
Logic Programming: Prolog

Table of Contents

Prolog

Key points

- **closed world assumption:** anything you haven't said to be true is assumed to be false
- **negation as failure:** to execute `\+G`, Prolog first tries to prove `G`. If it fails, `\+G` succeeds. Otherwise, it fails.
 - failing goals does not bind variables, so `\+G` cannot solve for variables
 - ensure all variables in a negated goal are bound before the goal is executed
- **terms:** all data structures are terms. They can be
 - **atomic:**
 - * integers, floating point numbers
 - * **atoms:** begins with lower case letter
 - **compound:** functor/function symbol followed by 0+ arguments
 - **variable:** denotes a single unknown term, begins with upper case letter
 - * single `_` specifies a different variable each time it appears
- **Datalog:** fragment of Prolog with only atomic terms and variables
- **single-assignment language:** a variable can only be bound once
- **ground term:** contains no variables
 - has only one instance
- **nonground term:** contains 1+ variable
 - has an infinite number of instances
- **substitution:** mapping from variables to terms
 - never replaces atomic/compound terms, only replaces variables
- **applying a substitution:** consistently replacing occurrences of each variable in the map with the term it is mapped to
- do arithmetic with the `is/2` predicate, infix e.g. `X is 6*7`
 - `is/2` only works if the 2nd argument is ground

Proper List

- either empty `[]` or not `[X|Y]`
- if not empty, the tail `Y` must be a proper list

```
1 proper_list([]).
2 proper_list([Head|Tail]) :- proper_list(Tail).
```

- this is a common pattern for recursive list traversal in Prolog

Append

```
1 append([], C, C).
2 append([A|B], C, [A|BC]) :- append(B, C, BC).
```

Member

```
1 member1(Elt, List) :- append(_, [Elt|_], List).
2
3 member2(Elt, [_|_]).
4 member2(Elt, [_|Rest]) :- member2(Elt, Rest).
```

- `member2` is more efficient: `member1` builds and ignores the list of elements before `Elt` in `List`. The second does not.

Logic and Resolution

Interpretations

- **interpretation:**
 - atomic terms stand for entities in the **domain of discourse (universe)**
 - each functor (function symbol of arity $n > 0$ stands for a function from n entities to one entity in the domain
 - each predicate of arity n stands for a particular relationship between n entities in the domain of discourse

Views of predicates

- predicate with n arguments can be viewed in different ways:

-
- function from all possible combinations of n terms to a truth value
 - set of tuples of n terms. Every tuple in the set is implicitly mapped to true, while every other tuple is mapped to false
 - predicate definition is then, for each view
 - define the mapping
 - define the set of tuples

Meaning of clauses

```
1 grandparent(A, C) :- parent(A,B), parent(B,C).
```

means “for all terms that A and C may stand for: A is a grandparent of C if there is a term B such that A is a parent of B, and B is a parent of C”

$$\forall A, C (\text{grandparent}(A, C) \leftarrow \exists B (\text{parent}(A, B) \wedge \text{parent}(B, C))).$$

- variables in the head are universally quantified over the entire clause
- variables appearing only in the body are existentially quantified over the body

Meaning of predicate definitions

- predicates are define by a finite number of clauses, each of which is in the form of an implication
- e.g. `parent(queen_elizabeth, prince_charles)` represents

$$\forall A, B (\text{parent}(A, B) \leftarrow (A = \text{queen_elizabeth} \wedge B = \text{prince_charles}))$$

- the meaning of the predicate is a disjunction of the bodies of all the clauses:

$$\forall A, B : \text{parent}(A, B) \leftarrow (A = \text{queen_elizabeth} \wedge B = \text{prince_charles}) \vee (A = \text{prince_philip} \wedge B = \text{prince_cha})$$

Closed world assumption

- to implement closed world assumption, make the implication biimplication

$$\forall A, B : \text{parent}(A, B) \iff (A = \text{queen_elizabeth} \wedge B = \text{prince_charles}) \vee (A = \text{prince_philip} \wedge B = \text{prince_ch})$$

- A is not a parent of B unless they are one of the listed cases

- adding reverse implication produces the **Clark completion** of the program

Semantics

- logic program P consists of a set of predicate definitions
- **semantics/meaning of P** : set of **logical consequences** as ground atomic formulae
- a ground atomic formula a is a logical consequence of a program P if P makes it true
- a negated ground atomic formula $\neg a$ ($\setminus +a$) is a logical consequence of P if a is not a logical consequence of P

Finding semantics

- you can determine the semantics of a logic program by working backwards: instead of reasoning from a query to find a satisfying substitution, you reason from the program to find what ground queries succeed.
- immediate consequence operator T_P takes a set of ground unit clauses C and produces the set of ground unit clauses implied by C together with the program P
- always includes all ground instances of all unit clauses in P
- for each clause $H \text{ :- } G_1, \dots, G_n$ in P

e.g.

```

1 % P
2 q(X,Z) :- p(X,Y), p(Y,Z).
3 % C
4 p(a,b).
5 p(b,c).
6 p(c,d).

```

Then:

$$T_P(C) = \{q(a,c).q(b,d).\}$$

- the semantics of P is always T_P applied infinitely many times to the empty set: $T_P(T_P(\dots())\dots)$

Procedural interpretation

```

1 grandparent(A, C) :- parent(A,B), parent(B,C).

```

Logical interpretation:

$$\forall A, C(\text{grandparent}(A, C) \leftarrow \exists B(\text{parent}(A, B) \wedge \text{parent}(B, C))).$$

Procedural interpretation: to show that A is a grandparent of C, it suffices to show A is a parent of B and B is a parent of C.

Selective Linear Definite Resolution

- consequences of a logic program are determined through **resolution**
- **SLD resolution** is an efficient version of resolution
- e.g. to determine if Queen Elizabeth is Prince Harry's grandparent:

```
1 ?- grandparent(queen_elizabeth, prince_harry).
```

With program:

```
1 grandparent(X,Z) :- parent(X,Y), parent(Y,Z).
```

- Unify query goal (`grandparent(queen_elizabeth, prince_harry)`) with clause head `grandparent(X,Z)`
- apply resulting substitution to the clause, to yield the **resolvent**
- the goal is identical to the resolvent head, so we can replace it with the resolvent body:

```
1 ?- parent(queen_elizabeth, Y), parent(Y, prince_harry).
```

- now pick one of these goals to resolve: say you choose `parent(Y, prince_harry)`
- only two clauses can resolve with it:

```
1 parent(prince_charles, prince_harry).
2 parent(princess_diana, prince_harry).
```

- choose the second clause. After resolution, we are left with the query:

```
1 ?- parent(queen_elizabeth, princess_diana).
```

- no clause unifies with this query: resolution fails. This can sometimes take many steps.
- now **backtrack** and try the first matching clause

```
1 ?- parent(queen_elizabeth, prince_charles).
```

- there is a matching program clause, so there is nothing more to prove. The query succeeds.

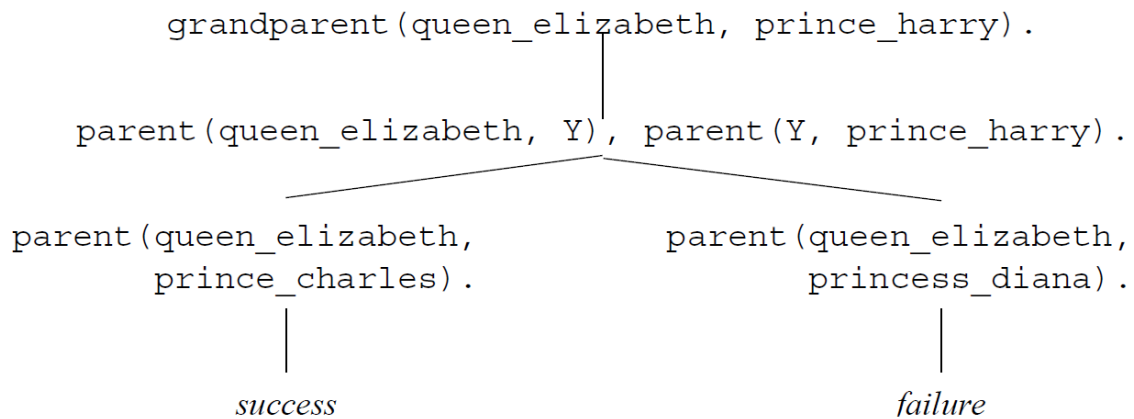


Figure 1: SLD Tree

Order of Execution

- order of goal resolution and trial of different clauses:
 - doesn't matter for correctness
 - does matter for efficiency
- Prolog always selects the first goal to resolve, and always selects the first matching clause to pursue first
 - allows programmer to control execution

Backtracking

- when there are multiple clauses matching a goal, Prolog leaves a **choicepoint** so that it can return to that state and try the next matching clause
- when a goal fail, Prolog **backtracks** to the most recent choicepoint
 - removes all variable bindings since the choicepoint
 - Prolog begins resolution with next matching clause
 - once all matching clauses are exhausted, the choicepoint is removed
 - subsequent failures then backtrack to the next choicepoint

Indexing

- **indexing** can improve efficiency

- Prolog systems automatically create an index for predicates with multiple clauses, where the heads have distinct constants/functors
- for a call with the first argument called, Prolog immediately jumps to the first matching clause
- SWI Prolog constructs indices for multiple arguments, meaning more queries benefit from indexing

Debugging

Prolog Debugger

- `trace` turns on the debugger
- `nodebug` turns off the debugger
- Byrd box model: goal execution is a box with ports for entry/exit




Figure 2: Prolog Debugger

Infinite Backtracking Loop

Initial version of reverse

```

1 % rev1(X,Y)
2 % rev1/2 holds when Y is the reverse of the list X
3 rev1([], []).
4 rev1([A|BC], CBA) :-
5     rev1(BC, CB),
6     append(CB, [A], CBA).
  
```

- Doesn't work if the first argument is free, e.g. `rev1(X, [a])`.
- Prolog enters an infinite backtracking loop:
 - `rev1(BC, CB)` has an infinite backtracking sequence of solutions $\{BC \rightarrow [Z], CB \rightarrow [Z]\}$, $\{BC \rightarrow [Y, Z], CB \rightarrow [Z, Y]\}$
 - `append([Z], [A], [a])` fails

-
- `append([Y,Z], [A], [a])` fails
 - ...
 - you could prevent this by swapping the body goals around, but then it won't work with the second argument free
 - solution: ensure that `rev1`'s first argument is always bound to a list when called
 - length of a list must always be the same as that of its reverse
 - when `same_length/2` succeeds, both arguments are bound to lists of the same fixed length

```
1 % rev3(X,Y)
2 % rev3/2 holds when Y is the reverse of the list X
3 rev3(ABC, CBA) :-
4     same_length(ABC, CBA),
5     rev1(ABC, CBA).
6
7 same_length([], []).
8 same_length(_|Xs, _|Ys) :-
9     same_length(Xs, Ys).
```

Managing nondeterminism

- when clauses succeed, but there are later clauses that may succeed, Prolog leaves a choicepoint so that it can later backtrack and try the later clause
- when efficiency matters: ensure recursive predicates don't leave choicepoints when they should be deterministic (i.e. leave no choice points) when there are no other solutions
- if choicepoints remain, it disables tail recursion optimisation

If-then-else (`->` ;).

- can be used to avoid choicepoints
- `->` is treated like conjunction, but any alternative solutions of the condition, and any alternatives of the else block will be forgotten.
- if the condition goal fails, the else goal is tried
- deterministic whenever the then/else blocks are deterministic
- indexing is preferable to ITE: avoid where possible
 - ITE often prevents code working in multiple modes

Tail Recursion

- **tail recursive:** the only recursive call is the last code executed before returning to the caller

Tail Recursion Optimisation (TRO)

- Prolog performs **tail recursion optimisation**, making recursive predicates behave as if they were loops
- more often applicable in Prolog than other languages

The Stack

- **stack frame:** stores local variables and where to return to when finished
- when **a** calls **b**, it creates a fresh stack frame for **b**, preserving **a**'s frame
- similarly when **b** calls **c**
- if all **b** does after calling **c** is return to **a**, there is no need to preserve **b**'s local variables




Figure 3: Stack frames

- **last call optimisation:** saves significant stack space
 - Prolog can release **b**'s frame before calling **c**
 - then when **c** finishes, it directly returns to **a**

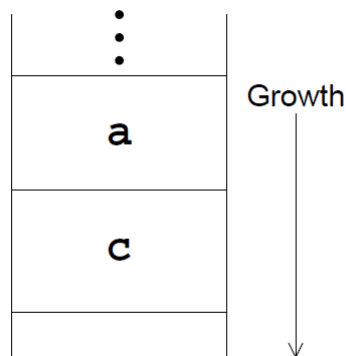


Figure 4: Last call optimisation

- TRO is a special case of last call optimisation, where the last call is recursive
- particularly beneficial, as recursion replaces looping
- without TRO: a new stack frame would be needed for each iteration, which would quickly exhaust the stack
- with TRO: tail recursive predicates execute in **constant** stack space, just like a loop

Choicepoints

- if **b** leaves a choicepoint, it sits on the stack above **b**'s frame
- this freezes it, and all earlier frames, meaning they can't be reclaimed
- this is necessary: when Prolog backtracks to the choicepoint, **b**'s arguments must be ready to try the next matching clause for **b**

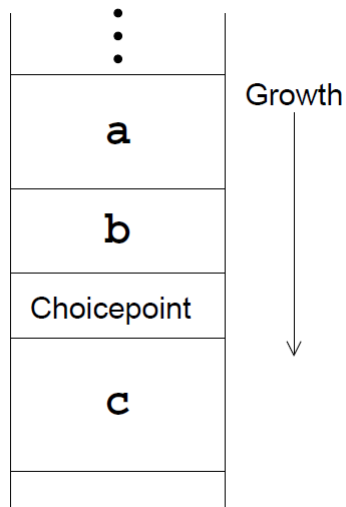


Figure 5: TRO and choicepoints

Accumulator

- make code tail recursive through an **accumulating parameter/accumulator**, an extra parameter to the predicate that holds a partially computed result
- base case: (usually) partially computed result is the actual result
- recursive clause: compute more of the partially computed result, and pass this to the recursive goal
- helpful approach: consider how you would implement it using a **while** loop, then implement it in Prolog

Accumulating Lists

- accumulators can make an order difference to efficiency
 - e.g. replacing `append/3` (linear time) with list construction (constant time)
- e.g. `rev1` defined earlier is $O(n^2)$:
 - for the n -th element from the end of we append a list of length $n - 1$ to a singleton list
 - doing this n times gives $\sim \frac{n(n-1)}{2}$

Tail Recursive `rev/2`

```
1 % rev(BCD, A, DCBA)
2 % DCBA is BCD reversed, with A appended
3 rev([], A, A).
4 rev([B|CD], A, DCBA) :-
5   rev(CD, [B|A], DCBA).
```

- at each step, an element is removed from the head of the input list, and added to the head of the accumulator
- the cost of each step is constant, so overall cost is linear in length of the list
- accumulator here works like a stack, last element is the first element added to it

Difference Pairs

- the tail recursive `rev` is a common pattern in Prolog
- a predicate that generates a list takes an extra argument specifying what should come after the list, avoiding the need to append to the list
- if you don't know what will come after the list at the time you call the predicate, you can pass an unbound variable
- then bind the variable when you do know what should come after
- many predicates intended to produce a list have 2 arguments
 - 1st argument: list produced
 - 2nd argument: what comes after
- **difference pair**: predicate that generates the difference between the first and second list

```
1 % inorder tree flatten
2 flatten(empty, List, List).
3 flatten(node(L,E,R), List, List0) :-
4   flatten(L, List, List1),
5   List1 = [E|List2],
6   flatten(R, List2, List0).
```

Predicates

- logic programming language based on predicate calculus
- build on **predicates** which define **relations** among their arguments
- e.g. relationships: parent/child
- predicates can be defined by 1+ **clauses**

-
- **fact/unit clause** *classes.pl*:

```
1 % the students we know about
2 student(alice).
3 student(bob).
4 student(claire).
5 student(don).
6 % who is enrolled in which subjects
7 enrolled(alice, logic).
8 enrolled(alice, maths).
9 enrolled(bob, maths).
10 enrolled(claire, physics).
11 enrolled(don, logic).
12 enrolled(don, art_history).
```

You can then load this file, and make queries:

```
1 ?- [classes].
2 true.
3 ?- student(bob).
4 true.
5 ?- student(sally).
6 false.
```

Variables

- **variables** in Prolog can only hold one value each time it exists, and refers to the same value each place it appears in that scope. Think of it as standing in for a value we don't yet know
- must begin with capital letter or underscore, containing only letters, digits, underscores If you pose queries with variables in them, Prolog looks for bindings that satisfy the query:

```
1 ?- student(X).
2 X = alice ;
3 X = bob ;
4 X = claire ;
5 X = don.
```

The semicolon is input by the user to move to the next possible binding. Enter accepts a binding.

```
1 ?- enrolled(alice, Subject).
2 Subject = logic ;
3 Subject = maths.
```

Special variable `_` matches anything, and each place you write it, it names a different variable.

```
1 % is alice enrolled in any subject?
2 ?- enrolled(claire, _).
```

```
3 true.
4 % is anyone enrolled in any subject?
5 ?- enrolled(_, _).
6 true ;
7 true ;
8 ...
9 true.
```

Compound queries

- queries can involve conjunctions (**AND**), disjunctions (**OR**), and negations (**NOT**)
- conjunction operator: `,`
- disjunction operator: `;`
- negation operator: `\+`

```
1 % who is taking both maths and logic?
2 ?- enrolled(S, maths), enrolled(S, logic).
3 S = alice ;
4 false.
5 % who is enrolled in either maths or logic?
6 ?- enrolled(S, maths) ; enrolled(S, logic).
7 S = alice ;
8 S = don ;
9 S = alice ;
10 S = bob.
11 % who is enrolled in logic but not maths?
12 ?- enrolled(S, logic), \+ enrolled(S, maths).
13 S = don.
```

Rules

Facts are clauses specifying that a relationship holds. **Rules** are clauses that specifies that a relationship holds under certain conditions.

```
1 head :- body
```

The rule specifies that `head` holds if `body` holds

```
1 % general syntax: the rule specifies that the relationship holds if
2 head :- body
3 % two people are classmates if they are enrolled in the same class
4 classmates(X, Y) :- enrolled(X, Class), enrolled(Y, Class)
```

Equality

```
1 % this shows bob is his own classmate
2 ?- classmates(bob, X).
3 X = alice ;
4 X = bob.
5 % we can use negation and equality to rectify:
6 ?- classmates(X, Y) :-
7     enrolled(X, Class),
8     enrolled(Y, Class),
9     \+ X = Y.
10 ?- classmates(bob, X).
11 X = alice ;
12 false.
```

Disequality and Negation as Failure

- `\=`: not equal predicate.. `X \= Y` behaves the same as `\+ X = Y`
- **negation as failure**: Prolog negates a query by attempting to find a solution: if it succeeds, the negation fails. If it fails, the negation succeeds. It doesn't bind any variables, so negations should be written *following* goals that do bind variables used in the negation.

```
1 % this doesn't work properly:
2 ?- X \= Y, X = bob, Y = alice.
3 false.
4 ?- X = bob, Y = alice, X \= Y.
5 X = bob,
6 Y = alice.
```

Terms

- Prolog is **dynamically typed**. All data are called **terms**.
- **atomic terms**: primitive types. Integers, floating point numbers, and atoms.
- atoms can begin with a lowercase letter and follow with letters, digits, or underscores, otherwise it begins/ends with a single quote `'` and can contain any characters
- the Prolog compiler will not identify type errors in the code
- any argument of any predicate you define can have any type

Compound Terms

- compound term: Prolog equivalent to C `struct`. Begins with a *functor* (an atom) and follows with 1+ terms as arguments

-
- e.g. compound term with functor `card`, arity 2, first argument is `clubs`, second argument is 3.

```
1 card(clubs, 3)
```

Lists

- `[]` : empty list
- `[E|Es]` : non empty list, `E`: head, `Es`: tail
- e.g. `[E1,E2,E3|Es]`

Unification

- variables in Prolog are a kind of data that stands for a currently unknown value, and continue to exist after the predicate that creates them finishes executing.
- variables become bound through **unification**, which takes two terms and tries to make them identical, binding variables as necessary
- if a set of *consistent* bindings cannot be found for all variables, unification fails
- unification happens at every predicate call: the call is unified with the head of the first clause for the predicate: if it succeeds, Prolog executes the body of the clause; if the unification fails, Prolog goes to the next clause for the predicate and tries the same thing
- the equality predicate `=` also unifies its two arguments

length

Here's an implementation of Haskell's `take` that returns the first `N` elements of a list

```
1 take(N, List, Front) :-
2     length(Front, N),
3     append(Front, _, List).
```

`Front` is the first `N` elements of `List` if the length of `Front` is `N` and you can append `Front` to something to produce `List`.

member

- `member(E, List)` holds when `E` is one of the elements of `List`
- use this to check whether `E` is an element of `List`, or to have Prolog produce elements of `List` one at a time

select

- `select(Elem, List1, List2)`: `List2` contains everything in `List1` except `Elem`
 - can use to remove single occurrence of `Elem` from `List1`, or insert `Elem` in any place in `List2`, or to select an element of `List1`, producing 1 element plus the rest of the list

nth0/3

- `nth0(Index, List, Elem)`: finds the n th element of `List` (0-based).
- can determine position of `Elem` in `List`
- can produce elements of `List` together with their positions.

nth0/4

- `nth0(N, List, Elem, Rest)`: same as `nth0/3`, but `Rest` is the list of elements other than `Elem`
- use it to remove an element from a list by position, by value while providing position, or to insert an element at a particular position

Documenting Modes

- document each Prolog predicate in a comment before the predicate definition
- give each argument a character indicating its mode
- `+`: input argument. Expected to be **bound** when the predicate is called
- `-`: output argument. Normally **unbound** when the predicate is called. If it is bound, it will be unified with the output
- `?`: the predicate may be input/output/both
- e.g. `append/3`

```
1 % append(+List1, +List2, -List3)
2 % append(-List1, -List2, +List3)
3 % List3 is a list of all the elements of List1 in order followed
4 % by the all the elements of List2 in order.
```

- documentation should indicate all intended modes of use, and be clear when it doesn't work

Arithmetic

- `is/2` is used to evaluate expressions:

```
1 ?- X is 6*7.  
2 X=42.
```

- the 2nd argument must be a ground term!

```
1 ?- X is 1*A.  
2 ERROR: Arguments are not sufficiently instantiated.  
3 ERROR: In:  
4 ...
```

- Prolog is not a symbolic computation system, with very limited ability to reason about arithmetic

Arithmetic Predicates

Predicates	Description
<code>V is Expr</code>	Unify V with the value of expression Expr
<code>Expr1 == Expr2</code>	Succeeds if Expr1 and Expr2 are equal
<code>Expr1 \= Expr2</code>	Succeeds if Expr1 and Expr2 are different
<code>Expr1 < Expr2</code>	Succeeds if Expr1 is strictly less than the value of Expr2
<code>Expr1 <= Expr2</code>	Succeeds if Expr1 is less or equal to the value of Expr2
<code>Expr1 > Expr2</code>	Succeeds if Expr1 is strictly greater than the value of Expr2
<code>Expr1 >= Expr2</code>	Succeeds if Expr1 is greater or equal to the value of Expr2

Arithmetic Expressions

Function	Description
<code>-X</code>	unary negation (integer or float)
<code>X + Y</code>	addition (integer or float)
<code>X - Y</code>	subtraction (integer or float)

Function	Description
$X * Y$	multiplication (integer or float)
X / Y	division, producing integer or float
$X // Y$	integer division, rounding toward zero
$X \text{ rem } Y$	integer remainder, same sign as X
$X \text{ div } Y$	integer division, rounding down
$X \text{ mod } Y$	integer modulus, same sign as Y
<code>integer(X)</code>	round X to the nearest integer
<code>float(X)</code>	floating point value of X
<code>ceil(X)</code>	smallest integer $\geq X$
<code>floor(X)</code>	largest integer $\leq X$
<code>max(X, Y)</code>	larger of X and Y (integer or float)
<code>min(X, Y)</code>	smaller of X and Y (integer or float)

Semantics

- **semantics** of a logic program: what does it make true?
- i.e. a program consisting of a set of ground facts
- to determine the semantics of a program containing rules, you start with an empty set of clauses, and then copy all facts into the semantics
- then for each `Head :- Body`, unify all goals in `Body` with every combination of facts in the semantics. Add each combination instance of `Head` to the semantics
- when this process reaches a **fixed point**, where no new clauses are added to the semantics, you are done.

Tail recursion

- make recursive calls operate more like iterative calls, using constant stack space instead of linear stack space
- this works when the last call in the predicate is recursive
- you can typically do this by adding an **accumulator** argument to the predicate that stores an intermediate result

e.g. we want to make `sumList2` tail recursive

```
1 sumlist([], 0).
2 sumlist([E|Es], N) :-
3     sumlist(Es, N1),
4     N is N1 + E.
```

As a for loop, this would be:

```
1 N = 0
2 for elem in list:
3     N = N + elem
4
5 return N
```

To translate this to Prolog, create a new predicate `sumlist/3` with an accumulator

```
1 % make sumlist/2 call sumlist/3 with initialised accumulator
2 sumlist(List, Sum) :-
3     sumlist(List, 0, Sum).
4
5 sumlist([], Sum, Sum).
6 sumlist([E|Es], Sum0, Sum) :-
7     Sum1 is Sum0 + E,
8     sumlist(Es, Sum1, Sum).
```

Determinism

Indexing

When you call a predicate with some arguments bound, Prolog looks at all the clauses for the predicate to see if some have non-variables in that position of the clause head. If so, it constructs an index on that argument position. Then Prolog can jump straight to the first clause that matches the query. If it is forced to backtrack, it will jump to the next clause that could match. When it knows there are no more clauses that could match, it removes the choicepoint.

if -> then ; else

- Prolog goal of the form $(p \rightarrow q ; r)$: first calls p . If that succeeds then call q and ignore r . Otherwise ignore q and call r
- used to produce determinism
- note that if p has more than one solution, Prolog commits to the first solution and throws away the others

-
- good practice to write the ; at the start of the line to distinguish it from the ,

```
1 ints_between(N0, N, List) :-  
2   ( N0 < N  
3   -> List = [N0|List1],  
4       N1 is N0 + 1,  
5       ints_between(N1, N, List1)  
6   ; N0 = N,  
7     List = [N]  
8   ).
```