

To appear in: *Integrated Models of Cognitive Systems (IMoCS)*, Wayne Gray (ed.).
New York, NY: Oxford University Press.

Lessons from Defining Theories of Stress

Frank E. Ritter and Andrew L. Reifers
Applied Cognitive Science Lab
School of Information Sciences & Technology
Penn State

Laura Cousino Klein
Biobehavioral Health
Penn State

Michael J. Schoelles
Dept. of Cognitive Science
Rensselaer Polytechnic Institute

25 January 2006

Target: 7,000 words: approximately 6,300 words in total.

Introduction

In this chapter we look for lessons applicable to cognitive architectures from several popular theories of stress. The goal is to specify mechanisms that can be implemented within the architecture or changes to current mechanisms to simulate the effect of stress on embodied cognition. We examine theories from Wickens, from Hancock and Warm, and from the biophysiology literature. We have chosen to incorporate these theories of stress into the ACT-R architecture because of its modular construction, but the intent is for the ideas presented here to be applicable to a wide range of cognitive architectures.

These theories of stress are typically not cast as additions to the knowledge necessary to perform a task (which would make implementing them be creating a cognitive model), but are described as changes to how people process information under stress. Thus, they make suggestions about process, about how the mechanisms of embodied cognition change across all tasks under stress. Implementing them thus becomes modifying the architecture or changing parameters of existing mechanisms. Unfortunately, not all the lessons that we have extracted from our survey are specific enough to be directly implemented. In these cases, the lessons serve as specification guidelines for new mechanisms.

Each theory of how stress influences embodied cognition, if specific enough, is presented as an “overlay” to the ACT-R cognitive architecture. These changes are “overlaid” onto the basic ACT-R theory, modifying how the mechanisms process information, and in one case, giving any ACT-R model a secondary task—to worry. The main idea of an overlay is to change the architecture in such a manner that the behavior of all models developed under that architecture will be affected. In a similar way, the changes to ACT-R to model fatigue by Gunzelmann and Gluck (this volume) should be applicable to other ACT-R models, and thus it is an example overlay.

Including these theories (overlying these theories) into cognitive architectures offers several advantages. We will be able to construct models that make stress theories more precise. Until now these theories of stress have typically only been verbally stated. We will be able to provide more realistic opponents in competitive environments that respond to stress. We also will extend the scope of cognitive architectures because most architectures have not been developed with these theories of stress in mind. And we will be able to explore more complex concepts, such as how situation awareness is influenced by stress.

The theories of cognition under stress can also benefit. These stress theories are not all complete. By implementing them in a cognitive architecture, and thus having them interact with other embodied cognitive mechanisms, the relationships among these theories can be made more explicit.

Before introducing the theories for implementation as overlays into ACT-R, we briefly review cognitive architectures, expand our definition of overlays, and describe a sample task to help explain the application of overlays. After describing the theories of stress we discuss why the theories of stress require time as an architectural mechanism. We conclude with a plan for testing the overlays including specifying what currently has been done and what remains to be accomplished.

Cognitive architectures and overlays

We start this review by explaining cognitive architectures and developing the idea of an overlay with respect to a cognitive architecture. We also describe an example task that we will use to illustrate and, in later work, to test these overlays.

Cognitive architectures

Cognitive architectures are an attempt to represent the mechanisms of cognition that do not change across tasks. Typically, cognitive architectures include a long-term procedural and declarative memory, a central processor and some working memory for that processor, and perceptual and motor components (e.g., see Anderson et al., 2004; Newell, 1990).

Gray (this volume) references three types of control mechanisms in cognitive architectures. Type 1 theories are those that have central control. A central executive processor represents this type of control. Type 2 mechanisms speak to peripheral processors and distributed control. It is likely that the human mind has both aspects. Current cognitive architectures typically include both types of control. Type 3 control mechanisms are task specific changes. We will also introduce the idea of distributed and physiological implementations, and label these potential changes as Type 4 mechanisms, which reflects the physiological aspects of other mechanisms.

Overlays

An overlay is a technique for including a theory of how a behavioral moderator, such as stress, influences cognition across all models within a cognitive architecture. An overlay, as we propose it, is an adjustment or set of adjustments to the parameters or mechanisms that influence all models implemented in the architecture to reflect changes due to an altered mental state or due to long term changes such as development (e.g., Jones, Ritter, & Wood, 2000). In many architectures there are a set of mechanisms and a number of global parameters that play a role in the model's functioning; an overlay modifies a combination of parameters and mechanisms to represent situation specific but relatively long acting changes to information processing. For example, an eyeglasses overlay would allow more inputs to be passed to the vision processor; a caffeine overlay could increase processing speed by 3% and improve vigilance by 30%.

This overlay approach keeps the architecture consistent across tasks and populations, but allows that there may be differences in processing mechanisms and capabilities for individuals or groups in certain contexts. The concept of an overlay and the associated software are useful concepts for developing theories in this area. It provides a way to describe these differences in a way that can be applied to all models in the architecture. With further work, of course, these overlays would migrate into the architecture.

There have been a few previous attempts to generate overlays, although most of their authors did not call them overlays. A brief review notes the scope of possible overlays. The earliest attempt of we are aware is by Jongman (1998), who attempted to model fatigue in ACT-R, and by Chong (1999) modeling fear in Soar. Later work on creating a model of stress and caffeine used simple adjustments of increasing the variability of applying procedural knowledge (Ritter, Avraamides, & Council, 2002). Another chapter in this volume (Gunzelmann et al.) presents a more complete overlay to model fatigue.

There have also been some complex overlays that add mechanisms rather than modify

parameters. Belavkin has worked on several overlays to ACT-R to make models in ACT-R more sensitive to local success and failures, creating a simple theory of how motivation influences problem solving (e.g., Belavkin & Ritter, 2004). Gratch and Marsella (2004) created a rather complex addition to Soar to model task appraisal. It is not clear that their overlay is truly portable, but their work suggests how such an appraisal process can be built in an architecture and provides an example overlay in Soar.

Application to ACT-R and other architectures

Implementing overlays is easier to do in architectures that are modular, where their source code is available, the implementation language easily allows extension (like Lisp), and the existing mechanisms have parameters. This chapter will use ACT-R as the architecture to implement the theories as overlays because ACT-R (Anderson et al., 2004) has all these features. However, the overlays description could be used fairly easily with other architectures that have these features or many of them. For example, COJACK (Norling & Ritter, 2004; Ritter & Norling, in press) was designed to support the creation of overlays, and EPIC (Kieras, Wood, & Meyer, 1995) should also support including overlays like these.

ACT-R is a hybrid embodied cognitive architecture (Anderson et al., 2004). ACT-R specifies constraints at a symbolic level and a Bayesian based sub-symbolic level. It has perceptual-motor components as well as memory and cognitive control components. All components or modules communicate through buffers. Examples of ACT-R models are available in **XX's, Y's, and Z's** chapters.

A sample task—Serial Subtraction

In this chapter we will use one of our tasks in the CafeNav suite (Ritter, Ceballos, Reifers, & Klein, 2005), repeated serial subtraction, as a task where performance is modified by stress. We have collected data on this task and have an ACT-R model. It shares components with tasks that use arithmetic knowledge, and a large literature in physiology and biophysiology has used it as a task to study stress (e.g., Kirschbaum, Pirke, & Hellhammer, 1993). Other tasks could be used, and, indeed, that is the point of overlays, that these changes would lead to different performance in other tasks as well.

In this task the subject is asked to repeatedly subtract a small number, like 7 or 13, from a running total that starts as a four-digit number, such as 5,342. If the answer is incorrect, the experimenter informs the subject that the answer is incorrect, and gives the correct starting number.

Subjects can be manipulated to appraise this task as threatening. When this occurs, the number of subtractions drops by about 10% from 55 per 4-minute block to 47, and the number correct only drops from 47 to 42 (Tomaka, Blascovich, Kelsey, & Leitten, 1993). Thus, threat appraisal appears to impair performance rate, but not performance accuracy.

We have created a model in ACT-R 5 and recently in ACT-R 6 to perform this task. The ACT-R 5 model has been modified with a simple stress and a simple caffeine overlay and the results compared to published data (Ritter, Reifers, Klein, Quigley, & Schoelles, 2004). In both versions of ACT-R, this model uses procedural and declarative knowledge and audio input and output. Overlays that modified these mechanisms would modify the model's performance.

The Theories of Stress and Cognition

We examine six theories of how stress influences performance. We have generated implementations of the first four theories as overlays to ACT-R¹. We have not generated implementations of the last two theories, but we pull lessons for the development of architectures, in particular, for ACT-R, from the last two theories.

Wickens' theories

We have taken Wickens' theories of stress from Wickens, Gordon, and Liu's textbook on human factors (1998, Chapter 13, pp. 324-349). The theories of how stress influences cognition are not yet very predictive. Wickens et al. note that the amount of stress is difficult to predict for a given situation, as differences can arise due to how the task is performed (allowing more or less time to appraise the situation), how the task is appraised (e.g., threatening versus challenging) due to level of expertise, and whether one perceives that they are in control of the situation. The authors do, however, go on to provide some theoretical statements that can be used to implement theories of stress as overlays. We examine three of them here, perceptual tunneling, cognitive tunneling, and changes to working memory.

Wickens-Perceptual Tunneling

Wickens et al. (1998; Wickens & Hollands, 2000) note perceptual narrowing (or tunneling) as a major effect of stress. We take perceptual narrowing to be where the effective visual perceptual field becomes smaller with stress, such that items in the periphery become less attended. Thus, the item focused on is typically the cause of stress or related to relieving a stressor and other items are less available to cognitive processing.

Perceptual tunneling can be implemented in ACT-R in several ways. One way is to decrease the default distance (from the screen) parameter. This effectively makes objects on the periphery less visible (we do not imply that under stress people lean into the screen, this is purely a way to implement this effect). Another way is to modify the visual attention latency parameter. This makes moving attention to the periphery slower, and some models and humans would use information less from the periphery because it is harder to get, missing more changes there. Another way is to limit what is in the perceptual field by decreasing its width; this is slightly more difficult to implement than the other approaches. Finally, an approach to create perceptual tunneling is to increase the saccade time. ACT-R 5 is based on a theory of moving visual attention and does not include a theory of eye movements. An extension to ACT-R called EMMA (Salvucci, 2001) adds a theory of eye-movements to ACT-R. In particular, EMMA computes eye movement times, which can be changed to implement perceptual tunneling.

All of these implementations modify Type 2 mechanisms of central cognition. These visual mechanisms, as implemented in ACT-R, are not used in the current model of the serial subtraction task, so these changes would have little effect on serial subtraction. For serial subtraction the inputs, corrections, and outputs are all auditory. Even if tunneling is applied to the auditory information, tunneling should only help performance on this task as tunneling focuses attention on the primary task and removes attention from a secondary task. The overlay of this popular theory will not be able explain what happens to the performance of serial

¹ Implementations of the model and these overlays are at acs.ist.psu.edu/cafenav/overlays/

subtraction under stress, as there is only one task. However, it would be usefully applied to many ACT-R models that use vision to interact with tasks, and predict how performance in more complex environments would change under stress. It may also lead to improved performance, as Wickens and Hollands (2000) review, in tasks where avoiding distractions improves performance.

Wickens-Cognitive Tunneling

Cognitive tunneling occurs where a limited number of options are considered by central cognition. To implement cognitive tunneling, the declarative and procedural retrieval thresholds in ACT-R can be modified by the overlays to represent a greater reliance on well-known and well-practiced knowledge. Alternatively, noise in the procedural rule application process can be decreased (i.e., activation noise and expected gain noise), which would lead to only the most well-practiced materials being retrieved and applied. These overlays modify central cognition, a Type 1 mechanism.

In the case of serial subtraction cognitive tunneling, as a form of focused attention, should improve performance. We might expect this overlay to hurt performance on a more complex task, but on serial subtraction cognitive tunneling should improve performance.

Wickens-Working Memory

Wickens and colleagues (1998, p. 385) are among the many that note that working memory capacity appears to decrease under stress. Under stress, working memory appears to be less available for storing and rehearsing information, and less useful when performing computations or other attention-demanding tasks. Wickens et al. go on to note that long-term memory appears to be little affected and may even be enhanced. This effect (or lack thereof) may be due to focusing on well-learned knowledge. Focusing on well-known knowledge to the exclusion of less on less well known knowledge would lead to using less of long-term memory.

The implementation of this theory is not completely straightforward. Lovett and colleagues (2000) showed how working memory can be modeled within the ACT-R architecture and proposed the W (goal activation) parameter as a measure of individual working memory capacity.

A promising approach is to modify the decay rate of working memory objects. This is interesting, as this approach is sensitive to many known effects on working memory, including that it would predict that the time to report objects in working memory will influence the measurable size of memory. Other ways to decrement working memory are to increase the declarative memory retrieval threshold parameter, which would make fewer memories available, and to decrease the base-level activation of all memory elements when stressed.

All of these overlays modify central cognition, and thus influence a Type 1 mechanism. Different implementations of this theory as an overlay would have different effects on the serial subtraction model. Changes to working memory should decrease the number of calculations per unit of time. The changes might also influence percent correct, but these overlays would probably need to be implemented to obtain accurate predictions.

Decreased attention

Hancock (1986) describe the effect of stress on cognition as being decreased attention, and provides a theory based on a review of work on stress and attention. Wickens et al. (1998) also note that there is decreased attention with stress.

ACT-R includes several ways that attention can be modified. Working memory capacity is a type of attention that is also a parameter that can be modified in ACT-R. We use it in the Wickens-WM overlay, discussed above, so we exclude it here.

Another way to implement an overlay that decreases attention to the task is to create a secondary task. This secondary task simulates worry. This approach to modeling stress is consistent with theories of math anxiety and other studies of anxiety that posit a dual task of worry as the cause of poor math performance in people with math anxiety (e.g., Ashcraft, 2002; Cadinu, Maass, Rosabianca, & Kiesner, 2005; Sarason, Sarason, Keefe, Hayes, & Shearin, 1986). The rules that create this secondary task might be seen as architectural productions, similar to Soar's default rules, but far less useful. The genesis and activation of these rules are likely tied into central cognition and emotion.

This secondary worry task has been implemented as a pair of productions that represent how worrying uses resources and decreases attention. When these rules are chosen, they take time from the primary task, allow the declarative memory of the primary task to decay over the time that they take to apply, and can leave the working memory in an incorrect state if they are interrupted themselves.

This overlay decreases performance rate on the serial subtraction task. If the secondary worry task was performed quite often and the math facts were poorly known, it could increase the error rate. If the worry task was performed only occasionally, it is possible that it would not change the error rate. We have used a secondary task as an overlay before to model the effects of worry as a type of stress with our serial subtraction task (Ritter et al., 2004). We have also applied the dual-task worry overlay to a driving model to predict how worry could increase lane deviation for a simple driving task (Ritter, Van Rooy, St. Amant, & Simpson, in press).

This overlay does not directly modify a mechanism, per se, but its effects are felt primarily in central cognition as the result of including another task. This overlay is classifiable as a Type 3 change (an addition to knowledge).

The task as a stressor

Hancock and Warm (1989) argued that tasks are themselves stressors. They describe the effect of stress on task performance as an individual's interest in maintaining an optimal amount of information flow. They note several negative effects, including that "physiological compensation is initiated at the point at which behavioral response reaches the exhaustive stage." So, they note, as stress increases, the agent modifies its cognitive efforts to maintain performance, then its physiology changes to support these efforts, and then, after a sufficient period of time, there will be a catastrophic collapse due to exhaustion of the physiological level. Hancock and Warm (1989) note that the task itself should be seen a source of stress, with sustained attention as the stress generator. For well-practiced, automatic tasks, which need little attention, there is little effect of stress on performance and performance of the task does not increase stress.

In their paper they explain why heat stress on a simple task can be better predicted than noise stress. Their approach is to understand the input, the adaptation, and the output. Heat stress has a fairly simple input description. Its effects on cognition, physiology and adaptation appear to be, they note (p. 532), direct and straightforward. There are few strategies to cope with heat and these strategies are simple and do not modify cognition. The empirical results they review show a very nice area where physiology directly affects cognition. Cognition directly suffers when thermoregulatory action can no longer maintain core body temperature.

Hancock and Warm (1989) point out that environmental auditory noise as a stressor is less easy to characterize, and that adaptation to noise is more complex. Adaptation to noise varies more across individuals and is more complex cognitively and physiologically than is adaptation to heat stress.

We have not yet created an overlay based on Hancock and Warm's proposal that tasks themselves are stressors. Such an overlay could be quite complex. They describe important aspects of this process, that of input, adaptation, and response. Although we have not modeled the effects of tasks as stressors themselves, we recognize that modeling tasks as stressors is an interesting and important next step in the effort to model the effects of stress. Hancock and Warm's (1989) theory also predicts that performance and resources will degrade faster when performing more complex tasks. We also gain a greater understanding about stress and other moderators: situations where people can use multiple coping strategies will be more difficult to model.

An overlay based on Hancock and Warm's (1989) theory, if created, would need to use central cognition to compute how to modify its behavior. The physiological implementation of the mechanisms and how they would tire from effort is perhaps a new type of mechanism (Type 4 noted above). This theory suggests that a detailed model of serial subtraction that includes coping strategies will be challenging to create. At the least, this theory will predict that errors will increase over time on the serial subtraction task because the system will fatigue.

Pre-task appraisal and stress

Some researchers, particularly biobehavioral health scientists have been interested in stress that occurs as a result of particular cognitive appraisal processes (e.g., Lazarus & Folkman, 1984). According to appraisal theories, before performing a task, the task is appraised as to how difficult the task will be, and how well the person thinks they will be able to cope with the task's demands (i.e., what coping resources they have). These appraisals are typically classified into two categories, either "threatening" or "challenging". Threatening tasks are those in which the person approaching the task believes that their resources (including task knowledge) and coping abilities are not great enough for the task, whereas tasks appraised as challenging indicates that their resources and coping abilities are great enough to meet the task demands.

Challenging appraisals give rise to better energy mobilization, and better performance in general than threatening appraisals. In these studies, such results are also found when the appraisal is manipulated and knowledge held constant, so it is not just a knowledge-based effect. This general result has been shown for a several different tasks, but performance on repeated serial subtraction has perhaps most often been used to study this effect especially with regard to cardiovascular physiological changes that occur as a function of these appraisals (e.g., Quigley, Feldman Barrett, & Weinstein, 2002). For example, in these serial subtraction tasks threatening appraisals lead to about a 25% slower performance than "challenging appraisal" conditions, but with the same accuracy (e.g., Tomaka et al., 1993).

This overlay, which would be much larger than the others reported here, has not been implemented in our set. Previous attempts to create an appraisal overlay for ACT-R have simply modified how accurately rules are applied, providing more accurate rule applications for challenged appraisals and less accurate application for threatened appraisals. This overlay appeared to match the limited data available so far (Ritter et al., 2004). A complete version of an appraisal overlay might be so extensive that it should not even be called an overlay, but rather a complete extension to the architecture.

The serial subtraction model would be particularly sensitive to this overlay. The difficulty will be translating the physiology results that either co-occur with or are reflective of changes to specific cognitive mechanisms. We note the next steps in this area below in steps towards testing these models.

Discussion

We have presented six interesting theories of how stress influences cognition and created implementations of four of them for applying to cognitive models in ACT-R. Creating this set of overlays provides a rich set of lessons for the further development of ACT-R and other cognitive architectures. Implementing the remaining theories and testing these overlays should provide further lessons about how cognition and performance vary on these tasks. The existing lessons include lessons on the development of these theories of stress, the important role that time has to play in cognitive architectures, and how to test these theories. After discussing these lessons we note directions for future work.

Theories of stress are not complete

The review above suggests the many ways that stress can be conceptualized as changes to cognition. However, most of the theories modify only one aspect of cognition.

In several cases there were multiple ways to implement these theories. Where this occurred we often included multiple variants. This is useful for the theories, as testing the multiple interpretations will help make them more concrete and accurate.

Some of the theories sound different, but become the same when we implement them (e.g., decreased attention and decreased working memory capacity). This may be due to the limitations of us as designers of models or of the architecture we are working with, but we can hope that it is because these multiple theories are attempting to describe the same phenomena.

Time: An additional aspect to mechanisms

This review suggests that there is an additional aspect to mechanisms that ACT-R needs, namely, how the mechanisms are sensitive to the passage of time. Many stress theories predict that people are sensitive to time. That is, the theories predict that information processing and thus the cognitive mechanisms vary as time passes. Hancock and Warm (1989), for example, explicitly note this effect. Most models in most architectures, including ACT-R, do not take account of time's influence on task performance. For a counter-example of a model that is influenced by time, see work by Gunzelmann, Gluck, and their colleagues (2005 and this volume). Their model modifies the value of the goal over time, representing fatigue.

These overlays will also need to incorporate aspects of how time is spent (e.g., sleeping) and how time-on-task for a single, particular task will influence task performance. Thus, we have proposed a Type 4 mechanism to add to Gray's (Chapter 3) taxonomy of control mechanisms. Type 4 is related to physiological mechanisms, how they support cognition, and how they become fatigued over time. (It could be called a sub-sub-symbolic level.)

Testing these overlays

These overlays need to be applied to a range of tasks, and the results compared to subjects' task performance. Fitting and testing the effects of these overlays are expected to refine them. A large battery of tasks will be important because on a single task some of these overlays may make similar predictions. Because this testing process requires a suite of tasks, a range of data, and several models, it will require reuse of models and tasks. Preliminary data (Ritter et al., 2005) on a selected task suite that uses different cognitive mechanisms shows potential differences in performance that should help differentiate and refine these overlays.

We are in the process of testing the overlays described in this chapter on a number of models in ACT-R of tasks: a working memory task (MODS, Lovett et al., 2000), a signal detection and simple reaction time task (Reifers, Ritter, Klein, & Whetzel, 2005), a serial subtraction task (Ritter et al., 2004), and a complex dynamic classification task (Schoelles & Gray, 2001). The current state is that the task environments have been developed and tested. The VSDT model was built from scratch, but the others are revisions to older ACT-R models. All of these models have been compared to normal data, either in their previous version or in their current version. The serial subtraction model and the VSDT model have also generated predictions with a caffeine overlay (Reifers et al., 2005; Ritter et al., 2004).

The perceptual tunneling, cognitive tunneling, working memory, and decreased attention overlays have been created in ACT-R 6 as well. These overlays are ready to be tested. One of them has been used before, however. We added the decreased attention overlay to a model of a simple driving task and generated predictions (Ritter et al., in press). The predictions leave the model a considerably poorer driver (as measured by average lane deviation and time between crashes in the game). We do not yet have data from stressed individuals that can be used to test these predictions.

The next phase is to run the models with and without the overlays and then to analyze the model and human data. Data from the CafeNav study (N=45) is now in hand for testing the overlays explained here. (A description of the pilot version study is available, Ritter et al., 2005.) This study measured subjects as they performed these tasks and measured their stress using a range of measures (including heart rate, blood pressure, and salivary cortisol). The next step is to run the models with the overlays included and to compare the resulting behavior to our data.

Further work on overlays

These stress overlays might also be further extended. They could be extended to model different levels of stress or individual differences such as personality type. They could also be refined. For example, the distracter thought overlay could have several different thoughts, some not focused on the task, which would slow down the model and thereby reduce the activation of working memory, and other thoughts that might be more task related, and would slow down task performance, but raise the activation of particular memories related to the current task, and thereby lead to different learning or performance.

Several theories of stress, not examined here in detail, make assumptions about processing that cannot yet be directly or routinely implemented in models or their architectures. Examples of these effects include how under stress, particularly high workload stress, tasks are deliberately shed (e.g., Parasuraman & Hancock, 2001; Wickens, Gordon, & Liu, 1998) and how negative information in particular may be lost or ignored under stress. Both of these require more complex models and architectures than are currently used to represent task load, and to be implemented they need the model to dynamically adjust its strategies, and to include a more complex representation of knowledge than is currently used.

These overlays developed here are static overlays, representing the state of processing in a challenged (or neutral) and in a threatened state, and are fixed across time. Stress is probably not a binary effect. It is clear that stress has a range of values and a more dynamic nature than the theories and these overlays provide. Future work will have to grapple with the additional difficulties of creating a model of stress that dynamically adjusts to the effect of the task itself and of a range of stress on the task.

Summary

We have reviewed the literature on the effects of stress on embodied cognition with the goal of being able to simulate these effects in cognitive models. We described the process of overlaying the ACT-R theory of embodied cognition with these theories of stress in order to develop more accurate cognitive models. The theories we examined as the basis of our architectural overlays were perceptual and cognitive tunneling, diminished working memory capacity, decreased attention, pre-task appraisals, and the task as a stressor.

Creating a set of overlays that start to model the effects of several of the major theories of stress was a first step. The next step here is to see how to compare the models with overlays to data, including the differences in error rates and error types that these overlays predict. After we test these overlays, there are several lines of research that can be investigated. First, we would hope to create other global overlays modeling other psychological phenomena, demonstrating that the concept of overlays in a cognitive architecture is indeed useful. For example, overlays could be extended to study extreme levels of stressors: of fatigue, of heat, and of dosages of caffeine. Granted, these goals are ambitious, but the possible implications and impact seem to be worth exploring.

Acknowledgements

Susan Chipman has provided several types of support for this work. Roman Belavkin, Wayne Gray, and Karen Quigley have helped our thinking in this area. Comments from Wayne Gray, Glenn Gunzelmann, Karen Quigley, and Dan Veksler have improved this chapter. This work was supported by the US Office of Naval Research, award number N000140310248.

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036-1060.
- Ashcraft, M. H. (2002). Math anxiety: Personal, educational, and cognitive consequences. *Current Directions*, *11*(5), 181-185.
- Belavkin, R. V., & Ritter, F. E. (2004). OPTIMIST: A new conflict resolution algorithm for

- ACT-R. In *Proceedings of the Sixth International Conference on Cognitive Modeling*, 40-45. Mahwah, NJ: Lawrence Erlbaum.
- Cadinu, M., Maass, A., Rosabianca, A., & Kiesner, J. (2005). Why do women underperform under stereotype threat? *Psychological Science*, *16*(7), 572-578.
- Chong, R. (1999). Towards a model of fear in Soar. In *Proceedings of Soar Workshop 19*, 6-9. U. of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html.
- Gratch, J., & Marsella, S. (2004). A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, *5*(4), 269-306.
- Hancock, P. A. (1986). The effect of skill on performance under an environmental stressor. *Aviation, Space, and Environmental Medicine*, *57*(1), 59-64.
- Hancock, P. A., & Warm, J. S. (1989). A dynamic model of stress and sustained attention. *Human Factors*, *31*(5), 519-537.
- Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. *Psychological Science*, *11*(2), 93-100.
- Jongman, G. M. G. (1998). How to fatigue ACT-R? In *Proceedings of the Second European Conference on Cognitive Modelling*, 52-57. Nottingham: Nottingham University Press.
- Kieras, D. E., Wood, S. D., & Meyer, D. E. (1995). Predictive engineering models using the EPIC architecture for a high-performance task. In *Proceedings of the CHI '95 Conference on Human Factors in Computing Systems*, 11-18. New York, NY: ACM.
- Kirschbaum, C., Pirke, K.-M., & Hellhammer, D. H. (1993). The Trier Social Stress Test—A tool for investigating psychobiological stress responses in a laboratory setting. *Neuropsychobiology*, *28*, 76-81.
- Lazarus, R. S., & Folkman, S. (1984). *Stress, appraisal and coping*. New York: Springer Publishing.
- Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory: Cross-task prediction of performance in ACT-R. *Journal of Cognitive Systems Research*, *1*, 99-118.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Norling, E., & Ritter, F. E. (2004). A parameter set to support psychologically plausible variability in agent-based human modelling. In *The Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS04)*, 758-765. New York, NY: ACM.
- Quigley, K. S., Feldman Barrett, L. F., & Weinstein, S. (2002). Cardiovascular patterns associated with threat and challenge appraisals: A within-subjects analysis. *Psychophysiology*, *39*, 292-302.
- Reifers, A., Ritter, F., Klein, L., & Whetzel, C. (2005). Modeling the effects of caffeine on visual signal detection (VSD) in a cognitive architecture: Poster presented at "Attention: From Theory to Practice" (A festschrift for Chris Wickens).
- Ritter, F. E., Avraamides, M., & Councill, I. G. (2002). An approach for accurately modeling the effects of behavior moderators. In *Proceedings of the 11th Computer Generated Forces Conference*, 29-40, 02-CGF-100. Orlando, FL: U. of Central Florida.
- Ritter, F. E., Ceballos, R., Reifers, A. L., & Klein, L. C. (2005). Measuring the effect of dental work as a stressor on cognition (Tech. Report No. 2005-1): Applied Cognitive Science Lab, School of Information Sciences and Technology, Penn State. acs.ist.psu.edu/acs-lab/reports/ritterCRK05.pdf.
- Ritter, F. E., & Norling, E. (in press). Extending a BDI architecture to make a better and more interesting team member: The case of JACK to COJACK. In *Cognition and multi-agent interaction: From cognitive modeling to social simulation*. Cambridge, UK: Cambridge University Press.
- Ritter, F. E., Reifers, A., Klein, L. C., Quigley, K., & Schoelles, M. (2004). Using cognitive modeling to study behavior moderators: Pre-task appraisal and anxiety. In *Proceedings of the Human Factors and Ergonomics Society*, 2121-2125. Santa Monica, CA: Human Factors and

- Ergonomics Society.
- Ritter, F. E., Van Rooy, D., St. Amant, R., & Simpson, K. (in press). Providing user models direct access to interfaces: An exploratory study of a simple interface with implications for HRI and HCI. *IEEE Transactions on System, Man and Cybernetics, Part A: Systems and Humans*.
- Salvucci, D. D. (2001). An integrated model of eye movements and visual encoding. *Cognitive Systems Research, 1*(4), 201-220.
- Sarason, I. G., Sarason, B. R., Keefe, D. E., Hayes, B. E., & Shearin, E. N. (1986). Cognitive interference: Situational determinates and traitlike characteristics. *Journal of Personality and social Psychology, 51*(215-226).
- Schoelles, M. J., & Gray, W. D. (2001). Argus: A suite of tools for research in complex cognition. *Behavior Research Methods, Instruments, & Computers, 33*(2), 130-140.
- Tomaka, J., Blascovich, J., Kelsey, R. M., & Leitten, C. L. (1993). Subjective, physiological, and behavioral effects of threat and challenge appraisal. *Journal of Personality and Social Psychology, 65*(2), 248-260.
- Wickens, C. D., Gordon, S. E., & Liu, Y. (1998). *An introduction to human factors engineering*. Addison-Wesley: New York, NY.
- Wickens, C. D., & Hollands, J. G. (2000). *Engineering psychology and human performance* (3rd ed.). Prentice-Hall: Upper Saddle River, NJ.