



Logical reduction of metarules

Andrew Cropper¹ · Sophie Touret²

Received: 15 February 2019 / Revised: 12 July 2019 / Accepted: 6 September 2019
© The Author(s) 2019

Abstract

Many forms of inductive logic programming (ILP) use *metarules*, second-order Horn clauses, to define the structure of learnable programs and thus the hypothesis space. Deciding which metarules to use for a given learning task is a major open problem and is a trade-off between efficiency and expressivity: the hypothesis space grows given more metarules, so we wish to use fewer metarules, but if we use too few metarules then we lose expressivity. In this paper, we study whether fragments of metarules can be logically reduced to minimal finite subsets. We consider two traditional forms of logical reduction: subsumption and entailment. We also consider a new reduction technique called *derivation reduction*, which is based on SLD-resolution. We compute reduced sets of metarules for fragments relevant to ILP and theoretically show whether these reduced sets are reductions for more general infinite fragments. We experimentally compare learning with reduced sets of metarules on three domains: Michalski trains, string transformations, and game rules. In general, derivation reduced sets of metarules outperform subsumption and entailment reduced sets, both in terms of predictive accuracies and learning times.

Keywords Inductive logic programming · Meta-interpretive learning · Logical reduction · Program induction · Inductive programming

1 Introduction

Many forms of inductive logic programming (ILP) (Albarghouthi et al. 2017; Campero et al. 2018; Cropper and Muggleton 2019; Emde et al. 1983; Evans and Grefenstette 2018; Flener 1996; Kaminski et al. 2018; Kietz and Wrobel 1992; Muggleton et al. 2015; De Raedt and Bruynooghe 1992; Si et al. 2018; Wang et al. 2014) use second-order Horn clauses, called

Editors: Dimitar Kazakov and Filip Zelezny.

✉ Andrew Cropper
andrew.cropper@cs.ox.ac.uk
Sophie Touret
stouret@mpi-inf.mpg.de

¹ University of Oxford, Oxford, UK

² Max Planck Institute for Informatics, Saarbrücken, Germany

metarules.¹ as a form of declarative bias (De Raedt 2012). Metarules define the structure of learnable programs which in turn defines the hypothesis space. For instance, to learn the *grandparent/2* relation given the *parent/2* relation, the *chain* metarule would be suitable:

$$P(A, B) \leftarrow Q(A, C), R(C, B)$$

In this metarule² the letters P , Q , and R denote existentially quantified second-order variables (variables that can be bound to predicate symbols) and the letters A , B and C denote universally quantified first-order variables (variables that can be bound to constant symbols). Given the *chain* metarule, the background *parent/2* relation, and examples of the *grandparent/2* relation, ILP approaches will try to find suitable substitutions for the existentially quantified second-order variables, such as the substitutions $\{P/\text{grandparent}, Q/\text{parent}, R/\text{parent}\}$ to induce the theory:

$$\text{grandparent}(A, B) \leftarrow \text{parent}(A, C), \text{parent}(C, B)$$

However, despite the widespread use of metarules, there is little work determining which metarules to use for a given learning task. Instead, suitable metarules are assumed to be given as part of the background knowledge, and are often used without any theoretical justification. Deciding which metarules to use for a given learning task is a major open challenge (Cropper 2017; Cropper and Muggleton 2014) and is a trade-off between efficiency and expressivity: the hypothesis space grows given more metarules (Cropper and Muggleton 2014; Lin et al. 2014), so we wish to use fewer metarules, but if we use too few metarules then we lose expressivity. For instance, it is impossible to learn the *grandparent/2* relation using only metarules with monadic predicates.

In this paper, we study whether potentially infinite fragments of metarules can be logically reduced to minimal, or irreducible, finite subsets, where a fragment is a syntactically restricted subset of a logical theory (Bradley and Manna 2007).

Cropper and Muggleton (2014) first studied this problem. They used Progol's entailment reduction algorithm (Muggleton 1995) to identify entailment reduced sets of metarules, where a clause C is entailment redundant in a clausal theory $T \cup \{C\}$ when $T \models C$. To illustrate entailment redundancy, consider the following first-order clausal theory T_1 , where p, q, r , and s are first-order predicates:

$$\begin{aligned} C_1 &= p(A, B) \leftarrow q(A, B) \\ C_2 &= p(A, B) \leftarrow q(A, B), r(A) \\ C_3 &= p(A, B) \leftarrow q(A, B), r(A), s(B, C) \end{aligned}$$

In T_1 the clauses C_2 and C_3 are entailment redundant because they are both logical consequences of C_1 , i.e. $\{C_1\} \models \{C_2, C_3\}$. Because $\{C_1\}$ cannot be reduced, it is a minimal entailment reduction of T_1 .

Cropper and Muggleton showed that in some cases as few as two metarules are sufficient to entail an infinite fragment of *chained*³ second-order dyadic Datalog (Cropper and Muggleton 2014). They also showed that learning with minimal sets of metarules improves predictive

¹ Metarules are also called *program schemata* (Flener 1996), *second-order schemata* (De Raedt and Bruynooghe 1992), and *clause templates* (Albarghouthi et al. 2017), amongst many other names.

² The fully quantified rule is $\exists P \exists Q \exists R \forall A \forall B \forall C P(A, B) \leftarrow Q(A, C), R(C, B)$.

³ A chained dyadic Datalog clause has the restriction that every first-order variable in a clause appears in exactly two literals and a path connects every literal in the body of C to the head of C . In other words, a chained dyadic Datalog clause has the form $P_0(X_0, X_1) \leftarrow P_1(X_0, X_2), P_2(X_2, X_3), \dots, P_n(X_n, X_1)$ where the order of the arguments in the literals does not matter.

accuracies and reduces learning times compared to non-minimal sets. To illustrate how a finite subset of metarules could entail an infinite set, consider the set of metarules with only monadic literals and a single first-order variable A :

$$\begin{aligned} M_1 &= P(A) \leftarrow T_1(A) \\ M_2 &= P(A) \leftarrow T_1(A), T_2(A) \\ M_3 &= P(A) \leftarrow T_1(A), T_2(A), T_3(A) \\ &\dots \\ M_n &= P(A) \leftarrow T_1(A), T_2(A), \dots, T_n(A) \\ &\dots \end{aligned}$$

Although this set is infinite it can be entailment reduced to the single metarule M_1 because it implies the rest of the theory.

However, in this paper, we claim that entailment reduction is not always the most appropriate form of reduction. For instance, suppose you want to learn the *father/2* relation given the background relations *parent/2*, *male/1*, and *female/1*. Then a suitable hypothesis is:

$$\textit{father}(A, B) \leftarrow \textit{parent}(A, B), \textit{male}(A)$$

To learn such a hypothesis one would need a metarule of the form $P(A, B) \leftarrow Q(A, B), R(A)$. Now suppose you have the metarules:

$$\begin{aligned} M_1 &= P(A, B) \leftarrow Q(A, B) \\ M_2 &= P(A, B) \leftarrow Q(A, B), R(A) \end{aligned}$$

Running entailment reduction on these metarules would remove M_2 because it is a logical consequence of M_1 . However, it is impossible to learn the intended *father/2* relation given only M_1 . As this example shows, entailment reduction can be too strong because it can remove metarules necessary to specialise a clause, where M_2 can be seen as a specialisation of M_1 .

To address this issue, we introduce *derivation reduction*, a new form of reduction based on derivations, which we claim is a more suitable form of reduction for reducing sets of metarules. Let \vdash represent derivability in SLD-resolution⁴ (Kowalski 1974), then a Horn clause C is derivationally redundant in a Horn theory $T \cup \{C\}$ when $T \vdash C$. A Horn theory is derivationally irreducible if it contains no derivationally redundant clauses. To illustrate the difference between entailment and derivation reduction, consider the metarules:

$$\begin{aligned} M_1 &= P(A, B) \leftarrow Q(A, B) \\ M_2 &= P(A, B) \leftarrow Q(A, B), R(A) \\ M_3 &= P(A, B) \leftarrow Q(A, B), R(A, B) \\ M_4 &= P(A, B) \leftarrow Q(A, B), R(A, B), S(A, B) \end{aligned}$$

Running entailment reduction on these metarules would result in the reduction $\{M_1\}$ because M_1 entails the rest of the theory. Likewise, running *subsumption reduction* (Plotkin 1971) (described in detail in Sect. 3.5) would also result in the reduction $\{M_1\}$. By contrast, running derivation reduction would only remove M_4 because it can be derived by self-resolving M_3 . The remaining metarules M_2 and M_3 are not derivationally redundant because there is no way to derive them from the other metarules.

⁴ We use \vdash to represent derivability of both first-order and second-order clauses. In practice we reason about second-order clauses using first-order resolution via encapsulation (Cropper and Muggleton 2014), which we describe in Sect. 3.3.

1.1 Contributions

In the rest of this paper, we study whether fragments of metarules relevant to ILP can be logically reduced to minimal finite subsets. We study three forms of reduction: subsumption (Robinson 1965), entailment (Muggleton 1995), and derivation (Cropper and Tourret 2018). We also study how learning with reduced sets of metarules affects learning performance. To do so, we supply Metagol (Cropper and Muggleton 2016b), a meta-interpretive learning (MIL) (Cropper and Muggleton 2016a; Muggleton et al. 2014, 2015) implementation, with different reduced sets of metarules and measure the resulting learning performance on three domains: Michalski trains (Larson and Michalski 1977), string transformations, and game rules (Cropper et al. 2019). In general, using derivation reduced sets of metarules outperforms using subsumption and entailment reduced sets, both in terms of predictive accuracies and learning times. Overall, our specific contributions are:

- We describe the logical reduction problem (Sect. 3).
- We describe subsumption and entailment reduction, and introduce derivation reduction, the problem of removing derivationally redundant clauses from a clausal theory (Sect. 3).
- We study the decidability of the three reduction problems and show, for instance, that the derivation reduction problem is undecidable for arbitrary Horn theories (Sect. 3).
- We introduce two general reduction algorithms that take a reduction relation as a parameter. We also study their complexity (Sect. 4).
- We run the reduction algorithms on finite sets of metarules to identify minimal sets (Sect. 5).
- We theoretically show whether infinite fragments of metarules can be logically reduced to finite sets (Sect. 5).
- We experimentally compare the learning performance of Metagol when supplied with reduced sets of metarules on three domains: Michalski trains, string transformations, and game rules (Sect. 6).

2 Related work

This section describes work related to this paper, mostly work on logical reduction techniques. We first, however, describe work related to MIL and metarules.

2.1 Meta-interpretive learning

Although the study of metarules has implications for many ILP approaches (Albarghouthi et al. 2017; Campero et al. 2018; Cropper and Muggleton 2019; Emde et al. 1983; Evans and Grefenstette 2018; Flener 1996; Kaminski et al. 2018; Kietz and Wrobel 1992; Muggleton et al. 2015; De Raedt and Bruynooghe 1992; Si et al. 2018; Wang et al. 2014), we focus on meta-interpretive learning (MIL), a form of ILP based on a Prolog meta-interpreter.⁵ The key difference between a MIL learner and a standard Prolog meta-interpreter is that whereas a standard Prolog meta-interpreter attempts to prove a goal by repeatedly fetching first-order clauses whose heads unify with a given goal, a MIL learner additionally attempts to prove a goal by fetching second-order metarules, supplied as background knowledge (BK), whose heads unify with the goal. The resulting meta-substitutions are saved and can be reused in

⁵ Although the MIL problem has also been encoded as an ASP problem (Kaminski et al. 2018).

later proofs. Following the proof of a set of goals, a logic program is formed by projecting the meta-substitutions onto their corresponding metarules, allowing for a form of ILP which supports predicate invention and learning recursive theories.

Most existing work on MIL has assumed suitable metarules as input to the problem, or has used metarules without any theoretical justification. In this paper, we try to address this issue by identifying minimal sets of metarules for interesting fragments of logic, such as Datalog, from which a MIL system can theoretically learn any logic program.

2.2 Metarules

McCarthy (1995) and Lloyd (2003) advocated using second-order logic to represent knowledge. Similarly, Muggleton et al. (2012) argued that using second-order representations in ILP provides more flexible ways of representing BK compared to existing methods. Metarules are second-order Horn clauses and are used as a form of declarative bias (Nédellec et al. 1996; De Raedt 2012) to determine the structure of learnable programs which in turn defines the hypothesis space. In contrast to other forms of declarative bias, such as modes (Muggleton 1995) or grammars (Cohen 1994), metarules are logical statements that can be reasoned about, such as to reason about the redundancy of sets of metarules, which we explore in this paper.

Metarules were introduced in the Blip system (Emde et al. 1983). Kietz and Wrobel (1992) studied generality measures for metarules in the RDT system. A generality order is necessary because the RDT system searches the hypothesis space (which is defined by the metarules) in a top-down general-to-specific order. A key difference between RDT and MIL is that whereas RDT requires metarules of increasing complexity (e.g. rules with an increasing number of literals in the body), MIL derives more complex metarules through SLD-resolution. This point is important because this ability allows MIL to start from smaller sets of primitive metarules. In this paper we try to identify such primitive sets.

Using metarules to build a logic program is similar to the use of refinement operators in ILP (Nienhuys-Cheng and de Wolf 1997; Shapiro 1983) to build a definite clause literal-by-literal.⁶ As with refinement operators, it seems reasonable to ask about completeness and irredundancy of a set of metarules, which we explore in this paper.

2.3 Logical redundancy

Detecting and eliminating redundancy in a clausal theory is useful in many areas of computer science. In ILP logically reducing a theory is useful to remove redundancy from a hypothesis space to improve learning performance (Cropper and Muggleton 2014; Fonseca et al. 2004). In general, simplifying or reducing a theory often makes a theory easier to understand and use, and may also have computational efficiency advantages.

2.3.1 Literal redundancy

Plotkin (1971) used subsumption to decide whether a literal is redundant in a first-order clause. Joyner (1976) independently investigated the same problem, which he called *clause condensation*, where a condensation of a clause C is a minimum cardinality subset C' of

⁶ MIL uses example driven test incorporation for finding consistent programs as opposed to the generate-and-test approach of clause refinement.

C such that $C' \models C$. Gottlob and Fermüller (1993) improved Joyner's algorithm and also showed that determining whether a clause is condensed is co-NP-complete. In contrast to removing redundant literals, we focus on removing redundant clauses.

2.3.2 Clause redundancy

Plotkin (1971) introduced methods to decide whether a clause is subsumption redundant in a first-order clausal theory. This problem has also been extensively studied in the context of first-order logic with equality due to its application in superposition-based theorem proving (Hillenbrand et al. 2013; Weidenbach and Wischniewski 2010). The same problem, and slight variants, has been extensively studied in the propositional case (Liberatore 2005, 2008). Removing redundant clauses has numerous applications, such as to improve the efficiency of SAT (Heule et al. 2015). In contrast to these works, we focus on reducing theories formed of second-order Horn clauses (without equality), which to our knowledge has not yet been extensively explored. Another difference is that we additionally study redundancy based on SLD-derivations.

Cropper and Muggleton (2014) used Progol's entailment-reduction algorithm (Muggleton 1995) to identify irreducible sets of metarules. Their approach removed entailment redundant clauses from sets of metarules. They identified theories that are (1) entailment complete for certain fragments of second-order Horn logic, and (2) irreducible. They demonstrated that in some cases as few as two clauses are sufficient to entail an infinite theory. However, they only considered small and highly constrained fragments of metarules. In particular, they focused on an *exactly-two-connected* fragment of metarules where each literal is dyadic and each first-order variable appears exactly twice in distinct literals. However, as discussed in the introduction, entailment reduction is not always the most appropriate form of reduction because it can remove metarules necessary to specialise a clause. Therefore, in this paper, we go beyond entailment reduction and introduce derivation reduction. We also consider more general fragments of metarules, such as a fragment of metarules sufficient to learn Datalog programs.

Cropper and Tourret (2018) introduced the derivation reduction problem and studied whether sets of metarules could be derivationally reduced. They considered the *exactly-two-connected* fragment previously considered by Cropper and Muggleton and a *two-connected* fragment in which every variable appears at least twice, which is analogous to our singleton-free fragment (Sect. 5.3). They used graph theoretic methods to show that certain fragments could not be completely derivationally reduced. They demonstrated on the Michalski trains dataset that the partially derivationally reduced set of metarules outperforms the entailment reduced set. In similar work Cropper and Tourret elaborated on their graph theoretic techniques and expanded the results to unconstrained resolution (Tourret and Cropper 2019).

In this paper, we go beyond the work of (Cropper and Tourret 2018) in several ways. First, we consider more general fragments of metarules, including *connected* and *Datalog* fragments. We additionally consider fragments with zero arity literals. In all cases we provide additional theoretical results showing whether certain fragments can be reduced, and, where possible, show the actual reductions. Second, Tourret and Cropper (2019) focused on derivation reduction modulo first-order variable unification, i.e. they considered the case where factorisation (Nienhuys-Cheng and de Wolf 1997) was allowed when resolving two clauses, which is not implemented in practice in current MIL systems. For this reason, although Section 5 in Tourret and Cropper (2019) and Sect. 5.1 in the present paper seemingly consider the same problem, the results are opposite to one another. Third, in addition to entailment and derivation reduction, we also consider subsumption reduction. We provide more theoretical

results on the decidability of the reduction problems, such as showing a decidable case for derivation reduction (Theorem 4). Fourth, we describe the reduction algorithms and discuss their computational complexity. Finally, we corroborate the experimental results of Cropper and Tourret on Michalski's train problem (Cropper and Tourret 2018) and provide additional experimental results on two more domains: real-world string transformations and inducing Datalog game rules from observations.

2.3.3 Theory minimisation

We focus on removing clauses from a clausal theory. A related yet distinct topic is theory minimisation where the goal is to find a minimum equivalent formula to a given input formula. This topic is often studied in propositional logic (Hemaspaandra and Schnoor 2011). The minimisation problem allows for the introduction of new clauses. By contrast, the reduction problem studied in this paper does not allow for the introduction of new clauses and instead only allows for the removal of redundant clauses.

2.3.4 Prime implicates

Implicates of a theory T are the clauses that are entailed by T and are called prime when they do not themselves entail other implicates of T . This notion differs from the subsumption and derivation reduction because it focuses on entailment, and it differs from entailment reduction because (1) the notion of a prime implicate has been studied only in propositional, first-order, and some modal logics (Bienvenu 2007; Echenim et al. 2015; Marquis 2000); (2) the generation of prime implicates allows for the introduction of new clauses in the formula.

3 Logical reduction

We now introduce the reduction problem: the problem of finding redundant clauses in a theory. We first describe the reduction problem starting with preliminaries, and then describe three instances of the problem. The first two instances are based on existing logical reduction methods: subsumption and entailment. The third instance is a new form of reduction introduced in Cropper and Tourret (2018) based on SLD-derivations.

3.1 Preliminaries

We assume familiarity with logic programming notation (Lloyd 1987) but we restate some key terminology. A *clause* is a disjunction of literals. A *clausal theory* is a set of clauses. A *Horn clause* is a clause with at most one positive literal. A Horn theory is a set of Horn clauses. A *definite clause* is a Horn clause with exactly one positive literal. A Horn clause is a *Datalog clause* if (1) it contains no function symbols, and (2) every variable that appears in the head of the clause also appears in a positive (i.e. not negated) literal in the body of the clause.⁷ We denote the powerset of the set S as 2^S .

⁷ Datalog also imposes additional constraints on negation in the body of a clause, but because we disallow negation in the body we omit these constraints for simplicity.

Table 1 Example metarules

Name	Metarule
Indent ₁	$P(A) \leftarrow Q(A)$
DIndent ₁	$P(A) \leftarrow Q(A), R(A)$
Indent ₂	$P(A, B) \leftarrow Q(A, B)$
DIndent ₂	$P(A, B) \leftarrow Q(A, B), R(A, B)$
Precon	$P(A, B) \leftarrow Q(A), R(A, B)$
Postcon	$P(A, B) \leftarrow Q(A, B), R(B)$
Curry	$P(A, B) \leftarrow Q(A, B, R)$
Chain	$P(A, B) \leftarrow Q(A, C), R(C, B)$

The letters P , Q , and R denote existentially quantified second-order variables. The letters A , B , and C denote universally quantified first-order variables

3.1.1 Metarules

Although the reduction problem applies to any clausal theory, we focus on theories formed of metarules:

Definition 1 (*Metarule*) A metarule is a second-order Horn clause of the form:

$$A_0 \leftarrow A_1, \dots, A_m$$

where each A_i is a literal of the form $P(T_1, \dots, T_n)$ where P is either a predicate symbol or a second-order variable that can be substituted by a predicate symbol, and each T_i is either a constant symbol or a first-order variable that can be substituted by a constant symbol.

Table 1 shows a selection of metarules commonly used in the MIL literature (Cropper and Muggleton 2015, 2016a, 2019; Cropper et al. 2015; Morel et al. 2019). As Definition 1 states, metarules may include predicate and constant symbols. However, we focus on the more general case where metarules only contain variables.⁸ In addition, although metarules can be any Horn clauses, we focus on definite clauses with at least one body literal, i.e. we disallow facts, because their inclusion leads to uninteresting reductions, where in almost all such cases the theories can be reduced to a single fact.⁹ We denote the infinite set of all such metarules as \mathcal{M} . We focus on *fragments* of \mathcal{M} , where a fragment is a syntactically restricted subset of a theory (Bradley and Manna 2007):

Definition 2 (*The fragment \mathcal{M}_m^a*) We denote as \mathcal{M}_m^a the fragment of \mathcal{M} where each literal has arity at most a and each clause has at most m literals in the body. We replace a by the explicit set of arities when we restrict the allowed arities further.

Example 1 $\mathcal{M}_2^{[2]}$ is a subset of \mathcal{M} where each predicate has arity 2 and each clause has at most 2 body literals.

Example 2 $\mathcal{M}_m^{[2]}$ is a subset of \mathcal{M} where each predicate has arity 2 and each clause has at most m body literals.

⁸ By more general we mean we focus on metarules that are independent of any particular ILP problem with particular predicate and constant symbols.

⁹ For instance, the metarule $P(A) \leftarrow$ entails and subsumes every metarule with a monadic head.

Example 3 $\mathcal{M}_m^{\{0,2\}}$ is a subset of \mathcal{M} where each predicate has arity 0 or 2 and each clause has at most m body literals.

Example 4 $\mathcal{M}_{\{1,2\}}^a$ is a subset of \mathcal{M} where each predicate has arity at most a and each clause has either 1 or 2 body literals.

Let T be a clausal theory. Then we say that T is in the fragment \mathcal{M}_m^a if and only if each clause in T is in \mathcal{M}_m^a .

3.2 Meta-interpretive learning

In Sect. 6 we conduct experiments to see whether using reduced sets of metarules can improve learning performance. The primary purpose of the experiments is to test our claim that entailment reduction is not always the most appropriate form of reduction. Our experiments focus on MIL. For self-containment, we briefly describe MIL.

Definition 3 (*MIL input*) An MIL input is a tuple (B, E^+, E^-, M) where:

- B is a set of Horn clauses denoting background knowledge
- E^+ and E^- are disjoint sets of ground atoms representing positive and negative examples respectively
- M is a set of metarules

The MIL problem is defined from a MIL input:

Definition 4 (*MIL problem*) Given a MIL input (B, E^+, E^-, M) , the MIL problem is to return a logic program hypothesis H such that:

- $\forall c \in H, \exists m \in M$ such that $c = m\theta$, where θ is a substitution that grounds all the existentially quantified variables in m
- $H \cup B \models E^+$
- $H \cup B \not\models E^-$

We call H a solution to the MIL problem.

The metarules and background define the hypothesis space. To explain our experimental results in Sect. 6, it is important to understand the effect that metarules have on the size of the MIL hypothesis space, and thus on learning performance. The following result generalises previous results (Cropper and Muggleton 2016a; Lin et al. 2014):

Theorem 1 (*MIL hypothesis space*) Given p predicate symbols and k metarules in \mathcal{M}_m^a , the number of programs expressible with n clauses is at most $(p^{m+1}k)^n$.

Proof The number of first-order clauses which can be constructed from a \mathcal{M}_m^a metarule given p predicate symbols is at most p^{m+1} because for a given metarule there are at most $m+1$ predicate variables with at most p^{m+1} possible substitutions. Therefore the set of such clauses S which can be formed from k distinct metarules in \mathcal{M}_m^a using p predicate symbols has cardinality at most $p^{m+1}k$. It follows that the number of programs which can be formed from a selection of n clauses chosen from S is at most $(p^{m+1}k)^n$. \square

Theorem 1 shows that the MIL hypothesis space increases given more metarules. The Blumer bound (Blumer et al. 1987),¹⁰ says that given two hypothesis spaces, searching the smaller space will result in fewer errors compared to the larger space, assuming that the target hypothesis is in both spaces. This result suggests that we should consider removing redundant metarules to improve the learning performance. We explore this idea in the rest of the paper.

¹⁰ The Blumer bound is a reformulation of Lemma 2.1 in Blumer et al. (1987).

3.3 Encapsulation

To reason about metarules (especially when running the Prolog implementations of the reduction algorithms), we use a method called encapsulation (Cropper and Muggleton 2014) to transform a second-order logic program to a first-order logic program. We first define encapsulation for atoms:

Definition 5 (*Atomic encapsulation*) Let A be a second-order or first-order atom of the form $P(T_1, \dots, T_n)$. Then $enc(A) = enc(P, T_1, \dots, T_n)$ is the encapsulation of A .

For instance, the encapsulation of the atom $parent(ann, andy)$ is $enc(parent, ann, andy)$. Note that encapsulation essentially ignores the quantification of variables in metarules by treating all variables, including predicate variables, as first-order universally quantified variables of the first-order enc predicate. In particular, replacing existential quantifiers with universal quantifiers on predicate variables is fine for our work because we only reason about the form of metarules, not their semantics, i.e. we treat metarules as templates for first-order clauses. We extend atomic encapsulation to logic programs:

Definition 6 (*Program encapsulation*) The logic program $enc(P)$ is the encapsulation of the logic program P in the case $enc(P)$ is formed by replacing all atoms A in P by $enc(A)$.

For example, the encapsulation of the metarule $P(A, B) \leftarrow Q(A, C), R(C, B)$ is $enc(P, A, B) \leftarrow enc(Q, A, C), enc(R, C, B)$. We extend encapsulation to interpretations (Nienhuys-Cheng and de Wolf 1997) of logic programs:

Definition 7 (*Interpretation encapsulation*) Let I be an interpretation over the predicate and constant symbols in a logic program. Then the encapsulated interpretation $enc(I)$ is formed by replacing each atom A in I by $enc(A)$.

We now have the proposition:

Proposition 1 [Encapsulation models (Cropper and Muggleton 2014)] *The second-order logic program P has a model M if and only if $enc(P)$ has the model $enc(M)$.*

Proof Follows trivially from the definitions of encapsulated programs and interpretations. \square

We can extend the definition of entailment to logic programs:

Proposition 2 [Entailment (Cropper and Muggleton 2014)] *Let P and Q be second-order logic programs. Then $P \models Q$ if and only if every model $enc(M)$ of $enc(P)$ is also a model of $enc(Q)$.*

Proof Follows immediately from Proposition 1. \square

These results allow us to reason about metarules using standard first-order logic. In the rest of the paper all the reasoning about second-order theories is performed at the first-order level. However, to aid the readability we continue to write non-encapsulated metarules in the rest of the paper, i.e. we will continue to refer to sets of metarules as second-order theories.

3.4 Logical reduction problem

We now describe the logical reduction problem. For the clarity of the paper, and to avoid repeating definitions for each form of reduction that we consider (entailment, subsumption, and derivability), we describe a general reduction problem which is parametrised by a binary relation \sqsubseteq defined over any clausal theory, although in the case of derivability, \sqsubseteq is in fact only defined over Horn clauses. Our only constraint on the relation \sqsubseteq is that if $A \sqsubseteq B$, $A \subseteq A'$ and $B' \subseteq B$ then $A' \sqsubseteq B'$. We first define a redundant clause:

Definition 8 (*\sqsubseteq -redundant clause*) The clause C is \sqsubseteq -redundant in the clausal theory $T \cup \{C\}$ whenever $T \sqsubseteq \{C\}$.

In a slight abuse of notation, we allow Definition 8 to also refer to a single clause, i.e. in our notation $T \sqsubseteq C$ is the same as $T \sqsubseteq \{C\}$. We define a reduced theory:

Definition 9 (*\sqsubseteq -reduced theory*) A clausal theory is \sqsubseteq -reduced if and only if it is finite and it does not contain any \sqsubseteq -redundant clauses.

We define the input to the reduction problem:

Definition 10 (*\sqsubseteq -reduction input*) A reduction input is a pair (T, \sqsubseteq) where T is a clausal theory and \sqsubseteq is a binary relation over a clausal theory.

Note that a reduction input may (and often will) be an infinite clausal theory. We define the reduction problem:

Definition 11 (*\sqsubseteq -reduction problem*) Let (T, \sqsubseteq) be a reduction input. Then the \sqsubseteq -reduction problem is to find a finite theory $T' \subseteq T$ such that (1) $T' \sqsubseteq T$ (i.e. $T' \sqsubseteq C$ for every clause C in T), and (2) T' is \sqsubseteq -reduced. We call T' a \sqsubseteq -reduction.

Although the input to a \sqsubseteq -reduction problem may contain an infinite theory, the output (a \sqsubseteq -reduction) must be a finite theory. We also introduce a variant of the \sqsubseteq -reduction problem where the reduction must obey certain syntactic restrictions:

Definition 12 (*\mathcal{M}_m^a - \sqsubseteq -reduction problem*) Let $(T, \sqsubseteq, \mathcal{M}_m^a)$ be a triple, where the first two elements are as in a standard reduction input and \mathcal{M}_m^a is a target reduction theory. Then the \mathcal{M}_m^a - \sqsubseteq -reduction problem is to find a finite theory $T' \subseteq T$ such that (1) T' is a \sqsubseteq -reduction of T , and (2) T' is in \mathcal{M}_m^a .

3.5 Subsumption reduction

The first form of reduction we consider is based on subsumption, which, as discussed in Sect. 2, is often used to eliminate redundancy in a clausal theory:

Definition 13 (*Subsumption*) A clause C subsumes a clause D , denoted as $C \preceq D$, if there exists a substitution θ such that $C\theta \subseteq D$.

Note that if a clause C subsumes a clause D then $C \models D$ (Robinson 1965). However, if $C \models D$ then it does not necessarily follow that $C \preceq D$. Subsumption can therefore be seen as being weaker than entailment. Whereas checking entailment between clauses is undecidable (Church 1936), Robinson (1965) showed that checking subsumption between

clauses is decidable [although in general deciding subsumption is a NP-complete problem (Nienhuys-Cheng and de Wolf 1997)].

If T is a clausal theory then the pair (T, \preceq) is an input to the \sqsubset -reduction problem, which leads to the *subsumption reduction* problem (S-reduction problem). We show that the S-reduction problem is decidable for finite theories:

Proposition 3 (Finite S-reduction problem decidability) *Let T be a finite theory. Then the corresponding S-reduction problem is decidable.*

Proof We can enumerate each element T' of 2^T in ascending order on the cardinality of T' . For each T' we can check whether T' subsumes T , which is decidable because subsumption between clauses is decidable. If T' subsumes T then we correctly return T' ; otherwise we continue to enumerate. Because the set 2^T is finite the enumeration must halt. Because the set 2^T contains T the algorithm will in the worst-case return T . Thus the problem is decidable. \square

3.6 Entailment reduction

As mentioned in the introduction, Cropper and Muggleton (2014) previously used entailment reduction (Muggleton 1995) to reduce sets of metarules using the notion of an entailment redundant clause:

Definition 14 (*E-redundant clause*) The clause C is entailment redundant (E-redundant) in the clausal theory $T \cup \{C\}$ whenever $T \models C$.

If T is a clausal theory then the pair (T, \models) is an input to the \sqsubset -reduction problem, which leads to the entailment reduction problem (E-reduction). We show the relationship between an E- and a S-reduction:

Proposition 4 *Let T be a clausal theory, T_S be a S-reduction of T , and T_E be an E-reduction of T . Then $T_E \models T_S$.*

Proof Assume the opposite, i.e. $T_E \not\models T_S$. This assumption implies that there is a clause $C \in T_S$ such that $T_E \not\models C$. By the definition of S-reduction, T_S is a subset of T so C must be in T , which implies that $T_E \not\models T$. But this contradicts the premise that T_E is an E-reduction of T . Therefore the assumption cannot hold, and thus $T_E \models T_S$. \square

We show that the E-reduction problem is undecidable for arbitrary clausal theories:

Proposition 5 (E-reduction problem clausal decidability) *The E-reduction problem for clausal theories is undecidable.*

Proof Follows from the undecidability of entailment in clausal logic (Church 1936). \square

The E-reduction problem for Horn theories is also undecidable:

Proposition 6 (E-reduction problem Horn decidability) *The E-reduction problem for Horn theories is undecidable.*

Proof Follows from the undecidability of entailment in Horn logic (Marcinkowski and Pacholski 1992). \square

The E-reduction problem is, however, decidable for finite Datalog theories:

Proposition 7 (E-reduction problem Datalog decidability) *The E-reduction problem for finite Datalog theories is decidable.*

Proof Follows from the decidability of entailment in Datalog (Dantsin et al. 2001) using a similar algorithm to the one used in the proof of Proposition 3. \square

3.7 Derivation reduction

As mentioned in the introduction, entailment reduction can be too strong a form of reduction. We therefore describe a new form of reduction based on derivability (Cropper and Tourret 2018; Tourret and Cropper 2019). Although our notion of derivation reduction can be defined for any proof system [such as unconstrained resolution as is done in Tourret and Cropper (2019)] we focus on SLD-resolution because we want to reduce sets of metarules, which are definite clauses. We define the function $R^n(T)$ of a Horn theory T as:

$$R^0(T) = T$$

$$R^n(T) = \{C \mid C_1 \in R^{n-1}(T), C_2 \in T, C \text{ is the binary resolvent of } C_1 \text{ and } C_2\}$$

We use this function to define the Horn closure of a Horn theory:

Definition 15 (*Horn closure*) The Horn closure $R^*(T)$ of a Horn theory T is:

$$\bigcup_{n \in \mathbb{N}} R^n(T)$$

We state our notion of derivability:

Definition 16 (*Derivability*) A Horn clause C is derivable from the Horn theory T , written $T \vdash C$, if and only if $C \in R^*(T)$.

We define a *derivationally redundant* (D-redundant) clause:

Definition 17 (*D-redundant clause*) A clause C is derivationally redundant in the Horn theory $T \cup \{C\}$ if $T \vdash C$.

Let T be a Horn theory, then the pair (T, \vdash) is an input to the \square -reduction problem, which leads to the *derivation reduction* problem (D-reduction problem). Note that a theory can have multiple D-reductions. For instance, consider the theory T :

$$C_1 = P(A, B) \leftarrow Q(B, A)$$

$$C_2 = P(A, B) \leftarrow Q(A, C), R(C, B)$$

$$C_3 = P(A, B) \leftarrow Q(C, A), R(C, B)$$

One D-reduction of T is $\{C_1, C_2\}$ because we can resolve the first body literal of C_2 with C_1 to derive C_3 (up to variable renaming). Another D-reduction of T is $\{C_1, C_3\}$ because we can likewise resolve the first body literal of C_3 with C_1 to derive C_2 .

We can show the relationship between E- and D-reductions by restating the notion of a SLD-deduction (Nienhuys-Cheng and de Wolf 1997):

Definition 18 [*SLD-deduction* (Nienhuys-Cheng and de Wolf 1997)] Let T be a Horn theory and C be a Horn clause. Then there exists a SLD-deduction of C from T , written $T \vdash_d C$, if C is a tautology or if there exists a clause D such that $T \vdash D$ and D subsumes C .

We can use the *subsumption theorem* (Nienhuys-Cheng and de Wolf 1997) to show the relationship between SLD-deductions and logical entailment:

Theorem 2 [SLD-subsumption theorem (Nienhuys-Cheng and de Wolf 1997)] *Let T be a Horn theory and C be a Horn clause. Then $T \models C$ if and only if $T \vdash_a C$.*

We can use this result to show the relationship between an E- and a D-reduction:

Proposition 8 *Let T be a Horn theory, T_E be an E-reduction of T , and T_D be a D-reduction of T . Then $T_E \models T_D$.*

Proof Follows from the definitions of E-reduction and D-reduction because an E-reduction can be obtained from a D-reduction with an additional subsumption check. \square

We also use the SLD-subsumption theorem to show that the D-reduction problem is undecidable for Horn theories:

Theorem 3 (D-reduction problem Horn decidability) *The D-reduction problem for Horn theories is undecidable.*

Proof Assume the opposite, that the problem is decidable, which implies that $T \vdash C$ is decidable. Since $T \vdash C$ is decidable and subsumption between Horn clauses is decidable (Garey and Johnson 1979), then finding a SLD-deduction is also decidable. Therefore, by the SLD-subsumption theorem, entailment between Horn clauses is decidable. However, entailment between Horn clauses is undecidable (Schmidt-Schauß 1988), so the assumption cannot hold. Therefore, the problem must be undecidable. \square

However, the D-reduction problem is decidable for any fragment \mathcal{M}_m^a (e.g. definite Datalog clauses where each clause has at least one body literal, with additional arity and body size constraints). To show this result, we first introduce two lemmas:

Lemma 1 *Let D , C_1 , and C_2 be definite clauses with m_d , m_{c_1} , and m_{c_2} body literals respectively, where m_d , m_{c_1} , and $m_{c_2} > 0$. If $\{C_1, C_2\} \vdash D$ then $m_{c_1} \leq m_d$ and $m_{c_2} \leq m_d$.*

Proof Follows from the definitions of SLD-resolution (Nienhuys-Cheng and de Wolf 1997). \square

Note that Lemma 1 does not hold for unconstrained resolution because it allows for factorisation (Nienhuys-Cheng and de Wolf 1997). Lemma 1 also does not hold when facts (bodyless definite clauses) are allowed because they would allow for resolvents that are smaller in body size than one of the original two clauses.

Lemma 2 *Let \mathcal{M}_m^a be a fragment of metarules. Then \mathcal{M}_m^a is finite up to variable renaming.*

Proof Any literal in \mathcal{M}_m^a has at most a first-order variables and 1 second-order variable, so any literal has at most $a + 1$ variables. Any metarule has at most m body literals plus the head literal, so any metarule has at most $m + 1$ literals. Therefore, any metarule has at most $((a + 1)(m + 1))$ variables. We can arrange the variables in at most $((a + 1)(m + 1))!$ ways, so there are at most $((a + 1)(m + 1))!$ metarules in \mathcal{M}_m^a up to variable renaming. Thus \mathcal{M}_m^a is finite up to variable renaming. \square

Note that the bound in the proof of Lemma 2 is a worst-case result. In practice there are fewer usable metarules because we consider fragments of constrained theories, thus not all clauses are admissible, and in all cases the order of the body literals is irrelevant. We use these two lemmas to show that the D-reduction problem is decidable for \mathcal{M}_m^a :

Theorem 4 (\mathcal{M}_m^a -D-reduction problem decidability) *The D-reduction problem for theories included in \mathcal{M}_m^a is decidable.*

Proof Let T be a finite clausal theory in \mathcal{M}_m^a and C be a definite clause with $n > 0$ body literals. The problem is whether $T \vdash C$ is decidable. By Lemma 1, we cannot derive C from any clause which has more than n body literals. We can therefore restrict the resolution closure $R^*(T)$ to only include clauses with body lengths less than or equal to n . In addition, by Lemma 2 there are only a finite number of such clauses so we can compute the fixed-point of $R^*(T)$ restricted to clauses of size smaller or equal to n in a finite amount of steps and check whether C is in the set. If it is then $T \vdash C$; otherwise $T \not\vdash C$. \square

3.8 k-Derivable clauses

Propositions 3 and 7 and Theorem 4 show that the \square -reduction problem is decidable under certain conditions. However, as we will shown in Sect. 4, even in decidable cases, solving the \square -reduction problem is computationally expensive. We therefore solve restricted k-bounded versions of the E- and D-reduction problems, which both rely on SLD-derivations. Specifically, we focus on resolution depth-limited derivations using the notion of *k-derivability*:

Definition 19 (*k-derivability*) Let k be a natural number. Then a Horn clause C is k-derivable from the Horn theory T , written $T \vdash_k C$, if and only if $C \in R^k(T)$.

The definitions for k-bounded E- and D-reductions follow from this definition but are omitted for brevity. In Sect. 4 we introduce a general algorithm (Algorithm 1) to solve the S-reduction problem and k-bounded E- and D-reduction problems.

4 Reduction algorithms

In Sect. 5 we logically reduce sets of metarules. We now describe the reduction algorithms that we use.

4.1 \square -Reduction algorithm

The reduce algorithm (Algorithm 1) shows a general \square -reduction algorithm that solves the \square -reduction problem (Definition 11) when the input theory is finite.¹¹ We ignore cases where the input is infinite because of the inherent undecidability of the problem. Algorithm 1 is largely based on Plotkin's clausal reduction algorithm (Plotkin 1971). Given a finite clausal theory T and a binary relation \square , the algorithm repeatedly tries to remove a \square -redundant clause in T . If it cannot find a \square -redundant clause, then it returns the \square -reduced theory. Note that since derivation reduction is only defined over Horn theories, in a \vdash -reduction input (T, \vdash) , the theory T has to be Horn. We show total correctness of the algorithm:

¹¹ In practice we use more efficient algorithms for each approach. For instance, in the derivation reduction Prolog implementation we use the knowledge gained from Lemma 1 to add pruning so as to ignore clauses that are too large to be useful to check whether a clause is derivable.

Proposition 9 (Algorithm 1 total correctness) *Let (T, \sqsubset) be a \sqsubset -reduction input where T is finite. Let the corresponding \sqsubset -reduction problem be decidable. Then Algorithm 1 solves the \sqsubset -reduction problem.*

Proof Trivial by induction on the size of T . □

Algorithm 1 \sqsubset -reduction algorithm

```

1  func reduce( $T, \sqsubset$ ):
2    if  $|T| < 2$ :
3      return  $T$ 
4    if there is a clause  $C$  in  $T$  such that  $T \setminus \{C\} \sqsubset C$ :
5      return reduce( $T \setminus \{C\}, \sqsubset$ )
6    else:
7      return  $T$ 

```

Note that Proposition 9 assumes that the given reduction problem is decidable and that the input theory is finite. If you call Algorithm 1 with an arbitrary clausal theory and the \models relation then it will not necessarily terminate. We can call Algorithm 1 with specific binary relations, where each variation has a different time-complexity. Table 2 shows different ways of calling Algorithm 1 with their corresponding time complexities, where we assume finite theories as input. We show the complexity of calling Algorithm 1 with the subsumption relation:

Proposition 10 (S-reduction complexity) *If T is a finite clausal theory then calling Algorithm 1 with (T, \preceq) requires at most $O(|T|^3)$ calls to a subsumption algorithm.*

Proof For every clause in T the algorithm checks whether any other clause in T subsumes C which requires at most $O(|T|^2)$ calls to a subsumption algorithm. If any clause C is found to be S-redundant then the algorithm repeats the procedure on the theory $(T \setminus \{C\})$, so overall the algorithm requires at most $O(|T|^3)$ calls to a subsumption algorithm. □

Note that a more detailed analysis of calling Algorithm 1 with the subsumption relation would depend on the subsumption algorithm used, which is an NP-complete problem (Garey and Johnson 1979). We show the complexity of calling Algorithm 1 with the k-bounded entailment relation:

Proposition 11 (k-bounded E-reduction complexity) *If T is a finite Horn theory and k is a natural number then calling Algorithm 1 with (T, \models_k) requires at most $O(|T|^{k+2})$ resolutions.*

Proof In the worst case the derivation check (line 4) requires searching the whole SLD-tree which has a maximum branching factor $|T|$ and a maximum depth k and takes $O(|T|^k)$ steps. The algorithm potentially does this step for every clause in T so the complexity of this step is $O(|T|^{k+1})$. The algorithm has to perform this check for every clause in T with an overall worst-case complexity $O(|T|^{k+2})$. □

The complexity of calling Algorithm 1 with the k-derivation relation is identical:

Proposition 12 (k-bounded D-reduction complexity) *Let T be a finite Horn theory and k be a natural number then calling Algorithm 1 with (T, \vdash_k) requires at most $O(|T|^{k+2})$ resolutions.*

Proof Follows using the same reasoning as Proposition 11. □

Table 2 Outputs and complexity of Algorithm 1 for different input relations and an arbitrary finite clausal theory T

Relation	Output	Complexity
\leq	S-reduction	$O(T ^3)$
\models	E-reduction	Undecidable
\models_k	k-E-reduction	$O(T ^{k+2})$
\vdash	D-reduction	Undecidable
\vdash_k	k-D-reduction	$O(T ^{k+2})$

The time complexities are a function of the size of the given theory, denoted by $|T|$

4.2 \mathcal{M}_m^a - \sqsubset -reduction algorithm

Although Algorithm 1 solves the \sqsubset -reduction problem, it does not solve the \mathcal{M}_m^a -reduction problem (Definition 12). For instance, suppose you have the following theory T in \mathcal{M}_4^2 :

$$\begin{aligned}
 M_1 &= P(A, B) \leftarrow Q(B, A) \\
 M_2 &= P(A, B) \leftarrow Q(A, A), R(B, B) \\
 M_3 &= P(A, B) \leftarrow Q(A, C), R(B, C) \\
 M_4 &= P(A, B) \leftarrow Q(B, C), R(A, D), S(A, D), T(B, C)
 \end{aligned}$$

Suppose you want to know whether T can be E-reduced to \mathcal{M}_2^2 . Then calling Algorithm 1 with (T, \models) (i.e. the entailment relation) will return $T' = \{M_1, M_4\}$ because: $M_4 \models M_2$,¹² $M_4 \models M_3$,¹³ and $\{M_1, M_4\}$ cannot be further E-reduced.

Although T' is an E-reduction of T , it is not in \mathcal{M}_2^2 because M_4 is not in \mathcal{M}_2^2 . However, the theory T can be \mathcal{M}_2^2 -E-reduced to $\{M_1, M_2, M_3\}$ because $\{M_2, M_3\} \models M_4$,¹⁴ and $\{M_1, M_2, M_3\}$ cannot be further reduced. In general, let T be a theory in \mathcal{M}_m^a and an T' be an E-reduction of T , then T' is not necessarily in \mathcal{M}_m^a .

Algorithm 2 overcomes this limitation of Algorithm 1. Given a finite clausal theory T , a binary relation \sqsubset , and a reduction fragment \mathcal{M}_m^a , Algorithm 2 determines whether there is a \sqsubset -reduction of T in \mathcal{M}_m^a . If there is, it returns the reduced theory; otherwise it returns false. In other words, Algorithm 2 solves the \mathcal{M}_m^a - \sqsubset -reduction problem. We show total correctness of Algorithm 2:

¹² Rename the variables in M_4 to form $M'_4 = P_0(X_1, X_2) \leftarrow P_1(X_2, X_3), P_2(X_1, X_4), P_3(X_1, X_4), P_4(X_2, X_3)$. Then $M'_4\theta = P(A, B) \leftarrow R(B, B), Q(A, A), Q(A, A), R(B, B)$ where $\theta = \{P_0/P, P_1/R, P_2/Q, P_3/Q, P_4/R, X_1/A, X_2/B, X_3/B, X_4/A\}$. It follows that $M'_4\theta \subseteq M_2$, so $M_4 \leq M_2$, which in turn implies $M_4 \models M_2$.

¹³ Rename the variables in M_4 to form $M'_4 = P_0(X_1, X_2) \leftarrow P_1(X_2, X_3), P_2(X_1, X_4), P_3(X_1, X_4), P_4(X_2, X_3)$. Then $M'_4\theta = P(A, B) \leftarrow R(B, C), Q(A, C), Q(A, C), R(B, C)$ where $\theta = \{P_0/P, P_1/R, P_2/Q, P_3/Q, P_4/R, X_1/A, X_2/B, X_3/C, X_4/C\}$. It follows that $M'_4\theta \subseteq M_3$, so $M_4 \leq M_3$, which in turn implies $M_4 \models M_3$.

¹⁴ Rename the variables in M_3 to form $M'_3 = P_0(X, Y) \leftarrow P_1(X, Z), P_2(Y, Z)$. Resolve the first body literal of M_2 with M_3 to form $R_1 = P(A, B) \leftarrow P_1(A, Z), P_2(A, Z), R(B, B)$. Rename the variables P_1 to P_3 , P_2 to P_4 , and Z to Z_1 in R_1 (to standardise apart the variables) to form $R_2 = P(A, B) \leftarrow P_3(A, Z_1), P_4(A, Z_1), R(B, B)$. Resolve the last body literal of R_2 with M'_3 to form $R_3 = P(A, B) \leftarrow P_3(A, Z_1), P_4(A, Z_1), P_1(B, Z), P_2(B, Z)$. Rename the variables Z_1 to D , Z to C , P_3 to R , P_4 to S , P_1 to Q , and P_2 to T in R_3 to form $R_4 = P(A, B) \leftarrow R(A, D), S(A, D), Q(B, C), T(B, C)$. Thus, $R_4 = M_4$, so it follows that $\{M_2, M_3\} \models M_4$.

Proposition 13 (Algorithm 2 correctness) *Let $(T, \sqsubset, \mathcal{M}_m^a)$ be a \mathcal{M}_m^a - \sqsubset -reduction input. If the corresponding \sqsubset -reduction problem is decidable then Algorithm 2 solves the corresponding \mathcal{M}_m^a - \sqsubset -reduction problem.*

Sketch Proof We provide a sketch proof for brevity. We need to show that the function `aux` correctly determines whether $B \sqsubset T$, which we can show by induction on the size of T . Assuming `aux` is correct, then if T can be reduced to B , the `mreduce` function calls Algorithm 1 to reduce B , which is correct by Proposition 9. Otherwise it returns false. \square

Algorithm 2 \mathcal{M}_m^a - \sqsubset -reduction

```

1  function mreduce( $T, \sqsubset, \mathcal{M}_m^a$ )
2       $B = \{C \mid C \in T \cap \mathcal{M}_m^a\}$ 
3      if aux( $T, \sqsubset, B$ ):
4          return reduce( $B, \sqsubset$ )
5      return false
6
7  function aux( $T, \sqsubset, B$ )
8      if  $|T| == 0$ :
9          return true
10     pick any  $C$  in  $T$ 
11      $T' = T \setminus \{C\}$ 
12     if  $B \sqsubset C$ :
13         return aux( $T', \sqsubset, B$ )
14     return false

```

5 Reduction of metarules

We now logically reduce fragments of metarules. Given a fragment \mathcal{M}_m^a and a reduction operator \sqsubset , we have three main goals:

G1: identify a \mathcal{M}_k^a - \sqsubset -reduction of \mathcal{M}_m^a for some k as small as possible

G2: determine whether $\mathcal{M}_2^a \sqsubset \mathcal{M}_\infty^a$

G3: determine whether \mathcal{M}_∞^a has any (finite) \sqsubset -reduction

We work on these goals for fragments of \mathcal{M}_m^a relevant to ILP. Table 3 shows the four fragments and their main restrictions. The subsequent sections precisely describe the fragments.

Our first goal (**G1**) is to essentially minimise the number of body literals in a set of metarules, which can be seen as trying to enforce an Occamist bias. We are particularly interested reducing sets of metarules to fragments with at most two body literals because $\mathcal{M}_2^{(2)}$ augmented with one function symbol has universal Turing machine expressivity (Tärnlund 1977). In addition, previous work on MIL has almost exclusively used metarules from the fragment \mathcal{M}_2^a . Our second goal (**G2**) is more general and concerns reducing an infinite set of metarules to \mathcal{M}_2^a . Our third goal (**G3**) is similar, but is about determining whether an infinite set of metarules has any finite reduction.

We work on the goals by first applying the reduction algorithms described in the previous section to finite fragments restricted to 5 body literals (i.e. \mathcal{M}_5^a). This value gives us a sufficiently large set of metarules to reduce but not too large that the reduction problem is intractable. When running the E- and D-reduction algorithms (both k -bounded), we use a

Table 3 The four main fragments of \mathcal{M} that we consider

Fragment	Description
\mathcal{C}	Connected clauses
\mathcal{D}	Connected Datalog clauses
\mathcal{K}	Connected Datalog clauses without singleton variables
\mathcal{U}	Connected Datalog clauses without duplicate variables

resolution-depth bound of 7, which is the largest value for which the algorithms terminate in reasonable time.¹⁵ After applying the reduction algorithms to the finite fragments, we then try to solve **G2** by extrapolating the results to the infinite case (i.e. \mathcal{M}_∞^a). In cases where $\mathcal{M}_2^a \not\sqsubseteq \mathcal{M}_\infty^a$, we then try to solve **G3** by seeing whether there exists any natural number k such that $\mathcal{M}_k^a \sqsubset \mathcal{M}_\infty^a$.

5.1 Connected (\mathcal{C}_m^a) results

We first consider a general fragment of metarules. The only constraint is that we follow the standard ILP convention (Cropper and Muggleton 2014; Evans and Grefenstette 2018; Gottlob et al. 1997; Nienhuys-Cheng and de Wolf 1997) and focus on connected clauses¹⁶:

Definition 20 (*Connected clause*) A clause is connected if the literals in the clause cannot be partitioned into two sets such that the variables appearing in the literals of one set are disjoint from the variables appearing in the literals of the other set.

The following clauses are all connected:

$$\begin{aligned} P(A) &\leftarrow Q(A) \\ P(A, B) &\leftarrow Q(A, C) \\ P(A, B) &\leftarrow Q(A, B), R(B, D), S(D, B) \end{aligned}$$

By contrast, these clauses are not connected:

$$\begin{aligned} P(A) &\leftarrow Q(B) \\ P(A, B) &\leftarrow Q(A), R(C) \\ P(A, B) &\leftarrow Q(A, B), S(C) \end{aligned}$$

We denote the connected fragment of \mathcal{M}_m^a as \mathcal{C}_m^a . Table 4 shows the maximum body size and the cardinality of the reductions obtained when applying the reduction algorithms to \mathcal{C}_5^a for different values of a . To give an idea of the scale of the reductions, the fragment $\mathcal{C}_5^{\{1,2\}}$ contains 77398 unique metarules, of which E-reduction removed all but two of them. Table 5 shows the actual reductions for $\mathcal{C}_5^{\{1,2\}}$. Reductions for other connected fragments are in Appendix “A.1”.

As Table 4 shows, all the fragments can be S- and E-reduced to \mathcal{C}_1^a . We show that in general \mathcal{C}_∞^a has a \mathcal{C}_1^a -S-reduction:

¹⁵ The entailment and derivation reduction algorithms often took 4–5 h to find a reduction. However, in some cases, typically where the fragments contained many metarules, the algorithms took around 12 h to find a reduction. By contrast, the subsumption reduction algorithm typically found a reduction in 30 min.

¹⁶ Connected clauses are also known as linked clauses (Gottlob et al. 1997).

Table 4 Cardinality and maximal body size of the reductions of \mathcal{C}_5^a

Arities a	S-reduction		E-reduction		D-reduction	
	Bodysize	Cardinality	Bodysize	Cardinality	Bodysize	Cardinality
0	1	1	1	1	2	2
1	1	1	1	1	2	2
2	1	4	1	1	5	6
0, 1	1	3	1	3	2	5
0, 2	1	6	1	3	5	21
1, 2	1	9	1	2	4	8
0, 1, 2	1	12	1	4	5	13

All the fragments can be S- and E-reduced to \mathcal{C}_1^a but they cannot all be D-reduced to \mathcal{C}_2^a

Table 5 Reductions of the connected fragment $\mathcal{C}_5^{\{1,2\}}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(B, A)$	$P(A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(A)$	$P(A, A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(B, A)$		$P(A, B) \leftarrow Q(B)$
$P(A, B) \leftarrow Q(A)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(B)$		$P(A, B) \leftarrow Q(B, B)$
$P(A, B) \leftarrow Q(A, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow Q(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow Q(C, A)$		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
$P(A, B) \leftarrow Q(C, B)$		

Theorem 5 (\mathcal{C}_∞^a S-reducibility) *For all $a > 0$, the fragment \mathcal{C}_∞^a has a \mathcal{C}_1^a -S-reduction.*

Proof Let C be any clause in \mathcal{C}_∞^a , where $a > 0$. By the definition of connected clauses there must be at least one body literal in C that shares a variable with the head literal of C . The clause formed of the head of C with the body literal directly connected to it is by definition in \mathcal{C}_1^a and clearly subsumes C . Therefore $\mathcal{C}_1^a \preceq \mathcal{C}_\infty^a$. \square

We likewise show that \mathcal{C}_∞^a always has a \mathcal{C}_1^a -E-reduction:

Theorem 6 (\mathcal{C}_∞^a E-reducibility) *For all $a > 0$, the fragment \mathcal{C}_∞^a has a \mathcal{C}_1^a -E-reduction.*

Proof Follows from Theorem 5 and Proposition 4. \square

As Table 4 shows, the fragment \mathcal{C}_5^2 could not be D-reduced to \mathcal{C}_2^2 when running the derivation reduction algorithm. However, because we run the derivation reduction algorithm with a maximum derivation depth, this result alone is not enough to guarantee that the output cannot be further reduced. Therefore, we show that \mathcal{C}_5^2 cannot be D-reduced to \mathcal{C}_2^2 :

Proposition 14 (\mathcal{C}_5^2 D-irreducibility) *The fragment \mathcal{C}_5^2 has no \mathcal{C}_2^2 -D-reduction.*

Proof We denote by $\mathcal{P}(C)$ the set of all clauses that can be obtained from a given clause C by permuting the arguments in its literals up to variable renaming. For example if $C = P(A, B) \leftarrow Q(A, C)$ then $\mathcal{P}(C) = \{(C), (P(A, B) \leftarrow Q(C, A)), (P(B, A) \leftarrow Q(A, C)), (P(B, A) \leftarrow Q(C, A))\}$ up to variable renaming.

Let C_I denote the clause $P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$. We prove that no clause in $\mathcal{P}(C_I)$ can be derived from \mathcal{C}_2^2 by induction on the length of derivations. Formally, we show that there exist no derivations of length n from \mathcal{C}_2^2 to a clause in $\mathcal{P}(C_I)$. We reason by contradiction and w.l.o.g. we consider only the clause C_I .

For the base case $n = 0$, assume that there is a derivation of length 0 from \mathcal{C}_2^2 to C_I . This assumption implies that $C_I \in \mathcal{C}_2^2$, but this clearly cannot hold given the body size of C_I .

For the general case, assume that the property holds for all $k < n$ and by contradiction consider the final inference in a derivation of length n of C_I from \mathcal{C}_2^2 . Let C_1 and C_2 denote the premises of this inference. Then the literals occurring in C_I must occur up to variable renaming in at least one of C_1 and C_2 . We consider the following cases separately.

- All the literals of C_I occur in the same premise: because of Lemma 1, this case is impossible because this premise would contain more literals than C_I (the ones from C_I plus the resolved literal).
- Only one of the literals of C_I occurs separately from the others: w.l.o.g., assume that the literal $Q(A, C)$ occurs alone in C_2 (up to variable renaming). Then C_2 must be of the form $H(A, C) \leftarrow Q(A, C)$ or $H(C, A) \leftarrow Q(A, C)$ for some H , where the H -headed literal is the resolved literal of the inference that allows the unification of A and C with their counterparts in C_1 .¹⁷ In this case, C_1 belongs to $\mathcal{P}(C_I)$ and a derivation of C_1 from \mathcal{C}_2^2 of length smaller than n exists as a strict subset of the derivation to C_I of length n . This contradicts the induction hypothesis, thus the assumed derivation of C cannot exist.
- Otherwise, the split of the literals of C_I between C_1 and C_2 is always such that at least three variables must be unified during the inference. For example, consider the case where $P(A, B) \leftarrow Q(A, C) \subset C_1$ and the set $\{R(A', D), S(B', C'), T(B', D), U(C', D)\}$ occurs in the body of C_2 (up to variable renaming). Then A', B' and C' must unify respectively with A, B and C for C_I to be derived (up to variable renaming). However the inference can at most unify two variable pairs since the resolved literal must be dyadic at most and thus this inference is impossible, a contradiction.

Thus C_I and all of $\mathcal{P}(C_I)$ cannot be derived from \mathcal{C}_2^2 . Note that, since $\mathcal{P}(C_I)$ is also neither a subset of \mathcal{C}_3^2 nor of \mathcal{C}_4^2 , this proof also shows that $\mathcal{P}(C_I)$ cannot be derived from \mathcal{C}_3^2 and from \mathcal{C}_4^2 . \square

We generalise this result to \mathcal{C}_∞^2 :

Theorem 7 (\mathcal{C}_∞^2 D-irreducibility) *The fragment \mathcal{C}_∞^2 has no D-reduction.*

Proof It is enough to prove that \mathcal{C}_∞^2 does not have a \mathcal{C}_m^2 -D-reduction for an arbitrary m because any D-reduced theory, being finite, admits a bound on the body size of the clauses it contains. Starting from C_I as defined in the proof of Proposition 14, apply the following transformation iteratively for k from 1 to m : replace the literals containing Q and R (i.e. at first $Q(A, C)$ and $R(A, D)$) with the following set of literals $Q(A, C_k), R(A, D_k), V_k(C_k, D_k), Q_k(C_k, C), R_k(D_k, D)$ where all variables and predicate variables labeled with k are new.

¹⁷ Those are the only options to derive C_I . Otherwise, e.g. with $C_2 = H(A', C') \leftarrow Q(A', D')$, the resulting clause is not C_I because D' is not unified with any of the variables in C_1 (whereas A' unifies with A and C' with C), e.g. the result includes the literal $Q(A, D')$ instead of $Q(A, C)$ hence it is not C_I .

Table 6 Existence of a S-, E- or D-reduction of \mathcal{C}_∞^a to \mathcal{C}_2^a

Arity	S	E	D
1	✓	✓	✓
2	✓	✓	×
> 2	✓	✓	×

The symbol ✓ denotes that the fragment does have a reduction. The symbol × denotes that the fragment does not have a reduction

Let the resulting clause be denoted C_{I_m} . This clause is of body size $3m + 5$ and thus does not belong to \mathcal{C}_m^2 . Moreover, for the same reason that C_I cannot be derived from any $\mathcal{C}_{m'}^2$, with $m' < 5$ (see the proof of Proposition 14) C_{I_m} cannot be derived from any $\mathcal{C}_{m'}^2$ with $m' < 3m + 5$. In particular, C_{I_m} cannot be derived from \mathcal{C}_m^2 . □

Another way to generalise Proposition 14 is the following:

Theorem 8 (\mathcal{C}_∞^a D-irreducibility) *For $a \geq 2$, the fragment \mathcal{C}_∞^a has no $\mathcal{C}_{a^2+a-2}^a$ -D-reduction.*

Proof Let C_a denote the clause

$$C_a = P(A_1, \dots, A_a) \leftarrow Q_{1,1}(A_1, B_{1,1}, \dots, B_{1,a-1}) \dots Q_{1,a}(A_1, B_{a,1}, \dots, B_{a,a-1})$$

$$\dots$$

$$Q_{a,1}(A_a, B_{1,1}, \dots, B_{1,a-1}) \dots Q_{a,a}(A_a, B_{a,1}, \dots, B_{a,a-1})$$

$$R_1(B_{1,1}, \dots, B_{a,1}), \dots, R_{a-1}(B_{1,a-1}, \dots, B_{a,a-1})$$

Note that for $a = 2$, the clauses C_a and C_I from the proof of Proposition 14 coincide. In fact to show that C_a is irreducible for any a , it is enough to consider the proof of Proposition 14 where C_a is substituted to C_I and where the last case is generalised in the following way:

- the split of the literals of C_a between C_1 and C_2 is always such that at least $a + 1$ variables must be unified during the inference, which is impossible since the resolved literal can at most hold a variables.

The reason this proof holds is that any subset of C_a contains at least $a + 1$ distinct variables. Since C_a is of body size $a^2 + a - 1$, this counter-example proves that \mathcal{C}_∞^a has no $\mathcal{C}_{a^2+a-2}^a$ -D-reduction. □

Note that this is enough to conclude that \mathcal{C}_∞^a cannot be reduced to \mathcal{C}_2^a but it does not prove that \mathcal{C}_∞^a is not D-reducible.

5.1.1 Summary

Table 6 summarises our theoretical results from this section. Theorems 5 and 6 show that \mathcal{C}_∞^a can always be S- and E-reduced to \mathcal{C}_1^a respectively. By contrast, Theorem 7 shows that \mathcal{C}_∞^2 cannot be D-reduced to \mathcal{C}_2^2 . In fact, Theorem 7 says that \mathcal{C}_∞^2 has no D-reduction. Theorem 7 has direct (negative) implications for MIL systems such as Metagol and HEXMIL. We discuss these implications in more detail in Sect. 7.

Table 7 Cardinality and maximal body size of the reductions of \mathcal{D}_5^a

Arities a	S-reduction		E-reduction		D-reduction	
	Bodysize	Cardinality	Bodysize	Cardinality	Bodysize	Cardinality
0	1	1	1	1	2	2
1	1	1	1	1	2	2
2	2	4	2	2	5	10
0, 1	1	2	1	2	2	5
0, 2	2	5	2	3	5	38
1, 2	2	10	2	3	5	11
0, 1, 2	2	11	2	4	5	14

All the fragments can be S- and E-reduced to \mathcal{D}_2^a but they cannot all be D-reduced to \mathcal{D}_2^a

5.2 Datalog (\mathcal{D}_m^a) results

We now consider Datalog clauses, which are often used in ILP (Albarghouthi et al. 2017; Cropper and Muggleton 2016a; Evans and Grefenstette 2018; Kaminski et al. 2018; Muggleton et al. 2015; Si et al. 2018). The relevant Datalog restriction is that if a variable appears in the head of a clause then it must also appear in a body literal. If we look at the S-reductions of $\mathcal{C}_5^{\{1,2\}}$ in Table 5 then the clause $P(A, B) \leftarrow Q(B)$ is not a Datalog clause because the variable A appears in the head but not in the body. We denote the Datalog fragment of \mathcal{C}_m^a as \mathcal{D}_m^a . Table 7 shows the results of applying the reduction algorithms to \mathcal{D}_5^a for different values of a . Table 8 shows the reductions for the fragment $\mathcal{D}_5^{\{1,2\}}$, which are used in Experiment 3 (Sect. 6.3) to induce Datalog game rules from observations. Reductions for other Datalog fragments are in Appendix “A.2”.

We show that \mathcal{D}_∞^2 can be S-reduced to \mathcal{D}_2^2 :

Proposition 15 (\mathcal{D}_∞^2 S-reducibility) *The fragment \mathcal{D}_∞^2 has a \mathcal{D}_2^2 -S-reduction.*

Proof Follows using the same argument as in Theorem 5 but the reduction is to \mathcal{D}_2^2 instead of \mathcal{D}_1^2 . This difference is due to the Datalog constraint that states: if a variable appears in the head it must also appear in the body. For clauses with dyadic heads, if the two head argument variables occur in two distinct body literals then the clause cannot be further reduced beyond \mathcal{D}_2^2 . □

We show how this result cannot be generalised to \mathcal{D}_∞^a :

Theorem 9 (\mathcal{D}_∞^a S-irreducibility) *For $a > 0$, the fragment \mathcal{D}_∞^a does not have a \mathcal{D}_{a-1}^a -S-reduction.*

Proof As a counter-example to a \mathcal{D}_{a-1}^a -S-reduction, consider $C_a = P(X_1, \dots, X_a) \leftarrow Q_1(X_1), \dots, Q_a(X_a)$. The clause C_a does not belong to \mathcal{D}_{a-1}^a and cannot be S-reduced to it because the removal of any subset of its literals leaves argument variables in the head without their counterparts in the body. Hence, any subset of C_a does not belong to the Datalog fragment. Thus C_a cannot be subsumed by a clause in \mathcal{D}_{a-1}^a . □

However, we can show that \mathcal{D}_∞^a can always be S-reduced to \mathcal{D}_a^a :

Theorem 10 (\mathcal{D}_∞^a to \mathcal{D}_a^a S-reducibility) *For $a > 0$, the fragment \mathcal{D}_∞^a has a \mathcal{D}_a^a -S-reduction.*

Table 8 Reductions of the Datalog fragment $\mathcal{D}_5^{\{1,2\}}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, B)$	$P(A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(B, A)$	$P(A, A) \leftarrow Q(A)$
$P(A) \leftarrow Q(B, A)$	$P(A, B) \leftarrow Q(A), R(B)$	$P(A, A) \leftarrow Q(A, A)$
$P(A, A) \leftarrow Q(B, A)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(A), R(B)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow Q(A), R(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow Q(A, B)$		$P(A, B) \leftarrow Q(B, C), R(A, D)$
$P(A, B) \leftarrow Q(B), R(A, C)$		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E)$
$P(A, B) \leftarrow Q(B, A)$		$P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
$P(A, B) \leftarrow Q(B, C), R(A, D)$		$P(A, B) \leftarrow Q(B, C), R(A, D), S(C, E), T(B, F), U(D, F)$
		$P(A, B) \leftarrow Q(B, C), R(B, D), S(C, E), T(A, F), U(D, F)$

Proof To prove that \mathcal{D}_∞^a has a \mathcal{D}_a^a -S-reduction it is enough to remark that any clause in \mathcal{D}_∞^a has a subclause of body size at most a that is also in \mathcal{D}_∞^a , the worst case being clauses such as C_a where all argument variables in the head occur in a distinct literal in the body. \square

We also show that \mathcal{D}_∞^a always has a \mathcal{D}_2^a -E-reduction, starting with the following lemma:

Lemma 3 For $a > 0$ and $n \in \{1, \dots, a\}$, the clause

$$P_0(A_1, A_2, \dots, A_n) \leftarrow P_1(A_1), P_2(A_2), \dots, P_n(A_n)$$

is \mathcal{D}_2^a -E-reducible.

Proof By induction on n .

- For the base case $n = 2$, by definition \mathcal{D}_2^a contains $P_0(A_1, A_2) \leftarrow P_1(A_1), P_2(A_2)$
- For the inductive step, assume the claim holds for $n - 1$. We show it holds for n . By definition \mathcal{D}_2^a contains the clause $D_1 = P(A_1, A_2, \dots, A_n) \leftarrow P_0(A_1, A_2, \dots, A_{n-1}), P_n(A_n)$. By the inductive hypothesis, $D_2 = P_0(A_1, A_2, \dots, A_{n-1}) \leftarrow P_1(A_1), \dots, P_{n-1}(A_{n-1})$ is \mathcal{D}_2^{a-1} -E-reducible, and thus also \mathcal{D}_2^a -E-reducible. Together, D_1 and D_2 entail $D = P_0(A_1, A_2, \dots, A_n) \leftarrow P_1(A_1), P_2(A_2), \dots, P_n(A_n)$, which can be seen by resolving the literal $P_0(A_1, A_2, \dots, A_{n-1})$ from D_1 with the same literal from D_2 to derive D . Thus D is \mathcal{D}_2^a -E-reducible.

\square

Theorem 11 (\mathcal{D}_∞^a E-reducibility) For $a > 0$, the fragment \mathcal{D}_∞^a has a \mathcal{D}_2^a -E-reduction.

Table 9 Existence of a S-, E- or D-reduction of \mathcal{D}_∞^a to \mathcal{D}_2^a

Arity	S	E	D
1	✓	✓	✓
2	✓	✓	×
> 2	×	✓	×

Proof Let C be any clause in \mathcal{D}_∞^a . We denote the head of C by $P(A_1, \dots, A_n)$, where $0 < n \leq a$. The possibility that some of the A_i are equal does not impact the reasoning.

If $n = 1$, then by definition, there exists a literal L_1 in the body of C such that A_1 occurs in L_1 . It is enough to consider the clause $P(A_1) \leftarrow L_1$ to conclude, because $P(A_1)$ is the head of C and L_1 belongs to the body of C , thus $P(A_1) \leftarrow L_1$ entails C , and this clause belongs to \mathcal{D}_2^a .

In the case where $n > 1$, there must exist literals L_1, \dots, L_n in the body of C such that A_i occurs in L_i for $i \in \{1, \dots, n\}$. Consider the clause $C' = P(A_1, \dots, A_n) \leftarrow L_1, \dots, L_n$. There are a few things to stress about C' :

- The clause C' belongs to \mathcal{D}_∞^a .
- Some L_i may be identical with each other, since the A_i s may occur together in literals or simply be equal, but this scenario does not impact the reasoning.
- The clause C' entails C because C' is equivalent to a subset of C (but this subset may be distinct from C' due to C' possibly including some extra duplicated literals).

Now consider the clause $D = P(A_1, \dots, A_n) \leftarrow P_1(A_1), \dots, P_n(A_n)$. For $i \in \{1, \dots, n\}$, the clause $P_i(A_i) \leftarrow L_i$ belongs to \mathcal{D}_2^a by definition, thus $\mathcal{D}_2^a \cup D \vdash D'$ where $D' = P(A_1, \dots, A_n) \leftarrow L_1, \dots, L_n$. Moreover, by Lemma 3, D is \mathcal{D}_2^a -E-reducible, hence D' is also \mathcal{D}_2^a -E-reducible. Note that this notation hides the fact that if a variable occurs in distinct body literals L_i in C' , this connection is not captured in D' where distinct variables will occur instead, thus there is no guarantee that $D' = C'$. For example, if $C' = P(A_1, A_2) \leftarrow Q(A_1, B, A_2), R(A_2, B)$ then $D' = P(A_1, A_2) \leftarrow Q(A_1, B, A'_2), Q(A_1, B, A'_2), R(A_2, B'), R(A_2, B')$. However, it always holds that $D' \models C'$, because D' subsumes C' . In our small example, it is enough to consider the substitution $\theta = \{B'/B, A'_2/A_2\}$ to observe this. Thus by transitivity of entailment, we can conclude that C is \mathcal{D}_2^a -E-reducible. □

As Table 7 shows, not all of the fragments can be D-reduced to \mathcal{D}_2^a . In particular, the result that \mathcal{D}_∞^2 has no \mathcal{D}_2^2 -D-reduction follows from Theorem 7 because the counterexamples presented in the proof also belong to \mathcal{D}_∞^2 .

5.2.1 Summary

Table 9 summarises our theoretical results from this section. Theorem 9 shows that \mathcal{D}_∞^a never has a \mathcal{D}_{a-1}^a -S-reduction. This result differs from the connected fragment where \mathcal{C}_∞^a could always be S-reduced to \mathcal{C}_2^a . However, Theorem 9 shows that \mathcal{D}_∞^a can always be S-reduced to \mathcal{D}_a^a . As with the connected fragment, Theorem 11 shows that \mathcal{D}_∞^a can always be E-reduced to \mathcal{C}_2^a . The result that \mathcal{D}_∞^2 has no D-reduction follows from Theorem 7.

Table 10 Cardinality and maximal body size of the reductions of \mathcal{K}_5^a

Arities a	S-reduction		E-reduction		D-reduction	
	Bodysize	Cardinality	Bodysize	Cardinality	Bodysize	Cardinality
0	1	1	1	1	2	2
1	1	1	1	1	2	2
2	4	3	2	3	5	7
0, 1	1	2	1	2	2	5
0, 2	5	4	2	3	5	23
1, 2	4	8	2	4	5	8
0, 1, 2	4	9	2	5	5	11

Table 11 Reductions of the singleton-free fragment $\mathcal{K}_5^{[2]}$

S-reduction	E-reduction	D-reduction
$P(A, B) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(B, A)$	$P(A, A) \leftarrow Q(A, A)$
$P(A, B) \leftarrow Q(B, A)$	$P(A, B) \leftarrow$ $Q(A, A), R(B, B)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$	$P(A, B) \leftarrow$ $Q(A, C), R(B, C)$	$P(A, A) \leftarrow Q(A, B), R(B, B)$
		$P(A, B) \leftarrow Q(A, A), R(B, B)$
		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D),$ $S(B, C), T(B, D), U(C, D)$

5.3 Singleton-free (\mathcal{K}_m^a) results

It is common in ILP to require that all the variables in a clause appear at least twice (Cropper and Muggleton 2014; Muggleton and Feng 1990; De Raedt and Bruynooghe 1992), which essentially eliminates singleton variables. We call this fragment the *singleton-free* fragment:

Definition 21 (*Singleton-free*) A clause is singleton-free if each first-order variable appears at least twice

For example, if we look at the E-reductions of the connected fragment $\mathcal{C}_5^{[1,2]}$ shown in Table 5 then the clause $P(A) \leftarrow Q(B, A)$ is not singleton-free because the variable B only appears once. We denote the singleton-free fragment of \mathcal{D}_m^a as \mathcal{K}_m^a . Table 10 shows the results of applying the reduction algorithms to \mathcal{K}_5^a . Table 11 shows the reductions of $\mathcal{K}_5^{[2]}$. Reductions for other singleton-free fragments are in Appendix “A.3”.

Unlike in the connected and Datalog cases, the fragment $\mathcal{K}_5^{[2]}$ is no longer S-reducible to $\mathcal{K}_2^{[2]}$. We show that \mathcal{K}_∞^2 cannot be reduced to \mathcal{K}_2^2 .

Proposition 16 (\mathcal{K}_∞^2 S-reducibility) *The fragment \mathcal{K}_∞^2 does not have a \mathcal{K}_2^2 -S-reduction.*

Table 12 Existence of a S-, E- or D-reduction of \mathcal{K}_∞^a to \mathcal{K}_2^a

Arity	S	E	D
1	✓	✓	✓
2	×	✓	×
> 2	×	✓	×

Proof As a counter-example, consider the clause:

$$C = P(A, B) \leftarrow Q(A, D), R(A, D), S(B, C), T(B, C)$$

Consider removing any non-empty subset of literals from the body of C . Doing so leads to a singleton variable in the remaining clause, so it is not a singleton-free clause. Moreover, for any other clause to subsume C it must be more general than C , but that is not possible again because of the singleton-free constraint. \square

We can likewise show that this result holds in the general case:

Theorem 12 (\mathcal{K}_∞^a S-reducibility) *For $a \geq 2$, the fragment \mathcal{K}_∞^a does not have a \mathcal{K}_{2a-1}^a -S-reduction.*

Proof We generalise the clause C from the proof of Proposition 16 to define the clause $C_a = P(A_1, \dots, A_a) \leftarrow P_1(A_1, B_1), P_2(A_1, B_1), \dots, P_{2a-1}(A_a, B_a), P_{2a}(A_a, B_a)$. The same reasoning applies to C_a as to $C(= C_2)$, making C_a irreducible in \mathcal{K}_∞^a . Moreover C_a is of body size $2a$, thus C_a is a counterexample to a \mathcal{K}_{2a-1}^a -S-reduction of \mathcal{K}_∞^a . \square

However, all the fragments can be E-reduced to \mathcal{K}_2^a .

Theorem 13 (\mathcal{K}_∞^a E-reducibility) *For $a > 0$, the fragment \mathcal{K}_∞^a has a \mathcal{K}_2^a -E-reduction.*

Proof The proof of Theorem 13 is an adaptation of that of Theorem 11. The only difference is that if $n = 1$ then $P(A_1) \leftarrow L_1, L_1$ must be considered instead of $P(A_1) \leftarrow L_1$ to ensure the absence of singleton variables in the body of the clause, and for the same reason, in the general case, the clause $D' = P(A_1, \dots, A_n) \leftarrow L_1, \dots, L_n$ must be replaced by $D' = P(A_1, \dots, A_n) \leftarrow L_1, L_1, \dots, L_n, L_n$. Note that C' is not modified and thus may or may not belong to \mathcal{K}_∞^a . However, it is enough that $C' \in \mathcal{D}_\infty^a$. With these modifications, the proof carries from \mathcal{K}_∞^a to \mathcal{K}_2^a as from \mathcal{D}_∞^a to \mathcal{D}_2^a , including the results in Lemma 3. \square

5.3.1 Summary

Table 12 summarises our theoretical results from this section. Theorem 12 shows that for $a \geq 2$, the fragment \mathcal{K}_∞^a does not have a \mathcal{K}_{2a-1}^a -S-reduction. This result contrasts with the Datalog fragment where \mathcal{D}_∞^a always has a \mathcal{D}_a^a -S-reduction. As is becoming clear, adding more restrictions to a fragment typically results in less S-reducibility. By contrast, as with the connected and Datalog fragments, Theorem 13 shows that fragment \mathcal{K}_∞^a always has a \mathcal{K}_2^a -E-reduction. In addition, as with the other fragments, \mathcal{K}_∞^a has no D-reduction for $a \geq 2$.

¹⁸ Note that this proof also shows that \mathcal{K}_∞^2 does not have a \mathcal{K}_3^2 -S-reduction.

5.4 Duplicate-free (\mathcal{U}_m^a) results

The previous three fragments are general in the sense that they have been widely used in ILP. By contrast, the final fragment that we consider is of particular interest to MIL. Table 1 shows a selection of metarules commonly used in the MIL literature. These metarules have been successfully used despite no theoretical justification. However, if we consider the reductions of the three fragments so far, the *identity*, *precon*, and *postcon* metarules do not appear in any reduction. These metarules can be derived from the reductions, typically using either the $P(A) \leftarrow Q(A, A)$ or $P(A, A) \leftarrow Q(A)$ metarules. To try to identify a reduction which more closely matches the metarules shown in Table 1, we consider a fragment that excludes clauses in which a literal contains multiple occurrences of the same variable. For instance, this fragment excludes the previously mentioned metarules and also excludes the metarule $P(A, A) \leftarrow Q(B, A)$, which was in the D-reduction shown in Table 5. We call this fragment *duplicate-free*. It is a sub-fragment of \mathcal{K}_m^a and we denote it as \mathcal{U}_m^a .

Table 13 shows the reductions for the fragment $\mathcal{U}_5^{(1,2)}$. Reductions for other duplicate-free fragments are in Appendix “A.4”. As Table 13 shows, the D-reduction of $\mathcal{U}_5^{(1,2)}$ contains some metarules commonly used in the MIL literature. For instance, it contains the *identity*₁, *didentity*₂, and *precon* metarules. We use the metarules shown in Table 13 in Experiments 1 and 2 (Sects. 6.1 and 6.2) to learn Michalski trains solutions and string transformation programs respectively.

Table 14 shows the results of applying the reduction algorithms to \mathcal{U}_5^a for different values of a . All the theoretical results that hold for the singleton-free fragments hold similarly for the duplicate-free fragments for the following reasons:

- (S) The clauses in the proofs of Proposition 16 and Theorem 12 belong to \mathcal{U}_∞^a .
- (E) If the clause C considered initially in the proof of Theorem 13 belongs to \mathcal{U}_∞^a , then all the subsequent clauses in that proof are also duplicate-free.
- (D) In the proof of Theorem 7, the C_{I_m} family of clauses all belong to \mathcal{U}_∞^a .

Thus Table 12 is also a summary of the S-, E- and D-reduction results of \mathcal{U}_∞^a to \mathcal{U}_2^a .

5.5 Summary

We started this section with three goals (**G1**, **G2**, and **G3**). Table 15 summarises the results towards these goals for fragments of metarules relevant to ILP (Table 3). For **G1**, our results are mostly empirical, i.e. the results are the outputs of the reduction algorithms. For **G2**, Table 15 shows that the results are all positive for E-reduction, but mostly negative for S- and D-reduction, especially for Datalog fragments. Similarly, for **G3** the results are again positive for E-reduction but negative for S- and D-reduction for Datalog fragments. We discuss the implications of these results in Sect. 7.

6 Experiments

As explained in Sect. 1, deciding which metarules to use for a given learning task is a major open problem. The problem is the trade-off between efficiency and expressivity: the hypothesis space grows given more metarules (Theorem 1), so we wish to use fewer metarules, but if we use too few metarules then we lose expressivity. In this section we experimentally explore this trade-off. As described in Sect. 2, Cropper and Muggleton (2014) showed that

Table 13 Reductions of the fragment $\mathcal{U}_5^{\{1,2\}}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, B), R(A, B)$	$P(A) \leftarrow Q(A)$
$P(A) \leftarrow Q(A, B), R(A, B)$	$P(A, B) \leftarrow Q(B, A)$	$P(A) \leftarrow Q(A), R(A)$
$P(A, B) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(A), R(B)$	$P(A) \leftarrow Q(A, B), R(B)$
$P(A, B) \leftarrow Q(B, A)$		$P(A) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow Q(A), R(B)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(B), R(A, C), S(A, C)$		$P(A, B) \leftarrow Q(A), R(B)$
$P(A, B) \leftarrow Q(A), R(B, C), S(B, C)$		$P(A, B) \leftarrow Q(A), R(A, B)$
$P(A, B) \leftarrow Q(B, C), R(A, D), S(A, D), T(B, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E), U(E)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E), U(C, E)$

Table 14 Cardinality and body size of the reductions of \mathcal{U}_5^a

Arities a	S-reduction		E-reduction		D-reduction	
	Bodysize	Cardinality	Bodysize	Cardinality	Bodysize	Cardinality
0	1	1	1	1	2	2
1	1	1	1	1	2	2
2	4	3	5	2	5	10
0, 1	1	2	1	2	2	5
0, 2	5	4	5	3	5	38
1, 2	4	8	2	3	5	12
0, 1, 2	4	9	2	4	5	16

Table 15 Existence of a S-, E- or D-reduction of \mathcal{M}_∞^a to \mathcal{M}_2^a

Arities a	\mathcal{E}_∞^a			\mathcal{D}_∞^a			\mathcal{H}_∞^a			\mathcal{U}_∞^a		
	S	E	D	S	E	D	S	E	D	S	E	D
1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
2	✓	✓	×	✓	✓	×	×	✓	×	×	✓	×
> 2	✓	✓	×	×	✓	×	×	✓	×	×	✓	×

The symbol ✓ denotes that the fragment does have such a reduction. The symbol × denotes that the fragment does not have such a reduction

Fig. 1 An example Michalski trains target program. In the Michalski trains domain, a carriage (*car*) can be long or short. A short carriage always has two wheels. A long carriage has either two or three wheels

```
f(X):-
    has_car(X,C1),
    long(C1),
    two_wheels(C1),
    has_car(X,C2),
    long(C2),
    three_wheels(C2).
```

learning with E-reduced sets of metarules can lead to higher predictive accuracies and lower learning times compared to learning with non-E-reduced sets. However, as argued in Sect. 1, we claim that E-reduction is not always the most suitable form of reduction because it can remove metarules necessary to learn programs with the appropriate specificity. To test this claim, we now conduct experiments that compare the learning performance of Metagol 2.3.0,¹⁹ the main MIL implementation, when given different reduced sets of metarules.²⁰ We test the null hypothesis:

Null hypothesis 1 There is no difference in the learning performance of Metagol when using different reduced sets of metarules

To test this null hypothesis, we consider three domains: Michalski trains, string transformations, and game rules.

6.1 Michalski trains

In the Michalski trains problems (Larson and Michalski 1977) the task is to induce a program that distinguishes eastbound trains from westbound trains. Figure 1 shows an example target program, where the target concept ($\mathcal{E}/1$) is that the train has a long carriage with two wheels and another with three wheels.

6.1.1 Materials

To obtain the experimental data, we first generated 8 random target train programs where the programs are progressively more difficult, where difficulty is measured by the number of literals in the generated program from the easiest task T_1 to the most difficult task T_8 . Figure 2 shows the background predicates available to Metagol. We vary the metarules given to Metagol. We use the S-, E-, and D-reductions of the fragment $\mathcal{W}_5^{\{1,2\}}$ (Table 13). In addition, we also consider the $\mathcal{W}_2^{\{1,2\}}$ fragment of the D-reduction of $\mathcal{W}_5^{\{1,2\}}$, i.e. a subset of the D-reduction consisting only of metarules with at most two body literals. This fragment, which we denote as D^* , contains three fewer metarules than the D-reduction of $\mathcal{W}_5^{\{1,2\}}$. Table 16 shows this fragment.

6.1.2 Method

For each train task t_i in $\{T_1, \dots, T_8\}$:

¹⁹ <https://github.com/metagol/metagol/releases/tag/2.3.0>.

²⁰ Experimental data is available at <http://github.com/andrewcropper/mlj19-reduce>.

Fig. 2 Background relations available in the trains experiment

has_car/2	has_load/2
short/1	long/1
two_wheels/1	three_wheels/1
roof_open/1	roof_closed/1
zero_load/1	one_load/1
two_load/1	circle/1
triangle/1	rectangle/1

Table 16 The D^* fragment, which is the D-reduction of the fragment $\mathcal{W}_5^{\{1,2\}}$ restricted to the fragment $\mathcal{W}_2^{\{1,2\}}$

$P(A) \leftarrow Q(A)$	$P(A, B) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(A), R(A)$	$P(A, B) \leftarrow Q(A), R(B)$
$P(A) \leftarrow Q(A, B), R(B)$	$P(A, B) \leftarrow Q(A), R(A, B)$
$P(A) \leftarrow Q(A, B), R(A, B)$	$P(A, B) \leftarrow Q(A, B), R(A, B)$
	$P(A, B) \leftarrow Q(A, C), R(B, C)$

Table 17 Predictive accuracies when using different reduced sets of metarules on the Michalski trains problems

Task	S	E	D	D^*
T_1	100 ± 0	100 ± 0	100 ± 0	100 ± 0
T_2	100 ± 0	100 ± 0	100 ± 0	100 ± 0
T_3	68 ± 5	62 ± 5	100 ± 0	100 ± 0
T_4	75 ± 6	75 ± 6	100 ± 0	100 ± 0
T_5	92 ± 4	78 ± 6	78 ± 6	100 ± 0
T_6	52 ± 2	50 ± 0	70 ± 6	100 ± 0
T_7	95 ± 3	65 ± 5	82 ± 5	100 ± 0
T_8	55 ± 3	52 ± 2	72 ± 6	98 ± 2
Mean	80 ± 1	73 ± 2	88 ± 2	100 ± 0

Bold values denote the best performing set of metarules

1. Generate 10 training examples of t_i , half positive and half negative
2. Generate 200 testing examples of t_i , half positive and half negative
3. For each set of metarules m in the S-, E-, D-, and D^* -reductions:
 - (a) Learn a program for task t_i using the training examples and metarules m
 - (b) Measure the predictive accuracy of the learned program using the testing examples

If a program is not found in 10 min then no program is returned and every testing example is deemed to have failed. We measure mean predictive accuracies, mean learning times, and standard errors over 10 repetitions.

6.1.3 Results

Table 17 shows the predictive accuracies when learning with the different sets of metarules. The D set generally outperforms the S and E sets with a higher mean accuracy of 88% versus 80% and 73% respectively. Moreover, the D^* set easily outperforms them all with

Table 18 Learning times in seconds when using different reduced sets of metarules on the Michalski trains problems

Task	S	E	D	D*
T_1	0 ± 0	0 ± 0	0 ± 0	0 ± 0
T_2	0 ± 0	0 ± 0	0 ± 0	0 ± 0
T_3	424 ± 59	461 ± 56	0 ± 0	0 ± 0
T_4	322 ± 64	340 ± 61	0 ± 0	0 ± 0
T_5	226 ± 48	320 ± 59	361 ± 59	5 ± 2
T_6	583 ± 17	600 ± 0	429 ± 51	7 ± 2
T_7	226 ± 44	446 ± 55	243 ± 61	6 ± 1
T_8	550 ± 35	570 ± 30	361 ± 64	183 ± 40
mean	292 ± 16	342 ± 17	174 ± 16	25 ± 5

Note that the values are rounded, so 0 represents that a solution was found in under half a second

Bold values denote the best performing set of metarules

$f(A) : \text{-has_car}(A, B), f1(A, B).$ $f1(A, B) : \text{-three_wheels}(B), \text{has_car}(A, C), f2(A, C).$ $f2(A, B) : \text{-roof_open}(B), \text{has_car}(A, C), \text{has_car}(A, C).$
S program

$f(A) : \text{-has_car}(A, B), f1(A, B).$ $f1(A, B) : \text{-f2}(A), \text{three_wheels}(B).$ $f2(A) : \text{-has_car}(A, B), \text{has_car}(A, B).$
E program

$f(A) : \text{-f1}(A), f2(A).$ $f1(A) : \text{-has_car}(A, B), \text{roof_open}(B).$ $f2(A) : \text{-has_car}(A, B), \text{three_wheels}(B).$
D program

$f(A) : \text{-f1}(A), f2(A).$ $f1(A) : \text{-has_car}(A, B), \text{three_wheels}(B).$ $f2(A) : \text{-has_car}(A, B), f3(B).$ $f3(A) : \text{-long}(A), \text{two_wheels}(A).$
D* program

Fig. 3 Example programs learned by Metagol when varying the metarule set. The target program is shown in Fig. 1

a mean accuracy of 100%. A McNemar's test²¹ on the D and D^* accuracies confirmed the significance at the $p < 0.01$ level.

²¹ A statistical test on paired *nominal* data https://en.wikipedia.org/wiki/McNemar%27s_test.

Table 18 shows the corresponding learning times when using different reduces sets of metarules. The D set outperforms (has lower mean learning time) the S and E sets, and again the D^* set outperforms them all. A paired t-test²² on the D and D^* learning times confirmed the significance at the $p < 0.01$ level.

The D^* set performs particularly well on the more difficult tasks. The poor performance of the S and E sets on the more difficult tasks is for one of two reasons. The first reason is that the S - and E -reduction algorithms have removed the metarules necessary to express the target concept. This observation strongly corroborates our claim that E -reduction can be too strong because it can remove metarules necessary to specialise a clause. The second reason is that the S - and E -reduction algorithms produce sets of metarules that are still sufficient to express the target theory but doing so requires a much larger and more complex program, measured by the number of clauses needed.

The performance discrepancy between the D and D^* sets of metarules can be explained by comparing the hypothesis spaces searched. For instance, when searching for a program with 3 clauses, Theorem 1 shows that when using the D set of metarules the hypothesis space contains approximately 10^{24} programs. By contrast, when using the D^* set of metarules the hypothesis space contains approximately 10^{14} programs. As explained in Sect. 3.2, assuming that the target hypothesis is in both hypothesis spaces, the Blumer bound (Blumer et al. 1987) tells us that searching the smaller hypothesis space will result in less error, which helps to explain these empirical results. Of course, there is the potential for the D^* set to perform worse than the D set when the target theory requires the three removed metarules, but we did not observe this situation in this experiment.

Figure 3 shows the target program for T_8 and example programs learned by Metagol using the various reduced sets of metarules. Only the D^* program is success set equivalent²³ to the target program when restricted to the target predicate $\text{f}/1$. In all three cases Metagol discovered that if a carriage has three wheels then it is a long carriage, i.e. Metagol discovered that the literal $\text{long}(C2)$ is redundant in the target program. Indeed, if we unfold the D^* program to remove the invented predicates then the resulting single clause program is one literal shorter than the target program.

Overall, the results from this experiment suggest that we can reject the null hypothesis, both in terms of predictive accuracies and learning times.

6.2 String transformations

In Lin et al. (2014) and Cropper and Muggleton (2019) the authors evaluate Metagol on 17 real-world string transformation tasks using a predefined (hand-crafted) set of metarules. In this experiment, we compare learning with different metarules on an expanded dataset with 250 string transformation tasks.

6.2.1 Materials

Each string transformation task has 10 examples. Each example is an atom of the form $f(x, y)$ where f is the task name and x and y are strings. Table 19 shows task $p6$ where the goal

²² A statistical test on paired *ordinal* data <http://www.biostathandbook.com/pairedttest.html>.

²³ The success set of a logic program P is the set of ground atoms $\{A \in \text{hb}(P) \mid P \cup \{\neg A\} \text{ has a SLD-refutation}\}$, where $\text{hb}(P)$ represents the Herbrand base of the logic program P . The success set restricted to a specific predicate symbol p is the subset of the success set restricted to atoms containing the predicate symbol p .

Table 19 Examples of the $p6$ string transformation problem input–output pairs

Input	Output
Arthur Joe Juan	AJJ
Jose Larry Scott	JLS
Kevin Jason Matthew	KJM
Donald Steven George	DSG
Raymond Frank Timothy	RFT

is to learn a program that filters the capital letters from the input. We supply Metagol with dyadic background predicates, such as `tail`, `dropLast`, `reverse`, `filter_letter`, `filter_uppercase`, `dropWhile_not_letter`, `takeWhile_uppercase`. The full details can be found in the code repository. We vary the metarules given to Metagol. We use the S-, E-, and D-reductions of the fragment $\mathcal{U}_5^{[2]}$. We again also use the D-reduction of the fragment $\mathcal{U}_5^{[2]}$ restricted to the fragment $\mathcal{U}_2^{[2]}$, which is again denoted as D^* .

6.2.2 Method

Our experimental method is:

1. Sample 50 tasks Ts from the set $\{p1, \dots, p250\}$
2. For each $t \in Ts$:
 - (a) Sample 5 training examples and use the remaining examples as testing examples
 - (b) For each set of metarules m in the S-, E-, D, and D^* -reductions:
 - i. Learn a program p for task t using the training examples and metarules m
 - ii. Measure the predictive accuracy of p using the testing examples

If a program is not found in 10min then no program is returned and every testing example is deemed to have failed. We measure mean predictive accuracies, mean learning times, and standard errors over 10 repetitions.

6.2.3 Results

Table 20 shows the mean predictive accuracies and learning times when learning with the different sets of metarules. Note that we are not interested in the absolute predictive accuracy, which is limited by factors such as the low timeout and insufficiency of the BK. We are instead interested in the relative accuracies. Table 20 shows that the D set outperforms the S and E sets, with a higher mean accuracy of 33%, versus 22% and 22% respectively. The D^* set outperforms them all with a mean accuracy of 56%. A McNemar's test on the D and D^* accuracies confirmed the significance at the $p < 0.01$ level.

Table 20 shows the corresponding learning times when varying the metarules. Again, the D set outperforms the S and E sets, and again the D^* set outperforms them all. A paired t-test on the D and D^* learning times confirmed the significance at the $p < 0.01$ level.

Overall, the results from this experiment give further evidence to reject the null hypothesis, both in terms of predictive accuracies and learning times.

Table 20 Experimental results on the string transformation problems

	S	E	D	D*
Mean predictive accuracy (%)	22 ± 0	22 ± 0	32 ± 0	56 ± 1
Mean learning time (seconds)	467 ± 1	467 ± 1	407 ± 3	270 ± 3

Bold values denote the best performing set of metarules

Table 21 IGGP games used in the experiments

GT attrition	GT chicken
GT prisoner	Minimal decay
Minimal even	Multiple buttons and lights
Scissors paper stone	Untwisty corridor

Fig. 4 Target solution for the `next` predicate for the *minimal decay* game

```

next_value(X):-
    true_value(Y),
    succ(X,Y),
    does_player(noop).
next_value(5):-
    does_player(pressButton).

```

6.3 Inducing game rules

The general game playing (GGP) framework (Genesereth et al. 2005) is a system for evaluating an agent's general intelligence across a wide range of tasks. In the GGP competition, agents are tested on games they have never seen before. In each round, the agents are given the rules of a new game. The rules are described symbolically as a logic program. The agents are given a few seconds to think, to process the rules of the game, and to then start playing, thus producing game traces. The winner of the competition is the agent who gets the best total score over all the games. In this experiment, we use the IGGP dataset (Cropper et al. 2019) which inverts the GGP task: an ILP system is given game traces and the task is to learn a set of rules (a logic program) that could have produced these traces.

6.3.1 Materials

The IGGP dataset contains problems drawn from 50 games. We focus on the eight games shown in Table 21 which contain BK compatible with the metarule fragments we consider (i.e. the BK contains predicates in the fragment \mathcal{M}_m^2). The other games contain predicates with arity greater than two. Each game has four target predicates `legal`, `next`, `goal`, and `terminal`, where the arities depend on the game. Figure 4 shows the target solution for the `next` predicate for the *minimal decay* game. Each game contains training/validate/test data, composed of sets of ground atoms, in a 4:1:1 split. We vary the metarules given to Metagol. We use the S-, E-, and D-reductions of the fragment $\mathcal{D}_5^{\{1,2\}}$. We again also use the D-reduction of the fragment $\mathcal{D}_5^{\{1,2\}}$ restricted to the fragment $\mathcal{D}_2^{\{1,2\}}$, which is again denoted as D^* .

Table 22 Experimental results on the IGGP data

	S	E	D	D*
Balanced accuracy (%)	66	66	72	73
Learning time (seconds)	316	316	327	296

Bold values denote the best performing set of metarules

6.3.2 Method

The majority of game examples are negative. We therefore use *balanced accuracy* to evaluate the approaches. Given background knowledge B , sets of positive E^+ and negative E^- testing examples, and a logic program H , we define the number of positive examples as $p = |E^+|$, the number of negative examples as $n = |E^-|$, the number of true positives as $tp = |\{e \in E^+ | B \cup H \models e\}|$, the number of true negatives as $tn = |\{e \in E^- | B \cup H \not\models e\}|$, and the balanced accuracy $ba = (tp/p + tn/n)/2$.

Our experimental method is as follows. For each game g , each task g_t , and each set of metarules m in the S-, E-, D-, and D^* -reductions:

1. Learn a program p using all the training examples for g_t using the metarules m with a timeout of 10 min
2. Measure the balanced accuracy of p using the testing examples

If no program is found in 10 min then no program is returned and every testing example is deemed to have failed.

6.3.3 Results

Table 22 shows the balanced accuracies when learning with the different sets of metarules. Again, we are not interested in the absolute accuracies only the relative differences when learning using different sets of metarules. The D set outperforms the S and E sets with a higher mean accuracy of 72%, versus 66% and 66% respectively. The D^* set again outperforms them all with a mean accuracy of 73%. A McNemar's test on the D and D^* accuracies confirmed the significance at the $p < 0.01$ level. Table 22 shows the corresponding learning times when varying the metarules. Again, the D set outperforms the S and E sets, and again the D^* set outperforms them all. However, a paired t-test on the D and D^* learning times confirmed the significance only at the $p < 0.08$ level, so the difference in learning times is insignificant. Overall, the results from this experiment suggest that we can reject the null hypothesis in terms of predictive accuracies but not learning times.

7 Conclusions and further work

As stated in Sect. 1, despite the widespread use of metarules, there is little work determining which metarules to use for a given learning task. Instead, suitable metarules are assumed to be given as part of the background knowledge, or are used without any theoretical justification. Deciding which metarules to use for a given learning task is a major open challenge (Cropper 2017; Cropper and Muggleton 2014) and is a trade-off between efficiency and expressivity: the hypothesis space grows given more metarules (Cropper and Muggleton 2014; Lin et al. 2014), so we wish to use fewer metarules, but if we use too few metarules then we lose expressivity. To address this issue, Cropper and Muggleton (2014) used E-reduction on sets

of metarules and showed that learning with E-reduced sets of metarules can lead to higher predictive accuracies and lower learning times compared to learning with non-E-reduced sets. However, as we claimed in Sect. 1, E-reduction is not always the most appropriate form of reduction because it can remove metarules necessary to learn programs with the necessary specificity.

To support our claim, we have compared three forms of logical reduction: S-, E-, and D-reduction, where the latter is a new form of reduction based on SLD-derivations. We have used the reduction algorithms to reduce finite sets of metarules. Table 15 summarises the results. We have shown that many sets of metarules relevant to ILP do not have finite reductions (Theorem 7). These negative results have direct (negative) implications for MIL. Specifically, our results mean that, in certain cases, a MIL system, such as Metagol or HEXMIL (Kaminski et al. 2018), cannot be given a finite set of metarules from which it can learn any program, such as when learning arbitrary Datalog programs. The results will also likely have implications for other forms of ILP which rely on metarules.

Our experiments compared learning the performance of Metagol when using the different reduced sets of metarules. In general, using the D-reduced set outperforms both the S- and E-reduced sets in terms of predictive accuracy and learning time. Our experimental results give strong evidence to our claim. We also compared a D^* -reduced set, a subset of the D-reduced metarules, which, although derivationally incomplete, outperforms the other two sets in terms of predictive accuracies and learning times.

7.1 Limitations and future work

Theorem 7 shows that certain fragments of metarules do not have finite D-reductions. However, our experimental results show that using D-reduced sets of metarules leads to higher predictive accuracies and lower learning times compared to the other forms of reduction. Therefore, our work now opens up a new challenge of overcoming this negative theoretical result. One idea is to explore whether special metarules, such as a currying metarule (Cropper and Muggleton 2016a), could alleviate the issue.

In future work we would also like to reduce more general fragments of logic, such as triadic logics, which would allow us to tackle a wider variety of problems, such as more of the games in the IGGP dataset.

We have compared the learning performance of Metagol when using different reduced sets of metarules. However, we have not investigated whether these reductions are optimal. For instance, when considering derivation reductions, it may, in some cases, be beneficial to re-add redundant metarules to the reduced sets to avoid having to derive them through SLD-resolution. In future work, we would like to investigate identifying an optimal set of metarules for a given learning task, or preferably learning which metarules to use for a given learning task.

We have shown that although incomplete the D^* -reduced set of metarules outperforms the other reductions. In future work we would like to explore other methods which sacrifice completeness for efficiency.

We have used the logical reduction techniques to remove redundant metarules. It may also be beneficial to simultaneously reduce metarules and standard background knowledge. The idea of purposely removing background predicates is similar to dimensionality reduction, widely used in other forms of machine learning (Skillicorn 2007), but which has been under researched in ILP (Fürnkranz 1997). Initial experiments indicate that this is possible (Cropper 2017; Cropper and Muggleton 2014), and we aim to develop this idea in future work.

Acknowledgements The authors thank Stephen Muggleton and Katsumi Inoue for discussions on this topic. We especially thank Rolf Morel for valuable feedback on the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

A Detailed reduction results

A.1 Connected (\mathcal{E}_m^a) reductions

See Tables 23 and 24.

Table 23 Reductions of the connected fragment $\mathcal{E}_5^{[2]}$

S-reduction	E-reduction	D-reduction
$P(A, B) \leftarrow Q(A, C)$	$P(A, B) \leftarrow Q(B, C)$	$P(A, A) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(B, C)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(C, A)$		$P(A, B) \leftarrow Q(B, B)$
$P(A, B) \leftarrow Q(C, B)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$

Table 24 Reductions of the connected fragment $\mathcal{E}_5^{\{1,2\}}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(B, A)$	$P(A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(A)$	$P(A, A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(B, A)$		$P(A, B) \leftarrow Q(B)$
$P(A, B) \leftarrow Q(A)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(B)$		$P(A, B) \leftarrow Q(B, B)$
$P(A, B) \leftarrow Q(A, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow Q(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow Q(C, A)$		$P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
$P(A, B) \leftarrow Q(C, B)$		

A.2 Datalog (\mathcal{D}_m^a) reductions

See Tables 25, 26 and 27.

Table 25 Reductions of the Datalog fragment $\mathcal{D}_5^{\{0,1,2\}}$

S-reduction	E-reduction	D-reduction
$P \leftarrow Q$	$P \leftarrow Q$	$P \leftarrow Q$
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, B)$	$P \leftarrow Q, R$
$P(A) \leftarrow Q(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(B, A)$	$P(A, B) \leftarrow$ $Q(A), R(B)$	$P(A, A) \leftarrow Q(A)$
$P(A, A) \leftarrow$ $Q(B, A)$		$P(A, A) \leftarrow Q(A, A)$
$P(A, B) \leftarrow$ $Q(A, B)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, A)$		$P(A, B) \leftarrow Q, R(A, B)$
$P(A, B) \leftarrow$ $Q(A), R(B)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow$ $Q(A), R(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow$ $Q(B), R(A, C)$		$P(A, B) \leftarrow Q(B, C), R(A, D)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D)$		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(C, E), T(B, F), U(D, F)$
		$P(A, B) \leftarrow$ $Q(B, C), R(B, D), S(C, E), T(A, F), U(D, F)$

Table 26 Reductions of the Datalog fragment $\mathcal{D}_5^{(1,2)}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, B)$	$P(A) \leftarrow Q(B, A)$
$P(A) \leftarrow Q(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, A) \leftarrow Q(A)$
$P(A) \leftarrow Q(B, A)$	$P(A, B) \leftarrow$ $Q(A), R(B)$	$P(A, A) \leftarrow Q(A, A)$
$P(A, A) \leftarrow$ $Q(B, A)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(A), R(B)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow$ $Q(A), R(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow$ $Q(A, B)$		$P(A, B) \leftarrow Q(B, C), R(A, D)$
$P(A, B) \leftarrow$ $Q(B), R(A, C)$		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E)$
$P(A, B) \leftarrow$ $Q(B, A)$		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D)$		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(C, E), T(B, F), U(D, F)$
		$P(A, B) \leftarrow$ $Q(B, C), R(B, D), S(C, E), T(A, F), U(D, F)$

Table 27 Reductions of the Datalog fragment $\mathcal{D}_5^{(2)}$

S-reduction	E-reduction	D-reduction
$P(A, A) \leftarrow$ $Q(B, A)$	$P(A, A) \leftarrow$ $Q(B, A)$	$P(A, A) \leftarrow Q(A, A)$
$P(A, B) \leftarrow$ $Q(A, B)$	$P(A, B) \leftarrow$ $Q(B, C), R(A, D)$	$P(A, A) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, A)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow Q(B, C), R(A, D)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(C, E), T(B, F), U(D, F)$
		$P(A, B) \leftarrow$ $Q(B, C), R(B, D), S(C, E), T(A, F), U(D, F)$

A.3 Singleton-free (\mathcal{K}_m^a) results

See Tables 28, 29 and 30.

Table 28 Reductions of the singleton-free fragment $\mathcal{K}_5^{[2]}$

S-reduction	E-reduction	D-reduction
$P(A, B) \leftarrow$ $Q(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, A) \leftarrow Q(A, A)$
$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow$ $Q(A, A), R(B, B)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$	$P(A, B) \leftarrow$ $Q(A, C), R(B, C)$	$P(A, A) \leftarrow Q(A, B), R(B, B)$ $P(A, B) \leftarrow Q(A, A), R(B, B)$ $P(A, B) \leftarrow Q(A, B), R(A, B)$ $P(A, B) \leftarrow Q(A, C), R(B, C)$ $P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C),$ $T(B, D), U(C, D)$

Table 29 Reductions of the singleton-free fragment $\mathcal{K}_5^{[1,2]}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, A)$	$P(A) \leftarrow Q(A, A)$
$P(A) \leftarrow$ $Q(A, B), R(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, A) \leftarrow Q(A)$
$P(A, B) \leftarrow$ $Q(A, B)$	$P(A, B) \leftarrow$ $Q(A), R(B)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow$ $Q(A, C), R(B, C)$	$P(A, A) \leftarrow Q(A, B), R(B, B)$
$P(A, B) \leftarrow$ $Q(A), R(B)$		$P(A, B) \leftarrow Q(A, A), R(B, B)$
$P(A, B) \leftarrow$ $Q(A), R(B, C), S(B, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow$ $Q(B), R(A, C), S(A, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C),$ $T(B, D), U(C, D)$

Table 30 Reductions of the singleton-free fragment $\mathcal{K}_5^{(0,1,2)}$

S-reduction	E-reduction	D-reduction
$P \leftarrow Q$	$P \leftarrow Q$	$P \leftarrow Q$
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, A)$	$P \leftarrow Q, R$
$P(A) \leftarrow$ $Q(A, B), R(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A) \leftarrow Q(A, A)$
$P(A, B) \leftarrow$ $Q(A, B)$	$P(A, B) \leftarrow$ $Q(A), R(B)$	$P(A, A) \leftarrow Q(A)$
$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow$ $Q(A, C), R(B, C)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(A), R(B)$		$P(A, A) \leftarrow Q(A, B), R(B, B)$
$P(A, B) \leftarrow$ $Q(A), R(B, C), S(B, C)$		$P(A, B) \leftarrow Q, R(A, B)$
$P(A, B) \leftarrow$ $Q(B), R(A, C), S(A, C)$		$P(A, B) \leftarrow Q(A, A), R(B, B)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C),$ $T(B, D), U(C, D)$

A.4 Duplicate-free (\mathcal{W}_m^a) results

See Tables 31, 32 and 33.

Table 31 Reductions of the fragment $\mathcal{W}_5^{\{2\}}$

S-reduction	E-reduction	D-reduction
$P(A, B) \leftarrow$ $Q(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$	$P(A, B) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow Q(A, B), R(A, C), S(A, C)$
		$P(A, B) \leftarrow Q(A, B), R(A, C), S(C, D), T(C, D)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(A, D), T(B, C)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, C), T(B, D)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(B, D), T(C, E), U(C, E)$
		$P(A, B) \leftarrow$ $Q(B, C), R(C, D), S(A, E), T(B, E), U(C, D)$

Table 32 Reductions of the fragment $\mathcal{W}_5^{\{1,2\}}$

S-reduction	E-reduction	D-reduction
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow Q(A, B), R(A, B)$	$P(A) \leftarrow Q(A)$
$P(A) \leftarrow Q(A, B), R(A, B)$	$P(A, B) \leftarrow Q(B, A)$	$P(A) \leftarrow Q(A), R(A)$
$P(A, B) \leftarrow Q(A, B)$	$P(A, B) \leftarrow Q(A), R(B)$	$P(A) \leftarrow Q(A, B), R(B)$
$P(A, B) \leftarrow Q(B, A)$		$P(A) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow Q(A), R(B)$		$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow Q(B), R(A, C), S(A, C)$		$P(A, B) \leftarrow Q(A), R(B)$
$P(A, B) \leftarrow Q(A), R(B, C), S(B, C)$		$P(A, B) \leftarrow Q(A), R(A, B)$
$P(A, B) \leftarrow Q(B, C), R(A, D), S(A, D), T(B, C)$		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E), U(E)$
		$P(A, B) \leftarrow Q(B, C), R(A, D), S(B, D), T(C, E), U(C, E)$

Table 33 Reductions of the fragment $\mathcal{Q}_5^{\{0,1,2\}}$

S-reduction	E-reduction	D-reduction
$P \leftarrow Q$	$P \leftarrow Q$	$P \leftarrow Q$
$P(A) \leftarrow Q(A)$	$P(A) \leftarrow$ $Q(A, B), R(A, B)$	$P(A) \leftarrow Q(A)$
$P(A) \leftarrow$ $Q(A, B), R(A, B)$	$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow Q(B, A)$
$P(A, B) \leftarrow$ $Q(B, A)$	$P(A, B) \leftarrow$ $Q(A), R(B)$	$P \leftarrow Q, R$
$P(A, B) \leftarrow$ $Q(A, B)$		$P(A) \leftarrow Q, R(A)$
$P(A, B) \leftarrow$ $Q(A), R(B)$		$P(A) \leftarrow Q(A), R(A)$
$P(A, B) \leftarrow$ $Q(A), R(B, C), S(B, C)$		$P(A) \leftarrow Q(A, B), R(B)$
$P(A, B) \leftarrow$ $Q(B), R(A, C), S(A, C)$		$P(A) \leftarrow Q(A, B), R(A, B)$
$P(A, B) \leftarrow$ $Q(B, C), R(A, D),$ $S(A, D), T(B, C)$		$P(A, B) \leftarrow Q, R(A, B)$
		$P(A, B) \leftarrow Q(A), R(B)$
		$P(A, B) \leftarrow Q(A), R(A, B)$
		$P(A, B) \leftarrow Q(A, B), R(A, B)$
		$P(A, B) \leftarrow Q(A, C), R(B, C)$
		$P(A, B) \leftarrow$ $Q(A, C), R(A, D), S(B, C), T(B, D), U(C, D)$
		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(B, D), T(C, E), U(E)$
		$P(A, B) \leftarrow$ $Q(B, C), R(A, D), S(B, D), T(C, E), U(C, E)$

References

- Albarghouthi, A., Koutris, P., Naik, M., & Smith, C. (2017). Constraint-based synthesis of Datalog programs. In J. C. Beck (Ed.), *Principles and practice of constraint programming—23rd international conference, CP 2017, Melbourne, VIC, Australia, August 28–September 1, 2017, Proceedings*, volume 10416 of *Lecture Notes in Computer Science* (pp. 689–706). Springer.
- Bienvenu, M. (2007). Prime implicates and prime implicants in modal logic. In *Proceedings of the twenty-second AAAI conference on artificial intelligence, July 22–26, 2007, Vancouver, BC, Canada* (pp. 379–384). AAAI Press.
- Blumer, A., Ehrenfeucht, A., Haussler, D., & Warmuth, M. K. (1987). Occam's razor. *Information Processing Letters*, 24(6), 377–380.
- Bradley, A. R., & Manna, Z. (2007). *The calculus of computation-decision procedures with applications to verification*. Berlin: Springer.
- Campero, A., Pareja, A., Klinger, T., Tenenbaum, J., & Riedel, S. (2018). Logical rule induction and theory learning using neural theorem proving. ArXiv e-prints, September 2018.
- Church, A. (1936). A note on the Entscheidungsproblem. *The Journal of Symbolic Logic*, 1(1), 40–41.
- Cohen, W. W. (1994). Grammatically biased learning: Learning logic programs using an explicit antecedent description language. *Artificial Intelligence*, 68(2), 303–366.

- Cropper, A. (2017). *Efficiently learning efficient programs*. Ph.D. thesis, Imperial College London, UK.
- Cropper, A., Evans, R., & Law, M. (2019). Inductive general game playing. ArXiv e-prints, [arXiv:1906.09627](https://arxiv.org/abs/1906.09627), Jun 2019.
- Cropper, A., & Muggleton, S. H. (2014). Logical minimisation of meta-rules within meta-interpretive learning. In J. Davis & J. Ramon (Eds.), *Inductive logic programming—24th international conference, ILP 2014, Nancy, France, September 14–16, 2014*. Revised selected papers, volume 9046 of Lecture Notes in Computer Science (pp. 62–75). Springer.
- Cropper, A., & Muggleton, S. H. (2015). Learning efficient logical robot strategies involving composable objects. In Yang, Q., & Wooldridge, M. (Eds.), *Proceedings of the twenty-fourth international joint conference on artificial intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25–31, 2015* (pp. 3423–3429). AAAI Press.
- Cropper, A., & Muggleton, S. H. (2016a). Learning higher-order logic programs through abstraction and invention. In Kambhampati, S. (Ed.), *Proceedings of the twenty-fifth international joint conference on artificial intelligence, IJCAI 2016, New York, NY, USA, 9–15 July 2016* (pp. 1418–1424). IJCAI/AAAI Press.
- Cropper, A., & Muggleton, S. H. (2016b). Metagol system. <https://github.com/metagol/metagol>. Accessed 1 July 2019.
- Cropper, A., & Muggleton, S. H. (2019). Learning efficient logic programs. *Machine Learning*, 108(7), 1063–1083.
- Cropper, A., Tamaddoni-Nezhad, A., & Muggleton, S. H. (2015). Meta-interpretive learning of data transformation programs. In Inoue, K., Ohwada, H., & Yamamoto, A. (Eds.), *Inductive logic programming—25th international conference, ILP 2015, Kyoto, Japan, August 20–22, 2015, revised selected papers*, volume 9575 of Lecture Notes in Computer Science (pp. 46–59). Springer.
- Cropper, A., & Tourret, S. (2018). Derivation reduction of metarules in meta-interpretive learning. In Riguzzi, F., Bellodi, E., & Zese, R. (Eds.), *Inductive logic programming—28th international conference, ILP 2018, Ferrara, Italy, September 2–4, 2018, proceedings*, volume 11105 of Lecture Notes in Computer Science (pp. 1–21). Springer.
- Dantsin, E., Eiter, T., Gottlob, G., & Voronkov, A. (2001). Complexity and expressive power of logic programming. *ACM Computing Surveys*, 33(3), 374–425.
- De Raedt, L. (2012). Declarative modeling for machine learning and data mining. In *Algorithmic learning theory—23rd international conference, ALT 2012, Lyon, France, October 29–31, 2012. proceedings* (p. 12).
- De Raedt, L., & Bruynooghe, M. (1992). Interactive concept-learning and constructive induction by analogy. *Machine Learning*, 8, 107–150.
- Echenim, M., Peltier, N., & Tourret, S. (2015). Quantifier-free equational logic and prime implicate generation. In A. P. Felty & A. Middeldorp (Eds.), *Automated deduction—CADE-25—25th international conference on automated deduction, Berlin, Germany, August 1–7, 2015, proceedings*, volume 9195 of Lecture Notes in Computer Science (pp. 311–325). Springer.
- Emde, W., Habel, C., & Rollinger, C.-R. (1983). The discovery of the equator or concept driven learning. In M. Alanbundy (Ed.), *Proceedings of the 8th international joint conference on artificial intelligence. Karlsruhe, FRG, August 1983* (pp. 455–458). William Kaufmann.
- Evans, R., & Grefenstette, E. (2018). Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61, 1–64.
- Flener, P. (1996). Inductive logic program synthesis with DIALOGS. In Muggleton, S. (Ed.), *Inductive logic programming, 6th international workshop, ILP-96, Stockholm, Sweden, August 26–28, 1996, selected papers*, volume 1314 of Lecture Notes in Computer Science (pp. 175–198). Springer.
- Fonseca, N. A., Costa, V. S., Silva, F. M. A., & Camacho, R. (2004). On avoiding redundancy in inductive logic programming. In R. Camacho, R. D. King & A. Srinivasan (Eds.), *Inductive logic programming, 14th international conference, ILP 2004, Porto, Portugal, September 6–8, 2004, proceedings*, volume 3194 of Lecture Notes in Computer Science (pp. 132–146). Springer.
- Fürnkranz, J. (1997). Dimensionality reduction in ILP: A call to arms. In *Proceedings of the IJCAI-97 workshop on frontiers of inductive logic programming* (pp. 81–86).
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. New York: W. H. Freeman.
- Genesereth, M. R., Love, N., & Pell, B. (2005). General game playing: Overview of the AAAI competition. *AI Magazine*, 26(2), 62–72.
- Gottlob, G., & Fermüller, C. G. (1993). Removing redundancy from a clause. *Artificial Intelligence*, 61(2), 263–289.
- Gottlob, G., Leone, N., & Scarcello, F. (1997). On the complexity of some inductive logic programming problems. In N. Lavrac & S. Dzeroski (Eds.), *Inductive logic programming, 7th international workshop*,

- ILP-97, Prague, Czech Republic, September 17–20, 1997, proceedings, volume 1297 of Lecture Notes in Computer Science (pp. 17–32). Springer.
- Hemaspaandra, E., & Schnoor, H. (2011). Minimization for generalized boolean formulas. In T. Walsh (Ed.), *IJCAI 2011, proceedings of the 22nd international joint conference on artificial intelligence, Barcelona, Catalonia, Spain, July 16–22, 2011* (pp. 566–571). IJCAI/AAAI.
- Heule, M., Järvisalo, M., Lonsing, F., Seidl, M., & Biere, A. (2015). Clause elimination for SAT and QSAT. *Artificial Intelligence Research*, 53, 127–168.
- Hillenbrand, T., Piskac, R., Waldmann, U., & Weidenbach, C. (2013). From search to computation: Redundancy criteria and simplification at work. In A. Voronkov, & C. Weidenbach (Eds.), *Programming logics - essays in memory of Harald Ganzinger*, volume 7797 of Lecture Notes in Computer Science (pp. 169–193). Springer.
- Joyner, W. H., Jr. (1976). Resolution strategies as decision procedures. *Journal of the ACM*, 23(3), 398–417.
- Kaminski, T., Eiter, T., & Inoue, K. (2018). Exploiting answer set programming with external sources for meta-interpretive learning. *TPLP*, 18(3–4), 571–588.
- Kietz, J.-U., & Wrobel, S. (1992). Controlling the complexity of learning in logic through syntactic and task-oriented models. In *Inductive logic programming*. Citeseer.
- Kowalski, R. A. (1974). Predicate logic as programming language. In *IFIP congress* (pp. 569–574).
- Larson, J., & Michalski, R. S. (1977). Inductive inference of VL decision rules. *SIGART Newsletter*, 63, 38–44.
- Liberatore, P. (2005). Redundancy in logic I: CNF propositional formulae. *Artificial Intelligence*, 163(2), 203–232.
- Liberatore, P. (2008). Redundancy in logic II: 2CNF and Horn propositional formulae. *Artificial Intelligence*, 172(2–3), 265–299.
- Lin, D., Dechter, E., Ellis, K., Tenenbaum, J. B., & Muggleton, S. (2014). Bias reformulation for one-shot function induction. In *ECAI 2014—21st European conference on artificial intelligence, 18–22 August 2014, Prague, Czech Republic—including prestigious applications of intelligent systems (PAIS 2014)* (pp. 525–530).
- Lloyd, J. W. (1987). *Foundations of logic programming* (2nd ed.). Berlin: Springer.
- Lloyd, J. W. (2003). *Logic for learning*. Berlin: Springer.
- Marcinkowski, J., & Pacholski, L. (1992). Undecidability of the Horn-clause implication problem. In *33rd annual symposium on foundations of computer science, Pittsburgh, Pennsylvania, USA, 24–27 October 1992* (pp. 354–362).
- Marquis, P. (2000). Consequence finding algorithms. In *Handbook of defeasible reasoning and uncertainty management systems* (pp. 41–145). Springer.
- McCarthy, J. (1995). Making robots conscious of their mental states. In *Machine intelligence 15, intelligent Agents [St. Catherine's College, Oxford, July 1995]* (pp. 3–17).
- Morel, R., Cropper, A., & Ong, C.-H. Luke (2019). Typed meta-interpretive learning of logic programs. In Calimeri, F., Leone, N., & Manna, M. (Eds.), *Logics in artificial intelligence—16th European conference, JELIA 2019, Rende, Italy, May 7–11, 2019, proceedings*, volume 11468 of Lecture Notes in Computer Science (pp. 198–213). Springer.
- Muggleton, S. (1995). Inverse entailment and Progol. *New Generation Computing*, 13(3&4), 245–286.
- Muggleton, S., De Raedt, L., Poole, D., Bratko, I., Flach, P. A., Inoue, K., et al. (2012). ILP turns 20—biography and future challenges. *Machine Learning*, 86(1), 3–23.
- Muggleton, S., & Feng, C. (1990). Efficient induction of logic programs. In *Algorithmic learning theory, first international workshop, ALT '90, Tokyo, Japan, October 8–10, 1990, proceedings* (pp. 368–381).
- Muggleton, S. H., Lin, D., Pahlavi, N., & Tamaddoni-Nezhad, A. (2014). Meta-interpretive learning: Application to grammatical inference. *Machine Learning*, 94(1), 25–49.
- Muggleton, S. H., Lin, D., & Tamaddoni-Nezhad, A. (2015). Meta-interpretive learning of higher-order dyadic Datalog: Predicate invention revisited. *Machine Learning*, 100(1), 49–73.
- Nédellec, C., Rouveirol, C., Adé, H., Bergadano, F., & Tausend, B. (1996). Declarative bias in ILP. *Advances in inductive logic programming*, 32, 82–103.
- Nienhuys-Cheng, S.-H., & de Wolf, R. (1997). *Foundations of inductive logic programming*. New York, Secaucus, NJ: Springer.
- Plotkin, G.D. (1971). *Automatic methods of inductive inference*. Ph.D. thesis, Edinburgh University, August 1971.
- Robinson, J. A. (1965). A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1), 23–41.
- Schmidt-Schauß, M. (1988). Implication of clauses is undecidable. *Theoretical Computer Science*, 59, 287–296.
- Shapiro, E. Y. (1983). *Algorithmic program debugging*. London: MIT Press.

- Si, X., Lee, W., Zhang, R., Albarghouthi, A., Koutris, P., & Naik, M. (2018). Syntax-guided synthesis of Datalog programs. In G. T. Leavens, A. Garcia, & C. S. Pasareanu (Eds.), *Proceedings of the 2018 ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04–09, 2018* (pp. 515–527). ACM.
- Skillicorn, D. (2007). *Understanding complex datasets: Data mining with matrix decompositions*. New York: Chapman and Hall/CRC.
- Tärnlund, S. Å. (1977). Horn clause computability. *BIT*, 17(2), 215–226.
- Tourret, S., & Cropper, A. (2019). SLD-resolution reduction of second-order Horn fragments. In F. Calimeri, N. Leone & M. Manna (Eds.), *Logics in artificial intelligence—16th European conference, JELIA 2019, Rende, Italy, May 7–11, 2019, proceedings*, volume 11468 of Lecture Notes in Computer Science (pp. 259–276). Springer.
- Wang, W. Y., Mazaitis, K., & Cohen, W. W. (2014). Structure learning via parameter learning. In Li, J., Wang, X. S., Garofalakis, M. N., Soboroff, I., Suel, T., & Wang, M. (Eds.), *Proceedings of the 23rd ACM international conference on conference on information and knowledge management, CIKM 2014, Shanghai, China, November 3–7, 2014* (pp. 1199–1208). ACM.
- Weidenbach, C., & Wischniewski, P. (2010). Subterm contextual rewriting. *AI Communications*, 23(2–3), 97–109.

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