

# *Exchanging Conflict Resolution in an Adaptable Implementation of ACT-R*

DANIEL GALL and THOM FRÜHWIRTH

*Faculty of Engineering and Computer Science, Ulm University, Germany*  
(e-mail: {daniel.gall,thom.fruehwirth}@uni-ulm.de)

*submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003*

---

## **Abstract**

In computational cognitive science, the cognitive architecture ACT-R is very popular. It describes a model of cognition that is amenable to computer implementation, paving the way for computational psychology. Its underlying psychological theory has been investigated in many psychological experiments, but ACT-R lacks a formal definition of its underlying concepts from a mathematical-computational point of view. Although the canonical implementation of ACT-R is now modularized, this production rule system is still hard to adapt and extend in central components like the conflict resolution mechanism (which decides which of the applicable rules to apply next).

In this work, we present a concise implementation of ACT-R based on Constraint Handling Rules which has been derived from a formalization in prior work. To show the adaptability of our approach, we implement several different conflict resolution mechanisms discussed in the ACT-R literature. This results in the first implementation of one such mechanism. For the other mechanisms, we empirically evaluate if our implementation matches the results of reference implementations of ACT-R.

**KEYWORDS:** computational cognitive modeling, computational psychology, ACT-R, Constraint Handling Rules, production rule systems, conflict resolution

---

## **1 Introduction**

Computational cognitive modeling is an approach in cognitive sciences which explores human cognition by implementing detailed computational models. This enables researchers to execute their models and simulate human behavior (Sun 2008). Due to their executability, computational models have to be defined precisely. Thereby ambiguities appearing in verbal-conceptual models can be eliminated. By conducting the same experiments with humans and an executable cognitive model, the plausibility of a model can be verified and gradually improved.

To implement cognitive models, it is helpful to introduce *cognitive architectures* which bundle well-investigated research results from several disciplines of psychology to a unified theory. On the basis of such an architecture, researchers are able to implement domain-specific computational models without having to deal with the

remodeling of fundamental psychological results. Additionally, cognitive architectures ideally constrain modeling to plausible models which facilitates the modeling process (Taatgen *et al.* 2006).

One of the most popular cognitive architectures is *Adaptive Control of Thought – Rational* (ACT-R), a production rule system introduced by John R. Anderson (Anderson and Lebiere 1998; Anderson *et al.* 2004). It has been used to model cognitive tasks like learning the past tense (Taatgen and Anderson 2002), but is also used in human-computer interaction or to improve educational software by simulating human students (Anderson *et al.* 2004, p. 1045 sqq.). Although providing a theory of the psychological foundations, ACT-R lacks a formal definition of its underlying concepts from a mathematical-computational point of view. This led to a reference implementation full of assumptions and technical artifacts beyond the theory making it difficult to overlook and inhibiting adaptability and extensibility. The situation improved with the modularization of the psychological theory, but it is still difficult to exchange more central parts of the implementation like conflict resolution (Stewart and West 2007).

To overcome these drawbacks, we have formalized parts of the implementation closing the gap between the psychological theory and the technical implementation. We describe an implementation of ACT-R which has been derived from our formalization using Constraint Handling Rules (CHR). Due to the power of logic programming, our implementation is very close to the formalization and leads to short and concise code covering the fundamental parts of the ACT-R theory. For the compilation of ACT-R models to CHR programs, source-to-source transformation is used. Our implementation is highly adaptable. In this paper, this is demonstrated by integrating four different conflict resolution strategies. Despite its proximity to the theory, the implementation can reproduce the results of the original implementation as exemplified in the evaluation of our work. The formalization may support the understanding of the details of our implementation, hence we refer to (Gall 2013) and the online appendix (Appendix A).

In section 2, we give an overview of the fundamental concepts of ACT-R and shortly describe their implementation in CHR. Section 3 describes the general conflict resolution process of ACT-R. Then the implementation of four different conflict resolution strategies proposed in the literature is presented. To evaluate our implementations, we use an example to compare the results of our implementation with those of the reference implementations where available in section 4. Eventually, in section 5 some related work is presented and a conclusion is given in section 6.

## 2 A CHR implementation of ACT-R

In the following, a short overview of the fundamental concepts of the ACT-R theory and their transfer to CHR is given. For reasons of space, we refer to the literature for an introduction to CHR (Frühwirth 2009). For a more detailed introduction to ACT-R, see (Anderson *et al.* 2004) and (Taatgen *et al.* 2006). The reference implementation of ACT-R is written in Lisp and can be obtained from the ACT-R website (ACT-R 2014). Details of our implementation including the formalization it

is based on can be found in (Gall 2013). Parts of the formalization are located in the online appendix (Appendix A).

## 2.1 Architecture

ACT-R is a production rule system which distinguishes two types of knowledge: *declarative knowledge* holding static facts and *procedural knowledge* representing processes controlling human cognition. For example, in a model of the game *rock, paper, scissors*, a declarative fact could be “The opponent played scissors”, whereas a procedural information could be that a round is won, if we played rock and the opponent played scissors. Declarative knowledge is represented as *chunks*. Each chunk consists of a symbolic name and labeled slots which hold symbolic values. The values can refer to other chunk names, i.e. chunks can be connected. Chunks are typed, i.e. the number and names of the slots provided by a chunk are determined by a type. As usual for production rule systems, procedural knowledge is represented as rules of the form IF *conditions* THEN *actions*. Conditions match values of chunks, actions modify them.

The psychological theory of ACT-R is modular: There are modules for each function of the human mind like a declarative module holding the declarative facts, a goal module taking track of the current goal of a task and buffering information and a procedural module holding the procedural information and controlling the cognitive process. There are also modules to interact with the environment like a visual module perceiving the visual field. The modules are independent from each other, i.e. there is no direct communication between them. Each module has a fixed number of *buffers* associated with it. The buffers can hold at most one single piece of information a time, i.e. one chunk. Modules can put chunks into their associated buffers.

The core of the system is the procedural module which can access the buffers of all other modules but does not have an own buffer. It consists of a *procedural memory* with a set of production rules. The conditions of a production rule refer to the contents of the buffers, i.e. they match the values of the chunk’s slots. The formal applicability condition of rules can be found in the online appendix (Appendix A).

There are three types of actions whose arguments are encoded as chunks as well: First of all, *buffer modifications* change the content of a buffer, i.e. the values of some of the slots of a chunk in a buffer. Secondly, the procedural module can state *requests* to other modules which then change the contents of their buffers. Eventually, *buffer clearings* remove the chunk from a buffer. Although our implementation can handle requests and clearings, we only regard buffer modifications in this work for the sake of simplicity.

### Example 1

Consider the following rule:

```
(p recognize-win
  =goal> isa game    me rock    opponent scissors
  ==>
  =goal> result win)
```

It recognizes a win situation in the game *rock, paper, scissors* if the model has realized that the opponent played scissors and the agent played rock (which could be accomplished by a corresponding production rule interacting with the visual module). The situation is represented by a chunk of type *game* providing the slots *me*, *opponent* and *result*. As a result, it adds the information that the round has been won by modifying the *result*-slot of the goal buffer.

Furthermore, the procedural module controls the *match-select-apply* cycle of the production rule system. It searches for matching rules. As soon as a matching rule has been selected to fire, it takes 50 ms for the rule to fire based on theories of human cognition (Anderson 2007, p. 54). During this time, the matching process is inhibited and no other rule can be selected until the selected rule is applied. Hence, the productions are executed serially. The production system is called *free*, if no rule is selected and waiting for execution. As long as the procedural module is free, it searches for matching rules.

The modules act in parallel. When a request is sent to a module by a production, the procedural module becomes free while the request is completed. Hence, new production rules can match while other modules might be busy with requests.

ACT-R can be extended by arbitrary modules communicating through buffers with the procedural system. However, to exchange more fundamental parts of the architecture it needs more than only architectural modules as shown in section 3.

## 2.2 The Procedural Module in CHR

The procedural module is the core of ACT-R's production rule system. Our implementation is based on the translation of production rule systems to CHR as presented in (Frühwirth 2009, chapter 6.1). However, we have to account for the concepts of chunks and buffers, since ACT-R differs in those particular points from other production systems. Details of the implementation can be found in (Gall 2013).

The set of chunks can be represented in CHR by a constraint *chunk(C,T)*, where *C* is the name of the chunk and *T* its type. The slots provided by this chunk and their values can be stored in constraints *chunk\_has\_slot(C,S,V)* denoting that chunk *C* has the value *V* in slot *S*. With special consistency rules it can be assured, that no chunk has two values in its slots and that it only provides the slots allowed by its type. Analogously, a buffer is represented by a constraint *buffer(B,M,C)* denoting that the buffer *B* is affiliated with the module *M* and holds chunk *C*. The formal definitions of chunks and buffers can be found in the online appendix (Appendix A).

A production rule can now match and modify the information of the buffer system. The actions are implemented by trigger constraints *buffer\_action(B,C)* which get the name of the buffer *B* and a chunk description *C* represented by a term *chunk(C,T,[(S,V),...])* which describes a chunk with name *C*, type *T* and a list of slot-value pairs representing the values of the chunk's slots. Note that such chunk descriptions can be incomplete in some arguments by simply letting them unspecified.

*Example 2*

The rule from example 1 can be translated to the following CHR rule:

```
buffer(goal,_,C), chunk(C,game),
  chunk_has_slot(C,me,rock), chunk_has_slot(C,opponent,scissors)
  ==> buffer_modification(goal,chunk(,_,[result,win])).
```

The name and type of the chunk in the modification are not specified in the original rule and therefore left blank as well as the me and opponent slots.

### 2.3 Timing and Phases

As mentioned before, the production system of ACT-R is occupied for 50 ms after a rule has been selected. To model such latencies, an event queue has to be added. It keeps track of the current time and holds an ordered set of events which can be dequeued one after another according to their scheduled times. In our implementation, the event queue is implemented as a priority queue sorting its elements after the time and a priority determining the order of application for simultaneous events. Events are arbitrary Prolog goals and can be added by `add_q(Time,Priority,Event)`. The current time can be queried by `get_time(Now)`.

To ensure that a production rule only matches when the module is free, we replace each CHR rule of the form  $C \Rightarrow A$  according to the following scheme consisting of two rules:

```
C \ match <=> add_q(Now + 0.05,0,apply_rule(rule(r,C))).
C \ apply_rule(rule(r,C)) <=> A, get_time(Now), add_q(Now,-10,match).
```

The constraint `match` indicates that the procedural module is free and searches for a matching rule. For the matching rule, an `apply_rule` event is scheduled 50 ms from the current time. This event will actually fire the rule. The actions `A` schedule their effects on the buffers at the current time with different priorities. Requests are only sent to the corresponding module. Its effects on the requested buffer are scheduled at a later time. Finally, a new `match` event is scheduled at the current time `Now` but with low priority of `-10`. This ensures that all current actions are performed before the next rule is scheduled to fire.

Otherwise, if no rule matches and the procedural module is free (i.e. a `match` constraint is present), a rule can only become matching if the content of the buffers change. Hence, a new `match` constraint is added directly after the next event in the queue. This models the fact that the procedural module is searching permanently for matching rules when it is free without adding unnecessary `match` events.

## 3 Conflict Resolution

Only one matching production rule can fire at a time. Hence, if there are multiple applicable productions, the system has to decide which to fire. This process is called *conflict resolution* (McDermott and Forgy 1977). In most implementations, CHR simply chooses the rule to fire by textual order, which is a valid conflict

resolution mechanism. However, in ACT-R a more advanced approach using subsymbolic concepts is needed to faithfully model human cognition.

### 3.1 General Conflict Resolution Process

In (Frühwirth 2009, p. 151) a general method to implement different conflict resolution mechanisms in CHR is given. This method is adapted to our CHR implementation of ACT-R. The first rule of each CHR rule pair from section 2.3 can be replaced by:

```
match, C ==> G | conflict_set(rule(r,C)).
```

Hence, the application of a matching production is delayed by adding the rule to the conflict set instead of choosing the first matching rule to be applied by scheduling `apply_rule/1` as explained in section 2.3. Thereby all matching rules are collected in `conflict_set/1` constraints which then can be reduced to one single constraint containing only the rule to be applied according to an arbitrary strategy.

As a last production rule, the rule `match <=> select.` occurs in the program. This rule will always be applied last (since rules are applied in textual order in CHR). It removes the remaining `match` constraint and adds a constraint `select` which triggers the selection process. This means that the conflict resolution is performed by choosing one rule from the conflict set constraints and removing all other such constraints. If no rule matches, a new `match` constraint is scheduled after the next event.

With the introduction of the `select` constraint, the system commits to the rule to be applied by scheduling the corresponding `apply_rule/1` event as explained in section 2.3. This leads the chosen production to perform its actions since its second CHR rule is applicable. After the actions are performed, the next matching phase is scheduled.

The strategy of how the conflict set is eliminated to one single rule which will be applied may vary and is exchangeable. In the following section, several strategies are presented and implemented.

### 3.2 Conflict Resolution Strategies

There have been several conflict resolution strategies proposed for ACT-R over time. To demonstrate the adaptability of our CHR implementation, we implement some of those strategies. In the reference implementation of ACT-R, such adaptations might need a lot of knowledge about its internal structures (Stewart and West 2007).

In general, ACT-R conflict resolution strategies usually use the subsymbolic concept of *production utilities*. The production utility for a production  $i$  is the function  $U_i : \mathbb{N} \rightarrow \mathbb{R}$  which expresses the value of utility of a particular production at its  $n$ th application which may be adapted according to a learning strategy. In the conflict resolution process, the current utility values are compared for all matching functions and the production with the highest utility is chosen. The production utility can therefore be seen as a dynamic rule priority which is adapted according to a certain strategy.

In the following, we present some different learning strategies to adapt the utility of a production. Eventually, the concept of rule refraction is introduced, which is a general conflict resolution concept and can be applied for all of the presented learning strategies.

### 3.2.1 Reinforcement-Learning-Based Utility Learning

The current implementation of ACT-R 6.0 uses a conflict resolution mechanism which is motivated by the Rescorla-Wagner learning equation (Rescorla and Wagner 1972). The basic concept is that there are special production rules which recognize a successful state (by some model-specific definition) and then trigger a certain amount of reward measured in units of time as a representation of the effort a person is willing to spend to receive a certain reward (Anderson 2007, p. 161). All productions which lead to the successful state, i.e. all productions which have been applied, receive a part of the triggered amount of reward which demounts the more time lies between the application of the production rule and the triggering of the reward. The utility  $U_i$  of a production  $i$  then is adapted as follows:

$$U_i(n) = U_i(n - 1) + \alpha(R_i(n) - U_i(n - 1)) \quad (1)$$

The reward  $R_i(n)$  for the  $n$ th application of the rule  $i$  is the difference of the external reward and the time between the selection of the rule and the triggering of the reward. The utility adapts gradually to the average reward a rule receives. Its calculation can be extended by noise to enable rules with initially low utilities to fire. This then may boost their utility values.

In CHR, this strategy can be implemented as follows: For each production rule, a `utility/2` constraint is stored holding its current utility value. For rules marked with a reward, a `reward/2` constraint holds the amount of reward. When a production rule is applied, this information is stored with the application time in a constraint by the rule `apply_rule(rule(P,_,_)) ==> get_time(Now), applied([(P,Now)])`. With a corresponding rule, the `applied/1` constraints are merged respecting the application time of the rules, since the adaptation strategy depends on the last utility value of a rule and rules might be applied more than once until they receive a reward. This leads to one `applied/1` constraint containing a sorted list of rules and their application time.

If a rule which is marked with a reward is going to be applied, the reward can be triggered by `apply_rule(rule(P,_)), reward(P,R) ==> trigger_reward(R)`. The triggering of the reward simply adapts the utilities according to equation 1 for all productions which have been applied indicated by the `applied/1` constraint respecting the order of application. Afterwards, this constraint is deleted because after a reward has been received, the rule is not considered in the next adaptation.

### 3.2.2 Success-/Cost-Based Utility Learning

In prior implementations of ACT-R, the utility learning is based on a success-/cost approach (Anderson *et al.* 2004; Taatgen *et al.* 2006). A detailed description can

be found in (ACT-R Tutorial 2004, unit 6). Each production rule  $i$  is associated to the values  $P_i$  denoting the success probability of the production and  $C_i$  denoting its costs. In this approach, the utility of a production rule is defined as:

$$U_i(n) = P_i(n)G - C_i(n) \quad (2)$$

Note that the current utility does not depend on the value of the last utility, but can be calculated by the current values of the parameters instead. Hence, the order of application does not play a role. Usually,  $C_i$  is measured in units of time to achieve a goal whereas  $G$  – the goal value – is an architectural parameter and usually set to 20s. The parameters  $P$  and  $C$  are obtained by the following equations:

$$P_i(n) = \frac{\#successes_i}{\#successes_i + \#failures_i} \quad C_i(n) = \frac{efforts_i}{\#successes_i + \#failures_i} \quad (3)$$

The values  $\#successes$  and  $\#failures$  count all applications of a rule which have been identified as a success or a failure respectively. Similarly to the reinforcement-based learning, some productions which identify a success or failure trigger an event which adapts the counters of successes or failures of all production rules which have been applied since the last triggering. The efforts are estimated by the difference of the time of the triggering and the selection of a rule. The values are initialized with  $\#successes = 1$ ,  $\#failures = 0$  and  $efforts = 0.05$  s which is the selection time of one firing. Analogously to the reward-based strategy, utilities can be extended by noise.

Similarly to the implementation of the reinforcement learning rule, the triggering of a success or failure can be achieved by a constraint `success(P)` or `failure(P)`, which encode that a production  $P$  is marked as success or failure respectively. Combined with the `apply_rule/2` constraint, a `success/0` or `failure/0` constraint can be propagated which trigger the utility adaptation. The following rules show the adaptation of  $\#successes_i$  and  $efforts_i$  when a success is triggered and rule  $i$  has been applied before:

```
success \ applied(P,T), efforts(P,E), successes(P,S) <=>
  get_time(Now), efforts(P,E+Now-T), successes(P,S+1).
success <=> true.
```

The number of successes or failures are stored in the respective binary constraints and if a success is triggered, they are incremented for all applied production rules and efforts are adjusted. The rules for failures are analogous. The adaptation of one of those parameters triggers the rules which replace the constraints holding the old  $P_i$  and  $C_i$  values by new values. When a  $P_i$  or  $C_i$  constraint is replaced, the calculation of the new utility value is triggered. To ensure that only one utility value is in the store, a destructive update rule is used.

### 3.2.3 Random Estimated Costs

In (Belavkin and Ritter 2004), a conflict resolution strategy motivated by research results in decision-making is presented. The current implementation varies slightly from this description (Belavkin 2005) and we stick to this most recent approach for a



better comparability of the results. The strategy is based on the success-/cost-based utility learning from section 3.2.2 and uses the same subsymbolic information (the counts of successes and failures and the efforts). However, instead of calculating the average cost  $C_i$ , the expected costs  $\theta_i$  of achieving a success by a rule are estimated:

$$\theta_i := E(C_i) \approx \frac{\text{efforts}_i}{\#\text{successes}_i} \quad (4)$$

From the expected costs  $\theta_i$  of a rule  $i$ , the *random estimated costs*  $\zeta_i$  are derived by drawing a random number  $r_i$  from a uniform distribution  $U(0, 1)$  and setting  $\zeta_i = -\theta_i \cdot \log(1 - r_i)$ . Eventually, production utilities are calculated analogously to the success-/cost-based strategy:  $U_i = P_i G - \zeta_i$ . The influence of the random estimated costs can be varied by adapting the parameter  $G$ . If  $G = 0$ , the production rule with minimal random estimated costs will be fired (as suggested in (Belavkin and Ritter 2004)).

Since this method uses the same parameters as the success-/cost-based variant, almost all of the code can be reused for an implementation. However, instead of the costs, the expected costs  $\theta_i$  are computed and saved in a constraint whenever the success/failure ratio changes. Additionally, the random costs must be calculated in every conflict resolution step and not only when the parameters change since they vary each time due to randomization. Hence, a rule must be added which calculates the utility value as soon as a production rule enters the conflict set:

```
conflict_set(rule(P,_), theta(P,T), succ_prob(P,SP) ==>
  random(R), Z is -T * log(1 - R), U is SP*20-Z,
  set_utility(P,U).
```

The rest of the implementation like the calculation of the success/failure counters, efforts or the pruning of the conflict set is identical to the success-/cost-based strategy.

### 3.2.4 Production Rule Refraction

In contrast to the previous strategies which only exchange the utility learning part, production rule refraction adapts the general conflict resolution mechanism and can be combined with all of the other presented strategies. It was first suggested in (Young 2003) to avoid over-programming of models in the sense that the order of application of a set of rules is fixed in advance by adding artificial signals to ensure the desired order. Rule refraction can avoid such operational concepts by inhibiting the application of the same rule instantiation more than once. To the best of our knowledge, our implementation is the first of its kind for ACT-R.

Refraction can be implemented by saving the instantiation of each applied production using the rule `apply_rule(R) ==> instantiation(R)`. When building the conflict set, the following rule eliminates all productions which already have been applied from the set: `instantiation(R) \ conflict_set(R) <=> true`. This pruning rule must be performed before the rule selection process, so that such productions are never considered as fire candidates.

## 4 Evaluation

After having implemented some different conflict resolution strategies, we test their validity with an example model of the game *rock, paper, scissors*. The idea is that the model simulates a player playing against three opponents with different preferences on the three choices in the game. We then want to observe, how the model adapts its strategy under the different conflict resolution mechanisms and test if the results of the ACT-R implementation and our CHR implementation match.

### 4.1 Setup

The player is basically modeled by the production rules `play-rock`, `play-paper` and `play-scissors` standing for the three choices a player has in the game. At the beginning, the production rules have equal utilities which are then adapted by the utility learning mechanisms of the three conflict resolution strategies. Since we only want to test our conflict resolution implementations, we try to rule out all other factors which could influence the behavior of our model. Hence, we only use the procedural module with the goal buffer and do not simulate any declarative knowledge or even perceptual and motor modules. I.e. the model is not a realistic psychological hypothesis of the game play, but only a test of our implementation. Furthermore, we disable noise where possible to better compare our results. In ACT-R, the canonical parameter setting is not recommended to change without justification (Stewart and West 2007, sec. 1.1). For our experiment, we used this setting.

The moves of the opponents are randomly generated in advance according to their defined preferences: Player 1 simply chooses rock for every move, player 2 chooses only between rock and paper and player 3 chooses equally between all three possibilities. For each player, we produced 20 samples of 20 moves (except for player 1 with only one sample of 20 moves). Their choices are put into the goal buffer one after another by host-language instructions (Lisp and Prolog/CHR). The game is played for 20 rounds until a restart with a new sample which corresponds to 2 s simulation time. Finally, the utility values  $U_{\{r,p,s\}}$  at the end of each run (for rock, paper and scissors respectively) are collected and compared to the reference implementation. We use the notation  $\bar{U}_{\{r,p,s\}}$  to denote the average of those values over all 20 samples. In the following the implementation of the production rule `play-rock`:

```
(p play-rock
  =goal> isa game    me nil      opponent nil
  ==>
  =goal>             me rock    opponent =x   !output! (rock =x) )
```

This rule simply puts the symbol `rock` into the goal buffer indicating that the model chose rock. The variable `=x` is set by built-in functions of the host language (omitted in the listing) modeling the choice of the opponent derived from a given list of moves. The rules for *paper* and *scissors* can be defined analogously. The model has

been translated to CHR by our compiler. We performed the translation of Lisp built-ins to Prolog built-ins by hand.

Furthermore, the model contains production rules detecting a win, draw or defeat situation (similar to example 1) and resetting the choices of the two players in the goal buffer to indicate that the next round begins. Those rules are marked with a reward (positive or negative) or as a success/failure respectively. In the case of a draw, no reward, success or failure will be triggered. Hence, the utility learning algorithms will adapt the values of the fired rules depending on their success.

If the highest utilities in the conflict set are equal, the strategy of ACT-R is undocumented. It depends on the order of the rules in the source code and may vary between the implementations (e.g. the strategy of ACT-R 6.0 differs from ACT-R 5.0 as we found in our experiments). We adapted the order of rules in our translated CHR model to match the strategy of ACT-R. Usually, noise would rule out such differences.

For the reference implementations, we used Clozure Common Lisp version 1.9-r15757. The CHR implementation has been run on SWI-Prolog version 6.2.6. The relevant data collected in our experiments can be found in the online appendix (Appendix B).

#### *4.2 Availability of the Strategies*

Our approach enables the user to exchange the complete conflict resolution strategy without relying on provided interfaces and hooks except for the very basic information that a rule is part of the conflict set or about to be applied. This information relies on the fundamental concept of the match-select-apply cycle of ACT-R. In the reference implementations of the strategies, there are deeper dependencies and assumptions on when and how subsymbolic information is adapted and stored.

This leads to incompatibilities: The reinforcement-learning-based strategy is only available for ACT-R 6.0. Although the success-/cost-based strategy is shipped with ACT-R 6.0, it was not executable for us and hence we had to use ACT-R 5.0 to run it. This leads to further incompatibility problems when using modules not available for ACT-R 5.0 (which is in general difficult to extend due to the lack of architectural modules). Since the method of random-estimated costs relies on the success-/cost-based strategy, it is also only available for ACT-R 5.0.

Our implementation of the refraction-based method is to the best of our knowledge the only existing implementation for ACT-R, although it has been suggested in (Young 2003).

#### *4.3 Reinforcement-Learning-Based Utility Learning*

For the reinforcement-learning-based strategy, we marked the win-detecting production rules with a reward of 2 and the defeat-detecting rules with 0 which leads to negative rewards for all applied rules when a defeat is detected. Draws do not lead to adjustments of the strategy in our configuration. We executed the model on ACT-R 6.0 version 1.5-r1451 and our CHR implementation.

Our implementation matches the results of the reference implementation exactly when rounded to the same decimal precision (see online appendix B.2). Differences of floating point precision did not influence the results, since ACT-R does round the final results to the one-thousandths. As expected, the model usually rewards the paper rule most when playing against player 1 and 2 (average utility at end of round for player 1:  $(\overline{U}_r, \overline{U}_p, \overline{U}_s) = (0, 1.87, -0.02)$ ; player 2:  $(0, 0.81, 0.49)$ ). Exceptions are rounds where the opponent chooses paper above average especially as first moves (e.g. sample 10: 75% rate of paper; first 9 moves;  $U_p = 0, U_s = 1.329$ ). In such cases, scissors has the highest utility. This is reinforced by the relatively high reward of successes compared to the punishment of defeats. However, the winning rate is still very high (15 wins, 5 defeats, no draws). Overall, the behavior of the model is very successful (average: 10.4 wins, 3.9 draws and 5.7 defeats in each sample). For player 3 – as expected – no unique result can be learned; wins, draws and defeats are very close in average (6.6 wins, 6.7 draws, 6.7 defeats).

#### 4.4 Success-/Cost-Based Utility Learning

For the success-/cost-based strategy, the production rules recognizing a win situation are marked as a success and analogously the production rules for the defeat situations as a failure. We used ACT-R 5.0 to test our implementation against the reference implementation, since it is not available for ACT-R 6.0. Again, noise is disabled for better comparability. Because the selection mechanism for rules with same utility differs from ACT-R 6.0, we adapted the order in which the rules appear in the source code.

Our implementation matches the results of the reference implementation exactly (see online appendix B.3). It can be seen that this strategy is not able to detect the optimal moves for player 1. Analyses showed that due to the order of the rules, the model first selects to play *rock*. This leads to a draw and hence no adaptation of the utilities. Hence, rock is played repeatedly. In real-world models, noise would help to overcome such problems. For player 2, the model correctly chose to play *paper* in average even for the samples where the opponent chooses paper more often than rock. However, in average, the model did only win 8.9 out of 20 rounds in a sample and produced 9.1 draws. For each of the samples, only two rounds were lost.

#### 4.5 Random Estimated Costs

Due to the randomness of this strategy, no exact matches of results can be expected. Hence, we executed the models on 3 samples (the first of each opponent) with 50 runs for each sample. The reference implementation has been run on ACT-R 5.0.

The average utilities are close to the reference implementation (error squares of average utilities player 1:  $(\Delta\overline{U}_r^2, \Delta\overline{U}_p^2, \Delta\overline{U}_s^2) = (0.145, 0.000, 0.000)$ ; player 2:  $(0.850, 0.000, 0.098)$ ; player 3:  $(2.823, 0.503, 0.003)$ , see online appendix B.4 for details). It can be seen that for most runs the production with the highest, medium and lowest utility value coincide. For player 1, the random estimated costs overcome the problem of the success-/cost-based implementation as discussed in section 4.4.

## 5 Related Work

There are several implementations of the ACT-R theory in different programming languages. First of all, there is the official ACT-R implementation in Lisp (ACT-R 2014) which we used as a reference. There are a lot of extensions to this implementation which partly have been included to the original package in later versions like the ACT-R/PM extension included in ACT-R 6.0 (Bothell, p. 264). The implementation comes with an experiment environment offering a graphical user interface to load, execute and observe models.

In (Stewart and West 2006; Stewart and West 2007), a Python implementation is presented which also has the aim to simplify and harmonize parts of the ACT-R theory by finding the central components of the theory. The architecture has been reduced to only the procedural and the declarative memory which are used to build other models combining and adapting them in different ways. However, there is no possibility to translate traditional ACT-R models automatically to Python code since the way of modeling differs too much from the original implementation.

Furthermore, there are two different implementations in Java: *jACT-R* (jACT-R b) and *ACT-R: The Java Simulation & Development Environment* (Salvucci b). The latter one is capable of executing original ACT-R models and offers an advanced graphical user interface. The focus of the project was to make ACT-R more portable with the help of Java (Salvucci a). In jACT-R, the focus was to offer a clean and exchangeable interface to all the components, so different versions of the ACT-R theory can be mixed (jACT-R a) and models are defined using XML. There is no compiler from original ACT-R models to XML models of jACT-R. Due to the modular design defining various interfaces which can be exchanged, jACT-R is highly adaptable to personal needs. However, both approaches are missing the proximity to a formal representation.

## 6 Conclusion

In this work, we have presented an implementation of ACT-R using Constraint Handling Rules which is capable of closing the gap between the theory of ACT-R and its technical realization. Our implementation abstracts from technical artifacts and is near to the theory but can reproduce the results of the reference implementation. Furthermore, the formalization itself enables implementations to check against this reference. The implementation of the different conflict resolution strategies has shown the adaptability of our approach. Most of the implemented strategies are not available for the current implementation of ACT-R and our implementation of production rule refraction is unique.

For the future, the implementation can be extended by other modules like the perceptive/motor modules provided by ACT-R. Currently, there is a running student project on implementing a temporal module which may be used to investigate time perception. The formalization and CHR translation pave the way to develop analysis tools (e.g. a confluence test) on the basis of the results for CHR programs.

### Supplementary material

To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/S1471068414000180>.

### References

- ACT-R 2014. The ACT-R Homepage. <http://act-r.psy.cmu.edu/>.
- ACT-R Tutorial 2004. The ACT-R 5.0 tutorial. <http://act-r.psy.cmu.edu/tutorials-5-0/>.
- ANDERSON, J. R. 2007. *How can the human mind occur in the physical universe?* Oxford University Press.
- ANDERSON, J. R., BOTHELL, D., BYRNE, M. D., DOUGLASS, S., LEBIERE, C., AND QIN, Y. 2004. An integrated theory of the mind. *Psychological Review* 111, 4, 1036–1060.
- ANDERSON, J. R. AND LEBIERE, C. 1998. *The Atomic Components of Thought*. Lawrence Erlbaum Associates, Inc.
- BELAVKIN, R. 2005. Optimist conflict resolution overlay for the ACTR cognitive architecture. <http://www.eis.mdx.ac.uk/staffpages/rvb/software/optimist/optimist-for-actr.pdf>.
- BELAVKIN, R. AND RITTER, F. E. 2004. Optimist: A new conflict resolution algorithm for act-r. In *ICCM*. 40–45.
- BOTHELL, D. *ACT-R 6.0 Reference Manual – Working Draft*. Department of Psychology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.
- FRÜHWIRTH, T. 2009. *Constraint Handling Rules*. Cambridge University Press.
- GALL, D. 2013. A rule-based implementation of ACT-R using constraint handling rules. *Master Thesis, Ulm University*.
- jACT-R. Benefits of jACT-R (part of the FAQ section of the homepage). <http://jactr.org/node/50>.
- jACT-R. The Homepage of jACT-R. <http://jactr.org/>.
- MCDERMOTT, J. AND FORGY, C. 1977. Production system conflict resolution strategies. *SIGART Bull.* 63 (June), 37–37.
- RESCORLA, R. A. AND WAGNER, A. W. 1972. *A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement*. Appleton-Century-Crofts, New York, Chapter 3, 64–99.
- SALVUCCI, D. About ACT-R: The Java Simulation & Development Environment. <http://cog.cs.drexel.edu/act-r/about.html>.
- SALVUCCI, D. ACT-R: The Java Simulation & Development Environment – Homepage. <http://cog.cs.drexel.edu/act-r/>.
- STEWART, T. C. AND WEST, R. L. 2006. Deconstructing ACT-R. In *Proceedings of the Seventh International Conference on Cognitive Modeling*. 298–303.
- STEWART, T. C. AND WEST, R. L. 2007. Deconstructing and reconstructing ACT-R: exploring the architectural space. *Cognitive Systems Research* 8, 3 (Sept.), 227–236.
- SUN, R. 2008. Introduction to computational cognitive modeling. In *The Cambridge Handbook of Computational Psychology*, R. Sun, Ed. Cambridge University Press, New York, 3–19.
- TAATGEN, N. A. AND ANDERSON, J. R. 2002. Why do children learn to say broke? a model of learning the past tense without feedback. *Cognition* 86, 2, 123–155.
- TAATGEN, N. A., LEBIERE, C., AND ANDERSON, J. 2006. Modeling paradigms in ACT-R. In *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press, 29–52.
- YOUNG, R. M. 2003. Should ACT-R include production refraction? In *Proceedings of 10th Annual ACT-R Workshop*.