Strategy use while learning to perform the Kanfer-Ackerman Air Traffic Controller ©[1] task

Bonnie E. John (BEJ@CS.CMU.EDU)
Yannick Lallement (YANNICK@CS.CMU.EDU)
School of Computer Science
Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

People chose different strategies for performing tasks, and that choice often plays a key role in performance. We investigate the use and evolution of strategic behavior in the Kanfer-Ackerman Air Traffic Controller© task, a fast-paced, dynamic task. We present strategies in two dimensions for one aspect of the task, examine how people use them and switch between them, and how their use relates to final performance. We also discuss the implications that the observed variety of strategic behavior has for cognitive modeling.

Introduction

Problem-solving and learning in the real world often occur in dynamic situations. The teacher speaks at her own pace while the student understands her words, decides what's important, and takes notes. The apprentice watches the fire chief assess the course of a burning factory and direct teams of fire-fighters. The nine-year old learns to play a seemingly manic new video-game in an afternoon.

Such domains have been studied in applied areas under the names "naturalistic decision-making" (Klein et al., 1993), "supervisory control" (Sheridan, 1987), and "highly-interative tasks" (Bauer & John, 1995). On the other hand, cognitive psychology has predominantly studied static environments, where the world only changes in response to the actions of a person, not of its own volition. However, the rich understanding of problem-solving and learning attained in static environments provides cognitive modelers with a firm foundation for making contributions to mechanistic models of how human perception, cognition, and action interact with a dynamic outside world allowing real-time performance of a task as well as learning to improve performance.

Prior to modeling, the characteristics of behavior in dynamic domains must be carefully laid out. Detailed performance data over time can give insights into what mechanisms are in play. As discussed in (Lee et al., 1995), the Kanfer-Ackerman Air Traffic Controller©[1] (KA-ATC©) task is an ideal vehicle for studying problem-solving and learning in a dynamic environment. As well as the task environment itself, timestamped keystroke data from over 3500 participants are available on a CD-ROM (Ackerman & Kanfer, 1994). These data are the basis for learning models that use different AI and cognitive architectures.

[1]The Kanfer-Ackerman Air Traffic Controller Task program is copyrighted software by Ruth Kanfer, Philip L. Ackerman, and Kim A. Pearson, University of Minnesota

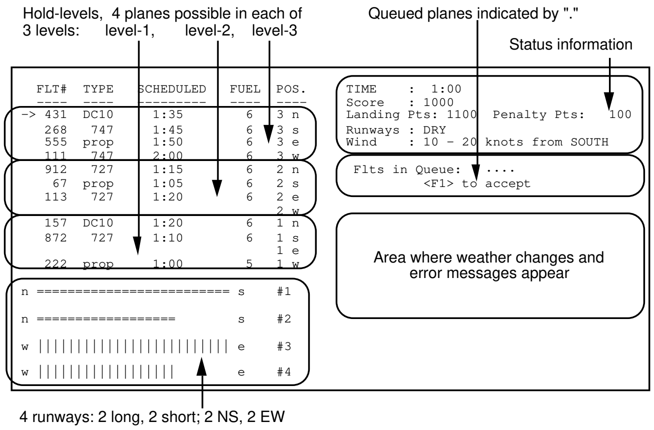

Figure 1: Startup screen of the KA-ATC© task, reconstructed from Study 2 on the Ackerman-Kanfer CD-ROM (1994), and annotated to show the hold-levels, the runways, and areas where information is given to the participants.

Ackerman analyzed these data with respect to independently-measured performance on a battery of cognitive, perceptual and psycho-motor tests in an effort to predict performance on this and other tasks (e.g., Ackerman, 1988). Lee et al. (1995) looked at aggregate data of performance and identified two efficient strategies that gradually increase with experience. In this paper, we look at the details of each individual's performance to identify several different strategies and strategy shifts which a cognitive model must be able to emulate if it wishes to reflect mechanisms employed by humans.

The KA-ATC© Task

The KA-ATC© task is a dynamic task where participants are presented with the start-up screen shown in Figure 1. Planes in a hold-pattern in the upper left corner of the screen must be moved down to runways in the lower-left corner before they run out of fuel. The planes are moved between adjacent hold-levels and from hold-level 1 to the runways using cursor-movement and function keys. A complex set of rules constrain which planes can land on which runways depending on the wind direction, wind speed, and weather conditions.

Once a plane is assigned to a runway, it takes 15 seconds to move across the runway before disappearing from the screen. As time passes, the planes use up their fuel (indicated in the FUEL column, in minutes until crash), the wind and weather

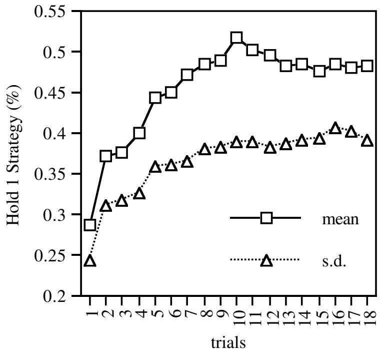

Figure 2: Mean hold 1 strategy and standard deviation (reprinted from Lee et al., 1995, with permission).

changes, and more planes queue up to be admitted to the hold-pattern. Planes are accepted into an empty hold-row from the queue by hitting the F1 key. Points are awarded for landing planes; points are substracted for crashing planes, landing planes with low fuel, and attempting to move planes to places that violate the rules (e.g., to an already occupied hold-row or to the wrong runway for the current weather conditions).

Anecdotally, in the first few minutes doing this task it feels as though the system is driving you. There seems to be a lot to pay attention to, many decisions to make, and time-pressure as you see the fuel counting down to a crash. However, after a few minutes, you gain control and seem to have time to plan your actions. In fact, very soon you find that you are always waiting while the plane moves slowly across the runway, so you can land another plane on that runway. How does this transition from system-driven performance to user-controlled performance happen? How much does simple speed-up of the psychomotor responses contribute, versus the learning of more efficient strategies? How can someone learn more efficient strategies while they are being driven by the system? To begin to answer these questions, we look at the performance of individuals on this task and then discuss the implications of their performance for cognitive modeling.

Previous Analyses

Lee et al. (1995) examined the performance of participants in Ackerman's Study 2 on the KA-ATC© CD-ROM (1994; reported in (Ackerman, 1988)). They looked at two different aspects of performance: which hold-level the participants used to bring planes in from the queue and how efficiently they used the runways when wind-direction changes occurred. We will look more closely at the first of these.

Lee et al. (1995) identified the "hold 1" strategy as the percentage of planes brought directly from the queue into hold level 1, bypassing levels 2 and 3 and saving 6 to 12 keystrokes, on average, to land each plane. Use of the hold-1 strategy increased over the first nine trials and then reached asymptote (Figure 2). But the variance remained high, which led Lee et al. to believe this to be an important source of individual differences. Indeed, prior research in strategy use in static tasks demonstrates that people use several different strategies even in relatively simple tasks (e.g., Reder, 1982, 1988; Siegler, 1996). Looking at the individual data itself will show us just how these individual differences are manifest in KA-ATC© performance.

Detailed behavior in queue acceptance

People perform very differently in the KA-ATC© task, with final scores ranging from 1860 to 4100. They improve substantially through the 18 trials from a mean of 176 to a mean of 3351 ($F = 1310$, $p=0.0001$). In this section we examine the different strategies they use that may produce these differences, and how they switch strategies throughout the trials.

Observed strategies

Looking at the individual performance data of accepting planes from the queue into the hold-pattern, Figure 3 shows timelines with seconds since the start of the trial (x-axis) and the 12 hold-rows (y-axis). The three hold-levels are separated by dotted horizontal lines. Each time a plane is brought in from the queue, a dot appears on the timeline at the hold-row in which the plane was accepted. We constructed timelines for the first 18 trials of 58 participants[2] in Ackerman's 1988 study to examine the different queue-acceptance strategies which contribute to the aggregate strategy shift reported by Lee et al. (1995). We have identified two dimensions of queue-acceptance strategies: the level into which the planes are brought (hereafter *level*), and the patterns of filling and emptying the hold rows (hereafter *pattern*).

Strategies involving levels In the KA-ATC© task, there are 3 hold-levels with 4 hold-rows apiece. There are seven possible combinations of hold-levels (i.e., seven level strategies): level 1-only, 2-only, 3-only, 1&2, 1&3, 2&3, and ALL levels. We assigned each trial to a single-level category if more than 90% of its dots were in that specific hold-level (e.g., 1-only). If a trial could not be assigned to a single-level category, if 90% of its dots were in two levels, it was assigned to a double-level category (e.g., 1&2). Finally, if a trial could not meet any of the previous criteria, it was assigned to the ALL category. For analyses using quarter-trials (e.g., Table 1), we used the same procedure and criteria at that smaller grainsize.

Strategies involving patterns We identified three pattern-dimension strategies: *stacked, sequential*, and *opportunistic*. *Stacked* indicates that the participant stacks up a series of planes one right after the other. This is evidenced by straight, almost vertical, lines of dots in the timelines, separated by blank areas indicating that the participant is landing the stacked-up planes (Figure 3, top). *Sequential* indicates that the participant is attending to one plane at a time, bringing it in to a particular hold-row and then landing it, bringing another one in and then landing it. This is evidenced by horizontal lines of dots (Figure 3, middle). *Opportunistic* indicates that the participant manipulates several planes at a time, interleaving acceptance from the queue and landing planes. This allows him or her to take advantage of slack time (for example, when the runways are busy) to bring new planes into the hold pattern. This is evidenced by seemingly random dots in

[2] 65 participants were involved in that study, but 5 did not complete 18 10-minutes trials and 2 could not be reconstructed from the KA-ATC© CD (Ackerman, 1994). This is the same data set used by Lee et al. to produce the graph in figure 2. These trials involve only fair weather; trials 19-27 add foul weather to the task.

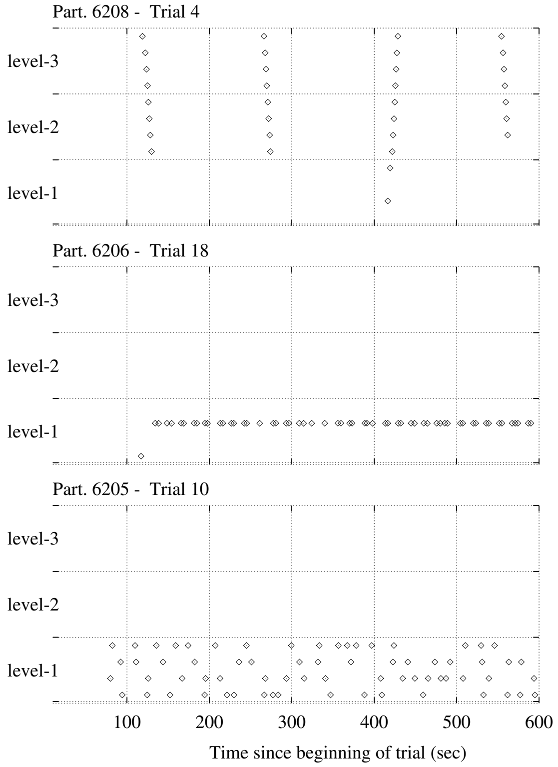

Figure 3: Strategies: stacked in level 2&3 (top), sequential in level 1 (middle), opportunistic in level 1 (bottom).

the timelines (Figure 3, bottom) as the participant alternates between accepting some planes and landing others.

To label the pattern strategies, we used a classifier based on the OC-1 decision tree algorithm (Murthy et al., 1994). We trained this classifier with a set of 424 hand-labelled trials from another study (Ackerman & Kanfer, 1994) in which participants performed the task in the same conditions as (Ackerman, 1988). The inter-rater reliability between the two authors for the hand-labelling was 85.7%.

To label a new trial, the classifier first tried to label the full trial, and, if no strategy was identified, then tried to label its two halves separately (the halves are measured from the time the first plane is brought into the hold pattern, not from the start of the trial). The process repeats to the level of quarter trials. If no strategy was identified at that level, the quarter trial was labelled *non identifiable strategy* (NIS). The classifier used three sets of 50 trees (one set for whole trials, one set for half-trials and one set for quarter trials). Each tree has been trained on a random subset of the training set (using the bagging method (Breiman, 1994)). At each level, the 50 trees voted, providing a confidence rate for the decision.

The agreement rate between the classifier and the authors' consensus on the 424 hand-labelled trials was 88.4% at the level of quarter trials. The average confidence rate was 90.5% for the agreement cases, and 68% for the disagreement cases. For confidence rates over 90%, the agreement was over 90%.

Distribution of the strategies

All of the strategies described above, in both level and pattern, were observed in the participants' performance (Table 1). Although these two dimensions could be orthogonal in theory, in practice they are not. The stacked pattern is used predominantly with long stacks that stretch further than a single level; 91% of all stacks are in levels 1&2, 2&3 or ALL levels. The sequential pattern is used exclusively in the bottom level (level 1). The opportunistic pattern is predominantly in the lower levels; 90% in levels 1 or 1&2. In contrast, the quarter-trials which do not fall into our identified pattern-strategies span all possible level-strategies fairly evenly.

	1	2	3	1&2	1&3	2&3	All	Tot.
NIS	4.3	1.2	2.7	6.8	1.2	2.9	3.3	22.4
STA	0.7	0.6	1.5	8.5	0.4	12.5	12.1	36.3
SEQ	6.4	0.0	0.0	0.0	0.0	0.0	0.0	6.4
OPP	22.7	2.6	0.6	8.8	0.0	0.1	0.1	34.9
Tot.	34.1	4.3	4.8	24.1	1.7	15.5	15.5	100

Table 1: Distribution in percentage of the level and pattern strategies at the quarter trial level for all the participants and all the trials.

These strategy combinations "make sense" in terms of the task environment. The keystroke-pattern for stacking planes involves repeating a series of three quick keystrokes for each plane brought in from the queue and no perception to make sure that the row is empty (because they are all empty when stacking begins). Thus, it is easy to get into a rhythm that produces long stacks. The level 1-only sequential strategy may result from the fact that the system places the cursor back on the hold-row from which the last plane was assigned to a runway. Since this row is necessarily empty, it only takes two keystrokes to accept a plane, and since planes can only be landed from level-1, the easiest sequential acceptance is always at level-1. It is not as obvious why the opportunistic pattern concentrates in the lower levels, but lower levels require fewer keystrokes to land the plane, so any participant rationally optimizing his or her score would gravitate to the bottom levels. Thus, all of the identifiable strategy combinations observed seem to reflect bounded rationality, taking advantage of the task environment to maximize score, while minimizing resources like perception and motor-movement.

Strategies and performance

A 2-factor ANOVA, pattern by level, with score on the 18th trial as the dependent variable, reveals a main effect of pattern ($p < 0.05$), but not level ($p = 0.34$), nor an interaction effect ($p = 0.19$). Bonferroni/Dunn post-hoc analysis reveals that the opportunistic strategy is significantly better than the other pattern strategies ($p < 0.01$), but stacked, sequential and not-identifiable are not significantly different from each other. Upon reflection, high scores are obtained by using both runways as much as possible. This requires having planes available to land on both the short and long runways. Sequential only has one plane available at a time, which may need to wait for the long runway to be free. Stacked usaully has several planes, of different types, available, except at the very end of landing the stack. In contrast, opportunistic always has several planes available and there is almost always a plane suitable for the short runway.

The results also indicate that it is not simply the number of keystrokes required to accept and land planes that accounts

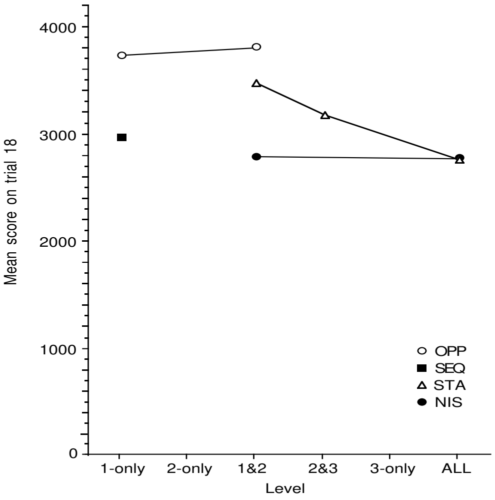

Figure 4: Distribution of the pattern strategies according to the average score and level in trial 18. (Four points representing only one participant in each of OPP-2, OPP-2&3, NIS-3 are not displayed to avoid clutter; they do not contradict the main trends of the diagram).

for the score. Since it takes, on average, 6 more keystrokes to land a plane from level-2 than from level-1 and 12 more keystrokes to land from level-3 than from level-1, we would expect a difference of level. However, as can be seen in Figure 4, a poor choice of pattern-strategy (sequential), though exclusively in level-1, scores poorly. Likewise, participants using the stacked or no identifiable pattern-strategy in levels 1&2 score below those using opportunistic in the same levels. This also provides a possible hint as to why the opportunistic pattern tends to be in the lower levels. The participants using this strategy have evolved to a very efficient pattern and have placed it in the most efficient levels; perhaps they are considering other aspects of the task beyond minimizing keystrokes and easy-to-find empty rows.

Strategy shifts

Although the variety of strategies observed in this simple task is interesting in itself, the shift between strategies is even more challenging for cognitive modeling. Concentrating on the pattern-strategies, we observed that many people shift strategies in the course of 18 trials.

Figure 5 summarizes the of pattern-strategies observed in the 58 participants along the 18 trials. Twenty-seven participants picked a constant identifiable strategy from the beginning, 25 needed several trials before adopting a constant identifiable strategy, and 6 never adopted any constant identifiable strategy. (Periods of no constant identifiable strategy (ncis) are represented in Figure 5 if they occurred at the beginning or the end of the 18 trials. If such a period appeared between two identifiable strategies, we considered that part of a shift from the previous strategy to the new one (see below).) Among the 51 eventually adopting a constant identifiable strategy, 31 shifted pattern 1 time, 6 shifted 2 times, and 1 shifted 3 times[3].

[3]This participant switched from sequential to opportunistic to a period where opportunistic was interspersed at fairly regular intervals with long stacks, and back to pure opportunistic. Although not strictly within our original definitions, this unusual pattern (opp w/stacks) appeared so regular in this participant's timeline that we encoded it as 3 pattern-strategy shifts.

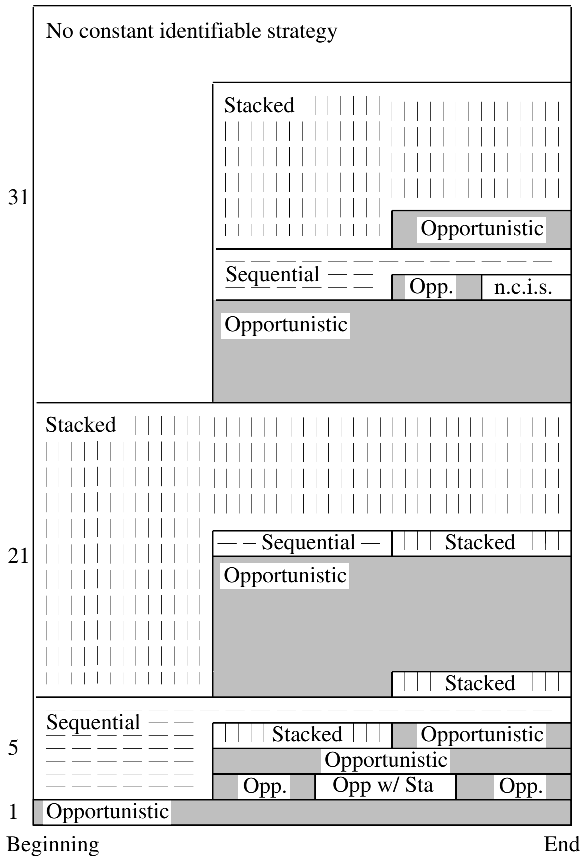

Figure 5: Evolution of the pattern-strategies for the 58 participants between the first and the 18th trial. Number of participants beginning in each pattern-strategy appears at the left; number ending in each pattern-strategy appears at the right. (The horizontal axis does not encode the shifting time, which varies between the participants, but simply that a shift occurred at some time during the 18 trials).

One participant returned to no constant identifiable strategy after 3 shifts. As for the content of the shifts, of the 40 participants who shifted at all, 25 shifted to the opportunistic pattern by the 18th trial. Only two persons who used the opportunistic pattern shifted away from that pattern. Since the opportunistic pattern is associated with the best score, again, the participants seem rational in their approach to improving their performance.

Types of strategy shifts We have identified two types of pattern strategy shifts: *gradual* and *abrupt*. Shifts are gradual when one strategy clearly appears in one portion of the timeline, a different strategy clearly appears in a later portion of the timeline, but there is no clear demarcation between the two. On the other hand, shifts are called *abrupt* when the onset of the new strategy can be pointed to as being at a particular trial and time. Out of the 50 strategy shifts, 39 are gradual and

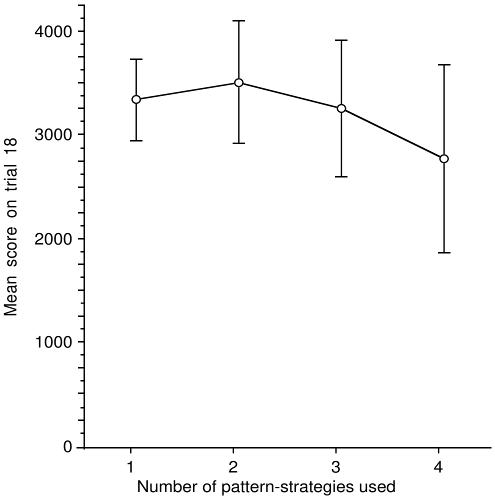

Figure 6: Mean score on the 18th trial vs. number of pattern-strategies used throughout the 18 trials (with standard deviation error bars). The 6 participants never using an identifiable pattern-strategy are excluded from this figure, although the results do not change if they are included.

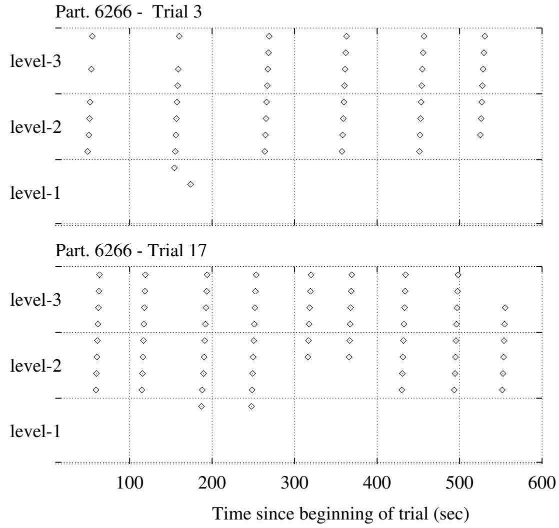

Figure 7: Improvement in speed from trial 3 to trial 17.

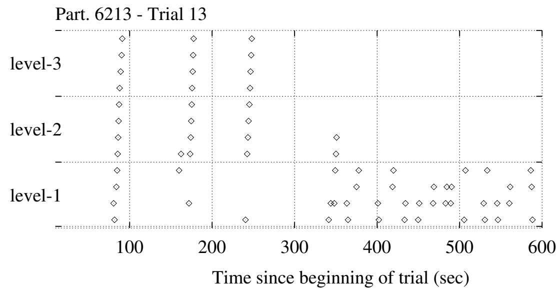

Figure 8: Abrupt shift within trial 13.

11 are abrupt. The implications of the existence of these two kinds of shifts will be examined in the next section.

Strategy shifts and performance Prior research in developmental tasks (Siegler, 1996) indicates that the more strategies children articulate about a task, the better they perform on that task. Although there are too few participants in this study to be confident, the analogous result (that the number of strategies used predicts end-performance) does not seem to be the case. Figure 6 shows a slightly U-shaped curve, and a 2-factor ANOVA (ending pattern-strategy by number of pattern-strategies explored, with score in trial 18 as a dependent variable), indicates that the ending pattern-strategy is highly significant ($p<0.005$) but the number of pattern-strategies explored is not ($p=0.18$), nor is the interaction ($p=0.999$). That is, if the participant ended up using the opportunistic strategy, he or she performed well, no matter how many other strategies were tried first.

Implications for cognitive modeling

The improvement in participants' performance, the variety of strategies observed, and the shifts between strategies, have many implications for cognitive modeling and raise many questions for future investigation.

Improved performance It is well documented that people speed up when they perform any task repeatedly (e.g., Newell, 1981). This speed-up happens as a function of practice, and may be attributed to simple speed-up of cognitive, perceptual and motor actions, or to changes in strategy. The Ackerman (1988) participants do both. For example, Figure 7 shows how a participant's stacking behavior got faster from trial 3 (46 planes brought in from the queue) to trial 17 (68 planes brought in). Figure 8 shows an abrupt shift from stacking in all-levels to opportunistic in level-1 in the middle of trial 13.

Any computational cognitive models of this task will need to produce both types of improvements. This is not to say that cognitive architectures must have several different architecturally-supported learning mechanisms, just that the mechanisms they do have must work together to display these different types of behavior. For instance, ACT-R (Anderson, 1993) has three different architecturally-supported learning mechanisms: strengthening of production and associative links which produces simple speed-up, analogical reasoning which can radically change strategy, and bringing new knowledge from the outside world into declarative memory, which could also radically change strategy. On the other hand, Soar (Newell, 1990) has only one architecturally-supported learning mechanism but that mechanism can interact with different knowledge to produce different learning behavior (Rosenbloom et al., 1993). Either solution is an acceptable path to an architecture of cognition, as long as attention is paid to producing the full range of human learning behavior.

Variety of strategies Although 28 combinations of pattern and level strategies can theoretically appear, 71% of the observed behavior falls into 6 pattern/level strategies: stacked in levels 1&2, 2&3 and ALL, sequential in level-1, and opportunistic in levels 1 and 1&2. Thus, a cognitive architecture which can model these six strategies will account for nearly three-quarters of observed behavior.

Gradual vs. abrupt strategy shifts The existence of gradual strategy-shifts implies than an architecture cannot have a dominant all-or-nothing learning algorithm that functions at the strategy level; it cannot display dogmatic behavior with respect to strategies. Rather, it must have mechanisms for tuning existing strategies and experimenting with new ones. It must be able to gain information about different strategies over time to migrate to a new one. Also, once it has experimented with, or been told, or otherwise discovered a new strategy, it cannot thereafter always choose to use that strategy.

The existence of abrupt strategy changes puts other constraint on a cognitive architecture. Gradual changes point towards tuning or experimentation mechanisms, but abrupt changes in strategy (in the absence of changes in the task environment) point towards reasoning mechanisms that do indeed discover that a new strategy is better (by some measure) and replace the behavior associated with the old strategy.

Abrupt changes can happen between trials (6 changes) or within trials (5 changes) and these have different implications for a cognitive architecture. Intra-trial abrupt strategy changes imply that the reasoning processes are happening while the task is being done, perhaps in some cognitive "slack time" (e.g., the 15 seconds while planes are moving across the runways). Thus, cognitive resources must be able to be applied both to the task-at-hand and to reasoning about the task. This implies simultaneity of thought, or rapid task-switching, or some other mechanism that can display such behavior.

Furthermore, our think-aloud protocols in other dynamic tasks indicate that some people have an awareness of the slack time itself, that they deliberately use that time to think about better ways to do the task. Thus, a cognitive architecture may need the capability to perceive and reason about time itself to fully model human behavior in dynamic tasks.

Inter-trial abrupt strategy changes implies reasoning about the task when not in the task environment. When strategy changes happen after the 5-minute breaks between 30-minute sessions, or on the first trial of the second day of a two-day experiment, it is plausible that the participant has been thinking about the task. If so, a cognitive architecture must have the capability to learn enough about the task, the display, the dynamics, etc. to do this reasoning without the environment in front of it. This may be evidence for the existence of some type of mental simulation in these participants.

Conclusions

This study of strategy use for the KA-ATC© task shows the importance of the strategy choice for the final performance on the task: many participants (26) eventually picked the pattern-strategy giving the best performance, but almost always (in 25 cases) after an exploratory phase involving the use of other pattern-strategies, or no identifiable strategy at all. We also believe that the great variety in strategy use puts some constraints on a cognitive architecture used to model this task.

As future work, we intend to determine whether if the strategy shifts are related to specific events arriving during the task (e.g., a plane crash). We also intend to explore further how level strategies evolve, and how they are connected to pattern strategies. Finally, we will examine how the strategies evolve during the subsequent 9 trials of the study, and other KA-ATC© studies, when bad weather conditions come into play and force the participants to deal with many more rules about landing planes.

Acknowledgements

Many thanks for Phil Ackerman and Ruth Kanfer for making their task and data available through the Office of Naval Research (ONR), and to the ATC group at CMU for their continued contributions to this work. This research is supported by the ONR, Cognitive Science Program, Contract Number N00014-89-J-1975N158. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ONR or the US Government.

References

Ackerman, P. L. (1988). Determinants of individual differences during skill acquisition: cognitive abilities and information processing. *Journal of experimental psychology: general*, (117).

Ackerman, P. L. & Kanfer, R. (1994). Air Traffic Controller Task CD-ROM database, data collection program and playback program. Office of Naval Research, Cognitive Science Program.

Anderson, J. (1993). *Rules of the mind*. Hillsdale, NJ:Lawrence Erlbaum Associates.

Bauer, M. I. & John, B. E. (1995). Modelling time-constrained learning in a highly interactive task. In *Proceedings of CHI*, Denver, Colorado.

Breiman, L. (1994). Bagging predictors. Technical Report 421, Dept. of statistics, University of California, Berkeley.

Klein, G. A., Orasanu, J., Calderwood, R. & C. E. Zsambok (Eds.) (1993). *Decision making in action: models and methods*. Norwood, NJ:Ablex publishing corp.

Lee, F. J., Anderson, J. R., & Matessa, M. P. (1995). Components of dynamic skill acquisition. In Moore, J. D. and Lehman, J. F., editors, *Seventeenth annual conference of the cognitive science society*. Hillsdale, NJ:Lawrence Erlbaum associates.

Murthy, S. K., Kasif, S., & Salzberg, S. (1994). A system for induction of oblique decision trees. *Journal of artificial intelligence research*, 2:1–32.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, Mass:Harvard University press.

Newell, A. & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In Anderson, J. R., editor, *Cognitive skills and their acquisition*, Hillsdale, NJ: Lawrence Erlbaum Associates.

Reder, L. M. (1982). Plausibility judgments vs. fact retrieval: Alternative strategies for sentence verification. In *Psychological Review*, number 89, pages 250–280.

Reder, L. M. (1988). Strategic control of retrieval strategies. In Bower, G., editor, *The psychology of learning and motivation*, volume 22. New York:Academic Press.

Rosenbloom, P. S., Laird, J. E., & Newell, A. (1993). Introduction. In Rosenbloom, P. S., Laird, J. E., and Newell, A., editors, *The Soar papers: research on integrated intelligence*, volume 1, Cambridge, Mass: MIT Press.

Sheridan, T. B. (1987). *Supervisory control*, chapter 9.6. New York: John Wiley & Sons.

Siegler, R. (1996). *Emerging minds: The process of change in children's thinking*. New York: Oxford University Press.