

Signature-Based Abduction for Expressive Description Logics

Patrick Koopmann¹, Warren Del-Pinto², Sophie Touret³ and Renate A. Schmidt²

¹Institute for Theoretical Computer Science, Technische Universität Dresden, Germany

²Department of Computer Science, The University of Manchester, United Kingdom

³Automation of Logic Group, Max Planck Institut für Informatik, Germany

Abstract

Signature-based abduction aims at building hypotheses over a specified set of names, the signature, that explain an observation relative to some background knowledge. This type of abduction is useful for tasks such as diagnosis, where the vocabulary used for observed symptoms differs from the vocabulary expected to explain those symptoms. We present the first complete method solving signature-based abduction for observations expressed in the expressive description logic \mathcal{ALC} , which can include TBox and ABox axioms. The method is guaranteed to compute a finite and complete set of hypotheses, and is evaluated on a set of realistic knowledge bases.

1 Introduction

Abductive reasoning is central to knowledge discovery and knowledge processing and has been intensely studied in artificial intelligence, computer science, cognitive science, philosophy and logic [Flach and Kakas, 2000; Kakas, Kowalski, and Toni, 1992; Ray, 2009]. Abduction is the process of explaining new observations using background knowledge. It is an important enabling mechanism for a variety of tasks that require explanations that go beyond what is already implied by existing knowledge, including scientific discovery, database update, belief expansion, diagnostics, planning, language interpretation and inductive learning [Hobbs et al., 1993; Flach and Kakas, 2000]. However, in the description logic (DL) literature, abduction has received much less attention, despite being recognised as important for ontology repair, query update and matchmaking [Elsenbroich, Kutz, and Sattler, 2006; Wei-Kleiner, Dragisic, and Lambrix, 2014; Calvanese et al., 2013; Di Noia, Di Sciascio, and Donini, 2007].

In this paper, the first complete and practical approach for a variant of abduction we call *signature-based abduction* is presented, which solves this problem for inputs expressed in the expressive DL \mathcal{ALC} . In general, the abduction problem considers an observation and a knowledge base (KB), and we are looking for an extension of the KB, called *hypothesis*, that would logically imply the observation. Without further constraints, this definition is too underspecified to be of practical use, as there may be many such extensions, including the observation itself. In signature-based abduction, we restrict the space of solutions by a set of names, the *signature of abducibles*: any hypothesis has to be constructed

using only names from this set. This restriction makes sense since in many applications of abduction, the vocabulary that we expect in an observation differs from the one we would wish for in a helpful explanation. To illustrate, consider the following KB \mathcal{K} :

$$\text{EbolaPatient} \equiv \text{Patient} \sqcap \exists \text{infectedBy.Ebola} \quad (1)$$

$$\exists \text{contactWith.EbolaBat} \sqsubseteq \text{EbolaPatient} \quad (2)$$

$$\text{EbolaPatient} \sqsubseteq \forall \text{infected.EbolaPatient} \quad (3)$$

$$\text{EbolaPatient}(p_1) \quad (4)$$

It states that an `EbolaPatient` is a patient infected by Ebola (1), individuals that were in contact with a bat carrying Ebola have Ebola (2), individuals infected by an Ebola patient have Ebola (3), and individual p_1 is an Ebola patient (4). Suppose it is now observed that patient p_2 also has Ebola; that is, we want to explain the observation `EbolaPatient(p_2)`. A sufficient extension of \mathcal{K} to imply this observation is $\{\text{Patient}(p_2), \exists \text{infectedBy.Ebola}(p_2)\}$. However, these axioms do not really explain anything. To obtain a more meaningful answer, we constrain the explanations to a specified signature related to known causes of Ebola:

$$\Sigma = \{\text{EbolaBat}, \text{infected}, \text{contactWith}\}.$$

An explanation using only names from Σ is for example:

$$\exists \text{contactWith.EbolaBat}(p_2) \vee \text{infected}(p_1, p_2),$$

stating that either p_2 had contact with an Ebola bat, or p_1 infected p_2 . Note that we allow names in the signature to be arbitrarily combined using the constructs of the language. Furthermore, we allow for the use of *disjunction*, which in our setting does not trivialise the problem, but allows for capturing more than one hypothesis in one solution of the abduction problem. Specifically, we are interested in a hypothesis that generalises every possible hypothesis, and is thus semantically minimal among all possible solutions. Ideally, we would present the optimal hypothesis as a disjunction of hypotheses that are *independent* of each other, in the sense that there are no logical relations between the hypotheses. This is not always possible in \mathcal{ALC} , even when extended with disjunctions: as we show, in general, also nominals, inverse roles and fixpoints can be required by such a solution.

Constraining hypotheses to those using only a set of allowed names, called *abducibles*, is a long-standing practice in abductive logic programming [Kakas, Kowalski, and Toni, 1992; Ray, 2009], and has recently also been investigated for first-order logic [Echenim, Peltier, and Tourret, 2017; Echenim, Peltier, and Sellami, 2018]. In the domain of DL knowledge bases, most research either does not consider abducibles, or only in restricted forms.

Elsenbroich, Kutz, and Sattler [2006] motivate abduction in DL ontologies through several use cases, give desirable properties of a sensible abduction operator, and distinguish between *ABox*, *TBox* and *knowledge base (KB) abduction*. The example above is an *ABox abduction* problem, where the observation and the hypotheses consist of facts about specific individuals. Methods for *ABox abduction* without signature-restriction were investigated by Klarman, Endriss, and Schlobach [2011] and Halland and Britz [2012], where the hypotheses computed cannot use concept disjunction and negation is restricted to concept names. Based on these methods, Pukancová and Homola [2017, 2018] proposed an approach for *ABox abduction* relying on minimal hitting sets. In ontology repair, often it is interesting to explain not only facts (*ABox assertions*), but also terminological axioms (*TBox axioms*). *TBox abduction* has been studied by Wei-Kleiner, Dragisic, and Lambrix [2014] and Du, Wan, and Ma [2017], though not the signature-based variant.

Abducibles were considered in Bienvenu [2008] for the light-weight DL \mathcal{EL} , but there hypotheses can only be conjunctions over those abducibles, and more complex combinations are not possible. In [Calvanese et al., 2013], signature-based *ABox abduction* for DL-Lite is investigated theoretically. Different to our work, solutions consider only flat *ABoxes* without complex concepts, but may introduce fresh individual names. The first work that uses signature-based abduction as considered here is by Del-Pinto and Schmidt [2019] for *ABox abduction* in \mathcal{ALC} , but with some restrictions: i) the observation can only consider a single individual, and ii) the set of abducibles has to contain all role names. Our approach allows both *TBoxes*, *ABoxes* and mixed observations, and poses no restrictions on the signature, and is thus the first work that solves signature-based abduction in the most general form. This means that both observations and explanations can be in mixed form. Consider the knowledge base $\{A \sqsubseteq \forall r.B, \neg C(b)\}$, observation $\forall r.C(a)$ and signature $\Sigma = \{A, B, C\}$. The solution returned by the approach in this paper is $\{A(a), B \sqsubseteq C\}$ containing both an *ABox* and a *TBox* axiom.

Finding abduction solutions is often implemented as inverse deduction. While deduction infers consequences from a given set of premises, abduction infers missing premises from which the input observations are consequences relative to some background knowledge. In signature-based abduction, the aim is finding the *most general set of hypotheses* over the given signature that entails the input observations. This is the reverse of looking for the *most specific consequences* of the input over a given signature, which is called the *uniform interpolation problem*. Since these are dual notions [Lin, 2001], in principle abduction and uniform interpolation problems reduce to each other

via negation. Specifically, for a KB \mathcal{K} and observation Ψ , we may perform signature-based abduction by computing a uniform interpolant for $\mathcal{K} \wedge \neg\Psi$. This idea has been used for signature-based abduction in first-order logic using second-order quantifier elimination [Doherty, Lukaszewicz, and Szalas, 2001; Gabbay, Schmidt, and Szalas, 2008], a concept closely related to uniform interpolation. However, so far complete methods for uniform interpolation only exist for DLs that are not closed under negation [Koopmann and Schmidt, 2015; Zhao and Schmidt, 2017; Lutz and Wolter, 2011]. Hence we cannot directly use existing approaches to uniform interpolation, but must develop new methods that can deal with negated KBs. This is why the method by Del-Pinto and Schmidt [2019], which uses the existing uniform interpolation tool *LETHE* [Koopmann, 2020], only supports a restricted abduction setting.

It turns out the seemingly small extension of adding negation brings significant new challenges to uniform interpolation, which can already be seen from the fact that abduction solutions may require a more expressive DL (disjunctive $\mathcal{ALCOI}\mu$) than uniform interpolants (disjunctive $\mathcal{ALC}\nu$) [Koopmann and Schmidt, 2015]. But even with a uniform interpolation method that solves this issue, there remains a practical problem. Uniform interpolation for \mathcal{ALC} is challenging, as solutions can in the worst case be of size triple-exponential in the size of the input [Lutz and Wolter, 2011]. This challenge does impact practical implementations, which usually only perform well with signatures that are either very small or very large [Koopmann and Schmidt, 2015; Zhao and Schmidt, 2017; Chen et al., 2019]. Moreover, as we would use $\mathcal{K} \wedge \neg\Psi$ as input, most of the computed consequences would only depend on \mathcal{K} and thus have no relevance to the abduction problem of explaining Ψ . In practice, we expect Ψ to be considerably smaller than \mathcal{K} , so that the majority of computed consequences would be useless. In [Del-Pinto and Schmidt, 2019], irrelevant consequences are removed using a filtering technique and post-processing. A more efficient solution, essential for larger KBs, is to not compute irrelevant consequences in the first place.

In this paper, we solve the signature-based abduction problem for \mathcal{ALC} using ideas from the uniform interpolation method for \mathcal{ALC} KBs presented in [Koopmann and Schmidt, 2015]. This method is generalised to deal with Boolean \mathcal{ALC} KBs as input, using a modification of the *set-of-support strategy* [Plaisted, 1994] to prune irrelevant inferences. The method by Koopmann and Schmidt [2015] uses a resolution-based calculus to compute relevant inferences for a given signature. Specific to this approach is that new concept names are introduced during the process, so-called *definers*, which are eliminated using simple unfolding operations in a final post-processing step. In our setting, the problem becomes more challenging, leading to a more complex resolution calculus and a more involved definer elimination step during which inverse roles and nominals are introduced. The extended calculus and the generalised form of the input lead to special challenges also in the implementation, which for practicality has to determine required rule applications in a goal-oriented way. Further post-processing is applied to compute abduction solutions that are not only semantically

minimal, but are also presented in the natural form of a disjunction of alternative hypotheses, which are *independent* in the sense that no hypothesis is logically implied by another.

To summarise, the contributions of this paper are: We solve the signature-based KB abduction problem for \mathcal{ALC} in the most general form, allowing for arbitrary signatures, as well as KBs, as input. We establish a minimal extension of \mathcal{ALC} that is sufficient to cover the set of all possible solutions to a signature-based abduction problem in a single disjunction of KBs. We extend the resolution-calculus in [Koopmann and Schmidt, 2015] to deal with Boolean \mathcal{ALC} -KBs. We devise a modification of the set-of-support strategy that can be used with this calculus. We develop a method to compute solutions to the signature-based abduction problem that are semantically minimal in the sense considered here. We introduce techniques to avoid irrelevant explanations and remove dependencies between explanations. We implemented and evaluated the method on a corpus of realistic ontologies, finding that solutions can be computed in practice and usually consist of few small hypotheses.

Additional details and proofs are provided in the extended version of the paper [Koopmann et al., 2020].

2 Description Logic Knowledge Bases

We recall the DLs \mathcal{ALC} , $\mathcal{ALCOI}\mu$, and $\mathcal{ALCOI}\mu^\nabla$ relevant to this paper [Baader et al., 2003]. Let N_V , N_C , N_R and N_I be pair-wise disjoint, countably infinite sets of respectively *concept variable*-, *concept*-, *role*- and *individual names*. A *role* is an expression of the form r , r^- (inverse role) or ∇ (universal role), where $r \in N_R$. $\mathcal{ALCOI}\mu^\nabla$ concepts are then built according to the following syntax rule, where $A \in N_C$, $X \in N_V$, $a \in N_I$ and R is a role:

$$C ::= A \mid \{a\} \mid X \mid \neg C \mid C \sqcup C \mid \exists R.C \mid \mu X.C.$$

We additionally require that for *least fixpoint concepts* $\mu X.C$, X occurs in C only under an even number of negations (\neg). This is a standard requirement to ensure that the semantics of the least fixpoint concept is always well-defined [Calvanese, De Giacomo, and Lenzerini, 1999]. We use $C_1[C_2 \mapsto C_3]$ to denote the result of replacing in C_1 every C_2 by C_3 . Further concepts are introduced as abbreviations: $\top = (A \sqcup \neg A)$ (for an arbitrary $A \in N_C$), $\perp = \neg \top$, $C_1 \sqcap C_2 = \neg(\neg C_1 \sqcup \neg C_2)$, $\forall R.C = \neg \exists R.\neg C$ and $\nu X.C[X] = \neg \mu X.\neg C[X \mapsto \neg X]$.

A concept is *closed* if every concept variable name X occurs in the scope of the least fixpoint operator μX . *Knowledge bases* are sets of *concept inclusions* (CIs) of the form $C_1 \sqsubseteq C_2$ and *assertions* of the form $C_1(a)$, $r(a, b)$, where C_1, C_2 are closed concepts, $a, b \in N_I$ and $r \in N_R$. CIs and assertions are collectively called *axioms*.

A *Boolean KB* is built according to the syntax rule

$$\mathcal{K} ::= \alpha \mid \neg \mathcal{K} \mid \mathcal{K} \wedge \mathcal{K} \mid \mathcal{K} \vee \mathcal{K},$$

where α is an axiom. We identify each KB with the Boolean KB that is the conjunction of all its axioms. If a (Boolean) KB/concept does not use the universal role, it is in $\mathcal{ALCOI}\mu$, if it does not use inverse roles, fixpoint operators and *nominals* $\{a\}$, it is in \mathcal{ALC} . For an expression E ,

$\text{sig}(E)$ denotes the concept and role names occurring in E . For a signature $\Sigma \subseteq N_C \cup N_R$, a Σ -*axiom* is an axiom α with $\text{sig}(\alpha) \subseteq \Sigma$.

The semantics is defined in terms of *interpretations* $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}, \cdot^{\mathcal{I}, \nu} \rangle$, with $\Delta^{\mathcal{I}}$ a set of *domain elements* and $\cdot^{\mathcal{I}}$ the *interpretation function* mapping every $a \in N_I$ to some $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, every $A \in N_C$ to some $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, every $r \in N_R$ to some $r \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and is extended to inverse and universal roles by $(r^-)^{\mathcal{I}} = (r^{\mathcal{I}})^-$ and $\nabla^{\mathcal{I}} = \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. A *valuation for \mathcal{I}* is a function $\nu : N_V \rightarrow 2^{\Delta^{\mathcal{I}}}$. Given a valuation ν for \mathcal{I} , $X \in N_V$ and $W \subseteq \Delta^{\mathcal{I}}$, $\nu[X \mapsto W]$ is the valuation identical to ν except $\nu(X) = W$. Concepts C and valuations ν are mapped to subsets of $C^{\mathcal{I}, \nu} \subseteq \Delta^{\mathcal{I}}$ by

$$\begin{aligned} X^{\mathcal{I}, \nu} &= \nu(X), & A^{\mathcal{I}, \nu} &= A^{\mathcal{I}}, & \{a\}^{\mathcal{I}, \nu} &= \{a^{\mathcal{I}}\}, \\ (\neg C)^{\mathcal{I}, \nu} &= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}, \nu}, & (C \sqcap D)^{\mathcal{I}, \nu} &= C^{\mathcal{I}, \nu} \cap D^{\mathcal{I}, \nu}, \\ (\exists R.C)^{\mathcal{I}, \nu} &= \{d \in \Delta^{\mathcal{I}} \mid \exists (d, e) \in R^{\mathcal{I}} : e \in C^{\mathcal{I}, \nu}\}, & \text{and} \\ (\mu X.C)^{\mathcal{I}, \nu} &= \bigcap \{W \subseteq \Delta^{\mathcal{I}} \mid C^{\mathcal{I}, \nu[X \mapsto W]} \subseteq W\}. \end{aligned}$$

Intuitively, concepts $\mu X.C[X]$ are equivalent to an unbounded disjunction of concepts:

$$C[\perp] \sqcup C[C[\perp]] \sqcup C[C[C[\perp]]] \sqcup C[C[C[C[\perp]]]] \sqcup \dots$$

$(\mu X.C)^{\mathcal{I}, \nu}$ is independent of the value of $\nu(X)$. Thus, for closed concepts C , $C^{\mathcal{I}, \nu}$ is independent of ν . We extend $\cdot^{\mathcal{I}}$ to closed concepts by setting $C^{\mathcal{I}} = C^{\mathcal{I}, \nu}$ for an arbitrary ν .

We define *satisfaction of Boolean KBs \mathcal{K} in \mathcal{I}* , in symbols $\mathcal{I} \models \mathcal{K}$, by $\mathcal{I} \models C \sqsubseteq D$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, $\mathcal{I} \models C(a)$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$, $\mathcal{I} \models r(a, b)$ if $(a, b) \in r^{\mathcal{I}}$, $\mathcal{I} \models \neg \mathcal{K}$ if $\mathcal{I} \not\models \mathcal{K}$, $\mathcal{I} \models \mathcal{K}_1 \wedge \mathcal{K}_2$ if $\mathcal{I} \models \mathcal{K}_1$ and $\mathcal{I} \models \mathcal{K}_2$, and $\mathcal{I} \models \mathcal{K}_1 \vee \mathcal{K}_2$ if $\mathcal{I} \models \mathcal{K}_1$ or $\mathcal{I} \models \mathcal{K}_2$. We then also say \mathcal{I} is a model of \mathcal{K} . We write $\mathcal{K}_1 \models \mathcal{K}_2$ if every model of \mathcal{K}_1 is a model of \mathcal{K}_2 .

3 Signature-Based Abduction

Our aim is to produce the least assumptive hypothesis within a given signature of allowed *abducible* names.

Definition 1. Let \mathcal{K} be an \mathcal{ALC} KB (the background knowledge), Ψ a set of CIs and assertions in \mathcal{ALC} (the observation), and Σ a signature (the set of abducibles). The signature-based abduction problem $\langle \mathcal{K}, \Psi, \Sigma \rangle$ is then to compute a hypothesis $\mathcal{H} = \bigvee_{i=1}^n \mathcal{K}_i$ that satisfies all of the following conditions:

- A1** $\mathcal{K} \wedge \mathcal{H} \not\models \perp$,
- A2** $\mathcal{K} \wedge \mathcal{H} \models \Psi$,
- A3** $\text{sig}(\mathcal{H}) \subseteq \Sigma$,
- A4** for any \mathcal{ALC} KB \mathcal{H}' satisfying **A2** and **A3**, $\mathcal{K} \wedge \mathcal{H}' \models \mathcal{H}$.

If \mathcal{H} satisfies Conditions **A1**, **A2** and **A3**, it is a hypothesis for Ψ in Σ . If it additionally satisfies Condition **A4**, it is an optimal hypothesis.

Conditions **A1** and **A2** are the standard conditions for the abduction problem: the hypothesis should be consistent with the background knowledge and should *explain* the observation. Condition **A3** is what makes this a *signature-based* abduction problem. Condition **A4** finally requires the solution to be *optimal*, in the sense that every possible explanation in \mathcal{ALC} is covered. This is often captured by the

notion of *semantic minimality*: \mathcal{H} should be the most general hypothesis, the one that makes the least assumptions, among all possible hypotheses. Represented as a disjunction, \mathcal{H} can be seen as a collection of possible hypotheses for the observation. The notion of semantic minimality is often considered without disjunctions [Klarman, Endriss, and Schlobach, 2011; Halland and Britz, 2012]. Note that in the current case, allowing for disjunctions does not trivialise the problem, as the signature of abducibles restricts the space of solutions. The problem of finding a hypothesis using a *minimal* number of *independent* solutions, i.e., using only disjunctions where necessary, is considered in Sect. 5.

It is possible that the only KBs that satisfy Conditions **A2**–**A3** are inconsistent. In this case, due to Condition **A1**, the signature-based abduction problem has no solution. In all other cases, an optimal hypothesis exists and is computed by our method. Condition **A4** ensures that, if there is a hypothesis, there is also an optimal solution unique up to logical equivalence.

Condition **A4** only refers to hypotheses in \mathcal{ALC} , and both observation and explanation are expressed in \mathcal{ALC} . However, for a single hypothesis to cover every possible hypothesis, further expressivity might be needed, as the set of possible hypotheses might be infinite. The first reason is that additional individuals may play a role in a hypothesis, the second reason is that the ontology may entail cyclic CIs. As a consequence, optimal hypotheses might need the use of inverse roles, nominals, or fixpoint expressions. We illustrate the need for further expressivity with the example from the introduction.

Inverse Roles. For every individual name $a \in N_I$, the following KB is a hypothesis in \mathcal{ALC} for the observation $\text{EbolaPatient}(p_2)$:

$$\{ \exists \text{contactWith.EbolaBat}(a), \text{infected}(a, p_2) \}.$$

This makes one hypothesis for each individual in the countably infinite set N_I . One way around this problem could be to alter the requirements in Def. 1, by enforcing individual names to be taken from a finite set. But then we would have a large number of solutions that are all identical modulo renaming of individual names, which is neither convenient nor insightful. Another solution is to use inverse roles, which allow coverage of all of the above hypotheses:

$$(\exists \text{infected}^- . \exists \text{contactWith.EbolaBat})(p_2).$$

Nominals. A similar problem occurs when additional individuals can connect two individuals from the observation and background knowledge. Hypotheses such as

$$\{ \text{infected}(p_1, a), \text{infected}(a, p_2) \},$$

where $a \in N_I$, cannot be captured using inverse roles alone, but can using nominals:

$$(\exists \text{infected} . \exists \text{infected} . \{p_2\})(p_1).$$

Least Fixpoints. Finally, due to the cyclic axiom (3), any chain of infections connecting a known Ebola patient to p_2 is a valid hypothesis. To cover all of these unbounded chains, we use the least fixpoint operator. Using all three constructs, we can represent the optimal hypothesis as follows:

$$\mu X . (\exists \text{contactWith.EbolaBat} \\ \sqcup \exists \text{infected}^- . \{p_1\} \sqcup \exists \text{infected}^- . X)(p_2),$$

which could be read as: “*One the following happened: 1) p_2 was in contact with an Ebola carrying bat, 2) p_2 was infected by p_1 , or 3) p_2 was infected by someone else to whom Conditions 1–3 apply.*”

From an optimal hypothesis with least fixpoints, explanations without fixpoints can be easily obtained by unfolding. Furthermore, inverse roles and nominals are only needed under existential role restrictions, so that we can often reconstruct the hypotheses without those constructs using additional individual names. We note that while least fixpoints are not supported by standard DL reasoners, satisfiability of KBs with *greatest* fixpoints can be decided by using auxiliary concept names [Koopmann and Schmidt, 2015]. The same technique can be used to decide entailment of KBs with least fixpoints.

4 Computing Optimal Hypotheses

The general idea for solving abduction problems $\langle \mathcal{K}, \Psi, \Sigma \rangle$ is to consider the Boolean KB $\mathcal{K} \wedge \neg \Psi$, and eliminate the names outside of Σ similar to how it is done in the resolution-based method for uniform interpolation for \mathcal{ALC} -KBs presented in [Koopmann and Schmidt, 2015]. The resulting KB can then be used to construct a hypothesis by negating it again—if this hypothesis is consistent with \mathcal{K} , it is an optimal abductive explanation, and otherwise, such a hypothesis does not exist.

In uniform interpolation, the aim is to compute, for a given KB \mathcal{K} and signature Σ , a KB \mathcal{K}_Σ such that $\text{sig}(\mathcal{K}_\Sigma) \subseteq \Sigma$ and for every Σ -axiom α , $\mathcal{K} \models \alpha$ iff $\mathcal{K}_\Sigma \models \alpha$. This KB \mathcal{K}_Σ is computed in [Koopmann and Schmidt, 2015] by eliminating each name in $\text{sig}(\mathcal{K}) \setminus \Sigma$ one after the other. To eliminate a name, a set of rules is used to perform all relevant inferences on that name. During this process, new names called *definers* are introduced. After a name $\text{sig}(\mathcal{K}) \setminus \Sigma$ has been successfully eliminated, these definers are eliminated again using simple rewriting rules.

While the general structure of our method is similar, there are several additional challenges we have to address.

1. In order to support *Boolean* KBs, the calculus in [Koopmann and Schmidt, 2015] has to be supplemented with further rules.
2. Existing methods for uniform interpolation perform best when the signature or its complement relative to the signature of the ontology are small [Koopmann and Schmidt, 2015; Zhao and Schmidt, 2017; Chen et al., 2019]. For signature-based abduction, we exploit the fact that the observation Ψ is usually small compared to \mathcal{K} . Our method focuses on inferences relevant to $\neg \Psi$ using a modified set-of-support strategy. Effectively, this means we eliminate names in $\neg \Psi$, but not in \mathcal{K} .

3. In uniform interpolation, we can eliminate names one after the other and do not have to reconsider previously eliminated names, which means we can encapsulate the elimination of each name without affecting termination. For abduction, this approach would cause a termination problem, because nothing is eliminated in \mathcal{K} , and previously eliminated names might get propagated back into the hypothesis. Thus, a more integrated approach is needed.
4. Eliminating definers turns out to be more challenging in our setting, and a simple set of rewriting rules is not sufficient anymore. In fact, it is only in this step that nominals and inverse roles are introduced. For uniform interpolation of \mathcal{ALC} -KBs, disjunction and fixpoints are the only required language extension.
5. Finally, we want to compute hypotheses that have the form of a disjunction of alternative hypotheses in $\mathcal{ALCOI}\mu$. It turns out that for the calculus, it is easier to compute the optimal hypothesis first in $\mathcal{ALCOI}\mu^\nabla$ and then reformulate it into a disjunction of $\mathcal{ALCOI}\mu$ -KBs.

Our method to solve the abduction problem $\langle \mathcal{K}, \Psi, \Sigma \rangle$ proceeds using four steps, which we describe in turn below.

Step 1 Transform $\mathcal{K} \wedge \neg\Psi$ into clausal normal form.

Step 2 Eliminate all names outside Σ using the calculus.

Step 3 Express the resulting clause set as a Boolean KB.

Step 4 Negate the result and eliminate universal roles.

4.1 Step 1: Normalisation

Our method uses the following clausal normal form.

Definition 2 (Normal form). Let $N_D \subset N_C$ be a special set of concept names, called definers, and $N_T = N_I \cup \{x\}$ be the set of terms, where x is a fresh symbol referring to a universally quantified variable. Literals are built according to the following syntax rule:

$$L ::= A(t) \mid \neg A(t) \mid \text{Qr}.D(t) \mid r(a, b) \mid \neg r(a, b),$$

where $t \in N_T$, $r \in N_R \cup \{\nabla\}$, $a, b \in N_I$, $A \in N_C$, $Q \in \{\exists, \forall\}$ and $D \in N_D$. A clause φ is an expression of the form $L_1 \vee \dots \vee L_n$, where each L_i is a literal and at most one literal is of the form $\neg D(x)$ where $D \in N_D$. We treat clauses as sets of literals by ignoring the order of the literals in a clause and assuming that clauses contain no duplicates.

A similar normal form is used in [Koopmann and Schmidt, 2015], however without universal roles, negated role assertions, or clauses that mix role assertions with other literals. Clauses $L_1(x) \vee \dots \vee L_n(x)$ are interpreted as $\top \sqsubseteq L_1 \sqcup \dots \sqcup L_n$, and clauses without variables as disjunctions of ABox assertions. Our method never introduces clauses that mix variable and constant terms.

To make the following more convenient, we also allow clause sets and Boolean KBs to be mixed: Specifically, a generalised KB is a set $\Phi \cup \{\mathcal{K}\}$ containing a set Φ of clauses and a Boolean KB \mathcal{K} , and an interpretation \mathcal{I} is a model of $\Phi \cup \{\mathcal{K}\}$ if both $\mathcal{I} \models \mathcal{K}$ and $\mathcal{I} \models \Phi$. Entailment of axioms from generalised KBs is defined accordingly.

In the rest of this section, we adopt the following naming conventions, where additional primes, sub- or superscripts

may be used: $a, b \in N_I$; $D \in N_D$; $t \in N_T$; $A, B \in N_C$; $r \in N_R \cup \{\nabla\}$; $Q \in \{\exists, \forall\}$; C is a concept; $L(t)$ is a literal with argument t ; and φ is a clause. A *definer-free* Boolean KB \mathcal{K} is such that $\text{sig}(\mathcal{K}) \cap N_D = \emptyset$.

We describe how $\mathcal{K} \wedge \neg\Psi$ is turned into a set of clauses that preserves all definer-free entailments. The observation Ψ is a conjunction of assertions and CIs, and the negated CI $\neg(C \sqsubseteq D)$ is equivalent to $(\exists \nabla.(C \sqcap \neg D))(a)$, where a can be any individual. Hence, $\neg\Psi$ can be equivalently represented as a disjunction of negated assertions. Using standard logical laws and the equivalence between $C_1 \sqsubseteq C_2$ and $\top \sqsubseteq \neg C_1 \sqcup C_2$, we ensure that every CI has \top on the left-hand side and negation occurs only in front of concept names. Definners are then used to replace concepts under role restrictions: every $\text{Qr}.C$ gets replaced by $\text{Qr}.D_C$, and we add the CI $\top \sqsubseteq \neg D_C \sqcup C$, where D_C is a definer uniquely associated with C . Using standard CNF transformations we can then make sure that every CI is of the form $\top \sqsubseteq L_1 \sqcup \dots \sqcup L_n$, equivalent to the clause $L_1(x) \vee \dots \vee L_n(x)$, and that the negated observation corresponds to a set of clauses without variables.

Example 1. Let $\mathcal{K} = \{A_1 \sqsubseteq B, B \sqsubseteq \exists r.B\}$, $\Psi = \{B(a), A_2 \sqsubseteq \exists r.B\}$ and $\langle \mathcal{K}, \Psi, \Sigma \rangle$ be an abduction problem for some Σ . We first represent $\neg\Psi$ as $\neg B(a) \vee \exists \nabla.(A_2 \sqcap \neg \exists r.B)(a)$. After normalisation, we obtain $\Phi = \Phi_{\mathcal{K}} \cup \Phi_{\neg\Psi}$, where $\Phi_{\mathcal{K}} = \{\neg A_1(x) \vee B(x), \neg B(x) \vee \exists r.D_1(x), \neg D_1(x) \vee B(x)\}$ and $\Phi_{\neg\Psi} = \{\neg B(a) \vee \exists \nabla.D_2(a), \neg D_2(x) \vee A_2(x), \neg D_2(x) \vee \forall r.D_3(x), \neg D_3(x) \vee \neg B(x)\}$.

4.2 Step 2: Elimination of Names Outside Σ

Names outside of Σ are eliminated using a resolution procedure with a special set-of-support strategy that is based on the calculus in Fig. 1. We first describe these rules and then discuss how they are used. Rules RA, RQ, R \forall -1 and R ∇ are taken as is, or adapted, from the calculus by Koopmann and Schmidt [2015] for \mathcal{ALC} KBs. The other rules are new.

A *substitution* is a function replacing all occurrences of x by a given term in a clause (or literal, or term) and a *unifier* of two terms t_1 and t_2 is a substitution σ s.t. $\sigma(t_1) = \sigma(t_2)$. The rules RA, RQ and R \forall -1 rely on the most general unifier (mgu) σ of t_1 and t_2 , i.e., the identity if $t_1 = t_2$ or a function mapping x to $a \in N_I$ if one of t_1 and t_2 is x and the other is a . If there is no mgu, a rule cannot be applied. Inferences are also forbidden if the resulting clause contains more than a single literal of the form $\neg D(x)$ where $D \in N_D$.

Let N_D^Φ denote the set of definers introduced in Φ in **Step 1**. Every subset \mathbf{D} of N_D^Φ is mapped to a unique definer $D_{\mathbf{D}} \in N_D$ s.t. $D_{\{D_1\}} = D_1$ for any $D_1 \in N_D^\Phi$. Intuitively, $D_{\mathbf{D}}$ represents $D_1 \sqcap \dots \sqcap D_n$ for $\mathbf{D} = \{D_1, \dots, D_n\}$. Every time RQ is applied on definers $D_{\mathbf{D}_1}$ and $D_{\mathbf{D}_2}$, clauses $\neg D_{\mathbf{D}_1 \cup \mathbf{D}_2}(x) \vee D_{\mathbf{D}_1}(x)$ and $\neg D_{\mathbf{D}_1 \cup \mathbf{D}_2}(x) \vee D_{\mathbf{D}_2}(x)$ are added, unless they already exist. This can make further inferences on another concept or role name possible. For instance, we cannot apply RA on the clauses $\neg D_1(x) \vee B(x)$ and $\neg D_3(x) \vee \neg B(x)$ in Ex. 1 since the resulting clause would contain two negative definers of the form $\neg D(x)$. But if we first apply RQ on $\neg B(x) \vee \exists r.D_1(x)$ and $\neg D_2(x) \vee$

RA	$\frac{\varphi_1 \vee A(t_1) \quad \varphi_2 \vee \neg A(t_2)}{(\varphi_1 \vee \varphi_2)\sigma}$
Rr	$\frac{\varphi_1 \vee r(a, b) \quad \varphi_2 \vee \neg r(a, b)}{\varphi_1 \vee \varphi_2}$
RQ	$\frac{\varphi_1 \vee (\text{Qr}.D_{\mathbf{D}_1})(t_1) \quad \varphi_2 \vee (\forall r.D_{\mathbf{D}_2})(t_2)}{(\varphi_1 \vee \varphi_2 \vee \text{Qr}.D_{\mathbf{D}_1 \cup \mathbf{D}_2}(t_1))\sigma}$
R \forall -1	$\frac{\varphi_1 \vee r(t_1, b) \quad \varphi_2 \vee (\forall r.D)(t_2)}{(\varphi_1 \vee \varphi_2 \vee D(b))\sigma}$
R \forall -2	$\frac{\varphi_1 \vee \neg D(a) \quad \varphi_2 \vee (\forall r.D)(b)}{\varphi_1 \vee \varphi_2 \vee \neg r(b, a)}$
R \exists	$\frac{\varphi_1 \vee (\exists r.D)(t)}{\varphi_1 \vee (\exists \nabla.D)(t)}$
R ∇	$\frac{\varphi \vee (\exists \nabla.D)(t) \quad \neg D(x)}{\varphi}$

Figure 1: Calculus for eliminating concept and role names.

$\forall r.D_3(x)$, this introduces the definer $D_{\{\mathbf{D}_1 \cup \mathbf{D}_3\}}$, that we denote D_{13} to lighten the notations, and the corresponding extra clauses. After applying RA on these clauses and the problematic ones, the obtained clauses $\neg D_{13} \vee B(x)$ and $\neg D_{13} \vee \neg B(x)$ can be resolved on B using RA.

To focus on inferences relevant to the observation, we use an extension of the *set-of-support* strategy [Plaisted, 1994]. The general idea of this strategy is to have a special set Φ_s of *supported clauses*. Inferences are then applied with the side condition that at least one premise is taken from Φ_s , while the newly inferred clauses are added to the set of supported clauses. In our context, we initialise Φ as the normal clausal form of \mathcal{K} ; and Φ_s as the normal form of $\neg\Psi$. The set-of-support strategy thus makes sure that we only perform inferences that are connected to the observation. However, some modification to the standard set-of-support strategy are necessary due to the special role of definers. Intuitively, definers represent complex concepts under role restrictions, which have been introduced either as part of the normalisation step or by rule applications. When adding a new clause to Φ_s that contains a definer, we have to make sure that the connections of this definer to other clauses are now considered in Φ_s as well.

To eliminate a name S , we perform all inferences on S using the above strategy. Afterwards, we can remove from Φ_s all clauses that use S . It is now possible that previously filtered clauses containing S get inferred again when eliminating a different name. To ensure termination, we thus have to keep track of previous inferences, for which we use the set Φ_a .

We describe the algorithm used in **Step 2**. All inferences are performed under the condition that at least one premise is

from Φ_s . The set Φ_a is initialized as $\Phi \cup \Phi_s$. We then repeat the following steps for some name $S \in \text{sig}(\Phi_s) \setminus (\mathbf{N}_D \cup \Sigma)$ as long as there is such an S .

- F1** Perform all possible inferences on S , all possible RA and R \forall -2 inferences on definers, and all RQ and R \forall -1 inferences that make inferences on S possible. Add all inferred clauses not already in Φ_a to Φ_s and Φ_a .
- F2** Remove from Φ_s all clauses φ with $S \in \text{sig}(\varphi)$.
- F3** Repeatedly move to Φ_s all clauses $\varphi \in \Phi$ in which a $D \in \text{sig}(\Phi_s)$ occurs until no more clauses are added.

The set Φ_a ensures the method terminates, as we may reintroduce a formerly eliminated name S in **F3**. Note that our normal form allows at most double-exponentially many distinct elements in Φ_a . The final state of Φ_s is denoted $\text{Sat}(\Phi, \Phi_s, \Sigma)$. (Sat is short for saturation.)

Example 2. Continuing with Ex. 1, let us assume $\Sigma_1 = \{A_1, A_2, r\}$. We obtain $\text{Sat}(\Phi, \Phi_s, \Sigma_1) = \{\neg A_1(a) \vee \exists \nabla.D_2(a), \neg D_2(x) \vee A_2(x), \neg D_2(x) \vee \exists r.D_3(x), \neg D_2(x) \vee \neg A_1(x) \vee \exists r.D_{13}(x), \neg D_{13}(x)\}$.

The following central theorem is proved in the extended version of the paper.

Theorem 1. Let Φ and Φ_s be normalised clause sets and Σ be a signature. Let $\mathbf{M}_\Sigma = \text{Sat}(\Phi, \Phi_s, \Sigma)$. Then, $\Phi \cup \Phi_s \models \mathcal{K}$ if and only if $\Phi \cup \mathbf{M}_\Sigma \models \mathcal{K}$, for every Boolean \mathcal{ALC} KB \mathcal{K} s.t. $\text{sig}(\mathcal{K}) \subseteq \Sigma$.

4.3 Step 3: Denormalisation

Next, we turn $\mathbf{M}_\Sigma = \text{Sat}(\Phi, \Phi_s, \Sigma)$ into a definer-free Boolean $\mathcal{ALC}\mathcal{O}\mathcal{L}\mu^\nabla$ KB that preserves all entailments in \mathcal{ALC} modulo definer names. Each definer represents a concept that occurs under a role restriction. To eliminate the definers, we compute CIs of the form $D \sqsubseteq C$ (the definition of D), so that occurrences of D can be replaced by C . While such a step is also performed in the resolution-based uniform interpolation method in [Koopmann and Schmidt, 2015], we have to do more in our current setting. For uniform interpolation of \mathcal{ALC} KBs, it is sufficient to look at clauses of the form $\neg D(x) \vee \varphi$ to build the definition $D \sqsubseteq C$ of D . In our setting, special care has to be taken also of clauses of the form $\neg D(a) \vee \varphi$: first, the definition of D has to refer to a , for which we use nominals. Second, if we have a clause $\varphi' \vee \forall r.D(t)$, we also have to consider substitutions of $\neg D(a)$ with $(\forall r^-.C)(a)$, where C is a concept that corresponds to the clause φ' .

Since we applied the RA-rule exhaustively on positive definer occurrences, clauses that contain literals of the form $D(t)$ are not needed anymore at this stage and are removed before further operations are performed. We first introduce a concept-representation of clauses. For every definer D occurring in \mathbf{M}_Σ , we introduce a fresh definer \bar{D} representing $\neg D$. Given a concept C , we denote by C^- the result of replacing every concept of the form $\neg D$ by \bar{D} . Given a clause $\varphi = L_1 \vee \dots \vee L_n$, we define the concept $\varphi^c = L_1^c \sqcup \dots \sqcup L_n^c$, where L^c is defined as:

- C1** C^- if $L = C(x)$,
- C2** $\exists \nabla.(\{a\} \sqcap C^-)$ if $L = C(a)$,
- C3** $\exists \nabla.(\{a\} \sqcap \exists r.\{b\})$ if $L = r(a, b)$,

C4 $\exists \forall. (\{a\} \sqcap \forall r. \neg \{b\})$ if $L = \neg r(a, b)$.

Every clause contains either only variables or only individual names as terms. In the former case, φ^c is the concept described by the clause. In the latter case, φ^c is such that for every interpretation \mathcal{I} and every $d \in \Delta^{\mathcal{I}}$, $\mathcal{I} \models \varphi$ if and only if $d \in (\varphi^c)^{\mathcal{I}}$, that is, either $(\varphi^c)^{\mathcal{I}} = \emptyset$ or $(\varphi^c)^{\mathcal{I}} = \Delta^{\mathcal{I}}$ depending on whether φ is entailed. Recall that all universal roles introduced here are eliminated in **Step 4**.

We then build a set of CIs giving meaning to the definers:

- D1** for every $\varphi \vee \forall r. D(x)$, we add $\overline{D} \sqsubseteq \forall r^- . \varphi^c$,
- D2** for every $\varphi \vee \forall r. D(a)$, we add $\overline{D} \sqsubseteq \forall r^- . (\neg \{a\} \sqcup \varphi^c)$,
- D3** every $\neg D(x) \vee \varphi$ gets replaced by $D \sqsubseteq \varphi^c$,
- D4** for every $\neg D(a) \vee \varphi$, we add $D \sqsubseteq \neg \{a\} \sqcup \varphi^c$,
- D5** in every clause, we replace literals $\neg D(a)$ by $\overline{D}(a)$, and
- D6** every remaining φ that is not a disjunction of (negated) ABox assertions is replaced by $\top \sqsubseteq \varphi^c$.

The transformation makes sure that the meaning of each definer D is captured by CIs $D \sqsubseteq C$, and that all remaining negative occurrences of D now refer to the definer \overline{D} . Addition of axioms happens first, then replacement of clauses.

We apply **D1–D2** on \mathbf{M}_Σ and denote the result by \mathcal{K}_0 .

Lemma 1. *Every model of \mathcal{K}_0 can be transformed into a model of $\mathcal{K}_0 \cup \{\overline{D} \equiv \neg D \mid D \in \text{sig}(\mathbf{M}_\Sigma)\}$ by only changing the interpretation of definers.*

Proof sketch. We show how to construct a model of $\mathcal{K}_0 \cup \{\overline{D} \equiv \neg D \mid D \in \text{sig}(\mathbf{M}_\Sigma)\}$ iteratively by working on one D at a time. Let \mathcal{I} be a model of \mathcal{K}_0 s.t. for some definer $D \in \text{sig}(\mathbf{M}_\Sigma)$, $\mathcal{I} \not\models \overline{D} \equiv \neg D$. The normalisation and the calculus make sure that D occurs either only under existential restrictions or only under universal restrictions. In the former case, we note that no CIs are introduced for \overline{D} , and we can obtain a model \mathcal{I}' with $\mathcal{I}' \models \overline{D} \equiv \neg D$ by setting $\overline{D}^{\mathcal{I}'} = (\neg D)^{\mathcal{I}}$ and keeping everything else the same. Otherwise, D occurs only under universal restrictions, and we transform \mathcal{I} into a model \mathcal{I}' of $\overline{D} \equiv \neg D$ by setting $D^{\mathcal{I}'} = D^{\mathcal{I}} \setminus \overline{D}^{\mathcal{I}}$ and $\overline{D}^{\mathcal{I}'} = \Delta^{\mathcal{I}} \setminus D^{\mathcal{I}'}$. Then we prove that \mathcal{I}' is still a model of \mathcal{K}_0 . \square

Let \mathcal{K}_1 be the result of applying **D1–D6**.

Lemma 2. *Every model of \mathcal{K}_1 can be turned into a model of \mathbf{M}_Σ by only changing the interpretation of definers, and every model of \mathbf{M}_Σ can be extended to a model of \mathcal{K}_0 by setting $\overline{D}^{\mathcal{I}} = \neg D^{\mathcal{I}}$ for all definers D occurring in \mathbf{M}_Σ .*

Proof. Looking at the axioms added in **D3–D6**, one obtains that $\mathcal{K}_1 \models \mathcal{K}_0$ and $\mathcal{K}_0 \cup \{\overline{D} \equiv \neg D \mid D \in \text{sig}(\mathbf{M}_\Sigma)\} \models \mathcal{K}_1$. By Lemma 1, we can thus transform any model \mathcal{I} of \mathcal{K}_1 into a model \mathcal{I}' of \mathcal{K}_0 s.t. $\mathcal{I}' \models \overline{D} \equiv \neg D$ for every $D \in \text{sig}(\mathbf{M}_\Sigma)$, by only changing the interpretation of the definers in \mathcal{I} . \mathcal{I}' is also a model of \mathcal{K}_1 . Inspection of \mathcal{K}_1 shows that every clause $\varphi \in \mathbf{M}_\Sigma$ has an axiom $\alpha \in \mathcal{K}_0$ s.t. $\mathcal{I}' \models \varphi$ if and only if $\mathcal{I}' \models \alpha$. It follows that \mathcal{I}' is a model of \mathbf{M}_Σ .

Let $\overline{\mathcal{I}}$ be a model of \mathbf{M}_Σ and let \mathcal{I}' coincide with $\overline{\mathcal{I}}$, except on $(\overline{D})^{\mathcal{I}'} = \neg D^{\overline{\mathcal{I}}}$ for every definer D occurring in \mathbf{M}_Σ . Again, an inspection of \mathcal{K}_1 reveals that, for every axiom $\alpha \in$

\mathcal{K}_0 , there is some clause $\varphi \in \mathbf{M}_\Sigma$ s.t. $\mathcal{I}' \models \varphi$ if and only if $\mathcal{I}' \models \alpha$. It follows that \mathcal{I}' is a model of \mathcal{K}_1 . \square

As a consequence of Lemma 2, \mathcal{K}_1 preserves all relevant entailments of \mathbf{M}_Σ not using definers, while now every negative occurrence of a definer is in a CI of the form $D \sqsubseteq C$. This allows us now to use the same technique as in [Koopmann and Schmidt, 2015] to eliminate definers. For each definer D occurring in \mathcal{K}_1 , including the introduced definers \overline{D} , we replace all axioms $D \sqsubseteq C_1, \dots, D \sqsubseteq C_n$ by a single axiom $D \sqsubseteq C_1 \sqcap \dots \sqcap C_n$. If there is no such axiom, we add $D \sqsubseteq \top$. We then eliminate each definer D as follows:

- E1** Each CI $D \sqsubseteq C[D]$ (D occurs on both sides of the CI), is replaced by $D \sqsubseteq \nu X. C[D \rightarrow X]$,
- E2** For every CI $\alpha = D \sqsubseteq C$, where $D \notin \text{sig}(C)$, remove α and replace D everywhere by C .

The result is then denoted by \mathcal{K}_2 . The following theorem is a Lemma 1 and 2 together with Ackermann's Lemma [Ackermann, 1935] and Generalised Ackermann's Lemma [Nonnengart and Szalas, 1999], which concern the corresponding transformation in second-order logics.

Theorem 2. *For every definer-free Boolean KB \mathcal{K} , $\mathcal{K}_2 \models \mathcal{K}$ if and only if $\mathbf{M}_\Sigma \models \mathcal{K}$.*

4.4 Step 4: Negate Result

The CIs in the Boolean KB after **Step 3** are only needed to eliminate definers and can be discarded. Thus, we obtain a conjunction of disjunctions of (possibly negated) assertions, which can be directly negated to obtain a disjunction of conjunctions of assertions. From this, we eliminate occurrences of universal roles and make sure that only least fixpoint expressions are used. Universal roles initially represent the negations of CIs. Further universal roles are introduced in **Steps 2 and 3**. In each case, universal roles are only used in existential role restrictions. Eliminating definers only introduces greatest fixpoints. Thus, by pushing negations inside, we can make sure that the hypothesis contains only least fixpoints, and universal roles occur only in universal role restrictions. To obtain the final hypothesis, we pull out universal roles using the following equivalences together with standard DNF transformations:

$$\begin{aligned} \exists r. (C_1 \sqcap \forall \nabla. C_2) &\Leftrightarrow \exists r. C_1 \sqcap \forall \nabla. C_2 \\ \forall r. (C_1 \sqcup \forall \nabla. C_2) &\Leftrightarrow \forall r. C_1 \sqcup \forall \nabla. C_2 \\ (\forall \nabla. C)(a) &\Leftrightarrow \top \sqsubseteq C \end{aligned}$$

Finally, to ensure that the hypothesis returned is consistent, the check $\mathcal{K} \wedge \mathcal{H} \not\models \perp$ is performed using an external reasoner. This eliminates false hypotheses in cases for which there is no suitable hypothesis to explain Ψ in the given signature of abducibles Σ .

The following theorem is a direct consequence of Th. 1 and 2.

Theorem 3. *Let $\langle \mathcal{K}, \Psi, \Sigma \rangle$ be an abduction problem. Applying **Steps 1–4** to this problem produces a disjunction of $ALCO\mathcal{I}\mu$ KBs that is a solution of the abduction problem.*

Example 3. For the clauses M_Σ from Ex. 2, we obtain the Boolean KB $\neg A_1(a) \vee \exists \nabla.(A_1 \sqcap \forall r. \top \sqcap (\neg A_2 \sqcap \exists r. \perp))(a)$, which can be simplified to $\neg A_1(a) \vee \exists \nabla.(A_1 \sqcap \neg A_2)(a)$. The final hypothesis is equivalent to $A_1(a) \wedge (A_1 \sqsubseteq A_2)$.

5 Spaces of Independent Explanations

The abduction problem in Def. 1 searches for semantically minimal hypotheses. However, without extra constraints, the result from the method in the last section may include redundant disjuncts. Consider the following example

$$\begin{aligned} \mathcal{K} &= \{F \sqsubseteq B, D \sqcap E \sqsubseteq \perp, C(a), E(a)\} \\ \psi &= (\exists r. A) \sqcap C \sqsubseteq (\exists r. B) \sqcup D \\ \Sigma &= \{A, B, C, D, E, F\} \end{aligned}$$

and the following three hypotheses $\mathcal{H}_1 = A \sqsubseteq B \vee C \sqsubseteq D$, $\mathcal{H}_2 = A \sqsubseteq B$ and $\mathcal{H}_3 = A \sqsubseteq F$. All of these satisfy **A1–A3** in Def. 1. $\mathcal{K} \wedge \mathcal{H}_3$ entails both \mathcal{H}_1 and \mathcal{H}_2 , and thus \mathcal{H}_3 does not satisfy **A4**. However, it is also the case that $\mathcal{K} \wedge \mathcal{H}_1 \equiv \mathcal{K} \wedge \mathcal{H}_2$. In fact, both \mathcal{H}_1 and \mathcal{H}_2 are semantically minimal. However, \mathcal{H}_2 is shorter, and would thus be preferred as an answer, while \mathcal{H}_1 contains a disjunct that is inconsistent with \mathcal{K} .

To account for these inter-disjunct redundancies, we extend Def. 1, following [Del-Pinto and Schmidt, 2019]. For this, we first need to make sure that disjunctions are pulled out where possible: a disjunction of KBs is in *disjunctive normal form* if every concept of the form $C \sqcup D$ only occurs in a CI or in an assertion under a universal role restriction.

Definition 3. Let $\langle \mathcal{K}, \Psi, \Sigma \rangle$ be an abduction problem and \mathcal{H} a solution to it satisfying Def. 1. Then $\mathcal{H} = \mathcal{K}_1 \vee \dots \vee \mathcal{K}_n$ is a space of independent explanations if it is in *disjunctive normal form* and there is no disjunct \mathcal{K}_i in \mathcal{H} such that:

$$\mathcal{K} \wedge \mathcal{K}_i \models \mathcal{K}_1 \vee \dots \vee \mathcal{K}_{i-1} \vee \mathcal{K}_{i+1} \vee \dots \vee \mathcal{K}_n.$$

In the above example, \mathcal{H}_2 satisfies Def. 3 while \mathcal{H}_1 does not, since $\mathcal{K} \wedge (C \sqsubseteq D) \models A \sqsubseteq B$. It is worth noting that this condition also constrains hypotheses to those that are consistent with the background KB, i.e., if a hypothesis \mathcal{H} satisfies Def. 3 then condition **A1** of Def. 1 will also be satisfied and does not need to be checked separately.

In practice, the entailment tests in Def. 3 are done using a DL reasoner that checks the satisfiability of $\mathcal{K} \wedge \mathcal{K}_i \wedge \neg \mathcal{K}_1 \wedge \dots \wedge \neg \mathcal{K}_{i-1} \wedge \neg \mathcal{K}_{i+1} \wedge \dots \wedge \neg \mathcal{K}_n$. Unfortunately, current DL reasoners do not support fixpoint operators. Positive occurrences of greatest fixpoints, and thus also least fixpoints occurring under the negation in $\neg \mathcal{K}_j$, $1 \leq j \leq n$, can be simulated using auxiliary concept names: for this, we replace $\nu X. C[X]$ by a fresh name D and add the CI $D \sqsubseteq C[D \rightarrow X]$. However, this does not work for positive occurrences of least fixpoints. Thus, in practice, we cannot detect whether a disjunct with fixpoints is redundant. We therefore keep these disjuncts and include them for the redundancy test of the other disjuncts.

6 Evaluation

To evaluate our approach, we implemented a prototype in Scala and Java using the OWL-API.¹ We added redundancy

elimination as in [Koopmann and Schmidt, 2013] to the saturation procedure, and implemented some equivalence-preserving syntactical rewritings to optimise the syntactical shape of the result. While we reused some code from the forgetting tool LETHE [Koopmann, 2020], most had to be reimplemented from scratch due to the set-of-support strategy, and because the calculus required different data structures, indexing and strategies for determining relevant rule applications efficiently. The prototype and an explanation on how to repeat the experiment are available online.² We created a corpus based on the ‘‘DL Instantiation’’ track from the OWL Reasoner Competition 2015 [Parsia et al., 2017], as it provides a balanced mix of different ontologies containing both TBox axioms and ABox assertions. Each ontology was restricted to its *ACC* fragment, where axioms such as domain restrictions or disjointness axioms were turned into corresponding CIs. We excluded inconsistent ontologies, as well as those of more than 50,000 and less than 100 axioms from the resulting set, leaving 202 ontologies. Statistics regarding the corpus are in Table 1.

For these ontologies, we created two sets of abduction problems, (A) and (K), which we describe in detail in the following. Both sets were split into larger and smaller problems to get a feeling of the impact of the observation size. The smaller problems contained observations of up to 10 axioms, the larger problems observations of up to 20 axioms. The exact number of axioms was picked from a uniform distribution. This way, we obtained four sets of observations, in the following denoted (A1), (A2), (K1) and (K2).

The aim of (A) was to generate random observations with little assumptions on their shape. (A) contained observations including randomly generated ABox assertions. For this, we randomly used 2–5 fresh individual names, for which we generated the required number of assertions of the form $A(a)$ and $r(a, b)$ by randomly selecting concept and role names from the background ontology. (A) thus contained ABox abduction problems, but using names that may have only little relation to each other. For (K), we generated less random observations for KB abduction problems that better reflect the structure of the ontology, and use names that are related to each other. To have a realistic mix of CIs and assertions that reflects the typical shape of the background ontology, observations in (K) were generated by selecting the given number of axioms from the background ontology, which were then removed from the background. Since in large ontologies, a fully random selection would result in an observation of unrelated axioms, we first extracted a subset of at least 100 axioms by combining randomly selected *genuine modules*: genuine modules are small subsets of the ontology that contain a given axiom and preserve all entailments over the signature of the subset, and thus contain only axioms that in some way interact with each other [Vescovo et al., 2011]. From these subsets of the ontology, which contained between 100 and 20,979 axioms (median 199), we generated the observations by random selection. For each ontology, we created 150 observations each for the sets (K1) and (K2), and 50 observations each for the sets (A1)

¹<http://owlapi.sourceforge.net/>

²<https://lat.inf.tu-dresden.de/evaluation-kr2020-dl-abduction/>

TBox size (axioms)				ABox size (axioms)			
Min	Max	Mdn	Mean	Min	Max	Mdn	Mean
48	36302	885	3146	32	42429	1424	4610

Table 1: Characteristics of the experimental corpus (Mdn: median).

Initial Hypothesis (s.)					Indep. Explanations (s.)					
Min	Mdn	Mean	P90	Max	Min	Mdn	Mean	P90	Max	
(K1)	0.2	0.6	2.4	1.8	293.2	0.0	0.0	1.2	0.8	292.5
(K2)	0.2	0.4	2.5	2.1	294.8	0.0	0.0	1.1	0.6	293.4
(A1)	0.2	1.9	16.2	51.7	292.0	0.0	0.1	4.6	2.2	282.1
(A2)	0.2	1.6	16.9	49.6	296.9	0.0	0.1	4.9	2.5	293.1

Table 2: Running times (Mdn: median, P90: 90th percentile).

and (A2). We ignored observations that were already entailed by the background ontology, which happened in 3.4% of cases for (A) and in 24.9% of cases for (K).

Signatures for (A) and (K) were generated by selecting respectively 50% and 60% of the background ontology’s signature, where we made sure at least one name from the observation was not an abducible. To reflect the differing relevance of names, each name was chosen with a probability proportional to its number of occurrences. For instance, for an ontology describing paronomies, it would be unlikely to compute hypotheses without the `hasPart` relation. Selecting names this way reduced the number of trivial solutions significantly.

The experiments were run on an Intel Core i5-4590 CPU machine with 3.30GHz and 32 GB RAM, using Debian/GNU Linux 9 and OpenJDK 11.0.5. For each abduction problem, the timeout was set to 5 minutes.

The hypotheses obtained after applying the method in Sect. 4 (computing optimal hypotheses) and those obtained after additionally applying the method in Sect. 5 (removing redundancies) are referred to as H1 and H2 respectively. Computation times are shown in Table 2. The success rate (no timeout nor out-of-memory exception) for (K1) was 96.3% for H1 and 95.3% for H2. For (A1), success rates were 91.3% for H1 and 90.1% for H2. For observations of size up to 20, the rates were very similar: (K2)/(A2) got success rates of 96.4%/91.1% for H1 and 95.4%/89.5% for H2. The size of the observation thus had only very little impact. Since we selected the signatures randomly, in a lot of cases they did not fit to the observation and no hypothesis according to Def. 1 existed. This was the case for 64.5% and 67.4% of the observations in (A1) and in (A2), and for 72.6% and 76.1% of the observations in (K1) and (K2). In the following, we focus on the remaining cases.

Table 3 shows statistics regarding the size of computed hypotheses. Eliminating redundant disjuncts for H2 often reduced the sizes of the hypotheses produced for both experiment (K) and experiment (A). Table 4 shows statistics about the disjuncts in each hypothesis. For (K1)/(K2), disjuncts were required to represent H1 and H2 in 18.0%/20.9% and 7.1%/7.7% of cases respectively. For (A1)/(A2), disjuncts were needed much more often: here the values were 47.2%/43.8% and 36.8%/32.4%. This can be explained by the fact that disjunctive hypotheses are more easy to obtain

Number of TBox axioms					Number of ABox axioms					
Min	Mdn	Mean	P90	Max	Min	Mdn	Mean	P90	Max	
(K1)										
H1	0	2	5.6	7	846	0	3	18.6	9	10294
H2	0	1	2.3	5	105	0	3	6.4	8	4374
(K2)										
H1	0	2	8.9	14	1848	0	6	42.9	19	28817
H2	0	2	3.9	8	1848	0	5	9.0	16	1212
(A1)										
H1	0	0	3.9	0	5762	0	9	30.5	51	8296
H2	0	0	0.0	0	5	0	8	17.1	36	1176
(A2)										
H1	0	0	2.4	0	2646	0	14	52.2	80	21146
H2	0	0	0.0	0	12	0	12	24.2	52	1536

Table 3: TBox and ABox axioms in the computed hypotheses (Mdn: median, P90: 90th percentile).

	Mdn.	Mean	P90.	Max	Mdn.	Mean	P90.	Max
(K1)					(A1)			
H1	1	7.1	3	4751	2	8.3	10	4184
H2	1	1.8	1	729	2	3.7	8	343
(K2)					(A2)			
H1	1	15.3	4	13126	2	11.9	10	7796
H2	1	1.8	1	243	1	3.5	7	256

Table 4: Number of disjuncts in the computed hypotheses (Mdn: median, P90: 90th percentile).

from ABox assertions than from CIs, and that many observations in (K) contained only CIs. Though our method may introduce inverse roles, nominals and fixpoint operators, this was only observed in the minority of cases: inverse roles and nominals were needed in 16.4%/16.6% of cases for (K) and in 1.5%/1.3% of cases for (A). Fixpoint expressions were needed in 3.1%/5.3% of cases for (K) and 0.8%/1.3% of cases for (A).

7 Conclusion

We presented the first general method for signature-based abduction on \mathcal{ALC} KBs. At its center lies a new resolution-based calculus for Boolean KBs, which is used in combination with a set-of-support strategy. We combined it with a simplification technique ensuring the generation of independent explanations and evaluated the overall technique in practice on realistic benchmarks. While the theoretical complexity of our problem remains open, we conjecture that the situation is the same as for uniform interpolation [Lutz and Wolter, 2011], and that optimal, fixpoint-free solutions can be of size triple-exponential in the input. A natural next step is to allow nominals and inverse roles not only in the output, but also in the input, which especially for nominals might prove challenging. Finally, we plan to use our approach for abduction-related tasks such as gentle repairs [Baader et al., 2018] and induction with iterative refinement [Lehmann and Hitzler, 2010].

Acknowledgements

Patrick Koopmann and Sophie Tourret are supported by DFG grant 389793660 of TRR 248. Warren Del-Pinto is a recipient of a RCUK Scholarship.

References

- Ackermann, W. 1935. Untersuchungen über das Eliminationsproblem der mathematischen Logik. *Mathematische Annalen* 110(1):390–413.
- Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2003. *The Description Logic Handbook*. Cambridge University Press.
- Baader, F.; Kriegel, F.; Nuradiansyah, A.; and Peñaloza, R. 2018. Making repairs in description logics more gentle. In *Proceedings of KR 2018*, 319–328. AAAI Press.
- Bienvenu, M. 2008. Complexity of abduction in the \mathcal{EL} family of lightweight description logics. In *Proceedings of KR 2008*, 220–230. AAAI Press.
- Calvanese, D.; Ortiz, M.; Simkus, M.; and Stefanoni, G. 2013. Reasoning about explanations for negative query answers in DL-Lite. *Journal of Artificial Intelligence Research* 48:635–669.
- Calvanese, D.; De Giacomo, G.; and Lenzerini, M. 1999. Reasoning in expressive description logics with fixpoints based on automata on infinite trees. In *Proceedings of IJCAI 99*, 84–89. Morgan Kaufmann.
- Chen, J.; Alghamdi, G.; Schmidt, R. A.; Walther, D.; and Gao, Y. 2019. Ontology extraction for large ontologies via modularity and forgetting. In *Proceedings of K-CAP'19*, 45–52. ACM.
- Del-Pinto, W., and Schmidt, R. A. 2019. ABox abduction via forgetting in \mathcal{ALC} . In *Proceedings of AAAI 2019*, 2768–2775. AAAI Press.
- Di Noia, T.; Di Sciascio, E.; and Donini, F. M. 2007. Semantic matchmaking as non-monotonic reasoning: A description logic approach. *Journal of Artificial Intelligence Research* 29:269–307.
- Doherty, P.; Lukaszewicz, W.; and Szalas, A. 2001. Computing strongest necessary and weakest sufficient conditions of first-order formulas. In *Proceedings of IJCAI 2001*, 145–154. Morgan Kaufmann.
- Du, J.; Wan, H.; and Ma, H. 2017. Practical TBox abduction based on justification patterns. In *Proceedings of AAAI 2017*, 1100–1106. AAAI Press.
- Echenim, M.; Peltier, N.; and Sellami, Y. 2018. A generic framework for implicate generation modulo theories. In *Proceedings of IJCAR 2018*, volume 10900 of *LNCS*, 279–294. Springer.
- Echenim, M.; Peltier, N.; and Tourret, S. 2017. Prime implicate generation in equational logic. *Journal of Artificial Intelligence Research* 60:827–880.
- Elsenbroich, C.; Kutz, O.; and Sattler, U. 2006. A case for abductive reasoning over ontologies. In *Proceedings of the OWLED'06*, volume 216 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Flach, P. A., and Kakas, A. C., eds. 2000. *Abduction and Induction: Essays on their Relation and Integration*, volume 18 of *Applied Logic Series*. Springer, Dordrecht.
- Gabbay, D. M.; Schmidt, R. A.; and Szalas, A. 2008. *Second-Order Quantifier Elimination*. College Publications.
- Halland, K., and Britz, K. 2012. ABox abduction in \mathcal{ALC} using a DL tableau. In *Proceedings of SAICSIT 2012*, 51–58. ACM.
- Hobbs, J. R.; Stickel, M. E.; Appelt, D. E.; and Martin, P. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1-2):69–142.
- Kakas, A. C.; Kowalski, R. A.; and Toni, F. 1992. Abductive logic programming. *Journal of Logic and Computation* 2(6):719–770.
- Klarman, S.; Endriss, U.; and Schlobach, S. 2011. ABox abduction in the description logic \mathcal{ALC} . *Journal of Automated Reasoning* 46(1):43–80.
- Koopmann, P., and Schmidt, R. A. 2013. Forgetting concept and role symbols in \mathcal{ALCH} -ontologies. In *Proceedings of LPAR-19*, volume 8312 of *LNCS*, 552–567. Springer.
- Koopmann, P., and Schmidt, R. A. 2015. Uniform interpolation and forgetting for \mathcal{ALC} ontologies with ABoxes. In *Proceedings of AAAI 2015*, 175–181. AAAI Press.
- Koopmann, P.; Del-Pinto, W.; Tourret, S.; and Schmidt, R. A. 2020. Signature-based abduction for expressive description logics – technical report. *CoRR* arXiv:abs/2007.00757 [cs.AI].
- Koopmann, P. 2020. LETHE: Forgetting and uniform interpolation for expressive description logics. *KI - Künstliche Intelligenz*.
- Lehmann, J., and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Machine Learning* 78(1-2):203–250.
- Lin, F. 2001. On strongest necessary and weakest sufficient conditions. *Artificial Intelligence* 128(1-2):143–159.
- Lutz, C., and Wolter, F. 2011. Foundations for uniform interpolation and forgetting in expressive description logics. In *Proceedings of IJCAI 2011*, 989–995. IJCAI/AAAI.
- Nonnengart, A., and Szalas, A. 1999. A fixpoint approach to second-order quantifier elimination with applications to correspondence theory. In *Logic at Work*. Springer. 307–328.
- Parsia, B.; Matentzoglou, N.; Gonçalves, R. S.; Glimm, B.; and Steigmiller, A. 2017. The OWL Reasoner Evaluation (ORE) 2015 competition report. *Journal of Automated Reasoning* 59(4):455–482.
- Plaisted, D. A. 1994. The search efficiency of theorem proving strategies. In *Automated Deduction (CADE-12)*, volume 814 of *LNCS*, 57–71. Springer.
- Pukancová, J., and Homola, M. 2017. Tableau-based ABox abduction for the \mathcal{ALCHO} description logic. In *Proceedings of DL 2017*, volume 1879 of *CEUR Workshop Proceedings*. CEUR-WS.org.

- Pukancová, J., and Homola, M. 2018. ABox abduction for description logics: The case of multiple observations. In *Proceedings of DL 2018*, volume 2211 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Ray, O. 2009. Nonmonotonic abductive inductive learning. *Journal of Applied Logic* 7(3):329–340.
- Vescovo, C. D.; Parsia, B.; Sattler, U.; and Schneider, T. 2011. The modular structure of an ontology: Atomic decomposition. In *Proceedings of IJCAI 2011*, 2232–2237. IJCAI/AAAI.
- Wei-Kleiner, F.; Dragisic, Z.; and Lambrix, P. 2014. Abduction framework for repairing incomplete \mathcal{EL} ontologies: Complexity results and algorithms. In *Proceedings of AAAI 2014*, 1120–1127. AAAI Press.
- Zhao, Y., and Schmidt, R. A. 2017. Role forgetting for $\mathcal{ALCOQH}(\nabla)$ -ontologies using an Ackermann-based approach. In *Proceedings of IJCAI 2017*, 1354–1361. AAAI Press/IJCAI.