

Adaptive training for aggression de-escalation

The ability to de-escalate confrontations with aggressive individuals is a useful skill, in particular within professions in public domains. Nevertheless, offering appropriate training that enables students to develop such skills is a nontrivial matter. As a complementary approach to real-world training, the STRESS project proposes a simulation-based environment for training of aggression de-escalation. The main focus of the current paper is to make this system adaptive to the performance of the trainee. To realize this, first a number of learning goals have been identified. Based on these, several levels of difficulty were established, as well as a mechanism to switch up and down between these levels based on the user's score. A preliminary evaluation demonstrated that the system successfully adapts its difficulty level to the performance of the user, and that users are generally positive about the adaptation mechanism.

This chapter has been published as:

Bosse, T., Gerritsen, C., Man, J. de & Tolmeijer, S. (2015). Adaptive Training for Aggression De-escalation. In *Proceedings of the Artificial Life and Intelligent Agents Symposium* (pp. 80-93). Springer Verlag.*

* Authors ordered alphabetically and contributed equally to this work.

10.1 Introduction

People working in the public sector (e.g. police officers, ambulance personnel, public transport employees) are often confronted with aggressive behavior. According to a recent study, around 60% of the employees in the public sector in the Netherlands have been confronted with such behavior in the last 12 months (Abraham et al., 2007). Being confronted with (verbal) aggression can have severe consequences and is closely associated with psychological distress, which in turn can have a negative impact on work performance (James et al., 2006). Responses to aggression range from emotions like anger and humiliation through intent to leave the profession, and verbal aggression by customers may even impair employees' recognition and working memory (Rafaeli et al., 2012). In case of extreme incidents, employees may even develop symptoms indicating post-traumatic stress syndrome (Bonner & McLaughlin, 2007).

To deal with aggression, a variety of techniques are available that may prevent escalation (Anderson & Clarke, 1996; 'BZK', 2008). These include communication skills (both verbal and non-verbal), conflict resolution strategies, and emotion regulation techniques. The current paper is part of a project (called STRESS)¹ that aims to develop a serious game (Ritterfeld et al., 2009) for aggression de-escalation training, based on Virtual Reality. VR-based training has proven to be a cost-effective alternative for real world training in a variety of domains, including military missions (Hulst et al., 2008), surgery (Graafland et al., 2012) and negotiation (Kim et al., 2009).

The core of the STRESS project is the development of an intelligent training system that is able to analyze the behavior of human trainees while they interact with aggressive virtual agents. Users of the system will be placed in front of a 3D Virtual Reality (VR) environment (see Figure 4.1, p.71) that is either projected on a computer screen or on a head-mounted display. During the training, users will be placed in a virtual scenario in a particular domain (e.g., issuing parking tickets, or selling tram 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 behavior by applying the appropriate communication techniques. Users will be able to communicate with the agents via multiple modalities (e.g., text, speech, facial expression). Meanwhile, they will be monitored by intelligent software that observes and analyzes the behavior and physiological state (e.g., heart rate, skin conductance, brain activity) of the trainee and provides tailored feedback (Christoph, 2006; Heuvelink & Mioch, 2008).

Feedback will consist of two categories, namely hints and prompts on the one hand, and run-time modification in the scenarios on the other hand. An

¹<http://stress.few.vu.nl> (accessed 28-10-2014).

example of the former would be to inform the trainee that (s)he should use a more empathic communication style, whereas an example of the latter would be to decrease the difficulty level in case the trainee makes many mistakes. In order to offer this feedback in a personalized manner, it is very important that the system *adapts* it to the needs of individual users. Adapting a task to the behavior of the trainee is a well-known training paradigm in a variety of domains. Hence, the current paper presents a mechanism for adaptive training of aggression de-escalation. The emphasis is on adapting the difficulty level of the scenarios offered to the performance of the trainee.

The remainder of this paper is organized as follows. In Section 10.2, the existing literature on aggression is discussed, as well as the prescribed approaches to de-escalate aggression. In Section 10.3, the state-of-the-art on adaptive training is reviewed. Next, a conceptual model for adaptive training and its implementation are presented in Sections 10.4 and 10.5, followed by some preliminary results in Section 10.6. In Section 10.7, the paper is concluded with a discussion.

10.2 Aggression de-escalation

Within psychological literature, a distinction is made between two important theories regarding the nature of aggression: aggression can be either *functional* (or *proactive*) or *emotional* (or *reactive*). One of the key differences between these two types is the absence or presence of anger (Miller & Lynam, 2006).

When the aggression is of a functional nature, the aggressive behavior is not a response to some negative event, but is used instrumentally to achieve a goal. The *social learning theory* states that aggressive behavior can be learned through positive reinforcement (Bandura & Walters, 1963). The essence of this theory is that if a person has used aggression to achieve a goal in the past, and if this behavior was successful, then by operant conditioning (s)he will be likely to follow the same behavioral pattern in the future.

In contrast with functional aggression, aggression can also have an emotional nature, meaning that it is an angry reaction to a negative event that frustrates a person's desires. The *frustration-aggression hypothesis* (Berkowitz, 1978) tells us that aggression flows forth from a person's goals being frustrated. Such a person is likely to be angry with respect to whatever stopped him from achieving his goal. By a carry-over effect, the anger can be transferred to new situations as well (Angie et al., 2011).

To de-escalate aggressive behavior, it is important that public service workers understand the specific type of aggression they are dealing with, as each type of aggression requires a different approach. In particular, in situations when dealing with a functional aggressor, a *directive* type of intervention is

assumed to be most effective, focusing on an alteration of the contingencies associated with the aggression. In this case it is necessary to show the aggressor that there is a limit to how far he can pursue his aggressive behavior, and making him aware of the consequences of this behavior. Instead, when dealing with an emotional aggressor, more *supportive* behavior from the de-escalator is required, for example by ignoring the conflict-seeking behavior, making contact with the aggressor and actively listening to what he has to say. According to Hubbard et al. (2002), such interventions should focus on reducing hostile attribution biases, i.e., the tendency to perceive others as threatening.

This distinction between functional and emotional aggression, as well as the associated de-escalation techniques (i.e., directive vs. supportive approaches), are some of the key assumptions underlying the training system developed in the STRESS project. The following sections will present a mechanism to train the relevant communication skills in an adaptive manner.

10.3 Adaptive training

Adapting the difficulty level of a task to the performance of the player has been a well-known paradigm in serious gaming (and learning in general) for many years. The main underlying idea is that players' learning experience is related to their level of motivation. In general, if an activity is more engaging, interesting and engrossing, motivation will be higher (Graham & Weiner, 1996). However, this does not always imply that the difficulty level of a certain task should be maximal. Instead, especially when it comes to digital games, there are also situations in which motivation can be increased by lowering the difficulty level. Hoogen et al. (2008) describe the relation between a difficulty level and a player's mental state as follows:

‘Through striking the balance between a person's skills and the challenges an activity offers, that person may arrive in a psychological state known as flow. [...] Flow may gradually increase over the course of the game in a homeostatic positive feedback loop, until either the challenge becomes too great (resulting in frustration) or the player's skill outpaces the challenges the game can offer (leading to boredom).’

This suggests that there exists something like an optimal level of difficulty (or challenge) that yields maximal learning experience, which is often used as an argument to develop flexible training games that adapt dynamically to the player's behavior.

Indeed, the recent literature shows a number of examples of such adaptive training systems, which in one way or another tune their internal parameters

to the user's state or behavior. For example, Holmes et al. (2009) have demonstrated that adaptive training may be used to overcome learning difficulties for people with impairments in working memory. Wickens et al. (2013) have shown 'increasing difficulty' to be a successful technique in knowledge transfer when implemented adaptively (but not when increased in fixed steps). Also, several authors focus on increasing players' affective experience by adapting the emotional content of a game (see e.g. Nogueira et al., 2013). Finally, Yannakakis et al. (2010) argue that adaptive serious gaming is an effective method for training of conflict resolution skills, focusing on children as their user group rather. Implementing adaptivity can be achieved using various approaches. For example, Westra et al. (2009) propose a multi-agent system where each agent controls a particular aspect of a serious game and proposes changes to a centralized adaptation system, while the work by Peeters et al. (2014) is concerned with developing an ontology for automated generation of training scenarios. A more generic overview of design principles to develop effective adaptive training systems is provided in Raybourn (2007).

10.4 Conceptual model

In this section, the proposed model for adaptive training is described at a conceptual level. First, a number of relevant learning goals are formulated. After that, the structure of the model is described, based on the notion of dialog system. Finally, the idea of difficulty levels is introduced, as well as a mechanism to switch between them.

10.4.1 Learning goals

As mentioned in the introduction, the main learning goal of the proposed system is to be able to de-escalate confrontations with (verbally) aggressive individuals, in order to prevent these individuals from becoming physically aggressive. Based on discussions with domain experts, the following sub-goals have been identified:

- *Recognizing the type of aggression:* are we dealing with a person that is showing emotional or functional aggression? To assess this, trainees need to observe the verbal as well as the non-verbal behavior of the aggressive individual. In general, emotionally aggressive people will show more arousal (e.g., flushed face, emotional speech) than functionally aggressive people. 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).

- *Selecting the appropriate response*: based on the type of aggression observed, the trainee needs to either show some empathy (in case of emotional aggression) or act more dominantly (in case of functional aggression). It is crucial for the training that these responses are not swapped; in other words, showing empathy in case of functional aggression, or acting dominantly in case of emotional aggression is undesired.
- *Being able to make decisions under time pressure*: after some practice, the trainee should be able to perform the tasks mentioned above within limited time, and without much cognitive effort.

10.4.2 Structure

To train users to acquire the above skills, the STRESS project works with a dialog system where users (playing the role of the public service worker) engage in a conversation with a virtual agent (playing the role of a difficult customer). Conversations are represented as simple decision trees where user and virtual agent exchange sentences according to a turn-taking protocol. In this paper, the conversations used are text-based only and consist of a short introduction, the latest response of the customer and a multiple-choice list of possible answers.² In the remainder of this paper, we will refer to this as a *question*, to which the trainee has to choose the most appropriate response considering the situation. An example of a question in the context of aggression de-escalation training for tram drivers is the following, taken from Bosse & Gerritsen (2014):

A passenger enters the tram and wants to check in, but the balance on his public transport card turns out to be insufficient. You tell him that his balance is too low to check in. The passenger reacts with surprise, and says: *No, are you kidding me?! Really?! O my god, something should have went wrong with those damn machines of yours! Can't you for once just take me with you? I am in an extreme hurry! What do you respond?*

- a. I'm sorry sir, I feel really bad for you. But don't worry, you can just buy a ticket from me. Or if you prefer, you quickly run to the machine; over there you can recharge your card in a second!
- b. Too bad sir, our policy states that we only take paying customers on board. There is no other option for you than to buy a ticket or leave the tram.
- c. Sir, that's the way it is, these are the rules. You will have to but a new ticket.

²Nevertheless, the project as a whole also explores other interaction modalities, such as speech, facial expressions and gestures.

Note that this example addresses a case of emotional aggression, where the customer shows clear signs of emotional arousal (most notably swearing), probably caused by the unexpected message that his balance is too low. In such cases, the most appropriate way to respond would be by showing empathy and offering potential solutions for the problem: this is represented by answer a). The other two answers, b) and c), have an increasingly dominant (and less empathic) tone, which makes them the less appropriate responses in this type of situation.

For each question, a database with potential answers is established (although during training only a few of them (e.g., three) are offered in the multiple-choice menu). To enable the system to assess automatically which answers should be considered as appropriate, for each question all answers are divided into three categories, namely *exemplary*, *acceptable* and *unacceptable*. For example, for the above question, answer a) is exemplary, answer b) is acceptable, and answer c) is unacceptable.

In addition, for each question, a variant is defined with the *other* type of aggression (e.g., instrumental instead of emotional). For instance, such a variant for the above example would involve a passenger that is not emotional at all, but simply wants to use intimidation as an instrument to get a free ride. Also for these other variants, exemplary, acceptable and unacceptable answers are included in the database.³ This allows the system to add extra difficulty by mixing up answers for both types of aggression. The types of answers that are included in the multiple-choice menu depend on the level the trainee is in, as explained in the next section.

10.4.3 Difficulty levels

In order to make the dialog system adapt to the performance of the trainee, different levels of difficulty need to be distinguished. This way, the trainee can climb in levels when (s)he is performing well, and decline when many errors are made. For the proposed system, six levels of difficulty are used (see Table 10.1 for an overview):

1. Here, the type of aggression that is applicable to the current question (i.e., emotional or functional) is already revealed to the user, so all (s)he needs to do is to decide upon the appropriate response. Three potential answers are offered: one answer that is exemplary and two answers that are clearly wrong (the ‘unacceptable’ answers).

³Note that in some cases, answers that are unacceptable for one type of aggression may be exemplary for the other type. However, this is not necessarily always the case.

2. Similar to level 1, but the applicable type of aggression is not revealed anymore.
3. Similar to level 2, but instead of two ‘unacceptable’ answers, this time two ‘acceptable’ answers are provided, in addition to the ‘exemplary’ one. Hence, the main challenge for the user is to distinguish the exemplary answer from the acceptable ones.
4. Similar to level 3, but now the list of potential answers also includes an answer that is applicable to the type of aggression that is not applicable.
5. Similar to level 4, but now the difference between the answers is again more subtle (see Table 10.1), with makes it even harder to select the ideal one.
6. Similar to level 5, but with an additional time limit included. If no answer is selected before that time, the answer is considered unacceptable.

In principle, the system determines at random whether it offers a case of emotional or instrumental aggression. However, if the trainee performs significantly worse on one type of aggression, that particular question type will be offered more often, to facilitate learning. Furthermore, also the order in which the answers are presented in the multiple-choice menu is determined randomly.

Level	Aggression type shown	Time Limit	Types of answers (<i>type of aggression</i>)
Level 1	Yes	No	2x unacceptable 1x exemplary
Level 2	No	No	2x unacceptable 1x exemplary
Level 3	No	No	2x acceptable 1x exemplary
Level 4	No	No	1x (un)acceptable (<i>incorrect type</i>) 1x (un)acceptable 1x exemplary
Level 5	No	No	1x acceptable/exemplary (<i>incorrect type</i>) 1x acceptable 1x exemplary
Level 6	No	Yes	1x acceptable/exemplary (<i>incorrect type</i>) 1x acceptable 1x exemplary

Table 10.1: Difficulty levels

10.4.4 Transitions between levels

To determine when the difficulty level needs to increase or decrease, the system needs to keep track of the user's performance. This can be done by keeping score. Because the training will consist of different questions for the two types of aggression, for each type a separate score will be kept (as some trainees could be good in de-escalating one type of aggression, but may have difficulties with the other type). In order to reach a higher level, the score for both types of aggression needs to be sufficiently high to meet the demands of a level. To be a bit lenient, one error can be made without directly falling back a level.

The first part of the training (level 1-3) will focus on training the correct approach per aggression type. Once this is mastered, the second part of the training (level 4-6) will give answers that match reactions for both types of aggression, to test if the trainee can tell them apart. Levels are determined 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. The transitions between levels are depicted in Figure 10.1.

Scores below 0 in level 1 will remain zero. To complete the training, a score of 36 or higher is needed for both types of aggression. Because each level has different combinations of answer types, the score for each type of answer is determined per level. The entire scoring mechanism is shown in Table 10.2 (note that this mechanism is the same for both types of aggression).

Level	Types of answers (<i>type of aggression</i>)	Score
Level 1	2x unacceptable	-1
	1x exemplary	+1
Level 2	2x unacceptable	-1
	1x exemplary	+1
Level 3	2x acceptable	-1
	1x exemplary	+1
Level 4	1x (un)acceptable (<i>incorrect type</i>)	-2
	1x (un)acceptable	-1
	1x exemplary	+1
Level 5	1x acceptable/exemplary (<i>incorrect type</i>)	-2
	1x acceptable	-1
	1x exemplary	+1
Level 6	1x acceptable/exemplary (<i>incorrect type</i>)	-2
	1x acceptable	-1
	1x exemplary	+1

Table 10.2: Scoring mechanism

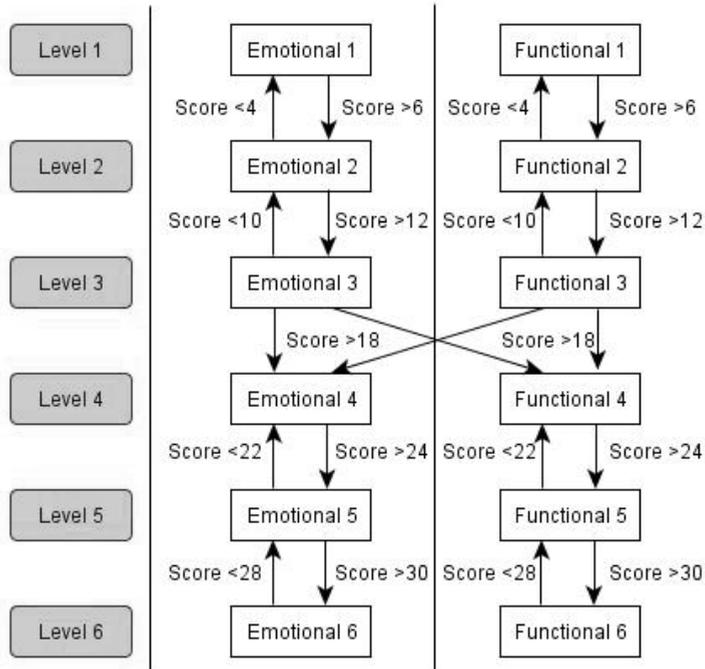


Figure 10.1: Transitions between levels

10.5 Implementation

In order to evaluate the conceptual model described above, it has been implemented using the Python programming language. For this implementation, abstract text-based questions have been used, as the intended VR was not yet fully functional. Nonetheless, this version implements the entire conceptual model and can be used within the VR environment with little effort.

To implement the model, a program has been written that loops through a number of different functions, as shown in Figure 10.2. The double line shows where the loop starts at the beginning of the training, and where it stops at the end of the training. It will determine the level the trainee is in, and the *question type* that will be given (i.e., emotional or functional aggression). Given the level, it will determine the list of potential answers and present the question to the user. After the user gives input, the input is processed and the score is determined based on the given answer. If the trainee has got enough points for both types of aggression, the training ends.

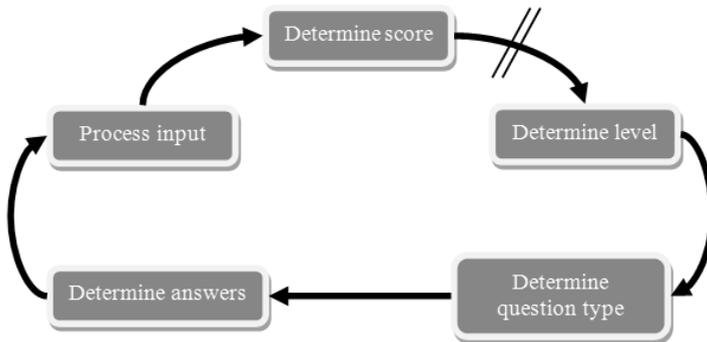


Figure 10.2: The main functions of the implementation of the adaptive model

Determine level. Before a question can be selected, first the level the trainee is in has to be determined. Based on the current level and score, a decision is made whether or not to advance (or demote) the trainee to another level.

Start training. Part of this function is a subfunction which initializes the training for the current setup. This is only done the first time the trainee starts a training session.

Determine question type. As mentioned before, the type of aggression can be either emotional or functional. Normally, the type of question is selected at random. However, when the difference between the scores is larger than some predefined value d , questions will be selected for the type with the lowest score. Currently, a value of 5 is used for d .

Determine answers. The options to appear in the multiple choice list presented to the trainee need to be carefully selected, depending on the current level of the trainee. Each conversation contains many possible responses to particular questions, categorized on aggression type and degree of correctness (unacceptable, acceptable, exemplary). Answers are selected for the current question according to the scheme depicted in Table 10.1. If there are more options for a particular answer, one is selected at random.

Process input. At this point, all information about the current question to be presented to the user is available. For the interaction with the user, two subfunctions are created; the first presents the question to the user, while the second waits for his or her input.

Present question. This function simply displays the current question, with the possible options displayed in the multiple choice list in random order. Only in the lowest level, the aggression type is shown to the user as well.

Process input. The trainee can select his or her choice of answer by pressing the corresponding key on the keyboard. Only when the trainee is currently in the highest level, a time limit is imposed on the user. Otherwise, the trainee can take as much time as desired in coming to a decision on which response to give.

Determine score. To determine the score, the given answer is evaluated based on the question type and correctness of the answer. Next, the score is updated based on the scheme presented in Table 10.2.

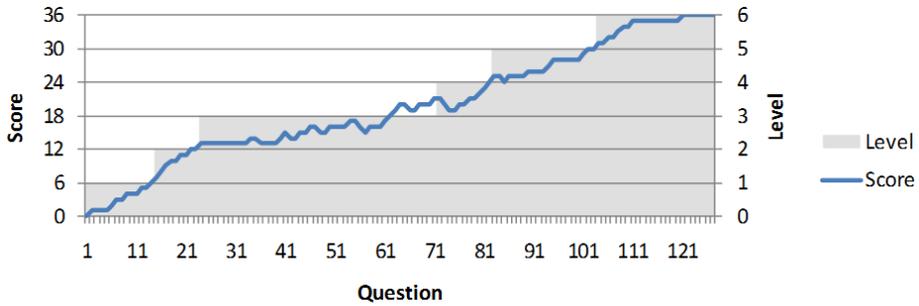
10.6 Preliminary evaluation

To test whether the implementation works as described in the conceptual model, a number of functional evaluations have been performed. By systematically running the program for a number of test sessions, a range of information has been obtained about how the training progressed in different situations. Below, an overview is presented of the results of one illustrative test session, and an analysis is made of whether they correspond to the expected behavior.

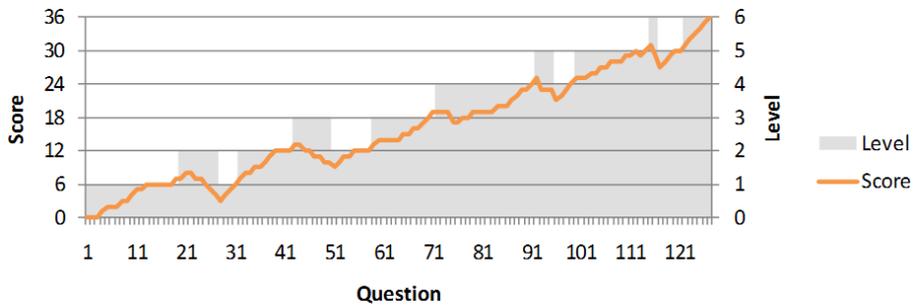
First, Figures 10.3(a) and 10.3(b) below the scores are shown for functional and emotional aggression respectively. In gray on the background, the current level of the trainee for that type of aggression is displayed. Here, it can be seen that the score increases throughout the training and drops in some cases where incorrect answers are provided. Although it is difficult to see in these graphs, a closer inspection of the data has shown that this changing score follows the scheme as described in Table 10.2.

Focusing on level progression, Figure 10.3(a) shows clearly that the levels increase if the trainee reaches the required score, except for level 4 which is only reached if the score for both types of aggression is sufficient. Looking at Figure 10.3(b), the trainee's level decreases a couple of times after (s)he made a number of mistakes. However, when the trainee has reached level 4, and again makes some mistakes, the level does not drop back down to three, which is consistent with the intended behavior of the model.

Another important aspect of this approach is the selection of the possible answers for each level. Figure 10.4 shows an overview of how the multiple choice options were distributed among the various categories of possible answers. As there are three options the trainee can choose from for each question, for each level always a third of the answers is from the category *exemplary*, meaning



(a) Functional aggression



(b) Emotional aggression

Figure 10.3: Score over time with the corresponding level of the trainee

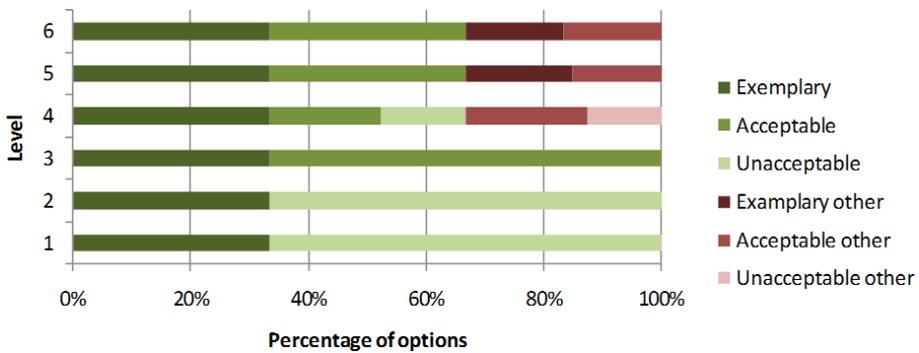


Figure 10.4: Distribution of multiple choice options per level

that there is always one exemplary option to choose from. In level 1 and 2, the other options come from the *unacceptable* category (with the difference between the levels being whether or not the type of aggression is given), while in level 3 *acceptable* answers are given as alternatives. Starting from level 4, options for the other type of aggression are also given in such a manner that there is always one non-exemplary alternative for the right type of aggression and another alternative for the wrong type of aggression. In level 4, both alternatives can be of the unacceptable kind, while in level 5 it is either acceptable or an exemplary answer, but for the other type of aggression. Between level 5 and 6, no difference can be seen, as the only addition for level 6 is the time pressure.

To obtain these results, a few people have already worked with the adaptive system. Although this is not yet an extensive subjective evaluation, feedback from these users was of such a positive nature that we consider it noteworthy to mention here. Nevertheless, as discussed below, a more extensive evaluation of this adaptive training method is planned for the near future.

10.7 Discussion

Being able to de-escalate confrontations with aggressive individuals is a useful skill, in particular within professions where such confrontations are likely to happen, e.g., in the domains of public transport or public safety. Nevertheless, offering appropriate training that enables students to develop such skills is a nontrivial matter. In particular, existing (real world) training approaches are limited in terms of personalization: since the training is typically offered to groups of students together, it is hard to tune the content of training scenarios to individual needs.

As a complementary approach to real-world training, the current paper proposed a simulation-based environment for training of aggression de-escalation. The environment consists of a dialog system that allows a trainee to engage in a conversation with a (possibly aggressive) virtual agent. The agent can show aggressive behavior in terms of emotional speech, gestures and facial expressions. By observing these cues, the trainee needs to assess the situation (specifically: assess which type of aggression is shown) and select an appropriate response via a multiple choice menu.

The main focus of the current paper was on a module to make the system adaptive to the performance of the trainee. To this end, first a number of separate *learning goals* were identified, such as ‘recognizing the type of aggression’ and ‘being able to make decisions under time pressure’. Based on these learning goals, a number of levels of difficulty were identified, as well as a mechanism to switch up and down between these levels based on the user’s score. A preliminary evaluation demonstrated that the system successfully adapted

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

Obviously, this finding should not be interpreted as a definitive proof that the adaptation mechanism results in quicker or better learning than a non-adaptive training system. To test this more specific hypothesis, future research will involve an experiment to systematically compare the effectiveness of the proposed training system with a non-adaptive one. Additionally, for follow-up research it will be interesting to compare the proposed (manual) adaptation mechanism with a mechanism in which the transitions between levels are learned automatically based on user performance data, and to explore possibilities of using data gathered across multiple trainees to improve the experience for any individual trainee.

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

References

- Abraham, M., Flight, S., & Roorda, W. (2007). Agressie en geweld tegen werknemers met een publieke taak. *Onderzoek voor Veilige Publieke Taak*, (pp. 2009–2011).
- Anderson, L. N., & Clarke, J. T. (1996). De-escalating verbal aggression in primary care settings. *The Nurse practitioner*, *21*(10), 95–107.
- Angie, A. D., Connelly, S., Waples, E. P., & Kligyte, V. (2011). The influence of discrete emotions on judgement and decision-making: A meta-analytic review. *Cognition & Emotion*, *25*(8), 1393–1422.
- Bandura, A., & Walters, R. H. (1963). *Social learning and personality development*. Holt, Rinehart and Winston.
- Berkowitz, L. (1978). Whatever happened to the frustration-aggression hypothesis? *The American Behavioral Scientist (pre-1986)*, *21*(5), 691–708.
- Bonner, G., & McLaughlin, S. (2007). The psychological impact of aggression on nursing staff. *British journal of nursing*, *16*(13), 810–814.
- Bosse, T., & Gerritsen, C. (2014). Scenario's gvb. Tech. rep., VU University Amsterdam.
- 'BZK' (2008). Handboek agressie en geweld - voorkomen, beperken, afhandelen. Tech. rep., Ministerie van Binnenlandse Zaken en Koninkrijksrelaties.

- Christoph, L. H. (2006). *The role of metacognitive skills in learning to solve problems*. Ph.D. thesis, University of Amsterdam.
- Graafland, M., Schraagen, J., & Schijven, M. (2012). Systematic review of serious games for medical education and surgical skills training. *British journal of surgery*, *99*(10), 1322–1330.
- Graham, S., & Weiner, B. (1996). Theories and principles of motivation. *Handbook of educational psychology*, *4*, 63–84.
- Heuvelink, A., & Mioch, T. (2008). Fega: A feedback-generating agent. In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, (pp. 567–572). IEEE Computer Society.
- Holmes, J., Gathercole, S. E., & Dunning, D. L. (2009). Adaptive training leads to sustained enhancement of poor working memory in children. *Developmental science*, *12*(4), F9–F15.
- Hoogen, W. v. d., IJsselsteijn, W., Kort, Y. d., & Poels, K. (2008). Toward real-time behavioral indicators of player experiences: Pressure patterns and postural responses. *Proceedings of Measuring Behaviour 2008*, (pp. 100–101).
- Hubbard, J. A., Smithmyer, C. M., Ramsden, S. R., Parker, E. H., Flanagan, K. D., Dearing, K. F., Relyea, N., & Simons, R. F. (2002). Observational, physiological, and self-report measures of childrens anger: Relations to reactive versus proactive aggression. *Child development*, *73*(4), 1101–1118.
- Hulst, A., Muller, T., Besselink, S., Coetsier, D., & Roos, C. (2008). Bloody serious gaming: experiences with job oriented training. In *Proceedings: Interservice/Industry Training, Simulation, and Education Conference (I/IT-SEC)*, (pp. 375–385).
- James, A., Madeley, R., & Dove, A. (2006). Violence and aggression in the emergency department. *Emergency Medicine Journal*, *23*(6), 431–434.
- Kim, J. M., Hill Jr, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S., Pynadath, D., & Hart, J. (2009). Bilat: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, *19*(3), 289–308.
- Miller, J. D., & Lynam, D. R. (2006). Reactive and proactive aggression: Similarities and differences. *Personality and Individual Differences*, *41*(8), 1469–1480.

- Nogueira, P. A., Rodrigues, R., Oliveira, E. C., & Nacke, L. E. (2013). Guided emotional state regulation: Understanding and shaping players' affective experiences in digital games. In *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, (pp. 51–57).
- Peeters, M. M., Bosch, K. V. d., Neerincx, M. A., & Meyer, J.-J. C. (2014). An ontology for automated scenario-based training. *International Journal of Technology Enhanced Learning*, 6(3), 195–211.
- Rafaeli, A., Erez, A., Ravid, S., Derfler-Rozin, R., Treister, D. E., & Scheyer, R. (2012). When customers exhibit verbal aggression, employees pay cognitive costs. *Journal of applied psychology*, 97(5), 931–950.
- Raybourn, E. M. (2007). Applying simulation experience design methods to creating serious game-based adaptive training systems. *Interacting with Computers*, 19(2), 206–214.
- Ritterfeld, U., Cody, M., & Vorderer, P. (2009). *Serious games: Mechanisms and effects*. Routledge.
- Westra, J., van Hasselt, H., Dignum, F., & Dignum, V. (2009). Adaptive serious games using agent organizations. In F. Dignum, J. Bradshaw, B. Silverman, & W. van Doesburg (Eds.) *Agents for Games and Simulations*, vol. 5920 of *Lecture Notes in Computer Science*, (pp. 206–220). Springer Berlin Heidelberg.
- Wickens, C. D., Hutchins, S., Carolan, T., & Cumming, J. (2013). Effectiveness of part-task training and increasing-difficulty training strategies a meta-analysis approach. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 55(2), 461–470.
- Yannakakis, G. N., Togelius, J., Khaled, R., Jhala, A., Karpouzis, K., Paiva, A., & Vasalou, A. (2010). Siren: Towards adaptive serious games for teaching conflict resolution. *Proceedings of the 4th European Conference on Games Based Learning*, (pp. 412–417).