

Extending Forgetting-Based Abduction Using Nominals

Warren Del-Pinto and Renate A. Schmidt

School of Computer Science, University of Manchester, Oxford Road, Manchester,
M13 9PL, United Kingdom

Abstract. Abductive reasoning produces hypotheses to explain new observations with respect to some background knowledge. This paper focuses on ABox abduction in ontologies, where knowledge is expressed in description logics and both the observations and hypotheses are ground statements. The input is expressed in the description logic \mathcal{ALC} and the observation can contain any set of \mathcal{ALC} concept or role assertions. The proposed approach uses forgetting to produce hypotheses in the form of a disjunctive set of axioms, where each disjunct is an independent explanation for the observation and the overall hypothesis is semantically minimal, i.e., makes the least assumptions required. Previous work on forgetting-based abduction is combined with the semantic forgetting method of the system FAME. The hypotheses produced are expressed in an extension of \mathcal{ALC} which uses nominals, role inverses and fixpoints: $\mathcal{ALCOI}\mu(\nabla)$. This combination overcomes the inability of the existing forgetting-based approach to allow role assertions in observations and hypotheses, and enables the computation of other previously unreachable hypotheses. An experimental evaluation is performed using a prototype implementation of the method on a corpus of real world ontologies.

1 Introduction

Abduction was first identified as a form of reasoning by C.S. Peirce, who likened it to a “flash of insight”. Like induction, and unlike deduction, abduction is ampliative: the conclusion of the reasoning process extends beyond what already follows from existing background knowledge. Abduction is often seen as the process of hypothesis generation, while induction can be seen as the process of hypothesis evaluation or generalisation. The use cases for abduction have led to a diverse range of investigations into the topic. These include complexity studies [10], applications in natural language interpretation [16], inductive and abductive logic programming [26, 29], statistical relational AI [28] and studies of the interaction between abduction and induction [12].

This paper focuses on abductive reasoning in description logics (DLs), which are fragments of first-order logic. In this setting, background knowledge is expressed in an ontology, which contains information regarding concepts and relations between entities. Ontologies are used in a wide variety of fields including bioinformatics, robotics and finance. Benefits of using ontologies include the

ability to clearly model, reuse, share and reason about existing knowledge. Most existing reasoning systems in ontologies are deductive. They can be used to derive consequences of the existing ontology that are not explicitly represented. However, they cannot be used directly to explain new observations that do not follow from the existing knowledge, which is required for tasks such as hypothesis generation, diagnostics and belief expansion. The importance of abduction in DLs has been recognised [11] and a variety of work exists on the topic, including complexity studies [3], applications to repair and query explanation [21, 5] and methods for different forms of TBox and ABox abduction [17, 8, 15, 27, 9].

One approach to performing abductive reasoning in DL ontologies uses *forgetting*. Forgetting aims to eliminate specified symbols in an ontology while preserving all entailments that can be represented in the restricted signature. The dual task is called *uniform interpolation*. Both of these are related to second-order quantifier elimination [13], which translates logical formulae expressed in second-order logic into equivalent formulae in first-order logic by eliminating existentially quantified predicate symbols. The use of second-order quantifier elimination for abduction has been proposed for relatively small theories expressed in propositional or classical logics [7, 13, 30], while forgetting has been proposed for TBox abduction in DLs [20]. More recently, a method for performing ABox abduction in the DL \mathcal{ALC} was developed [6], which utilises contrapositive reasoning and the resolution-based forgetting system LETHE [19]. This approach produces hypotheses that consist of a disjunctive set of ABox assertions. Each disjunct is an independent explanation [18], resulting in a space of possible explanations. The method has been shown to be practical over large ontologies. However, a limitation to the method is the absence of role assertions in observations and hypotheses, which restricts its use of existing information contained within the ABox. Given an ABox observation, the approach cannot use relationships between individuals to provide a more specific explanation. These explanations would be useful in many applications, such as those involving the use of large knowledge graphs which have seen increasing interest in recent years.

The primary aim of this work is to overcome this limitation by combining the method in [6] with another forgetting system: FAME [32, 33]. The key characteristic of FAME is its ability to perform forgetting in $\mathcal{ALCOI}\mu(\nabla)$, which includes nominals. As suggested in [6], nominals can be used to overcome the limitations of the abduction method. This is explored and confirmed in this paper.

2 Problem Definition

In this work, knowledge is expressed in the description logic \mathcal{ALC} [2]. The signature of \mathcal{ALC} is defined by disjoint sets N_c , N_r and N_I containing atomic concept names, role names and individual names respectively. Concepts in \mathcal{ALC} can take the following forms: $\perp \mid \top \mid A \mid \neg C \mid C \sqcup D \mid C \sqcap D \mid \exists r.C \mid \forall r.C$, where A is any atomic concept name, C and D are any \mathcal{ALC} concepts and r is a role name.

An ontology \mathcal{O} expressed in \mathcal{ALC} takes the form $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$, where \mathcal{T} is a TBox and \mathcal{A} is an ABox. The TBox contains information about concepts rep-

resented as general concept inclusions (GCIs) of the form $C \sqsubseteq D$ or equivalence axioms of the form $C \equiv D$, which can also be expressed as the two GCIs $C \sqsubseteq D$ and $D \sqsubseteq C$. The ABox contains (ground) assertions about specific individuals of the form $C(a)$ or $r(a, b)$ where a, b are arbitrary individual names.

The semantics of \mathcal{ALC} is defined in terms of an interpretation \mathcal{I} as a pair $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$, where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain* and $\cdot^{\mathcal{I}}$ is an *interpretation function* mapping each individual $a \in N_I$ to a single element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each concept to a subset of $\Delta^{\mathcal{I}}$ and each role to a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. This is extended to \mathcal{ALC} concepts as follows:

$$\begin{aligned} \perp^{\mathcal{I}} &= \emptyset & \top^{\mathcal{I}} &= \Delta^{\mathcal{I}} & \neg C &= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (C \sqcap D)^{\mathcal{I}} &= C^{\mathcal{I}} \cap D^{\mathcal{I}} & (C \sqcup D)^{\mathcal{I}} &= C^{\mathcal{I}} \cup D^{\mathcal{I}} \\ (\exists r.C)^{\mathcal{I}} &= \{x \in \Delta^{\mathcal{I}} \mid \exists y.(x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \\ (\forall r.C)^{\mathcal{I}} &= \{x \in \Delta^{\mathcal{I}} \mid \forall y.(x, y) \in r^{\mathcal{I}} \rightarrow y \in C^{\mathcal{I}}\} \end{aligned}$$

For TBox axioms, the GCI $C \sqsubseteq D$ is true in \mathcal{I} iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ holds. A model of a TBox is an interpretation for which all axioms in the TBox are true, and if a TBox has a model then it is *satisfiable*.

For this paper, it is also necessary to consider the DL $\mathcal{ALCOI}\mu(\nabla)$, which extends \mathcal{ALC} with *nominals*, *role inverses*, the *top role* and in some cases *fixpoints*. For each individual $a \in N_I$, the corresponding nominal $\{a\}$ is interpreted as a concept containing only a . Using nominals, ABox assertions can be expressed as equivalent TBox axioms, i.e., $C(a)$ and $r(a, b)$ can be expressed as $\{a\} \sqsubseteq C$ and $\{a\} \sqsubseteq \exists r.\{b\}$ respectively. The role inverse of a role r is denoted by r^- , and the top role is denoted by ∇ . The semantics of these are defined as follows: $(r^-)^{\mathcal{I}} = \{(y, x) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \mid (x, y) \in r^{\mathcal{I}}\}$ and $\nabla^{\mathcal{I}} = \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. In some specific cases, described later in this paper, fixpoints may be used to represent cyclic results. We refer to [4] for a full description and the semantics of fixpoints.

Here the focus is on ABox abduction over \mathcal{ALC} ontologies. Our aim is to produce hypotheses that satisfy the following:

Definition 1. *Let \mathcal{O} be an ontology and ψ be a set of ABox assertions, both expressed in \mathcal{ALC} , where $\mathcal{O}, \psi \not\models \perp$ and $\mathcal{O} \not\models \psi$. Let \mathcal{S}_A be a set of abducible symbols, containing any subset of the symbols in the signature of \mathcal{O}, ψ . The ABox abduction problem is to find a hypothesis \mathcal{H} as a disjunction of ABox assertions, containing only the symbols in \mathcal{S}_A , which satisfies the following conditions:*

- (i) $\mathcal{O}, \mathcal{H} \not\models \perp$, (ii) $\mathcal{O}, \mathcal{H} \models \psi$
- (iii) *Each disjunct α_i in \mathcal{H} is an independent explanation for ψ : i.e. for every α_i in \mathcal{H} , $\mathcal{O}, \alpha_i \not\models \alpha_1 \sqcup \dots \sqcup \alpha_{i-1} \sqcup \alpha_{i+1} \sqcup \dots \sqcup \alpha_n$.*
- (iv) *If there exists a \mathcal{H}' which satisfies conditions (i)–(iii), such that \mathcal{H}' contains only symbols in \mathcal{S}_A and $\mathcal{O}, \mathcal{H} \models \mathcal{O}, \mathcal{H}'$, then $\mathcal{O}, \mathcal{H}' \models \mathcal{O}, \mathcal{H}$.*

Conditions (i), *consistency*, and (ii), *explanation*, are standard conditions on abductive hypotheses, requiring that the hypothesis computed explains the observation using the information in \mathcal{O} without contradicting it. Condition (iii)

requires that there are no redundant disjuncts in the hypothesis \mathcal{H} . A redundant disjunct is one that contradicts the information in \mathcal{O} , or provides an explanation that is simply stronger than one that is already contained within the rest of the hypothesis \mathcal{H} . From this, it can be seen that condition (i) is a consequence of condition (iii). However, condition (i) is still included here for clarity as it is a key requirement. Condition (iv) requires that the overall hypothesis computed is the one that makes the least assumptions required to entail the observation, and is referred to as semantic minimality [15]. As noted in [6], in settings where the hypothesis can contain disjunctions it is necessary to consider the redundancy of individual disjuncts prior to checking for semantic minimality.

The above definition extends the problem defined in [6] by lifting the restrictions on ψ , \mathcal{S}_A and \mathcal{H} . Specifically, both ψ and \mathcal{H} may contain any combination of \mathcal{ALC} ABox assertions, including role assertions, and the set of abducibles \mathcal{S}_A is no longer required to contain all role symbols in \mathcal{O} , ψ .

In addition, in this work the produced hypothesis consists of a disjunction of $\mathcal{ALCOI}\mu(\nabla)$ axioms. Thus, certain hypotheses that can only be expressed using nominals and inverse roles are also reachable using this extended approach. The exact form of these are discussed alongside the proposed method.

3 Forgetting-Based Abduction

Forgetting eliminates symbols, i.e., concept and role names from an ontology while preserving the entailments that are representable in the restricted signature. The symbols to eliminate are specified by a forgetting signature \mathcal{F} , where \mathcal{F} can contain any subset of symbols in the signature of the ontology.

Forgetting can be utilised for abduction via contraposition: $\mathcal{O}, \mathcal{H} \models \psi$ if and only if $\mathcal{O}, \neg\psi \models \neg\mathcal{H}$. For an ontology \mathcal{O} and observation ψ , both expressed in \mathcal{ALC} , the steps in forgetting-based abduction are given in Figure 1.

- | |
|---|
| <ol style="list-style-type: none"> 1. Eliminate a specified set of symbols \mathcal{F} from $(\mathcal{O}, \neg\psi)$. The result of this step is a new ontology, $\mathcal{V} = \{\beta_1, \dots, \beta_n\}$, which is called the <i>forgetting solution</i> of $(\mathcal{O}, \neg\psi)$ with respect to \mathcal{F}. 2. Extract a <i>reduced forgetting solution</i> \mathcal{V}^* from \mathcal{V}. This is done by eliminating all axioms in \mathcal{V} that violate the dual of Definition 1(iii), i.e., those $\beta_i \in \mathcal{V}$ such that $\mathcal{O}, \beta_1, \dots, \beta_{i-1}, \beta_{i+1}, \dots, \beta_n \models \beta_i$. 3. Negate \mathcal{V}^* to obtain a hypothesis \mathcal{H}, in the form of a disjunctive set of axioms, which satisfies Definition 1. |
|---|

Fig. 1. Steps in forgetting-based abduction [6].

There are several important features of forgetting that make forgetting-based abduction promising [6]. (1) Given that forgetting preserves all remaining entailments, the forgetting solutions can be seen as strongest necessary conditions of an input ontology in the restricted signature. This means that the forgetting solution \mathcal{V} is the strongest necessary condition [22] of $(\mathcal{O}, \neg\psi)$ in the signature \mathcal{S}_A . As

a result, the negation of \mathcal{V} is a weakest sufficient condition, as strongest necessary and weakest sufficient conditions are dual notions. Weakest sufficient conditions correspond to a weakest abduction result [22, 7], i.e., a semantically minimal hypothesis. When combined with the negation step 3 in Figure 1, this results in a hypothesis \mathcal{H} that satisfies both conditions (iii) and (iv). (2) The use of a forgetting signature \mathcal{F} provides a goal-oriented method for specifying the abducible symbols \mathcal{S}_A : for an ontology \mathcal{O} and an observation ψ , $\mathcal{S}_A = \text{sig}(\mathcal{O} \cup \{\psi\}) \setminus \mathcal{F}$. It may be the case that a user does not have sufficient information to manually choose the abducibles from a large set of available symbols. In this case, by inspecting the signature of the observation ψ , \mathcal{F} can be defined by simply setting it equal to a subset of the symbols in ψ . This guarantees that inferences are made between \mathcal{O} and ψ in order to eliminate the symbols in \mathcal{F} , in turn guaranteeing that a non-trivial hypothesis $\mathcal{H} \neq \psi$ is obtained. (3) Forgetting can be applied iteratively. For example, eliminating a set of symbols \mathcal{F}_1 from an ontology \mathcal{O} results in a forgetting result \mathcal{V}_1 . If a second set of symbols \mathcal{F}_2 is then eliminated from \mathcal{V}_1 , the result obtained will be the same as eliminating all of the symbols in $\mathcal{F}_1 \cup \mathcal{F}_2$ from \mathcal{O} . This provides a method for hypothesis refinement: the steps in Figure 1 can be repeated by eliminating symbols that occur in the current hypothesis to obtain a stronger hypothesis, perhaps based on heuristics or external data. (4) There exist several forgetting systems that have been shown to be efficient across large real-world ontologies, for example [23, 19, 33].

The abduction method proposed in [6] utilises the resolution-based forgetting approach implemented in the system LETHE [19], which performs forgetting over \mathcal{ALC} ontologies with ABoxes. LETHE adopts the *uniform interpolation* perspective on forgetting [24], meaning that it preserves all consequences of the input ontology in the restricted signature. Thus, the forgetting solution \mathcal{V} computed by LETHE is called a *uniform interpolant*.

By using LETHE for the forgetting step 1 in Figure 1, the abduction approach in [6] can compute hypotheses which satisfy a restricted form of Definition 1, and has been shown to be sound and complete for this problem. The first main restriction, however, is that the set of abducibles \mathcal{S}_A must contain all role symbols in $(\mathcal{O}, \neg\psi)$. This is due to the fact that the form of role forgetting currently implemented in LETHE is not complete for the abduction problem. This was noted in [6] and an example was provided. Here is another example:

Example 1. Consider the following ontology: $\mathcal{O} = \{B \sqsubseteq \exists r.B\}$. Let the observation be $\psi = \exists r.B(a)$. The expected hypothesis given $\mathcal{S}_A = \{B\}$ is $\mathcal{H} = B(a)$.

The result of applying LETHE's calculus [19] to forget r is an empty uniform interpolant. While this is sufficient for the uniform interpolation problem, for abduction the expected hypothesis $B(a)$ is not computed.

The second restriction is that the observations and hypotheses cannot contain role assertions. Consider the following example:

Example 2. Consider the following ontology $\mathcal{O} = \{\exists r.B \sqsubseteq A, B(b)\}$. Let the observation be $\psi = A(a)$. Case 1: given a set of abducibles $\mathcal{S}_A = \{B, r\}$, the hy-

pothesis satisfying Definition 1 would be $\mathcal{H}_1 = \exists r.B(a)$. Case 2: given $\mathcal{S}_A = \{r\}$, the hypothesis would instead be $\mathcal{H}_2 = r(a, b)$.

The first hypothesis \mathcal{H}_1 is reachable using LETHE. However, the second hypothesis \mathcal{H}_2 is not. This is due to the fact that LETHE’s calculus does not support deriving negated role assertions of the form $\neg r(a, b)$, as it is designed to preserve all entailments of the input that are expressible in \mathcal{ALC} and $\neg r(a, b)$ is not expressible in \mathcal{ALC} . Thus, these are not obtained in the reduced forgetting solution \mathcal{V}^* and role assertions will be absent from the hypothesis when \mathcal{V}^* is negated in step 3. This restriction means that the system is not able to utilise existing relationships between individuals in the observation and those in the ABox of the background ontology when generating hypotheses. Thus, many of the more specific hypotheses such as \mathcal{H}_2 in the example above are not reachable.

In summary, the existing forgetting-based abduction approach in [6] takes as **input** an \mathcal{ALC} ontology \mathcal{O} , an observation of the form $\psi = \{C_1(a_1), \dots, C_k(a_k)\}$ where each C_i is an \mathcal{ALC} concept and each a_i is an individual, and a set of abducibles \mathcal{S}_A containing all role symbols in $(\mathcal{O}, \neg\psi)$. The negation of the observation takes the form $\neg\psi = \neg C_1(a_1) \sqcup \dots \sqcup \neg C_k(a_k)$. The final hypothesis produced via steps 1–3 in Figure 1 takes the form $\mathcal{H} = \alpha_1(a_1) \sqcup \dots \sqcup \alpha_n(a_n)$ where each α_i is an \mathcal{ALC} concept. It turns out the final output may need to be expressed in \mathcal{ALC} extended with disjunctive ABox assertions $\mathcal{ALC}\nu$.

4 Extending the System

In this work, the forgetting-based approach is extended via the use of nominals. However, producing hypotheses with nominals requires the use of a forgetting method that can compute forgetting solutions for ontologies expressed in \mathcal{ALCO} .

For this purpose, the abduction approach above is combined with the forgetting system FAME [32, 33]. As opposed to the uniform interpolation perspective taken by LETHE, FAME frames the problem of forgetting in terms of *semantic forgetting* [32, 33]. This view is closely related to second-order quantifier elimination [7, 13], where the forgetting result must be equivalent to the original formula in second-order logic. From the viewpoint of abduction, while the steps taken in the approach are conceptually the same when utilising FAME, each step must be extended in several ways compared to the previous approach.

First it is helpful to consider how FAME handles the forgetting process [33]. Any ABox assertions in the input are translated to equivalent TBox axioms involving nominals. Then, the ontology is transformed into a set of clauses \mathcal{N} . Before the process of eliminating the symbols in the forgetting signature \mathcal{F} can take place, the set \mathcal{N} must be transformed into the appropriate reduced forms.

For role forgetting, \mathcal{N} is transformed into *r-reduced* form. Given a role symbol to be forgotten r , every clause in \mathcal{N} that contains r must be of the form $C \sqcup \forall r.D$ or $C \sqcup \neg\forall r.D$ where C and D are possibly complex concepts that do not contain r . During this transformation, definer symbols may need to be introduced. These are fresh symbols that do not appear in the input ontology and

are used to incrementally replace the concept symbols C and D in clauses such as the one above, until neither of them contain r . For example, given a forgetting signature $\mathcal{F} = \{r\}$ and a clause $\forall r.A \sqcup \forall r.B$, a definer is introduced to replace $\forall r.A$. This results in the two clauses $D_1 \sqcup \forall r.B$ and $\neg D_1 \sqcup \forall r.A$ respectively.

For concept forgetting, the reduced form is the A -reduced form. To forget a concept A , every clause must be of the form $A \sqcup C$ where C is a possibly complex concept which does not contain A or contains only negative occurrences of A . For example, the clauses $A \sqcup B$ and $A \sqcup \forall r.\neg A$ are in A -reduced form.

When forgetting the symbols in \mathcal{F} , role symbols are eliminated first followed by concept symbols. In both cases, the set \mathcal{N} is first transformed into r -reduced and A -reduced forms respectively [33]. Once the appropriate reduced form is obtained, rules based around Ackermann's Lemma [1] are used to forget the symbols in \mathcal{F} . Given a symbol A to be eliminated, the essence of Ackermann's lemma is to construct a definition of A , which does not contain it, from the existing clauses in which it occurs. This definition can then be used to replace every instance of A , thereby eliminating it from the original ontology. The Ackermann^R and Purify^R rules are used to eliminate role symbols, while the Ackermann^C and Purify^C rules, shown in Figure 2, are used to eliminate concept symbols [33]. Finally any definer symbols are eliminated, via the use of the Ackermann^C and Purify^C rules, resulting in the forgetting solution \mathcal{V} .

Full details of this process, including the rules for obtaining the reduced forms and the forgetting rules, can be found in the relevant papers [32, 33]. The concept forgetting rules have been included for reference in Figure 2, since these are utilised in illustrative examples throughout this paper.

An important aspect of FAME for this work is the fact that it is sound for forgetting in $\mathcal{ALCCOI}\mu(\nabla)$, as expressed in this theorem.

Theorem 1. *For any $\mathcal{ALCCOI}\mu(\nabla)$ ontology \mathcal{O} and any signature $\mathcal{F} \subseteq \text{sig}(\mathcal{O})$, where $\text{sig}(\mathcal{O})$ is the set of concept and role symbols in \mathcal{O} , FAME always terminates and returns a set \mathcal{V} of clauses. If \mathcal{V} does not contain any symbols in \mathcal{F} , then the symbols in \mathcal{F} were successfully forgotten and the set \mathcal{V} is a solution of forgetting the symbols in \mathcal{F} from \mathcal{O} .*

This theorem is a weaker form of the theorem in [33], specifically focussed on the description logic $\mathcal{ALCCOI}\mu(\nabla)$, as this is the language required in the setting of this paper. The original theorem holds for the logic $\mathcal{ALCCOI}\mathcal{H}\mu^+(\nabla, \sqcup)$, which also includes role hierarchies of the form $r \sqsubseteq s$ and role conjunction. However, here all inputs are expressed in $\mathcal{ALCCO}(\nabla)$. For this setting, role hierarchies are excluded. Role conjunctions are also excluded since they are only needed in the solution when the input is expressed in $\mathcal{ALCCOI}\mathcal{H}$ [31].

One limitation of using FAME to compute the forgetting result is that FAME is not complete. However, it is worth noting that LETHE is not complete for role forgetting in the context of abduction either [6]. Thus, the incompleteness drawback of FAME is offset by the fact that additional hypotheses can be reached. This is illustrated by examples 4, 5 and 7 later in this paper.

Via the use of FAME, the forgetting step 1 of the abduction method can now take any ontology \mathcal{O} and observation ψ for which the combination (\mathcal{O}, ψ)

<p>Non-cyclic Ackermann^C: $\frac{\mathcal{N}, C_1 \sqcup A, \dots, C_n \sqcup A}{\mathcal{N}_{C_1 \sqcup \dots \sqcup C_n}^A}$ <i>where (i) A does not occur in any C_i (ii) N is negative with respect to A</i></p> <p>Cyclic Ackermann^C: $\frac{\mathcal{N}, C_1[A] \sqcup A, \dots, C_n[A] \sqcup A}{\mathcal{N}_{\mu X.(\neg C_1 \sqcup \dots \sqcup \neg C_n)[X]}^A}$ <i>where (i) each C_i is negative with respect to A, (ii) N is negative w.r.t A</i></p> <p>Purify^C: $\frac{\mathcal{N}}{\mathcal{N}_{(\neg)\top}^A}$ <i>where N is positive with respect to A</i></p>

Fig. 2. Rules used in the forgetting system FAME to forget a concept A , where \mathcal{N}_C^A is the set of clauses obtained from \mathcal{N} by replacing every occurrence of A with C [33].

is expressible in \mathcal{ALCO} , where the signature of abducibles \mathcal{S}_A can be any set of concept or role symbols occurring in (\mathcal{O}, ψ) . Note, the set of abducibles \mathcal{S}_A includes all nominals, since no form of nominal forgetting is utilised in this work.

However, it is first necessary to decide how to represent the negation of the observation ψ and the form of the hypothesis \mathcal{H} , since FAME operates on and produces \mathcal{ALCO} TBoxes rather than ABox assertions. To illustrate this, as well as the procedure for concept forgetting in FAME, recall Example 2:

Example 3. The \mathcal{ALCO} reformulation of the ontology considered in Example 2 is: $\mathcal{O} = \{\exists r.B \sqsubseteq A, \{b\} \sqsubseteq B\}$ and the observation is $\psi = \{a\} \sqsubseteq A$, where $\{a\}$ and $\{b\}$ are nominals. Given a set of abducibles $\mathcal{S}_A = \{B, r\}$, the hypothesis obtained using FAME in step 1 of Figure 1 is $\mathcal{H}_1 = \{a\} \sqsubseteq \exists r.B$. If instead $\mathcal{S}_A = \{r\}$, the hypothesis is $\mathcal{H}_2 = \{a\} \sqsubseteq \exists r.\{b\}$. Both \mathcal{H}_1 and \mathcal{H}_2 satisfy Definition 1.

However, in cases where either the observation or the hypothesis take the form of a conjunction or disjunction of ABox assertions, the reformulation is less obvious. In this case, it is possible to take advantage of the fact that FAME can perform forgetting in the presence of the top role ∇ .

Example 4. Consider the following ontology $\mathcal{O} = \{\exists r.B \sqsubseteq A, C \sqsubseteq D, \{b\} \sqsubseteq B\}$ and the observation $\psi = A(a) \sqcap D(c)$. Let $\mathcal{S}_A = \{r, C\}$. The expected hypothesis under Definition 1 should be equivalent to $\mathcal{H} = r(a, b) \sqcap C(c)$. The negation of ψ can be represented as: $\neg\psi = \top \sqsubseteq \forall\nabla.(\neg\{a\} \sqcup \neg A) \sqcup \forall\nabla.(\neg\{c\} \sqcup \neg D)$. Following the steps in Figure 1 using FAME, in step 1, where $\mathcal{F} = \{A, B, D\}$, the hypothesis obtained can be represented as: $\mathcal{H} = \top \sqsubseteq \exists\nabla.(\neg\{a\} \sqcup \exists r.\{b\}) \sqcup \exists\nabla.(\neg\{c\} \sqcup C)$. This is equivalent to the expected hypothesis $r(a, b) \sqcap C(c)$.

This leads to the general form used for the negated observations and the hypotheses produced. These are shown below:

$$\neg\psi = \top \sqsubseteq \forall\nabla.(\neg\{a_1\} \sqcup \neg C_1) \sqcup \dots \sqcup \forall\nabla.(\neg\{a_k\} \sqcup \neg C_k) \quad (1)$$

$$\mathcal{H}_F = \top \sqsubseteq \exists\nabla.D_1 \sqcup \dots \sqcup \exists\nabla.D_n \quad (2)$$

where each C_i is an \mathcal{ALCO} concept and each D_i is an $\mathcal{ALCOI}\mu(\nabla)$ concept. From here \mathcal{H}_F is used to denote the hypothesis obtained using FAME in step 1 in Figure 1 and \mathcal{H}_L is used to refer to the one obtained using LETHE.

The following lemma relates \mathcal{H}_F to the disjunctive form in Definition 1:

Lemma 1. *The hypothesis \mathcal{H}_F is expressed as $\mathcal{H}_F = \alpha_1 \sqcup \dots \sqcup \alpha_n$ where each disjunct α_i is of the form $\top \sqsubseteq \exists\nabla.D_i$ and each D_i is an $\mathcal{ALCOI}\mu(\nabla)$ concept.*

It is still possible to satisfy conditions (iii) and (iv) of Definition 1 using this representation. However, it is first necessary to adapt the filtering method of step 2 in Figure 1 to obtain the reduced forgetting solution \mathcal{V}^* , as this is an important part of the feasibility of the approach in practice [6].

An annotation concept ℓ is used to efficiently trace any dependencies on the negated observation $\neg\psi$ in the forgetting result \mathcal{V} . Any axioms which do not contain the concept ℓ are removed from \mathcal{V} , thereby removing the majority of the axioms that are redundant with respect to Definition 1. Fortunately, extending this approach to the current setting is straightforward. Here, the negated observation provided in step 1 of Figure 1 is annotated as follows:

$$\neg\psi = \top \sqsubseteq \forall\nabla.(\neg\{a_1\} \sqcup \neg C_1 \sqcup \ell) \sqcup \dots \sqcup \forall\nabla.(\neg\{a_k\} \sqcup \neg C_k \sqcup \ell)$$

where as before, ℓ is a fresh concept symbol that does not occur in (\mathcal{O}, ψ) , nor in the signature \mathcal{F} . The soundness of this filtering approach is expressed below.

Theorem 2. *Let \mathcal{O} be an $\mathcal{ALCOI}\mu(\nabla)$ ontology, ψ an observation as a set of axioms, \mathcal{F} a forgetting signature and ℓ an annotator concept appended disjunctively to each disjunct in $\neg\psi$, where $\ell \notin \text{sig}(\mathcal{O})$ and $\ell \notin \mathcal{F}$. For each axiom β in the forgetting result \mathcal{V} obtained by forgetting all symbols in \mathcal{F} , if $\ell \notin \text{sig}(\beta)$ then β is redundant under the dual of Definition 1(iii), and should be removed in the extraction of the reduced forgetting result \mathcal{V}^* .*

Proof sketch: The proof is by induction over the construction of a derivation using the calculus of FAME [33], and takes the same form as the proof in [6]. The annotation concept ℓ does not appear in the signature \mathcal{F} . Thus, ℓ is not eliminated and if a clause in the normal form of $(\mathcal{O}, \neg\psi)$ contains the annotation concept ℓ , then any clause derived via inferences on this clause under FAME's forgetting calculus will also contain ℓ . Therefore, any axiom β in the forgetting result \mathcal{V} that does not contain ℓ was derived purely using axioms in the background ontology \mathcal{O} , i.e., $\mathcal{O} \models \beta$. Since under Definition 1, $\mathcal{O} \not\models \psi$, such a β will not contribute to the explanation of ψ required by abduction, and should be omitted from \mathcal{H}_F to satisfy Definition 1(iii).

As proposed in [6], the filtering step 3 in Figure 1 can be performed in an *approximate* or *full* manner. The approximate filtering utilises the annotation-based method to inexpensively remove all redundancies that can be captured using this approach. The result is an approximation of the reduced forgetting

result \mathcal{V}^* , denoted by \mathcal{V}_{app}^* . This can be negated in step 3 to return an approximate hypothesis. Alternatively, the full filtering setting further performs the dual entailment check of Definition 1(iii) over each axiom in \mathcal{V}_{app}^* using an external reasoner. This eliminates any remaining redundancies that cannot be captured using annotations, an example of which appears in [6]. The result is then \mathcal{V}^* , which is negated to return a hypothesis satisfying Definition 1.

It is worth noting that, for fixpoints to occur in the hypothesis, a cycle would need to occur both over the symbols in \mathcal{F} and also not be redundant under Theorem 2. As found in [6] this is rare in practice.

Example 5. To illustrate the full procedure, for the ontology \mathcal{O} and the observation ψ from Example 4, the input given to FAME is:

$$\begin{array}{ll} \exists r.B \sqsubseteq A & C \sqsubseteq D \\ \{b\} \sqsubseteq B & \top \sqsubseteq \forall \nabla.(\neg\{a\} \sqcup A \sqcup \ell) \sqcup \forall \nabla.(\neg\{c\} \sqcup D \sqcup \ell). \end{array}$$

The set of abducibles is $\mathcal{S}_A = \{C, r\}$ and thus the forgetting signature is $\mathcal{F} = \{A, B, D\}$. In step 1, $(\mathcal{O}, \neg\psi)$ is first transformed into A -reduced form:

$$\begin{array}{ll} \forall r.\neg B \sqcup A & \neg C \sqcup D \\ \neg\{b\} \sqcup B & \forall \nabla.(\neg\{a\} \sqcup \neg A \sqcup \ell) \sqcup \forall \nabla.(\neg\{c\} \sqcup \neg D \sqcup \ell) \end{array}$$

This is also in D -reduced form. Forgetting the concepts A and D results in:

$$\neg\{b\} \sqcup B \quad \forall \nabla.(\neg\{a\} \sqcup \forall r.\neg B \sqcup \ell) \sqcup \forall \nabla.(\neg\{c\} \sqcup \neg C \sqcup \ell)$$

Forgetting the concept B then produces:

$$\forall \nabla.(\neg\{a\} \sqcup \forall r.\neg\{b\} \sqcup \ell) \sqcup \forall \nabla.(\neg\{c\} \sqcup \neg C \sqcup \ell),$$

which is the forgetting result \mathcal{V} . In the filtering step 3 of Figure 1, the axiom is retained and the annotation concept ℓ is set to \perp . Neither disjunct in this hypothesis is redundant with respect to the dual of Definition 1(iii) and thus both are retained in the reduced forgetting result \mathcal{V}^* , which is then negated in step 3 to produce the hypothesis: $\mathcal{H}_F = \top \sqsubseteq \exists \nabla.(\{a\} \sqcap \exists r.\{b\}) \sqcup \exists \nabla.(\{c\} \sqcap C)$. This is equivalent to the suggested hypothesis $\mathcal{H} = r(a, b) \sqcap C(c)$.

5 Comparing Hypotheses

Since the main aim of abductive reasoning is to produce an *explanation*, the form taken by the hypotheses is important. This is in contrast to the problem of forgetting, where restricting the original ontology while preserving all representable entailments [19] or obtaining an equivalent set of formulae [32] is the main goal. Thus, the readability of the forgetting result has so far received little attention. For abduction, aside from the conditions in Definition 1, the readability of the hypotheses should be considered to provide insight into unseen observations.

Therefore, it is useful to compare the hypotheses produced by both approaches to forgetting-based abduction: the first using the resolution-based approach of LETHE, and the second using the Ackermann approach of FAME. Consider the following example:

Example 6. Let the background ontology \mathcal{O} contain the following axioms:

$$\begin{aligned} Pogona &\sqsubseteq \exists \text{livesIn}.(Arid \sqcap Woodlands) & Woodlands &\sqsubseteq Habitat \\ EucalyptForest &\sqsubseteq Woodlands & EucalyptForest(SpringbrookPark) & \end{aligned}$$

and consider the observation $\psi = \exists \text{livesIn}.Woodlands(Gary)$. Case (1): let \mathcal{S}_A include all symbols in \mathcal{O} except *Woodlands*, i.e. $\mathcal{F} = \{Woodlands\}$. The hypotheses obtained using LETHE and FAME respectively are:

$$\begin{aligned} \mathcal{H}_L &= Pogona \sqcup \exists \text{livesIn}.EucalyptForest(Gary) \\ \mathcal{H}_F &= \top \sqsubseteq \exists \nabla.(Pogona \sqcap \forall \text{livesIn}^-(\neg Arid \sqcup \neg Habitat \sqcup \exists \text{livesIn}^-. \{Gary\}) \\ &\quad \sqcup \exists \nabla.(\{Gary\} \sqcap \exists \text{livesIn}.EucalyptForest), \end{aligned}$$

where livesIn^- denotes the inverse of the role *livesIn*.

Example 6 illustrates a potential drawback of utilising a more expressive forgetting approach: the hypothesis produced can be more difficult to interpret, as seen by the additional syntactic redundancy in the first disjunct of \mathcal{H}_F . Despite this, the extra expressivity in the target language of FAME can be useful. Since FAME's solution preserves additional entailments compared to LETHE's, it may lead to additional explanations (disjuncts) in the final hypothesis. In Example 6, if \mathcal{F} is extended to $\mathcal{F} = \{Woodlands, EucalyptForest\}$, then $\mathcal{H}_L = Pogona(Gary)$, whereas $\mathcal{H}_F = \top \sqsubseteq \exists \nabla.(Pogona \sqcap \forall \text{livesIn}^-(\neg Arid \sqcup \neg Habitat \sqcup \exists \text{livesIn}^-. \{Gary\}) \sqcup \exists \nabla.(\{Gary\} \sqcap \exists \text{livesIn}^-. \{SpringbrookPark\}))$. The second disjunct in \mathcal{H}_F is equivalent to $\text{livesIn}(Gary, SpringbrookPark)$, an explanation that is absent from \mathcal{H}_L .

In Example 6, the following relations hold: $\mathcal{O}, \mathcal{H}_L \models \mathcal{H}_F$ and $\mathcal{O}, \mathcal{H}_F \not\models \mathcal{H}_L$. This indicates that the hypotheses obtained by using FAME in the forgetting step 1 in Figure 1 can be weaker than those obtained using LETHE. This is to be expected, since the forgetting solution computed by FAME can be stronger than the uniform interpolant produced by LETHE due to the extended language of FAME's solution. Thus, \mathcal{H}_F can be weaker than \mathcal{H}_L under the background ontology, since these are obtained by negating the reduced forgetting solutions.

6 Experimental Evaluation

To perform a preliminary evaluation of the new forgetting-based abduction method, a prototype was implemented in Java using the OWL-API¹. Since one of the primary aims of this work is to assess the benefit of utilising FAME for abduction rather than LETHE, the forgetting module in the abduction method utilises either of the two tools: LETHE² or FAME [32, 33].

Since no benchmarks exist for abduction problems in DLs, a challenging aspect of experimentally evaluating tools for abduction is the generation of appropriate observations. These observations should not violate the conditions in

¹ <http://owlapi.sourceforge.net/>

² <http://www.cs.man.ac.uk/koopmanp/lethe/index.html>

Ontology Name	DL	TBox Size	ABox Size	Num. Concepts	Num. Roles
BFO	\mathcal{EL}	52	0	35	0
LUBM	\mathcal{EL}	87	0	44	24
HOM	\mathcal{EL}	83	0	66	0
DOID	\mathcal{EL}	7892	0	11663	15
SYN	\mathcal{EL}	15352	0	14462	0
ICF	\mathcal{ALC}	1910	6597	1597	41
Semintec	\mathcal{ALC}	199	65189	61	16
OBI	\mathcal{ALC}	28888	196	3691	67
NATPRO	\mathcal{ALC}	68565	42763	9464	12

Table 1. Characteristics of the experimental corpus.

Definition 1, i.e. they should be consistent with the corresponding background ontology, but should also not be entailed by it. However, it is also necessary to consider the forms the observations take. While it is not possible to know exactly what forms the observations may take outside of case studies, it is important to try to emulate information that may be seen in practice. To do this, in this work the observations were generated randomly using existing information in each background ontology, as in [6]. Specifically, to generate a set of candidate observations for a background ontology, the concepts occurring in the axioms of the ontology were scanned and stored. These were selected at random and combined with \mathcal{ALC} operators, also at random, to encourage variety. Each candidate observation was checked using HermiT to determine if it satisfied the conditions in Definition 1. If it did not, it was discarded. This process was repeated until the required number of observations was obtained.

For the first experiment, the aim was to compare the performance of the abduction method using LETHE and FAME in terms of time and the hypotheses obtained. The set of observations was restricted to those that can be handled by the abduction system using LETHE as in [6]. These included any \mathcal{ALC} concept assertion, with at least one concept symbol that is not \top or \perp , over a single individual. For ontologies with an ABox, each individual in the observations was an existing individual, while for those without an ABox the individual was a fresh one. The restriction to one individual was performed because in the OWL API disjunctive assertions cannot be expressed for \mathcal{ALC} . For each observation, the forgetting signature \mathcal{F} was set to one random concept symbol in the observation ψ . In this way, the results are indicative of a single step of the abduction procedure, assuming that the user has no additional information that would lead them to further restrict the set of abducibles \mathcal{S}_A . Thus, the hypothesis obtained is one of the weakest possible hypotheses (least assumptions). It is assumed that the user would proceed to further refine the hypothesis by forgetting symbols from the hypotheses obtained. The time limit in this experiment was 300s for both the forgetting and filtering steps respectively.

Number of	Mean	Median	90th Percentile	Maximum
TBox Axioms	1374	328	3830	8535
ABox Assertions	1014	26	2472	10889
Concepts	783	221	2232	6446
Roles	54	21	76	1043
Individuals	558	23	1605	8220

Table 2. Characteristics of the experimental corpus used in experiment 2.

The corpus used in experiment 1 is the same as the one used in [6], which consists of ontologies taken from NCBO Bioportal ³, OBO Foundry ⁴, the LUBM benchmark [14] and the Semintec⁵ financial ontology. The choice of corpus is detailed in [6]. The statistics of this corpus are shown in Table 1.

The aim of the second experiment was to assess the performance of FAME with the approximate and full filtering settings of the abduction approach in the less restrictive setting of this paper. The corpus was extracted from a snapshot of NCBO Bioportal [25]. The observations were generated in the same way as in experiment 1, but without the restrictions which excluded role assertions. The forgetting signature in each case included at least one symbol from the observation, including role symbols. Again the assumption is that, unlike for forgetting, the aim is not to restrict a background ontology to a portion of the original, but to produce a space of independent explanations that does not make too many assumptions without sufficient prior knowledge about the observation. Thus, the forgetting signature was set to small portions of the symbols in the ontology. The timeout for the method was set to 1000 seconds in total. The success rates reported include cases for which FAME failed to forget at least one symbol, or one of these two steps exceeded the time limit.

The requirements for the ontologies in the extracted corpus were as follows. (1) They should be parsable by the OWLAPI, FAME and Hermit. This excludes cases for which there were errors in loading the ontology into any of these systems. For this reason, the ontologies were also restricted to those containing at most 100,000 axioms. (2) The observation generation method should succeed after 2,000 attempts. This was done to exclude ontologies for which it is not possible to generate a sufficient number of non-entailed, consistent observations for the given ontology. (3) The ontology should contain an ABox, since the main benefit of this less restrictive setting is that information in the ABox can be used to produce hypotheses that utilise local information about an individual and its relationships with other individuals in the ABox. The final corpus contained 50 ontologies, the statistics of which are summarised in Table 2.

All ontologies were preprocessed into their \mathcal{ALC} fragments, since this is the setting of this work. To do this, axioms not representable in \mathcal{ALC} were removed. Others that were representable in \mathcal{ALC} were translated using simple conversions.

³ <https://bioportal.bioontology.org/>

⁴ <http://www.obofoundry.org/>

⁵ <http://www.cs.put.poznan.pl/alawrynowicz/semintec.htm>

Ont. Name	Mean Time /s \mathcal{H}_L	Mean Time /s \mathcal{H}_F	Max Time/s \mathcal{H}_L	Max Time/s \mathcal{H}_F	Mean Disjuncts \mathcal{H}_L	Mean Disjuncts \mathcal{H}_F	$\mathcal{O}, \mathcal{H}_L \equiv \mathcal{O}, \mathcal{H}_F$ %	Success % \mathcal{H}_L	Success % \mathcal{H}_F
BFO	0.05	0.04	0.64	0.26	1.73	1.73	100.0	100.0	100.0
LUBM	0.08	0.06	0.67	0.30	2.53	2.96	60.8	100.0	86.7
HOM	0.06	0.05	0.65	0.26	2.5	2.5	100.0	100.0	100.0
DOID	3.35	3.07	9.97	10.26	4.77	4.77	100.0	100.0	100.0
SYN	6.18	2.84	16.12	13.92	5.6	5.6	100.0	100.0	100.0
ICF	0.96	0.67	3.56	2.16	1.93	1.93	100.0	100.0	100.0
Sem.	2.89	3.09	6.70	6.39	1.10	1.63	58.3	96.7	100.0
OBI	34.47	32.97	120.05	108.85	43.45	42.2	91.3	96.7	100.0
NAT.	46.04	138.24	301.27	688.87	10.61	4.17	62.5	76.7	76.7

Table 3. Results for the first experiment. \mathcal{H}_L (\mathcal{H}_F) indicates results for the abduction system using LETHE (FAME). The time limit for forgetting and filtering was 300 seconds each. For the equivalence check, only cases where both LETHE and FAME computed a hypothesis were compared. For the success rate, failures took into account both times exceeding the timeout and, in the case of FAME, results for which the concept could not be forgotten and results containing definer symbols.

Since fixpoint operators are not utilised in the implementation of FAME, these were not present in the results. Thus, cases requiring fixpoints are deemed to be a failure case and count against the reported success rates. However, these are unlikely to have a significant impact as they are rare in practice [6].

Both experiments were performed on a machine using a 2.8GHz Intel Core i7-7700HQ CPU and 12GB of RAM.

The results for experiment 1 are shown in Table 3. Over most ontologies, utilising FAME resulted in a shorter mean runtime. Two exceptions were the Semintec and NATPRO ontologies. The maximum runtime was longer when using FAME in a few cases, most noticeably over NATPRO, for which it was over double that obtained using LETHE. These differences could be due to the computation of additional explanations requiring the expressivity of FAME’s solution, which would necessitate additional entailment checks during the filtering step. Also, for ontologies with large ABoxes, a significant number of axioms need to be transformed to TBox axioms using nominals, which may increase the time taken. In most cases, the success rate when using LETHE was 100%. The same is true using FAME. In LETHE’s case, failures occurred over the larger and more expressive ontologies, Semintec, OBI and NATPRO. These are due to timeouts, indicating that LETHE took longer than 300 seconds to produce a solution. For FAME, failures can occur due to the incompleteness of FAME’s calculus: all of the failures over the LUBM ontology were due to this characteristic. For the NATPRO ontology, all of the failures observed using FAME were instead due to timeouts. In most cases, the hypotheses \mathcal{H}_L and \mathcal{H}_F were equivalent under the corresponding ontology. This indicates that it should often be possible to express \mathcal{H}_F in \mathcal{ALC} , which may help to improve the readability issue discussed in Example 6 in these cases. Over the LUBM, Semintec, OBI and NATPRO ontologies, a number of the hypotheses produced using FAME were weaker than

those returned using LETHE. This is expected: the forgetting result returned by FAME may be stronger than the uniform interpolant produced by LETHE, and in some cases there may be hypotheses that cannot be expressed without the extra expressivity of FAME’s result. The following is an example taken from the LUBM experiments, demonstrating the benefit of this in practice.

Example 7. For the observation $\psi = \neg Organization(a)$, where a is a fresh individual, the key axioms in the LUBM ontology were:

$$\begin{aligned} Person \sqcap \exists worksFor.Organization &\sqsubseteq Employee \quad College \sqsubseteq Organization \\ Employee &\sqsubseteq Person \sqcap \exists worksFor.Organization \end{aligned}$$

For the forgetting signature $\mathcal{F} = \{Organization\}$, the hypothesis was:

$$\mathcal{H}_F = \top \sqsubseteq \forall \nabla . (\neg \{a\} \sqcup \exists worksFor^-. (\neg Employee \sqcap Person))$$

Other explanations, such as those equivalent to $\neg College(a)$, are redundant with respect to Definition 1(iii) and are removed by the filtering in step 3 of Figure 1. Using LETHE, no hypothesis was produced as the above hypothesis requires the use of the inverse role $worksFor^-$, which cannot be produced by LETHE.

The results for experiment 2 are shown in Table 4. As expected, the approximate filtering took less time than the full filtering across all cases, as it does not perform the additional, expensive entailment checks. The maximum time for the approximate filtering for an \mathcal{F} size of 1 is particularly high. It is likely that for this single case the forgetting solution was particularly large, indicating that the forgotten symbol occurred frequently in the given ontology. The mean number of redundant axioms removed from the forgetting results by the approximate filtering was 2444.6, 2510.4 and 2873.3 for \mathcal{F} sizes of 1, 5% and 10% respectively. The mean additional redundancies removed by the full filtering setting was 11.7, 11.3 and 9.7 axioms respectively. This indicates that in many cases the approximate filtering may be sufficient to obtain a space of explanations that is largely free of redundancies. The success rates indicate that the full filtering setting caused a number of additional timeouts for each size of \mathcal{F} . However, the majority of failures were the result of FAME failing to forget at least one symbol in \mathcal{F} . For the approximate filtering cases, 100%, 100% and 99.5% of failures occurred due to the forgetting step for \mathcal{F} sizes 1, 5% and 10% respectively. For the full filtering cases, the corresponding values were 88.8%, 94.8% and 94.8% respectively. FAME’s failure rates for these abduction experiments are higher than those reported for forgetting experiments [33, 31]. This may be due to the frequency of role symbols occurring in ABox observations for abduction, many of which included role assertions or complex concepts involving roles.

7 Conclusion and Future Work

In this paper the expressivity of forgetting-based abduction was extended using the forgetting system FAME. Role symbols can now be excluded from the

\mathcal{F} Size	Forgetting Time		Approx. Filter Time		Full Filter Time		Successes %	
	Mean	Max	Mean	Max	Mean	Max	Approx.	Full
1	0.05	1.02	0.74	869.63	7.40	880.11	90.3	89.3
5%	0.13	11.15	0.09	28.25	8.29	878.05	81.7	80.9
10%	1.04	75.09	0.06	5.52	6.45	975.24	70.9	70.6

Table 4. Results for experiment 2. Percentages for \mathcal{F} are relative to $sig(\mathcal{O}, \psi)$. All times are in seconds.

abducibles, and observations and hypotheses can now contain role assertions, including conjunctions and disjunctions of these. Hypotheses requiring role inverses can also now be computed. These extensions are useful in practice, as data in the ABox of an ontology can be used to compute more specific hypotheses.

One limitation of the approach is the lack of completeness, due to the fact that FAME uses semantic forgetting, which is not complete. A possible solution is to combine the use of FAME and LETHE, enabling LETHE to forget definer or forgetting symbols in FAME’s result. Future work will also include further, fine-grained experimental evaluation and applications such as concept learning. These will benefit significantly from the enhanced expressivity of the approach.

References

1. Ackermann, W.: Untersuchungen über das Eliminationsproblem der mathematischen Logik. *Mathematische Annalen* **110**, 330–413 (1935)
2. Baader, F., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F.e.: *The description logic handbook*. Cambridge University Press (2003)
3. Bienvenu, M.: Complexity of abduction in the \mathcal{EL} family of lightweight description logics. In: Proc. KR’08. pp. 220–230. AAAI Press (2008)
4. Calvanese, D., De Giacomo, G., Lenzerini, M.: Reasoning in expressive description logics with fixpoints based on automata on finite trees. In: Proc. IJCAI’99. pp. 84–89. AAAI Press (1999)
5. Calvanese, D., Ortiz, M., Simkus, M., Stefanoni, G.: Reasoning about explanations for negative query answers in DL-Lite. *J. Artificial Intelligence Research* **48**, 635–669 (2013)
6. Del-Pinto, W., Schmidt, R.A.: ABox abduction via forgetting in \mathcal{ALC} . In: Proc. AAAI’19. AAAI Press (2019)
7. Doherty, P., Łukaszewicz, W., Szalas, A.: Computing strongest necessary and weakest sufficient conditions of first-order formulas. *Artificial Intelligence* **128**, 143–159 (2001)
8. Du, J., Qi, G., Shen, Y., Pan, J.Z.: Towards practical ABox abduction in large OWL DL ontologies. In: Proc. AAAI’11. pp. 1160–1165. AAAI Press (2011)
9. Du, J., Wan, H., Ma, H.: Practical TBox abduction based on justification patterns. In: Proc. AAAI’17. pp. 1100–1106. AAAI Press (2017)
10. Eiter, T., Gottlob, G.: The complexity of logic-based abduction. *J. ACM* **42**, 3–42 (1995)
11. Elsenbroich, C., Kutz, O., Sattler, U.: A case for abductive reasoning over ontologies. In: Proc. OWL: Experiences and Directions. CEUR workshop proceedings, vol. 216. CEUR-WS.org (2006)

12. Flach, P.A., Kakas, A.C.: On the relation between abduction and inductive learning. *Handbook of Defeasible Reasoning and Uncertainty Management Systems* **4**, 155–196 (2000)
13. Gabbay, D.M., Schmidt, R.A., Szałas, A.: *Second-order quantifier elimination: Foundations, computational aspects and applications*. College Publications (2008)
14. Guo, Y., Pan, Z., Heflin, J.: LUBM: A benchmark for OWL knowledge base systems. *J. Web Semantics* **3**, 158–182 (2005)
15. Halland, K., Britz, K.: ABox abduction in \mathcal{ALC} using a DL tableau. In: *Proc. SAICSIT'12*. pp. 51–58. ACM (2012)
16. Hobbs, J.R., Stickel, M., Martin, P., Edwards, D.: Interpretation as abduction. *Artificial Intelligence* **63**, 69–142 (1993)
17. Klarman, S., Endriss, U., Schlobach, S.: ABox abduction in the description logic \mathcal{ALC} . *J. Automated Reasoning* **46**, 43–80 (2011)
18. Konolige, K.: Abduction versus closure in causal theories. *Artificial Intelligence* **53**, 255–272 (1992)
19. Koopmann, P., Schmidt, R.A.: Uniform interpolation and forgetting for \mathcal{ALC} ontologies with ABoxes. In: *Proc. AAAI'15*. pp. 175–181. AAAI Press (2015)
20. Koopmann, P., Schmidt, R.A.: LETHE: Saturation based reasoning for non-standard reasoning tasks. In: *Proc. ORE'15. CEUR Workshop Proceedings*, vol. 1387, pp. 23–30. CEUR-WS.org (2015)
21. Lambrix, P., Dragisic, Z., Ivanova, V.: Get my pizza right: Repairing missing is-a relations in \mathcal{ALC} ontologies. In: *Proc. JIST'12. LNCS*, vol. 7774, pp. 17–32. Springer (2012)
22. Lin, F.: On strongest necessary and weakest sufficient conditions. *Artificial Intelligence* **128**, 143–159 (2001)
23. Ludwig, M., Konev, B.: Practical uniform interpolation and forgetting for \mathcal{ALC} TBoxes with applications to logical difference. In: *Proc. KR'14*. pp. 318–327. AAAI Press (2014)
24. Lutz, C., Wolter, F.: Foundations for uniform interpolation and forgetting in expressive description logics. In: *Proc. IJCAI'11*. pp. 989–995. AAAI Press (2011)
25. Matentzoglou, N., Parsia, B.: Biportal snapshot 30.03.2017 [data set] (2017), Zenodo, <http://doi.org/10.5281/zenodo.439510>
26. Muggleton, S., Bryant, C.: Theory completion using inverse entailment. In: *Proc. ILP'00. LNCS*, vol. 1866, pp. 130–146. Springer (2000)
27. Pukancová, J., Homola, M.: Tableau-based ABox abduction for the \mathcal{ALCHO} description logic. In: *Proc. DL'17. CEUR workshop proceedings*, vol. 1879. CEUR-WS.org (2017)
28. Raghavan, S., Mooney, R.: Bayesian abductive logic programs. In: *AAAI'10 Workshop on Statistical Relational AI*. pp. 82–87. AAAI Press (2010)
29. Ray, O.: Nonmonotonic abductive inductive learning. *J. Applied Logic* **7**, 329–340 (2009)
30. Wernhard, C.: Abduction in logic programming as second-order quantifier elimination. In: *Proc. FroCoS'13. LNCS*, vol. 8152, pp. 103–119. Springer (2013)
31. Zhao, Y.: *Automated Semantic Forgetting for Expressive Description Logics*. Ph.D. thesis, The University of Manchester, UK (2018)
32. Zhao, Y., Schmidt, R.A.: Concept forgetting in \mathcal{ALCOI} -ontologies using an Ackermann approach. In: *Proc. ISWC'15. LNCS*, vol. 9366, pp. 587–602. Springer (2015)
33. Zhao, Y., Schmidt, R.A.: Forgetting concept and role symbols in $\mathcal{ALCOI}\mu^+(\nabla, \sqcap)$ -ontologies. In: *Proc. IJCAI'16*. pp. 1345–1352. AAAI Press (2016)