Lecturecise 19
Proofs and Resolution
Compactness for Propositional Logic
Resolution for First-Order Logic

2013

Proof Systems

- Proof rules are computatable relations on finite sequences formulas.
- Given some number of assumptions, a proof rule produces new conclusions.
- A proof tree describes the application of proof rules
- $\Gamma \vdash F$ means that there is a proof tree with leaves Γ that derives F
- Proof steps should be computable: must be able to decide whether a rule applies and can produce a given conclusion.
- A system of logical rules is **sound** iff every conclusion that it only derives is a consequence.
- A proof system is **complete** when it can prove all properties that are true.

Proof System for Propositional Logic

- Fix a countable set of propositional variables V e.g. $p_0, p_1, \ldots$.
  All formulas have variables from V
- Propositional interpretation is a map $I : V \to \mathbb{B}$, $\mathbb{B} = \{true, false\}$
- We write $I \models F$ if formula F is true in model m
- Let Γ be a set of formulas
- $I \models \Gamma$ means $\forall F \in \Gamma.\ I \models F$
- Γ is consistent (satisfiable) if there exists I for which $I \models \Gamma$, else it is contradictory
- $\Gamma \models F$ means $\forall I.\ (I \models \Gamma) \to (I \models F)$
- Proof system "$\vdash$" is **sound** iff $\Gamma \vdash F$ implies $\Gamma \models F$
- Proof system "$\vdash$" is **complete** iff $\Gamma \models F$ implies $\Gamma \vdash F$

Propositional Resolution

$$\frac{A \vee L \qquad \neg L \vee B}{A \vee B}$$

Soundness proof:

- Let I be an interpretation in which both $I(A \vee L) = \mathit{true}$ and $I(\neg L \vee B) = \mathit{true}$
- if $I(L) = \mathit{true}$ then from $I(\neg L \vee B) = \mathit{true}$ we conclude $I(B) = \mathit{true}$, so $I(A \vee B) = \mathit{true}$
- if $I(L) = \mathit{false}$ then from $I(A \vee L) = \mathit{true}$ we conclude $I(A) = \mathit{true}$, so $I(A \vee B) = \mathit{true}$
- In any case $I(A \vee B) = \mathit{true}$.

Propositional Resolution on Clauses

Rule on formulas:

$$\frac{A \vee L \qquad \neg L \vee B}{A \vee B}$$

When we represent disjunctions as sets of literals becomes:

$$\frac{A \cup \{L\} \qquad \{\neg L\} \cup B}{A \cup B}$$

To prove that a formula is valid, we prove that its negation is contradictory by deriving an empty clause (which represents false).

Example Proof of Contradiction by Resolution

$\{\neg p, \neg q, r\}$ $\{\neg q, p\}$ $\{q, r\}$ $\{\neg r\}$

Thus, the original set of assumptions is contradictory.

Example Proof of Contradiction by Resolution

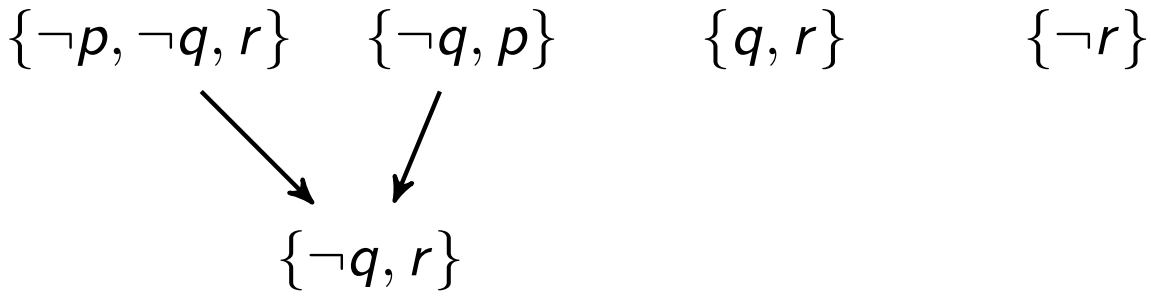

Thus, the original set of assumptions is contradictory.

Example Proof of Contradiction by Resolution

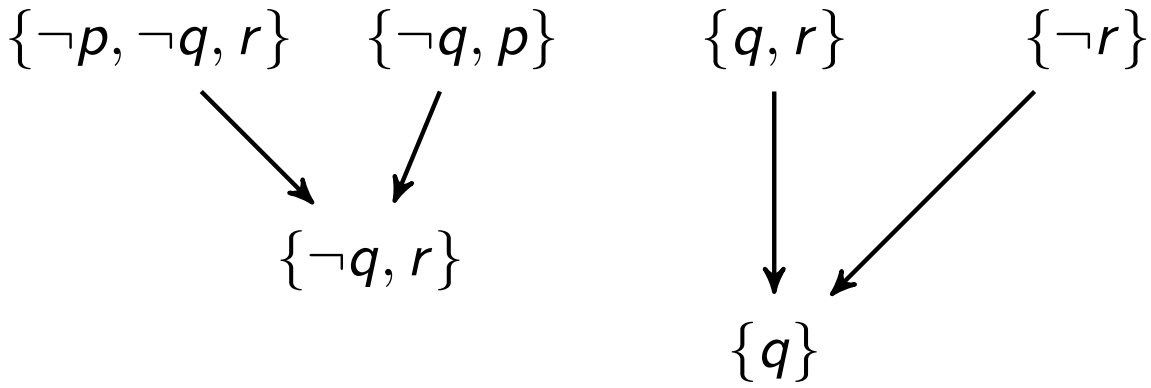

Thus, the original set of assumptions is contradictory.

Example Proof of Contradiction by Resolution

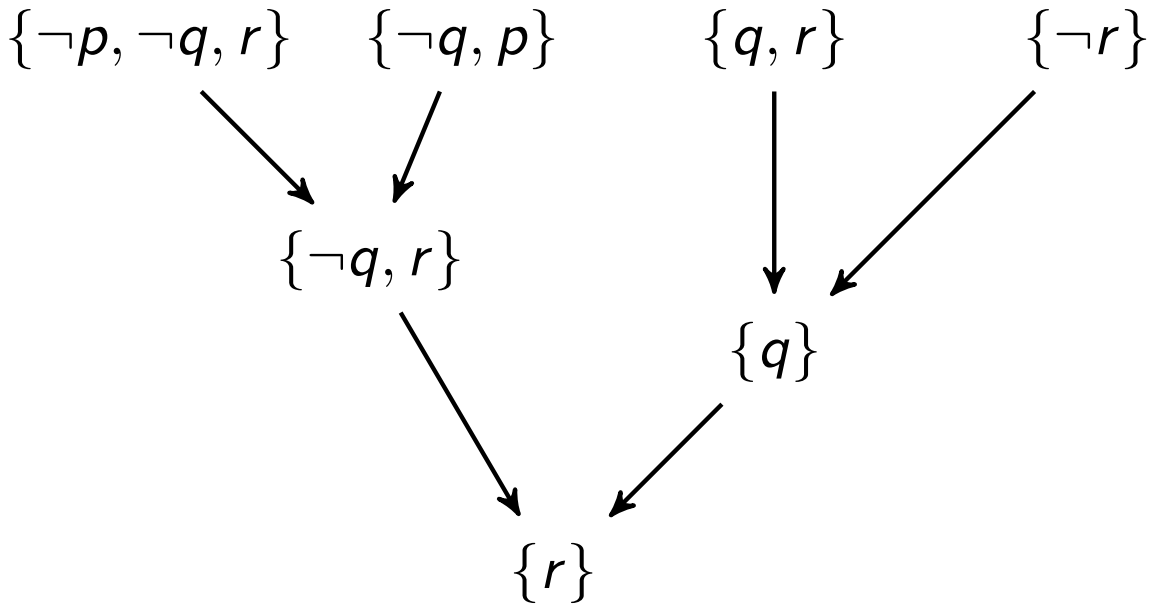

Thus, the original set of assumptions is contradictory.

Example Proof of Contradiction by Resolution

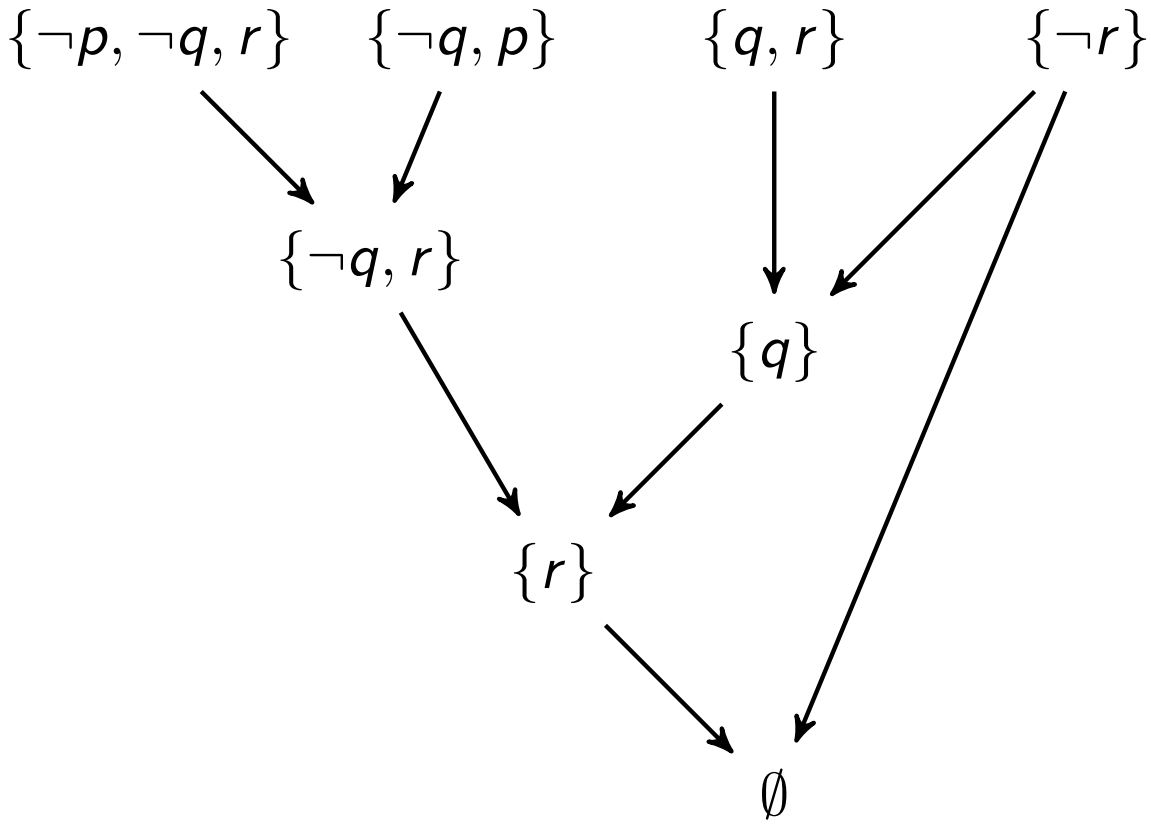

Thus, the original set of assumptions is contradictory.

Consistency by Absence of Contradiction

Conversely, if the set is contradictory, then existentially quantifying over all variables yields false, so applying resolution exhaustively also yields false.
Resolution is **complete**.
Therefore, if resolution does not yield false, the set is consistent.

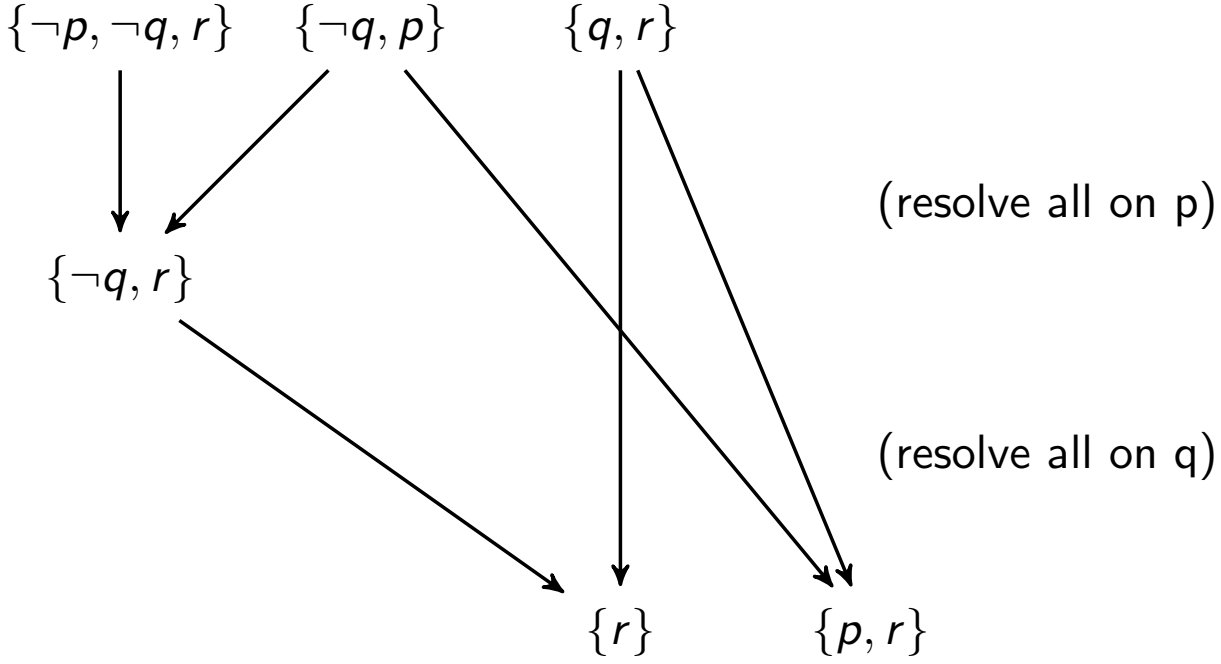

Further resolution attempts would only yield clauses that are subsumed (their subsets, which are stronger, are already derived). No empty clause is generated, so the original set is consistent (a model: $p \mapsto \mathit{true}, q \mapsto \mathit{true}$)

Compactness

Infinite set of Formulas

Suppose that we have a countably infinite set of formulas, with countably many propositional variables

Apply resolution exhaustively to larger and larger prefixes of this infinite set

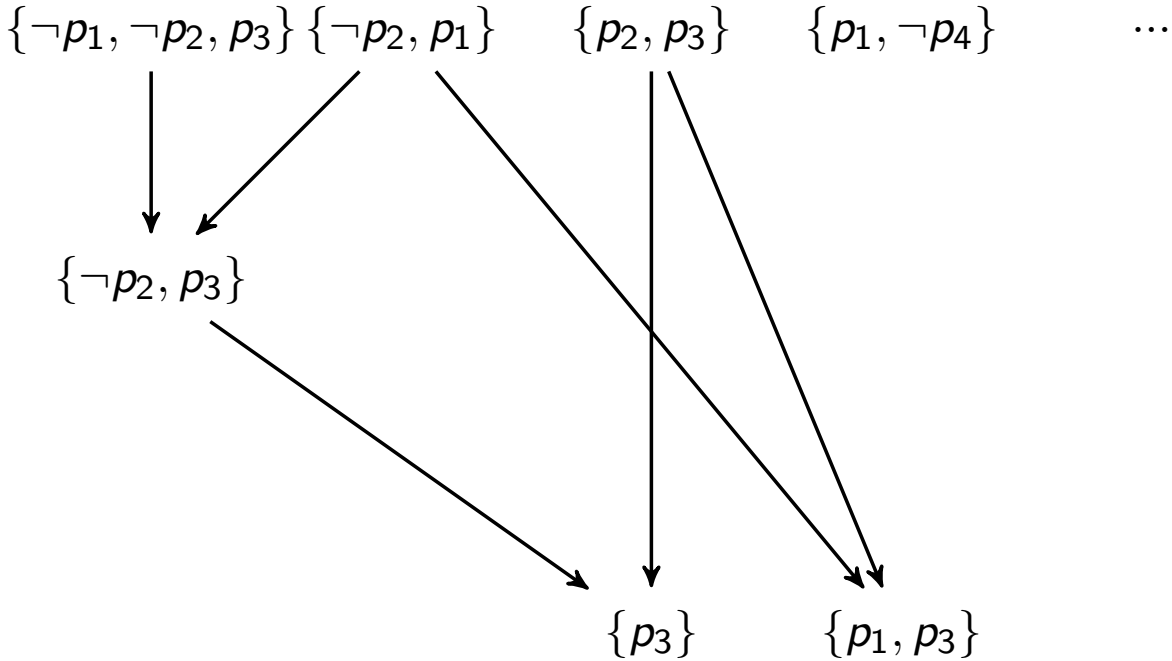

Suppose we are not finding a contradiction in such way. Is the entire infinite set consistent?

Equivalently: if a countable set is contradictory, is there always a finite subset that is contradictory? (Note: there are ∞ many variables.)

Infinite set of Formulas

Suppose that we have a countably infinite set of formulas, with countably many propositional variables

Apply resolution exhaustively to larger and larger prefixes of this infinite set

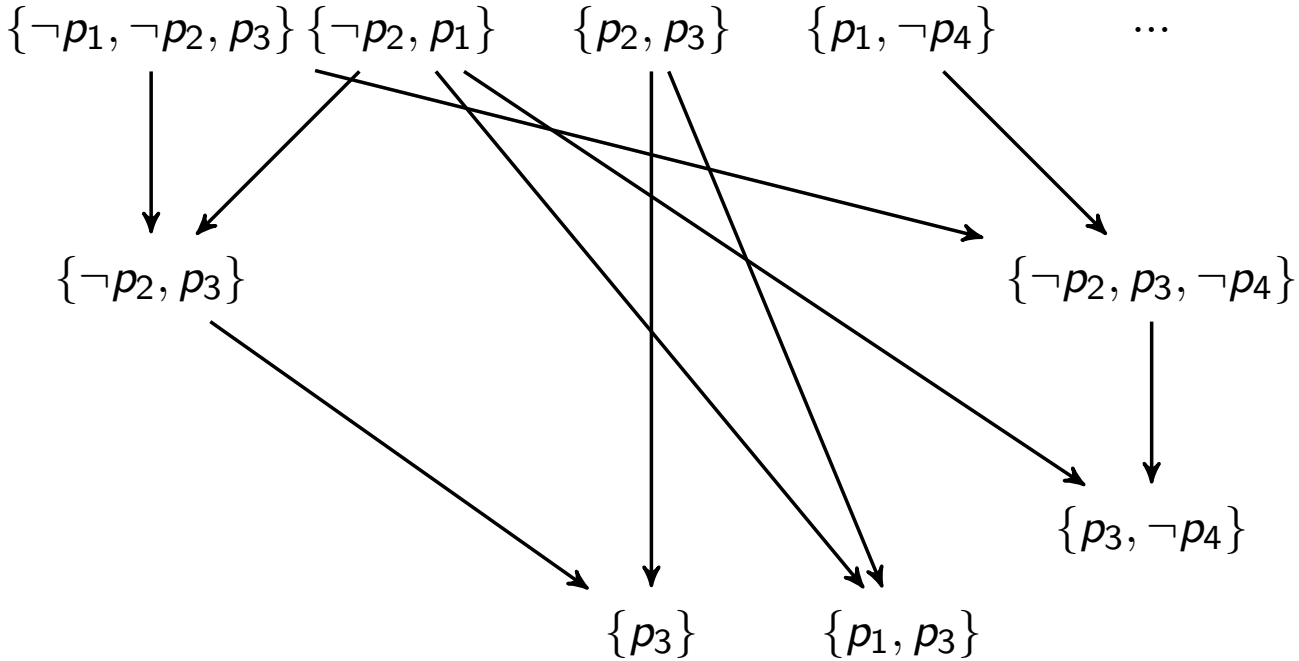

Suppose we are not finding a contradiction in such way. Is the entire infinite set consistent?

Equivalently: if a countable set is contradictory, is there always a finite subset that is contradictory? (Note: there are ∞ many variables.)

Infinite set of Formulas

Suppose that we have a countably infinite set of formulas, with countably many propositional variables

Apply resolution exhaustively to larger and larger prefixes of this infinite set

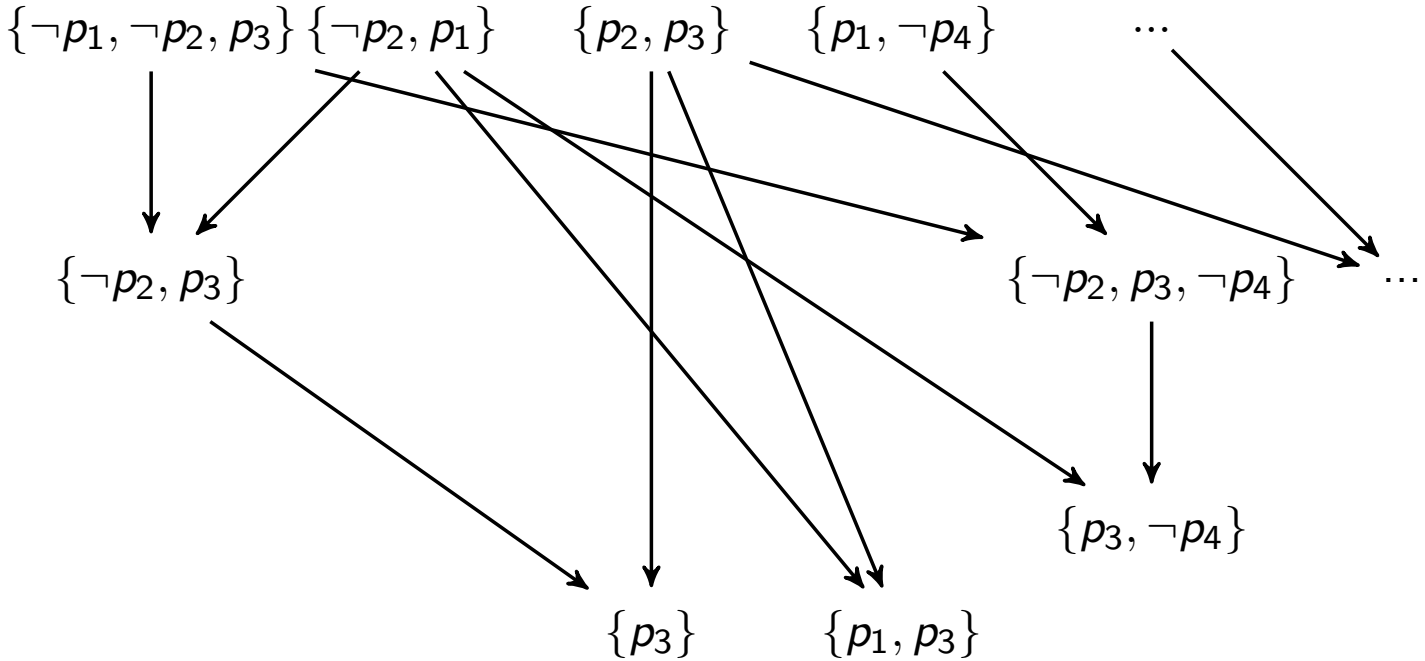

Suppose we are not finding a contradiction in such way. Is the entire infinite set consistent?

Equivalently: if a countable set is contradictory, is there always a finite subset that is contradictory? (Note: there are ∞ many variables.)

Compactness

Theorem (Compactness for Propositional Logic.)

Let S be a set of propositional formulas. Then S is satisfiable iff every finite subset of S is satisfiable.

Equivalently: S is contradictory iff some finite subset of S is contradictory

Remark: Compactness is a non-trivial property. In logic with infinite disjunctions it does not hold. In such *infinitary logic* we could take $S = \{D, p_1, p_2, p_3, \ldots\}$ where $D = \bigvee_{i=1}^{\infty} \neg p_i$, that is, D is equivalent to $\exists i \geq 0. \neg p_i$. In this example, every finite subset of S is satisfiable, but S itself is not.

Proof of Compactness

One direction is trivial: if S is satisfiable then there exists I such that $I \models S$. Then for every finite subset $T \subseteq S$ we have $I \models T$, so T is satisfiable. So, the point is to show the converse.

Intuition: A finitely satisfiable set has "all finite pieces" satisfiable (using potentially different interpretations). The question is whether we can somehow assemble interpretations for all finite pieces T into one large interpretation for the entire infinite set S. We will define such interpretation by extending it, variable by variable, while preserving finite satisfiability for interepretations that begin with values for propositional variables chosen so far.

Let S be finitely satisfiable. Let $V = \{p_1, p_2, \ldots\}$ be the sequence of all propositional variables for formulas in S (this set is countable by our assumption on syntax of formulas, but can be infinite).

Given a sequence of boolean values $u_1, u_2, \ldots, u_n \in \mathcal{B}$ of length $n \geq 0$, by an $(u_1, u_2, \ldots, u_n)$-interpretation we mean an interpretation $I : V \to \mathcal{B}$ such that $I(p_1) = u_1, \ldots, I(p_n) = u_n$.

Proof: Constructing Interpretation

We will define interpretation $I^*(p_k) = v_k$ where the sequence of values $v_1, v_2, \ldots$ is given as follows:

$$v_{k+1} = \begin{cases} \mathit{false}, & \text{if for every finite } T \subseteq S, \text{ there exists a} \\ & (v_1, \ldots, v_k, \mathit{false}) - \text{interpretation } I \text{ such that } I \models T \\ \mathit{true}, & \text{otherwise} \end{cases}$$

We next show by induction the following.

FIRST PART.
Claim: For every non-negative integer k, every finite subset $T \subseteq S$ has a $(v_1, \ldots, v_k)$-interpretation I such that $I \models T$.
Base case: For $k = 0$ the statement reduces to claim that every finite subset of S is satisfiable, which is an assumption of the theorem.

Inductiveness and the Model

Inductive step: Assume the claim for k: every finite subset $T \subseteq S$ has a $(v_1, \ldots, v_k)$-interpretation I such that $I \models T$, we show that the statement holds for $k+1$.
If $v_{k+1} = \mathit{false}$, the inductive statement holds by definition of v_{k+1}. Let $v_{k+1} = \mathit{true}$.
Then by definition of v_{k+1}, there exists a finite set $A \subseteq S$ that has no $(v_1, \ldots, v_k, \mathit{false})$ interpretation. We wish to show that every finite set $B \subseteq T$ has a $(v_1, \ldots, v_k, \mathit{true})$-interpretation such that $I \models B$. Take any such set B. Consider the set $A \cup B$. This is a finite set, so by inductive hypothesis, it has a $(v_1, \ldots, v_k)$-interpretation I. Because $I \models A$, which has no $(v_1, \ldots, v_k, \mathit{false})$-interpretation, we have $I(p_{k+1}) = \mathit{true}$. Therefore, I is a $(v_1, \ldots, v_k, \mathit{true})$-interpretation for $A \cup B$, and therefore for B. This completes the inductive proof.

From Sequence of Interpretations to One

We have shown that for every non-negative integer k, every finite subset $T \subseteq S$ has a $(v_1, \ldots, v_k)$-interpretation I such that $I \models T$. We have defined $I^*(p_k) = v_k$.

SECOND PART.
We finally show that $I^* \models S$. Let $F \in S$. Let $FV(F) = \{p_{i_1}, \ldots, p_{i_k}\}$ and $M = \max(i_1, \ldots, i_k)$. Then $FV(F) \subseteq \{p_1, \ldots, p_M\}$. The set $\{F\}$ is finite, so, by the Claim, it has a $v_1, \ldots, v_M$-interpretation I such that $I \models F$. Because I^* is also a $v_1, \ldots, v_M$-interpretation, and it agrees with I on all variables in F, we have $I^* \models F$.

We have therefore shown that I^* makes all formulas in S true, as desired.

Why did this work

How does this proof break if we allow infinite disjunctions? Consider the above example $S = \{D, p_1, p_2, p_3, \ldots\}$ where $D = \bigvee_{i=1}^{\infty} \neg p_i$. The inductively proved claim still holds, and the sequence defined must be $true, true, true, \ldots$. Here is why the claim holds for every k. Let k be arbitrary and $T \subseteq S$ be finite. Define

$$m = \max(k, \max\{i \mid p_i \in T\})$$

Then consider interpretation that assigns to true all p_j for $j \leq m$ and sets the rest to false. Such interpretation makes D true, so if it is in the set T, then interpretation makes it true. Moreover, all other formulas in T are propositional variables set to true, so the interpretation makes T true. Thus, we see that the inductively proved statement holds even in this case. What the infinite formula D breaks is the second part, which, from the existence of interpretations that agree on an arbitrarily long finite prefix derives an interpretation for infinitely many variables. Indeed, this part explicitly refers to a finite number of variables in the formula.

Resolution for First-Order Logic

Automating First-Order Logic

First-order logic allows arbitrary relations and functions (they are defined only through their axioms)

Useful for modeling all of math (e.g. through set theory axioms), and thus in principle applies to all program verification problems as well.

To prove whether a property holds:

- describe the property using a formula F
- describe the functions and relations in F using a sequence of axioms S

Check if the sequence $(\neg F; S)$ is contradictory. If yes, then F follows from S

Completeness: if F follows from S, then there is a procedure that will in finite time find this (in general we do not know how long it will take).

- semantic notion $S \models F$ (in all interpretation of axioms S formulas F is true) can, in first-order logic, too, be replaced with syntactic notion $S \vdash F$ (F can be derived from S)

We give a complete syntactic inference procedure for first-order logic

First-Order Logic

Set of first-order variables $x_1, x_2, \ldots$

Set of function symbols $f \in \mathcal{L}$ of arity $ar(f_i)$. Constants are of arity zero.
Used to build terms. If $ar(f) = n$ and $t_1, \ldots, t_n$ are terms, then
$f(t_1, \ldots, t_n)$ is a term

Set of relation symbols $R \in \mathcal{L}$ of arity $ar(R_i)$
Used to build artomic formulas. If $ar(R) = n$ and $t_1, \ldots, t_n$ are terms, then
$R(t_1, \ldots, t_n)$ is an atomic formula.

From atomic formulas we build quantifier-free formulas using $\wedge, \vee, \neg$

From quantifier-free formulas we build quantified formulas by quantifying over first-order variables using $\forall x_i, \exists x_i$