Human-Task Adaptations: The Next Step for Cognitive Modeling

Michael Freed, Michael Matessa, John Rehling, Roger Remington, and Alonso Vera
{mfreed, mmatessa, jrehling, rremington, avera}@arc.nasa.gov
NASA-Ames Research Center, Mailstop 262-4
Moffett Field, CA 94035-1000

Influences on behavior range in a continuum from task-specific (e.g. there are seven buttons on the interface) to architecture-specific (e.g., human working memory can span 7+/- 2 items). The study of this continuum has unfortunately been focused more on the ends than the center. Situated cognition approaches suggest that all behavior is specific to the situation in which it takes place. In contrast, traditional cognitive views suggest that there is a common set of architectural resources that underlie behavior consistently across situations. Somewhere between the implementation of a cognitive architecture and the analysis of a specific task, lies a set of psychological phenomena that determine how people actually perform in that task given a set of architectural constraints. These phenomena, which we will call *Human-Task Adaptations*, are consistent characteristics of the way the cognitive architecture shapes itself around and adapts to the specific characteristics and demands of a given task. Only by understanding this class of phenomena are we likely to make any serious progress in cognitive modeling.

There are ways in which humans multitask, plan reactively, use working memory, etc. that are relative, but not specific, to the kinds of tasks they are working on. That is, there are certain predictable ways in which people behave in given tasks and environments (e.g., air traffic control, flight deck, shuttle control, etc.) that are not specifically determined by the cognitive architecture nor uniquely determined by the specifics of the task. Computational cognitive modeling has demonstrated its value over the past decade but little progress has been made on the key issues that would allow us to predict error, cognitive load, resource allocation and so on. These phenomena, we argue, are a consequence of an interaction between the cognitive architecture and specific tasks.

If a task has frequent interruption, like ATC, then the behavior of experts might be expected to adapt to the patterns of interruptions in consistent ways. This may be achieved by evolving unit tasks with shallow goal stacks that allow for frequent checks of the environment for new information while maintaining a reasonable completion rate for ongoing unit-tasks. Nevertheless, performance on such tasks has typically been studied by cognitive scientists at the level of the architecture rather than relative to specific task domains. This is, of course, for good reason. There is little value in describing, in great detail, how the human architecture interacts with an interface such as that on a microwave oven. There is little to generalize from it with respect to the cognitive architecture.

Efforts to model human performance have traditionally focused on fundamental human characteristics such as limited attention and working memory. For brief tasks carried out in controlled environments, this focus is appropriate since performance often depends directly on fundamental capabilities. The issue is more complicated for more complex tasks that may be carried out in any of several ways. Since these may vary in important measures of performance (e.g. required time, likelihood of error), predicting performance depends on knowing how the task will be carried out. Experts will have learned how to minimize the constraints imposed by cognition and other potential sources of limitation. In some cases, native human limits will become largely irrelevant, influencing performance only indirectly by having shaped behavioral strategies.

A situated cognition perspective would prescribe that humans constantly undergo unique processes of adaptation to every aspect of new tasks, with little of the cognitive architecture affecting performance from one trial and one task to another. A strict cognitive architectural approach, on the other hand, might argue that the architecture should define every aspect of performance in any task. We have attempted to define a space between these two possibilities in which aspects of the architecture adapt to specific characteristics of tasks in ways that are consistent and predictable.

The key theoretical point behind *Human-Task Adaptations* is the argument that, although they are not directly determined by the cognitive architecture, they are nevertheless well-defined phenomena that are *generalizable* to tasks that impose similar demands on the architecture. We are not claiming that *Human-Task Adaptations* are learned explicitly and abstracted from situations in the way that using means-ends analysis or analogy might be considered to be. They are acquired implicitly and remain specific to the task from which they resulted. They are generalizable only in the sense that we, as scientist, can predict how an individual will adapt to the demands of a particular task.

HCI provides classes of problems that share relevant demand characteristics: Air-traffic control is similar to radar operation on AEGIS ships, shuttle operation, satellite operations, nuclear power-plant control, and so on. All of these tasks involve experts in dynamic, real-time, complex environments. Expertise makes their interactions with dynamically changing situations *routine*. Experts come to adapt their planning and decision-making in consistent ways with respect to the rapidly evolving characteristics of the situation. Specifically, their behavior adapts to variables like the patterns of interruption in the task, the number of tasks to be carried out in parallel and the amount of perceptual support provided by the interface to working memory. These are the human-task adaptations that must be understood and built into our computational cognitive models if we are to succeed in simulating complex human behavior. Complex HCI tasks provide an ideal test-bed for exploring psychological idioms exactly because they involve experts deploying routine behavior in rapidly changing environments.

Different kinds of tasks force the emphasis of an investigation into different points along a continuum. At one extreme (A), the kinds of tasks used in psychophysics experiments necessarily probe the fundamental abilities of the human cognitive architecture. At the other extreme (Z), a task like solving the missionaries-and-cannibals problem tells us very much about the task, quite independent of the cognitive architecture of the problem-solver; the constraints of the task ensure that a human, a martian, and Prolog would all produce the same solution. We argue that what is needed to elevate cognitive modeling to the next level is the development of a vocabulary of cognitive activities and the ways to model them, striving for a higher level of complexity than the findings (with a large corpus of them already uncovered) at (A), but generalizable, unlike those at (Z). We observe that our ability to model many domains has little to gain from investigations focusing on either end of the continuum.

This is perhaps illustrated better by means of analogy to tasks that involve not only human cognition, but also human physiognomy. Consider that we were trying to enrich the field of human modeling with the goal in mind that future modelers might undertake modeling a task such as bear hunting. (We mean the task of hunting a bear without modern weaponry; people in pre-technological civilizations found a number of ways to go about this.) We argue that studies of tasks of this kind would uncover basic human strategies that would generalize to other tasks. For example, the use of several bowmen solve the problem in one way. The dead-fall trap, which uses bait to lure a bear underneath a log which can be released to break the bear's back, may be of use in other tasks in which human strength is a serious constraint. We observe that ***no*** amount of study of human abilities alone (A) is sufficient to generate the strategies of use in bear hunting. Likewise, studies with an excessive focus on the task (Z) will not generalize sufficiently. Even if a great deal is known about bears, that will not generalize to **hunting** bears (without taking human constraints into account, one will not focus on the importance of avoiding the use of sheer strength). In contrast, the study of a task that is very different from hunting and does not involve bears (such as how to split very large rocks) may generalize to bear hunting in useful ways - more than intensive and extensive fields of study dedicated to human capabilities, and to bears, respectively.

Similarly, studying a task such as Air Traffic Control, where genuine experts work on a complex, time-critical, highly interactive task, similarly sheds light on the unique ways in which the human architecture adapts its constraints to meet the demands of such tasks.