrow fa(P), \mathbf{not} \text{ factor}(P).
\end{aligned}$$

Complete Factoring Encoding BWD-A.

$$\begin{aligned}
& fa(P) \leftarrow fai(P). \\
& \text{factorCluster}(c(P, S_2, O_2), c(P, S_1, O_1)) \leftarrow fa(c(P, S_1, O_1)), fa(c(P, S_2, O_2)), eq(S_1, S_2), \\
& \quad eq(O_1, O_2), c(P, S_1, O_1) < c(P, S_2, O_2). \\
& \text{factorClusterAbove}(A) \leftarrow \text{factorCluster}(A, _). \\
& \text{factorVia}(A, B) \leftarrow \text{factorCluster}(A, B), \\
& \quad \mathbf{not} \text{ factorClusterAbove}(B). \\
& \text{factor}(P) \leftarrow \text{factorVia}(P, _). \\
& \text{abduce}(P) \leftarrow fa(P), \mathbf{not} \text{ factor}(P).
\end{aligned}$$

Complete Encoding FWD-A. Each axiom of form (2) is rewritten into the following set of ASP rules:

$$\begin{aligned}
& \text{infer}(c(q, Y_1, \dots, Y_m)) \leftarrow \text{true}(c(p_1, X_1^1, \dots, X_{k_1}^1)), \dots, \text{true}(c(p_r, X_1^r, \dots, X_{k_r}^r)). \\
& \text{pot}(c(p_i, X_1^i, \dots, X_{k_1}^i)) \leftarrow Z_1 = s(a, 1, Y_1, \dots, Y_m), \dots, Z_v = s(a, v, Y_1, \dots, Y_m), \\
& \quad \text{pot}(c(q, Y_1, \dots, Y_m)). \quad \text{for } i \in \{1, \dots, r\}
\end{aligned}$$

where a is a unique identifier for that particular axiom and $Z_1, \dots, Z_v = \mathcal{X} \setminus \mathcal{Y}$.

The encoding FWD-A then contains the following rules.

$$\begin{aligned}
& \text{pot}(X) \leftarrow \text{goal}(X). \\
& \{ fai(X) : \text{pot}(X) \} \leftarrow . \\
& \text{true}(X) \leftarrow fai(X). \\
& \text{true}(X) \leftarrow \text{infer}(X). \\
& \quad \leftarrow \text{goal}(A), \mathbf{not} \text{ true}(A). \\
& fa(X) \leftarrow fai(X), \mathbf{not} \text{ infer}(X). \\
& hu(X) \leftarrow \text{pot}(c(_, X, _)). \\
& hu(X) \leftarrow \text{pot}(c(_, _, X)). \\
& uhu(X) \leftarrow hu(X), \mathbf{not} \text{ sortname}(X). \\
& \{ eq(A, B) : uhu(A), uhu(B), A \neq B \} \leftarrow . \\
& eq(A, A) \leftarrow hu(A). \\
& \quad \leftarrow eq(A, B), \mathbf{not} eq(B, A). \\
& \quad \leftarrow eq(A, B), eq(B, C), A \neq B, B \neq C, A \neq C, \mathbf{not} eq(A, C). \\
& \text{factorCluster}(c(P, S_2, O_2), c(P, S_1, O_1)) \leftarrow fa(c(P, S_1, O_1)), fa(c(P, S_2, O_2)), eq(S_1, S_2), \\
& \quad eq(O_1, O_2), c(P, S_1, O_1) < c(P, S_2, O_2). \\
& \text{factorClusterAbove}(A) \leftarrow \text{factorCluster}(A, _). \\
& \text{factorVia}(A, B) \leftarrow \text{factorCluster}(A, B), \\
& \quad \mathbf{not} \text{ factorClusterAbove}(B).
\end{aligned}$$

$$\begin{aligned} factor(P) &\leftarrow factorVia(P, _). \\ abduce(P) &\leftarrow fa(P), \mathbf{not} factor(P). \end{aligned}$$

9 Appendix: Running Example ASP Encoding and Answer Set

We next give the ASP input instance for our running example (Example 1) when used with BWD encodings. We then give representative parts of an answer set describing the abductive explanation shown in Figure 1.

9.1 Rewriting of Axioms, Goal, and Sortnames

As described in Section 3, given the an abduction instance $A = (B, O, S)$ of Example 1 we create the following ASP code.

We represent the goal in terms of facts (21).

$$\begin{aligned} goal(c(name, m, mary)). & & goal(c(lost, m, f)). & & goal(c(fatherof, f, m)). \\ goal(c(inst, s, female)). & & goal(c(is, s, depressed)). & & \end{aligned}$$

We represent sort names as follows.

$$\begin{aligned} sortname(depressed). & & sortname(dead). & & sortname(person). \\ sortname(female). & & sortname(male). & & \end{aligned}$$

According to the rewriting (35) in Section 3.2, we rewrite axioms of Example 1 as follows.

Axiom (4) ‘ $inst(X, male) \Leftarrow fatherof(X, Y)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_4, c(inst, X, male), l(Y)) &\leftarrow pot(c(inst, X, male)), Y = s(r_4, \text{“}Y\text{”}, X). \\ inferenceNeeds(c(inst, X, male), r_4, c(fatherof, X, Y)) &\leftarrow \\ &mayInferVia(r_4, c(inst, X, male), l(Y)). \\ numberOfBodyAtoms(r_4, 1). \end{aligned}$$

Note the Skolemization of variable Y which exists only in the body of (4) but not in the head. Also note that we use rule identifiers that are synchronized with the original axiom numbers, i.e., for (4) we use r_4 .

Axiom (5) ‘ $inst(X, female) \Leftarrow name(X, mary)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_5, c(inst, X, female), l) &\leftarrow pot(c(inst, X, female)). \\ inferenceNeeds(c(inst, X, female), r_5, c(name, X, mary)) &\leftarrow \\ &mayInferVia(r_5, c(inst, X, female), l). \\ numberOfBodyAtoms(r_5, 1). \end{aligned}$$

Note that in this axiom there is no Skolemization, so l has no arguments (however we still need it to keep the arity of $mayInferVia$ the same throughout the encoding).

Axiom (6) ‘ $importantfor(Y, X) \Leftarrow fatherof(Y, X)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_6, c(importantfor, Y, X), l) &\leftarrow pot(c(importantfor, Y, X)). \\ inferenceNeeds(c(importantfor, Y, X), r_6, c(fatherof, Y, X)) &\leftarrow \\ &mayInferVia(r_6, c(importantfor, Y, X), l). \\ numberOfBodyAtoms(r_6, 1). \end{aligned}$$

Axiom (7) ‘ $inst(X, person) \Leftarrow inst(X, male)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_7, c(inst, X, person), l) &\leftarrow pot(c(inst, X, person)). \\ inferenceNeeds(c(inst, X, person), r_7, c(inst, X, male)) &\leftarrow mayInferVia(r_7, c(inst, X, person), l). \end{aligned}$$

$numberOfBodyAtoms(r_7, 1)$.

Axiom (8) ‘ $is(X, depressed) \Leftarrow inst(X, pessimist)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_8, c(is, X, depressed), l) &\leftarrow pot(c(is, X, depressed)). \\ inferenceNeeds(c(is, X, depressed), r_8, c(inst, X, pessimist)) &\leftarrow \\ &mayInferVia(r_8, c(is, X, depressed), l). \\ numberOfBodyAtoms(r_8, 1). \end{aligned}$$

Axiom (9) ‘ $is(X, depressed) \Leftarrow is(Y, dead) \wedge importantfor(Y, X)$ ’ is rewritten as follows.

$$\begin{aligned} mayInferVia(r_9, c(is, X, depressed), l(Y)) &\leftarrow pot(c(is, X, depressed)), Y = s(r_9, \text{“Y”}, X). \\ inferenceNeeds(c(is, X, depressed), r_9, c(importantfor, Y, X)) &\leftarrow \\ &mayInferVia(r_9, c(is, X, depressed), l(Y)). \\ inferenceNeeds(c(is, X, depressed), r_9, c(is, Y, dead)) &\leftarrow \\ &mayInferVia(r_9, c(is, X, depressed), l(Y)). \\ numberOfBodyAtoms(r_9, 2). \end{aligned}$$

Axiom (10) ‘ $lost(X, Y) \Leftarrow is(Y, dead) \wedge importantfor(Y, X) \wedge inst(Y, person)$ ’ is rewritten into the following ASP rules.

$$\begin{aligned} mayInferVia(r_{10}, c(lost, X, Y), l) &\leftarrow pot(c(lost, X, Y)). \\ inferenceNeeds(c(lost, X, Y), r_{10}, c(inst, Y, person)) &\leftarrow mayInferVia(r_{10}, c(lost, X, Y), l). \\ inferenceNeeds(c(lost, X, Y), r_{10}, c(importantfor, Y, X)) &\leftarrow mayInferVia(r_{10}, c(lost, X, Y), l). \\ inferenceNeeds(c(lost, X, Y), r_{10}, c(is, Y, dead)) &\leftarrow mayInferVia(r_{10}, c(lost, X, Y), l). \\ numberOfBodyAtoms(r_{10}, 3). \end{aligned}$$

9.2 Example Answer Set

We next give parts of the answer set representing the abductive explanation depicted in Figure 1. We show an answer set of the encoding BWD-G which does not perform any canonicalization and therefore can produce the proof graph in Figure 1 (other encodings will produce larger proof graphs with the same cost for this instance).

Note that the encoding produces Skolem terms $s(r_9, \text{“Y”}, s)$ and $s(r_4, \text{“Y”}, f)$, which, for practical reasons, have been displayed in Figure 1 as n_1 and n_2 , respectively.

Deterministically determined Atoms. Based on the rewriting of axioms (previous section) and the goal atoms, the truth of atoms of form $pot(\cdot)$, $mayInferVia(\cdot, \cdot, \cdot)$, and $inferenceNeeds(\cdot, \cdot, \cdot)$, is deterministically determined via rules (23) and (34).

In our example we obtain the following true atoms in the answer set.

$$\begin{array}{lll} pot(c(fatherof, f, s(r_4, \text{“Y”}, f))) & pot(c(fatherof, f, m)) & \\ pot(c(fatherof, s(r_9, \text{“Y”}, s), s)) & pot(c(importantfor, f, m)) & \\ pot(c(importantfor, s(r_9, \text{“Y”}, s), s)) & pot(c(inst, f, female)) & \\ pot(c(inst, f, male)) & pot(c(inst, f, person)) & pot(c(inst, s, female)) \\ pot(c(is, s(r_9, \text{“Y”}, s), dead)) & pot(c(inst, s, pessimist)) & pot(c(is, f, dead)) \\ pot(c(is, s, depressed)) & pot(c(lost, m, f)) & pot(c(name, f, mary)) \\ pot(c(name, m, mary)) & pot(c(name, s, mary)) & \\ \\ mayInferVia(r_4, c(inst, f, male), l(s(r_4, \text{“Y”}, f))) & mayInferVia(r_5, c(inst, f, female), l) & \\ mayInferVia(r_5, c(inst, s, female), l) & mayInferVia(r_6, c(importantfor, f, m), l) & \\ mayInferVia(r_6, c(importantfor, s(r_9, \text{“Y”}, s), s), l) & mayInferVia(r_7, c(inst, f, person), l) & \\ mayInferVia(r_8, c(is, s, depressed), l) & & \\ mayInferVia(r_9, c(is, s, depressed), l(s(r_9, \text{“Y”}, s))) & mayInferVia(r_{10}, c(lost, m, f), l) & \end{array}$$

inferenceNeeds(*c*(*inst*, *f*, *male*), *r*₄, *c*(*fatherof*, *f*, *s*(*r*₄, “*Y*”, *f*)))
inferenceNeeds(*c*(*inst*, *f*, *female*), *r*₅, *c*(*name*, *f*, *mary*))
inferenceNeeds(*c*(*inst*, *s*, *female*), *r*₅, *c*(*name*, *s*, *mary*))
inferenceNeeds(*c*(*importantfor*, *f*, *m*), *r*₆, *c*(*fatherof*, *f*, *m*))
inferenceNeeds(*c*(*importantfor*, *s*(*r*₉, “*Y*”, *s*), *s*), *r*₆, *c*(*fatherof*, *s*(*r*₉, “*Y*”, *s*), *s*))
inferenceNeeds(*c*(*inst*, *f*, *person*), *r*₇, *c*(*inst*, *f*, *male*))
inferenceNeeds(*c*(*is*, *s*, *depressed*), *r*₈, *c*(*inst*, *s*, *pessimist*))
inferenceNeeds(*c*(*is*, *s*, *depressed*), *r*₉, *c*(*is*, *s*(*r*₉, “*Y*”, *s*), *dead*))
inferenceNeeds(*c*(*is*, *s*, *depressed*), *r*₉, *c*(*importantfor*, *s*(*r*₉, “*Y*”, *s*), *s*))
inferenceNeeds(*c*(*lost*, *m*, *f*), *r*₁₀, *c*(*inst*, *f*, *person*))
inferenceNeeds(*c*(*lost*, *m*, *f*), *r*₁₀, *c*(*importantfor*, *f*, *m*))
inferenceNeeds(*c*(*lost*, *m*, *f*), *r*₁₀, *c*(*is*, *f*, *dead*))

From these atoms, *hu*(·), and *uhu*(·) are deterministically determined using rules (24)–(26).

<i>hu</i> (<i>m</i>)	<i>hu</i> (<i>female</i>)	<i>hu</i> (<i>person</i>)	<i>hu</i> (<i>depressed</i>)
<i>hu</i> (<i>f</i>)	<i>hu</i> (<i>male</i>)	<i>hu</i> (<i>pessimist</i>)	<i>hu</i> (<i>s</i> (<i>r</i> ₄ , “ <i>Y</i> ”, <i>f</i>))
<i>hu</i> (<i>s</i>)	<i>hu</i> (<i>mary</i>)	<i>hu</i> (<i>dead</i>)	<i>hu</i> (<i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>))
<i>uhu</i> (<i>f</i>)	<i>uhu</i> (<i>m</i>)	<i>uhu</i> (<i>mary</i>)	<i>uhu</i> (<i>pessimist</i>)
<i>uhu</i> (<i>s</i>)	<i>uhu</i> (<i>s</i> (<i>r</i> ₄ , “ <i>Y</i> ”, <i>f</i>))	<i>uhu</i> (<i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>))	

Goal atoms are deterministically defined as true by (36).

<i>true</i> (<i>c</i> (<i>name</i> , <i>m</i> , <i>mary</i>))	<i>true</i> (<i>c</i> (<i>inst</i> , <i>s</i> , <i>female</i>))	<i>true</i> (<i>c</i> (<i>lost</i> , <i>m</i> , <i>f</i>))
<i>true</i> (<i>c</i> (<i>fatherof</i> , <i>f</i> , <i>m</i>))	<i>true</i> (<i>c</i> (<i>is</i> , <i>s</i> , <i>depressed</i>))	

Truth Justification. The BWD-G encodings requires a justification for each true value, this justification is nondeterministically guessed to be either *infer*(·), *factor*(·), or *abduce*(·) by rules (37) and (40). Justifying an atom by inference performs another nondeterministic guess with (38) about the concrete axiom used for that inference, which is represented in *inferVia*(·, ·) and defines more atoms of form *true*(·) to be true (and hence their need to be justified) with (39).

In the answer set representing Figure 1, this guess contains the following true atoms (we omit *fai*(·)).

<i>infer</i> (<i>c</i> (<i>inst</i> , <i>f</i> , <i>male</i>))	<i>inferVia</i> (<i>r</i> ₄ , <i>c</i> (<i>inst</i> , <i>f</i> , <i>male</i>))
<i>infer</i> (<i>c</i> (<i>inst</i> , <i>s</i> , <i>female</i>))	<i>inferVia</i> (<i>r</i> ₅ , <i>c</i> (<i>inst</i> , <i>s</i> , <i>female</i>))
<i>infer</i> (<i>c</i> (<i>importantfor</i> , <i>f</i> , <i>m</i>))	<i>inferVia</i> (<i>r</i> ₆ , <i>c</i> (<i>importantfor</i> , <i>f</i> , <i>m</i>))
<i>infer</i> (<i>c</i> (<i>inst</i> , <i>f</i> , <i>person</i>))	<i>inferVia</i> (<i>r</i> ₇ , <i>c</i> (<i>inst</i> , <i>f</i> , <i>person</i>))
<i>infer</i> (<i>c</i> (<i>is</i> , <i>s</i> , <i>depressed</i>))	<i>inferVia</i> (<i>r</i> ₉ , <i>c</i> (<i>is</i> , <i>s</i> , <i>depressed</i>))
<i>infer</i> (<i>c</i> (<i>lost</i> , <i>m</i> , <i>f</i>))	<i>inferVia</i> (<i>r</i> ₁₀ , <i>c</i> (<i>lost</i> , <i>m</i> , <i>f</i>))
<i>factor</i> (<i>c</i> (<i>name</i> , <i>s</i> , <i>mary</i>))	<i>factor</i> (<i>c</i> (<i>fatherof</i> , <i>f</i> , <i>s</i> (<i>r</i> ₄ , “ <i>Y</i> ”, <i>f</i>)))
<i>factor</i> (<i>c</i> (<i>is</i> , <i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>), <i>dead</i>))	<i>factor</i> (<i>c</i> (<i>importantfor</i> , <i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>), <i>s</i>))
<i>abduce</i> (<i>c</i> (<i>name</i> , <i>m</i> , <i>mary</i>))	<i>abduce</i> (<i>c</i> (<i>fatherof</i> , <i>f</i> , <i>m</i>))
<i>abduce</i> (<i>c</i> (<i>is</i> , <i>f</i> , <i>dead</i>))	
<i>true</i> (<i>c</i> (<i>name</i> , <i>s</i> , <i>mary</i>))	<i>true</i> (<i>c</i> (<i>inst</i> , <i>f</i> , <i>person</i>))
<i>true</i> (<i>c</i> (<i>inst</i> , <i>f</i> , <i>male</i>))	<i>true</i> (<i>c</i> (<i>fatherof</i> , <i>f</i> , <i>s</i> (<i>r</i> ₄ , “ <i>Y</i> ”, <i>f</i>)))
<i>true</i> (<i>c</i> (<i>is</i> , <i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>), <i>dead</i>))	<i>true</i> (<i>c</i> (<i>is</i> , <i>f</i> , <i>dead</i>))
<i>true</i> (<i>c</i> (<i>importantfor</i> , <i>f</i> , <i>m</i>))	<i>true</i> (<i>c</i> (<i>importantfor</i> , <i>s</i> (<i>r</i> ₉ , “ <i>Y</i> ”, <i>s</i>), <i>s</i>))

Term Equivalence. Independent from justification of true atoms, an equivalence relation over all potential terms is nondeterministically guessed by means of rules (27)–(30).

The answer set representing Figure 1 contains the following true atoms for $eq(\cdot, \cdot)$. These atoms represent two equivalence classes $\{m, s, s(r_4, \text{“Y”}, f)\}$ and $\{f, s(r_9, \text{“Y”}, s)\}$ that have more than one element, and singleton equivalence classes for sort names and other constants.

$eq(male, male)$	$eq(female, female)$	$eq(dead, dead)$
$eq(mary, mary)$	$eq(person, person)$	$eq(pessimist, pessimist)$
$eq(f, f)$	$eq(f, s(r_9, \text{“Y”}, s))$	$eq(depressed, depressed)$
$eq(m, m)$	$eq(m, s)$	$eq(m, s(r_4, \text{“Y”}, f))$
$eq(s, m)$	$eq(s, s)$	$eq(s, s(r_4, \text{“Y”}, f))$
$eq(s(r_4, \text{“Y”}, f), m)$	$eq(s(r_4, \text{“Y”}, f), s)$	$eq(s(r_4, \text{“Y”}, f), s(r_4, \text{“Y”}, f))$
$eq(s(r_9, \text{“Y”}, s), f)$	$eq(s(r_9, \text{“Y”}, s), s(r_9, \text{“Y”}, s))$	

Factoring. For each atom that was guessed as factored, rules (41)–(46) define $factorVia(\cdot, \cdot)$ and $factorOk(\cdot)$ for atoms that can be factored with inferred or abduced atoms while respecting $eq(\cdot, \cdot)$ and acyclicity. For acyclicity, the partial order defined by the graph is represented in $below(\cdot, \cdot)$, which is true for every pair of atoms such that the first atom is reachable via arcs from the second atom. The representation of Figure 1 contains the following true atoms.

$factorVia(c(name, s, mary), c(name, m, mary))$	
$factorVia(c(is, s(r_9, \text{“Y”}, s), dead), c(is, f, dead))$	
$factorVia(c(importantfor, s(r_9, \text{“Y”}, s), s), c(importantfor, f, m))$	
$factorVia(c(fatherof, f, s(r_4, \text{“Y”}, f)), c(fatherof, f, m))$	
$factorOk(c(name, s, mary))$	$factorOk(c(is, s(r_9, \text{“Y”}, s), dead))$
$factorOk(c(importantfor, s(r_9, \text{“Y”}, s), s))$	$factorOk(c(fatherof, f, s(r_4, \text{“Y”}, f)))$
$below(c(inst, s, female), c(name, m, mary))$	$below(c(inst, s, female), c(name, s, mary))$
$below(c(name, s, mary), c(name, m, mary))$	$below(c(lost, m, f), c(inst, f, person))$
$below(c(lost, m, f), c(importantfor, f, m))$	$below(c(lost, m, f), c(is, f, dead))$
$below(c(lost, m, f), c(inst, f, male))$	$below(c(lost, m, f), c(fatherof, f, m))$
$below(c(lost, m, f), c(fatherof, f, s(r_4, \text{“Y”}, f)))$	$below(c(inst, f, person), c(inst, f, male))$
$below(c(inst, f, person), c(fatherof, f, m))$	$below(c(is, s, depressed), c(is, f, dead))$
$below(c(inst, f, male), c(fatherof, f, s(r_4, \text{“Y”}, f)))$	$below(c(inst, f, male), c(fatherof, f, m))$
$below(c(is, s, depressed), c(importantfor, f, m))$	$below(c(is, s, depressed), c(fatherof, f, m))$
$below(c(importantfor, f, m), c(fatherof, f, m))$	
$below(c(fatherof, f, s(r_4, \text{“Y”}, f)), c(fatherof, f, m))$	
$below(c(is, s(r_9, \text{“Y”}, s), dead), c(is, f, dead))$	
$below(c(inst, f, person), c(fatherof, f, s(r_4, \text{“Y”}, f)))$	
$below(c(is, s, depressed), c(is, s(r_9, \text{“Y”}, s), dead))$	
$below(c(is, s, depressed), c(importantfor, s(r_9, \text{“Y”}, s), s))$	
$below(c(importantfor, s(r_9, \text{“Y”}, s), s), c(fatherof, f, m))$	
$below(c(importantfor, s(r_9, \text{“Y”}, s), s), c(importantfor, f, m))$	

For representing the cost of the solution wrt. objective WA, (72) defines the following goal costs.

$pcost(c(name, m, mary), 100)$	$pcost(c(lost, m, f), 100)$	$pcost(c(fatherof, f, m), 100)$
$pcost(c(inst, s, female), 100)$	$pcost(c(is, s, depressed), 100)$	

Cost propagation via inferences is done by (73) which makes the following atoms true.

$$\begin{array}{ll}
pcost(c(name, s, mary), 120) & pcost(c(inst, f, person), 40) \\
pcost(c(importantfor, f, m), 40) & pcost(c(is, f, dead), 40) \\
pcost(c(inst, f, male), 48) & pcost(c(fatherof, f, s(r_4, "Y", f)), 57) \\
pcost(c(is, s(r_9, "Y", s), dead), 60) & pcost(c(importantfor, s(r_9, "Y", s), s), 60) \\
pcost(c(fatherof, f, m), 48) & pcost(c(fatherof, f, m), 72)
\end{array}$$

Note that the cost of 72\$ is not shown in Figure 1 because it is not contained in the definition of proof graph (Definition 2). In the Proof of Proposition 7 we argue that including these costs in ASP only adds further (higher) costs to each atom, hence the minimal costs and therefore the optimal solution remain unchanged. Note that we include these costs (i.e., we omit one more minimization step) for efficiency reasons.

Cost propagation via factoring is realized in (74) and makes the following atoms true.

$$\begin{array}{ll}
pcost(c(name, m, mary), 120) & pcost(c(fatherof, f, m), 57) \\
pcost(c(is, f, dead), 60) & pcost(c(importantfor, f, m), 60)
\end{array}$$

Actual cost is defined from potential cost via (75) which yields the following true atoms.

$$cost(c(name, m, mary), 100) \quad cost(c(fatherof, f, m), 48) \quad cost(c(is, f, dead), 40)$$

From these atoms, the constraint (76) obtains the overall cost of 188 for this answer set.

10 Appendix: Proofs for ASP Encoding Correctness

Proof of Lemma 1 (Sketch):

$P_{bpt}(A)$ does not contain **not** and its rules are not expanding term depth except for the first rule of (35). As B is acyclic, by construction of $P_{bpt}(A)$ there are no loops over (35), $grnd(P_{bpt}(A))$ is finite, and $AS(P_{bpt}(A)) = \{I\}$ has a single answer set I . The representation of proof graphs is achieved as follows: atoms $A = p(a, b)$ that can be back-chained are represented as $pot(c(p, a, b)) \in I$: We proceed by induction on the distance d of back-chaining from observations. (Base: $d=0$) Due to (23) all atoms $p(a, b) \in O$ are true as $pot(c(p, a, b)) \in I$. (Step: $d \Rightarrow d+1$) Assuming that $A = q(a, b)$, $A \in H$, $H \in \hat{\mathcal{H}}$, is represented as $pot(c(q, a, b)) \in I$, this causes an instantiation of the first rule in (35), which defines $mayInferVia(r, c(q, a, b), l(z_1, \dots, z_v))$ true in I . This represents potential backchaining from $q(a, b)$ over an axiom identified by r using substitution $\theta = \{Z_1 \mapsto z_1, \dots, Z_v \mapsto z_v\}$ where Z_1, \dots, Z_v are variables occurring only in the body of r . Truth of $mayInferVia(r, c(q, a, b), l(z_1, \dots, z_v))$ causes truth of $inferenceNeeds(c(q, a, b), r, c(p_i, x_1^j, x_2^j))$ where p_i is the predicate and x_i^j are the substituted variables at position i in body atom j of axiom r , analogous to backward chaining in Def. 1. Due to truth of $inferenceNeeds(c(q, a, b), r, c(p_i, x_i^j, x_{i+1}^j))$ and due to (34), all body atoms $\theta(p_i(X_1^i, \dots, X_{k_i}^i))$ added to some $H \in \hat{\mathcal{H}}$ in Def. 1 become represented as $pot(c(p_i, x_i^j, x_{i+1}^j)) \in I$. (Conclusion) We conclude that I contains all atoms $p(a, b)$ in hypotheses generated from observations represented as $pot(c(p, a, b)) \in I$. Moreover, $mayInferVia(\dots)$ represents potential back-chaining and $inferenceNeeds(\dots)$ represents the body atoms that are added to a hypothesis by a particular backchaining. \square

Proof of Proposition 2 (Sketch):

We write P_{bpt} for $P_{bpt}(A)$. P_{bpt} is a bottom of $P_{bpt} \cup P_{gpt}$ and therefore each $I \in AS(P_{bpt} \cup P_{gpt})$ is such that $I = I' \cup I_g$ where $I' \in AS(P_{bpt})$ and I_g contains only predicates $infer, fai, inferVia, true$. (36) defines for all $o \in O$, $o = p(a, b)$, that $true(c(p, a, b)) \in I_g$. (37) defines that every atom P with $true(P) \in I_g$, either $infer(P) \in I_g$ or $fai(P) \in I_g$ (i.e., two answer set candidates are generated). (38) defines that every atom P with $infer(P) \in I_g$ is marked as inferred via a particular axiom along an edge $mayInferVia(R, P, Z) \in I'$ of the potential proof graph, and represents this inference as $inferVia(R, P)$ where R is the axiom identifier. (39) defines that for each $inferVia(R, P)$ the corresponding required body atoms $X_1^i, \dots, X_{k_i}^i$ represented as $inferenceNeeds(P, R, c(p_i, X_1^i, \dots, X_{k_i}^i)) \in I'$ are defined as $true(c(p_i, X_1^i, \dots, X_{k_i}^i)) \in I_g$. As defining these as true, they again must be represented as $infer$ or fai due to (37). Due to minimality

of answer sets, all atoms marked as *true*, *infer*, and *fai* are reachable from objectives $o \in O$ which are also represented as $true(\cdot) \in I_g$. The set of hypotheses is inductively built from observations and any combination of backchaining over axioms from these observations and atoms obtained from backchaining, exactly as the set of atoms marked as *true* in I_g is inductively defined from observations and any choice about back-chaining in (37). Therefore, answer sets and proof graphs are in 1-1 correspondence. \square

Proof of Lemma 3 (Sketch):

We write P_{bpt} for $P_{bpt}(A)$. P_{eq} does not define predicates present in P_{bpt} , hence P_{bpt} is a bottom wrt. $P_{bpt} \cup P_{eq}$ (see [LT94]) and $I' \in AS(P_{bpt} \cup P_{eq})$ is such that $I' = I \cup I_e$ where $I \in AS(P_{bpt})$ and I_e contains only predicates *eq*, *hu*, and *uhu*. All constants c in argument positions of hypotheses are represented in $hu(c)$ due to (24)–(25), those that are not sort names are additionally represented in $uhu(c)$ due to (26). (27) guesses a relation $eq(c, c')$ a solution candidate for all pairs (c, c') of constants with $c \neq c'$ that do not contain sort names. Finally, (28) defines *eq* to be reflexive for all constants (including sort names), and (29)/(30) exclude answer sets where the represented relation *eq* is not symmetric/transitive. Therefore, only those relations remain that are reflexive, symmetric, and transitive, i.e., equivalence relations. \square

Proof of Proposition 6 (Sketch):

Let $P_{gpeq}(A) = P_{bpt}(A) \cup P_{gp} \cup P_{eq}$, then $P_{bpt}(A) \cup P_{gp}$ and $P_{bpt}(A) \cup P_{eq}$ are bottoms wrt. $P_{gpeq}(A)$, and P_{gp} and P_{eq} do not have common head atoms, hence both Prop. 2 and Prop. 3 apply to answer sets $I \in AS(P_{gpeq}(A))$, viz. each I is in 1-1 correspondence with some proof graph G and hypothesis $H \in \hat{\mathcal{H}}(A)$ and moreover represents some equivalence relation over all constants of the proof graph in I , moreover all proof graphs originating in back-chaining are covered. Furthermore, $P_{gpeq}(A)$ is a bottom wrt. $P_{BWD-A}(A)$, therefore each answer set $I' \in AS(P_{BWD-A}(A))$ is such that $I' = I_{gpeq} \cup I$ where $I_{gpeq} \in AS(P_{gpeq}(A))$ and I contains predicates defined by P_c based on I_{gpeq} . P_c contains only stratified negation. (57) defines $fa(P)$ to be true iff $fai(P) \in I_{gpeq}$, hence iff atom P is not inferred in H . For all atoms $P, Q \in H$ that are not inferred, (52) defines $factorCluster(P, Q)$ to be true if P and Q unify under *eq* represented in I_{gpeq} , Q is lexicographically smaller than P , and neither P nor Q were back-chained in G , i.e., they would be abducted unless they can be factored. Note that, given a set X of atoms that unify wrt. *eq* (called a ‘cluster’), (52) defines a relation that contains $factorCluster(s_1, s_2)$ for all $s_1, s_2 \in X$ where $s_2 < s_1$. (53) represents all constants in all clusters that can be factored with a smaller element. (54) uses constants that have no such smaller element as representatives and defines $factorVia(s, s')$ for all $s, s' \in X$ where s' is the smallest element of the respective cluster X . Finally, every atom s that was unified with a representative s' in $factorVia(s, s')$ is represented as factored $factor(s)$ by (55) and those atoms that are neither factored nor inferred are defined as $abduce(s)$ by (56). Hence, in the answer set $I' \in AS(P_{BWD-A}(A))$, all atoms $s \in H$ that (i) are not back-chained on, i.e., are not marked as inferred, and that (ii) can be unified with a lexicographically smaller atom s' , are marked as $factor(s) \in I'$, and the factoring edge $(s, s') \in E(G)$ is represented as $factorVia(s, s') \in I'$. Those atoms $s \in H$ that are neither factored nor inferred are marked as $abduce(s) \in I'$. As $I \in AS(P_{gpeq}(A))$ is in 1-1 correspondence with proof graphs based on backchaining and some equivalence relation, and P_c adds to that the representation of factored and abducted atoms ($factor(\cdot)$ and $abduce(\cdot)$) and factoring edges $factorVia(\cdot, \cdot)$, $I' \in AS(P_{BWD-A}(A))$ is in 1-1 correspondence with proof graphs based on back-chaining and unification, and the equivalence relation that supports this unification. \square

Proof of Proposition 7 (Sketch):

(CARD) (66) causes cost $|\{a \mid abduce(a) \in I\}|$ and $abduce(a) \in I$ iff a is neither inferred nor factored. As being abducted is equivalent with the absence of edges $factorVia(a, a') \in I$ and edges $inferVia(r, a) \in I$, (66) produces cost $CARD(G)$ for an answer set I representing G .

(COH) (67) seeds a definition of reachability from goal nodes in predicate *reach*, (68) defines reachability across those inference edges that correspond with inferences actually done in the proof graph represented in I , and (69) defines reachability across factoring edges. As a result $reach(a, o)$ is true for atom $a \in H$ and observation $o \in O$ iff o is reachable from a in G . (70) defines $reachFromBoth(o, o') \in I$ iff $o, o' \in O$, $o < o'$, and there is some atom $a \in H$ such that a is reachable from o and o' . The if direction is ensured by rule satisfaction, the only if direction is ensured by answer set minimality. Finally, weak constraint (71) attaches cost 1 for each distinct pair o, o' of observation nodes where no node reachable from both o and o' exists in G . This exactly corresponds to the definition of COH.

(WA) (72) defines potential cost of 100 for objective nodes $o \in O$. (73) defines potential cost of $1.2 \cdot c/n$ for each body atom of back-chaining, given that the back-chained atom had potential cost c and was

Algorithm 1: ONMODELWITHOUTOPTIMIZATION(Model m , Clasp control object c)

```
1 violations := FINDONDEMANDVIOLATIONS( $m$ )
2 foreach  $v \in \textit{violations}$  do add nogood forbidding violation  $v$  to  $c$ 
3 if  $\textit{violations} = \emptyset$  then print  $m$ 
```

back-chained over a rule with n body atoms. (74) defines that for atoms p, q where p was factored with q , q obtains all potential costs from p . Potential cost includes all costs obtained from reachable nodes, including the minimum cost in case of factoring. Therefore, the minimum costs propagated for unification as described in Def. 3 is represented as potential cost, along with bigger costs. Rule (75) represents for each abduced atom, i.e., for each $a \in A(G)$ for G represented in I , the minimum over all potential costs of a . Hence, $\textit{cost}(p, c) \in I$ iff $p \in A(G)$, and c is the cost of abduced atom $p \in H$ according to WA. (76) sums up costs of distinct atoms p , hence $\textit{cost WA}(G)$ is assigned to I . \square

11 Appendix: Realizing On-Demand Constraints with Optimization

If we solve a problem without weak constraints, i.e., without optimization, realizing on-demand constraints is simple: we register the callback ONMODELWITHOUTOPTIMIZATION (Algorithm 1) to Clasp. This callback checks on-demand constraints in FINDONDEMANDVIOLATIONS (which is an application-specific algorithm), adds nogoods for violated constraints, and prints (or otherwise processes) the answer set if no constraint was violated. The first model we print this way is the first model that does not violate on-demand constraints.

However, in the presence of optimization, Clasp has two modes that are both unsuitable with on-demand constraints: in mode (opt) each answer set updates the optimization bound and subsequent answer sets must be better; in mode (enum) we have to explicitly specify a bound and answer set candidates of same or better quality are found. With (opt) it can happen that the first answer set with optimal cost violates an on-demand constraint, so we discard that model, but the solver will not find further models with same cost, but no models with better cost exist, so we will not find any models (even if some exist). With (enum) the search is blind as better models are found only by chance. A straightforward and more elegant solution would be, to update the bound only for good answer sets in the on_model callback, but the API currently does not allow this.

To solve this problem we created algorithm FINDOPTIMALMODEL (Algorithm 2) for finding an optimal model with on-demand constraints. This algorithm first grounds the program P in line 1, then uses (opt) mode to find an optimistic bound (*optimisticCost*) for the cost of the optimal model in lines 2–4 using callback ONMODELFINDBOUND (Algorithm 3). This callback records the cost of the best encountered model candidate and the best model that does not violate on-demand constraints (*bestModel*). This search aggressively updates the bound and also uses on-demand constraints. If no optimistic cost is set, the callback was never called and we return UNSAT (line 5). If the cost of the best found feasible model is the optimal cost, we directly return this model as optimal solution (line 6). Otherwise, we enter a loop that enumerates models using callback ONMODELFINDSOLUTION (Algorithm 4). The loop increases the optimization bound of the solver by one in each iteration, until an answer set that does not violate on-demand constraints can be found. Our abduction instances always have some solution, so we will find that solution and leave the endless loop in line 13. To make the algorithm terminate in the general case where on-demand constraints might eliminate all solutions, we can obtain the worst-case cost from the instantiation of all weak constraints, and abort the loop once we increment *tryingCost* to that cost.

Algorithm 2: FINDOPTIMALMODEL (ASP Program P)

```
global: optimisticCost, bestModelCost, bestModel
1 Gringo.ground( $P$ )
2 optimisticCost, bestModelCost, bestModel := undef, undef, undef
3 Clasp.mode := opt // models must be strictly better than previously found models
4 Clasp.solve(on_model=ONMODELFINDBOUND)
5 if optimisticCost  $\neq$  undef then return (UNSAT, -1)
6 if bestModelCost = optimisticCost then return (OPT, bestModel)
7 tryingCost, bestModel := bestModelCost, undef
8 repeat
9   Clasp.mode := enum // finds models with equal or better cost
10  Clasp.opt_bound := tryingCost // cost bound for models
11  Clasp.models := 0 // find all models, not only the first one
12  Clasp.solve(on_model=ONMODELFINDSOLUTION)
13  if bestModel  $\neq$  undef then return (OPT, bestModel)
14  else tryingCost := tryingCost + 1
15 until forever
```

Algorithm 3: ONMODELFINDBOUND(Model m , Clasp control object c)

```
global: optimisticCost, bestModelCost, bestModel
1 violations := FINDONDEMANDVIOLATIONS( $m$ )
2 foreach  $v \in$  violations do add nogood forbidding violation  $v$  to  $c$ 
3 optimisticCost :=  $m$ .cost
4 if violations =  $\emptyset$  then
5   bestModelCost :=  $m$ .cost
6   bestModel :=  $m$ 
```

Algorithm 4: ONMODELFINDSOLUTION(Model m , Clasp control object c)

```
global: bestModelCost
1 violations := FINDONDEMANDVIOLATIONS( $m$ )
2 foreach  $v \in$  violations do add nogood forbidding violation  $v$  to  $c$ 
3 if violations =  $\emptyset$  then
4   bestModel :=  $m$ 
5   add nogood forbidding  $\emptyset$  to  $c$  // make problem inconsistent, abort search
```
