Kampweg 55
3769 DE Soesterberg
P.O. Box 23
3769 ZG Soesterberg
The Netherlands

www.tno.nl

T +31 88 866 15 00
F +31 34 635 39 77

TNO report

TNO 2019 R11607

Virtual Characters that Work: fit-for-purpose behaviour of virtual entities in training simulations

Date	October 2019
Author(s)	Dr. K. van den Bosch T.A.J. Schoonderwoerd, MSc R.A.M. Blankendaal, MSc Dr. M.M.M. Peeters
Classification report	TNO Public
Classified by	maj. E. de Boer
Classification date	9 oktober 2019
Title	TNO Public
Managementuittreksel	TNO Public
Abstract	TNO Public
Report text	TNO Public
Appendices	TNO Public
Number of pages	48 (incl. appendices, excl. RDP & distribution list)
Number of appendices	1

Samenvatting

Achtergrond
Moderne simulatoren en games bieden steeds meer en betere mogelijkheden om militairen te trainen voor allerlei missies en taken. Het benutten van dit potentieel vereist dat de virtuele spelers (de Non-Playing Characters, of NPCs) in de simulatie zich op een geloofwaardige manier gedragen en dat hun gedrag in dienst staat van de te behalen leerdoelen. In de praktijk wordt dit doel lang niet altijd bereikt. Er zijn gevallen waarbij de instructeurs het gedrag van NPCs zo onrealistisch vinden dat zij de simulatie of game niet bruikbaar achten voor de trainingsdoelen. Helaas is het moeilijk om vast te stellen hoe realistisch een NPC moet zijn, aangezien dat betrekking kan hebben op verschillende facetten van gedrag. Het kan bijvoorbeeld betrekking hebben op de emotie die een NPC ervaart; op de manier waarop hij die emotie tot uitdrukking brengt; op het vermogen om emoties van anderen te herkennen; hoe de NPC er uit zet; hoe die beweegt; welke beslissingen die neemt; op de wijze van communiceren; enzovoort. Een leerdoel vereist zelden dat een NPC zich op álle aspecten realistisch moet gedragen. Afhankelijk van de leerdoelen moeten sommige aspecten van gedrag wél geloofwaardig zijn, terwijl andere aspecten ook best weg gelaten kunnen worden, of vereenvoudigd kunnen worden gesimuleerd. De vraag wanneer het gedrag van een NPC realistisch genoeg is om effectief te kunnen worden ingezet voor trainingsdoeleinden is belangrijk voor de Nederlandse Defensieorganisatie, omdat zij in hun opleidingen veel gebruik maken van trainingssimulatoren.

Doel en aanpak
Het doel van dit project is het ontwikkelen van een werkwijze die helpt te bepalen aan welke eisen van geloofwaardigheid de virtuele spelers in een game of simulatie moeten voldoen om geschikt te zijn voor een bepaalde training. Voor het vaststellen van *fit-for-purpose* gedrag moeten contextuele omstandigheden in beschouwing worden genomen, zoals de leerdoelen, gebeurtenissen in een scenario, het vaardigheidsniveau van de lerende, en nog verschillende andere factoren. Dat betekent dat het bepalen van *fit-for-purpose* gedrag van NPCs ingebed moet zijn in het proces van de behoeftestelling en ontwikkeling van een simulator voor een trainingsprogramma.

Uitkomsten en aanbevelingen
Om een NPC *fit-for-purpose* te laten zijn moet het gedrag overeenstemmen met relevante voorschriften, processen, en normen en waarden van de omgeving, en het gedrag van de NPC moet een response ontlokken bij de trainee die geschikt is als zo’n situatie zich in het echt zou voordoen.
Ten behoeve van de ondersteuning van een team dat werkt aan de behoeftestelling en ontwikkeling van een trainingsimulator is een werkwijze ontwikkeld, gebaseerd op wetenschappelijke kennis over de samenhang tussen het realisme van virtuele spelers in simulatoren en leeropbrengst, en door gebruik te maken van bestaande ontwerpmethoden voor trainingen. De werkwijze omvat drie delen, elk bestaande uit een systematische verzameling van opdrachten en vragen aan het team. Het eerste deel is gebaseerd op de *Event-Based Approach to Training* (EBAT) en beoogt de verschijning van een NPC en diens gedrag te identificeren die essentieel zijn voor het behalen van de leerdoelen door de trainee(s).

Het tweede deel ondersteunt het zodanig specificeren van gedrag dat het is afgestemd op de context, het trainingsdoel, en op de regels van geloofwaardigheid. Het derde deel ondersteunt een kritische evaluatie van de verkregen resultaten. De werkwijze is bedoeld als hulpmiddel voor een team. Het biedt ondersteuning bij het bepalen van het belang en bij het afwegen van opties, zodat het team onderbouwde beslissingen kan nemen over welke gedragingen van NPCs in de simulator nodig zijn om de leerdoelen te kunnen behalen.

Summary

Background
Moderns simulators and games provide more and better opportunities to train the military for all kinds of missions and tasks. Realizing this potential demands that the virtual players (the Non Playing Characters, or NPCs) in the simulation behave in a believable fashion and in conjunction with the training objectives. In practice, however, this is not always achieved. There are occasions that instructors evaluate the behavior of NPCs in training simulators as too unrealistic, and therefore disregard the simulation as a valuable training tool. This emphasizes the importance of determining what realism is required in NPC behavior in order to render a simulation valuable for a particular training objective. However, determining appropriate levels of realism is not easy, as realism of behavior may pertain to many aspects. It may, for example, refer to an NPC's emotion expression and emotion recognition, to the way it looks, how it navigates, its style of communication, its decision making, and many more aspects. A particular training objective seldom requires that an NPC acts with a high level of realism on all aspects of its behavior. More typically, it requires that some aspects of behavior are simulated realistically, while for others aspects the realism may be less important, or perhaps be even indifferent. The question what constitutes fit-for-purpose realism of NPCs is of importance to the Netherlands Defense organization, as they make ample use of simulations for training purposes.

Goal & approach
The objective of this project is to develop an approach for determining the required realism of virtual entities in a given training simulation. Determining the fit-for-purpose level of realism of NPC behavior needs to take contextual factors into account, like the learning objective, the scenario, the context of the scenario-event, the competency of the trainee, and many other factors. Thus, determining an appropriate and suitable level of realism should be embedded in the entire cycle of developing a simulation for a training program.

Results and recommendations
In order for an NPC to have 'fit-for-purpose realism', it should adhere to relevant protocols, processes, norms and values of the situational context ("believability"), and it should elicit a response in the trainee that is considered appropriate when the situation would be encountered in the real world. In order to provide support to a team tasked with the development of a training simulation, a working approach was developed to address the need for realism of NPCs. It consists of three parts, each with a systematic series of assignments and questions to the team members. The first part is based upon the Event-Based Approach to Training (EBAT) and aims to identify the NPC behaviors that are essential for the learning goals of the trainee. The second part involves shaping the NPC behaviors to attune them to the situational context, the purpose of training, and to the laws of believability. The third part consists of a guided critical evaluation of the behavior specifications. The working approach is meant to be used as *guidance*, assisting with the consideration of options, enabling the team to make well-founded decisions about which behaviors are necessary to achieve the learning goals of the trainee.

Contents

1 Introduction

The Netherlands Army uses the game "Virtual Battle Space" (VBS) for the tactical training of its soldiers. The game has been in use for more than a decade and is considered a valuable tool for presenting representative tactical problems to the trainees, and to let them experience the outcomes of various course of action. However, training with the game does not always go smoothly. A couple of years back, the school of military tactics prepared a scenario in VBS, and assigned a commander-in-training to lead a reconnaissance team. The mission was to infiltrate enemy area, and to secretly observe troops, means and movements.
The commander-in-training controlled his own avatar in VBS; the avatars of his six team members were controlled by the behavior models: algorithms in the game that govern their behavior. The commander-in-training guided his team along the prepared route and managed to successfully take concealed positions in the enemy area (see Figure 1). From that location, they were able to collect observations, unnoticedly. At one point in the scenario, one enemy soldier (also automatically controlled by algorithms in the game) happened to be passing relatively close to the concealed position of the team. To the surprise and dismay of the commander-in-training, his team members suddenly placed their guns into position and opened fire immediately. Of course, the very next moment the reconnaissance team was discovered and taken down by the enemy. What happened here? The algorithms generating the virtual team members made them fire automatically upon any enemy entity closer than 25 meters. The algorithm did not provide exclusion conditions that would have prevented firing when oneself is being unnoticed, and when on a reconnaissance mission. As a result of the undesired actions of the team mates, the commander-in-training was not able to continue the exercise, and to achieve all the learning objectives. In other words, the models governing the virtual team mates were not fit for their purpose, namely generating the behavior that would enable the trainee to achieve the learning objectives of the exercise.

Figure 1 Scene from VBS, used for the training of a commander to lead a reconnaissance team.

Of course, the incident provoked hilarity among the participants, accompanied with feelings of frustration. Although the game provided an appealing, dynamic, and realistic context for practicing tactical skills, the instructors considered the behavior of the Non-Playing Characters (NPCs)[1] as too unrealistic, and dismissed the game as a training tool. It is beyond doubt that in this particular case, lack of realism thwarted the intended use of the game. But establishing that simulated behavior is too unrealistic is one thing; defining what realism *is* required, is another. For one thing, the incident does not show that high levels of realism are always required. Realism of behavior may pertain to many aspects of behavior. It may, for example, refer to whether emotions and facial expressions are rendered in a true-to-life fashion; it may refer to the way a virtual character responds to unexpected events in the environment (e.g., an alarm, or a loud sound); it may refer to the route a character takes when going from one location to another; it may refer to the naturalness and smoothness of communication with other entities in the simulation; the probability of its physical performance; and it may also refer to the logicality of planning and decision making. The purpose of a simulation seldom requires that a virtual entity acts with a high level of realism on all aspects of its behavior.
More typically, a purpose of a simulation requires some aspects of behavior to be realistic while for others aspects the realism may be less important, or perhaps be even indifferent. To make matters even more complex, a simulation may even require that a virtual entity behaves in a more simplified manner than it would in reality. For example, if a simulation is used to enable a beginning aircraft pilot to practice the procedures for communication with air traffic control, then it may be more fruitful to let the ATC-operator respond in a plain and basic fashion, rather than adding optional additions, even though the latter may perhaps be more in line with reality. So, determining what level of realism is required needs to take into account the purpose of the simulation, the context and the properties of the participants. This is called ‘fit-for-purpose realism’ and is the subject of the present report.

This work is part of the Research Program “Artificial Intelligence for Military Simulations” (AIMS) that TNO is conducting for the Netherlands Ministry of Defense. In order to plan and execute current and future military missions in a safe and effective manner, there is a growing need for simulations that enable the training and preparation of military personnel. Current developments in artificial intelligence make it possible to develop simulations that reflect the appropriate aspects of the mission in a representative and realistic manner. The program AIMS aims to investigate these opportunities.
The question into what constitutes fit-for-purpose realism of virtual entities is of importance to the Netherlands Defense organization, as they make ample use of simulations for many purposes, as for example: training, mission support, doctrine development, system acquisition and validation, and for testing interface designs. The incident described above illustrates that an appropriate level of realism determines whether the objectives can be achieved or not. In this report we focus on the realism of the behavior of virtual entities in training simulations.

[1] Different terms for virtual entities are used in the literature and also in this report, such as non-playing characters; intelligent agents; virtual characters.

2 Project Goal and Scope

The objective of this project is to develop an approach for determining the required realism of virtual entities in a given training simulation. The approach should assist a team that develops a simulation, by pointing out how realism in the behavior of virtual entities relates to achieving the objectives of training. In that way, the modelers of the team are helped to recognize the relevant considerations, to weigh the options, and to make well-founded decisions.

Determining the fit-for-purpose level of realism of NPC behavior is complex. There is no one-answer to the question. The realism of an NPC depends upon the scenario, the learning objective, the context of the scenario-event, the competency of the trainee, and many other factors. These may dynamically change from one stage of the training to the next, from one trainee to the next, from one scenario to the next, and so on. A team tasked with the development of a training simulation wants to be able to address the ultimate need for realism of NPCs by considering the realism requirements for individual training objectives, and then evaluate the set of requirements for all training objectives of the entire training program. This enables the team to make sound decisions regarding the to be developed functionalities of the NPCs.

Thus, determining an appropriate and suitable level of realism is a process that should be embedded in the entire cycle of developing training simulations. Figure 2 illustrates the principal stages involved in the development of simulation-based training.

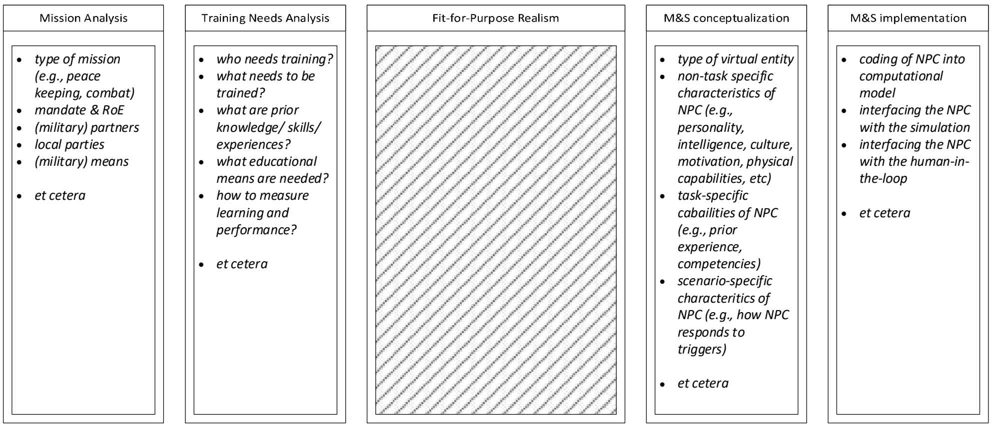

Figure 2 Scoping the determination of behavioral realism of NPCs in training simulations.

On the very left there is the stage of mission analysis. In this stage, the operational objectives are analysed; the strategies that are being used to accomplish the mission objectives; what means are being used; and how the means will be employed (Farmer, Van Rooij, Riemersma, & Jorna, 2017). Having a clear picture of the mission is needed to conduct a Training Needs Analysis (TNA), the second column in the Figure. A TNA often requires a team that includes designers, developers, subject matter experts, and the instructor. Frequently, educational psychologists, instructional designers, and human factors psychologists are also involved because of their expertise in the psychology of learning, skill acquisition, and assessment.

Determining valid, clear and objective learning goals are of vital performance for being able to define NPCs and their behavior, and ultimately, for the success of a training simulator. Learning objectives, in this context, are statements that define the desired outcomes of a training, or a training program, with a simulator. They are defined in terms of skills or knowledge that will be acquired by the trainee. There are many hierarchies that may be used to categorize learning objectives, like e.g., Bloom's taxonomy that classifies learning objectives into levels of complexity and specificity, and distinguishes between the cognitive, affective and sensory domains (Bloom, Engelhart, Furst, Hill, & Krathwohl, 1956). More recently, other taxonomy has been proposed that also cover psychomotor skills (Ferris & Aziz, 2005; Kyllonen & Shute, 1988). Learning goals, when achieved, should be observable or measurable by observable behaviors or actions. Once the learning objectives of the training have been identified, further analysis is performed to address specific tasks, the trainees, and training needs. The TNA helps to distinguish trainees in terms of differences in prior knowledge, skills, and abilities and thus helps to identify the requirements for attaining the training objectives. More important, the TNA also establishes the conditions and metrics for verifying whether the training objectives have been met (Scerbo & Dawson, 2007). Finally, the TNA delivers the scenarios that will enable the trainees to achieve the learning objectives. The literature reports ample guidelines, recommendations, and papers to support these analyses (Farmer et al., 2017).

In theory, the mission and TNA should produce the information that enables a development team to unequivocally specify and develop the simulation (e.g., terrain, objects, weather; et cetera) and the virtual characters (e.g., team mates, opponents, bystanders) that together will produce the a learning environment in which trainees can successfully achieve the defined learning objectives. Unfortunately, in practice this is not always the case, as the incident in the introduction amusingly demonstrates. In fact, designers, modelers, and programmers of training simulators are often unsure as to what their colleagues of the earlier stages may have meant with their outcomes, or they establish that specifications are missing, or are contradictory. Sometimes it is possible to solve such unclarity or undefinedness through consultation later on, but that certainly does not always work. The observation that designers, modelers and programmers often need to make additional assumptions and choices in order to produce a training simulation, does not imply that officers in the earlier stages can be hold responsible. It is, after all, a very difficult task to specify in detail what functionalities of an environment will be needed to present trainees with scenarios that will enable them to successfully achieve the learning objectives. This specification is typically carried out by domain experts and instructors. Exactly because of their expertise they tend to consider some properties of a learning environment as so self-evident that they do not recognize it as a functionality that needs to be explicitly defined (e.g., in the opening anecdote, the functionality that members of a reconnaissance team do not fire when undetected is in the eyes of experts so self-evident that it is conceivable that they missed to include it in the behavioral specification of the NPCs). This difficulty to be aware of one’s own expertise is called tacit knowledge (Reber, 1989).

Given the above observations, it is needed to support the specification of NPCs, shown in Figure 2 of the process as a shaded column. A team should receive assistance in identifying the functional specifications of the virtual entities in a training simulator. Domain experts of the team should be assisted in making their motivations and considerations for particular behavior of NPCs explicit, thereby revealing essential information for the simulation builders. This assistance can take the form of a systematic approach consisting of assignments and questions that induce the domain experts and training experts of the team to think about aspects of NPC-behavior that they would otherwise not pay attention to. In addition, the systematic approach should provide the team expert's with information, knowledge, practical evidence, and existing procedures concerning the relationship between realism of NPC-behavior and training effectiveness. The challenge of developing a first version of such an approach is the topic of this report.

3 Virtual Entities for Simulations

The use of embodied virtual characters and symbolically represented entities within simulated environments has increased considerably in the past two decades. They have been deployed in different application domains such as: crowd simulation (e.g., Pan, Han, Dauber, & Law, 2007); language acquisition, (e.g., Morton & Jack, 2005); education (e.g., Baylor, 2003), and defense (e.g., Li, 2003; Van den Bosch, Harbers, Heuvelink, & van Doesburg, 2009; Van den Bosch, Kerbusch, & Schram, 2012). Moreover, agents have also been used to study social phenomena at a group level (e.g., emotion contagion (Bosse, Duell, Memon, Treur, & van der Wal, 2009)). The use of sophisticated human-like characters is especially relevant when interpersonal interactions are important for the purpose of the simulation.
There are several advantages of using virtual characters in simulations (Korteling, Van den Bosch, & Van der Pal, 2015). First, modeling virtual characters enables accurate control over the type and timing of behaviors that agents can express, and over the interactions they elicit. This control may be used to generate consistently the same behavior of a virtual character across a series of scenario's with slight variations (which may be helpful to drill a particular response in a trainee); control may also be helpful to have the virtual character generate adaptive behavior to variations across a series of scenarios with varying contextual conditions (which may be helpful to let the trainee experience how contextual conditions influence the response of the agents involved) (e.g., Arellano, Varona, & Perales, 2008; McQuiggan, Robison, Phillips, & Lester, 2008).

Second, simulation using NPCs has the potential to be more cost-effective than real life, which often requires big investment in terms of time, budget, materials, and people. It also offers more sustainability because its design can be relatively adjusted to new requirements or purposes.

Finally, there is evidence that virtual characters in simulations show and elicit responses that are representative for real life. This is important because a simulation is the imitation of a system over time, by running (a set of) models that represent the system's components. And any model is a simplification of the real world, hence inherently departing from the real phenomenon on some aspects. The challenge is therefore to include the behavior components that matter to the purpose, and omit or simplify the components that do not. Research has shown that when a virtual agent interacts with a human-in-the-loop, it can elicit the same social-cognitive behavior and performance as a real human being would do (e.g., Bailenson, Yee, Merget, Koslow, & Brave, 2007; Baylor, 2003). However, this similarity is achieved only when the human perceives the agent as sufficiently realistic and credible (Bailenson et al., 2005; Rosenberg-Kima, Baylor, Plant, & Doerr, 2007) and drives the human to perceive and treat the agent as a social entity.

4 Virtual Entities in Training Simulations

Simulation has become a critical technology for military training. It is an indispensable tool, which has benefited from vast improvements in computational power over the last decades. Simulation is especially valuable when reality does not provide suitable and attractive opportunities to train because of risks of casualties, costs of using ‘real’ subjects or systems, legislation, or the limited availability of role players (Korteling et al., 2015). In order to realize the potential of training simulations for the military, modeling the behavior of humans has become more and more important as modern military operations tend to be staged in urban areas, amidst the local population. Soldiers need to learn how to assess the nature and intentions of individuals and groups, and have to learn to predict likely reactions to decisions and actions. Such assessment skills are needed to decide upon appropriate action, which may include a wide range of options from hostile engagement to social communication. Simulation can be a very valuable tool for the training of soldiers in these competencies. A prerequisite is, however, that the behavior of the human(s) involved is adequately modelled for its purpose in the simulation, i.e. that the models are fit for use, or “valid” (Van Hemel, MacMillan, & Zacharias, 2008).

The need for models that generate behavior of NPCs that is adequate and plausible for the specific context, and that is supportive for achieving the purpose of the simulation (i.e., achieving the learning goals) emphasize the importance of the human factor. Although we have become proficient in developing accurate models of physical systems for a simulation, developing models that generate human behavior accurately and realistically has proven to be quite another matter. It is often not known what contextual factors influence someone’s behavior (often subconsciously). Furthermore, there tends to be much more individual variation among humans than in physical systems.

The difficulty to develop *fit-for-purpose* models of human behavior drive the military to utilize another option, namely to have domain experts, instructors, or fellow-students to act as a role player in the simulation. In this way, the role player controls the behavior of the virtual character (e.g., an enemy, a bystander, a team mate). This can be a successful solution, but it has also significant drawbacks. One practical issue is that using SMEs as role players elevates costs of training, and staff is generally scarcely available. Having fellow pupils act as role player is an often used alternative. However, fellow-student role players often fail to act in a fit-for-purpose fashion: their behavior has been reported as often erroneous, too simple, stereotyped, and it often does not fit the context of the simulated scenario (Van den Bosch, Bosse, & de Jong, 2015). Furthermore, the behavior of a virtual entity in a training simulation should be driven by didactic considerations, which are dependent upon the specific trainee and the specific scenario. What is desirable for achieving a particular learning goal for one trainee doesn’t necessarily need to be desirable for another. Role players are often unaware of such considerations, and they generally do not have the expertise to adjust their behavior accordingly. This illustrates the need of investing in the development of human behavior models that contain the expertise and considerations to use them in a fit-for-purpose fashion in training simulations.

5 Realism in Training Simulations

A simulation is a representation of some part of reality. An important demand is that it is representative for the real thing (Peeters, 2014).Yet for a scenario to be representative to a learner it need not necessarily be fully realistic (e.g., Houtkamp, 2012). The simulation should offer support for the learner or decision maker to understand the relations between the situation, their own actions, and the consequences thereof. This is often labelled as "functional fidelity" (e.g., Allen, Buffardi & Hays, 1991; Korteling, Helsdingen & Baeyer, 2000; Neubauer, Khooshabeh & Campbell, 2017). Others have been criticizing this terminology, and propose the term "functional task alignment" (Hamstra, Brydges, Hatala, Zendejas & Cook, 2014). Anyhow, to accomplish that learners acquire a proper understanding about the task and context, and learn how to act appropriately, the simulation should incorporate high levels of realism only for those aspects that are relevant to the task performance and/or situation assessment (Appleton & Lovett, 2003; Korteling et al., 2000; Smallman & John, 2005). All other aspects should be abstracted, so the amount of unnecessary information is reduced, thereby allowing the learner to focus on what is important. This, in turn, also motivates the learner to actively and persistently engage in the scenario (Maehr & Meyer, 1997; Mitchell, 1993; Schunk & Pajares, 2009). Especially in cases where simulation-based training revolves around storylines that follow a narrative structure, the need for realistic behaviour of the characters involved becomes paramount to the effectiveness of the simulation. A narrative structure for a scenario is beneficial for several reasons (Peeters, 2014):

- it guides attention and helps to draw relations between situation-action-effect triplets in a natural and understandable way
- it adds to the entertainment value and engagement
- scenarios facilitate the application of acquired skills in the perspective of the task, thus allowing learners to develop contextualized and integrated competencies.

All of these advantages are valid only if the storylines are representative and believable (Brown, Collins & Duguid, 1989; Kirschner, Beers, Boshuizen & Gijselaers, 2008; McQuiggan et al., 2008; Merrill, 2002; Winn, 1990). However, the behaviour demonstrated by a virtual character does not always necessarily has to be exactly as it would be in real life. What counts in training is whether, and to what degree, the model generates the behavior that helps the trainee to achieve the learning objectives. For example, assume the learning goal: "detecting and correcting errors made by team mates". The scenarios for training this skill should enable the trainee with ample and appropriate practice. In real life, team mates may not make errors very frequently. A virtual character in a training simulation may therefore be modeled in such a way that it makes errors much more often than is likely under real life conditions. This facilitates the trainee to learn. Of course, the nature of the errors that the virtual character makes, and the conditions under which they arise, do need to be in accordance with reality.

Developers of virtual training environments often make great efforts to create realistic simulations. However, in order for a virtual character to have ‘fit-for-purpose realism’, the behavior does not necessarily need to be fully realistic. The behavior should adhere to relevant protocols, processes, norms and values of the situational context, and it should elicit a response in the trainee that is considered appropriate when the situation would be encountered in the real world.

6 Realism and Believability

Realism and believability are closely related terms that should be considered when modeling the behavior of virtual characters, because they have different psychological effects on a human observer (e.g., a trainee). A simulation of human behavior is considered as realistic when it corresponds to the psychological theories and empirical data, and when the behavior is accurately and naturally expressed. That is, realism requires models that generate and express human behavior in accordance with cognitive, social, and emotional theories (Bates, 1994). For example, a virtual character that makes frequent errors under conditions of high workload, and that starts stuttering when confronted with a verbally-aggressive commander, behaves realistically, as the psychological literature predicts that such behavior is likely to occur under such circumstances.

Believability refers to the coherence and consistency of behavior within an individual, and the compatibility of behavior with personal end behavioral characteristics (Avradinis, Panayiotopoulos & Anastassakis, 2013). A virtual character behaves in a believable fashion if it is consistent with its assigned personality; with its motivational and emotional state; with the context and conditions of the world in which the agent is situated, and with the characteristics of the real-life character that it is supposed to represent (Avradinis et al., 2013; Ortony, 2002). Behavior that is believable does not necessarily convey realism. For example, the animated virtual characters appearing in movies of *Walt Disney* and *Pixar Studios* are obviously not realistic. They are nevertheless accepted by the audience as representations of humans because they consistently act in accordance with personality traits, rationales, and emotions that are ‘human-like’. They bring about in people a *suspension of disbelief*. A person has a suspension of disbelief towards a virtual character when he or she does not question its existence as a real entity (Hall, Ap Cenydd & Headleand, 2016). Importantly, the behavior of such unrealistic but highly believable animated characters is sufficiently strong to elicit cognitive and emotional reactions in human observers, and these are similar in type and magnitude to those caused by behavior of real humans (e.g., Bailenson et al., 2007; Baylor, 2003).

Believability *is* required to facilitate immersion and for the desired acceptance of the character. It creates face validity in the perception of the user: the belief that the simulated environment is a relevant representation (Korteling et al., 2000). Believability is a necessary condition to create a deeper sense of presence in the learner. The believability of a virtual character in a simulation may be disrupted in several ways. For example when a character demonstrates an unintelligible lack of situational awareness. Consider two virtual soldiers on guard duty. Both are realistically modelled in terms of their appearance, movements, and communication, and they follow a realistic route for their patrol. Observing their behavior in the virtual world, a human trainee is likely to accept these characters as believable representations of human beings. However, when the soldiers return to their rendezvous point, one of them is shot by an enemy sniper. To the bewilderment of the trainee, he sees the remaining soldier not to have noticed that his mate has been killed, and continues to chat with him. Such incongruent, non-believable behavior is expected, nor accepted by the trainee and will disrupt the suspension of disbelief.

This simulation will no longer be accepted as a representative training environment, thus undermining learning (Lampton, Bliss & Morris, 2002).

To summarize: *Realism* concerns whether the demonstrated appearance and behavior represents that of humans *in general*, under the given circumstances. *Believability* concerns whether the behavior is consistent with the personality, affect, beliefs, and goals of the particular virtual character.

7 Developing Virtual Characters that work

One purpose of this project is to contribute to the development of virtual characters that are fit-for-purpose, often denoted as "functional validity" (Korteling, Van den Bosch & Van Emmerik, 1997; Sanders, 1991). In this case, the development of virtual characters that, when used in training simulations, demonstrate the behavior that is sufficiently realistic and believable for a trainee to achieve the learning goals. In other words: developing virtual characters that work!

In order to develop a systematic working approach that can be of assistance to simulation development teams (see chapter 9), we review and discuss in this chapter the evidence and methods available in the scientific literature and in simulation design.

One might think that virtual characters representing human beings in a simulated environment should ideally be able to demonstrate the full spectrum of human behavior (e.g., movement, speech, posture, facial expressions). However, creating high fidelity virtual agents and agent behaviors is often not at all feasible neither required in order to create a suitable learning environment for simulation-based training or decision support (e.g., Hamstra et al., 2014; Choi et al., 2017). Moreover, it is known that if the appearance and behavior of a virtual character is almost, but not exactly like that of real human beings, this evokes uncanny, feelings of eeriness in the observer or interlocutor. This phenomenon is known as the 'uncanny valley' (Mori, 1970), see Figure 3.

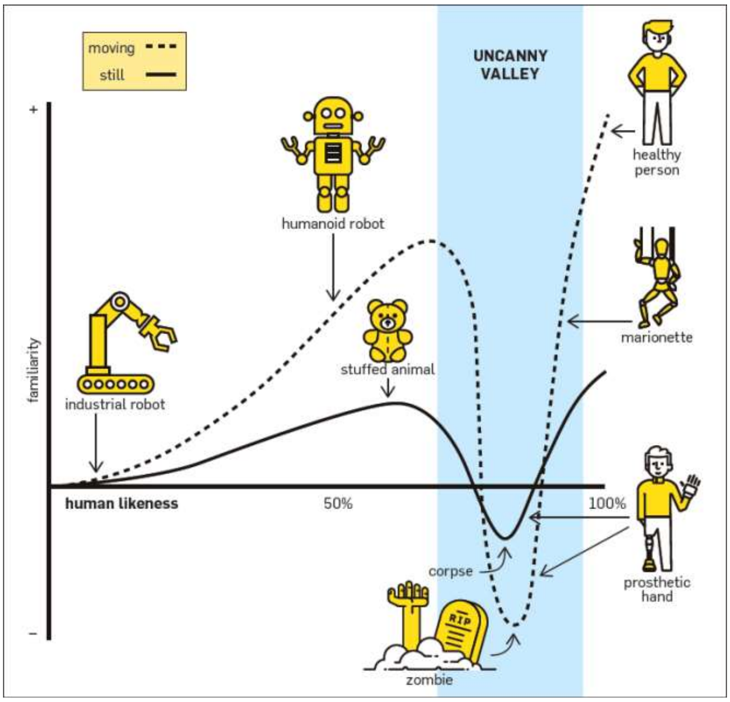

Figure 3 The uncanney valley: a human's response to a virtual character or a robot is positive as its appearance becomes increasingly human. However at some point, as the character appears close to humans, but not quite, it triggers a strong revulsion (the uncanny valley). Only when the appearance becomes practically indistinguishable from that of a human being, the emotional response becomes positive again.

Furthermore, highly realistic behaviors are very difficult to create, and do not necessarily facilitate learning. In fact there is evidence that high-fidelity simulations can sometimes affect transfer of learned material negatively (e.g., Feinstein & Cannon, 2002; Scerbo & Dawson, 2007). For example, 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 (e.g., Martin & Waag, 1978). 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) (Kenny et al., 2007; Korteling et al., 2000). Achieving the optimal level of fidelity is difficult because it requires (1) determining what behaviors are essential to the purpose of the simulation and what behaviors are not, and (2) how to model these human behaviors in such a way that they appear realistic in the perception of the trainee(s).

This is a difficult task, because even domain experts often do not fully understand what task behaviors are essential, and what behaviors can be ignored in a simulation (Harmon et al., 2002). Moreover, behaviors that are not directly related to the task to be trained can nevertheless also be important to the simulation. This may refer to any non-task specific behavior of a virtual character, like e.g., taking an appropriate distance to a conversational partner; selecting an appropriate volume; providing non-verbal signals (*backchanneling*, see Smith et al., 2010) when the partner talks, and so forth. A virtual character needs to display all these subtle behaviors in a way that is consistent with conventions and social norms in order to be accepted as believable by the trainee (a 'willing suspension of disbelief'). Thus, a virtual character should consistently demonstrate believable behavior throughout the entire simulation, including not only task-specific behaviors, but also more 'mundane' behaviors (e.g., walking and looking around, reacting to the player's proximity and other task-unspecific events).

Thus, a virtual character needs to respond in a believable fashion to the situational context of the training scenario. The behavior should adhere to the social norms and structure of the situation (Sunstein, 1996) and take into account any changes in the situational context. This may, for instance, imply reacting suitably to environmental cues (e.g., frightening and ducking away upon the sound of a sudden explosion). It is, however, not always clear as to what and what not constitutes a believable response, as human behavior is the result of complex underlying processes that involve both rational and emotional systems, much of which are not yet understood (Kennedy, 2011).

Developing models that pursue realism om all dimensions of behavior is neither feasible nor necessary. There exist multiple frameworks within the field of Human Behavior Representation that offer feasible alternatives to a full-blown behavioral model, like PRESTO (Busetta & Dragoni, 2015), or PECS (Schmidt, 2000). PRESTO for instance, uses a taxonomy (see Figure 4) of different instantiations of a virtual character (e.g., a young physically fit soldier displays other behavior that an older corpulent soldier; a skilled and highly experienced operator behaves differently than a newbie).

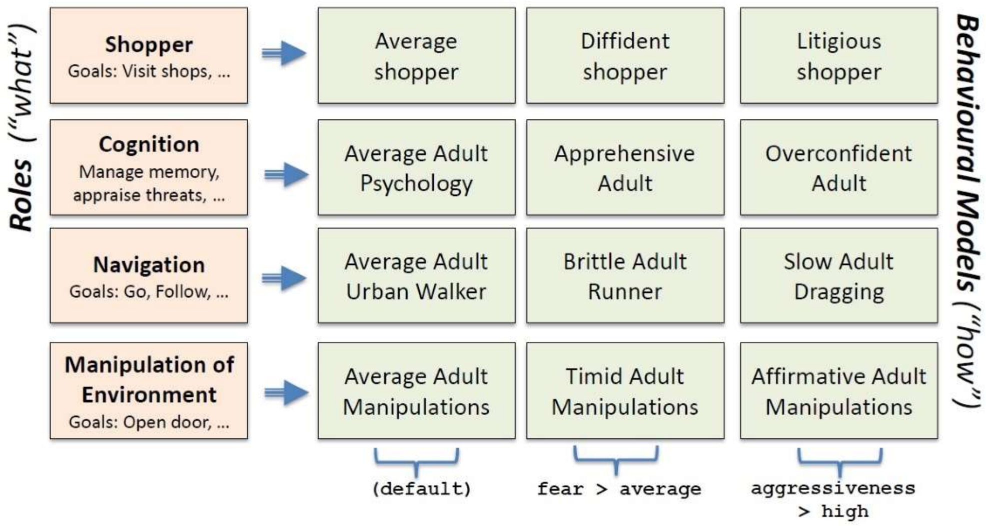

Figure 4 Example of a taxonomy for developing role-appropriate behavior of an NPC

The objective of such frameworks is to define the essential characteristics of humans that are required to accomplish a particular purpose while maintaining practicality and feasibility (Harmon, Hoffman, Gonzalez, Knauf & Barr, 2002).

7.1 Design of Fictional Characters

When it comes to defining virtual characters that behave realistically and that will be perceived as believable by trainees, and that furthermore will bring about an effective and engaging learning situation in the simulation, there is inspiration from the arts. The discussion below is relevant to story-like (training) simulations, such as for example, a simulation that aims to train military commanders in assessing complex mission situations including many parties and individuals. When it comes to defining characters and groups to populate such a simulation, quite a few tips and tricks can be found on dedicated websites about character creation. The most important guideline seems to be that an interesting and believable fictional character features a complex mix of -potentially conflicting- personality traits (Kress, n.d.; Wong Ken, n.d.). Especially the main characters in a scenario should be charismatic, both ordinary and extraordinary; when such a character enters a room, it should draw everyone's attention (Korsmo, 2019). When creating a main character, it can be helpful to create a character profile, describing:

- The character's appearance, e.g. dressing style, face, body;
- The character's body language, e.g. posture, mannerisms;
- The character's psychology, e.g. objectives, fears, desires, needs, personality traits (Big Five), emotions, pride, shame;
- The character's philosophy of life, e.g. religion, spirituality, norms, values, existentialism;
- The character's verbal language, e.g. stock phrases, sound, vocabulary;
- The character's background, e.g. birth place, childhood, history;
- The character's relations with other characters, e.g. what do they value in others and what do they find annoying or abhorrent in others?
- How the character develops during the story.

When picking the traits for a main character that serves as a role model for a trainee, it is recommended to render it with traits that people generally find likeable, such as being reliable, trustworthy, honest, uncondescending, helpful, humoristic, brave, courageous, prepared to self-sacrifice for the greater good, morally conscientious, level-headed, smart in a street-wise fashion, not overly intelligent, kind, compassionate, non-complaining, cool-headed (Now Novel, n.d.). Of course, when the character is an antagonist or an adversary, giving them some unlikeable traits can add to their character and role in the story, e.g. being unreliable, dishonest, condescending, unhelpful, selfish, etc. Either way, 'flat' characters should be omitted, i.e. characters that are "all good" or "all bad": it is better to add a few humanizing flaws and/or likeable traits to make characters more relatable. Characters are considered more interesting when they have conflicting motivations or values leading to inner conflict. However, each character should remain coherent as a whole, or else it will no longer be believable (Now Novel, n.d.).

It is advised to make characters stand out: they should be identifiable in a police line-up. This can be achieved by adding details of dress, gait, and personality (Now Novel, n.d.). Picking the names of your characters with care, by paying attention to symbolism or origin (cf. the characters in the Harry Potter novels), can add extra layers to your characters as well (Now Novel, n.d.). Other ways to make virtual characters interesting is by giving them ethically grey motivations, lines of reasoning, or ways of behaving (Now Novel, n.d.). When a character only plays a minor role in the story, ensure that their behavior is well-described, and equip the character with some unique or memorable feature, yet don't make the character overly complex, as depth is not needed for support characters (Sambruchino, 2013). Their function is primarily to "dress the scene" and should not draw a lot of attention; in fact, they should blend in (Sambruchino, 2013).

When there are multiple characters in a story, one should carefully consider the dynamics between the characters. In general, all characters in the story are there because they impact the protagonist in one way or another (Sambruchino, 2013). For instance, it can be a good idea to add a side-kick to the protagonist that is in stark contrast to the protagonist when it comes to a certain trait. An adversary or antagonist (or both) may be added to introduce obstacles for the protagonist. Interactions between characters, such as yelling, whispering, ignoring, confessing, fighting, or eye contact are useful to reveal your characters' motives without explicitly describing their thoughts and feelings. Having characters reflect on one another can further introduce both their own internal psychology as well as offer a broader view on a character as it is described from different perspectives (Now Novel, n.d.).

7.2 Determining Realism of Virtual Characters' components

The presence and behavior of an agent in a simulation can be expressed in terms its appearance as a whole, consisting of an integration of both physical and behavioral aspects. By identifying these elements it is possible to define a framework (or taxonomy) in which behaviors can be decomposed and described when talking about the realism of virtual agents.

Table 1 Descriptive Categories and Aspects of Virtual Agent Behavior

Virtual agent	**Behavioral aspects**
Physical appearance	Proportions; posture; colors; clothing; facial details.
Physical behavior	Movement (speed, fluency, rotation, direction); speech (vocabulary, intonation, pitch, rate, volume); physical fitness; physiological state (e.g., fatigue).
Cognitive behavior	Knowledgebase (beliefs), long term memory capacity, working memory capacity, perception, situational context (awareness) on different levels (e.g., physical, tactical, civil), doctrines, context-specific skills and drills, implicit and explicit behavioral rules.
Affective behavior	(Expression of) basic emotions such as joy, anger, fear, sadness, surprise, disgust; more complex emotions such as pride, love, jealousy, annoyance, indignation, compassion, interest; motivation.
Interactive behavior	Holding a conversation one-to-one or one-to-many (e.g., addressing a team); adapt communication channel, content, and expression to interaction partner; non-verbal behaviors; working in teams, reacting to entities that are near (e.g., avoiding collisions, making eye-contact).
Personality/Attitudes/Traits	Extraversion/Introversion; emotional stability; risk assessment and sensitivity; sociability; rationality; consciousness.

Table 1 lists the categories in which virtual agent behavior can be decomposed, along with the behavioral aspects that can be distinguished when talking about the realism of agents. These categories are inspired by the taxonomies of Schmidt (Schmidt, 2000), PRESTO (Busetta & Dragoni, 2015), and on our own previous work, e.g., (Van den Bosch et al., 2012; Van den Bosch & Korteling, 2016).
The perceived realism of the virtual agent is determined by these aspects and their interactions. It is important to note that these aspects refer to the expression of behavior, and not necessarily to the underlying processes that are required to enable these behavioral elements. That is, these processes are in service of the behavioral complexity that may be required or preferred in a simulation (i.e., the processes can bring about the necessary variation and flexibility between and within behaviors), and their implementation can vary from rigid scripts, to relatively simple rule-based systems, and even to deep learning methods that may convey subtle details of the modelled processes.

Using the elements described in Table 1 it is possible to define virtual agents that demonstrate a desired level of realism on some or all of the behavioral categories, and their interactions. However, what will be perceived as realistic and what not, is to a large extent determined by the context in which the behavior takes place. People tend to categorize the outside world into schemata (Schank & Abelson, 1975) that request a fairly determined pattern of behavior (e.g., a wedding, dinner at a restaurant; going to the soccer games). Behavior in such schematic situations is largely driven by implicit norms (i.e., cognitive scripts). These norms also produce expectations about the type, nature, and sequence of events and behaviors that are expressed by others (i.e., a virtual character). In order for a virtual agent to behave realistic, the specification of the behavioral aspects need to be considered in relation to the specific nature of the situational context.

In addition, It is known that behavior is only partly produced by conscious and rational processes. Subconscious processes and emotions exert most likely a bigger influence on behavior.

The degree to which rational and perceptive-emotional systems affect behavior differs between people (e.g., Kahneman, 2011; Kennedy, 2011). Importantly, variability in human behavior is not random; there is variability across individuals, groups, and situations (Kennedy, 2011). Thus, an important question when considering virtual agent behavior is what (combinations of) behavioral aspects and behaviors are required to imitate these causal relations in a simulation environment?

Although it may intuitively be tempting to strive for realism on as many behavioral aspects as possible, this would most likely not lead to '*fit for purpose*' behavior of a virtual character, because it does not facilitate and can even decrease effectiveness of learning in the simulation by distracting from learning goals, invoking cognitive overload, and may result in a disruption of the suspension of disbelief (Muckler, 2017; Scerbo & Dawson, 2007). Instead, the learning requirements and context of the simulation should be carefully considered in order to determine what (combinations of) behavioral aspects and underlying processes are essential to be modelled to obtain a fit-for-purpose virtual character. That is, the type of agent behaviors and desired level of realism depend upon the learning goals and on the context of the simulation.

Example of an analysis of fit-for-purpose realism for a virtual agent in a training simulation

Consider a simulated training environment with the goal to teach trainees how to recognize, analyze, and cope with teammates (represented by virtual agents) that are underperforming on a task as a result of an excessive workload.

Figure 5 Impression of a game-based training in how to respond to team mates with high workload stress.

The learning objective implies that a virtual agent should be capable of showing different behaviors that are a function of its current workload. In order to model this behavior, we need to discover whether specific behaviors can be expected from someone that experiences a high workload and how this (and other) behavior changes as a result of an increased workload. Moreover, the behavior occurring in such situations depend upon the underlying process that are responsible for causing high workload in an individual. That is, increased workload can for example be the result of external stressors (e.g., acting within a battle scenario or under time pressure), low working memory capacity of the particular individual, low resilience to stressors, low task-related skills, a high self-awareness, or low self-efficacy, all of which can, in principle, result in (subtle) behavioral differences. As stated in the learning objectives of this example, the trainee has to learn to identify differences in causes of workload in order to be able to choose an effective approach to resolve the situation. These causes should be made derivable from the behaviors or behavioral cues of the one who experiences this workload, so in this training environment it is essential that these behaviors are accurately demonstrated by a virtual agent. Thus, for all training simulations it is important to take careful consideration of the learning goals when determining what level of detail is required for the behavior of virtual characters, and what relations have to be modelled in order to enable this level of detail.

Apart from which types of behavior an agent should be able to demonstrate, and what underlying processes should be modelled for the agent to be able to do so, it is important that agents behave consistently during the simulation in order to be believable in the eyes of the trainee. That is, they should act in accordance with expectations forthcoming from the situation presented in the simulation. This means that agent behavior should fall within the range of what is physically and cognitively possible for humans (e.g., movement speed, perception, reasoning skills, memory), and that they should act according to the cognitive scripts that are appropriate for: the presented training situation (e.g., agents should not carelessly walk around during an intense battle scenario); the agent’s assumed personality (e.g., an agent should not suddenly change from being relatively quiet to very forthcoming during group conversations); the agent’s social and hierarchical role (e.g., interactions with a high-ranked officer should differ from interactions with a petty officer); and other implicit behavioral norms that are relevant in the situational context.

7.3 Approaches to Modeling the Behavior of Virtual Characters

A model (or: Human Behavior Model, HBM) is needed to generate the behavior of a virtual character in a simulation. The behavior should be a function of the character’s individual properties, the training situation, and their interactions.
In many training simulations, virtual character behavior is controlled by defining a list of rules and contingencies. This is a successful approach for the training of tasks that are straightforward (e.g. procedural tasks) and in simulated worlds that can be strictly controlled, so that no situations emerge for which the model of the virtual character cannot produce adequate behavior. The entertainment gaming industry has been successfully using this approach to develop elaborated and complex scripts of input-output rules to control the behavior of virtual characters in their games. Using input-output contingencies make it possible to let a virtual character behave in a quasi-intelligent fashion, provided that the developer anticipated the situation and developed a behavioral rule for it.

The advantage of this approach is extended control over the game, and it has proven to be a robust, error-resistant technology. A disadvantage is, however, that the behavior of virtual players tend to become fairly predictable (especially in the eyes of experienced and skilled human players), hence characters lose their believability and credibility as a (virtual) person.

Some training programs, however, require more freedom on the part of the trainee. This can, for example, be for the training of complex tasks, like e.g., tactical planning and decision making in military operations. An example of such a training simulation in the Netherlands Army would be TACTIS and the “Commandanten Gevechts Trainer” (CGT). In these training simulations, the trainee exerts influence upon the course of the exercise through a number of successive assessments and decisions. And the virtual characters should be able to respond appropriately to the situations that emerge as a result of the trainee’s decisions. It is considered as very hard or even impossible to create a ‘spanning set’ of input-output contingencies specifying appropriate behavior for all possible states that may occur during a scenario (Silverman, 2001). Even in relatively simple tactical scenarios the number of states tend to be very high (Klein, 1998). An approach that is more suitable for this type of applications is to model the behavior of the virtual characters involved as a function of fundamental underlying processes (Zachary, Ryder & Hicinbothom, 1998). Such a model represents the knowledge and processes of an individual or entity in a certain domain, task or scenario.

Most HBMs follow the architectural structure of a behavior engine interacting with a knowledge base to update an internal representation of a (simulated) world (see Figure 6, Harmon et al., 2002).

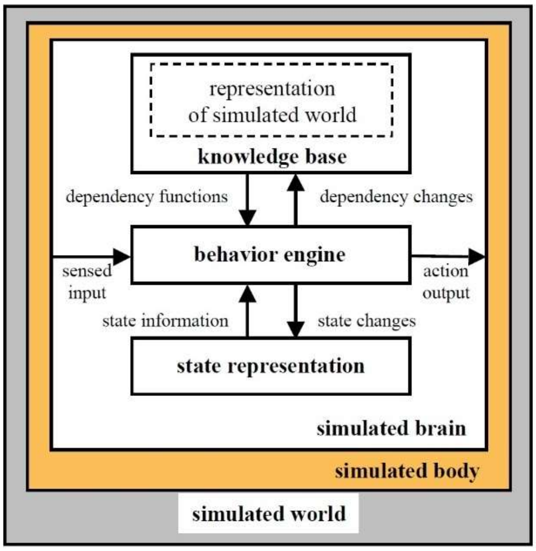

Figure 6 Canonical Model of Human Behavior (Harmon et al., 2002).

There exist different approaches to develop behavior models, each producing behavior with different levels of realism. Often used approaches are the Belief, Desire, Intentions (BDI) approach, and the cognitive modeling approach.

BDI-models: A popular approach is to model behavior as a function of beliefs, desires and intentions (Bratman, 1987; Rao & Georgeff, 1995). BDI (see Figure 7) is fundamentally different from modeling behavior as input-output contingencies

In BDI models, a virtual character is not instructed to act upon a certain state in the scenario, but rather upon *the interpretation* of that state. An event in the world brings about a belief in "the mind" of a character (e.g. hearing a fire alarm creates the belief that the house is on fire). The belief triggers a goal. What goal is triggered by the event depends upon the context and the role of the character: a mother, for instance, may adopt the goal to search for her child; another person may adopt the goal to leave the building quickly; a fire fighter may adopt the goal to locate and fight the fire. The advantage of BDI over input-output contingencies is that BDI-models are more flexible and reusable. For example, using input-output contingencies requires specifying separate actions for each of the following events: "person threatening with rifle"; "person threatening with knife"; "person threatening with hand grenade", and so on for all thinkable threats. In BDI however, all these events activate the belief "I am in danger" that subsequently invokes a common goal, e.g., "escape".

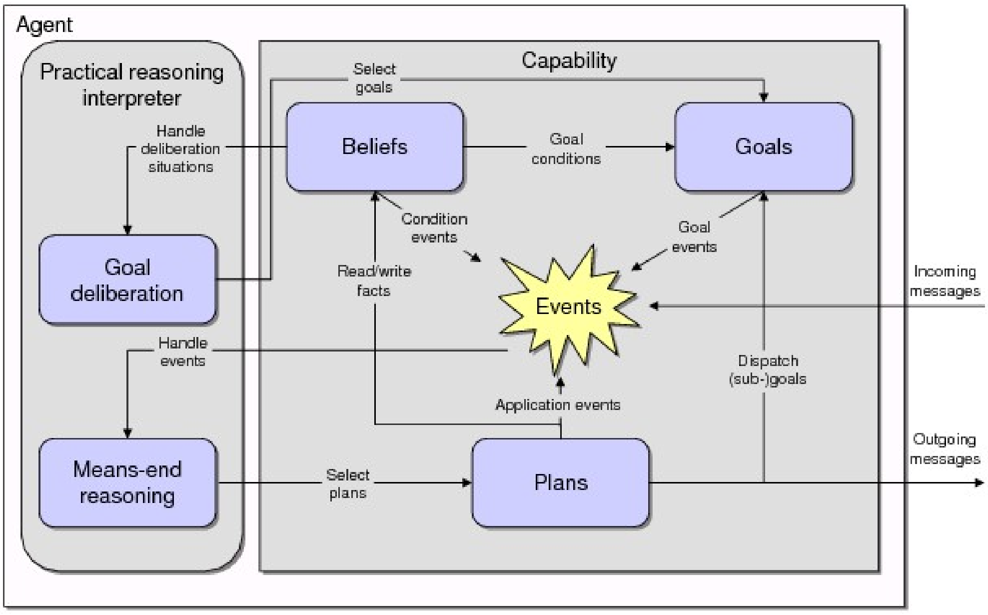

Figure 7 JADEX, an example of a BDI-architecture (Figure taken from: (Pokahr, Braubach & Lamersdorf, 2005)).

Although people have the feeling that their behavior is indeed a function of their beliefs and goals, psychology learns that behavior is actually caused and influenced by many other (often subconscious) factors (e.g. emotion, bias, fatigue, stress, et cetera). And these causal and moderating factors are typically not included in BDI-models (although there have been attempts to include emotion in BDI-models (e.g., Jiang, Vidal & Huhns, 2007). A more fundamental approach than BDI at modeling human behavior is 'cognitive modeling'.

Cognitive models: A cognitive model is a representation of human cognitive processes for the purposes of comprehending and predicting that behavior. Cognitive models may focus on a single cognitive phenomenon or process (e.g., pattern recognition), how two or more processes interact (e.g., pattern recognition and decision making), or to make behavioral predictions for a specific task (e.g., performance in a tactical picture compilation task). Cognitive models can function independently, or they may be embedded in a cognitive architecture. A cognitive architecture represents the conceptual and structural properties of the human mind. There exist many different theories about how the human mind functions; and these theories are associated with different cognitive architectures. Cognitive architectures can be symbolic (e.g., SOAR (Laird, Newell & Rosenbloom, 1987)), connectionist (e.g., PDP, (McClelland, Rumelhart & Group, 1986)), or hybrid (e.g., CLARION, (Sun, 2007)).

Machine Learning: An approach within the field of AI that currently receives much attention is called Machine Learning. It consists of the more advanced techniques (e.g., reinforcement learning, deep learning) and models that enable computers to figure things out from the data and deliver AI applications. Machine Learning is the science of getting computers to act without being explicitly programmed (Genç, 2019). Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to perform the task. Early applications of machine learning were the automatic execution of complex functions, such as email spam filtering and automatic picture recognition. Currently, machine learning also address the development of behavior models (e.g., Knox, 2017).

Whatever modeling approach is selected, a model of human behavior must always be viewed in the context of its purpose (Van Hemel et al., 2008, p.302). If a simple model serves the intended purpose (i.e., if it is “fit for use”), then it should be preferred over a complex one.

7.4 Evaluating Behavior Models of Virtual Characters

In order to achieve an effective training simulation, it is important to evaluate whether the behavior of a virtual character, generated by its Human Behavior Model, is fit-for-its-purpose. Establishing the degree of *face validity* is perhaps the most often used technique for evaluation. In this technique, a subject matter expert (SME) is presented with a series of training scenarios, observes the behavior of the virtual character, and judges whether that behavior will enable the trainee to achieve the learning objective. Determining whether a training simulation has face validity may also be conducted with a trainee as evaluator. This approach evaluates whether the simulation appeals to the intended users (Cronbach, 1949), and is important because this will affect their expectations and the way they will make use of the simulation. Depending on the experience and competency of the evaluator, face validity may provide satisfying answers. However, a practical drawback is that the behavior of the virtual character is often evaluated against implicit and indefinite expectations on realism in the evaluator’s mind. Another drawback is that evaluating SMEs tend to focus on the experienced realism of the virtual character’s behavior, rather than on the question whether the demonstrated behavior supports the intended purpose (e.g. learning complex relationships in a particular domain, or acquiring a particular skill) (Caro, 1977).

Thus, testing face validity is important, but does by itself not provide a complete coverage of the fit-for-purpose validation process.

Another approach is to test in a more explicit and systematic fashion whether the behavior of the virtual character meets the requirements for fit-for-purposeness. This is called *application validity* (Van den Bosch & Korteling, 2016) for a more elaborated discussion on this subject). It evaluates whether a virtual character's model can be used or generalized to the training situations for which the model is intended (Aronson, Wilson, Akert & Fehr, 2004). One way to do this is to administer a training using the model that is assumed to generate fit-for-purpose behavior of a virtual character in a simulation. The effects of training should be evaluated by measuring whether the learning objectives are being achieved (and how fast), whether the trainee is able to use the mastered competencies in new scenarios on the training simulator, and whether the trainee is able to demonstrate the appropriate behavior on the task in real life conditions (transfer). These are the principal measures of application validity. In addition, it may also be measured how the trainee perceives the training, e.g.,: is the behavior of the virtual considered as natural and logical? Is the interaction with the virtual character smooth and clear?; is the training scenario perceived as realistic and useful? Is the training experienced as engaging? For assessing these variables, surveys, questionnaires, ratings, and checklists may be used in combination with quantitative measurements (e.g., time, speed, error).

To evaluate whether a virtual character behaves in a fit-for-purpose fashion in a trainings simulator, it is advocated to combine the methods of face validity and application validity.

8 Towards an Approach for Determining Fit-for-Purpose Realism

The international community of professional organizations that utilize simulation as a technology to train their employees for operational missions experience that the realism of virtual characters is very important, and that inappropriate behavior of virtual characters often renders a simulation as less or not effective. The literature has shown that when a virtual character does show representative behavior, it elicits the same social-cognitive behavior and performance in a trainee as a real human being would do (Bailenson et al., 2007; Baylor, 2003). The use of sophisticated human-like characters is especially relevant when interpersonal interactions are important for the purpose of the simulation (Van den Bosch, Kerbusch & Schram, 2012).

Teams that are involved with the development of training simulations are faced with the task of defining the level of realism of virtual characters' behavior needed for trainees to achieve the learning objectives. This is difficult, as determining what behavior is *fit-for-purpose* differs for each training event: it is dependent upon the nature of the specific learning goal, the skill level of the trainee, and the scenario in the simulation. A development team therefore has to analyse the requirements for the entire set of learning objectives, scenario's and anticipated properties of trainees. Based upon that analyses a team should be able to make founded decisions concerning the required realism of virtual characters. In this report we argue that this analysis requires detailed output ideally forthcoming from the Training Needs Analysis (see Chapter 2). However, in current practice, this output is often not sufficiently good, nor detailed to conduct a proper determination of *fit-for-purpose* realism. There is a need of a method or approach that assists a development team in identifying the specifications that makes a virtual entity of a training simulation *fit-for-purpose*. In this project we developed such an approach by using the knowledge obtained from the literature (as documented in the present report) and combined it with our experience acquired from developing simulation based training programs. The result is a systematic approach consisting of assignments and questions. Furthermore, the approach provides scientific information, knowledge, and practical evidence concerning the relationship between realism of NPC-behavior and training effectiveness. It is important to realize that the approach does not pretend that when it is punctually followed, it guarantees to deliver the final set of requirements that makes a virtual character fit-for-purpose. The present report makes clear that there are too many situational factors involved that would warrant such a claim. However, we hope that the approach induce the domain experts and training experts of the development team to reflect upon the needs of virtual characters that they otherwise would possibly not have paid attention to. The approach is addressed in the next chapter.

9 Approach for Determining Fit-for-Purpose Realism

This working approach is intended to assist a training simulation development team in determining the required realism of behaviors as shown by virtual entities (or Non-Playing Characters; NPCs) in a given training simulation. The approach assists the modeler (and his team) in identifying NPC behaviors, and in specifying these behaviors to obtain the level of realism that increases the likelihood of achieving the training objectives (i.e., realism that is 'fit-for-purpose'). The approach helps the modeler to consider the options for NPC behaviors and make well-founded decisions about which behaviors are necessary to achieve the learning goals of the trainee. The approach is meant to be used as *guidance*. That is, in the end the behavior specifications can only be made by the modeler.

By following this approach, the modeler works towards identifying:

1 The NPC appearance, and its behaviors that are essential to the learning goals (i.e., training objectives) of the trainee, broken down for each (type of) NPC;
2 The characteristics of these behaviors in order to be sufficiently believable and realistic for achieving the learning goals of the training simulation (i.e., fit-for-purpose).

The working approach guides modelers through three parts (see Figure 8).

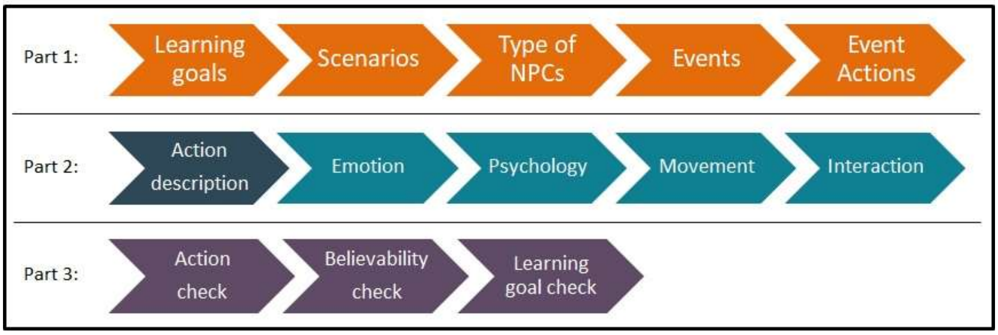

Figure 8 Overview of parts of working approach.

The first part aims to help the modeler to identify the NPC behaviors that are essential for the learning goals of the trainee. This part is based upon the Event-Based Approach to Training (EBAT). The principle of EBAT (see Figure 9) is to define explicit learning events to achieve learning goals, making such events essential for a training simulation (Fowlkes, Dwyer, Oser & Salas, 1998; Oser, Gualtieri, Cannon-Bowers & Salas, 1999) Moreover, EBAT helps to define the requirements for such events, and to specify how they elicit the desired trainee behavior(s).

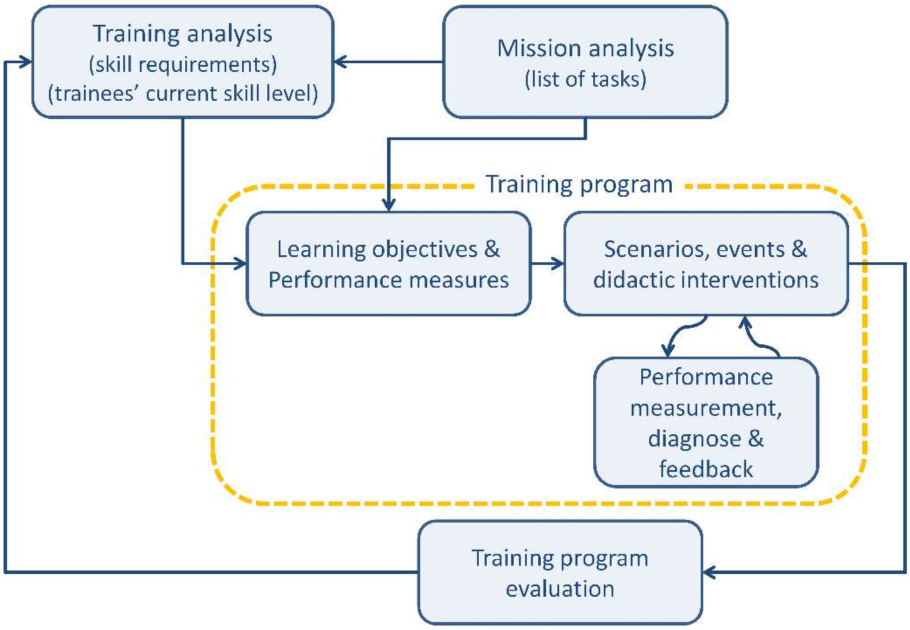

Figure 9 EBAT approach for training (inspired by Fowlkes et al., 1998).

The present approach utilizes the main principle of EBAT: to define the events that in a simulation will bring about the achievement of the learning goals.

In the second part, modelers are incited to think about the characteristics of the behaviors that they have identified as required in part 1. For this part, we analyzed, aggregated and combined available frameworks for human behavior classification. The behavior categories that we identified and which we use in our method are summarized in Table 1. To make the approach workable in practicable settings, we omitted fine-grained details in an attempt to distinguish behavior categories that are distinctive, understandable and useful for simulation development teams.

The third part consists of a verification. The objective of this part is to make modelers aware that first drafts of behavior specifications are rarely complete and oftentimes need further refinement and discussion within the team. The third part of the approach therefore encourages its users to critically evaluate the created behavior specifications in terms of believability, and successively, in terms of whether they support the training objectives.

The approach requires three files that in combination assist the users of the development team: an explanatory 'readme' document, an interactive PowerPoint, and a worksheet documentation[2]. The 'readme'-file (see Appendix 1) contains a 2-page document providing general instructions that explains why, when, and how the approach needs to be used, what results can be expected from using it, and what prior information is required in order to commence the work. It also provides general information concerning the distinction between realism and believability, and tips for the creation of realistic and believable behaviors.

[2] Files can be obtained through the authors

The Powerpoint-file contains the approach itself. It guides users through a structured set of questions that invite them to consider the behavior of NPCs in the simulation (e.g., about the function, causes, timing, and the expression of behavior). The answers that users come up with are documented in the third file of the approach: the *worksheets* document.

The argument to use PowerPoint as the carrying tool of the approach is that it belongs to the standard software collection of most organizations, and people are already familiar with using it. Furthermore, Powerpoint allows for a quick and easy implementation of simple interactions through the use of hyperlinks. A drawback of using this software is that all questionnaire logic should manually be implemented, which limits the interaction that can be achieved. For example, it is not possible to implement control logic that enables a user to go back to a question from which (s)he has been redirected because of a particular answer that was provided. Thus, when users want to go back to a previous question, they have to navigate through the PowerPoint (by using the arrow keys) until they find the question that they have answered.

Assumptions for use: as noted earlier in the report (see chapter 2), the working approach assumes that results of initial analyses of the simulation development process are already available. In particular, it is assumed that the team has a completed Training Needs Analysis available including the following information:
- The learning goals for the trainees;
- A global description of the simulation;
- The scenarios that are to be presented during the simulation;
- The instruction possibilities (e.g., providing feedback, logging information, an after-action review etc.);
- Operational knowledge that should be present in the simulation;
- A general description of the trainees that are expected to participate (e.g., prior experience and domain knowledge);
- A specification of domain-specific features that should be present in the simulation.

9.1 Part 1: Identifying the actions that NPCs should perform

The purpose of the first part of the approach is to assist the team in considering what actions of a NPC are absolutely essential to include in the simulation because they are directly relevant to the learning goals for the trainee. For example, if a learning goal concerns the support and protection of civilians during a military engagement, it is likely to be important that the civilians respond in a certain manner to the events taking place. This part of the approach assists users in defining what that may entail.

This part contains five series of questions (see Figure 10), that in combination should help to fully identify the needed actions of NPCs. The approach specifically encourages the users NOT to take an overall impression of the situation in mind and then use this to generate behaviors of NPCs associated with this type of situation, but instead use the learning goals as a guiding beacon to identify the behavioral requirements of NPC’s. This enables the team to carefully determine what is really needed (and what not) to achieve the learning objectives in the simulation.

It structurally helps the team to identify the relationships between the learning goals, the scenarios, the type of NPCs present, the events taking place in the scenario, and the behaviors of NPCs in this setting that are required to enable trainees to achieve the learning goals.

Figure 10 Flowchart of part 1 of the working method.

Learning goals
The approach requests the user to document the learning goals in the worksheet. The learning goals can be copied from the Training Needs Analysis, but in order to emphasize its central importance in this endeavor, the learning goal is taken as pivot in this part of the approach. It is assumed that the learning goals are in line with the adopted educational strategy of the school and the instructor (see also chapter 2). Some argue that training should involve practice in task situations that gradually increase in complexity (e.g., Munro & Mavin, 2012); others argue that it helps to engage the learner by starting with presenting situations that involve the task domain in its full complexity (Korteling, Helsdingen & Theunissen, 2013). It is important that a development team has the educational strategy in mind, but what strategy should be adopted in a particular simulation is not covered in the present work.

Every question in this part of the working approach is presented from the perspective of the learning goals. This helps users to distinguish between behavior that is essential to the simulation (i.e., fit-for-purpose) and behavior that may add to the realism but is not essential for achieving the learning objective. In addition, the users are asked to describe *at what stage* in the scenario the learning is designed to take place (in the example: prior to the military engagement, immediately after the initiation, or when the combat has stabilized or ended).

Scenarios
The approach request users to define in detail the following characteristics of the scenario (based on Fowlkes, Dwyer, Oser & Salas, 1998):

(1) the task that has to be performed by the trainee
    a. (e.g., reconnaissance of a hostile area)
(2) the starting situation of the scenario
    a. (e.g., trainee (and his team) enters village in hostile area)
(3) the main elements (events) of the scenario
    a. (e.g., blue forces investigate a small village, villagers start coming out of their houses to question the blue forces, red forces suddenly pop up and open fire at people and blue forces, villagers quickly return to their houses, some are hit by enemy fire)
(4) the desired end state
    a. (e.g., player has remained situational awareness and successfully instructed his/her team to lead the villagers to safety while eliminating the opposing forces, resulting in minimal casualties).

Structuring a training scenario in terms of these elements encourages (or requires) modelers to think about what exactly is needed in the scenario to enable a trainee to practice the actions associated with the learning goals (e.g., walking or running towards the village, holding a conversation, shooting, et cetera).

Type of *NPCs*
A training situation in a simulator may involve one or more different virtual characters, or NPCs. All NPCs should demonstrate the behavior that fits their role in the scenario (e.g., allied or enemy soldier, lieutenant, civilian), and the situational circumstances of the moment. All NPCs included in a simulation should have a clear reason as to why they are present. This reason defines what kind of behavior they should demonstrate (e.g., a civilian NPC needs to be able to walk around, run away when event X happens; an allied soldier NPC needs to be able to scan the area for enemies and show offensive and defensive behaviors when the situational circumstances require him to do so). It is important to identify what properties an NPC should have in order to demonstrate the behavior that helps the trainee to achieve his learning goals. For example, an NPC may, or not may not be, resilient to stress, which would result in divergent behavior of the NPC. Depending on the specific learning goal, the property of stress-resilience of the NPC may be important to include in the simulation. Likewise, in the interest of achieving the learning goals, it may be desired to equip an NPC with properties like: level of experience, emotional state; level of fatigue; et cetera.

In order to assist users to differentiate between NPCs, we ask them to define NPCs *not* at a level of individual instantiation (e.g., Lieutenant Bob, Bystander 1), but at the level of the class of an NPC (e.g., NPC from the class of Lieutenant or Bystander). Importantly, the approach supports the definition of NPCs at both individual and group level. That is, for some simulations it could be preferred to define NPCs at group level, for example when simulating a large and diverse crowd of people. Interestingly, humans often use identical terms to describe behavior of individuals (e.g., the man is angry and plans to wreck the market) and behavior of groups (e.g., the crowd is angry and plans to wreck the market) (e.g., Kass & Leake, 1987; O'laughlin & Malle, 2002; Susskind, Maurer & Thakkar, 1999). This suggests that humans consider both individuals and groups to be *agents* and similarly ascribe causes and intentions of behaviors.

Events
The working approach requests the user to indicate, for each NPC, in which training events it is involved. Events are task-specific or contextual triggers or circumstances that are purposely introduced in the scenario to elicit particular behaviour of a trainee (Van den Bosch & Riemersma, 2017). The approach asks the modelers to reflect upon what events are relevant to an NPC (i.e., because the NPC behavior is expected to affect the behavior of the NPC). Thus, events are used as a filter to narrow the behaviors that NPCs should be able to demonstrate. Events can bring about particular NPC actions; events may also be the result of NPC actions; they may also arise as result of other triggers (for example, resulting from an action performed by the trainee, or by another NPC; or it may arise after a particular other event happens in the simulation, or an event may be automatically introduced by design in the learning situation at a particular moment).

Event Actions
When the training scenario has been defined in terms of a series of events, the working approach identifies the responses to be performed by the NPCs. The term ‘actions’ rather than ’behaviors’ is used because it is estimated that users find this easier to think about behaviors of the NPC, as this term is more concrete (‘action’ implies more agency than ‘behavior’). Modelers can simply ask themselves what they think that the NPC should (be able to) *do* in response to a particular event. Modelers are asked to focus on the NPC-actions that are observable to the trainee: those that can be observed by the player. That is, some actions of the NPC will not directly be observable to the trainee (e.g., an NPC that creates a plan, or conducts an action outside the perception range of the trainee. Of course, it can still be important to identify such actions, as the outcomes of the actions *can* be observed by the trainee (e.g., the NPC starts to explain the plan, or provides a report of his actions).

It is possible that a development team, for purposes that may relate to the adopted educational strategy, to define variations in NPC-response to a particular event in a scenario. For example, in one scenario a particular event (e.g., the sound of a loud fire alarm) an NPC may respond calmly and deliberate, while in another version of the scenario an NPC may respond in a stressed and chaotic fashion.
Apart from actions of the NPC that are relevant for the trainee to learn the task, the working approach also helps to specify actions of the NPC that are not directly related to the task to be learned, but are nevertheless to be conducted in a realistic manner in order to create a believable NPC. In particular, such activities need to adhere to the social and domain-specific context of the scenario.

9.2 Part 2: Specifying the behaviors of NPCs

The objective of the first part of the approach was to support modelers in defining the actions of NPCs in the learning events. In this second part, the approach supports modelers in shaping these actions according to the purposes of training, and to the laws of believability. This second part involves five sections. Figure 11 shows the components of this part.

Figure 11 Flowchart of part two of the working method.

Part two supports the analysis of specified NPC-actions in order to develop a better understanding of what processes and conditions affect whether and how the behavior is executed. The distinguished factors of analysis are: emotion, psychology, movement and interaction. Each section step in this part begins with the question whether the factor is relevant for the NPC-behavior, given the purpose of the training simulation. This is achieved by asking “should”-questions, i.e.,: “Should the NPC show one or more particular emotions during this action [*in order for the trainee to achieve its learning objective*]?”.

Action description
In part one, all actions of each NPC in the simulation scenario are defined. In this section of phase 2, the modeler is asked, for each NPC-action, what aspects of the NPC-behavior are perceived by the trainee. For example, an NPC-action may be the driving of a vehicle, and stopping it in front of the trainee-player. The outcome of this part of the approach may yield that it is important that the trainee-player is able to observe *that* the NPC is performing this action (and its consequences), but not necessarily *how* this action is performed. That is, some aspects of the driving need not always be modeled (e.g., steering the wheel, or using the pedals to accelerate and stop the vehicle), as they take place outside the perception of the trainee. Such actions may be described on a more global level, leaving out the specific acts of the NPC.

Emotion
The behavior of people is to a large extent initiated and shaped by affect and emotion, (e.g., Epstein, 1994; Slovic, Fischhoff & Lichtenstein, 1980). NPCs that do not show the influence of these factors on their behavior, tend to be perceived as rigid, and unbelievable. In many cases, training simulation will therefore be better accepted as a plausible environment when trainees recognize the influence of emotion and affect. For example, the trainee expects that the emotional state of a virtual character being present at a huge fire or an explosion, is affected by such events. If the NPCs fail to demonstrate any sign of emotions in their behavior, then the trainee is likely to discard such characters as believable representations of real people. Emotion and affect may also be directly relevant to tasks to be learned. For example, when a trainee has to learn to assess the emotional state of someone, and subsequently to select an action that fits the assessed emotional state, then a virtual character in a training simulation needs to be able to show the relevant emotions, and to express the influence of these emotions in an appropriate manner in its appearance and behavior.
The approach guides the users in asking whether any emotions of an NPC are fixed or adaptive. If they are fixed, then the NPC preserves this state throughout the scenario. If they are adaptive, the NPC’s emotional state may be affected by the nature of the events in the scenario, including the behavior of the trainee.

Psychology
Human behavior is to a large extent determined by the environment in interaction with the unique combination of an individual’s personality, past experiences, intellect, sentiment, knowledge, and motivation. That explains why there is so much variation in the behavior of people. People are aware of this, and they expect other people to be influenced by such psychological factors, and they often explain their behavior in such terms (for example: “He likes to be the center of attention, because he is very extraverted.”). Moreover, people often assume that such characteristics are relatively stable over time and generally affect all behaviors of an individual.

To enhance the believability of an NPC’s behavior in a simulation, it is important that trainees experience the behavior as being consistent with the given or assumed characteristics of the NPC. This is particularly important when the properties assigned to the NPC are relevant for achieving the learning goals (for example, when a trainee has to learn to identify different types of task-induced stress in other people).

There are many theories on how individual differences affect the behavior and performance of people (e.g., Motowildo, Borman & Schmit, 1997). It is beyond the scope of the current working approach to systematically reveal the implications of the vast literature on this topic for the specification of NPC behavior, but Table 2 gives an impression, and should be treated for illustration purposes only.

Table 2 Illustration of how some psychological factors influence behaviour.

Psychological aspect	**Effects on behavior**
Extraversion / introversion	Extraverted: more enthusiastic, talkative, assertive, energized, seeks social contact. Introverted: more reserved and docile, less talkative, think before they speak, seeks solidarity
Knowledge & skills	Impacts task performance, reasoning, decision-making, and problem-solving, movement certainty
Memory capacity	Impacts learning, task performance, remembering past events/agreements
Stress & perceived workload	High: lower task performance, poorer memory and attention, jerky movements, tight facial expression Low: relaxed movements, decreases task performance (when too low), boredom, decreased attention
Attention span	Low attention span: easily distracted (e.g., from sounds, visuals), poorer learning and memory, … High attention span: focused, higher task performance, high listening skills,
Motivation	Low: less enthusiastic, less concentrated, movement show minimal effort to complete action, .. High: enthusiastic, pro-active behavior, focused, …

A table like this intends to support a modeler to define how behavior and performance should be affected by individual characteristics, and to select the levels that should produce the desired behavior of an NPC (e.g., an NPC that prior to an event has a low stress level and thus behaves in relaxed attitude, and becomes more stressed after the introduction of a stressful event, resulting in hectic movements, and busy presence).

Appearance and Movement

If an NPC is embodied (an intelligent agent that interacts with the environment through a virtual physical body within the simulated environment), then the trainee is able to perceive the NPC in the simulation. The NPC navigates and moves through the simulated environment, and should do so in a manner that it not disrupts the trainee’s willful suspension of disbelief in the training simulation (see chapter 6).

The approach requests the modeler to make a distinction between intentional and unintentional movements of the NPC. Intentional movements are initiated by the NPC itself, such as walking, looking around or writing. Unintentional movements are caused by external factors, such as making a turn in a vehicle or the rocking of a boat.

Interaction
An NPC in a training simulation is often an active and responsive entity, purposely designed to be a component of the simulation in order to enable the trainee to achieve its learning objectives. This implies that an NPC is capable of interacting with its environment, and that these interactions may affect the nature and state of the environment. The working approach arranges the modeler's attention to these interactions by requesting the modeler to make a distinction between NPC-object interaction (e.g., an NPC using a computer, or driving a vehicle) and NPC-human interaction (e.g., NPC communicating with the trainee, or with other NPCs). For NPC-object interactions, the approach requests modelers to develop a description of the interaction based upon information on what the object is, how the interaction will be perceived by the trainee-player, and whether this interaction should be consistent with domain-specific procedures. For NPC-human interactions, the approach supports the modeler to develop a description of the verbal (vocabulary, intonation, pitch, rate, et cetara) and non-verbal behavior (eye-contact, hand movements, et cetera).

9.3 Part 3: Evaluation

In the first part, the actions for each NPC are determined. In the second part, these actions are described and further specified with possible factors that could influence the performance of the action, such as emotion, psychology, movement and interaction. The output of these parts are behavioral specifications of NPCs, documented in the worksheets.
As has been already argued, the approach intends to assists a development team in the process; to help modelers in defining virtual characters that are fit-for-the-purpose-of-training. This final part of the approach entails a reflective evaluation of the collected information and decisions. The modeler is guided to think critically whether the provided information really produced the required behaviors for all learning goals of the simulation. In addition, the approach also encourages to reflect upon whether less necessary behaviors have been identified also. Finally, the check also involves the believability and consistency of the actions, and whether behavior will be believable over the course of a simulation scenario.
This part has three sections (see Figure 12).

Figure 12 Flowchart of part three of the working method.

Action check
The modeler is requested to reflect upon whether the described behavior will support the achievement of the corresponding learning goal. If this brings about doubt in the development team, the user is referred to appropriate sections of part 2 of the approach, for refinement or repair of the NPC-specifications.

Believability check
In this part, the modeler is requested to evaluate whether the defined NPC-behavior is expected to come across as believable to a trainee. This involves, for example, the believability of emotional and psychological states of the NPC over the course of a training scenario (e.g., an NPC that is sad, is unlikely to be very happy and enthusiastic at an adjacent moment in the scenario). In addition, emotions should be represented in a believable fashion as well. If they change during the simulation as a result of events and (inter)actions that take place, it should be described how and at what pace the emotional state of the NPC transforms. Important in this step is that any changes to psychological states of the NPC are logical and defendable from the perspective of the scenario. If the modeling team has the feeling that these requirements are not met, they are referred to appropriate sections of part 2 of the approach.

Learning goal check
The most important check of this approach is the question whether the behaviors of the NPC are indeed in service of the purpose of the simulation: Learning. The approach supports this by addressing the question whether all learning objectives are covered by associated learning events in the training scenarios, and whether the NPCs respond to these learning events in a manner that will enable the trainee to achieve its learning objectives. If this checks points out that the description of NPCs do not fully cover the requirements, then the modeler team is referred to appropriate sections elsewhere in the approach, where the necessary refinements and corrections can be made. In the end, the behaviors that are absolutely essential to achieve the learning goals are described in more detail compared to the less relevant behaviors.

10 Conclusion

Military simulations and games provide high potential for training. They make it possible to present situations in a realistic and dynamic fashion, enabling trainees to learn the knowledge and skills required for performance in the operational world. To achieve this potential, the virtual characters in such simulations and games (the Non-Playing Characters, NPCs) should behave in a realistic and purposeful manner, in other words: in a *fit-for-purpose* manner. However, defining what behavior is fit-for-purpose is not easy. Yet the question is important to the Netherlands Defense organization, as they make ample use of simulations for training purposes. Determining fit-for-purpose behavior needs to take into account the purpose of the simulation training, the context, and the properties of the participants.

This report provide a working approach to be used for determining fit-for-purpose behavior of NPCs, given a particular simulation training. It is intended to be used by a team, tasked with requirement and development. This assistance has the form of a systematic series of questions and assignments that provide the team expert's with information, knowledge, practical evidence, and existing procedures concerning the relationship between realism of NPC-behavior and training effectiveness. When procuring, developing, or updating training simulations, the Netherlands Defence is recommended to utilize this approach as it will guide the team in considering the options for NPC behaviors and to make well-founded decisions about which behaviors are necessary to achieve the learning goals.

11 References

Allen, J., Buffardi, L. & Hays, R. (1991). *The relationship of simulator fidelity to task and performance variables*. Technical Report. George Mason Univ Fairfax Va, Dept of Psychology.

Appleton, K. & Lovett, A. (2003). GIS-based visualisation of rural landscapes: defining 'sufficient' realism for environmental decision-making. *Landscape and Urban Planning*, *65*(3), 117–131.

Arellano, D., Varona, J. & Perales, F. J. (2008). Generation and visualization of emotional states in virtual characters. *Computer Animation and Virtual Worlds*, *19*(3-4), 259–270.

Aronson, E., Wilson, T. D., Akert, R. M. & Fehr, B. (2004). Social psychology (2nd Canadian ed.). *Toronto, ON, Canada: Pearson Education*.

Avradinis, N., Panayiotopoulos, T. & Anastassakis, G. (2013). *Behavior believability in virtual worlds: Agents acting when they need to*. https://doi.org/10.1186/2193-1801-2-246

Bailenson, J. N., Swinth, K., Hoyt, C., Persky, S., Dimov, A. & Blascovich, J. (2005). The Independent and Interactive Effects of Embodied-Agent Appearance and Behavior on Self-Report, Cognitive, and Behavioral Markers of Copresence in Immersive Virtual Environments. *Presence: Teleoperators and Virtual Environments*, *14*(4), 379–393. https://doi.org/10.1162/105474605774785235

Bailenson, J. N., Yee, N., Merget, D., Koslow, D. & Brave, S. (2007). Virtual Interpersonal Touch: Expressing and Recognizing Emotions Through Haptic Devices. *HUMAN-COMPUTER INTERACTION*, *22*, 325–353. https://doi.org/10.1080/07370020701493509

Bates, J. (1994). *The Role of Emotion in Believable Agents*. Retrieved from https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.47.8186&rep=rep1&type=pdf

Baylor, A. L. (2003). The impact of three pedagogical agent roles. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems* (pp. 928-929). ACM.

Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H. & Krathwohl, D. R. (1956). *Taxonomy of Educational Objectives: The Classification of Educational Goals*. McKay Publishing, New York.

Bosse, T., Duell, R., Memon, Z. A., Treur, J. & van der Wal, C. N. (2009). *A Multi-agent Model for Emotion Contagion Spirals Integrated within a Supporting Ambient Agent Model BT - Principles of Practice in Multi-Agent Systems* (J.-J. Yang, M. Yokoo, T. Ito, Z. Jin & P. Scerri, Eds.). Springer Berlin Heidelberg.

Bratman, M. (1987). *Intention, plans, and practical reason* (Vol. 10). Harvard University Press Cambridge, MA.

Brown, J. S., Collins, A. & Duguid, P. (1989). Situated Cognition and the Culture of Learning. *Educational Researcher*, *18*(1), 32–42. https://doi.org/10.3102/0013189X018001032

Busetta, P. & Dragoni, M. (2015). *Composing Cognitive Agents from Behavioural Models in PRESTO*. Retrieved from http://ceur-ws.org/Vol-1382/paper13.pdf

Caro, F. G. (1977). *Readings in evaluation research*. Russell Sage Foundation.

Cronbach, L. J. (1949). *Essentials of psychological testing.*

Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American psychologist*, *49*(8), 709.

Farmer, E., Van Rooij, J., Riemersma, J. & Jorna, P. (2017). *Handbook of simulator-based training*. Routledge.

Ferris, T. L. & Aziz, S. (2005). *A psychomotor skills extension to Bloom's taxonomy of education objectives for engineering education* (PhD Thesis). National Cheng Kung University Tainan.

Fowlkes, J., Dwyer, D. J., Oser, R. L. & Salas, E. (1998). Event-Based Approach to Training (EBAT). *The International Journal of Aviation Psychology*, *8*(3), 209–221. https://doi.org/10.1207/s15327108ijap0803_3

Genç, Ö. (2019). Notes on Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL)…. Retrieved May 10, 2019, from https://towardsdatascience.com/notes-on-artificial-intelligence-ai-machine-learning-ml-and-deep-learning-dl-for-56e51a2071c2

Hall, J. A., Ap Cenydd, L. & Headleand, C. J. (2016). Identifying Human-Like Qualities for Non-player Characters. *Artificial Life and Intelligent Agents Symposium*, 117–129. Springer.

Hamstra, S. J., Brydges, R., Hatala, R., Zendejas, B. & Cook, D. A. (2014). Reconsidering fidelity in simulation-based training. *Academic Medicine*, *89*(3), 387–392.

Harmon, S. Y., Hoffman, C. W. D., Gonzalez, A. J., Knauf, R. & Barr, V. B. (2002). Validation of human behavior representations. *Foundations for V&V in the 21st Century Workshop (Foundations' 02)*, 22–24.

Houtkamp, J. M. (2012). *Affective appraisal of virtual environments* (Utrecht University). Retrieved from https://dspace.library.uu.nl/handle/1874/243556

Jiang, H., Vidal, J. M. & Huhns, M. N. (2007, May). EBDI: an architecture for emotional agents. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems* (p. 11). ACM.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kass, A. & Leake, D. (1987). *Types of Explanations.* Technical Report No. YALEU/CSD/RR-523). Yale Univ New Haven; Ct Dept Of Computer Science.

Kennedy, W. G. (2011). The Roots of Trust: Cognition Beyond Rational. *BICA*, 188–193.

Kenny, P., Hartholt, A., Gratch, J., Swartout, W., Traum, D., Marsella, S. & Piepol, D. (2007). Building interactive virtual humans for training environments. *Proceedings of i/Itsec*, *174*, 911–916.

Kirschner, P. A., Beers, P. J., Boshuizen, H. P. A. & Gijselaers, W. H. (2008). Coercing shared knowledge in collaborative learning environments. *Computers in Human Behavior*, *24*(2), 403–420.

Klein, G. (1998). Sources of Power: How people make decisions. 1998. *MIT Press, ISBN*, *13*, 978–0.

Knox, B. (2017). Building character AI through machine learning – MIT MEDIA LAB – Medium. Retrieved May 10, 2019, from https://medium.com/mit-media-lab/building-character-ai-through-machine-learning-7a3159dc4940

Korsmo, J. (2019). How to Write a Character Analysis. Retrieved from: https://www.wikihow.com/Write-a-Character-Analysis, on July 8th, 2019.

Korteling, H., Helsdingen, A. & Theunissen, N. C. (2013). Serious gaming@ work: Learning job-related competencies using serious gaming. In *The psychology of digital media at work* (pp. 129–150). Psychology Press.

Korteling, J. E., Bosch, van den, K. & Pal, van der, J. (2015). *Effectiviteitsfactoren van simulatiemiddelen voor opleiding en training* (Memo No. 14M11842). Soesterberg: TNO.

Korteling, J. E., Helsdingen, A. S. & Baeyer, A. von. (2000). Handbook of low-cost simulation for military training. *ELSTAR-EUCLID RTP 11.8, ELS-DEL/5-HB, Index 8, Issue Dated 02/05/00*.

Korteling, J. E., Van den Bosch, K. & Van Emmerik, M. L. (1997). Low-cost simulators 1a: Literature review, analysis of military training, and selection of task donamins. *The Netherlands: TNO Human Factors Research Institute. TNO-Report TM-97-A035*.

Kress, N. (n.d.). 4 Ways to Motivate Characters and Plot | Writer's Digest. Retrieved May 17, 2019, from https://www.writersdigest.com/online-editor/4-ways-to-motivate-characters-and-plot

Kyllonen, P. C. & Shute, V. J. (1988). *Taxonomy of learning skills*. Technical Report. Universal Energy Systems, Inc Dayton Oh.

Laird, J. E., Newell, A. & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, *33*(1), 1–64.

Lampton, D. R., Bliss, J. P. & Morris, C. S. (2002). Human performance measurement in virtual environments. In *Handbook of Virtual Environments* (pp. 741–760). CRC Press.

Li, Z. (2003). *The potential of America's Army, the video game as civilian-military public sphere* (Doctoral dissertation, Massachusetts Institute of Technology).

Livingstone, D. (2006). Turing's test and believable AI in games. *Computers in Entertainment* (CIE), 4(1), 6.

Maehr, M. L. & Meyer, H. A. (1997). Understanding motivation and schooling: Where we've been, where we are, and where we need to go. *Educational Psychology Review*, *9*(4), 371–409.

McClelland, J. L., Rumelhart, D. E. & Group, P. R. (1986). Parallel distributed processing. *Explorations in the Microstructure of Cognition*, *2*, 216–271.

McQuiggan, S. W., Robison, J. L., Phillips, R. & Lester, J. C. (2008). Modeling parallel and reactive empathy in virtual agents: An inductive approach. *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 167–174. International Foundation for Autonomous Agents and Multiagent Systems.

Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development*, *50*(3), 43–59.

Mitchell, M. (1993). Situational interest: Its multifaceted structure in the secondary school mathematics classroom. *Journal of Educational Psychology*, *85*(3), 424–424.

Morton, H. & Jack, M. A. (2005). Scenario-based spoken interaction with virtual agents. *Computer Assisted Language Learning*, *18*(3), 171–191.

Motowildo, S. J., Borman, W. C. & Schmit, M. J. (1997). A theory of individual differences in task and contextual performance. *Human Performance*, *10*(2), 71–83.

Muckler, V. C. (2017). Exploring suspension of disbelief during simulation-based learning. *Clinical Simulation in Nursing*, *13*(1), 3–9.

Munro, I. & Mavin, T. J. (2012). Crawl-walk-run. *10th International Symposium of the Australian Aviation Psychology Association, Sydney, Australia*.

Neubauer, C., Khooshabeh, P. & Campbell, J. (2017). When Less is More: Studying the Role of Functional Fidelity in a Low Fidelity Mixed-Reality Tank Simulator. *International Conference on Applied Human Factors and Ergonomics*, 220–229. Springer.

Novel Writing Help. (n.d.). The First Rule of Creating Fictional Characters | Novel Writing Help. Retrieved May 17, 2019, from https://www.novel-writing-help.com/creating-fictional-characters.html

NowNovel. (n.d.). How to Write a Novel Now and Finish | Now Novel. Retrieved May 10, 2019, from https://www.nownovel.com/

O'laughlin, M. J. & Malle, B. F. (2002). How people explain actions performed by groups and individuals. *Journal of Personality and Social Psychology*, *82*(1), 33–33.

Ortony, A. (2002). On making believable emotional agents believable. *Trappl et al.(Eds.)(2002)*, 189–211.

Oser, R. L., Gualtieri, J. W., Cannon-Bowers, J. A. & Salas, E. (1999). Training team problem solving skills: An event-based approach. *Computers in Human Behavior*, *15*(3–4), 441–462.

Pan, X., Han, C. S., Dauber, K. & Law, K. H. (2007). A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations. *Ai & Society*, *22*(2), 113–132.

Peeters, M. (2014). *Personalized educational games-developing agent-supported scenario-based training* (PhD Thesis). SIKS, the Dutch Graduate School for Information and Knowledge Systems.

Pokahr, A., Braubach, L. & Lamersdorf, W. (2005). Jadex: A BDI reasoning engine. In *Multi-agent programming* (pp. 149–174). Springer.

Rao, A. S. & Georgeff, M. P. (1995). BDI agents: From theory to practice. *ICMAS*, *95*, 312–319.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*(3), 219.

Rosenberg-Kima, R. B., Baylor, A. L., Plant, E. A. & Doerr, C. E. (2007). The importance of interface agent visual presence: Voice alone is less effective in impacting young women's attitudes toward engineering. *International Conference on Persuasive Technology*, 214–222. Springer.

Sambuchino, C. (2013). How to write a novel. Retrieved from: https://www.wikihow.com/Write-a-Character-Analysis, on July 9th, 2019.

Sanders, A. F. (1991). Simulation as a tool in the measurement of human performance. *Ergonomics*, *34*(8), 995–1025.

Scerbo, M. W. & Dawson, S. (2007). High Fidelity, High Performance?: *Simulation In Healthcare: The Journal of the Society for Simulation in Healthcare*, 2(4), 224–230. https://doi.org/10.1097/SIH.0b013e31815c25f1

Schank, R. C. & Abelson, R. P. (1975). Scripts, plans, and knowledge. *IJCAI*, 151–157.

Schmidt, B. (2000). *The modelling of human behaviour: The PECS reference models*. SCS-Europe BVBA Delft.

Schunk, D. H. & Pajares, F. (2009). Self-efficacy theory. *Handbook of Motivation at School*, 35–53.

Silverman, B. G. (2001). *More Realistic Human Behavior Models for Agents in Virtual Worlds: Emotion, Stress, and Value Ontologies*. Retrieved from http://repository.upenn.edu/hmshttp://repository.upenn.edu/hms/34

Slovic, P., Fischhoff, B. & Lichtenstein, S. (1980). Facts and fears: Understanding perceived risk. In *Societal risk assessment* (pp. 181–216). Springer.

Smallman, H. S. & John, M. (2005). Naive realism: Limits of realism as a display principle. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *49*, 1564–1568. Sage Publications Sage CA: Los Angeles, CA.

Smith, C., Crook, N., Boye, J., Charlton, D., Dobnik, S., Pizzi, D., … Turunen, M. (2010). Interaction strategies for an affective conversational agent. *International Conference on Intelligent Virtual Agents*, 301–314. Springer.

Sun, R. (2007). The importance of cognitive architectures: An analysis based on CLARION. *Journal of Experimental & Theoretical Artificial Intelligence*, *19*(2), 159–193.

Susskind, J., Maurer, J. & Thakkar, K. (1999). Perceiving individuals and groups: Expectancies, dispositional inferences, and causal attributions. *Journal Journal of Personality and Social Psychology*, *76*(2). https://doi.org/10.1037/0022-3514.76.2.181

Van den Bosch, K., Harbers, M., Heuvelink, A. & van Doesburg, W. (2009). Intelligent agents for training on-board fire fighting. *International Conference on Digital Human Modeling*, 463–472. Springer.

Van den Bosch, K., Kerbusch, P. & Schram, J. (2012). Modeling cultural behavior for military virtual training. *Proceedings of the Interservice/Industry Training, Simulation & Education Conference (I/ITSEC'12)*.

Van den Bosch, K. & Korteling, J. E. (2016). *Validating models of human behavior* (No. R11848). Soesterberg, the Netherlands: TNO.

Van den Bosch, K. & Riemersma, J. B. (2017). Reflections on scenario-based training in tactical command. In *Scaled Worlds: Development, validation and applications* (pp. 11–31). Routledge.

Van den Bosch, K., Bosse, T. & Jong, S. de. (2015). Trainen met gesimuleerde mensen: Effectief inzetten van gedragsmodellen voor militaire training. *Militaire Spectator, September, 9, 184, 358-373*.

Van den Bosch, K., Kerbusch, P. & Schram, J. (2012). Modeling cultural behavior for military virtual training. *Proceedings of the Interservice/Industry Training, Simulation & Education Conference (I/ITSEC'12)*.

Van Hemel, S. B., MacMillan, J. & Zacharias, G. (2008). *Behavioral modeling and simulation: From individuals to societies*. National Academies Press.

Winn, W. (1990). Some implications of cognitive theory for instructional design. *Instructional Science*, *19*(1), 53–69.

Wong Ken, S. (n.d.). 4 Ways to Create a Realistic Fiction Character - wikiHow. Retrieved May 17, 2019, from https://www.wikihow.com/Create-a-Realistic-Fiction-Character

Zachary, W., Ryder, J. M. & Hicinbothom, J. H. (1998). Cognitive task analysis and modeling of decision making in complex environments. *Making Decisions under Stress: Implications for Individual and Team Training*, 315–344.

A Read-Me for Working Approach

Working Approach for Specifying Non-Player Characters Behaviors in Training Simulations

Why using this working approach?
The behavior of virtual humans in training simulations, also called non-player characters (NPCs), is not always viewed as realistic or believable by the trainee (for example, an NPC does not respond appropriately to danger obviously present in the environment, e.g., a soldier continuous his walk while being shot at). Unrealistic behavior of NPCs negatively influence the learning process and learning opportunities of a trainee. However, at the same time, it is not feasible, nor desirable, to develop a simulation in which every aspect of the behavior of a virtual character is realistic. It is therefore important that the specifications of an NPC refer to the types of behavior that support, rather than interfere with the learning goals of the trainee. This approach aims to support determining such behavior specifications.

Who should use this working approach?
The team tasked with determining the requirements of a training simulator, and with the development of training simulators.

When to use this working approach?
This approach is intended to be used based upon the findings obtained through a Training Needs Analysis, more specifically based upon the identified learning goals and training scenarios. The approach is to be used for specific training simulations; it is considered NOT feasible to use this for simulation development in a general sense; for situations where there are no known details about the training objectives of the simulator.

What result may be expected of the working approach?
The approach delivers support to the user(s) in determining the specifications of NPCs that will bring about those learning situations that contribute to achieving the learning objectives by the trainee(s) (i.e., NPCs that show *fit-for-purpose* behavior). Most importantly, the instructions and guidelines of the approach will support the users' awareness about the relationships between NPC behaviors, their significance for the learning situations in the simulation, and the desired responses of the trainee. In some cases, the approach may support the entire process of NPC behavior specification; in other cases additional refinement may be needed (e.g., the addition of domain-knowledge).

Disclaimer: not *all* aspects of NPC-behavior will be addressed in this approach. The approach supports the user to determine the NPC behavior requirements, but in the end it needs to be the user who makes the final decisions regarding the behavior specifications.

What is the difference between Believability and Realism?

Believability and realism are both requirements for the behavior of NPCs in order to create an effective learning environment *Realism* concerns whether the demonstrated behavior represents that of humans *in general*, under the given circumstances. *Believability* concerns whether the behavior is consistent with the personality, affect, beliefs, and goals of the particular virtual character.

What information is needed for the working approach?

The working approach uses a structured set of questions and assignments that call upon the following information (mostly obtained through an earlier conducted Training Needs Analysis):

Table 3 Information required for conducting the working approach.

Material	**Description**
Learning goals	Not just learning goals about task procedures, but also goals about knowledge and skills required to develop the competencies
Global simulation description	What type of simulation is necessary/possible?
Scenarios	Relative global description about the flow of the scenarios and what should be addressed.
Instruction possibilities	Measuring behavior, provide feedback, logging, after-action review input, log analysis, use of events, creating or adapting scenarios, availability own resources, steering scenario.
Operational knowledge	Context/ world knowledge – strategies, resources of the enemy, ROEs, typical behavior for a particular person, domain specific information.
Description of the user	Description of the trainee, type of trainee, tasks for the trainee, knowledge/skill level of the trainee. What is the gap needs to be bridged by the trainee?
Domain specific knowledge	Type of virtual entities (e.g. rank). Includes relevant knowledge regarding the domain which is not mentioned in the learning goals, but are necessary for the virtual entities in order to behave natural.

How should this working approach be used?

This approach makes uses of three files that should be used:

1 **Readme file:** the file you are currently reading.
2 **Interactive PowerPoint:** The PowerPoint for administering this approach requires the use of a computer mouse and arrow buttons.
3 **Worksheets:** Answers and solutions to questions and assignments should be documented in worksheets. Which (part of the) worksheet is to be used for a particular part of the approach is shown in the right upper corner in the PowerPoint. You can fill in these worksheets digitally in real-time, or you may also print an empty copy and fill this in by pen. After completing the approach, the worksheets contains a summary of all the options and decisions made regarding the required behavior of the NPC(s).

General tips for specifying realistic and believable behaviors

- Remember that full-scale reality is seldom necessary: players are inclined to accept lower levels of realism, as long as the NPC shows believable behavior. *The focus should be on determining the NPC-behaviors that fit the learning goals, and not to make every aspect of the NPC's behavior realistic.*
- In order for an NPC to behave in a fashion that is perceived as believable by humans (Livingstone, 2006), the NPC should demonstrate the capacity to:
  - Plan (for example, not unnecessary repeat an action),
  - Act (according to human-like reaction times and abilities), and
  - React (to actions and presence of others and to changes in the environment).

REPORT DOCUMENTATION PAGE

(MOD-NL)

1. DEFENCE REPORT NO (MOD-NL)	2. RECIPIENT'S ACCESSION NO	3. PERFORMING ORGANIZATION REPORT NO
-	-	TNO 2019 R11607
4. PROJECT/TASK/WORK UNIT NO	5. CONTRACT NO	6. REPORT DATE
060.31925		October 2019
7. NUMBER OF PAGES	8. NUMBER OF REFERENCES	9. TYPE OF REPORT AND DATES COVERED
48 (incl 1 appendix, excl RDP & distribution list)	87	Final

10. TITLE AND SUBTITLE

Virtual Characters that Work:
fit-for-purpose behavior of virtual entities in training simulations

11. AUTHOR(S)

Dr. K. van den Bosch, T.A.J. Schoonderwoerd MSc, R.A.M. Blankendaal MSc, Dr. M.M.M. Peeters

12. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

TNO, P.O. Box 23, 3769 ZG Soesterberg, The Netherlands

13. SPONSORING AGENCY NAME(S) AND ADDRESS(ES)

Ministerie van Defensie

14. SUPPLEMENTARY NOTES

The classification designation Ongerubriceerd is equivalent to Unclassified, Stg. Confidentieel is equivalent to Confidential and Stg. Geheim is equivalent to Secret.

15. ABSTRACT (MAXIMUM 200 WORDS (1044 BYTE))

Realizing the potential of simulation for training demands that virtual players (the Non Playing Characters, or NPCs) behave in a believable fashion and in conjunction with the training objectives. To demonstrate ‘fit-for-purpose realism’, an NPC should adhere to relevant protocols, norms and values of the situational context (“believability”), and it should elicit a response in the trainee that is considered appropriate when the situation would be encountered in the real world. To support teams tasked with the development of a training simulation, a working approach was developed to address the requirements of NPCs. It consists of three systematic series of assignments and questions. The first part supports the team with identifying the NPC behaviors that are essential for the learning goals of the trainee. The second part supports tuning the NPC behaviors to the situational context. The third part consists of a guided critical evaluation of the determined behavior specifications. The working approach is meant to be used as guidance, assisting a team with the consideration of options, enabling making well-founded decisions.

16. DESCRIPTORS	IDENTIFIERS
Simulation; Artificial Intelligence; learning; training; fidelity	

17a. SECURITY CLASSIFICATION (OF REPORT)	17b. SECURITY CLASSIFICATION (OF PAGE)	17c. SECURITY CLASSIFICATION (OF ABSTRACT)
TNO Public	TNO Public	TNO Public

18. DISTRIBUTION AVAILABILITY STATEMENT	17d. SECURITY CLASSIFICATION (OF TITLES)
Unlimited Distribution/Subject to approval MOD-NL	TNO Public

Distribution list

DEFENSIE

hardcopy	NLDA/Projectbureau K&I, Defensie Programma procesbegeleider drs C.M.P. van de Meulenhof
hardcopy	Defensie Programmabegeleider Majoor E.W. (Evert) de Boer, JIVC/KIXS
hardcopy	Defensie Projectbegeleider Maj. H. de Groot, CLAS/OTCO/LWC/AFD KIOT
pdf	PLANNEN/K&I, KTZ C.M. van den Berg
pdf	DMO/Joint IV Commando/C4I&I/InformatieBeheer/PDB

TNO

pdf	TNO Referent, Roadmap Directeur R. le Fèvre
pdf	TNO Programmaleider (PGL) P.J.M. (Philip) Kerbusch, MsC.
pdf	Research manager TNO PGL dr. E.W. Boot
pdf	TNO Medewerkers op aangeven van de TNO PGL Dr. K. van den Bosch T.A.J. Schoonderwoerd, MSc R.A.M. Blankendaal, MSc Dr. M.M.M. Peeters Dr. J.E. Korteling
hardcopy	TNO Archief (locatie Den Haag)