

An account of cognitive flexibility and inflexibility for a complex dynamic task

Alberto De Obeso Orendain (a.de-obeso-orendain@sussex.ac.uk)
Cognitive & Languages Processing Systems Group, School of Informatics
University of Sussex, Falmer, Brighton, BN1 9QJ, UK

Sharon Wood (s.wood@sussex.ac.uk)
Cognitive & Languages Processing Systems Group, School of Informatics
University of Sussex, Falmer, Brighton, BN1 9QJ, UK

Abstract

Problem solving involves adapting known problem solving methods and strategies to the task at hand (Schunn & Reder, 2001) and cognitive flexibility is considered to be “the human ability to adapt the cognitive processing strategies to face new and unexpected conditions of the environment” (Cañas et al., 2005, p. 95). This work presents an ACT-R 6.0 model of complex problem solving behavior for the dynamic microworld game FireChief (Omodei & Wearing, 1995) that models the performance of participants predisposed to behave either more or less flexibly based on the nature of previous training on the task (Cañas et al., 2005). The model exhibits a greater or lesser degree of cognitive inflexibility in problem solving strategy choice reflecting variations in task training. The model provides an explanation of dynamic task performance compatible with the Competing Strategies paradigm (Taatgen et al., 2006) by creating a second layer of strategy competition that renders it more flexible with respect to strategy learning, and provides an explanation of cognitive inflexibility based on reward mechanisms.

Keywords: complex problem solving; cognitive inflexibility; dynamic tasks; strategy use; adaptation.

Introduction

Problem solving involves adapting known problem solving methods and strategies to the task at hand (Schunn & Reder, 2001) and cognitive flexibility is considered to be “the human ability to adapt the cognitive processing strategies to face new and unexpected conditions of the environment” (Cañas et al., 2005, p. 95). When approaching a new problem, it is thought that problem solvers with higher levels of cognitive flexibility will outperform those who are less flexible because the former tend to consider alternative ways to solve the problem (Stewin & Anderson, 1974) rather than rigidly adhering to well-used methods. In their study of cognitive flexibility, Cañas et al. (2005) found that participants became predisposed to behave either more or less flexibly based on the nature of previous training on the task. Those trained repeatedly on the same problem scenario developed a preference for how they solved the task, becoming faster and more fluid in their actions over time. When subsequently tested on a different scenario their behavior was inflexible in adapting to the new test conditions and performance suffered. In contrast, those trained on a series of varying problem solving scenarios demonstrated an ability to adapt their problem solving behavior flexibly to the challenges presented by the new test scenario. The work presented here describes an ACT-R

model for the Cañas et al. (2005) problem solving task that demonstrates varying degrees of cognitive flexibility depending on the training regime it undergoes. Analysis of the model provides an explanation of cognitive inflexibility based on reward mechanisms.

Background

There are several cognitive modeling paradigms (Taatgen et al., 2006) for problem solving involving strategy selection. In the Competing Strategies paradigm (ibid.), several strategies are implemented in a cognitive architecture and then compete for use in solving a problem. According to Taatgen et al. (2006) utility learning can be used to determine the best strategy. This paradigm has been successfully applied in modeling problem solving behavior for static tasks (Lovett & Anderson, 1996; Peebles & Bothell, 2004) and tasks in dynamically changing situations such as Air Traffic Control (Schunn & Reder, 2001; Schoelles & Gray, 2000).

Dynamic problem solving tasks pose an added layer of complexity. In dynamic situations the problem solver needs to execute not only the appropriate action but also to implement it at the right time: a good decision at one moment could be ineffective the next. In order to obtain good performance both selection and execution of the chosen strategy must be effective.

Problem solvers must also be ready to change strategy as and when the situation demands (Gonzalez et al., 2004); they must continuously process feedback in order to select appropriate actions within an ever-changing situation (Brehmer & Dörner, 1993). Underlying this ability, according to Schunn & Reder (2001), strategy choice is influenced by overall success and “Dynamic tasks bring to the forefront the importance of the ability to adapt to changing success rates” (p. 61). They argue that although participants may use a similar set of strategies they can differ in their ability to opportunistically apply those strategies in response to the situation.

This ability to adapt behavior may be affected by factors such as cognitive inflexibility, which can be produced as a consequence of the way problem solvers interact with the task at hand. As skill in a task improves and becomes more automatic so cognitive inflexibility may increase, particularly in tasks with a high level of consistency (Ackerman, 1988). For example, in a fire-fighting task, Cañas et al. (2005) found evidence of cognitive inflexibility in participants trained repeatedly on the same problem scenario who, having found an effective strategy, failed to

relinquish it despite situational changes that reduced its effectiveness. This contrasted with participants trained on a variety of different problem scenarios.

However, studies investigating cognitive inflexibility have not always drawn consistent results. For example, Schunn & Reder (2001) found no evidence for cognitive inflexibility in their study involving training on an Air Traffic Control task when situational changes affecting success on the task were introduced.

The work presented here implements an ACT-R model of the Cañas et al. (2005) study to elucidate the mechanisms of cognitive inflexibility further in an attempt to reconcile these disparate findings.

The FireChief Microworld

The Cañas et al. (2005) study used a dynamic microworld game called FireChief (Omodei & Wearing, 1995) for the problem solving task. Figure 1 shows the FireChief display.

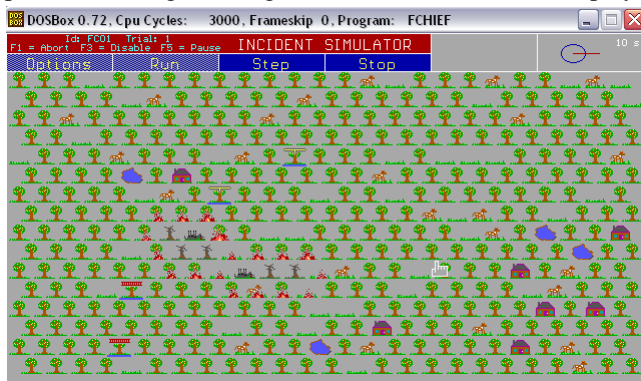


Figure 1: The FireChief microworld display

Players combat fires spreading in a landscape using truck and copter fire-fighting units. A FireChief problem scenario depicts a landscape comprising forest, clearings and property, the position of initial fires, fire-fighting units, and the direction and strength of the wind. Copter and trucks can be moved between landscape grid cells and Drop Water (DW) over cells to extinguish fires. Copters move three times faster than trucks and cannot be destroyed by fire, but a truck's water tanks have double capacity and are able to Control Fire (CF) by creating a fire-break. Commands are issued through a combination of mouse and keyboard operations and their execution takes a fixed amount of time (4 seconds to DW; 2 seconds to CF) and a variable amount of time to Move a unit depending on distance and type of unit. Wind strength and direction are in the upper right-hand corner of the display. Task performance is inversely proportional to the number of cells destroyed by fire at the end of the trial.

The FireChief problem state changes both independently and as a consequence of the participant's actions and time pressure is directly related to fire development, which depends heavily on wind strength.

The Cañas et al. (2005) study

Each trial for the FireChief task lasts 260 seconds. The experimental data comprises a list of commands executed

during each trial that is indexed to a detailed description of the changing scenario. The first 16 trials comprise the training phase and the last 8 trials the testing phase. There were two training conditions: constant and variable.

In the constant training (CT) condition the problem scenario is exactly the same for each trial and wind strength and direction remains fixed. In the case of variable training (VT) a different scenario is presented in each of the sixteen trials. Trials vary in landscape composition, initial position of fire-fighting units and fires and, importantly, wind direction and strength varies throughout the trial.

There are also two test conditions. In the Wind Direction Change (WDC) condition the wind changes direction every 60 seconds. These shifts in wind direction have a dramatic impact on fire development. In the second Efficiency Reduction (ER) test condition, appliances deliver less water and are therefore less effective in extinguishing fires.

As previously hypothesized, Cañas et al. (2005) found participants in the CT condition improved performance as the number of trials increased; however, during the test phase this same group demonstrated a distinct lack of flexibility in adapting their problem solving strategy to the new task demands. In contrast, participants in the VT condition demonstrated a greater facility for changing strategies under test conditions. The findings were consistent across both WDC and ER test conditions

The Model

The ACT-R 6.0 (Anderson et al., 2004) model interfaces to a LISP version of the FireChief microworld (De Obeso Orendain & Wood, 2010). Task knowledge comprises both procedural (condition-action) rules that produce behavior according to four high level strategies: *Barrier*, *Non-Barrier*, *Stop*, and *Follow* (ibid.) and three declarative knowledge components that impact this behavior: (1) the *goal chunk*, the main task objective is to extinguish the fire; (2) the *strategy specification chunk*, which defines whether the model will use a mixture of DW and CF commands, whether or not a barrier will be created, and which method of attacking the fire is preferred (attack weak fires, attack strong fires or attack the strongest fire); and (3) the *intention chunk*, used to track the current intention (stored in the ACT-R imaginal buffer, Anderson et al., 2004). Intentions emanate from steps in pursuit of the main goal, according to the chosen strategy.

The model identifies its preferred strategy by comparing the utility of its four strategy rules, combined with a situation assessment, and retrieves the corresponding *strategy specification chunk*. This chunk remains unaltered throughout the entire trial, unless there is a strategy change.

Overall the model behavior reflects the use of procedural knowledge over declarative knowledge: it is constructed in such a way that it is mainly controlled by the utility learning mechanism. The content of the three declarative chunks determine which rules are applicable in different situations, but there is always more than one eligible rule, so the decision about what to do next is taken in terms of utility.

ACT-R's utility learning mechanism

Utility designates the perceived value of implementing a procedural rule, and thereby its associated behavior, and is updated via a reward mechanism reflecting task success. Throughout runtime, Rule utilities are compared during the process of conflict resolution where only the rule with the highest utility is selected and thereby acted upon. In ACT-R when a reward is triggered the utility values of all rules that have fired since the last reward are updated. The actual reward allocated depends on the absolute value of the reward and the length, in time, between the giving of the reward and the execution of that rule. The consolidation of strategies and the existence of cognitive inflexibility discussed here are explained in terms of utility variations in the set of rules identified as key in implementing a strategy. A key rule is one that enters the conflict set during ACT-R conflict resolution and hence competes in determining the next intention or action of the model.

Achieving adaptivity

The considerable variability observed in participants' protocols suggests that for the FireChief task there is variation not only in strategy choice, but also in the chosen method of execution. The dynamic nature of FireChief introduces a dynamic component into the execution of strategies that forces a second layer of competition between alternative courses of action within the same strategy. For this reason a paramount feature of the model is to enable this kind of competition. In the FireChief task there are four fire-fighting units (Copter, Truck), three commands (DW, CF, Move) and four hundred locations. From a very broad perspective the model's operations are devoted to determining the agent, type and spatial location of the next command and a strategy functions as a mechanism for helping the model to constrain this decision. Two types of control coexist within the model. The current representation of the task (the *strategy specification chunk*) guides actions through top-down control. Nevertheless bottom-up control is particularly relevant when considering dynamic tasks therefore feedback from the environment is used to guide the further selection of actions by triggering a wider variety of rules than those specified in the strategy chunk.

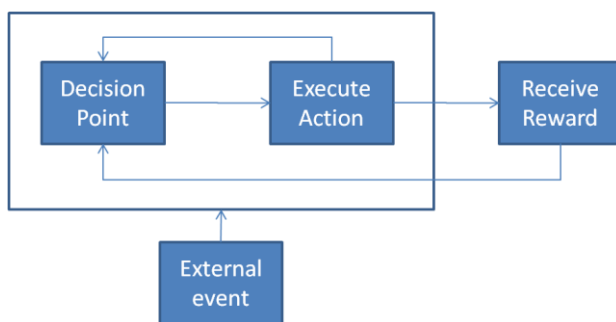


Figure 2: The basic cycle of the model comprising a second layer of within-strategy competition

The model uses an adaptation of the Competing Strategies paradigm (Taatgen et al., 2006): the core of the model is the *Decision Point/Action/Reward* cycle shown in figure 2.

The basic cycle starts with a *Decision Point* (identifying eligible rules) continues with the *Execution* of an action (rule-firing), and finishes with the awarding of a *Reward*. The branching factor at every *Decision Point* is variable and there are *External Events* that can interrupt the flow of actions in the cycle such as alarms and visible changes in the environment that prevent the effects of an action taking place, for example, a cell catching fire before a CF command is completed. The model is designed in such a way that *Decision Points* occur frequently. In this way the model is mainly governed by the utility values of its rules. This bottom-up control feature results in the emergence of interesting behaviors (observed in participants) such as “waiting behavior”: when a truck is Moved to a cell with the intention of issuing a CF command, if the movement's length is shorter than 2 cells, the model tends to wait for the unit to arrive (incurring in a waste of time but increasing the probability of issuing the CF command as soon as the unit arrives, rendering its success more likely). The description and analysis of emergent behaviors is outside the scope of this paper.

The same set of rules is used for modeling performance of the task under both training conditions from the Cañas et al. (2005) study. However, rewards for task performance and thus specific rule utility values will vary according to the unique experience of the model on any given trial (model run). Furthermore, these utility values will accumulate over both training and testing phase.

Rewarding the execution of commands

Within the model positive rewards are received for successfully completing commands and negative rewards for failing to execute commands successfully or for wasting time (this means that the utility of a rule can be negative). In this way, any action that contributes to the successful completion of a command is rewarded predisposing the model to continually issue commands. External reward: final performance

In addition to built-in ACT-R utility learning mechanisms a further external reward mechanism affects the utility of the four strategy rules. The strategy rule invoked for a given trial is modified at the end of each trial based on final performance (the amount of non-destroyed terrain remaining at the end of the trial). For instance, if the rule that selects the *Stop* strategy is fired and the final performance achieved during the trial is high, the rule's utility is increased. Manipulating rule utilities outside the standard ACT-R mechanism, has also been used elsewhere (e.g., Schoelles & Gray, 2000).

Results

Data fitting: The model was fitted to the Cañas et al. (2005) study participant data as described in De Obeso Orendain & Wood (2010).

Performance: During the *training* phase the average performance of participants in the CT group is 78.7 while

the average performance of the model for CT is 77.1. In the VT group, the average performance of participants is 78.45 versus 81.2 for the model. The fit of the model is better for the *Barrier* and *Stop* strategies ($r^2=.987$) which are the most structured strategies (De Obeso Orendain & Wood, 2010).

Strategy use: For the CT *training* scenario the *Barrier* strategy using CF commands to construct a fire-break (ibid.) is a good option because the fire develops quickly and soon reaches an intensity that surpasses the capability of the fire-fighting units. In the CT condition both participants and the model use the *Barrier* strategy increasingly more frequently, by trial 16 participants use the *Barrier* strategy 71% of the time while the model is using it 79% of the time.

Strategy change: During the *training* phase participants in the VT group change strategy with more frequency than participants in the CT group, the model captures this tendency ($r^2=.93$ RMSD=1.43). The fact that both participants and model use the *Barrier* strategy more frequently, and there is less strategy change, during CT facilitates the consolidation of this strategy in the CT group.

Learning in CT: A significant performance increment was obtained by comparing the first (1-4) and last four *training* trials (12-16) for both participants and the model. ($F(1,33)=4.417$, $p<.05$ and $F(1,33)=5.17$ $p<.05$ respectively). This means that consolidating the use of the *Barrier* strategy is beneficial by objective criteria.

Cognitive inflexibility: After the training period both participants and the model undergoing the CT condition exhibit inflexibility on two levels: strategy choice and strategy implementation. Both kinds of inflexibility can be traced to variations in key rule utility values induced by the two training conditions.

The set of rules available for use are exactly the same for both training conditions (a single model undergoes either of the training conditions). However, the pattern of change in utility values varies as a consequence of the training received. As shown in figure 3 for the *Barrier* strategy: over the sixteen training trials average utility values of *Barrier* strategy rules for the CT group (TOP-DOWN CT) far exceed those for the VT group (TOP-DOWN VT).

This contributes towards an explanation of cognitive inflexibility in strategy choice. As a consequence of the CT condition, the reward function shapes the utility values of the model's rules in such a way that it becomes relatively insensitive to changes in reward. The high utility values of rules for the preferred *Barrier* strategy in the CT group shield the model from relatively small variations in success. When creating a barrier is no longer the best approach, such as occurs during the test phase, the model will eventually change its behavior through repeated negative reward after the utility values of the rules for the preferred strategy have reduced sufficiently in comparison to the rules for alternative strategies. But this takes time, giving rise to the observable phenomenon of cognitive inflexibility.

In contrast, the model subjected to the VT condition is more sensitive to changes in reward during the test phase because its rules for implementing alternative strategies are more evenly weighted; because the differences between their utility values is smaller, a small amount of negative reward is able to trigger a switch to an alternative strategy.

Differences in utility also contribute towards an explanation of cognitive inflexibility in strategy implementation, again discussed here in relation to the *Barrier* strategy.

There are a range of actions that might be involved in constructing a barrier by a variety of methods represented as the set of rules available whenever the *Barrier* strategy is selected. One subset of rules comprises methods that implement the *Barrier* strategy in a structured top-down manner. For example, top-down *Barrier* strategy rules systematically identify the next section of the barrier to be constructed by locating CF commands in grid cells adjacent to that section of the barrier just formed.

In comparison, other rules involve a greater degree of bottom-up control in implementing actions. For example, a bottom-up strategy rule might locate the next section of the barrier to be constructed by looking to see where the fire is before making a decision about where to put the next section of the barrier. These top-down and bottom-up rules compete throughout the creation of a barrier (while the *Barrier* strategy is selected) and those selected by ACT-R give rise to the final form of the barrier.

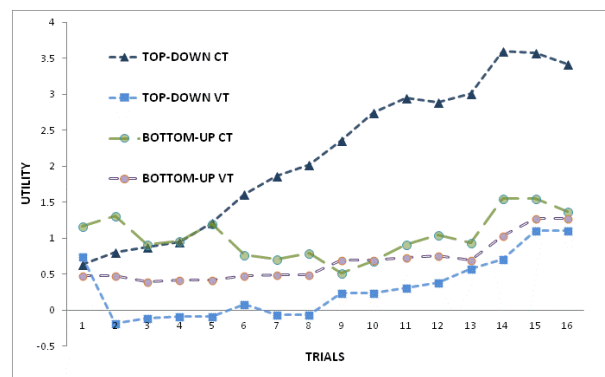


Figure 3: Changes in top-down strategy and bottom-up responsive *Barrier* rule utilities during training

Figure 3 shows the average utility values for these two notional subsets of rules over the sixteen *training* trials: the utility of the rules implementing the strategy top-down increases as more trials are completed during CT (TOP-DOWN CT) as their repeated use is continuously rewarded. This phenomenon occurs only when the problem scenario does not vary dramatically between trials so that there is no significant variation in the effectiveness (and thus reward value) of the actions being executed on repeated trials. In comparison the bottom-up responsive rules involve many more perceptual actions to locate the spread of fire, taking longer to construct the barrier, consequently receiving a relatively lower reward (BOTTOM-UP CT). Over time, this serves to increase the probability of using the top-down subset of rules in the CT group producing the divergence shown in Figure 3. The utility values for the same notional subsets of rules for the VT group, again, remain more evenly balanced owing to the variability in training rewarding the top down implementation of the strategy less consistently.

As in the case of strategy choice, CT leads to cognitive inflexibility in strategy implementation, with potentially insufficient regard given to sensing the environment over top-down construction of the barrier, when conditions change, as witnessed for the CT group under test conditions.

Testing phase: Comparisons were made to determine the impact of cognitive inflexibility on performance in the first *testing* trial. The average performance in the 17th trial in the ER condition is better for participants/model after CT (86.09/78.14) than after VT (72.19/69.83). Both participants and the model in the CT group use the *Barrier* strategy more effectively than the VT group in the ER condition, an indication that these participants have consolidated the Barrier strategy following a top-down approach. The CT group does not need to change strategy because using CF commands is the only sure way to stop the fire in the ER condition (and constructing a barrier using CF is the best approach and therefore has an advantage). The average performance in the 17th trial in the WDC condition is better for participants/model after VT (78.51/78.14) than after CT (71.38/74.87). This is because shifts in wind direction make fire behavior unpredictable so flexible behavior is required. This flexibility is best achieved using more situation-sensitive responsive rules such as those contributing more bottom-up control in the creation of the barrier.

Control of behavior: Figure 3 shows that the model trained in the CT condition has a clear preference for the use of top-down control while the model trained in the VT condition has no such preference. This difference has an impact in the WDC test phase when the wind changes direction in trial 17 at second 60. In a model trained in the VT condition the bottom-up rules are more easily able to win the competition through small variations in utility values following negative reward. Therefore, when the change in the wind occurs, the model will probably select the next target cell based on the location of the fire. On the other hand, the behavior of a model trained in the CT condition will reflect its high utility rules implementing the top-down approach to the creation of the barrier so it will continue to place the next section of barrier without recourse to observing the fire. The risk is that when the form of the barrier is constructed without considering the actual shape of the fire it may not be effective. In this sense the *automation* of the strategy (cf. Ackerman, 1988) runs the risk of deterring the problem solver from extracting relevant information about the problem state to guide behavior.

To validate the results obtained with the cognitive model, further evidence to support this interpretation was sought from the spatial distribution of CF commands in the Cañas et al. (2005) study data for participants during the WDC testing phase: groups CT-WDC and VT-WDC to determine whether the semicircle pattern, a top-down control outcome, was present. These test groups were chosen because the wind direction change test condition alters the path of the fire in such a way as to make the top-down control implementation of the Barrier strategy less effective than a bottom-up more responsive mode of barrier construction. It was found that the CT-WDC group data presents a semicircle pattern of barrier, evidence of top-down

application of the *Barrier* strategy, whilst the VT-WDC group does not. This indicates that the semicircular pattern does not emerge in the VT group behavior because the variability of both the VT condition and the WDC testing phase does not reward the rules implementing it.

Discussion

The model captures the behavior of both training groups with a single set of rules for implementing all four strategies either or both bottom-up and top-down control. Participants in the CT group have the opportunity to consolidate their strategies and hence generate quick, fluid actions; while those in the VT group execute more controlled, albeit flexible, actions. When the testing phase begins people in the CT group are less (cognitively) flexible in adapting to the new demands of the task. In general terms, participants in the VT condition changed strategy more often and showed more cognitive flexibility during the testing phase. The model demonstrates how cognitive inflexibility can be traced to the utility values of rules governing behavior indicating the potential role of reward feedback learning mechanisms in complex problem solving in dynamic domains.

The CT condition presents to the model more stable feedback from the environment (in the form of rewards) to its actions in comparison with the VT condition. In the CT condition the model tends to respond by executing CF commands in a fashion that resembles a barrier. As experience in the task is gained, the model learns how to deploy this strategy with more efficiency.

The ACT-R reinforcement learning mechanism is able to capture the phenomenon of cognitive inflexibility but in order to achieve this it was necessary to provide the model with adequate responsiveness. Rather than following a recipe to implement a strategy, the approach used in this research was the *Decision Point/Action/Reward* cycle which (using standard ACT-R mechanisms) maximizes the number of decision points during strategy execution and thereby enforces competition between rules in selecting the next action at almost every time step so that the model can find the best way of implementing a strategy. This reflects the model's dependence on ACT-R's sub-symbolic processes. In this way, the model was able to capture critical aspects of the data including interesting phenomena such as waiting behavior. This indicates that in complex dynamic tasks participants may be aware of the consequences of their actions over relatively small time intervals.

This study contributes to our understanding about strategy use in complex dynamic tasks: which strategies are used, how they are selected, and how strategy execution changes as experience is gained. Good performance is linked to an effective combination of strategic control with attention to changing task demands.

The cognitive model also prescribes a mechanism in which environmental feedback controls how actions are selected in a highly dynamic task. Through the implementation of the cognitive model it was found, for example, that strategy execution depends on the fine-tuning of ACT-R production rule utilities as a consequence of

environmental rewards. Selecting actions based on utility comparisons facilitates a fluid and quick selection of actions that is instrumental in obtaining good performance, particularly in dynamic and time pressured situations. In dynamic tasks there is a continuous competition between top-down and bottom-up control. This competition is mediated by the characteristics of the learning process such as those exemplified in the Cañas et al. (2005) study, for which in the CT condition the top-down form of control dominates. The account provided by the model is that rules implementing top-down strategic control come to dominate behavior increasingly over rules implementing bottom-up responsive behavior during the CT phase owing to task consistency. This phenomenon increases the probability of performing well in the CT problem scenario but also produces cognitive inflexibility.

As mentioned in the introduction, Schunn & Reder (1996) found no evidence for cognitive inflexibility in their ATC study regarding strategy selection (choice of runway – long or short – on which to land aircraft) despite a long training period. However, we can learn from the work presented here; this would indicate that rules involved in the selection of choices in behaviour (for example, choosing between runways on which to land aircraft) have similar utility. A critical factor that enabled the dominance of certain rules in FireChief was the high consistency of the CT trial. In this respect, the ATC task is only partially consistent. An examination of the Ackerman (1988) study, from which the data for the second experiment of Schunn & Reder (1991) was extracted, reveals that weather conditions (wind speed, wind direction, and ground condition) varied randomly about twice a minute, and also that within each trial aircraft type, of which there are four, are randomly drawn from the queue. It seems that this experimental design shares more similarity with the VT condition in the Cañas et al. (2005) study rather than the CT condition, so when experimental changes are introduced no subset of rules has become dominant.

This research also provides an explanation of how dynamic tasks can be modeled using the Competing Strategies paradigm by incorporating an additional layer of within-strategy execution competition, enabling the bottom-up manifestation of strategies, such as that described here.

Acknowledgments

This research is funded by CONACYT. Special thanks to Hewlett-Packard GADSC MEX DEV Center for a conference grant.

References

Ackerman, P. L. (1988). Determinants of individual differences during skill acquisition: Cognitive abilities and information processing. *Journal of Experimental Psychology: General*, 117, 288-318.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review* 111, (4), 1036-1060.

Brehmer, B., & Dörner, D. B. (1993) Experiments With Computer-Simulated Microworlds: Escaping Both the Narrow Straits of the Laboratory and the Deep Blue Sea

of the Field Study. *Computers in Human Behavior*; Vol. 9, pp. 171-184.

Cañas, J.J., Antolí, A., Fajardo, I., Salmerón, L. (2005) Cognitive inflexibility and the development and use of strategies for solving complex dynamic problems: effects of different types of training. *Theoretical Issues in Ergonomic Science*, 6 (1) 95-108.

De Obeso Orendain, A., & Wood, S. (2010). Cognitive Modeling of Strategies in Dynamic Tasks. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 52-57). Philadelphia, PA: Drexel University.

Gonzalez, C., Thomas, R., & Vanyukov, P. (2004). The relationships between cognitive ability and dynamic decision making. *Elsevier Intelligence* 33 (2005) 169–186.

Lovett, M. C. & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology*, 31(2), 168-217.

Omodei, M.M., Wearing, A.J. (1995). The FireChief microworld generating program: An illustration of computer simulated microworlds as an experimental paradigm for studying complex decision-making behavior. *Behav. Res. Meth. Instr. & Comp.*, 27, 303-316.

Peebles, D. & Bothell, D. (2004). Modelling performance in the sustained attention to response task. In *Proceedings of the sixth International Conference on Cognitive Modeling* (pp. 231-236). Pittsburgh, PA: Carnegie Mellon University/University of Pittsburgh.

Schoelles, M. J., & Gray, W. D. (2000). Argus Prime: Modeling emergent microstrategies in a complex simulated task environment. In N. Taatgen, & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modeling* (pp. 260-270). Veenendaal, NL: Universal Press.

Schunn, C. D., & Reder, L. M. (2001). Another source of individual differences: Strategy adaptivity to changing rates of success. *Journal of Experimental Psychology: General*, 130(1), 59-76.

Stewin, L., & Anderson, C. (1974). Cognitive complexity as a determinant of information processing. *Alberta Journal of Educational Research*, 20(3), 233-243.

Taatgen, N. A., Lebiere, C. & Anderson, J.R. (2006). Modeling Paradigms in ACT-R. In R. Sun (ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press; 29-52.