Preliminary extended version of a paper that appeared at KR 2016

Minimality Postulates for Ontology Revision

Özgür L. Özçep
Institute of Information Systems (IFIS)
University of Lübeck, Germany
oezcep@ifis.uni-luebeck.de

Abstract

In many scenarios where the integration of information into a knowledge base (KB) leads to inconsistencies there is a need to change the KB minimally. In belief revision, relevance postulates meet the minimality requirement by restricting the elimination of KB elements to those that are relevant for the incoming information. This paper focuses on two minimality postulates in an ontology revision scenario in which conflicts are caused by ambiguous use of symbols: a relevance postulate and a generalized inclusion postulate which limits the creativity of the operators. Both postulates exploit the (satisfiably) equivalent representation of a first-order logic KB by its prime implicates, which, intuitively, represent the most atomic logical components of the KB. The paper shows that reinterpretation operators (which are ontology revision operators) fulfill both postulates.

Introduction

Not long after the seminal papers of Alchourrón, Gärdenfors and Makinson (AGM) (Alchourrón, Gärdenfors, and Makinson 1985) it was realized that belief-revision (BR) techniques could be fruitfully applied to different types of ontology change (OC) (Flouris et al. 2008) such as ontology evolution (Kharlamov, Zheleznyakov, and Calvanese 2013), ontology merge, ontology debugging etc. Most of the work exploiting BR for OC (Meyer, Lee, and Booth 2005; Flouris et al. 2006; Ribeiro and Wassermann 2009; Eschenbach and Özçep 2010; Özçep 2008) follows the dual approach of classical BR of, on the one hand, defining axiomatic specifications in the form of postulates and, on the other hand, constructing operators fulfilling them.

In this paper, I propose postulates that are intended to specify a minimal change of a knowledge base (or more concretely an ontology) and show that so-called reinterpretation operators fulfill them. The intended revision scenario of this paper for which the minimality postulates are going to be developed can be described as follows. A receiver agent holds an ontology which is formally described by a knowledge base (KB) in an expressive formal language such as first-order logic (FOL). In particular, a KB is a finite set of sentences in FOL (or a fragment of it). She receives information from a sender agent with possibly different ontology and she wants to integrate the information into her ontology.

I assume that both the sender's KB and the receiver's KB are well developed ontologies over the same application domain (e.g., ontologies for an online library system in universities, see examples below). Further it is assumed that the semantics of the same symbols in ontologies are strongly related. Nonetheless, there may be symbols that are used in differently by the sender and the receiver (ambiguity). The receiver is assumed to give priority to the sender's meanings so the integration result will contain the trigger and result in a revision of the receiver's ontology to preserve consistency. But, as the ontology of the receiver is assumed to be well developed, she is interested in changing her ontology only minimally, i.e., se wants to delete sentences of his KB and add additional sentences to it only as much as needed.

In belief revision the theme of minimality is mainly discussed within the context of relevance postulates (Hansson 1993; Parikh 1999). But also inclusion postulates can be seen as contributions to a minimal-change specification as they limit the operators's "creativity" by prescribing an upper bound to the result. In this paper, I start from these postulates for classical BR, argue why they are not proper minimality specifications for the intended revision scenario and formulate radically adapted versions that exploit the fine grained structure of ontologies by the notion of prime implicates. This is an extended version a paper that appeared at KR 16.

Logical Preliminaries

A first-order logic (FOL) vocabulary $\mathcal{V}$ consists of constants, predicate symbols and function symbols. For a FOL formula or set of formulas X let $\mathcal{V}(X)$ be the set of non-logical symbols occurring in X. I use the usual notions from Tarskian semantics based on FOL interpretations $\mathcal{I}$. The set of sentences containing only non-logical symbols in the vocabulary $\mathcal{V}$ is denoted $\mathrm{Sent}(\mathcal{V})$. The set of sentences in $\mathrm{Sent}(\mathcal{V})$ following from a set of sentences X (over a perhaps larger vocabulary) is denoted by $\mathrm{Cn}^{\mathcal{V}}(X)$. If two sets of FOL sentences X_1, X_2 are equivalent, I write $X_1 \equiv X_2$.

A non-logical symbol $s \in \mathcal{V}$ *properly occurs in a sentence* $\alpha \in \mathrm{Sent}(\mathcal{V})$ iff there are $\mathcal{I}_1, \mathcal{I}_2 \in \mathrm{Int}(\mathcal{V})$, s.t.: $\mathcal{I}_1$ and $\mathcal{I}_2$ differ only in the denotation of s and $\alpha^{\mathcal{I}_1} \neq \alpha^{\mathcal{I}_2}$. Let $P \in \mathcal{V}$ be an n-ary predicate symbol in $\mathcal{V}$. It occurs *syntactically positive (negative)* in a FOL formula iff it occurs in the scope of an even (uneven) number of negations—assuming that only the propositional truth functions $\wedge, \vee, \neg$ are used.

For $P \in \mathcal{V}(\alpha)$ we say that P *occurs semantically positive in sentence* α, $pos(P, \alpha)$ for short, iff: For all $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ and for subsets $D_1, D_2 \subseteq (\Delta^{\mathcal{I}})^n$ one has: If $D_1 \subseteq D_2$ and $\mathcal{I}_{[P \mapsto D_1]} \models \alpha$, then also $\mathcal{I}_{[P \mapsto D_2]} \models \alpha$. P *occurs semantically negative in sentence* α, $neg(P, \alpha)$ for short, iff $pos(P, \neg\alpha)$. P occurs *mixed in* α, $mix(P, \alpha)$ for short, iff it properly occurs in α but neither $pos(P, \alpha)$ nor $neg(P, \alpha)$. We write $pos0(P, \alpha)$ (resp. $neg0(P, \alpha)$) iff $pos(P, \alpha)$ (resp. $neg(P, \alpha)$) or P does not occur syntactically in α.

The *dual remainder sets modulo* α, $B \top \alpha$, consists of inclusion maximal subsets X of B that are consistent with α, i.e., $X \in B \top \alpha$ iff $X \subseteq B$, $X \cup \{\alpha\}$ is consistent and for all $\bar{X} \subseteq B$ with $X \subset \bar{X}$ the set $\bar{X} \cup \{\alpha\}$ is not consistent. The notion of dual remainders is extended to arbitrary KBs B_1 as second argument by defining $B \top B_1$ as $B \top \bigwedge B_1$.

Minimality in Belief Revision

The AGM (Alchourrón, Gärdenfors, and Makinson 1985) postulates do not constrain the revision result in the interesting case of conflict between KB and trigger. In fact, the amnesic operator defined by $B * \alpha = \mathrm{Cn}(\alpha)$ fulfills all AGM postulates though it is not minimal. The relevance postulates of Hansson (Hansson 1993) and of Parikh (Parikh 1999) are two different possibilities that remedy the unwanted property of amnesic revision. These kinds of postulates constrain the revision result by an approximation from below in the sense that they say which set of sentences X have to be in the (set of consequences of the) revision result: $X \subseteq \mathrm{Cn}(B * \alpha)$.

The relevance postulate of Hansson (Hansson 1993) is formulated for arbitrary, i.e. not necessarily logically closed, sets of sentences B called belief bases.

(Rel-H) If $\beta \in B$ and $\beta \notin B * \alpha$, then there is a set B', such that: 1. $B * \alpha \subseteq B' \subseteq B \cup \{\alpha\}$; 2. B' is consistent; 3. $B' \cup \{\beta\}$ is not consistent.

(Rel-H) it is not an adequate postulate for the intended revision scenario. In this scenario, it is not individual sentences that cause a conflict but different uses of symbols in B and β. And indeed, the reinterpretation based operators defined below do not fulfill this postulate.

Example 1 *Let B be a KB according to which media p_1, p_2, which are published in some proceedings, are articles: $B = \{Article(p_1), Article(p_2)\}$. The trigger $\alpha = \neg Article(p_1)$ stems from an agent with a different understanding of 'article' according to which only publications in journals are articles. An appropriate revision result $B * \alpha$ would not only delete $Article(p_1)$ but also $Article(p_2)$; because the next time the sender sends a trigger containing $Article$ negatively, namely $\neg Article(p_2)$, a conflict will occur. But this operator $*$ does not fulfill (Rel-H).*

A different relevance postulate called (Rel-P), which is more symbol-oriented, was formulated by Parikh (Parikh 1999). Criterion (Rel-P) is formulated for propositional KBs and so cannot be used directly for FOL KBs as assumed in this paper. Hence we define a different relevance postulate.

The relevance postulates cover only one aspect of minimality, but completely miss the other aspect of minimality which is to constrain the (consequences of the) revision result from above. That is, one has to prescribe a set X such that $\mathrm{Cn}(B * \alpha) \subseteq X$. In belief base revision this aspect is handled by the so-called inclusion postulate.

(Incl) $B * \alpha \subseteq B \cup \alpha$.

But for the revision scenario of this paper, belief base revision is not the means of choice as its results depend on the syntactic representation of the belief bases.

Adapted Minimality Postulates

For the following I will assume that B is a predicate logical KB without the identity and function symbols. The new relevance postulate adapts Hansson's relevance postulate (Rel-H). The main technical tool for the adaptation is the concept of a prime implicate, which roughly represents a most atomic component of the KB. While the notion of prime implicate is omnipresent for propositional logic (Armstrong et al. 1998) and has been exploited for the definitions of propositional revision operators (Bienvenu, Herzig, and Qi 2008; Zhuang, Pagnucco, and Meyer 2007), there is no real semantic notion of prime implicate for FOL and there is no approach that uses prime implicates in the postulates—except for (Özçep 2012) for propositional logic.

The core idea of the new relevance theorem is this: A sentence β entailed by B may be eliminated from the revision result if there is a related sentence ϵ of the normal form of B that together with other formulas of the normal form leads to a contradiction. The kind of relatedness between β and ϵ is explicated technically as follows.

Definition 1 *β and ϵ are called* related w.r.t. P *iff a) either $mix(P, \epsilon)$ or $mix(P, \beta)$; or b) $pos(P, \epsilon)$ and $pos(P, \beta)$; or c) $neg(P, \epsilon)$ and $neg(P, \beta)$.*

For the normal form representation I use prime implicates. A FOL formula α is universal iff α is equivalent to a formula in prenex form containing only all-quantifiers $\forall$. A universal formula of the form $\forall x_1 \ldots \forall x_n (li_1 \vee \cdots \vee li_m)$, where the li_j are literals with variables in $\{x_1, \ldots, x_n\}$, is a *FOL clause*. A FOL clause $\alpha_1 = \forall x_1 \ldots \forall x_n \beta$ is a *(proper) subclause* of a FOL clause α_2 iff α_2 is of the form $\alpha_2 = \forall y_1 \ldots \forall y_n \delta$, where all x_i are among the y_j and the set of literals in β is a (proper) subset of the literals in δ.

Let X be a set of universal formulas. The set of *FOL clauses of X w.r.t. a vocabulary $\mathcal{V}$*, $\mathrm{Cl}^{\mathcal{V}}(X)$, is the set of all FOL clauses in $\mathrm{Sent}(\mathcal{V})$ entailed by X. If X is an arbitrary set of FOL sentences, let X^* be the result of skolemizing every sentence in X. Let $\mathcal{V}_{\mathrm{sk}}$ be the set of used skolem symbols and let $\hat{\mathcal{V}} = \mathcal{V} \cup \mathcal{V}_{\mathrm{sk}}$. The set of *FOL clause of X w.r.t. $\mathcal{V}$ and skolem symbols $\mathcal{V}_{\mathrm{sk}}$* is defined by $\mathrm{Cl}^{\hat{\mathcal{V}}}(X^*)$.

The set of *FOL prime implicates* of a set of universal formula X w.r.t. $\mathcal{V}$ consists of non-tautological clauses of X for which there is no proper subclause in $\mathrm{Cl}^{\mathcal{V}}(X)$.

$$\mathrm{PI}^{\mathcal{V}}(X) = \{pr \in \mathrm{Cl}^{\mathcal{V}}(X) \mid pr \text{ is non-tautological and has no proper subclauses in } \mathrm{Cl}^{\mathcal{V}}(X)\}$$

This notion leads to a logically equivalent characterisation of sets X containing only universal formulas.

Proposition 1 *For every set X of universal formulas with $\mathcal{V}(X) \subseteq \mathcal{V}$: $X \equiv \mathrm{PI}^{\mathcal{V}}(X)$.*

The postulate of *reinterpretation relevance* (Rel-R) reads as follows:

(Rel-R) Let be given a vocabulary $\mathcal{V}$, a FOL KB B over $\mathcal{V}$, an FOL sentence α over $\mathcal{V}$ and a FOL clause β over $\mathcal{V}$. Let B^* be a skolemization of $\bigwedge B$ with skolem constants from $\mathcal{V}_{\text{sk}}$ and again $\hat{\mathcal{V}} = \mathcal{V} \cup \mathcal{V}_{\text{sk}}$.

If $B \models \beta$ and $B * \alpha \not\models \beta$, then there is a set X and a sentence $\epsilon \in X$ s.t.: 1. $X \subseteq \text{PI}^{\hat{\mathcal{V}}}(B^*)$; 2. $X \cup \{\alpha\}$ is inconsistent; 3. $(X \setminus \{\epsilon\}) \cup \{\alpha\}$ is consistent and 4. ϵ is related with β w.r.t. a predicate symbol P.

Example 2 *Consider the following KB B and trigger α: $B = \{Article(p_1), Article(p_2), \neg Article(bo_1)\}$ and $\alpha = \neg Article(p_1)$. Clearly $\text{PI}^{\hat{\mathcal{V}}}(B^*) = B^* = B$. Let $\beta = \neg Article(bo_1)$. Then $B \models \beta$ and $B * \alpha \not\models \beta$. But for $Article$ there is no $X \subseteq \text{PI}^{\hat{\mathcal{V}}}(B^*)$ that fulfills the conditions of (Rel-R) because the only β-related prime implicate is $\neg Article(bo_1)$ which is not involved in a conflict.*

Prime implicates can be further exploited to define a postulate that captures the other aspect of minimal revision where one constrains the result from above. The idea is to enrich the given KB B to an equivalent set $\text{Enr}(B)$ that contains enough consequences of B in order to identify the real potential culprits in the revision process. The general schema of the extended inclusion axiom is the following:

(Incl-ES) For all α there is $X \subseteq \text{Enr}(B)$ s.t. $X \cup \{\alpha\} \not\models \bot$, and for all β: If $B * \alpha \models \beta$, then $X \cup \{\alpha\} \models \beta$.

This schema says: There is a subset of the enrichment of B such that all sentences β entailed by the revision result follow from a subset X of the enrichment together with the trigger α. The enrichment operator Enr that I use in the following is defined as: $\text{Enr}(B) = B \cup \text{PI}^{\hat{\mathcal{V}}}(B^*)$.

In fact, though the enriched KB $\text{Enr}(B)$ is not equivalent to B, it is at least equivalent w.r.t. the non-skolem symbols.

Proposition 2 $\text{Cn}^{\mathcal{V}}(B) = \text{Cn}^{\mathcal{V}}(\text{Enr}(B))$

I call the postulate that results from (Incl-ES) by instantiating the parameter Enr by $\text{Enr}(B) = B \cup \text{PI}^{\hat{\mathcal{V}}}(B^*)$ the *extended inclusion postulate (Incl-E)*.

Reinterpretation Operators

The extended relevance postulate and inclusion postulate are intended to specify minimal changes of operators which are used in a particular semantic integration scenario described in the introduction. In this section, we recapitulate the definition of operators of this kind (Eschenbach and Özçep 2010; Özçep 2008) and show that they fulfill the new postulates. (For other postulates see (Özçep 2008)). The construction of the operators mimics the construction of the propositional revision operators of (Delgrande and Schaub 2003).

The revision operator defined in the following is denoted by $\circ$ and is called a reinterpretation operator. $\circ$ is a binary operator with a finite FOL KB as left and a FOL sentence α as right argument. Before giving the technical definition, the main construction idea will be illustrated with an example.

Example 3 *Define the knwoledge base B as*

$$B = \{Article(p_1), Article(p_2), \neg Article(bo_1)\}$$

and the trigger $\alpha = \neg Article(p_1)$. The reinterpretation operator $\circ$ results in the following KB:

$$B \circ \alpha = \{Article'(p_1), Article'(p_2), \neg Article'(bo_1), \neg Article(p_1), \forall x(Article(x) \rightarrow Article'(x))\}$$

The conflict between B and α is traced back to ambiguous use of symbols. As I assume that only predicate symbols (and not constants) may be used ambiguously, the conflict can only be caused by different uses of the unary predicate $Article$. The receiver (holder of B) gives priority to the sender's use of $Article$ over her use of $Article$, and hence puts $\neg Article(p_1)$ in $B \circ \alpha$. Her own use of $Article$ is internalized, i.e., all occurrences of $Article$ in B are substituted by a new symbol $Article'$. The receiver adds hypotheses on the semantical relatedness (bridging axioms) of his and the sender's use of $Article$. The hypothesis in this case is $\forall x(Article(x) \rightarrow Article'(x))$ which says that articles in the sender's sense are also articles in the receiver's sense. Note that because of this hypothesis the result $B \circ \alpha$ entails the assertion $\neg Article(bo_1)$ from the initial KB B.

Technically the disambiguation is realized by uniform substitutions called *ambiguity compliant resolution substitutions*, $\text{AR}(\mathcal{V}, \mathcal{V}')$ for short. Here, $\mathcal{V} \cap \mathcal{V}' = \emptyset$ where $\mathcal{V}'$ is the set of symbols used for internalization. The substitutions in $\text{AR}(\mathcal{V}, \mathcal{V}')$ get as input a non-logical symbol in $\mathcal{V}$ and map it either to itself or to a new non-logical symbol (of the same type) in $\mathcal{V}'$. I only consider the substitution of predicate symbols. $\text{supp}(\sigma) = \{s \in \mathcal{V} \mid \sigma(s) \neq s\}$ is called the support of σ. A substitution with support S is also denoted by σ_S. For $\sigma_1, \sigma_2 \in \text{AR}(\mathcal{V}, \mathcal{V}')$ let $\sigma_1 \leq \sigma_2$ iff $\text{supp}(\sigma_1) \subseteq \text{supp}(\sigma_2)$. A *disambiguation schema* Φ picks for every S a substitution $\Phi(S) \in \text{AR}(\mathcal{V}, \mathcal{V}')$ with support S.

In general, there may be more than one predicate symbol which has to be disambiguated; and there may be many different sets of symbols for which a disambiguation leads to consistency. So I define the minimal conflict symbol sets:

Definition 2 *Let B be a FOL KB over $\mathcal{V}$ and α a FOL sentence over $\mathcal{V}$. The set of* minimal conflicting symbols sets, $\text{MCS}(B, \alpha)$, *is defined by:*

$$\text{MCS}(B, \alpha) = \{\ S \subseteq \mathcal{V} \mid \textit{There is } \sigma_S \in \text{AR}(\mathcal{V}, \mathcal{V}'),\ \textit{s.t. } B\sigma_S \cup \{\alpha\} \textit{ is consistent, and for all } \sigma_R \in \text{AR}(\mathcal{V}, \mathcal{V}') \textit{ with } \sigma_R < \sigma_S\ B\sigma_R \cup \{\alpha\} \textit{ is not consistent.}\ \}$$

As no symbol set in $\text{MCS}(B, \alpha)$ is a better candidate than the other, we assume that a selection function γ_1 selects the good candidates: $\gamma_1(\text{MCS}(B, \alpha)) \subseteq \text{MCS}(B, \alpha)$. In the end, the symbol set $S^{\#} = \bigcup \gamma_1(\text{MCS}(B, \alpha))$ is the set of symbols which will be internalized.

In the second step, the disambiguated symbols of $S^{\#}$ are related by bridging axioms. Depending on what kind of bridging axioms are chosen, different integration operators result. Here, I consider so-called *simple bridging axioms*. (For other types of bridging axioms see (Özçep 2008).) Let $\sigma = \sigma_S \in \text{AR}(\mathcal{V}, \mathcal{V}')$ be a substitution with support $S \subseteq \mathcal{V}$.

Let P be an n-ary predicate symbol in S, $\sigma(P) = P'$ and let $\vec{x} = x_1, \ldots, x_n$. Then define $\overrightarrow{P} = \forall \vec{x}(P(\vec{x}) \rightarrow P'(\vec{x}))$ and $\overleftarrow{P} = \forall \vec{x}(P'(\vec{x}) \rightarrow P(\vec{x}))$.

Definition 3 *For $\sigma = \sigma_S \in \mathrm{AR}(\mathcal{V}, \mathcal{V}')$ for $S \subseteq \mathcal{V}$ the* simple bridging axioms *w.r.t. σ are* $\mathrm{BA}(\sigma) = \{\overrightarrow{P}, \overleftarrow{P} \mid P \in S\}$

In case of conflict, not all bridging axioms of $\mathrm{BA}(S^{\#})$ can be added to the revision result. Again a selection function (γ_2) is used.

Definition 4 *Let $\mathcal{V} \cap \mathcal{V}' = \emptyset$; Φ be a disambiguation scheme; $\overline{\gamma} = (\gamma_1, \gamma_2)$ be a pair of selection functions. For any FOL KB B and FOL sentence α over $\mathcal{V}$ let $S^{\#} = \bigcup \gamma_1(\mathrm{MCS}(B, \alpha))$ and $\sigma = \Phi(S^{\#})$. Then the* reinterpretation operator $\circ = \circ^{\overline{\gamma}}$ *is defined by*

$$B \circ \alpha \quad = \quad \sigma(B) \cup \{\alpha\} \cup \bigcap \gamma_2\big(\mathrm{BA}(\sigma) \top (\sigma(B) \cup \{\alpha\})\big)$$

It can easily be checked that this definition of $\circ$ gives the results in Ex. 3 (for any pair of selection functions γ_1, γ_2).

Minimality of Reinterpretation

Reinterpretation operators fulfill the reinterpretation relevance postulate and the extended inclusion postulate.

Theorem 1 *The reinterpretation operators fulfill the postulates (Rel-R) and (Incl-ES).*

The main proof idea is to explicate the interaction of the internalization and of the bridging axioms with the prime implicates entailed by B as done by the following propositions.

Proposition 3 *Let $\mathcal{V}$ and $\mathcal{V}'$ be disjoint vocabularies; B be a set of universal formula in FOL (without identity and function symbols) over $\mathcal{V}$; σ be a substitution of predicate symbols P by new symbols $\sigma(P) \in \mathcal{V}'$ and $\mathrm{PI}(\cdot) = \mathrm{PI}^{\mathcal{V} \cup \mathcal{V}'}(\cdot)$. Then: $\mathrm{Cn}^{\mathcal{V}}(\mathrm{PI}(B\sigma)) = \mathrm{Cn}^{\mathcal{V}}(\mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V}))$*

The proposition says that in order to discover losses of $\mathcal{V}$-consequences of B one can stick to its prime implicates.

While this proposition hints to the interaction of prime implicates with the internalization, the following one talks about their interaction with simple bridging axioms. Let $B^* = \forall \tilde{x}_1 \ldots \forall \tilde{x}_m \tilde{B}$ be a skolemization of B with skolem constants not in $\mathcal{V}(B \cup B\sigma)$. Then $B^*\sigma = \forall \tilde{x}_1 \ldots \forall \tilde{x}_m \tilde{B}\sigma$ is a skolemization of $B\sigma$. Let $\forall z \tilde{ba}$ be a prenex form of some set of bridging axioms $ba \subseteq \mathrm{BA}(\sigma)$. Then $(B\sigma \cup ba)^*$ is called an *B^*-admissible skolemization* of $B\sigma \cup ba$ iff it has the form $(B\sigma \cup ba)^* = \forall z \forall \tilde{x}_1 \ldots \forall \tilde{x}_m (\tilde{B}\sigma \wedge \tilde{ba})$.

Proposition 4 *Let $\mathcal{V}$, $\mathcal{V}'$, $\mathcal{V}_{\mathrm{sk}}$ be pairwise disjoint vocabularies. Let B be a KB over $\mathcal{V}$ and σ be a substitution of predicate symbols P by new predicate symbols $\sigma(P) \in \mathcal{V}'$; $ba \subseteq \mathrm{BA}(\sigma)$ be a subset of bridging axioms; and $(B\sigma \cup ba)^*$ be a B^*-admissible skolemization of $B\sigma \cup ba$ with skolem constants from $\mathcal{V}_{\mathrm{sk}}$; then:*

$$\mathrm{PI}^{\hat{\mathcal{V}} \cup \mathcal{V}'}((B\sigma \cup ba)^*) \cap \mathrm{Sent}(\mathcal{V} \cup \mathcal{V}(B^*)) \subseteq \mathrm{PI}^{\hat{\mathcal{V}}}(B^*)$$

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Armstrong, T.; Marriott, K.; Schachte, P.; and Søndergaard, H. 1998. Two classes of boolean functions for dependency analysis. *Science of Computer Programming* 31(1):3–45.

Bienvenu, M.; Herzig, A.; and Qi, G. 2008. Prime implicate-based belief revision operators. In Ghallab, M.; Spyropoulos, C. D.; Fakotakis, N.; and Avouris, N. M., eds., *ECAI*, volume 178, 741–742. IOS Press.

Delgrande, J. P., and Schaub, T. 2003. A consistency-based approach for belief change. *Artificial Intelligence* 151(1–2):1–41.

Eschenbach, C., and Özçep, Ö. L. 2010. Ontology revision based on reinterpretation. *Logic Journal of the IGPL* 18(4):579–616. First published online August 12, 2009.

Flouris, G.; Huang, Z.; Pan, J. Z.; Plexousakis, D.; and Wache, H. 2006. Inconsistencies, negations and changes in ontologies. In Gil, Y., and Mooney, R., eds., *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, 1295–1300. AAAI Press.

Flouris, G.; Manakanatas, D.; Kondylakis, H.; Plexousakis, D.; and Antoniou, G. 2008. Ontology change: classification and survey. *The Knowledge Engineering Review* 23(2):117–152.

Hansson, S. O. 1993. Reversing the Levi identity. *Journal of Philosophical Logic* 22:637–669.

Kharlamov, E.; Zheleznyakov, D.; and Calvanese, D. 2013. Capturing model-based ontology evolution at the instance level: The case of dl-lite. *J. Comput. Syst. Sci.* 79(6):835–872.

Meyer, T.; Lee, K.; and Booth, R. 2005. Knowledge integration for description logics. In Veloso, M. M., and Kambhampati, S., eds., *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-2005)*, 645–650. AAAI Press/The MIT Press.

Özçep, Ö. L. 2008. Towards principles for ontology integration. In Eschenbach, C., and Grüninger, M., eds., *FOIS*, volume 183, 137–150. IOS Press.

Özçep, Ö. L. 2012. Knowledge-base revision using implications as hypotheses. In Glimm, B., and Krüger, A., eds., *KI 2012*, LNCS, 217–228. Springer Berlin Heidelberg.

Parikh, R. 1999. Beliefs, belief revision, and splitting languages. In Moss, L.; Ginzburg, J.; and de Rijke, M., eds., *Logic, Language and Computation*, volume 2. CSLI Publications. 266–278.

Ribeiro, M. M., and Wassermann, R. 2009. Base revision for ontology debugging. *J. Log. Comput.* 19(5):721–743.

Zhuang, Z.; Pagnucco, M.; and Meyer, T. 2007. Implementing iterated belief change via prime implicates. In Orgun, M., and Thornton, J., eds., *AI 2007: Advances in Artificial Intelligence*, volume 4830 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 507–518.

Appendix: Proofs of Results

Proof of Proposition 1

Every formula in X is equivalent to a prenex formula containing only $\forall$-quantifiers and having a matrix which is in CNF. Hence, X is equivalent to a set of FOL clauses w.r.t. $\mathcal{V}$. Consequently, $\mathrm{Cl}^{\mathcal{V}}(X) \equiv X$. We show $\mathrm{PI}^{\mathcal{V}}(X) \equiv \mathrm{Cl}^{\mathcal{V}}(X)$. First, $\mathrm{PI}^{\mathcal{V}}(X) \subseteq \mathrm{Cl}^{\mathcal{V}}(X)$ and hence $\mathrm{Cl}^{\mathcal{V}}(X) \models \mathrm{PI}^{\mathcal{V}}(X)$. Second, we show $\mathrm{PI}^{\mathcal{V}}(X) \models \mathrm{Cl}^{\mathcal{V}}(X)$ by showing that for $kl \in \mathrm{Cl}^{\mathcal{V}}(X)$ one can find a $pr \in \mathrm{PI}^{\mathcal{V}}(X)$ with $pr \models kl$. If kl is tautological, then already $\emptyset \models kl$. Otherwise, if $kl \in \mathrm{PI}^{\mathcal{V}}(X)$, then let $pr = kl$. If $kl \notin \mathrm{PI}^{\mathcal{V}}(X)$, then there is a $kl' \in \mathrm{Cl}^{\mathcal{V}}(X)$, s.t. kl' is a proper subclause of kl. As the matrices are finite, one can find a $kl' \in \mathrm{Cl}^{\mathcal{V}}(X)$ which is minimal w.r.t. the subclause relation and which is a proper subclause of kl. But then, $pr = kl' \in \mathrm{PI}^{\mathcal{V}}(X)$ and $pr \models kl$.

Proof of Proposition 2

Let $\beta \in \mathrm{Sent}(\mathcal{V})$. If $B \models \beta$, then by definition $\mathrm{Enr}(B) \models \beta$. If $\mathrm{Enr}(B) \models \beta$, then $\mathrm{PI}^{\hat{\mathcal{V}}}(B^*) \models \bigwedge B \rightarrow \beta$. Because of Proposition **??** it follows that $B \models B \rightarrow \beta$ and hence $B \models \beta$.

Proof of Proposition 3

We let $\mathrm{PI} = \mathrm{PI}^{\mathcal{V}\cup\mathcal{V}'}$. The $\supseteq$ relations holds because of the monotonicity of $\mathrm{Cn}^{\mathcal{V}}$. I show the $\subseteq$ relation. I use the following abbreviations: $\Gamma_a = \mathrm{PI}(B\sigma)$ and $\Gamma_b = \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V})$. Let $\mathrm{supp}(\sigma) = \{P_1, \ldots, P_n\}$ and $P_i' = \sigma(P_i)$.

Let be given a formula β with $\beta \in \mathrm{Sent}(\mathcal{V})$ and let $\Gamma_b \not\models \beta$. Then there is an interpretation $\mathcal{J}' \models \Gamma_b \cup \{\neg\beta\}$. We skolemize $\neg\beta$ into a universal formula δ. As $\Gamma_b \cup \{\neg\beta\}$ is consistent, $\Gamma_b \cup \{\delta\}$ is consistent, too, and has a herbrand model $\mathcal{J} \models \Gamma_b \cup \{\delta\}$. Let for short $\Delta = \Delta^{\mathcal{J}}$. All other interpretations to be constructed will be herbrand models and will have the same domain Δ. Because all are herbrand models we have the following fact: For all formula α and all terms $t_1, \ldots, t_n \in \Delta^{\mathcal{J}}$:

$$\mathcal{J}_{[x_1 \mapsto t_1, \ldots, x_n \mapsto t_n]} \models \alpha \text{ iff } \mathcal{J} \models \alpha[x_1/t_1, \ldots, x_n/t_n]$$

Starting with $\mathcal{J}$ we construct a sequence of interpretations $\mathcal{I}_i$ s.t.:

$$\begin{array}{rcl} \mathcal{J} = \mathcal{I}_0 & \models & \Gamma_b \cup \{\delta\} \\ \mathcal{I}_1 & \models & \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1'\}) \cup \{\delta\} \\ & \ldots & \\ \mathcal{I}_n & \models & \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_n'\}) \cup \{\delta\} \\ & & = \Gamma_a \cup \{\delta\} \end{array}$$

The interpretation $\mathcal{I}_n$ then shows that $\Gamma_a \cup \{\delta\}$ and so $\Gamma_a \cup \{\neg\beta\}$ is consistent. This will prove $\Gamma_a \not\models \beta$, as desired.

The sequence $(\mathcal{I}_i)_{1 \leq i \leq n}$ is defined as follows. Suppose, $\mathcal{I}_{i-1}$ (for $1 \leq i \leq n$) has been constructed and it fulfils

$$\mathcal{I}_{i-1} \models \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_{i-1}'\}) \cup \{\delta\}$$

We have to construct a model $\mathcal{I}_i$ with:

$$\mathcal{I}_i \models \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\}) \cup \{\delta\}$$

$\mathcal{I}_i$ is built out of $\mathcal{I}_{i-1}$ by changing only the denotation of P_i' in a minimal fashion. The proofs are given for the the case that P_i' is a unary predicate symbol (but all proofs are easily generalizable to predicate symbols of arbitrary arity). Let H be the set of formulas in $\mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$, in which P_i' occurs at most positively in a literal $P_i'(t)$. As $\mathcal{I}_i$ changes only the denotation of P_i' in comparison with $\mathcal{I}_{i-1}$, we will also have $\mathcal{I}_i \models \delta$.

The following lemma shows the existence of a minimal model $\mathcal{I}_i$.

Lemma 1 *There is a model $\mathcal{I}_i \models H$ s.t. $\mathcal{I}_i$ assigns to P_i' an inclusion minimal subset of Δ, i.e., for all modifications $\mathcal{I}'$ of $\mathcal{I}_i$ with $(P_i')^{\mathcal{I}'} \subset (P_i')^{\mathcal{I}_i}$ there is an $\alpha \in H$ with $\mathcal{I}' \not\models \alpha$.*

Proof this lemma: H is the set of formula in $\mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$, in which$P_i'$ occurs at most positively in a literal of the form $P_i'(t)$. Let $(e_j)_{j\in\mathbb{N}}$ be an enumeration of $\Delta^{\mathcal{I}_{i-1}} = \Delta$. All e_j are grounded terms (i.e. do not contain variables). We set $\Delta^{\mathcal{I}_i} = \Delta^{\mathcal{I}_{i-1}} = \Delta$ and construct with a sequence $(M_i^k)_{k\in\mathbb{N}}$ the denotation of P_i' in $\mathcal{I}_i$. The set M_i^{k+1} is constructed out of M_i^k by eliminating a grounded term e_{k+1} only if not formula $\alpha \in H$ becomes false by the elimination. We start for $k = 0$ by setting $M_i^0 = \Delta$. Then $(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^0]} \models H$, as $\mathcal{I}_{i-1}$ makes all prime implicate true, in which P_i' does not occur , and as in H there are at most prime implicates with positive occurrence of P_i'. Because of the definition of M_i^k all interpretation $(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^k]}$ are models of H.

$$M_i^{k+1} = \begin{cases} M_i^k & \text{if there is an } \alpha \in H \text{ s.t.} \\ & (\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^k \setminus \{e_{k+1}\}]} \not\models \alpha \\ M_i^k \setminus \{e_{k+1}\} & \text{else} \end{cases}$$

At the end let $(P_i')^{\mathcal{I}_i} = \bigcap_{k\in\mathbb{N}} M_i^k$.

We first show that $\mathcal{I}_i$ is a model of all $\alpha \in H$ and then show its minimality. If $\alpha \in H$ does not contain any of the P_i', then $\mathcal{I}_{i-1} \models \alpha$ entails $\mathcal{I}_i \models \alpha$. Otherwise α is of the form $\alpha = \forall x_1 \ldots \forall x_m P_i'(t_1) \vee \cdots \vee P_i'(t_p) \vee M$, where M is composed of literals, in which P_i' does not occur. Now assume for contradiction that $\mathcal{I}_i \models \neg\alpha$. Then there are $d_1, \ldots, d_m \in \Delta$ and a substitution $\rho = [x_1/d_1, \ldots, x_m/d_m]$, such that

$$\mathcal{I}_i \models \neg P_i'(t_1\rho) \wedge \cdots \wedge \neg P_i'(t_p\rho) \wedge \neg M\rho \qquad (1)$$

But all arguments of the P_i' are terms that are contained in the enumeration $(e_k)_{k\in\mathbb{N}}$. Hence there are $e_{k_1}, \ldots, e_{k_p} \in \Delta$ with

$$e_{k_1} = t_1\rho, \ldots, e_{k_p} = t_p\rho$$

Let $r = max\{k_1, \ldots, k_p\}$ the index of the ground term among $e_{k_1}, \ldots, e_{k_p}$ that was the last eliminated. By definition $M_i^r \cap \{e_{k_1}, \ldots, e_{k_p}\} = \emptyset$, hence

$$(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^r]} \models \neg P_i'(t_1\rho) \wedge \cdots \wedge \neg P_i'(t_p\rho)$$

By definition also $(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^r]} \models \alpha$. So we have $(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^r]} \models M\rho$. As M does not contain P_i' this implies $\mathcal{I}_i \models M\rho$ which contradicts (1).

Now we prove the minimality. It suffices to show that for all $e_k \in \Delta$ and all interpretations $\mathcal{I}'$ with $(P_i')^{\mathcal{I}'} = (P_i')^{\mathcal{I}_i} \setminus \{e_k\}$ there exists a formula $\alpha \in H$ for which $\mathcal{I}' \models \neg\alpha$. (As i P_i' occurs only negatively in $\neg\alpha$ for all other interpretation $\mathcal{I}''$ assigning P_i' an even smaller denotation, i.e. $(P_i')^{\mathcal{I}''} \subset (P_i')^{\mathcal{I}'}$, one has $\mathcal{I}'' \models \neg\alpha$.)

That $e_k \in (P_i')^{\mathcal{I}_i}$ means that $(\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^{k-1}]} \models H$ and that for the modification $\mathcal{J}' = (\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^{k-1} \setminus \{e_k\}]} = (\mathcal{I}_{i-1})_{[P_i' \mapsto M_i^k]}$ there exists $\alpha \in H$ s.t. $\mathcal{J}' \not\models \alpha$, i.e. $\mathcal{J}' \models \neg\alpha$. By definition $(P_i')^{\mathcal{I}'} = (P_i')^{\mathcal{I}_i} \setminus \{e_k\} \subseteq M_i^k = (P_i')^{\mathcal{J}'}$. As P_i' occurs only negatively in $\neg\alpha$, the fact that $\mathcal{J}' \models \neg\alpha$ implies $\mathcal{I}' \models \neg\alpha$ and hence also $\mathcal{I}' \not\models \alpha$. This ends the proof of the lemma.

Now we show that the minimal model $\mathcal{I}_i$ of Lemma 1 also makes all those prime implicates of $\mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$ true in which P_i' occurs at least once (syntactically) negatively. This can be proven by induction on the number k of negative occurrences of P_i'.

Base case ($k = 0$): If $pr \in \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$ and pr contains no negative occurrences of P_i', then $pr \in H$ and we already have $\mathcal{I}_i \models pr$ from what was proven above.

Induction hypothesis: For all prime implicates $pr \in \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$, which have less than k negative occurrences of P_i', it holds that $\mathcal{I}_i \models pr$.

Induction step: Let $\alpha \in \mathrm{PI}(B\sigma) \cap \mathrm{Sent}(\mathcal{V} \cup \{P_1', \ldots, P_i'\})$ be a prime implicate over the vocabulary $\mathcal{V} \cup \{P_1', \ldots, P_i'\}$ that has k negative occurrences of P_i':

$$\alpha := \forall x_1 \ldots \forall x_u \neg P_i'(t_1) \vee \cdots \vee \neg P_i'(t_k) \vee M$$

Here M consist of literals, in which P_i' does not occur negatively (but perhaps positively). In order to show $\mathcal{I}_i \models \alpha$ assume that

$$\mathcal{I}_i \models P_i'(t_1) \wedge \cdots \wedge P_i'(t_k)[x_1/d_1, \ldots, x_u/d_u] \quad (2)$$

for $d_1, \ldots, d_u \in \Delta^{\mathcal{I}_i}$. We have to show:

$$\mathcal{I}_i \models M[x_1/d_1, \ldots, x_u/d_u] \quad (3)$$

In particular, (2) entails $\mathcal{I}_i \models P_1'(t_1)[x_1/d_1, \ldots, x_u/d_u]$. Let $d := t_1[x_1/d_1, \ldots, x_u/d_u]$. As we can reorder the terms, we can assume that there is a n_0, $1 \leq n_0 \leq u$ s.t.:

$$\begin{aligned} d = t_1[x_1/d_1, \ldots, x_u/d_u] &= t_2[x_1/d_1, \ldots, x_u/d_u] = \ldots \\ \ldots &= t_{n_0}[x_1/d_1, \ldots, x_u/d_u] \end{aligned}$$

The minimality of $\mathcal{I}_i$ implies that there is a prime implicate $\overline{\alpha} \in H$ (containing only positive occurrences of P_i'), ground terms $\overline{d_1}, \ldots \overline{d}_m \in \Delta^{\mathcal{I}_i}$ and some n_1 with $1 \leq n_1 \leq v$ s.t.:

$$\overline{\alpha} = \forall \overline{y}_1 \ldots \overline{y}_m P_i'(\overline{t}_1) \vee \cdots \vee P_i'(\overline{t}_v) \vee N \quad (4)$$

where N does not contain P_i' and

$$\begin{aligned} \mathcal{I}_i \models \neg(P_i'(\overline{t}_{n_1+1}) \vee \cdots \vee P_i'(\overline{t}_v) \vee N) \\ [\overline{y}_1/\overline{d}_1, \ldots, \overline{y}_m/\overline{d}_m] \end{aligned} \quad (5)$$

and

$$\begin{aligned} &d = \\ &t_1[x_1/d_1, \ldots, x_u/d_u] = \cdots = t_{n_0}[x_1/d_1, \ldots, x_u/d_u] = \\ &\overline{t}_1[\overline{y}_1/\overline{d}_1, \ldots, \overline{y}_m/\overline{d}_m] = \cdots = \overline{t}_{n_1}[\overline{y}_1/\overline{d}_1, \ldots, \overline{y}_m/\overline{d}_m] \end{aligned}$$

By possibly renaming variables we can assume that all x_i are different from all $\overline{y}_j$. The last equation can thus be written as

$$t_1\rho = \cdots = t_{n_0}\rho = \overline{t}_1\rho = \cdots = \overline{t}_{n_1}\rho \quad (6)$$

The substitution ρ is given by

$$\rho := [x_1/d_1, \ldots, x_u/d_u, \overline{y}_1/\overline{d}_1, \ldots, \overline{y}_m/\overline{d}_m]$$

The equation chain (6) entails that $\{t_1, \ldots, t_{n_0}, \overline{t}_1, \ldots \overline{t}_{n_1}\}$ is unifiable. Hence there is a most general unifier σ^* s.t.

$$t_1\sigma^* = \cdots = t_{n_0}\sigma^* = \overline{t}_1\sigma^* = \cdots = \overline{t}_{n_1}\sigma^*$$

σ^* can be chosen s.t. in the range of σ^* at most variables from $\{x_1, \ldots, x_u, \overline{y}_1, \ldots, \overline{y}_m\}$ occur. In particular there is a substitution σ' such that $\rho = \sigma^*\sigma'$. σ' substitutes variables by ground terms because ρ does. Hence the clauses α and $\overline{\alpha}$ are resolvable to the clause kl given by

$$\begin{aligned} &\forall x_1 \ldots x_u \forall \overline{y}_1 \ldots \overline{y}_m (\neg P_i'(t_{n_0+1}) \vee \cdots \vee \neg P_i'(t_k) \vee M)\sigma^* \\ &\vee (P_i'(\overline{t}_{n_1+1}) \vee \cdots \vee P_i'(\overline{t}_v) \vee N)\sigma^*) =: kl \end{aligned}$$

There is a prime implicate pr which is a proper subclause of kl. In particular $pr \models kl$. pr has less than k negative occurrences of P_i'. But according to the induction hypothesis the interpretation $\mathcal{I}_i$ makes pr and hence also kl true.

$$\mathcal{I}_i \models kl \quad (7)$$

Now, (5) entails $\mathcal{I}_i \models \neg(P_i'(\overline{t}_{n_1+1}) \vee \cdots \vee P_i'(\overline{t}_v) \vee N)\rho$, which means $\mathcal{I}_i \models \neg(P_i'(\overline{t}_{n_1+1}) \vee \cdots \vee P_i'(\overline{t}_v) \vee N)\sigma^*\sigma'$. Hence there are $e_1, \ldots, e_{u+m} \in \Delta^{\tilde{\mathcal{I}}}$ s.t.:

$$\begin{aligned} &\mathcal{I}_{i[x_1 \mapsto e_1, \ldots, x_u \mapsto e_n, \overline{y}_1 \mapsto e_{u+1}, \ldots, \overline{y}_m \mapsto e_{u+m}]} \models \\ &\qquad \neg(P_i'(\overline{t}_{n_1+1}) \vee \cdots \vee P_i'(\overline{t}_v) \vee N)\sigma^* \end{aligned} \quad (8)$$

Because of (2) and $P_i'(t_h)[x_1/d_1, \ldots, x_u/d_u] = P_i'(t_h)\rho$ (for $n_0 + 1 \leq h \leq a$) we get:

$$\begin{aligned} &\mathcal{I}_{i[x_1 \mapsto e_1, \ldots, x_u \mapsto e_n, \overline{y}_1 \mapsto e_{n+1}, \ldots, \overline{y}_m \mapsto e_{u+m}]} \models \\ &\qquad \neg(\neg P_i'(t_{n_0+1}) \vee \cdots \vee \neg P_i'(t_k))\sigma^* \end{aligned} \quad (9)$$

With (9), (7) and (8) it follows that:

$$\mathcal{I}_{i[x_1 \mapsto e_1, \ldots, x_u \mapsto e_n, \overline{y}_1 \mapsto e_{u+1}, \ldots, \overline{y}_m \mapsto e_{u+m}]} \models M\sigma^*$$

But this relation means nothing else than $\mathcal{I}_i \models M\sigma^*\sigma'$, hence: $\mathcal{I}_i \models M\rho = M[x_1/d_1, \ldots, x_u/d_u]$. And this was the relation (3) which we wanted to prove.

Proof of Proposition 4

Let $pr \in \mathrm{PI}^{\hat{\mathcal{V}} \cup \mathcal{V}'}((B\sigma \cup B)^*) \cap \mathrm{Sent}(\hat{\mathcal{V}})$. We have to show $pr \in \mathrm{PI}^{\hat{\mathcal{V}}}(B^*)$. Clearly, $pr \in Cl^{\hat{\mathcal{V}}}(B^*)$. Assume $pr \notin \mathrm{PI}^{\hat{\mathcal{V}}}(O^*)$. Then there is a clause $kl \in Cl^{\hat{\mathcal{V}}}(B^*)$, which is a proper subclause of pr. There are two principal cases: a) $kl \in \mathrm{Cl}^{\hat{\mathcal{V}} \cup \mathcal{V}'}((B\sigma \cup ba)^*)$; b) $kl \notin \mathrm{Cl}^{\hat{\mathcal{V}} \cup \mathcal{V}'}((B\sigma \cup ba)^*)$. Both cases are contradictory as shown below.

Case a): This contradicts the fact that pr is prime w.r.t. $(B\sigma \cup ba)^*$.

Case b): In this case $(B\sigma \cup ba)^* \not\models kl$ and $kl \models pr$. The first fact and Lemma **??** imply that kl contains an P such that i) there is no bridging axiom for it in ba or ii) there is a bridging axiom in the wrong direction.

Subcase i): Because of Lemma **??** there is for pr a clause kl' s.t.: $kl' \in Cl^{\hat{\mathcal{V}}}((B\sigma \cup ba)^*)$; P is not contained in kl' and $kl' \models pr$. Let pr' be constructed from pr by deleting all literals with P. Then $(B\sigma \cup ba)^* \models kl' \models pr'$, contradicting the primeness of pr w.r.t. $(B\sigma \cup ba)^*$.

Subcase ii): W.l.o.g. assume $\overleftarrow{P} \in ba$. Then kl contains P at least once negatively. Because of Lemma **??** there is for pr a clause $kl' \in Cl^{\hat{\mathcal{V}}}((B\sigma \cup ba)^*)$ which contains P only positively and for which $(B\sigma \cup ba)^* \models kl' \models pr$. The predicate symbol P can occur in pr at most positively. Otherwise we could construct a smaller pr' with $(B\sigma \cup ba)^* \models kl' \models pr'$ contradicting the primeness of pr w.r.t. $(B\sigma \cup ba)^*$. Now $kl \models pr$ and so also $kl[s/\bot] \models pr[s/\bot]$. As P occurs syntactically negative in kl, $kl[s/\bot]$ is a tautology; but then also $pr[s/\bot]$is a tautology—contradicting the primeness of pr w.r.t. $(B\sigma \cup ba)^*$.