A PROOF-THEORETIC ANALYSIS OF GOAL-DIRECTED PROVABILITY

JAMES HARLAND
Department of Computer Science, University of Melbourne
Parkville, 3052, Victoria, Australia

Abstract

One of the distinguishing features of logic programming seems to be the notion of *goal-directed* provability, i.e. that the structure of the goal is used to determine the next step in the proof search process. It is known that by restricting the class of formulae it is possible to guarantee that a certain class of proofs, known as *uniform proofs*, are complete with respect to provability in intuitionistic logic. In this paper we explore the relationship between uniform proofs and classes of formulae more deeply. Firstly we show that uniform proofs arise naturally as a normal form for proofs in first-order intuitionistic sequent calculus. Next we show that the class of formulae known as hereditary Harrop formulae are intimately related to uniform proofs, and that we may extract such formulae from uniform proofs in two different ways. We also give results which may be interpreted as showing that hereditary Harrop formulae are the largest class of formulae for which uniform proofs are guaranteed to be complete, along the lines of an interpolation theorem.

1 Introduction

It has long been known that there are efficient implementation techniques which make Horn clauses, a particular fragment of first-order logic, able to be used as a programming language [12)], and that this class of formulae forms the semantic basis for the programming language Prolog [2,19)]. It has also been shown that this computational paradigm is as powerful as that of Turing machines [20)]. Thus we may think of Horn clauses as incorporating some form of algorithmic knowledge. As Horn clauses are not a particularly large fragment of first-order logic, it is perhaps not surprising that this class of formulae has such a relatively strong property. There have been various schemes proposed for logic programming languages which are extensions of Horn clauses [4,13,14,16,17,18)]. Given these various extensions, it seems natural to ask whether there is a maximal class of formulae which may be used as a programming language. Moreover, there does not seem to be a universally agreed criterion which may be used to determine what constitutes a logic programming language, without which any such notion of maximality would seem premature. A criterion of this nature has been proposed by Miller et al. [16)], in that they identify various first-order and higher-order fragments as logic programming languages by showing that these fragments satisfy a completeness property for a certain class of proofs. However it would seem that a general criterion should be strong enough not only to verify that certain fragments may be used as programming languages, but also to discover such fragments in the first place. Thus it seems natural to use the criterion of [16)], namely the completeness of *goal-directed provability*, to investigate this question of maximality for (first-order) logic programming languages.

A useful notion this context is that of a *uniform proof* [16]. A uniform proof is one in which the principal connective of the formula is introduced in the last step of the proof; in other words, when searching for a proof of a given formula, we need only consider the immediate subformulae of the desired conclusion. Hence we may think of uniform proofs as *goal-directed*, in that when searching for a uniform proof of a given goal, we may use the structure of the goal to determine the structure of the proof. We will denote uniform provability by $\vdash_u$. Such proofs lead to an identification of formulae with operations in a search space, and hence have a natural interpretation as instructions, thus establishing a direct relationship between proof and computation. This restriction also allows a more feasible implementation of the proof search process than may be done in the case of arbitrary (intuitionistic) proofs.

Uniform proofs may be used as a basis for logic programming [16,15)], and there are several interesting investigations along these lines. However, it is not the case that all intuitionistic proofs are uniform. Hence, one of the features of this approach is that the formulae involved are restricted to a class for which uniform proofs are complete with respect to intuitionistic logic, i.e. $F_1 \vdash_I F_2$ iff $F_1 \vdash_u F_2$ when F_1 and F_2 belong to a certain class of formulae. One such restriction is that the formulae which may be used as assertions (i.e. those which may appear on the left of $\vdash$) are Horn clauses, and the goals (i.e. those formulae which may appear on the right of $\vdash$) are conjunctions of atoms, so that when F_1 is a set of Horn clauses and F_2 is a conjunction of atoms, then $F_1 \vdash_I F_2$ iff $F_1 \vdash_u F_2$. Larger classes of formulae for which the existence of uniform proofs is guaranteed may also be given [16)], and the largest class of first-order formulae for which this property has been established is known as *hereditary Harrop formulae*. Intuitively, these formulae may be thought of as those which contain no negative occurrences of $\exists$ or $\vee$.

In this paper we examine the relationship between uniform proofs and hereditary Harrop formulae, and we give several results which may be interpreted as establishing the maximality of this class of formulae. As uniformity is a property of proofs rather than formulae, it is not strictly possible to establish that a given class of formulae is the largest one for which uniform proofs are complete. For example, given that $F_1 \vdash_u F_2$, where F_1 and F_2 satisfy some restriction, for any formula F whatsoever, $F_1, F \vdash_u F_2$. However, as we shall see, there is a natural relationship between uniform proofs and hereditary Harrop formulae. Essentially this is that whilst $F_1, F \vdash_u F_2$, there is a hereditary Harrop formula D such that $D \vdash_u F_2$, and D is related to F and F_1 in such a way that D is the formula "doing the work" in the uniform proof. The relationship between the formulae is made precise in later sections.

An important insight which arises from this analysis is that the class of hereditary Harrop formulae arises naturally out of the *permutation properties* of the rules of intuitionistic logic [3,11] This occurs by determining when it is possible to permute certain combinations of inference rules so that an arbitrary proof may be converted into a uniform proof. Thus we may identify hereditary Harrop formulae as a logic programming language purely from the notion of a uniform proof and the proof theory of intuitionistic logic; no prior knowledge of logic programming languages *per se* is needed. This suggests that the strategy of studying permutation rules in order to investigate the completeness of goal-directed provability may be used to identify logic programming languages independently of the logic in use; such a strategy has been used to identify logic programming languages in linear logic [8,9] In this way we may think of the permutation properties of the proof theory of the logic in question (in conjunction with the notion of goal-directed search) as determining what fragments of the logic may be used as a logic programming language.

2 Preliminaries

First we define hereditary Harrop formulae [16]. We assume the existence of a finite set of constant and function symbols, and a countable set of variables. We refer to the set of all ground terms as the *Herbrand universe*, denoted by $\mathcal{U}$.

Definition 2.1 *D and G formulae are given by the grammar*

$$D ::= A \mid \forall x D \mid D_1 \wedge D_2 \mid G \supset A$$
$$G ::= A \mid \exists x G \mid \forall x G \mid G_1 \wedge G_2 \mid G_1 \vee G_2 \mid D \supset G$$

where A is an atom.

We refer to D formulae as definite *formulae, and to G formulae as* goal *formulae. The set of all definite formulae will be referred to as $\mathcal{D}$, and the set of all goal formulae as $\mathcal{G}$.*

A program *is a set of closed definite formulae, and a goal is any closed goal formula.*

We will often refer to the above classes of formulae as hereditary Harrop formulae. Note that we do not allow negations here. We will refer to the formulae which do not contain any negations as *negation-free* formulae.

It was shown in [16] that an operational notion of proof $\vdash_o$ may be given for the above class of formulae in such a way that for a program P and a goal G, $P \vdash_o G$ iff $P \vdash_I G$ where $\vdash_I$ denotes intuitionistic provability, so that $P \vdash_o G$ iff there is a proof in intuitionistic logic of the sequent $P \longrightarrow G$. Below we give a slightly different definition, which we will denote as $\vdash_u$. The rules for the standard sequent calculus for intuitionistic logic are given in the Appendix. We will often refer to the rules $\wedge$-L, $\vee$-L, $\exists$-L, $\forall$-L and $\supset$-L as left rules, and the rules $\wedge$-R, $\vee$-R, $\exists$-R, $\forall$-R and $\supset$-R as right rules. The $\perp$-R rule will be of little interest, as we will not be dealing with formulae which may contain $\perp$.

Definition 2.2 *We define the* uniform rule *for a formula F as follows:*

- *The uniform rule for an atom is* $\supset$-L
- *The uniform rule for* $F_1 \wedge F_2$ *is* $\wedge$-R
- *The uniform rule for* $F_1 \vee F_2$ *is* $\vee$-R
- *The uniform rule for* $\exists x F$ *is* $\exists$-R
- *The uniform rule for* $\forall x F$ *is* $\forall$-R
- *The uniform rule for* $F_1 \supset F_2$ *is* $\supset$-R

We say that a formula F is compound *if F is not an atom.*

Definition 2.3 *A proof Ξ is* uniform *if for each non-initial sequent $\Gamma \longrightarrow F$ in Ξ where F is a compound formula, the rule used to derive $\Gamma \longrightarrow F$ is the uniform rule for F.*

It should be clear that the following proposition holds.

Proposition 2.1 *Let F be a formula containing no negations, and let Γ be a set of such formulae.*

Then

- $\Gamma \vdash_u F_1 \vee F_2$ *iff* $\Gamma \vdash_u F_1$ *or* $\Gamma \vdash_u F_2$
- $\Gamma \vdash_u F_1 \wedge F_2$ *iff* $\Gamma \vdash_u F_1$ *and* $\Gamma \vdash_u F_2$
- $\Gamma \vdash_u \exists x F$ *iff* $\Gamma \vdash_u F[t/x]$ *for some* $t \in \mathcal{U}$
- $\Gamma \vdash_u \forall x F$ *iff* $\Gamma \vdash_u F[y/x]$ *where y is not free in Γ or F*
- $\Gamma \vdash_u F_1 \supset F_2$ *iff* $\Gamma, F_1 \vdash_u F_2$

It is not hard to show that $P \vdash_o G$ iff $P \vdash_u G$; for more details, the reader is referred to [16,6]).

Our interest in hereditary Harrop formulae is due to the fact that uniform proofs are complete with respect to intuitionistic logic for this class of formulae, rather than due to a desire to implement a particular style of theorem prover for intuitionistic logic. The notion of uniform proof is a stronger requirement than intuitionistic proof; for example, $\exists x p(x) \vdash_I \exists x p(x)$, but there is no uniform proof of the sequent $\exists x p(x) \longrightarrow \exists x p(x)$. In this way we are more interested in the strength of our conclusions than a particular proof system.

3 Uniform Provability and Definiteness

Whilst the restriction to hereditary Harrop formulae is sufficient to guarantee the existence of uniform proofs, a natural question to ask is whether this restriction is necessary. As mentioned above, it was shown in [16] that if the antecedent is a set of definite formulae and the consequent a goal formula, then the sequent has a proof iff it has a uniform proof. The converse to this result is not (strictly) true. For example, $p(a) \vee p(b), (\exists x p(x) \supset q) \vdash_u q$, but the antecedent is not a set of definite formulae. Similarly, $\exists x p(x), \forall x(p(x) \supset q) \vdash_u q$, but the antecedent is not a set of definite formulae.

Hence, it is not strictly true that for $F_1 \vdash_u F_2$ to hold we must have that F_1 is a definite formula. However it seems that the first uniform proof above relies on the fact that $p(a), (\exists x p(x)) \supset q \vdash_u q$ and $p(b), (\exists x p(x)) \supset q \vdash_u q$, in which both the antecedents are definite formulae. Similarly the second uniform proof above is dependent on the fact that the universally quantified variable may be replaced by any term, and hence the proof may be thought of as a template for a number of proofs of sequents of the form $p(t), \forall x(p(x) \supset q) \longrightarrow q$ for any term t. In this way there seems to be a more subtle relationship between uniform proofs and hereditary Harrop formulae. Indeed, as mentioned above, it is not possible to give a strict classification of the largest class of formulae for which uniform proofs are complete, but it does seem that there is a relationship between uniform proofs and hereditary Harrop formulae which may be elucidated.

A result reported in [15] is that for sequents of the form $\Gamma \longrightarrow G$ where Γ is a set of definite formulae, there are no occurrences of the $\exists$-L or $\vee$-L rules. Hence, if there is an intuitionistic proof of a sequent in which the antecedent is a definite formula and the consequent a goal formula, then there are no occurrences of the $\exists$-L and $\vee$-L rules and the sequent has a uniform proof. Thus we may conject that if a uniform proof of $F_1 \longrightarrow F_2$ contains no occurrences of either of these rules, then F_1 is a definite formula and F_2 is a goal formula. This again is not true, as there may be parts of the formula F_1 which ensure that F_1 is not a definite formula, but are not used in the proof. For example, $\exists x q(x), p(a) \vdash_u p(a) \vee p(b)$, due to the fact that $p(a) \vdash_u p(a) \vee p(b)$, and hence $F, p(a) \vdash_u p(a) \vee p(b)$ for any formula F. This means that the relationship between a sequent $F_1 \longrightarrow F_2$ and some "equivalent" sequent $D \longrightarrow G$ will require more investigation. In particular, the role of the rules $\exists$-L and $\vee$-L need examination.

Note that apart from $\supset$-L, the left rules may be thought of as converting the antecedent into a desired form so that the appropriate right rules may be used. Hence, from the point of view of goal-directed provability, it will often be useful to perform these manipulations before starting the "main" proof, as it were. This will be the case if we can interchange the order of the rules when a right rule precedes a left one. It turns out that the nature of the $\exists$-L and $\vee$-L rules may make this difficult, and so it may not always be possible to re-arrange a given proof so that all the manipulation of the assertions can be done prior to the proof search process. However, there are some conditions under which this can be done.

For these reasons we introduce below the concept of a *definite* proof.

Definition 3.1 *A proof* Ξ *is* definite *if* Ξ *contains no occurrences of either the* $\exists$*-L rule or the* $\vee$*-L rule. We denote definite provability by* $\vdash_d$*.*

For this reason we will sometimes refer to the $\exists$-L and $\vee$-L rules as *indefinite* rules.

As mentioned above, it was shown in [15] that definite proofs are complete with respect to intuitionistic provability for a large fragment of hereditary Harrop formulae. Below we state the generalisation of this result for hereditary Harrop formulae.

Proposition 3.1 *Let* Γ *be a set of definite formulae, and let* G *be a goal. Then any proof* Ξ *of* $\Gamma \longrightarrow G$ *is definite.*

Note that it is not true that $\Gamma \vdash_u F \Rightarrow \Gamma \vdash_d F$, as when the succedent is just an atom we may use either $\vee$-L or $\exists$-L without violating the uniformity property. However, the converse is true, i.e. that if $\Gamma \vdash_d F$, then $\Gamma \vdash_u F$. In other words, a sequent with a definite proof has a uniform proof, but a uniform proof need not be definite.

Theorem 3.2 *Let* F *be an negation-free formula, and let* Γ *be a set of negation-free formulae. Then*

$$\Gamma \vdash_d F \Rightarrow \Gamma \vdash_u F$$

This result may be established by using the permutation properties of intuitionistic logic, as determined by Kleene [11]; space prevents us from giving the proof here.

The above theorem may be thought of as showing that if we ignore the $\exists$-L and $\vee$-L rules, then we need only consider uniform proofs. Note also the strength of the contrapositive of the theorem, i.e. that if $\Gamma \longrightarrow F$ has a proof but no uniform proof, then all proofs of $\Gamma \longrightarrow F$ contain an occurrence of an indefinite rule. Thus an obvious way to ensure the completeness of uniform proofs is to restrict the class of formulae so that the indefinite rules become redundant.

One such class of formulae are definite formulae, and the redundance of the indefinite rules for definite formulae is precisely why definite formulae are interesting. Definite formulae seem very apt in this context, as they force the programmer to present his or her knowledge in a relatively strong way. We may think of an indefinite formula as conveying less information than a definite one. For example, the formula $\exists x p(x)$ carries less information than the formula $p(t)$, which may be used to derive the former one. Indeed, if we may imagine an intuitionistic programmer asserting that $\exists x p(x)$ is true, we may expect him to be able to construct a term t such that $p(t)$ is true. In fact this is a requirement if we insist upon goal-directed provability, as $\exists x p(x) \vdash_I \exists x p(x)$, but we cannot derive the truth of any instance of $p(x)$. Hence we may imagine a compiler taking as input a set of formulae, and retaining only the definite parts of the formulae, as the indefinite parts do not provide us with enough information to make them useful.

In this way it seems that there is a strong connection between definite proofs and definite formulae, which is that given a definite proof of $\Gamma \longrightarrow F$, we may extract a set of definite formulae and a goal formula from the sequent, in the manner briefly described above. A more precise description is given below. We denote by $\top$ the formula "true".

Definition 3.2 *Let F be an negation-free formula. Then we define*

$$
\begin{array}{ll@{\qquad}ll}
\mathrm{def}(A) & = A & \mathrm{goal}(A) & = A \\
\mathrm{def}(F_1 \wedge F_2) & = \mathrm{def}(F_1) \wedge \mathrm{def}(F_2) & \mathrm{goal}(F_1 \wedge F_2) & = \mathrm{goal}(F_1) \wedge \mathrm{goal}(F_2) \\
\mathrm{def}(F_1 \vee F_2) & = \top & \mathrm{goal}(F_1 \vee F_2) & = \mathrm{goal}(F_1) \vee \mathrm{goal}(F_2) \\
\mathrm{def}(\exists x F) & = \top & \mathrm{goal}(\exists x F) & = \exists x\, \mathrm{goal}(F) \\
\mathrm{def}(\forall x F) & = \forall x\, \mathrm{def}(F) & \mathrm{goal}(\forall x F) & = \forall x\, \mathrm{goal}(F) \\
\mathrm{def}(F_1 \supset F_2) & = \mathrm{goal}(F_1) \supset \mathrm{def}(F_2) & \mathrm{goal}(F_1 \supset F_2) & = \mathrm{def}(F_1) \supset \mathrm{goal}(F_2)
\end{array}
$$

We also define def$(\{F_1, \ldots F_n\}) = \bigcup_{i=1}^{n} \{\mathrm{def}(F_i)\}$.

Note that def(F) is either $\top$ or a definite formula, and that goal(F) is a goal formula. Note also that goal(F) can never be $\top$. We thus arrive at the following useful lemma.

Lemma 3.3 *Let F be an negation-free formula. Then*

1. $F \vdash_I \mathrm{def}(F)$
2. $\mathrm{goal}(F) \vdash_I F$

Hence we see that def(F) and goal(F) preserve certain information, in that anything deducible from def(F) is deducible from F, and that anything deducible from F is deducible from goal(F). In addition, as shown below, the converse relationships hold for uniform provability.

Proposition 3.4 *Let F be an negation-free formula, and let Γ be a set of negation-free formulae.*

If $\Gamma \vdash_d F$, then there is a set of definite formula Γ' and a goal formula G such that

1. $\Gamma \vdash_I \bigwedge \Gamma'$
2. $G \vdash_I F$
3. $\Gamma' \vdash_u G$

Thus if $\Gamma \longrightarrow F$ has a definite proof, then not only does the same sequent have a uniform proof, but also we may extract a set of definite formulae Γ' from Γ such that $\Gamma \vdash_I \Gamma'$ and $\Gamma' \vdash_u F$, and a goal formula G from F such that $G \vdash_I F$ and $\Gamma \vdash_u G$. In this way we may think of this result as a version of Craig's Interpolation theorem [1]), in that given a proof of $\Gamma \longrightarrow F$, then provided that there are no occurrences of $\exists$-L or $\vee$-L in the proof, then we can interpolate a definite formula D such that $\Gamma \vdash_I D$ and $D \vdash_I F$. Thus given Γ, we can derive a definite formula which is provable from Γ and has the same consequences, provided that we consider only definite proofs. Hence, definite formulae arise naturally out of consideration of definite proofs, which in turn arise naturally out of consideration of the permutability of the left and right rules in intuitionistic logic.

4 Maximality of Information and Definiteness

The result above may be interpreted as showing what efficiencies we can make in the process of searching for a proof provided that we restrict our attention to definite proofs. As described above, we may think of this in a similar manner to the Interpolation theorem. A criticism which may be made of this approach is that whilst indefinite formulae may contain less information than definite ones, that information is lost when the indefinite parts of the formulae are ignored. Also, the requirement that the proof be definite is a stronger one than merely requiring the proof to be uniform. Hence it may be interesting to examine what may be done to preserve (or strengthen) the original information rather than weakening it, and to see if uniform proofs are still sufficient in these circumstances.

An obvious alternative approach to extracting definite information from a proof is to find a definite formula of which the premise is a consequence, rather than a definite formula which is a consequence of the premise. We may think of this approach as attempting to supply sufficient information in order to make the formula definite, rather than ignoring the indefinite parts of the formula, and hence we will be suggesting hypotheses which will make the formula true. This leads us to the concept of a *definite condition* and a *definite consequence.*

Definition 4.1 *A* definite condition *of a formula F is a formula which is the same as F except that*

1. *Every positive occurrence of a subformula $\exists x F'$ in F is replaced by $F'[t/x]$ for some term t in which all variables of t appear universally quantified elsewhere in F outside the scope of $\exists x$.*

2. *Every positive occurrence of a subformula $F_1 \vee F_2$ in F is replaced by one of F_i, $i = 1, 2$*

We denote by defprem(F) the set of all definite conditions of F.

If Γ is a set of formulae, then D is a definite condition of Γ if D is a conjunction of definite conditions of each element of Γ.

Definition 4.2 *A* definite consequence *of a formula F is a formula which is the same as F except that*

1. *Every negative occurrence of a subformula $\exists x F'$ in F is replaced by $F'[t/x]$ for some term t in which all variables of t appear universally quantified elsewhere in F outside the scope of $\exists x$.*

2. *Every negative occurrence of a subformula $F_1 \vee F_2$ in F is replaced by one of F_i, $i = 1, 2$*

We denote by defconc(F) the set of all definite consequences of F.

If Γ is a set of formulae, then G is a definite consequence of Γ if G is a conjunction of definite consequences of each element of Γ.

Note that a definite condition of $\exists x F$ cannot contain any occurrence of x, and hence must produce a "ground witness" for x. For example, the only definite conditions of $\exists x\ p(x)$ are atoms of the form $p(t)$ where t is a ground term.

For existentially quantified variables appearing within the scope of a universally quantified variable, we may use the universally quantified variable to construct the witness. For example, one of the definite conditions of $\forall x \exists y\ p(x, y)$ is $\forall x\ p(x, f(x))$.

It should be clear that for negation-free formulae, definite conditions and definite consequences are definite and goal formulae respectively.

It is not hard to show that definite conditions and definite consequences behave in the expected manner.

Proposition 4.1 *Let F be an negation-free formula. Then for any definite condition D of F and definite consequence G of F*

1. $D \vdash_I F$

2. $F \vdash_I G$

We may think of this as stating that D has more explicit information than F, so that if we were to consider an ordering of formulae in which $F_1 \leq F_2$ iff $F_2 \vdash_I F_1$, then the above proposition ensures that for any F, there is always a definite formula D such that $F \leq D$. Similar remarks apply to G, in that there is always a G such that $G \leq F$. In this way if we think of a lattice of formulae in which the partial order is (intuitionistic) provability, then any chain has a least upper bound which is a definite formula, and a greatest lower bound which is a goal formula. Thus we extrapolate from the formula to a more definite statement.

It is not hard to show that definite conditions and definite consequences preserve uniform provability.

Proposition 4.2 *Let F be a negation-free formulae, and let Γ be a set of negation-free formulae. Then for any definite condition D of Γ and any definite consequence G of F*

1. $\Gamma \vdash_u F \Rightarrow D \vdash_u F$
2. $\Gamma \vdash_u F \Rightarrow \Gamma \vdash_u G$

We may think of the above proposition as a form of "extrapolation" result, in that given a uniform proof of $\Gamma \longrightarrow F$ we can find a definite formula D and a goal formula G such that $D \vdash_I \bigwedge\Gamma$, $D \vdash_u F$, $F \vdash_I G$ and $\Gamma \vdash_u G$, and as a consequence, $D \vdash_u G$. Thus given any uniform proof, we can find a definite formula D and a goal formula G which preserve the appropriate provability relationships. Hence we may conclude that this result supports our contention that hereditary Harrop formulae are maximal with respect to uniform proofs, in that any sequent which has a uniform proof may be thought of as an incomplete specification of a sequent $D \longrightarrow G$ which preserves the provability properties of the original sequent.

5 Conclusions and Further Work

We have seen how restricting first-order intuitionistic proofs in certain ways leads to some results which ensure that the task of searching for a proof is made more feasible than in the general case. We may think of the restrictions as ensuring that the information contained in the formulae is presented in a maximal way, so that we do not need to waste time discovering this information during the computation process. This may be thought of as requiring that we only consider proofs in a "normal form".

One way to think of this maximal class is to consider it as a "definite" or "uniform" sub-logic of intuitionistic logic, with a more restricted notion of provability. In particular, this gives us a notion of *constructive consequence*, i.e. that the following properties hold:

$$\Gamma \vdash \exists x F \Leftrightarrow \Gamma \vdash F[t/x]$$
$$\Gamma \vdash F_1 \vee F_2 \Leftrightarrow \Gamma \vdash F_1 \text{ or } \Gamma \vdash F_2$$

Note that intuitionistic logic alone is not sufficient to guarantee these equivalences (unless, of course, Γ is empty). However the following equivalences do hold in intuitionistic logic:

$$\Gamma \vdash \forall x F \Leftrightarrow \Gamma \vdash F[y/x]$$
$$\Gamma \vdash F_1 \wedge F_2 \Leftrightarrow \Gamma \vdash F_1 \text{ and } \Gamma \vdash F_2$$
$$\Gamma \vdash F_1 \supset F_2 \Leftrightarrow \Gamma, F_1 \vdash F_2$$

where y is not free in Γ or F.

Hence we see that imposing constructive consequence on intuitionistic logic gives us precisely goal-directed provability. Alternatively, imposing goal-directed provability on intuitionistic logic gives us constructive consequence. Thus we may think of hereditary Harrop formulae as an important fragment of intuitionistic logic, in that they seem to be the largest class of formulae for which the notion of constructive consequence, and hence goal-directed provability, can be guaranteed. In fact, the natural logic in which to interpret hereditary Harrop formulae is slightly stronger than intuitionistic logic; see [5,7] for details.

We have also seen some relationship between the restricted classes of proofs and formulae and the more general classes, and in particular how definite formulae and goals may be extracted from an arbitrary uniform proof, and that the extracted formulae preserve uniform provability. It is possible that this result may be useful for program specification, in that if a specification is given as a first-order formula (without negation), then the extraction process described above may be thought of as finding a definite formula (i.e. a program) which satisfies the specification.

6 Acknowledgements

My thanks go to Dale Miller for many interesting discussions and weighty deliberations. Discussions with David Pym have also been enlightening, and comments from some anonymous referees were very helpful.

This work has been made possible by a grant of the Australian Research Council through the Machine Intelligence Project.

7 References

1. G.S. Boolos and R.C. Jeffrey, *Computability and Logic*, Cambridge University Press, 1980.

2. W.F. Clocksin and C.S. Mellish, *Programming in Prolog*, Springer-Verlag, 1984.

3. H.B. Curry, The Permutability of Rules in the Classical Inferential Calculus, *Journal of Symbolic Logic* 17, 245-8, 1952.

4. D. Gabbay and U. Reyle, N-Prolog: An Extension of Prolog with Hypothetical Implications. I., *Journal of Logic Programming* 1:319-355, 1984.

5. J. Harland, An Intermediate Logic for Logic Programs, Technical Report 90/29, Department of Computer Science, University of Melbourne, 1990.

6. J. Harland, A Proof-Theoretic Analysis of Logic Programming, Technical Report 90/21, Department of Computer Science, University of Melbourne, 1990.

7. J. Harland, *On Hereditary Harrop Formulae as a Basis for Logic Programming*, Ph.D. Thesis, Department of Computer Science, University of Edinburgh, July, 1991.

8. J. Harland and D. Pym, The Uniform Proof-theoretic Foundation of Linear Logic Programming, *Proceedings of the International Logic Programming Symposium*, San Diego, October, 1991.

9. J. Harland and D. Pym, The Uniform Proof-theoretic Foundation of Linear Logic Programming, Report ECS-LFCS-90-124. University of Edinburgh, 1990. Also published as Technical Report 90/26, Department of Computer Science, University of Melbourne.

10. S.C. Kleene, *Introduction to Metamathematics*, North-Holland, 1952.

11. S.C. Kleene, Permutability of Inferences in Gentzen's Calculi LK and LJ, *Memoirs of the American Mathematical Society* 10, 1952.

12. J.W. Lloyd, *Foundations of Logic Programming*, Springer-Verlag, Berlin, 1984.

13. L.T. McCarty, Clausal Intuitionistic Logic I. Fixed Point Semantics, *Journal of Logic Programming* 5:1:1-32, 1988.

14. L.T. McCarty, Clausal Intuitionistic Logic II. Tableau Proof Procedures, *Journal of Logic Programming* 5:2:93-132, 1988.

15. D.A. Miller, A Logical Analysis of Modules in Logic Programming, *Journal of Logic Programming* 6:79-108, 1989.

16. D.A. Miller, G. Nadathur, F. Pfenning and A. Scedrov, Uniform Proofs as a Foundation for Logic Programming, *Annals of Pure and Applied Logic* 51:125-157, 1991.

17. G. Nadathur and D.A. Miller, Higher-Order Horn Clauses *Journal of the Association for Computing Machinery* 37:4:777-814, October, 1990.

18. P. Schroeder-Heister, Hypothetical Reasoning and Definitional Reflection in Logic Programming, *Extensions of Logic Programming: International Workshop, Tübingen FRG, December 1989*, P. Schroeder-Heister (ed.), Lecture Notes In Artificial Intelligence 475, 327-340, Springer-Verlag, 1991.

19. L. Sterling and E. Shapiro, *The Art of Prolog*, MIT Press, 1986.

20. S.-A. Tärnlund, Horn Clause Computability, *BIT* 17:215-226, 1977.

A Intuitionistic Sequent Calculus

$$\frac{B, C, \Gamma \longrightarrow F}{B \wedge C, \Gamma \longrightarrow F} \wedge\text{-L} \qquad \frac{\Gamma \longrightarrow B \quad \Gamma \longrightarrow C}{\Gamma \longrightarrow B \wedge C} \wedge\text{-R}$$

$$\frac{B, \Gamma \longrightarrow F \quad C, \Gamma \longrightarrow F}{B \vee C, \Gamma \longrightarrow F} \vee\text{-L} \qquad \frac{\Gamma \longrightarrow B}{\Gamma \longrightarrow B \vee C} \; \frac{\Gamma \longrightarrow C}{\Gamma \longrightarrow B \vee C} \vee\text{-R}$$

$$\frac{\Gamma \longrightarrow B \quad C, \Gamma \longrightarrow F}{B \supset C, \Gamma \longrightarrow F} \supset\text{-L} \qquad \frac{B, \Gamma \longrightarrow C}{\Gamma \longrightarrow B \supset C} \supset\text{-R}$$

$$\frac{\Gamma, B[t/x] \longrightarrow F}{\Gamma, \forall x B \longrightarrow F} \forall\text{-L} \qquad \frac{\Gamma \longrightarrow B[y/x]}{\Gamma \longrightarrow \forall x B} \forall\text{-R}$$

$$\frac{\Gamma, B[y/x] \longrightarrow F}{\Gamma, \exists x B \longrightarrow F} \exists\text{-L} \qquad \frac{\Gamma \longrightarrow B[t/x]}{\Gamma \longrightarrow \exists x B} \exists\text{-R}$$

$$\frac{\Gamma \longrightarrow \perp}{\Gamma \longrightarrow B} \perp\text{-R}$$

The rules $\forall$-R and $\exists$-L have the side condition that y is not free in Γ, B or F.

An *initial sequent* is a sequent $\Gamma \longrightarrow F$ where F is either an atomic formula or $\perp$ and $F \in \Gamma$. A *proof* for the sequent $\Gamma \longrightarrow F$ is a finite tree, constructed using the above rules, whose root is $\Gamma \longrightarrow F$ and whose leaves are initial sequents.

As is done in [15], we omit the interchange and contraction rules by considering the antecedents of sequents to be sets. Note also that thinning is not necessary due to the way an initial sequent is defined.