ELEMENTS OF LATENT LEARNING IN A MAZE ENVIRONMENT

Dale M. McNulty
and
Richard H. Granger

Irvine Computational Intelligence Project
and
Center for the Neurobiology of Learning and Memory
University of California
Irvine, CA 92717
CSNet Address: MCNULTY@UCI
714-856-6360
January 3, 1985

85-03

Topic: Cognitive Modelling

Length: 4456 words

Keywords: Latent learning, machine learning, conflict resolution, conceptual clustering.

ABSTRACT

A general purpose learning program is described which demonstrates a latent learning ability by operating at two separate goal pursuit levels. At one level are the constant, implicit goals associated with the system's memory management mechanisms. At the higher level are the dynamic, explicit behavioral goals which the implicit goals enable by manipulating memory representations to conform to the external surroundings. The program is shown to negotiate a simulated maze environment by the step-wise refinement of its latently learned experiences.

This research was supported in part by the National Science Foundation under grant IST-81-20685 and the Naval Ocean Systems Center under contracts N00123-81-C-1078 and N66001-83-C-0255.

1.0 The Latent Learning Problem

Learning has been characterized as a goal-directed, state-space search problem [Newell, 1980; Simon, 1973, 1983]. However, rats allowed to wander in mazes unrewarded will later, under reward conditions, demonstrate the same maze knowledge as control groups which are immediately rewarded [Blodgett 1929; Tolman and Honzik 1930]. Clearly something is being learned in the trials which lack goal reward. The learning is taking place prior to any measurement of the learning and is, thus, described by psychologists as latent.

CEL-0 (pronounced like cello) is a learning system capable of demonstrating a latent learning ability comparable to a rat in a maze because the model distinguishes between the explicit (higher) goal level of food reward and the implicit (lower) goal level where goals remain constant from one problem domain to another. At the explicit level, CEL-0's goals are dynamic to reflect the different situations CEL-0 might encounter. At the lower level, CEL-0's operators encompass a constant set of goals that manipulate memory to convert it into an accurate predictor of the external environment. Many of the implicit goal based mechanisms operate independently of the explicit goal level.

This paper describes the CEL-0 program and presents a simple example which demonstrates CEL-0's maze performance. During the goal pursuit phase of the experiment CEL-0 learns a direct solution path which did not exist in memory. CEL-0 does this by reconstructively piecing together independently clustered pieces of previously learned episodes. Some results and a brief explanation of a more complex latent learning example are presented.

2.0 Introduction to CEL and the CEL-0 Program

The CEL (Components of Experiential Learning) theory [Granger, 1982, 1983a, 1983b] is a process model which identifies low level mechanisms of learning and memory conforming to results from psychology and neurophysiology. The model, building on previous machine learning concepts, is composed of twelve parallel and semi-independent primitives which build and manipulate hierarchical memories of episodic schemata from experiential events. CEL-0 is a program written to demonstrate various aspects of the CEL theory.

CEL-0's operators constitute five classes of memory manipulation which can be summarized as follows:

Reception Operators
DETECT - sensory events from the various modalities (sight, touch, etc.).
SELECT - ively attend to sensory inputs on the basis of past experience.
NOTICE - input which matches the desirable and undesirable states

and trigger COLLECT.
COLLECT - sensory events into packets (episodic schemata) for memory storage.
INDEX - schemata into the memory hierarchy for future reference.

Retrieval Operators
REMIND - the system of past experiences which are related to the present situation.
ACTIVATE - REMINDed schemata based on predictive values of schemata and behavior desired. ACTIVATE triggers reconstruction.

Reconstruction Operators
ENACT - the appropriate efferent actions in the active schema; tune SELECT to attend to predicted afferent events.
SYNTHESIZE - new schemata by matching inputs against predicted events; trigger refinement operators to modify schema based on matches or mismatches.

Refinement Operators
REINFORCE - strengthen successful schemata to reflect their predictive success.
BRANCH - within a schemata to indicate mismatches
DETOUR - within a schemata to indicate branches to be avoided.

3.0 Latent Learning in the Maze Environment

CEL-0 operates in a maze environment simulated interactively by MAZEWORLD. MAZEWORLD interprets CEL-0's actions and provides CEL-0 with the sensory events to reflect those actions. Watching over CEL-0's behavior, SUPERVISOR starts and removes CEL-0 from the environment depending on the design of the experiment. Just as in a 'real' maze experiment neither MAZEWORLD nor SUPERVISOR guide or direct CEL-0 in any way during maze negotiation.

Section 6.0 of this paper contains sample output from CEL-0's experience with a simple four component T maze.

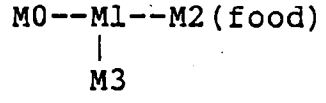

In a latent learning paradigm (henceforth referred to as complete latent learning to distinguish it from the simple example CEL-0 uses here) the food would not be included until CEL-0 experienced the environment a number of times. At that point the reward would be introduced and the progress of an unsatiated CEL-0 would be noted while it learned, trial by trial, to consistently move directly to the food. In this simplified example, food has been introduced immediately. This allows CEL-0 to learn the direct path to the food in two trials, one satiated and one unsatiated. The principles of CEL-0's operation are the same in both experiments, but the simpler of the two is covered here to conserve space. In Section 5.0 some results of CEL-0 in the complete example are discussed.

3.1 Exploration in the Maze

In the first trial, SUPERVISOR starts a satiated CEL-0 at location M0. CEL-0 explores the maze (order: M0-M1-M3-M1-M2), discovers the reward nine moves later (at M2), but does not eat since it is not hungry. Exploration appears to be random because, to the observer, no explicit goal seems to exist. However, CEL-0 has implicit goals determined by the inherent mechanisms of memory management. CEL-0 moves from spot to spot by matching previously stored or innate representations with currently experienced sensory input and deciding which representations (schemata) are appropriate predictions to pursue. The implicit goals create mental representations which serve as additional predictors of CEL-0's environment during the second trial when it has an explicit goal.

SUPERVISOR starts CEL-0 facing the north wall of M0. CEL-0 is REMINDed, by the opening to its right, that it can turn right and move forward. ENACTing the efferent acts of Episodic Schema 4 (EP4) moves CEL-0 to M1. The successful completion of EP4 indicates that it is a useful predictor, so the strength of EP4 is REINFORCEd and the most recent experiences are COLLECTed and INDEXed (learned).

At M1, CEL-0 is REMINDed of four schemata; EP5 (turn right and move forword) was just strengthened, so it gets selected over the others as a more useful predictor.

CEL-0 explores the entire maze in this manner. At first it relies a great deal on the innate, move-from-spot-to-spot schemata. As it gains experience, sensory events at each location REMIND it of previous experiences and CEL-0 chooses to pursue one of those predictions. If a prediction proves to be accurate, CEL-0 marks that success and builds new schemata to incorporate the specific sensory experiences which lead to that success.

When CEL-0 arrives at M2 the sight of food is NOTICEd. A NOTICE takes place when events on the Desirable State List (DSL) match sensory input. The resultant COLLECTion includes all of the events beginning with CEL-0's view from M3 and ending with the sight of food. Sensitivity Analysis (discussed in Section 4.3) creates a number of schemata. All of these schemata end with the goal 'view food'. Some of these will have important application in the next trial.

The SUPERVISOR then notices that CEL-0 has perused the entire maze and removes it as an experimenter would remove a subject when it completes a maze experiment. SUPERVISOR then places CEL-0 at M0 for another trial.

3.2 Explicit Goal Pursuit in the Maze

As CEL-0 begins the second trial, two things have changed. First, CEL-0 has all of the memories it accumulated during exploration. Second, CEL-0 is not satiated; it has the food drive.

CEL-0 starts as before, but as it begins to turn to its right (using EP5) it is REMINDed of a previous time when, after moving from M3 to M1, it turned right and sighted food. CEL-0 ACTIVATEs this schema, matches the first event (a right turn), and notices that EP63 no longer matches the environment. Simultaneously, the sight of M1 forward REMINDs CEL-0 of its last time through the maze when it saw M1 from M0 and eventually found food (EP57). ACTIVATing EP57 will lead CEL-0 to the food, but it includes an excursion to M3. Of course, CEL-0 doesn't know that the side trip is useless, it's only predicting that it will work. EP57 is ACTIVATEd and the unmatched portion of EP63 is deactivated.

When CEL-0 arrives at M1 the sight of M2 ahead and M3 to the right causes a REMIND of EP65. CEL-0 also remembers the deactive EP63. CEL-0 now has three accurate schemata to select from. The deactive EP63 is equivalent to EP65 because the first event in EP63 has been matched off. EP63 had some predictive success before it was deactivated, so EP65 is eliminated. Of the two remaining schemata, EP63 is shorter so CEL-0 ACTIVATES the remainder of EP63 and, as a result, finds its way directly to the food (order: M0-M1-M2), skipping the useless excursion it made to M3 in the exploratory phase.

Upon NOTICing the food at M2, CEL-0 is REMINDed of EP13 (the eating schema). The result (goal) of EP13 is eating. This is on the DSL, so ACTIVATE will give this schema a very high priority if it matches the current sensory experiences. It does match, so EP13 is ACTIVATEd, and CEL-0 satisfies its hunger. NOTICE matches the act of eating with its spot on the DSL and triggers a COLLECT and INDEX. This is a very important COLLECTion because at this point CEL-0's recent history includes the events which led directly to eating the food. The next time CEL-0 begins the maze it does not have to piece together a route from parts of schemata. It has learned an episode that leads effectively and efficiently from the start of the maze to goal satisfaction.

When CEL-0 eats, SUPERVISOR considers the trial completed, and removes CEL-0 from the maze.

4.0 CEL-0 Attributes Which Enable Latent Learning

CEL-0 is modelled to be a domain independent learning machine, and as such, finds it useful to employ most of its mechanisms in every domain. There are, however, four particular aspects of CEL-0 which enable it to demonstrate this latent learning ability.

4.1 COLLECTion Points, or When To Learn

CEL-0's COLLECT operator groups events into bundles, called episodic schemata, and hands them over to INDEX to create addressable memory units. The points where COLLECTion occur are critical because only the seven most recent sensory experiences are available for COLLECTion. Each new action and environment dictate a new set of events available for COLLECTion. During the exploratory phase, the COLLECTion points determine the contents of the schemata which will be available during the second trial. The COLLECTion which occurs while exploring M2 is the reason CEL-0 can learn to find the food in two trials. Lacking this COLLECTion, CEL-0 would take longer in the pursuit phase to piece together, from other schemata, the actions necessary to find the reward.

4.2 CEL-0's INDEX Operator

Part of CEL-0's implicit goal based repertoire is the INDEXing which creates clusters of schemata. The INDEX operator creates Experiential and Goal Indices from the first and last (respectively) event of every COLLECTed sequence. This episode is grouped, by index, with existing schemata such that episodes of similar characteristics will be REMINDed simultaneously. This is similar to conceptual clustering as reported in the machine learning literature [Dietterich and Michalski, 1983; Michalski and Stepp, 1983]. CEL-0's clustering differs from the clustering techniques reported by Michalski and Stepp [1983] in that a simple categorization (first and last event of each episode) of episodic data is performed.

CEL-0's indices perform two functions. Experiential Indices cause REMINDs of all those events associated with particular sensory experiences. Goal Indices allow REMINDs of every schema associated with goals. The latter grouping allows CEL-0 to search for schemata that possibly satisfy its goals, and the former offers up possible paths to those goals.

The REMIND of EP63 in Subsection 44 demonstrates the usefulness of the Experiential Index. EP63 is one of many episodes clustered about the 'right turn' index. A right turn will REMIND CEL-0 of all of these schemata and ACTIVATE can decide if any are useful.

4.3 Activation of Schemata

As CEL-0 gains experience, the numbers of clusters and the numbers associated with each cluster increase. Since REMINDs are constantly occurring in response to sensory input, CEL-0 must have some way to choose among the potentially large number of possibilities. ACTIVATE must decide when a REMIND should be ACTIVATEd and which REMIND is the best to ACTIVATE. In the later respect, ACTIVATE is similar to machine learning's conflict

resolution [McDermott and Forgy, 1978]. ACTIVATE performs this selection process by applying a set of evaluation metrics to the group of REMINDs. The application of these metrics in the maze can be seen in Section 3.

4.4 Sensitivity Analysis

The newly INDEXed schemata contain many more potentially useful predictions than the single sequence of events which get COLLECTed. When CEL-0 COLLECTs all of the most recent experiences it doesn't know which events actually contributed to the result which caused the COLLECTion. Of the seven events COLLECTed, only the last three or four may have contributed to seeing the food (e.g. EP57 and EP65). If that were the case the first three events are extraneous. CEL-0 tests for this sensitivity to event inclusion by performing Sensitivity Analysis when the schema is created. Instead of creating the single episode of experience, CEL-0 forms multiple schemata by repeatedly removing the initial event from the first schema and INDEXing the resultant episode. Sensitivity Analysis creates many possibly predictive schemata where only one was COLLECTed.

CEL-0 uses this to advantage in situations where a COLLECT didn't occur at just the right time. EP63 is an example since it is a subset formed by performing Sensitivity Analysis on the COLLECTion triggered by the sighting of food during exploration. The full set of COLLECTed experiences (EP67) includes the excursion to M3. When CEL-0 gets to M1, ACTIVATing EP63 enables it to bypass that wasted effort.

5.0 Summary and Future Directions

As a learning system, CEL-0 is driven by goals at two levels. At the lower level, CEL-0's operators are driven by a set of constant, implicit goals to adjust its internal representation to be an accurate predictor of the external surroundings. At the explicit level, the goals are dynamic, reflecting CEL-0's changing desires, such as satisfying hunger and thirst. CEL-0 demonstrates a latent learning ability because many of the mechanisms operate independently of the explicit goals. For example, as CEL-0 explores the maze unrewarded, COLLECT is triggered at points unrelated to explicit drives, and Sensitivity Analysis, without regard for any other operators, is creating predictive possibilities for use during explicit goal acquisition.

CEL-0's success might seem to be an artifact of the specific maze configuration or the early introduction of reward. While it is true that these affect the observed behavior, CEL-0's ability to eventually learn a direct route to the food does not seem to be affected by either. This is because CEL-0 utilizes a step-wise refinement process to construct new schemata from pieces of existing episodes. The process works as follows.

Sensory experiences REMIND CEL-0 of previously learned episodes. CEL-0 ACTIVATEs one of those REMINDs and ENACTs the efferent events of those schemata as long as the learned episodes match the external environment. When they no longer match, CEL-0 chooses a new schema to ACTIVATE from the new set of REMINDs. Eventually CEL-0 will discover the goal. The sequence of events leading up to goal acquisition will form new clustered schemata which can in turn be refined by the same process. CEL-0 benefits from exploration because it has a chance to construct a rich set of clustered memories to step-wise refine during explicit goal pursuit.

Preliminary work with the complete latent learning problem indicate that CEL-0 requires 3 to 6 rewarded trials to become consistent in its performance. This seems to agree favorably with data attributed to rats in similar experiments [Bower and Hilgard, 1981, page 338], but more work is needed to determine the equivalence of the two paradigms.

Despite the apparent 'sense an event, respond' action which CEL-0 is constantly performing, CEL-0 is far from a stimulus-response (S-R) model. In the first place, CEL-0's actions build crucial memory structures which in turn contribute to CEL-0's later behavior. Second, CEL-0 will not necessarily respond to the same sensory event in the same way each time. CEL-0's behavior will depend on various internal states and conditions which change dynamically. Third, and most important, CEL-0 demonstrates a latent learning ability which is not generally thought of as a stimulus-response behavior.

5.1 Future Work

Given a complex theory like CEL, it is difficult to single out any one part for attention, but the future direction is to address two broad areas (only efforts directly related to the maze paradigm are discussed here since work on the CEL theory is taking many forms):

1. Investigation of useful conflict resolution strategies in the form of ACTIVATE metrics and INDEXing (clustering) criteria. The larger the maze the greater the number of sensory experiences and, thus, the greater the number of stored schemata. Although not an exponential problem, ACTIVATE does have the increased burden of selecting from the many more REMINDs which occur in the later stages of maze negotiation. It should be noted, however, that the entire burden for selecting schemata does not fall on ACTIVATE. The selection process is a tradeoff between the metrics ACTIVATE employs and the hierarchical clustering which REMIND works with. At the moment, efforts are being focused on the ACTIVATE metrics in order to identify the results which the selection process must provide. Some selection criteria may then be placed at the front end of the memory process. That is, effective clustering in the memory hierarchy could effectively restrict

the number of possibly useful REMINDs that ACTIVATE must select from.

2. Comparison of CEL-0's behavior in the complete latent learning paradigm with observed rat performance. This is holding with the basic CEL philosophy of attempting to make the model conform to experimental results wherever possible and reasonable.

6.0 Sample CEL-0 Output

The following is an abbreviated listing of CEL-0's experience in the simple four component T maze. Some portions have been summarized for brevity. Actor's responses have been delineated and numbered (e.g. Subsection 1,2,34,etc.) for clarity.

```
[DSKLOG started: 1-3-85 9:07 AM]
*(MAZE)
Initializing CEL-0's memory
Indexing innate schemata: (EP1 EP10 EP11 EP12 EP13 EP14 EP2 EP3 EP4 EP5 EP6 EP7 E
8 EP9)
Set food satiated
**** SUPERVISOR (Subsection 1) ****
Begin exploration phase (i.e. CEL-0 is not hungry)
Place CEL-0 at M0
****  MAZEWORLD (Subsection 2) ****
CEL-0 at M0 facing NORTH
CEL-0's view is: ((FORWARD (WALL)) (RIGHT (M1)) (BACK (WALL)) (LEFT (WALL)))
****    CEL-0 (Subsection 3)    ****
DETECTed: (VIEW (FORWARD (WALL))
                (RIGHT (M1))
                (BACK (WALL))
                (LEFT (WALL)))
REMINDed of: (EP5)
Choosing (EP5); most predictive of group
ACTIVATing EP5
Enter reconstruction mode using: (VIEW (FORWARD (*VAR* W))
                                       (RIGHT (*OPENING* X))
                                       (BACK (*VAR* Y))
                                       (LEFT (*VAR* Z)))
                                 (TURN (DIRECTION (RIGHT))
                                       (ACTOR (ME)))
                                 (MOVE (DIRECTION (FORWARD))
                                       (ACTOR (ME)))
SYNTHESIZE match on: (VIEW (FORWARD (WALL))
                           (RIGHT (M1))
                           (BACK (WALL))
                           (LEFT (WALL)))
ENACTing: (TURN (DIRECTION (RIGHT))
                (ACTOR (ME)))
****  MAZEWORLD (Subsection 4) ****
CEL-0 at M0 facing NORTH
CEL-0's view is: ((FORWARD (WALL)) (RIGHT (M1)) (BACK (WALL)) (LEFT (WALL)))
CEL-0's action number 1 is (TURN (DIRECTION (RIGHT))
                                 (ACTOR (ME)))
****    CEL-0 (Subsection 5)    ****
DETECTed: (TURN (DIRECTION (RIGHT))
                (ACTOR (ME)))
No REMINDs on this event
A schema is active, REMINDed of: EP5
Choosing (EP5); most predictive of group
Continue SYNTHESIS with EP5 since already active
SYNTHESIZE match on: (TURN (DIRECTION (RIGHT))
```

```
                              (ACTOR (ME)))
****  MAZEWORLD (Subsection 6) ****
CEL-0 at M0 facing EAST
CEL-0's view is: ((FORWARD (M1)) (RIGHT (WALL)) (BACK (WALL)) (LEFT (WALL)))
****    CEL-0 (Subsection 7)    ****
DETECTed: (VIEW (FORWARD (M1))
                (RIGHT (WALL))
                (BACK (WALL))
                (LEFT (WALL)))
REMINDed of: (EP3)
A schema is active, REMINDed of: EP5
Choosing (EP5 EP3); most predictive of group
Continue SYNTHESIS with EP5 since already active
ENACTing: (MOVE (DIRECTION (FORWARD))
                (ACTOR (ME)))
****  MAZEWORLD (Subsection 8) ****
CEL-0 at M0 facing EAST
CEL-0's view is: ((FORWARD (M1)) (RIGHT (WALL)) (BACK (WALL)) (LEFT (WALL)))
CEL-0's action number 2 is (MOVE (DIRECTION (FORWARD))
                                 (ACTOR (ME)))
****    CEL-0 (Subsection 9)    ****
DETECTed: (MOVE (DIRECTION (FORWARD))
                (ACTOR (ME)))
REMINDed of: (EP4 EP6 EP8 EP10 EP12)
A schema is active, REMINDed of: EP5
     Removing (EP6 EP8 EP10 EP12); not viable
Choosing (EP5 EP4); most predictive of group
Continue SYNTHESIS with EP5 since already active
SYNTHESIZE match on: (MOVE (DIRECTION (FORWARD))
                           (ACTOR (ME)))
Successful match on EP5: (VIEW (FORWARD (*VAR* W))
                               (RIGHT (*OPENING* X))
                               (BACK (*VAR* Y))
                               (LEFT (*VAR* Z)))
                         (TURN (DIRECTION (RIGHT))
                               (ACTOR (ME)))
                         (MOVE (DIRECTION (FORWARD))
                               (ACTOR (ME)))
Inexact match - REINFORCE by 1
COLLECT triggered
Indexing: (EP15 EP16 EP17 EP18 EP19 EP20)
Terminate reconstruction of EP5
    .
    . CEL-0 continues in this fashion employing REMINDed schemata to
    . explore the entire maze (order: M0-M1-M3-M1-M2).
    .
****  MAZEWORLD (Subsection 38) ****
CEL-0 at M2 facing EAST
CEL-0's view is: ((FORWARD (OBJECT (FOOD))) (RIGHT (WALL)) (BACK (M1))
(LEFT (WALL)))
****    CEL-0 (Subsection 39)    ****
DETECTed: (VIEW (FORWARD (OBJECT (FOOD)))
                (RIGHT (WALL))
                (BACK (M1))
                (LEFT (WALL)))
```

```
NOTICE triggered on the last event
COLLECT triggered
 Indexing: (EP57 EP58 EP59 EP60 EP61 EP62 EP63 EP64 EP65 EP66 EP67 EP68)
REMINDed of: (EP9 EP11 EP13 EP58 EP60 EP62 EP64 EP66 EP68)
     Removing (EP58 EP62 EP66); not viable
     EP13 deleted as an unrequired goal: (TASTE (OBJECT (FOOD))
                                               (ACTOR (ME)))
     (EP60 EP64 EP68) deleted as related but unrequired
 Choosing (EP9); most predictive of group
 ACTIVATing EP9
Enter reconstruction mode using: (VIEW (FORWARD (*VAR* W))
                                       (RIGHT (*VAR* X))
                                       (BACK (*OPENING* Y))
                                       (LEFT (*VAR* Z)))
                                  (TURN (DIRECTION (LEFT))
                                       (ACTOR (ME)))
                                  (TURN (DIRECTION (LEFT))
                                       (ACTOR (ME)))
                                  (MOVE (DIRECTION (FORWARD))
                                       (ACTOR (ME)))
SYNTHESIZE match on: (VIEW (FORWARD (OBJECT (FOOD)))
                           (RIGHT (WALL))
                           (BACK (M1))
                           (LEFT (WALL)))
ENACTing: (TURN (DIRECTION (LEFT))
                (ACTOR (ME)))
 **** SUPERVISOR (Subsection 40) ****
Exploration completed, CEL-0 has seen the entire maze
 CEL-0 used 9 moves to explore the maze
 Now CEL-0 will do it again, with a goal (i.e. hungry)
Place CEL-0 at M0
****  MAZEWORLD (Subsection 41) ****
 CEL-0 at M0 facing NORTH
CEL-0's view is: ((FORWARD (WALL)) (RIGHT (M1)) (BACK (WALL)) (LEFT (WALL)))
****    CEL-0 (Subsection 42)   ****
 DETECTed: (VIEW (FORWARD (WALL))
                 (RIGHT (M1))
                 (BACK (WALL))
                 (LEFT (WALL)))
 REMINDed of: (EP5 EP15)
Choosing (EP5); most predictive of group
 ACTIVATing EP5
 Enter reconstruction mode using: (VIEW (FORWARD (*VAR* W))
                                        (RIGHT (*OPENING* X))
                                        (BACK (*VAR* Y))
                                        (LEFT (*VAR* Z)))
                                  (TURN (DIRECTION (RIGHT))
                                        (ACTOR (ME)))
                                  (MOVE (DIRECTION (FORWARD))
                                        (ACTOR (ME)))
YNTHESIZE match on: (VIEW (FORWARD (WALL))
                          (RIGHT (M1))
                          (BACK (WALL))
                          (LEFT (WALL)))
NACTing: (TURN (DIRECTION (RIGHT))
```

```
                (ACTOR (ME)))
****  MAZEWORLD (Subsection 43) ****
 CEL-0 at M0 facing NORTH
 CEL-0's view is: ((FORWARD (WALL)) (RIGHT (M1)) (BACK (WALL)) (LEFT (WALL)))
 CEL-0's action number 1 is (TURN (DIRECTION (RIGHT))
                                  (ACTOR (ME)))
****    CEL-0 (Subsection 44)    ****
 DETECTed: (TURN (DIRECTION (RIGHT))
                 (ACTOR (ME)))
 REMINDed of: (EP17 EP21 EP29 EP53 EP63)
 A schema is active, REMINDed of: EP5
      EP63 related to DSL item above threshold
 Choosing (EP63); goal related
 Choosing (EP63); most predictive of group
 ACTIVATing EP63
  Deactivate: EP5
 Enter reconstruction mode using: (TURN (DIRECTION (RIGHT))
                                        (ACTOR (ME)))
                                  (VIEW (FORWARD (M2))
                                        (RIGHT (M3))
                                        (BACK (M0))
                                        (LEFT (WALL)))
                                  (MOVE (DIRECTION (FORWARD))
                                        (ACTOR (ME)))
                                  (VIEW (FORWARD (OBJECT (FOOD)))
                                        (RIGHT (WALL))
                                        (BACK (M1))
                                        (LEFT (WALL)))
 SYNTHESIZE match on: (TURN (DIRECTION (RIGHT))
                            (ACTOR (ME)))
****  MAZEWORLD (Subsection 45) ****
 CEL-0 at M0 facing EAST
 CEL-0's view is: ((FORWARD (M1)) (RIGHT (WALL)) (BACK (WALL)) (LEFT (WALL)))
****    CEL-0 (Subsection 46)    ****
 DETECTed: (VIEW (FORWARD (M1))
                 (RIGHT (WALL))
                 (BACK (WALL))
                 (LEFT (WALL)))
 REMINDed of: (EP3 EP19 EP23 EP43 EP47 EP57)
 A schema is active, REMINDed of: EP63
     Removing EP63; not viable
 Choosing EP5 from deactivated schemata
      EP57 related to DSL item above threshold
 Choosing (EP57); goal related
 Choosing (EP57); most predictive of group
 ACTIVATing EP57
 Deactivate: EP63
 Enter reconstruction mode using: (VIEW (FORWARD (M1))
                                        (RIGHT (WALL))
                                        (BACK (WALL))
                                        (LEFT (WALL)))
                                  (MOVE (DIRECTION (FORWARD))
                                        (ACTOR (ME)))
                                  (VIEW (FORWARD (WALL))
                                        (RIGHT (M2))
```

```
                                        (BACK (M3))
                                        (LEFT (M0)))
                                  (TURN (DIRECTION (RIGHT))
                                        (ACTOR (ME)))
                                  (VIEW (FORWARD (M2))
                                        (RIGHT (M3))
                                        (BACK (M0))
                                        (LEFT (WALL)))
                                  (MOVE (DIRECTION (FORWARD))
                                        (ACTOR (ME)))
                                  (VIEW (FORWARD (OBJECT (FOOD)))
                                        (RIGHT (WALL))
                                        (BACK (M1))
                                        (LEFT (WALL)))
SYNTHESIZE match on: (VIEW (FORWARD (M1))
                           (RIGHT (WALL))
                           (BACK (WALL))
                           (LEFT (WALL)))
 ENACTing: (MOVE (DIRECTION (FORWARD))
                 (ACTOR (ME)))
****  MAZEWORLD (Subsection 49) ****
CEL-0 at M1 facing EAST
 CEL-0's view is: ((FORWARD (M2)) (RIGHT (M3)) (BACK (M0)) (LEFT (WALL)))
****    CEL-0 (Subsection 50)   ****
DETECTed: (VIEW (FORWARD (M2))
                (RIGHT (M3))
                (BACK (M0))
                (LEFT (WALL)))
 REMINDed of: (EP27 EP55 EP65 EP9 EP11 EP5 EP3)
 A schema is active, REMINDed of: EP57
     Removing EP57; not viable
Choosing EP63 from deactivated schemata
 Choosing EP5 from deactivated schemata
     (EP63 EP65) related to DSL item above threshold
Choosing (EP63 EP65); goal related
 Choosing (EP65 EP63); most predictive of group
ACTIVATing EP63
Deactivate: EP57
 Reactivating EP63 continue SYNTHESIS with (VIEW (FORWARD (M2))
                                                 (RIGHT (M3))
                                                 (BACK (M0))
                                                 (LEFT (WALL)))
                                           (MOVE (DIRECTION (FORWARD))
                                                 (ACTOR (ME)))
                                           (VIEW (FORWARD (OBJECT (FOOD)))
                                                 (RIGHT (WALL))
                                                 (BACK (M1))
                                                 (LEFT (WALL)))
SYNTHESIZE match on: (VIEW (FORWARD (M2))
                           (RIGHT (M3))
                           (BACK (M0))
                           (LEFT (WALL)))
ENACTing: (MOVE (DIRECTION (FORWARD))
                (ACTOR (ME)))
***  MAZEWORLD (Subsection 51) ****
```

```
CEL-0 at M1 facing EAST
CEL-0's view is: ((FORWARD (M2)) (RIGHT (M3)) (BACK (M0)) (LEFT (WALL)))
CEL-0's action number 3 is (MOVE (DIRECTION (FORWARD))
                                 (ACTOR (ME)))
****    CEL-0 (Subsection 52)    ****
DETECTed: (MOVE (DIRECTION (FORWARD))
                (ACTOR (ME)))
REMINDed of: (EP4 EP6 EP8 EP10 EP12 EP16 EP18 EP20 EP22 EP24 EP25 EP26 EP28 EP30
P32 EP33 EP34 EP36 EP38 EP40 EP42 EP44 EP46 EP48 EP49 EP50 EP52 EP54 EP56 EP59 EP
7)
A schema is active, REMINDed of: EP63
    Removing (EP8 EP16 EP20 EP24 EP32 EP36 EP40 EP44 EP48 EP52); not viable
Choosing EP5 from deactivated schemata
    (EP59 EP67 EP63) related to DSL item above threshold
Choosing (EP59 EP67 EP63); goal related
Choosing (EP63 EP67 EP59); most predictive of group
Continue SYNTHESIS with EP63 since already active
SYNTHESIZE match on: (MOVE (DIRECTION (FORWARD))
                           (ACTOR (ME)))
****  MAZEWORLD (Subsection 53) ****
CEL-0 at M2 facing EAST
CEL-0's view is: ((FORWARD (OBJECT (FOOD))) (RIGHT (WALL)) (BACK (M1))
(LEFT (WALL)))
****    CEL-0 (Subsection 54)    ****
DETECTed: (VIEW (FORWARD (OBJECT (FOOD)))
                (RIGHT (WALL))
                (BACK (M1))
                (LEFT (WALL)))
NOTICE triggered on the last event
COLLECT triggered
Indexing: (EP69 EP70 EP71 EP72 EP73 EP74 EP75 EP76 EP77 EP78 EP79 EP80)
REMINDed of: (EP9 EP11 EP13 EP58 EP60 EP62 EP64 EP66 EP68 EP70 EP72 EP74 EP76 EP7
 EP80)
A schema is active, REMINDed of: EP63
    Removing (EP78 EP74 EP70 EP66 EP62 EP58); not viable
Choosing EP5 from deactivated schemata
    EP13 matches DSL item above threshold
    (EP60 EP64 EP68 EP72 EP76 EP80 EP63) related to DSL item above threshold
Choosing (EP13); goal match
Choosing (EP13); most predictive of group
ACTIVATing EP13
Deactivate: EP63
Enter reconstruction mode using: (VIEW (FORWARD (OBJECT (FOOD)))
                                       (RIGHT (*VAR* X))
                                       (LEFT (*VAR* Y))
                                       (BACK (*VAR* Z)))
                                 (TASTE (OBJECT (FOOD))
                                        (ACTOR (ME)))
SYNTHESIZE match on: (VIEW (FORWARD (OBJECT (FOOD)))
                           (RIGHT (WALL))
                           (BACK (M1))
                           (LEFT (WALL)))
ENACTing: (TASTE (OBJECT (FOOD))
                 (ACTOR (ME)))
***  MAZEWORLD (Subsection 55) ****
```

```
CEL-0 at M2 facing EAST
CEL-0's view is: ((FORWARD (OBJECT (FOOD))) (RIGHT (WALL)) (BACK (M1))
 (LEFT (WALL)))
CEL-0's action number 4 is (TASTE (OBJECT (FOOD))
                                  (ACTOR (ME)))
****    CEL-0 (Subsection 56)    ****
DETECTed: (TASTE (OBJECT (FOOD))
                 (ACTOR (ME)))
NOTICE triggered on the last event
COLLECT triggered
Indexing: (EP81 EP82 EP83 EP84 EP85 EP86 EP87 EP88 EP89 EP90 EP91 EP92)
REMINDed of: (EP14 EP82 EP84 EP86 EP88 EP90 EP92)
A schema is active, REMINDed of: EP13
    Removing (EP88 EP84); not viable
Choosing EP63 from deactivated schemata
Choosing EP5 from deactivated schemata
    EP63 related to DSL item above threshold
    (EP14 EP82 EP86 EP90 EP92 EP13) matches DSL item above threshold
Choosing (EP14 EP82 EP86 EP90 EP92 EP13); goal match
Choosing (EP13 EP92 EP90 EP86 EP82 EP14); most predictive of group
Continue SYNTHESIS with EP13 since already active
SYNTHESIZE match on: (TASTE (OBJECT (FOOD))
                            (ACTOR (ME)))
Successful match on EP13: (VIEW (FORWARD (OBJECT (FOOD)))
                                (RIGHT (*VAR* X))
                                (LEFT (*VAR* Y))
                                (BACK (*VAR* Z)))
                          (TASTE (OBJECT (FOOD))
                                 (ACTOR (ME)))
Inexact match - REINFORCE by 1
COLLECT triggered
Indexing: (EP100 EP101 EP102 EP103 EP104 EP93 EP94 EP95 EP96 EP97 EP98 EP99)
Terminate reconstruction of EP13
**** SUPERVISOR (Subsection 57) ****
CEL has successfully negotiated the maze
It took CEL-0 4 moves to find the food
NIL
(LOGOUT)
[DSKLOG finished: 1-3-85 9:11 AM]
```

References

Blodgett, H.C., "The Effect of the Introduction of Reward Upon the Maze Performance of Rats", Univ. Calif. Psychol., 4:113-34.

Bower, G.H. and Hilgard E.R., Theories of Learning, Englewood Cliffs, Prentice-Hall, 1981.

Dietterich, T.G., and Michalski, R.S., "A Comparative Review of Selected Methods for Learning from Examples", in Machine Learning, an Artificial Intelligence Approach, Michalski, R.S., Carbonell, J.G., and Mitchell, T.M. (eds.), Tioga Publishing, Palo Alto, CA, pp. 331-364, 1983.

Granger, R.H., "Identification of Mental Operators Underlying the Acquisition of Simple Predictive Behavior," Department of Computer Science Technical Report 191, University of California, Irvine, November 1982.

Granger, R.H., "Identification of Components of Episodic Learning: The CEL Process Model of Early Learning and Memory," Journal of Cognition and Brain Theory, 6,1 (February 1983a), 5-38.

Granger, R.H., "An Artificial-Intelligence Model of Learning and Memory that Provides a Theoretical Framework for the Interpretation of Experimental Data in Psychology and Neurobiology", Department of Computer Science Technical Report 210, University of California, Irvine, September, 1983b.

Granger, R.H. and McNulty, D.M., "The CEL-0 system: Experience with a computer model that learns to run a maze", Department of Computer Science Technical Report 220, University of California, Irvine, March, 1984a.

Granger, R.H. and McNulty, D.M., "Learning and Memory in Machines and Animals: An AI Model That Accounts for Some Neurobiological Data", in Proceedings of the Sixth Annual Conference of the Cognitive Science Society, pp. 161, 1984b. (also in Memory , Experience, and Reasoning, Kolodner, J.L. and Riesbeck, C.K. (eds.), Erlbaum and Associates, Hillsdale, NJ, (to appear) 1985.)

McDermott, J. and Forgy, C., "Production System Conflict Resolution Strategies", in Pattern-Directed Inference Systems, Waterman, D.A. and Hayes-roth, F. (Eds.), Academic Press, New York, NY, pp. 177-199, 1978.

Michalski, R.S. and Stepp, R.E., "Learning From Observation: Conceptual Clustering", in Machine Learning, an Artificial Intelligence Approach, Michalski, R.S., Carbonell, J.G., and Mitchell, T.M. (eds.), Tioga Publishing, Palo Alto,

CA, pp. 331-364, 1983.

Newell, A., "Reasoning, Problem Solving, and Decision Processes: The Problem Space Hypothesis", in Attention and Performance, Nickerson, R. (ed.), Hillsdale, NJ, Erlbaum, 1980.

Simon, H., "Does Scientific Discovery Have a Logic?", Philosophy of Science, 1983, 4, 471-480.

Simon, H., "Why Should Machines Learn?", in Machine Learning, an Artificial Intelligence Approach, Michalski, R.S., Carbonell, J.G., and Mitchell, T.M. (eds.), Tioga Publishing, Palo Alto, CA, pp. 331-364, 1983.

Tolman, E.C. and Honzik, C.H., "Introduction and Removal of Reward, and Maze Performance in Rats", Univ. Calif. Publ. Psychol., 4:257-75.