



# Consequence-Based Axiom Pinpointing

Ana Ozaki and Rafael Peñaloza<sup>(✉)</sup>

KRDB Research Centre, Free University of Bozen-Bolzano, Bolzano, Italy  
penaloza@inf.unibz.it

**Abstract.** Axiom pinpointing refers to the problem of finding the axioms in an ontology that are relevant for understanding a given entailment or consequence. One approach for axiom pinpointing, known as *glass-box*, is to modify a classical decision procedure for the entailments into a method that computes the solutions for the pinpointing problem. Recently, consequence-based decision procedures have been proposed as a promising alternative for tableaux-based reasoners for standard ontology languages. In this work, we present a general framework to extend consequence-based algorithms with axiom pinpointing.

**Keywords:** Consequence-based reasoning · Non-standard reasoning  
Consequence management

## 1 Introduction

Ontologies are now widely used in various domains such as medicine [22–24], biology [26], chemistry [10], geography [17, 18] and many others [29], to represent conceptual knowledge in a formal and easy to understand manner. It is a multi-task effort to construct and maintain such ontologies, often containing thousands of concepts. As these ontologies increase in size and complexity, it becomes more and more challenging for an ontology engineer to understand which parts of the ontology cause a certain consequence to be entailed. If, for example, this consequence is an error, the ontology engineer would want to understand its precise causes, and correct it with minimal disturbances to the rest of the ontology.

To support this task, a technique known as *axiom pinpointing* was introduced in [25]. The goal of axiom pinpointing is to identify the minimal sub-ontologies (w.r.t. set inclusion) that entail a given consequence; we call these sets *MinAs*. There are two basic approaches to axiom pinpointing. The *black-box approach* [20] uses repeated calls to an unmodified decision procedure to find these MinAs. The *glass-box approach* [4, 5], on the other hand, modifies the decision algorithm to generate the MinAs during one execution. In reality, glass-box methods do not explicitly compute the MinAs, but rather a compact representation of them known as the *pinpointing formula*. In this setting, each axiom of the ontology is labelled with a unique propositional symbol. The pinpointing formula is a (monotone) Boolean formula, satisfied exactly by those valuations

which evaluate to true the labels of the axioms in the ontology which cause the entailment of the consequence. Thus, the formula points out to the user the relevant parts of the ontology for the entailment of a certain consequence, where disjunction means alternative use of the axioms and conjunction means that the axioms are jointly used.

Axiom pinpointing can be used to enrich a decision procedure for entailment checking by further presenting to the user the axioms which cause a certain consequence. Since glass-box methods modify an existing decision procedure, they require a specification of the decision method to be studied. Previously, general methods for extending tableaux-based and automata-based decision procedures to axiom pinpointing have been studied in detail [4,5]. Classically, automata-based decision procedures often exhibit optimal worst-case complexity, but the most efficient reasoners for standard ontology languages are tableaux-based. When dealing with pinpointing extensions one observes a similar behaviour: the automata-based axiom pinpointing approach preserves the complexity of the original method, while tableau-based axiom pinpointing is not even guaranteed to terminate in general. However, the latter are more goal-directed and lead to a better run-time in practice.

A different kind of reasoning procedure that is gaining interest is known as the consequence-based method. In this setting, rules are applied to derive explicit consequences from previously derived knowledge. Consequence-based decision procedures often enjoy optimal worst-case complexity and, more recently, they have been presented as a promising alternative for tableaux-based reasoners for standard ontology languages [8,9,14,15,27,28,30]. Consequence-based algorithms have been previously described as simple variants of tableau algorithms [6], and as syntactic variants of automata-based methods [12]. They share the positive complexity bounds of automata, and the goal-directed nature of tableaux.

In this work, we present a general approach to produce axiom pinpointing extensions of consequence-based algorithms. Our driving example and use case is the extension of the consequence-based algorithm for entailment checking for the prototypical ontology language  $\mathcal{ALC}$  [15]. We show that the pinpointing extension does not change the ExpTime complexity of the consequence-based algorithm for  $\mathcal{ALC}$ .

## 2 Preliminaries

We briefly introduce the notions needed for this paper. We are interested in the problem of understanding the causes for a consequence to follow from an ontology. We consider an abstract notion of ontology and consequence relation. For the sake of clarity, however, we instantiate these notions to the description logic  $\mathcal{ALC}$ .

## 2.1 Axiom Pinpointing

To keep the discourse as general as possible, we consider an *ontology language* to define a class  $\mathfrak{A}$  of *axioms*. An *ontology* is then a finite set of axioms; that is, a finite subset of  $\mathfrak{A}$ . We denote the set of all ontologies as  $\mathfrak{O}$ . A *consequence property* (or *c-property* for short) is a binary relation  $\mathcal{P} \subseteq \mathfrak{O} \times \mathfrak{A}$  that relates ontologies to axioms. If  $(\mathcal{O}, \alpha) \in \mathcal{P}$ , we say that  $\alpha$  is a *consequence* of  $\mathcal{O}$  or alternatively, that  $\mathcal{O}$  *entails*  $\alpha$ .

We are only interested in relations that are monotonic in the sense that for any two ontologies  $\mathcal{O}, \mathcal{O}' \in \mathfrak{O}$  and axiom  $\alpha \in \mathfrak{A}$  such that  $\mathcal{O} \subseteq \mathcal{O}'$ , if  $(\mathcal{O}, \alpha) \in \mathcal{P}$  then  $(\mathcal{O}', \alpha) \in \mathcal{P}$ . In other words, adding more axioms to an ontology will only increase the set of axioms that are entailed from it. For the rest of this paper whenever we speak about a c-property, we implicitly assume that it is monotonic in this sense.

Notice that our notions of ontology and consequence property differ from previous work. In [4,5], c-properties are defined using two different types of statements and ontologies are allowed to require additional structural constraints. The former difference is just syntactic and does not change the generality of our approach. In the latter case, our setting becomes slightly less expressive, but at the benefit of simplifying the overall notation and explanation of our methods. Our results can be easily extended to the more general setting from [4,5].

When dealing with ontology languages, one is usually interested in deciding whether an ontology  $\mathcal{O}$  entails an axiom  $\alpha$ ; that is, whether  $(\mathcal{O}, \alpha) \in \mathcal{P}$ . In axiom pinpointing, we are more interested in the more detailed question of *why* it is a consequence. More precisely, we want to find the minimal (w.r.t. set inclusion) sub-ontologies  $\mathcal{O}' \subseteq \mathcal{O}$  such that  $(\mathcal{O}', \alpha) \in \mathcal{P}$  still holds. These subsets are known as *MinAs* [4,5], *justifications* [13], or *MUPS* [25]—among many other names—in the literature. Rather than enumerating all these sub-ontologies explicitly, one approach is to compute a formula, known as the pinpointing formula, that encodes them.

Formally, suppose that every axiom  $\alpha \in \mathfrak{A}$  is associated with a unique propositional variable  $\text{lab}(\alpha)$ , and let  $\text{lab}(\mathcal{O})$  be the set of all the propositional variables corresponding to axioms in the ontology  $\mathcal{O}$ . A *monotone Boolean formula*  $\phi$  over  $\text{lab}(\mathcal{O})$  is a Boolean formula using only variables in  $\text{lab}(\mathcal{O})$  and the connectives for conjunction ( $\wedge$ ) and disjunction ( $\vee$ ). The constants  $\top$  and  $\perp$ , always evaluated to true and false, respectively, are also monotone Boolean formulae. We identify a propositional valuation with the set of variables which are true in it. For a valuation  $\mathcal{V}$  and a set of axioms  $\mathcal{O}$ , the  $\mathcal{V}$ -*projection* of  $\mathcal{O}$  is the set  $\mathcal{O}_{\mathcal{V}} := \{\alpha \in \mathcal{O} \mid \text{lab}(\alpha) \in \mathcal{V}\}$ . Given a c-property  $\mathcal{P}$  and an axiom  $\alpha \in \mathfrak{A}$ , a monotone Boolean formula  $\phi$  over  $\text{lab}(\mathcal{O})$  is called a *pinpointing formula* for  $(\mathcal{O}, \alpha)$  w.r.t  $\mathcal{P}$  if for every valuation  $\mathcal{V} \subseteq \text{lab}(\mathcal{O})$ :

$$(\mathcal{O}_{\mathcal{V}}, \alpha) \in \mathcal{P} \text{ iff } \mathcal{V} \text{ satisfies } \phi.$$

## 2.2 Description Logics

Description logics (DLs) [3] are a family of knowledge representation formalisms that have been successfully applied to represent the knowledge of many application domains, in particular from the life sciences [29]. We briefly introduce, as a prototypical example,  $\mathcal{ALC}$ , which is the smallest propositionally closed description logic.

Given two disjoint sets  $N_C$  and  $N_R$  of *concept names* and *role names*, respectively,  $\mathcal{ALC}$  concepts are defined through the grammar rule:

$$C ::= A \mid \neg C \mid C \sqcap C \mid \exists r.C,$$

where  $A \in N_C$  and  $r \in N_R$ . A *general concept inclusion* (GCI) is an expression of the form  $C \sqsubseteq D$ , where  $C, D$  are  $\mathcal{ALC}$  concepts. A *TBox* is a finite set of GCIs.

The semantics of this logic is given in terms of *interpretations* which are pairs of the form  $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$  where  $\Delta^{\mathcal{I}}$  is a finite set called the *domain*, and  $\cdot^{\mathcal{I}}$  is the *interpretation function* that maps every concept name  $A \in N_C$  to a set  $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$  and every role name  $r \in N_R$  to a binary relation  $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ . The interpretation function is extended to arbitrary  $\mathcal{ALC}$  concepts inductively as shown in Fig. 1. Following this semantics, we introduce the usual abbreviations  $C \sqcup D := \neg(\neg C \sqcap \neg D)$ ,  $\forall r.C := \neg(\exists r.\neg C)$ ,  $\perp := A \sqcap \neg A$ , and  $\top := \neg \perp$ . That is,  $\top$  stands for a (DL) tautology, and  $\perp$  for a contradiction. The interpretation  $\mathcal{I}$  *satisfies* the GCI  $C \sqsubseteq D$  iff  $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ . It is a *model* of the TBox  $\mathcal{T}$  iff it satisfies all the GCIs in  $\mathcal{T}$ .

$$\begin{aligned} (\neg C)^{\mathcal{I}} &:= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (C \sqcap D)^{\mathcal{I}} &:= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (\exists r.C)^{\mathcal{I}} &:= \{d \in \Delta^{\mathcal{I}} \mid \exists e \in \Delta^{\mathcal{I}}. (d, e) \in r^{\mathcal{I}} \text{ and } e \in C^{\mathcal{I}}\} \end{aligned}$$

**Fig. 1.** Semantics of  $\mathcal{ALC}$

One of the main reasoning problems in DLs is to decide *subsumption* between two concepts  $C, D$  w.r.t. a TBox  $\mathcal{T}$ ; that is, to verify that every model of the TBox  $\mathcal{T}$  also satisfies the GCI  $C \sqsubseteq D$ . If this is the case, we denote it as  $\mathcal{T} \models C \sqsubseteq D$ . It is easy to see that the relation  $\models$  defines a c-property over the class  $\mathfrak{A}$  of axioms containing all possible GCIs; in this case, an ontology is a TBox.<sup>1</sup>

The following example instantiates the basic ideas presented in this section.

<sup>1</sup>  $\mathcal{ALC}$  ontologies often include also an *ABox* with facts about individuals. We disregard that part, as it is irrelevant for our example setting.

*Example 1.* Consider for example the  $\mathcal{ALC}$  TBox  $\mathcal{T}_{\text{exa}}$  containing the axioms

$$\begin{array}{ll} A \sqsubseteq \exists r.A : \mathbf{ax}_1, & \exists r.A \sqsubseteq B : \mathbf{ax}_2, \\ A \sqsubseteq \forall r.B : \mathbf{ax}_3, & A \sqcap B \sqsubseteq \perp : \mathbf{ax}_4, \end{array}$$

where  $\mathbf{ax}_i, 1 \leq i \leq 4$  are the propositional variables labelling the axiom. It is easy to see that  $\mathcal{T}_{\text{exa}} \models A \sqsubseteq \perp$ , and there are two justifications for this fact; namely, the TBoxes  $\{\mathbf{ax}_1, \mathbf{ax}_2, \mathbf{ax}_4\}$  and  $\{\mathbf{ax}_1, \mathbf{ax}_3, \mathbf{ax}_4\}$ . From this, it follows that  $\mathbf{ax}_1 \wedge \mathbf{ax}_4 \wedge (\mathbf{ax}_2 \vee \mathbf{ax}_3)$  is a pinpointing formula for  $A \sqsubseteq \perp$  w.r.t.  $\mathcal{T}_{\text{exa}}$ .

### 3 Consequence-Based Algorithms

Abstracting from particularities, a *consequence-based algorithm* works on a set  $\mathcal{A}$  of *consequences*, which is expanded through rule applications. Algorithms of this kind have two phases. The *normalization* phase first transforms all the axioms in an ontology into a suitable normal form. The *saturation* phase initializes the set  $\mathcal{A}$  of *derived consequences* with the normalized ontology and applies the rules to expand it. The set  $\mathcal{A}$  is often called a *state*. As mentioned, the initial state  $\mathcal{A}_0$  contains the normalization of the input ontology  $\mathcal{O}$ . A *rule* is of the form  $\mathcal{B}_0 \rightarrow \mathcal{B}_1$ , where  $\mathcal{B}_0, \mathcal{B}_1$  are finite sets of consequences. This rule is *applicable* to the state  $\mathcal{A}$  if  $\mathcal{B}_0 \subseteq \mathcal{A}$  and  $\mathcal{B}_1 \not\subseteq \mathcal{A}$ . Its *application* extends  $\mathcal{A}$  to  $\mathcal{A} \cup \mathcal{B}_1$ .  $\mathcal{A}$  is *saturated* if no rule is applicable to it. The method *terminates* if  $\mathcal{A}$  is saturated after finitely many rule applications, independently of the rule application order chosen. For the rest of this section and most of the following, we assume that the input ontology is already in this normal form, and focus only on the second phase.

Given a rule  $R = \mathcal{B}_0 \rightarrow \mathcal{B}_1$ , we use  $\text{pre}(R)$  and  $\text{res}(R)$  to denote the sets  $\mathcal{B}_0$  of premises that trigger  $R$  and  $\mathcal{B}_1$  of consequences resulting of its applicability, respectively. If the state  $\mathcal{A}'$  is obtained from  $\mathcal{A}$  through the application of the rule  $R$ , we write  $\mathcal{A} \rightarrow_R \mathcal{A}'$ , and denote  $\mathcal{A} \rightarrow \mathcal{A}'$  if the precise rule used is not relevant. As usual,  $\rightarrow^*$  denotes the transitive and reflexive closure of  $\rightarrow$ .

Consequence-based algorithms derive, in a single execution, several axioms that are entailed from the input ontology. Obviously, in general they cannot generate *all* possible entailed axioms, as such a set may be infinite (e.g., in the case of  $\mathcal{ALC}$ ). Thus, to define correctness, we need to specify for every ontology  $\mathcal{O}$ , a finite set  $\delta(\mathcal{O})$  of *derivable consequences* of  $\mathcal{O}$ . This set is assumed to be provided as part of the consequence-based algorithm.

**Definition 2 (Correctness).** *A consequence-based algorithm is correct for the consequence property  $\mathcal{P}$  if for every ontology  $\mathcal{O}$ , the following two conditions hold: (i) it terminates, and (ii) if  $\mathcal{O} \rightarrow^* \mathcal{A}$  and  $\mathcal{A}$  is saturated, then for every derivable consequence  $\alpha \in \delta(\mathcal{O})$  it follows that  $(\mathcal{O}, \alpha) \in \mathcal{P}$  iff  $\alpha \in \mathcal{A}$ .*

That is, the algorithm is correct for a property if it terminates and is sound and complete w.r.t. the finite set of derivable consequences  $\delta(\mathcal{O})$ .

Notice that the definition of correctness requires that the resulting set of consequences obtained from the application of the rules is always the same, independently of the order in which the rules are applied. In other words, if  $\mathcal{O} \rightarrow^* \mathcal{A}$ ,  $\mathcal{O} \rightarrow^* \mathcal{A}'$ , and  $\mathcal{A}, \mathcal{A}'$  are both saturated, then  $\mathcal{A} = \mathcal{A}'$ . This is a fundamental property that will be helpful for showing correctness of the pinpointing extensions in the next section.

**Table 1.**  $\mathcal{ALC}$  consequence-based algorithm rules  $\mathcal{B}_0 \rightarrow \mathcal{B}_1$ .

$\mathcal{B}_0$	$\mathcal{B}_1$
1: $\emptyset$	$H \sqcap A \sqsubseteq A$
2: $H \sqcap \neg A \sqsubseteq N \sqcup A$	$H \sqcap \neg A \sqsubseteq N$
3: $H \sqsubseteq N_1 \sqcup A_1, \dots, H \sqsubseteq N_n \sqcup A_n$ $A_1 \sqcap \dots \sqcap A_n \sqsubseteq N$	$H \sqsubseteq \bigsqcup_{i=1}^n N_i \sqcup N$
4: $H \sqsubseteq N \sqcup A$ $A \sqsubseteq \exists r.B$	$H \sqsubseteq N \sqcup \exists r.B$
5: $H \sqsubseteq M \sqcup \exists r.K, K \sqsubseteq N \sqcup A$ $\exists r.A \sqsubseteq B$	$H \sqsubseteq M \sqcup B \sqcup \exists r.(K \sqcap \neg A)$
6: $H \sqsubseteq M \sqcup \exists r.K, K \sqsubseteq \perp$	$H \sqsubseteq M$
7: $H \sqsubseteq M \sqcup \exists r.K, H \sqsubseteq N \sqcup A$ $A \sqsubseteq \forall r.B$	$H \sqsubseteq M \sqcup N \sqcup \exists r.(K \sqcap B)$

A well-known example of a consequence-based algorithm is the  $\mathcal{ALC}$  reasoning method from [27]. To describe this algorithm we need some notation. A *literal* is either a concept name or a negated concept name. Let  $H, K$  denote (possibly empty) conjunctions of literals, and  $M, N$  are (possibly empty) disjunctions of concept names. For simplicity, we treat these conjunctions and disjunctions as sets. The normalization phase transforms all GCIs to be of the form:

$$\prod_{i=1}^n A_i \sqsubseteq \prod_{j=1}^m B_j, \quad A \sqsubseteq \exists r.B, \quad A \sqsubseteq \forall r.B, \quad \exists r.A \sqsubseteq B.$$


For a given  $\mathcal{ALC}$  TBox  $\mathcal{T}$ , the set  $\delta(\mathcal{T})$  of derivable consequences contains all GCIs of the form  $H \sqsubseteq M$  and  $H \sqsubseteq N \sqcup \exists r.K$ . The saturation phase initializes  $\mathcal{A}$  to contain the axioms in the (normalized) TBox, and applies the rules from Table 1 until a saturated state is found. After termination, one can check that for every derivable consequence  $C \sqsubseteq D$  it holds that  $\mathcal{T} \models C \sqsubseteq D$  iff  $C \sqsubseteq D \in \mathcal{A}$ ; that is, this algorithm is correct for the property [27].

*Example 3.* Recall the  $\mathcal{ALC}$  TBox  $\mathcal{T}_{\text{exa}}$  from Example 1. Notice that all axioms in this TBox are already in normal form; hence the normalization step does not modify it. The consequence-based algorithm starts with  $\mathcal{A} := \mathcal{T}_{\text{exa}}$  and applies the rules until saturation. One possible execution of the algorithm is

$$\mathcal{A}_0 \rightarrow_5 A \sqsubseteq B \rightarrow_3 A \sqsubseteq \perp \rightarrow_7 A \sqsubseteq \exists r.(A \sqcap B) \rightarrow^* \dots,$$

where  $\mathcal{A}_0$  contains  $\mathcal{T}_{\text{exa}}$  and the result of adding all the tautologies generated by the application of Rule 1 over it (see Fig. 2). Since rule applications only extend

the set of consequences, we depict exclusively the newly added consequence; e.g., the first rule application  $\mathcal{A}_0 \rightarrow_1 A \sqsubseteq A$  is in fact representing  $\mathcal{A}_0 \rightarrow_1 \mathcal{A}_0 \cup \{A \sqsubseteq A\}$ . When the execution of the method terminates, the set of consequences  $\mathcal{A}$  contains  $A \sqsubseteq \perp$ ; hence we can conclude that this subsumption follows from  $\mathcal{T}_{\text{exa}}$ . Notice that other consequences (e.g.,  $A \sqsubseteq B$ ) are also derived from the same execution.



**Fig. 2.** An execution of the  $\mathcal{ALC}$  consequence-based algorithm over  $\mathcal{T}_{\text{exa}}$  from Example 3. Arrows point from the premises to the consequences generated by the application of the rule denoted in the subindex.

For the rest of this paper, we consider an arbitrary, but fixed, consequence-based algorithm, that is correct for a given c-property  $\mathcal{P}$ .

## 4 The Pinpointing Extension

Our goal is to extend consequence-based algorithms from the previous section to methods that compute pinpointing formulae for their consequences. We achieve this by modifying the notion of states, and the rule applications on them. Recall that every axiom  $\alpha$  in the class  $\mathfrak{A}$  (in hence, also every axiom in the ontology  $\mathcal{O}$ ) is labelled with a unique propositional variable  $\text{lab}(\alpha)$ . In a similar manner, we consider sets of consequences  $\mathcal{A}$  that are labelled with a monotone Boolean formula. We use the notation  $\mathcal{A}^{\text{pin}}$  to indicate that the elements in the set  $\mathcal{A}$  are labelled in this way, and use  $\alpha : \varphi_\alpha \in \mathcal{A}^{\text{pin}}$  to express that the consequence  $\alpha$ , labelled with the formula  $\varphi_\alpha$ , belongs to  $\mathcal{A}^{\text{pin}}$ . A *pinpointing state* is a set of labelled consequences. We assume that each consequence in this set is labelled with only one formula. For a set of labelled consequences  $\mathcal{A}^{\text{pin}}$  and a set of (unlabelled) consequences  $X$ , we define  $\text{fm}(X, \mathcal{A}^{\text{pin}}) := \bigwedge_{\alpha \in X} \varphi_\alpha$ , where  $\varphi_\alpha = \perp$  if  $\alpha \notin \mathcal{A}$ .

A consequence-based algorithm  $A$  induces a pinpointing consequence-based algorithm  $A^{pin}$  by modifying the notion of rule application, and dealing with pinpointing states, instead of classical states, through a modification of the formulae labelling the derived consequences.

**Definition 4 (Pinpointing Application).** *The rule  $R = \mathcal{B}_0 \rightarrow \mathcal{B}_1$  is pinpointing applicable to the pinpointing state  $\mathcal{A}^{pin}$  if  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) \not\equiv \text{fm}(\mathcal{B}_1, \mathcal{A}^{pin})$ . The pinpointing application of this rule modifies  $\mathcal{A}^{pin}$  to:*

$$\{\alpha : \varphi_\alpha \vee \text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) \mid \alpha \in \mathcal{B}_1, \alpha : \varphi_\alpha \in \mathcal{A}^{pin}\} \cup (\mathcal{A}^{pin} \setminus \{\alpha : \varphi_\alpha \mid \alpha \in \mathcal{B}_1\}).$$

*The pinpointing state  $\mathcal{A}^{pin}$  is pinpointing saturated if no rule is pinpointing applicable to it.*

*Example 5.* Consider again the TBox  $\mathcal{T}_{\text{exa}}$  from Example 1. At the beginning of the execution of the pinpointing algorithm, the set of consequences is the TBox, with each axiom labelled by the unique propositional variable representing it; that is  $\mathcal{T}^{pin} = \{A \sqsubseteq \exists r.A : \text{ax}_1, \exists r.A \sqsubseteq B : \text{ax}_2, A \sqsubseteq \forall r.B : \text{ax}_3, A \sqcap B \sqsubseteq \perp : \text{ax}_4\}$ . A pinpointing application of Rule 1 adds the new consequence  $A \sqsubseteq A : \top$ , where the tautology  $\top$  labelling this consequence arises from the fact that rule 1 has no premises. At this point, one can pinpointing apply Rule 5 with

$$\mathcal{B}_0 = \{A \sqsubseteq \exists r.A, A \sqsubseteq A, \exists r.A \sqsubseteq B\}, \quad \mathcal{B}_1 = \{A \sqsubseteq B\}$$

(see the solid arrow in Fig. 2). In this case,  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) = \text{ax}_1 \wedge \top \wedge \text{ax}_2$  and  $\text{fm}(\mathcal{B}_1, \mathcal{A}^{pin}) = \perp$  because the consequence  $A \sqsubseteq B$  does not belong to  $\mathcal{A}^{pin}$  yet. Hence  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) \not\equiv \text{fm}(\mathcal{B}_1, \mathcal{A}^{pin})$ , and the rule is indeed pinpointing applicable. The pinpointing application of this rule adds the new labelled consequence  $A \sqsubseteq B : \text{ax}_1 \wedge \text{ax}_2$  to  $\mathcal{A}^{pin}$ . Then, Rule 3 becomes pinpointing applicable with

$$\mathcal{B}_0 = \{A \sqsubseteq B, A \sqsubseteq A, A \sqcap B \sqsubseteq \perp\},$$

which adds  $A \sqsubseteq \perp : \text{ax}_1 \wedge \text{ax}_2 \wedge \text{ax}_4$  to the set of consequences. Then, Rule 7 over the set of premises

$$\mathcal{B}_0 = \{A \sqsubseteq \exists r.A, A \sqsubseteq A, A \sqsubseteq \forall r.B\},$$

yields the new consequence  $A \sqsubseteq \exists r.(A \sqcap B) : \text{ax}_1 \wedge \text{ax}_3$ .

Notice that, at this point Rule 6 is not applicable in the classical case over the set of premises  $\mathcal{B}_0 = \{A \sqsubseteq \exists r.(A \sqcap B), A \sqcap B \sqsubseteq \perp\}$  because its (regular) application would add the consequence  $A \sqsubseteq \perp$  that was already derived. However,


$$\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) = \text{ax}_1 \wedge \text{ax}_3 \wedge \text{ax}_4 \not\equiv \text{ax}_1 \wedge \text{ax}_2 \wedge \text{ax}_4 = \varphi_{A \sqsubseteq \perp} = \text{fm}(\mathcal{B}_1, \mathcal{A}^{pin});$$

hence, the rule is in fact pinpointing applicable. The pinpointing application of this Rule 6 substitutes the labelled consequence  $A \sqsubseteq \perp : \text{ax}_1 \wedge \text{ax}_2 \wedge \text{ax}_4$  with the consequence  $A \sqsubseteq \perp : (\text{ax}_1 \wedge \text{ax}_2 \wedge \text{ax}_4) \vee (\text{ax}_1 \wedge \text{ax}_3 \wedge \text{ax}_4)$ . The pinpointing



extension will then continue applying rules until a saturated state is reached. This execution is summarized in Fig. 3. At that point, the set of labelled consequences will contain, among others,  $A \sqsubseteq \perp : (\mathbf{ax}_1 \wedge \mathbf{ax}_2 \wedge \mathbf{ax}_4) \vee (\mathbf{ax}_1 \wedge \mathbf{ax}_3 \wedge \mathbf{ax}_4)$ . The label of this consequence corresponds to the pinpointing formula that was computed in Example 1.

We denote as  $\mathcal{A}^{pin} \rightarrow_R \mathcal{B}^{pin}$  the fact that  $\mathcal{B}^{pin}$  is obtained from the pinpointing application of the rule  $R$  to  $\mathcal{A}^{pin}$ . As before, we drop the subscript  $R$  if the name of the rule is irrelevant and write simply  $\mathcal{A}^{pin} \rightarrow \mathcal{B}^{pin}$ . The pinpointing extension starts, as the classical one, with the set of all normalized axioms. For the rest of this section, we assume that the input ontology is already normalized, and hence each axiom in the initial pinpointing state is labelled with its corresponding propositional variable. In the next section we show how to deal with normalization.



**Fig. 3.** Pinpointing application of rules over  $\mathcal{T}_{\text{exa}}$  in Example 5. Arrows point from the premises to the consequences generated by the pinpointing application of the rule denoted in the subindex.

Notice that if a rule  $R$  is applicable to some state  $\mathcal{A}$ , then it is also pinpointing applicable to it. This holds because the regular applicability condition requires that at least one consequence  $\alpha$  in  $\text{res}(R)$  should not exist already in the state  $\mathcal{A}$ , which is equivalent to having the consequence  $\alpha : \perp \in \mathcal{A}^{pin}$ . Indeed, we used this fact in the first pinpointing rule applications of Example 5. If the consequence-based algorithm is correct, then it follows by definition that for any saturated state  $\mathcal{A}$  obtained by a sequence of rule applications from  $\mathcal{O}$ ,  $\mathcal{O} \models \alpha$  iff  $\alpha \in \mathcal{A}$ . Conversely, as shown next, every consequence created by a pinpointing rule application is also generated by a regular rule application. First, we extend the notion of a  $\mathcal{V}$ -projection to sets of consequences (i.e., states) in the obvious manner:  $\mathcal{A}_{\mathcal{V}} := \{\alpha \mid \alpha : \varphi_{\alpha} \in \mathcal{A}^{pin}, \mathcal{V} \models \varphi\}$ .

**Lemma 6.** *Let  $\mathcal{A}^{pin}, \mathcal{B}^{pin}$  be pinpointing states and let  $\mathcal{V}$  be a valuation. If  $\mathcal{A}^{pin} \rightarrow^* \mathcal{B}^{pin}$  then  $\mathcal{A}_{\mathcal{V}} \rightarrow^* \mathcal{B}_{\mathcal{V}}$ .*

*Proof.* We show that if  $\mathcal{A}^{pin} \rightarrow_R \mathcal{B}^{pin}$  then  $\mathcal{A}_\mathcal{V} \rightarrow_R \mathcal{B}_\mathcal{V}$  or  $\mathcal{A}_\mathcal{V} = \mathcal{B}_\mathcal{V}$ , where  $R = \mathcal{B}_0 \rightarrow \mathcal{B}_1$  is a rule. If  $\mathcal{V}$  does not satisfy  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin})$  then  $\mathcal{A}_\mathcal{V} = \mathcal{B}_\mathcal{V}$  since the labels of the newly added assertions are not satisfied by  $\mathcal{V}$ , and the disjunction with  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin})$  does not change the evaluation of the modified labels under  $\mathcal{V}$ . On the other hand, if  $\mathcal{V}$  satisfies  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin})$  then  $\mathcal{B}_0 \subseteq \mathcal{A}_\mathcal{V}$ . If  $\mathcal{B}_1 \not\subseteq \mathcal{A}_\mathcal{V}$  then  $\mathcal{A}_\mathcal{V} \rightarrow_R \mathcal{B}_\mathcal{V}$ . Otherwise, again we have  $\mathcal{A}_\mathcal{V} = \mathcal{B}_\mathcal{V}$ .  $\square$

Since all the labels are monotone Boolean formulae, it follows that the valuation  $\mathcal{V}_\top = \text{lab}(\mathcal{O})$  that makes every propositional variable true satisfies all labels, and hence for every pinpointing state  $\mathcal{A}^{pin}$ ,  $\mathcal{A}_{\mathcal{V}_\top} = \mathcal{A}$ . Lemma 6 hence entails that the pinpointing extension of the consequence-based algorithm  $A$  does not create new consequences, but only labels these consequences. Termination of the pinpointing extension then follows from the termination of the consequence-based algorithm and the condition for pinpointing rule application that entails that, whenever a rule is pinpointing applied, the set of labelled consequences is necessarily modified either by adding a new consequence, or by modifying the label of at least one existing consequence to a weaker (i.e., more general) monotone Boolean formula. Since there are only finitely many monotone Boolean formulas over  $\text{lab}(\mathcal{O})$ , every label can be changed finitely many times only.

It is in fact possible to get a better understanding of the running time of the pinpointing extension of a consequence-based algorithm. Suppose that, on input  $\mathcal{O}$ , the consequence-based algorithm  $A$  stops after at most  $f(\mathcal{O})$  rule applications. Since every rule application must add at least one consequence to the state, the saturated state reached by this algorithm will have at most  $f(\mathcal{O})$  consequences. Consider now the pinpointing extension of  $A$ . We know, from the previous discussion, that this pinpointing extension generates the same set of consequences. Moreover, since there are  $2^{|\mathcal{O}|}$  possible valuations over  $\text{lab}(\mathcal{O})$ , and every pinpointing rule application that does not add a new consequence must generalize at least one formula, the labels of each consequence can be modified at most  $2^{|\mathcal{O}|}$  times. Overall, this means that the pinpointing extension of  $A$  stops after at most  $2^{|\mathcal{O}|}f(\mathcal{O})$  rule applications. We now formalize this result.

**Theorem 7.** *If a consequence-based algorithm  $A$  stops after at most  $f(\mathcal{O})$  rule applications, then  $\mathcal{A}^{pin}$  stops after at most  $2^{|\mathcal{O}|}f(\mathcal{O})$  rule applications.*

Another important property of the pinpointing extension is that saturatedness of a state is preserved under projections.

**Lemma 8.** *Let  $\mathcal{A}^{pin}$  be a pinpointing state and  $\mathcal{V}$  a valuation. If  $\mathcal{A}^{pin}$  is pinpointing saturated then  $\mathcal{A}_\mathcal{V}$  is saturated.*

*Proof.* Suppose there is a rule  $R$  such that  $R$  is applicable to  $\mathcal{A}_\mathcal{V}$ . This means that  $\mathcal{B}_0 \subseteq \mathcal{A}_\mathcal{V}$  and  $\mathcal{B}_1 \not\subseteq \mathcal{A}_\mathcal{V}$ . We show that  $R$  is pinpointing applicable to  $\mathcal{A}^{pin}$ . Since  $\mathcal{B}_0 \subseteq \mathcal{A}_\mathcal{V}$ ,  $\mathcal{V}$  satisfies  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin})$ . As  $\mathcal{B}_1 \not\subseteq \mathcal{A}_\mathcal{V}$ , there is  $\alpha \in \mathcal{B}_1$  such that either  $\alpha \notin \mathcal{A}$  or  $\alpha : \varphi_\alpha \in \mathcal{A}^{pin}$  but  $\mathcal{V}$  does not satisfy  $\varphi_\alpha$ . In the former case,  $R$  is clearly pinpointing applicable to  $\mathcal{A}^{pin}$ . In the latter,  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin}) \not\models \text{fm}(\mathcal{B}_1, \mathcal{A}^{pin})$  since  $\mathcal{V}$  satisfies  $\text{fm}(\mathcal{B}_0, \mathcal{A}^{pin})$  but not  $\text{fm}(\mathcal{B}_1, \mathcal{A}^{pin})$ .  $\square$

We can now show that the pinpointing extension of a consequence-based algorithm is indeed a pinpointing algorithm; that is, that when a saturated pinpointing state  $\mathcal{A}^{pin}$  is reached from rule applications starting from  $\mathcal{O}^{pin}$ , then for every  $\alpha : \varphi_\alpha \in \mathcal{A}^{pin}$ ,  $\varphi_\alpha$  is a pinpointing formula for  $\alpha$  w.r.t.  $\mathcal{O}^{pin}$ .

**Theorem 9 (Correctness of Pinpointing).** *Let  $\mathcal{P}$  be a c-property on axiomatized inputs for  $\mathfrak{J}$  and  $\mathfrak{A}$ . Given a correct consequence-based algorithm  $A$  for  $\mathcal{P}$ , for every axiomatized input  $\Gamma = (\mathcal{O}, \alpha) \in \mathcal{P}$ , where  $\mathcal{O}$  is normalized, then*

*if  $\mathcal{O}^{pin} \xrightarrow{*} \mathcal{A}^{pin}$ ,  $\alpha : \varphi_\alpha \in \mathcal{A}^{pin}$ , and  $\mathcal{A}^{pin}$  is pinpointing saturated, then  $\varphi_\alpha$  is a pinpointing formula for  $\Gamma$  w.r.t.  $\mathcal{P}$ .*

*Proof.* We want to show that  $\varphi_\alpha$  is a pinpointing formula for  $\mathcal{P}$  and  $(\alpha, \mathcal{O})$ . That is, for every valuation  $\mathcal{V} \subseteq \text{lab}(\mathcal{O})$ :  $(\alpha, \mathcal{O}_\mathcal{V}) \in \mathcal{P}$  iff  $\mathcal{V}$  satisfies  $\varphi_\alpha$ .

Assume that  $(\alpha, \mathcal{O}_\mathcal{V}) \in \mathcal{P}$ , i.e.,  $\mathcal{O}_\mathcal{V} \models \alpha$ , and let  $\mathcal{A}_0 = \mathcal{O}_\mathcal{V}$ . Since  $A$  terminates on every input, there is a saturated state  $\mathcal{B}$  such that  $\mathcal{A}_0 \xrightarrow{*} \mathcal{B}$ . Completeness of  $A$  then implies that  $\alpha \in \mathcal{B}$ . By assumption,  $\mathcal{A}_0^{pin} \xrightarrow{*} \mathcal{A}^{pin}$  and  $\mathcal{A}^{pin}$  is pinpointing saturated. By Lemma 6 it follows that  $\mathcal{O}_\mathcal{V} \xrightarrow{*} \mathcal{A}_\mathcal{V}$ , and by Lemma 8,  $\mathcal{A}_\mathcal{V}$  is saturated. Hence, since  $A$  is correct,  $\mathcal{A}_\mathcal{V} = \mathcal{B}$ . This implies that  $\mathcal{V} \models \varphi_\alpha$  because  $\alpha \in \mathcal{A}_\mathcal{V}$ .

Conversely, suppose that  $\mathcal{V}$  satisfies  $\varphi_\alpha$ . By assumption,  $\alpha : \varphi_\alpha \in \mathcal{A}^{pin}$ ,  $\mathcal{O}^{pin} \xrightarrow{*} \mathcal{A}^{pin}$ , and  $\mathcal{A}^{pin}$  is saturated. By Lemma 6,  $\mathcal{O}_\mathcal{V} \xrightarrow{*} \mathcal{A}_\mathcal{V}$ . Since  $\mathcal{V}$  satisfies  $\varphi_\alpha$ ,  $\alpha \in \mathcal{A}_\mathcal{V}$ . Then, by soundness of  $A$ ,  $\mathcal{O}_\mathcal{V} \models \alpha$ .  $\square$

As it was the case for classical consequence-based algorithms, their pinpointing extensions can apply the rules in any desired order. The notion of correctness of consequence-based algorithms guarantees that a saturated state will always be found, and the result will be the same, regardless of the order in which the rules are applied. We have previously seen that termination transfers also the pinpointing extensions. Theorem 9 also shows that the formula associated to the consequences derived is always equivalent.

**Corollary 10.** *Let  $\mathcal{A}^{pin}, \mathcal{B}^{pin}$  two pinpointing saturated states,  $\mathcal{O}$  an ontology, and  $\alpha$  a consequence such that  $\alpha : \varphi_\alpha \in \mathcal{A}^{pin}$  and  $\alpha : \psi_\alpha \in \mathcal{B}^{pin}$ . If  $\mathcal{O}^{pin} \xrightarrow{*} \mathcal{A}^{pin}$  and  $\mathcal{O}^{pin} \xrightarrow{*} \mathcal{B}^{pin}$ , then  $\varphi_\alpha \equiv \psi_\alpha$ .*

To finalize this section, we consider again our running example of deciding subsumption in  $\mathcal{ALC}$  described in Sect. 3. It terminates after an exponential number of rule applications on the size of the input TBox  $\mathcal{T}$ . Notice that every pinpointing rule application requires an entailment test between two monotone Boolean formulas, which can be decided in non-deterministic polynomial time on  $|\mathcal{O}|$ . Thus, it follows from Theorem 7 that the pinpointing extension of the consequence-based algorithm for  $\mathcal{ALC}$  runs in exponential time.

**Corollary 11.** *Let  $\mathcal{T}$  be an  $\mathcal{ALC}$  TBox, and  $C, D$  two  $\mathcal{ALC}$  concepts. A pinpointing formula for  $C \sqsubseteq D$  w.r.t.  $\mathcal{T}$  is computable in exponential time.*

## 5 Dealing with Normalization

Throughout the last two sections, we have disregarded the first phase of the consequence-based algorithms in which the axioms in the input ontology are transformed into a suitable normal form. In a nutshell, the normalization phase takes every axiom in the ontology and substitutes it by a set of simpler axioms that are, in combination, equivalent to the original one w.r.t. the set of derivable consequences. For example, in  $\mathcal{ALC}$  the axiom  $A \sqsubseteq B \sqcap C$  is *not* in normal form. During the normalization phase, it would then be substituted by the two axioms  $A \sqsubseteq B$ ,  $A \sqsubseteq C$ , which in combination provide the exact same constraints as the original axiom.

Obviously, in the context of pinpointing, we are interested in finding the set of *original* axioms that cause the consequence of interest, and not those in normal form; in fact, normalization is an internal process of the algorithm, and the user should be agnostic to the internal structures used. Hence, we need to find a way to track the original axioms.

To solve this, we slightly modify the initialization of the pinpointing extension. Recall from the previous section that, if the input ontology is already in normal form, then we initialize the algorithm with the state that contains exactly that ontology, where every axiom is labelled with the unique propositional variable that represents it. If the ontology is not originally in normal form, then it is first normalized. In this case, we set as the initial state the newly normalized ontology, but every axiom is labelled with the disjunction of the variables representing the axioms that generated it. This disjunction is used to represent the fact that the normalized axiom may be generated in more than one way. The following example explains this idea.

*Example 12.* Consider a variant  $\mathcal{T}'_{\text{exa}}$  of the  $\mathcal{ALC}$  TBox from Example 1 that is now formed by the three axioms

$$\begin{array}{ll} A \sqsubseteq \exists r.A \sqcap \forall r.B : \mathbf{ax}'_1, & \exists r.A \sqsubseteq B : \mathbf{ax}'_2, \\ A \sqsubseteq \forall r.B : \mathbf{ax}'_3 & A \sqcap B \sqsubseteq \perp : \mathbf{ax}'_4. \end{array}$$

Obviously, the first axiom  $\mathbf{ax}'_1$  is not in normal form, but can be normalized by substituting it with the two axioms  $A \sqsubseteq \exists r.A$ ,  $A \sqsubseteq \forall r.B$ . Thus, the normalization step yields the same TBox  $\mathcal{T}_{\text{exa}}$  from Example 1. However, instead of using different propositional variables to label these two axioms, they just inherit the label from the axiom that generated them; in this case  $\mathbf{ax}'_1$ . However, the axiom  $A \sqsubseteq \forall r.B$  is also caused by axiom  $\mathbf{ax}'_3$ . Thus, the pinpointing algorithm is initialized with

$$\mathcal{A}^{pin} = \{A \sqsubseteq \exists r.A : \mathbf{ax}'_1, \exists r.A \sqsubseteq B : \mathbf{ax}'_2, A \sqsubseteq \forall r.B : \mathbf{ax}'_1 \vee \mathbf{ax}'_3, A \sqcap B \sqsubseteq \perp : \mathbf{ax}'_4\}.$$

Following the same process as in Example 5, we see that we can derive the consequence  $A \sqsubseteq \perp : \mathbf{ax}'_1 \wedge \mathbf{ax}'_4$ . Hence  $\mathbf{ax}'_1 \wedge \mathbf{ax}'_4$  is a pinpointing formula for  $A \sqsubseteq \perp$  w.r.t.  $\mathcal{T}'_{\text{exa}}$ . It can be easily verified that this is in fact the case.

Thus, the normalization phase does not affect the correctness, nor the complexity of the pinpointing extension of a consequence-based algorithm.

## 6 Conclusions

We presented a general framework to extend consequence-based algorithms with axiom pinpointing. These algorithms often enjoy optimal upper bound complexity and can be efficiently implemented in practice. Our focus in this paper and use case is for the prototypical ontology language  $\mathcal{ALC}$ . We emphasize that this is only one of many consequence-based algorithms available. The completion-based algorithm for  $\mathcal{EL}^+$  [2] is obtained by restricting the assertions to be of the form  $A \sqsubseteq B$  and  $A \sqsubseteq \exists r.B$  with  $A, B \in N_C \cup \{\top\}$  and  $r \in N_R$ , and adding one rule to handle role constructors. Other examples of consequence-based methods include LTUR approach for Horn clauses [19], and methods for more expressive and Horn DLs [8, 14, 16].

Understanding the axiomatic causes for a consequence, and in particular the pinpointing formula, has importance beyond MinA enumeration. For example, the pinpointing formula also encodes all the ways to *repair* an ontology [1]. Depending on the application in hand, a simpler version of the formula can be computed, potentially more efficiently. This idea has already been employed to find good approximations for MinAs [7] and lean kernels [21] efficiently.

As future work, it would be interesting to investigate how algorithms for query answering in an ontology-based data access setting can be extended with the pinpointing technique. The pinpointing formula in this case could also be seen as a provenance polynomial, as introduced by Green et al. [11], in database theory. Another direction is to investigate axiom pinpointing in decision procedures for non-monotonic reasoning, where one would also expect the presence of negations in the pinpointing formula.

## References

1. Arif, M.F., Mencia, C., Marques-Silva, J.: Efficient axiom pinpointing with EL2MCS. In: KI, pp. 225–233 (2015)
2. Baader, F., Brandt, S., Lutz, C.: Pushing the  $\mathcal{EL}$  envelope. In: IJCAI, pp. 364–369 (2005)
3. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, Cambridge (2003)
4. Baader, F., Peñaloza, R.: Automata-based axiom pinpointing. J. Autom. Reason. **45**(2), 91–129 (2010)
5. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. J. Log. Comput. **20**(1), 5–34 (2010)
6. Baader, F., Peñaloza, R., Suntisrivaraporn, B.: Pinpointing in the description logic  $\mathcal{EL}^+$ . In: Hertzberg, J., Beetz, M., Englert, R. (eds.) KI 2007. LNCS (LNAI), vol. 4667, pp. 52–67. Springer, Heidelberg (2007). [https://doi.org/10.1007/978-3-540-74565-5\\_7](https://doi.org/10.1007/978-3-540-74565-5_7)
7. Baader, F., Suntisrivaraporn, B.: Debugging SNOMED CT using axiom pinpointing in the description logic  $\mathcal{EL}^+$ . In: Proceedings of the 3rd Knowledge Representation in Medicine (KR-MED 2008): Representing and Sharing Knowledge Using SNOMED. CEUR-WS, vol. 410 (2008)

8. Bate, A., Motik, B., Grau, B.C., Simancik, F., Horrocks, I.: Extending consequence-based reasoning to SRIQ. In: KR, pp. 187–196 (2016)
9. Cucala, D.T., Grau, B.C., Horrocks, I.: Consequence-based reasoning for description logics with disjunction, inverse roles, and nominals. In: Proceedings of the 30th International Workshop on Description Logics, Montpellier, France, 18–21 July 2017 (2017)
10. Degtyarenko, K., et al.: ChEBI: a database and ontology for chemical entities of biological interest. *Nucleic Acids Res.* **36**(suppl 1), D344–D350 (2008)
11. Green, T.J., Karvounarakis, G., Tannen, V.: Provenance semirings. In: Proceedings of the Twenty-Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Beijing, China, 11–13 June 2007, pp. 31–40 (2007)
12. Hutschenreiter, L., Peñaloza, R.: An automata view to goal-directed methods. In: Drewes, F., Martín-Vide, C., Truthe, B. (eds.) *LATA 2017*. LNCS, vol. 10168, pp. 103–114. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-53733-7\\_7](https://doi.org/10.1007/978-3-319-53733-7_7)
13. Kalyanpur, A., Parsia, B., Horridge, M., Sirin, E.: Finding all justifications of OWL DL entailments. In: Aberer, K. (ed.) *ASWC/ISWC -2007*. LNCS, vol. 4825, pp. 267–280. Springer, Heidelberg (2007). [https://doi.org/10.1007/978-3-540-76298-0\\_20](https://doi.org/10.1007/978-3-540-76298-0_20)
14. Kazakov, Y.: Consequence-driven reasoning for Horn SHIQ ontologies. In: Boutilier, C. (ed.) *IJCAI 2009*, pp. 2040–2045 (2009)
15. Kazakov, Y., Klinov, P.: Bridging the gap between tableau and consequence-based reasoning. In: *Informal Proceedings of the 27th International Workshop on Description Logics*, Vienna, Austria, 17–20 July 2014, pp. 579–590 (2014)
16. Kazakov, Y., Krötzsch, M., Simancik, F.: The incredible ELK - from polynomial procedures to efficient reasoning with  $\mathcal{EL}$  ontologies. *JAR* **53**(1), 1–61 (2014)
17. Kuipers, B.J.: An ontological hierarchy for spatial knowledge. In: Proceedings of the 10th International Workshop on Qualitative Reasoning About Physical Systems, Fallen Leaf Lake, California, USA (1996)
18. McMaster, R.B., User, E.L.: *A Research Agenda for Geographic Information Science*, vol. 3. CRC Press, Boca Raton (2004)
19. Minoux, M.: LTUR: a simplified linear-time unit resolution algorithm for Horn formulae and computer implementation. *Inf. Process. Lett.* **29**(1), 1–12 (1988)
20. Parsia, B., Sirin, E., Kalyanpur, A.: Debugging OWL ontologies. In: WWW, pp. 633–640 (2005)
21. Peñaloza, R., Mencia, C., Ignatiev, A., Marques-Silva, J.: Lean kernels in description logics. In: Blomqvist, E., Maynard, D., Gangemi, A., Hoekstra, R., Hitzler, P., Hartig, O. (eds.) *ESWC 2017*. LNCS, vol. 10249, pp. 518–533. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-58068-5\\_32](https://doi.org/10.1007/978-3-319-58068-5_32)
22. Price, C., Spackman, K.: SNOMED clinical terms. *Br. J. Healthc. Comput. Inf. Manag.* **17**(3), 27–31 (2000)
23. Rector, A.L., Solomon, W.D., Nowlan, W.A., Rush, T.W., Zanstra, P.E., Claassen, W.M.: A terminology server for medical language and medical information systems. *Methods Inf. Med.* **34**(1–2), 147–157 (1995)
24. Ruch, P., Gobeill, J., Lovis, C., Geissbühler, A.: Automatic medical encoding with SNOMED categories. *BMC Med. Inform. Decis. Mak.* **8**(Suppl. 1), S6 (2008)
25. Schlobach, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI 2003, pp. 355–360. Morgan Kaufmann Publishers Inc. (2003)

26. Sidhu, E.S., Dillon, T.S., Chang, E., Sidhu, B.S.: Protein ontology development using OWL. In: OWL Experiences and Directions Workshop, OWLED, p. 188 (2005)
27. Simancik, F., Kazakov, Y., Horrocks, I.: Consequence-based reasoning beyond Horn ontologies. In: IJCAI 2011, pp. 1093–1098. IJCAI/AAAI (2011)
28. Simancik, F., Motik, B., Horrocks, I.: Consequence-based and fixed-parameter tractable reasoning in description logics. *Artif. Intell.* **209**, 29–77 (2014)
29. Smith, B.: The OBO foundry: coordinated evolution of ontologies to support biomedical data integration. *Nat. Biotechnol.* **25**, 1251–1255 (2007)
30. Wang, C., Hitzler, P.: Consequence-based procedure for description logics with self-restriction. In: Semantic Web and Web Science - 6th Chinese Semantic Web Symposium and 1st Chinese Web Science Conference, CSWS 2012, Shenzhen, China, 28–30 November 2012, pp. 169–180 (2012)