Normal form results for default logic

V. Wiktor Marek
Mirosław Truszczyński
Department of Computer Science
University of Kentucky
Lexington, KY 40506–0027
marek@ms.uky.edu, mirek@ms.uky.edu

Abstract In this paper we continue investigations of proof theory of default logic. It turns out that, similarly to classical logic, default theories can be represented in normal forms.

1 Introduction and preliminaries

In this paper we develop a representation theory for default logic of Reiter ([Rei80]). The question is whether one can find "normal forms" for default theories, that is, if there are syntactical constraints which can be imposed on default theories without changing extensions. In this section we introduce basic definitions and recall fundamental concepts of default logic. We introduce the notion of representability of default theories in Section 2 and prove a number of results of both positive and negative nature. A weaker notion, semi-representability, is studied in Section 3. We prove that with this weaker notion we can represent every default theory by a default theory with all rules either monotonic (that is justification-free) or semi-normal ([Eth88]). In Section 4 we discuss another structure associated with default logic, a weak extension. We show that every finite family of finitely generated theories can be represented as a family of weak extensions of a suitably constructed default theory. A result on autoepistemic expansions is given as a corollary.

We use standard logical notation. The reader is referred to [Fit90] and [Men64] for the unexplained concepts. The presentation of default logic follows one we gave in [MT93].

Definition 1.1 (Basic definitions)

1. A *default rule* is a triple $r =< \alpha, L, \omega >$ where α, ω are sentences of a propositional language $\mathcal{L}$, and $L = \{\beta_1, \ldots, \beta_m\}$ a finite set of formulas of $\mathcal{L}$. Such rule is usually written as $\frac{\alpha:\beta_1,\ldots,\beta_m}{\omega}$. Elements of L are called *justifications* of r. When a rule has no justifications we call r *monotonic*.

2. The formula ω is called a *consequent* of r. The set of consequents of rules in D is denoted by $CONS(D)$.

3. A *default theory* is a pair (D, W) where D is a set of default rules, and $W \subseteq \mathcal{L}$.

4. A default rule $r =< \alpha, L, \omega >$ is *normal* if $L = \{\omega\}$. That is, a normal rule is of the form $\frac{\alpha:\omega}{\omega}$.

5. A *normal* theory is a theory (D, W) where D consists of normal rules only.

6. A *semi-normal* default rule is a rule of the form $\frac{\alpha:\gamma\wedge\omega}{\omega}$.

7. A *semi-normal* theory is a theory (D, W) in which every default is semi-normal.

8. A *weakly semi-normal* theory is a default theory (D, W) such that every rule in D is either monotonic or semi-normal.

Definition 1.2 (S-derivations, extensions for default theories)

1. Let $S \subseteq \mathcal{L}$. Let (D, W) be a default theory. An S-derivation of a formula φ in (D, W) is any finite sequence $< \psi_1, \ldots, \psi_n >$ such that $\psi_n = \varphi$ and for all $i \leq n$ at least one of the following holds:

   (a) ψ_i is a tautology or

   (b) $\psi_i \in W$ or

   (c) ψ_i is the result of applying modus ponens to some ψ_j, ψ_k, $j, k < i$ or

   (d) There is a rule $\frac{\alpha:\beta_1,\ldots,\beta_m}{\psi_i} \in D$ such that $\alpha = \psi_j$ for some $j < i$ and $\neg\beta_1, \ldots, \neg\beta_m \notin S$.

2. $Cn^D_S(W)$ is the set of all formulas possessing an S derivation in (D, W).

3. S is called an *extension* of (D, W) if $S = Cn^D_S(W)$.

Definition 1.3 (Generating defaults, weak extensions)

1. Let $S \subseteq \mathcal{L}$. A default rule $r = \frac{\alpha:\beta_1,\ldots,\beta_m}{\omega}$ is called a *generating default* for S if $\alpha \in S$ and $\neg\beta_1, \ldots, \neg\beta_m \notin S$.

2. If $S \subseteq \mathcal{L}$ and (D, W) is a default theory then S is called a *weak extension* of (D, W) if $S = Cn(W \cup C)$ where C consists of consequents of generating defaults for S.

Theorem 1.4 (Reiter) *Every extension is a weak extension. In particular, if S is an extension of (D, W) then for some $C \subseteq CONS(D)$, $S = Cn(W \cup C$.*

Theorem 1.5 (Reiter) *1. If S_1, S_2 are two different extensions of a default theory (D, W), then $S_1 \subseteq S_2$ implies that $S_1 = S_2$.*

2. If S_1, S_2 are two different extensions of a normal theory (D, W), then S_1, S_2 are incompatible. That is $S_1 \cup S_2$ is inconsistent.

Definition 1.6 (Operator associated with default theories) Given a default theory (D, W) and a theory $S \subseteq \mathcal{L}$, the operator R^D_S is the mapping from subsets of $\mathcal{L}$ to subsets of $\mathcal{L}$ defined by:

$$R^D_S(T) = Cn(T \cup \{\omega : \quad \textit{For some } r \in D, r = \frac{\alpha : \beta_1, \ldots, \beta_m}{\omega} \textit{ and } \alpha \in T, \neg\beta_1, \ldots, \neg\beta_m \notin S\}$$

Proposition 1.7 (A characterization of extensions) *A theory $S \subseteq \mathcal{L}$ is an extension of (D, W) if and only if S is the least fixpoint of R^D_S over W.*

Definition 1.8 (Double disjunctive normal form) A formula φ is in *double disjunctive form* if φ is of the form $\bigvee \psi_i$ where each ψ_i is the negation of a clause.

Proposition 1.9 *For every formula Θ there is a formula φ such that φ is in double disjunctive form, and $\Theta \equiv \varphi$ is a tautology.*

2 Representation of Default Theories

One way of looking at the normal form theorems in propositional logic is this. Assign to a formula the set of all its models. It is quite obvious that $\varphi \equiv \psi$ is a tautology if and only if φ and ψ possess exactly the same models. Then the normal form theorems say that for every formula φ there exists a formula ψ in, say, disjunctive normal form with precisely the same set of models. We lift this way of looking at normal form to the present context. To this end we introduce the notion of *equivalent* default theories and the associated notion of representability.

Definition 2.1 (Representability)

1. Let (D, W), (D', W') be default theories. We say that (D, W) is *equivalent* to (D', W') (in symbols $(D, W) \approx (D', W')$) if (D, W) and (D', W') have exactly the same extensions.

2. Let $\mathcal{R}$ be a class of default theories, let (D, W) be a default theory. We say that (D, W) is *representable* in $\mathcal{R}$ if there is $(D', W') \in \mathcal{R}$ such that (D, W) is equivalent to (D', W').

3. A class $\mathcal{R}$ of default theories *represents* default logic if every default theory is representable in $\mathcal{R}$.

It is quite obvious that $\approx$ is an equivalence relation. This, in turn, implies that if (D, W) is representable in $\mathcal{R}$ then every theory equivalent to (D, W) is also representable in $\mathcal{R}$.

Example 2.1 Let $W_1 = \{p\}$, $D_1 = \{\frac{p:\neg r}{q}, \frac{q:s}{t}\}$, $W_2 = \{p, q\}, D_2 = \{\frac{q:\neg u}{t}\}$. Then theories $(D_1, W_1), (D_2, W_2)$ are equivalent.

We first prove a negative result. Let $\mathcal{N}$ be a class of all normal default theories. One can ask if $\mathcal{N}$ represents default logic. The answer to this question is negative. We base it on Theorem 1.5 which says that not only are extensions of a normal default theory inclusion-incompatible, but they are, actually, pairwise inconsistent.

Proposition 2.2 *The class $\mathcal{N}$ of all normal default theories does not represent all default theories.*

Proof: Let $p, q \in At$, $p \neq q$, $D = \{\frac{:\neg p}{q}, \frac{:\neg q}{p}\}$, $W = \emptyset$. Theory (D, W) possesses two extensions $Cn(\{p\})$, and $Cn(\{q\})$. The theory $Cn(\{p\}) \cup Cn(\{q\})$ is consistent. Therefore (D, W) is not $\mathcal{N}$-representable.

The proof of Proposition 2.2 tells us that the reason for non-representability of (D, W) in $\mathcal{N}$ is that two of its extensions are consistent each with the other. In a sense it is the only reason.

Definition 2.3 A theory T is *finitely generated* if there exists a finite set of sentences $\{\varphi_1, \ldots, \varphi_k\}$ such that $Cn(T) = Cn(\{\varphi_1, \ldots, \varphi_k\})$.

We will now prove the following result.

Proposition 2.4 *Let $\langle T_\alpha : \alpha < \gamma \rangle$ be a collection of pairwise inconsistent, finitely generated consistent theories (γ can be finite or infinite). Then there exists a normal default theory (D, W) such that $\langle T_\alpha : \alpha < \gamma \rangle$ are precisely the extensions of (D, W).*

Proof: Let $\{\varphi_{\alpha,0}, \ldots \varphi_{\alpha,k_\alpha}\}$ be generators of theory T_α. Let $\psi_\alpha = \varphi_{\alpha,0} \wedge \ldots \wedge \varphi_{\alpha,k_\alpha}$. Then $T_\alpha = Cn(\{\psi_\alpha\})$. Our assumption is that for $\alpha \neq \beta$, $T_\alpha \cup T_\beta$ is an inconsistent theory, that is $\{\psi_\alpha \wedge \psi_\beta\}$ is an inconsistent theory.

Let $d_\alpha = \frac{:\psi_\alpha}{\psi_\alpha}$ for $\alpha < \gamma$. Let (D, W) be a default theory where $D = \{d_\alpha : \alpha < \gamma\}$, $W = \emptyset$. It is clear that (D, W) is a normal default theory. We claim that T_α, $\alpha < \gamma$, are precisely extensions of (D, W). By Theorem 1.4 the only candidates for extensions of (D, W) are of the form $T = Cn(\{\psi_{\alpha_0}, \ldots \psi_{\alpha_\xi} \ldots : \xi < \eta\})$. But if $\eta > 1$ that is if more than one generator is selected, then such theory T is inconsistent. But then $Cn^D_T(W) = Cn(\emptyset) \neq T$. Thus the only possible candidates for extensions are T_α's. Since ψ_α, ψ_β are pairwise inconsistent therefore, for $\alpha \neq \beta$, $\neg\psi_\alpha \in T_\beta$. Therefore for every $\alpha < \gamma$, the only rule applicable with respect to T_α is d_α. But then:

$$Cn^D_{T_\alpha}(W) = Cn(\{\psi_\alpha\}) = T_\alpha.$$

Thus T_α's are precisely extensions of (D, W).

Corollary 2.5 *If all extensions of (D, W) are finitely generated, consistent, and pairwise inconsistent, then (D, W) is representable in $\mathcal{N}$.*

One can read Proposition 2.4 and similar results as saying that once the desired images of the world are known then it is "easy" to design a default theory determining them. Of course it is precisely what does *not* happen. In practice the process is different; we usually have a collection of facts and rules and then look for images of reality, not conversely.

We have a simple but useful property of the relation $\approx$.

Proposition 2.6 *(a) If $Cn(W_1) = Cn(W_2)$ then for every D, $(D, W_1) \approx (D, W_2)$.*
(b) If $W \vdash \alpha \equiv \alpha'$, $W \vdash \psi_1 \equiv \theta_1, \ldots, W \vdash \psi_k \equiv \theta_k$, $W \vdash \gamma \equiv \gamma'$, then for every D and W, $(D \cup \{\frac{\alpha:\psi_1,\ldots\psi_k}{\gamma}\}, W) \approx (D \cup \{\frac{\alpha':\theta_1,\ldots\theta_k}{\gamma'}\}, W)$.

Proof: In both (a) and (b) we prove that for every $S \subseteq \mathcal{L}$, formulas S-provable in each of the two default theories are precisely the same. This implies that these theories have precisely same extensions.

We will prove now a simple representability result. Given $W \subseteq \mathcal{L}$, let D_W be the following collection of rules:

$$D_W = \{\frac{:}{\varphi} : \varphi \in W\}$$

That is, every $\varphi \in W$ is transformed to a rule r_φ which has no prerequisite and no justifications and has φ as the consequent.

Proposition 2.7 *Let $T \subseteq \mathcal{L}$. Then T is an extension of (D, W) if and only if T is an extension of $(D \cup D_W, \emptyset)$.*

Proof: Let $R_T^{D \cup D_W} \uparrow n(\emptyset)$ denote the result of the n-fold iteration of the operator $R_T^{D \cup D_W}$ on the empty set. By induction on n we prove that for all n, $R_T^{D \cup D_W} \uparrow n(\emptyset) \subseteq R_T^D \uparrow n+1(W)$ and $R_T^D \uparrow n(W) \subseteq R_T^{D \cup D_W} \uparrow n+1(\emptyset)$. This implies that the least fixpoint of $R_T^D(\cdot)$ above W and the least fixpoint of $R_T^{D \cup D_W}(\cdot)$ are identical. Therefore extensions of (D, W) and of $(D \cup D_W, \emptyset)$ coincide.

Corollary 2.8 *Every default theory is representable in the class of default theories (D, W) with $W = \emptyset$.*

The practical role of Corollary 2.8 is restricted. The reason is that often we think about default rules as the part of the reasoning mechanism that does not change (or at least changes rarely), and of W as the data, that can change frequently. In this setting Proposition 2.7 says that in case of data that changes very rarely, we can build it into the reasoning system.

We prove now two technical results. They play the role of distributive laws and will be used below to prove a subtler normal form result.

Proposition 2.9 *Let $d = \frac{\varphi:\beta_1,\ldots,\beta_k}{\gamma_1 \wedge \gamma_2}$ be a default rule and let $d_1 = \frac{\varphi:\beta_1,\ldots,\beta_k}{\gamma_1}$ $d_2 = \frac{\varphi:\beta_1,\ldots,\beta_k}{\gamma_2}$ be two rules arising from d by weakening the consequent of r. Then for every set of rules D and every $W \subseteq \mathcal{L}$, $(D \cup \{d\}, W) \approx (D \cup \{d_1, d_2\}, W)$.*

Proof: Notice that for every context $S \subseteq \mathcal{L}$, d is applicable with respect to s if and only if both d_1 and d_2 are applicable with respect to S. All three of these rules possess the same prerequisite. This implies that any time d is used in an S-derivation, we can use d_1 and d_2, and vice versa. Therefore for any formula ψ, any S proof of ψ from W using rules from $D \cup \{d\}$ can be transformed to an S-proof from W using rules from $D \cup \{d_1, d_2\}$ and conversely. This means that $Cn_S^{D\cup\{d\}}(W) = Cn_S^{D\cup\{d_1,d_2\}}(W)$. Hence the default extensions of both theories $(D \cup \{d\}, W)$ and $(D \cup \{d_1, d_2\}, W)$ are the same.

In Proposition 2.9 we proved distributivity of consequents with respect to conjunction. Next, we prove the distributivity of justifications with respect to alternative.

Proposition 2.10 *Let* $d = \frac{\varphi:(\delta_1 \vee \delta_2),\ldots,\beta_k}{\gamma}$ *be a default rule and let* $d_1 = \frac{\varphi:\delta_1,\ldots,\beta_k}{\gamma}$ $d_2 = \frac{\varphi:\delta_2,\ldots,\beta_k}{\gamma}$ *be two rules arising from d by strengthening justifications of d. Then for every set of rules D and every* $W \subseteq \mathcal{L}$, $(D \cup \{d\}, W) \approx (D \cup \{d_1, d_2\}, W)$.

Proof If S is closed under consequence then $\neg(\delta_1 \vee \delta_2) \notin S$ if and only if $\neg\delta_1 \notin S$ or $\neg\delta_2 \notin S$. Indeed, $\neg(\delta_1 \vee \delta_2)$ is equivalent to $\neg\delta_1 \wedge \neg\delta_2$ and the equivalence follows.

Next, observe that all the rules d, d_1, and d_2 possess the same prerequisite (premise) and the same consequent. This means that for every formula ψ an S-derivation of ψ from W using rules from $D \cup \{d\}$ *is*, actually an S-derivation of ψ from W using rules from $D \cup \{d_1, d_2\}$. Indeed, the sequence itself does not change, only the reason why a formula is put in that sequence change! Similarly, every S-proof of ψ using rules of $D \cup \{d_1, d_2\}$ is an S-derivation of ψ using rules in $D \cup \{d\}$. As in the proof of 2.9 this implies that $Cn_S^{D\cup\{d\}}(W) = Cn_S^{D\cup\{d_1,d_2\}}(W)$. Hence the default extensions of both theories $(D \cup \{d\}, W)$ and $(D \cup \{d_1, d_2\}, W)$ are the same.

Recall, that a clause is a formula of the form $s_1 \vee \ldots \vee s_k$ where $s_1 \ldots s_k$ are literals, that is atoms or negated atoms. *Clause* is the set of all clauses of the language $\mathcal{L}$.

We will present now a much deeper representability result (due to Yang and others [YBB92]). First, recall that every propositional theory T has precisely the same consequences as the theory $T' = Cn(T) \cap Clause$. That is, every theory is faithfully represented by a set of clauses. We show an analogous result for default logic.

Definition 2.11 1. A default rule $d = \frac{\varphi:\beta_1,\ldots,\beta_k}{\gamma}$ is called *clausal* if:

   (a) φ is a conjunction of clauses.

   (b) Each β_j is negation of a clause.

   (c) γ is a clause.

2. A default theory (D, W) is *clausal* if W consists of clauses, and every rule in D is clausal.

Theorem 2.12 (Yang, [YBB92]) *For every default theory (D, W) there exists a clausal default theory (D', W') such that $(D, W) \approx (D', W')$.*

Proof: We construct our desired theory (D', W) in three stages, enforcing, respectively, conditions (a), (b), and (c). By Proposition 2.6 we can, of course, assume that W consists of clauses.

First, for every default rule $r = \frac{\varphi:\beta_1,\ldots,\beta_k}{\gamma}$ in D consider any conjunctive normal form φ' of φ. Substitute the default rule $r' = \frac{\varphi':\beta_1,\ldots,\beta_k}{\gamma}$ for r. A direct application of Proposition 2.6 shows that the resulting theory $(D_1, W) = (\{r' : r \in D\}, W)$ has precisely the same extensions as (D, W).

Next, again by Proposition 2.6 we can assume that every consequent γ of a default rule in D_1 is in a conjunctive normal form. Then the repeated application of Proposition 2.9 produces a theory (D_2, W) such that $(D_2, W) \approx (D_1, W)$ (hence $(D_2, W) \approx (D, W)$) and the consequents of rules in D_2 are clauses. In this fashion both the conditions (a) and (c) are enforced.

To make sure that the condition (b) is satisfied we use Proposition 1.9 and Proposition 2.6. Thus each β_i can be assumed to be in double disjunctive normal form, that is $\bigvee \theta_j$ where each θ_j is a negated clause. Now we apply repeatedly Proposition 2.10. In this fashion we get a set D' of clausal rules such that $(D', W) \approx (D, W)$.

Notice that if each β_i has in its double disjunctive form s_i terms then one default rule d produces $s_1 \times \ldots \times s_k$ rules. Hence the size of D' may be large relative to size of D.

Example 2.2 Let $W = \{p\}$, $D = \{\frac{p:(r\vee s),(\neg u\vee w)}{q}, \frac{p\wedge q:r,(w\vee t)}{s}\}$. Then $W' = W$, and the new family of default rules is:

$D' = \{\frac{p:r,\neg u}{q}, \frac{p:r,w}{q}, \frac{p:s,\neg u}{q}, \frac{p:s,w}{q}, \frac{p\wedge q:r,w}{s}, \frac{p\wedge q:r,t}{s}\}$.
Theory (D', W') is clausal. By Theorem 2.12 theories (D, W) and (D', W') possess exactly the same extensions.

Theorem 2.12 yields a theory that consists of simpler objects (clausal default rules), but may be much larger than the original theory (D, W). In a sense the situation is similar to one we get when we represent a theory $T \subseteq \mathcal{L}$ by the set of clauses. But here we have an additional step: after all formulas α, β_j, γ are put in disjunctive normal form, the theory D' is obtained in the process of splitting conclusions and justifications that once again increases the size considerably.

Another fact worth mentioning is that the transformation described in Theorem 2.12 does not preserve normality. That is, if (D, W) is a normal default theory, and D' is computed as above, then (D', W) usually will not be normal.

3 Semi-representability of default theories

We will prove now another representability result. This time, however, the language in which we will construct suitable default theories will be an extension of the language $\mathcal{L}$ by means of new constants.

Definition 3.1 We say that a default theory (D, W) is *semi-representable* in the class of default theories $\mathcal{D}$ if there is an extension of the language $\mathcal{L}$, $\mathcal{L}'$, and a theory $(D', W') \in \mathcal{D}$ such that (D', W') is a theory in the language $\mathcal{L}'$ and for every theory $T \subseteq \mathcal{L}$, T is an extension of (D, W) if and only if there is an extension T' of (D', W') such that $T = T' \cap \mathcal{L}$.

We turn our attention to semi-normal rules. The concept of a semi-normal rule is a natural generalization of a normal default rule. The idea is that the process of derivation here is weaker than in the case of normal default theories. Instead of deriving $\psi \wedge \gamma$ out of φ and consistency of $\psi \wedge \gamma$, one cautiously derives only γ. This mode of reasoning does not, in general, guarantee the existence of extensions.

Example 3.1 Let $W = \emptyset$, $D = \{\frac{:(p\wedge\neg q)}{p}, \frac{:(q\wedge\neg r)}{q}, \frac{:(r\wedge\neg p)}{r}\}$ where p,q,r are distinct atoms. Then (D, W) is a semi-normal default theory without extensions.

A weakly semi-normal theory is very closed to semi-normal. Besides of cautious derivations, as described above, we allow derivation steps which employ rules that do not require justifications at all. These rules, in the provability paradigm we use are applicable with respect to any context.

We will prove now a result on semi-representability of default theories.

Theorem 3.2 *Every default theory is semi-representable in the class of weakly semi-normal default theories.*

Proof: Let (D', W') be a default theory. For every default rule

$$r = \frac{\alpha : \beta_1^r, \ldots, \beta_{k_r}^r}{\gamma}$$

in D introduce k_r new constants $c_1^r, \ldots, c_{k_r}^r$. The language $\mathcal{L}'$ is the extension of the language $\mathcal{L}$ by all these constants (for all the rules of D).

Now, for every rule r as above consider $k_r + 1$ new rules:
First k_r rules are of the form:

$$d_{r,i} = \frac{\alpha : (\beta_i^r \wedge c_i^r)}{c_i^r}$$

$(i = 1, \ldots, i_r)$ and one more rule:

$$e_r = \frac{c_1^r \wedge \ldots \wedge c_{k_r}^r :}{\gamma}$$

When r has no justifications then $e_r = r$ and there are no rules $d_{r,i}$ at all. Now, define $R_r = \{d_{r,1}, \ldots, d_{r,k_r}, e_r\}$, and finally:

$$D' = \bigcup_{r \in D} R_r$$

Clearly (D', W) is a weakly seminormal theory.
We will prove that:
(1) Every extension of (D, W) extends to an extension of (D', W).
(2) For every extension T' of (D', W), the theory $T' \cap \mathcal{L}$ is an extension of (D, W).
(3) If T', T'' are extensions of (D', W) and $T' \cap \mathcal{L} = T'' \cap \mathcal{L}$ then $T' = T''$.
These three facts together clearly imply our theorem.

We prove now (1). First we make the following simple observation:

Claim 3.3 *Let $T \subseteq \mathcal{L}$ be consistent theory, and let $\{c_1, \ldots, c_s\}$ be a set of new atoms. Then $T^+ = T \cup \{c_1, \ldots c_s\}$ is consistent. Moreover $Cn(T^+) \cap \mathcal{L} = Cn(T)$.*

Proof of claim: Every valuation V of $\mathcal{L}$ can be extended to a valuation V' of $At \cup \{c_1, \ldots, c_s\}$ making $c_1, \ldots, c_s$ true. Since V' and V coincide on At, they coincide on all the formulas of $\mathcal{L}$ as well. Therefore, if T is consistent, so is T'. Also, every formula of $\mathcal{L}$ unprovable from T is still unprovable from T'. Thus the second part of the claim is valid. Claim

Continuing our argument for (1) we first find an extension T' of (D', W) which contains T. Consider the case when $\mathcal{L}$ is an extension of (D, W). Since we admit the possibility of justification free-rules in D, we know that $\mathcal{L}$ is an extension of (D, W) if and only if the closure of W under the justification-free rules in D is inconsistent. Notice that the justification-less rules in D' are of two types: those whose prerequisites are in $\mathcal{L}$ and new rules of the form e_r. Since $W \subseteq \mathcal{L}$, the closure of W under justification-free rules of D' and the consequence operation of $\mathcal{L}'$ is generated by the closure of W under justification-free rules in D and tautologies of $\mathcal{L}'$. This implies that $\mathcal{L}$ is an extension of (D, W) if and only if $\mathcal{L}'$ is an extension of (D', W).
Therefore we can assume, from now on, that (D, W) has only consistent extensions. So, let T be a consistent extension of (D, W). In order to see what a desired T' is let us look at a rule r in D. If

$$r = \frac{\alpha : \beta_1^r, \ldots, \beta_k^r}{\gamma}$$

two cases are possible. If $T \not\vdash \alpha$ set $C_r = \emptyset$. But if $T \vdash \alpha$ define:

$$C_r = \{c_i^r : \neg\beta_i \notin T\}$$

Now set $T' = Cn(T \cup \bigcup_{r \in D} C_r)$. By Claim 3.3, $T = T' \cap \mathcal{L}$. Moreover T' is closed under consequence.

Now, it is easy to see that T is an extension of (D, W) if and only if T is the closure of W under the propositional rules of proof and the rules from D_T, where $D_T = \{\frac{\alpha :}{\gamma} : \frac{\alpha : \beta_1, \ldots, \beta_k}{\gamma} \in D \wedge \neg\beta_1 \notin T, \ldots \neg\beta_k \notin T\}$.

Let us now look at both D_T and $D_{T'}$. We notice that if $r \in D$ has no justification then $r \in D_{T'}$ and if $r = \frac{\alpha : \beta_1, \ldots, \beta_k}{\gamma}$ then whenever β_i is consistent with T then the rule $\frac{\alpha :}{c_i^r}$ belongs to $D_{T'}$. The rule e_r is justification-free and is always in $D_{T'}$. It follows that unless all the formulas β_i are consistent with T, the rule e_r is not applicable. This implies that if T is an extension of (D, W) then for T' as defined above, all the constants from the set $\bigcup_{r \in D_T} C_r$ can be T'-proved and so, for every rule $r \in D_T$, its consequent γ is proved using $D_{T'}$ as well. This implies that T' coincides with the closure of W under the rules from $D_{T'}$ and so T' is an extension of (D', W).

Next, we prove (2). Let T' be an extension of (D', W). Consider the collection of rules actually used in the reconstruction of T'. These rules are of three types: Either its consequent is a new constant (when the rule is of the form $d_{r,i}$, or it is a justification-free rule $\frac{\alpha:}{\gamma}$ with $\gamma \in \mathcal{L}$, or, finally, it is of the form $\frac{c_1^r \wedge \ldots \wedge c_k^r:}{\gamma}$ again with $\gamma \in \mathcal{L}$.
It is now clear that $\gamma \in \mathcal{L}$ is the consequent of an applicable default rule (with respect to T') if and only if $\alpha \in T' \cap \mathcal{L}$ or if $c_1^r, \ldots, c_k^r$ all belong to T'. But this last case happens precisely when $\beta_1^r, \ldots, \beta_k^r$ are all consistent with respect to T'. But $\beta_1^r, \ldots, \beta_{k_r}^r$ all belong to $\mathcal{L}$. Therefore $\beta_1^r, \ldots, \beta_{k_r}^r$ are consistent with respect to T if and only if they are consistent with respect to $T' \cap \mathcal{L}$. This implies that all the consequents of these applicable defaults that are in $\mathcal{L}$, are provable from W using rules of D with $T' \cap \mathcal{L}$ as a context. It is also clear that no other consequent is provable. Therefore $T' \cap \mathcal{L}$ is an extension of (D, W).

Finally, we prove (3). If T is an extension of (D, W) and T' is an extension of (D', W) and $T \subseteq T'$ then $T \subseteq T' \cap \mathcal{L}$, and by (2) $T' \cap \mathcal{L}$ is an extension of (D, W). Therefore $T = T' \cap \mathcal{L}$ since different extensions of (D, W) cannot be included one in the other (Theorem 1.5). Consequently, if T_1, T_2 are two extensions of (D', W), and T is an extension of (D, W), and $T \subseteq T_1$, $T \subseteq T_2$ then $T_1 \cap \mathcal{L} = T = T_2 \cap \mathcal{L}$. Now let us look at the form of the rules in D'. Those of these rules which are used to derive new constants have as a justification a (single) formula of $\mathcal{L}$. Therefore its consistency with T_i depends only on $T_i \cap \mathcal{L}$. But for both $i = 1, 2$ this intersections $T_1 \cap \mathcal{L}$ and $T_2 \cap \mathcal{L}$ coincide. Therefore T_1 and T_2 contain precisely the same new atoms. Since T_i is generated by $T_i \cap \mathcal{L}$ and a collection of new atoms, T_1 and T_2 must be identical.

4 Representation for weak extensions

Finally, we prove a representability result for weak extensions. We will show that every finite set of finitely generated theories is precisely the collection of weak extensions of a suitably chosen default theory.

Theorem 4.1 *Let $T_1, \ldots, T_k$ be a finite set of theories in $\mathcal{L}$, each T_j finitely generated. Then there exists a finite default theory (D, W) such that $T_1, \ldots,$ T_k are precisely the weak extensions of (D, W).*

Proof: We can assume that $T_i = Cn(\{\varphi_i\})$. Define a preordering in the set $\{\varphi_i : i \leq k\}$ as follows:

$$\varphi_i \prec \varphi_j \ \equiv \vdash \varphi_j \supset \varphi_i$$

Clearly: $\varphi_i \prec \varphi_j \ \equiv \ T_i \subseteq T_j$.
Define a default theory as follows: put $W = \emptyset$, and to define D first define $Z_i = \{\varphi_j : \neg(\varphi_j \prec \varphi_i)\}$. Thus $\varphi_j \in Z_i$ if and only if $\not\vdash \varphi_i \supset \varphi_i$. Define $\neg Z = \{\neg\psi : \psi \in Z\}$. Now we are ready to define our set of default rules D. When φ_i is minimal in the ordering $\prec$, set

$$d_i = \frac{: \neg Z_i}{\varphi_i}$$

and for φ_i that is not minimal define:

$$d_i = \frac{\varphi_i : \neg Z_i}{\varphi_i}$$

Then set $D = \{d_i : 1 \leq i \leq k\}$. Clearly, (D, W) is a finite default theory.
Firstly, we claim that every T_i is a weak extension of (D, W). Indeed, it is easy to see that d_i is the only applicable rule with respect to T_i. Notice that d_i is the only applicable rule and the consequent of d_i is φ_i and $T_i = Cn(\{\varphi_i\})$. Therefore T_i is a weak extension – because it is generated by W and the conclusion of (the only) generating rule.

Next, we prove that T_i's are the only weak extensions of (D, W). So let T be a weak extension of (D, W). Then T is generated by W and the consequents of applicable default rules. Let d be such rule.
Case 1. $d = \frac{:\neg Z_i}{\varphi_i}$. Then all the formulas $\varphi_j, j \neq i$ do not belong to T. But $\varphi_i \in T$. Since T is generated by some of formulas φ_n, $T = Cn(\{\varphi_i\})$.
Case 2: $d_i = \frac{\varphi_i:\neg Z_i}{\varphi_i}$. Here all the formulas φ_j such that $\not\vdash \varphi_i \supset \varphi_j$ do not belong to T, but $\varphi_i \in T$. As above $T = Cn(\{\varphi_i\})$.

It is perhaps worth mentioning that for minimal $\varphi_i i$, T_i is also an extension of (D, W).

Corollary 4.2 *Let $T_1, \ldots T_k$ be a finite set of stable theories in the language $\mathcal{L}_L$ such that for every $i \leq k$, the theory $S_i = T_i \cap \mathcal{L}$ is finitely generated. Then there is a finite theory $I \subseteq \mathcal{L}_L$ such that T_i are precisely autoepistemic expansions of I.*

Proof: As shown in [MT89] under the translation $\frac{\varphi:\beta_1,\ldots,\beta_m}{\gamma} \mapsto L\varphi \wedge M\beta_1 \wedge \ldots \wedge M\beta_m \supset \gamma$, weak exensions of a default theory are precisely objective parts of the autoepistemic expansions of the translation. Since for a stable T, $T = St(T \cap \mathcal{L})$, corollary follows.

Acknowledgements

This work was partially supported by Army Research Office under grant DAAL03-89-K-0124, and by National Science Foundation and the Commonwealth of Kentucky EPSCoR program under grant RII 8610671.

References

[Eth88] D. W. Etherington. *Reasoning with Incomplete Information.* Pitman, London, 1988.

[Fit90] M. C. Fitting. *First Order Logic and Automated Theorem Proving.* Springer Verlag, 1990.

[Men64] E. Mendelson. *Introduction to Mathematical Logic.* Van Nostrand Reinhold, 1964.

[MT89] W. Marek and M. Truszczyński. Relating autoepistemic and default logics. In *Principles of Knowledge Representation and Reasoning*, pages 276–288, San Mateo, CA., 1989. Morgan Kaufmann.

[MT93] W. Marek and M. Truszczyński. Nonmonotonic Logics; Context-Dependent Reasoning, *to appear.*

[Rei80] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[YBB92] F. Yang, H. Blair, and A. Brown. Programming in default logic. University of Syracuse, 1992.