http://acs.ist.psu.edu/papers/ritterAC02.pdf
In *Proceedings of the 11th Computer-Generated Forces and Behavior Representation Conference*
pp. 29-40, 02-CGF-002. Orlando, FL: U. of Central Florida.

Validating Changes to a Cognitive Architecture to More Accurately Model the Effects of Two Example Behavior Moderators

Frank E. Ritter
School of Information Sciences and Technology
ritter@ist.psu.edu

Marios N. Avraamides
Department of Psychology
marios@psu.edu

Isaac G. Councill
igc2@psu.edu

The Pennsylvania State University
University Park, PA16802
USA

Keywords:
Behavior moderators, cognitive architectures, cognitive models

ABSTRACT: *The success of simulation environments will depend partly on how realistically the models mimic human behavior. While human behavior is affected by various moderators (see Pew and Mavor [28], for an initial list), cognitive models typically do not take into account the effects of many of these moderators. We propose that cognitive models can be augmented to account for such effects by modifying either their knowledge or the parameters of the architecture that they are built with. To provide an example of the two ways in which cognitive models can be modified to capture the effects of behavior moderators, we present an ACT-R model that performs a cognitive task while being affected by the moderators of as anxiety and pre-task appraisal. These changes are validated in a preliminary way by comparison with human data, which shows us where these models can be improved and provides lessons for further work. Most importantly, we argue that more realistic models of human behavior reflecting these moderators and individual differences can be achieved by implementing similar modifications within other cognitive models and by reusing these modifications for an existing architecture as an overlay.*

1. Introduction

In the 1950's Alan Turing suggested that intelligent machines could be built by modeling the human brain. Years later, a fair number of computational models have been built to either answer psychological questions or replace humans in various tasks. Despite their successes at predicting human behavior in a variety of tasks, cognitive models have almost totally ignored the fact that cognitive activity is often moderated by factors that are not directly related to the ongoing task. As a result, cognitive models have rarely included the effects of behavior-moderating factors such as noise, temperature, stress, excitement and so on. Even more rarely have these models been validated or tested.

How to include behavior moderators into models of cognition is a crucial issue if the goal is to build high-fidelity cognitive models. Building such models might not be very important for some types of psychological research. After all, it is a common practice for empirical researchers, who are not interested in behavior moderators per se, to average data across subjects to avoid the contaminating effects of such extraneous variables. Indeed, for most cognitive psychologists, behavior moderators are nothing but confounds that need to be controlled. Similarly, cognitive models aimed to be used as surrogate users in some situations need not include the effects of behavior moderators. In some cases it even seems more desirable to avoid the effects of moderators. A model-tutor [1], for example, does not need to get impatient or angry just like a human tutor might get.

However, models that account for the effects of behavior moderators are essential when the situation requires that the models be so realistic that they, optimally, cannot be distinguished from humans performing the same task. Such situations include military simulations. While in the past, a military exercise would require an extensive amount of human resources, nowadays training can take place with synthetic environments. Cognitive models can populate these environments representing some or all of the entities involved in real combats. Using cognitive models as intelligent agents enables the use of realistic environments for training purposes [28, 34].

Using synthetic environments to train military personnel requires that the cognitive models driving the friendly and enemy entities engage in actions that are expected by human pilots in comparable combat situations. However, until now, the behavior of such agents is not affected by many factors that are very likely to affect the behavior of humans. Factors such as stress, time of day, level of training and so on, influence human behavior in the battlefield [28] and should be built-in any cognitive model that is aimed to be realistic enough to provide high-quality training.

2. Cognitive Architectures and CGF's

Newell [26] defined cognitive architectures as those aspects of cognition that are task-independent and relatively constant. That is, cognitive architectures represent the set of fixed mechanisms that mediate human cognition.

In response to Newell's [26, 27] call for unified theories of cognition, a number of cognitive architectures have emerged. Soar [19], ACT-R [2], and EPIC [18] are the ones used more often to guide the construction of cognitive models. JACK is also being increasingly used [6].

In some cases, models based on cognitive architectures were built for use for training purposes in military simulations. The TacAir-Soar system [35] is a notable example. In TacAir-Soar, cognitive models developed with the Soar cognitive architecture [22, 17] simulate the behavior of military personnel in fixed-wing aircraft missions. TacAir-Soar was successfully used in Stow '97, a large-scale simulation exercise in which up to 3,700 computer-generated forces were involved as both friendly and enemy entities [14]. A next step for all of these models is to include more aspects of human behavior.

3. Behavior Moderators

We are using the term behavior moderators to refer to those variables that affect human performance in a given task. Furthermore, we adopt Pew and Mavor's ([28] chapter 9) taxonomy for distinguishing types of moderators into external and internal. External moderators are inputs from the environment that influence how the person performs the task. These moderators originate outside the person and include physiological stressors (i.e., environmental factors such as temperature, noise, etc), physical and cognitive workload, and fatigue. Internal moderators are those originating inside the person. Examples of internal moderators are the intelligence, expertise levels, expectancies, etc.

It should be noted, however, that dividing moderators into external and internal is not always easy. While, for example, noise and altitude can be easily identified as environmental stressors and therefore placed in the external moderator category, deciding where to place dehydration is not as obvious. Although dehydration is an internal state of a person, it is also the result of external factors such heat and water shortage.

Behavior moderators that affect human behavior in the battlefield are discussed in detail by Pew and Mavor ([28] chapter 9). Therefore, we only present a few examples of moderators that have been studied empirically (Appendix).

4. Modeling Behavior Moderators in Cognitive Architectures

The effects of behavior moderators can be included in cognitive models in two ways. One way is by modifying the content of the model itself and the other by augmenting the cognitive architecture that was used to build the model. The two options exist because models rely on both the specific knowledge that are provided with and the fixed mechanisms contained in the cognitive architecture that are built with [11, 34].

We will briefly discuss the two options and then we will present work that illustrates both approaches.

4.1 Modifying the cognitive model

A number of individual-difference factors can be simulated by varying the amount of knowledge that is put into the model. This represents varying the level of expertise, typically representing the level of education or training. This change can be done either directly (i.e., by providing the model with fewer or different rules) or indirectly by training the model with life-time simulations of different length and type. This approach is mostly applicable for factors that are based on content.

Jones, Ritter, and Wood [15] provide an illustrative example of modifying the model. They were attempting to classify the differences between adult and children's performance on a task. One of the approaches considered was simply to reduce the knowledge available to the model by removing production rules. Another way was to reduce the number of memory elements that could be processed simultaneously by deleting and splitting the production rules that contained a great number of memory elements.

The first change is direct, and models directly differences in knowledge. The second change relates to how information is represented across and within rules. Both of these types of changes can reflect the effects of moderators as well as individual differences.

4.2 Modifying the Cognitive Architecture

Other moderators can be implemented by producing overlays to the architecture. That is, the effect of the factors can be modeled by adjusting the architecture itself. For example, moderators that affect the processing speed of working memory can be simulated by simply adjusting the value of the architectural parameter that corresponds to working memory processing speed.

Jones et al. [15] tried out two such architectural changes to simulate the performance of children on the Tower of Nottingham task. One such change was to limit the number of elements that were active in working memory by increasing the retrieval threshold parameter of the ACT-R architecture. Alternatively, they increased the value of the expected gain noise parameter to influence the strategy choice procedure and therefore increase the stochasticity of the model. The latter manipulation produced the best match to empirically collected data from children.

There are numerous, interesting behavior moderators that can be implemented this way that are of interest to those modeling synthetic forces (see the Appendix for a few examples).

5. An Example Implementation of these Two Approaches

We now turn into an example of how the two ways of including behavior moderators into models of cognition can be applied and tested. We present a cognitive model of serial subtraction that performs the task under the effects of task-appraisal and worry[1].

[1] In fact, the model also includes the effects of caffeine on performance. However, caffeine will not be discussed here. The interested reader is directed to [33]

The serial subtraction task is to start with a large 4-digit number and to repeatedly subtract from it a specified 1- or 2-digit number. For example, the number 1,396 can be the starting number from which the number 7 should be subtracted repeatedly.

Serial subtraction was chosen because it is a task that is often used to measure the effects of stressors on cognitive performance (e.g., [36]). It was not chosen because of surface validity for synthetic forces tasks, although it and related tasks will have to be included in models of air traffic controllers and other operators that do navigation tasks and mental arithmetic as part of their problem solving. As we shall see, we have data on how serial subtraction is affected by various moderators. These moderators can give rise to large individual differences on task performance.

The cognitive model was built using the ACT-R 4.0 cognitive architecture [2]. ACT-R is a production system-based cognitive architecture that combines a symbolic with a sub-symbolic level. The symbolic level represents knowledge as rules (i.e., productions) and as an associative network of interconnected nodes, also called chunks. The distinction between productions and network nodes maps into a distinction between procedural and declarative knowledge. The sub-symbolic level describes the processes that support knowledge, with activation being a central concept at this level.

In ACT-R, production rules are considered the atomic components of thought, that is, they are the most basic unit by which thought processes. Therefore, a production must be selected at each step of performing a task. When more than one production matches the current goal, the systems selects one of them via a process called conflict resolution. In general, the conflict resolution mechanism selects a production by weighing the cost and benefits for each of the matching productions and then selecting the best candidate. When, however, there is a lot of noise in the process, the mechanism can sometimes select less optimal productions. More information on ACT-R is available at act.psy.cmu.edu.

Our ACT-R model performs a serial subtraction task in the same format that is used to study performance under stress. The model's declarative knowledge consists solely of arithmetic facts and goal-related information, and its procedural knowledge by rules to retrieve subtraction results from memory. The task can be performed using two strategies. One strategy is to perform subtractions by counting back from the starting number as may times as indicated by the second number (e.g., 7 times). The other strategy is to retrieve subtraction results directly from memory.

Figure 1 presents the graphical interface of the model. The two main windows are the *Control Panel* and the *Model Behavior* windows. The Control Panel window contains several options for selecting the model's conditions, run control, and some advanced output options. This window allows the model' s moderators to be set.

The Model Behavior window displays aspects of the model's behavior, such as the current result and whether it is a correct or erroneous result, as well as the declarative memory chunks that are being used to solve the problem. Summary statistics (number of attempts, number of errors, and task latency) are also displayed in this window.

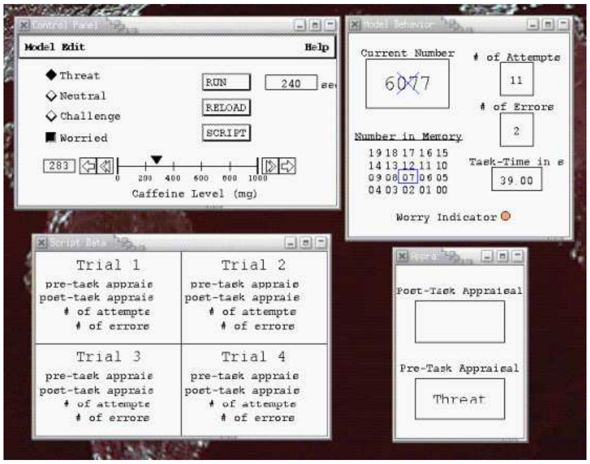

Figure 1: The graphical interface of the serial subtraction model.

5.1 Modifying the architecture

The behavior moderator we chose to include in the serial-subtraction model is task-appraisal [23, 24]. Task-appraisal is considered an internal moderator as it represents an individual's subjective evaluation of a stressful event. Based on the evaluation, appraisal can be of a challenging or a threatening form. A challenging appraisal is made when the individual deems her abilities high enough to cope with the stressful event, while a threatening appraisal arises when the stressfulness of the task is judged to surpass the coping abilities of the individual. Task-appraisals can be distinguished further into pre-task appraisals and post-task appraisals based on whether they are formed before or after the execution of the task.

Empirical evidence suggests a link between the form of task-appraisal (i.e., threatening vs. challenging) and performance on arithmetic tasks such as the serial-subtraction task we use in our model. While threatening appraisals have been associated with fewer solution attempts and poorer performance, challenging appraisals have been related with better performance and more solution attempts than neutral situations [17, 31, 36].

We have attempted to model these results by varying the level of the Expected Gain Noise (EGN) parameter of ACT-R. EGN represents the level of randomness present in the conflict resolution process; that is, the process by which ACT-R decides which rule will fire when more than one rule matches the goal of the system.

This noise parameter has been previously varied to capture the irrationality present in the thought process of children [15]. We have simply set this parameter to a small value (0.1) to model a "clear-head" in the case of challenging appraisal and to a greater value (1.0) in order to provide greater stochasticity in the strategy selection process under a threatening appraisal state.

By varying the default values of the EGN parameter of the ACT-R architecture we have been able to model the effects of the pre-task appraisal moderator. The model also takes into account post-task appraisal. Post-appraisal simply inherits the parameters of pre-task appraisal at the end of each running cycle.

Data suggest that it is not this simple; there is almost a resetting that occurs such that the appraisals are not exactly the same (Task 1 post is not identical to Task 1 pre). This simplification is a working assumption that can be refined later.

As can be seen in the top two sections of Table 1, the model produces a pattern of results that is similar with that reported at the group level in an empirical study using the same serial-subtraction task [36].

As can been see in the upper section of Table 1, the model performs more attempts and it is also more accurate under neutral than under challenging appraisal. This raises questions about ACT-R's default value of the EGN parameter (default value is "Nil").

The model performance provides a very close fit to the empirical data for threatening appraisals. For challenging appraisals it does not attempt as many subtractions as subjects did. As the number of model runs (N) can be increased, the difference is reliable but not a terrible flaw. However, the model produces very precisely the percentage of correct responses out of the total number of attempts for both conditions.

5.2 Modifying the knowledge of the model

The previous subsection described how we modified the EGN parameter of the cognitive architecture to capture the effects of task-appraisal in our serial-subtraction model. However, we believe that some behavior moderators might be built into cognitive models without varying the values of architectural parameters. Instead, the knowledge provided to the model can be modified to incorporate the effects of behavior moderators. As an example, we have used the same serial-subtraction model and we have modified its knowledge in order to simulate the effects of worry on performance.

For the purposes of our current work, we defined worry as the anxiety that is specific to the task to be performed but that is processed in a non-task specific and non-productive way. Because our task is of an arithmetic nature, worry may be equivalent to the term math anxiety that is used by Ashcraft and Kirk [4]. This type of effect is likely to be found in other stressful, anxiety-producing tasks that might be found is synthetic environments or their real-world analogue.

Previous research has associated math anxiety with performance decrements on somewhat complex arithmetic tasks. Particularly, lower accuracy and longer latencies have been observed in solving arithmetic problems that involve a carry operation, such as multicolumn addition [3, 9]. Ashcraft and Kirk [4] suggest that the effect of math anxiety on arithmetic performance is caused by an on-line reduction of working memory resources. In line with Eysenck and Calvo's [8] processing efficiency theory, they propose that math anxiety produces intrusive thoughts that compete with the main task for cognitive resources. Because of this, the amount of cognitive resources that remains available for the arithmetic task is diminished under high math anxiety. Indeed, participants with high levels of math anxiety report the presence of such intrusive thoughts when solving arithmetic tasks (Faust, 1992, cited in [4]).

We have simulated the experience of intrusive thoughts by modifying the knowledge of the serial-subtraction model to enable the model to "worry". Specifically, we added into the model's procedural knowledge a simple rule that can fire any time while the model is performing the serial-subtraction task. In essence, math anxiety here is modeled as a secondary task that is performed concurrently with serial subtraction. It directly implements distracting thoughts. These thoughts thus lead to a decrease in working memory -- due to the serial nature of rule-firing in ACT-R, whenever the worry rule fires, it results into a slowing down of the execution of the subtraction task. In addition to producing an increase in total solution time, the occasional firing of the worry production affects the content of working memory. Because the processing of the main task is halted when the worry rule fires, there is more time for task-relevant declarative information to decay from working memory. The decay of memory information produces more frequent retrievals of inappropriate arithmetic facts. This results in performance that is not only slower, but more errorful as well when the task is performed under high anxiety conditions.

As shown in Table 1, when the model performs the serial-subtraction task with math-anxiety "turned-on", it makes more errors and takes more time.

To the best of our knowledge, there are no available data that examine the effects of math anxiety on performance in a serial subtraction task. Therefore, we have not yet been able to compare directly the performance of our model with human data. Nevertheless, the model seems to capture the effects reported by studies that use multicolumn addition (e.g., [4]). The middle section of Table 1 shows the average performance of our model with math-anxiety turned on and off under different levels of task-appraisal.

Table 1. Comparison of model with human data.

Pre-task appraisal	Number of	CH		TH		NE
Model	Attempts	57.6	>	46.2	<	70.7
	Correct	53.2	>	42.1	<	70.7
(N=100)						
% correct		92%		91%		100%
Model with Worry	Attempts	43.2	>	37.2	<	59.3.
(N=100)	Correct	37.9	>	32.6	<	59.3
% correct		88%		88%		
Tomaka et al. (1993)	Attempts	61	>	46		n.a.
	Correct	56	>	42		n.a.
% correct		92%		91%		

Note Human data taken from Tomaka et al. [36];
< and > denote significant differences at the p<.01 level, N= number of simulation runs, CH=challenging appraisal, TH=threatening appraisal, NE=neutral

6. Conclusions

The cognitive model that we have presented implements the two approaches we suggested for including the effects of behavior moderators. First, we have varied a parameter that is provided by the ACT-R architecture to model performance under different task-appraisals. Second, we have modified the knowledge of the model to simulate the influence of anxiety in general as knowledge that has secondary effects and applied it to a math task.

In both cases we were able to produce the pattern of results that are documented by empirical research. We did this by using very simple techniques that could be easily adopted and used in cognitive models of other tasks by reusing our overlay in ACT-R. This approach of adding a reusable overlay could be applied to other architectures, moderators, and models. We believe that a greater number of moderators should be explored and their effects should be modeled by using reusable techniques that can be shared among modelers.

These changes do not appear to be model specific. As they were created within a cognitive architecture, they should be reusable by other models. We can now create nearly immediate predictions of the effects of task appraisal and worry on all the tasks that have models provided in the ACT-R model library, including driving, phone dialing, and interface use. All the source files for the ACT-R model of serial subtraction and the two types of overlays we discussed here are available at acs.ist.psu.edu/papers/serial-sub/

Including the effects of behavior moderators into computational models of cognition will give power to cognitive modelers as it will provide them with the capability of designing models that can capture more realistically human behavior. The design of high-fidelity models is particularly important for models that can be used for training purposes. In particular, cognitive models that populate synthetic environments in military simulations should incorporate the effects of behavior moderators in order to achieve training environments that closely match the characteristics of real combat. Ideally, entities in synthetic environments should perform their missions should very similarly when they are operated by human users and when they are driven cognitive models.

References

[1] Anderson, J. R., Corbett, A. T., Koedinger, K. R., and Pelletier, R.: Cognitive tutors: Lessons learned. Journal of the Learning Sciences, Vol. 4, pp. 167-207, 1995.

[2] Anderson, J. R, and Lebiere, C.: The atomic components of thought. Mahwah, NJ: Erlbaum, 1998

[3] Ashcraft, M., H, and Faust, M. W.: Mathematics anxiety and mental arithmetic performance: An exploratory investigation. Cognition and Emotion, Vol. 8, pp. 97-125, 1994

[4] Ashcraft, M. H, and Kirk, E. P.: The relationship among working memory, math anxiety, and performance. Journal of Experimental Psychology: General, Vol. 13, pp. 224-237, 2001.

[5] Bartl, C., and Dörner, D.: PSI: A theory of the integration of cognition, emotion and motivation. In F. E. Ritter & R. M. Young (Eds.), Proceedings of the 2nd European Conference on Cognitive Modelling, pp. 66-73, Thrumpton, Nottingham, UK: Nottingham University Press, 1998.

[6] Busetta, P., Rönnquist, R., Hodgson, A., and Lucas, A.: JACK intelligent agents - Components for intelligent agents in JAVA. AgentLink News Letter, *2*(Jan.), 1999, www.agent-software.com/white-paper.pdf.

[7] Cian, C., Koulmann, N., Barraud, P. A, Raphel, C., Jimenez, C., and Melin, B.: Influence of Variations in Body Hydration on Cognitive Function: Effect of Hyperhydration, Heat Stress, and Exercise-Induced Dehydration. Journal of Psychophysiology, Vol. 14, 2000.

[8] Eysenck, M. W, and Calvo, M. G.: Anxiety and performance: the processing efficiency theory. Cognition and Emotion, Vol. 6, pp. 409-434, 1992.

[9] Faust, M. W., Ashcraft, M. H, and Fleck, D. E.: Mathematics anxiety effects in simple and complex addition. Mathematical Cognition, Vol. 2, pp. 25-62, 1996.

[10] Hammond, K., R, Hamm, R., M, Grassia, J., and Pearson, T.: Direct comparison of the efficacy of intuitive and analytical cognition in expert judgment. IEEE Transactions on Systems, Man and Cybernetics, Vol. 17, pp. 753-770, 1987.

[11] Howes, A., and Young, R.: Learning consistent, interactive, and meaningful task-action mappings: A computational model. Cognitive Science, Vol. 20, pp. 301-356, 1996.

[12] Hudlicka, E. Modeling behavior moderators in military performance models (Technical report 9716). Lincoln, MA: Psychonometrix Associates Inc, 1997.

[13] Hudlicka, E., and McNeese, M. C.: Assessment of user affective & belief states for interface adaptation: Application to an Air Force pilot task. Journal of User Modeling and User Adapted Interaction, Special Issue on User Modeling and Adaptation of Affective Computing, 2002.

[14] Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., and Koss, F. V.: Automated intelligent pilots for combat flight simulation. AI Magazine, Vol. 20, pp. 27-41, 1999.

[15] Jones, G., Ritter, F. E, and Wood, D. J.: Using a

cognitive architecture to examine what develops. Psychological Science, Vol. 11, pp. 93-100, 2000.

[16] Jongman, G. M.: How to fatigue ACT-R? In Proceedings of the Second European Conference on Cognitive Modelling, pp. 52-57, Nottingham: Nottingham University Press, 1998.

[17] Kelsey, R. M., Blascovich, J., Leitten, C. L., Schneider, T. R., Tomaka, J., and Wiens, S.: Cardiovascular reactivity and adaptation to recurrent psychological stress: The moderating effects of evaluative observation. Psychophysiology, Vol. 37, pp. 748-756, 2000.

[18] Kieras, D. E, and Meyer, D. E.: An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. Human-Computer Interaction, Vol. 12, 391-438, 1997.

[19] Laird, J. E, Newell, A., and Rosenbloom, P. S.: Soar: An architecture for general intelligence. Artificial Intelligence, Vol. 47, pp. 289-325, 1991.

[20] Koelega, H. S, and Brinkman, J.-L.: Noise and vigilance: An evaluative review. Human Factors, Vol. 28, 465-481, 1986.

[21] Koelega, H. S, Brinkman, J.-L., and Bergman, H.: No effect of noise on vigilance performance? Human Factors, Vol. 28, pp. 581-593, 1986.

[22] Laird, J., E, Newell, A., and Rosenbloom, P. S.: Soar: An architecture for general intelligence. Artificial Intelligence, Vol. 47, pp. 289-325, 1991.

[23] Lazarus, R. S.: Psychological stress and the coping process. New York: McGraw-Hill, 1966.

[24] Lazarus, R. S., and Folkman, S.: Stress, appraisal and coping. New York: Springer Publishing, 1984.

[25] Lovett, M. C., Daily, L. Z., and Reder, L. M.: A source activation theory of working memory: cross-task prediction of performance in ACT-R. Journal of Cognitive Systems Research, Vol. 1, pp. 99-118, 2000.

[26] Newell, A.: You can' t play 20 questions with nature and win. In W. G. Chase (Ed.), Visual Information Processing: Academic Press, 1973.

[27] Newell, A.: Unified theories of cognition. Cambridge, MA: Harvard University Press, 1990.

[28] Pew, R. W., and Mavor, A. S. (Eds.): Modeling human and organizational behavior. Washington, D.C: National Academy Press, 1998.

[29] Mertens, H., W, and Collins, W., E.: The effects of age, sleep deprivation, and altitude on complex performance. Human factors, Vol. 28, pp. 541-551, 1986.

[30] Quigley, K. S., Barret, L. F., and Weinstein, S.: Cardiovascular patterns associated with threat and challenge appraisals: Individual responses across time. Psychophysiology, in press.

[31] Ritter, F. E., Avraamides, M. N., Councill, I., van Rooy, D., Quigley, K. S., Klein, L. C., McNeese, M. D., Stine, M. M., and Rodrigues, I. M.: Pre-task appraisal and caffeine: An architectural overlay for ACT-R. In Air Force Workshop on ACT-R Models of Human-System Interaction, Mesa, AZ, January 2002.

[32] Ritter, F. E., Quigley, K. S., Klein, L. C., McNeese, M. D., Rooy, D. V., Councill, I., Avraamides, M. N., Stine, M. M., and Rodrigues, I. M.: Including the effects of pretask appraisal and caffeine in the ACT-R cognitive architecture: Creating a dynamic architectural overlay, submitted.

[33] Ritter, F. E., Shadbolt, N. R., Elliman, D., Young, R., Gobet, F., and Baxter, G. D.: Techniques for modeling human performance in synthetic environments: A supplementary review. Wright-Patterson Air Force Base, OH: Human Systems Information Analysis Center, (in press).

[34] Ritter, F. E., & Young, R. M.: Embodied models as simulated users: Introduction to this special issue on using cognitive models to improve interface design. International Journal of Human-Computer Studies, Vol. *55*, pp. 1-14, 2001.

[35] Tambe, M., Johnson, W. L, Jones, R., M, Koss, F., V, Laird, J., E, Rosenbloom, P. S, and Schwamb, K., B.: Intelligent agents for interactive simulation environments. AI Magazine, Vol. 16, pp. 15-39, 1995.

[36] Tomaka, J., Blascovich, J., Kelsey, R. M., and Leitten, C. L.: Subjective, physiological, and behavioral effects of threat and challenge appraisal. Journal of Personality and Social Psychology, Vol. 65, pp. 248-260, 1993.

[37] Van Dijk, F., Souman, A., and de Vries, F.: Non-auditory effects of noise in industry.VI . A final field study in industry. International Archive of Occupational Environmental Health, Vol. 59, pp. 133-145, 1987

[38] Webb, W. B, and Levy, C., M.: Age, sleep deprivation and performance. Psychophysiology, Vol. 19, pp. 272-276, 1982.

[39] Webb, W. B.: A further analysis of age and sleep deprivation effects. Psychophysiology, Vol. 22, pp. 156-161, 1985.

[40] Williams, P. S.: Processing demands, training, and the vigilance decrement. Human Factors, Vol. 28, pp. 567-579, 1986.

Acknowledgements

This project was supported by the US Office of Navy Research, award number N000140110547 and by the Space and Naval Warfare Systems Center, San Diego. The views expressed in this article do not necessarily reflect the positions or the policies of the U.S. Government, and no official endorsement should be inferred. Discussions with Roman Belavkin, Wayne Gray, John Anderson, and our colleagues on this project, McNeese, Klein, and Quigley, have greatly benefited our work.

Author Biographies

DR FRANK E. RITTER is one of the founding faculty of the School of Information Sciences and Technology, a new interdisciplinary academic unit at Penn State to study how people process information using technology and to train leaders for the digital economy. Ritter works on the development, application, and methodology of cognitive models, particularly as applied to interfaces and emotions. Ritter is a member of the editorial board of Human Factors, and is on the board of the UK' s Society for the Study of AI and Simulation of Behaviour (AISB). His review (with others) on applying models in synthetic environments will be published as a book this year by HSIAC as a State of the Art Report.

MARIOS AVRAAMIDES is a doctoral candidate in Cognitive Psychology and a research assistant for the School of Information Sciences and Technology at the Pennsylvania State University. He has previously obtained an MS in Cognitive Psychology from Penn State, and a BA in Psychology from the University of Texas at Austin. His current research examines how people update spatial information provided in texts. In the School of IST, Avraamides has worked on several projects building and supporting cognitive models, including helping update the Soar FAQ.

ISAAC COUNCILL is a doctoral student and research assistant in the School of Information Sciences and Technology at the Pennsylvania State University. He has previously received a BA in Psychology from the University of North Carolina at Asheville. Currently, he is developing methodologies that will allow agents created within the Soar architecture to be articulate regarding the reasons for their decisions. Councill is the author of the ACT-R/A/C serial subtraction model and user interface.

Appendix

Behavior Moderator	**Effect on cognition/behavior**
Background Noise	Reduces vigilance and deteriorates attention [37]
Intermittent Noise	Increases sensitivity in vigilance tasks ([21]; see [20] for an evaluative review).
Dehydration	Impairs perceptive discrimination, psycho-motor skills, and short-term memory [7]
Expertise Level	Correlates with ability to perform mental what-if simulations, ease of adopting multiple perspectives, ability to extract relevant information ([10]; Deckert et al., 1994, cited in [28]; Badre, 1978, cited in [28]; see [12] for an elaborate discussion on expertise-level differences).
Anxiety	Narrows the focus of attention, biases interpretation of ambiguous stimuli as threatening [40]. Reduces working memory span and harms performance on math-related tasks [4].
Sleep deprivation	Reduces performance on the Civil Aeromedical Institute's Multiple Task Performance Battery (MTPB), which includes tracking, monitoring of warning stimuli, mental arithmetic, target identification and problem-solving [30].
Altitude	Interacts with sleep deprivation to harm performance on the MTPB. The effects of sleep deprivation are exaggerated in higher altitudes [30].
Age	Interacts with sleep deprivation to affect performance on psychological tasks. Sleep deprivation effects are exaggerated in older populations [38, 39]