Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS04), 758-765, New York, NY: ACM Press, which will be available from the ACM web site: http://doi.acm.org/

Towards Supporting Psychologically Plausible Variability in Agent-Based Human Modelling

Emma Norling
Department of Computer Science and Software Engineering
The University of Melbourne
norling@acm.org

Frank E. Ritter
School of Information Sciences and Technology
The Pennsylvania State University
frank.ritter@psu.edu

Abstract

We describe the initial steps in developing an agent-based cognitive architecture designed to support psychologically plausible human variability. The new architecture, COJACK, is based on JACK, a BDI-based agent language. It will constrain the agents to reason and act in a psychologically plausible manner. Their information processing will be adjusted by a set of parameters that moderate the agent's reasoning and actions, combined with a set of guidelines for developing plans and beliefs for the agents. This set of parameters will also support varying the agents' performance, both in terms of differences across agents as well as differences that arise within an individual due to internal and external factors. We conclude that other architectures will want to include a similar set, including representing a body, its interaction with the environment, and the passage of time.

1. Introduction

Put two people in the same situation, and chances are they will not do the same thing. Put two *models* of people into the same simulated world, and the rare cases that they will not do the same thing occur when the models have been provided with different knowledge. Human variability arises from more than just different knowledge, however, and indeed sometimes helps account for the variations in knowledge between two people. Human variability can also lead to different behaviours within an individual, when internal and/or external factors moderate their behaviour.

This paper describes some initial steps towards the development of an agent architecture that will support psychologically plausible models of human variability, accounting for individual differences and moderator effects. We draw upon studies in psychology, physiology and cognitive science to develop a framework that captures many of the mechanisms that are known to influence behaviour. Some existing architectures go some way towards this, and indeed we draw upon parameter sets used by other architectures. While there is some work towards incorporating behaviour moderators into existing architectures, as discussed in Section 3.4, we believe that our proposed architecture, COJACK, is one of the first that has been deliberately designed to support the modelling of a wide range of behaviour moderators.

The architecture itself will take the form of a "cognitive overlay" for JACK Intelligent Agents, a BDI-based agent programming language [1]. The cognitive overlay will introduce a range of parameters that will influence the reasoning and actions of JACK agents, constraining their performance to be within psychologically plausible bounds. These parameters influence both mental and physical aspects represented in the overlay. In most cases, as well as having an initial value which can vary from one agent to another, the parameter can also vary over time. The variations in initial parameter values will account for individual differences, whereas variations over time (which could arise from things such as fatigue, boredom, or ingestion of a drug such as caffeine) will account for the variations that occur within an individual.

Such an architecture will lead to more realistic simulations, as the agents within the simulation will display more of the variations that occur naturally in people. It will be possible to study, for example, the suitability of a procedure for a range of people, rather than just the average person, or the effects of different combinations of individuals in team performance. It will also allow the study of the effects of behaviour moderators on the behaviour of individuals and groups of people, by simulating how a particular moderator influences the reasoning, perception, action and physiology of an individual, and subsequently, their behaviour.

The paper first looks at individual differences and the types of parameters that are needed to capture them. It then moves on to consider behaviour moderators, and how to model them using the same set of parameters. We then discuss the development of the architecture: how these parameters can be applied to JACK. Finally, we present the implications for other architectures.

The importance of the parameters we include here is dependant on the tasks of interest, the grain size of the analysis, and the uses to which the model will be applied. In our situation, the parameter set is being developed to support cognitive tasks and work in military synthetic environments; the parameter set would likely have a different flavour if it was being developed with, for example, synthetic characters for entertainment in mind.

2. Individual Differences

Knowledge is a key component giving rise to differences between individuals, but there are many other factors that also contribute. Our initial survey has identified sixty parameters that would be particularly useful when modelling individual differences [16]. We have broadly classified these parameters into four groups: cognition, perception, action and physiology. The set of parameters that we have selected is not exhaustive — it would certainly be possible to find more that influence human reasoning and action — but we believe that the our set is sufficient to capture the main elements that contribute to human variability.

Most parameters have a default value, giving rise to “average” behaviour, as well as a range of plausible values, between which the bulk of the human population lies. There are, of course, co-dependencies between parameters (e.g. a high value for processing speed may correlate with working memory capacity). We do not attempt to capture these relationships here, but they will become important when creating a “random” individual using COJACK.

We do not attempt to examine each of the sixty identified parameters here, but discuss each of the main parameter groups including examples.

2.1. Cognition

The parameters that we have selected to capture variability in cognition are largely taken from the parameter set used by ACT-R 5.0 [2], a prominent cognitive architecture used in psychology. While still not perfect, and complex because it models human performance fairly well, the ACT-R memory and processing system and its associated parameters have been validated across use in numerous tasks and data sets and has shown itself to be able to account for a wide range of behaviour. For example, its working memory capacity parameter has been used to predict strategy choice between strategies using different amounts of working memory [9].

In addition to the parameters used by ACT-R, we have identified some additional high-level parameters that affect cognition, outlined in Table 1. These parameters are actually labels encompassing properties that emerge out of lower-level cognition, but the low level constructs taken from ACT-R may lack the complexity to accurately represent them. It will be important to ensure that these parameters are represented in the architecture, either implicitly, through a combination of low level parameters, or directly as separate parameters.

2.2. Action

Existing models that predict human behaviour including motor output have typically done so at the level of hand movements and typing, including both speed and accuracy. For many simulations, however, the agent must specify larger scale behaviour, and there must be some way of mapping between the granularity provided by the simulation environment and the data that is available. This behaviour must be aligned with the simulation. For example, if the agent must specify movement in the world via two parameters — walking speed and direction — but if COJACK represents walking in terms of stride length and direction, there must be a means for translating between the two.

Most action types have two components — speed and accuracy — but it should also be noted that speed particularly is not constant, with a degree of variance under standard conditions that can itself be affected by moderators.

Action parameters are important for modelling behaviour in games and simulations. They influence agent grouping and dispersion. They are also important for modelling fatigue. The parameters will be influenced by fatigue, but more importantly fatigue increases through extended motor output.

The most commonly modelled parameters related to action are related to mouse and keyboard inputs. While we expect that COJACK will contain a much wider range of parameters relating to action designed to suit the military simulation applications of the sponsor, we will nevertheless include small scale action parameters to help test the architecture. A number of previous studies have looked at simulated eyes and hands [10, 15, 18], providing data that can be used for comparison in COJACK.

2.3. Perception

The majority of synthetic environments provide most perceptual data as visual data. Sound is sometimes also pro-

Parameter	Description
Number of parallel tasks	The number of tasks that the agent can perform in parallel. This is likely to be supported by a combination of working memory capacity, processing speed, and procedural memory strength.
Ease of engagement	Ability to switch between tasks (as opposed to fixating on a task).
Level of training	A highly trained individual will spend less time deliberating. This can probably be directly implemented through changes to the strength of procedural memory elements, but it may be useful to offer this as a more global parameter.
Acquiescence	A more acquiescent individual is more likely to follow orders. It may be possible to implement this in knowledge, but this will require further investigation if this parameter is going to be used in COJACK.
Stability	More stable individuals are less influenced by unexpected events and losses and gains. This can be implemented by changing the constants in the physiological parameters, providing a slower pace of change.
Humour	Individuals with more 'humour' are able to absorb and dissipate own and others losses and shocks. A complete theory here may be difficult to defend, but a simple implementation would be to provide a fast decay in the relevant physiology variables.

Table 1. High level parameters relating to cognition

vided, but data for other senses such as touch and smell are less often included. We focus here on visual perception, but have considered a similar parameter set for aural perception.

Table 2 shows a number of parameters associated with visual perception. Small changes in these parameters can be significant. For example, user's strategy choices are sometimes sensitive to millisecond differences [7].

These parameters (and the visual mechanisms that they influence) are assumed to be separate from other cognitive mechanisms. This approach is at times called impenetrable [12]. Differences in what is perceived are assumed to arise from how the percepts are interpreted in cognition, not how they are received from the perceptual apparatus. Thus, perception is impenetrable to cognition's ability to help or hinder it. This approach represents a modular approach that is useful because it makes it easier to create cognitive agents. There are already suggestions that this approach is too modular when taken to this extreme, and should only be seen as a useful working hypothesis.

In addition to the parameters presented in this table, an important aspect of human visual behaviour is search, scanning, and monitoring — vigilance tasks. Signal detection theory (SDT) provides a theoretical way to describe behaviour in this area [19]. As shown in Figure 1, SDT represents the signal to be recognised and noise that represents distracting stimuli and noise in the decision process of the observer. Both signal and noise are modelled as a signal that has a mean and Gaussian distribution. SDT represents observers as differing on two parameters, as shown in the figure. The first parameter is called d' (d-prime). It represents the distance between the distributions of the noise and the signal. Better observers have a greater d', a greater separation between noise and the signal, making the signal clearer. Observers also have a threshold, lambda. Measurements taken from the noise and signal distribution (the observer does not know which one) above the threshold are reported as a signal event, those below are not reported or are reported as noise.

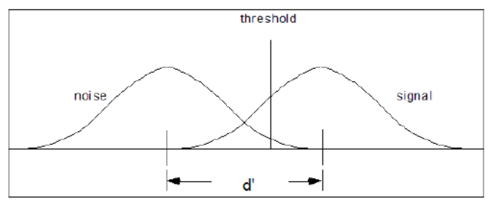

Figure 1. The major components of signal detection theory

Observers who wish to avoid missing signals will move the threshold towards the noise distribution to capture more of the signal; those who wish to reduce their false alarms will move it towards the signal distribution (thereby classifying more of the signal as noise). This approach abstracts many things, as the measures of d' across tasks are clearly tapping different mechanisms across some pairs of tasks, but for simple perceptual tasks it is a useful approach, as both parameters are known to be sensitive to many moderators.

Parameter	**Default (Variance)**	**Description**
Saccade time	120 ms (10 ms)	Time taken to move the eye from one location to another
Fovea size	3° (0.2°)	The size of the cone of vision for which full visual detail is available
Parafovea size	15° (3°)	The size of the cone for which limited visual detail (e.g. position and rough shape, but not exact size) is available
Peripheral vision	165° (10°)	The size of the cone in which objects can be detected
Visual working memory	3 (0.5)	Number of items that can be stored in the visual buffer
Visual acuity	Task dependent	Amount of detail that can been seen
Number of finsts	3 (1)	Number of moving objects that can be kept track of in a display

Table 2. Sample parameters affecting visual perception

2.4. Physiology

Physiological parameters are used to represent fundamental aspects of the agent's body. Most, if not all, of these parameters will influence the agent via their interaction with other parameters, rather than directly influencing the reasoning/action of the agent, and as such, they can themselves be seen as behaviour moderators. Some may be applicable to such a wide range of scenarios that they should be included within the COJACK architecture, rather than as optional behaviour moderators.

Possible parameters in this set include parameters such as heart rate, blood pressure, body temperature, and various hormone levels. One of the difficulties of including them in the architecture at this stage is that the effects of many of these variables on cognition have not been analysed with cognitive architectures in mind, giving limited data to build from. As a result, it is likely that the initial version of the architecture will only contain place holders for these parameters, without attempting to capture their full influence on the agent.

3. Behaviour Moderators

Behaviour moderators are factors, both external and internal to the entities, that cause them to deviate from their "normal" behaviour. It is not intended to include behaviour moderators in the COJACK architecture; rather, the intention is to provide support for the addition of moderators required by particular simulations. The architecture will provide this support via a set of hooks to the parameters described in Section 2. Incorporating a moderator into a simulation will then require an understanding of how that moderator affects these parameters so that the appropriate variables can be controlled.

Moderators can be organised in many ways, one way that has been proposed is to group them into three classes based on the source of the moderators [13]: external — those which arise outside the entity, internal — those which arise out of internal changes in the entity, and task-based — those that arise from information processing. Strictly speaking, task-based moderators are a special sub-class of internal moderators. They are a particularly significant group in themselves however, with important implications for modelling behaviour, so they are kept separate here.

Some moderators do not clearly fall exactly into one of these groups. Appraisal is one such example. This is the agent's evaluation of the current situation, and the most obvious group to put it in is the group of internal moderators. However it is also influenced by task-based parameters, and also environmental parameters.

3.1. External Moderators

External moderators are external events/conditions that can affect the entity's behaviour. Table 3 provides examples. Because the range of possible external moderators is so large, this table cannot provide a comprehensive list of environmental moderators; it is only intended to give representative examples. The external moderators that can and should be modelled for a given task and model will vary based on the model, the task, what is available in the environment, and, importantly, the use of the model.

These moderators influence the agent's body, and will have to be implemented as changes to intermediate, physiological parameters that are time dependent. For example the effect of temperature is a cumulative function. These parameters can then be used to moderate internal parameters.

3.2. Internal Moderators

Internal moderators are those which arise out of changes in the individual. Variations in the values of the entity's parameters (see Section 2) can themselves lead to variations in other parameters. Task-based moderators (discussed in Section 3.3) are a special sub-class of internal moderators that

Moderator	Description
Temperature	Exposure to excessive heat decrements performance on several dimensions, including visual tracking, vigilance, and complex tasks. With exposure, reaction increases as a function of time and temperature. [5, Ch. 10.601]. Cold has similar effects.
Humidity	A factor contributing to physical fatigue.
Noise	In addition to contributing to fatigue, noise can lead to poor communication (misunderstandings, or even lack of information transfer), resulting in reduced situation awareness.
Vibration	Vibration can reduce accuracy - e.g. firing from a platform over rough terrain. It also contributes to physical fatigue.
Time of day	Circadian rhythm affects cognitive processing speed and attention. Quality of sleep is affected by when it occurs.
Precipitation	As well as affecting humidity and temperature, this can have direct effects: e.g. moisture can affect the ease with which weapons can be manipulated.
Incoming fire	Direct hit causes physical damage, but even indirect fire can cause increased arousal at low levels and at higher levels incapacitation and, over time, stress.

Table 3. Sample external moderators

we have chosen to present separately. Other types of internal moderators include changes in chemical levels, such as caffeine, and sleep and fatigue-related factors.

Chemical moderators such as caffeine are in a way like external moderators. These moderators originate outside the body, but it is their affect on the body (and subsequently on the brain) that produce the changes in behaviour. Typically, an initial dose is absorbed into the body, which may take a short time, and then over time decays. Depending on the level of the chemical in the body, various aspects of cognition, perception and/or action may be affected.

Other internal moderators, such as those related to sleep and fatigue, are more obviously internal, but still have some dependence on the task and environment. For example long periods of wakefulness will degrade various parameters associated with cognition, perception, and action.

Appraisal-based moderators are another type of internal moderator. These moderators, also referred to as emotive moderators, are those that arise from an entity's judgement (of events, things, people, people's actions, etc.). For example, if an individual perceives a high threat to self, they may revert to instinctive rather than trained behaviours, or may abandon goals otherwise achievable.

3.3. Task-based Moderators

Task-based moderators are those associated with the information being processed and the passage of time. Most existing cognitive architectures are static — their mechanisms that give rise to behaviour are fixed across time; however there are many elements of the task that can moderate behaviour, including time itself.

A task that is performed over an extended period of time can lead to changes in behaviour via a number of routes. If the task has no frequent changes or challenges, boredom can lead to longer reaction times and decreased task performance, from 85% correct to 65% correct detections on a vigilance task [5, Ch. 7.403]. Physical fatigue may lead to reduced muscular strength, and mental fatigue affects cognition in a number of ways, degrading performance. Over time, the success/failure of subtasks may result in changes in motivation. This has been modelled by changes in parameters to the decision process [4].

3.4. Related Work

There have been many other attempts to model behaviour moderators, falling into two categories: those where the moderators have been explicitly incorporated into the modelling framework, and those where a model has been built as an extension to an existing framework. To the best of our knowledge, no other framework has been developed that was designed to explicitly support the addition of as wide a range of behaviour moderators as this parameter set.

Cognitive architectures that include explicit representations of behaviour moderators typically include a fixed set of moderators that they represent. Examples include PSI [3], which includes a number of physiological drives. Sloman et al. [17] provide a tool for developing cognitive architectures that are influenced by a variety of moderators.

Performance architectures predict how well humans will perform a task, not the detailed actions. IMPRINT (as reviewed in [11, pp. 259-260]) is an example performance architecture that includes the effects of environmental stressors such as climatic conditions and the type of clothes being worn. There are many more performance architectures that model the effects of emotions.

There are also many examples of architectures being extended to add models of behaviour moderators, particularly for Soar and ACT-R. Gratch and Marsella's appraisal model [6] is one of the more advanced examples. Their system maintains an appraisal of its situation which in turn influences its behaviour. In ACT-R, several projects have modelled individual differences, such as the work by Lovett et al on differences in working memory capacity [9], Jones et al's work on theories of development [8], and the model of pre-task appraisal developed by Ritter et al [14].

Although these examples are concerned with *specific* parameters and moderators, rather than the general support for moderators that COJACK will provide, they provide insights for both the development of architectures, and for developing and testing moderators for the architecture. In particular, the mechanisms used to achieve the moderation are of interest, and the way in which model data is matched with human data.

4. Development of the COJACK Architecture

Figure 2 presents a schematic of the COJACK architecture. The BDI-based JACK agent architecture provides the core mechanisms. Data from the environment must pass through the perception module, and actions are output via the action module. These mechanisms are moderated and constrained by the four sets of parameters described above, which in interact with each other through the mechanisms.

The two sample moderators in Figure 2, one external and one internal, only interact with parameters in COJACK, but it is anticipated that certain moderators may also interact with other moderators. For example, caffeine is a time-dependent moderator which is known to moderate the effects of fatigue. The figure also shows that whereas external moderators have data from the environment as their inputs, internal moderators (and also task-based moderators) take data from the agent's parameter set (and possibly other moderators) as inputs.

4.1. Moderators versus Parameters

The arrows in Figure 2 represent interactions between parameters and moderators, and it is immediately apparent that the distinction between the two is somewhat blurred. This is particularly the case in the physiology module, where most parameters moderate other parameters, rather than directly constraining the core agent. It is not yet clear where to draw the line for such parameters: it could be argued that they should be left out of the architecture and implemented as plug-in moderators, but it could also be argued that the more broadly-applicable mod-

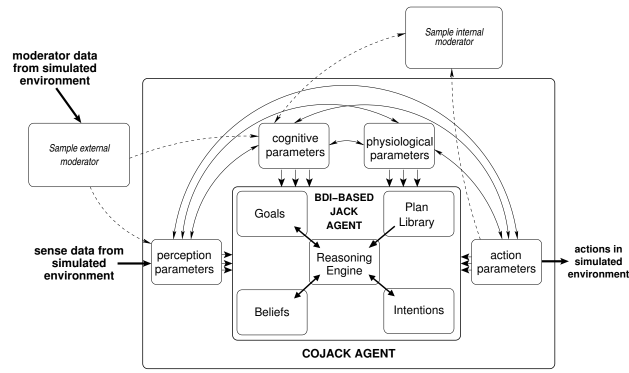

Figure 2. A COJACK architecture schematic

erators/parameters should be included in the architecture. It is likely that the initial version of the architecture will contain few, if any, of these parameters, but as the architecture is used, the more commonly used moderators of this type will be incorporated as parameters in the architecture.

4.2. The Importance of Time

Historically, cognitive architectures have put little emphasis on the passage of time, but time is an important factor in the effects of many moderators. Some moderators arise through the passage of time, such as decreased vigilance as time on task increases, whereas others disappear over time, such as anger dissipating.

Time influences behaviour on several scales. The scale of greatest importance and interest is over hours or perhaps days, as this is the time scale of the most important moderators and because most simulations and exercises do operate over longer periods than this.

Not only will it be necessary for the architecture to track the passage of time, but it is anticipated that displaying the changes to parameters and moderators over time will be of considerable importance to users, and support for this should be inherent in the architecture.

4.3. Representation of Moderators

Moderators will be represented using a scripting language, which allows the user to specify their inputs, and their influences on other parameters/moderators via complex time-dependent functions. Such functions may consist of a series of segments: taking caffeine again as an example, there is one function to represent the initial uptake of the drug, a second function to represent the period of maximum effect, and a final function to represent the decay period.

4.4. Constraints on Agent Representation

In order for the parameters to be meaningfully applied to the core JACK agent in the architecture, there must be some constraints placed on the representation that the agent uses, particularly for its beliefs and plans. For example, currently beliefs can be represented in JACK through various mechanisms, including using *any* JAVA class. However, if JACK beliefs are to be constrained by parameters relating to memory access and decay, the beliefs must be captured in a form that can be translated to memory elements.

We are documenting these constraints on programming style, rather than enforcing them, leaving it to the user to ensure that their agent is represented in a cognitively plausible manner. If the guidelines are followed, the behaviour of the agent, including the effects of moderators, will be plausible, but otherwise there will be no guarantees on behaviour. Eventually, it may be possible to warn about inappropriate representations.

4.5. Interface Needs

There are a number of issues related to the interface design that will be important in the development of COJACK. Firstly, the parameter set is large, considerably larger than that of similar architectures (such as ACT-R). The interface must provide a means for dealing with this parameter set, both for initialising it, and for monitoring it, bearing in mind that a single simulation may include a considerable number of these agents.

With such a large parameter set, it would be useful to be able to specify typical agents (as in Section 2), or even typical agents from certain populations, where a population is specified by having different means and standard deviations for particular parameters.

While ultimately the results of moderators should be visible through the actions of the agents, it is also desirable to be able to monitor the variation over time of individual parameters. A historical view of changes in beliefs, and possibly also the reasons for changes, will be a powerful tool for debugging and validation.

Another issue related to the interface design is related to time. We have already stressed that the passage of time is important when modelling behaviour moderators. However for many moderators, the effects are played out over a very long time period. While it would be possible to simply play out the simulation to observe the effects, it is desirable to be able to specify time and its effects — that is, to be able to specify something like "the agent ingested 200 mg of caffeine 20 minutes ago."

5. Future Work

The development of COJACK is in progress. This includes the development of the overlay containing the parameters and their influences on JACK, the development of a scripting language for representing behaviour moderators, and the documentation of the constraints on agent representation when writing JACK agents for use with COJACK. As this work progresses, it is likely that the parameter set will itself evolve, but we expect the final set to remain roughly the same size — about sixty parameters.

Once COJACK is operational, the initial testing will use a series of well-studied tasks, so that the architecture can be verified. Such tasks include serial subtraction and air traffic control tasks. Extensive data is available for these tasks, both with and without the effects of moderators. As part of this testing, we will be developing some initial moderators, not as core components of the architecture, but as samples. Likely candidates include pre-task appraisal, fatigue, and caffeine. These moderators have been extensively studied, providing sound data.

Ultimately, a large scale demonstration will be created using a military simulation environment. The parameter set supports including further moderators in this demonstration if required. Finding moderators that are both relevant to a military scenario *and* well-studied enough to implement is challenging. Reviews are very useful, for gathering such data is costly, time-consuming, and sometimes, particularly for militarily-relevant moderators, a dangerous task!

6. Conclusions

This paper has described the development of a parameter set that will be used to support psychologically plausible human variability for BDI-based agents. There are already some general lessons for other architectures.

The parameter set (and the effects of variations in the parameter values) form the basis for COJACK, which will allow the representation of individual differences (through variations in the values of these parameters) and the addition of behaviour moderators (through specifications of their influences on parameters). This parameter set suggests that multiple parameters will be necessary to model human variability across individuals and across time.

Unlike most cognitive architectures, COJACK agents will include a representation of a body, including perception and action modules, and a number of parameters related to physiology. The interaction of the model's body with the environment will require a richer representation of interaction with the environment, keeping track of time and the tasks the model performs, and also more information from the environment, such as temperature and humidity.

While we expect that the parameter set that we have described will support a wide range of behaviour moderators, it is important that the models that are developed of specific moderators be psychologically plausible. While CO-JACK itself will be psychologically plausible, the effects of parameters on behaviour must be known in order to develop a valid model. This parameter set was chosen with that problem in mind. Without specific knowledge of moderators, it is possible to build a models that exhibit variability, but unless they are built from effects known qualitatively and quantitatively, or until it is verified against empirical data, the models cannot be said to represent human behaviour.

Acknowledgements

The investigation which is the subject of this paper was initiated by the Director of Technology Development, Ministry of Defence, Metropole Building, Northumberland Ave, London WC2N 5BP and was carried out under the terms of Contract No RT/COM/2/007. The authors would like to thank the anonymous reviewers for their feedback, as well as Mark Cohen, and members of the project team at Agent Oriented Software.

References

[1] Agent Oriented Software Pty. Ltd. JACK Intelligent Agents. http://www.agent-software.com/shared/products/index.html.

[2] J. R. Anderson, D. Bothell, M. D. Byrne, and C. Lebiere. An integrated theory of the mind. *Psychological Review*, (submitted), 2002.

[3] C. Bartl and D. Dörner. PSI: A theory of the integration of cognition, emotion and motivation. In F. E. Ritter and R. M. Young, editors, *Proceedings of the 2nd European Conference on Cognitive Modelling*, pages 66–73, Thrumpton, Nottingham, UK, 1998. Nottingham University Press.

[4] R. V. Belavkin and F. E. Ritter. The use of entropy for analysis and control of cognitive models. In F. Detje, D. Doerner, and H. Schaub, editors, *Proceedings of the Fifth International Conference on Cognitive Modeling*, pages 21–26, Bamberg, Germany, 2003.

[5] K. R. Boff and J. E. Lincoln, editors. *Engineering Data Compendium: Human Perception and Performance*. Harry G. Armstrong Aerospace Medical Research Laboratory, Wright-Patterson Airforce Base, Ohio, 1988.

[6] J. Gratch and S. Marsalla. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, In press.

[7] W. D. Gray and D. A. Boehm-Davis. Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied*, 6(4):322–335, 2000.

[8] G. Jones, F. E. Ritter, and D. J. Wood. Using a cognitive architecture to examine what develops. *Psychological Science*, 11(2):93–100, 2000.

[9] M. C. Lovett, L. Z. Daily, and L. M. Reder. A source activation theory of working memory: Cross-task prediction of performance in act-r. *Journal of Cognitive Systems Research*, 1:99–118, 2000.

[10] E. Norling and F. E. Ritter. Embodying the JACK agent architecture. In M. Stumptner, D. Corbett, and M. Brooks, editors, *AI 2001: Advances in Artificial Intelligence*, volume 2256 of *Springer Lecture Notes in Artificial Intelligence*, pages 368–377. Springer, 2001.

[11] R. W. Pew and A. S. Mavor, editors. *Modelling Human and Organizational Behavior: Application to Military Simulations*. National Academy Press, 1998.

[12] Z. Pylyshyn. Is vision continuous with cognition? The case for cognitive impenetrability of visual perception [lead article and responses]. *Behavioural and Brain Sciences*, 22(3):341–365, 1999.

[13] F. E. Ritter. Three types of emotional effects that will occur in cognitive architectures. In *Proceedings of the Workshop on Architectures Underlying Motivation and Emotion (WAUME93)*, School of Computer Science and Centre for Research in Cognitive Science, University of Birmingham, UK, 1993. http://acs.ist.psu.edu/papers/ritter93e.pdf.

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

[15] F. E. Ritter, G. D. Baxter, G. Jones, and R. M. Young. Supporting cognitive models as users. *ACM Transactions on Computer-Human Interaction*, 7(2):141–173, 2000.

[16] F. E. Ritter and E. Norling. The causal relationship between behavior moderators and entity behaviour. Technical Report RT/COM/3/006-D3, Agent Oriented Software Limited, 2003.

[17] A. Sloman. Architectural requirements for human-like agents both natural and artificial (What sorts of machines can love?). In K. Dautenhahn, editor, *Human Cognition and Social Agent Technology. Advances in Consciousness Research*, pages 163–195. John Benjamins, Amsterdam, 2000.

[18] R. St. Amant, T. E. Horton, and F. E. Ritter. Model-based evaluation of cell phone menu interaction. In *Proceedings of CHI'04*, pages 343–350, New York, NY, 2004. ACM.

[19] J. A. Swets. The relative operating characteristic in psychology. *Science*, 182(7 December 1973):990–1000, 1973.