Training Tactical Decision Making Using Cognitive Models

Karel VandenBosch
TNO Defence, security and safety
karel.vandenbosch@tno.nl

Willem A. VanDoesburg
TNO Defence, security and safety
willem.vandoesburg@tno.nl

ABSTRACT

Simulation-based tactical training can be made more effective by using cognitive software agents to play key roles (e.g. team mates, adversaries). For successful use in training, agents should show tactically representative behavior that support trainees in achieving the learning objectives. For a tactical maritime scenario, we developed such a cognitive agent, playing the enemy force. We used a BDI-architecture to develop a cognitive model of expertise, capable of generating reactive and proactive behavior. This model was implemented in COGNET and then integrated in a training simulation. Evaluation with naive naval officers as subjects showed that the agent successfully supported achievement of the learning goals.

Keywords

Cognitive agents, Intelligent agents, Simulation based training, Cognitive modeling, Decision making.

INTRODUCTION

After almost two decades of studying complex decision making in natural settings there is common agreement that a decision maker's knowledge and skills are not isolated capabilities, but instead are inextricably associated with the task and context. The implication of this understanding for training is that students need ample opportunities to experience typical decision making situations and to practice the strategies for interpreting and handling such problems.

Scenario-based simulator training is an appropriate approach for training decision making in complex environments (Oser, 1999). In order to achieve the training goals effectively and efficiently, the events in the scenario and the behavior of the key players (e.g. team members, adversaries) need to be carefully controlled (Cannon-Bowers, Burns, Salas, & Pruitt, 1998; Fowlkes, Dwyer, Oser, & Salas, 1998). However, the issue of control is still a problem in current training programs. This is caused by the open character of the decision making task. Depending on his assumptions and inferences, a trainee can often choose among many different actions at any point in a scenario. The control problem lies in the fact that other players in the tactical scenario (e.g. opposing forces, team mates and other own forces) must be able to respond in a tactically viable manner to any situation emerging as a result of the trainee's action. And, as argued, there can be many such situations.

The Royal Netherlands Navy (RNLN) makes ample use of simulators for training their warfare officers in tactical decision making. In order to maintain control over the training scenario, they use Subject Matter Experts (SMEs) of the training staff to play the role of adversaries and own forces. In addition, fellow-trainees are required to play the role of team member (=warfare officer's assistant). The advantage of this solution is that SMEs have the expertise to take the context into account when evaluating (on-line) the appropriateness of trainee behavior. They can also assess whether the training scenario develops in the intended direction, and make adjustments, if necessary. Thus, SMEs make it possible to deliver tactical training scenario that represent reality in terms of dynamics and complexity, whilst maintaining a high level of control. However, this solution has disadvantages as well. One problem is that experts tend to evaluate trainee performance intuitively, without being able to precisely point out which cues (or absence of cues) they use for diagnosing trainee behavior. Furthermore, SMEs often differ in opinion on what is to be considered appropriate and inappropriate behavior. It is clear that this hampers transparency of performance measurement and feedback. Finally, the need for SMEs to deliver training elevates costs of training and requires high organizational and logistic efforts. Using fellow-trainees to play the role of team member has its drawbacks too. Firstly, these fellow-students are in training to become a warfare officer and not to become the assistant. They are therefore likely to display unrepresentative behavior. It would be better if all available training time could be used for practising the role of warfare officer, and not for playing supplementary roles. The RNLN is interested in methods that can improve the quality and efficiency of training.

A solution would be to use virtual humans (*agents*) to play the supporting roles autonomously. If we can develop agents that in training scenarios produce intelligent and tactically realistic behavior of the individual or entity that they represent, we would be able to make training more traceable, more systematic, and more cost-efficient. In most computer games, characters are almost always controlled by defining list of rules and contingencies, also called *Finite State*

Machines. This is suitable for procedural and other constrained tasks. However, for ‘open’ tasks like tactical decision making, it is hard or even impossible to create a ‘spanning set’ decision matrix specifying appropriate behavior for all entities for all possible states that may occur during a scenario (Silverman, 2001), as in even relatively simple scenarios the number of states tends to be very high (Klein, 1998). A more promising approach is to develop agents whose behavior is a function of simulated cognitive processes (e.g. beliefs, intentions, goals) (Zachary, Ryder & Hicinbothom, 1998). The heart of such agents is a cognitive model. A cognitive model represents the knowledge and cognitive processes of an individual or entity in a certain domain, task or scenario. This representation needs to be so specific that, when provided with input, the cognitive model produces realistic behavior as output. There is a growing conviction that cognitive modeling be can used successfully to improve tactical training (Pew & Mavor, 1998).

For the Royal Netherlands Navy (RNLN) we are currently investigating how we can use cognitive architectures and cognitive models to open up new opportunities for tactical training. This paper reports a study on the development, deployment, and validation of a cognitive agent representing the enemy force in a simulation-based tactical training scenario.

THE TRAINING

Learning objectives

The development of scenario-based tactical training starts with the formulation of learning objectives. In a broad sense, the goal of tactical training is to learn and practice the cognitive activities that are fundamental to adequate tactical command, such as identifying and interpreting cues; collecting and selecting information; making inferences and judgments; deciding; developing, implementing, and communicating plans. The problem is that these activities take place inside the head of a person, usually unavailable for observation. Therefore, it is necessary to define the learning objectives in terms of observable activities that relate to the activities of interest. For instance: “trainee re-orders his task group formation to obtain information that would otherwise remain unknown”. In general, learning goals should focus on those activities that are necessary to (a) gain and improve insight in the exact nature of the tactical situation at hand, and (b) to examine the opportunities and restrictions of different tactical plans. The following learning goals were selected: (a) when to split up a task group of ships in order to gather information for threat assessment, and (b) how to make an optimal selection and formation of ships based on the task, the capabilities and availability of each ship.

Scenario

A tactical training scenario consists of a mission, a situation description and anticipated events. The mission is described in terms of operational goals. The situation is described in terms of environmental conditions, enemy intent and capacities, intelligence information, logistic and organizational constraints, and other relevant types of information. A scenario contains opportunities for events to take place. An event can be defined as a critical moment in task execution, marked by (the absence of) a significant cue, requiring a response by the trainee as formulated in the learning goals and resulting in reactions of other players. However, trainees may or may not perform the desired behavior. As the actions of the trainee obviously influence the course of the scenario, it is difficult to plan the training fully ahead. A scenario can offer the opportunity for these situations to occur, but whether they actually do depends on the actions of the trainee.

For the present study, the following scenario was constructed in collaboration with an RNLN instructor: The trainee plays the role of commander of a task group consisting of an Amphibious Transport Dock (the High Value Unit, or HVU) and 4 escorting frigates. His task is to bring the HVU to its destination within a tight timeframe. The commander has intelligence information about the type and number of enemy forces present in the area (see Figure 1 for a map of the situation).

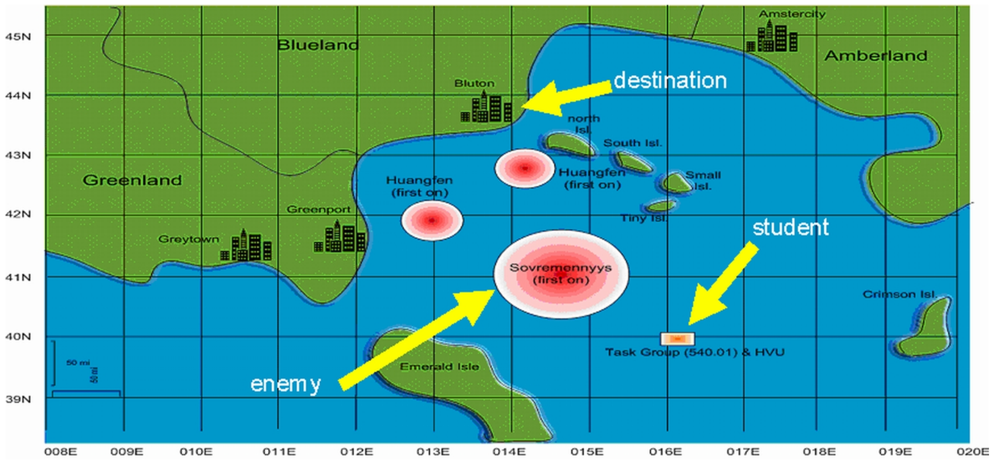

Figure 1. Scenario map, showing the position of the ships of the trainee (task group), the two ships of the enemy (sovremennyys) and the destination of the HVU.

The provided intelligence information should lead to the conclusion that a detour in order to escape the enemy is not possible. The trainee has to formulate a plan how to clear the fastest route for the HVU and executes that plan in an interactive simulation. Learning objective one is achieved when the trainee takes all information into account and comes to the conclusion that information concerning the location of the opponent can best be gained by splitting up the task group. Learning objective two is achieved when the trainee takes the specific capabilities of the different frigates into account to form an adequate picket and main group.

THE AGENT

In current training the enemy forces are played by SMEs. The merit of SMEs is that for any emerging situation they can generate options for enemy behavior that are not only tactically appropriate, but that they can also select from these options the one that best supports the trainee in achieving the learning objectives. So, if we want to develop an agent replacing the SME, then the agent needs to be able to (a) generate tactically representative behavior, and (b) select actions that facilitates trainee learning. Furthermore, the cognitive agent needs to be able to do so in unpredictable environments. After all, the trainee may at any time decide to do something unexpected.

For the present study we developed an agent meeting the requirements above. The agent was named TACOP, standing for TActical COgnitive OPponent. Expert instructors were asked to take the perspective of the enemy in the scenario, and were interviewed thoroughly about each key event in the scenario. The goal of the interview was to identify which cues the expert considers as relevant and which are not; what knowledge they bring in to assess the nature of the situation, what assumptions they make and why; what assumptions they reject and why, which strategies they consider appropriate and which are inappropriate. The results were coded in terms of Beliefs, Desires and Intents (see Figure 2). A BDI-architecture (Rao & Georgeff, 1997) was used because it is a well-known paradigm for generating reactive and proactive behavior (e.g. Norling, 2003). The next five subsections will briefly explain the behavior generation process.

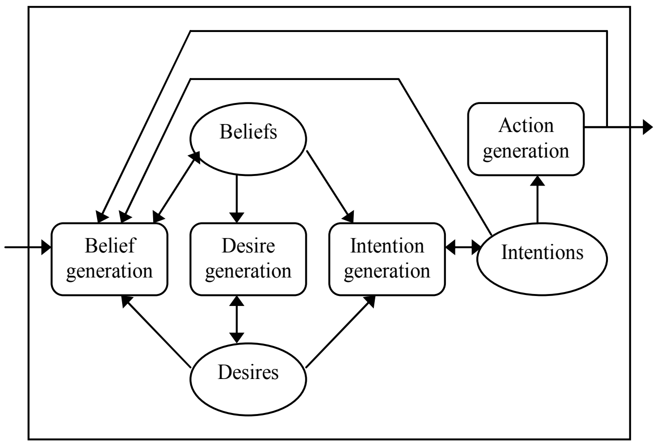

Figure 2: BDI-model of TACOP

Belief generation

The agent's beliefs define his knowledge and reasoning. Simple beliefs get formed passively through sensor perception; e.g., when the radar sensor fires, it triggers the belief that a track is detected. Complex beliefs get actively formed when the agent is in a certain state of mind (formed by its beliefs, desires, and intentions) and reasons about it. The belief about which radar track is the nearest is such a belief; it is only generated when there is an intention to shoot at a track. Note that beliefs have to be generated, but also have to be updated and (sometimes even) deleted (see Table 1 for an example).

Belief generation	Belief deletion
at any point in time, *if* a radar detects track x with bearing y, range z, heading u, and speed v *and* it is not already believed that there is a track x *then* the belief that there is a track x *with* bearing y, range z, heading u, *and* speed v will be generated	at any point in time, *if* track x, with bearing y, range z, heading u, and speed v, disappears from the radar *then* the belief that there is a track x, *with* bearing y, range z, heading u, *and* speed v will be withdrawn

Table 1: Example of belief generation and deletion

Desire generation

The desires of the agent are formed by the agent's goals. Static desires (i.e., primary and always activated) and dynamic desires (i.e., temporary and only activated under certain circumstances) can be distinguished. For our cognitive agent TACOP, the static desires are "Self Defense", "Disable the High Value Unit (HVU)", and "Return to Base". A dynamic desire is, for example, the desire to fire at the HVU. This desire gets activated when the belief is present that the HVU is within range, and lasts as long as that belief persists and it is not believed that the HVU is disabled. Although multiple desires can be activated at the same time, only one desire can be in focus, depending on the beliefs and other desires an agent has.

Intention generation

An intention is a step in the plan towards reaching the final goal, and will be executed as soon as possible. When a desire of an agent is in focus, intentions will be generated. For example, if the desire to engage the HVU is in focus, then the intention to attack gets generated. Beliefs about its weaponry (e.g. range, accurateness, destructive power) determine which weapon becomes selected.

Action generation

Actions are necessary in order to make sure that beliefs, desires and intentions actually produce behavior. Observations or actions can be generated from intentions (see example in Table 2).

Action generation
at any point in time, *if* there is an intention of own ship to sail to bearing b, over a distance d, with speed s *then* own ship will sail to bearing b, over a distance d, with speed s

Table 2: example of action generation

External

In order to make sure that the agent's actions effect the environment, a link can be made between the actions of an agent and the external (real) world (see example in Table 3).

External link	
at any point in time,	
if	*there is a ship in the world with bearing y, range z, heading u, and speed v, which is in range of a radar*
then	*that radar will observe that there is a track x with bearing y, range z, heading u, and speed v*

Table 3: example of an external link

Didactic function

Instructors are familiar with typical errors of trainees. They know common gaps in tactical knowledge, what kind of erroneous assumptions trainees tend to make, and so on. We asked the instructors to explore the consequences of these knowledge gaps and errors for this particular scenario. Together we identified what actions of TACOP would make the consequences clear to the trainees. This information on typical errors and associated actions was incorporated into the cognitive model.

An example from the scenario illustrates how this works out. If the trainee decides to keep the ships of his task group together (which experts consider as inappropriate for this situation) rather than splitting them up, the task group will pass within radar range of the vessels controlled by TACOP. TACOP will then detect the radar of the HVU and will attempt to attack and destroy the HVU. In terms of the model: TACOP *desires* to find the HVU in order to destroy it and is searching for it. A *belief* is generated that the HVU is close due to the radar signal. A *desire* to attack the ship from which the signal originates is generated. The *intentions* to change course, turn on radar and prepare to attack are generated and result in the corresponding *actions*. However, if the trainee chooses to split up the task group and form one or more scouting escorts, then TACOP will encounter these scouts and consequently generates different *desires* (to avoid the scouts) and different *intentions* (circumnavigate). This demonstrates the proactive capabilities of the model, the manner in which it can deal with different emerging situations and its correspondence to human instructor behavior.

From cognitive model to cognitive agent

The cognitive agent model outlined in the previous section was implemented using the COGNET Architecture and IGEN Toolset (http://www.chisystems.com/). COGNET's main components are: (1) a *blackboard* that stores the declarative information of an agent, (2) *tasks* that represent the agent's procedural knowledge, (3) *perceptual demons* that sense the external world, and (4) *actions* that can be performed on the external world.

The blackboard, demons, and tasks (with associated priorities) form the cognitive model. The COGNET processor manages the attention and time of the cognitive agent; i.e. it manages the focus and execution of the active desires and certain reasoning steps. For example, normally the global goal "Disable the HVU" has the highest priority. However, as soon as the belief exists that a missile is approaching, the global goal "Self Defence" gets the highest priority and is executed.

THE SIMULATION

The properties and effectiveness of a cognitive agent can only be studied in an environment. An agent needs to be able to "perceive" the environment, and actions of the agent have to affect it. For training purposes, agents are used in simulated environments. The simulation provides a representation of the physical aspects of the environment, like environment and entities (e.g. distance, weapons, sensors, motion, damage, and inventory). Both the trainee and the opponent agent form a part of each others environment. The simulation forms the shared environment of both.

VR-forces was used to produce the scenario and simulation for the present study. The COGNET agent TACOP was attached to this simulation using a HLA-network connection. See Figure 4 for an impression.

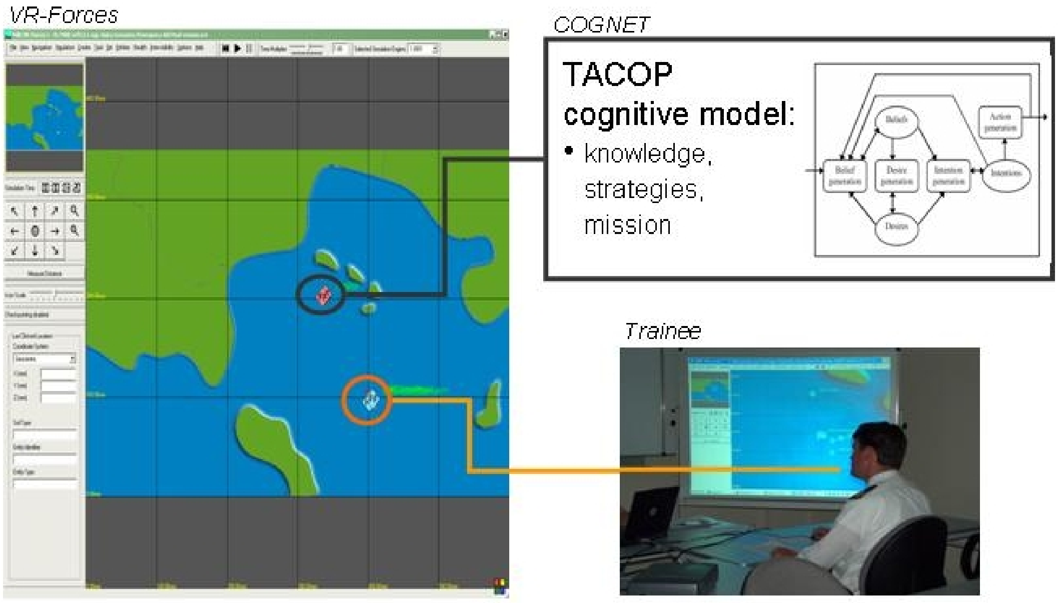

Figure 4: Impression of TACOP simulation-based training

EVALUATION

In order to evaluate this prototype, two officers in training to be instructor, were asked to participate. Both officers were not informed about the nature of this project and were unfamiliar with this particular scenario. The evaluation was done in three successive phases:

- Trainees performed the exercise while thinking aloud
- Trainees filled out a questionnaire after scenario completion
- Trainees were interviewed after scenario completion

The think-aloud protocols indicate that the tactical situation presented in the scenario elicited trainees to reflect on the issues specified in the learning objectives. In situations emerging as a result of inappropriate trainee actions, TACOP generates actions that demonstrate the consequences. For example, a key decision point in the scenario is when the picket of the Task Group and the enemy ships (TACOP) come within visual range. TACOP will notice that the HVU is not among these ships and decides immediately to split his group of ships and change course, each at right angle to the picket's course, thus quickly sailing out of sight again. If the trainee decides to ignore TACOP's action and move on, then TACOP will circumnavigate the picket to engage the following HVU. If, however, the trainee decides to attack the enemy ships, then TACOP's desire for "self-defence" becomes active, resulting in the action to engage the picket. The agent's behavior thus support trainees to develop insight in the causal relationships between environment, decision and actions. The reactions of the trainees when realizing the effects of TACOP's actions ("F*ck, he sailed around me!") indicates that the agent's behavior contributed to this understanding.

The results of the questionnaire show that the participating officers are positive of the potential learning value of this type of training. The interview was structured around three themes: *instructional value* of the agent, *tactical representatively* of the agent, and *requirements for simulation fidelity*. Both participants confirmed that the actions of the opponent (generated by the agent) caused them to reflect upon their choices and to consider alternative tactical strategies. The participants agreed and responded positively on questions about tactical representatively: all actions taken by the agents were considered as tactically logical. Finally, although the environment and entities were modeled on a functional level only, participants evaluated the necessary requirements for simulation fidelity satisfied. They indicated that the simulation contained the relevant aspects of the tactical situation to perform the assignment.

DISCUSSION

The study reported in this paper demonstrates that it is possible to effectively use a cognitive software agent as adversary in a tactical simulation, thus creating a functionally realistic training environment. This environment is *open* because participants (agent and trainee) were free to develop their own tactics and pursue their own line of reasoning. The environment is *complex* because participants had to manage information about ships, systems, sensors as well as the decisions of their opponent. The environment is *dynamic* because the actual state of the tactical situation continuously changed as function of decisions. Thus, the training recreates the conditions in which naval officers make decisions during surface warfare.

Although enemy behavior in the present simulation resembles that of computer generated forces (CGFs) in (entertainment) games, TACOP generates its actions in a fundamentally different manner. CGFs can be described as *reacting* to (constellation of) events by determining an appropriate response in a contingency table. Their response is thus scripted: a predefined association between situations and actions. The CGF can therefore be considered as the executor of the contingency, without “understanding” the grounds behind it. A cognitive agent, however, does not respond by association, but as a result of a simulation of human cognitive processes operating upon relevant domain knowledge. Cognitive models can be considered as detailed, explicit and validated descriptions of expert knowledge. This transparency of cognitive agents makes them especially suited for instructional purposes (e.g. guidance, explaining feedback, etc).

In the present study we interviewed experts to develop the cognitive opponent model TACOP. TACOP has shown to function as intended in the tested scenario. However, in the future, we want to ensure that TACOP can function equally well in other scenarios. This calls for the conceptual model to be adapted such that it can incorporate temporal aspects (Jonker & Treur, 1998). This will enable the specification and analysis of dynamic properties required to display adaptive behavior.

The literature on cognitive modeling emphasizes the importance of validity (Pew & Mavor, 1998; Ritter et al., 2002). Usually, a model is considered valid if it represents behavior in the real world. This is validity in an *operational* sense. In the present study we argued that our cognitive opponent agent should display representative behavior. Although strongly related, the two concepts are not the same. An operationally valid agent always shows the behavior approximating as closely as possible to that of the represented individual or entity in the real world. A representative agent shows the behavior that is typical for the represented individual or entity, but the frequency or trigger conditions may be adjusted to fit the goal of the user. Let us use an example to illustrate this point. When in charge of tactical command, commanding officers are highly dependent on their assistants. Among others, the assistants provide essential information, make checks for the officers and monitor equipment readiness. One important learning objective of the officers is therefore to keep an eye on the assistant’s task performance and to correct any errors. However, in normal life, assistants do their job accurately: they make few errors. If we would develop an operationally valid cognitive agent of the assistant, this agent would not provide the officer with many opportunities to practice error detection and correction. However, if we would develop an agent that makes ample errors, this agent would provide the officer with plenty practice opportunities. In conclusion, an agent must represent an assistant; however for the purpose of training the assistant is augmented with erroneous behavior. For the purpose of the model -in this case: training-, a representative agent is better than an operationally valid agent.

The departure point of our research is that simulation-based tactical training can be made more effective by using cognitive software agents to play key roles (e.g. team mate, adversaries, instructor). In the present study we showed that a cognitive agent playing the enemy force in a tactical training scenario successfully supported achievement of the learning goals. The focus of our current and future research on modeling is on expanding the reach of cognitive agents for training applications.

ACKNOWLEDGMENTS

We gratefully acknowledge the RNLN Operational School, especially Brian Saal, for their close cooperation, and Jan Treur and Annerieke Heuvelink for their advice on the formal design of TACOP.

REFERENCES

1. Cannon-Bowers, J.A., Burns, J.J., Salas, E., & Pruitt, J.S. (1998). Advanced technology in scenario-based training. In: J. A. Cannon-Bowers, & E. Salas (Eds.), *Making decisions under stress: implications for individual and team training* (pp. 365-374). Washington, DC: APA.
2. Fowlkes, J., Dwyer, D.J., Oser, R.L., & Salas, E. (1998). Event-based approach to training (EBAT). *The International Journal of Aviation Psychology, 8*(3), 209-222.
3. Jonker, C.M., & Treur, J. (1998). Compositional verification of multi-agent systems: a formal analysis of pro-activeness and reactiveness. In: *Proceedings of the International Workshop on Compositionality, COMPOS9 7* (pp. 350-380). Springer-Verlag.
4. Klein, G. (1998). *The source of power: how people make decisions*. Cambridge, MA: MIT Press.

5. Norling, E. (2003). Capturing the Quake Player: Using a BDI Agent to Model Human Behavior. In: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*. Held at: Melbourne, Australia.

6. Oser, R.L. (1999). A structured approach for scenario-based training. *43rd Annual meeting of the Human Factors and Ergonomics society*. Held at: Houston, TX.

7. Pew, R.W., & Mavor, A.S. (1999). *Modeling Human and Organizational Behavior*. Washington, D.C.: National Academy Press.

8. Rao, A.S., & Georgeff, M.P. (1997). Modeling rational agents with a BDI-structure. In: A.S. Rao, & M.P. Georgeff (Eds.), *Readings in Agents* (pp. 317-328). San Francisco, USA: Morgan Kaufmann Publishers Inc.

9. Ritter, F.E., Shadbolt, N.R., Elliman, D., Young, R.M., Gobet, F., & Baxter, G.D. (2002). *Techniques for modeling human performance in synthetic environments: a supplemental review.* Wright Patterson Air force base, OH: Human Systems Information Analysis Center.

10. Silverman, B.A. (2001). *More realistic human behavior models for agents in virtual worlds: emotion, stress and value ontologies.* (Report No. Technical Report). Philadelphia, PA: Univ. of Penn/ACASA.

11. Zachary, W.W., Ryder, J.M., & Hicinbothom, J.H. (1998). Cognitive task analysis and modeling of decision making in complex environments. In: J. A. Cannon-Bowers, & E. Salas (Eds.), *Making decisions under stress: implications for individual and team training* (pp. 315-344). Washington, DC: APA.