



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2004

Vigilance Performance modeled as a Complex Adaptive System

Wellbrink, Joerg; Zyda, Mike; Hiles, John

Monterey, California: Naval Postgraduate School.

<https://hdl.handle.net/10945/37864>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. As such, it is in the public domain, and under the provisions of Title 17, United States Code, Section 105, is not copyrighted in the U.S.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Vigilance Performance modeled as a Complex Adaptive System

Joerg Wellbrink, Ph.D.
Federal Office of the Bundeswehr
for Information Management
and Information Technology
Ferdinand Sauerbruch Strasse 1
56073 Koblenz, Germany
+49(0)261/896-8311
JoergWellbrink@bundeswehr.org

Mike Zyda, Ph.D.
John Hiles
The MOVES Institute
Naval Postgraduate School
Dyer Road 422
Monterey, CA, 93943
831-656-3733
mzyda@nps.navy.mil, jhiles@nps.navy.mil

Keywords:

Cognitive Architectures, Human Performance Models, Vigilance, Complex Adaptive System, Attention

ABSTRACT: *This research has addressed the need for modeling human performance more realistically. It developed a computational model for vigilance performance, embedded in a new cognitive framework that utilizes recent advances in system neuroscience, evolutionary psychology, and complexity theory. A computational model of vigilance is needed—for example to simulate airport security screeners, radar screen operators, sonar operators, and intelligence analysts. The developed model allows the simulation of realistic human errors in monitoring tasks; it can thereby generate surprises in simulation programs that might show weaknesses of security systems. After studying human performance especially vigilance, experiments were conducted to establish correlations between personality and performance and to collect data for calibrating and validating the model. The robust model shows a reasonable range of individual behaviors and represents a tool well suited for gaining insights into vigilance theories. The insights can potentially be used to improve existing theories and monitoring procedures, minimizing errors that might lead to catastrophic outcomes*

1. Introduction

This research represents a multidisciplinary effort that led to a Ph.D. in MOVES (modeling, virtual environments and simulation) at the Naval Postgraduate School's MOVES Institute in Monterey, California. The disciplines of complexity theory, computer science, psychology, social science, and operations research were amalgamated to yield a computational treatment of reduced human performance—specifically, vigilance performance—as a complex adaptive system, or CAS. The work falls under the rubric of modeling and simulation.

This paper summarizes the dissertation produced [1] and discusses research motivation, approach, and overall contribution to modeling and simulation, focusing on design considerations, especially the use of loosely coupled components that promote interoperability and reuse of code. Event graphs were used to describe the design of the multi-agent system, as they facilitated discussion between modeler and expert and enabled us to transform psychological into computational models.

1.1. Motivation

The attacks of September 11, 2001, showed, not for the first time in Western history, a need for simulative models that are capable of generating or revealing surprises, unintended consequences, and blind spots [2](Smith 2002). One axiom of this research is that the modeling of surprise demands realistic simulation of reduced human performance, which is a primary cause of less-than-ideal behavior. In such a simulation, manlike errors should lead to surprising or unexpected outcomes—for example, through a cascade of errors due to laxity or irrationality that analysts did not predict.

The authors will discuss their assumptions and show a proof-of-concept implementation of their main hypothesis: *Human performance can be modeled as a complex adaptive system.*

The National Research Council's report in 1998[3], and follow-on research [4], [5] on modeling human and organizational behavior described the status of cognitive modeling, broadly indicating strengths and weaknesses of current models. The NRC recommended "... continued efforts to improve the quality of existing modeling approaches that will result in architectures as yet unconceived." [3,p.111]

Our research focused to overcome two difficulties in modeling realistic human performance:

- 1) Current cognitive models generate neither adaptation nor emergent behavior, which are essential features in the modeling of human behavior.
- 2) Current cognitive models do not model the individual's reduction in performance over time, thus producing homogenous, predictable, and brittle modeled behavior.

We hypothesize these weaknesses can be overcome by using CAS theory as the foundation for a new cognitive architecture, and that use of multi-agent systems is the ideal implementation for a CAS. We anticipate the advantages of autonomy, emergence, flexibility, adaptivity, dynamism, robustness, and self-explanation in modeled behaviors.

1.1 Research Scope and Approach

It is beyond the scope of this research to design, implement, and validate a new cognitive architecture; such work requires the cooperation of many research communities over a number of years (the States Operators and Rules [SOAR] community, for example,

has been operating for more than two decades). But even a partial implementation of a cognitive architecture should be embedded in a framework using proper design techniques, so that the model can be enlarged at any time.

We used psychological models (i.e., the information-stage processing model) as blueprints for the computational model. While we will not describe the entire design of the agent-based model (called the reduced-human-performance model, or RHPM), we will show an event graph of the short-term memory store to manifest our design approach.

Our implementation focuses on cognitive resource modeling with respect to vigilance tasks. This is meant as a proof-of-concept implementation and should add validity to our hypothesis. Considerable future work is required to implement the full framework.

Figure 1 shows the approach for this research, an approach cited by McKelvey as a "third way" of doing science, namely, by fitting an agent-based model to a real system.[6]

We construct a model that produces humanlike experimental data (including variability) by using the appropriate theories and models for the social model (in our case, a man or woman). We collect data in human experiments and configure the fit of the computational and social models using genetic algorithms to generate a parameter setup that meets human data.

Of course, this in itself would not be sufficient to validate the model; we also test it by maintaining the parameters but introducing unfamiliar scenarios, then measuring output. Only if the model's output approximates human results can we claim our model is robust. If so, we hypothesize that the model can be used a surrogate of the social system, producing testable predictions and insights into the problem domain.

We strive to create a cognitive model that can identify weaknesses in organizations by modeling the effects of reduced human performance. Because the decisions and policies of imperfectly rational actors can be exploited towards some malevolent goal, agents must be forced to operate and decide with imperfect knowledge and restricted cognitive resources.

Computational Psychology

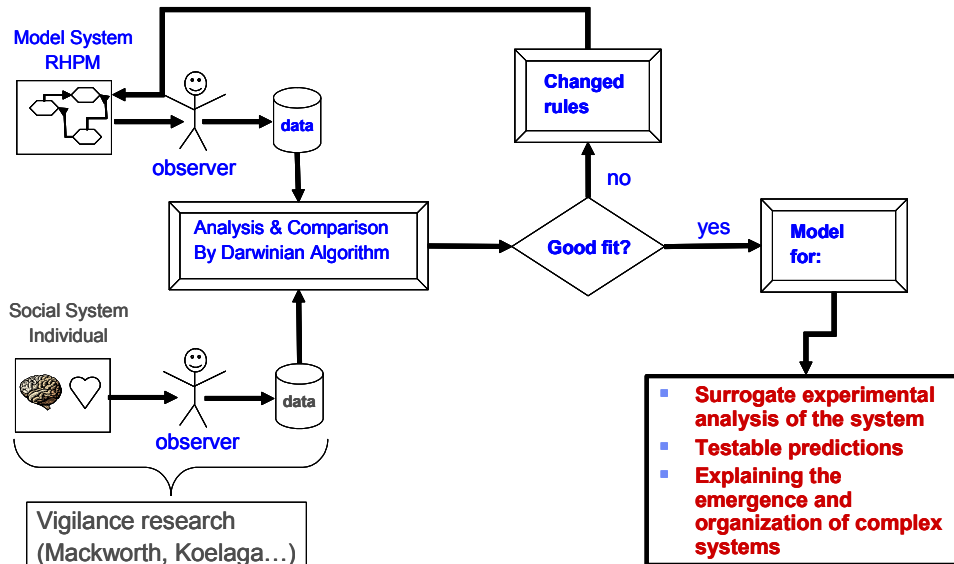


Figure 1 Fitting a Model to a Real System (adapted from [6])

2. Human Performance and Complex Adaptive Systems

This chapter is a cursory overview of the definition and features of human performance and complex adaptive systems. [1] provides a minute discussion for the interested reader.

2.1 Defining Human Performance

Our definition of human performance is derived from performance psychology: “Human beings are born to perform. In a broader sense, we perform every time we engage in a goal-directed activity.” [7, p.1]

While cognitive models assume ideal behavior, real human performance routinely suffers breakdown and failure. The possibility of error, which plays a major role in accidents such as car or airplane crashes, should always be predicated. At the same time, the possibility of enhanced performance is also an integral part of the picture; thus our framework and the cognitive model take performance variability into consideration.

2.2 Defining Complex Adaptive Systems

There is no standard definition for complex adaptive systems. CASs provide insights into a problem domain,

but these insights do not necessarily forecast certain behaviors or behavioral ranges. Thus CASs do not function as weather-forecasting tools; rather, they show possible interactions, producing emergent behavior that might occur at some point.

Hence we formulated a working definition that enables us to discern whether a system is a CAS:

A CAS consists of many autonomous agents that act in parallel with decentralized control. The nonlinear interaction between these agents leads to adaptive, emergent behavior. The agents are organized in dynamically rearranging structures that achieve equilibria while never maintaining one particular equilibrium. In many systems, the CAS builds an internal (implicit or explicit) model of the future. There is a strong sense of path dependency in CASs, built upon the interaction of autonomous active entities and the nonlinearity of their mutual impact. As the system’s structure evolves, it incorporates information that can serve as the foundation for new interaction and behavior.

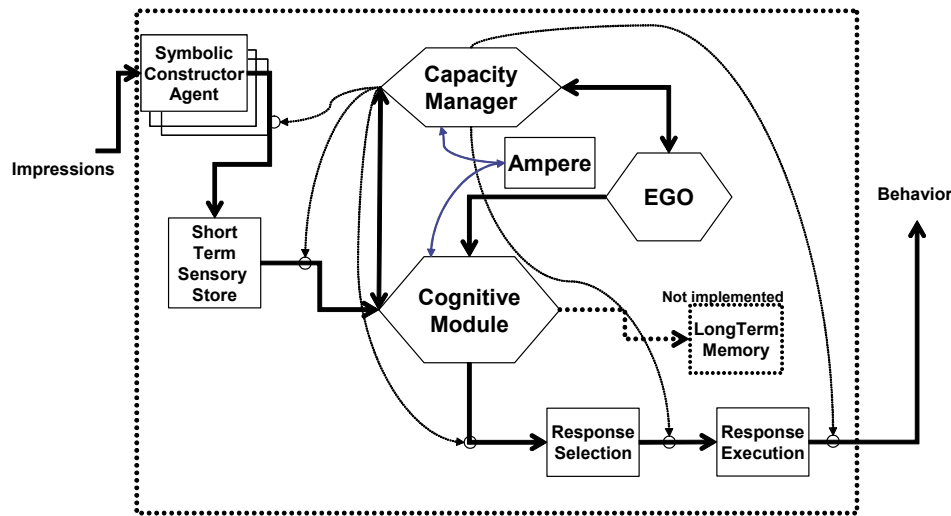


Figure 2 Design of the Reduced Human Performance Model

3. Design of RHPM

3.1 Description of Main Components

Figure 2 shows the principle design of the RHPM. This is a transformation of the human information-processing model, after Wickens [8]. Symbolic constructor agents (SCAs) perceive information (impressions) and relay them to the cognitive module, which holds a symbolic interpretation of the environment. The symbolic representation depends on the inner state of the system. A highly aroused person, for example, may perceive background noise as a threat, whereas someone used to the noise might not even register it. The cognitive module is a multi-agent system itself and contains several diverse composite agents. This module coordinates intentions with actions and creates behavior. The capacity manager is a multi-agent system, based on Wickens' multiple-resource model, that determines the current arousal level and introduces noise into the system. It can also interrupt or disturb transitions and access the cognitive module to suppress processes. The impression stream is analyzed and, if appropriate, a capacity decrease is initiated. The capacity manager also evaluates capacity demands of planned activities, determining whether they will be executed. The Ampere component conducts all computations, based on the idea that CapacityManager provides resources analogously with resource distributions in electrical circuits.

The Ego module represents personality, emotions and goals. GoalAgent deals with conflicting goals and actions, using a weighting scheme based on personality traits to determine action in the face of opposing goals. An example might be an airport-security screener

before a long line of travelers. He wants to decrease the queue, but also to find weapons. At one point he detects an item he cannot identify but which does not look like a weapon. What will he do? The answer will lie in his personality.

Personality plays a major role in human performance. To return to our example, evidence shows that introverts outperform extraverts when it comes to routine screening [9], [10], [11]. Evidence also indicates that extroverts outperform introverts under conditions of high arousal [12], and that in such a case an extrovert will probably examine the item. A realistic cognitive model should capture this interplay between conditions and personality. The conflict and resource managers use personality traits to set parameters that determine whether the screener inspects the item or lets it pass.

3.2 Design Decisions

Although this research focuses on building a computational model of vigilance, the big picture, of embedding the model in a cognitive framework, is of ultimate concern. This research suggests that a future cognitive architecture should consist of interoperable subcomponents.

Loosely coupled components (LCCs) are currently under investigation at the Naval Postgraduate School and should assist modelers in rapidly prototyping and utilizing software components as building blocks; fittingly, the latest research paper describes these components as LEGOs (listener-event-graph objects). "The name is also a metaphor for how complex models can be built by rapidly linking simpler component sub-models" [13, p.732]. Since 1995, many projects have

used this design strategy successfully in the domain of military simulation (see [14],[15],[16],[17]).

LCC design philosophy is based on the “observer” design pattern of the “gang of four” [18]. LCCs use weaker criteria on the coupling mechanism and therefore call it the “listener pattern.” For example, the observer pattern uses interfaces for attaching and detaching observer objects. LCCs use no coupling between components, though initially they used mediators. This research employs software routers embedded in the listener pattern, in a manner analogous to networking.

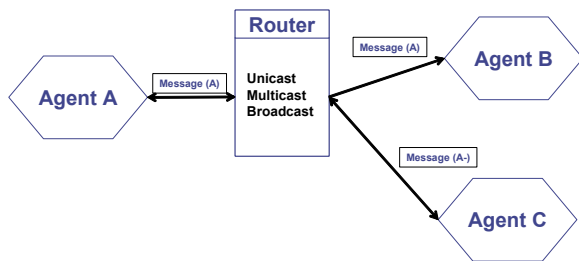


Figure 3 Message Routing between Agents

Figure 3 illustrates why the structure is considered loosely coupled: Agent A could be any agent in a simulation system that provides a message (A) as output. The agent sends this message out even if nobody listens, continuing to work independently of acknowledgement (in contrast to the observer pattern). In this example, the router listens to Agent A’s messages. It can use several information modes: unicast (one-to-one connection); multicast (one-to-many); or broadcast (one-to-all). The router can also filter the information and transport the entire message (A) to Agent B and a reduced message (A-) to Agent C. Unidirectional arrows indicate the listener pattern: the arrowhead points to the listener and the shaft connects to the sender. The message object can be formatted using typical agent-communication protocols (e.g., using Knowledge-Queering Modeling Language [19]).

Bidirectional arrows show that the entities communicate two ways, acting as receivers or senders. The listener pattern seems very apt for a next-generation cognitive architecture. Some advantages include:

- An architecture based on the listener pattern is dynamically extensible. Its structure can be changed during run-time, which is essential for CAS modeling.
- Components can be exchanged at any time (event run-time) without creating a new system.
- The listener pattern facilitates reuse of software.

- The pattern lends itself to a plug-and-play approach, like exchanging hardware components via USB.

To exploit these benefits, we designed our research using the listener pattern for most components of the simulation system.

3.3 Using Event Graphs

Event graphs depict discrete-event simulation models. Also known as “simulation graphs,” they have a primitive design, with a single type of node and two types of edges with up to three options. Despite this simplicity, event graphs are extremely powerful. The event graph is the only graphical paradigm that directly models event-list logic. There is no limitation to the event graph’s ability to create a simulation model for any circumstance. Their simplicity and extensibility make them an ideal tool for rapid construction and prototyping of simulation models [20].

Event graphs contain several important features:

1. They visually describe the logic of a model.
2. Event graphs are simple and extensible.
3. They are ideal for rapid prototyping.
4. They help identify important state variables.
5. Event graphs help anticipate problems with distributed (simultaneous) events.
6. They help streamline a model by eliminating unnecessary event routines.

There is evidence that event graphs are equivalent to stochastic Petri nets [21]. Petri nets have a decided reachability problem. This can be expensive. “In the worst case, the time and memory (computational complexity) needed to analyze a Petri net grows exponentially with the size of a net” [22, p.1402] However, since a Petri net can be analyzed, by induction so also can an event graph. The following references explain the features of Petri nets and event graphs in more detail: [23],[24],[25].

In cognitive modeling, event graphs have other key advantages:

- Event graphs are easy to read and there is no need to go into implementation details to explain how the model works. They provide a transparent look into the model, avoiding black boxes.
- Since event graphs are extensible, it is easy to create new relationships or simulation entities (e.g., introducing emotions into a cognitive architecture), thus allowing flexibility in the mathematical model.

ShortTermSensoryStore

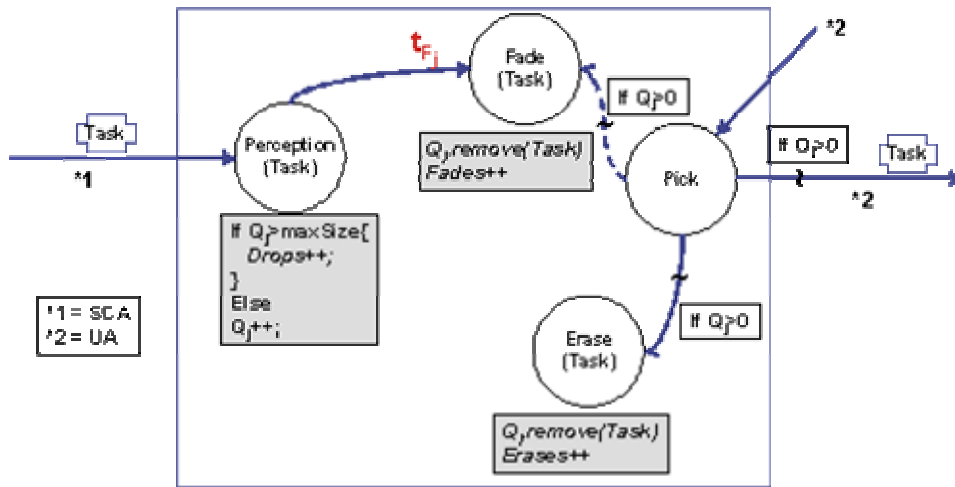


Figure 4: Event Graph for the Short Term Sensory Store

Example: Short Term Sensory Store (STSS)

At first the parameters of an entity (an object or agent) are defined; these parameters can be changed or varied. They represent the buttons used to calibrate a model.

1. Parameters:

$t_{F_j} \in \mathcal{R}^+$, where $j \in B$, fadetime for a specific buffer

$\kappa_j \in \mathcal{Z}^+$ where $j \in B$ storage capacity;

The important parameters for STSS are:

- the time when a stimulus fades away (a positive real number),
- and its storage capacity (positive integer).
- The index j belongs to a set B (the set of all agents or objects that use a capacitated queue).

Next the system's state variables in this entity are defined:

2. State Variables

$Q_j \in \mathcal{Z}$ where $j \in B$ num. of elements in FIFO queue

STSS only contains a state variable for its queue. Q has a positive integer value. Next the statistical variables are declared.

3. Statistic Variables

$Fades \in \mathcal{Z}$ number of faded observations

$Drops \in \mathcal{Z}$ number of dropped observations

$Erases \in \mathcal{Z}$ number of erased (picked) observations

Fades are observations for which the time for storage has expired. Drops are observations that couldn't enter the system due to its limited capacity. Erases are

stimuli that made it into the system as percepts. Finally, the event graph shows the logic of STSS inner workings.

4. Event Graph

See Figure 4.

5. Event-Graph Description

STSS receives a task object from the SymbolicConstructorAgent SCA (perception event) and checks whether its storage capacity is sufficient. If so, it stores the observation; if not, it drops it. A time ticker is instantiated on this task. If it expires, the observation fades away (fade event). The pick event interrupts (dashed line) the fade event, given a task in the queue. It then erases this task from the queue and relays it to the UpdateAgent (UA).

This brief example shows how a system's components can be described in a rigid way. Its design is transparent and facilitates discussion. The implementation's particulars can be changed. In our case, we implemented (after intense discussion) the STSS as a FIFO queue; however, a LIFO queue or prioritized FIFO queue are alternatives that can be tried.

5. Experiments and Results

Laboratory data on vigilance performance of some fifty participants was collected in four experiments using performance test-and-evaluation software (SynWinGenerator from Activity Research). The

experiments lasted thirty minutes and varied workloads (experiment 1: low workload; experiment 2: high workload; experiment 3: switch from low to high, and back to low (LHL) in ten-minute intervals; and experiment 4: switch from high low high (HLH) in 10-minute intervals). See [1] for a more detailed description of the methodology. Because we expected that time would influence performance, we collected reaction time, hit rate, and false-alarm rate every ten minutes during the thirty-minute trials.

The tasks of the SynWinGenerator were simulated and used as input events for RHPM; for example, the auditory task was simulated by generating random numbers from two normal distributions with different means. RHPM has a decision process based on signal-detection theory to decide whether the stimulus is a signal. The computational output of a specific parameter configuration was compared with the experimental data (hit rates, false-alarm rates, and reaction times) of low- and high-workload conditions after ten, twenty, and thirty minutes. A fitness value for the specific parameter configuration was computed using the least square method for every MOP. RHPM was calibrated using data from low- and high-workload conditions. We used a genetic algorithm to find a robust parameter setup for the model; this configuration was used to evaluate whether RHPM could match experimental data without recalibration. These runs were called validation runs because the model's output was validated to experimental data that was not used for calibration beforehand. The performance outcomes from the RHPM show a very close match with the human experimental data for most MOPs.




Figure 5 Comparison of Reaction Times

Figure 5 shows that the model's output for a previously unknown scenario closely matches human data. Because the differences were statistically insignificant, we could not conclude that the data from the model stems from a distribution different from the experimental data.

Overall, RHPM produced matches in sixteen of eighteen possibilities. The differences in two cases can be traced to theoretical implications such as warm-up and perceptual-learning effects. Another important question was whether the model would be brittle or produce variability.




Figure 6 False-Alarm Variability RHPM

Figure 6 shows that RHPM is a stochastic model. The curves are individual (random) outcomes that would converge to the mean MOEs. RHPM produces outcomes with almost no false alarms (yellow triangle) or with an unusually high number in the last ten minutes (red diamond).

Another important research question to be answered was whether RHPM would have enough variability to mimic human variation.

Standard Deviation LHL		
	FalseAlarm RHPM	FalseAlarm Human
10	1.21	2.04
20	1.29	2.76
30	1.32	2.72
	Misses RHPM	Misses Human
10	1.38	2.30
20	1.33	1.75
30	1.88	1.92

Table 1 Standard Deviation Comparison

Table 1 shows a comparison of the MOEs' standard deviations of RHPM and human subjects in the LHL condition. While human data is more dispersed, the differences, especially in misses, are small. Together with the good approximation to the mean MOEs, RHPM is shown to be neither mechanistic nor brittle. The random-number generation allows repeatability of runs and also pseudo-randomness, making it very difficult to predict the next outcome precisely. We can

introduce even more variability by using more random numbers in different situations.

The validation-run results of RHPM match our expectations. The model showed reliable behavior during “normal” simulation runs. It also generated insights into applying theories to the vigilance-decrement phenomenon (i.e., the importance of the detectability of a signal). RHPM could be validated against two previously unseen scenarios. RHPM can also demonstrate the pitfalls of certain theories. For example, it is well known that an increase in signal probability leads to an improvement of miss- and false-alarm rates. Signal-detection theory does not address this phenomenon; consequently, RHPM increases, rather than decreases, the miss rate. However, by looking at the design of the model and how different modules work together, possible avenues to improving model performance become apparent. These improvements could lead to improvements in theory.

The multiple-resource model’s implementation influences reaction time in a normal case. As soon as the main energy level decreases, error rates increase to the point that the model no longer processes signals. This effect should be encouraged by, for example, introducing other stressors to the model than time on task or change of workload.

Some experimental results also indicated the need to introduce perceptual learning (better distinguishing signals) into the model. These mechanisms are not well documented in vigilance research, and further research is required to introduce perceptual learning to a computational model of vigilance.

One issue in the proof-of-concept implementation is how well it would fit scenarios from a different experimental setup. It happens that the time required to program new scenarios is minimal, and due to the open architecture, they can be connected swiftly. However some task characteristics (e.g., signal salience) should be adjusted before RHPM should be used.

6. Summary

This research suggested a new cognitive model that simulates individual reduced human performance.

Using a discrete-event simulation with an event-graph design opens a cognitive architecture to design discussions with domain experts (in our case, human-factor specialists and psychologists).

There is convincing evidence that a paradigm shift in human-behavior modeling to take vagary into account is suggestive. The proposed framework for a next-

generation cognitive architecture has shown advantages in terms of robustness and adaptivity. The open and flexible architecture shows a possible path of cooperation between modelers. The implemented parts of the cognitive framework show their contribution by modeling the challenging problem of vigilance decrement.

RHPM has been validated with quantitative and qualitative analyses. Limitations to the model and possible improvements have been identified. These improvements should occur in cooperation with vigilance researchers.

This research suggests directions towards improving signal-detection theory. The model’s behavior and the theory’s predicted behavior are coherent. However, differences in outcome between human experimentation and RHPM lead to the assumption that there are perceptual-learning effects in signal detection that affect sensitivity. RHPM can fit the data better with a sensitivity increase based on number of signals. If the number reaches a threshold, detection seems to become easier. The next step should be human-vigilance experiments that try to find a relationship between signal quantity and sensitivity increment.

Two further achievements deserve mentioning:

- 1) RHPM appears to be the first computational vigilance model composed of multi-agent systems.
- 2) The implementation of Wickens’ multiple-resource model also seems to be a first try for a computational model on multiple cognitive resources.

Thus, this research contributes to the modeling of human behavior as well as to cognitive psychology, especially vigilance research. It is difficult to make comparisons with current cognitive architecture since this research has not the same level of sophistication. However, it has shown potential by modeling an important phenomenon that has been inadequately explored. It also showed that a multi-agent system based on CAS theory can be used to produce desired results that are within range of human performance.

6. References

- [1] Wellbrink, Joerg. (2003). "Modeling Reduced Human Performance as a Complex Adaptive System." Dissertation MOVES Institute. Monterey, California, Naval Postgraduate School
- [2] Smith, R. (2002). "Counter Terrorism Simulation: A New Breed of Federation." Simulation Interoperability Workshop Spring, Orlando, FL..
- [3] Pew, R. W. and A. S. Mavor (1998). Modeling human and organizational behavior: application to military simulation. Washington D.C., National Academy of Science.
- [4] Ritter, F. E., N. R. Shadbolt, et al. (1999). "Techniques for Modelling Human Performance in Synthetic Environments: A Supplementary Review." Nottingham, UK, ESRC Centre for Research in Development, Instruction and Training: 1-92.
- [5] Tenney, Y. J., D. E. Diller, et al. (2003). The AMBR Project: A Case-Study in Human Performance Model Comparison. Behavior Relation in Modeling and Simulation (BRIMS) 2003, Scottsdale, Arizona.
- [6] McKelvey, B. (2000). "Complexity Theory in Organization Science: Seizing the Promise or Becoming a Fad?" *Emergence* 1(1): 5-32.
- [7] Matthews, G., D. Davies, et al. (2000). Human Performance Cognition, Stress and Individual Differences. East Sussex, UK, Psychology Press.
- [8] Wickens, C. (1992). Engineering Psychology and Human Performance. New York, Harper Collins.
- [9] Methot, L. L. and B. E. Huitema (1998). "Effect of Signal Probability on Individual Differences in Vigilance." *Human Factors* 40(1): 102-110.
- [10] Gusev, A. N. and S. A. Schapkin (2001). Individual Differences in Auditory Signal Detection Task: Subject-Oriented Study. Fechner Day 2001. Proceedings of the Seventeenth Annual Meeting of the International Society of Psychophysics, Leipzig Germany, Lengerich.
- [11] Schapkin, S. A. and A. N. Gusev (2001). Effect of Hemispheric Asymmetry on Performance in Auditory Vigilance Task. Fechner Day 2001: Proceedings of the Seventeenth Annual Meeting of the International Society of Psychophysics, Leipzig Germany, Lengerich.
- [12] Matthews, G., D. Davies, et al. (1990). "Arousal, Extraversion and Individual Differences in Resource Availability." *Journal of Personality and Social Psychology*.
- [13] Buss, A. H. and P. J. Sanchez (2002). Building Complex Models With LEGOs (Listener Event Graph Objects). Winter Simulation Conference.
- [14] Arntzen, A. (1998). Software Components for Air Defense Planning. Operations Research. Monterey, Naval Postgraduate School.
- [15] Bohmann, W. (1999). STAFFSIM, An Interactive Simulation For Rapid, Real Time Course of Action Analysis by U.S. Army Brigade Staffs. Operations Research. Monterey, Naval Postgraduate School
- [16] Le, H. B. (1999). Advanced Naval Surface Fire Support Weapon Employment Against Mobile Targets. Operations Research. Monterey, Naval Postgraduate School.
- [17] Schrepf, N. (1999). Visual Planning Aid for Movement of Ground Forces in Operations other than War. Operations Research. Monterey, Naval Postgraduate School.
- [18] Gamma, E., R. Helm, et al. (1995). Design Patterns: Elements of Reusable Object-oriented Software. Indianapolis, IN, Addison-Wesley.
- [19] Flores-Mendez, R. A. (1999). "Towards a Standardization of Multi-Agent System Frameworks." *ACM Crossroads Student Magazine* 5(4).
- [20] Buss, A. H. (1996). Modeling with Event Graphs. 1996 Winter Simulation Conference.
- [21] Schruben, L. and E. Yucesan (1994). Transforming Petri Nets into Event Graph Models. Winter Simulation Conference 1994 (WSC '94), Lake Buena Vista, FL, USA.
- [22] Ralston, A., E. D. Reilly, et al. (2000). Encyclopedia of Computer Science. London, UK, Nature Publishing Group.
- [23] Schruben, L. (1992). Graphical Model Structures for Discrete Event Simulation. Winter Simulation Conference 1992, Arlington, VA, USA.
- [24] Balbo, G., J. Desel, et al. (2000). Introductory Tutorial Petri Nets. 21st International Conference on Application and Theory of Petri Nets, Aarhus, Denmark,.

Author Biographies

JOERG WELLBRINK is a German Army officer who completed the modeling, virtual environments and simulation (MOVES) Ph.D. program at the MOVES Institute of the Naval Postgraduate School (NPS) in Monterey, California, in September 2003. He received his M.S. in electrical engineering from the German Armed Forces University, Munich, in 1985 and in 1998 received his M.S. in operations research at the Naval Postgraduate School. His research covers artificial-life techniques, human-performance modeling, cognitive modeling, and complexity theory. He currently works as a modeling and simulation scientist for the Federal Office of the Bundeswehr for Information Management and Information Technology.

MICHAEL ZYDA is the director of the MOVES Institute at NPS. He supervised the presented work on modeling reduced human performance. His research interests include computer graphics, large-scale, networked, 3D virtual environments, agent-based simulation, modeling human and organizational

behavior, interactive computer-generated stories, computer-generated characters, video production, and entertainment/defense collaboration. He received his B.A. in bioengineering from the University of California, San Diego, La Jolla, his M.S. in computer science from the University of Massachusetts, Amherst, and his DSc in computer science from Washington University, St. Louis.

JOHN HILES is the technical director for computer-generated autonomy at the MOVES Institute and a research professor at NPS. He advised the principal author on the presented research. Professor Hiles's focus is in using agent-based simulation and entertainment technology to model human and organizational behavior. He has a B.A. in creative writing from the University of California, Santa Barbara.