

Skill as the Fit Between Performer Resources and Task Demands: A Perspective from Software Use and Learning

Catherine A. Ashworth

Institute of Cognitive Science
Campus Box 345
University of Colorado
Boulder, CO 80309-0345
ashworth@clipr.colorado.edu

Abstract

This paper goes beyond the routine vs. adaptive expertise distinction seen most recently in Holyoak (1991) by offering a framework which locates skill in the fit between performer resources and task demands. Empirical support for this framework is derived from a review of the literature about "real world" software learning and usage.¹

Introduction

When we speak of skill or expertise in some endeavor or domain we are typically referring to both depth and breadth of knowledge and performance. In addition, we usually expect an expert to perform more quickly, accurately, and with less effort than a novice. Despite this consensus, some researchers have asserted the existence of more than one type of expertise. The goal of this paper is to sketch a framework that embraces both routine and adaptive expertise, explains how each is acquired, and sketches the knowledge base required for each.

Routine and Adaptive Expertise

A recent treatment of this position (Holyoak, 1991) presents a list of established properties of expertise with challenging findings for each. Holyoak concluded that there are two types of expertise, "routine" and "adaptive." The canonical view of expertise, well modeled by production system architectures such as ACT* (Anderson, 1987) and

Soar (Newell, 1990), is termed "routine." In contrast, the contrary findings are collected under the rubric "adaptive." In brief, routine expertise is rapid and accurate on highly stereotyped tasks but its inflexibility and sparse meaning precludes transfer to novel tasks or creativity. Adaptive expertise can also be swift, but is characterized primarily by knowledge breadth and the invention of new procedures.

Holyoak's list contains six important and contentious features of skill: 1) constant improvement with practice, 2) automaticity of responses, 3) domain specificity, 4) ease of task solution, 5) superior domain memory, and 6) forward search. Exceptions to these generalizations about expertise include: 1) improvement plateaux (Ericsson & Krampe, 1991), 2) successful remapping of conditions and actions (Singley & Anderson, 1989; Allard & Starkes, 1991), 3) transfer (Singley & Anderson, 1989; Dörner & Scholkopf, 1991), 4) more effortful or elaborate solution (Scardamalia & Bereiter, 1991; Jeffries, Turner, Polson, & Atwood, 1981), 5) poorer memory for some information (Adelson, 1984; Patel & Groen, 1991), and 6) flexible or breadth-first search (Jeffries, Turner, Polson, & Atwood, 1981; Dörner & Scholkopf, 1991). A review of the literature involved in the disagreement on these six issues supports the following characterization of two families of tasks -- routine and adaptive -- which differ in duration, variety, and knowledge requirements.

Routine tasks are fairly well constrained. The given-solution organization of these tasks renders them fairly brief in duration, residing in Newell's (1990) cognitive timeband. This is due to a lack of extended reasoning such as constraint elaboration, coherence checking among subplans, et cetera.² In addition, there is very little task variety. Consider a

This work was conducted while the author was being supported by grant MDA 903-89-K-0025 from the Army Research Institute

¹A much more thorough treatment of these and other related topics can be found in the author's predissertation paper. Please contact.

²From another perspective, tasks which require strategy, internal coherence checks, or elaboration of constraints cannot be represented in a straightforward set of <condition> <action> links, which can be rapidly executed.

typical task such as solving algebra problems. A certain type of problem (perhaps a dilution problem) is always properly solved with a certain method. Because this ability is rooted in recognition of problem types (beneath varying surface features) and the application of specified solutions, improvement means increasing speed and accuracy in the application of these schemas. That is, schema construction during practice essentially eliminates search during problem solving on these tasks. Skilled performance on this type of task gives the impression of extreme ease, automaticity, forward reasoning, constant (inevitable) improvement, and recall of details (for there is no higher structure).

Adaptive expertise may occur on tasks which are marked by longer duration, fundamental problem variety, and greater knowledge requirements. These tasks reside in Newell's (1990) rational time band. Example tasks are essay writing, computer programming, or medical diagnosis. Because these tasks are complex and varied, successful solution requires strong metacognition such as a detailed knowledge of plans and organizational requirements. These tasks can be regarded as problem solving tasks on which search and planning cannot be severely reduced. Therefore these tasks cannot be successfully completed through the application of so-called automatic responses. This family of tasks provides many of the exceptional results which challenge the canonical description of expertise. On these tasks, skilled performance is marked by improvement discontinuities, great effort, lack of automaticity, varied reasoning and search direction, and memory for more abstract information than for specifics.

In short, one might map the distinction between routine and adaptive expertise onto the distinction between well-defined and ill-defined problems or between closed and open tasks. Although the degree of external constraint on a task is important, this paper asserts that the sort of expertise observed is not determined solely by this factor. Many and most tasks lie between the two extremes of the routine-adaptive continuum and can be accomplished with either approach. On these tasks the performer is a major determinant of the nature of the performance. That is, the choices and style of a performer is another factor, crossed with external constraints, that must be considered. The following review of a portion of the human-computer interaction literature strongly supports this assertion.

Software Usage and Learning

The human-computer interaction (HCI) literature contains intriguing findings which challenge current thinking about skill. From a skill perspective, one of the most provocative findings in HCI is found in

studies of software learning and usage. Specifically, long term users of a piece of software (application) frequently know very few commands (Carroll & Rosson, 1987). For many users this is the result of a conscious choice. Rare additions to the small repertoire are driven by tasks that cannot be accomplished with known methods. Also, knowledge of commands may be noncumulative such that a given repertoire may not include all commands of a smaller repertoire. This pattern challenges the notion that skill necessarily and almost artifactually increases through practice in an orderly manner.

The decoupling of knowledge breadth and experience or a noncumulative progression of command acquisition has been documented in various applications. Experienced library database searchers employ a small proportion of available, useful commands (Fenichel, 1981). Bank clerks also use a small portion of available information retrieval commands, often using a sequence of two or three known commands to avoid relearning a more direct function (Kraut, Hanson, & Farber, 1983). Even computer professionals with broad technical and software experience press known functions and software into service to avoid learning (Nielsen, Mack, Bergendorff, & Grischkowsky, 1986).

Experienced users with small command repertoires are also the rule in operating system usage. The number of unique operating system commands issued by users at a large university computing site was uncorrelated with the length of the computing tradition of their various disciplines (Anstey, 1988). A study of electronic mail (e-mail) with experimental tasks revealed that long term ("regular") users issued significantly fewer commands than did system managers or consultants (Akin & Rao, 1985). Interestingly, each of these user groups employed unique commands. This noncumulativity of commands was similarly observed in an operating system command logging study which also found most users had small command sets (Draper, 1984).

Not only do most users have small, idiosyncratic command sets, they also resist complex or highly specific commands. Experienced database searchers use very simple query structures which do not exploit the system's power (Fenichel, 1981). E-mail commands which were unique to the "regular" users were more general than those unique to the system managers (Akin & Rao, 1985). In addition, the use of powerful, tailorable editing functions such as macros can be as strongly correlated to years of programming experience as to experience with the editor (Rosson, 1984). Even when self tailoring of software and workstations is aggressively taught and supported, heavy users of the software are remarkably conservative (MacLean, Carter, Lovstrand, & Moran, 1990). In summary, long term users of software typically know only a handful of general commands.

There are two basic arguments offered by those who use a small command set and routinize their work. One is that their goal is to accomplish a bit of work such as drawing a graph or reading mail from a colleague, and not to learn a piece of software. The software is strictly regarded as a tool. Following from this perspective is the strategy of attempting to achieve maximum coverage from a minimal investment, which is the other argument often made. Most users with a small command set are quite satisfied with that repertoire and feel prepared to deal with anything that might arise (Anstey, 1988). Individual users acquire a new command or function only when a specific task forces them to, which produces the observed noncumulativity of command repertoire across users.

It is not clear that users who maintain a small command set suffer. In fact, they avoid an entire family of problems for it turns out that increasing skill breadth with an application may yield low returns or even handicaps and is effortful and error prone. A lack of performance advantage from a larger command repertoire was obtained in the e-mail study: In the end, those users who employed more commands were no speedier nor more accurate (Akin & Rao, 1985). From one perspective, those users who employed fewer unique commands were more efficient. Spreadsheet based experiments reveal that learning or performance costs can accompany knowing more than one method for a task (Olson & Nilson, 1987; Ashworth, 1992) which is highly likely in a moderate or larger command set. Learning and performance costs are observed in error rates and increased planning times. One must learn an almost overwhelming amount of information of many different sorts to profit from the work (Olson & Polson, 1992; Doane, Pellegrino, & Klatzky, 1990).

Moreover, learning new functions requires venturing into poorly understood, error prone territory. Errors during exploration are particularly problematic because 1) they are difficult to identify and 2) it is not clear how to recover from an action which itself was poorly understood. In fact, a user has no guarantee that the target method will work, will be appropriate, or will save time. In addition, in a domain such as software in which the action to initiate a command is fairly arbitrary, there may be severe memory problems associated with using a large command set. Thus, acquiring enough knowledge to be adaptive is very effortful, risky, may entail costs, and may not yield benefits.

In summary, most users purposely (and perhaps wisely) maintain a small command repertoire. These users operationalize incoming tasks so that they can be accomplished with known methods, minimizing their knowledge overhead. Although this description characterizes the majority of users an interesting exception is a minority group of users, often called

"wizards," who are usually system managers or informal troubleshooters. These users possess an extensive repertoire of immediately usable commands and also a large pool of commands which can be quickly reconstructed or derived. These wizards demonstrate that a given task a user's style can be routine or adaptive. Interestingly, wizards can make more errors and access on-line help or manuals more often than do other users (Draper, 1984; but see Vaubel & Gettys, 1990). These findings represent another (apparent) conflict in the skill literature for we expect experts to make fewer errors.

A small command repertoire and a low frequency of errors both derive from an approach which routinizes and avoids additional learning. The user who avoids learning new commands or functions and instead assimilates tasks to a small well known set of functions avoids three things: 1) the error prone process of identifying and employing a new command, 2) the effort of learning and remembering a new command which may well have to be relearned for its next use, and 3) selection among known commands. However, this strategy may also reduce the ease with which new functions can be acquired (should the user select this option). The contrasting approach, knowledge seeking, produces the opposite pattern of an extensive command repertoire and a high rate of errors and on-line help accesses. These users bear the memory, learning, and performance liabilities of a large repertoire, but can do more things and may be better able to use new commands. The software learning and usage literature supports the assertion that a given task, in and of itself may not require routine or adaptive expertise. Instead, for many tasks, a performer can elect one style or the other.

A Conceptual Framework

The combination of these results from the HCI literature and the distinction between routine and adaptive expertise (discussed above) supports the following framework embracing both routine and adaptive expertise. The central tenet of this framework is that the performer determines whether she will accomplish a given task in a routine or adaptive manner. This decision is made in reference to the performer's task operationalization which incorporates the personal goal to seek or avoid additional knowledge. Although the HCI literature emphasizes the influence of the performer, expertise is located in the fit between two factors: 1) constraints inherent in the task (such as grainsize) or the situation (such as time pressure) and 2) the performer's resources. Concepts important in this argument include: performer resources, personal learning goals, and task operationalization

Real world tasks typically have an external definition based on the final goal, such as "produce a manuscript in APA format." Although such definitions may be constrained by specific process instructions or time pressure, they often contain many choices. It is this task freedom which permits the influence of performer choices. For instance, the manuscript preparer decides how and when to achieve the required margins. Because the task definition is rarely completely specified for the performer, it is detailed by the performer. The performer's operationalization of the task is developed in reference to two strongly related criteria: 1) achieving personal goals and 2) achieving a fit between performer resources and task plan demands.

Personal learning goals are best understood by example. For example, in addition to satisfying the manuscript requirements, the preparer may be pursuing personal goals including: completing the task using currently known methods only, using the opportunity to learn how to do related (potential) tasks, or using the task simply to explore the software. These personal goals fall into two general styles: 1) learning avoidance and 2) knowledge seeking. The performer considers these personal learning goals when operationalizing the task. The plan generated to accomplish the task therefore describes the particulars of what the performer is willing to do. These particulars include the knowledge that the performer has committed to bringing to the task. The relationship of this knowledge to the performer's current resources determines the nature of the fit, routine or adaptive. Although the performer also chooses the degree of fit achieved, the focus here is on the manner in which it is achieved.

When a task is acquired there may appear to be a gap between performer resources and task demands. To bridge this gap, the performer can change: 1) performer resources or 2) task demands. These two ways of bridging the gap are tied to personal learning goals just discussed. When a task seems to require unknown methods, the performer may acquire that knowledge (increase user resources) and complete the task. In this case, the personal goal is of the knowledge or understanding seeking variety. Another approach is to accomplish the task while avoiding resource expansion. This occurs when the performer introduces the personal goal of accomplishing the task with known methods. In accordance with this personal goal, the performer creates an operational definition of the task that can be achieved with current resources. This assimilates the task to known methods and avoids learning. Of course, there are many mixtures of these two extremes.

In summary, a performer has personal learning goals which shape the task demands to those that the performer is willing to accomplish. In this way, the

performer selects the degree to which, and manner in which, performer resources will fit task demands. Additionally, each task episode can shape the performer's resources for the next. As discussed below, both approaches accomplish the task, both have costs and benefits, and neither is inherently superior or more skilled than the other.

Discussion

Current skill theories which espouse the canonical view of expertise are rooted primarily in accounting for the power law of practice. In fact, investigating how these theories model the power law is a useful way to become acquainted with them for it exposes most of their mechanisms and processes. And, the power law of practice summarizes several core beliefs about skill including: constant improvement, increased accuracy, specificity, and automaticity. The performance of one who has practiced a fairly closed task of brief duration is well modeled by these theories. Example tasks include: transcription typing, naming state capitals, recalling the contents of a chess board, or sorting and solving physics problems. All of these tasks are fairly constrained by their timeband, their lack of variety, and the extremely tight guidance of mature problem schemas.

According to the power law, improvement is the reduction of problem solving processes in favor of proceduralized processes. That is, skill is regarded as a strengthening or sharpening process. Tasks which are not well captured in current theories are tasks on which improved performance or increased skill requires an effortful broadening of available procedures forced by a diet of task variety. A review of the literature cited by Holyoak (1991) makes it clear that the rogue results which he gathers together under the umbrella of adaptive expertise come from these sorts of tasks. Examples include social problem solving, writing about one's summer vacation, designing software, or controlling a complex computer simulation.

It is not clear how a theory of skill based on the power law and its entailments could explain adaptive skill. Conceptually, performance gains accrue in two loci. First, single procedures must become speeded. Second, the knowledge of pre- and post-conditions for every procedure must be represented and processed in a way that supports rapid creation of novel chains. A constraint based approach which does not model skill through the creation of large precompiled rules appears more useful (Kitajima & Polson, 1992; Mannes & Kintsch, 1989).

Although task constraints can be strong, there is another way for a performance to appear either routine or adaptive. On many tasks, the performer has the option of either assimilating or accommodating the

task. The software usage literature convincingly demonstrates that performers determine where their tasks fall along a continuum from extremely routine to extremely novel. The majority of users assimilate incoming tasks to known methods. From the perspective of efficient tool use (and controlling surprises during tool use) this is a reasonable approach. There is also an identifiable group of users who accommodate incoming tasks, taking the opportunity to learn new software functions and facts. This approach also has costs and benefits. Thus, the performer is also an important factor in the nature of the performance. This paper has sketched a framework which explicitly acknowledges and explores the consequences of the performer's choice.

This framework goes beyond the distinction between routine and adaptive expertise which focuses solely on task constraints by identifying the learner or performer as another important factor in the sort of performance observed. The relation of these two factors, the degree of external task constraints and the performer's approach (internal performer constraints), is the location of skill. That is, skill is located in the performance of a task which is the product of task and performer constraints.

This perspective has several entailments. First, in general, a given task does not inherently require or call forth one sort of expertise or the other. (See above for a discussion of the exception to this statement: tasks at the ends of the continuum.) Second, even within one domain a performer may not use solely a routine or an adaptive approach. For example, a performer may assimilate one task to current knowledge but be willing to learn a little something for another. Routine and adaptive expertise also coexist in an individual in that it is not clear how adaptive expertise could develop without the support of routine expertise.³ Consistent practice in these different styles, routine or adaptive, engenders knowledge representations which differ in content and size. Because there are costs and benefits to both ways of accomplishing a task it is inappropriate to label one superior to the other.

This framework explains some recently mentioned, apparently anomalous, findings in the area of skill acquisition and skilled performance. For instance, we have discovered how expert performance can be lengthier, more effortful, and more errorful than that of less skilled performers. In addition, we have employed the same simple ideas to understand how a performer can engage in an activity for thousands of tasks and not appear to improve in expertise. This framework is currently being

³In contrast, it is entirely possible for a performer who typically routinizes not to possess the knowledge required to support a more adaptive, flexible approach.

elaborated and expanded to account for other variations in experimental results.

Acknowledgements

The author expresses gratitude to Gene Gollin for championing the "system perspective" and to Marita Franzke and Peter Polson for useful discussions of these ideas. Appreciation to Marita Franzke, Evelyn Ferstl, Susan Davies, Cathleen Wharton, Julia Moravcsik, Adrienne Lee, and Abby Harrison who provided guidance, encouragement, and proofreading.

References

- Adelson, B. (1984). When novices surpass experts: the difficulty of a task may increase with expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 483-495.
- Akin, O., & Rao, D.R. (1985). Efficient computer-user interface in electronic mail systems. *International Journal of Man-Machine Studies*, 25, 557-572.
- Allard, F. & Starks, J.L. (1991). Motor skill experts in sports, dance, and other domains. In K.A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge: Cambridge University Press.
- Anderson, J.R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review*, 94, 192-210.
- Anstey, P. (1988). How much is enough? A study of user command repertoires. In D.M. Jones & R. Winder (Eds.), *People and Computers IV, Proceedings of the Fourth Conference of the British Computer Society*. Cambridge: Cambridge University Press.
- Ashworth, C.A. (1992). Specialized methods do not increase efficiency. In *CHI '92 Poster and Short Talks Proceedings of Human Factors in Computing Systems*. New York: ACM.
- Carroll, J.M. & Rosson, M.B. (1987). Paradox of the active user. In J.M. Carroll (Ed.), *Interfacing Thought: Cognitive aspects of human-computer interaction*. Cambridge, MA: MIT Press.
- Doane, S.M., Pellegrino, J.W., & Klatzky, R.L. (1990). Expertise in a computer operating system: conceptualization and performance. *Human-Computer Interaction*, 5, 267-304.
- Dorner, D & Scholkopf, J. (1991). Controlling complex systems or expertise as "grandmother's

- know-how". In K.A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge: Cambridge University Press.
- Draper, S.W. (1984). The nature of expertise in UNIX. In B. Shackel, (Ed.), *INTERACT '84 - IFIP Conference on Human-Computer Interaction*. Amsterdam: Elsevier-Science.
- Ericsson, K. A. & Krampe, R. (1991) The role of deliberate practice in the acquisition of expert performance. Technical Report of The Institute of Cognitive Science, University of Colorado, Boulder, CO.
- Fenichel, C.H. (1981). Online searching: measures that discriminate among users with different types of experiences. *Journal of the American Society for Information Science*, ?, 23-32.
- Holyoak, K.J. (1991). Symbolic connectionism: toward third-generation theories of expertise. In K.A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge: Cambridge University Press.
- Jeffries, R., Turner, T., Polson, P., & Atwood, M. (1981). Processes involved in designing software. In J.R. Anderson (Ed.), *Cognitive Skills and Their Acquisition*. New Jersey: Hillsdale.
- Kitajima, M. & Polson, P.G. (1992). A computational model of skilled use of graphical user interfaces. In *CHI '92 Proceedings of Human Factors in Computing Systems*. New York: ACM.
- Kraut, R. E., Hanson, S.J., & Farber, J.M. (1983). Command use and interface design. In *CHI '83 Proceedings of Human Factors in Computing Systems*. New York: ACM.
- MacLean, A., Carter, K., Lovstrand, L. & Moran, T. (1990). User-tailorable systems: pressing the issue with buttons. In *CHI '90 Proceedings of Human Factors in Computing Systems*. New York: ACM.
- Mannes, S.M. & Kintsch, W. (1989). Planning routine computing tasks: Understanding what to do. ICS Tech Report 89-09. Boulder, Colorado: Institute of Cognitive Science, University of Colorado.
- Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Nielsen, J., Mack, R.L., Bergendorff, K.H., & Grischkowsky. (1986). Integrate software usage in the professional work environment: Evidence from questionnaires and interviews. In *CHI '86 Proceedings of Human Factors in Computing Systems*. New York: ACM.
- Olson, J.R. & Polson, P. G. (1992). The cognitive cost paradigm. Manuscript in preparation.
- Olson, J.R., & Nilsen, E. (1987). Analysis of the cognition involved in spreadsheet software. *Human-Computer Interaction*, 5, 221-265.
- Patel, V.L. & Groen, G.G. (1991). The general and specific nature of medical expertise: a critical look. In K.A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge: Cambridge University Press.
- Rossen (1984). The role of experience in editing. In B. Shackel, (Ed.), *INTERACT '84 - IFIP Conference on Human-Computer Interaction*. Amsterdam: Elsevier-Science.
- Scardamalia, M. & Bereiter, C. (1991). Literate expertise. In K.A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge: Cambridge University Press.
- Singley, M.K. & Anderson, J.R. (1989). *The Transfer of Cognitive of Skill*. Cambridge, MA: Cambridge University Press
- Vaubel, K.P. & Gettys, C.F. (1990). Inferring user expertise for adaptive interfaces. *Human-Computer Interaction*, 5, 95-117.