Designing an Expressive Avatar of a Real Person

Sangyoon Lee, Gordon Carlson, Steve Jones, Andrew Johnson,
Jason Leigh, and Luc Renambot

Electronic Visualization Laboratory, University of Illinois at Chicago,
851 S. Morgan St. Chicago, IL 60607, USA
{slee14,gcarls4,sjones,ajohnson,spiff,renambot}@uic.edu

Abstract. The human ability to express and recognize emotions plays an important role in face-to-face communication, and as technology advances it will be increasingly important for computer-generated avatars to be similarly expressive. In this paper, we present the detailed development process for the Lifelike Responsive Avatar Framework (LRAF) and a prototype application for modeling a specific individual to analyze the effectiveness of expressive avatars. In particular, the goals of our pilot study ($n = 1,744$) are to determine whether the specific avatar being developed is capable of conveying emotional states (Ekmans six classic emotions) via facial features and whether a realistic avatar is an appropriate vehicle for conveying the emotional states accompanying spoken information. The results of this study show that happiness and sadness are correctly identified with a high degree of accuracy while the other four emotional states show mixed results.

Fig. 1. Dr. Alex Schwarzkopf and His Virtual Representation

1   Introduction

An avatar, a human-like computer interface, has been actively developed in various forms and is becoming more and more prevalent in diverse target applications [1]. With widespread adoption of advanced computer graphics, artificial intelligence and various

sensing technologies, a lifelike avatar becomes capable of increasingly natural interaction. The human ability to express and to recognize emotions plays an important role in face-to-face communication. As an avatar interface relies on the naturalness of peoples everyday communication with others, understanding its nature is a fundamental key to success. We understand people better as well as convey our ideas more appropriately given with such a capability in everyday life [2] . Our work makes three different contributions. First, it presents our experiences in creating the avatar framework. Second, we demonstrate an avatar creation process in detail. Finally, we conduct a user study with a prototype application to analyze the effectiveness of avatar emotion compared to that of a human. The rest of the paper is organized as follows. Section 2 reviews studies related to our work. In section 3, we describe an implemented Lifelike Responsive Avatar Framework (LRAF) system including a framework and avatar creation method. User study results are presented in section 4 and the conclusion is in section 5.

2   Related Work

Research on avatars or virtual humans has been ongoing for decades. Yee et al. conducted a meta-analysis of the effectiveness of avatar usage across 40 research projects. They found that visual human-like embodied agent interfaces and their realistic behavior showed significant effects in subjective responses but not in task performance [1]. In contrast to this, a more natural behavior model such as a relational agent provided users with a better experience not only in subjective but also in behavioral measures with a virtual nurse application [3] and a long-term daily exercise application [4]. This implies that the application domain can affect the effectiveness of an avatar. One counterexample in the early study is a poker game with emotional avatars by Koda et al. Some subjects believed that a poker player should not show any emotion during game play [5]. Our research focuses on an avatar's expressive capability. To this end, Ekmans classic study [6] has been widely adopted to model human emotion in avatar interfaces. Diana et al. proposed an affective model for the generation of emotional states using MPEG-4 Facial Animation Parameters and evaluated its recognition rate [7]. A series of studies by Wallraven et al. used a 4D-scanner to capture human expressions and create a computer-animated head. They compared perception accuracy with various rendering techniques such as texture and shape blurring [8] and stylizations of rendering [9]. They found that more realistic rendering was better at providing subjective certainty and quality of the conveyed expression. More recent studies of expressive virtual humans with various graphical features such as dynamic wrinkles, blushing, and tears also confirmed these results [10,11]. However, creation of realistic graphics to resemble a human subject requires either high fidelity hardware or a fair amount of a designers effort. In this study, we present our robust method to create an avatar of a person without unique hardware or a high degree of human intervention.

3   Design and Implementation

Creating a compelling avatar requires many complex steps involving programming, observation, design, and evaluation. In this section, we present our implementation of

a framework, and the method of developing an avatar based on a specific target person. Our first prototype application is designed to allow the user to interact with a senior National Science Foundation program director to retrieve domain specific knowledge about the program which he led (Figure 2).

Fig. 2. LRAF prototype application. A user interacts with a life-size avatar on a 50" display. Our first prototype application is a life-like embodiment of a particular person, retiring National Science Foundation program director Dr. Alex Schwarzkopf. The import aspect of our design goal is to ensure a user's experience as natural and realistic as possible. To this end, a deployed large display matches a target persons physical measures in the realized representation and an ambient microphone enables a user to interact with an avatar more naturally. It is designed to be a life-size avatar simulating face-to-face communication with a real person.

3.1 LifeLike Responsive Avatar Framework (LRAF)

The LifeLike Responsive Avatar Framework is composed of multiple modules to accommodate the various features of a lifelike avatar. Figure 3 shows the high-level framework architecture. In the following subsections, we will give more details of the framework.

Graphics Library and Animation. LRAF relies on the Object-Oriented Graphics Rendering Engine (OGRE) library for its low-level graphics modules [12]. OGRE is an open-source platform independent rendering engine that supports most modern graphics features such as skeletal / shape animation, shader languages, and flexible plug-in architecture. Custom rendering techniques are implemented in addition to the standard set of features.

LRAF avatar animation is a combination of a full body skeletal animation and shape-based facial animation. The Facial Expression Synthesizer controls verbal and nonverbal expressions together - verbal expressions for lip synchronization and nonverbal expression for emotion and other behavioral aspects. The synthesizer integrates
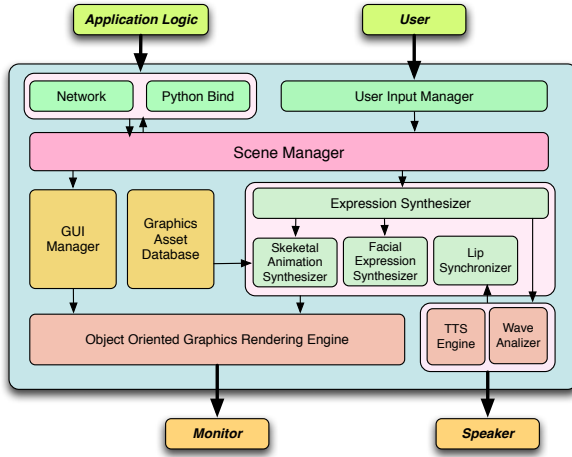
Fig. 3. LRAF architecture diagram

all the weight parameters to blend a facial expression. The Skeletal Animation Synthe-sizer manages full body animation. Avatars intended to mimic human behavior need to behave somewhat non-deterministically or else they will appear unnatural and mecha-nistic. To accomplish this we devised the concept of a Semi-Deterministic Hierarchical Finite State Machine (SDHFSM). An SDHFSM is a hierarchical finite state machine where a sub-state is chosen either based on a set of constraints or randomly given mul-tiple possible options (Figure 4).


Fig. 4. Semi-Deterministic Hierarchical Finite State Machine

Voice and Lip Synchronization. There are two main approaches to generating an avatar voice, a Text-To-Speech (TTS) synthesizer, and a prerecorded human voice. LRAF sup-ports both functionalities to synchronize lip animation. TTS can synthesize speech from unrestricted text. Since Microsofts SAPI TTS engine generates a discrete viseme event while synthesizing a voice in real time, we used a linear blending algorithm to fill the gaps between the distinct lip shapes. The TTS voice module constantly monitors and en-queues a viseme event to the lip synchronizer. The lip synchronizer computes the weight parameters for each viseme blendshape by multiplying the elapsed frame time, the current

Fig. 5. LRAF application using Python binding. A tour guide astronaut avatar describes features of the high-resolution Carina Nebula image on a tile display system.

amplitude of the audio mixer and a constant factor. This is not a precise mouth coarticulation model but it is a feasible solution. We are looking into a better algorithm to improve this in the future, such as King and Parents blendshape based coarticulation model [13]. A recorded voice can be more realistic than a synthesized voice as it uses a real persons voice. We have used a FFT waveform analyzer to monitor several frequency bands. Each band or group of bands can be used to drive a specific lip shape. As the recorded voice plays, the lip synchronizer continuously computes the weight parameters based on the power level of the assigned frequency bands.

External Communication. An external module or application serves as the brain for the avatar. This separation encapsulates the LRAF implementation from the application. LRAF supports two types of communication that expose most of the LRAF interfaces to control avatar behavior such as speech content, animation states, and expressions. One is synchronous network communication and the other is a binding to the Python script language. The network communication model is used in our first prototype application [14].

The Python binding extends the LRAF C++ abstract class Activity to Python. Application developers can inherit this class and implement it in a script based on application requirements. The Activity class defines a hierarchical task model with an internal FSM template. The script can then connect LRAF to an external application that benefits from an avatar-mediated interface (Figure 5).

3.2 Avatar Design

Developing a fully expressive conversational agent takes multiple steps to incorporate various aspects of human characteristics. We mainly used commodity software packages to reproduce a real person as the target model of our prototype agent.

Fig. 6. Skin Texture Enhancement. Top row: default texture (512x512 resolution) from Face-Gen and its rendering result. Bottom row: high-resolution skin texture (4096x4096 resolution) acquired by projecting high quality photos of a target person onto 3D mesh model.

Head Generation. An avatar head is the most complicated piece of the avatar. Primarily, it should support fundamental facial expressions and look reasonably like our target person, Dr. Schwarzkopf. We used Singular FaceGen Modeler to create a base mesh model of the avatar head. It provided an easy way to generate a head model with two photos of the real person, front and side. This model features 39 blendshapes including 7 emotions, 16 modifiers and 16 phonemes. A modifier is similar to a FACS Action Unit that controls small parts of the face such as a blink and eyebrow up / down.

While this is an efficient way to begin, a generated base mesh model and its texture is often not detailed enough or not similar enough to the target model. In particular we noticed that the accuracy of resemblance is lower if the target person is elderly and has many skin wrinkles in the photos. The default texture resolution is also not enough to fully realize the detailed facial features in a high-resolution display. Fine tuned adjustment of facial proportion and texture reconstruction are necessary to achieve the best resemblance. Figure 6 illustrates this realism enhancement by applying photo-based

Fig. 7. Effedt of Normal Map for Skin Details. Left image rendered with color map only and right image used the normal map technique.

Fig. 8. Dynamic texturing on glasses. Live webcam image (left) is blended on the reflective surface of glasses(right).

texture projection. In our current pipeline, a designer manually applies texture projection and blends multiple resulting images into one high-resolution texture. Even with a higher resolution color texture map, the resulting rendering of the skin is rather plain and too smooth. To realize human skin characteristics further, it is necessary to add more subtle features such as wrinkles and pores. The most widely adopted method to implement those details in graphics is to use a tangent space normal map. This technique utilizes the vast computation power of graphics hardware in pixel space without losing frame rate compared to a high-density polygonal mesh model approach. A normal map texture is generated by the color texture based normal extraction method [11]. The final rendering result of a normal map is presented in Figure 7.

A small graphical gadget can improve the user experience and embodiment in interacting with the avatar. One example to fill the gap between the virtual and real worlds is the use of a live video feed on a reflective material in the virtual scene (Figure 8). During our preliminary internal review, most users noted higher engagement with this feature. It will be useful to evaluate this effect in a future study.

Full body Modeling and Motion Acquisition. An avatar full body model was designed in Autodesk Maya software. All physical dimensions of the model matched the target person. Since we use a skinned mesh to animate an avatar body, we built the body skeleton with roughly 70 bones including all of the finger bones and rigged it to the body mesh. Finally, the head model was combined with the body. As we digitized the person, all behaviors and gestures were based on the target person. We utilized an optical motion capture system, Vicon MX-F40, to acquire full body motion and mannerisms. Then a segmentation of each individual motion clip was processed to construct an internal motion database. Each motion clip corresponds to a unit action such as look left / right, various pointing, idle, and other actions. The animation module in LRAF selects proper motion clips and blends them in real time.

Personification and Review. As we develop a specific person's virtual representation, the realism of the agent relies not only on visual realism but also on behavioral characteristics. So far we have discussed how we create a visually compelling agent with the help of advanced graphics. The use of a motion capture system partially solves the problem of simulating the gestural aspects of the agent as it precisely replicates a target persons kinematic movement. For instance, an 82 year old man's keyboard typing may just be pressing keys with one or two index fingers and his natural posture of sitting on a chair is slightly leaning toward one side rather than straight-up. However, developing a more fine-grained mannerism model is still a very active research area and requires a significant amount of effort in observation, modeling, and evaluation. In our first prototype application, we started with several obvious and most distinct mannerisms. For instance, the subject's blink rate is more frequent than the rate of an average male adult. His often used expressions in conversation such as "Hello, friend." and "Keep the peace." are included in the speech corpus. After establishing the design process of this avatar, we applied the same method to produce several virtual representations for other target persons (i.e., the space astronaut in Figure 5). Our method successfully managed to produce a realistic avatar with a few days work. This work involved taking photographs of a person, head generation, texture re-touching, and fitting to a full body model. We either purchased a commercial body models or exported a similar one from a popular modeling package such as Poser. Assuming each individual has unique gestures and full body expression, motion acquisition process remains as labor-intensive effort without much automation.

4   Pilot Study

The goals of the pilot study were twofold: First, determine whether the specific avatar being developed was capable of conveying emotional states; and second, determine, more generally, whether realistic avatars are good vehicles for conveying emotional states accompanying spoken information. In this study we used still renderings of an avatar and photos of the human the avatar was based on to determine whether users identified the emotional states comparably between the avatar and the human upon which it was based. The rendering images used in this study are the intermediate result of

graphical enhancement as the survey was conducted while we continuously improved our visual representation techniques.

4.1   Design and Procedure

The human model of the avatar, Dr. Alex Schwarzkopf, was chosen for our study. Our work here is based on Ekmans [6,15,16] approach to expressing emotions which have shown positive results in using comparative images to study human emotion recognition. The study draws on methods from two studies [17,18] and merges them together to focus specifically on the avatar, and to incorporate a larger pool of research subjects available online. Images were used for two reasons: 1) videos create significant issues of reliability in large online studies and 2) previous studies using this approach have yielded useful results. To extend this work into the avatar world, photographs were taken of Alex exhibiting six classic emotional states: anger, fear, disgust, happiness, sadness, and surprise. Three photos of each emotional state were selected based on how well they corresponded to the elements of Ekmans emotional characteristics (18 total human images). Images of the avatar were rendered to mimic the photos of Alex as closely as possible by manipulating key facial variables (18 total avatar images) in a process to similar studies listed previously. Eyeglasses were removed to avoid interfering with facial features. The avatar renderings used were not photorealistic but had the prominent facial features necessary. In some cases, we appended phoneme shape, head orientation, and eye gaze in addition to the modifier shape to obtain the best match (Figure 9).

Subjects were directed to an online survey tool where they were shown the 36 images, randomly ordered, and asked to identify which of the six emotional states the face was displaying. Subjects were only allowed to pick from the six emotional states and there was no other or none option. Subjects were recruited from across the student population of a major research university with approximately 25,000 students. Mass emails, posters, and word of mouth in classrooms were used in conjunction with recruiting incentives. After removing erroneous or invalid data 1,744 subjects participated (n=1744). Gender was split almost evenly: 864 males and 867 females, and ages ranged from 18 to 64 (*mean*=23.5, *median*=22, *mode*=20).

Fig. 9. A Sample Happiness Emotion. Avatar emotion is rendered with weight parameters as follows: Smile(1.0), BlinkLeft(0.2), BlinkRight(0,1), BrowUpLeft(0.2), BrowUpRight(0.1), and Phoneme B(0.65).

4.2   Result and Discussion

We sought two measures: (1) did the subjects correctly identify the emotion displayed and (2) did the subjects match the emotion for each human/avatar pair? Subjects did not identify anger in either the human or avatar to a useful degree. In four of the six anger


**Fig. 10.** Disgust emotion pair of images and results

Fig. 11. Happiness emotion pair of images and results

Table 1. Percentage of valid responses identifying the emotion displayed. Three pairs of images (each pair made up of one avatar rendering and one photo of the real human) for each emotion tested. Left column indicates what emotion was intended to be depicted. Columns represent what percentage of subjects identified each emotion in the image. Percentages in bold represent the most popular selection made by subjects regarding that pair. The far right column represents the number of valid responses from subjects (n). Highlighted pairings are the most successful based on Paired Samples T Test for each human/avatar pairing (threshold is $p < 0.05$). Thus, the table illustrates that subjects seem to recognize happiness and sadness between the human picture and avatar rendering.

Emotions	Anger	Disgust	Fear	Happiness	Sadness	Surprise	n
	41.9 / 17.6	27.4 / **34.1**	2.4 / 12.3	17.6 / 8.5	3.3 / 22.5	7.5 / 5.1	1597
Anger	**30.5** / **31.8**	18.3 / 19.3	10.2 / 3.1	2.9 / 27.7	13.1 / 10.6	25.1 / 7.4	1624
	49.1 / 17.0	28.0 / 21.4	3.5 / 12.1	0.4 / 10.0	14.0 / **32.5**	4.9 / 6.9	1587
	6.4 / 17.1	35.3 / **27.3**	0.6 / 3.4	**49.5** / 26.5	2.8 / 19.4	5.4 / 6.3	1587
Disgust	7.8 / 13.6	**38.0** / **57.9**	21.6 / 12.1	0.3 / 0.9	11.2 / 3.6	21.0 / 11.9	1684
	13.6 / **44.1**	**42.3** / 30.5	7.4 / 3.9	0.2 / 5.3	34.7 / 15.0	1.9 / 1.3	1593
	3.2 / 3.0	12.5 / 8.3	39.6 / 16.4	0.3 / 1.1	**40.2** / **59.9**	4.2 / 11.3	1592
Fear	3.9 / **78.8**	22.6 / 9.5	25.6 / 4.7	0.4 / 0.7	**42.4** / 5.4	5.0 / 0.9	1612
	30.8 / **69.1**	**32.7** / 11.0	14.8 / 2.1	0.4 / 12.0	13.6 / 1.6	7.6 / 4.3	1601
	0.0 / 0.7	0.1 / 0.8	0.1 / 0.3	**98.7** / **93.9**	0.2 / 0.9	0.9 / 3.3	1599
Happiness	0.2 / 1.1	0.2 / 0.5	0.4 / 0.5	**93.5** / **89.1**	0.2 / 0.5	5.5 / 8.2	1685
	0.3 / 1.2	0.2 / 0.4	0.1 / 0.2	**98.6** / **94.9**	0.1 / 0.4	0.7 / 2.9	1600
	0.7 / 20.9	20.2 / 13.5	1.8 / 5.8	0.8 / 9.5	**74.7** / **46.7**	1.7 / 3.6	1595
Sadness	0.9 / 1.9	2.7 / 4.7	2.2 / 6.3	0.2 / 0.6	**93.6** / **85.5**	0.4 / 1.1	1610
	1.1 / 1.6	7.7 / 4.1	3.2 / 3.7	0.2 / 1.8	**85.6** / **87.6**	2.2 / 1.2	1586
	5.3 / 4.1	13.2 / 10.5	23.3 / 27.7	0.9 / 1.9	8.6 / 27.6	**48.5** / **28.2**	1666
Surprise	5.6 / 8.8	7.3 / 6.3	18.4 / 2.7	2.7 / 21.1	**36.4** / 2.6	29.6 / **58.5**	1604
	5.9 / 7.2	18.4 / 6.2	8.8 / 5.2	1.9 / 31.5	29.2 / 10.3	**35.8** / **39.6**	1594

images the most common response was anger but it was never the majority answer. Disgust did not fare well either though in one pair (Figure 10) subjects did correctly identify the emotion as disgust even if less overwhelmingly in the avatar image. Subjects could not correctly identify the human or avatar images with regard to fear, indicating that perhaps the human images were not sufficiently prototypical. The images indicating surprise were also met with mixed results.

The largest successes were the emotions happiness and sadness. In all six happiness images (three human, three avatar) the results were overwhelmingly correct (Figure 11) and sadness was also identified with a high degree of accuracy. It appears that happiness and sadness are the easiest emotions to artificially indicate on the human face and the easiest to accurately replicate on the avatar. This pilot study provides useful feedback for our work and informs the decisions we will make in the next phases. It appears the current avatar is capable of successfully indicating happiness and sadness. Our avatar indicates happiness to roughly the same degree as the human upon which it was based;

the same is true of sadness (Table 1). The other four emotions anger, fear, surprise, and disgust are not currently indicated by our avatar to any useful degree.

While the avatar did not successfully display the other four emotions, the human photos did not achieve reliable levels of emotional indication either. In fact, there were several pairs where the avatar and human photos were identified in the same incorrect way (e.g. confusing sadness and disgust). One interpretation is that the avatar was sometimes being interpreted the same way as the human, but the human image was not a good prototypical indication of the given emotion. It is possible that the avatar was fundamentally recreating the human expression but that we chose the wrong human face on which to base the avatar. Further research needs to be done to determine whether the remaining four emotions can be better indicated by the human model and, if not, whether we may want to choose a new human to serve as the basis for the avatar.

5   Conclusion and Future Work

In this research, we have developed the Lifelike Responsive Avatar Framework and presented how we created a prototype application for modeling a specific person. With this method we were able to reproduce another avatar in several days without extensive manual intervention. Pilot study results showed that the current avatar is capable of successfully indicating two emotions, happ iness and sadness, but not necessarily the other four. Similarly the human photos in our pilot study showed mixed results for these four emotions.

One aspect we did not consider in the pilot study is that a human recognizes emotions within a context accompanying temporal changes over time and/or spoken information. Further investigation of how those factors affect on avatar emotion is necessary in the future. One good example of a context sensitive task oriented study is[19]. Stern et al performed a user study to examine persuasiveness and social perception of synthetic versus human voice in the context of the persuasive speech. Given all visual realism and higher modality an avatar model studied in small scale should also be evaluated with respect to application quality and performance as a whole in the future. Further dynamic face features such as wrinkle generation [10,11,20] will also be considered in later study.

Considering the results discussed in this work, we conclude that our framework and avatar design method is at least partially capable of successfully conveying human emotions and its accuracy is close to that of the target person.

References

1. Yee, N., Bailenson, J., Rickertsen, K.: A meta-analysis of the impact of the inclusion and realism of human-like faces on user experiences in interfaces. In: Proceedings of the SIGCHI conference on Human factors in computing systems (2007)

2. Ekman, P.: Emotions revealed. Henry Holt, New York (2003)
3. Bickmore, T., Pfeifer, L., Jack, B.: Taking the time to care: empowering low health literacy hospital patients with virtual nurse agents. In: Proceedings of the 27th international conference on Human factors in computing systems (2009)
4. Bickmore, T., Picard, R.: Establishing and maintaining long-term human-computer relationships. ACM Trans. Comput. Hum. Interact. 12(2), 293–327 (2005)
5. Koda, T., Maes, P.: Agents with faces: the effect of personification. In: 5th IEEE International Workshop on Robot and Human Communication, pp. 189–194 (November 1996)
6. Ekman, P.: Emotion in the Human Face. Pergamon Press, New York (1972)
7. Diana, A., Javier, V., Francisco, J.: Generation and visualization of emotional states in virtual characters. Comput. Animat. Virtual Worlds 19(3-4), 259–270 (2008), 1410372
8. Wallraven, C., Breidt, M., Cunningham, D., Bülthoff, H.: Psychophysical evaluation of animated facial expressions. In: APGV 2005: Proceedings of the 2nd symposium on Applied perception in graphics and visualization (August 2005)
9. Wallraven, C., Fischer, J., Cunningham, D., Bartz, D., Bülthoff, H.: The evaluation of stylized facial expressions. In: APGV 2006: Proceedings of the 3rd symposium on Applied perception in graphics and visualization (July 2006)
10. Courgeon, M., Buisine, S., Martin, J.C.: Impact of expressive wrinkles on perception of a virtual character's facial expressions of emotions. In: Ruttkay, Z., Kipp, M., Nijholt, A., Vilhjálmsson, H.H. (eds.) IVA 2009. LNCS, vol. 5773, pp. 201–214. Springer, Heidelberg (2009)
11. de Melo, C.M., Gratch, J.: Expression of emotions using wrinkles, blushing, sweating and tears. In: Ruttkay, Z., Kipp, M., Nijholt, A., Vilhjálmsson, H.H. (eds.) IVA 2009. LNCS, vol. 5773, pp. 188–200. Springer, Heidelberg (2009)
12. The OGRE Team: Ogre - open source 3d graphics engine, `http://www.ogre3d.org`
13. King, S.A., Parent, R.: Creating speech-synchronized animation. IEEE Transactions on Visualization and Computer Graphics 11(3), 341–352 (2005)
14. DeMara, R.F., Gonzalez, A.J., Hung, V., Leon-Barth, C., Dookhoo, R.A., Jones, S., Johnson, A., Leigh, J., Renambot, L., Lee, S., Carlson, G.: Towards interactive training with an avatar-based human-computer interface. In: The Interservice/Industry Training, Simulation & Education Conference, I/ITSEC (December 2008)
15. Ekman, P., Davidson, R.J.: The nature of emotion. Oxford Press, New York (1994)
16. Scherer, K.R., Ekman, P.: Handbook of methods in nonverbal behavior research, p. 593. Cambridge University Press, Cambridge (1982)
17. Mendolia, M.: Explicit use of categorical and dimensional strategies to decode facial expressions of emotion. Journal of Nonverbal Behavior (January 2007)
18. Ellison, J., Massaro, D.: Featural evaluation, integration, and judgment of facial affect. Journal of Experimental Psychology: Human Perception and Performance (January 1997)
19. Stern, S.E., Mullennix, J.W., Yaroslavsky, I.: Persuasion and social perception of human vs. synthetic voice across person as source and computer as source conditions. Int. J. Hum. Comput. Stud. 64(1), 43–52 (2006)
20. Maddock, S., Edge, J., Sanchez, M.: Movement realism in computer facial animation. In: 19th British HCI Group Annual Conference, Workshop on Human-animated Characters Interaction (January 2005)