Appeared in *Proceedings of ITALLC'98.*

Characterization Results for d-Horn Formulas,

or On formulas that are true on Dual Reduced Products

Carlos Areces	Verónica Becher	Sebastián Ferro
UvA	UBA	UBA
The Netherlands	Argentina	Argentina
carlos@wins.uva.nl	vbecher@dc.uba.ar	sf6o@dc.uba.ar

February, 1998

Abstract

We provide two different model theoretic characterizations of a new fragment of first-order logic which we call d-Horn formulas. This fragment is dual to the well know Horn fragment and have the same complexity for provability. The method used in the characterization (syntactic translation functions between formulas which are mimicked by translation functions between models) might be applied to characterize other first-order restrictions.

Keywords: Horn Formulas, d-Horn Formulas, Model Theoretic Characterizations.

1 Introduction

The Horn restriction of first-order logic (FO) is known to be a relevant fragment specially for Computer Science, its main import being that Horn theorem proving is polynomial via the SLD resolution method. For instance, the PROLOG programming language is based on the Horn fragment.

In this work we identify a new restriction of FO, the d-Horn set, that possesses the same complexity for theorem proving. We present a syntactic definition and different model-theoretic characterizations of d-Horn formulas.

While not disjoint, Horn and d-Horn seem to be dual FO restrictions. Some examples of sentences in these classes are: reflexivity $(\forall x)P(x,x)$, irreflexivity $(\forall x)\neg P(x,x)$ and symmetry $(\forall xy)(P(x,y) \to P(y,x))$ are all Horn as well as d-Horn sentences; transitivity $(\forall xyz)(P(x,y) \wedge P(y,z) \to P(x,z))$ is a Horn sentence not equivalent to any d-Horn sentence; but there are also d-Horn formulas which are not Horn formulas, witness the sentence for connectedness $(\forall xy)(P(x,y) \vee P(y,x))$.

We define a syntactic translation function (of linear complexity) that maps Horn formulas to d-Horn formulas, and conversely, and possesses the crucial property of preserving satisfiability. This is the key fact to be used in showing that theorem provability for the d-Horn fragment has also polynomial complexity.

The purely model theoretic characterization of d-Horn formulas (without translation functions) is as follows: the set of d-Horn formulas can be characterized as the set of formulas whose validity is preserved under *dual reduced products of models* (an operation defined as a variant of the reduced product [1].)

The sort of duality explored in this work suggests studying a similar behavior for other sets of FO formulas, as for example the set of positive formulas characterized by preservation under homomorphisms.

2 The Propositional Case

We start with a simple result in Propositional Logic (PL) which motivates our further research in FO.

Conditional sentences in PL are defined in [1] and a model theoretic characterization is also known:

Definition 2.1 (Conditional Sentences) A *conditional sentence* is a conjunction $\varphi_1 \wedge \ldots \wedge \varphi_n$ in PL such that each φ_i is either

- a propositional symbol S,
- a disjunction of negated propositional symbols $\neg S_1 \vee \ldots \vee \neg S_n$, or
- a disjunction of negated propositional symbols and a propositional symbol $\neg S_1 \vee \ldots \vee \neg S_n \vee S_{n+1}$.

Theorem 2.2 (Characterization of Conditional Sentences) A theory Γ of PL is preserved under intersections if and only if Γ has a set of conditional axioms.

Where "to be preserved under intersections" means that for a non empty index set I, $\{\mathcal{A}_i \models \Gamma\}_{i\in I} \Rightarrow \bigcap_{i\in I} \mathcal{A}_i \models \Gamma$. Since models in PL can be thought simply as sets of propositional variables (those which are true in the model), the concept of intersecting models is well defined.

Consider now the dual pattern: conjunction of sentences where only at most one negative propositional symbol appeared in disjunction with positive propositional symbols (which we called quasi-positive sentences.) A similar characterization for this set is not hard to find.

Theorem 2.3 (Characterization of Quasi-Positive Sentences) A theory Γ of PL is preserved under unions if and only if Γ has a set of quasi-positive axioms.

Proof. To prove the right to left implication it suffices to show that quasi-positive sentences are preserved by unions. Namely, let φ be quasi-positive and take two models such that $\mathcal{A} \models \varphi$ and $\mathcal{B} \models \varphi$ we have to show that $\mathcal{A}\cup\mathcal{B} \models \varphi$.

If φ is a positive sentence then clearly from $\mathcal{A} \models \varphi$ we can infer $\mathcal{A}\cup\mathcal{B} \models \varphi$. Suppose then that $\varphi = (\neg S_i \vee \psi)$ for $S_i \in \mathbb{P}$ and ψ a positive sentence. There are two cases:

- If $S_i \in \mathcal{A}$ then $\mathcal{A} \models (\neg S_i \vee \psi)$ implies $\mathcal{A} \models \psi$. Then $\mathcal{A}\cup\mathcal{B} \models \psi$, since ψ is positive, and $\mathcal{A}\cup\mathcal{B} \models \neg S_i \vee \psi$.
- If $S_i \notin \mathcal{A}$ then again there are two cases
  - If $S_i \in \mathcal{B}$ then $\mathcal{B} \models \psi$. Then $\mathcal{A}\cup\mathcal{B} \models \psi$, and $\mathcal{A}\cup\mathcal{B} \models \neg S_i \vee \psi$.
  - If $S_i \notin \mathcal{B}$ then $S_i \notin \mathcal{A}\cup\mathcal{B}$ and $\mathcal{A}\cup\mathcal{B} \models \neg S_i$. Hence $\mathcal{A}\cup\mathcal{B} \models \neg S_i \vee \psi$.

Let's see now the hard direction. Suppose Γ is preserved under unions. Let Δ the set of quasi-positive consequences of Γ. It suffices to show that if $\mathcal{B} \models \Delta$ then $\mathcal{B} \models \Gamma$.

Let $\mathcal{B}$ a model for Δ. For each $S_i \in \mathcal{B}$ we define Σ_{S_i} the set of sentences $S_i \wedge \psi$ where ψ is negative, that are true in $\mathcal{B}$. Let's notice that the finite conjunction of elements of Σ_{S_i} is equivalent to a sentence in Σ_{S_i}.

Let $\varphi \in \Sigma_{S_i}$. Clearly $\neg\varphi$ is equivalent to a quasi-positive sentence θ that is falsified in $\mathcal{B}$. Then, $\neg\varphi$ is not a consequence of Γ, therefore $\Gamma \cup \{\varphi\}$ is satisfiable. Hence, $\Gamma \cup \Sigma_{S_i}$ is satisfiable.

Let $\mathcal{A}_{S_i}$ be a model of $\Gamma \cup \Sigma_{S_i}$. If $S_i \in \mathcal{B}$ then $S_i \in \mathcal{A}_{S_i}$. $\mathcal{B} \subseteq \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$.

We show $\bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i} \subseteq \mathcal{B}$. Namely $S_j \in \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$ implies $S_j \in \mathcal{B}$. By contraposition, $S_j \notin \mathcal{B}$ implies $S_j \notin \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$.

If $S_j \notin \mathcal{B}$ then $\mathcal{B} \models \neg S_j$ and $S_i \wedge \neg S_j \in \Sigma_{S_i}, \forall S_i \in \mathcal{B}$. Thus, $S_j \notin \mathcal{A}_{S_i}, \forall S_i \in \mathcal{B}$. Thus $S_j \notin \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$. Hence $\mathcal{B} = \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$.

Now, each $\mathcal{A}_{S_i}$ is a model of $\Gamma \cup \Sigma_{S_i}$ and, a fortiori, a model of Γ. However, Γ is preserved by unions, then $\bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$ is a model of Γ. As $\mathcal{B} = \bigcup_{S_i \in \mathcal{B}} \mathcal{A}_{S_i}$, $\mathcal{B}$ is a model of Γ. QED

This result in no more than a simple exercise in basic model theory. But what is interesting is that given the characterization for positive formulas, we can give a much simpler proof, by using *translation functions*.

Definition 2.4 (Translation Function $t^{\mathcal{L}}$) Let $t^{\mathcal{L}} : \mathcal{L} \longrightarrow \mathcal{L}$ be defined recursively as:

$t^{\mathcal{L}}(S) = \neg S$, for S a propositional symbol.
$t^{\mathcal{L}}(\neg S) = S$, for S a propositional symbol.
$t^{\mathcal{L}}(\varphi \vee \psi) = t^{\mathcal{L}}(\varphi) \vee t^{\mathcal{L}}(\psi)$.
$t^{\mathcal{L}}(\varphi \wedge \psi) = t^{\mathcal{L}}(\varphi) \wedge t^{\mathcal{L}}(\psi)$.

It is clear that if φ is a conditional sentence then $t^{\mathcal{L}}(\varphi)$ is a quasi-positive sentence, and vice versa. Suppose furthermore that we define the following translation between PL models.

Definition 2.5 (Translation Function $t^{\mathcal{M}}$) Let $t^{\mathcal{M}} : \mathcal{M} \mapsto \mathcal{M}$ be defined simply as $t^{\mathcal{M}}(\mathcal{A}) = \mathcal{A}^c$.

Again, by the fact that PL models are sets of propositional symbols the translation is well defined. We will soon drop the superscripts $^{\mathcal{L}}$ and $^{\mathcal{M}}$ when no confusion arises. An important characteristic of this translations is that they preserve validity.

Proposition 2.6 Let φ be a conditional sentence or a quasi-positive sentence and $\mathcal{A}$ a PL model, then $\mathcal{A} \models \varphi$ iff $t(\mathcal{A}) \models t(\varphi)$.

Now we can reprove the characterization theorem as follows.

Theorem 2.7 (Characterization of Quasi-Positive Sentences) A theory Γ of PL is preserved under unions if and only if Γ has a set of quasi-positive axioms.

PROOF.

$\Rightarrow$) Let Δ be the set of quasi-positive axioms for Γ. Let also $\mathcal{A}_i, i \in I$ be PL models such that $\mathcal{A}_i \models \Gamma$. By definition, $t(\Delta)$ is a set of conditional sentences. By Proposition 2.6 above, $t(\mathcal{A}_i) \models t(\Delta)$. By Theorem 2.2, $\bigcap_{i\in I} t(\mathcal{A}_i) \models t(\Delta)$. Iff, $t(\bigcap_{i\in I} t(\mathcal{A}_i)) \models t(t(\Delta))$. Iff $\bigcup_{i\in I} \mathcal{A}_i \models \Delta$.

$\Leftarrow$) Now, let Γ be preserved under unions. We claim that $t(\Gamma)$ is preserved under intersections. To prove it, take $I, \mathcal{A}_i, i \in I$ such that $\mathcal{A}_i \models t(\Gamma)$. Iff $t(\mathcal{A}_i) \models \Gamma$. As Γ is preserved under unions, $\bigcup_{i\in I} t(\mathcal{A}_i) \models \Gamma$, iff $t(\bigcup_{i\in I} t(\mathcal{A}_i)) \models t(\Gamma)$. That is, $\bigcap_{i\in I} \mathcal{A}_i \models t(\Gamma)$. Applying now Theorem 2.2 we have that $t(\Gamma)$ has a set of conditional axioms Δ. But then $t(\Delta)$ is a set of quasi-positive axioms for Γ. QED

The reader might wonder which is the relation between these results and Horn/d-Horn formulas. If attention is given to the definition of conditional sentence, it will be noted that they are "the Horn formulas of PL" while we will define d-Horn formulas as the FO equivalent of quasi-positive sentences. Furthermore, the techniques used in the last characterization result for quasi-positive formulas can be lifted (mutatis mutandis) directly to FO. This is the topic of the next section.

3 Horn and d-Horn Formulas

In this section, $\mathcal{L}$ denotes a first order language, with the usual notational conventions. We use $\equiv$ as the identity symbol. A first order *model* is a tuple $\mathcal{A} = \langle A, \{R_i \mid i \in I_1\}, \{f_i \mid i \in I_2\}, \{c_i \mid i \in I_3\}\rangle$ such that A is a non-empty domain, $R_i, i \in I_1$ are relations over A, $f_i, i \in I_2$ functions on A and $c_i, i \in I_3$ constants in A. We start by stating the characterization of Horn formulas as presented in [1].

Definition 3.1 (Horn formulas) A formula φ of $\mathcal{L}$ is said to be a *basic Horn formula* iff φ is a disjunction of formulas θ_i, $\varphi = (\theta_1 \vee \ldots \vee \theta_m)$ where at most one of the formulas θ_i is an atomic formula, the rest being negations of atomic formulas.

Horn formulas are built up from basic Horn formulas with the connectives $\wedge$, $\exists$ and $\forall$. A *Horn sentence* is a Horn formula with no free variables.

The characterization uses the notion of reduced products.

Definition 3.2 (Filter, Ultrafilter and Reduced Product) Let I be a nonempty set, and $\mathcal{P}(I)$ the set of all subsets of I. A *filter* D over I is a set $D \subseteq \mathcal{P}(I)$ such that:

(i) $I \in D$.

(ii) If $X, Y \in D$ then $X \cap Y \in D$.

(iii) If $X \in D$ and $X \subseteq Z \subseteq I$ then $Z \in D$.

D is a *proper filter* iff it is not the improper filter $\mathcal{P}(I)$. D is said to be an *ultrafilter* over I iff D is a filter over I such that for all $X \in \mathcal{P}(I)$, $X \in D$ iff $(I \setminus X) \notin D$ (ultrafilters are always proper.) Given models $\mathcal{A}_i, i \in I$, the *reduced product over a nonempty proper filter* D (not. $\Pi_D \mathcal{A}_i$) is the model described as follows:

(i) The domain of $\Pi_D \mathcal{A}_i$ is $\Pi_D A_i$ the reduced product of A_i modulo D.

(ii) Let R be an n-ary relation symbol. The interpretation of R in $\Pi_D \mathcal{A}_i$ is the relation $R^{\Pi_D \mathcal{A}_i}(f^1_D \dots f^n_D)$ if and only if $\{i \in I |\ R^{\mathcal{A}_i}(f^1(i) \dots f^n(i))\} \in D$.

(iii) Let F be an n-ary function symbol. Then F is interpreted in $\Pi_D \mathcal{A}_i$ by the function $F^{\Pi_D \mathcal{A}_i}(f^1_D \dots f^n_D) = \langle F^{\mathcal{A}_i}(f^1(i) \dots f^n(i)) : i \in I \rangle_D$.

(iv) Let c be a constant. Then c is interpreted by the element $c^{\Pi_D \mathcal{A}_i} = \langle c^{\mathcal{A}_i} : i \in I \rangle_D$.

Theorem 3.3 ([1]. Theorem 6.2.5. Horn Characterization) Let $\mathcal{A}_i$ for $i \in I$ be models. Let D be a proper filter on I and $f^1, \dots, f^n$ be elements of $\Pi_{i \in I} A_i$. A formula $\varphi(x_1 \dots x_n)$ is equivalent to a Horn formula iff

$$\{i \in I |\ \mathcal{A}_i \models \varphi[f^1(i) \dots f^n(i)]\} \in D \Rightarrow \Pi_D \mathcal{A}_i \models \varphi[f^1_D \dots f^n_D].$$

As in the propositional case, a good definition of translation functions $t^{\mathcal{L}}$ and $t^{\mathcal{M}}$ lets us capture, using the above result, the set of d-Horn formulas. We start by defining this set. Intuitively, d-Horn formulas are built exactly the same as Horn formulas but with negated atoms where Horn formulas allow for positive atoms and vice versa, but special care has to be taken with equality (which has a fixed meaning on model and cannot be "adjusted" by the model translation $t^{\mathcal{M}}$.)

Definition 3.4 (d-Horn Formulas) A formula φ of $\mathcal{L}$ is said to be a *basic d-Horn formula* iff $\varphi = (\theta_1 \vee \dots \vee \theta_m)$, where at most one atomic identity formula (those of the form $t_1 \equiv t_2$) appears non negated and all atomic non identity formulas appear non negated; or (exclusive) all atomic identity formulas appears negated and at most one atomic non identity formula appears negated.

d-Horn formulas are built up from basic d-Horn formulas with the connectives $\wedge$, $\exists$ and $\forall$. *A d-Horn sentence* is a d-Horn formula with no free variables.

It is easy to see now that the sets of Horn formulas and d-Horn formulas are not disjoint. Any atomic formula is both Horn and d-Horn; so is the negation of an atomic formula, and so is an implication of a positive atomic antecedent and a positive atomic consequent. Identity formulas are a special case, being Horn if and only if they are d-Horn. While Horn formulas include all the negative formulas (namely, formulas just involving negated atoms) d-Horn formulas include all positive formulas free of identity.

As it was said before, the formulas for reflexivity $(\forall x)P(x,x)$, irreflexivity $(\forall x)\neg P(x,x)$ and symmetry $(\forall xy)(P(x,y) \rightarrow P(y,x))$ are all Horn as well as d-Horn sentences. In contrast, the formula for connectedness $(\forall xy)(P(x,y) \vee P(y,x))$ is d-Horn but not Horn. The formula for transitivity $(\forall xyz)(P(x,y) \wedge P(y,z) \rightarrow P(x,z))$ is Horn but not d-Horn. The formula $(\forall xyz)(P(x,y) \wedge P(y,z) \rightarrow (P(x,z) \vee P(z,x))$ is neither Horn nor d-Horn. It is quite straightforward to define a translation function that takes Horn formulas to d-Horn and conversely.

Definition 3.5 (Translation Function $t^{\mathcal{L}}$) Let $t^{\mathcal{L}} : \mathcal{L} \longrightarrow \mathcal{L}$ be defined recursively as:

$t^{\mathcal{L}}(\varphi) = \varphi$ if φ is an atomic or negated atomic identity formula.
$t^{\mathcal{L}}(\varphi) = \neg\varphi$ if φ is an atomic non identity formula.
$t^{\mathcal{L}}(\neg\varphi) = \varphi$ if φ is an atomic non identity formula.
$t^{\mathcal{L}}(\varphi \vee \psi) = t^{\mathcal{L}}(\varphi) \vee t^{\mathcal{L}}(\psi)$.
$t^{\mathcal{L}}(\varphi \wedge \psi) = t^{\mathcal{L}}(\varphi) \wedge t^{\mathcal{L}}(\psi)$.
$t^{\mathcal{L}}((\forall x)\varphi) = (\forall x)t^{\mathcal{L}}(\varphi)$.
$t^{\mathcal{L}}((\exists x)\varphi) = (\exists x)t^{\mathcal{L}}(\varphi)$.

$t^{\mathcal{L}}$ is defined in such a way that the image of a Horn formula is a d-Horn formula and vice versa. The translation function for models $t^{\mathcal{M}}$ is simpler.

Definition 3.6 (Translation Function $t^{\mathcal{M}}$) Let $t^{\mathcal{M}} : \mathcal{M} \mapsto \mathcal{M}$ be defined simply as $t^{\mathcal{M}}(\langle A, \{R_i \mid i \in I_1\}, \{f_i \mid i \in I_2\}, \{c_i \mid i \in I_3\}\rangle) = \langle A, \{R_i^c \mid i \in I_1\}, \{f_i \mid i \in I_2\}, \{c_i \mid i \in I_3\}\rangle$. Hence, $t^{\mathcal{M}}(\mathcal{A})$ is identical with $\mathcal{A}$ but has as relations the complements of the relations in $\mathcal{A}$.

Now we must check that satisfiability is preserved by the translations.

Proposition 3.7 (Satisfiability Preservation) Let $\mathcal{A}$ be a model and φ be a Horn or d-Horn formula, then $\mathcal{A} \models \varphi[a_1 \ldots a_n]$ iff $t(\mathcal{A}) \models t(\varphi)[a_1 \ldots a_n]$.

PROOF. The proof is by induction on φ. The only interesting cases being when φ is atomic or negated atomic as in all the other cases t commutes over the formula.

Suppose φ is an atomic identity formula. Then $\varphi = (t_1 \equiv t_2)$. By the translation function $t(\varphi) = \varphi = (t_1 \equiv t_2)$. Then as $\mathcal{A}$ and $t(\mathcal{A})$ give the same interpretation to functional and constant symbols, $\mathcal{A} \models (t_1 \equiv t_2)[a_1 \ldots a_n]$ iff $t(\mathcal{A}) \models t(t_1 \equiv t_2))[a_1 \ldots a_n]$. And the same is true for a negated atomic identity formula.

Suppose φ is $\varphi = R_i(t_1 \ldots t_m)$ for some relational symbol R_i and terms t_j. Then $t(\varphi) = \neg R_i(t_1 \ldots t_m)$. Now, $\mathcal{A} \models \varphi[a_1 \ldots a_n]$ iff $(t_1^{\mathcal{A}}[a_1 \ldots a_n] \ldots t_m^{\mathcal{A}}[a_1 \ldots a_n]) \in R^{\mathcal{A}}$ iff $(t_1^{\mathcal{A}}[a_1 \ldots a_n] \ldots t_m^{\mathcal{A}}[a_1 \ldots a_n]) \notin (R^c)^{\mathcal{A}}$. By the fact that $\mathcal{A}$ and $t(\mathcal{A})$ give the same interpretation to functional and constant symbols and the definition of $t^{\mathcal{L}}$, $t^{\mathcal{M}}$, $t(\mathcal{A}) \models t(\varphi)[a_1 \ldots a_n]$. QED

Now we can derive some useful properties of the translations.

Proposition 3.8 (Properties of the translations)

1. $t^{\mathcal{L}}$ and $t^{\mathcal{M}}$ are involutive: $t^{\mathcal{L}}(t^{\mathcal{L}}(\varphi)) = \varphi$ and $t^{\mathcal{M}}(t^{\mathcal{M}}(\mathcal{A})) = \mathcal{A}$.
2. $\models (\varphi \to \psi)$ iff $\models (t^{\mathcal{L}}(\varphi) \to t^{\mathcal{L}}(\psi))$.
3. For any $\varphi \in \mathcal{L}$, φ is equivalent to a Horn formula iff $t^{\mathcal{L}}(\varphi)$ is equivalent to a d-Horn formula.

(Above we apply t to the formula φ where negation appears only for atoms. Every first-order formula has an equivalent which satisfies this condition.)

PROOF.

1. Trivial.

2. We prove the left to right implication. The other is similar.

Suppose not. Then $\models (\varphi \to \psi)$ but $\not\models (t^{\mathcal{L}}(\varphi) \to t^{\mathcal{L}}(\psi))$. Hence there exists some model $\mathcal{A}$ such that $\mathcal{A} \not\models (t^{\mathcal{L}}(\varphi) \to t^{\mathcal{L}}(\psi))$ iff $t^{\mathcal{L}}(\mathcal{A}) \not\models t^{\mathcal{L}}(t^{\mathcal{L}}(\varphi) \to t^{\mathcal{L}}(\psi))$ iff $t^{\mathcal{L}}(\mathcal{A}) \not\models t^{\mathcal{L}}(t^{\mathcal{L}}(\varphi)) \to t^{\mathcal{L}}(t^{\mathcal{L}}(\psi))$ iff $t^{\mathcal{L}}(\mathcal{A}) \not\models \varphi \to \psi$ contradicting the hypothesis.

3. Let φ be equivalent to a Horn formula ψ. Then $\models \varphi \leftrightarrow \psi$ iff $\models t^{\mathcal{L}}(\varphi) \leftrightarrow t^{\mathcal{L}}(\psi)$ and $t^{\mathcal{L}}(\psi)$ is a d-Horn formula. QED

These results let us extend the characterization results of Horn formulas to d-Horn formulas in the following way.

Theorem 3.9 (Indirect d-Horn Characterization) Let $\mathcal{A}_i$ for $i \in I$ be models of $\mathcal{L}$. Let D be a proper filter on I and $f^1, \ldots, f^n$ be elements of $\Pi_{i\in I} A_i$. A formula $\varphi(x_1 \ldots x_n)$ is equivalent to a d-Horn formula iff

$$\{i \in I |\ \mathcal{A}_i \models \varphi[f^1(i) \ldots f^n(i)]\} \in D \Rightarrow \Pi_D t(\mathcal{A}_i) \models t(\varphi)[f^1_D \ldots f^n_D].$$

PROOF.

$\Rightarrow$) Suppose that φ is a d-Horn formula and that $\{i \in I |\ \mathcal{A}_i \models \varphi[f^1(i) \ldots f^n(i)]\} \in D$ holds. If and only if by Proposition 3.7 $\{i \in I |\ t(\mathcal{A}_i) \models t(\varphi)[f^1(i) \ldots f^n(i)]\} \in D$.

As $t(\varphi)$ is a Horn formula, it is preserved by reduced products: $\Pi_D t(\mathcal{A}_i) \models t(\varphi)[f^1_D \ldots f^n_D]$

$\Leftarrow$) Suppose that φ is such that whenever $\{i \in I |\ \mathcal{A}_i \models \varphi[f^1(i) \ldots f^n(i)]\} \in D$ then $\Pi_D t(\mathcal{A}_i) \models t(\varphi)[f^1_D \ldots f^n_D]$, and $\{i \in I |\ t(\mathcal{A}_i) \models t(\varphi)[f^1(i) \ldots f^n(i)]\} \in D$. Finally, $\Pi_D t(\mathcal{A}_i) \models t(\varphi)[f^1_D \ldots f^n_D]$.

Hence, $t(\varphi)$ is equivalent to a Horn formula. Now $t(t(\varphi))$ is equivalent to a d-Horn formula by Proposition 3.8 and by involution, φ is equivalent to a d-Horn formula. QED

It is possible yet to simplify the expression of Theorem 3.9 to one just containing the model translation using involution: $t(\Pi_D t(\mathcal{A}_i)) \models t(t(\varphi))[f^1_D, \ldots, f^n_D]$ iff $t(\Pi_D t(\mathcal{A}_i)) \models \varphi[f^1_D, \ldots, f^n_D]$. The theorem would then read:

Theorem 3.10 (Indirect d-Horn Characterization) Let $\mathcal{A}_i$ for $i \in I$ be models of $\mathcal{L}$. Let D be a proper filter on I and $f^1, \ldots, f^n$ be elements of $\Pi_{i\in I} A_i$. A formula $\varphi(x_1 \ldots x_n)$ is equivalent to a d-Horn formula iff

$$\{i \in I |\ \mathcal{A}_i \models \varphi[f^1(i) \ldots f^n(i)]\} \in D \Rightarrow t(\Pi_D t(\mathcal{A}_i)) \models \varphi[f^1_D \ldots f^n_D].$$

Even though this last formulation involves only translations between models, we are interested now in a direct characterization of d-Horn formulas; namely, one that also eliminates the translation for models. To make a parallel with the propositional case we are now in a case where we prove the result for $(\bigcap_{i\in I} \mathcal{A}^c_i)^c$ and we are after a results which proves $(\bigcap_{i\in I} \mathcal{A}^c_i)^c = \bigcup_{i\in I} \mathcal{A}$.

We introduce a new construction on models, called dual reduced product, notated as Π^*_D.

Definition 3.11 (Dual Reduced Product) Given models $\mathcal{A}_i, i \in I$, and D a nonempty proper filter, the *dual reduced product* $\Pi^*_D \mathcal{A}_i$ is the model described as follows:

(i) The domain of $\Pi^*_D \mathcal{A}_i$ is $\Pi_D A_i$.

(ii) Let R be an n-ary relation symbol. The interpretation of R in $\Pi^*_D \mathcal{A}_i$ is the relation $R^{\Pi^*_D \mathcal{A}_i}(f^1_D \dots f^n_D)$ if and only if $(\exists U$ an ultrafilter$)(D \subseteq U \wedge \{i \in I |\ R_i(f^1(i) \dots f^n(i))\} \in U)$.

(iii) Let F be an n-ary function symbol. Then F is interpreted in $\Pi^*_D \mathcal{A}_i$ by the function $F^{\Pi^*_D \mathcal{A}_i}(f^1_D \dots f^1_D) = \langle F_i(f^1(i) \dots f^n(i)) : i \in I \rangle_D$.

(iv) Let c be a constant. Then c is interpreted by the element $c^{\Pi^*_D \mathcal{A}_i} = \langle a_i : i \in I \rangle_D$.

Hence, dual reduced products are exactly like reduced products in their clauses for universes, functions and constants. What changes is the condition for the relations (this is in accordance to the translation between models we were using.) Perhaps surprisingly, the definition of dual reduced products is existential in nature ("there exists an ultrafilter U ..."), but if we consider again the propositional case we see that for an element to be in the union set, it needs to be just in one of the sets, contrasting with the universal definition of intersection.

To prove that this construction is the one we where looking for, we have to check that it coincides with the construction obtained through the translation function.

Proposition 3.12 Let $\mathcal{A}_i$ be models for $\mathcal{L}$, D be a proper filter on I, then $\Pi^*_D \mathcal{A}_i = t(\Pi_D t(\mathcal{A}_i))$.

Proof. As the translation does not change the universe of the model, or the interpretation of function and constant symbols these elements are in $t(\Pi_D t(\mathcal{A}_i))$ the same as in $\Pi_D \mathcal{A}_i$ and the same is true for $\Pi^*_D \mathcal{A}_i$.

It rests to check the interpretation of the relation symbols.

Suppose that R is an n-ary relational symbol of $\mathcal{L}$. We have to check that given $f^1, \dots, f^n \in \Pi^*_D A_i$, $R(f^1_D \dots f^n_D) \in \Pi^*_D \mathcal{A}_i$ iff $(R f^1_D \dots f^n_D) \in t(\Pi_D t(\mathcal{A}_i))$.

Suppose that $R(f^1_D \dots f^n_D) \in t(\Pi_D t(\mathcal{A}_i))$ iff $R(f^1_D \dots f^n_D) \notin \Pi_D t(\mathcal{A}_i)$ iff
$\{i \in I |\ R^{t(\mathcal{A}_i)}(f^1(i) \dots f^n(i))\} \notin D$ iff
$\{i \in I |\ (R_i)^c(f^1(i) \dots f^n(i))\} \notin D$ iff $\{i \in I |\ \neg R_i(f^1(i) \dots f^n(i))\} \notin D$.

$\Leftarrow$) But then it is consistent to extend D to a set D^{Ext} in such a way as to get $\{i \in I |\ R_i(f^1(i) \ldots f^n(i))\} \in D^{Ext}$ and furthermore, we can make this extension maximal and take an ultrafilter U.

$\Rightarrow$) Now suppose there exists an ultrafilter U extending D such that $\{i \in I |\ R_i(f^1(i) \ldots f^n(i))\} \in U$, we want to prove that $\{i \in I |\ \neg R_i(f^1(i) \ldots f^n(i))\} \notin D$.

Suppose not, then as U extends D we have $\{i \in I |\ \neg R_i(f^1(i) \ldots f^n(i))\} \in U$. But by hypothesis $\{i \in I |\ R_i(f^1(i) \ldots f^n(i))\} \in U$ and these two sets are complementary, arriving to a contradiction with the choice of U as an ultrafilter. QED

It would be very rewarding to find a direct proof for d-Horn characterization in terms of the dual reduced product operation, hopefully simpler than the proofs known for Horn characterization. Such a proof could provide an indirect Horn characterization following the same process used in this work but from d-Horn to Horn. The "hope" for a simpler proof comes from the similarities of the d-Horn fragment with the positive formulas (remember than the propositional set was called quasi-positive) and the easy characterization result of positive formulas via homomorphisms.

The dual reduced product of a set of models is a new model theoretic construction that might have interesting properties besides those presented in this work. We will briefly comment on the connections between reduced products and dual reduced products. It is easy to observe that the satisfaction of a formula $R(t_1 \ldots t_m)$ in a given reduced product $\Pi_D \mathcal{A}_i$ and valuation $[a_1 \ldots a_n]$ implies satisfaction in the corresponding dual reduced product.

Proposition 3.13 Let $\mathcal{A}_i, i \in I$ be models, D be a proper filter on I, $f_1, \ldots, f_m$ be elements of $\Pi_D A_i$ then $R^{\Pi_D \mathcal{A}_i}(t_1 \ldots t_n)[f_D^1 \ldots f_D^m]$ implies $R^{\Pi_D^* \mathcal{A}_i}(t_1 \ldots t_n)[f_D^1 \ldots f_D^n]$.

The proof only relies on the fact that every proper filter can always be extended to an ultrafilter. It is interesting to check when the converse holds. Apparently, only when D is an ultraproduct, which is a trivial case. Suppose we define a relation $E(I, D)$, where I is a set and D is a proper filter on I and that reflects the property we are studying now, namely, that for this special proper filter D on set I, the constructions Π_D and Π_D^* are equivalent. This condition is really a condition about the relations because the only difference between the two constructions is how the relational symbols in the language are interpreted. Actually, $E(D, I)$ iff $(\forall \{A_i\}_{i \in I})$ $(\forall n \in I\!N)$ $(\forall R)$ $(\forall f_D^1, \dots f_D^n \in \Pi_D A_i) R^{\Pi_D \mathcal{A}_i}(f_D^1 \dots f_D^n) \leftrightarrow R^{\Pi_D^* \mathcal{A}_i}(f_D^1 \dots f_D^n)$.

We know that $R^{\Pi_D A_i}$ is always included in $R^{\Pi_D^* A_i}$. So in order to satisfy $E(D, I)$ it is necessary and sufficient that $R^{\Pi_D^* A_i} \subseteq R^{\Pi_D A_i}$ holds. But as we saw in the proof of Proposition 3.12 for this to be true we need D to be an ultrafilter. Perhaps, setting further conditions on R or $\{\mathcal{A}_i\}, i \in I$ we can obtain a less strict condition on D, which we leave outside the scope of this work.

4 The Interest of the d-Horn Restriction

Given the great celebrity that the Horn restriction enjoys, the examples we have given here suggest that the d-Horn restriction might be of interest. As we already remarked, the formula for connectedness is a proper d-Horn sentence and so are (identity free) positive sentences. Hence, the expressive power of the d-Horn restriction is different from Horn's. In this short and quite undeveloped section we reveal the possible significance of the d-Horn restriction. The ideas we are presenting here are part of our future work.

Complexity Remarks of d-Horn Fragment The syntactic translation we presented maps Horn formulas to d-Horn formulas preserving satisfiability. Therefore, due to the soundness and completeness of FO, theoremhood in Horn and in d-Horn are in direct correspondence. Theorem proving for the Horn fragment has polynomial complexity, via SLD resolution [3]. Since the complexity of our translation function is linear, theorem proving is also polynomial for d-Horn. For instance, let Γ be a d-Horn theory and φ a d-Horn formula, and suppose we want to check if $\Gamma \models \varphi$ holds. By Corollary 3.7, this is the case if and only if $t(\Gamma) \models t(\varphi)$ holds. Given that $t(\Gamma)$ is equivalent to a Horn theory and $t(\varphi)$ a Horn formula, a derivation can be found in polynomial time by SLD resolution.

The Translation Function and Deduction Patterns Following C. S. Peirce [2], given a theory and a sentence two different sorts of logical inference can be performed, deduction and abduction, being dual forms of inference with respect to Modus Ponens. For example, consider the sentence $(\varphi \wedge \psi) \rightarrow \mu$. Given $(\varphi \wedge \psi)$, sentence μ can be deduced. However, given μ, via abduction $(\varphi \wedge \psi)$ is obtained. Numerous Artificial Intelligence logic oriented applications are based on abductive reasoning, like causation, explanation, language and image interpretation. From what we have already seen it is plain that the translation function carries through deduction. It also comes into obvious to realize that the translation function t carries through the abductive inference too. For any theory Γ and formulas φ, θ. $\Gamma \cup \{\theta\} \models \varphi$ iff $t(\Gamma) \cup \{t(\theta)\}) \models t(\varphi)$ iff $t(\Gamma) \cup \{t(\neg\varphi)\} \models t(\neg\theta)$. This last expression reveals a special behavior when φ, θ are atomic non-identity formulas, obtaining $\Gamma \cup \{\theta\} \models \varphi$ iff $t(\Gamma) \cup \{\varphi\} \models \theta$. This is a curious correlation since models for Γ and models for $t(\Gamma)$ in many cases are disjoint.

4.1 Illustrating d-Horn characterization results

This section deals with examples of FO sentences illustrating (proper) membership in the Horn and d-Horn restrictions.

Let's give a concrete example to illustrate the translation function over the first order language, the translation function over models and the preservation result for d-Horn formulae. Let φ be the characteristic axiom for connectivity, which is a d-Horn sentence.

$$\varphi = (\forall xy)(P(x, y) \vee P(y, x)) = (\forall xy)\psi(x, y)$$

(observe that φ implies reflexivity of the relation involved.)

Let the index set be $I = \{1, 2, 3\}$ and let $\mathcal{A}_i, i \in I$ be models of $\mathcal{L}$. We follow the usual convention that in models $\mathcal{A}_i$ relation symbols P are interpreted by relations R_i. Let $\mathcal{A}_i, \mathcal{A}_2, \mathcal{A}_3$ three models satisfying φ:

$\mathcal{A}_1 = \langle\{a\}, R_1\rangle$, where $R_1 = \{(a, a)\}$,
$\mathcal{A}_2 = \langle\{a, b\}, R_2\rangle$, where $R_2 = \{(a, a), (b, b), (a, b)\}$, and
$\mathcal{A}_3 = \langle\{a, b\}, R_3\rangle$, where $R_3 = \{(a, a), (b, b), (b, a)\}$.

Hence we have $\mathcal{A}_i \models \varphi$.

Now let's turn to the translation function over φ, which yields a Horn sentence. ($t(\varphi)$ implies irreflexivity of the relation it denotes.)

$$t(\varphi) = (\forall xy)(\neg P(x, y) \vee \neg P(y, x)).$$

Let's apply the translation function over the $\mathcal{A}_i$ models:

$t(\mathcal{A}_1) = \langle\{a\}, R_1^c\rangle$, where $R_1^c = t(R_1) = \{\}$,
$t(\mathcal{A}_2) = \langle\{a, b\}, R_2^c\rangle$, where $R_2^c = t(R_2) = \{(b, a)\}$, and
$t(\mathcal{A}_3) = \langle\{a, b\}, R_3^c\rangle$, where $R_3^c = t(R_3) = \{(a, b)\}$.

We want to illustrate our result that says that a formula $\varphi(x_1 \ldots x_n)$ is equivalent to a d-Horn formula then

$$\{i \in I : \mathcal{A}_i \models \varphi[a^1(i), \ldots, a^n(i)]\} \in D \Rightarrow \Pi_D t(\mathcal{A}_i) \models t(\varphi)[a^1_D \ldots a^n_D].$$

Let's turn now to the reduced product of models $t(\mathcal{A}_i)$. Consider the proper filter $D = \{\{2,3\},\{1,2,3\}\}$ on $I = \{1,2,3\}$.

$$\Pi_D t(\mathcal{A}_i) = \langle \Pi_D A_i, R\rangle$$

$$\text{where } \Pi_D A_i = \{[(aaa)], [(aba)], [(aab)], [(abb)]\}$$

$$\text{and } R = \{[(aba)], [(aab)]\} \text{ since } \{2,3\} = \{i \in I |\ R^c_i(f^1(i), f^2(i)\}.$$

$\Pi_D(t(\mathcal{A}_i))$ satisfies $t(\varphi)$, namely, $\Pi_D(t(\mathcal{A}_i)) \models (\forall xy)(\neg P(x,y) \vee \neg P(y,x))$.

We can also use the example to see that $t(\Pi_D t(\mathcal{A}_i)) \models \varphi[a^1_D \ldots a^n_D]$. Given that φ is a sentence, we can drop the assignment.

$$t(\Pi_D t(\mathcal{A}_i)) = t(\langle \Pi_D A_i, \{([aba],[aab])\}\rangle) = \langle \Pi_D A_i, S\rangle,$$

$$\text{where } S = \{(x,y)|\ x,y \in \Pi_D A_i\} \setminus \{([aba],[aab])\}.$$

S contains every pair of elements but $\{([aba],[aab])\}$. Clearly, $\langle \Pi_D A_i, R\rangle \models (\forall xy)(P(x,y) \vee P(y,x))$.

We can use the example further to illustrate that the pair $([aba],[aab])$ is not in the relation R for the dual reduced product construction. By definition of relations in a dual reduced product, $\{([aba],[aab])\} \in R$ iff there exists an ultrafilter $U \supseteq D$ such that $\{i \in I : R_i(f^1(i), \ldots, f^n(i))\} \in U$. Specifically,

$$\{([aba],[aab])\} \in R \iff \exists U s.t. \{i \in I : R_i([aba](i),[aab](i))\} = \{1\} \in U$$

Given that gen$\{2\}$ and gen$\{3\}$ are the only ultrafilters extending D, clearly $\{1\} \notin$ gen$\{2\}$ and $\{1\} \notin$ gen$\{3\}$. Hence $([aba],[aab]) \notin R$.

A d-Horn formula which is not Horn We know from its syntactic form that the characteristic axiom for connectedness is d-Horn. Let's prove now that there is no equivalent Horn formula. We will use the characterization result for Horn formulae that says that a formula $\varphi(x_1 \ldots x_n)$ is equivalent to a Horn formula iff

$$\{i \in I : \mathcal{A}_i \models \varphi[a^1(i), \ldots, a^n(i)]\} \in D \Rightarrow \Pi_D \mathcal{A}_i \models \varphi[a^1_D \ldots a^n_D].$$

We can reuse models $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and the filter $D = \{\{2,3\},\{1,2,3\}\}$ in order to show that their reduced product does not validate the formula of connectedness.

$$\Pi_D \mathcal{A}_i \not\models \psi[([aba],[aab])].$$

As $\Pi_D \mathcal{A}_i$ is not a model for φ, φ is not equivalent to a Horn formula.

d-Horn formulae are not preserved under direct products As expected, d-Horn formulae in general are not preserved under direct products. Again let's reuse the example to prove it.

$$\mathcal{A}_2 \times \mathcal{A}_3 = \langle A_2 \times A_3, R\rangle,$$

where $R = \langle$ $([aa], [aa])$, $([ab], [ab])$, $([ba], [ba])$, $([bb], [bb])$, $([aa], [ba])$, $([ab], [aa])$, $([ab], [bb])$, $([bb], [ba])$ $\rangle$. It suffices to show that there is an instance where ψ is not satisfied. In particular, $\mathcal{A}_2 \times \mathcal{A}_3 \not\models \psi([aa],[bb])$.

Neither Horn nor d-Horn Let's consider the following sentence

$$\varphi \equiv (\forall xyz)(P(x,y) \wedge P(y,z) \rightarrow P(x,z) \vee P(z,x)) = (\forall xyz)\psi(x,y,z).$$

We prove that φ is not equivalent to a Horn formula. Take models $\mathcal{A}_1 = \langle\{a,b,c\}, R_1\rangle$ such that $R_1 = \{(a,b),(b,c),(a,c)\}$ and $\mathcal{A}_2 = \langle\{a,b,c\}, R_2\rangle$ such that $R_2 = \{(a,b),(b,c),(c,a)\}$. Clearly $\mathcal{A}_i \models \varphi$. Let's consider the trivial filter $D = \{1,2\}$. Now,

$$\Pi_D \mathcal{A}_i \not\models \psi[[(a,a)],[(b,b)],[(c,c)]].$$

Therefore φ is not equivalent to any Horn formula.

Let's see now that φ is not equivalent to a d-Horn formula. Consider two new models $\mathcal{A}_1 = \langle\{a,b,c\}, R_1\rangle$ such that $R_1 = \{(a,b)\}$ and $\mathcal{A}_2 = \langle\{a,b,c\}, R_2\rangle$ such that $R_2 = \{(b,c)\}$. Clearly $\mathcal{A}_i \models \varphi$. Consider again the trivial filter $D = \{1,2\}$. Now,

$$\Pi^*_D \mathcal{A}_i \not\models \psi[[(a,a)],[(b,b)],[(c,c)]].$$

Therefore φ is not equivalent to any d-Horn formula.

5 Conclusions and Future Work

In this study we have identified a FO restriction that we named d-Horn, for which we have provided syntactic and model theoretic characterizations. The Horn and d-Horn are linked through translation functions that preserve satisfiability and have been used to carry on Horn properties over to the d-Horn set. In particular, the complexity of d-Horn theorem proving is polynomial via SLD resolution. d-Horn formulas include all the positive other than identity, resembling disjunctive databases in Computer Science.

A number of problems remain as future work. The same dual pattern explored in this study seems applicable to other sets of FO as the positive formulas, preserved via homomorphisms. As already remarked it can be of considerable interest to obtain a direct simple proof of our d-Horn characterization result. Then following the reverse path we have used in this work we could provide for a simpler indirect Horn characterization result A quite different line of research is to study the Horn and d-Horn restrictions of Modal Logic (ML), requiring the appropriate translation functions. Given that the Deduction Theorem does not hold in general for ML, it would be interesting to consider the correlation between abductive and deductive inference over theories and translated theories.

Acknowledgments: We would like to thank Dr. Guillermo Martinez for his corrections, directions and all the valuable discussions.

References

[1] Chang, C. and Keisler, H. *Model Theory.* Elsevier Science Publishers B. V., 1977.

[2] Rescher, N. *Peirce's Philosophy of Science: Critical Studies in his Theory of Induction and Scientific Method.* University of Notre Dame Press. Notre Dame, 1978.

[3] Schöning, U. *Logic for Computer Scientists.* Birkhäuser, 1989.