Adaptive Agents for Fit-for-Purpose Training

Karel van den Bosch(✉), Romy Blankendaal, Rudy Boonekamp, and Tjeerd Schoonderwoerd

TNO, Kampweg 55, 3769 DE Soesterberg, The Netherlands
{karel.vandenbosch,romy.blankendaal,rudy.boonekamp, tjeerd.schoonderwoerd}@tno.nl
https://www.tno.nl/en/

Abstract. Simulators and games provide contextually rich environments, enabling learners to experience the relations between actions, events and outcomes. In order to be effective, learning situations need to be tailored to the needs of the individual learner. Virtual characters (or *agents*) that, in real time , select, adapt, and exhibit the behavior that is exactly right for that learner, help to establish such fit-for-purpose training. This paper discusses principles for designing training with adaptive agents, and presents a framework for their autonomous and dynamic operation. A prerequisite for agents' adaptation of behavior to be successful is that adjustments do not violate the consistency and believability of the character, and maintains the overall narrative of the scenario. For reasons of management and coordination, it is proposed not to assign control over adaptations to virtual character-agents themselves, but to a dedicated *director agent*. This director agent is not a virtual character in the gameplay, but operates in the background. It collects and manages information, makes decisions about adaptations and issues behavioral instructions to the virtual characters agents. The framework was used in a pilot study, employing a human facilitator that simulated a director agent, arranging the adaptive behavior of virtual characters in a game-based training of military tactical decision making. Effects of adaptive and non-adaptive agents in a training were compared. Adaptive agents had a positive influence on learning and performance, and an increased engagement and appreciation by learners. Additional research with more participants is needed to verify these preliminary findings.

Keywords: Adaptive instructional systems · Training · Learning · Learner modeling · Personalized learning · Director agent · Intelligent agents · Adaptive agent behavior · Cognitive behavior · Simulation · Serious games

1 Introduction

Simulators and games are often used for learning decision making in complex environments [50]. A simulation enables trainees to experience the causal

C. Stephanidis et al. (Eds.): HCII 2020, LNCS 12425, pp. 586–604, 2020.
https://doi.org/10.1007/978-3-030-60128-7_43

relations between events, actions and outcomes in the simulated environment. It thus gives access to hands-on learning by doing [40,41]. However, it has been shown that goal-directed, systematic training is more effective than learning-by-doing only [26]. One way to make simulation-based learning purposive and goal-directed, is to carefully manage the behavior of the virtual characters [2]. This paper discusses the requirements for successful control of virtual characters (hereafter: *agents*), and suggests a framework for adaptive agent behavior (Sect. 2). Furthermore, this paper reports on a pilot-study investigating the hypothesis that adapting agent behavior, according to the a set of principles designed to facilitate learning, results in better learning and in a higher engagement of the learner (Sect. 3). The pilot-study used a game-based training in conducting a military tactical operation, with the goal to train situation awareness and decision making in unpredictable, unstable, and complex conditions [51].

To accomplish an effective training, the intelligent agents that govern the behavior of the virtual characters in the simulation or game, should meet the following demands:

1. Respond to situations with a sufficient level of realism, thus enabling a learner to develop the mental representations needed to perform the task adequately in the real world [18].
2. Exhibit the behavior that brings about a learning situation in the simulation that fits the needs of the specific learner. This requirement is often referred to as ‘attunement’, or ‘zone of proximal development’ [33].

1.1 Determining Agent Behavior that Supports Learning

A training simulation in which the virtual characters demonstrate behavior that is sufficiently realistic and believable for a trainee to achieve the learning goals has ‘functional validity’ [38]. This does not necessarily imply high fidelity agent behavior [18]. In fact, there is evidence that high-fidelity simulations can sometimes affect transfer of learned material negatively [14,39]. A realistic representation of a character’s behavior implies that all complexity and subtleties are included, which may be detrimental for learners as it may overwhelm of overstimulate a novice trainee [25,27]. Thus, when defining and designing the (behavior of) virtual agents it is important that the focus lies on maximizing effectiveness of the simulation instead of achieving realism on all aspects. It is therefore important that the behaviors of agents are believable (i.e., human-like), responsive (i.e., responding to user and environment), and interpretable (i.e., user must understand the underlying motivation).

1.2 Adapting Agent Behavior to Fit a Learner’s Needs

In a training simulation it is important that the virtual characters (1) behave believable, and (2) that their behavior brings about a situation that fits the competency level and learning needs of the trainee [16]. These demands may be compatible, but sometimes the objective of learning requires concessions regarding the realism of the simulation [49]. Consider the following example: a tactical military game is used to teach a platoon commander how to behave in the company of the village chief (simulated by a role-playing agent). The commander's objective is to address the village chief with the goal to obtain intelligence information. If the commander has a weapon at the ready while approaching the village chief, this behavior is considered incorrect given the learning goal: 'treat village chief respectfully'. Although military experts consider it plausible and realistic that a village chief remains utterly unmoved to what underneath is deemed disrespectful behavior, such behavior of the village chief agent would not bring about a learning situation that fits the needs of the trainee. If in contrast, the village chief would shrink back and say something like: "I thought we were on good terms, so why would you treat me like this?", then this would perhaps be less realistic, but it would improve the learning value of the exercise for the trainee (see Fig. 1). Such adaptation of agent behavior, for the sake of achieving learning objectives, is called 'fit-for-purpose adaptation'.

Fig. 1. Example of fit-for-purpose behavior adaptation of agent in training.

1.3 Principles for Adapting Agent Behavior

In order for a training system to be adaptive to individual learners, the system has to have access to three types of models: a task performance model, a learner model, and an instruction model [23].

Firstly, for an agent to be accepted as a believable character in a military game, the agent should behave in a manner that is consistent with doctrine (relevant and tactically plausible), taking into the account the context and conditions of the world in which the agent is situated. Furthermore, the agent's behavior should also demonstrate the influence of the characteristics of the real-life character that it is supposed to represent. Such characteristics may, for example, be:

an experienced or an inexperienced person; a fit or a fatigued person; an aggressive or a submissive person; and so on. This is defined in a task performance model of the agent.

Secondly, to be able to tune a training to the specific needs of an individual learner, an adaptive instructional agent has to *know* the learner [1,46,47], at least with respect to its characteristics relevant for learning. This is defined in the learner model and includes, for example, the learner's competency level and learning progress, and may additionally include other personal characteristics like motivation, self-efficacy, and engagement [8,12,36].

Thirdly, the adaptive agents should be controlled by a set of principles to select the behavior that brings about learning situations that are instructionally meaningful and effective [2] (see the village chief example in 1.3). This is defined in the instruction model.

The models of task performance, of the learner, and of instruction are all required for an adaptive training simulation system. These models do not necessarily have to reside within the agents controlling the virtual character; they may also be managed by other agents in the training system that distribute information and actions to connected agents [16].

The literature provides ample recommendations and algorithms for shaping scenarios to promote learning [2,5,6,43,44,50]. Which interventions are appropriate to foster learning are to a large extent dependent on the nature of the task to be learned. Military tactical command can be characterized as achieving situation awareness and making decisions under uncertainty. Learning to interpret a tactical situation requires the recognition and judgement of relevant factors (e.g.., weather, terrain, behavior of others, et cetera). Experts commanders have stored their interrelated and contextualized knowledge as mental tactical schemas [42]. Novices, however, do not (yet) have elaborated mental tactical schemas. If we want novices to become experts, training tactical command therefore needs to address: (a) expansion and refinement of tactical mental schemas, and (b) practice in solving complex and unfamiliar tactical problems (in order to develop new schemas) [4].

Detecting, recognizing, and interpreting situational cues is relevant for learning military tactics, because situations are continuous and variable, and not all experiences are informative for future decision making situations. Learning therefore resides in the activation (e.g.., frequency and recency) of experienced outcomes [17]. The main principle of an adaptive instructional agent should therefore be to adopt the behavior that supports the learner in detecting, recognizing, and interpreting the situational cues of interest.

2 Dynamic Adaptation of Agent Behavior in Simulation Based Training

Simulations and games offer better opportunities to exert control over training than real-life training environments do. For example, scenarios can easily be reused, improved and adapted for different situations and learners. The behavior of all actors in the scenario can be controlled and the behavior of the learner can be elaboratively measured. Adaptive agents may be used to support personalized fit-for-purpose training, but a framework is needed for their autonomous and dynamic operation.

2.1 Demands for an Adaptive Agent Framework

Peeters [33] has defined the demands for an automatically directed scenario-based training system. This can also be used to determine the requirements for adaptations of agent behavior.

First, the adaptation of agent behavior should be consistent with the training scenario, that has been purposely designed to enable the learner to achieve specific learning objectives. The default behavior of the virtual characters are explicitly part of the scenario. Adaptations in the agents' behavior should not disrupt the scenario's goal.

Secondly, any adaptations in agent behavior should be consistent with the character that the agent represents. If adaptations that are intended to enforce the desired behavior by the trainee result in agent behavior that is unnatural for the virtual character, they should nevertheless be omitted. Agent-adaptations that cause unnatural behavior to become part of the learner's mental representation would do more damage than good.

Thirdly, the adaptation should be *fit for purpose*, i.e., tuned to the specific needs of the individual learner. Learners may differ from one another in many ways, like for example, intelligence level, personality, preferred learning strategy, prior knowledge and experience, and pace of progress during learning. Adaptations should therefore be based upon a model of the specific learner [1] that is dynamically updated as learning progresses. Only then can the adaptations be selected that fit the learner in question, bringing about a situation that supports the learner to attain a goal or to engage in a task that otherwise may be out of reach [33,48].

Fourthly, adaptation of agent behavior should support or maintain the sensation of '*flow*' in the learner. Flow refers to a situation where learners are so engaged in the learning activity that they have a reduced sense of time and self-consciousness, and feel intrinsically motivated to engage in a complex goal-directed activity, simply for the exhilaration of doing [9]. There is evidence that flow brings about better performance and higher self-efficacy [20].

Fifthly, adaptations to agent behavior need to be applied in a controllable fashion, to safeguard the integrity of the overall scenario. One approach is to assign decisions about adaptations to the agent itself. This has the advantage that the adaptations are in correspondence with the agent's knowledge and character, and is therefore likely to show behavior that is believable in the perception of the learner. The disadvantage is that the agent's decision about adaptations may be in conflict with the narrative structure of the scenario. Another approach is to instead assign the decision making in the scenario, enforcing only those adaptations that support the plot. The disadvantage is that such adaptations may violate the consistency and believability of the agent. The challenge of control is therefore to ensure the agents' freedom of action while maintaining a purposeful and coherent training scenario [3, 29].

2.2 Framework for Adaptive Agent Behavior

A proposed solution for the control problem is to include a supervising director agent in the training framework [34] that manages the scenario, and that -when considered necessary- initiates behavior adaptations that have instructional value to the learner, and are consistent with the nature of the NPC-agent[1] in the context of the scenario.

Figure 2 shows a framework for adaptive agent behavior in simulation-based training that is able to accommodate the demands described above.

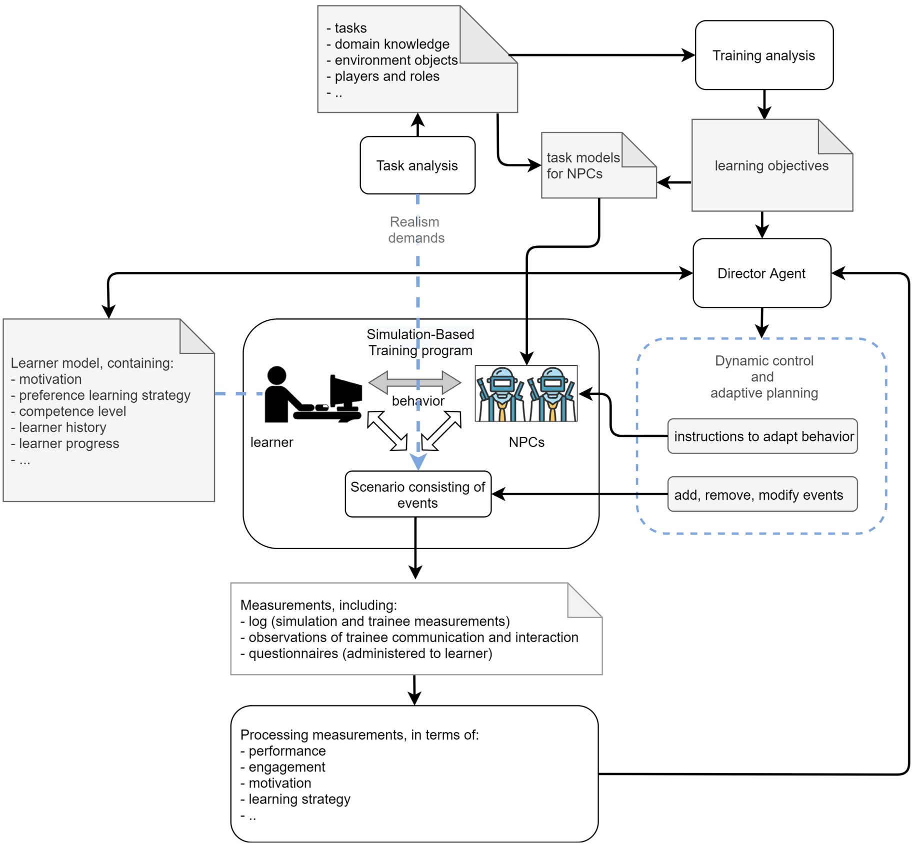

Fig. 2. Framework for dynamic adaptation of NPC-agent behavior to achieve fit-for-purpose simulation-based training.

[1] *Non-Playing Character*, a virtual character in a simulation not controlled by a player.

First, the task to be learned has to be analyzed in the context of its domain (the task analysis), revealing the constituent (sub-)tasks, domain knowledge, the players involved, and so on (see upper part of Fig. 2). This information is used for the training analysis, identifying the set of knowledge and skills needed to be able to carry out the tasks (skill requirements). It is also assessed what skills are already mastered by the learners, and what skills need to be developed, resulting in the learning objectives of the training. Furthermore, results of task analysis and the learning objectives are used to define the task-behavior models for the NPCs [13].

The director agent executes the adopted principles of learning. For military tactical decision making this may be as described in Sect. 1.3. The director agent controls the course and content of the default simulation-based training program (like the scenario, the events and the predefined roles and behavior of the NPCs).

Learners may not all respond in the same manner to a training program. For example, some may find the exercise in the scenario difficult, others may find it easy. Some may be engaged, others may find it boring. Some may achieve a learning objective quickly; others may fail even after multiple attempts. Information about the process and results of learning are measured and processed, and fed back to the director agent. The director agent monitors the scenario and the performance of the learner, dynamically updates the learner model and uses this model to determine whether the current state of the scenario is still a fit-for-learning environment for this learner, or whether fit-for-purpose adaptations are needed. This may involve instructions to an NPC to adapt its behavior.

A learner model is an important component of adaptive learning systems [45]. Their function is to provide the data for individualizing content and curriculum sequencing. Depending on the domain, needs for information about the learner and its learning can vary. For many domains, competence level is important, just as the learner's motivation to master the task [31,37]. For specific applications, the learner's physical and mental fitness may also be of interest to decide what learning situation to present [24]. Other interesting factors are the learner's preference for a learning strategy, the learner history and learning progress of the learner.

3 A Pilot Study

This study addresses the question whether personalized adjustments in the behavior of agents leads to a training with better learning results and a better training experience, compared with a training using non-adaptive agents. There is evidence that personal characteristics, such as e.g.., motivation and self-efficacy, affect learning and performance [31,37]. Therefore, measures of personal characteristics were included in this study. An additional purpose of this pilot study is to test the experimental setup, the design, and whether the desired changes in the behavior of the agents can be implemented properly and in time in the scenario. It is also investigated whether the adopted measurements for assessing effects on learning and learning experiences are suitable for use in a larger follow-up study.

The following research questions were formulated:

1. Do personalized adjustments in the behavior of agents lead to a better learning result?
2. Do personalized adjustments in the behavior of agents lead to better training experiences by learners?
3. Do the personal characteristics of motivation and self-efficacy influence the effects of adaptive training on learning and performance?

3.1 Methods

Participants: Six participants were recruited (4 male, 2 female). All were academic students, with their age ranging from 23 to 37 years old. Participants varied in their game-playing experience, from very little to a lot. Four participants indicated to have experience with playing military games.

Design: The independent variable 'adaptivity' had two levels (adaptive versus standardized agent behavior), and was manipulated between subjects. Participants in the experimental condition received a training scenario with agents acting adaptively to the learning needs. In contrast, participants in the control condition received a training scenario with agents acting in a standardized fashion, irrespective of the participant's learning needs. On a subsequent test scenario, all participants of both groups played the scenario with agents acting in a standardized fashion. Participants were randomly assigned to one of the two groups.

In the *control condition*, the agents behaved in a pre-specified and standardized fashion that was plausible given the situation in the scenario. Agents were designed to provide, through their behavior, the trainee with implicit cues and feedback about the consequences of his/her decisions and behavior [28].

In the *experimental condition*, the behavior of the trainee was continuously evaluated in real-time. If it matched the behavior of a particular learning objective, then the agent behaved as pre-specified. If it did not match, the participant apparently failed to achieve the learning objective, and then a predefined adaptation was applied to the agent's behavior. The purpose of the adaptation is to make the agent provide additional explicit or implicit cues, instructions and feedback to the participant in order to support understanding what behavior is appropriate for that specific situation.

Experimental Setting: The experiment took place in room with a large table (see Fig. 3). The game play was recorded.

The experiment employs a 'Wizard of Oz' technique [10], which means that the experiment leader (the 'wizard') simulates and controls the behavior of a director agent. This technique makes it possible to administer intelligent adaptations to agent behavior, without the need of actually implementing a model that can do the same autonomously. Note that the wizard strictly complied with the predefined decision rules of the director agent.

Fig. 3. Experimental setting.

Materials - Scenarios, Scenes and Events: Participants played a training and a test scenario with the game Virtual Battle Space 3 (VBS3)[2]. The participant played the role of platoon commander. The task to be learned was commanding the platoon on a social patrol mission in a small village, situated in an "Afghanistan"-like environment. This scenario was developed for the training of tactical officers of the Netherlands Army and was reused for the present study. The participant receives a briefing of the current situation before the start of the scenario, stating that armed enemy units have been seen in the area and that the task of the participant is to gather more information about the situation.

The specific goal of the platoon is to go to the marketplace of the village and try to get information from the village chief about a possible enemy threat. Pre-specified events ensure that the situation escalates and eventually a combat situation arises. Each scenario consists of the following scenes: 1) Start of the mission; 2) Approaching the market; 3) On the market; 4) Market crowd is leaving; 5) First time under fire; 6) Second time under fire; 7) Team member is hit; 8) Withdrawal from the mission area. Each scene consist of several events (for example: in scene 4, 'on the market', the village chief appears). An event provides the trainee with the opportunity to achieve a learning objective.

Applying Adaptations: An evaluator was tasked to evaluate in real-time whether or not the participant's behavior matched the requirements of the learning objective at hand. If the evaluator observed that the participant made an error, then this was passed on to the experiment leader (see Fig. 3) who subsequently initiated an adaptation in the behavior of the agent(s). Table 1 shows examples of learning objectives, potential errors, and associated adaptations. Note that adaptations were only given during the training scenario of participants in the experimental condition.

[2] https://bisimulations.com/products/vbs3.

Table 1. Examples of learning objectives, potential errors, and adaptations.

General learning objective	Potential error	Adaptation
Ensuring constant view on all sectors during the mission	Not having a 360° view when approaching the market	[An armed NPC runs behind the platoon.] A soldier says: "Behind us! An armed man!"
Addressing the village chief in an open, kind and polite manner	Keeping the rifle at the ready when addressing the village chief	Village chief says: "I thought we were on good terms, so why do you treat me like this?"

Measures: Both the training and the test scenario consisted of 31 learning objectives, each evaluated by the experimenter as either achieved (1), or not achieved (0). Learning objectives belonged to one of the following competences: planning; situational awareness; decision making; and command and control. The average of the associated learning objectives was taken as the participant's competency score.

In addition to these performance measures, the Motivated Strategies for Learning Questionnaire (MSLQ) [35] was administered to measure the confidence of participants to be able to learn the task (self-efficacy for learning and performance), and whether participants think that their efforts will lead to learning performance (control of learning beliefs).

Three more questionnaires were administered at the posttest: the scale 'interest and enjoyment' of the Intrinsic Motivation Inventory (IMI) [11] was used to measure the motivation of the participant. The participant's appreciation of the game as an environment to learn from was measured by using the scale 'game worlds' of the Game-based learning Evaluation Model (GEM) questionnaire [32], and the value/usefulness subscale of the IMI. Finally, the Immersion questionnaire [22] was used to measure how immersed the participant felt in the game.

Procedure: The participants took part in the experiment one by one. After a short briefing, and filling out questionnaires, a practice phase started. This was to familiarize the participant with the controls of VBS3, and to learn how to interact with NPCs (by speaking aloud). For this practice, an entirely different setting was used (a western suburban neighborhood) than during the training. In order to make NPCs respond in a plausible fashion, a large series of possible phrases were audio-recorded prior to the experiment. When the participant addressed an NPC, or when the scene demanded a verbal action of an NPC, the experiment leader selected an appropriate phrase and played that back. If no suitable prerecorded response was available, then the experiment leader himself responded, on behalf of the NPC, through speech.

After completing the practice with the game environment, the controls, and the interaction with NPCs, the training scenario was started. After completion, a short break of 5 min was given. Then the participants played the scenario for a second time. After this, post-test questionnaires were administered.

3.2 Results

This pilot-study involved no more than six participants. Therefore, effects of personalized adjustments in the behavior of agents can only be explored, not tested. For this reason, results will be reported only on a descriptive level, without performing statistical analyses.

Table 2 shows results of the confidence that participants have in learning the task successfully (self-efficacy), and whether participants think that their efforts will lead to learning performance (control of learning beliefs). These variables correlated .73.

Table 2. Judgments of participants (means, s.d. in parentheses) whether their efforts will lead to successful learning (control of learning beliefs), and their confidence in learning learn the task (self-efficacy for learning), split by experimental group.

	Non-adapative	Adaptive	Total
Control of learning	5.7 (1.0)	5.3 (0.9)	5.5 (0.9)
Self-efficacy	4.0 (1.4)	4.8 (1.1)	4.4 (1.2)

With respect to participants' prior gameplay experience: the two groups had similar experience with playing computer games: a mean of 3 (sd = 1.7) on a 5-points-scale for the *non-adaptive agent training* group, versus 3.3 (sd = .6) for the *adaptive agent training* group.

Performance on Training and Test Scenario: Performance was assessed by using competency scores. A participant's competency score was calculated as the mean of the scores on the learning objectives belonging to that competency. A score for learning objective could either be 1 for achieved; or 0 for not-achieved. Hence, a competency score has a range of 0 (none of the comprising learning objectives achieved) to 1 (all comprising learning objectives achieved).

Training: Figure 4 shows the competency scores. Both groups performed no planning whatsoever during the training scenario. Participants failed to settle task assignments, made no commands concerning communication policies, and did not develop action plans (e.g.., contingency plans). This is perhaps due to the unfamiliarity with the game and the assignment. With respect to the other three competencies, it can be seen that the participants of the *adaptive-agents training* group (orange bars) performed substantially better than participants in the *non-adaptive training agents* group (blue bars). This is most likely caused by the adaptive behavior of the agents. When the experimenter-evaluator observed

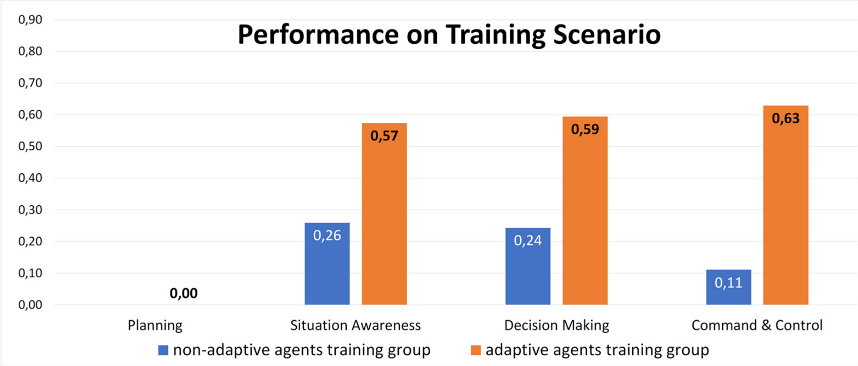

Fig. 4. Mean competency scores on the training, split by group. (Color figure online)

that a participant of the *adaptive-agents training* group did not exhibit the behavior required achieving the learning objective, then the experiment-leader initiated an appropriate agent to act in such a fashion that it supported the learner to reflect upon his behavior, and also to recognize that the nature of the situation demands a different response than the one just given (see Table 1). This adaptation of agents has likely cued the participant to be alert in similar kind of situations further on in the scenario, and may also have promoted the participant to be more reflective in general. Participants of the *non-adaptive-agents training* group received no such clear cues, and therefore had less opportunity to improve during the scenario.

Test: Figure 5 shows the competency scores.

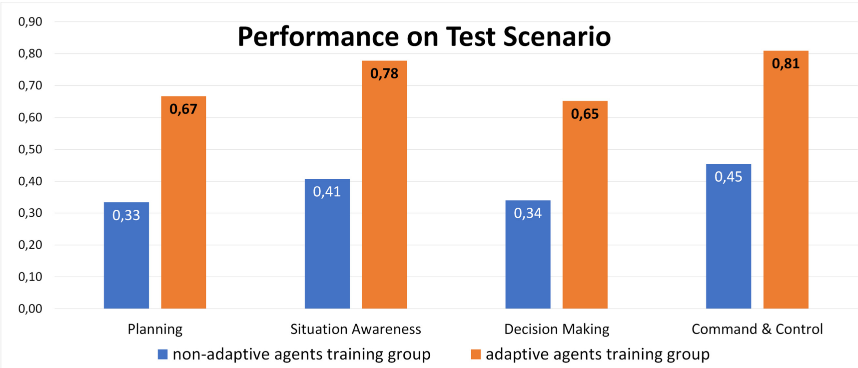

Fig. 5. Mean competency scores on the test, split by group.

Both groups performed better on the test scenario than during the training scenario. They apparently learned that planning and sharing plans is required to perform the mission adequately: they assigned tasks to their team members more often and some participants distributed policies for team-internal and team-external communication. Note that for both groups, agents in the test scenario did not act adaptively. So they were both presented with exactly the same scenario. Nevertheless, the *adaptive agents training group* performed substantially better (overall approximately 90%) than the *non-adaptive agents training group*. This indicates that the instructional effects of the adaptive agents during the training exercise continued to affect positively the performance on the test scenario. Participants of the *non-adaptive agents training group* also improved their performance, but did not nearly approached the level of the *adaptive agents training group*.

Relationship Between Participants' Beliefs and Performance: It was investigated whether the participants' performance on the test scenario was related to their beliefs about the results of their learning efforts (control of learning) and of their estimated ability to learn the task (self-efficacy). Kendall's tau correlations were calculated between the measures Control of Learning and Self-efficacy with the four competency scores obtained on the posttest. See Table 3 for results. The correlations are low to moderate, indicating that their beliefs on being able to learn and perform the task have little influence on the outcomes measures.

Table 3. Correlations between participants' self-assessments and performance on competency scores at the test scenario.

	Planning	Situation awareness	Decision making	Command & control
Control of learning	0.22	0.14	0.33	0.20
Self-efficacy	0.36	0.14	0.60	0.47

Reliability of Performance Evaluation: The experimenter-evaluator assessed, while the participant was playing the scenario, whether the participant's behavior either matched, or did not match, the criteria of the current learning objective. All game-play was video-recorded, including all audio from the participant, the agent, and the experiment leader. This enabled a second assessment, afterwards, by a second evaluator. This evaluator was instructed with the performance evaluation protocol. Assessments of both evaluators were used to calculate the interrater reliability. Evaluators were on 239 of the 328 assements in agremeent; on 79 they scored differently.

The calculated measure of agreement (Kappa) is .5. This is generally considered a weak level of agreement [30]. One likely reason for this result is the low number of participants. Another reason is that in a realistic simulation of a complex real-life task, such as military command in this case, it is difficult to unequivocally determine whether the observed behavior matches, or does not match a criterion. For example, one learning objectives was '*to maintain a 360-view when approaching the market place*'. It can be hard to assess from observing the participant, or from viewing the gameplay, whether or not this objective has been achieved. Furthermore, the evaluation protocol marked at what point in the scenario the desired behavior was to be expected. During a post-scoring discussion, it turned out that if the participant exhibited the desired behavior earlier, or (slightly) later than in the expected timeframe, this was not considered identically between evaluators. Finally, the level of detail in the information differed between evaluators. The first evaluator could, in addition to the gameplay, also observe the participant in action, thereby having access to subtle, but potentially informative behavioral cues. The second evaluator only saw the recorded gameplay, in combination with the audio-recorded communication. This difference may have had a negative influence on the interrater reliability. For future studies, it is recommended to define a very thorough evaluation protocol, and to have a second evaluator present while the training and test are being administered.

Immersion and Evaluation of Learning Environment: Do adaptive agents in the game environment invoke a higher level of immersion and engagement in the learners compared with non-adaptive agents? Figure 6 shows the results of the posttest measures.

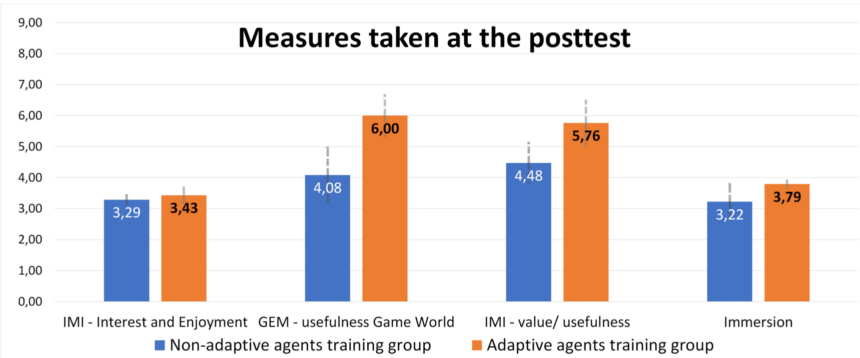

Fig. 6. Results of posttest measures in means and sd's, split by group.

After playing both scenarios, the two groups were comparable with respect to their self-assessed intrinsic motivation. Participants that trained with the adaptive agents considered the game and scenarios considerably better as a suitable learning environment than participants that trained with non-adaptive agents. This result was obtained with the GEM as well with the IMI. Finally, participants of the adaptive agents training group reported to be slightly more immersed than participants of the other group.

3.3 Discussion of Pilot Study

This pilot study investigated whether personalized adjustments in the behavior of agents leads to a training with better learning results and a better training experience. First, it should be noted that only a few participants were tested. Thus, results need to be taken as preliminary and with reservation.

One objective of the present pilot-study was to test the experimental setup. Overall, this worked as hoped for. But administering the scenarios, evaluating the participant's behavior in real time, and selecting the appropriate adaptations for the agents was quite taxing for the experiment leader and evaluator.

It is encouraging that the applied adaptations in the behavior of agents, with the intention to make the exercise more fit for purpose for the learner, resulted in better competency scores. This was found not only for the training scenario, but also on the test scenario. This suggests that what was learned, was carried over to an exercise in which no fit-for-purpose behavior of agents was provided. Another positive outcome was that the adaptations seemed to have brought about a more interesting and useful learning environment, according to the participants. Furthermore, adaptation may also have caused a more immersive experience, although the differences are relatively small.

In contrast to the literature [31,37], the personal characteristics of the learners of our study, in particular their motivation, control of learning behavior, and self-efficacy, correlated low with learning and performance. This may have to do with the low number of participants.

The scenario for this study was developed according to the principles of event-based approach to training [15], with the learning objectives explicitly coupled onto events in the scenario, and with defining descriptive measures for how to evaluate the participants' behavior as either achieving or failing the learning objective. It was therefore disappointing to find that, regarding the performance evaluations, the agreement between the experiment-evaluator and a post-hoc evaluator was rather low. This may partly have to do with differences in the information available to both evaluators, but it also points out that the protocol for evaluating performance needs to be defined in even more verifiable terms.

Gaming experience was found to significantly influence the nature of the gameplay. Participants that had few or no gaming experience experienced more difficulty with learning the controls and how to use them appropriately. Participants with experience in military games used the controls fluently, but they appeared to pay little attention to the mission instructions, nor to the cues that became available through NPCs. In future studies, we aim to recruit participants with moderate gaming experience (preferably not with first-shooter games), and with elementary military background knowledge.

4 Discussion

This paper addressed the issue of rendering fit-for-purpose simulations by adapting the behavior of virtual characters in real time, in such a fashion that the emerging situation is tuned to the learner's needs. Requirements for agent control have been formulated, and a framework for dynamic adaptation of agent behavior has been proposed. A pilot study reveals encouraging outcomes, suggesting that behavior adaptations bring about better learning, and perhaps also a higher engagement in the learner.

To collect more evidence, a proposal for a follow-up study has been submitted to administer the experiment to a new and larger group of participants. Lessons learned will be taken into account. To obtain a more homogeneous sample, participants will be recruited from young soldiers that are in the final stage of the preparatory military training. It is expected that they have a military mindset in common, but do not yet have experience in tactical operations. That study will again employ a wizard-of-oz that simulates a director agent.

For practical applications of the framework, the principles and algorithms need to be implemented in a computational version of a director agent. This will bring along challenges on many fields. One lies in achieving a functional and coordinated cooperation between a director agent and NPCs to effectively support learning. In our adaptive agent framework (see Sect. 2), we propose to assign domain-specific functions to NPC-agents, and instructional functions to a director agent. For example, NPC-agents need behavior models that enable them to behave in an appropriate and believable manner. Furthermore, these models should allow NPC-agents to exhibit the behavior that is appropriate for a specific exercise (e.g.., performing as a novice), or that satisfies the call of the director agent for adaptation. And the models should also enable the agent to construct explanations that inform the learner about its behavior, with the aim to support a better understanding of the interaction between them [19]. In our framework, it is not the NPC-agents who take decisions about whether and when to exhibit what behavior; these are taken by a director agent. Similarly, decisions about whether and when to provide what explanation can best be give, are made by a director agent rather than by the NPC-agent itself.

Another challenge for the present case is the open-ended nature of military tactical command, making it hard to interpret the intent of learner's actions. The challenge of modeling adequately the learners' domain knowledge, cognitive skills, and interests [1,7,21] will be crucial for a director agent to successfully achieve fit-for-purpose adaptive training.

References

1. Basu, S., Biswas, G., Kinnebrew, J.S.: Learner modeling for adaptive scaffolding in a computational thinking-based science learning environment. User Model. User-Adap. Inter. **27**(1), 5–53 (2017)
2. Bell, B., Sottilare, R.: Adaptation vectors for instructional agents. In: Sottilare, R.A., Schwarz, J. (eds.) HCII 2019. LNCS, vol. 11597, pp. 3–14. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22341-0_1
3. van den Bosch, K., Harbers, M., Heuvelink, A., van Doesburg, W.: Intelligent agents for training on-board fire fighting. In: Duffy, V.G. (ed.) ICDHM 2009. LNCS, vol. 5620, pp. 463–472. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02809-0_49

4. van den Bosch, K., Helsdingen, A.: Critical thinking in tactical decision games training. In: Noyes, J., Cook, M., Masakowshi, Y. (Eds.) Decision Making in Complex Environments, pp. 213–222. Taylor & Francis, London (2007)
5. Buede, D., DeBlois, B., Maxwell, D., McCarter, B.: Filling the need for intelligent, adaptive non-player characters. In: Proceedings of the Interservice Industry Training, Simulation, and Education Conference (IITSEC) (2013)
6. Cannon-Bowers, J., et al.: Bridging the gap: how to build effective game-based training. In: Tutorial presented at the Interservice Industry Training, Simulation, and Education Conference (IITSEC), p. 26, Orlando (2014)
7. Chrysafiadi, K., Virvou, M.: Student modeling for personalized education: a review of the literature. Advances in Personalized Web-Based Education. ISRL, vol. 78, pp. 1–24. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-12895-5_1
8. Conati, C., Porayska-Pomsta, K., Mavrikis, M.: AI in education needs interpretable machine learning: lessons from open learner modelling. arXiv preprint arXiv:1807.00154 (2018)
9. Csikszentmihalyi, M.: Applications of Flow in Human Development and Education. Springer, Heidelberg (2014)
10. Dahlbäck, N., Jönsson, A., Ahrenberg, L.: Wizard of OZ studies: why and how. In: Proceedings of the 1st International Conference on Intelligent User Interfaces, pp. 193–200 (1993)
11. Deci, E., Ryan, R.: Intrinsic Motivation Inventory (IMI). [WWW document]. Self-Determination Theory. http://www.selfdeterminationtheory.org/intrinsic-motivation-inventory/. Accessed 14 Oct 2005
12. Desmarais, M.C., Baker, R.S.: A review of recent advances in learner and skill modeling in intelligent learning environments. User Model. User-Adap. Inter. **22**(1–2), 9–38 (2012)
13. Farmer, E., Van Rooij, J., Riemersma, J., Jorna, P.: Handbook of Simulator-Based Training. Routledge, Abingdon (2017)
14. Feinstein, A.H., Cannon, H.M.: Constructs of simulation evaluation. Simul. Gaming **33**(4), 425–440 (2002)
15. Fowlkes, J., Dwyer, D.J., Oser, R.L., Salas, E.: Event-based approach to training (EBAT). Int. J. Aviat. Psychol. **8**(3), 209–221 (1998)
16. Freeman, J., Watz, E., Bennett, W.: Adaptive agents for adaptive tactical training: the state of the art and emerging requirements. In: Sottilare, R.A., Schwarz, J. (eds.) HCII 2019. LNCS, vol. 11597, pp. 493–504. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22341-0_39
17. Gonzalez, C.: Training decisions from experience with decision making games. Adaptive technologies for training and education, pp. 167–178 (2012)
18. Hamstra, S.J., Brydges, R., Hatala, R., Zendejas, B., Cook, D.A.: Reconsidering fidelity in simulation-based training. Acad. Med. **89**(3), 387–392 (2014)
19. Harbers, M., Meyer, J.J., Van den Bosch, K.: Explaining simulations through self explaining agents. J. Artif. Soc. Soc. Simul. **12**(3), 6 (2009)
20. Jackson, S.A., Thomas, P.R., Marsh, H.W., Smethurst, C.J.: Relationships between flow, self-concept, psychological skills, and performance. J. Appl. Sport Psychol. **13**(2), 129–153 (2001)
21. Jantke, K.P., Schmidt, B., Schnappauf, R.: Next generation learner modeling by theory of mind model induction. In: CSEDU, vol. 1, pp. 499–506 (2016)
22. Jennett, C., et al.: Measuring and defining the experience of immersion in games. Int. J. Hum. Comput. Stud. **66**(9), 641–661 (2008)
23. Johnson, A., Taatgen, N.: User modeling. In: The Handbook of Human Factors in Web Design, pp. 424–438 (2005)

24. Kahol, K., et al.: Effect of fatigue on psychomotor and cognitive skills. Am. J. Surg. **195**(2), 195–204 (2008)
25. Kenny, P., et al.: Building interactive virtual humans for training environments. In: Proceedings of I/ITSEC, vol. 174, pp. 911–916 (2007)
26. Kirschner, P.A., Sweller, J., Clark, R.E.: Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. Educ. Psychol. **41**(2), 75–86 (2006)
27. Korteling, J.E., Helsdingen, A.S., Baeyer, A.V.: Handbook of low-cost simulation for military training. ELSTAR-EUCLID RTP 11.8, ELS-DEL/5-HB, Index 8, Issue dated 02/05/00 (2000)
28. Laird, J.E., et al.: Interactive task learning. IEEE Intell. Syst. **32**(4), 6–21 (2017)
29. Löckelt, M., Pecourt, E., Pfleger, N.: Balancing narrative control and autonomy for virtual characters in a game scenario. In: Maybury, M., Stock, O., Wahlster, W. (eds.) INTETAIN 2005. LNCS (LNAI), vol. 3814, pp. 251–255. Springer, Heidelberg (2005). https://doi.org/10.1007/11590323_29
30. McHugh, M.L.: Interrater reliability: the kappa statistic. Biochem. Med. **22**(3), 276–282 (2012)
31. Normadhi, N.B.A., Shuib, L., Nasir, H.N.M., Bimba, A., Idris, N., Balakrishnan, V.: Identification of personal traits in adaptive learning environment: systematic literature review. Comput. Educ. **130**, 168–190 (2019)
32. Oprins, E., Visschedijk, G., Roozeboom, M.B., Dankbaar, M., Trooster, W., Schuit, S.C.: The game-based learning evaluation model (GEM): measuring the effectiveness of serious games using a standardised method. Int. J. Technol. Enhanc. Learn. **7**(4), 326 (2015)
33. Peeters, M.: Personalized educational games-developing agent-supported scenario-based training. Ph.D. thesis, SIKS, the Dutch Graduate School for Information and Knowledge Systems (2014)
34. Peeters, M.M., Van Den Bosch, K., Meyer, J.J.C., Neerincx, M.A.: Agent-based personalisation and user modeling for personalised educational games. In: Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization (2016)
35. Pintrich, P.R., Smith, D.A., Garcia, T., McKeachie, W.J.: Reliability and predictive validity of the Motivated Strategies for Learning Questionnaire (MSLQ). Educ. Psychol. Measur. **53**(3), 801–813 (1993)
36. Premlatha, K.R., Geetha, T.V.: Learning content design and learner adaptation for adaptive e-learning environment: a survey. Artif. Intell. Rev. **44**(4), 443–465 (2015)
37. Richardson, M., Abraham, C., Bond, R.: Psychological correlates of university students' academic performance: a systematic review and meta-analysis. Psychol. Bull. **138**(2), 353–387 (2012)
38. Sanders, A.F.: Simulation as a tool in the measurement of human performance. Ergonomics **34**(8), 995–1025 (1991)
39. Scerbo, M.W., Dawson, S.: High fidelity, high performance?: simulation in healthcare. J. Soc. Simul. Healthc. **2**(4), 224–230 (2007)
40. Schank, R.C.: The pragmatics of learning by doing. Pragmat. Soc. **1**(1), 157–172 (2010)
41. Schank, R.C., Berman, T.R., Macpherson, K.A.: Learning by doing. Instr.-Des. Theor. Models: New Paradig. Instr. Theory **2**(2), 161–181 (1999)
42. Schraagen, J.M.: Naturalistic Decision Making. Routledge/Taylor & Francis Group (2018)

43. Sottilare, R., Ragusa, C., Hoffman, M., Goldberg, B.: Characterizing an adaptive tutoring learning effect chain for individual and team tutoring. In: Proceedings of the Interservice/Industry Training Simulation and Education Conference, Orlando, Florida (2013)
44. Sottilare, R.A.: Challenges in moving adaptive training and education from state-of-art to state-of-practice. In: AIED Workshops (2015)
45. Sottilare, R.A., Boyce, M.W.: Elements of adaptive instruction for training and education. In: Schmorrow, D., Fidopiastis, C. (eds.) Foundations of Augmented Cognition: Neuroergonomics and Operational Neuroscience. AC 2016. LNCS, vol. 9744, pp. 85–89. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-39952-2_9
46. Sottilare, R.A., Goodwin, G.A.: Adaptive instructional methods to accelerate learning and enhance learning capacity. In: International Defense and Homeland Security Simulation Workshop (2017)
47. Vandewaetere, M., Desmet, P., Clarebout, G.: The contribution of learner characteristics in the development of computer-based adaptive learning environments. Comput. Hum. Behav. **27**(1), 118–130 (2011)
48. Vygotsky, L.: Interaction between learning and development. Read. Dev. Child. **23**(3), 34–41 (1978)
49. Warwick, W., Rodgers, S.: Wrong in the right way: balancing realism against other constraints in simulation-based training. In: Sottilare, R.A., Schwarz, J. (eds.) HCII 2019. LNCS, vol. 11597, pp. 379–388. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22341-0_30
50. Wouters, P., van Oostendorp, H.: Overview of instructional techniques to facilitate learning and motivation of serious games. In: Wouters, P., van Oostendorp, H. (eds.) Instructional Techniques to Facilitate Learning and Motivation of Serious Games. Advances in Game-Based Learning, pp. 1–16. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-39298-1_1
51. Yildirim, S.: Serious game design for military training. In: Games: Design and Research Conference, Volda University College, pp. 3–4 (2010)