Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2 / MoSIG M1   Second Semester 2010/2011

Lesson 8   24 February 2011

Control of Reasoning and Decision Trees

Production System Architecture

The system implements an "inference engine" that operates as a 3-phase cycle:

The cycle is called the "recognize act" cycle.
The phases are:

MATCH: match facts in Short Term memory to rules

SELECT: Select the correspondence of facts and rules to execute

EXECUTE: Execute the action par of the rule.

```
(defrule <rule-name> [<comment>]
     [<declaration>]            ; Rule Properties
     <conditional-element>*     ; Left-Hand Side (LHS)
=>
     <action>*)                 ; Right-Hand Side (RHS)
```

The Agenda

The agenda is a list of activations of rules. It provides the rule name, index of the fact that matches each CE, and variable bindings.   There are a number of different control regimes.

Control regimes following different principles.
A fundamental principle is "Refraction" .

Refraction: A unique activation is only executed once.
         Activation is removed from the agenda on execution.

By default, the agenda acts as a stack (LIFO).
Other general principles include:

recency: Recent activations are given priority
MEA: A variation of recency where the index of the fact matching the FIRST CE determines the priority of the activation.
specificity: Rules with more CEs are given priority.

For example:

```
   (defrule example
      (item ?x ?y ?x)
      (test (and (numberp ?x) (> ?x (+ 10 ?y)) (< ?x 100)))
      =>)
```

has specificity 2

OPS 5 provided 2 control regimes: LEX and MEA.
CLIPS has 7.

1) "Depth Strategy" - Agenda acts as a list of stacks (LIFO).
         There is a separate stack for each salience.

2) "Breadth Strategy" - Agenda acts as a list queue (FIFO) with a separate stack for each salience

3) LEX strategy (Lexographic). (Compatibility with OPS). The agenda is a list of sorted activations. Activations for each saliency are sorted by Recency and then by specificity.

4) MEA strategy (Means-Ends-Analysis). (Compatibility with OPS). Activations for each saliency are sorted based on Fact-Index of the FIRST CE, then by specificity.

5) Complexity Strategy: Rules are sorted by Specificity with most complex rules given priority.

6) Simplicity: Rules are sorted by specificity with simplest rules given priority.

7) Random: using a random seed determined at start of execution.

Depth strategy is recommended (Agenda acts as a stack).

Examples of Strategies

```
 (set-strategy depth)
 (get-strategy)

(defrule rule-A
     ?f <- (a)
=>
     (printout t "Rule A fires with  " ?f  crlf)
)

(defrule rule-B
     ?f <- (b)
=>
     (printout t "Rule B fires with  " ?f crlf)
)

(defrule rule-A-and-B
     ?f1 <- (a)
     ?f2 <- (b)
=>
     (printout t "Rule B fires with  A =" ?f1 " and B = " ?f2  crlf)
)

(assert (a))
(assert (a))
(assert (b))

(set-strategy depth)
(set-strategy breadth)
(set-strategy lex)
(set-strategy mea)
(set-strategy complexity)
(set-strategy simplicity)
(set-strategy random)
(set-strategy depth)
```

Salience

The salience property for a rule determines its priority.
Salient rules are given higher priority in the agenda.

Salience is "declared" in the [<declaration>] part of the LHS, before the CE's

```
(defrule <rule-name> [<comment>]
      [<declaration>]                 ; Rule Properties
      <conditional-element>*          ; Left-Hand Side (LHS)
=>
      <action>*)                      ; Right-Hand Side (RHS)
```

(declare (salience S)) where   -10 000 ≤ S ≤ 10 000
by default S is 0.

example:

```
(defrule example
        (declare (salience 999))
        (initial-fact)
    =>
        (printout "I am an important rule! Salience= 999" crlf)
)
```

There is a tendency for beginners to abuse salience in order to force the order of rule execution. Don't! Rules should be structured with contexts.
If the system is well constructed, rule execution order is not important and only a few saliencies are needed.

A well-constructed program should need only 3 or 4 salience. At most 7 may be needed.

<u>Salience Hierarchy:</u>

Different styles of programs can require different hierarchies of salience.
A good practice is to declare the hierarchy in advance, using multiples of 100.
An example is the following:

<u>Level</u>	<u>Salience</u>	
Constraints	300	;; Rules that eliminate hypotheses
Expertise	200	;; Domain knowledge
Query	100	;; Rules that interrogate the user
Control	0	;; Context transitions

Using Context to Structure Rules

It is possible to organize an expert system as a finite state machine, where each state corresponds to a rule "context" (also called a phase). Phases (or contexts) can be organized into cycles, networks or trees. The transition between phases can be coded reactively (as rules) or declaratively (as facts). Declarative control is slightly slower but provides the advantages of being easy to inspect, easy to change, and easier to debug. Examples include:

Diagnostic Systems (e.g. MYCIN) - tree of phases (or contexts)
Systems for self-monitoring and self repair – cycles

Within each phase, a set of rules encode the systems knowledge.

In general it is good practice to group rules into contexts (or phases). Each context (or phase) concerns some sort of calculation coded as a body of rules that may fire in any order.

Phases and Control Elements

Phases are indicated by the existence of a token: an element in the facts list that indicates the current phase.

    (phase <Name of the phase>)

A classic example is a self-monitoring and self-repair system used for satellites and space applications. Such systems typically operate in cycle with 3 phases:

Fault-Detection: A set of rules that test the integrity of subsystems
Fault-Diagnosis: A set of rules that determine the origin of an error.
Fault-Repair: A set of rules that reconfigure the system to repair a fault.

| Fault Detection | → | Fault Diagnosis | → | Fault Repair |

One can imagine many ways to code control.

For example,

1) It is possible to code the fault condition in each rule. However this is expensive and can allow multiple interacting faults to interfere with each other.

2) It is possible to use salience to define a hierarchy for the rules. This is a very bad idea that leads to complex un-maintainable code.

    example :   Fault Detection :   Salience 3
                              Fault Diagnosis :   Salience 2
                              Fault Repair : Salience 1.

3) Use a "phase" element to mark each phase and then use a rule to manage the transitions . This is the preferred solution.

The transition rules encode the control of the system. These can be reactive or declarative. For example:

```
(defrule detection-to-diagnosis
    (declare (salience 10))
    ?phase <- (phase detection)
    (fault ?f detected)
=>
    (retract ?phase)
    (assert (phase diagnosis))
    (printout t "Fault " ?f " detected!" crlf)
)

(defrule diagnosis-to-repair
    (declare (salience 10))
    ?phase <- (phase diagnosis)
    (fault ?f detected)
    (fault ?f diagnosis ?d)
=>
    (retract ?phase)
    (assert (phase repair))
    (printout t " Fault " ?f " diagnosis " ?d crlf)
)

(defrule repair-to-detection
    (declare (salience 10))
```

```
    ?phase <- (phase repair)
    (fault ?f diagnosis ?d repair ?r)
=>
    (retract ?phase)
    (repair ?f ?d ?r)
    (assert (phase detection))
    (printout t "Fault repaired" crlf)
)
```

Within each phase, are a collection of rules for domain knowledge that diagnosis and suggest a repair for the fault.

Declarative Control Structures

An alternative to coding the phase transitions in explicit rules, is to encode the phase transitions in a declarative data structure and use a single generic transition rule.

```
(deffacts control-list
    (phase detection)
    (next-phase detection diagnosis)
    (next-phase diagnosis repair)
    (next-phase repair detection)
)

(defrule phase transition rule.
    (declare (salience -10))
    ?P <- (phase ?phase)
    (next-phase? phase ?next)
=>
    (retract ?P)
    (assert (phase ?next))
)
```

Decision Trees

A decision trees can be used for diagnosis and classification problems.
They require that :

    The set of answers be finite and known in advance
    The space of problems can be reduced to a series of yes/no tests

A decision tree is composed of decision nodes and leaves.
The decision nodes compose the tree. The answers are found in the leaves.

Example: Learn to guess the animal

The following example is a system that learns to guess an animal by asking questions.

The decision tree is composed of decision-nodes and answer-nodes.


```
(deffacts tree
(node  root decision "Is it warm blooded?" n1 n2)
(node  n1 decision "Does it purr ?" n3 n4)
(node  n2 answer snake)
(node  n3 answer cat)
(node  n4 answer dog)
)

;; initialization rule

(defrule init
    (initial-fact)
=>
 (assert (current-node root))
)

(defrule init
=>
```

```
  (assert (current-node root))
)
```

;; rule to request a decision node

```
(defrule make-decision
    ?N <- (current-node    ?name)
    (node   ?name decision ?q ?yes ?no)
=>
    (retract ?N)
    (format t "%s (yes or no) " ?q)
    (bind ?answer (read))
    (if (eq ?answer yes)
       then (assert (current-node ?yes) )
       else (assert (current-node ?no))
    )
)
```

;; rule to give answers

```
(defrule give-answer
    ?N <- (current-node    ?name)
    (node   ?name answer ?r)
=>
    (printout t "I guess that it is a " ?r crlf)
    (printout t "Am I right? (yes or no) ")
    (bind ?rep (read))
    (if (eq ?rep yes)
       then (assert (phase play-again))
           (retract ?N)
       else (assert (phase correct-answer))
    )
)
```

;; rule to play again

```
(defrule play-again
    ?phase <- (phase play-again)
=>
    (retract ?phase)
    (printout t "play again? (yes or no) ")
    (bind ?rep (read))
    (if (eq ?rep yes)
       then (assert (current-node   root))
       else (save-facts "animal.dat")
    )
)
```

```
;; rule to learn the correct answer

(defrule learn-correct-answer
    ?P <- (phase correct-answer)
    ?N <- (current-node    ?name)
    ?D <- (node  ?name answer ?r)
=>
    (retract ?P ?N ?D)   ;;     Ask the correct answer
    (printout t "What animal was it? ")
    (bind ?new (read))
    (printout t  "What question should I ask to tell a "
        ?new " from a " ?r "? ")
    (bind ?question (readline))
    (bind ?newnode1 (gensym*))
    (bind ?newnode2 (gensym*))
    (assert (node  ?newnode1 answer ?new))
    (assert (node  ?newnode2 answer ?r))
    (assert
    (node ?name decision ?question ?newnode1 ?newnode2))
    (assert (phase play-again))
)



;; Rule to open the file animal.dat

(defrule init
    (initial-fact)
=>
 (assert (file (open "animal.dat" data "r")))
)
```

```
;; rule to close the file animal.dat

(defrule no-file
   ?f <- (file FALSE)
=>
   (retract ?f)
   (assert (current-node root))
)

;; rule to read the file.

(defrule init-file
   ?f <- (file TRUE)
=>
 (bind ?in (readline data))
 (printout t ?in crlf)
 (if (eq ?in EOF) then (assert (eof))
  else
    (assert-string ?in)
    (retract ?f)
    (assert (file TRUE))
   )
)

(defrule eof
   (declare (salience 30))
   ?f <- (file TRUE)
   ?eof <- (eof)
=>
   (retract ?f ?eof)
   (close data)
   (assert (current-node root))
)
```