COMPLETENESS THEOREMS FOR SEMANTIC RESOLUTION IN CONSEQUENCE-FINDING

J.R. Slagle, C.L. Chang, R.C.T. Lee
National Institutes of Health
Bethesda, Maryland

Summary

The object of consequence-finding is to deduce logical consequences from a set of axioms. The theory of J.R. Slagle's semantic resolution principle, an inference rule for first-order predicate calculus, is extended to consequence-finding. Given an interpretation I, it is proved that any prime (non-trivial) consequence, which is false in I, can be derived from a set of axioms by applying I-(semantic) resolution.

1. Introduction

The consequence-finding problem which has been described by Lee {2} may have the following form: Given statements $A_1, \ldots, A_n$, find a statement B such that B follows from $A_1, \ldots, A_n$. It may well be that $A_1, \ldots, A_n$ are the axioms or postulates of some theory, and B is a new logical consequence of that theory possible not found before. A proof-finding problem is a problem which is given a theorem and asked to find a proof for the theorem [VJ. In this paper, we shall restrict ourselves to first-order predicate calculus, and consider only the cases where $A_1, \ldots, A_n$ and B are (quantifier-free) clauses [1,4,5,6]. Clearly, under this formulation, if B is the empty clause, consequence-finding reduces to proof-finding. Hence, consequence-finding may be considered more general than proof-finding.

Because of the relation between consequence-finding and proof-finding, most of the strategies [3-7,9,10] used in proof-finding can also be used in consequence-finding. In [2] Lee proved that Robinson's ordinary resolution principle [4], which is for proof-finding is also complete for consequence-finding. The present paper wil_l extend the semantic resolution of Slagle [7] to consequence-finding, and prove that given an interpretation I, any prime consequence B, which is false in I, can be derived from a set of axioms (or postulates) $A_1, \ldots, A_n$ by applying I-(semantic) resolution. In order to be more in keeping with standard practice, we use "interpretation" rather than "model," which was used in [7]. Our theorems are more general than Lee's results. The definitions, theorems, and proofs in this paper follow closely those given in [7]. It is also assumed that the reader is familiar with Robinson's review article [6].

2. Prime (Non-trivial) Consequences

In this section, we shall define a prime consequence, and use this concept to explain completeness theorems in consequence-finding. Roughly speaking, we shall prove below completeness theorems which show that ordinary resolution will generate all the prime consequences of a set of clauses and that semantic resolution will generate all the prime consequences which are false in the interpretation used in the semantic deduction. It is possible that both kinds of resolutions may or may not generate some non-prime consequences.

Definition* An interpretation of a set S of clauses consists of a nonempty set U called a universe (sometimes, to emphasize the universe U, we say an interpretation over U.) and an assignment of "values" as follows to each individual symbol, function symbol and predicate symbol occurring in S: (i) to each individual symbol, we assign an element in U; (ii) to each function symbol with n arguments, we assign a function whose n variables range over U and whose values are in Uj (iii) to each predicate symbol with n arguments, we assign a function whose n variables range over U and whose values are the truth values, true and false. When S is a set of ground (variable-free) clauses, we call an interpretation of S a ground interpretation.

Definition. A ground instance of a clause C over a universe U is a ground clause obtained by substituting elements of U for all the variables in C.

We note that if U is the Herbrand universe H of S, then an interpretation of S over H may be represented by a sequence $m_1, m_2, \ldots, m_i, \ldots$ in which m_i is either A_i or $\sim A_i$ where $A_1, A_2, \ldots$ are the distinct ground instances of atoms of S over H.

Definition. An interpretation I of a set S of clauses over a universe U is said to satisfy a clause C in S (or C is true in I) iff every ground instance of C over U is true in I. Otherwise, I is said to falsify C (or C is false in I).

Definition* A set S of clauses is said to be satisfiable iff there is an interpretation I of S such that I satisfies all the clauses in S. Otherwise, S is said to be unsatisfiable.

Definition* A clause C is said to be a logical consequence of a set S of clauses $C_1, \ldots, C_n$ (or C follows from $C_1, \ldots, C_n$), iff for every interpretation I of (S & C) if I satisfies $C_1, \ldots, C_n$, I also satisfies C. When no confusion is possible, we shall sometimes say that C is a logical consequence of $C_1, \ldots, C_n$. We note that if C is a logical consequence of S, every instance of C is also a logical consequence of S.

It is noted that if a clause B is a logical consequence of a clause A, and if B and A have the corresponding well-formed formulas B' and A' respectively, then B' is a logical consequence of A', but not conversely. For example, let $A' = \exists x \forall y Pxy$ and $B' = \forall y \exists x Pxy$. The corresponding respective clauses of A' and B' are A = Pay and B = Pf(y)y. Clearly, B' is a logical consequence of A', but B is not a logical consequence of A.

We now prove a generalization of Herbrand's theorem.

Theorem 1. A ground clause C is a logical consequence of a set S of clauses iff there is a finite set S' of ground instances of clauses in S over the Herbrand universe of (S & C) such that C is a logical consequence of S'.

Proof. C is a logical consequence of S.
$\Leftrightarrow$ Every interpretation which satisfies S must satisfy C.
$\Leftrightarrow$ Every interpretation must falsify ($S \,\&\, \sim C$), where $\sim C$ is the negation of C.
$\Leftrightarrow$ ($S \,\&\, \sim C$) is unsatisfiable.
$\Leftrightarrow$ By the Herbrand theorem, there is a finite set S' of ground instances of clauses in S over the Herbrand universe of (S & C) such that ($S' \,\&\, \sim C$) is unsatisfiable.
$\Leftrightarrow$ Every interpretation which satisfies S' must falsify $\sim C$ (or satisfy C).
$\Leftrightarrow$ C is a logical consequence of S'.

We note that Theorem 1 cannot be further generalized to the case where C is a clause. A counterexample is that $S = \{P(x)\}$ and $C = P(f(x))$.

Definition. Let S be a set of clauses. A clause C is said to be a prime* (non-trivial) consequence of S iff C is a logical consequence of S and there exists no other logical consequence D of S such that C is a logical consequence of D. A logical consequence of S is a non-prime (trivial) consequence of S iff it is not a prime consequence of S.

We note that if a set S of clauses is unsatisfiable, the empty clause is the only prime consequence of S.

Definition. If C and D are two clauses, we say that C subsumes D iff there is a substitution σ such that $C\sigma \subseteq D$. Clearly, if a ground clause C subsumes a ground clause D, then $C \subseteq D$.

We note that if C subsumes D, D is a logical consequence of C. We also note that there is a procedure which can decide whether or not a clause C subsumes a clause D.

* We note that a prime consequence of a set of ground clauses is a prime implicate of the set and is the dual of a prime implicant of the dual of the set [8].

3. Completeness Theorem for Propositional Calculus

We now state and prove our basic results in Theorem 2 and 3. Theorem 2 and 3 of the present paper are the respective extensions of Theorem 5 and 6 of Slagle's paper [7].

Theorem 2. If a ground clause C is a logical consequence of a finite set S of ground clauses but is not subsumed by any one clause in S and if C contains no negative literals after a renaming r and if the first atoms in an ordering A of the atoms in S & C are the atoms in C, then there exists an unresolved maximal Ar-clash $\{E_1, E_2, \ldots, E_q, N\}$ in S with nucleus N such that any set consisting of N and one or more of the electrons $E_1, E_2, \ldots, E_q$ is an unresolved Ar-clash in S.

Proof. The following proof follows closely that of Slagle [7].

Let R be the renamable subset corresponding to the renaming r. Because of renaming, we can without loss in generality assume that R is the positive subset of S, that C is a positive ground clause, and that r is the empty renaming. Let A be the sequence $A^1, A^2, \ldots, A^k$. The literals in C are the first p atoms in A. Let I_p be the sequence $\sim A^1, \ldots, \sim A^p$. We note that no R-clause consists entirely of complements of literals in I_p because no S-clause subsumes C. For $j = p + 1, \ldots, k$, form the sequence I_j as follows:

$$I_j = \begin{cases} I_{j-1}, A^j & \text{if some R-clause consists entirely of complements of literals in the sequence } I_{j-1}, \sim A^j; \\ I_{j-1}, \sim A^j & \text{if no R-clause consists entirely of complements of literals in the sequence } I_{j-1}, \sim A^j. \end{cases}$$

Let I be I_k. From the construction, it is obvious that the ground interpretation I satisfies ($R \,\&\, \sim C$). In addition it is clear that, for every positive literal A^i which appears in I, there is an R-clause F_i such that A^i is the largest atom in F_i and such that all F_i-literals other than A^i are false in I. Since C is a logical consequence of S, ($S \,\&\, \sim C$) is truth-functionally unsatisfiable. Since I satisfies ($R \,\&\, \sim C$), I must falsify one or more clauses in S − R. Of all the S-clauses which I falsifies, let N be one of those which contains the fewest negative literals. Let $\sim B_1, \sim B_2, \ldots, \sim B_q$ be all the negative literals in N. Let $i = 1, 2, \ldots, q$. For each i, B_i is true in I. Hence for each B_i, there is an R-clause E_i such that B_i is the largest atom in E_i and such that all E_i − literals other than B_i are false under I. (Note that for each i, $B_i = A^j$ for some j; hence we can let E_i be F_j.) Hence no literal in E_i −

$\{B_i\}$ is B_j for $j \neq i$ (since such B_j are true in I) nor is it the complement of any literal in $N-\{\sim B_1, \sim B_2, \ldots, \sim B_q\}$ (since such literals are false in I). It follows that any set whose members are N and one or more of the clauses $E_1, E_2, \ldots, E_q$ is an AI-clash in S whose resolvent is false under I and contains fewer negative literals than does N. Therefore the resolvent of any such clash is not in S and so the clash is unresolved in S. The proof is now completed by noting that the clash $\{E_1, E_2, \ldots, E_q, N\}$ is maximal.

By Theorem 4 of Slagle [7], the following theorem is also true.

Theorem 3. If a ground clause C is a logical consequence of a finite set S of ground clauses but is not subsumed by any one clause in S and if C is false in a ground interpretation I of (S & C) and if the first atoms in an ordering A of the atoms in S & C are the atoms in C, then there exists an unresolved maximal AI-clash $\{E_1, E_2, \ldots, E_q, N\}$ in S with nucleus N such that any set consisting of N and one or more of the electrons $E_1, E_2, \ldots, E_q$ is an unresolved AI-clash in S.

Although the theory can now be extended to semantic and renamable resolution, we restrict ourselves to semantic resolution in the remainder of this paper.

Theorem 4. If a ground clause C is a logical consequence of a finite set S of ground clauses and is false in a ground interpretation I of (S & C), and if the first atoms in an ordering A of the atoms in S & C are the atoms in C, then there is a ground maximal AI-deduction from S of some ground clause T which subsumes C.

Proof. Let S_o be S. In general, for $j \geq o$, if there is no ground clause T in S_j which subsumes C, let S_{j+1} be the result of adding to S_j the ground maximal AI-resolvent obtained as in Theorem 3. This gives a nested sequence $S_o, S_1, \ldots,$ of sets of ground clauses. The sequence cannot continue to grow since the successive resolvents are all constructed from the same fixed finite set of literals, namely, those which occur in members of S. Hence there are only finitely many such resolvents which can be added. Therefore, for some $j \geq o$ we have a ground clause T in S_j such that C is subsumed by T. Tracing the ancestry of T from S then gives the required deduction.

There is a difficulty in applying Theorem 4 to deduce logical consequences from S because C is generally unknown beforehand and we cannot have an interpretation of (S & C). However, if we consider prime consequences of S, this difficulty can be easily overcome. To see this, let C be a prime consequence of S in Theorem 4. Thus, T = C, i.e., C can be derieved from S. Hence, every symbol which occurs in C occurs in S. In this case every ground interpretation of S is a ground interpretation of (S & C), and vice versa. Therefore, we have the following completeness theorem which follows from Theorem 4.

Theorem 5. If a ground clause C is a prime consequence of a finite set S of ground clauses and is false in a ground interpretation I of S, and if the first atoms in an ordering A of the atoms in S are atoms in C, then there is a ground maximal AI-deduction of C from S.

We note that if a finite set S of ground clauses is unsatisfiable, the empty clause $\square$ is the only prime consequence of S. Since $\square$ is false in every interpretation, Theorem 5 reduces to a theorem from which the previous completeness theorems given in [3-7, 9, 10] for resolution (including ordinary, hyper-, and set-of-support resolution) in proof-finding can be derived [7].

4. Completeness Theorem for First-order Predicate Calculus

The above theorems in Section 3 hold for propositional calculus. We now extend them to first-order predicate calculus.

Theorem 6. If a clause C is a logical consequence of a finite set S of clauses and is false in an interpretation I of (S & C), then there is an I-(semantic) deduction from S of some clause T which subsumes C.

Proof. Let $x_1, \ldots, x_m$ be the variables occurring in C. To concentrate on the variables in C, we shall represent C as $C(x_1,\ldots,x_m)$. Let U be the universe over which I is defined. Since $C(x_1, \ldots,x_m)$ is false in I, there are elements $a_1,\ldots, a_m$ in U such that the ground instance $C(a_1,\ldots,a_m)$ of C is false in I. Let $b_1,\ldots,b_m$ be new distinct individual symbols not occurring in (S & C) and let $H(b_1,\ldots,b_m)$ be the Herbrand universe of $[S \& C(b_1,\ldots,b_m)]$. We define an interpretation I_H of (S & C) over $H(b_1,\ldots,b_m)$ as follows: A ground literal L over $H(b_1,\ldots,b_m)$ is in I_H iff the ground literal obtained from L by substituting a_i for b_i, for $i = 1,\ldots,m$, is true in I. (We note that the definition for I_H is self-consistent.) Clearly, by this definition, $C(b_1,\ldots,b_m)$ is false in I_H. Since $C(x_1,\ldots,x_m)$ is a logical consequence of S, $C(b_1,\ldots,b_m)$ is a logical consequence of S. By Theorem 1, there is a finite set S' of ground instances of clauses in S over $H(b_1,\ldots,b_m)$ such that $C(b_1,\ldots,b_m)$ is a logical consequence of S'. Let I_H' be the ground interpretation of $[S' \& C(b_1, \ldots,b_m)]$ corresponding to the interpretation I_H. Since $C(b_1,\ldots,b_m)$ is a logical consequence of S' and is false in I_H' and there exists an ordering

A as required by Theorem 4, by Theorem 4 there is a ground maximal AI_H-deduct ion from S' of some ground clause T'which subsumes $C(b_1,\ldots,b_m)$. For this deduction tree Tr, we attach to each node of Tr a further clause over and above the clause already there as follows: To each initial node of Tr, attach the corresponding S-clause of which the clause already there is an instance. Then for each non-initial node of Tr, if clauses have been attached in this way to each of its immediate predecessor nodes and constitute the set W, attach to it that I_H-resolvent of the latent I_H-clash W of which the clause already there is an instance. In this fashion, a clause is attached to each node of which the clause already at the node is an instance. Let the clause attached to the terminal node be T. Since each I_H-clash is an I-clash, the deduction of T is an I-deduction. Since T does not contain the individual symbols $b_1,\ldots,b_m$ for T is derived from S, and T' which subsumes $C(b_1,\ldots,b_m)$ is an instance of T, there must exist a substitution $\sigma(b_1,\ldots,b_m)$ such that $T\sigma(b_1,\ldots,b_m) = T' \subseteq C(b_1,\ldots,b_m)$. So that, $T\sigma(x_1,\ldots,x_m) \subseteq C(x_1,\ldots,x_m)$, i.e. T subsumes C. This completes our proof.

As in ground case, in Theorem 6, if C is a prime consequence of S, then T = C. Since C can be derived from S, every symbol which occurs in C must occur in S. Therefore, we can use interpretations of S, instead of interpretations of (S & C), in the following completeness theorem for first-order predicate calculus.

Theorem 7. If a clause C is a prime consequence of a finite set S of clauses, and is false in an interpretation I of S, then there is an I-(semantic) deduction of C from S.

The following theorem follows directly from Theorem 7:

Theorem 8, If a clause C is a prime consequence of a finite set S of clauses and contains no negative (positive) literals, then there is a positive (negative) hyper-deduction of C from S.

We now consider the theorems for unit clauses because prime consequences which are unit clauses are interesting. In this case, we can use PI-deduction instead of I-deduction as stated in the following.

Theorem 9. If a unit clause C is a prime consequence of a finite set S of clauses and is false in an interpretation I of S and if the first predicate symbol in an ordering P of the predicate symbols occurring in S is the predicate symbol in D, then there is a Pi-deduction of C from S.

The proof is the same as the proof for Theorem 6 except we carry the ordering of the predicate symbols from the ground case to the latent case.

The following theorem is a corollary of Theorem 9:

Theorem 10. If a positive (negative) unit clause C is a prime consequence of a finite set S of clauses and if the first predicate symbol in an ordering P of the predicate symbols in S is the predicate symbol in C, then there is a positive (negative) P-hyper-deduction of C from S.

We now give an example to illustrate our results.

Example. In an associative system, if there are left and right solutions, then we can derive that there is a right identity.

The system is expressed in three clauses:

AI: $\sim P(x,y,u) \lor \sim P(y,z,v) \lor \sim P(x,v,w) \lor P(u,z,w)$ part of associativity

A2: $P(g(x,y),x,y)$ Existence of a left solution

A 3 $P(x,h(x,y),y)$ tence of a right solution

Since we know beforehand that the theorem concerning the right identity is a positive unit, we use a negative interpretation to guide our deduction. The electrons in the deduction will be positive clauses and the deduction uses positive hyper-resolution.

R1: $\sim P(y,z,v) \lor \sim P(g(y,u),v,w) \lor P(u,z,w)$ from $A1_1$ and $A2_1$

R2: $\sim P(v,z,v) \lor P(w,z,w)$ from $R1_2$ and $A2_1$

R3: $P(w,h(v,v),w)$ from $R2_1$ and $A3_1$

R3 means that there is a right identity.

5. Conclusion

In the preceding sections, we have given several completeness theorems for consequence-finding. Using I-(semantic) resolution, from a finite set of axioms, we can generate at least the prime consequences which are false in I. Semantic resolution reduces the number of clauses we have to generate in order to derive these prime consequences •

In consequence-finding, it is difficult to find a criterion for interesting consequences. In this paper, we have given the definition of a prime consequence. We believe that many interesting consequences may be obtained from prime consequences. However, much work needs to be done in this area.

In addition to generating interesting consequences, consequence-finding may also be relevant to proof-finding. The theoretical framework given tere reveals part of the behavior of semantic resolution. It shows that semantic resolution generates prime consequences of any subset of a set of clauses. This gives us more insight into semantic resolution and may prove to be useful in proof-finding. For example, in proof-finding, we are interested in deducing a pair of contradictory units. From Theorem 10, these contradictory units can be deduced by using positive and negative P-hyper-resolution simultaneously. We may call the strategy that uses positive and negative P-hyper-resolution simultaneously the "bridging strategy." In most cases, positive hyper-resolution corresponds to thinking forward from axioms whereas negative hyper-resolution corresponds to thinking backward from conclusion. This is a good sign because people reason in both directions when proving difficult theorems. Therefore, a program that implements the bridging strategy corresponds to a theorem prover who thinks forward and backward simultaneously. The bridging strategy may be useful for theorems which require long proofs. Since unit clauses are either positive or negative the unit section of the unit preference strategy, which is the most efficient strategy ever implemented on computers, is a special case of the bridging strategy.

References

1. Davis, M., and Putnam, H. A computing procedure for quantification theory. J. ACM 7 (1960), 201-215.

2. Lee, R.C.T. A completeness theorem and a computer program for finding theorems derivable from given axioms, Ph.D. dissertation, University of California, Berkeley, 1967.

3. Meltzer, B. Theorem proving for computers: some results on resolution and renaming. Computer J. 8(1966), 341-343.

4. Robinson, J.A. A machine-oriented logic based on the resolution principle. J. ACM 12, l(Jan. 1965), 23-41.

5.———. Automatic deduction with hyper-resolution. Int. J. Computer Math. 1 (1965), 227-234.

6. —.-, A review of automatic theorem proving. Proc. Symp. in Appl. Math., Amer. Math. Soc., Providence, R.I., 1967.

7. Slagle, J.R. Automatic theorem proving with renamable and semantic resolution. J. ACM 12, 4(0ct. 1967), 687-697.

8. Slagle, J.R., Chang, C.L., and Lee, R.C.T. A new algorithm for generating prime implicants. (September 1968) to appear.

9. Wos, L., Carson, D.F., and Robinson, G.A. The unit preference strategy in theorem proving. Proc. AFIPS 1964 Fall Comput. Conf., Vol. 26, 616-621.

10.———. Robinson, G.A., and Carson, D.F. Efficiency and completeness of the set of support strategy in theorem proving. J. ACM 12, 4 (Oct. 1965), 536-541.