

An Intelligent System for Aggression De-Escalation Training

Tibor Bosse¹ and Charlotte Gerritsen² and Jeroen de Man¹

Abstract. Artificial Intelligence techniques are increasingly being used to develop smart training applications for professionals in various domains. This paper presents an intelligent training system that enables professionals in the public domain to practice their aggression de-escalation skills. The system is one of the main products of the STRESS project, an interdisciplinary research project involving partners from academia, industry and society. The system makes use of a variety of AI-related techniques, including simulation, virtual agents, sensor fusion, model-based analysis and adaptive support. A preliminary evaluation of the system has been conducted with two groups of potential end users, namely tram conductors and police academy students.

1 INTRODUCTION

“A train conductor was assaulted at Nuneaton railway station after asking to see a man’s ticket. The 48-year-old victim was working on a service from Crewe to London when he asked to see a passenger’s ticket before he boarded the train in Nuneaton. The man became abusive and started to push the conductor before leaving the station.” [28]

Although it is just one example, this incident illustrates the vulnerability of employees in the public sector to aggressive behaviour of people like customers, patients, travellers or other citizens. Aggressive behaviour against public service workers (e.g. police officers, ambulance personnel, public transport employees) is an ongoing concern in many countries [14, 21]. According to a national safety investigation in the Netherlands in 2011, almost 60% of the employees is confronted with unwanted behaviour on a daily basis [1]. Most incidents of aggression are of a verbal nature (e.g., insulting, swearing, intimidating), but in about 10% of the cases the conflicts escalate into physical aggression (e.g., threatening, abusing, robbing).

To better prepare them for these incidents, professionals in the public domain often receive dedicated resilience training. Such training is typically performed in a group setting based on role-play, where employees learn to communicate with aggressive clients in a de-escalating manner. Although this form of training has shown to be successful, it is quite expensive with respect to both money and time. Furthermore, the training is not always easy to control or repeat systematically.

As a complementary approach, the aim of the STRESS project [26] was to develop a simulation-based training system for aggression de-escalation. This is in line with a number of recent initiatives that show promising results regarding the possibility to

train social and communicative skills based on simulated environments involving virtual humans [2, 12, 13, 17, 18, 19, 22]. The main idea of the current system is that employees in the public domain can practice their aggression de-escalation skills by engaging in conversations with aggressive virtual characters. By designing the scenarios in such a way that the characters calm down if they are being approached correctly, but become more aggressive if they are being treated inappropriately, trainees will receive immediate feedback on their performance. Meanwhile, they are monitored by intelligent software that observes and analyses their behaviour and physiological state (e.g., heart rate, skin conductance, brain activity) and provides tailored feedback. Feedback consists of two categories, namely hints and prompts on the one hand, and run-time modifications in the scenarios on the other hand. By using such a system, employees have the ability to practice their aggression de-escalation skills in a cost-effective, personalised and systematic manner.

The purpose of this paper is to describe the overall architecture of the system, some details of the various components of which it consists, and the preliminary results of evaluation studies with potential end users. Many of these elements have been published in previous papers [4-11], but this is the first time in which all of them are combined into our coherent description. Hence, the main contribution of the paper is in explaining how the various parts of the STRESS project come together in a practical application.

The remainder of the article is structured as follows. In Section 2, some background knowledge is presented about aggressive behaviour and the prescribed techniques to de-escalate aggression. Section 3 describes the architecture of the overall training system, as well as some details about its separate components. Section 4 summarises two case studies that have been conducted to test the system: one in the domain of public transport, and one in the domain of law enforcement. Finally, Section 5 concludes the paper with a discussion.

2 AGGRESSION DE-ESCALATION

To design an effective training tool, a first question to be asked is what should be the learning goals of the system. For the current context, these learning goals are similar to the ones used in real world aggression de-escalation training, and are related to the development of *emotional intelligence*: employees should learn to recognize the emotional state of the (virtual) conversation partner, and choose the communication style that suits this emotional state.

More specifically, when it comes to aggressive behaviour, it is important that employees learn to recognize the nature of the aggression. Here, two main categories can be distinguished: aggression can be either *emotional* (or *reactive*) or *instrumental* (or *proactive*) [15]. In case of emotional aggression, the aggressive

¹ Department of Computer Science, Vrije Universiteit Amsterdam, The Netherlands, email: t.bosse@vu.nl, jeroendeman@gmail.com

² Netherlands Institute for the Study of Crime and Law Enforcement, Amsterdam, The Netherlands, email: cgerritsen@nscr.nl

behaviour typically is caused by an angry reaction to a negative event that frustrates a person's desires [3]. Such a person is likely to be angry with respect to whatever stopped him from achieving his goal. An example in the public transport domain is a traveller getting angry because the tram is late while he has to attend an important meeting. When dealing with an emotional aggressor, *supportive* behaviour from the de-escalator is required, for example by ignoring the conflict-seeking behaviour, calmly making contact with the aggressor, actively listening to what he has to say, showing empathy, and suggesting solutions to his problems.

In contrast, in case of instrumental aggression, the aggressive behaviour is only used 'instrumentally', to achieve a certain goal. Such behaviour is not a direct response to a negative event and is less strongly related to heavy emotions. A well-known example in the domain of public transport involves someone who wants to travel without paying for his ticket. This type of aggression often starts with an attempt to persuade the conversation partner, e.g. "Oh, I forgot my wallet, can I just come along for two stops?". Often, in case the employee does not give the aggressor what he wants, the aggressive behaviour will reveal itself through more threatening remarks like "I'll be back tomorrow with my friends". To de-escalate instrumental aggression, a *directive* response is assumed to be most effective. It is necessary to show the aggressor that there is a limit to how far he can pursue his aggressive behaviour, and to make him aware of its consequences [10].

To conclude, the presented training environment is centred around two main learning goals, namely 1) *recognizing* the type of aggression of the conversation partner (i.e., emotional or instrumental), and 2) selecting the appropriate *communication style* towards the conversation partner (i.e., supportive or directive). To assess the type of aggression, employees need to carefully observe the verbal and non-verbal behaviour of the aggressive virtual character. In general, reactive aggressors will show more arousal (e.g., flushed face, emotional speech) than proactive aggressors. Also, the context should be taken into account (e.g., someone who just finds out that he lost his ticket will be more emotional than someone who knew this all along, and just tries to intimidate the tram driver to ride for free).

3 TRAINING SYSTEM

The main aim of the STRESS project is to develop an intelligent training system that is able to analyse the trainee's behaviour during confrontations with aggressive individuals, and provide appropriate feedback, enabling trainees to improve their performance. During the training, users will be placed in a virtual scenario in a particular domain (e.g., selling tram tickets, or issuing parking tickets), which involves one or more virtual agents that at some point in time start behaving aggressively (e.g., insulting the tram driver because he is late). The user's task is to de-escalate the aggressive behaviour of the virtual agents by applying the appropriate communication techniques. An important asset of the system is that it can adapt various aspects of the training (e.g., scenarios, difficulty level, feedback) at runtime on the basis of its estimation of the trainee's physiological state and performance.

3.1 System Overview

Figure 1 depicts the global architecture of the system [6]. The rounded rectangles denote components of the system, and the

arrows denote information flows. The normal rectangles indicate clusters of components that have the same function (i.e. the *analysis* and *support* layer). In the training environment, the trainee will be engaged in a *virtual reality environment* shown on a computer screen (or possibly on a head-mounted display), while being monitored by an intelligent *training agent*. The virtual scenario is generated by a separate module within the agent, which contains knowledge about relevant scenarios in a particular domain. The trainee observes the events that happen in the scenario, and has to act in the scenario this by selecting the most appropriate action (currently this is simply implemented by means of a multiple choice menu). During training, the user is connected to (non-intrusive) devices that measure physiological states related to arousal and stress; in particular heart rate, skin conductance and electroencephalogram (EEG) signals. The data measured by these devices are then used by the agent as input for a computational model that integrates them at runtime, to assess the trainee's levels of stress and (negative) emotions (the *affective model*). This assessment of the trainee's affective state is combined with information about the status of the task (e.g., the actions performed by the trainee), and used by another computational model (the *decision making model*) to assess whether (and why) the trainee made certain mistakes. The outputs of both models are analyses of the trainee's emotional state (e.g., how much stress does (s)he experience?) and decision making behaviour (e.g., are any mistakes made?), respectively. This information is used for two purposes: by the *scenario development module*, to modify the running scenario (e.g., to repeat a certain scenario that is considered difficult), and by the *feedback determination module*, to provide feedback to the trainee (e.g., advice to change the conversation style).

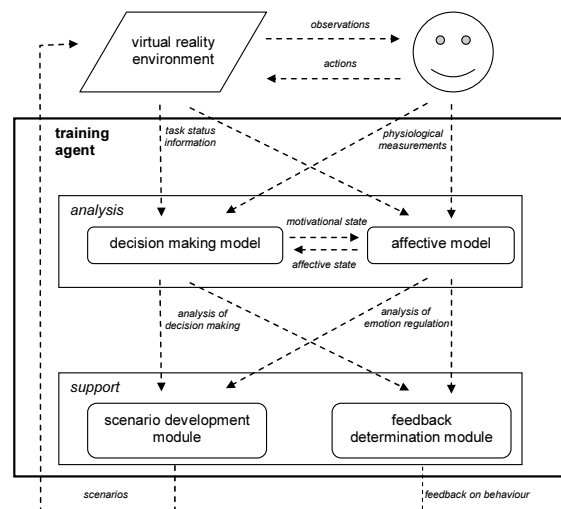


Figure 1. Global architecture of the training system

3.2 Virtual Reality Environment

The virtual reality environment is based on the InterACT software [30], developed by the company IC3D Media [25]. InterACT is a software platform that has been specifically designed for simulation-based training of interpersonal skills. Unlike most existing software, it focuses on smaller situations, with high realism and detailed interactions with virtual characters. True-to-life animations and photo-realistic characters are used to immerse the player in the game. An example screenshot of a training

scenario for the public transport domain is shown in Figure 2. In this example, the user plays the role of a tram conductor that has the task of calming down an aggressive virtual traveller.



Figure 2. Screenshot of a training scenario for tram conductors

To enable users to engage in a conversation with an embodied conversational agent (ECA), a dialogue system based on conversation trees is used. The system assumes that a dialogue consists of a sequence of spoken sentences that follow a turn-taking protocol. That is, first the ECA says something (e.g. “I forgot my public transport card. You probably don’t mind if I ride for free?”). After that, the user can respond, followed by a response from the ECA, and so on. In InterACT, these dialogues are represented by conversation trees, where vertices are either atomic ECA behaviours or decision nodes (enabling the user to determine a response), and the edges are transitions between nodes.

The atomic ECA behaviours consist of pre-generated fragments of speech, synchronised with facial expressions and possibly extended with gestures. Scenario developers can generate their own fragments using a motion sensing input device such as the Microsoft Kinect camera and a commercial software package FaceShift [29]. As the recorded fragments are independent from a particular avatar, they can be projected on arbitrary characters.

Each decision node is implemented as a multiple choice menu. Via such a menu, the user has the ability to choose between multiple sentences. Hence, the emphasis of the current system is on the verbal aspects of aggression de-escalation. In most of the scenarios, three options are available within every decision node. Typically, these options have been created in such a way that one of them is clearly *supportive*, another one is clearly *directive*, and the third option is neutral. Here, the supportive and directive option relate to the communication styles explained in Section 2. Figure 2 illustrates how these three options can be instantiated in terms of concrete sentences (where A=neutral, B=directive, C=supportive). The user’s choice determines how the scenario continues, by triggering a corresponding branch in the tree.

Although this approach works well, there is a risk that the behaviour of the ECAs becomes predictable in the long term. For example, in the situation shown in Figure 2, choosing option B (the ‘directive’ option) will always result in the ECA becoming irritated, no matter how often the scenario is played, or what has happened before. To overcome this problem, in [11] an approach was put forward to endow the agent with *internal states* that are either set beforehand (e.g. whether the agent is a reactive or a proactive aggressor) or are the result of earlier interactions (e.g. a state of anger that gradually increases during the scenario). This means that the agent is equipped with a cognitive model of aggression, replacing the direct connections between user choices and ECA

responses. As shown in [11], the resulting conversations indeed provided more variation, and were perceived as less predictable.

3.3 Physiological Measurements

Simulation-based training can only be effective if the virtual scenarios trigger emotional responses that are comparable to the reactions people show to the same stimuli in real world scenarios. To investigate to what extent this is the case, in [4] an experiment was performed in which the impact of an aggressive virtual agent was compared with that of an aggressive human. By randomly distributing a group of 28 participants over two conditions (virtual and human) and measuring their physiological and subjective emotional state before and after an aggressive outburst of their conversation partner, the difference between virtual and human aggression was studied. The ‘outburst’ was realized by having the (virtual or human) conversation partner suddenly get extremely angry towards the participant, while shouting and accusing him or her of not paying attention. The results showed that both conditions induced a substantial stress response, but that the impact of the human aggression was stronger than that of the virtual aggression.

Part of these results are illustrated by Figure 3. This figure depicts the dynamics of skin conductance (also called electrodermal activity, EDA, one of the most common indicators for arousal) over time in microSiemens during the relevant part of the experiment, averaged over all participants in each condition. The horizontal axis denotes a period of 2 minutes, i.e., 1 minute before the start of the aggressive outburst and 1 minute after it. The vertical line indicates the moment the outburst started. More details about the study can be found in [4].

Although these results show that there is still a gap between the intensity of the stress response caused by a virtual character and by a real human, they are nevertheless promising, because they indicate that virtual characters can at least elicit responses that are comparable to the ones triggered in reality.




Figure 3. Dynamics of electrodermal activity over time

Inspired by these results, an interface was developed to connect the training system to two devices that measure physiological states related to arousal and stress. In particular, the Plux wireless biosensors toolkit [27] is used to measure heart rate and skin conductance, and the Myndplay Brainband [31] is used to measure EEG signals. Additionally, within InterACT a visualisation window was implemented in which the measurements of these three sensors can be displayed at run-time during training (see upper right corner of Figure 2). From top to bottom, this window shows a user friendly interpretation of the trainee’s current heart rate (in beats per minute), skin conductance (in microSiemens) and EEG signals (in terms of ‘meditation value’, one of the outputs of the Myndplay Brainband that is correlated with a state of

relaxation). This allows trainees to receive instant (bio)feedback on their physiological state while training, thereby helping them to stay calm during aggressive confrontations.

3.4 Analysis Layer

As explained in Section 3.1, the purpose of the analysis layer is to process data about the user (in particular: physiological data and task status information from the virtual environment), in order to draw high-level conclusions about the user's state. It consists of two sub-models, the affective model and the decision making model, which are described in the following sub-sections. Both models have been formalised using the LEADSTO language [9], which enables modellers to describe mental processes in terms of transitions between states that are expressed in terms of logical and/or numerical variables.

3.4.1 Affective Model

The affective model was inspired by the work of Gross [16], and is explained in detail in [6]. A high-level overview is shown in Fig. 4. The circles represent different states, which are all formalised in a numerical manner, in terms of a variable with a real value between 0 and 1. In an actual application, real world data should be mapped to values in this interval. For instance, a very threatening stimulus in the training system (e.g., an aggressive virtual character) could be represented as a *world state* with value 0.9. Similarly, a moderately intensive feeling of fear (e.g., as measured by the physiological devices) could be represented as a *feeling* with value 0.5.

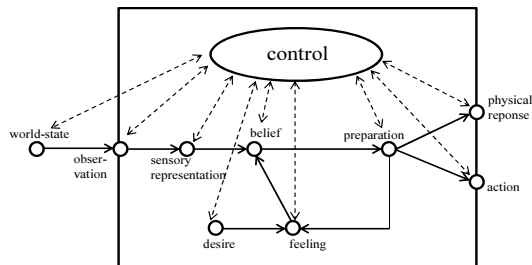


Figure 4. Overview of affective model

Arrows in Figure 4 denote the influence of one state on another state. The model that represents *generation* of emotions is depicted by using solid arrows. Additionally, *regulation* of emotions is represented by the *control* state. Each regular state has a positive effect on the control state (representing a monitoring process), but can in turn be suppressed by the control state (representing a regulation process), as indicated by the dashed arrows.

3.4.2 Decision Making Model

The decision making model is explained in detail in [10], and is shown graphically in Figure 5. The circles on the left denote observations made by the user, the circles on the right (communicative) actions, and the remaining circles internal states.

Roughly, the dynamics of the model can be split into three sub-processes. First, as shown in the lower part of the figure, the emotional state of the user is updated based on the observed (verbal and non-verbal) behaviour of the conversation partner, and has in turn an impact on his or her own non-verbal behaviour. Next, as shown in the upper left part of Figure 5, there is a sub-process

related to the evaluation of (both the nature and the intensity of) the conversation partner's emotional state. More specifically, this boils down to deciding whether we are dealing with reactive or proactive aggression (or no aggression). Finally, as shown in the upper right part of the figure, the evaluation of the conversation partner's emotional state serves as input for a decision about which 'de-escalation approach' to select. For this, domain-specific knowledge is used about which approach works best in which situation (as explained in Section 2).




Figure 5. Overview of decision making model

To conclude, the affective model and decision making model enable the system to draw conclusions about the user's emotional state and decisions (and errors) made, respectively. Preliminary evaluations of the models are discussed in [6] and [10]. In the next section we explain how the output of the models can be used to provide dedicated support to increase training effectiveness.

3.5 Support Layer

Like the analysis layer, the support layer also consists of two sub-models, namely the scenario development module and the feedback determination module. Both modules are described below.

3.5.1 Scenario Development Module

The main purpose of the scenario development module is to generate interesting training scenarios that fit to the learning goals of the trainee. In particular, the concept of adaptive training (or scaffolding) is used, where the difficulty level of the scenarios adapts to the performance of the trainee. To this end, based on the learning goals identified in Section 2, a score was introduced to keep track of how well the goals were achieved. This score was calculated based on the output of the analysis layer. Next, a number of difficulty levels were established, as well as a mechanism to navigate up and down between these levels based on the user's score. This mechanism is visualised in Figure 6. As can be seen, a separate score is maintained for cases of emotional (or reactive) aggression as well as for cases of functional (or pro-active) aggression. The main idea is that the user's score for a particular type of aggression needs to be sufficiently high to reach a higher level. However, after losing two points, the user falls back a level. In the first part of the training (level 1-3), the type of aggression is already given. Instead, in the second part of the training (level 4-6) the trainee needs to identify the type of aggression by him- or herself. Levels are traversed per aggression type separately, with one exception: after the first part of the training (i.e., level 1-3), the trainee needs to have sufficient knowledge of both types of aggression before (s)he can continue.

Details of this module are presented in [8]. An initial evaluation reported in that paper demonstrated that the system successfully

adapted its difficulty level to the user’s performance, and that users were positive about the effect of this adaptation mechanism.

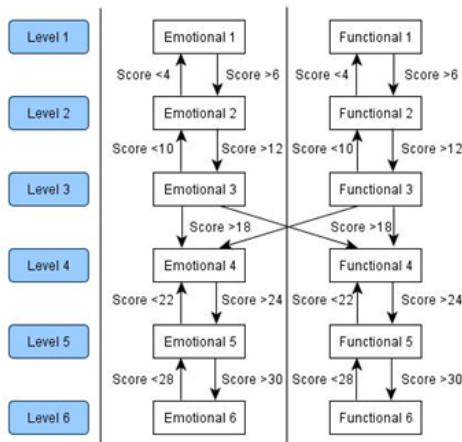


Figure 6. Transitions between difficulty levels

3.5.2 Feedback Determination Module

The feedback determination module is described in detail in [11]. This model uses the output of the analysis layer to generate appropriate feedback on the user’s performance in terms of after-session hints. Essentially it checks whether the situation was successfully de-escalated or not, and in the latter case, it analyses what the cause of this unsuccessful de-escalation was. In this analysis, several types of mistakes are distinguished such as 1) the user failed to judge the type of aggression correctly (i.e. reactive or proactive), 2) the user failed to apply the appropriate communication style (supportive or directive), and 3) the user failed to control his or her own emotional state. The decision tree that is used by the module is shown in Figure 7.

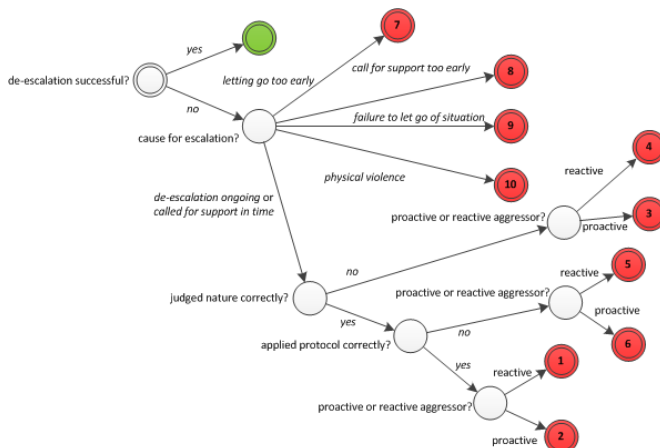


Figure 7. Overview of decision making model

Here, a green end state indicates successful de-escalation, whereas a red end state indicates unsuccessful de-escalation. Based on the specific end state a corresponding feedback message is generated, represented by the numbers in the figure. As an example, in case a scenario is classified as category (6), the following feedback is presented:

6. *User applies wrong approach towards a proactive aggressor.*
 ”You correctly judged the nature of the aggression, but you used the wrong verbal approach. A proactive aggressor should always be approached in a directive manner.

To test the module, a specific scenario has been worked out in the context of a man who has no cash money to pay for a tram ticket. A group of users have extensively played the scenario under varying the parameter settings. These user tests pointed out that the module indeed offered the desired support at the appropriate times (see [11] for more details).

4 CASE STUDIES

As explained in the previous section, most components of the system were tested separately based on user studies in a laboratory setting. As a follow up, additional evaluation studies have been performed with potential end users from two domains of interest: public transport and law enforcement. In these experiments, the focus was on testing these users’ experience with respect to the virtual reality environment (and the underlying dialogue system); hence, the training agent was disabled during these tests. The two experiments are briefly described in the following sub-sections.

4.1 Public Transport

For this evaluation study (see [7]), a number of scenarios have been developed, in collaboration with (and approved by) domain experts of GVB, the public transport company in Amsterdam. All scenarios are perceived from the perspective of a tram conductor. In total, they address 9 different situations in which a conflict may arise, such as ‘traveler is not allowed to take hot coffee on board’ and ‘tram arrives 10 minutes late’. Moreover, for each scenario, different variants have been written: some in which the virtual character shows emotional aggression, and some in which it shows instrumental aggression. In total, there were 40 scenarios.

Twenty-four people participated in the experiment (13 male and 11 female), all of which were employees of GVB. Their average age was 45,4 ($\sigma = 12.0$). The experiment was executed in a computer room at GVB. Participants had to play all 40 scenarios during 4 different sessions distributed over 4 weeks (i.e., 10 scenarios per session). After the last session, the participants were asked to fill out a usability questionnaire, consisting of 20 statements about which the participants had to express their opinion on a 5-point Likert scale. The questionnaire was inspired by Witmer and Singer [24], and included statements about issues such as user experience, presence, and perceived effectiveness. In the end, the statements were grouped into 4 categories, namely *content*, *interaction*, *engagement*, and *effect*, to obtain an average score on these aspects. The *content* category contained statements about the perceived realism of the scenarios and the characters (e.g., ‘the scenarios were representative for real world situations’). The *interaction* category contained statements about how natural it was to interact with the characters (e.g., ‘I felt that my answers had an influence in the behaviour of the virtual characters’). The *emotional* category addressed the perceived sense of presence of the participants (e.g., ‘during training I felt engaged in the scenarios’). Finally, the *effect* category contained statements asking the participants for their opinion about the effectiveness of the training (e.g., ‘I think this type of training is a useful addition to real world training’).

The aggregated results are summarised in Figure 8 (on a scale from -2 to 2). As shown there, participants were generally positive about the content of the scenarios, the interaction possibilities of the system, and (in particular) about its perceived learning effect,

but were less enthusiastic about the system's emotional impact. More details can be found in [7].




Figure 8. User experience results for the Public Transport case

4.2 Law Enforcement

To investigate how stable these results are across different application domains, a second evaluation study was conducted (see [5]), in the domain of law enforcement. This experiment was executed in collaboration with the Dutch Police Academy, and focused on the module 'Noodhulp' (Emergency Assistance), which is part of their education program. As part of this module, students have to learn to correctly handle the so-called 'Door Scene'. This is a situation in which a police officer has just been informed about an incoming emergency call. For the current experiment, we focused on the domain of domestic violence (e.g., a call from a crying woman who claims that her boyfriend is abusing her). The scenario starts at the moment that the police officer (together with his or her partner) arrives at the address in question, and rings at the door.

The setup of this experiment was very similar to the previous one, with some minor differences: instead of 40 scenarios, only 4 (slightly longer) scenarios were used. Moreover, these scenarios were not played during 4 different sessions, but during one single session. The experiment was executed in a computer room at the Police Academy. Also the demographics of the participants were different: in total, 41 Police Academy students participated in the experiment (31 male and 10 female), and their average age was 27.1 ($\sigma = 6.5$). The questionnaire consisted of the same 20 statements as the previous one; however, instead of a 5-point Likert scale, a 7-point scale was used.

The aggregated results of this study are summarised in Figure 9 (on a scale from -3 to 3). As shown, the outcomes are similar to those for the public transport domain, with the main difference that the participants were more positive about the system's interaction possibilities, but believed a bit less in the effectiveness of the tool. More details of this study can be found in [5].




Figure 9. User experience results for the Law Enforcement case

5 DISCUSSION

The current paper introduced a prototype of an intelligent simulation-based training system that enables professionals in the public domain to practice their verbal aggression de-escalation skills during face-to-face conversations. The system has been designed in a modular fashion, and integrates various AI-related techniques, including simulation, virtual agents, sensor fusion, model-based analysis and adaptive support. The various modules have been tested separately in a lab setting. Additionally, case studies in two real world domains have been used to obtain feedback on the virtual reality environment from potential end users.

The presented system has similarities with several recent approaches to train social skills through conversations with virtual humans. These projects have addressed a variety of tasks in different domains, including job interviews [2], police interviews [13], leadership training for naval officers [17], medical consultations [18], negotiation exercises [12, 19], and manager-employee conversations [22]. The current system differs from these systems in the sense that it focuses on a domain (aggression de-escalation) in which the stimuli from the virtual environment are mainly negative. As a result, more effort was put into creating visually and behaviourally believable characters, and into measuring the users' physiological response to the behaviour of these characters. Additionally, the system was enhanced with an intelligent 'training agent' that gives adaptive personalised feedback based on the user's state and behaviour (cf. [20, 23]).

The results from the case studies indicate that with respect to user satisfaction, participants were generally positive about the content of the virtual scenarios, the mechanisms to interact with the characters, and the potential of the system as a learning tool. Nevertheless, also a number of points for improvement were identified, which mainly have to do with the emotional aspect of the system: for several participants, the perceived sense of presence was limited because they did not 'feel' the emotion in the virtual conversation partner. One interesting way to improve this situation is to combine the scenarios with haptic feedback. Another possible extension would be to use a head-mounted display instead a flat video screen. Our expectation is that such extensions will lead to a more engaging, and therefore more effective training tool.

ACKNOWLEDGEMENTS

This research was supported by funding from the National Initiative Brain and Cognition, coordinated by the Netherlands Organisation for Scientific Research (NWO), under grant agreement No. 056-25-013.

REFERENCES

- [1] M. Abraham, S. Flight, and W. Roorda. Aggressie en geweld tegen werknemers met een publieke taak. Research for the program 'Veilige Publieke Taak 2007 - 2009 - 2011'. Amsterdam: DSP.
- [2] A. Ben Youssef, M. Chollet, H. Jones, N. Sabouret, C. Pelachaud, and M. Ochs. Towards a Socially Adaptive Virtual Agent. In: Intelligent Virtual Agents, Springer LNCS, vol. 9238, pp. 3-16, 2015.
- [3] L. Berkowitz. Whatever Happened to the Frustration-Aggression Hypothesis? *American Behavioral Scientist*, 21, 691-708, 1978.
- [4] R.A.M. Blankendaal, T. Bosse, C. Gerritsen, T. de Jong, and J. de Man. Are Aggressive Agents as Scary as Aggressive Humans? In: Proceedings of AAMAS'15. ACM Press, 2015, pp. 553-561.

- [5] T. Bosse, and C. Gerritsen. Towards Serious Gaming for Communication Training - A Pilot Study with Police Academy Students. In: Proceedings of Intetain 2016. Springer Verlag, 2016.
- [6] T. Bosse, C. Gerritsen, and J. de Man. Agent-Based Simulation as a Tool for the Design of a Virtual Training Environment. In: Proceedings of the 14th International Conference on Intelligent Agent Technology, IAT'14. IEEE Computer Society Press, pp. 40-47, 2014.
- [7] T. Bosse, C. Gerritsen, and J. de Man. Evaluation of a Virtual Training Environment for Aggression De-escalation. In: Proceedings of Game-On 2015. Eurosis, 2015.
- [8] T. Bosse, C. Gerritsen, J. de Man, and S. Tolmeijer. Adaptive Training for Aggression De-escalation. In: C.J. Headleand et al. (Eds.), Proc. of ALIA'14. Springer CCIS, vol. 519, 2015, pp. 80-93.
- [9] T. Bosse, C.M. Jonker, L. van der Meij, and J. Treur. A Language and Environment for Analysis of Dynamics by Simulation. Int. J. of AI Tools, volume 16, issue 3, 2007, pp. 435-464, 2007.
- [10] T. Bosse, and S. Provoost. Towards Aggression De-escalation Training with Virtual Agents: A Computational Model. In: Proceedings of HCI'14. Springer Verlag, pp. 375-387, 2014.
- [11] T. Bosse, and S. Provoost. Integrating Conversation Trees and Cognitive Models within an ECA for Aggression De-escalation Training. In: Proceedings of PRIMA'15. Springer Verlag, Lecture Notes in Artificial Intelligence, 2015.
- [12] J. Broekens, M. Harbers, W.P. Brinkman, C.M. Jonker, K. van den Bosch, and J.J. Meyer. Virtual reality negotiation training increases negotiation knowledge and skill. In: Intelligent Virtual Agents, Springer LNCS, vol. 7502, pp. 218-230, 2012.
- [13] M. Bruijnes, J.M. Linssen, H.J.A. op den Akker, M. Theune, S. Wapperom, C. Broekema, and D.K.J. Heylen. Social Behaviour in Police Interviews: Relating Data to Theories. In: Conflict and Multimodal Communication, Springer Verlag, pp. 317-347, 2015.
- [14] J. DelBel. De-escalating workplace aggression. Nursing Management (USA) 34, 31-34, 2003.
- [15] K.A. Dodge. The structure and function of reactive and proactive aggression. In D. Pepler and H. Rubin, (eds.), The development and treatment of childhood aggression (pp. 201-218). Hillsdale, NJ: Erlbaum, 1990.
- [16] J.J. Gross. Emotion regulation in adulthood: timing is everything. Current directions in psychological science, vol. 10, issue 6, pp. 214-219, 2001.
- [17] M. Hays, J. Campbell, M. Trimmer, J. Poore, A. Webb, C. Stark, and T. King. "Can Role-Play with Virtual Humans Teach Interpersonal Skills?". In Interservice/Industry Training, Simulation and Education Conference (IITSEC), 2012.
- [18] J. Jeurig, F. Grosfeld, B. Heeren, M. Hulsbergen, R. Ijntema, V. Jonker, N. Mastenbroek, M. van der Smagt, F. Wijmans, M. Wolters, and H. van Zeijts. Communicate! – A Serious Game for Communication Skills. In: G. Conole et al. (Eds.): EC-TEL 2015, Springer LNCS, vol. 9307, pp. 513-517, 2015.
- [19] J. Kim, R.W. Hill, P. Durlach, H.C. Lane, E. Forbell, C. Core, S. Marsella, D. Pynadeth, and J. Hart. BiLAT: A game-based environment for practicing negotiation in a cultural context. Int. J. of AI in Education, vol. 19, issue 3, pp. 289-308, 2009.
- [20] M. Peeters. Personalized Educational Games. Developing Agent-Supported Scenario-Based Training. PhD Thesis, Utrecht University, 2014.
- [21] M.J. Smith and R.V. Clarke. Crime and public transport. In R. Clarke (Ed.), Crime and Justice: A Review of Research (pp. 169-234). Chicago, IL: The University of Chicago Press, 2000.
- [22] F. Vaassen and J. Wauters. deLearyous: Training interpersonal communication skills using unconstrained text input. In: Proceedings of ECGBL. pp. 505-513, 2012.
- [23] J. Westra, H. van Hasselt, F. Dignum, and V. Dignum. Adaptive Serious Games Using Agent Organizations. In: F. Dignum et al. (eds.), Agents for Games and Simulations, Springer LNCS, vol. 5920, pp. 206-220, 2009.
- [24] B.G. Witmer. and M.J. Singer. Measuring presence in virtual environments: a presence questionnaire. Presence: Teleoperators and Virtual Environments 7, pp. 225-240, 1998.
- [25] <http://ic3dmedia.com/>.
- [26] <http://stress.few.vu.nl/>.
- [27] <http://www.biosignalsplux.com/>.
- [28] <http://www.coventrytelegraph.net/news/local-news/conductor-assaulted-platform-railway-station-10757452>.
- [29] <http://www.faceshift.com/>.
- [30] <http://www.interact-training.nl/>.
- [31] <http://www.myndplay.com/>.