Capture and animation of human-like figures

Translated in 2016-2017 by author from the Finnish language original from 2007

Juho Kunsola

Helsinki 2007-12-16
Bachelor’s Thesis
University of Helsinki
Department of Computer Science

Table of Contents

Abstract....3
Motivation....4
History recap....4
Capture of the appearance of a human-like....5
- Capture of geometry and textures....5
  - Stereo-camera capture....5
  - 3D scanning....5
  - Capture from a single photo....6
- Reflectance capture....7
- Models....10
- Synthesizing animations....10
- Physics-based generation of animation parameters....11
  - Acquisition of animation parameters with motion capture-based....11
  - Physics / motion capture hybrid....11
- Skin texture analysis, synthesis and rendering....11
- Case: Talking muscle-driven head....12
  - Introduction....12
  - Data acquisition....13
  - Animation and rendering....13
  - Synchronization of speech....14
  - Results and conclusions....14
- Case: The digital look-alikes of the Matrix sequels....15
  - Capture of the performance and the appearance of the actor....15
  - Capturing and simulation of reflectance....15
  - Rendering....15
- Conclusion....16
  - Image copyright acknowledgements....16
- Literature list....16

Abstract

The thesis attempts to give an overall view of different methods for capturing the appearance of a human into a computer model and the generation of animations by parametrizing the movement of the models from a human actor or fully automated.

Motivation

Computer-generated human-like images woo their watchers' be it avatars of 3D worlds or the synthetic heroes and villains of the silver screen. Science fiction movies have for ages been full of futuristic user interfaces, where computer has gotten the image and voice of a human. Surely talking heads are of concrete help too: they have applications for treating various cognitively or sensorily handicapped people and their daily lives. People who suffer from partial hearing loss can read from the lips of a talking head.

The techniques used in capturing human appearance also have applications in automatization of CCTV surveillance. Aging simulations are useful to the police when looking for people who disappeared a long time ago. Reconstructing human facial appearance from skull shapes is useful in forensic science. Texture analysis is useful in automatic detection of dermatological illnesses.

In the start of the 2000's human-like computer animations have been developed to such level of convincingness that they are ***digital look-alike***s, and distinguishing them from a real imaged human is often impossible. This opens up opportunities for evil besides just more and more visually impressive entertainment. Humankind has been split in three:

1) People who don't know about the existence of digital look-alikes. This ignorance is both a blessing and curse. They will not worry about such which they do not know to exist, but if someone of these is targeted with a digital look-alike attack it is highly likely that his/her own and their loved ones lives are in ruins soon with high probability.

2) People who know that these exist but they or their close circle lacks the technological means to produce these. These people will no longer swallow without chewing all they see in the Internet or hear of been seen there, but on the other hand can be very worried about the potential danger, that someone or someones could wield them as weapons against them and their loved ones in case they step on some wrong toes.

3) The digicratic elite, who have at their disposal the technology that enables making digital look-alikes. Vast possibilities are available to them to exert terror on other people and to pervert the human kind's view of the human kind by e.g. making available computer-generated "child porn" or other porn containing interactions that are not possible with traditional human actors.

History recap

In 1971 Henri Gouraud made the first digital human facial representation. The model for this image (https://interstices.info/jcms/c_25256/images-de-synthese-palme-de-la-longevite-pour-lombrage-de-gouraud) was his wife Sylvie Gouraud. We can determine that Gouraud used a manual camera capture setup i.e. took 2 photos of the face utilizing markers on the face. This photometric capture method where combining information from several images appearances can be captured fully automatically or semi-automatically is still going strong regardless that 3D scanners are available. Stereo camera- or even multi-camera setups are considerably less inexpensive than 3D scanners and the developments in the area of pattern recognition have lead to the situation where in many captures markers on the skin are no longer required but the facial form and gesture capture can be done by applying machine vision techniques.

In 1972 Frederik I. Parke in a project partially funded by DARPA made the first computer animation of human face resemblance by using Gourdaud's method where curving surfaces are approximated with polygons and also the shader made by Gouraud. [Par1972]

In 1981 Plaat and Badler built the first model that was not completely manually animatable but contained some logic within it. Their model contained some springs and masses to simulate the forces generated by muscles and it utilized the Facial Action Coding System developed by Ekman and Friesen in 1978. [Sif2005]

Capture of the appearance of a human-like

To capture someone's appearance at least 2 things must be captured: Topology and textures. To reach photo-realistic rendering additionally one needs to capture the reflectance of the target [PiL2005]. I will try to explain each method briefly in this chapter.

Particularly difficult targets are ears for their complex topology, hair because understandably every hair cannot be modeled[HaK2001]. Also teeth and tongue are problematic areas because they cannot be captured with conventional methods for understandable reasons.

Capture of geometry and textures

The geometry and textures can be captured with a stereo-camera that consists of more than one camera that are fired simultaneously, a 3D scanner, which is a depth finding instrument and even from a single picture by sculpting a sculptable 3D model onto it.

Stereo-camera capture

Photogrammetry means techniques where geometric properties of an object are determined from a photograph or photographs. When applying 2 or more digital cameras simultaneously it is possible to automatically find the geometrical information including depth data by utilizing photogrammetric algorithms.

Stereo-camera capture can be used both capture of shape and capture of movement.

3D scanning

Cyberware-scanners use laser-measurement to find the topology and a distinct measurement system for texture capture and are able to scan 30,000 points per second into X, Y, Z and R, G, B components. Cyberware scanners [as of 2007] reach only 400µm resolution for the Y-axle and 50-150µm for the Z-axle. Cyberware is an privately funded American business[Cyberw99].

Arius-3d scanners are multicolor laser scanners that capture both the color and the location of a point by sending 3 different wavelength laser beams in a centralized beam. Color can be determined from the measuring the intensities of the returning laser beams. The X-axle data is read by calculating the scanning mirror's position and and angle in relation to the scanning camera. The Y-axle is calculated from the movement data received from the camera movement system. The Z-axle coordinate i.e. distance from the object is accomplished by triangulation inside of the camera utilizing a galvanometer-guided dual sided mirror's both sides as follows: The first side reflects the beam to the surface of the object and the returning ray is directed onto the other side of the mirror whereby the angular difference when turning the mirror does not alter the point of impact in the

CCD but only the distance of the object being measured. Arius3d is Canadian privately owned company that has received funding from the Government of Canada in order to develop its scanner[arius3d]. Arius3d scanners can provide 100µm precision, which according to George Borshukov (page 106) is the size of the smallest visible objects in the human face, such as pores and thin wrinkles [HaD2004].

Capture from a single photo

Volker Blanz and Thomas Vetter have devised a sculptable head model by utilizing a large database of 3D scanned heads. Their system is able to produce a look-alike 3D model of a human head *from a single facial photo* (See image 1.) requiring only a modest amount of interactive work. First the user coarsely determines the position of the head in the photo. After this the system automatically matches the sculptable model to the photo in an “analysis by synthesis”-loop, wherein the differences between the original photo and the photo made from the synthesized texture mapped 3D model yield the parameters for the next iteration. The appearance of the modeled digital look-alike can be altered e.g. made thinner or fatter [BlV1999].

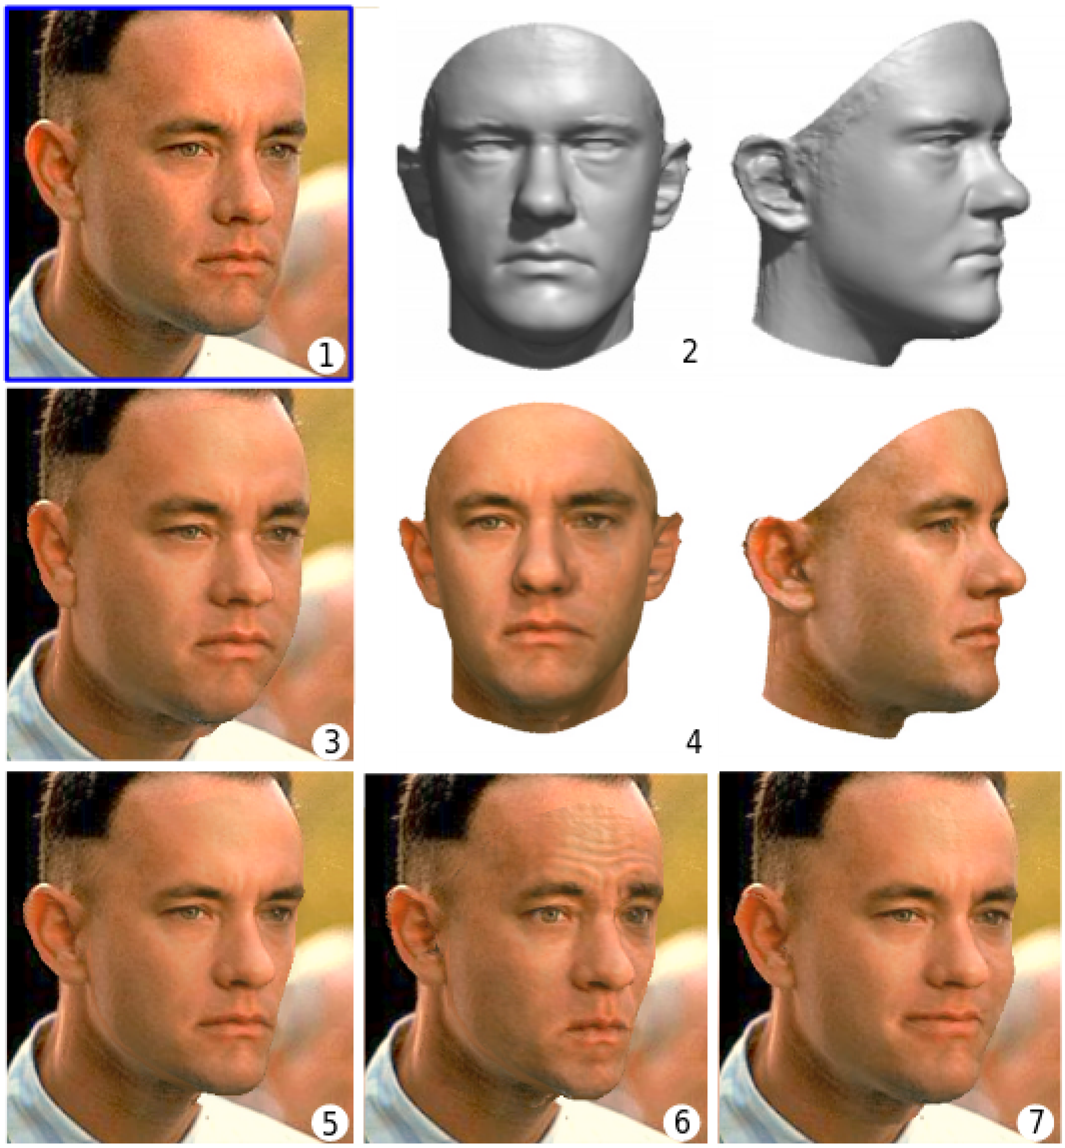

Image 1 (low resolution rip) **Sculpting a morphable model to one single picture** yields:

Sculpting a morphable model to one single picture (1)
Produces 3D approximation (2)
Texture capture (4)
The 3D model is rendered back to the image with weight gain (3)
With weight loss (5)
Looking annoyed (6)
Forced to smile (7)

Reflectance capture

Reflectance of surfaces is a property of material where the perceived color and brightness is dependent on both direction of lighting and direction of the watcher. It can be described with a bidirectional reflectance distribution function (BRDF) which describes the reflectability of a given point.

Bidirectional texture function is a 6-dimensional function that contains the plane coordinates, viewing and lighting angles. BTF is formed by taking thousands of photos of the object from different angles of illumination and observation. It is set of BRDFs over some area. (correction added 2017-01-01: BRDF is a set of BTFs over some area, not the other way around.)

For reflectance capture a light stage is utilized [Dev2000], that creates different illumination situations, based on which they can build an analytic BRDF, which enables rendering the 3D-model to virtual or real (sampled) environments and illumination conditions from any angle. Results are presented in image 3 [HaD2004].

The reflected light can be separated into specular light, which is reflected directly from the the fat of the skin and to the diffused light, that has traveled inside the skin before exiting, by applying a polarizer in front of the light source and the camera. Specular light that reflected straight off the surface retains its polarity meanwhile the light scattering from subsurfaces loses its polarization[Deb2000].

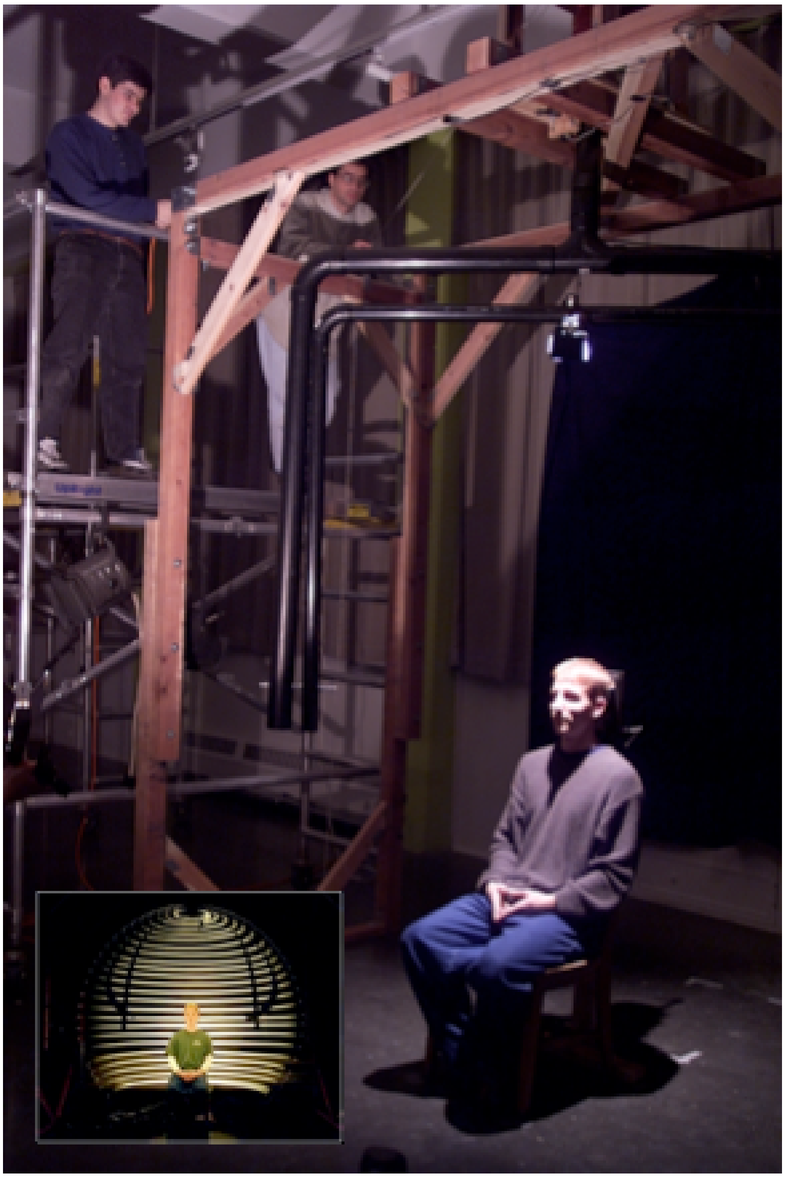

Image 2 (low res rip of the original): **Light stage**

Pictured is a light stage which is used to make measurements required for forming a BRDF. In the smaller picture is long exposure time photo showing the light stage in action.

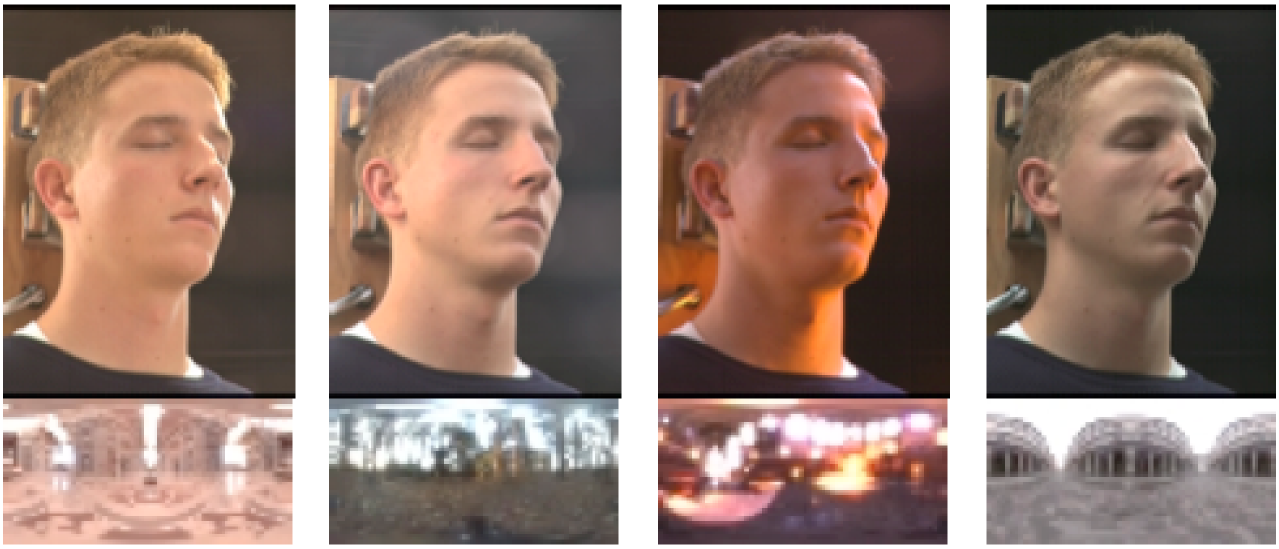

Image3: **Faces rendered in sampled lighting**

Each image images a face in synthesized lighting. The lower images represent the captured illumination map. The images are generated taking a dot product of each pixel's reflectance function with the illumination map.

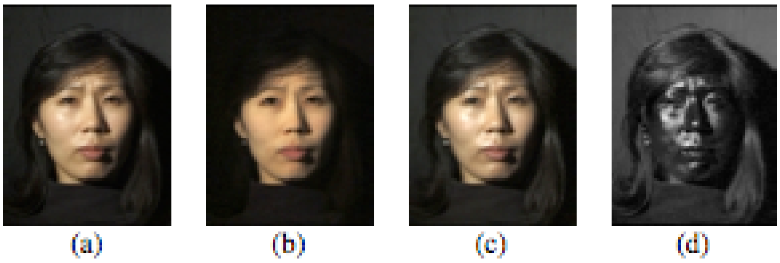

Image 4: **Separating specular and diffuse reflected light**

Normal image in dot lighting (a)

Image of the diffuse reflection which is caught by placing a vertical polarizer in front of the light source and a horizontal in the front the camera (b)

Image of the highlight specular reflection which is caught by placing both polarizers vertically (c)

Subtraction of c from b, which yields the specular component(d)

Images are scaled to seem to be the same luminosity.

Models

Gouraud's and Parke's models where polygon representations that contained only the surface i.e. "the skin". Modern models include a skeleton, muscles and other sub-dermal tissues and skin. The joints define the maximums of the movement of the bones and defining these empirically is simple and straightforward. Normal muscles (i.e. muscles that connect at 2 ends to the skeleton) can easily be approximated as ellipsoids, that thicken, shorten and straighten when muscles are flexed. For special cases like orbicularis oris (the muscles around the mouth) ja the tongue (connected only in one end) this does not obviously apply. When being tensed and relaxed the muscles affect other interconnected muscles. This interaction between muscles can be taken into account in the models by tying ideal strings between the muscles [HaK2001]. In anatomically correct models the muscles and other membranes can be modeled as tetrahedron volumes [Sif2005].

Synthesizing animations

The first animations of human-like faces were done with interpolation i.e. by manually defining location of each vertex in the key frames and the in-between frames were generated by interpolating from the key frames [Par1972].

After this people moved on to parametric animation, wherein the vertexes of the face were grouped into functional groups for which transformation operations were defined. These models can be driven with higher abstraction level commands such as tell the model to “smile” or say ‘A’. This reduces the amount of data transferral needing to animate the model.

In modern animation of human-like figures there are two different styles of producing the parameters, physics-based and motion capture-based, both integrate the results of the other style method.

Physics-based generation of animation parameters

With physics-based generation of animation parameters the aim is 100% automated systems with applications ranging from the avatars of virtual realities to the talking heads made for special needs. Interactions of virtual figures i.e. event based animation is outside of the scope of this work and I will focus on describing making visual lip-sync synthesis fully automatic by presenting of talking head in the section “Case: Talking muscle-driven head”.

In physics-based visual speech synthesis the skin and especially the deformations of the speech organs are accomplished by moving select facial muscles, chin bone and the tongue to reach the desired shape of the mouth. The upper teeth are naturally locked onto the skull’s position. Besides that lower teeth rotate around the jawbone axle they are also able to move forward, backward and sideways. The tongue can move with the jawbone but also independently[HaK2001].

Acquisition of animation parameters with motion capture-based

In sparse capture markers are placed on the face of the target and following the movement of those markers we get sparse data of geometry deformations by following the markers in a series of images.

In dense capture markers are not needed but the changes in geometry are acquired using optical flow.

Optical flow approximates movement of targets in a visual representation. In the real world physical movement does not necessarily manifest in color or gray scale changes, which complicates the implementation of the approximation. Often the movement is described as vectors that start or end in pixels in the digital series of pictures.

Physics / motion capture hybrid

Sifakis et al. presented in 2005 an approach that combines both physics-based and motion capture based generation of animation parameters[Sif2005]. This development is slightly frightening because a physics-based model that is parametrized / taught with motion capture data can be fed back to the motion capture based animation, which reduces the need for interactive work and reduces the costs involved in making photo-realistic or at least believable animations. For example in making the digital look-alikes for the Matrix sequels required manual labor to eliminate the accumulation errors in the optical flow could possibly be fully automated by using the hybrid approach.

Skin texture analysis, synthesis and rendering

Generating textures synthetically is a problem that has been researched extensively. When the aim is to produce textures that imitate human skin, the methods devised to produce other natural textures don’t necessarily work satisfactorily and methods that are specifically aimed at skin texture analysis, like Cula et al. image-based method for describing skin texture.

The skin is a complex topology and its modeling is difficult for various reasons. Reflection and traveling of light are affected by complex optical properties and also the micro-geometry of the skin; pores and wrinkles. When light hits the surface from front or from the side the skin looks different from same viewing point.

Cula mentions that skin simulation is a demanding challenge for yet do not exist methods that produce believable results, but he assumes that the approach that combines geometric transformations and image-based rendering could provide a satisfactory result[CuD2005].

The softness and translucency of skin and non-linearity of skin stretching and changes in tone caused by it pose vast challenges to reaching a believable result. The translucency of the skin adds to the challenge of rendering, because part of the light enters the skin and scatters back out,,[HaD2004]

In rendering static wrinkles and skin pores are described as a static bump map whereas dynamical wrinkles , that appear and disappear upon alteration of facial expressions, must be modeled and isolated from texture maps and layer on top of the bump map[HaD2004].

Bump mapping is a technique in computer graphics where the depth difference of each pixel is measured from a hight map and this data is taken into account when calculating the illumination of a pixel.

Recently the most impressive leaps in this have been done by the entertainment industry while the role of academia has been smaller. ESC Entertainment Ltd produced the digital look-alikes used in the Matrix-movies and Sony Pictures Imageworks was responsible for the digital look-alikes of the Spiderman sequels[Sag2004].

Case: Talking muscle-driven head

In this case I will go through the talking head article “Face to Face: From Real Humans to Realistic Facial Animation” by Haber et al. the head knows the phonemes and the transcript with timing data embedded to it and is able to produce visual speech[HaK2001.

Introduction

To model and animate a human face is very difficult because the human brain has been specialized to detect even the slightest facial gestures. The complexity of the face: skull, muscles and skin furthermore increase the challenge. Target’s facial geometry and textures must be acquired to a computer, select proper animation parameters and the deformations arising from these must be calculated and rendered. This forms the basis of forming facial gestures and speech movement formation.

The team presents the process step-by-step which principally goes like this:

1. Data (geometry and textures) are acquired and transformed into polygon representation.

2. Low level animation is based on the physics-based approach, where muscle contractions yield the animation parameters.

3. Building a virtual head is based on human anatomy and they have developed a tool with which the model can be sculpted and linked to the facial geometry.

4. Higher level animation synchronization produces the speaking movements.

Data acquisition

The method presented in the paper the geometry and the textures are acquired separately. Data is acquired with a triangulation-method using distance measurements. Distance measurements form a point cloud from which a surface-finding algorithm makes a polygon representation. When capturing the geometry the human model has his mouth closed and a neutral facial expression, so the mouth has to be cut open (in the polygon representation) to make a functional model. Additionally the parts of the surface that represent the eyes are flattened so that the model can be implanted with synthetic simulated eyes.

The research team noticed that satisfactory quality can be reached with a polygon network that consists of approx. 3000-4000 triangles, though it must be noted that critical areas like eyes and mouth require higher resolution mesh than other parts.

The textures are captured by taking 4-11 photos with a high resolution camera in a well dispersed lighting. The photos are registered and combined into an OpenGL-renderable model. It is not necessary to lock each node with own texture input (especially inside of the ears and behind them) but a suitable texture can be interpolated from surrounding nodes.

The team noticed that eyes, teeth and tongue are targets which are difficult to capture and they decided to use generic models of these body parts. For eye texture they made three pupil versions, blue, green and brown.

The generic eye-model offers the following animation possibilities:

1. Direction of sight

2. Size of pupil

3. Appearance of the eyelids

Position of the eyelids affects the brightness of the eyes.

Upper teeth are naturally locked onto the skull. In addition to the lower teeth rotating with the chin axis they are also able to protrude, retract and move sideways.

The tongue can move independently besides just moving with the jawbone moments.

In the physics-based simulation skin deformations are calculated from muscle contractions. For speech animation the team built a 12 muscle group of the most important muscles of the lower face by interactively modeling them.

Animation and rendering

Traditional muscles' (connected from 2 points) physiology was approximated as ellipsoids that thickens, shortens and straightens as the muscle is contracted.

Orbicularis oris (group of muscles around the mouth) was approximated as a ring. When talking, the shape of the mouth is mostly determined by muscles in and connected to the orbicularis oris. In the case of the mouth defining a central point of tension is not a workable solution, because it would enable only changing the diameter of the mouth and this is why they decided to use a tension axle that comes from the throat, between the teeth and out. This enables mouth form and shape manipulation (viewed from front) and additionally moving lips forwards and backwards.

Because orbicularis oris does not protrude homogenically, but the parts closest to the middle of the face protrude and contract more than the sides of the mouth (other facial muscles control the sides of the mouth) they used a protrusion gradient where these extremes protrusion/contraction was interpolated linearly.

Muscle interconnection in such a way that the tenseness state of another muscle and location affect was approximated by tying them to each other with ideal springs and by iterating these forces a satisfactory result was reached.

Orbicularis oris was divided into two parts concerning upper lip and lower lip to reach more powerful protrusion and contraction.

Head was made more natural by adding random eye blinking and random sine-wave rotation of the head in X, Y and Z axis

The system utilizes two processors effectively because the simulation and the rendering are done in two threads. Rendering gets the key frames from a ring buffer where simulation writes them. Then rendering interpolates the changes between the key frames.

Synchronization of speech

The most crucial faction in speech synchronization is co-articulation. This refers to that the shape of the mouth and the tongue in different combinations of phonemes may be radically different. For example in the words “molasses” and “middle” the starting letter ‘m’ causes a different starting mouth position as the letter ‘o’ doesn’t affect the pronunciation of ‘m’ but ‘i’ does. The affecting phonemes may be very far from each others and even cross word boundaries thus it is difficult to know where an acoustic segment starts and where it ends.

Speech synchronization requires a transcript that contains the phonemes and their timing data. Co-articulation is modeled as dominance function. They describe speech segment’s effect on vocal tract shape over time. The team’s physics based model they use contraction and protrusion parameters of 5 key muscles and jawbone rotation to model articulation. Each parameter is controlled by a distinct function that is calculated as weighed average of the dominance function over all segments and multiplied with target phonemes. Target for the muscle for a certain phoneme is acquired from the state of the muscles when a phoneme is orated distinct from other speech.

Results and conclusions

The team presented a system for modeling and animation of faces. Clearly most arduous work (4-7hrs) was forming a polygon mesh from the point cloud produced by the depth metering. The team is researching a generic head model and fitting it to the measured point data to reduce interactive work. The topologies of ears and insides of lips caused problems (bugs in appearance) but with approximation relatively good results were achieved. The team is confident that also hair could be incorporated into the system by finding the geometry and textures from photographs.

In the context of speech synchronization their method was straightforward and once the parameters of the dominance function had been defined ja each muscles' target contraction had been interactively defined for each phoneme the key frames could be calculated automatically.

Case: The digital look-alikes of the Matrix sequels

Matrix-movies sequels required photo-realistic digital look-alikes of known actors and the team lead by George Borshukov developed for this a method they call "universal capture". It combines two very powerful machine vision techniques: optical flow and photogrammetry.

Capture of the performance and the appearance of the actor

The action was captured with 5 synchronized cameras in well diffused lighting. After this each pixel's (sic.) movement over time in each camera was measured using optical flow and the product of this process was combined into a neutral expression 3D model of the actor and to the camera positions photogrammetric reconstruction. The algorithm works by projecting each vertex of the model to each camera vie and then tracking this vertex movement in 2D views with optical flow and for every vertex the 3D location is approximated by triangulation.

In optical flow the errors can accumulate over time causing unwanted drifting in the 3D reconstruction. In this production to eliminate the errors manual key frame adjustment where geometric errors are corrected manually when the error becomes annoyingly bad and the corrections are fed back to previous frames using reverse optical flow.

Believable facial rendering cannot be reached without taking into account that the texture changes with facial expressions. Since the team did not use markers attached to the faces they were able to combine views from multiple cameras over time and thus able to produce a seamless animated UV-map, from which they could isolate essential textural changes like fine and bigger wrinkles and changes in color caused by stretching skin. UV-mapping is used to describe a 3D object onto a texture map.

Even as the animation isolated from the actor's performance included most movement nuances it was missing tiny details of miniature scale like pores and wrinkles. Dynamical wrinkles were identified and isolated from the texture maps and layered onto the static bump map.

Capturing and simulation of reflectance

In making the film the facial BRDF was isolated based on images. Actors were photographed with many simultaneous cameras in various different illumination conditions utilizing a light stage. The capture set-up was carefully color calibrated and camera locations and head's location were reconstructed with photogrammetric techniques. The acquired footage was brought to common UV-space by re-projecting it to a cyberscan-model. Using this space the team could analyze reflectance properties in various lighting conditions ja light exit angle and armed with this data the team isolated parameters for an approximal analytic BRDF, which was composed of a lambertian diffuse component and modified Phong-type of specular component, which had a Fresnel-type of effect.

Rendering

As the production went on it became clear that to realistically render "skin" without subsurface scattering simulation in not possible. The existing methods for rendering translucent materials were deemed too demanding processor-wise and unsatisfactory in the end product so the team devised their own method for producing subsurface scattering which was easy to implement. Diffused light in the direction of the camera was stored in a 2D light map and then approximally simulate the dispersion of the light in the picture. To improve the result they adjusted the diffusion parameters for different wavelengths. For targets where light can pass right through, e.g. ears, they used a more traditional ray tracing approach to provide the wanted translucency effect [HaD2004]

Conclusion

The thesis presented capture of human-like figures into computer models and animating those models with different methods and for different uses. The industry has moved from clearly machine-like figures to digital look-alikes that are extremely hard to tell apart from real images of real people. Development in the field has mostly moved to the commercial side leaving academia in the sidelines, which has caused that getting information of the new technologies and techniques becomes difficult as they tend to be business secrets.

The thesis brought to view the current situation where even from a single photograph it is possible to make a facial model almost without human help. From the point of view of Computer Science these things were addressed on a non-advanced level, but bringing these things into the popular knowledge is an important issue so that the destruction and constant threat of it and detoriation of good manners could one day be on the agenda of the legislators or even public discussion. Handling the situation as an issue requiring jurisprudence would be most welcome.

In theological examination it is very difficult to not notice the various conjunctions and similarities of the beasts in the Book of Revelations (written over 1800 years ago) and the digital look-alikes. Division of the human kind to over-humans, sub-humans and somewhere between there is a pressing issue for philosophers, psychologists and social scientists too.

Image copyright acknowledgements

Literature list

arius3d	Arius3d, Adrius3d digital imaging solutions – 3d color laser scanner http://www.arius3d.com/main.html?contentId=1 [Accessed 2007-12-16, January 2017 it is a redirect to 404 on http://www.ariustechnology.com/]

BlV1999	Blanz, V. and Vetter, T., A morphable model for the synthesis of 3d faces. *SIGGRAPH'99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, New York, NY, USA, 1999, ACM Press/Addison-Wesley Publishing Co., pages 187-194
Cyberw1999	Cyberware, Head faces scanner specifications http://cyberware.com/products/scanners/pxSpecs.html [Accessed 2007-12-16 and again on 2017-01-01]
Deb2000	Debevec, P., Hawkins T., Tchou, C., Duiker, H-P., Sarokin, W and Sagar, M., Acquiring the reflectance field of a human face. SIGGRAPH '00: *Proceedings on the 27th annual conference on Computer graphics and interactive techniques*, New York, NY, USA, 2000, ACM Press/Addison-Wesley Publishing Co., pages 145-156
HaK2001	Haber, J., Kähler, K., Albrecht, I., Yamauchi, H and Seidel, H-P., Face to Face: From Real Humans to Realistic Facial Animation. *Proceedings of the 3rd Israel-Korea Binanational Conference on Geometrical Modeling and Computer Grahphics*, Seoul, Korea, October 2001, Kyung Moon, pages 73-82
HaD2004	Haber, J. and Terzopoulous, D., Facial modeling and animation. *SIGGRAPH '04 ACM SIGGRAPH Course Notes*, New York, NY, USA, 2004, ACM Press, page 6
CuD2005	Cula, Oana G., Dana, Kirstin J., F.P.M. and Rao, B. K., Skin texture modeling. *International Journal of Computer Vision*, 62,1 (2005)
Par1972	Parke, F. I.., Computer generated animation of faces. *ACM '72: Proceedings of the ACM annual conference*, New York, NY, USA, 1972, ACM Press, pages 451-457
PiL2005	Pighin, F. and Lewis, J.., Introduction *SIGGRAPH '05 ACM SIGGRAPH 2005 Courses*, New York, NY, USA, 2005 ACM, page 1