PEER-MEDIATED LEAP TO EFFICIENCY: COST-BENEFIT ANALYSIS IN THE SELECTION OF EFFICIENT STRATEGIES.

Franklin P. Tamborello, II
Rice University
Houston, Texas

S. Camille Peres
University of Houston,
Clear Lake
Houston, Texas

Michael D. Fleetwood
Adaptive Cognitive Systems
Boulder, Colorado

There is ample evidence that computer users often do not progress from novice to expert levels of performance, particularly when efficiency is included in the definition of performance. This paper describes a theoretical argument that one of the pieces missing in the understanding of this process is an accurate assessment of how people calculate a cost/benefit analysis (CBA) of learning and using new techniques and strategies. While there are many explanations for why people often do not use accurate methods of calculating this ratio, there is little discussion describing why some people do. We suggest that an important and predictable influence on whether an individual uses an accurate CBA is the observation of others using efficient techniques. We propose that with a more full understanding of the CBA calculation process, it will be possible to predict when users will and will not utilize a more efficient technique.

Expertise and the achievement of expert performance is a broad field. Although much of the previous work in expertise has provided important descriptions of expertise, the findings are equivocal regarding the processes associated with achieving this level of performance. The process of obtaining expert performance with software programs is no exception, particularly when efficiency is included in the definition of expert performance (Bhavnani et al., 1996; Carroll & Rosson, 1987). This is important as one of the hallmarks of expert levels of performance is often efficient performance, i.e., performing tasks with the least amount of effort and time necessary. Although there is evidence that people can initiate using efficient techniques with software (Charman & Howes, 2003) there is considerable evidence that given no guidance, they do not (Lane, Napier, Peres, & Sándor, 2005).

The goal of this paper is to offer a theoretical argument that one of the pieces missing in the understanding of the mechanisms associated with the utilization of efficient techniques is an accurate assessment of how people calculate the cost/benefit analysis (CBA) of learning and/or using particular techniques with software. To describe our argument, we will first illustrate that much of the previous literature has either implied or overtly discussed that CBA is inherently part of the process of acquiring new procedures with software. We will further present that although previous work has mentioned CBA, there has been little to no work to overtly investigate or test the mechanisms associated with how CBA occurs. Finally, we will present the beginnings of a plan of research that is designed to investigate explicitly how people perform CBA and most specifically how observing peers influences this analysis. An important note is that we are building on top of and explaining current evidence and theoretical perspectives. Our position is not necessarily in contrast to other theories, but is instead accounting for holes in the current literature.

CBA is discussed in the GOMS (Goals Operators, Methods, and Selection rules) family of techniques and associated tools, which are widely used to evaluate the efficiency of users of different systems (e.g., John & Kieras, 1994). GOMS tools/techniques are employed to predict and describe the interaction of users and systems in terms of time,

strategy, and possibly errors. Selection rules specify which Method should be used to satisfy a given Goal, based on the context of the task. An assumption of Selection rules is that users select the most efficient method to complete the goal. However, it is often unclear what Selection rules people are using. Further, users do not always select the most efficient method available to them despite the assumption in the selection rule aspect of the GOMS methodology that people are performing a cost/benefit analysis, i.e. people are evaluating the relative efficiency of the different methods.

An area where CBA is implied, although not directly discussed, is in Carrol and Rosson's "production paradox" (1987). This is the notion that people generally do not want to stop their work to learn new or more efficient techniques even though it may be beneficial in the end, implying that users weigh the benefits of acquiring new knowledge against the costs to stopping their work. Some people nevertheless are motivated to learn efficient methods and do take the time to do so (Peres, Tamborello, Fleetwood, Chung, & Paige-Smith, 2004). Those people who learn efficient methods see something beneficial in acquiring more knowledge than the people who do not learn. However, it remains an open question regarding how much benefit a person must perceive before considering learning something new. An investigation of the mechanics of CBA would allow for the explanation and prediction of what is necessary for users to become efficient.

Fu and Gray (2004) attempt to resolve the production paradox and suggest that people are cognitive misers by nature, so they tend to: 1) use methods and strategies that have cues and offer immediate feedback on the user's progress through the problem space; and 2) that apply in multiple situations. For instance, previous work has shown that it is faster to use the keyboard to issue commands (KICs), yet people tend to use other, slower, methods for issuing equivalent commands (Lane et al., 2005). In the alternate methods, menus and icon toolbars, there is a general procedure (move the mouse to locate/search for the command) and interaction (the computer shows you commands and you see the mouse cursor on the command as you issue it). With few exceptions, KICs tend to be specific and do not offer incremental feedback. People may develop inaccurate CBA

calculations because they overvalue the generalizability of procedures and incremental feedback and undervalue the efficiency that can be achieved with the use of more efficient work strategies.

Bhavnani and his colleagues suggested that people are inefficient with software because they do not know how to integrate efficient methods of completing a task into their current knowledge of the program (Bhavnani & John, 1997). For instance, users may know that objects can be changed and that more than one object can be selected at a time, but they may not understand that it is possible to select multiple objects and then change all of those objects simultaneously. To state the problem in GOMS's terms, people may not possess the selection rules that allow them to operate efficiently, even though they possess knowledge of the methods. As a remediation for this, Bhavnani and colleagues developed a course that explicitly taught students how to employ more efficient methods of using software programs. The course has been successful at increasing individuals' use of efficient methods across several programs (Thomas & Foster, 2001), but the authors did not specify what change in the users' selection rules might have taken place to move them from inefficient usage to efficient usage. Similar to the work done by Carrol and Rosson, the behavior (and even the remediation) have been well described with Bhavnani's work, but the cognitive processes or mechanics that would allow for the prediction of movement from inefficient to efficient performance has not been outlined. It is conceivable that through the course that explicitly taught efficient methods, users' CBA were influenced sufficiently to create selection rules that resulted in more efficient performance. Once the mechanisms associated with the development of efficient selection rules have been mapped, it may be possible to eventually apply them in predictive models that can instantiate a theory of what is sufficient and necessary for someone to make the transition from inefficient strategy selection to efficient strategy selection.

One mechanism that has been associated with increased efficiency is the observation of efficient users, often termed peer learning. Evidence for learning from peers has been obtained in several studies (Bhavnani et al., 1996) including some recent work (Peres et al., 2004) finding that people who use keyboard shortcuts are more likely to work with and/or around others who use KICs. Given these findings, it is conceivable that observing others using efficient techniques may affect a user's CBA and thus influence their selection rules. If the findings that observing others influences a user's selection rules can be replicated and experimentally manipulated, it could provide important insight into the mechanisms associated with CBA and selection rules. One possible mechanism could be that when a person observes another person using a new specific technique, the observer learns the new strategy. Alternatively, observing others may reinforce previously known but perhaps forgotten knowledge. Another facet may be that when people observe others using efficient techniques, the learning process occurs just-in-time, i.e. just when obtaining/reinforcing such knowledge is most salient to the user's current goals. If this is the case, then the timing of the knowledge gained may be as important as the conveying of the knowledge itself.

Reinforcing knowledge, particularly knowledge about the costs and benefits of behaviors, appears to play an important role when people change their behaviors. Indeed Janis and Mann (1977) found that the weight people give to the benefits of a behavior often influences their adoption of that behavior more so than the weight assigned to the costs. This approach suggests that in order to change their behavior, people may need to raise their estimate of the benefits of engaging in that behavior. Our study investigating the weightings of the pros and the cons as a function of the use of KICs found that people who were more likely to report using KICs also weighed the pros more than the cons (Peres, Fleetwood, Yang, Tamborello, & Paige-Smith, 2005). The peer learning paradigm described previously may be an important way this shift in relative weighting of pros and cons occurs. To test our theoretical position that observing and learning efficient techniques from peers causes a unique and substantial adjustment of the CBA, we are currently investigating how observing others affects people's weighting of the pros and cons of using efficient techniques. We expect that after observing others using efficient techniques, those who do not normally use efficient techniques will adjust their weighting of pros and cons to more closely mirror those who normally do use efficient techniques. Furthermore, we will experimentally manipulate individuals' observation of KICs to determine whether a causal relationship exists between the observation of and subsequent use of these techniques.

We will also employ cognitive modeling as methodology for evaluating hypotheses about how knowledge of efficient techniques, such as KICs, can be acquired by users and put into practice. Cognitive architectures such as ACT-R (Anderson & Lebiere, 1998) force the strategy selection aspect of a model of human behavior to be explicitly accounted for. In doing so one can elucidate the complex interaction between learning, memory, and strategy selection on a particular task with a given interface.

We plan to use ACT-R to provide a sufficiency proof that the just-in-time hypothesis would produce the behavior patterns seen in people. ACT-R is a computational cognitive architecture which takes as inputs knowledge (both procedural and declarative about how to do the task of interest) and a simulated environment in which to run. It contains a variety of computational mechanisms and the ultimate output of the model is a time stamped series of behaviors including individual attention shifts and saccades, speech output, button presses, and the like.

In order to model such a process as selection rule learning in ACT-R, we must take advantage of several of ACT-R's learning mechanisms, e.g. production compilation, instance-based learning, production utility learning, base-level learning, and a newly developed mechanism, production induction (Best, 2006). Other than production induction, these learning mechanisms are well-established and have been applied in a variety contexts and educational applications, such as list memory (e.g. Altmann, 2000), skill acquisition (e.g. Taatgen and Lee, 2003), category learning (e.g. Anderson and Betz,

2001), and the development of intelligent tutoring systems (e.g. Anderson and Gluck, 2001).

One of the benefits of embodying a theory in a computational architecture, such as ACT-R, is that it allows researchers to develop and test concrete, quantitative hypotheses and it forces the theorist to make virtually all assumptions explicit. To the extent that the model is able to simulate human-like performance, the model provides a sufficiency proof of the theory. As a first step towards embodying our theory of efficient strategy selection, the following section describes an outline for the ACT-R model that we are proposing to build.

A model of technique learning and selection will operate at two distinct levels, learning of new techniques and the application of those techniques in practice. In the case of learning KICs, we believe that the observation of others is an important method of learning, and hence, the learning stage of our model will focus on this type of observational learning.

During the learning stage of the process (Figure 1), as the model observes an "actor" (be it another model or a human interacting with a system) performing some task, it will follow along with the task by going through the same process (in ACT-R terms, the same sequence of "productions" will fire). When the actor demonstrates a new method of performing a task for which the model has no corresponding process, the model will create the declarative and procedural knowledge corresponding to the new method. In ACT-R, this corresponds to creating a new "chunk" (e.g. the shortcut for pasting text is to push the "control" and the "v" keys), and a new production to recall that process (chunk) the next time the same task is encountered (a process known as "production induction" in ACT-R).

Using production induction in this manner implies an important benefit of learning via observation. In the just-described scenario, the new procedural knowledge corresponding to how to use a new method of completing a task (e.g. use shortcut "control + v") is inserted into the larger decision process just at the moment when the need to apply a method for completing the task is encountered. Hence, the next time the model encounters the need to paste text, and must determine a method for completing the task, there are two methods associated with that need. In contrast, if the model, or a person, were told of a shortcut either well-before or well-after it was needed, that knowledge would not necessarily be inserted into the model's decision structure at
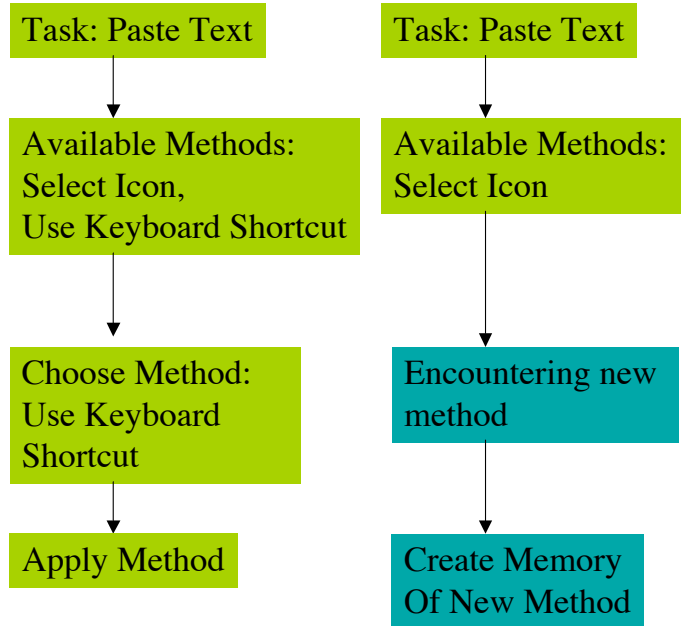

Figure 1. The Learning stage of the proposed model.

the exact moment where it was needed. Interestingly, this lack of timeliness in learning efficient software techniques may be why many people fail to learn them as discussed previously.

At the action stage of the process, when the model is required to perform the task, there now is an additional method of completing the task that it learned from observation. As new methods of performing the task are learned, there may be multiple methods available to the model to perform the task (Figure 2). These methods may compete in ACT-R's "conflict resolution" process. Which method is chosen is based on the time cost of the method, with the method with the lowest time cost generally being chosen (some noise in the system lends stochasticity to the conflict resolution process).

The time cost associated with a method is calculated as the time from when the model determines that its next goal is to paste text to when that goal is successfully completed. The

When the model has multiple methods For completing a task, they compete for selection.
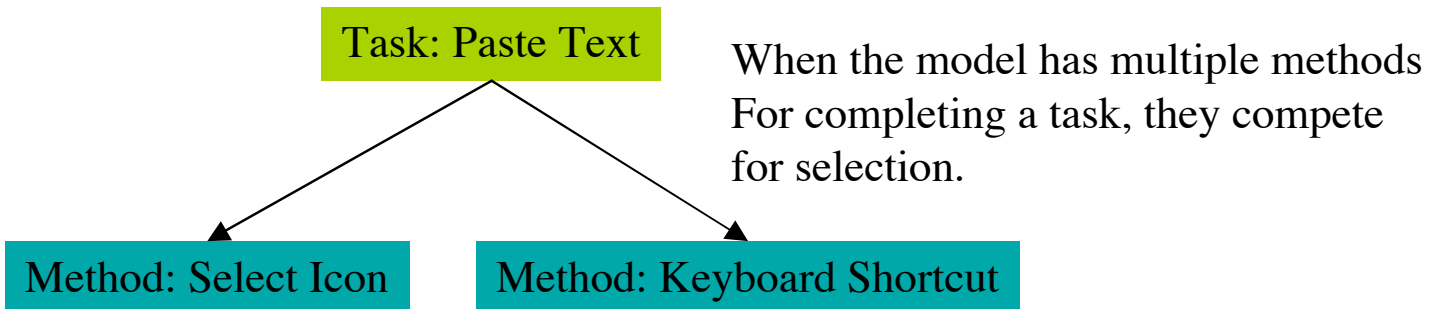
Figure 2. When multiple methods are available to complete a task.

first time a method is observed the time cost associated with that method is the observed time, i.e. the time it took the actor to complete the method. This provides a high likelihood that the method will be tried at least one other time if the observed method is faster than an older, competing method. Successive uses of the method will then contribute to an average time cost associated with that method. Hence, if the method is indeed faster than alternative methods, it should become the preferred method of completing the task. It should be noted that time cost may not be the only "cost" associated with learning a new method. There may also be other cognitive or physical costs associated with learning new methods. However, time cost provides a simple, quantifiable metric on which to base our first version of a model of efficient strategy selection and learning via observation.

With repeated use, a method can be further streamlined via ACT-R's production compilation mechanism, thereby increasing its efficiency and its likelihood of being selected in the conflict resolution process (Figure 3). For instance, a newly observed method for performing a task via a keyboard shortcut will require the model to recall the shortcut from declarative memory before performing the shortcut. Through production compilation it is possible that if the method is used repeatedly, then the step of retrieving the shortcut from long-term memory may be skipped, and once the need to perform the shortcut is recognized by the model, it may quickly perform the required key presses. It is important to note that the model, as described, is theoretical at this point. The description provided here is the basis for work currently underway to implement it in ACT-R.

To the extent that an ACT-R modeling effort is successful in developing an experimental paradigm that can elicit some of the inefficient behaviors reported in the literature reviewed above, that paradigm will be used to collect data from human users. Once an experimental manipulation is developed that can reliably influence the efficiency of user behavior, that paradigm will then be adapted for use in an fMRI study. Brain imaging will allow for the collection of converging evidence, such as perhaps differing degrees of blood oxygen level dependent response (BOLD) in the anterior cingulate cortex, which previously has been implicated in action monitoring and error detection (Magno, Foxe, Molholm, Robertson, & Garavan, 2006). It is possible that people who perform inefficiently never compute an error signal to the effect that they could be performing better, i.e. more efficiently.

Furthermore, since work performed by Anderson and colleagues (2004) has already provided strong evidence for functional relationships between ACT-R's modules and human brain regions, it is conceivable that a strong theory of efficient human performance could be built using converging behavioral, modeling, and imaging data.

In conclusion, we acknowledge that other mechanisms could be associated with the development of efficient selection rules. For example, it could be that the cognitive cost associated with remembering to use a new technique has more influence on the selection rule than the benefit of a more efficient technique. Another mechanism that may be associated with this is how well-learned the inefficient technique is that the efficient technique would be replacing. We submit that there is sufficient evidence to support our position that people often utilize faulty CBA when using software programs. We further submit that this is reflected in faulty selection rules as described by the GOMS modeling paradigm. Thus, the mechanism of peer-learning is the first mechanism we plan to investigate and with further research, we plan to incorporate other potential mechanisms that facilitate, and in fact, predict, the transition from inefficient to efficient use of software.

Thus, the essence of our theoretical position is that much of the findings in the literature to date on the acquisition of expert performance with software (or rather the lack thereof) may be more fully explained with the incorporation of how the findings relate and predict a user's CBA. From Carroll and Rosson's seminal work on the Paradox of the Active user (1987) to Gray and Fu's (2004) most recent work on the selection strategies associated with sub-optimal selection rules, the results strongly suggest that a sub-optimal analysis of the costs and benefits of a particular behavior is occurring. The recurring theme of improved levels of performance being associated with peer learning suggest that there are cognitive mechanisms involved in this type of learning that are not fully understood. As we develop a theory of how people transition from novice to expert levels of performance, we hope to encourage computer users to select more efficient strategies for interface interaction, potentially through the development of training programs, intelligent system tutors, or even through clever interface design.

REFERENCES

Altmann, E. M. (2000). Memory in chains: Modeling primacy and recency effects in memory for order. In *Proceedings of the twenty second annual conference of the Cognitive Science Society* (pp. 31-36).

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin and Review*, 8, 629-647.

Anderson, J. R. & Gluck, K. (2001). What role do cognitive architectures play in intelligent tutoring systems? In D. Klahr & S. M. Carver (Eds.) *Cognition & Instruction: Twenty-five years of progress*, 227-262. Erlbaum.

Anderson, J. R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Quin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
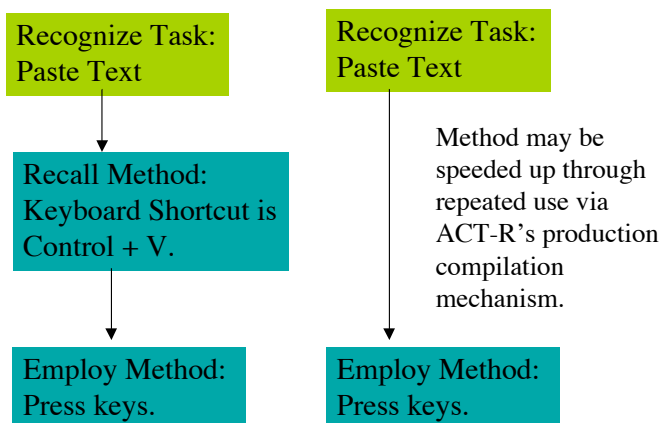
Figure 3. Example of production compilation in the model.

Best, B. J. (2006). Using the EPAM Theory to Guide Cognitive Model Rule Induction. In *Proceedings of the 13th Conference on Behavior Representation In Modeling and Simulation*.

Bhavnani, S. K., Flemming, U., Forsythe, D. E., Garrett, J. H., Shaw, D. S., & Tsai, A. (1996). CAD usage in an architectural office: from observation to active assistance. *Automation in Construction*, 5(3), 243-255.

Bhavnani, S. K., & John, B. E. (1997). From sufficient to efficient usage: An analysis of strategic knowledge. Paper presented at the *SIGCHI conference on Human Factors in Computing Systems*, Atlanta, GA.

Carroll, J. M., & Rosson, M. B. (1987). Paradox of the active user. In J. M. Carroll (Ed.), *Interface thought: Cognitive aspects of human-computer interaction* (pp. 80-111). Cambridge: Bradley Books.

Charman, S. C., & Howes, A. (2003). The adaptive user: An investigation into the cognitive and task constraints on the generation of new methods. *Journal of experimental psychology: Applied, 9(4),* 236-248.

Fu, W.-T., & Gray, W. D. (2004). Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive science, 28,* 901-935.

Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. London, UK: Cassel & Collier Macmillan.

John, B. E., & Kieras, D. E. (1994). The GOMS family of analysis techniques: Tools for design and evaluation (Technical report No. CMU-HCII-94-106). Human-Computer Interaction Institute, Carnegie Mellon University.

Lane, D. M., Napier, H. A., Peres, S. C., & Sándor, A. (2005). The hidden costs of graphical user interfaces: The failure to make the transition from menus and icons tool bars to keyboard shortcuts. *International journal of human computer interaction, 18(2),* 133-144.

Magno, E., Foxe, J. J., Molholm, S., Robertson, I. H., & Garavan, H. (2006). The anterior cingulate and error avoidance. *Journal of Neuroscience*, 26(18), 4769-73.

Peres, S. C., Fleetwood, M. D., Tamborello II, F. P., Yang, M., & Paige-Smith, D. L. (2005). Pros, cons, and changing behavior: An application in using the keyboard to issue commands. *Proceedings of Human Factors and Ergonomics Society 49th Annual Meeting*, Orlando, FL.

Peres, S. C., Tamborello, F. P., II, Fleetwood, M. D., Chung, P., & Paige-Smith, D. L. (2004). Keyboard shortcut usage: The roles of social factors and computer experience. *Proceedings of Human Factors and Ergonomics Society 48th Annual Meeting*, New Orleans, LA, pp. 803-807, Human Factors and Ergonomics Society.

Taatgen, N.A. & Lee, F.J. (2003). Production Compilation: A simple mechanism to model Complex Skill Acquisition. *Human Factors*, 45(1), 61-76.

Thomas, R. C., & Foster, M. R. K. (2001). A pilot study of teaching the strategic use of common computer applications. *Proceedings of The Australasian User Interface Conference*, Gold Coast, Australia, pp. 85-92, IEEE Press.