

6.034 Notes: Section 11.1

Slide 11.1.1

We've now spent a fair bit of time learning about the language of first-order logic and the mechanisms of automatic inference. And, we've also found that (a) it is quite difficult to write first-order logic and (b) quite expensive to do inference. Both of these conclusions are well justified. Therefore, you may be wondering why we spent the time on logic.

We can motivate our study of logic in a variety of ways. For one, it is the intellectual foundation for all other ways of representing knowledge about the world. As we have already seen, the Web Consortium has adopted a logical language for its Semantic Web project. We also saw that airlines use a language not unlike FOL to describe fare restrictions. We will see later when we talk about natural language understanding that logic also plays a key role.

There is another practical application of logic that is reasonably widespread namely **logic programming**. In this section, we will look briefly at logic programming. Later, when we study natural language understanding, we will build on these ideas.

Rules and Logic Programming

6.034 - Spring 03 - 1

Logic in Practice

- Language of logic is extremely powerful.
- Say what's true, not how to use it.
 - $\forall x, y (\exists z \text{ Parent}(x,z) \wedge \text{Parent}(z,y)) \leftrightarrow \text{GrandParent}(x,y)$
 - Given parents, find grandparents
 - Given grandparents, find parents

6.034 - Spring 03 - 2

Slide 11.1.2

We have seen that the language of logic is extremely general, with much of the power of natural language. One of the key characteristics of logic, as opposed to programming languages but like natural languages, is that in logic you write down what's true about the world, without saying how to use it. So, for example, one can characterize the relationship between parents and grandparents in this sentence without giving an algorithm for finding the grandparents from the grandchildren or a different algorithm for finding the grandchildren given the grandparents.

Slide 11.1.3

However, this very power and lack of specificity about algorithms means that the general methods for performing computations on logical representations (for example, resolution refutation) are hopelessly inefficient for most practical problems.

Logic in Practice

- Language of logic is extremely powerful.
- Say what's true, not how to use it.
 - $\forall x, y (\exists z \text{ Parent}(x,z) \wedge \text{Parent}(z,y)) \leftrightarrow \text{GrandParent}(x,y)$
 - Given parents, find grandparents
 - Given grandparents, find parents
- But, resolution theorem-provers are too inefficient!

6.034 - Spring 03 - 3

Logic in Practice

- Language of logic is extremely powerful.
- Say what's true, not how to use it.
 - $\forall x, y (\exists z \text{ Parent}(x,z) \wedge \text{Parent}(z,y)) \leftrightarrow \text{GrandParent}(x,y)$
 - Given parents, find grandparents
 - Given grandparents, find parents
- But, resolution theorem-provers are too inefficient!
- To regain practicality:
 - Limit the language
 - Simplify the proof algorithm
- Rule-Based Systems
- Logic Programming

6.034 - Spring 03 - 4

Slide 11.1.4

There are, however, approaches to regaining some of the efficiency while keeping much of the power of the representation. These approaches involve both limiting the language as well as simplifying the inference algorithms to make them more predictable. Similar ideas underlie both **logic programming** and **rule-based systems**. We will bias our presentation towards logic programming.

Slide 11.1.5

In logic programming we will also use the clausal representation that we derived for resolution refutation. However, we will limit the type of clauses that we will consider to the class called **Horn clauses**. A clause is Horn if it has at most one positive literal. In the examples below, we show literals without variables, but the discussion applies both to propositional and first order logic.

There are three cases of Horn clauses:

- A **rule** is a clause with one or more negative literals and exactly one positive literal. You can see that this is the clause form of an implication of the form $P_1 \wedge \dots \wedge P_n \rightarrow Q$, that is, the conjunction of the P's implies Q.
- A **fact** is a clause with exactly one positive literal and no negative literals. We generally will distinguish the case of a **ground fact**, that is, a literal with no variables, from the general case of a literal with variables, which is more like an unconditional rule than what one would think of as a "fact".
- In general, there is another case, known as a **consistency constraint** when the clause has no positive literals. We will not deal with these further, except for the special case of a **conjunctive goal clause** which will take this form (the negation of a conjunction of literals is a Horn clause with no positive literal). However, goal clauses are not rules.

Horn Clauses

- A clause is **Horn** if it has at most one positive literal
 - $\neg P_1 \vee \dots \vee \neg P_n \vee Q$ (Rule)
 - Q (Fact)
 - $\neg P_1 \vee \dots \vee \neg P_n$ (Consistency Constraint)
- We will not deal with Consistency Constraints

6.034 - Spring 03 - 5

Horn Clauses

- A clause is **Horn** if it has at most one positive literal
 - $\neg P_1 \vee \dots \vee \neg P_n \vee Q$ (Rule)
 - Q (Fact)
 - $\neg P_1 \vee \dots \vee \neg P_n$ (Consistency Constraint)
- We will not deal with Consistency Constraints
- Rule Notation
 - $P_1 \wedge \dots \wedge P_n \rightarrow Q$ (Logic)
 - If $P_1 \dots P_n$ Then Q (Rule-Based System)
 - $Q \text{ :- } P_1, \dots, P_n$ (Prolog)
- P_i are called **antecedents** (or body)
- Q is called the **consequent** (or head)

6.034 - Spring 03 - 6

Slide 11.1.6

There are many notations that are in common use for Horn clauses. We could write them in standard logical notation, either as clauses, or as implications. In rule-based systems, one usually has some form of equivalent "If-Then" syntax for the rules. In Prolog, which is the most popular logic programming language, the clauses are written as a sort of reverse implication with the ":-" instead of "<".

We will call the Q (positive) literal the **consequent** of a rule and call the P_i (negative) literals the **antecedents**. This is terminology for implications borrowed from logic. In Prolog it is more common to call Q the **head** of the clause and to call the P literals the **body** of the clause.

Slide 11.1.7

Note that not every logical statement can be written in Horn clause form, especially if we disallow clauses with zero positive literals (consistency constraints). Importantly, one cannot have a negation on the right hand side of an implication. That is, we cannot have rules that conclude that something is not true! This is a reasonably profound limitation in general but we can work around it in many useful situations, which we will discuss later. Note that because we are not dealing with consistency constraints (all negative literals) we will not be able to deal with negative facts either.

Limitations

- Cannot conclude negation
 - $P \rightarrow \neg Q$
 - $\neg P \vee \neg Q$: Consistency constraint
 - $\neg P$: Consistency constraint

6.034 - Spring 03 - 7

**Limitations**

- Cannot conclude negation
 - $P \rightarrow \neg Q$
 - $\neg P \vee \neg Q$: Consistency constraint
 - $\neg P$: Consistency constraint
- Cannot conclude (or assert) disjunction
 - $P_1 \wedge P_2 \rightarrow Q_1 \vee Q_2$
 - $Q_1 \vee Q_2$
 - These are not Horn

6.034 - Spring 03 - 8

**Slide 11.1.8**

Similarly, if we have a disjunction on the right hand side of an implication, the resulting clause is not Horn. In fact, we cannot assert a disjunction with more than one positive literal or a disjunction of all negative literals. The former is not Horn while the latter is a consistency constraint.

Slide 11.1.9

It turns out that given our simplified language, we can use a simplified procedure for inference, called **backchaining**, which is basically a generalized form of Modus Ponens (one of the "natural deduction" rules we saw earlier).

Backchaining is relatively simple to understand given that you've seen how resolution works. We start with a literal to "prove", which we call C. We will also use Green's trick (as in Chapter 6.3) to keep track of any variable bindings in C during the proof.

We will keep a stack (first in, last out) of goals to be proved. We initialize the stack to have C (first) followed by the Answer literal (which we write as Ans).

Inference: Backchaining

- To "prove" a literal C
 - Push C and an Ans literal on a stack

6.034 - Spring 03 - 9

**Inference: Backchaining**

- To "prove" a literal C
 - Push C and an Ans literal on a stack
 - Repeat until stack only has Ans literal or no actions available.
 - Pop literal L off of stack

6.034 - Spring 03 - 10

**Slide 11.1.10**

The basic loop is to pop a literal (L) off the stack until either (a) only the Ans literal remains or (b) there are no further actions possible. The first case corresponds to a successful proof; the second case represents a failed proof.

A word of warning. This loop does not necessarily terminate. We will see examples later where simple sets of rules lead to infinite loops.

Slide 11.1.11

Given a literal L, we look for a fact that unifies with L or a rule whose consequent (head) unifies with L. If we find a match, we push the antecedent literals (if any) onto the stack, apply the unifier to the entire stack and then rename all the variables to make sure that there are no variable conflicts in the future. There are other ways of dealing with the renaming but this one will work.

In general, there will be more than one fact or rule that could match L; we will pick one now but be prepared to come back to try another one if the proof doesn't work out. More on this later.

Inference: Backchaining

- To "prove" a literal C
 - Push C and an Ans literal on a stack
 - Repeat until stack only has Ans literal or no actions available.
 - Pop literal L off of stack
 - Choose [with backup] a rule (or fact) whose consequent unifies with L
 - Push antecedents (in order) onto stack
 - Apply unifier to entire stack
 - Rename variables on stack

6.034 - Spring 03 - 11

**Inference: Backchaining**

- To "prove" a literal C
 - Push C and an Ans literal on a stack
 - Repeat until stack only has Ans literal or no actions available.
 - Pop literal L off of stack
 - Choose [with backup] a rule (or fact) whose consequent unifies with L
 - Push antecedents (in order) onto stack
 - Apply unifier to entire stack
 - Rename variables on stack
 - If no match, fail [backup to last choice]

6.034 - Spring 03 - 12

**Slide 11.1.12**

If no match can be found for L, we fail and backup to try the last choice that has other pending matches.

Slide 11.1.13

If you think about it, you'll notice that backchaining is just our familiar friend, resolution. The stack of goals can be seen as negative literals, starting with the negated goal. We don't actually show literals on the stack with explicit negation but they are implicitly negated.

At every point, we pair up a negative literal from the stack with a positive literal (the consequent) from a fact or rule and add the remaining negative literals (the antecedents) to the stack.

Backchaining and Resolution

- Backchaining is just resolution
- To prove C (propositional case)
 - Negate $C \Rightarrow \neg C$
 - Find rule $\neg P_1 \vee \dots \vee \neg P_n \vee C$
 - Resolve to get $\neg P_1 \vee \dots \vee \neg P_n$
 - Repeat for each negative literal
- First order case introduces unification but otherwise the same.

6.034 - Spring 03 - 13

**Proof Strategy**

- Depth-First search for a proof
- Order matters
 - Rule order
 - try ground facts first
 - then rules in given order
 - Antecedent order
 - left to right
- More predictable, like a program, less like logic

6.034 - Spring 03 - 14

**Slide 11.1.14**

When we specified backchaining we did it with a particular search algorithm (using the stack), which is basically depth-first search. Furthermore, we will assume that the facts and rules are examined in the order in which they occur in the program. Also that literals from the body of a rule are pushed onto the stack in reverse order, so that the one that occurs first in the body will be the first popped off the stack.

Given these ordering restrictions, it is much easier to understand what a logic program will do. On the other hand, one must understand that what it will do is not what a general theorem prover would do with the same rules and facts.

Example

1. Father(A,B) ; ground fact
2. Mother(B,C) ; ground fact
3. GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)
4. Parent(?x,?y) :- Father(?x,?y)
5. Parent(?x,?y) :- Mother(?x,?y)

- Prove: GrandP(?g,C), Ans(?g)
 - [3,?x/?g,?z/C; ?y=?y₁,?g=?g₁]
- Parent(?g₁,?y₁), Parent(?y₁,C), Ans(?g₁)
 - [4,?x/?g₁,?y/?y₁; ?y₁=?y₂,?g₁=?g₂]
- Father(?g₂,?y₂), Parent(?y₂,C), Ans(?g₂)
 - [1,?g₂/A,?y₂/B]

```

Value=[1,C], Ans[A]
<?x>
Value=[1,C], Ans[A]
x
    
```

6.034 - Spring 03 • 18

Slide 11.1.18

The first Parent goal literal unifies with the consequent of rule 4 with the unifier shown. The antecedent (the Father literal) is pushed on the stack, the unifier is applied and the variables are renamed.

Note that there are two Parent rules, we use the first one but we have the other one available should we fail with this one.

Slide 11.1.19

The Father goal literal matches the first fact, which now unifies the ?g variable to A and the ?y variable to B. Note that since we matched a fact, there are no antecedents to push on the stack (as in resolution with a unit-length clause). We apply the unifier, rename and proceed.

Example

1. Father(A,B) ; ground fact
2. Mother(B,C) ; ground fact
3. GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)
4. Parent(?x,?y) :- Father(?x,?y)
5. Parent(?x,?y) :- Mother(?x,?y)

- Prove: GrandP(?g,C), Ans(?g)
 - [3,?x/?g,?z/C; ?y=?y₁,?g=?g₁]
- Parent(?g₁,?y₁), Parent(?y₁,C), Ans(?g₁)
 - [4,?x/?g₁,?y/?y₁; ?y₁=?y₂,?g₁=?g₂]
- Father(?g₂,?y₂), Parent(?y₂,C), Ans(?g₂)
 - [1,?g₂/A,?y₂/B]
- Parent(B,C), Ans(A)
 - [2,?x/B,?y/C]

```

Value=[1,C], Ans[A]
<?x>
Value=[1,C], Ans[A]
x
    
```

6.034 - Spring 03 • 19

Example

1. Father(A,B) ; ground fact
2. Mother(B,C) ; ground fact
3. GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)
4. Parent(?x,?y) :- Father(?x,?y)
5. Parent(?x,?y) :- Mother(?x,?y)

- Prove: GrandP(?g,C), Ans(?g)
 - [3,?x/?g,?z/C; ?y=?y₁,?g=?g₁]
- Parent(?g₁,?y₁), Parent(?y₁,C), Ans(?g₁)
 - [4,?x/?g₁,?y/?y₁; ?y₁=?y₂,?g₁=?g₂]
- Father(?g₂,?y₂), Parent(?y₂,C), Ans(?g₂)
 - [1,?g₂/A,?y₂/B]
- Parent(B,C), Ans(A)
 - [2,?x/B,?y/C]
- Father(B,C), Ans(A)
 - [4,?x/B,?y/C]
- **<fail>**
 - [5,?x/B,?y/C]

```

Value=[1,C], Ans[A]
x
    
```

6.034 - Spring 03 • 20

Slide 11.1.20

Now, we can match the Parent(B,C) goal literal to the consequent of rule 4 and get a new goal (after applying the substitution to the antecedent), Father(B,C). However we can see that this will not match anything in the database and we get a failure.

Slide 11.1.21

The last choice we made that has a pending alternative is when we matched Parent(B,C) to the consequent of rule 4. If we instead match the consequent of rule 5, we get an alternative literal to try, namely Mother(B,C).

Example

1. Father(A,B) ; ground fact
2. Mother(B,C) ; ground fact
3. GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)
4. Parent(?x,?y) :- Father(?x,?y)
5. Parent(?x,?y) :- Mother(?x,?y)

- Prove: GrandP(?g,C), Ans(?g)
 - [3,?x/?g,?z/C; ?y=?y₁,?g=?g₁]
- Parent(?g₁,?y₁), Parent(?y₁,C), Ans(?g₁)
 - [4,?x/?g₁,?y/?y₁; ?y₁=?y₂,?g₁=?g₂]
- Father(?g₂,?y₂), Parent(?y₂,C), Ans(?g₂)
 - [1,?g₂/A,?y₂/B]
- Parent(B,C), Ans(A)
 - [2,?x/B,?y/C]
- Father(B,C), Ans(A)
 - [4,?x/B,?y/C]
- **<fail>**
 - [5,?x/B,?y/C]
- Mother(B,C), Ans(A)
 - [2,?x/B,?y/C]

```

Value=[1,C], Ans[A]
x
    
```

6.034 - Spring 03 • 21

Example

1. Father(A,B) ; ground fact
2. Mother(B,C) ; ground fact
3. GrandP(?x,?z) :- Parent(?x,?y), Parent(?y,?z)
4. Parent(?x,?y) :- Father(?x,?y)
5. Parent(?x,?y) :- Mother(?x,?y)

- Prove: GrandP(?g,C), Ans(?g)
 - [3,?x/?g,?z/C; ?y=?y₁,?g=?g₁]
- Parent(?g₁,?y₁), Parent(?y₁,C), Ans(?g₁)
 - [4,?x/?g₁,?y/?y₁; ?y₁=?y₂,?g₁=?g₂]
- Father(?g₂,?y₂), Parent(?y₂,C), Ans(?g₂)
 - [1,?g₂/A,?y₂/B]
- Parent(B,C), Ans(A)
 - [4,?x/B,?y/C]
- Father(B,C), Ans(A)
 - [5,?x/B,?y/C]
- <fail>
- Mother(B,C), Ans(A)
 - [2]
- Ans(A)

6.034 - Spring 03 + 22

Slide 11.1.22

This matches fact 2. At this point there are no antecedents to add to the stack and the Ans literal is on the top of the stack. Note that the binding of the variable ?g to A is in fact the correct answer to our original question.

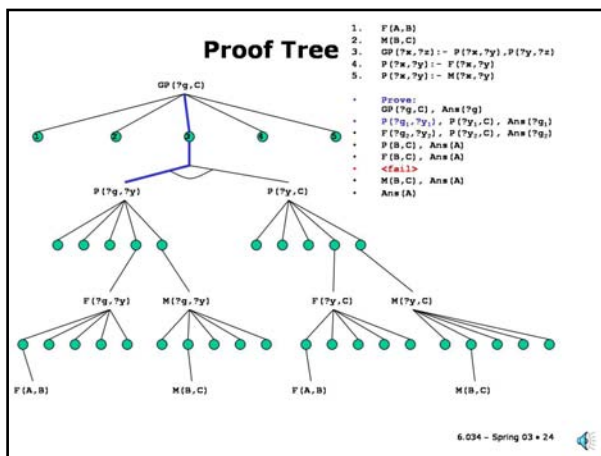
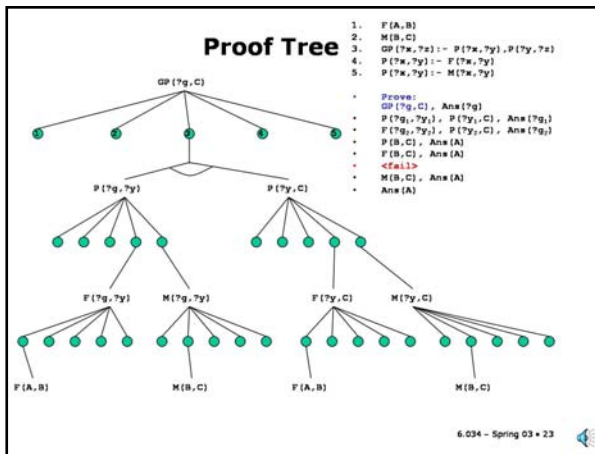
Slide 11.1.23

Another way to look at the process we have just gone through is as a form of tree search. In this search space, the states are the entries in the stack, that is, the literals that appear on our stack. The edges (shown with a green dot in the middle of each edge) are the rules or facts. However, there is one complication: a rule with multiple antecedents generates multiple children, each of which must be solved. This is indicated by the arc connecting the two descendants of rule 3 near the top of the tree.

This type of tree is called an AND-OR tree. The OR nodes come from the choice of a rule or fact to match to a goal. The AND nodes come from the multiple antecedents of a rule (all of which must be proved).

You should remember that such a tree is **implicit** in the rules and facts in our database, once we have been given a goal to prove. The tree is not constructed explicitly; it is just a way of visualizing the search process.

Let's go through our previous proof in this representation, which makes the choices we've made more explicit. We start with the GrandP goal at the top of the tree.

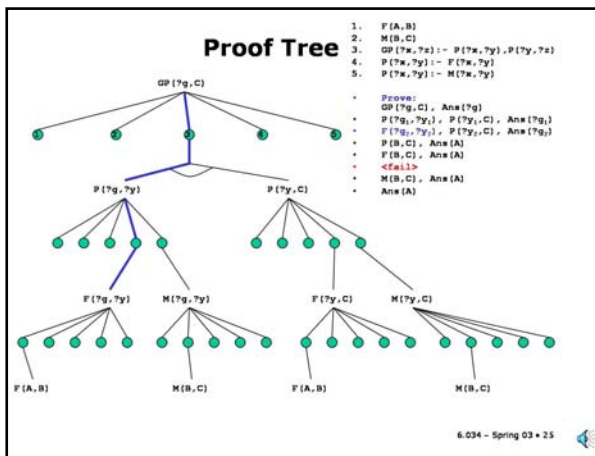


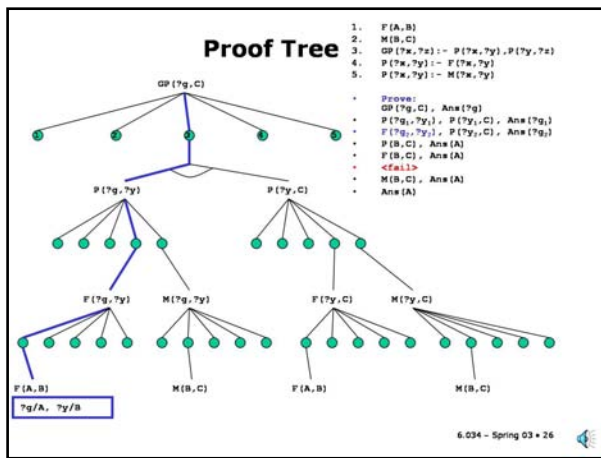
Slide 11.1.24

We match that goal to the consequent of rule 3 and we create two subgoals for each of the antecedents (after carrying out the substitutions from the unification). We will look at the first one (the one on the left) next.

Slide 11.1.25

We match the Parent subgoal to the rule 4 and generate a Father subgoal.



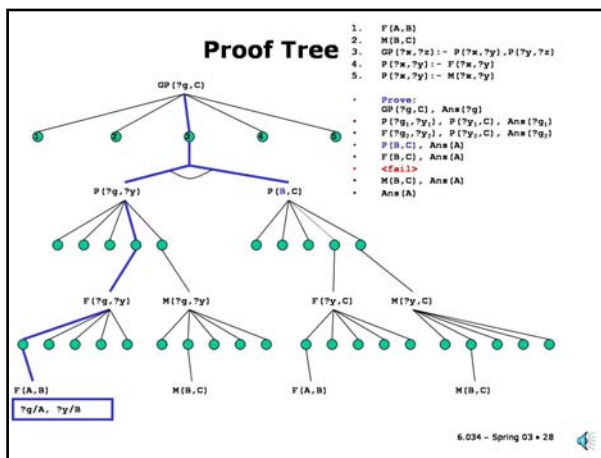
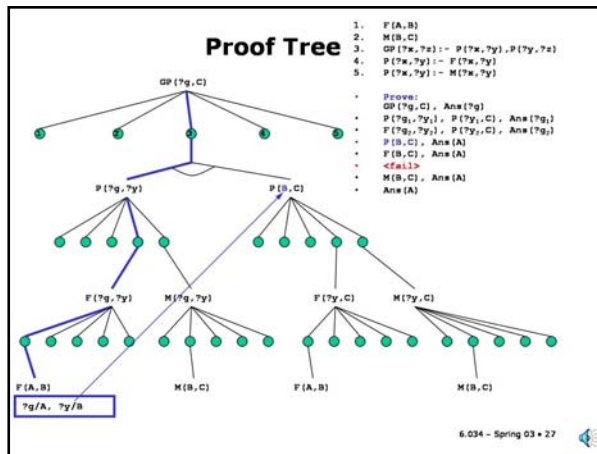


Slide 11.1.26

Which we match to fact 1 and create bindings for the variables in the goal. In all our previous steps we also created variable bindings but they were variable to variable bindings. Here, we finally match some variables to constants.

Slide 11.1.27

We have to apply this unifier to all the pending goals, including the pending Parent subgoal from rule 3. This is the part that's easy to forget when using this tree representation.

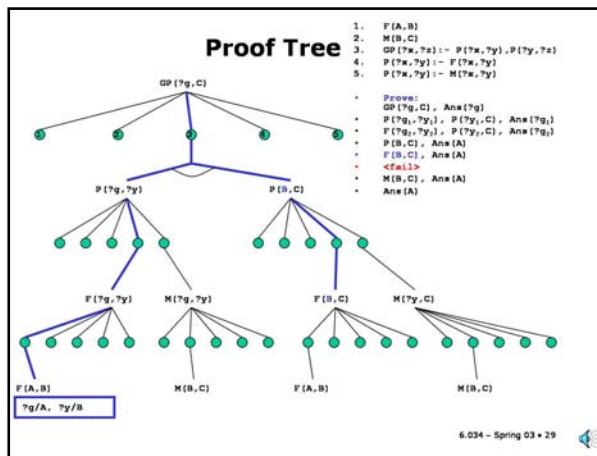


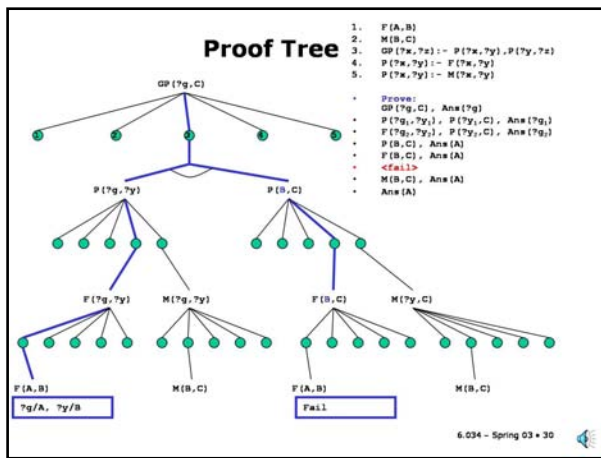
Slide 11.1.28

Now, we tackle the second Parent subgoal ...

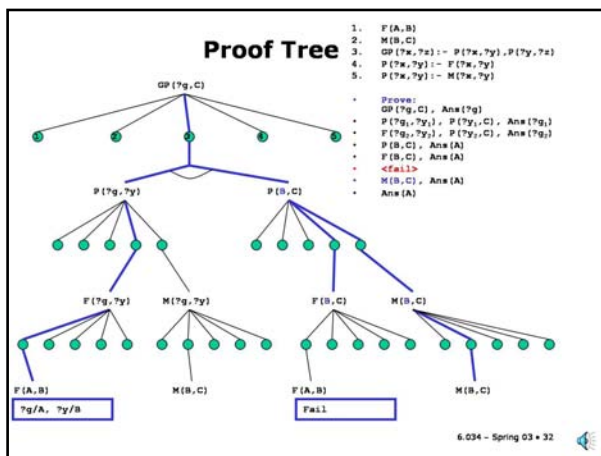
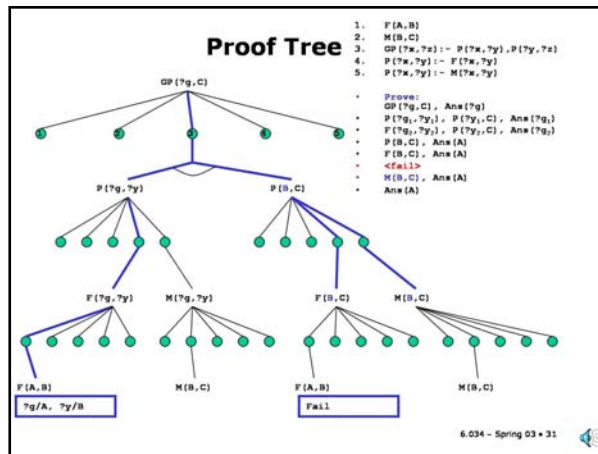
Slide 11.1.29

... which proceeds as before to match rule 4 and generate a Father subgoal, Father(B, C) in this case.





Slide 11.1.31
So, instead, we look at the other alternative, matching the second Parent subgoal to rule 5, and generate a Mother(B,C) subgoal.

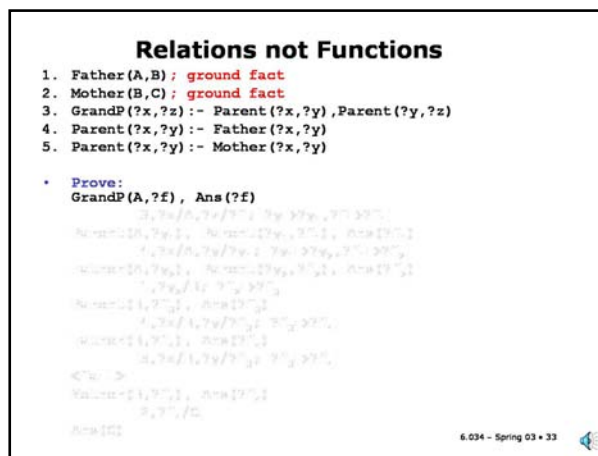


Slide 11.1.32
This matches the second fact in the database and we succeed with our proof since we have no pending subgoals to prove.

This view of the proof process highlights the search connection and is a useful mental model, although it is too awkward for any big problem.

Slide 11.1.33
At the beginning of this section, we indicated as one of the advantages of a logical representation that we could define the relationship between parents and grandparents without having to give an algorithm that might be specific to finding grandparents of grandchildren or vice versa. This is still (partly) true for logic programming. We have just seen how we could use the facts and rules shown here to find a grandparent of someone. Can we go the other way? The answer is yes.

The initial goal we have shown here asks for the grandchild of A, which we know is C. Let's see how we find this answer.



Relations not Functions

1. **Father(A,B) ; ground fact**
2. **Mother(B,C) ; ground fact**
3. **GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)**
4. **Parent(?x,?y) :- Father(?x,?y)**
5. **Parent(?x,?y) :- Mother(?x,?y)**

- **Prove:**
 GrandP(A,?f), Ans(?f)
 - [3,?x/A,?z/?f: ?y=?y₁,?f=?f₁]
- **Parent(A,?y₁), Parent(?y₁,?f₁), Ans(?f₁)**

6.034 - Spring 03 • 34

Slide 11.1.34

Once again, we match the GrandP goal to rule 3, but now the variable bindings are different. We have a constant binding in the first Parent subgoal rather than in the second.

Slide 11.1.35

Once again, we match the Parent subgoal to rule 4 and get a new Father subgoal, this time involving A. We are basically looking for a child of A.

Relations not Functions

1. **Father(A,B) ; ground fact**
2. **Mother(B,C) ; ground fact**
3. **GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)**
4. **Parent(?x,?y) :- Father(?x,?y)**
5. **Parent(?x,?y) :- Mother(?x,?y)**

- **Prove:**
 GrandP(A,?f), Ans(?f)
 - [3,?x/A,?z/?f: ?y=?y₁,?f=?f₁]
- **Parent(A,?y₁), Parent(?y₁,?f₁), Ans(?f₁)**
 - [4,?x/A,?y/?y₁: ?y₁=?y₂,?f₁=?f₂]
- **Father(A,?y₂), Parent(?y₂,?f₂), Ans(?f₂)**

6.034 - Spring 03 • 35

Relations not Functions

1. **Father(A,B) ; ground fact**
2. **Mother(B,C) ; ground fact**
3. **GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)**
4. **Parent(?x,?y) :- Father(?x,?y)**
5. **Parent(?x,?y) :- Mother(?x,?y)**

- **Prove:**
 GrandP(A,?f), Ans(?f)
 - [3,?x/A,?z/?f: ?y=?y₁,?f=?f₁]
- **Parent(A,?y₁), Parent(?y₁,?f₁), Ans(?f₁)**
 - [4,?x/A,?y/?y₁: ?y₁=?y₂,?f₁=?f₂]
- **Father(A,?y₂), Parent(?y₂,?f₂), Ans(?f₂)**
 - [1,?y₂/B: ?f₂=?f₃]
- **Parent(B,?f₃), Ans(?f₃)**

6.034 - Spring 03 • 36

Slide 11.1.36

Then, we match the first fact, namely Father(A,B), which causes us to bind the ?x variable in the second Parent subgoal to B. So, now, we look for a child of B.

Slide 11.1.37

We match the Parent subgoal to rule 4 and generate another Father subgoal, which fails. So, we backup to find an alternative.

Relations not Functions

1. **Father(A,B) ; ground fact**
2. **Mother(B,C) ; ground fact**
3. **GrandP(?x,?z) :- Parent(?x,?y),Parent(?y,?z)**
4. **Parent(?x,?y) :- Father(?x,?y)**
5. **Parent(?x,?y) :- Mother(?x,?y)**

- **Prove:**
 GrandP(A,?f), Ans(?f)
 - [3,?x/A,?z/?f: ?y=?y₁,?f=?f₁]
- **Parent(A,?y₁), Parent(?y₁,?f₁), Ans(?f₁)**
 - [4,?x/A,?y/?y₁: ?y₁=?y₂,?f₁=?f₂]
- **Father(A,?y₂), Parent(?y₂,?f₂), Ans(?f₂)**
 - [1,?y₂/B: ?f₂=?f₃]
- **Parent(B,?f₃), Ans(?f₃)**
 - [4,?x/B,?y/?f₃: ?f₃=?f₄]
- **Father(B,?f₄), Ans(?f₄)**
- **<fail>**

6.034 - Spring 03 • 37

Order Revisited

- Given
 - parent(A,B)
 - parent(B,C)
 - ancestor(?x,?z) :- parent(?x,?z)
 - ancestor(?x,?z) :- parent(?x,?y), ancestor(?y,?z)
- Prove: ancestor(?x,C), Ans(?x)
- Ans(A)
- How about:
 - parent(A,B)
 - parent(B,C)
 - ancestor(?x,?z) :- ancestor(?y,?z), parent(?x,?y)
 - ancestor(?x,?z) :- parent(?x,?z)
- Prove: ancestor(?x,C), Ans(?x)
- <error: stack overflow>
- Clauses examined top to bottom and literals left to right. This is not logic!

6.034 - Spring 03 - 42

Slide 11.1.42

This type of behavior is what you would expect from a recursive program if you put the recursive case before the base case. The key point is that logic programming is half way between traditional programming and logic and exactly like neither one.

Slide 11.1.43

It is often the case that we want to have a condition on a rule that says that something is not true. However, that has two problems, one is that the resulting rule would not be Horn. Furthermore, as we saw earlier, we have no way of concluding a negative literal. In logic programming one typically makes a **closed world** assumption, sometimes jokingly referred to as the "closed mind" assumption, which says that we know everything to be known about our domain. And, if we don't know it (or can't prove it), then it must be false. We all know people like this...

Negation

- We cannot have a rule such as
 - $P_1 \wedge \neg P_2 \rightarrow Q$
 - $\neg P_1 \vee P_2 \vee Q$ - not Horn (two pos literals)
 - Cannot have rule that concludes a negation
- In logic programming, we assume we have complete information about the world (**closed-world assumption**)

```

?- We use "failure to prove" as negation - a dangerous assumption.
?- Prove: ; in empty KB
   not P(?x), Ans(?x)
?- Ans(?x) ; success

```

6.034 - Spring 03 - 43

Negation

- We cannot have a rule such as
 - $P_1 \wedge \neg P_2 \rightarrow Q$
 - $\neg P_1 \vee P_2 \vee Q$ - not Horn (two pos literals)
 - Cannot have rule that concludes a negation
- In logic programming, we assume we have complete information about the world (**closed-world assumption**)
- We use "failure to prove" as negation - a dangerous assumption.
 - Prove: ; in empty KB
 - not P(?x), Ans(?x)
 - Ans(?x) ; success

```

?- We use "failure to prove" as negation - a dangerous assumption.
?- Prove: ; in empty KB
   not P(?x), Ans(?x)
?- Ans(?x) ; success

```

6.034 - Spring 03 - 44

Slide 11.1.44

Given we assume we know everything relevant, we can simulate negation by failure to prove. This is very dangerous in general situations where you may not know everything (for example, it's not a good thing to assume in exams)...

Slide 11.1.45

... but very useful in practice. For example, we can write rules of the form "if there are no other acceptable flights, accept a long layover" and we establish this by looking over all the known flights.

Negation

- But often very useful in finite domains, e.g. flights database, products of a company, etc.
- For example:


```
Layover_not_too_long(?f1, ?f2) :-
  Arrival_time(?f1, ?t1),
  Departure_time(?f2, ?t2),
  not Alternative_connection(?f1, ?t1, ?f2, ?t2)
```
- Will succeed if the Alternative_connection literal fails.

6.034 - Spring 03 - 45

6.034 Notes: Section 11.2

Slide 11.2.1

So far, what we have seen of logic programming may not seem much like programming. Now, we will look at a number of list processing examples that will look more like the examples that you are used to writing in Scheme.

What we will see is essentially a subset of the logic programming language Prolog, which is used fairly widely. There are a number of open-source and commercial versions of Prolog available. We will use a very simple home-brew system implemented in Scheme rather than one of these systems so that there are no mysteries in the implementation. However, we will pay a substantial performance penalty for this choice.

Logic Programming

- So far, not much like programming
- But, this framework can be used as the basis of a general purpose programming language
- Prolog is the most widely used logic programming language
- For example:
 - Gnu Prolog <http://www.gnu.org/software/prolog/prolog.html>
 - SWI Prolog <http://www.swi-prolog.org/>
 - SICStus Prolog <http://www.sics.se/sicstus/>
 - Visual Prolog <http://www.visual-prolog.com/>
 - ...

6.034 - Spring 03 • 1

List Processing: length

```
(define (length y)
  (if (null? y)
      0
      (+ 1 (length (cdr y)))))
```

6.034 - Spring 03 • 2

Slide 11.2.2

Let's start with a very simple Scheme program to compute the length of a list. It's composed of two "cases", the base case when the list is null and the recursive case, in which we reduce the problem into a simpler instance of the same problem (getting the length of the cdr of the list) and compute the final result by adding one to the result of the recursive call.

Slide 11.2.3

This would be a Prolog-like solution to the same problem. It has essentially the same structure as the Scheme program. We use a predicate "length" that has two arguments, one is the list and the other its length.

The first "fact" handles the base case; it defines the length of the null list as 0.

The second rule handles the recursive case. The consequent of the rule (the left-hand-side) is what will match a pending subgoal. Note the form of the first argument of the consequent: it is a Scheme dotted pair. It is set up to match the variable ?h to the car of a list and the variable ?x to the cdr of the list. The second argument of the consequent expresses the length of the list as a function of the length of the cdr of the list.

The right hand side of the rule is the IF part. It sets up a simpler subgoal to solve. Once we solve it, we will have bound ?l to the length of the cdr and we will know the length of the full list (including the car).

Let's look at an example.

List Processing: length

```
(define (length y)
  (if (null? y)
      0
      (+ 1 (length (cdr y))
         )))

• length((), 0)
• length((?h . ?x), ?l+1) :- length(?x, ?l)
```

Recall "dotted pair" notation $(x . y)$ means x is car of list and y is cdr of list. $(cons\ x\ y)$ returns $(x . y)$. In general something like $(a\ b . x)$ indicates that x is rest of list.
 $(a . ()) \equiv (a)$
 $(a\ b . (c\ d)) \equiv (a\ b\ c\ d)$

6.034 - Spring 03 • 3

List Processing: length

1. `length((),0)`
 2. `length(?h . ?x), ?l+1):- length(?x,?l)`
- **Prove:**
`length((a b),?a),Ans(?a)`
`[2,?h/a,?x/(b),?a/?l+1]`
 - a. `length((a b),?l),Ans(?l+1)`
`[rename,l->?l1]`
 - b. `length((b),?l1),Ans(?l1+1)`
`[2,?h/b,?x/(),?l1/?l+1]`
 - c. `length((b),?l),Ans(?l+1+1)`
`[rename,?l->?l2]`
 - d. `length((),?l2),Ans(?l2+1+1)`
`[1,?l2/0]`
 - e. `Ans(0+1+1)`

Slide 11.2.4

You can see the operation of this little program here. The operation is very like that of the corresponding Scheme program. The sequence of subgoals corresponds to the recursive calls to the program.

We have separated the unifier substitution step from the renaming to make things a little clearer. Note that without the renaming we would be hopelessly confused with the bindings of ?l.

In practice, in a Prolog system, the arithmetic expressions would be evaluated by the system and we would get Ans(2).

Slide 11.2.5

This is the same operation but combining the unifier substitution and renaming steps. You can see the sequence of subgoals more clearly here.

List Processing: length

1. `length((),0)`
 2. `length(?h . ?x), ?l+1):- length(?x,?l)`
- **Prove:**
`length((a b),?a),Ans(?a)`
`[2,?h/a,?x/(b),?a/?l+1;`
`?l->?l1]`
 - a. `length((b),?l1),Ans(?l1+1)`
`[2,?h/b,?x/(),?l1/?l+1;`
`?l->?l2]`
 - b. `length((),?l2),Ans(?l2+1+1)`
`[1,?l2/0]`
 - c. `Ans(0+1+1)`

List Processing: length?

1. `length((),0)`
 2. `length(?x,?l-1):- length((?h . ?x),?l)`
- **Prove:**
`length((a b),?a),Ans(?a)`
`?l=?a/0`
`argchk(0,0,0):- (a b),?l,Ans(?l)`
`rename,?l->?l1`
`argchk(0,0,0):- (a b),?l,Ans(?l)`
`?l=?a/0`
`argchk(0,0,0):- (?h . a b),?l,Ans(?l)`
`rename,?l->?l1`
`argchk(0,0,0):- (?h . a b),?l,Ans(?l)`
`?l=?a/0`
`fail`

Slide 11.2.6

You may be wondering whether this formulation of length would also work. Certainly, it seems just as valid as the one we used. Let's trace it through. We start with the same goal as before, finding the length of the list (a b).

Slide 11.2.7

We match the goal to the consequent of rule 2 and do the substitution to get a new subgoal. Note, however, that this is not a simpler subgoal. It's actually trying to find the length of a longer, not completely specified, list. If we knew the length of such a list then we could know the length of our input list. Can you smell trouble brewing?

List Processing: length?

1. `length((),0)`
 2. `length(?x,?l-1):- length((?h . ?x),?l)`
- **Prove:**
`length((a b),?a),Ans(?a)`
`- [2,?x/(a b),?a/?l-1]`
 - `length((?h . (a b)),?l),Ans(?l-1)`
`rename,?l->?l1`
`argchk(0,0,0):- (a b),?l,Ans(?l)`
`?l=?a/0`
`argchk(0,0,0):- (?h . a b),?l,Ans(?l)`
`rename,?l->?l1`
`argchk(0,0,0):- (?h . a b),?l,Ans(?l)`
`?l=?a/0`
`fail`

List Processing: length?

```

1. length((),0)
2. length(?x,?l-1) :- length((?h . ?x),?l)

• Prove:
  length((a b),?a),Ans(?a)
  - [2,?x/(a b),?a/?l-1]
• length((?h . (a b)),?l),Ans(?l-1)
  - [rename,l=>?l1,?h=>?h1]
• length((?h1 a b),?l1),Ans(?l1-1)
  - [2,?x/(?h1 a b),?l1/?l1-1]
• length((?h . (?h1 a b)),?l),Ans(?l-1-1)
  - [rename,?l=>?l2,?h=>?h2,?h1=>?h3]
• length((?h2 ?h3 a b),?l2), Ans(?l2-1-1)
• etc

```

6.034 - Spring 03 • 8

Slide 11.2.8

Sure enough; we've coded an infinite loop. The moral of the story is that you want to write these recursive rules in the form of "complex consequent :- simple antecedent" and not the other way around.

Slide 11.2.9

Here's another formulation of length that does work. Here we've separated the updating of the length into a separate equality statement. The software that we will be using will require this particular form.

List Processing: length

- Another equivalent formulation
 - length((),0)
 - length((?h . ?x),?l) :- length(?x,?lw),?l=?lw+1

6.034 - Spring 03 • 9

List Processing: append

```

(define (append x y)
  (if (null? x)
      y
      (cons (car x) (append (cdr x) y))))

append([],?y,?y)
append([?x . ?x],?y,[?z . ?z]) :- append(?x,?y,?z)

```

6.034 - Spring 03 • 10

Slide 11.2.10

Let's look at another example, which is quite parallel to length. Here is a Scheme implementation of an append function. It too consists of two cases. The base case handles the case of the first argument being null, in which case the answer is simply the second list. The recursive case involves computing the solution to a simpler case (append of the cdr of x to y) and updating it to the final answer by consing the car of x to the result.

Slide 11.2.11

The logic program is completely analogous. The append predicate has three arguments, the lists to be appended and the result list. The first fact just says that the output of appending the null list and any list is just the second list. The second rule looks more complicated but it is just like the Scheme program. We pick out the car and the cdr in the consequent (note the use of dotted pair notation) and bind them to ?h and ?x respectively. Then we define a subgoal involving ?x and ?y and bind the result to ?z. We can then construct the result for the original list by consing ?h to ?z (using dotted pair notation).

List Processing: append

```

(define (append x y)
  (if (null? x)
      y
      (cons (car x) (append (cdr x) y))))

• append((),?y,?y)
• append((?h . ?x),?y,(?h . ?z)) :- append(?x,?y,?z)

```

Recall "dotted pair" notation (x . y) means x is car of list and y is cdr of list. (cons x y) returns (x . y). In general something like (a b . x) indicates that x is rest of list.

(a . ()) ≡ (a)
(a b . (c d)) ≡ (a b c d)

6.034 - Spring 03 • 11

List Processing: append

1. `append((), ?y, ?y)`
2. `append(?h . ?x, ?y, (?h . ?z)) :- append(?x, ?y, ?z)`

- **Prove:**
`append(a b, (c d), ?l), Ans(?l)`
`[2, ?h/a, ?x/(b), ?y/(c d), ?l/(a . ?z); ?z=>?z1]`
- a. `append(b, (c d), ?z1), Ans((a . ?z1))`
`[2, ?h/b, ?x/(), ?y/(c d), ?z1/(b . ?z); ?z=>?z2]`
- b. `append((), (c d), ?z2), Ans((a . (b . ?z2)))`
`[1, ?y/?z2, ?z2/(c d)]`
- c. `Ans((a . (b . (c d))))`
 - Note that `(a . (b . (c d)))` is `(a b c d)`

6.034 - Spring 03 • 12

Slide 11.2.12

Here you can trace the operation of these rules in a very simple example. Once again note that the renaming is crucial for keeping things straight.

Slide 11.2.13

Thus far, the structure of the logic programs we have seen is very similar to that of the natural Scheme programs. But that's not always the case. Let's look at an alternative way of representing lists that leads to very different looking programs. The representation is called **difference lists**. You can see some examples of this representation of a simple list with three elements here. The most important one is `diff((a b c) . ?x, ?x)`, which says that any list starting with `(a b c)` and followed by anything can be used to represent the list.

Difference Lists

- A list can be represented as the difference between two lists, which we will write `diff(L1, L2)`
- For example, `(a b c)` can be written as
 - `diff((a b c), ())`
 - `diff((a b c d), (d))`
 - `diff((a b c d e), (d e))`
 - `diff((a b c . ?x), ?x)`
- The empty list is any list of the form
 - `diff(?x, ?x)`

6.034 - Spring 03 • 13

Difference Lists

- A list can be represented as the difference between two lists, which we will write `diff(L1, L2)`
- For example, `(a b c)` can be written as
 - `diff((a b c), ())`
 - `diff((a b c d), (d))`
 - `diff((a b c d e), (d e))`
 - `diff((a b c . ?x), ?x)`
- The empty list is any list of the form
 - `diff(?x, ?x)`
- In `diff(L1, L2)` think of `L1` as a pointer to the beginning of the list and `L2` as a pointer to the end of the list.

6.034 - Spring 03 • 14

Slide 11.2.14

The basic idea is that we can represent a list by a pair of pointers into a bigger list, one to the beginning and the other to the end of the list.

Slide 11.2.15

In this representation we can code `append` as a single fact! The picture shows the intuition behind the definition. On first viewing this seems like we're cheating. It's easy to see that this statement is true, but how does it actually compute anything? Partly, one has to think carefully about the representation of the input.

List Processing: dappend

1. `dappend(diff(?x, ?y), diff(?y, ?z), diff(?x, ?z))`

`?x = (1 2 3 4 5 6 7 8 9 10 11 12 13)`
`?y = (6 7 8 9 10 11 12 13)`
`?z = (12 13)`

6.034 - Spring 03 • 15

List Processing: dappend

1. `dappend(diff(?x,?y),diff(?y,?z),diff(?x,?z))`

• **Prove:**
`dappend(diff((a b . ?p),?p),diff((c d . ?q),?q),?w)`



6.034 - Spring 03 • 16

Slide 11.2.16

Here we see how a goal would be phrased in this representation. We have used the most general representation of the input lists.

Slide 11.2.17


Now, we unify. Note that unification ends up equating `?y`, the end of the first list in the rule, with `?p`, the end of the first input list. Then `?p` (and therefore `?y`) is matched to the start of the second list. This is basically what carries out the append operation.

Yes, it looks like magic. We will be using difference lists a fair bit when we do natural language processing, so it is worth spending a bit of time understanding them.

List Processing: dappend

1. `dappend(diff(?x,?y),diff(?y,?z),diff(?x,?z))`

• **Prove:**
`dappend(diff((a b . ?p),?p),diff((c d . ?q),?q),?w)`
 - `[?x/(a b . ?p),`
 `?y/?p,`
 `?p/(c d . ?q),`
 `?z/?q,`
 `?w/diff((a b . (c d . ?q)), ?q)]`
 • **Note that `diff((a b . (c d . ?q)), ?q)` is equivalent to `diff((a b c d . ?q), ?q)` - which is correct.**



6.034 - Spring 03 • 17

List Processing: reverse

```
(define (reverse l)
  (define (reversel x y)
    (if (null? x)
        y
        (reversel (cdr x) (cons (car x) y))))
  (reversel l '()))
```

```
reverse([? , ?rev]) :- reverse([? , [] , ?rev]
reverse([], ?y, ?y)
reverse([?h . ?r], ?y, ?z) :-
reverse([?r, [?h . ?y], ?z)
```

6.034 - Spring 03 • 18

Slide 11.2.18

Let's look at another example, first without difference lists and then with.

This is Scheme for a list reverse operation. It's a bit more complicated than the cases we've seen so far. To reverse the list, we need a temporary value to serve as an accumulator for the reversed list. That's what the `y` argument to the inner procedure is. `y` starts with the null list and we cons each of the elements of the input onto this list. When the first argument is null, we return the accumulated list.

Slide 11.2.19

We follow the same pattern in the logic program. We define the predicate `reverse`, with two arguments, the input and output lists, in terms of a three-place auxiliary predicate `reversel`, which introduces the accumulator and initializes it to nil. Note that if you reverse the first argument of `reversel` and append it to the second argument of `reversel` then that gives the answer to the original query.

`reversel` is defined by two rules: in the base case when the first argument is nil, we simply equate the output list to the accumulator. In the general case, we set up a recursive subgoal with the `cdr` of the list, but we cons the car of the input list to the accumulator.

List Processing: reverse

```
(define (reverse l)
  (define (reversel x y)
    (if (null? x)
        y
        (reversel (cdr x) (cons (car x) y))))
  (reversel l '()))
```

• `reverse(?l, ?rev) :- reversel(?l, (), ?rev)`
 • `reversel([], ?y, ?y)`
 • `reversel([?h . ?r], ?y, ?z) :- reversel(?r, (?h . ?y), ?z)`

6.034 - Spring 03 • 19

List Processing: reverse

```
1. reverse(?l, ?rev) :- reversel(?l, (), ?rev)
2. reversel((), ?y, ?y)
3. reversel(?h . ?r, ?y, ?z) :-
   reversel(?r, (?h . ?y), ?z)
```

• Prove:

```
reverse( (a b), ?v), Ans(?v)
[1,?l/(a b),?rev/?v]
a. reversel( (a b), (), ?v), Ans(?v)
[3,?h/a,?r/(b),?y/(),?z/?v]
b. reversel( (b), (a . ()), ?z), Ans(?v)
[3,?h/b,?r/(),?y/(a),?z/?v]
c. reversel( (), (b . (a)), ?v), Ans(?v)
[2,?v/?y,?y/(b a)]
d. Ans( (b a))
```

6.034 - Spring 03 • 20



Slide 11.2.20

You can see the operation on a simple example here.

Slide 11.2.21

Here is the implementation using a difference list instead of an explicit accumulator. Note that the end of the difference list essentially behaves as the accumulator variable did in the previous implementation.

Hopefully, this has given you some flavor for logic programming. We will see more examples of these types of rules in the next chapter.

List Processing: dreverse

```
1. dreverse(?l, ?rev) :- dreversel(?l, diff(?rev, ()))
2. dreversel((), diff(?y, ?y))
3. dreversel(?h . ?r, diff(?ys, ?ye)) :-
   dreversel(?r, diff(?ys, (?h . ?ye)))
```

6.034 - Spring 03 • 21

