A Situated Cognitive Model of the Routine Evolution of Skills

J. Michelle Moon and Wai-Tat Fu
Human Factors Division
University of Illinois at Urbana-Champaign

This paper presents a cognitive model for the routine evolution of skills in a photocopying task. In the classic study by Agre and Shrager (1990), it was found that although speedup in performance conformed to the power law of practice, microgenetic activities of the subject showed multiple stages of qualitative changes as the subject repeated the same task. The evolution of skills was shown to be more appropriately characterized by a "situated", dialectic relationship between the subject and the environment than by a set of general mental mechanisms as in traditional theories of skill acquisition. This paper attempts to bridge the gap between these two perspectives of skill acquisition by showing how the dialectic relationship can be characterized by a traditional information processing model by making a few additional assumptions of the relationship between the human and the environment. A model based on the ACT-R architecture was constructed, which provided good fits, both qualitatively and quantitatively, to the data collected by Agre and Shrager. The model demonstrated how multiple mechanisms could be implemented in a cognitive architecture to explain both qualitative and quantitative changes at both the macro and micro levels of analysis. Implications to human factors research are discussed.

INTRODUCTION

One of the most robust findings in human factors and psychology is the effect of practice on performance. A substantial literature shows that when subjects are asked to perform a task repeatedly, the time to perform a task will decrease with the amount of practice according to a power function. Many cognitive accounts have been proposed for this "power law of practice". Many researchers suggest that the speedup can be explained by faster processing of information (e.g., Shiffrin & Schneider, 1977), the substitution of algorithmic processing by direct retrieval of operators (e.g., Newell & Rosenbloom, 1981; Logan, 1988), strengthening or concatenation of procedural subunits (e.g., Anderson, 1983), or the use of better strategies across time (Anzai & Simon, 1979; Delaney, Reder, Staszewski, & Ritter, 1998).

There are two major categories of complaints to these theories of skill acquisition. First, it is argued that speedup in performance cannot be cast as a unitary phenomenon that can be explained by one or a few general mechanisms. Instead, improvement is likely composed of a set of complex changes that interact with each other over time. Besides, these approaches are sometimes criticized for running the danger of suppressing the qualitative changes of activities by just showing how the general mechanisms may explain the speedup in aggregated data.

The second complaint is that speedup in performance is likely dependent on the larger context under which the activities occur, and that behavior cannot be explained by "decontextualizing" actions from the environment. Specifically, these "situativity" theories (e.g., Suchman, 1987) claim that control of behavior is mostly driven by the specific circumstances or environmental characteristics that the person is situated at, and a pure mental account of performance speedup is inadequate in addressing how the systems or environments should be improved to assist human activities.

Despite the fact that a lot of studies have been done in skill acquisition, there has been little attempt to model the detailed dialectic skill development of an individual in a natural setting. The current model is a meaningful attempt to model routine skill acquisition not only with its quantitative nature but also its qualitative, interactive aspects. The main goal of this paper is to show what additional assumptions are needed to construct cognitive models of routine skill acquisition in a natural setting. First, we will show that multiple general mechanisms can be implemented in a cognitive architecture to characterize the complex changes of behavior at both the micro and macro levels. Second, we will show that the situational aspects of skill acquisition can be approximated by assuming certain characteristics of the connections between human perception and cognition in a routine task. Specifically, we will demonstrate the importance of including some "situated" measures between the human and the environment, as well as the use of multiple learning mechanisms to account for the routine evolution of skills, and how both qualitative and quantitative changes may contribute to the overall performance speedup.

ROUTINE EVOLUTION OF SKILL

We have chosen a classic study by Agre & Shrager (1990). In their study, a subject was shown to a room with Xerox 1075 copier (Figure 1) and was given four types of copying tasks. One of those tasks was to make three copies of a paper which was 17 pages in length. During this photocopying task, the subject was videotaped by a camera mounted in the ceiling with microphone hidden near the copier. This setup was intended to minutely document and track the subject's activities without her self-conscious optimization of performance. Observation of subject's behavior throughout four minutes revealed speedup in performance over time corresponding to the traditional practice effect. Microgenetic analysis adopted to study changing process of activity across time showed that speedup of performance was clearly more than a unitary phenomenon claimed in traditional theories of skill acquisition. Instead, behavior was found to evolve through a series of complex activities involving both the person and the environment.

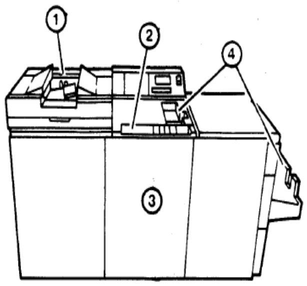

Figure 1. The Xerox 1075 photocopier used in the study: 1. Recirculating Document Handler (RDH), 2. Control Panel, 3. Copier Body, 4. Output Tray (adapted from Agre & Shrager).

Two major phenomena were observed as the subject repeated the task. At the macro level, the subject's performance showed a graph close to classical power law of practice: Performance time decreased as the same task was repeated over time. At the micro level, it was found that the changes in performance times could be attributed to a number of qualitative and quantitative changes in behavior. Some actions dropped out and some actions were merged with other actions over the session. It was also notable that the subject's approach to the copier was initially circumspect as characterized by cautious behaviors such as readjusting controls, looking around, or checking the output. Over repeated cycles, the subject settled down to a more stable and regular routine patterns of actions. For example, the subject initially checked the output of copy, but stopped checking it after a couple of sessions. Other circumspect behaviors such as readjusting controls or looking around observed in the earlier stages did not present in the later part. The subject also learned to expect the third flash of light from the copier at the end of each trial, which is deemed as an adaptation to the responses made by the copier. Another prominent aspect was the changing patterns of hand-over-hand moves, which became an evident example of situated routine evolution of skills. We will elaborate on these changes as well as how they were explained by the model in the next section.

THE MODEL AND THE DATA

Macro Changes: Speedup of Performance

Figure 2 shows the improvement in performance times as the subject repeatedly copied the pages. The subject initially took 53s to copy the first two pages, but it quickly dropped to 25s in the third session and became relatively stable from there. The speedup approximated a power law relationship. The model provided good fit to the original data (R^2=0.99). However, instead of using a single mechanism to produce this power curve of learning, we combined a set of situated assumptions of actions with mechanisms that better match the qualitative changes of activities as observed by Agre and Shrager. We will next describe a number of contributing factors and mechanisms to speedup in the first three sessions.

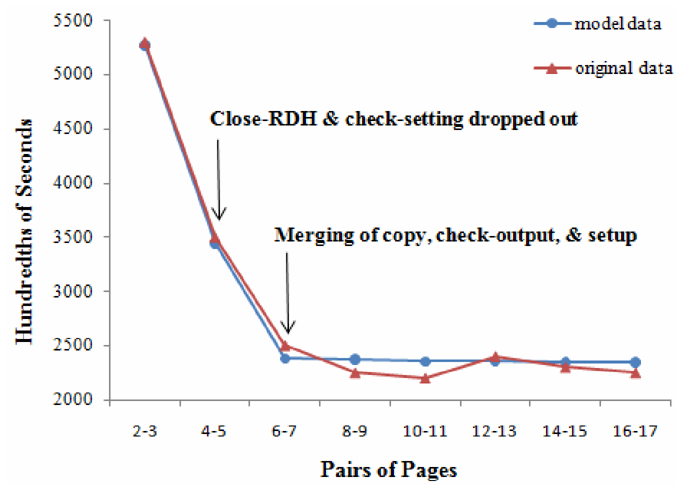

Figure 2. The elapsed time for each successive pair of copies (original and model data).

From the first to the second session. Agre and Shrager reported that starting from the second session (Figure 3), the subject stopped closing the Recirculating Document Handler (RDH) cover. Although checking the settings of the copier was not explicitly mentioned in Agre and Shrager, they reported that the subject engaged in markedly circumspect behaviors such as carefully settling the book into place, readjusting the controls, and looking around in the first session. In the model, these circumspect behaviors were represented in a single discrete production called check-setting which fired after close-RDH production and assumed to disappear over time. This reflected the observation that the subject was in a "cautious mode" when she first approached the copier, but it was quickly changed starting from the second session.

The model explained this by assigning a high initial utility and low reward to the close-RDH and check-setting productions. These two productions fired in the first session because of their high utility. However, because these two actions had high time costs (approx. 6.4s and 12s respectively) and low rewards, they stopped firing in the second session. When the close-RDH and check-setting productions were dropped, total performance times dropped by 18.4 (6.4+12) seconds.

From the second to the third session. The subject initially walked to the output tray of the copier to check the output copy. In later sessions, perhaps satisfied with the settings of the copier, the subject stopped checking the output. This was again modeled by assigning a high initial utility and a low reward to the check-output production. Since it was not clear how long it took for the subject to check output and when exactly this activity started or ended, we assumed that checking output initially was a discrete activity over a relatively long time but eventually it quickened and merged with other activities. For instance, it was likely that the subject was able to finish checking output during the three flashes of light from

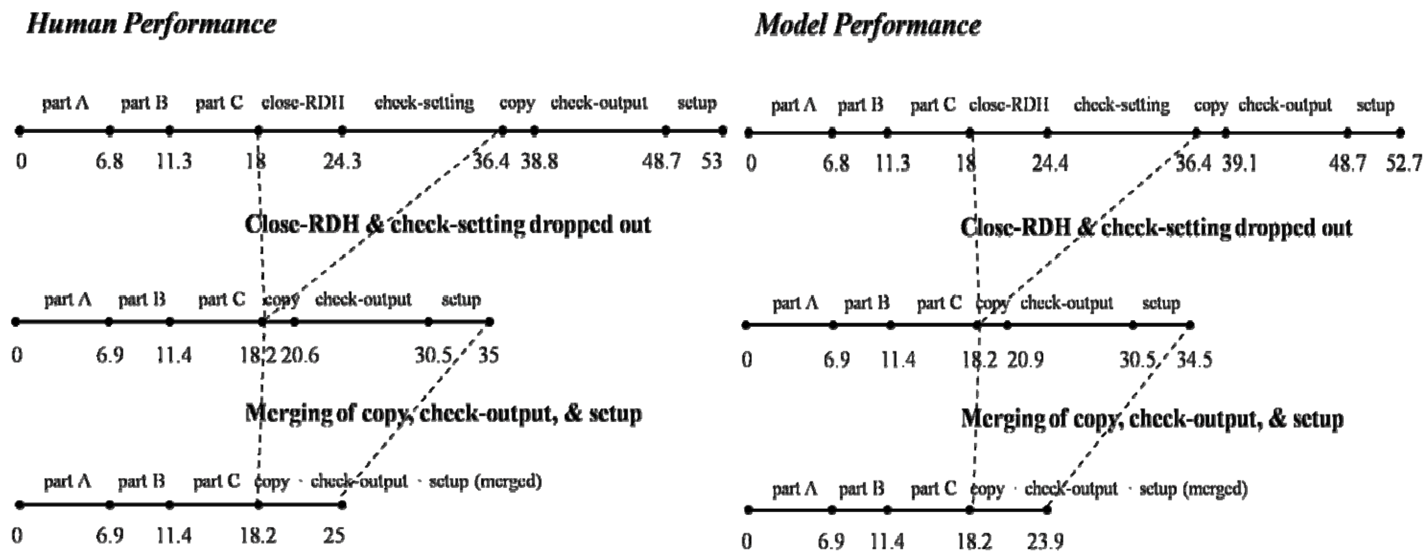

Figure 3. Human and model performance (in seconds) in the first three sessions (from top to bottom).

the copier as the anticipation of the third flash rendered free time for other activities to fill in. After the drop-out of check-output, performance times decreased by 9.6 seconds. Drop-out of checking output activity also corresponded to transition from cautious to regular mode.

In addition to the change from a cautious to a confident mode of operation, the subject also learned to anticipate certain responses from the copier. The subject learned to anticipate the third flash from the copier as an indication that the copying was about to finish. Thus, instead of waiting for the third flash, the subject started to prepare for the next copy after the second flash. This represented a merging of activities as the subject gained knowledge about the machine and learned to anticipate event sequences.

The model characterized this function merging of activities by first observing the output of its own actions, and then forming declarative representations (chunks) of these observations. These chunks were later retrieved to allow the model to anticipate the future event. Specifically, as the model pressed the "copy" button and observed a flash every 0.8s, a chunk was created to represent the events relating to what would happen to the action of pressing the button, and what action should be performed next after the event. The model would actively anticipate the next event by trying to retrieve from experience what would happen next, and what action should be performed after an event had occurred. Initially these newly created chunks were not strong enough to be retrieved. But each occurrence of the same events would reinforce the activations of the chunks (see appendix for the base-level learning equation). As a result, starting from the third session, the new chunk was strong enough to be retrieved. With the retrieval of this new chunk, the setup production fired right after the second-flash production fired. The chunk created from experience thus served as a short-cut bridging between the encoding of the second flash and the setup for the next session without waiting for the third flash. This merging of the copy and setup activity and the parallel execution of check-output saved 10.6 (1+9.6) seconds by reducing the total time devoted to the encoding of the flashes of light.

Micro Changes: Change of hands

Part A: Unsystematic Change. In part A (Figure 4), the subject lifted the book with the hand closer to right edge of the book and flipped it face up with assistance of the other hand. Part A initially showed quite stable RLR (Right-Left-Right) pattern, but LLR (Left-Left-Right) pattern came into place in pages following 11 and 15.The choice of hand in these two sessions was mainly affected by the side effect of previous activities. The copier flashed three times before the start of part A since she was making three copies. The third flash of light from the copier during copying served as the important marker in organizing further actions. The subject learned to anticipate the third flash and put her hands on the book so that she could start the next session right after the flash. She also engaged in other activities during the slack time of flashes. She walked to the right of the copier perhaps to check the output of copy during the flashes for pages 5, 7, and 9, and took a short break during the flashes for page 13. These activities changed her physical relationship to the copier, which in turn determined which hand is closer to the book edge. Though the third flash of light contributed to at least partially systematic change, a number of other factors such as concerns about the output or physical position influenced part A resulting the initially stable RLR pattern break down to LLR pattern, which illustrates the "situated" nature of action selection.

To model changes in part A, a number of situated assumptions were made. First, distance between each of hands and the edge of the book was deemed to serve as an important factor in the decision of hands. We identified that in most cases the subject used the hand whichever closer to the book than the other hand. The distance in turn was determined by the subject's lateral position to the copier and previous actions affecting the position. For example, activities such as moving to the other side of copier to see the output changed the distance and made one hand closer to the book than the other one. In the model, new distance values were given when the subject changed her position so that the model could make the situated decision of which hand to use.

Last Page Copied

Functions	1	3	5	7	9	11	13	15	
lift to flip face up	R	R	R	R	R	L	R	L	Part A
flip book face up	L	L	L	L	L	L	L	L	
flip assist	R	R	R	R	R	R	R	R	
lift page	L	R	L	R	R	R	R	R	part B
turn page	L	R	L	L	L	L	L	L	
lift to flip face down	R	R	R	R	R	R	R	R	part C
flip book face down	L	L	L	L	L	L	L	L	
flip assist	R	R	R	R	R	R	R	R	

Figure 4. Which hand was used in each of eight functions in the page turning activity over the session (adapted from Agre & Shrager).

In the original ACT-R's utility learning mechanism, time is the only measure for cost of action, and the model has to experience the action (i.e., fire the production) in order to learn the cost of action to influence future selection. In our model, we assumed that judgment of distance is well practiced and can be "automatically" inferred through perception and incorporated as a variable in the utility equation to influence the choice of productions. Incorporating distances as costs in the utility equation therefore allowed the model to be situated in the environment.

Besides distance, physical efforts to bring one hand to the desired place also seemed to determine the choice of hands. We simulated this factor by assigning higher reward to the production that presumably minimizes physical efforts of moving hands. It was also assumed that, this perceived effort is automatic and can be readily inferred from the environment as the person is situated in a natural environment.

Part B: Systematic Change. In part B, the subject lifted and turned the right-hand page of the book lying down on the copier. Part B showed systematically evolving pattern from either LL or RR to stabilized RL pattern. The subject used the same hand to lift and turn the page with guidance of the other hand in the first three sessions. In later sessions with stable RL pattern, the subject moved both hands to the bottom of the right hand page, used the right hand to lift the page, and used the left hand to turn it over. Right after the completion of part A, page turning activity in part B started from physically standardized setting with both of her hands on the book. As the subject incrementally adjusted the position of her hands along the edge of the book, the activity later evolved into settled routine pattern forming the standardized physical setting between the subject and the environment.

To model changes in part B, the availability of hands was assumed to be the major controlling variable for the choice of hands. Throughout the whole sequence of actions, there was a clear tendency that the subject alternated between different hands for each function and used the free hand that was not engaged in the previous action. This was especially true when the function was to assist the other main function, or the successive functions were forming one set of action. For example, if right hand was used to flip the book, left hand was selected to assist the flipping. Although this alternation of hands was not significant in the first three sessions, it seemed to finally settle down to a stable pattern in later sessions.

In the model, lower initial utility and higher reward were given to the production for selecting different hand in lifting and flipping function. This was based on the assumption that people are accustomed to alternate hands as they manipulate objects. In the first three sessions in part B, the model fired the productions for the same hand reflecting higher noise value in the selection. The noise factor was set to decrease over sessions to the level that the model settled into consistent firing of production for different hand. Slowly decreasing noise was consistent with the finding that people tend to be more exploratory when they first approach unfamiliar device but this randomness tends to decrease as actions become more practiced over time.

Part C: Settled Pattern without Change. In part C, the subject lifted the book again and flipped it face down to start copying. Part C maintained highly stable RLR pattern without evolution throughout the sessions. Despite the unsystematic changes in part A, part B provided settled physical setting rendering part C start in very standardized way, which in turn led to settled choice of hand. One of the possible reasons pointed out to explain the difference between systematic change in part B and lack of change in part C was that activity in part C had become already highly evolved since it corresponded to the task the subject had to perform throughout the copying session. On the other hand, activity in part B entailed a number of physical settings that called for different operation of hands from the previous tasks.

CONCLUSIONS

We have shown how situated assumptions can be combined with a set of general mechanisms of ACT-R to simulate the qualitative and quantitative changes at both the macro and micro levels as observed by Agre and Shrager during the routine evolution of skills in a natural environment. At this point, our goal is of course not to claim that these assumptions can be directly applied to other situations. Quite the contrary, the goal of this modeling effort was precisely to show, in principle, what are the major assumptions needed to reconcile these seemingly distinct "worldviews" of skill acquisition. The idea is similar to the argument made by Vera and Simon (1993), who stated that extra assumptions are all that are needed to make the connection between symbolic representations and situativity theories, and the intricate interactions between the human and the environment can only be understood if characteristics of both the human head and the environment are given equal weight in the model.

One specific situated assumption we found necessary was that utilities of actions could come directly from perception, such that the decisions on which hand to use depended critically on where the hands were before, and the least-effort principle still seemed to apply in the decisions. This was in contrast with the current ACT-R utility learning mechanism, which assumes that actions were chosen independent of contexts (e.g., if using the right hand was better then its likelihood of use in the future will be higher regardless of situations). We showed that the direct perception of effort was important in characterizing action selection in a natural setting. It is therefore more realistic to include mechanisms that allow ACT-R to directly infer utilities of actions from direct perception. It seems possible that, for example, distance between the hand and the object to be manipulated could be directly perceived and thus its utilities dynamically updated based on information received from perception.

On the practical side, a clearer understanding of the dialectic relationship between the human and the environment is of essential importance to design and performance predictions, as changes in the design will always entail changes in the various variables that affect the coupling of the human operators and the system. These human-environment specific variables are important in the development of human performance models that predict how skills may evolve with time.

REFERENCES

Agre, P. E., & Shrager, J. (1990). Routine evolution as the microgenetic basis of skill acquisition. *Proceedings of the 12th Annual Conference of the Cognitive Science Society,* Cambridge, MA.

Anderson, J. R. (1983). *The Architecture of Cognition.* Cambridge, MA: Harvard University Press.

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

Anzai, Y., & Simon, H. (1979). The theory of learning by doing. *Psychological Review, 86* (2), 124-140.

Delaney, P. F., Reder, L. M., Staszewski, J. J., & Ritter F. E. (1998). The strategy specific nature of improvement: The power law applies by strategy within task. *Psychological Science* 9(1), 1-7.

Fu, W.-T., Anderson, J. R. (2006), From recurrent choice to skilled learning: A reinforcement learning model. *Journal of Experimental Psychology: General, 135*(2), 184-206.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.

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

Shriffrin, R., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84*, 127-190.

Suchman, L. (1987). *Plans and Situated Actions: The Problem of Human-machine Communication.* Cambridge University Press, MA: Boston.

Vera, A. & Simon, H. (1993). Situated action: A symbolic interpretation, *Cognitive Science, 17*, 7-48.

APPENDIX: ACT-R

ACT-R is a cognitive architecture (Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004) that tries to explain cognitive processes of human being by developing a computational model based upon assumptions from psychological experiments. Chunks, elements of declarative knowledge, are basic units of knowledge that a person is expected to have to solve a problem or perform a task. Productions are if-then rules that execute actions when they fire.

Utility Learning

In ACT-R, each production has utility value and random noise associated with it. Execution of production is serial, which forces only one production to fire at a given moment. When multiple productions compete to fire, a production with the highest utility is selected to fire. The probability of choosing production i when a number of productions competing with expected utility values U_j is computed by the following equation.

$$\text{Probability}(i) = \frac{e^{U_i/\sqrt{2}s}}{\sum_j e^{U_j/\sqrt{2}s}}$$

Utility $U_i(n)$ after nth application of reward is computed by the following utility learning equation when α is the learning rate.

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

$R_i(n)$, the reward that a production receives for its nth application, is determined by the external reward received and the time from the production's selection to the reward. The reward can work as either cost or gain. A production with higher benefit and lower cost is more likely to fire.

Utility learning enables model to decide which production is optimal among multiple productions that are available at given time. It is one of the generally employed mechanisms in explaining speedup of performance (Fu & Anderson, 2006).

Declarative Learning

Each chunk in declarative memory has an activation value associated with it. If there are multiple number of chunks which match the request, the one with the highest activation value above threshold value will be retrieved. If no matching chunk has an activation above the threshold, request of retrieval results in failure.

Base-level activation reflects the recency and frequency of use of the chunk. If :ol (optimized learning parameter) is set to **t,** base-level activation is computed by the following equation when *n* is the number of presentations of chunk *i, L* is the lifetime of chunk *i, d* is the decay parameter, and βi is a constant offset determined by :blc or the chunk's :base-level.

$$B_i = \ln(n/(1-d)) - d * \ln(L) + \beta_i$$