Talking Virtual Characters for the Internet

Igor S. Pandzic

igor@isy.liu.se

Department of Electrical Engineering

Linköping University

SE-581 83 Linköping, Sweden

Abstract

Talking Virtual Characters are graphical simulations of real or imaginary persons capable of human-like behavior, most importantly talking and gesturing. They may find applications on the Internet as newscasters, customer service representatives, sales representatives, guides etc. We review the potential applications of Virtual Characters on the Internet, their technical requirements and current solutions, and present our approach to the problem. The presented system is an MPEG-4 Facial Animation decoder built in a Java applet and capable of rendering a talking Virtual Character in a standard Web browser without plug-ins.

1. Introduction

The technology of modeling and animating Virtual Humans has been a fascinating research topic for more than two decades. With the recent boom in the Internet services, new applications for Virtual Humans technologies can be identified. We believe in particular that talking Virtual Characters [1][2][3][4][5][6][7][37] can be useful when integrated in Web sites to provide novel services such as virtual hosts, salespersons, newscasters and other. In this article we attempt to identify classes of applications that we believe could emerge in this field. We base our analysis on research work as well as emerging commercial attempts that are currently multiplying on the Internet.

After presenting the applications, we analyze the technical requirements for bringing such applications to life. We then identify five main technical approaches to the problem. We analyze the advantages of each approach and present concrete examples.

In the final sections we present our approach to the problem of bringing Virtual Characters to the Web – an MPEG-4 Facial Animation decoder built in a Java applet.

In the concluding section we attempt to draw conclusions concerning technologies and applications currently emerging in this exciting new area, summarize our approach and discuss future work.

2. Applications

We believe that placing a human-like Virtual Character on a Web site can bring the following benefits:

- Give a personality to the Web site.
- Enable to talk to each person visiting the site; people like to be talked to.
- Make visitors remember main messages better.
- A talking person is more persuasive than written text.

Additionally, as outlined in [8], the Virtual Character may in certain situations alleviate the (sometimes unavoidable) Internet waiting times by “entertaining” the user. The same study has also shown that the end users show a preference for a service enhanced with a Virtual Character to a plain text service.

If correctly implemented, Virtual Characters have low bandwidth requirements and high interaction capacity, making them a natural replacement for video streaming in “talking head” scenarios.

We believe that we are just beginning to recognize possible practical applications, and what is presented here may well be just the tip of the iceberg. We can broadly classify the emerging application areas in the following categories: entertainment, personal communications, navigation aid, broadcasting, commerce and education.

2.1. Entertainment

Jokes and amusing content seem to be among the first targets of many creative technologies and Web-based Virtual Characters are no exception. Animated caricatures of famous politicians (e.g. Figure 1) have already made their appearance, as well as other amusing creatures. Virtual Characters are suitable for delivering this kind of content because they generally do not require very high bandwidth (compared to video) and at the same time they can offer interactivity, i.e. instead of delivering a linear story, the user can trigger various reactions from the character.

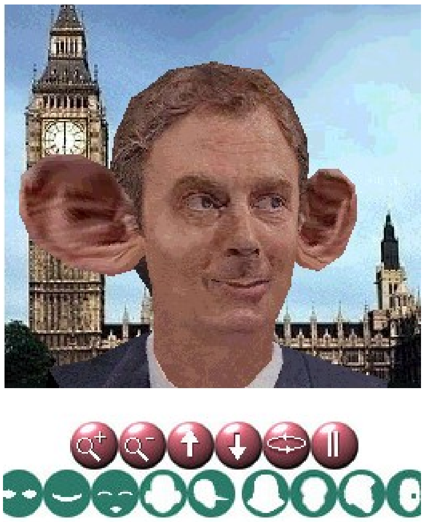

Figure 1. Animated caricature of a politician [9]; viewer can trigger funny actions (© BioVirtual Ltd.)

2.2. Personal communications

Invitations, birthday wishes, jokes, all kinds of personal messages can be delivered either on personal Web pages or through email by appealing, talking Virtual Characters. Electronic greeting cards with 3D Virtual Characters have already been offered by at least one company [10].

2.3. Navigation aid

In the simplest form, a Virtual Character can be used simply to welcome visitors to a Web site by appearing on the home page and delivering a short welcome address to each visitor. More interestingly, a virtual guide can accompany the visitor navigating around numerous pages of a big Web site, always appearing in a corner of the screen and providing guidelines, tips, or even a full guided tour where the Virtual Guide shows various Web pages in particular order and comments them while the visitor sits back and watches the tour.

2.4. Broadcasting

Virtual newscasters will deliver daily or hourly news updates. Unlike the TV newscaster, the one on the Web can be personalized and remember that a particular user is interested in sports, stock market and weather and likes to hear the related news in this particular order (see Figure 2).

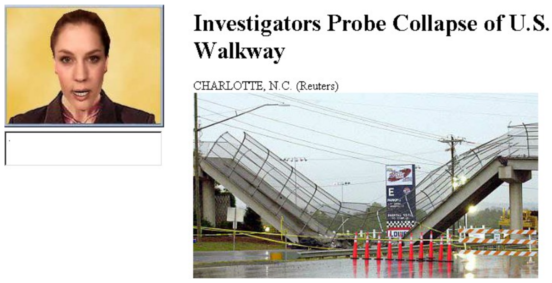

Figure 2. A virtual newscaster; images on the right appear as she talks (© W Interactive SARL)

This may bring an interesting new service integrating advantages of both TV (dynamic audio-visual content) and the Internet (personalized service on demand).

Figure 3. A text based conversational character (© eGain Communications Corp.)

2.5. Commerce

Some companies specializing in Customer Relationship Management (CRM) applications already offer "intelligent" conversational characters that a company can use as the front-line customer support on their web site. Currently most of these characters are purely text based. The customer types in questions in plain language and gets the answer in form of text, sometimes also displaying an appropriate web page. The characters are programmed with an extensive database of answers to all kinds of questions that customers may ask, as well as conversational rules enabling them to perform a semblance of natural conversation. While relatively simple, these characters can answer many "standard" questions in a satisfying way, in many cases avoiding the need for human intervention.

Currently it is usual for such characters to display images of a person (see Figure 3) as a means of giving it an identity. An animated, talking virtual character is the logical next step for such applications, giving them even more personality and making them more user-friendly. Since the characters' answers are mostly pre-programmed in form of text, using text-to-speech and lip synchronization techniques to bring them to life is just a small step away.

Similar approach can be employed for sales. Instead of just showing an image of the product and main information, a virtual salesperson simultaneously provide a “sales pitch” and answer any additional question the potential buyer may have. Alternatively, the customer can ask for a tour of the product, obtaining a rich audio-visual experience with the virtual sales representative explaining the qualities of the product while at the same time showing appealing images, as the example in Figure 4 shows.

Figure 4. Virtual sales representative presenting a car model (© W Interactive SARL)

2.6. Education

Even though we did not find existing examples of using Virtual Characters for educational purposes on the Web, we did find a certain interest in the education community for using this technology in online learning, e.g. staff training, online degrees, distance education.

In the following sections we discuss the technologies needed to deploy all these possible applications by first analyzing the technical requirements defined by the application scope, then presenting and analyzing the currently available technologies.

3. Requirements

Based on the proposed applications, in this section we analyze the technical requirements on the implementations of Virtual Characters for the web, and their possible implications. We have identified the following requirements: visual quality, easy installation, fast access, interactivity, Web integration.

3.1. Visual quality

It is obvious that Virtual Characters featured in any of the mentioned applications must have a visual appeal. They must look nice, and animations should look natural. This does not necessarily mean photo-realistic models, or even very high-resolution models. A fairly simple cartoon character, if cleverly designed, can carry loads of appeal. The real implication of this requirement is that any system hoping for success must provide means for visual artists to design models and animations, preferably using the tools that they are already used to.

3.2. Easy installation

In many of the mentioned applications the virtual character is not in itself the main attraction of the proposed service, but rather an extra bonus, an improvement. The users who never used such a novel application might be reticent to install software on their computers only to experience an unknown improvement of the service. This creates a barrier. To make the user jump over the barrier, one can either provide strong incentive in form of attractive marketing, or lower the barrier by making it easy to install the virtual character support. Ideally, no installation should be necessary.

3.3. Fast access

Although broadband-for-everyone is being promised by many, the reality is that most users do not have fast Internet access. Therefore the new applications involving Virtual Characters should not require high additional bandwidth. This implies several things. First, the virtual human models should not be too complex. Second, they need to be compressed for download. Third, both audio and animation data should not only be compressed, but also streamed, rather than downloaded and played. This allows for a faster response time.

3.4. Interactivity

When asked why Virtual Characters are better than video streaming, the easy answer is: they save bandwidth. But that is just the beginning. The main advantage is that they can interact with the user, i.e. provide immediate responses with respect to users’ actions (mouse clicks, typing, in the future possibly speech, expressions and gestures). These interactions are unique for each individual user. This kind of interaction is impossible with video streaming, and this is the real advantage of Virtual Characters.

To achieve interactivity, several issues are important.

First, there needs to be an easy, preferably automatic way to generate contents. This is because no interesting interaction can be achieved without enough variety of contents. In simple words, a virtual character that can just say “yes” and “no” is not very interesting. If generating contents is expensive, designing the whole application will be very expensive. It is preferable to have the possibility to generate contents automatically. In the ideal case, the speech and behavior of the virtual character is completely generated on-the-fly, putting no limits to the variety of interaction. Unfortunately, this is in most cases opposing the requirement on the quality (best content is still created with a lot of manual work), and a compromise must be found.

Second, it is desirable to have a fairly sophisticated decision-making mechanism, possibly based on AI algorithms, to deliver meaningful interactions.

Third, there must be some kind of user input. It can be within the virtual character application itself, or it can come through the rest of the Web page through standard Web page elements. This brings us to the next requirement, the Web integration.

3.5. Web integration

The virtual character will in most cases not appear alone on a Web page. There will be other contents: text, graphics, forms, buttons etc. It is important for the virtual character to be well integrated with all this, e.g. to be able to react to users' actions in the Web page, and to control and change the contents of the page. This implies that an open interface should exist to make logical links between the Web page elements and the virtual character, in both directions.

Having presented the potential applications and technical requirements, in the next sections we will review the current efforts in implementing this kind of systems. We will first briefly analyze the underlying technology of 3D graphics on the Web, in order to support the next section which reviews virtual human implementations on the Web.

4. Virtual Characters on the Web

The current efforts in Virtual Characters on the Web can broadly be categorized in the following classes: text based bots, 2D animation, video streaming, 3D plug-ins, 3D without plug-in.

4.1. Text based bots

These applications, often called chatterbots, use artificial intelligence algorithms ranging from simple keyword search to more complex pattern matching or neural networks algorithms in order to perform a conversation with a human. The users are usually typing their input to the dialogue in a text field or prompt, and the answers come back as plain text. The earliest and probably most famous example of this technology is Eliza [18], the computer program capable of engaging a human in a conversation.

Currently companies like eGain [19], Artificial Life [20] and Native Minds [21] offer bot-based products aimed at providing customer service on the Web. The bot is programmed with a database corresponding to a relatively narrow subject field (e.g. product information) making it realistic to obtain very reasonable responses, especially to the questions that are frequently asked.

4.2. 2D animation

In this category we find implementations that provide animation by using a set of pre-rendered images. The images correspond to key postures of a face: visemes, expressions and movements. A viseme is the visual effect of pronouncing a phoneme and it represents a particular mouth shape. As mouth shapes for certain groups of phonemes (e.g. p, b, m) correspond to highly similar mouth shapes, there are less visemes then phonemes. It is therefore possible to obtain a crude speech animation with a limited set of viseme images, typically between 8 and 15. To this are added a few images for expressions (e.g. smiling, surprise) and movement (e.g. nod). Therefore the total volume of data required for a simple but effective animation stays low, which makes it suitable for the Internet applications today, when broadband Internet access for the masses is still some time ahead. Additionally, fast manipulation of images is possible in current Web browsers using standard tools like Java or Flash. Low bitrate and standard tools mean that this approach is perfectly realistic and these applications easily accessible by virtually all Web users. The tradeoff is in accepting a lower quality and variety of animation then it would be possible with a full 3D solution (e.g. the face can not turn or smile while talking).

Examples of this approach include W Interactive [22], with fully interactive, personalized Java-based Virtual Characters, and Virtual Personalities [23] offering an animation-enhanced chatterbot apparently based on Flash.

4.3. Video streaming

This approach consists in rendering a video sequence containing a virtual character off-line, then using a de-facto-standard video streaming technology (RealVideo) to deliver the video sequence to the user. It inherits the disadvantages of the video streaming on the Internet of today, namely the small image size and low quality, unless using a broadband connection, which is out of reach for most users. Even worse is the lack of interactivity. The user can only play the video or stop it. The potential of Virtual Characters to interact, deliver answers to specific questions, or provide customized contents is completely lost.

The example of this approach is Ananova [24], "the first virtual newscaster", used as a gimmick in a brilliant marketing campaign to promote an otherwise conventional news site.

4.4. 3D plug-ins

This is currently the most wide-spread approach to delivering Virtual Characters on the web. In some cases, a general-purpose 3D plug-in is used to deliver the characters. In other cases, the plug-in is designed specifically for the virtual character animation. We will review these two approaches separately.

4.4.1 General-purpose plug-ins

Although none of the currently available 3D plug-ins is truly widespread among Web users, these plug-ins offer a wider functionality and it is therefore more likely that they will be widely accepted and become de-facto standards.

Several plug-ins exist for displaying VRML contents in browsers. This enables us to exploit the functions of VRML and use them to display and animate Virtual Characters. Babski and Thalmann [25] present various techniques for doing this, including sensors, interpolators, scripts and Java applets using the External Application Interface. However, same authors point to several problems in using current implementations of VRML plug-ins. They do not support all VRML functions in exactly the same way, and it is therefore difficult to design an application that will run in all configurations (usually creators recommend a particular plug-in). Also, the plug-in implementations are sometimes not very stable, and may demand a lot of resources.

A 3D plug-in that has rapidly gained popularity is Pulse3D. The Pulse3D production system allows easy export of animated objects from the widely used 3D design software 3DS Max. These objects and animations can then be played in Web pages using the Pulse3D plug-in. They are compressed in a proprietary format for relatively fast delivery, and the rendering capabilities of the plug-in allow for very good graphical quality. The use of a familiar tool, interactive production system and appealing results have attracted the computer animation community and Pulse3D has become quite popular. Several independent productions [28] [29] [9] [30] illustrate the use of Pulse3D. In terms of visual quality, these are certainly among the finest animated Virtual Characters on the Web today. However, it is also obvious that producing these animations is relatively expensive in terms of artists' time.

4.4.2 Plug-ins for Virtual Characters

These implementations require the user to download a particular plug-in that is used only to deliver their virtual character contents. This is supposed to guarantee a superior quality of animation and allows the developer to build in other features like streaming, interactivity etc. The disadvantage is that the user is required to download and install it just to access this particular content.

The examples of Virtual Character plug-ins include the Virtual Friend by Haptek Inc. [10], NoDNA [31], Famous3D [32] and Mendel3D Avatar [16].

4.5. 3D without plug-in.

Having mentioned the downside of the plug-in approach, it is attractive to look for a solution that would work on most users' desktops immediately and without any installation. Currently such solutions are possible using Java applets, which are supported on the vast majority of desktops today.

This solution also has its downsides. First, an applet providing full 3D capability is necessarily fairly big, on the order of 100K, which is acceptable but not very convenient. Second, the performance of a Java applet is still lower than that of a well-implemented plug-in, limiting the complexity (and therefore visual quality) of the Virtual Characters that can be used.

On the positive side, beside almost universal portability, the applet programmer can implement streaming capabilities, as well as interaction support for seamless integration with the rest of the Web page.

Red Ted offers Javahead [33], an applet that displays a talking face and uses a proprietary streaming audio format and automatic lip sync with audio. This make for the easy creation of contents based on natural speech recording, but the lip sync quality is quite poor.

Another example of this approach is our own implementation, described in more detail in the following sections.

5. An MPEG-4 Facial Animation system in an applet

The requirements on Web-based Virtual Characters outlined earlier in this paper were used as design criteria for implementing our Facial Animation system for the Web. In the next subsections we discus how this implementation satisfies each of the given criteria.

5.1. Visual quality

It is doubtless that a graphical system capable of rendering fairly complex 3D models at interactive frame rates is a necessity in order to obtain visually pleasing animations While a Java-based implementation does not provide the best performance available, the recent results are quite encouraging and can approach very reasonable levels of quality.

Furthermore, a fairly simple model, if carefully designed and animated, can be as appealing as a more complex one. We therefore believe that the most important requirement for achieving visual quality is the openness of the system for visual artists. It should be convenient for them to design face models with the tools they are used to. While numerous algorithmic facial animation systems have been developed, the best-looking animations in current productions are done manually by artists or by facial tracking equipment and performing talent. This manual creation is painstakingly time-consuming, but on the other hand the fast algorithmic animations may not give satisfying results. We therefore propose a compromising solution by letting the artist create the basic building blocks of the facial animation system, i.e. design the facial movement corresponding to individual Facial Animation Parameters. Once all the animation building blocks are designed, they are used to produce animations automatically.

5.2. Easy installation

We believe that installation of particular software plug-ins in order to view character animation will deter many users from using it, especially in the (numerous) situations where the Virtual Character is an added bonus, but not the main attraction of the application. Our strong design goal was to create a system that will instantly run on a majority of today's platforms and browsers, without any installation. The only solution satisfying this requirement is the Java applet and therefore we chose this implementation.

5.3. Fast access

Although broadband is the big promises, it is not available on an average desktop yet, and therefore applications have to be sparing in their use of bandwidth. The face model in our implementation is compressed efficiently in the Shout3D format. Animation uses the MPEG-4 FBA encoding and streaming. The MPEG-4 standard provides specification for the modeling and animation of human faces and bodies [4][6][38], commonly called FBA (Face and Body Animation) after the MPEG-4 subgroup that developed the specification. For facial animation high- and low-level Facial Animation Parameters (FAPs) are available. The high-level FAPs consist of visemes and facial expressions. The low-level FAPs express movements of defined feature points of the face. They are normalized with respect to the face geometry, meaning that the same set of parameters can be reasonably applied on different face models. MPEG-4 is also a very efficient way of encoding facial animations, with bitrates ranging from 2 to 5 kbit/s. Audio streaming has been implemented as well.

5.4. Interactivity and Web integration

Java implementation makes it easy to integrate the Virtual Character in the Web page and interact with it. The Virtual Character can be controlled using JavaScript commands, allowing to link interactions to all elements of the Web page (links, buttons, forms etc.). At the same time it is easy to control the Web page presentation diretly from the Java applet, allowing the Virtual Character to present various contents in the Web page while it speaks, in a synchronized manner (bearing in mind the possible delay in downloading the contents over Internet which may cause certain loss of synchronization).

6. Implementation

The current implementation of the MPEG-4 Facial Animation player is written entirely in Java and based on the Shout3D rendering engine [17]. The off-line content production tool is written in C++ and permits to compose a facial model with animation based on a series of separate VRML facial models, each model containing the face with a different expression corresponding to high- and low-level MPEG-4 Facial Animation Parameters. The final model that is used by the player is a standard VRML file containing the animation information in Interpolator nodes. This information is registered in a compact way that can be interpreted by the player. This means that the model will be visible in any VRML player, albeit static; in order to view animations the Facial Animation player applet has to be used. The content production tool also includes a Text-to-Speech capability with the encoding of an FBA bitstream synchronized with the audio file, thus enabling the fast production of facial animation contents with speech.

7. Results

We have successfully tested the MPEG-4 Facial Animation player with several facial models as illustrated in Figure 5. We can achieve interactive frame rates with models of up to 3000 polygons. The player has correctly interpreted the test FBA bitstreams (Marco, Wow, Emotions) as well as the bitstreams produced by the Text-to Speech system. As the demonstration page [34] shows, the applet is fully controllable from the Web page by Javascript, making all interactions possible.

Table 1 summarizes the bandwidth requirements of our implementation. The size of the applet is problematic, due mostly to the size of Shout3D renderer implementation (the Facial Animation implementation, including the MPEG-4 FBA decoder, is fairly small). This can possibly be improved by a more rigorous selection of necessary Shout3D classes. The face model sizes, as well as facial animation bitstreams, are very reasonable. However, current implementation uses no compression for audio and is therefore very inefficient.

Total applet size	223K
Applet size (Shout3D)	187K
Applet size (FA implementation)	36K
Viseme-encoded FBA bitstreams	~0.3 kbit/s
Low-level FAPs FBA bitstreams	2-6 kbit/s
Audio (.au)	64 kbit/s
Face models (medium complexity)	~50K

Table 1: Bandwidth requirements

Table 2 shows the performance measurements for different face models. Performance tests were done on two different computer configurations. A third measurement was performed using the graphics accelleration hardware. In order to use the graphics accelleration, a special plug-in for the browser needs to be installed (currently available only for IE); this plug-in allows the Shout3D renderer to use the graphics accelleration. The results show that the implementation can animate facial models of medium complexity at interactive frame rates on average hardware. Tests were made in the Internet Explorer browser; most of the measurements were confirmed in Netscape Navigator as well.

Model	Polygons	Size (KB)	Frames/second		
			C1	C2	C3
Demy	2800	32	11,5	16	21
Dummy	1362	50	16,5	21	32
Jörgen	168	40	22	30	60
Candide	168	4	35	60	60
MIRAface	3692	67	10,5	16	32
Cm. Lake	16917	284	0,97	1	1,5

Table 2: Performance measurements for different face models on different computer configurations. C1 = P3/600 laptop; C2 = P3/1000; C3 = C2 with graphics accelleration hardware (ELSA GLoria II-64)

If we analyze our implementation with respect to the earlier outlined requirements, we can conclude that most of them can be satisfied. The visual quality can be achieved by importing artist-designed face models; although it has to be admitted that a Java implementation will not be capable of rendering very complex models, even fairly simple models can be appealing if carefully designed. Installation is not necessary with the applet approach. Facial models are compressed for fast access, and MPEG-4 FBA bitstreams are probably the most efficient way to stream facial animation. Interactivity and Web integration are achieved through an open JavaScript interface is be used to interact with the Virtual Character directly from the Web page.

Figure 5. Examples of face models experimentally animated using the player, from left to right: *dummy*, a model built using 3D modeling software; *Miraface*, a model donated by MIRALab, University of Geneva to ISO as MPEG-4 reference software; *Demy*, designed by Sasa Galic; *Jörgen*, *Candide*, source Linköping University; *Commander Lake*, source 3DS Max

8. Conclusions

We have attempted to predict the possible applications for Virtual Characters on the Web based on current examples and research. We have analyzed the technical requirements for these applications and their implications. Based on these criteria, we have reviewed a large (though not exhaustive) body of current implementations of Virtual Characters specifically aimed at the Web. Finally, we presented our own approach to the problem.

We must conclude that no current implementation satisfies all the requirements we identified. However, all requirements are addressed by one implementation or another and it is certainly possible today to implement practical applications based on some implementations, compromising on certain issues (e.g. fast, portable implementation but not the best visual quality). We have presented our own approach to the problem, based on an MPEG-4 FBA decoder in an applet, and showed that it comes fairly close to satisfying most given requirements. Our future work in this field will concentrate on trying to improve the visual appeal of the Virtual Characters by further enhancing the possibility for artists to produce them using standard tools.

We believe that we are seeing a beginning of an emerging new class of interactive Web applications and services that will benefit from Virtual Humans technologies.

Acknowledgments

This research is partly supported by the VISIT programme of the Swedish Foundation for Strategic Research (SSF) and the European Commission's IST project Interface [35]. We would also like to thank Gael Sannier for his help in certain aspects of implementation. Particular acknowledgment goes to Sasa Galic who designed the face model *Demy* as a part of his student project at the University of Zagreb.

References

[1] M.M.Cohen and D.W.Massaro, "Modeling Coarticulation in Synthetic Visual Speech." In M.Thalmann & D.Thalmann (Eds.) Computer Animation' 93. Tokyo: Springer-Verlag.

[2] Cosatto E., Graf H.P., "Sample-Based Synthesis of Photo-Realistic Talking Heads", Proc. Computer Animation '98, Philadelphia, USA, pp. 103-110.

[3] P. Eisert, S. Chaudhuri and B. Girod, "Speech Driven Synthesis of Talking Head Sequences," 3D Image Analysis and Synthesis, pp. 51-56, Erlangen, November 1997.

[4] Facial Deformations for MPEG-4, M. Escher, I.S. Pandzic, N. Magnenat-Thalmann, Computer Animation 98, Philadelphia, USA, pp. 138-145, IEEE Computer Society Press, 1998.

[HANIM] H-ANIM Humanoid Animation Working Group, Specification for a Standard Humanoid Version1.1, ece.uwaterloo.ca/~h-anim/spec1.1

[5] Kalra P., Mangili A., Magnenat-Thalmann N., Thalmann D., Simulation of Facial Muscle Actions based on Rational Free Form Deformation", Proceedings Eurographics 92, pp. 65-69

[6] Tekalp M.A., Ostermann J., "Face and 2-D Mesh Animation in MPEG-4", Image Communication Journal, Tutorial Issue on MPEG-4 Standard, Elsevier, 2000.

[7] "Computer Facial Animation", F.I. Parke, K. Waters, A K Peters Ltd. 1996., ISBN 1-56881-014-8

[8] "Synthetic Faces: What are they good for?" Igor S. Pandzic, Joern Ostermann, David Millen, The Visual Computer, 1999.

[9] Tony Blair caricature using Pulse3D www.phonyblair.com

[10] Virtual Friend, Haptek Inc., www.haptek.com

[11] VRML, ISO/IEC 14772-1:1999, www.web3d.org/fs_specifications.htm

[12] blaxxun interactive, www.blaxxun.com

[13] Cosmo Software, Computer Associates International, Inc., www.cai.com/cosmo

[14] Pulse 3D, www.pulse3d.com

[15] Cycore Cult 3D, www.cult3d.com

[16] Mendel3D, Mendel 3D Avatar, www.mendelbox.com

[17] Shout 3D, Eyematic Interfaces Incorporated, http://www.shout3d.com/

[18] Weizenbaum, J., "ELIZA - A computer program for the study of natural language communication between man and machine", Communications of the ACM 9(1):36-45, 1966.

[19] eGain Assistant, eGain Communications Corp., www.egain.com

[20] ALife, Artificial Life, Inc., www.artificial-life.com

[21] vRep, NativeMinds, Inc., www.nativeminds.com

[22] WebFace, W Interactive SARL, www.winteractive.fr

[23] Saphire, Virtual Personalities Inc., www.verbot.com

[24] Ananova, Ananova Ltd., www.anannova.com

[25] "3D Graphics Define Virtual Humans on the Web", C. Babski, D. Thalmann, Software Focus 1(1), pp. 6-14, 2000.

[28] www.cynthiaknows.com

[29] Ask Sal, Excite Inc., arcadia.excitextreme.com/asksal

[30] Abraham Lincoln caricature using Pulse3D www.shererdigital.com

[31] Nodna, Nodna AG,www.nodna.com, demonstration: www.popkomm.de/popclash/pophead

[32] famous3D, famous3D Pty. Ltd., www.famous3d.com

[33] Javahead, Red Ted, www.redted.com

[34] ”A Web-Based MPEG-4 Facial Animation System”, I.S. Pandzic, Proc. ICAV 3D 2001 (to be published), demonstration at www.icg.isy.liu.se/~igor/MpegWeb

[35] The Interface project, IST-1999-10036, www.ist-interface.org

[36] A.L.I.C.E. natural language A.I. parser and chat robot, www.alicebot.org

[37] “From Photographs to Interactive Virtual Characters on the Web”, Igor S. Pandzic, Gael Sannier, Proc. Scanning 2000, Paris, France

[38] ISO/IEC 14496 - MPEG-4 International Standard, Moving Picture Experts Group, www.cselt.it/mpeg