Embedding Non-Ground Logic Programs into Autoepistemic Logic for Knowledge-Base Combination*

Jos de Bruijn[1], **Thomas Eiter**[2], **Axel Polleres**[3], **and Hans Tompits**[2]
[1]Digital Enterprise Research Institute (DERI), Leopold-Franzens Universität Innsbruck, Austria
[2]Institut für Informationssysteme 184/3, Technische Universität Wien, Austria
[3]Universidad Rey Juan Carlos, Madrid, Spain
jos.debruijn@deri.org, {eiter,tompits}@kr.tuwien.ac.at, axel.polleres@urjc.es

Abstract

In the context of the Semantic Web, several approaches to the combination of ontologies, given in terms of theories of classical first-order logic, and rule bases have been proposed. They either cast rules into classical logic or limit the interaction between rules and ontologies. Autoepistemic logic (AEL) is an attractive formalism which allows to overcome these limitations, by serving as a uniform host language to embed ontologies and nonmonotonic logic programs into it. For the latter, so far only the propositional setting has been considered. In this paper, we present several embeddings of normal and disjunctive non-ground logic programs under the stable-model semantics into first-order AEL, and compare them in combination with classical theories, with respect to stable expansions and autoepistemic consequences. Our results reveal differences and correspondences of the embeddings and provide a useful guidance in the choice of a particular embedding for knowledge combination.

1 Introduction

In the context of the ongoing discussion around combinations of rules and ontologies for the Semantic Web, there have been several proposals for integrating classical knowledge bases (ontologies) and rule bases (logic programs). Generally speaking, all these approaches try to define a reasonable semantics for a combined knowledge base consisting of a classical component and a rules component.

Two trends are currently observable. On the one hand, approaches such as SWRL [Horrocks and Patel-Schneider, 2004] extend the ontology with Horn formulas in a classical framework. This approach is straightforward, but prohibits nonmonotonic rules. On the other hand, existing approaches which do allow nonmonotonic rules either (a) distinguish between "classical" and "rules" predicates and limit the domain of interpretation (e.g., [Rosati, 2006]) or (b) restrict the interaction to ground entailment (e.g., [Eiter *et al.*, 2004]). The main distinction between these approaches is the type of interaction between the classical knowledge base on the one hand and the rule base on the other (cf. de Bruijn *et al.* [2006] for an examination of this issue).

As for combination, a classical theory and a logic program should be viewed as complementary descriptions of the same domain. Therefore, a syntactic separation between predicates defined in these two components should not be enforced. Furthermore, it is desirable to neither restrict the interaction between the classical and the rules components nor impose any syntactic or semantic restrictions on the individual components. That is, the classical component may be an arbitrary theory Φ of some first-order language with equality, and the rules component may be an arbitrary non-ground normal or disjunctive logic program P, interpreted using, e.g., the common stable-model semantics [Gelfond and Lifschitz, 1988].

The goal is a combined theory, $\iota(\Phi, P)$, in a uniform logical formalism. Naturally, this theory should amount to Φ if P is empty, and to P if Φ is empty. Therefore, such a combination must provide faithful embeddings $\sigma(\Phi)$ and $\tau(P)$ of Φ and P, respectively, in this formalism, given by $\sigma(\Phi) = \iota(\Phi, \emptyset)$ and $\tau(P) = \iota(\emptyset, P)$, respectively. In turn, knowledge combination may be carried out on top of such embeddings $\sigma(\cdot)$ and $\tau(\cdot)$, where in the simplest case one may choose $\iota(\Phi, P) = \sigma(\Phi) \cup \tau(P)$.

This raises the questions (a) which uniform formalism is suitable and (b) which embeddings are suitable and, furthermore, how do embeddings relate to each other and how do they behave under knowledge combination?

Autoepistemic logic (AEL) [Moore, 1985], which extends classical logic with a modal belief operator, is an attractive candidate for a uniform formalism. In fact, embedding a classical theory in AEL is trivial, and several embeddings of logic programs in AEL have been described [Gelfond and Lifschitz, 1988; Marek and Truszczyński, 1993; Lifschitz and Schwarz, 1993; Chen, 1993; Przymusinski, 1991]. However, all these embeddings have been developed for the propositional case only, whereas we need to deal with non-ground theories and programs. This requires us to consider *first-order autoepistemic logic* (FO-AEL) [Konolige, 1991; Kaminski and Rey, 2002; Levesque and Lakemeyer, 2000], and non-ground versions of these embeddings. Our main contributions are as follows.

We define several embeddings of non-ground logic pro-

*This work was funded by the European Commission projects KnowledgeWeb (IST 507482), DIP (IST 507483), and REWERSE (IST 506779), by the Austrian Science Fund (FWF) project P17212-N04, and by the CICyT of Spain project TIC-2003-9001.

grams into FO-AEL, taking into account subtle issues of quantification in FO-AEL. We show that these embeddings are faithful in the sense that the stable models of the program and the sets of objective ground atoms in the stable expansions of the embeddings are in a one-to-one correspondence. However, the embeddings behave differently on formulas beyond ground atoms, and when combined with classical theories, even when considering propositional formulas.

Motivated by these differences, we compare the embeddings along two dimensions:

1. We determine correspondences between the stable expansions of different possible embeddings, with respect to various classes of formulas. This is done for the embeddings themselves, as well as for combinations with theories from different fragments of classical logic which are important in ontology representation.
2. We present inclusion relations between the sets of autoepistemic consequences of the embeddings.

Compared to other well-known nonmonotonic formalisms like Reiter's default logic, FO-AEL offers a uniform language in which (nonmonotonic) rules themselves can be expressed at the object level. This conforms with the idea of treating an ontology and a logic program together as a unified theory.

Arguably, none of the embeddings can a priori be considered to be superior to the others. Our results give useful insight into the properties of the different embeddings, both on its own right and for knowledge combination. They provide a helpful guidance for the selection of an embedding for a particular scenario.

Proofs of all results are available in an extended version of this paper.

2 Preliminaries

First-Order Logic A first-order (FO) language $\mathcal{L}$ consists of all formulas over a signature $\Sigma = (\mathcal{F}, \mathcal{P})$, where $\mathcal{F}$ and $\mathcal{P}$ are countable sets of *function* and *predicate symbols*, respectively. Function symbols with arity 0 are called *constants*. $\mathcal{V}$ is a countably infinite set of *variable symbols*. Terms and atomic formulas (atoms) are constructed as usual for first-order logic with equality. Ground terms are also called *names*; $\mathcal{N}_\Sigma$ is the set of names of a given signature Σ. Complex formulas are constructed as usual using the symbols $\neg$, $\wedge$, $\vee$, $\supset$, $\exists$, $\forall$,(, and). A sentence is a formula with no free variables. The universal closure of a formula ϕ is denoted by $\forall\phi$. $\mathcal{L}_g$ is the restriction of $\mathcal{L}$ to ground formulas; $\mathcal{L}_{ga}$ is the restriction of $\mathcal{L}_g$ to atomic formulas. An *FO theory* $\Phi \subseteq \mathcal{L}$ is a set of sentences.

An *interpretation* of a language $\mathcal{L}$ is a tuple $w = \langle U, \cdot^I \rangle$, where U is a nonempty set, called the *domain*, and $\cdot^I$ is a mapping which assigns a function $f^I : U^n \to U$ to every n-ary function symbol $f \in \mathcal{F}$ and a relation $p^I \subseteq U^n$ to every n-ary predicate symbol $p \in \mathcal{P}$. A *variable assignment* B for w is a mapping which assigns an element $x^B \in U$ to every variable $x \in \mathcal{V}$. The interpretation of a term t, denoted $t^{w,B}$, is defined as usual; if t is ground, we write t^w instead of $t^{w,B}$.

An individual k with at least one name $t \in \mathcal{N}$ such that $t^w = k$ is called a *named* individual, and *unnamed* otherwise. In case names are interpreted distinctly, the *unique-names assumption* applies. If, additionally, every individual is named, the *standard-names assumption* applies.

A *variable substitution* β is a set $\{x_1/t_1, ..., x_k/t_k\}$, where $x_1, ..., x_k$ are distinct variables and $t_1, ..., t_k$ are names. β is *total* if it contains some x/n for every variable $x \in \mathcal{V}$. Given variable assignment B and substitution β, if $\beta = \{x/t \mid x \in \mathcal{V}, t^w = x^B$, for some name $t\}$, then β is *associated with* B. The *application* of a variable substitution β to some term, formula, or theory, denoted by appending β to it, is defined as syntactical replacement, as usual. Clearly, if the unique-names assumption applies, each variable assignment has a unique associated substitution; if the standard-names assumption applies, each associated substitution is total.

Example 1. *Consider a language $\mathcal{L}$ with constants $\mathcal{F} = \{a, b, c\}$, and an interpretation $w = \langle U, \cdot^I \rangle$ with $U = \{k, l, m\}$ such that $a^w = k$, $b^w = l$, and $c^w = l$, and the variable assignment B: $x^B = k$, $y^B = l$, and $z^B = m$. B has two associated variable substitutions, $\beta_1 = \{x/a, y/b\}$ and $\beta_2 = \{x/a, y/c\}$, which are not total.*

Logic Programs A *disjunctive logic program* P consists of rules of the form

$$h_1 \mid \ldots \mid h_l \leftarrow b_1, \ldots, b_m, \, not\, c_1, \ldots, \, not\, c_n, \quad (1)$$

where $h_1, \ldots, h_l, b_1, \ldots, b_m, c_1, \ldots, c_n$ are (equality-free) atoms. $H(r) = \{h_1, \ldots, h_l\}$ is the set of *head atoms* of r, $B^+(r) = \{b_1, \ldots, b_m\}$ is the set of *positive body atoms* of r, and $B^-(r) = \{c_1, \ldots, c_n\}$ is the set of *negative body atoms* of r. If $l = 1$, then r is *normal*. If $B^-(r) = \emptyset$, then r is *positive*. If every variable in r occurs in $B^+(r)$, then r is *safe*. If every rule $r \in P$ is normal (resp., positive, safe), then P is normal (resp., positive, safe).

By a *first-order signature*, Σ_P, we understand a superset of the function and predicate symbols which occur in P. Let $\mathcal{L}_P$ denote the first-order language based on Σ_P. We assume that Σ_P contains at least one 0-ary function symbol or only 0-ary predicate symbols. The *Herbrand base*, B_H, of $\mathcal{L}_P$ is the set of ground atomic formulas of $\mathcal{L}_P$. Subsets of B_H are called *Herbrand interpretations*.

The *grounding* of a logic program P, denoted $gr(P)$, is the union of all possible ground instantiations of P, obtained by replacing each variable in a rule r with a name in $\mathcal{N}_{\Sigma_P}$, for each rule $r \in P$.

Let P be a positive program. A Herbrand interpretation M of P is a *model* of P if, for every rule $r \in gr(P)$, $B^+(r) \subseteq M$ implies $H(r) \cap M \neq \emptyset$. A Herbrand model M is *minimal* iff for every model M' such that $M' \subseteq M$, $M' = M$.

Following Gelfond and Lifschitz [1991], the *reduct* of a logic program P with respect to an interpretation M, denoted P^M, is obtained from $gr(P)$ by deleting (i) each rule r with $B^-(r) \cap M \neq \emptyset$, and (ii) *not* c from the body of every remaining rule r with $c \in B^-(r)$. If M is a minimal Herbrand model of P^M, then M is a *stable model* of P.

3 First-order Autoepistemic Logic

We adopt the definition of first-order autoepistemic logic (FO-AEL) under the any- and all-name semantics following

Konolige [1991], using a novel characterization with *associated variable substitutions*. The benefit of these semantics is that they allow quantification over arbitrary domains and generalize classical first-order logic with equality, thereby allowing a trivial embedding of first-order theories. Other approaches [Kaminski and Rey, 2002; Levesque and Lakemeyer, 2000] restrict the domains of interpretations to unique or standard names and therefore do not allow such direct embeddings.

An FO-AEL language $\mathcal{L}_{\mathsf{L}}$ is defined relative to a first-order language $\mathcal{L}$: (i) any atomic formula in $\mathcal{L}$ is a formula in $\mathcal{L}_{\mathsf{L}}$; (ii) if ϕ is a *formula* in $\mathcal{L}_{\mathsf{L}}$, then $\mathsf{L}\phi$, called a *modal atom*,[1] is a formula in $\mathcal{L}_{\mathsf{L}}$; and (iii) *complex formulas* are constructed as in first-order logic. A formula without modal atoms is an *objective* formula. *Standard autoepistemic logic* is FO-AEL without variables.

An *autoepistemic interpretation* is a pair $\langle w, \Gamma \rangle$, where $w = \langle U, \cdot^I \rangle$ is a first-order interpretation and $\Gamma \subseteq \mathcal{L}_{\mathsf{L}}$ is a set of sentences, called the *belief set*. Satisfaction of objective atomic formulas in w is as in first-order logic.

Satisfaction of a formula $\mathsf{L}\phi$ in an interpretation $\langle w, \Gamma \rangle$ with respect to a variable assignment B under the *any-name semantics* (resp., *all-name semantics*) is defined as follows:

> $w, B \models_\Gamma \mathsf{L}\phi$ iff, for some (resp., all) variable substitution(s) β, associated with B, $\phi\beta$ is closed and $\phi\beta \in \Gamma$.

This extends to complex formulas in the usual way. Notice that in case the unique-names assumption applies, the any- and all-name semantics coincide.

$\langle w, \Gamma \rangle$ is a *model* of ϕ, denoted $w \models_\Gamma \phi$, if $w, B \models_\Gamma \phi$ for every variable assignment B of w. This extends to sets of formulas in the usual way. A set $A \subseteq \mathcal{L}_{\mathsf{L}}$ of formulas *entails* a sentence ϕ with respect to a belief set Γ, denoted $A \models_\Gamma \phi$, if for every interpretation w such that $w \models_\Gamma A$, $w \models_\Gamma \phi$.

Example 2. *Consider a language with constant symbols a, b and unary predicate symbol p, and an interpretation $\langle w, \Gamma \rangle$ with $w = \langle \{k\}, \cdot^I \rangle$ and $\Gamma = \{p(a)\}$. Under the any-name semantics, $w \models_\Gamma \exists x.\mathsf{L}p(x)$; under the all-name semantics, $w \not\models_\Gamma \exists x.\mathsf{L}p(x)$, because $b^w = a^w = k$, but $p(b) \notin \Gamma$.*

We deem this behavior of the all-name semantics counterintuitive; so, following Konolige [1991], we use the any-name semantics in what follows, unless stated otherwise.

Example 3. *Consider the formula $\phi = \forall x(p(x) \supset \mathsf{L}p(x))$ and some interpretation $\langle w, \Gamma \rangle$. Then: $w \models_\Gamma \phi$ iff for every variable assignment B, $w, B \models_\Gamma p(x) \supset \mathsf{L}p(x)$ iff $w, B \not\models_\Gamma p(x)$ or $w, B \models_\Gamma \mathsf{L}p(x)$. Now, $w, B \models_\Gamma \mathsf{L}p(x)$, with $x^B = k$, iff for some $t \in \mathcal{N}_\Sigma$, $t^w = k$, and $p(t) \in \Gamma$. Thus, ϕ is false in any interpretation where p^I contains unnamed individuals.*

A belief set $T \subseteq \mathcal{L}_{\mathsf{L}}$ is a *stable expansion* of a base set $A \subseteq \mathcal{L}_{\mathsf{L}}$ iff $T = \{\phi \mid A \models_T \phi\}$. We use the following notation in the remainder: $T_o = T \cap \mathcal{L}$, $T_{og} = T \cap \mathcal{L}_g$, and $T_{oga} = T \cap \mathcal{L}_{ga}$.

A formula ϕ is an *autoepistemic consequence* of A if ϕ is included in every stable expansion of A. $Cons(A)$ denotes the set of all autoepistemic consequences of A. $Cons_o(A)$ denotes the restriction of $Cons(A)$ to objective formulas.

[1] $\mathsf{L}\phi$ is usually read as "ϕ is known" or "ϕ is believed."

Every stable expansion T fulfills the following properties: (a) T is closed under first-order entailment, (b) if $\phi \in T$ then $\mathsf{L}\phi \in T$, and (c) if $\phi \notin T$ then $\neg\mathsf{L}\phi \in T$. If T is consistent, the converses of (b) and (c) also hold.

Konolige [1991] shows that a stable expansion T of a base set A is determined by its objective subset T_o, called the *kernel* of T. If A does not have nested modal operators, then, additionally, $T_o = \{\phi \in \mathcal{L} \mid A \models_{T_o} \phi\}$ iff T_o is the kernel of a stable expansion T of A. We extend this result as follows:

Proposition 1. *Given a base set $A \subseteq \mathcal{L}_{\mathsf{L}}$ with only objective atomic formulas in the context of modal atoms, and a set of objective formulas $\Gamma_o \subseteq \mathcal{L}$, with $\Gamma_{ga} = \Gamma_o \cap \mathcal{L}_{ga}$, then $\Gamma_o = \{\phi \in \mathcal{L} \mid A \models_{\Gamma_{ga}} \phi\}$ iff $\Gamma_o = T \cap \mathcal{L}$ for some stable expansion T of A.*

4 Embedding Non-Ground Logic Programs

We define an embedding as a function which takes a logic program P as its argument and returns a set of sentences in the FO-AEL language obtained from Σ_P.

Since the unique-names assumption does not hold in FO-AEL in general, it is necessary to axiomatize default uniqueness of names (as introduced by Konolige [1991]). By UNA_Σ we denote the set of axioms

$$\neg\mathsf{L}(t_1 = t_2) \supset t_1 \neq t_2, \quad \text{for all distinct } t_1, t_2 \in \mathcal{N}_\Sigma.$$

4.1 Embedding Normal Logic Programs

The first embedding we consider is an extension of the one which originally led Gelfond and Lifschitz to the discovery of the stable model semantics [Gelfond and Lifschitz, 1988]. The second and third embedding are extensions of the embeddings due to Marek and Truszczyński [1993]. The third was independently developed by Lifschitz and Schwarz [1993], and Chen [1993]. The original motivation for the second and third embedding was the possibility to directly embed programs with strong negation and disjunctive programs.

Definition 1. *Let r be a rule of form* (1) *with $l = 1$. Then:*

$$\tau_{HP}(r) = \forall \textstyle\bigwedge_i b_i \wedge \bigwedge_j \neg\mathsf{L}c_j \supset h;$$
$$\tau_{EB}(r) = \forall \textstyle\bigwedge_i (b_i \wedge \mathsf{L}b_i) \wedge \bigwedge_j \neg\mathsf{L}c_j \supset h;$$
$$\tau_{EH}(r) = \forall \textstyle\bigwedge_i (b_i \wedge \mathsf{L}b_i) \wedge \bigwedge_j \neg\mathsf{L}c_j \supset h \wedge \mathsf{L}h.$$

For a normal logic program P, we define:

$$\tau_x(P) = \{\tau_x(r) \mid r \in P\} \cup UNA_{\Sigma_P}, \ \ x \in \{HP, EB, EH\}.$$

In the above embeddings, "HP" stands for "Horn for positive rules" (positive rules are translated to objective Horn clauses); "EB" stands for "epistemic rule bodies" (the body of a rule can only become true if it is *known* to be true); and "EH" stands for "epistemic rule heads" (if the body of a rule is true, the head is *known* to be true). For all three embeddings, we assume $\Sigma_{\tau_x(P)} = \Sigma_P$ (here and henceforth we use "x" as a meta-variable to range over HP, EB, and EH). Furthermore, by τ_x^- we denote the embedding τ_x *without* the UNA axioms. In the examples of embeddings in the remainder of the paper, we do not write the UNA axioms explicitly.

A notable distinction between the embedding τ_{HP} on the one hand and the embeddings τ_{EB}, τ_{EH} on the other is that the contrapositive of the rules in P is included in stable expansions of τ_{HP}, but not in stable expansions of τ_{EB}, τ_{EH}:

Example 4. *Consider $P = \{p \leftarrow q\}$. The stable expansion of $\tau_{HP}(P) = \{p \supset q\}$ includes $\neg q \supset \neg p$; the expansion of $\tau_{EB}(P) = \{p \wedge \mathsf{L}p \supset q\}$ includes $\neg q \supset \neg \mathsf{L}p \vee \neg p$, but not $\neg q \supset \neg p$, and neither does $\tau_{EH}(P)$.*

For the case of standard autoepistemic logic and ground logic programs, the following correspondence result holds:

Proposition 2 ([Gelfond and Lifschitz, 1988; Marek and Truszczyński, 1993]). *A Herbrand interpretation M is a stable model of a ground normal logic program P iff there is a consistent stable expansion T of $\tau_x^-(P)$ in standard autoepistemic logic such that $M = T \cap \mathcal{L}_{ga}$.*

Now consider the case of non-ground programs. The following example illustrates the embeddings:

Example 5. *Consider $P = \{q(a);\ p(x);\ r(x) \leftarrow \mathit{not}\ s(x), p(x)\}$, having a single stable model $M = \{q(a), p(a), r(a)\}$. Likewise, each of the embeddings $\tau_{HP}(P)$, $\tau_{EB}(P)$, and $\tau_{EH}(P)$ has a single consistent stable expansion:*

$$
\begin{aligned}
T^{HP} = \{&q(a), p(a), \mathsf{L}p(a), \neg \mathsf{L}s(a), r(a),\\
&\forall x(\neg p(x) \supset \neg q(x)), \neg \mathsf{L}\forall x(\mathsf{L}p(x)), ...\},\\
T^{EB} = \{&q(a), p(a), \mathsf{L}p(a), \neg \mathsf{L}s(a), r(a), \neg \mathsf{L}\forall x(\mathsf{L}p(x)), ...\},\\
T^{EH} = \{&q(a), p(a), \mathsf{L}p(a), \neg \mathsf{L}s(a), r(a), \forall x(\mathsf{L}p(x))...\}.
\end{aligned}
$$

The stable expansions in Example 5 agree on objective ground atoms, but not on arbitrary formulas. We now extend Proposition 2 to the non-ground case.

Lemma 1. *Given a set $A \subseteq \mathcal{L}_{ga}$ of objective ground atoms, there exists a stable expansion T of $\tau_x(P)$ under the any- or all-name semantics with $T_{oga} = A$ iff there exists a stable expansion T' of $\tau_x(gr(P))$ with $T'_{oga} = A$. Moreover, the same result holds for τ_{HP}^- under the all-name semantics.*

Theorem 1. *A Herbrand interpretation M of a normal logic program P is a stable model of P iff there is a consistent stable expansion T of $\tau_x(P)$ under the any- or all-name semantics such that $M = T \cap \mathcal{L}_{ga}$. Moreover, the same result holds for τ_{HP}^- under the all-name semantics.*

Note that this result does not hold for τ_{HP}^- under the any-name semantics. Consider $P = \{p(n_1); r(n_2); q \leftarrow \mathit{not}\ p(x)\}$ such that Σ_P has only two names, n_1 and n_2. P has one stable model, $M = \{p(n_1), r(n_2), q\}$. $\tau_{HP}^-(P) = \{p(n_1); r(n_2); \forall x(\neg \mathsf{L}p(x) \supset q)\}$ has one stable expansion, $T = \{p(n_1), r(n_2), \mathsf{L}p(n_1), \mathsf{L}r(n_2), \neg \mathsf{L}p(n_2), ...\}$. T does *not* include q. To see this, consider an interpretation w with only one individual k. $\mathsf{L}p(x)$ is trivially true under the any-name semantics, because there is some name for k such that $p(t) \in T$ (viz. $t = n_1$). In the all-name semantics, this situation does not occur, because for $\mathsf{L}p(x)$ to be true, $p(t)$ must be included in T for *every* name ($t = n_1$ and $t = n_2$) for k. One can similarly verify that the result does not apply to the embeddings τ_{EB}^- and τ_{EH}^- under the all-name semantics, by the positive modal atoms in the antecedents.

4.2 Embedding Disjunctive Logic Programs

The embeddings τ_{HP} and τ_{EB} cannot be straightforwardly extended to disjunctive logic programs, even in the propositional case. Consider the program $P = \{a \mid b \leftarrow\}$. P has two stable models: $M_1 = \{a\}$ and $M_2 = \{b\}$. However, a straightforward extension of τ_{HP}, $\tau_{HP}^{\vee}(P) = \{a \vee b\}$, has one stable expansion $T = \{a \vee b, \mathsf{L}(a \vee b), \neg \mathsf{L}a, \neg \mathsf{L}b, ...\}$. In contrast, τ_{EH} can be straightforwardly extended because of the modal atoms in the consequent of the implication: $\tau_{EH}^{\vee}(P) = \{(a \wedge \mathsf{L}a) \vee (b \wedge \mathsf{L}b)\}$ has two stable expansions $T_1 = \{a \vee b, a, \mathsf{L}a, \neg \mathsf{L}b, ...\}$ and $T_2 = \{a \vee b, b, \mathsf{L}b, \neg \mathsf{L}a, ...\}$.

The so-called *positive introspection axioms* (PIAs) [Przymusinski, 1991] remedy this situation for $\tau_{HP}^{\vee}$ and $\tau_{EB}^{\vee}$. Let PIA_{Σ} be the set of axioms

$$\alpha \supset \mathsf{L}\alpha, \quad \text{for every objective ground atom } \alpha \text{ of } \Sigma.$$

The PIA $\alpha \supset \mathsf{L}\alpha$ ensures that every consistent stable expansion contains either α or $\neg\alpha$.

It would have been possible to define the PIAs in a different way: $\forall\, \phi \supset \mathsf{L}\phi$ for any objective atomic formula ϕ. This would, however, effectively close the domain of the predicates in Σ (see Example 3). We deem this aspect undesirable in combinations with FO theories.

Definition 2. *Let r be a rule of form* (1). *Then:*

$$
\begin{aligned}
\tau_{HP}^{\vee}(r) &= \forall \textstyle\bigwedge_i b_i \wedge \bigwedge_j \neg \mathsf{L}c_j \supset \bigvee_k h_k;\\
\tau_{EB}^{\vee}(r) &= \forall \textstyle\bigwedge_i (b_i \wedge \mathsf{L}b_i) \wedge \bigwedge_j \neg \mathsf{L}c_j \supset \bigvee_k h_k;\\
\tau_{EH}^{\vee}(r) &= \forall \textstyle\bigwedge_i (b_i \wedge \mathsf{L}b_i) \wedge \bigwedge_j \neg \mathsf{L}c_j \supset \bigvee_k (h_k \wedge \mathsf{L}h_k).
\end{aligned}
$$

For a disjunctive logic program P, we define:

$$
\begin{aligned}
\tau_{HP}^{\vee}(P) &= \{\tau_{HP}^{\vee}(r) \mid r \in P\} \cup PIA_{\Sigma_P} \cup UNA_{\Sigma_P};\\
\tau_{EB}^{\vee}(P) &= \{\tau_{EB}^{\vee}(r) \mid r \in P\} \cup PIA_{\Sigma_P} \cup UNA_{\Sigma_P};\\
\tau_{EH}^{\vee}(P) &= \{\tau_{EH}^{\vee}(r) \mid r \in P\} \cup UNA_{\Sigma_P}.
\end{aligned}
$$

As before, by $\tau_x^{\vee -}$ we denote the embedding $\tau_x^{\vee}$ *without* the UNA part. We do not write the UNA and PIA parts explicitly in the examples below.

For the case of standard autoepistemic logic and ground disjunctive logic programs, the correspondence between the stable expansions of the embeddings $\tau_{HP}^{\vee}(P)$ and $\tau_{EH}^{\vee}(P)$ and the stable models of P is known:

Proposition 3 ([Przymusinski, 1991; Marek and Truszczyński, 1993]). *A Herbrand interpretation M of a ground disjunctive logic program P is a stable model of P iff there is a consistent stable expansion T of $\tau_{HP}^{\vee -}(P)$ (resp., $\tau_{EH}^{\vee -}(P)$) in standard autoepistemic logic such that $M = T \cap \mathcal{L}_{ga}$.*

We generalize this result to the case of FO-AEL and non-ground programs, and additionally for $\tau_{EB}^{\vee}$:

Theorem 2. *A Herbrand interpretation M of a disjunctive logic program P is a stable model of P iff there is a consistent stable expansion T of $\tau_x^{\vee}(P)$ under the any- or all-name semantics such that $M = T \cap \mathcal{L}_{ga}$. Moreover, the same result holds for τ_{HP}^- under the all-name semantics.*

A notable distinction between the embeddings $\tau_{HP}^{\vee}$ and $\tau_{EB}^{\vee}$ on the one hand and $\tau_{EH}^{\vee}$ on the other is the presence and absence of the PIAs, respectively:

Example 6. *Consider $P = \{p \mid q \leftarrow\}$, $\tau_{HP}^{\vee}(P) = \{p \vee q\} \cup PIA_{\Sigma_P}$, and $\tau_{EH}^{\vee}(P) = \{(p \wedge \mathsf{L}p) \vee (q \wedge \mathsf{L}q)\}$. The stable expansions of $\tau_{HP}^{\vee}(P)$ are $T_1^{HP} = \{p, \neg q, \mathsf{L}p, \neg \mathsf{L}q, ...\}$ and $T_2^{HP} = \{q, \neg p, \mathsf{L}p, \neg \mathsf{L}p, ...\}$; the expansions of $\tau_{EH}^{\vee}(P)$ are $T_1^{EH} = \{p, \mathsf{L}p, \neg \mathsf{L}q, ...\}$ and $T_2^{EH} = \{q, \mathsf{L}p, \neg \mathsf{L}p, ...\}$. The expansions T_1^{EH} and T_2^{EH} include neither $\neg q$ nor $\neg p$.*

$\Phi\backslash P$	$\mathcal{P}rg$	$\mathcal{S}afe$	$\mathcal{G}rnd$
$\mathcal{T}hr$	$\iota_{EH} \equiv \iota_{EH}^{\vee}$	$\iota_{EB} \equiv \iota_{EH}$	$\iota_{HP}^{\vee} \equiv \iota_{EB}^{\vee}$
$\mathcal{U}ni$	$\iota_{EB} \equiv_g \iota_{EH}$		
$g\mathcal{H}orn$			$\iota_{HP} \equiv_{ga} \iota_{EB}$
$\mathcal{H}orn$	$\iota_{HP} \equiv_{ga} \iota_{EB}$		
$\mathcal{P}rop$	$\iota_{HP}^{\vee} \equiv_g \iota_{EB}^{\vee}$		
$\{\emptyset\}$	$\iota_{HP} \equiv_{ga} \iota_{EB} \equiv_{ga}$ $\iota_{EH} \equiv_{ga} \iota_{HP}^{\vee} \equiv_{ga}$ $\iota_{EB}^{\vee} \equiv_{ga} \iota_{EH}^{\vee}$		

Table 1: Correspondence between expansions of combinations; $\iota_x^{(\vee)}$ is short for $\iota_x^{(\vee)}(\Phi, P)$.

5 Relations between the Embeddings

In this section, we explore correspondences between the embeddings presented in the previous section in combinations with FO theories. In our simple setting, we define the combination of a program P and an FO theory Φ as

$$\iota_x^{(\vee)}(\Phi, P) = \Phi \cup \tau_x^{(\vee)}(P) \subseteq \mathcal{L}_{\mathsf{L}},^2$$

where $\Sigma_{\mathcal{L}_{\mathsf{L}}}$ is the union of the signatures Σ_Φ and Σ_P. Recall that we consider the any-name semantics, because of its more intuitive behavior (cf. 3).

In the following, we compare (i) the stable expansions of such combinations and (ii) the sets of autoepistemic consequences of the individual embeddings. To this end, we introduce the following notation:

Let A_1 and A_2 be FO-AEL theories. We write $A_1 \equiv A_2$ iff A_1 and A_2 have the same stable expansions. Moreover, for $\alpha \in \{g, ga\}$, we write $A_1 \equiv_\alpha A_2$ iff

$$\{T \cap \mathcal{L}_\alpha \mid T \text{ is a stable expansion of } A_1\} = \{T' \cap \mathcal{L}_\alpha \mid T' \text{ is a stable expansion of } A_2\}.$$

Note that, by definition, $A_1 \equiv A_2$ implies $A_1 \equiv_g A_2$, and $A_1 \equiv_g A_2$ implies $A_1 \equiv_{ga} A_2$.

In our analysis, we furthermore use the following classes of programs and theories:

- the classes $\mathcal{P}rg$, $\mathcal{S}afe$, and $\mathcal{G}rnd$ of arbitrary, safe, and ground logic programs, respectively; and
- the classes $\mathcal{T}hr$, $\mathcal{U}ni$, $g\mathcal{H}orn$, $\mathcal{H}orn$, $\mathcal{P}rop$, and $\{\emptyset\}$ of arbitrary, universal, generalized Horn,[3] Horn, propositional, and empty FO theories.

Observe the following inclusions:

$$\mathcal{G}rnd \subset \mathcal{S}afe \subset \mathcal{P}rg; \quad \{\emptyset\} \subset \left| \begin{array}{c} \mathcal{P}rop \subset \mathcal{U}ni \\ \mathcal{H}orn \subset \mathcal{U}ni \\ \mathcal{H}orn \subset g\mathcal{H}orn \end{array} \right| \subset \mathcal{T}hr.$$

Theorem 3. *Let P be a normal (disjunctive, resp.) logic program and Φ be a first-order theory. Then, the relations depicted in Table 1 (with the respective provisos) hold, providing P and Φ belong to the classes listed there.*

[2]One could imagine other, non-trivial, embeddings of the classical theory. Such embeddings are a topic for future investigations.

[3]Generalized Horn formulas are Horn formulas which additionally allow existentially quantified variables in the consequent of the material implication.

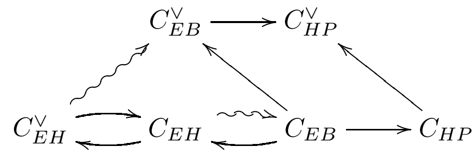

Figure 1: Relationships between sets of consequences; $C_x^{(\vee)}$ stands for $Cons_o(\tau_x^{(\vee)}(P))$, $\rightarrow$ stands for $\subseteq$, and $\rightsquigarrow$ stands for $\subseteq$ in case P is safe.

Consider the logic program P from Example 5. P is neither safe nor ground: to determine correspondence between embeddings, we need to use the first column of Table 1. Since P is normal, all equations in this column are applicable. We have that $\tau_{EB}(P) \equiv_g \tau_{EH}(P)$ and $\tau_{HP}^{\vee}(P) \equiv_g \tau_{EB}^{\vee}(P)$. Let Φ be a Horn theory, then $\iota_{HP}(\Phi, P) \equiv_{ga} \iota_{EH}(\Phi, P) \equiv_{ga} \iota_{EB}(\Phi, P)$ and $\iota_{EH}(\Phi, P) \equiv \iota_{EH}^{\vee}(\Phi, P)$.

Additionally, since autoepistemic consequence is defined through the intersection of all stable expansions, we can conclude that $\tau_{EB}(P)$ and $\tau_{EH}(P)$, and also $\tau_{HP}^{\vee}(P)$ and $\tau_{EB}^{\vee}(P)$, agree on objective ground autoepistemic consequence and that $\iota_{HP}(\Phi, P)$, $\iota_{EH}(\Phi, P)$, and $\iota_{EB}(\Phi, P)$ agree on objective ground atomic autoepistemic consequence.

We now consider the relative behavior of the embeddings with respect to autoepistemic consequences.

Theorem 4. *Let P be a (safe) normal (disjunctive, resp.) logic program, and let $\tau_x^{(\vee)}$ and $\tau_y^{(\vee)}$ be embedding functions, for $x, y \in \{HP, EB, EH\}$. Then, relations $Cons_o(\tau_x^{(\vee)}(P)) \subseteq Cons_o(\tau_y^{(\vee)}(P))$ hold as depicted in Figure 1 (with the respective provisos).*

Most of the relations given in Figure 1 do not hold for combinations with FO theories. Consider, e.g., $P = \{r \leftarrow not\, p, not\, q\}$ and $\Phi = \{p \vee q\}$. Then, $\tau_{HP}(P) = \{\neg \mathsf{L}p \wedge \neg \mathsf{L}q \supset r\}$ and $\tau_{HP}^{\vee}(P) = \{\neg \mathsf{L}p \wedge \neg \mathsf{L}q \supset r\} \cup PIA_{\Sigma_P}$ both have one stable expansion, each containing $\neg \mathsf{L}p$, $\neg \mathsf{L}q$, and r. The combination $\tau_{HP}(P) \cup \Phi$ has one stable expansion which includes $\neg \mathsf{L}p$, $\neg \mathsf{L}q$, and r; $\tau_{HP}^{\vee}(P) \cup \Phi$ has two stable expansions $\{p, \mathsf{L}p, \neg \mathsf{L}q, ...\}$ and $\{q, \mathsf{L}q, \neg \mathsf{L}p, ...\}$, neither of which includes r. Thus, r is an autoepistemic consequence of $\iota_{HP}(\Phi, P)$, but not of $\iota_{HP}^{\vee}(\Phi, P)$. Therefore, $Cons_o(\iota_{HP}(\Phi, P)) \not\subseteq Cons_o(\iota_{HP}^{\vee}(\Phi, P))$.

Using the results in this section, we can make a number of observations about the embeddings:

(1) Few correspondences between embeddings with PIAs and those without hold. However, we can note that the former are stronger in terms of the number of objective autoepistemic consequences (cf. Figure 1 and Example 6).

(2) The embeddings τ_{HP} and $\tau_{HP}^{\vee}$ are generally the strongest in terms of consequences (see Figure 1). They allow to derive the contrapositive of rules (cf. Example 4) and the bodies of rules are applicable to unnamed individuals, whereas the antecedents of the axioms in the other embeddings are only applicable to named individuals, because of the positive modal atoms in the bodies.

(3) For unsafe programs, the embeddings τ_{EH} and $\tau_{EH}^{\vee}$ are generally not comparable with the others; embeddings of un-

safe rules result in axioms of form $\forall x \mathrm{L}p(x)$ (cf. Example 5), which require all individuals to be named.
(4) In case the programs are safe, or one assumes that all individuals are named, τ_{EB} and τ_{EH} coincide.

We conclude this section with an example which demonstrates possibly unexpected effects of the UNA axioms in their interaction with an FO theory.

Example 7. *Consider $P = \{p(a); p(b)\}$ and $\Phi = \{a \neq b \supset r\}$. Then, r is included in any stable expansion of $\Phi \cup \tau_x(P)$, for any τ_x, in view of the UNA axioms.*

6 Related and Future Work

In this paper, we have studied the combination of logic programs and ontologies (FO theories) using embeddings in a unifying formalism (FO-AEL). One could imagine, in contrast, extensions of semantics for logic programs or ontologies to incorporate (parts of) the other formalism. One such extension of logic programming semantics is that of open domains [Gelfond and Przymusinska, 1993; Van Belleghem *et al.*, 1997; Heymans *et al.*, 2005]. Such extended semantics can be used to accommodate incomplete knowledge, an important aspect of ontology languages. Nonmonotonic extensions of description logics (an FO-based formalism suitable for ontologies) have been presented in the literature [Baader and Hollunder, 1995; Donini *et al.*, 2002; Bonatti *et al.*, 2006]. Such approaches might be extended to accommodate logic programs.

We have investigated basic correspondences between different embeddings of non-ground programs in FO-AEL, and simple combinations with FO theories. Choosing different embeddings for logic programs, but also possibly different embeddings for first-order theories, will give rise to different properties of such combinations [de Bruijn *et al.*, 2006]. In future work, we will investigate these properties, as well as the relationship with existing approaches to combine logic programs and classical theories [Horrocks and Patel-Schneider, 2004; Eiter *et al.*, 2004; Rosati, 2006].

So far, we have only considered equality-free logic programs. We conjecture that equality in rule bodies poses no problems, since still only the trivial equalities are derivable. Allowing equality in rule heads is a topic for further research.

We expect that the proposed combinations of rules and ontologies based on FO-AEL will give rise to the definition of novel decidable fragments and for sound (but possibly incomplete) algorithms for specific reasoning tasks for such combinations. Additionally, we will consider other nonmonotonic logics (e.g., default logic and circumscription) as formalisms for combining logic programs and classical knowledge bases.

References

[Baader and Hollunder, 1995] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reas.*, 14:149–180, 1995.

[Bonatti *et al.*, 2006] P. Bonatti, C. Lutz, and F. Wolter. Expressive non-monotonic description logics based on circumscription. In *Proc. KR 2006*.

[de Bruijn *et al.*, 2006] J. de Bruijn, T. Eiter, A. Polleres, H. Tompits. On representational issues about combinations of classical theories with nonmonotonic rules. In *Proc. KSEM 2006*.

[Chen, 1993] J. Chen. Minimal knowledge + negation as failure = only knowing (sometimes). In *Proc. LPNMR'93*.

[Donini *et al.*, 2002] F. M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *ACM ToCL*, 3(2):177–225, 2002.

[Eiter *et al.*, 2004] T. Eiter, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining answer set programming with description logics for the semantic web. In *Proc. KR 2004*.

[Gelfond and Lifschitz, 1988] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proc. ICLP'88*.

[Gelfond and Lifschitz, 1991] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Gen. Computing*, 9(3/4):365–386, 1991.

[Gelfond and Przymusinska, 1993] M. Gelfond and H. Przymusinska. Reasoning on open domains. In *Proc. LPNMR'93*.

[Heymans *et al.*, 2005] S. Heymans, D. Van Nieuwenborgh, and D. Vermeir. Guarded Open Answer Set Programming. In *Proc. LPNMR 2005*.

[Horrocks and Patel-Schneider, 2004] I. Horrocks and P. F. Patel-Schneider. A proposal for an OWL rules language. In *Proc. WWW 2004*.

[Kaminski and Rey, 2002] M. Kaminski and G. Rey. Revisiting quantification in autoepistemic logic. *ACM ToCL*, 3(4):542–561, 2002.

[Konolige, 1991] K. Konolige. Quantification in autoepistemic logic. *Fund. Informaticae*, 15(3–4):275–300, 1991.

[Levesque and Lakemeyer, 2000] H. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2000.

[Lifschitz and Schwarz, 1993] V. Lifschitz and G. Schwarz. Extended logic programs as autoepistemic theories. In *Proc. LPNMR'93*.

[Marek and Truszczyński, 1993] V. Marek and M. Truszczyński. Reflexive autoepistemic logic and logic programming. In *Proc. LPNMR'93*.

[Moore, 1985] R. C. Moore. Semantical considerations on nonmonotonic logic. *Art. Intell.*, 25(1):75–94, 1985.

[Przymusinski, 1991] T. Przymusinski. Stable semantics for disjunctive programs. *New Gen. Comp.*, 9(3–4):401–424.

[Rosati, 2006] R. Rosati. $\mathcal{DL}$+log: Tight integration of description logics and disjunctive datalog. In *Proc. KR 2006*.

[Van Belleghem *et al.*, 1997] K. van Belleghem, M. Denecker, and D. De Schreye. A strong correspondence between description logics and open logic programming. In *Proc. ICLP'97*.