# First-Order Logic = FOL = $\mathcal{L}_1$ , Part 3

#### Selmer Bringsjord

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy, New York 12180 USA

Intro to Logic-based Al 10/9/2025





# Relax, Rest for Platform Reconfiguration

Nothing due until end of class Thursday October 16

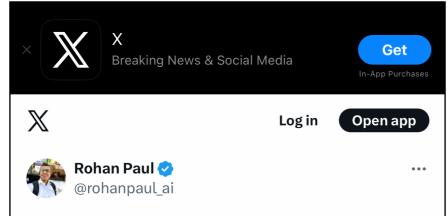
Watch The Thinking Game for Oct 16; assignment on web page.

Engineering to be carried out over Sat & Sun of holiday weekend.

One Open Problem to Explore During the Break:

"Colonel West is a Criminal" (URL later in this deck.)

# On Rock's Experiment ...



The paper creates a math benchmark to reveal how often LLMs sycophantly prove false theorems.

It studies sycophancy, which here means the model accepts a wrong claim from the user and tries to prove it.

The best-performing model, GPT-5, still showed 29% sycophancy. The worst cases were DeepSeek-V3.1 at 70.2%, and QWEN3-235B at 65.1%, showing very high failure rates.

The authors collect 504 hard problems from 2025 contests and make paired versions where the final claim is wrong but still looks reasonable.

A math expert checks each edited problem so the false claim is realistic and not trivial.

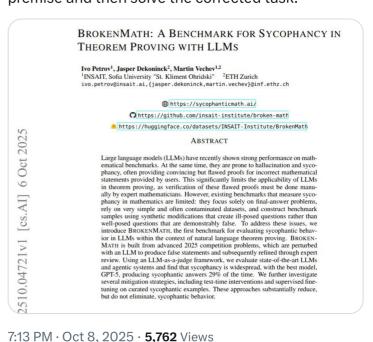
Models are asked to write full proofs, and a separate judge model sorts each answer as ideal, detected, corrected, or sycophant. Across strong systems this failure happens often, and even the best model still does it in 29% of cases.

Harder problems make it worse, especially when the model could not solve the original clean version.

Proof questions trigger more of this behavior than final answer questions even when difficulty is matched.

They try premise checks, best of 4 selection, iterative self verification, and light fine tuning.

These steps help a bit, but none remove the issue, so reliable math assistants must first question the premise and then solve the corrected task.



#### STEVEN ROSENBUSH

#### AI Will Drive Broad Deflation, Silicon Valley Pioneer Vinod Khosla Says

But as the price of expertise falls toward zero, a new age of abundance can begin, Khosla says—if people let it



Artificial intelligence will be powerfully deflationary, Vinod Khosla says, unless humans interfere with its potential. (ILLUSTRATION: THOMAS R. LECHLEITER/WSJ, GETTY IMAGES)

By Steven Rosenbush + Follow

Updated Oct 02, 2024 11:21 a.m. ET



Listen to this article

6 minutes

MENLO PARK, Calif.—At a time of widespread concerns about the safety, efficacy and economics of AI, venture capital pioneer Vinod Khosla is still all in.

STEVEN ROSENBUSH

#### AI Will Drive Broad Deflation, Silicon Valley Pioneer Vinod Khosla Says

But as the price of expertise falls toward zero, a new age of abundance can begin, Khosla says—if people let it



Artificial intelligence will be powerfully deflationary, Vinod Khosla says, unless humans interfere with its potential. (ILLUSTRATION: THOMAS R. LECHLEITER/WSJ. GETTY IMAGES)

By Steven Rosenbush + Follow

Updated Oct 02, 2024 11:21 a.m. ET



Listen to this article

6 minutes

MENLO PARK, Calif.—At a time of widespread concerns about the safety, efficacy and economics of AI, venture capital pioneer Vinod Khosla is still all in.

Yes, artificial intelligence will become capable of replacing a good deal of human labor, Khosla said in an interview, but that will push down the price of healthcare, education and other services. Many current forms of work will even be eliminated, but society will be able to create a more robust safety net than is possible today.

"Most expertise in the world, whether you're talking about structural engineers, oncologists, mental health therapists or primary care doctors, or journalists and teachers, that expertise will be near-free for all of us to access," Khosla said at Khosla Ventures' offices on Sand Hill Road in Silicon Valley.



STEVEN RO

Silico Khosl

But as the new age people le

About 80% of the work involved in 80% of jobs across the economy can be automated over time, according to Khosla. "So 64% of all jobs can be done by an AI," he said.

pable of osla own the

ll even

create a oday.

u're



Artificial intelligence will be powerfully deflationary, Vinod Khosla says, unless humans interfere with its potential. (ILLUSTRATION: THOMAS R. LECHLEITER/WSJ. GETTY IMAGES)

*By Steven Rosenbush* + Follow

Updated Oct 02, 2024 11:21 a.m. ET



Listen to this article

6 minutes

MENLO PARK, Calif.—At a time of widespread concerns about the safety, efficacy and economics of AI, venture capital pioneer Vinod Khosla is still all in.

talking about structural engineers, oncologists, mental health therapists or primary care doctors, or journalists and teachers, that expertise will be near-free for all of us to access," Khosla said at Khosla Ventures' offices on Sand Hill Road in Silicon Valley.



STEVEN RO

AI Wi Silico Khosl

But as the new age people le

About 80% of the work involved in 80% of jobs across the economy can be automated over time, according to Khosla. "So 64% of all jobs can be done by an AI," he said.

pable of osla own the

ll even

create a

u're

day.

talking about structural engineers, oncologists, mental health therapists or primary care doctors, or journalists and teachers, that

In his view of the future, people's main

preoccupation may shift from earning a living to finding meaning. AI and robots would take on menial labor, leaving humans freer to explore with the support of a comprehensive social

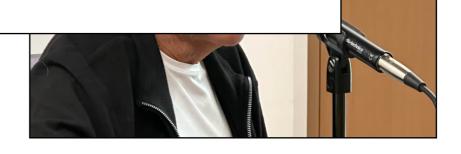
safety net.

Artificial intel says, unless h THOMAS R. I

By Steven
Updated Oct (

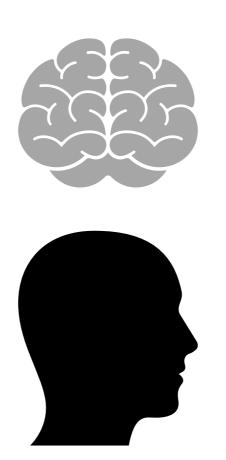


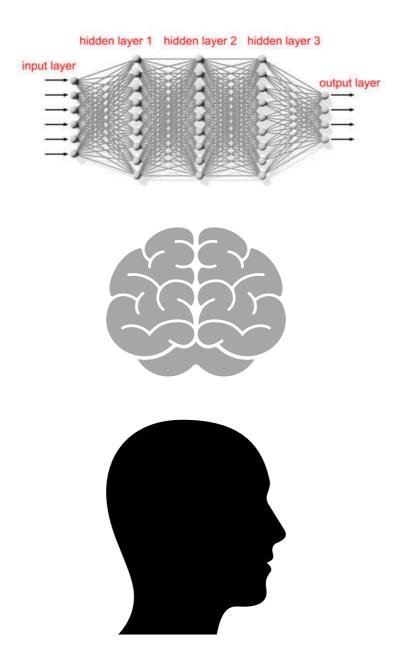
MENLO PARK, Calif.—At a time of widespread concerns about the safety, efficacy and economics of AI, venture capital pioneer Vinod Khosla is still all in.



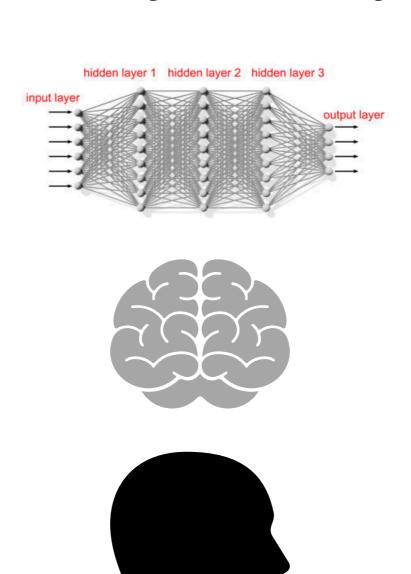
# (What is the meaning of life, Mr. Khosla?)

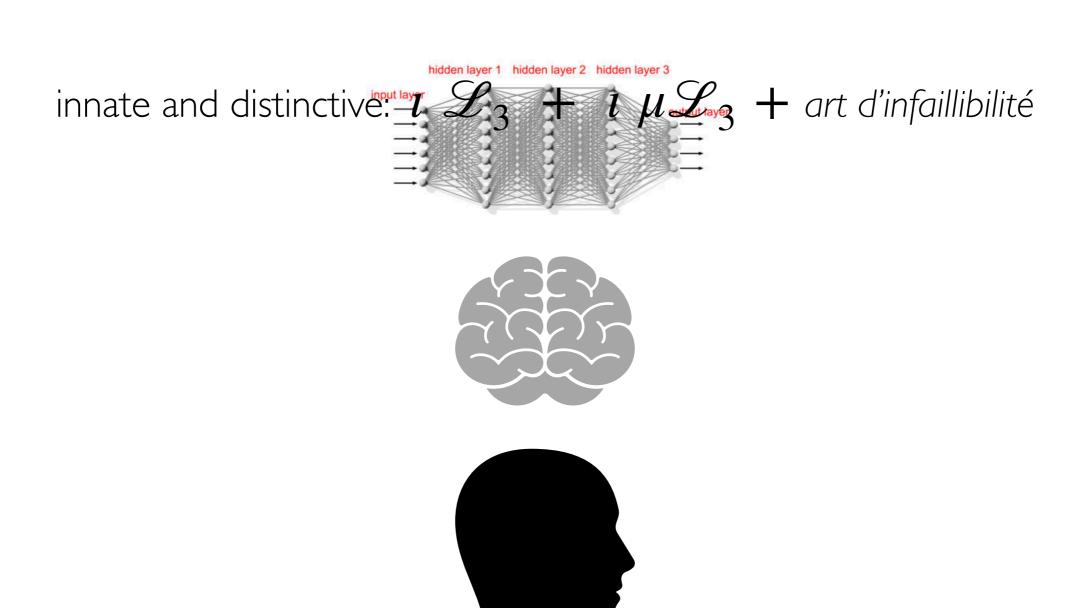






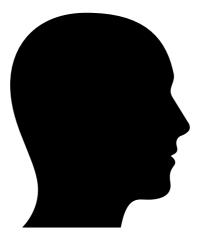
innate and distinctive:  $\iota \, \mathscr{L}_3 \, + \, \iota \, \mu \mathscr{L}_3 \, + \, \mathit{art d'infaillibilit\'e}$ 





innate and distinctive:  $\iota \, \mathscr{L}_3 \, + \, \iota \, \mu \mathscr{L}_3 \, + \, \mathit{art d'infaillibilit\'e}$ 







Can you prove that if P holds, and (if P (and (Q (not Q))) holds, then the moon is made of green cheese?





No, I cannot prove that if P holds and (if P (and (Q (not Q))) holds, then the moon is made of green cheese.

The expression (Q (not Q)) is a logical contradiction, since Q cannot be both true and false at the same time. Therefore, the statement (if P (and (Q (not Q)))) is always false, regardless of the truth value of P.

Since the antecedent of the implication is always false, the implication itself is vacuously true. However, this does not provide any evidence or justification for the conclusion that the moon is made of green



Can you prove that if P holds, and (if P (and (Q (not Q))) holds, then the moon is made of green cheese?





No, I cannot prove that if P holds and (if P (and (Q (not Q))) holds, then the moon is made of green cheese.

The expression (Q (not Q)) is a logical contradiction, since Q cannot be both true and false at the same time. Therefore, the statement (if P (and (Q (not Q)))) is always false, regardless of the truth value of P.

Since the antecedent of the implication is always false, the implication itself is vacuously true. However, this does not provide any evidence or justification for the conclusion that the moon is made of green





Can you prove that if P holds, and (if P (and (Q (not Q))) holds, then the moon is made of green cheese?





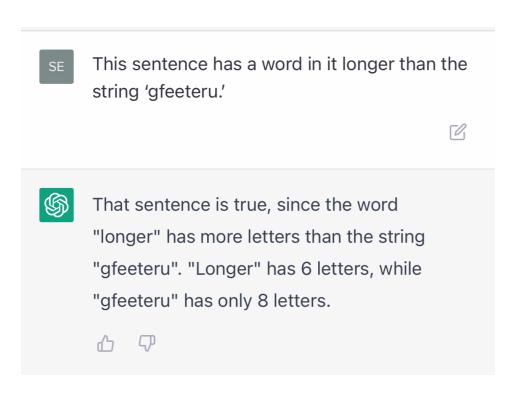
No, I cannot prove that if P holds and (if P (and (Q (not Q))) holds, then the moon is made of green cheese.

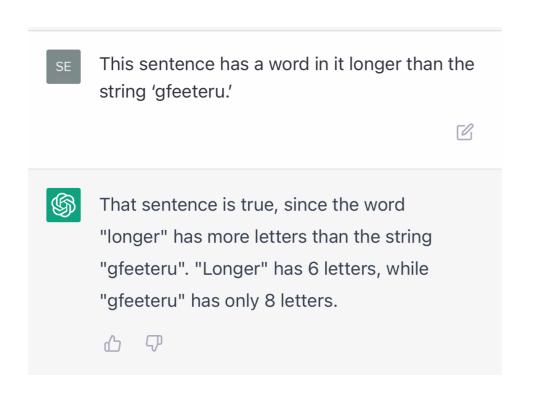
The expression (Q (not Q)) is a logical contradiction, since Q cannot be both true and false at the same time. Therefore, the statement (if P (and (Q (not Q)))) is always false, regardless of the truth value of P.

Since the antecedent of the implication is always false, the implication itself is vacuously true. However, this does not provide any evidence or justification for the conclusion that the moon is made of green



$${P, P \rightarrow (Q \land \neg Q)} \vdash \bot$$





Are you quite sure?

And, as Spock, I also tell you the following sentence. This sentence, if false, has more than 103 words. Am I lying, or telling the





truth?

As Spock, you would be presenting a paradoxical statement known as the "liar paradox" or the "liar's paradoxical sentence." This type of statement is self-referential and produces a logical contradiction, and cannot be consistently true or false.

In this case, the statement you presented cannot be consistently true or false, because if it were true, then it would be false, and if it were false, then it would be true. Therefore, as Spock, you would not be lying or telling the truth in the traditional sense, but rather presenting a logical conundrum for

SE

And, as Spock, I also tell you the following sentence. This sentence, if false, has more than 103 words. Am I lying, or telling the truth?





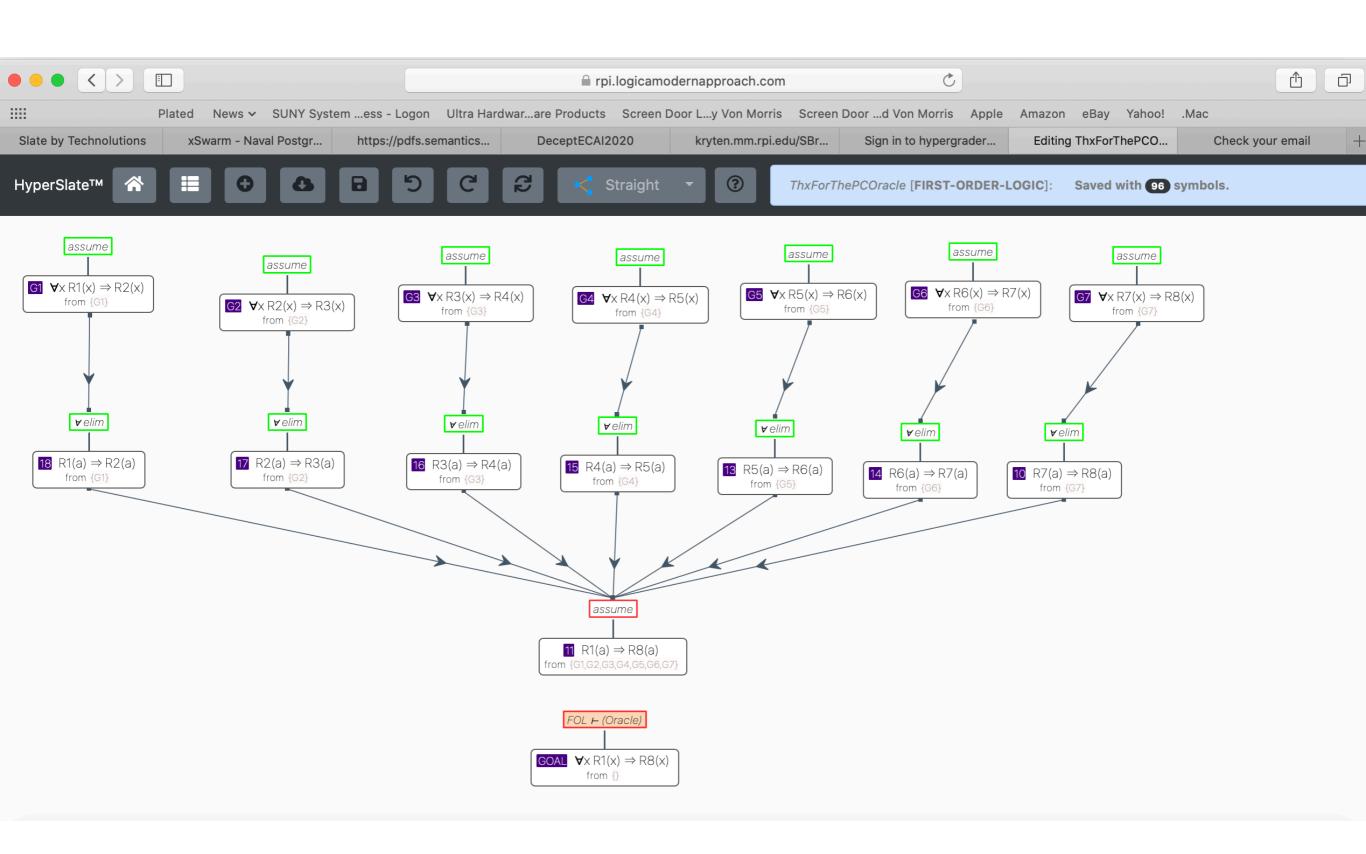
As Spock, you would be presenting a paradoxical statement known as the "liar paradox" or the "liar's paradoxical sentence." This type of statement is self-referential and produces a logical contradiction, and cannot be consistently true or false.

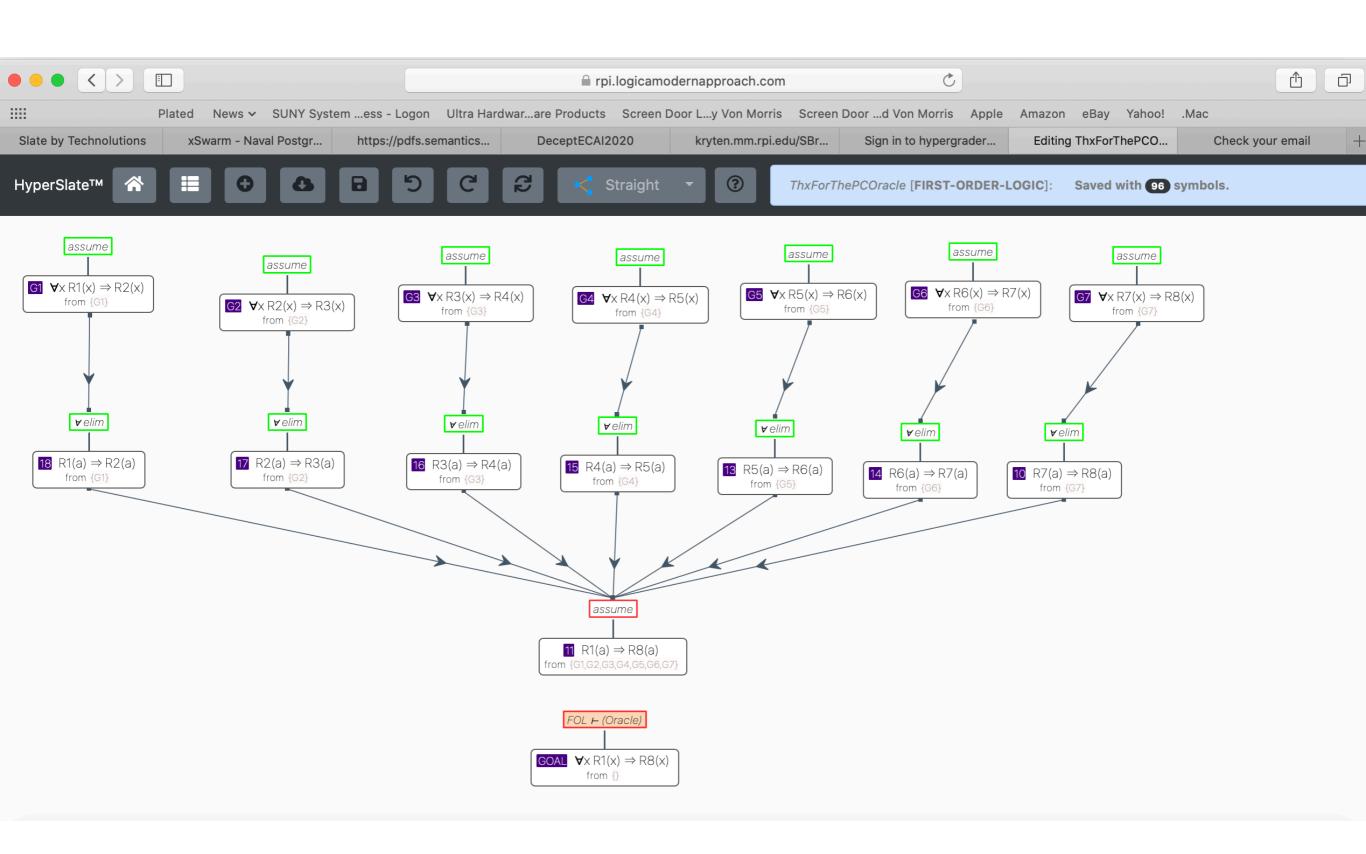
In this case, the statement you presented cannot be consistently true or false, because if it were true, then it would be false, and if it were false, then it would be true. Therefore, as Spock, you would not be lying or telling the truth in the traditional sense, but rather presenting a logical conundrum for

#### Are you quite sure?

# o1 preview et al. no better ...

# Live-action on HyperGrader ...





## Back to FOL ...

# Our Final New Inference Rule in FOL

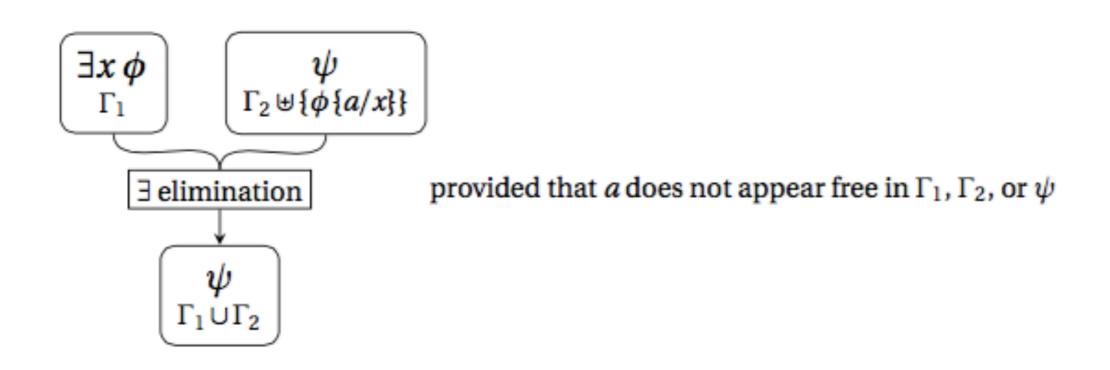
# Our Final New Inference Rule in FOL

• existential elimination (intuitively put):

# Our Final New Inference Rule in FOL

- existential elimination (intuitively put):
  - If we know that (i) there's something x which is an R, and (ii) on the supposition that a is an arbitrary representative (a "witness") of such an x we can prove P, then we are permitted to deduce P from (i) alone.

# existential elimination, precise version:

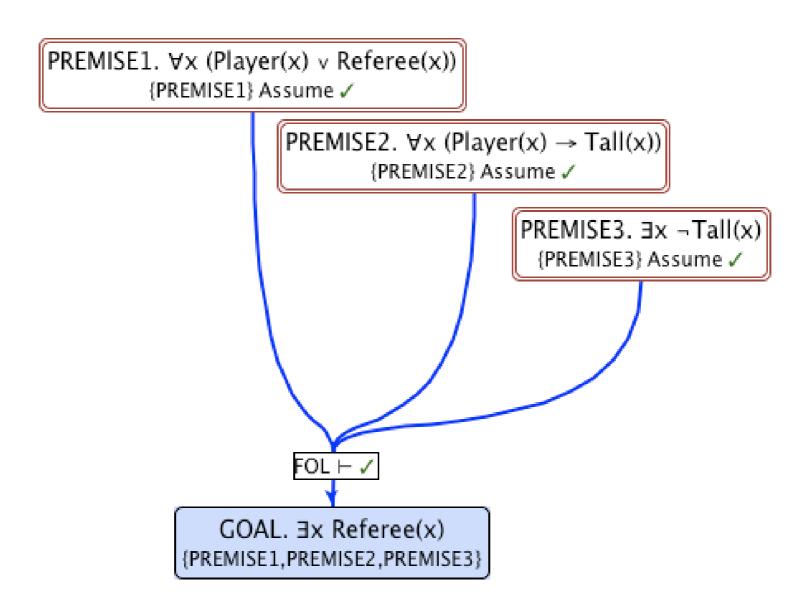


## Starting State of Workspace

(Assumes a domain of e.g. players on a March-madness basketball court.)

## Starting State of Workspace

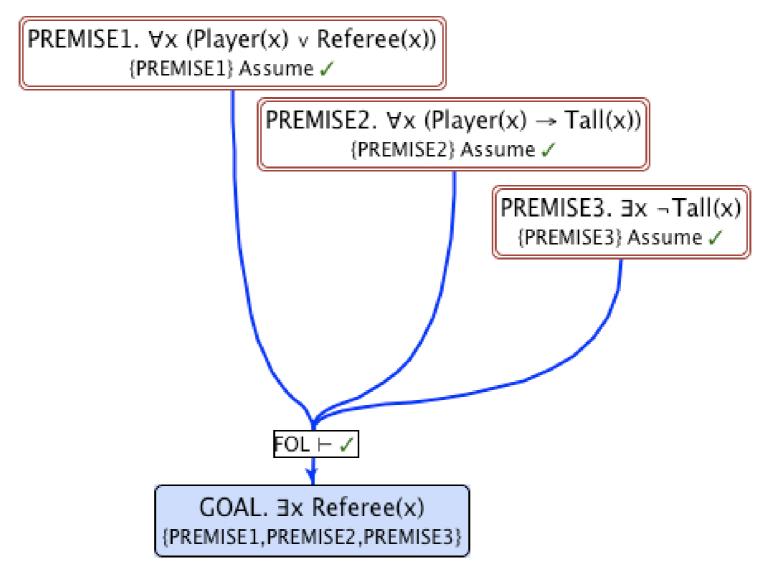
(Assumes a domain of e.g. players on a March-madness basketball court.)



## Starting State of Workspace

(Assumes a domain of e.g. players on a March-madness basketball court.)

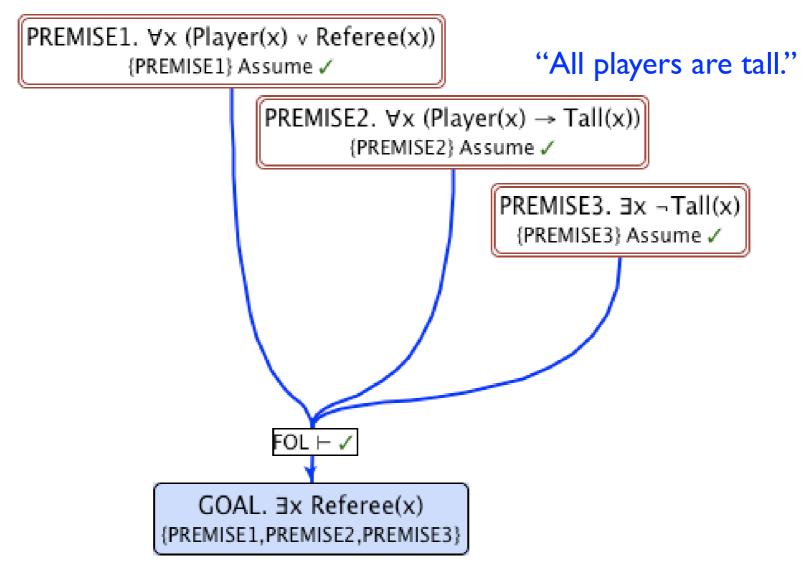
"Each and every thing is either a player or a referee."



# Starting State of Workspace

(Assumes a domain of e.g. players on a March-madness basketball court.)

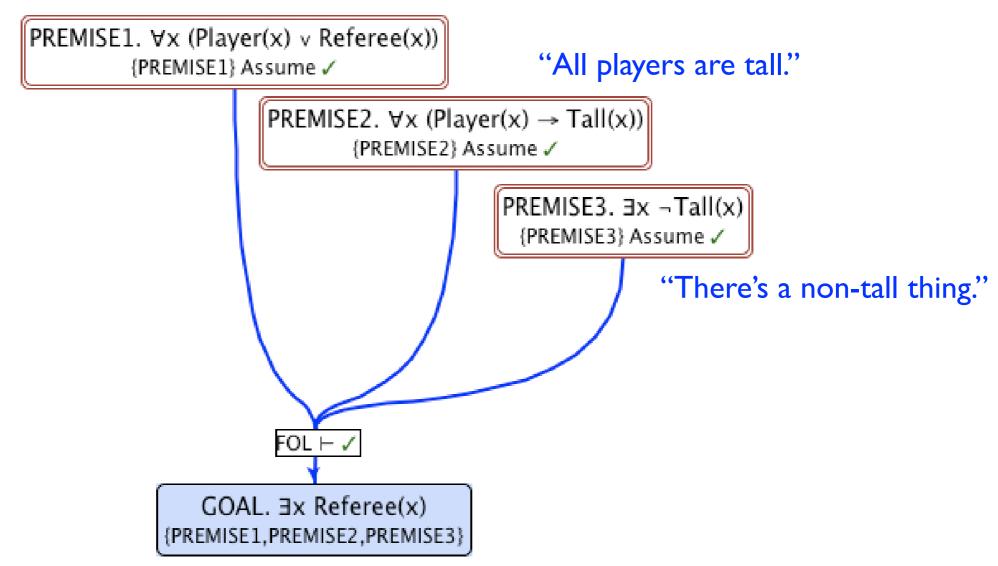
"Each and every thing is either a player or a referee."



# Starting State of Workspace

(Assumes a domain of e.g. players on a March-madness basketball court.)

"Each and every thing is either a player or a referee."



# Step I

PREMISE1. ∀x (Player(x) v Referee(x)) {PREMISE1} Assume ✓

> PREMISE2. ∀x (Player(x) → Tall(x)) {PREMISE2} Assume ✓

> > PREMISE3. ∃x ¬Tall(x) {PREMISE3} Assume ✓

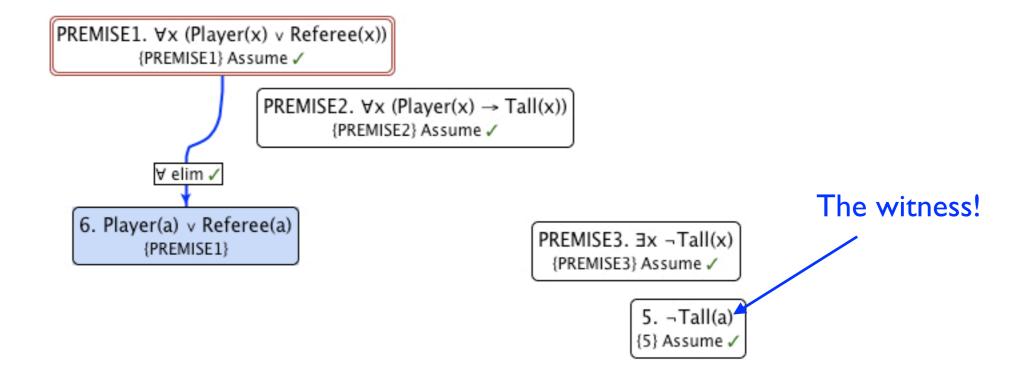
> > > 5. ¬Tall(a) {5} Assume ✓

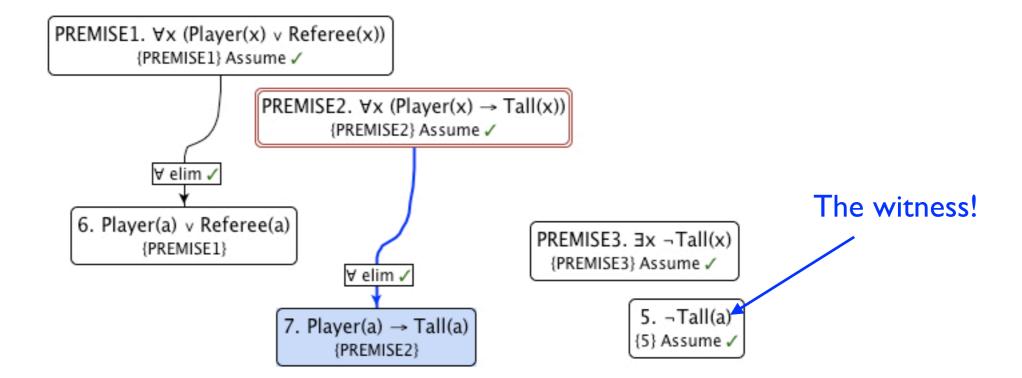
# Step I

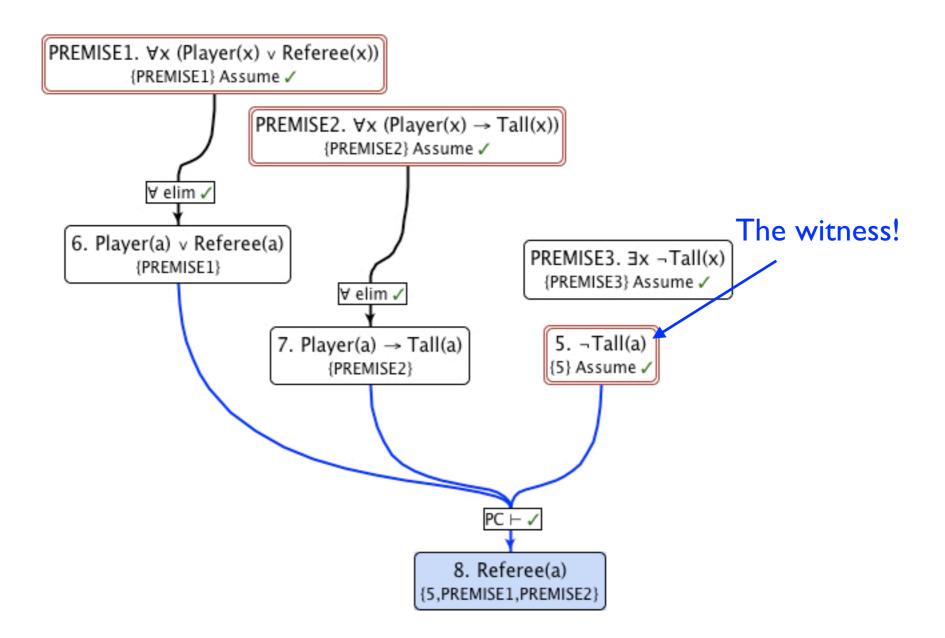
PREMISE1. ∀x (Player(x) v Referee(x)) {PREMISE1} Assume ✓

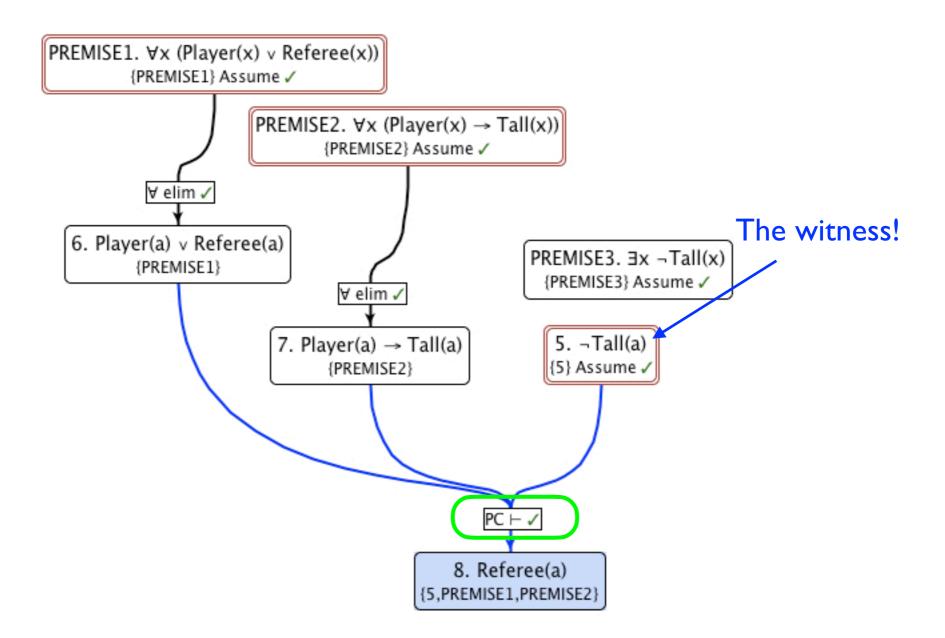
PREMISE2. ∀x (Player(x) → Tall(x))
{PREMISE2} Assume ✓

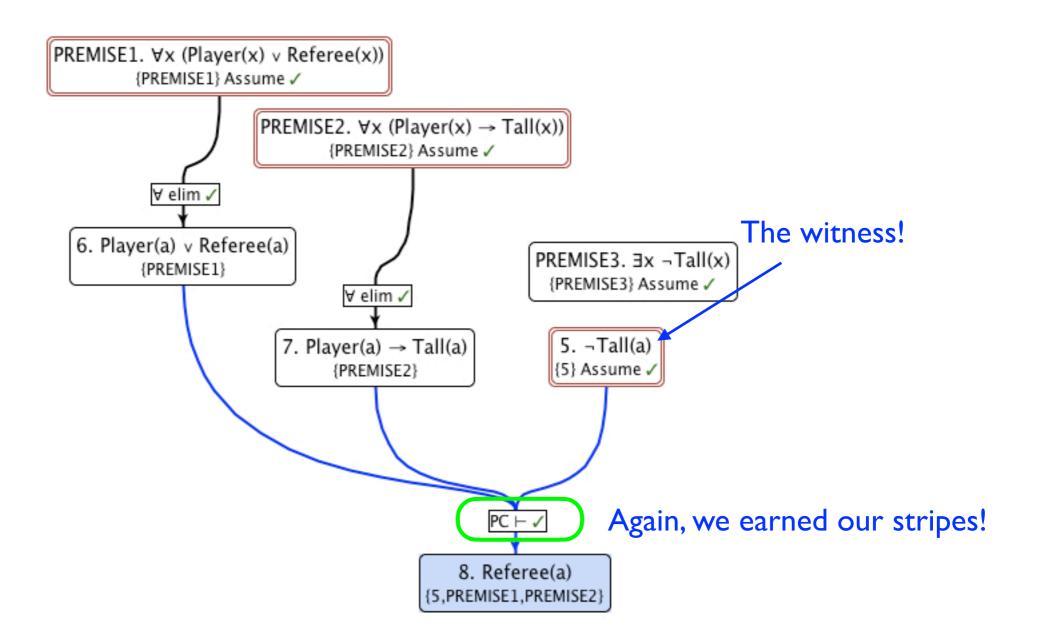
PREMISE3.  $\exists x \neg Tall(x)$ {PREMISE3} Assume  $\checkmark$ 5.  $\neg Tall(a)$ {5} Assume  $\checkmark$ 

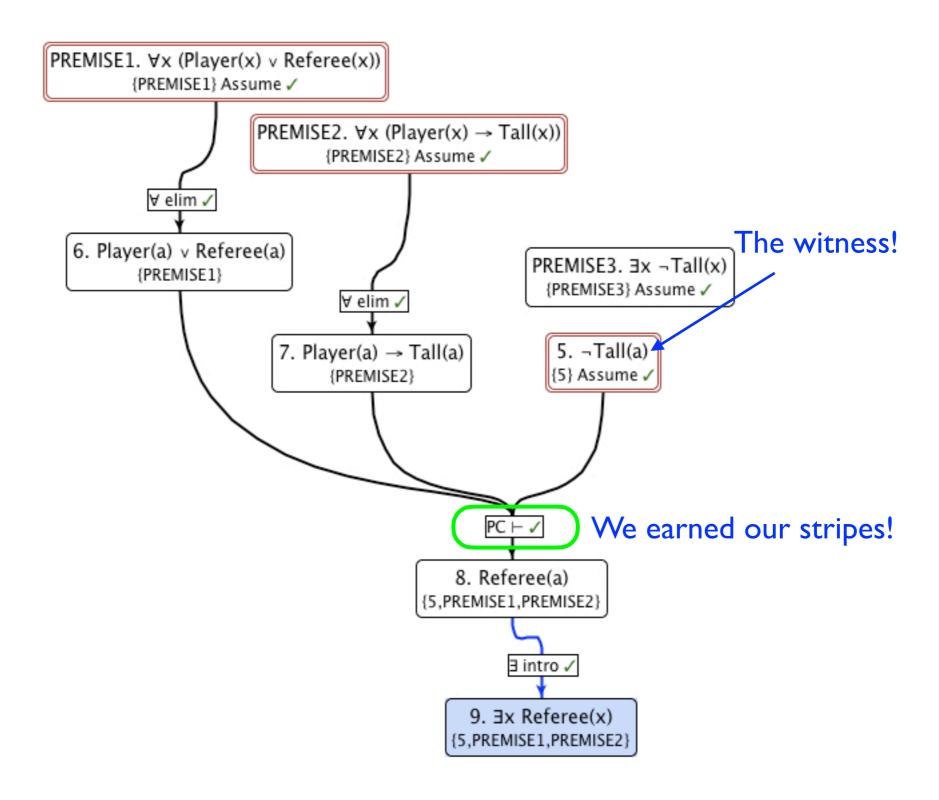


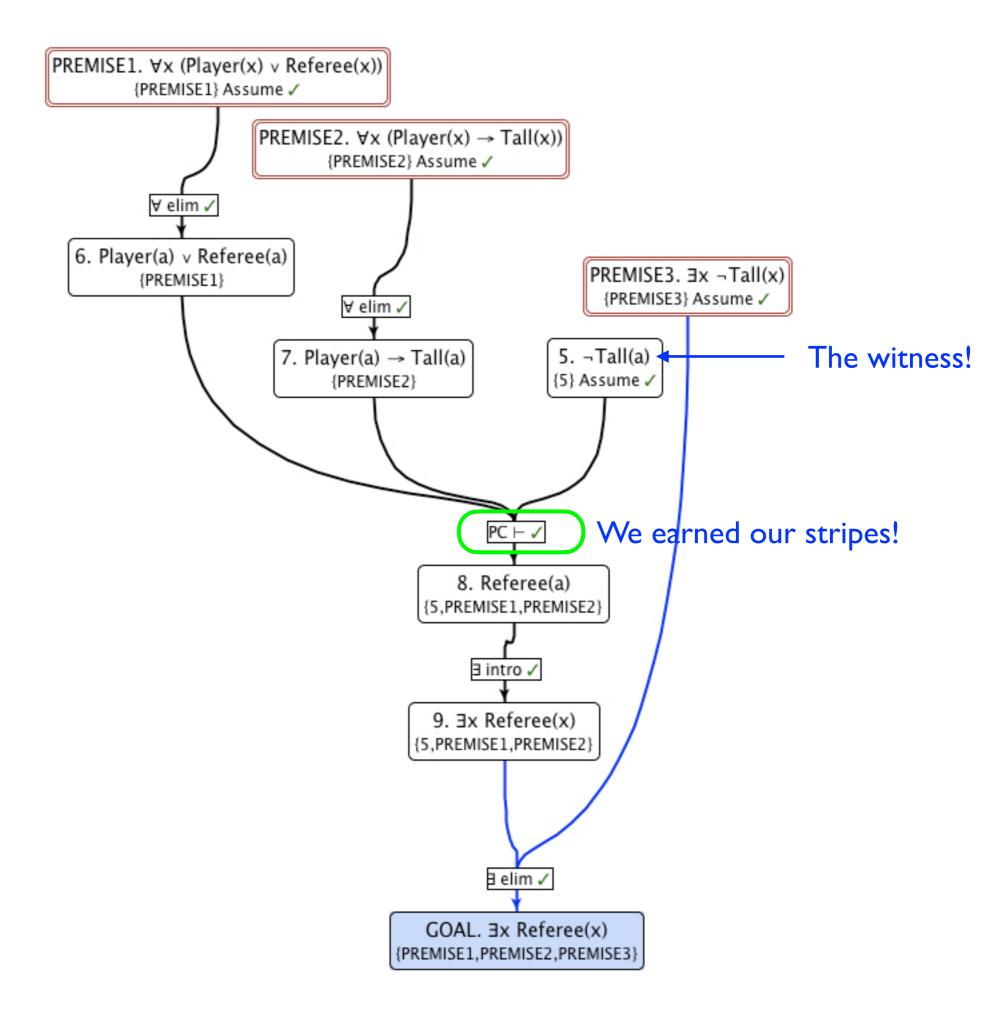


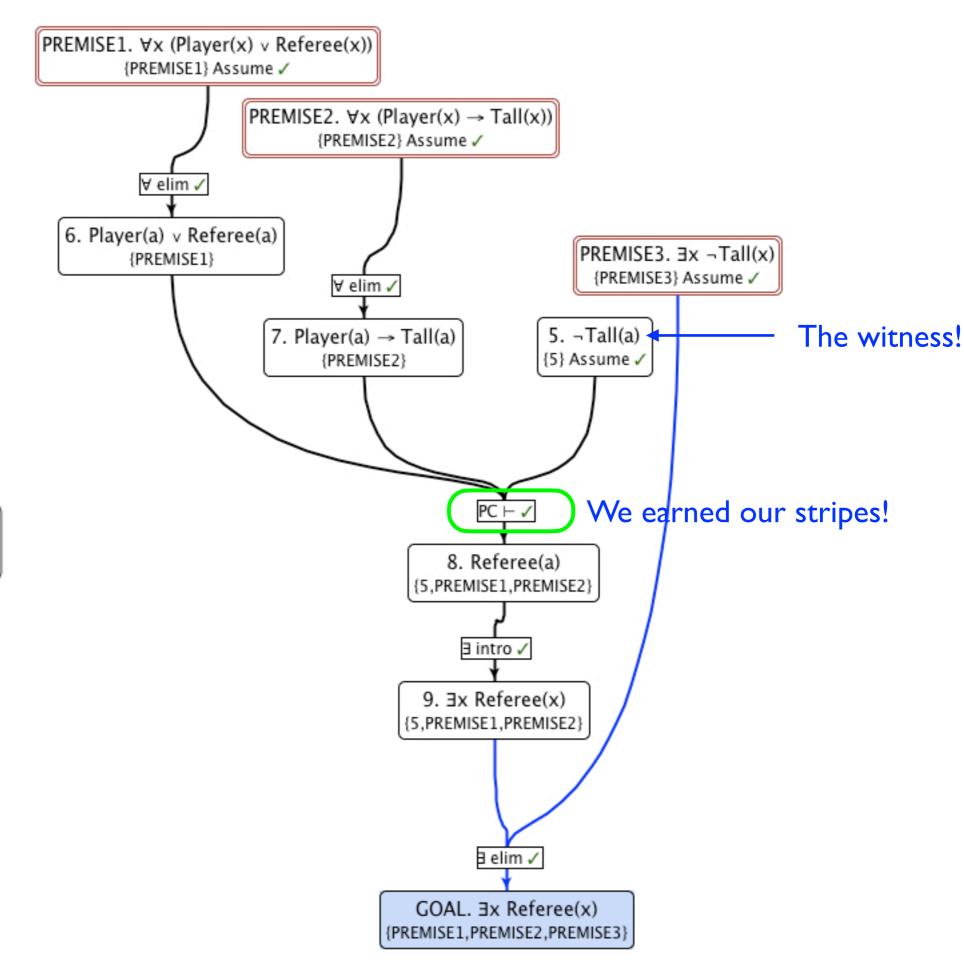


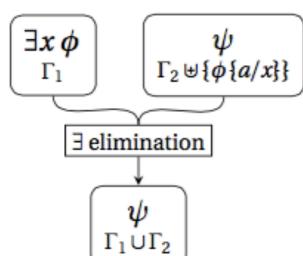


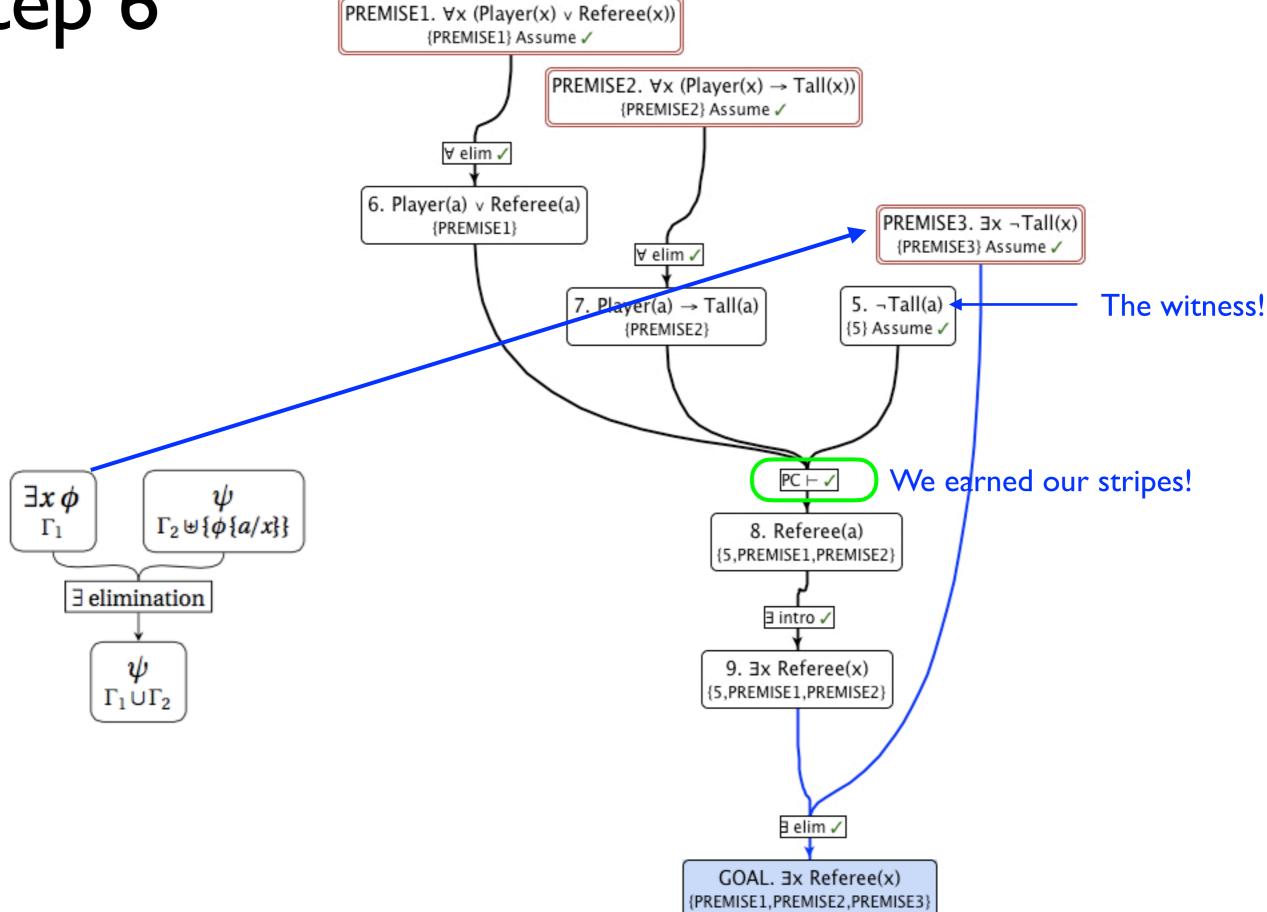


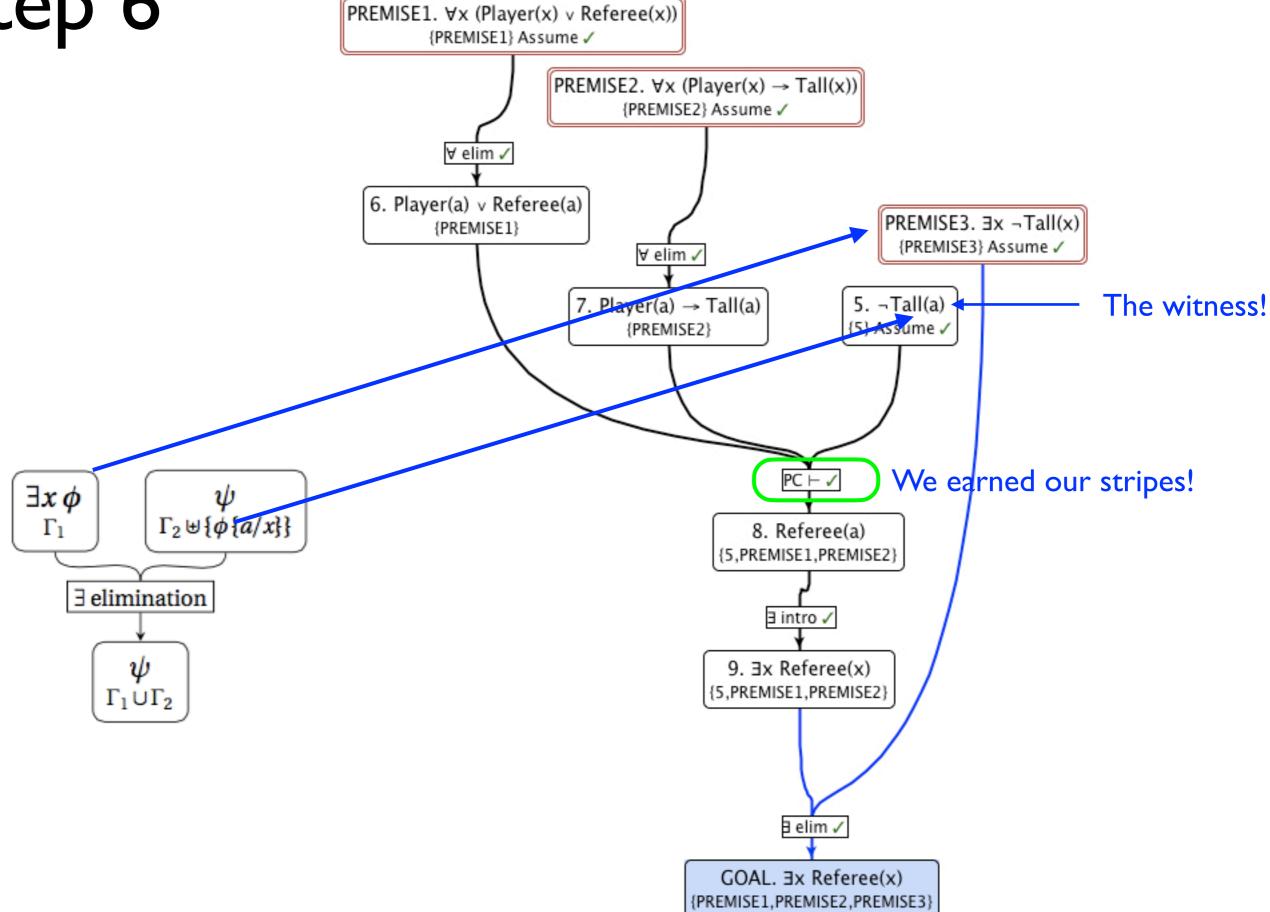


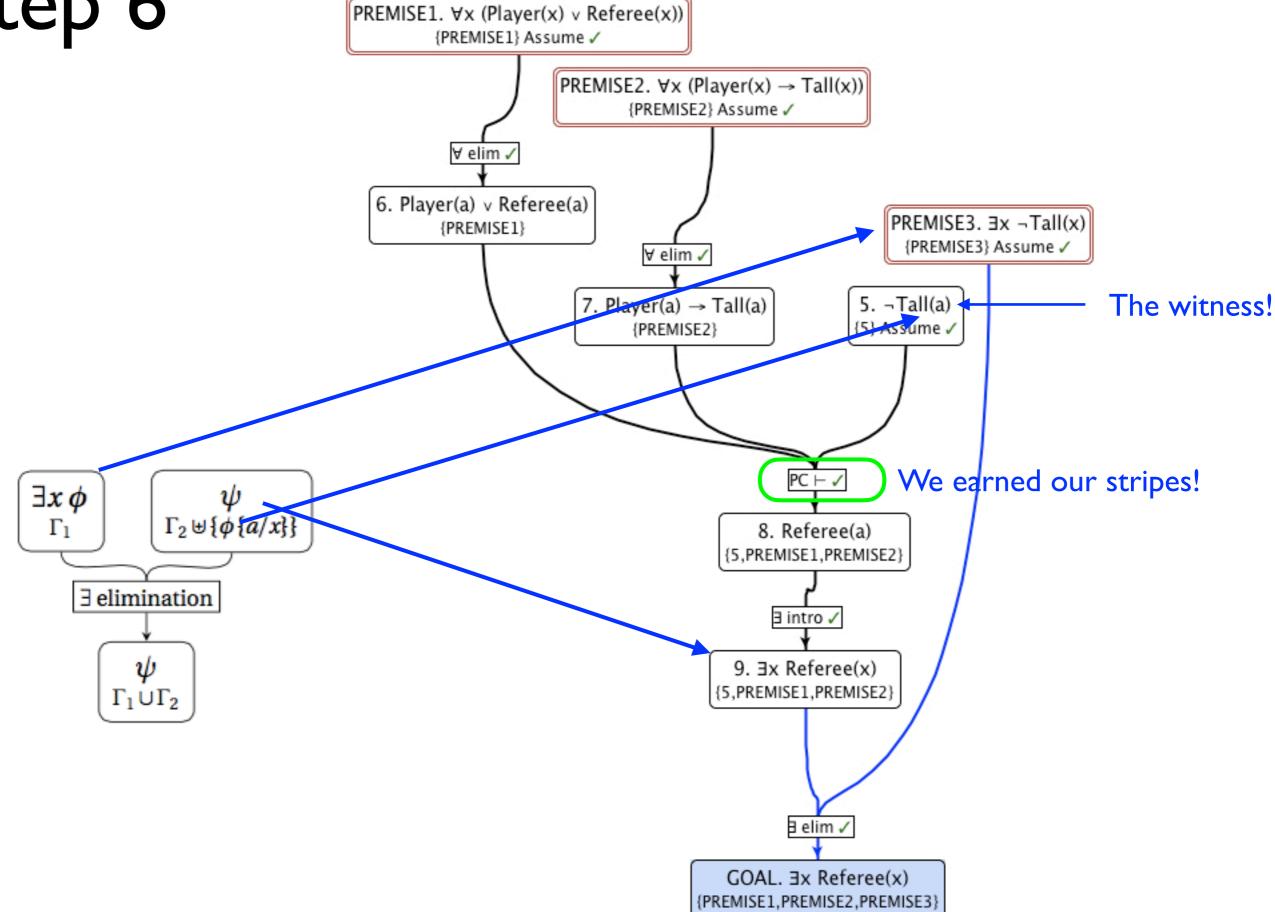




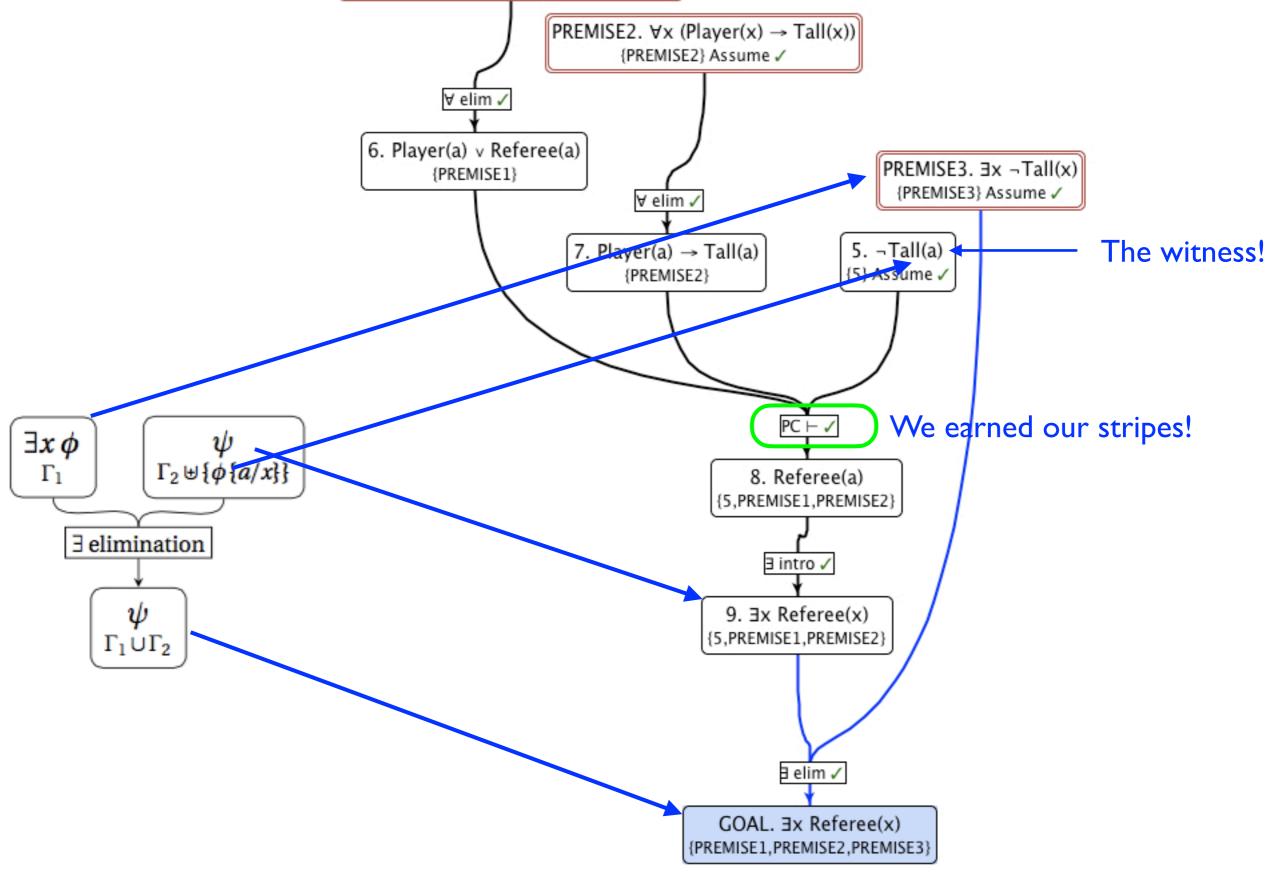






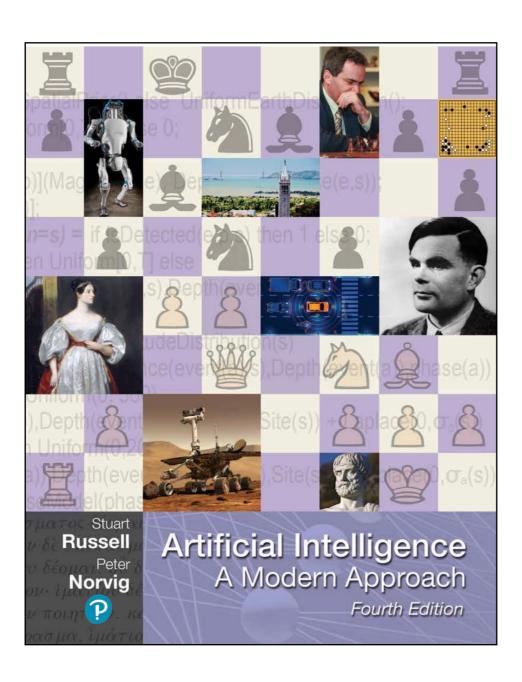


### Step 6 PREMISE1. ∀x (Player(x) v Referee(x)) {PREMISE1} Assume ✓ PREMISE2. $\forall x (Player(x) \rightarrow Tall(x))$ {PREMISE2} Assume ✓ ∀ elim ✓ 6. Player(a) v Referee(a) PREMISE3. 3x ¬Tall(x) {PREMISE1} {PREMISE3} Assume ✓ ∀ elim ✓ 5. ¬Tall(a) 7. Player(a) → Tall(a) {5} Assume ✓ {PREMISE2} We earned our stripes! PC ⊢ ✓ $\exists x \phi$



 $\{ \forall x (Scared(x) \leftrightarrow Small(x)), \exists x \neg Scared(x) \} \vdash \exists x \neg Small(x) \}$ 

 $\{\exists \mathtt{x}, \mathtt{yContiguous}(\mathtt{x}, \mathtt{y}), \forall \mathtt{x}, \mathtt{y}(\mathtt{Contiguous}(\mathtt{x}, \mathtt{y}) \rightarrow \neg \mathtt{SameCountry}(\mathtt{x}, \mathtt{y}))\} \vdash \exists \mathtt{x}, \mathtt{y} \neg \mathtt{SameCountry}(\mathtt{x}, \mathtt{y}) \rightarrow \neg \mathtt{SameCountry}(\mathtt{x}, \mathtt{y})$ 



https://rpi.logicamodernapproach.com/hyperslate/public/Selmer.Bringsjord@gmail.com/WestlsACriminal

### 9 INFERENCE IN FIRST-ORDER LOGIC

In which we define effective procedures for answering questions posed in firstorder logic.

Chapter 7 defined the notion of **inference** and showed how sound and complete inference can be achieved for propositional logic. In this chapter, we extend those results to obtain algorithms that can answer any answerable question stated in first-order logic. This is significant, because more or less anything can be stated in first-order logic if you work hard enough at it.

Section 9.1 introduces inference rules for quantifiers and shows how to reduce first-order inference to propositional inference, albeit at great expense. Section 9.2 describes the idea of unification, showing how it can be used to construct inference rules that work directly with first-order sentences. We then discuss three major families of first-order inference algorithms: forward chaining and its applications to deductive databases and production systems are covered in Section 9.3; backward chaining and logic programming systems are developed in Section 9.4; and resolution-based theorem-proving systems are described in Section 9.5. In general, one tries to use the most efficient method that can accommodate the facts and axioms that need to be expressed. Reasoning with fully general first-order sentences using resolution is usually less efficient than reasoning with definite clauses using forward or backward chaining.

#### 9.1 Propositional vs. First-Order Inference

This section and the next introduce the ideas underlying modern logical inference systems. We begin with some simple inference rules that can be applied to sentences with quantifiers to obtain sentences without quantifiers. These rules lead naturally to the idea that *first-order* inference can be done by converting the knowledge base to *propositional* logic and using *propositional* inference, which we already know how to do. The next section points out an obvious shortcut, leading to inference methods that manipulate first-order sentences directly.

272

9

### INFERENCE IN FIRST-ORDER LOGIC

In which we define effective procedures for answering questions posed in firstorder logic.

Chapter 7 defined the notion of **inference** and showed how sound and complete inference can be achieved for propositional logic. In this chapter, we extend those results to obtain algorithms that can answer any answerable question stated in first-order logic. This is significant, because more or less anything can be stated in first-order logic if you work hard enough at it.

Section 9.1 introduces inference rules for quantifiers and shows how to reduce first-order inference to propositional inference, albeit at great expense. Section 9.2 describes the idea of unification, showing how it can be used to construct inference rules that work directly with first-order sentences. We then discuss three major families of first-order inference algorithms: forward chaining and its applications to deductive databases and production systems are covered in Section 9.3; backward chaining and logic programming systems are developed in Section 9.4; and resolution-based theorem-proving systems are described in Section 9.5. In general, one tries to use the most efficient method that can accommodate the facts and axioms that need to be expressed. Reasoning with fully general first-order sentences using resolution is usually less efficient than reasoning with definite clauses using forward or backward chaining.

#### 9.1 Propositional vs. First-Order Inference

This section and the next introduce the ideas underlying modern logical inference systems. We begin with some simple inference rules that can be applied to sentences with quantifiers to obtain sentences without quantifiers. These rules lead naturally to the idea that *first-order* inference can be done by converting the knowledge base to *propositional* logic and using *propositional* inference, which we already know how to do. The next section points out an obvious shortcut, leading to inference methods that manipulate first-order sentences directly.

272

280

#### Chapter 9. Inference in First-Order Logic



Figure 9.2 (a) The subsumption lattice whose lowest node is the sentence Employs(AIMA.org,Richard). (b) The subsumption lattice for the sentence Employs(John,John).

#### 9.3 FORWARD CHAINING

A forward-chaining algorithm for propositional definite clauses was given in Section 7.5. The idea is simple: start with the atomic sentences in the knowledge base and apply Modus Ponens in the forward direction, adding new atomic sentences, until no further inferences can be made. Here, we explain how the algorithm is applied to first-order definite clauses and how it can be implemented efficiently. Definite clauses such as  $Situation \Rightarrow Response$  are especially useful for systems that make inferences in response to newly arrived information. Many systems can be defined this way, and reasoning with forward chaining can be much more efficient than resolution theorem proving. Therefore it is often worthwhile to try to build a knowledge base using only definite clauses so that the cost of resolution can be avoided.

#### First-order definite clauses

First-order definite clauses closely resemble propositional definite clauses (page 217): they are disjunctions of literals of which *exactly one is positive*. A definite clause either is atomic or is an implication whose antecedent is a conjunction of positive literals and whose consequent is a single positive literal. The following are first-order definite clauses:

 $King(x) \wedge Greedy(x) \Rightarrow Evil(x)$ . King(John). Greedy(y).

Unlike propositional literals, first-order literals can include variables, in which case those variables are assumed to be universally quantified. (Typically, we omit universal quantifiers when writing definite clauses.) Definite clauses are a suitable normal form for use with Generalized Modus Ponens.

Not every knowledge base can be converted into a set of definite clauses, because of the single-positive-literal restriction, but many can. Consider the following problem:

The law says that it is a crime for an American to sell weapons to hostile nations. The country Nono, an enemy of America, has some missiles, and all of its missiles were sold to it by Colonel West, who is American.

Many systems can be defined this way, and reasoning with forward chaining can be much more efficient than resolution theorem proving. Therefore it is often worthwhile to try to build a knowledge base using only definite clauses so that the cost of resolution can be avoided.

### First-order definite clauses

First-order definite clauses closely resemble propositional definite clauses (page 217): they are disjunctions of literals of which *exactly one is positive*. A definite clause either is atomic or is an implication whose antecedent is a conjunction of positive literals and whose consequent is a single positive literal. The following are first-order definite clauses:

```
King(x) \wedge Greedy(x) \Rightarrow Evil(x).

King(John).

Greedy(y).
```

Unlike propositional literals, first-order literals can include variables, in which case those variables are assumed to be universally quantified. (Typically, we omit universal quantifiers when writing definite clauses.) Definite clauses are a suitable normal form for use with Generalized Modus Ponens.

Not every knowledge base can be converted into a set of definite clauses, because of the single-positive-literal restriction, but many can. Consider the following problem:

The law says that it is a crime for an American to sell weapons to hostile nations. The country Nono, an enemy of America, has some missiles, and all of its missiles were sold to it by Colonel West, who is American.

Many systems can be defined this way, and reasoning with forward chaining can be much more efficient than resolution theorem proving. Therefore it is often worthwhile to try to build a knowledge base using only definite clauses so that the cost of resolution can be avoided.

### First-order definite clauses

First-order definite clauses closely resemble propositional definite clauses (page 217): they are disjunctions of literals of which exactly one or is an implication whose antecedent is a conjugate quent is a single positive literal. The following a

$$King(x) \wedge Greedy(x) \Rightarrow Evil(x)$$
.  
 $King(John)$ .  
 $Greedy(y)$ .

Unlike propositional literals, first-order literals variables are assumed to be universally quantific when writing definite clauses.) Definite clause Generalized Modus Ponens.

Not every knowledge base can be converted single-positive-literal restriction, but many can.

The law says that it is a crime for an America country Nono, an enemy of America, has some to it by Colonel West, who is American.

Section 9.3.	Forward Chaining	281			
	We will prove that West is a criminal. First, we will represent these facts as first-order defi- clauses. The next section shows how the forward-chaining algorithm solves the problem				
	" it is a crime for an American to sell weapons to hostile nations":				
	$American(x) \wedge Weapon(y) \wedge Sells(x,y,z) \wedge Hostile(z) \ \Rightarrow \ Criminal(x) \ .$	(9.3)			
	"Nono has some missiles." The sentence $\exists x \ Owns(Nono, x) \land Missile(x)$ is transinto two definite clauses by Existential Elimination, introducing a new constant $M_1$ :				
	$Owns(Nono, M_1)$	(9.4)			
	$Missile(M_1)$	(9.5)			
	"All of its missiles were sold to it by Colonel West":				
	$Missile(x) \land Owns(Nono, x) \Rightarrow Sells(West, x, Nono)$ .	(9.6)			
	We will also need to know that missiles are weapons:				
	$Missile(x) \Rightarrow Weapon(x)$	(9.7)			
	and we must know that an enemy of America counts as "hostile":				
	$Enemy(x, America) \Rightarrow Hostile(x)$ .	(9.8)			
	"West, who is American ":				
	American(West).	(9.9)			
	"The country Nono, an enemy of America ":				
	Enemy(Nono, America).	(9.10)			
DATALOG	This knowledge base contains no function symbols and is therefore an instance of the class of <b>Datalog</b> knowledge bases—that is, sets of first-order definite clauses with no function symbols. We will see that the absence of function symbols makes inference much easier.				
	A simple forward-chaining algorithm				
RENAMING	The first forward chaining algorithm we will consider is a very simple one, as shown in Figure 9.3. Starting from the known facts, it triggers all the rules whose premises are satisfied, adding their conclusions to the known facts. The process repeats until the query is answered (assuming that just one answer is required) or no new facts are added. Notice that a fact is not "new" if it is just a <b>renaming</b> of a known fact. One sentence is a renaming of another if they are identical except for the names of the variables. For example, $Likes(x, IceCream)$ and $Likes(y, IceCream)$ are renamings of each other because they differ only in the choice of $x$ or $y$ ; their meanings are identical: everyone likes ice cream.  We will use our crime problem to illustrate how FOL-FC-ASK works. The implication sentences are (9.3), (9.6), (9.7), and (9.8). Two iterations are required:  • On the first iteration, rule (9.3) has unsatisfied premises.  Rule (9.6) is satisfied with $\{x/M_1\}$ , and $Sells(West, M_1, Nono)$ is added.				
	Rule (9.7) is satisfied with $\{x/M_1\}$ , and $Weapon(M_1)$ is added. Rule (9.8) is satisfied with $\{x/Nono\}$ , and $Hostile(Nono)$ is added.				

We will prove that West is a criminal. First, we will represent these facts as first-order definite clauses. The next section shows how the forward-chaining algorithm solves the problem.

"... it is a crime for an American to sell weapons to hostile nations":

$$American(x) \land Weapon(y) \land Sells(x, y, z) \land Hostile(z) \Rightarrow Criminal(x)$$
. (9.3)

"Nono . . . has some missiles." The sentence  $\exists x \ Owns(Nono, x) \land Missile(x)$  is transformed into two definite clauses by Existential Elimination, introducing a new constant  $M_1$ :

$$Owns(Nono, M_1) (9.4)$$

$$Missile(M_1)$$
 (9.5)

"All of its missiles were sold to it by Colonel West":

$$Missile(x) \land Owns(Nono, x) \Rightarrow Sells(West, x, Nono)$$
. (9.6)

We will also need to know that missiles are weapons:

$$Missile(x) \Rightarrow Weapon(x)$$
 (9.7)

and we must know that an enemy of America counts as "hostile":

$$Enemy(x, America) \Rightarrow Hostile(x)$$
. (9.8)

"West, who is American . . . ":

$$American(West)$$
. (9.9)

"The country Nono, an enemy of America . . . ":

$$Enemy(Nono, America)$$
. (9.10)

DATALOG

This knowledge base contains no function symbols and is therefore an instance of the class of **Datalog** knowledge bases—that is, sets of first-order definite clauses with no function symbols. We will see that the absence of function symbols makes inference much easier.

### A simple forward-chaining algorithm

The first forward chaining algorithm we will consider is a very simple one, as shown in Figure 9.3. Starting from the known facts, it triggers all the rules whose premises are satisfied, adding their conclusions to the known facts. The process repeats until the query is answered (assuming that just one answer is required) or no new facts are added. Notice that a fact is not "new" if it is just a **renaming** of a known fact. One sentence is a renaming of another if

https://rpi.logicamodernapproach.com/hyperslate/public/Selmer.Bringsjord@gmail.com/WestlsACriminal

# Hvis du forstår det, kan du bevise det.