



# Lecture 7

## Implementing Prolog

unification, backtracking with coroutines

Ras Bodik  
Shaon Barman  
Thibaud Hottelier

*Hack Your Language!*

CS164: Introduction to Programming  
Languages and Compilers, Spring 2012  
[UC Berkeley](#)

# Where are we heading today?

---

Today, we'll go deeper into the territory of *programming under abstraction* –

developing abstractions that others can conveniently use.

Previously in cs164, we built constructs with yield

iterators (L4),

lazy list concatenation (HW2),

regexes based on backtracking (HW2)

Today, we will build Prolog, an entirely new language

PA3 is assigned today: Prolog on top of your PA2 coroutines

# Today

---

Find a partner. Get a paper and pencil.

You will solve a series of exercises leading to a Prolog interpreter.

# Prolog refresher

---

## Program:

```
eat(thibaud, vegetables).  
eat(thibaud, fruits).  
eat(lion, thibaud).
```

## Queries:

```
eat(thibaud, lion)?  
eat(thibaud, X)?
```

→ fails  
→ vegetables;  
fruits;  
no more solutions

no solution  
two solutions

# Structure of Programs

---

works(ras).



Fact (Axiom)

works(thibaud) :- works(ras).

Rule

works(X)?

Query



Clause

Variable

Constant

In a rule:

*bound (not free)*  
*clause*

a(X, Y) :- b(X, Z), c(Z, Y)

Free Variable

Head

Body



# Unification

Unification is what happens during matching.

What does it mean to be compatible?

$a(1, Y)$		$a(X, 2)$ yes	$\left\{ \begin{array}{l} Y \mapsto 2 \\ X \mapsto 1 \end{array} \right.$
$a(X)$		$b(X)$ no $a \neq b$	
$a(1, Y)$		$a(2, X)$ no $1 \neq 2$	
$a(1, Y)$		$a(1, X)$ yes	$X \mapsto Y$

*more general than*

$X = Y = 1$

A call to `unify(term1, term2)` yields *most general unifier* (mgu)

# Two exercises


Find mgu for these two unifications:

$$a(X, Y) \quad | \quad a(b(Y), c(Z))$$

mgu:

$$\begin{aligned} X &\mapsto b(c(-z)) \\ Y &\mapsto c(-z) \end{aligned}$$

same  $\uparrow$  ? Yes



$$a([1|X]) \quad | \quad a(X)$$

mgu:  ~~$X = [1..]$~~  infinite stream of 1's

not in Prolog. The substitution for  $X$  must be finite

# Unification algorithm

---


See the simple description in *The Art of Prolog*

Chapter 4.1, pages 88-91.

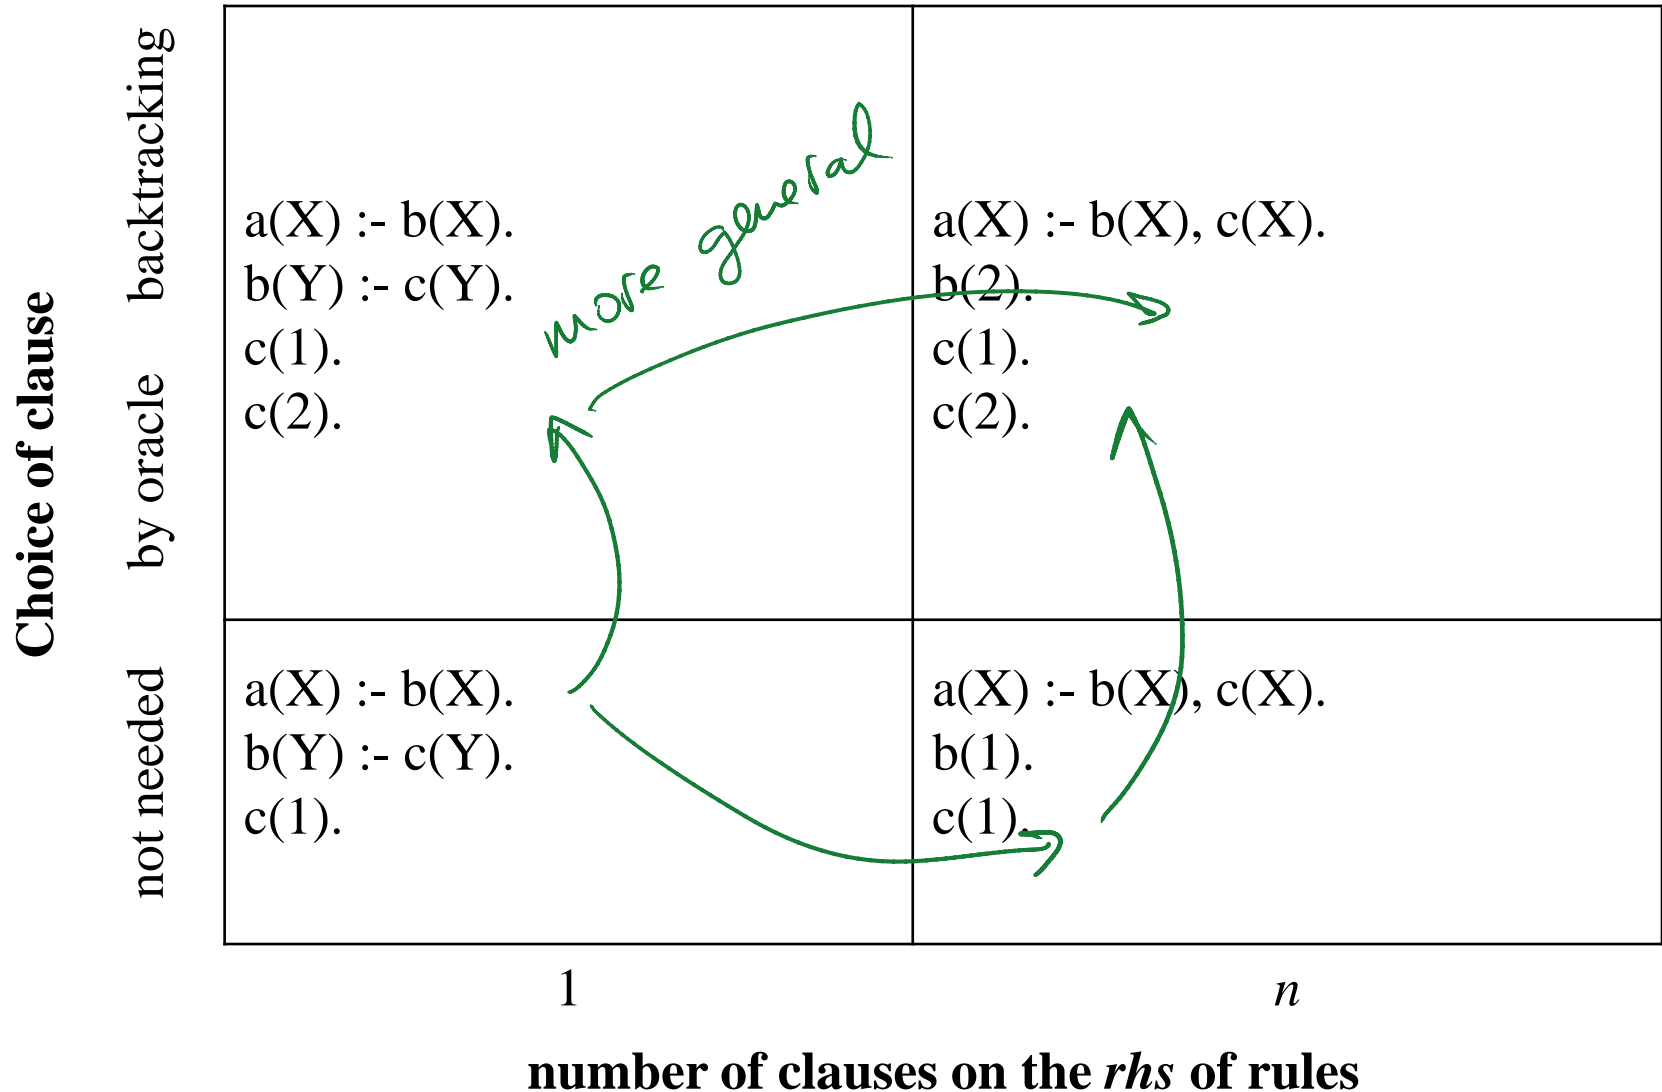
part of required reading



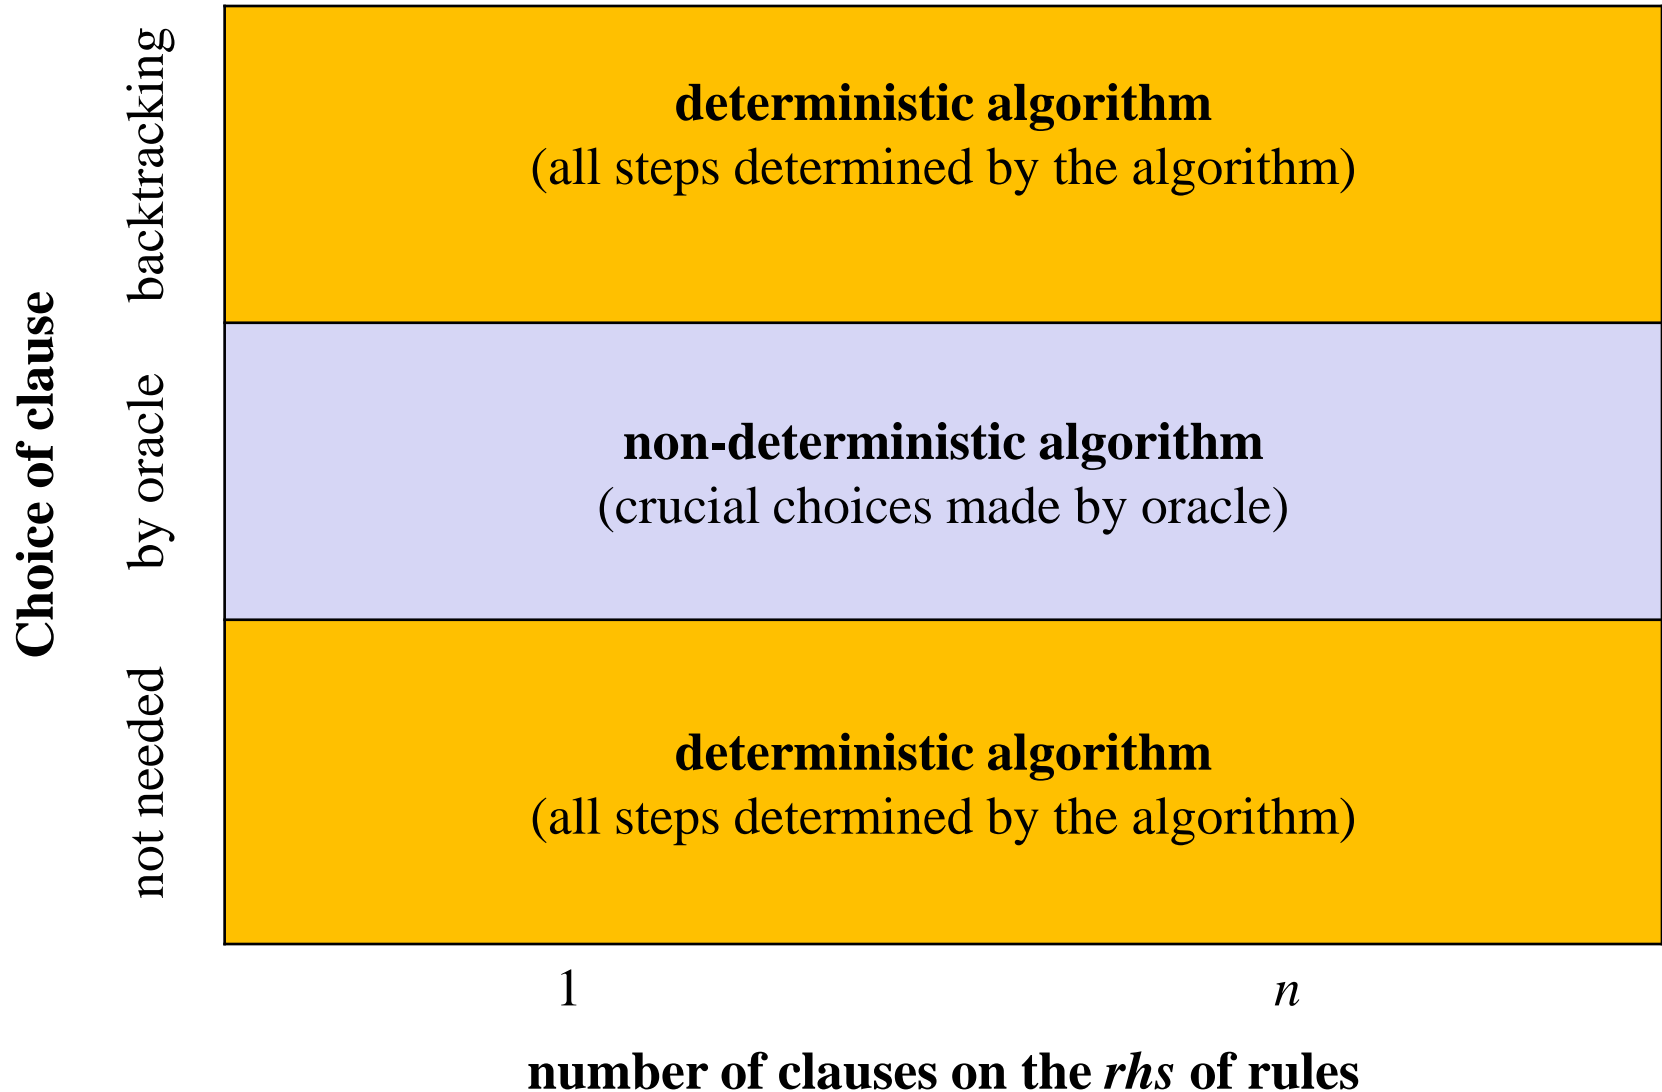
# Today, you will design a series of algorithms



# We will start with subsets of Prolog



# Some algorithms will use “magic”



# Algorithm (1, no choice)

<b>Choice of clause</b>	backtracking		
	by oracle		
	not needed	$a(X) :- b(X).$ $b(Y) :- c(Y).$ $c(1).$	
		1	$n$
		<b>number of clauses on the <i>rhs</i> of rules</b>	

# Prolog execution is finding a proof of query truth

---

Program:

$a(X) :- b(X).$

$b(Y) :- c(Y).$

$c(1).$

Goal (query):

$?- a(Z).$

Answer:

true

$Z = 1$

Proof that the query holds:

$c(1)$  base fact, implies that ...  
 $c(Y)$  holds, which implies that ...  
 $b(Y)$  holds, which implies that ...  
 $b(X)$  holds, which implies that ...  
 $a(X)$  holds, which implies that ...  
 $a(Z)$  holds.

The last one is the query

so the answer is true!

Recall “ $c(Y)$  holds” means

exists value for  $Y$  such that  $C(Y)$  holds.

# Proof tree

---

Program:

`a(X) :- b(X).`

`b(Y) :- c(Y).`

`c(1).`

Goal (query):

`?- a(Z).`

Answer:

`true`

`Z = 1`

These steps form a proof tree

`a(Z)`

`a(X)`

`b(X)`

`b(Y)`

`c(Y)`

`c(1)`

`true`

*N.B. this would be a proof tree, rather than a chain, if rhs's had multiple goals.*

# Let's trace the process of the computation

Program:

$a(X) :- b(X).$   
 $b(Y) :- c(Y).$   
 $c(1).$

Goal (query):

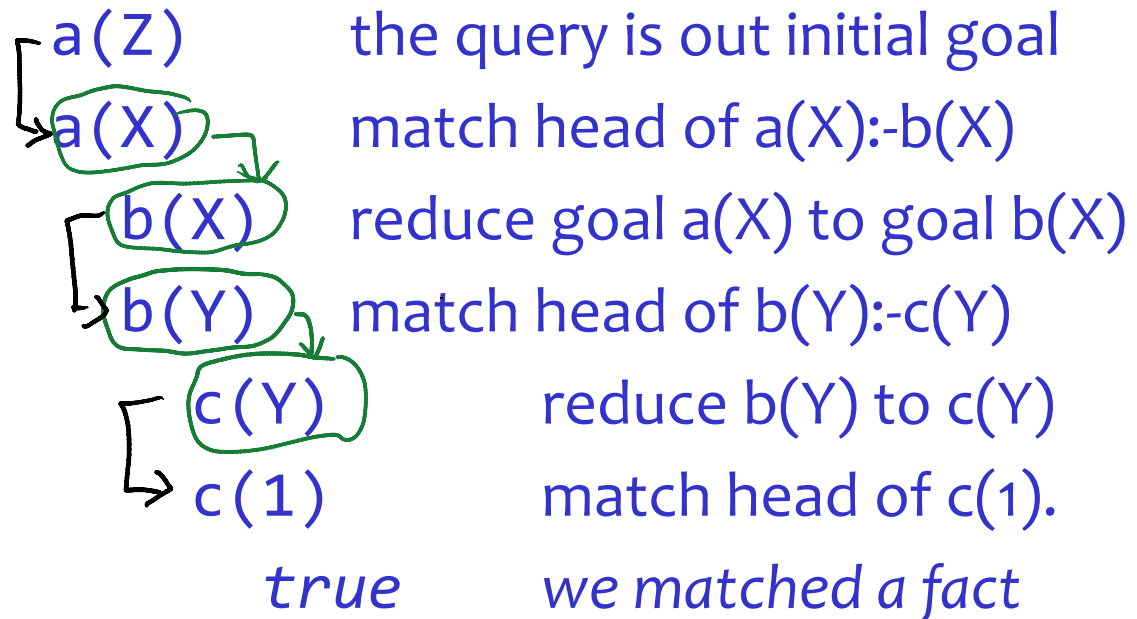
$?- a(Z).$

Answer:

true

$Z = 1$

Two operations do all the work:



The operations:

- 1) match goal to a head of clause C
- 2) reduce goal to rhs of C

↳ matches produce mgu ←  
rhs rewrite goals to new subgoals

# Now develop an outline of the interpreter

---

Student answer:

process (A: goal)

for each head H of a clause C

if unify (A, H):

new-goal = H

→ new-goal = C.rhs

return process (new-goal)

if TRMS  
return  
OK.

return fail



# Algorithm (1,no choice) w/out handling of mgus

---

```
def solve(goal):
    match goal against the head C.H of a clause C
    // how many matches are there? Can assume 0/1
    if no matching head found:
        return FAILURE // done
    if C has no rhs:
        return SUCCESS // done, found a fact
    else // reduce the goal to the rhs of C
        return solve_goal(C.rhs)
```

**Note:** we ignore the handling of mgus here, to focus on how the control flows in the algorithm. We'll do mgus next ...

# Concepts: Reduction of a goal. Unifier.

---

We reduce a goal to a subgoal

If the current goal matches the head of a clause  $C$ , then we reduce the goal to the rhs of  $C$ .

Result of solving a subgoal is a unifier (mgu)

or false, in the case when the goal is not true

But what do we do with the unifiers?

are these mgus merged? If yes, when?

# An algorithmic question: when to merge mgus

Program:

```
a(X) :- b(X).  
b(Y) :- c(Y).  
c(1).
```


Goal (query):

```
?- a(Z).
```

Answer:

```
true  
Z = 1
```

Unifications created in matching



Result is conjunction of these mgus:

$Z=X, X=Y, Y=1$

So, the answer is  $Z=1$

internal variables  $X, Y$  are suppressed

# Design question: How do MGUs propagate?

Down the recursion? or ...

a(Z)

a(X)

b(X)

b(Y)

c(Y)

c(1)

true

Z=X

Z=X,

X=Y

Z=X,

X=Y,

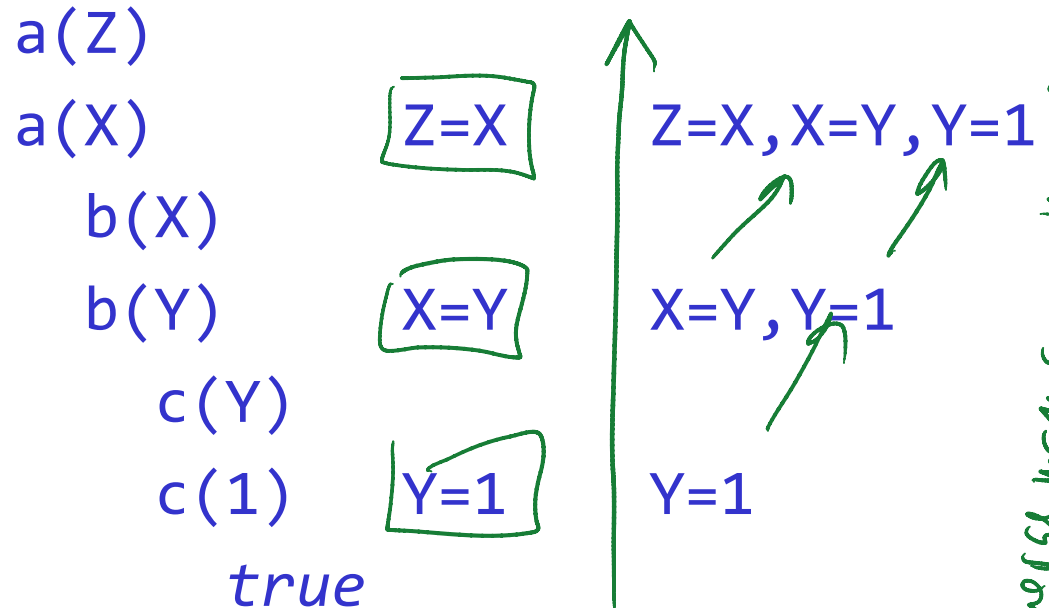
Y=1

merge mgus  
on the way down?



# MGUs propagate the answer

... up the recursion or ...



merge mgus on the way up?

# MGUs propagate the answer

---

... or both?

a(Z)

a(X)

b(X)

b(Y)

c(Y)

c(1)

*true*

Z=X

Z=X, X=Y

Z=X, X=Y, Y=1

Z=X, X=Y, Y=1

Z=X, X=Y, Y=1

Z=X, X=Y, Y=1

↓ accumulate  
mgus

↑ propagate  
constraints up

# Both up and down propagation is needed

---

Consider program:

$a(X, Y, Z) \text{ :- } b(X, Y, Z).$

$b(A, B, C) \text{ :- } c(A, B), d(C).$

$c(1, 2).$

$d(1).$

Down propagation: needed to propagate constraints

given query  $a(X, X, Z)?$ , goal  $c(X, Y)$  must be reduced to  $c(X, X)$  so that match with  $c(1, 2)$  fails

Up propagation: needed to compute the answer to  $q$ .

given query  $a(X, Y, Z)?$ , we must show that  $Z=1$  is in the result. So we must propagate the mgus up the recursion.

# Algorithm (1,no choice) with unification, style 1

---

```
solve(goal, mgu):
    // match goal against the head C.H of a
    // clause C, producing a new mgu.
    // unify goal and head wrt constraints in mgu
mgu = unify(goal, head, mgu)
if no matching head found:
    return nil // nil signifies FAILURE
if C has no rhs:
    return mgu // this signifies SUCCESS
else
    // solve and return the updated mgu
    return solve(C.rhs, mgu)
```



# Algorithm (1,no choice) with unification, style 2

---

```
solve(goal, mgu):
```

```
    // mgus've been substituted into goal and head
```

```
    mgu = unify(goal,head)
```

```
    if no matching head found:
```

```
        return nil // nil signifies FAILURE
```

```
    if C has no rhs:
```

```
        return mgu // this signifies SUCCESS
```

```
    else
```

```
        sub_goal = substitute(mgu,C.rhs)
```

```
        sub_mgu = solve(sub_goal)
```

```
        return merge(mgu, sub_mgu)
```

# Summary of Algorithm for (1, no choice)

---

The algorithm is a simple recursion that reduces the goal until we answer true or fail.

the match of a goal with a head produces the mgu

The answer is the most general unifier

if the answer is true

mgus are unified as we return from recursion

This algorithm is implemented in the PA3 starter kit

# Discussion

---

## Style 1:

unify() performs the substitution of vars in goal, head based on the mgu argument. This is expensive.

## Style 2:

mgus are substituted into new goals. This is done just once. But we need to merge the mgus returned from goals.

This merge always succeeds (conflicts such as  $X=1, X=2$  can't arise)

PA3 uses the second style.

In the rest of the lecture, we will abstract mgus.

You'll retrofit handling of mgus into algorithms we'll cover.

# Unify and subst used in PA3

---

**unify:** Are two terms compatible? If yes, give a unifier  
 $a(X, Y) \mid a(1, 2) \rightarrow \{X \rightarrow 1, Y \rightarrow 2\}$

**subst:** Apply Substitution on clauses  
 $\text{subst}[a(X, Y), \{X \rightarrow \text{ras}, Y \rightarrow Z\}] \rightarrow a(\text{ras}, Z)$

# Example executed on PA3 Prolog

---

a(X) :- b(X).

b(Y) :- c(Y).

c(1).

a(I)?

Goal: a(I)

Unify: a(X\_1) and a(I)

Unifier: {X\_1->I }

Goal: b(I)

Unify: a(X\_2) and b(I)

Unifier: null

Unify: b(Y\_3) and b(I)

Unifier: {Y\_3->I }

Goal: c(I)

Unify: a(X\_4) and c(I)

Unifier: null

Unify: b(Y\_5) and c(I)

Unifier: null

Unify: c(1) and c(I)

Unifier: {I->1 }

I = 1

Asking for solution 2

Unify: c(1) and b(I)

Unifier: null

Unify: b(Y\_8) and a(I)

Unifier: null

Unify: c(1) and a(I)

Unifier: null

None

# Algorithm (n, no choice)

<b>Choice of clause</b>	backtracking		
	by oracle		
	not needed	New concepts: <b>unifier, proof tree</b> Implementation: <b>reduce a goal and recurse</b>	<b><math>a(X) :- b(X), c(X).</math></b> <b><math>b(1).</math></b> <b><math>c(1).</math></b>
		1	<i>n</i>
		<b>number of clauses on the <i>rhs</i> of rules</b>	

# Resolvent

---

**Resolvent:** the set of goals that need to be answered  
with one goal on rhs, we have always just one pending goal

Resolvent goals form a stack. The algorithm:

- 1) pop a goal
- 2) finds a matching clause for a goal, as in (1, no choice)
- 3) if popped goal answered, goto 1
- 4) else, push goals from rhs to the stack, goto 1

This is a conceptual stack.

Need not be implemented as an explicit stack

# Algorithm

---

For your reference, here is algorithm (1,no choice)

```
solve_goal(goal):  
    match goal against the head C.H of a clause C  
  
    if no matching head found:  
        return FAILURE  
  
    if C has no rhs:      // C is a fact  
        return SUCCESS  
  
    else                // reduce the goal to the rhs of C  
        return solve(C.rhs)
```



# Student algorithm

---

# What to change in (n, no choice)?

---

```
solve(goal):
```

```
    match goal against a head C.H of a clause C
```

```
    if no matching head found:
```

```
        return FAILURE
```

```
    if C has no rhs: // C is a fact
```

```
        return SUCCESS
```

```
    else // reduce goal to the goals in the rhs of C
```

```
        for each goal in C.rhs
```

```
            if solve(goal) == FAILURE
```

```
                // oracle failed to find a solution for goal
```

```
                return FAILURE
```

```
        end for
```

```
        // goals on the rhs were solved successfully
```

```
        return SUCCESS
```

# Your exercise

---

Add handling of mgus to (n, no choice)

# Summary

---

The for-loop across rhs goals effectively pops the goals from the top of the conceptual resolver stack

# Example executed on PA3 Prolog

---

a(X) :- b(X), c(X).

b(1).

c(1).

a(1)?

Asking for solution 1

Goal: a(I)

Unify: a(X\_1) and a(I)

Unifier: {X\_1->I }

Goal: b(I)

Unify: a(X\_2) and b(I)

Unifier: null

Unify: b(1) and b(I)

Unifier: {I->1 }

Goal: c(1)

Unify: a(X\_4) and c(1)

Unifier: null

Unify: b(1) and c(1)

Unifier: null

Unify: c(1) and c(1)

Unifier: {}

I = 1

Asking for solution 2

Unify: c(1) and b(I)

Unifier: null

Unify: b(1) and a(I)

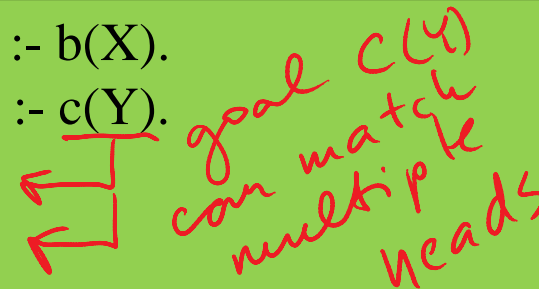
Unifier: null

Unify: c(1) and a(I)

Unifier: null

None

# Algorithm (1, oracular choice)

Choice of clause	backtracking		
	by oracle	$a(X) :- b(X).$ $b(Y) :- c(Y).$ $c(1).$ $c(2).$ 	
	not needed	New concepts: <b>unifier, proof tree</b> Implementation: <b>reduce a goal and recurse</b>	Concept: <b>resolvent</b> Implementation: <b>recursion deals with reduced goals; iteration deals with rhs goals</b>

1

$n$

number of clauses on the *rhs* of rules

# Search tree

---

First, assume we want just one solution (if one exists)

- ie, no need to enumerate all solutions in this algorithm

We'll visualize the space of choices with a search tree

- Node is the current goal
- Edges lead to possible reductions of the goal

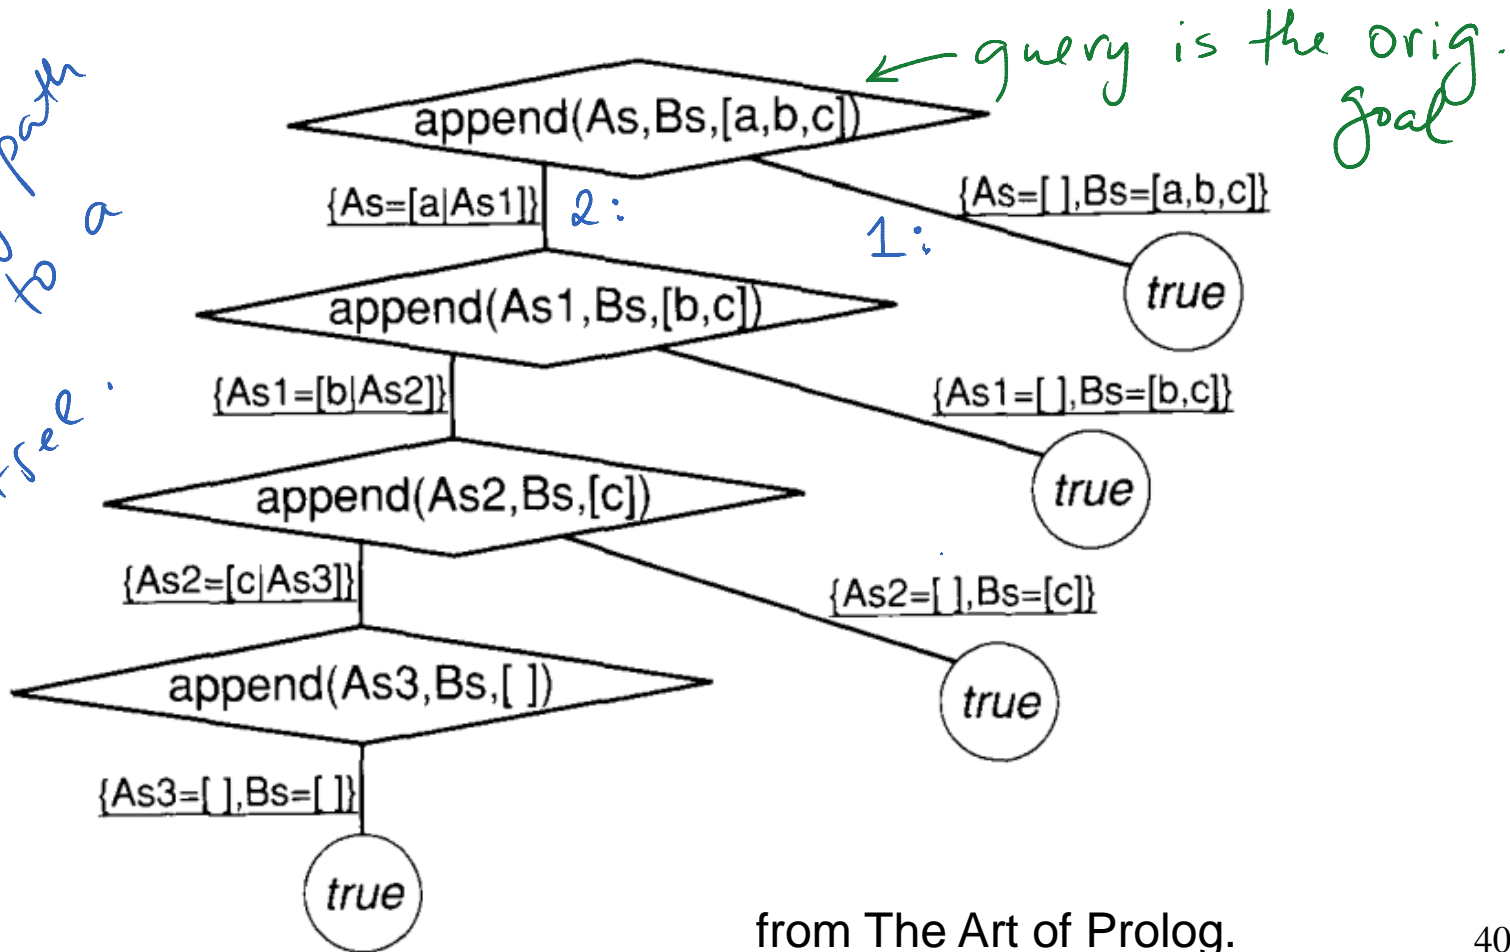
Number of children of a node depends on \_\_\_\_\_

your answer: *number of heads matching the goal*

# Example search tree (for Append)

1: `append([], Ys, Ys).`

2: `append([X|Xs], Ys, [X|Zs1]) :- append(Xs, Ys, Zs).`





# Algorithm

---

student answer:

# Algorithm for (1,oracle choice)

---

solve(goal):

    match goal against a head C.H of a clause C

**if multiple matches exist: ask the oracle to pick one**

    if no matching head found:

        return FAILURE

    if C has no rhs:

        return SUCCESS

    else

        solve(C.rhs)

Oracle is guaranteed to pick a head that is part of a proof tree  
assuming a solution exists

# Summary

---

We relied on an oracle to make just the right choice

The choice is clairvoyant: takes into consideration choices to be made by oracles down the search tree

Asking an oracle is known as non-determinism. It simplifies explanations of algorithms.

We will have to implement the oracle with backtracking in (1, backtracking)

# Algorithm (n, oracular choice)

<b>Choice of clause</b> by oracle backtracking not needed	1	
	n	New concept: <b>search tree</b> Implementation: <b>ask oracle for the right choice.</b>
	New concepts: <b>unifier, proof tree</b> Implementation: <b>reduce a goal and recurse</b>	Concept: <b>resolvent</b> Implementation: <b>recursion deals with reduced goals; iteration deals with rhs goals</b>
	1	n
	<b>number of clauses on the <i>rhs</i> of rules</b>	

# Analysis of this problem

---

Nothing too different from (1,oracle), except that we are dealing with a resolvent (ie, 2+ pending goals)

We deal with them as in (n, no choice), by reducing the goal on top of the conceptual stack

As in (1,oracular choice), which of the alternative matches to take is up to the oracle.

# What to change in (n, no choice)?

---

```
solve(goal):
```

```
    match goal against a head C.H of a clause C
```

```
    if multiple matches exist: ask the oracle to pick one
```

```
    if no matching head found:
```

```
        return FAILURE
```

```
    if C has no rhs: // C is a fact
```

```
        return SUCCESS
```

```
    else // reduce goal to the goals in the rhs of C
```

```
        for each goal in C.rhs
```

```
            if solve(goal) == FAILURE
```

```
                // oracle failed to find a solution for goal
```

```
                return FAILURE
```

```
        end for
```

```
        // goals on the rhs were solved successfully
```

```
        return SUCCESS
```

# Algorithm (1, backtracking)

<b>Choice of clause</b>	backtracking	$a(X) :- b(X).$ $b(Y) :- c(Y).$ $c(1).$ $c(2).$	
	by oracle	New concept: <b>search tree</b> Implementation: <b>ask oracle for the right choice.</b>	as below, with oracular choice
	not needed	New concepts: <b>unifier, proof tree</b> Implementation: <b>reduce a goal and recurse</b>	Concept: <b>resolvent</b> Implementation: <b>recursion deals with reduced goals; iteration deals with rhs goals</b>
		1	$n$
		<b>number of clauses on the <i>rhs</i> of rules</b>	

# Implementing the oracle

---

We can no longer ask the oracle which of the (potentially multiple) matching heads to choose.

We need to iterate over these matches, testing whether one of them solves the goal. If we fail, we return to try the next match. This is backtracking.

Effectively, backtracking implements the oracle.

The backtracking process corresponds to dfs traversal over the search tree. See *The Art of Prolog*.



# Algorithm for (1,backtracking)

---

```
solve(goal):  
  for each match of goal with a head C.H of a clause C  
    // this match is found with unify(), of course  
    current_goal = C.rhs  
    res = solve(current_goal)  
    if res == SUCCESS:  
      return res  
  end for  
  return FAILURE
```

Again, this algorithm ignores how mgus are handled.  
This is up to you to figure out.

# Example

---

```
a(X) :- b(X).  
b(Y) :- c(Y).  
b(3).  
c(1). c(2).  
?- a(Z)
```

When it reaches `c(1)`, the interpreter call stack is:

*bottom*

solve `a(Z)`: matched the single `a(X)` head

solve `b(Z)`: matched head `b(Y)`; head `b(3)` still to explore

solve `c(Z)`: matched head `c(1)`; head `c(2)` still to explore

# The implementation structure

---

Recursion solves the new subgoal.

For loop iteration iterates over alternative clauses.

Backtracking is achieved by returning to higher level of recursion and taking the next iteration of the loop.

# Example executed on PA3 Prolog

---

a(X) :- b(X).

b(Y) :- c(Y).

c(1).

c(2).

a(I)?

Asking for solution 1

Goal: a(I)

Unify: a(X<sub>1</sub>) and a(I)

Unifier: {X<sub>1</sub>->I }

Goal: b(I)

Unify: a(X<sub>2</sub>) and b(I)

Unifier: null

Unify: b(Y<sub>3</sub>) and b(I)

Unifier: {Y<sub>3</sub>->I }

Goal: c(I)

Unify: a(X<sub>4</sub>) and c(I)

Unifier: null

Unify: b(Y<sub>5</sub>) and c(I)

Unifier: null

Unify: c(1) and c(I)

Unifier: {I->1 }

I = 1

Asking for solution 2

Unify: c(2) and c(I)

Unifier: {I->2 }

I = 2

Asking for solution 3

Unify: c(1) and b(I)

Unifier: null

Unify: c(2) and b(I)

Unifier: null

Unify: b(Y<sub>10</sub>) and a(I)

Unifier: null

Unify: c(1) and a(I)

Unifier: null

Unify: c(2) and a(I)

Unifier: null

None

# Algorithm (n, backtracking)

<b>Choice of clause</b>	backtracking	<p><u>Concept</u>: backtracking is dfs of search tree.</p> <p><u>Implementation</u>: b/tracking remembers remaining choices in a for loop on the call stack.</p>	<p>a(X) :- b(X), c(X).                  b(2).                  c(1).                  c(2).</p>
	by oracle	<p>New concept: <b>search tree</b></p> <p>Implementation: <b>ask oracle for the right choice.</b></p>	<p>as below, with oracular choice</p>
	not needed	<p>New concepts: <b>unifier, proof tree</b></p> <p>Implementation: <b>reduce a goal and recurse</b></p>	<p>Concept: <b>resolvent</b></p> <p>Implementation: <b>recursion deals with reduced goals; iteration deals with rhs goals</b></p>
		1	<i>n</i>
		<b>number of clauses on the <i>rhs</i> of rules</b>	

# Algorithm (n,backtracking) is the key task in PA3

---

You will design and implement this algorithm in PA3

here, we provide useful hints

Key challenge: having to deal with a resolver

we no longer have a single pending subgoal

This will require a different backtracking algo design

one that is easier to implement with coroutines

We will show you an outline of algo (2, backtracking)

you will generalize it to (n,backtracking)

# Example

---

This example demonstrates the difficulty

$a(X) \text{ :- } b(X,Y), c(Y).$

$b(1,1).$

$b(2,2).$

$c(2).$

The subgoal  $b(X,Y)$  has two solutions.

Only the second one will make  $c(Y)$  succeed.

We need a way to backtrack to the “solver” of  $b(X,Y)$

and ask it for the next solution

# Algorithm (2, backtracking)

---

Restriction: we have exactly two goals on the rhs

call them `rhs[0]` and `rhs[1]`

`solutions(goal)` returns a solution iterator

it uses `yield` to provide the next solution to `goal`

(2,backtracking):

```
for sol0 in solutions(rhs[0])
```

```
    for sol1 in solutions(rhs[1])
```

```
        if sol0 and sol1 “work together”: return SUCCESS
```

```
return FAILURE
```

Again, we are abstracting the propagation of `mgus`

as a result, we need to use the informal term “goals work together”;

it means that, after `mgus` found in `sol0`, there exists a valid `sol1`.



## Algorithm (2,backtracking), cont.

---

solve() must be adapted to work as a coroutine.

Key step: replace return with yield.

```
solve(goal):  
    for each match of goal with a head C.H of a clause C  
        current_goal = C.rhs  
        res = solve(current_goal)  
        if res == SUCCESS:  
            yield res return res  
    return FAILURE    // think whether this needs to be yield, too
```

# The complete view of control transfer

iterate over alternative rules

recursive function conjunction()

head(A,B)....






head(X,Y) :- a(X), b(X,Z), c(Y, Z, D)

head(X,Y)....

coroutine process()

coroutine process()

coroutine process()



# Example executed on PA3 Prolog

---

a(X) :- b(X), C(X).

b(2).

c(1).

c(2).

a(1)?

Asking for solution 1

Goal: a(1)

Unify: a(X\_1) and a(1)

Unifier: {X\_1->1 }

Goal: b(1)

Unify: a(X\_2) and b(1)

Unifier: null

Unify: b(2) and b(1)

Unifier: {1->2 }

Goal: c(2)

Unify: a(X\_4) and c(2)

Unifier: null

Unify: b(2) and c(2)

Unifier: null

Unify: c(1) and c(2)

Unifier: null

Unify: c(2) and c(2)

Unifier: {}

1 = 2

Asking for solution 2

Unify: c(1) and b(1)

Unifier: null

Unify: c(2) and b(1)

Unifier: null

Unify: b(2) and a(1)

Unifier: null

Unify: c(1) and a(1)

Unifier: null

Unify: c(2) and a(1)

Unifier: null

None

# Algorithm ( $n$ , backtracking)

<b>Choice of clause</b> backtracking by oracle not needed	<p><u>Concept</u>: backtracking is dfs of search tree.</p> <p><u>Implementation</u>: b/tracking remembers remaining choices on the call stack.</p>	<p>You will design and implement this algorithm in PA3</p>
	<p>New concept: <b>search tree</b></p> <p>Implementation: <b>ask oracle for the right choice.</b></p>	<p>as below, with oracular choice</p>
	<p>New concepts: <b>unifier, proof tree</b></p> <p>Implementation: <b>reduce a goal and recurse</b></p>	<p>Concept: <b>resolvent</b></p> <p>Implementation: <b>recursion deals with reduced goals; iteration deals with rhs goals</b></p>
	1	$n$
	<b>number of clauses on the <i>rhs</i> of rules</b>	

# Reading

---

## Required

The Art of Prolog, Chapters 4, 6, and search trees in Ch 5.  
(on reserve in Kresge and in Google Books.)

## Recommended

HW2: backtracking with coroutines (the regex problem)

## Insightful

Logic programming via streams in CS61A textbook (SICP).