VU Research Portal

Cognitive Models for Training Simulations

Heuvelink, A.

2009

document version
Publisher's PDF, also known as Version of record

citation for published version (APA)
Heuvelink, A. (2009). *Cognitive Models for Training Simulations*.

Chapter 8

Conclusion

The presented study was motivated by the gap between current scenario-based simulator training for which many people are required, and the desired situation in which students can train at any time by themselves. In order to narrow this gap, we developed methods and techniques for the modeling of intelligent agents that can perform roles in simulator training. It is likely that in the future such agents replace the humans currently required for training; by this our research contributes to future independent training.

We started this research with three research questions (Section 1.2.2). To answer these questions we developed in this study 1) components for cognitive agents that enable them to display rational as well as biased behavior, 2) a framework to describe cognitive capabilities, and 3) a feedback generating system based on multiple agents that diagnoses task performance. Our major research effort has been in the first area, so in developing *content* of cognitive agents. Our second research effort, in the *development* of cognitive agents, and third, in the *application* of cognitive agents, are at a more initial stage.

In the following sections we discuss for each research question the opportunities and limitations of the developed methods and techniques, as well as additional research possibilities. In the last section we reflect on the relevance of our work, and identify the need for future research.

8.1   Modeling Human-Like Behavior

The first research question of this study was: 'How can a cognitive agent display human-like behavior with a varying degree of biasedness?' Existing methods and techniques are commonly developed to bring about either rational or biased behavior, and not both types. We developed mechanisms that enable an agent to display rational as well as, to

a larger or smaller extent, biased behavior. For this, we extended the scope of Artificial Intelligence to more human-like behaviors inspired by Cognitive Science.

We investigated our first research question in the context of a typical military task: situational assessment. We started with modeling a rational, expert agent for this task. To achieve an assessment of a complex situation, an agent must be able to deal with factual, but also with uncertain information. In addition, it must be able to integrate information received from different sources, and relate new information to information previously gathered. These things are self-evident for humans, since most real-world tasks require these capabilities.

Unfortunately, reasoning about uncertain information received over time and from different sources is as yet not supported by integrated architectures. In general, they do not require information to be stored with a degree of belief, a time stamp, or source label, and therefore also do not offer support for the handling of these values. Usually, architectures do support that their basic knowledge entities can embed strings and numerical values, and often also offer operators that can compare them for equality and ordering. However, there exists a large gap between such generic processes and *cognitive* processes operating on these values.

A reason for the architectures stemming from Artificial Intelligence to not support the explicit reasoning over uncertain beliefs from different sources over time might be that these architectures are usually used to model agents for tasks that are better structured than the open, dynamic, complex tasks investigated here. Architectures stemming from Cognitive Science might not support these processes because they are usually used to model tasks at a smaller scale, and with more detail for the cognitive validity of the underlying processes, than the current research task.

We do not claim that the methods and techniques we developed are most suited to model human-like behavior for all types of tasks. For many tasks existing architectures possess all what is required for their modeling, and frequently more than we offer. However, they lack the required cognitive processes (e.g., reasoning about uncertainty or about time) for modeling dynamic, open, complex tasks like situational assessment. We hope that the mechanisms developed in this study can bridge the gap between the processes required to model complex tasks and those usually offered by integrated architectures.

8.1.1   Developed Cognitive Agent Capabilities

In order to perform its task in a simulated world, an agent must be able to reason about this world. We decided to model an agent's knowledge about what is going on by means of beliefs. For modeling a rational, expert agent in situational assessment we first had to develop methods and techniques that would enable agents to reason about uncertain

information received over time and from different sources. Next, for modeling behavior with a varying degree of biasedness, we had to determine which cognitive processes could become biased, as well as how and when this would happen.

In Chapter 3 we presented methods for an agent to transform observed information into beliefs, to integrate beliefs from different sources and/or times, and to reason about beliefs over time. These methods enable an agent to do this in a rational, or in a to a smaller or larger extent biased way. The implemented biases were selected by examining literature from Cognitive Science and by talking to task experts. The extent to which the biases influence the agent's behavior is determined by a parameter. This parameter was labeled 'stress' because of the context in which we want biases to emerge: under circumstances of stress and exhaustion. In order to model the agent's (biased) belief formation processes we proposed a belief notation that includes information about the time, source, and certainty of the information. At this stage, memory was modeled as an unlimited database of always available beliefs. This is not very human-like, and also turns out to be computationally undesirable.

In Chapter 4 we presented a more human-like memory model that distinguishes working/short-term memory from long-term memory. This memory model supports rational behavior due to the storage, maintenance, and perfect retrieval of beliefs, but also biased behavior, among others by forming abstract versions of beliefs, and by attaching an availability value to beliefs that enables biased retrieval. In particular, the memory model enables the inherent human memory aspect that sometimes specific details about an event can not be remembered anymore, but the event itself can. In addition, the model enables task-specific effects on memory, like the fact that beliefs required for the current task are readily available, and that not relevant beliefs are quickly forgotten.

In Chapter 5 we presented two control models that determine the behavior of the agent during task execution. The first model focuses on controlling the execution of the agent's cognitive processing components, and determines on-line whether rational or biased behavior emerges. The parameter that mediates between rational and biased behavior is labeled (cognitive) exhaustion. The mechanism that we developed to implement this dynamic exhaustion value was inspired by the idea of Hancock and Meshkati (1988) of 'tasks as stressors'. The essence of this idea is that humans get stressed when the tasks they have to execute lie above their cognitive capacity. This idea has, to our knowledge, not yet been implemented in cognitive agents to model stress. The second control model focuses on controlling the decision whether required information will be acquired by memory retrieval, or by sensing the world. For this model it is determined off-line which type of behavior emerges by the selection of one of the implemented, to a smaller of larger degree heuristic, task-strategies.

8.1.2   Points of Discussion

Belief Component

We decided to extend beliefs with a time stamp, source label and certainty value. Multiple aspects of this choice can be discussed. From an implementation point of view it is debatable whether the combination of these arguments with the belief term in one tuple was optimal, or whether it would have been better to have used multiple predicates. We selected the first option because of the increased transparency, although on multiple occasions it also demanded the specification of unused belief arguments.

Of another order is the question whether the added belief arguments are complete and sufficient. For the selected task it turned out they are. However, it is likely that for other tasks not all the arguments are relevant. In addition, for modeling other types of tasks and biases other arguments might be required, e.g., an emotion value.

Related to this point is the selection of the implemented biases, which are only a subset of all known biases. The reason to select the chosen biases was because they were known to influence situational assessment. The reason why we are satisfied with the implementation of this subset is that in order to create more human-like behavior it is, although desirable, not necessary to implement all possible emerging biases.

This point stresses the fact that we conduct *design* science: we 'devise artifacts to attain goals' and do not 'explain how and why things are'. This is also important when discussing the validity of the behavior generated by the agents embedding the developed belief component. We use the term validity to express the requirement that the behavior of the agents has observational fidelity, and not that this behavior is the result of cognitive valid processes. We validated this observational fidelity using experts (Section 3.3), but this validation study was unable to provide us with conclusive evidence on whether the belief framework is suited for modeling valid human-like behavior. The reason was that other factors (the agent's speed and task control) already impaired the face validity of the agent's behavior. Therefore, we cannot conclude with high certainty that the developed belief framework enables the modeling of human-like belief maintenance. However, the fact that we could use it to implement the situational assessment task and that it was successful in mimicking a specific case of biased human behavior, stems us hopeful.

The last discussion point follows from Section 3.1.2 in which we state that the belief arguments can be used to maintain beliefs in a variety of ways. Although we showed some ways in our research, we did not investigate how possible it is to implement other techniques using the same notation. It is likely that adjustments need to be made, e.g., for implementing rational source integration using Dempster's rule of combination, the certainty value should be transformed to a probability interval.

Memory Component

The structure of the proposed declarative memory model mimics, to our knowledge, none of the memories embedded in known integrated architectures. The memory model is inspired by the proposed belief representation and is unique in that it is purely episodic in nature. We propose that over time semantic memory can emerge from this episodic memory, due to aggregations that abstract from the episodic details of memories and combine their content. Most integrated architectures only embed semantic declarative memories, with Clarion and Soar as noticeable exceptions that embed episodic memory in addition to semantic memory. We acknowledge that Cognitive Science, starting from Tulving (1972), considers episodic and semantic memory to be two parallel, partially overlapping, but distinct information-processing systems. Nevertheless, semantic knowledge is not innate, but acquired during the lifetime of a person. In our view it is therefore defendable to model semantic memories as emerging from episodic memories, and to model them in one system.

The incentive to develop the memory model was to decrease the query time of the agent's belief base. We proposed to do this by deducing aggregated beliefs that are required for task execution, and by making them more available than the basic beliefs. Whether the memory model will indeed prevent the agent from slowing down awaits evaluation. An important point for this will be how frequently it is possible to have the beliefs that are required during task execution readily available. This is likely to vary from task to task.

Control Component

The developed agent reasoning control model determines on-line, by taking the task circumstances and its internal state into account, which reasoning components to execute. We developed this component to enable the modeling of agents that display rational behavior under normal task circumstances, but start showing biased behavior under stress, like humans do. In order to model this behavior the agent has to reason in a declarative way about many aspects, such as the possible actions and their relevancy. Most integrated cognitive architectures do not support such explicit, declarative, on-line reasoning over actions, let alone by taking into account an agent's internal state. We think that it is required for an agent to reason in a declarative way over its actions to execute the variety of behavior required for training simulations. One interesting question concerning the developed control for an agent's reasoning process, is its computational scalability. To determine which cognitive processing components are best executed it reasons over many aspects, which costs much processing time.

In the control model for information acquisition, the agent's behavior was determined by the selected task strategy. Besides the so-called 'rational' task strategy that reasoned about the costs and benefits of the various actions to choose between them, we implemented several 'heuristic' task strategies. The actions that could be selected by the latter were limited, and the heuristic strategies varied in the amount of information they retrieved in order to make a choice between the actions. The developed task model reasonably fitted the actions of humans executing the task. This is interesting, because the heuristic strategies were, although inspired by the participant's behaviors, derived by a meta-model on how to form heuristic strategies: by varying the number (and order) of retrieval actions humans are willing to take to come to a decision.

Of the participant's reaction times, only two out of four correlated with those of the model. The reaction times of the model that correlated the strongest with human reaction times was the model that followed a rational strategy. This, together with the fact that the behavior of the heuristic strategies appeared to be more restricted than the behavior actually shown by humans, makes it questionable whether the heuristic strategies truly capture the way humans operate. It might be more valid to adapt the rational strategy to fit various personalities. Currently, the strategies only consider the costs of the actions as they really are. However, it might be that people's personality influences how these costs are perceived, e.g., some people are very sensitive to making mistakes, so for them the costs of making a mistake should be weighted more. In addition, people's current internal state might influence these costs, e.g., when stressed or exhausted they might be slower in belief retrieval.

8.1.3   Additional Research

In this section we elaborate on ideas for future research that directly follow from the research presented in this study. At the end of this chapter, in Section 8.4, we elaborate on interesting future research directions of a more general order.

In this study we developed several agent components. A challenge for the future is to combine these components: the model that dynamically determines the agent's current exhaustion level can be used to make the fixed stress level embedded in the belief framework dynamic. In addition, the control model that decides whether rational or biased cognitive processing components execute, can be applied to the memory model to determine whether beliefs are perfectly retrieved, or 'quick and dirty'.

When the agent components are combined in one cognitive agent, it becomes more feasible to validate them by asking experts their opinion about the face validity of the agent's behavior. When the behavior is not judged to be observationally valid, it has to be established which of the embedded components causes that. Of course, it is also possible

that not these processes, but the embedded task knowledge causes it. This might be circumvented by implementing a task of which a well-established model exists (although it is not very likely to come across such a task-model for an open, dynamic, and complex task). An additional benefit of implementing a new task using the developed components is that this exercise will highlight how suited the developed mechanisms and techniques are for modeling other types of complex tasks.

By combining the components as proposed above, the biasedness of their processes will be tunable through one parameter. The advantage of this approach is that when a training instructor does not want agents to display biased behavior, or only wants them to do so, he or she can fix this parameter. This way the biasedness of all processes are controlled in one step. A disadvantage might be that this parameter combines a variety of cognitive states, like stress, cognitive exhaustion (workload), and fatigue. These states may have different effects, e.g., Harris et al. (2005) found that extended stress deteriorated performance, but not fatigue. On the other hand, Hancock and Desmond (2001) state that stress, workload, and fatigue are not distinct and separate phenomena, but actually only different facets of the same phenomenon: they are all reflections of the energetic state of an individual. For modeling observational valid behavior for training simulations, it is probably unnecessary to disentangle these states and their effects.

In future research, it would be useful to model more capabilities in such a way that they can display rational as well as biased behavior. An interesting question is how these biases will reinforce each other when the capabilities are combined. For example, at this moment an agent can display biased belief formation because of its trust in the source of the information. In the future, this trust in a specific source might also influence the decision to use, or to not use that source for acquiring information. When capabilities are combined and multiple biases operate at the same time, it is important to research how they reinforce each other and whether their effect needs to be proportionalized in order to keep the behavior believable.

It may be hard to determine the validity of the possible emerging interactions, because these have not been studied much. Teachman et al. (2007) investigated three different information processing biases to determine how they inter-relate, but found that the biases showed little relationship to one another. They state that the results of the few studies into multiple biases general suggest no significant correlations. For example, Lundh et al. (1999) found no correlation between memory and attention biases, and Lundh et al. (1997) found no relationship between measures of explicit memory and implicit memory biases. On the other hand, Hirsch et al. (2006) introduce the 'combined cognitive biases hypothesis' and propose that biases do not operate in isolation, but influence each other and interact.

A last point for future research is implementing the pattern-matching capability that enables experts to recognize a current situation as similar to a previous one (Klein, 1998). This capability inspired us to attach a time stamp to beliefs: time stamps enable the recognition of belief patterns in time. In addition, the memory model enables the formation of generic belief-pattern representations by abstracting from specific details of beliefs. These generic belief-pattern representations can be used to compare current specific beliefs with in order to determine a match. Therefore, we think the developed agent components are very suited for implementing pattern-matching.

8.2   Describing Agent Components

The second research question of this study was: 'How can cognitive agent capabilities be described?' We investigated this question because a uniform manner to describe capabilities, the typical content of cognitive agent components, can be used to label these components. These labels can, in turn, be used when searching for specific components to reuse when developing a new cognitive agent. So with this research, we hope to contribute to the cost-efficient modeling of cognitive agents.

8.2.1   Developed Capability Description Framework

We were motivated to develop a framework for describing capabilities of cognitive agents because in the literature there exists no consensus on what constitutes capabilities, or on how to describe them. We explained the wide variety of capability descriptions found to be a result of whether they are inspired by the agent's underlying cognitive theory, or by practical, task specific design choices. The first leads to definitions of so-called horizontal (agent-specific) capabilities such as 'reasoning', the second to definitions of so-called vertical (task-specific) capabilities such as 'multiply'.

The goal of the Capability Description Framework (CaDeF) is to be able to describe all these capability variants. For this, CaDeF 1) defines a method for describing (variants of) agent capabilities, and 2) provides definitions of generic agent capabilities. The method prescribes that capabilities are defined by specifying the *functional*, *system* and *dynamic* properties for three types of capability entities: *means*, *processes on means*, and *control of processes*. The combination of these entities make up the capability. The generic agent capability definitions can be used to describe horizontal as well as vertical variants of these generic capabilities. Specific capability variants are defined by specifying values for properties of the capabilities entities, or by specifying addition properties.

Our research is still at an initial stage: only two generic cognitive capabilities are

defined by studying the literature on cognitive agents and incorporating the common ideas shared. Also the evaluation of CaDeF is at a starting point. We evaluated CaDeF by using it to describe two variants of the defined definitions implemented in BOA. For these capabilities CaDeF turned out to be satisfactory: their definitions were enclosed by the generic definitions, and could be expressed by extending them.

8.2.2 Points of Discussion

We determined that each capability can be described using three types of entities. The main reason why we are confident that all capabilities can be described by defining properties for these entities on three different levels, is that these entities are defined by their functionality, and that each entity is allowed to embed multiple entities itself. However, it is a point of discussion whether we will be able to formally express the embedding of multiple entity types in another entity. Our major concern is whether it is possible to define generic rules for the inheritance of properties.

Another issue are the terms used to define properties. Because a wide variety of capability-variants exists, the terms used to specify properties should be free. An advantage of not predefining concepts is that this offers the freedom to integrate an arbitrary model (capability) that determines the required concept in a way suited for the task at hand. However, this freedom delivers a risk for the discoverability of components. For example, when an agent designer is interested in the capability to make decisions based on emotions, he or she might search for a decision-making capability with the functional property 'takes emotions into account'. When another agent designer has developed an agent that bases its decisions on its mood (which can be considered an emotional state) and has defined this agent's decision-making capability differently, e.g., as 'takes mood into account', this agent component will not be discovered. To deal with such situations it is useful to attach an ontology of cognitive terms to the search mechanism. An ontology enables the discovery of cognitive agent components based on description similarity. There exists a variety of ways in which ontologies can be formed and their content mapped to a search term, for an overview see Shvaiko and Euzenat (2005).

8.2.3 Additional Research

In future research we want to provide definitions for all the generic (cognitive) capabilities that cognitive agents embed. We suspect that only a limited amount of generic capability definitions is required, because of the freedom to add arbitrary properties to define specific variants, and the possibility to include one capability in another.

Each of these new, as the current, generic capability definitions need to evaluated

on their ability to enclose and capture all possible variants. For this, a wide variety of (implemented) cognitive agent models should be described using CaDeF. This effort will also inform us about the usability and effectiveness of CaDeF.

A last point we would like to investigate is the ability of CaDeF to describe the horizontal capabilities embedded in specific integrated architectures, and to use these descriptions to select an architecture for implementing a specific task model. When developing BOA and Boar, we encountered the difficulty of implementing an established task model in an established cognitive architecture. This undertaking made us aware of an important factor when developing agents: before selecting an architecture in which to implement an agent model, check whether the vertical capabilities the model requires are compatible with the horizontal capabilities of the architecture. In the future, CaDeF could deliver the means to check this compatibility.

8.3   Generating Cognitive Feedback

The third research question of this study was: 'How can an agent generate cognitive feedback to the behavior of a trainee?' Because our research concerns open, dynamic, complex tasks, the generation of feedback on the level of cognitive processes is hard.

8.3.1   Developed Feedback System

The approach we took for generating cognitive feedback was to build a robust system by combining several methods to diagnose performance. The feedback system is based on multiple agents. The *feedback-generating* agent (FeGA) provides feedback and diagnoses performance using the other agents. FeGA starts it diagnosis by comparing the performance of the student with that of *expert* and *deficient* agents (model-tracing). When multiple matches exist, it attempts to clear the confusion by comparing the strategies of these agents with the deduced strategy of the student (plan-recognition). The latter is provided to FeGA by the *student-behavior* agent that deduces this strategy from the interaction of the student with the training environment. To support this interaction, we proposed to extend the training environment with non-intrusive 'task-support-buttons'. When at this stage the student's performance is still undiagnosed, FeGA diagnoses it by comparing the student's result with the expert result (constraint-based modeling). The diagnosed performance of the student is stored in FeGA's *student-performance model*. This model denotes the student's performance for each relevant task aspect using fuzzy sets, which are updated after each training session. On the basis of its student-performance-model, FeGA generates cognitive feedback.

8.3.2 Points of Discussion

A principle question here is whether the feedback system will be suitable for training tasks that are more complex than the current one.

The task for which the feedback system is developed possesses important aspects for training real-world situational assessment tasks, namely a criterion (threat level) that the student must learn to predict based on cues (sea lane, speed, distance). However, the relation between the criterion values and cue values was straightforward in our study: only three cues had to be taken into account, their values were known, and these values were available at the same moment. In real-world situational assessment tasks there are often more cues relevant, and their values are usually uncertain and dynamic.

The first question that emerges for more complex tasks is whether we will be able to capture the required task knowledge in an expert agent model. The current expert and deficient agents are formed by simple models with limited performance. When the task becomes more complex, the performance of these models will as well, as will the number of deficient agents. Even when we are able to correctly capture (biased) task knowledge in agent models it is a question whether under such circumstances it is computationally feasible to 1) run all these agents on one machine, and 2) compare their performances on-line to the performance of the student. Another question is whether for more complex tasks and cues, the three fuzzy sets that the feedback system uses to denote the diagnosed performance of a student on a specific task aspect (cue) are suited. A last question is whether it is for all tasks possible to extend the training environment with facilities (e.g., task-support-buttons) that extract additional knowledge about the student's behavior during task execution, and, more importantly, whether it is also possible for more complex tasks to deduce cognitive strategies from that knowledge.

8.3.3 Additional Research

Foremost, it has to be validated whether the proposed feedback system is indeed capable of supporting the threat-assessment training. We evaluated FeGA's capability to diagnose performance by letting it evaluate and diagnose software agents representing possible students, which it did satisfactory. However, for the representation of the students certain simplifications were made, e.g., their reasoning process was considered to be static. A real indication of the validity of the feedback-generation method is when students are able to learn the task based on the feedback that is generated. Therefore, we are planning to test whether students that receive the feedback generated by FeGA learn the task quicker than students that receive simple result-based feedback.

When it is found that the feedback system is capable of supporting training for this

simplified situational assessment task, it can be researched whether it is also capable of training more complex tasks, i.e., with more uncertainty and higher dynamics. For the modeling of a feedback system for real complex, open, dynamic tasks, the cognitive agent components developed in this study might be useful. For the current task, the expert and deficient agents are formed by simple models which did not require, e.g., the reasoning over uncertain beliefs over time. However, for modeling more complex tasks this might be required, and than the developed components can be used. For this, first an expert agent has to be modeled that incorporates the developed capabilities whose biasedness can be tuned. Then, a set of agents can be formed, ranging from expert to deficient agents, by simply setting the stress/exhaustion parameter to a variety of values. However, it is a question whether the use of such agents to diagnose performance is scalable.

8.4   Future Research

In the previous sections we discussed how our study contributes to future independent training by developing methods and techniques for the 1) content, 2) development, as well as 3) application of cognitive agents. In this section we discuss future research concerning these aspects of cognitive agents.

8.4.1   Cognitive Agent Content

Sandercock (2004) identified several areas in which current computer generated forces consistently show weaknesses compared to human players: environment awareness, human variance, persistence, vengeance, anticipation, learning, and teaming. In the current study we have developed content for cognitive agents that, of these 7 weaknesses, will decrease the inability to show human variance and environment awareness the most.

In Section 1.4.4 we listed the processes a cognitive agent should be capable of in order to show human-like behavior for the situational assessment task. When we would have listed the capabilities required to show human behavior in all its facets, the list would have been considerably longer and include aspects such as empathy, communication, and learning. In order to support the independent training of all possible types of tasks, eventually all the known human capabilities need to be modeled, i.e., those listed by Gordon (2005) and Langley et al. (2006). This will support the development of agents that can validly represent human behavior in all its facets, for all types of tasks.

Situational assessment was selected as task because it is important within the military. It is a task that underlies many other military tasks, and it is potentially subject to a wide variety of cognitive biases, which makes it an important task to train. However, sit-

uational assessment is a task which can be executed by an individual (unlike many other military tasks), and this limited the requirements for the cognitive capabilities (content) of the agent. In future research the progression of human-like behavior in training simulations can be supported by modeling additional content for cognitive agents, e.g., content that enables cognitive agents to display varied, human-like behavior in team tasks. For team tasks it is required that agents are able to communicate; to reason about the beliefs, goals, and intentions of others (theory of mind reasoning); and to explain their own behavior. Because the cognitive agent components developed in this study incorporate mentalistic notions, and declaratively represent and explicitly reason about their reasoning rules, they are well suited to support these new capabilities. Other capabilities, e.g., learning, might require more research to implement using the developed components.

8.4.2 Cognitive Agent Development

Modeling cognitive agents following a component-based design approach is cost-effective when components are reused. To foster reusability, cognitive agent components should be placed in a repository that can be queried for useful components. Useful components are those that embed capabilities and properties required for modeling the task of the agent under construction. In this study we started the development of a Capability Description Framework (CaDeF) for capabilities of cognitive agents, whose descriptions can be used to label components so that they can be discovered for reuse.

Unfortunately, a capability description is not enough to determine whether a component can be reused. Tracz (1990) proposed that a component's description should contain: 1) the *concept* or abstraction the component represents; 2) the *content* of the component, or its implementation, and; 3) the *context* that component is defined under, or what is needed to complete the definition of a concept or content within a certain environment. CaDeF supports the description of the *concept* and *content* of a component by embedding functional and system properties of capabilities, respectively. However, CaDeF offers no support for denoting the *context* of a component. This aspect is interesting to investigate in future research, e.g., by specifying the relations between capabilities.

Future research into component-based development of cognitive agents can benefit from studying the modeling of agents within AI; these agents are frequently based on separate components and a coordination mechanism (see, e.g., Brazier et al., 2002; Bosse et al., 2007). In addition, much can be learned from the research fields of Component-Based Software Engineering (CBSE), and Web Services. For example, CBSE defines three common stages in the process of developing a system based on components, labeled component qualification, adaptation, and integration (Brown and Wallnau, 1996), and has developed methods and techniques for each of these stages. Similarly, Web Services are

concerned with the specification, discovery, and combination of specific services on the web, and have also developed many possibly useful methods and techniques.

In this study we only made a small step towards a structured, cost-efficient development methodology for cognitive agents. But recognizing that this is important aspect for their future application, and therefore keeping it in our mind when developing agents, is a gain compared to modeling agents in an ad-hoc fashion.

8.4.3 Cognitive Agent Applications

In this study we developed cognitive agents for displaying human-like behavior within a simulated environment, and for executing a training task in parallel with, but invisible to, a student. In the last case the agent's task performance was used for comparison with the performance of the student, which leads to a diagnosis of the student's task performance that is required for feedback generation.

The generation of feedback to a student's task performance is one of the functionalities defined by VanLehn (2006) for Intelligent Tutoring Systems (ITSs). Other functions of ITSs are constructing an individual-tailored curriculum, and answering questions about the exercises or domain in general. In this study we focused on using cognitive agents to generate feedback. Below we elaborate on how in future research cognitive agents might also aid the other two functionalities.

For constructing an individual-tailored curriculum, it is important to diagnose the student's task deficiencies, and to know which training scenarios will require the application of these deficiencies. When the latter is known, the curriculum can be tailored to the student by offering these scenarios. For open, dynamic, complex tasks it is hard to determine which specific task knowledge the scenario will draw upon. Possibly cognitive agents can aid to determine this, by performing the student's task in such a scenario.

For answering questions about the exercises or domain in general, cognitive agents can be useful, especially when they are extended with the capabilities to explain their own behavior and to perform theory-of-mind reasoning. The first capability enables the expert agent to provide generic answers to questions about the task knowledge it embeds, the second capability to tailor this explanation to the student. For example, when it is known that a student desires to know about one particular task aspect, the agent can especially elaborate on this aspect.

The methods and techniques for creating agents capable of showing varied human-like behavior could also be used in other applications than simulator training. They might be useful for developing and testing military doctrine for which nowadays, among others, Computer Generated Forces are used that do not always behave very human-like. In addition, they can be used for decision-support systems. Decision-support systems that

embed a rational task model can provide the system operator with suggestions on how to execute the task. Moreover, decision-support systems that embed knowledge about false task behavior might not only be capable of detecting that the operator's behavior is not rational, but also diagnose which mistake he or she made. When it is possible to diagnose false task behavior it is possible to extend the support suggestion with an explanation tailored to the operator, or to perform other types of actions that bring the diagnosed mistake to the operator's attention. The cognitive agent components might also be useful to implement agents for validating system design, e.g., to investigate which design decreases the emergence of biases the most, or for modeling more human-like companion agents.

In spite of these many directions for future research, we can envisage applications for cognitive models in the near future. This study focused on the development of cognitive models for complex tasks that would enable single persons to train by themselves, and this is still a challenge. However, the development of cognitive models for more procedural tasks is well possible, and could already enhance current simulation-based training. After all, any part of training in which a cognitive agent can replace a person is a gain, especially now the military faces a strong reduction in personnel. Similarly, cognitive models that perhaps cannot replace, but aid instructors so that multiple trainees can be trained at any one time are already possible. However, for training domains that face less pressure on available man-hours, a careful cost-benefit analyses should be made to determine whether the time and effort put in the development of cognitive models pays back by freeing up personnel.

8.5   Concluding Remark

The proverb 'time is money' dates back 2300 years: the favorite saying of the Greek philosopher Theophrastus ($\pm$ 300 B.C.) was 'time is the most valuable thing a man can spend', or more precisely: '$\pi o \lambda \upsilon \tau \epsilon \lambda \epsilon \varsigma \ \alpha \nu \alpha \lambda \omega \mu \alpha \ \epsilon \iota \nu \alpha \iota \ \tau o \nu \ X \rho o \nu o \nu$'. Although many centuries later, also today much time and effort is spent on developing artifacts that in time will save us time. Indeed, more than that: no other century has yielded so many time-saving inventions as the previous one, and this trend is likely to continue in many years to come. A specific case in point is this study of which the objective is, in a nutshell, to decrease the man-hours required for training students in complex tasks.

We hope that this dissertation contributes to this time-saving trend by providing a clear story concerning the relevant questions, possible difficulties, as well as directions toward possible solutions to these questions and difficulties, when developing *Cognitive Models for Training Simulations*.