AgeRegression: Rejuvenating 3D-Facial Scans

Katharina Legde
BTU
Cottbus-Senftenberg
Platz der Deutschen
Einheit 1,
03046 Cottbus, Germany
Legdekat@b-tu.de

Susana Castillo
BTU
Cottbus-Senftenberg
Platz der Deutschen
Einheit 1,
03046 Cottbus, Germany
castillo@b-tu.de

Douglas W. Cunningham
BTU
Cottbus-Senftenberg
Platz der Deutschen
Einheit 1,
03046 Cottbus, Germany
cunningham@b-tu.de

ABSTRACT

The majority of virtual agents have adult bodies. There are, however, a number of reasons for using younger avatars. For example, an adult interface agent usually leads users to expect adult-level communicational and social skills. As a result, users tend to be rather intolerant when the interface agent makes obvious mistakes (e.g., incorrect grammar) or uses inappropriate behavior (e.g., looking away from the interlocutor). Since computer social-skills are still under-developed, it seems reasonable to use a body model that reflects this: child avatars. Unfortunately, the use of database-driven techniques for creating a variable-aged animation system would require a very large number of scans of children at different ages, making such a system impractical for technical and ethical reasons. As an alternative, this paper develops and validates a method for synthetically and systematically altering the apparent age of a virtual character. The here proposed technique is able to create younger and older versions of a facial scan and guarantees that the resulting meshes can be animated. Starting with a three-dimensional, adult facial scan, we use a physiologically-inspired, trigonometric polynomial to age-regress the model to a desired age. Quantitative measurements show that the technique can reconstruct the correct anthropometric proportions of 2-10 year-old children. A perceptual experiment provides an initial mapping of the technique's parameters onto the perceived age and realism.

Keywords

Age Regression, Cardioidal Transformation, Face Models

1 INTRODUCTION

Human faces convey an impressive amount of information. We intuitively and automatically detect static facial cues like gender, identity, mental state, health, and age [Cah90, Ram09, Fis16] and use them to determine how to interact with an interlocutor (e.g., [Rya86, Gle75, Yar09]). Synthesizing the human face is especially challenging. Still, there are a wide range of applications using an artificial human face including advertisement, entertainment, education and user interfaces. User interfaces are supposed to ease the communication in between the user and the machine with the help of e.g. an intuitive graphical user interface, speech recognition or with a virtual agent. Virtual agents are granting the computer human-like communication and social abilities (for examples see [Kop05, Kro17, Mor15, Nie09]). In most cases they are represented with human adult bodies. Especially when the computer "looks" like an adult we expect it to also act like an adult, including the use of appropriate social behavior, proper grammar, intonation, and body language [Vin06, McD12]. We also expect the virtual agent to understand what we say at an adult level. These expectations are often not met. This conflict between appearance and behavior usually leads – at the very least – to a frustrated and confused user [Vin06, McD12] who no longer wishes to interact with the system.

This paper combines the suggestion that the appearance of a virtual agent should match its behavior, the underlying task, and the system constraints [Kru12] with the observation that we adjust our behavior (including what we say and how we say it) and our expectations to the other person's age [Gle75, Rya86, Yar09]. Since people tend to be very tolerant to social and communication mistakes made by children as well as to their constrained level of general knowledge we propose that

virtual agents that look like children can improve the interactions with computers.

Creating a virtual child poses a number of technical and ethical difficulties. Among the core concerns is that modern performance-driven animation techniques require a large number of three-dimensional (3D) facial scans and motion capture information (e.g., for different expressions). Furthermore, if the interface is to be able to provide a range of apparent ages, then scans and recordings from many ages will have to be made. Long recording sessions of carefully presented facial expressions is already difficult with adults. With small children, such recording sessions are nearly impossible. These technical concerns (and the related ethical concerns of using a real-child's face as an interface agent) can largely be side-stepped by creating and animating an adult face, and then rejuvenating it to the desired age.

The age-rejuvenation technique proposed in this paper uses a trigonometric polynomial inspired by Todd et al.'s revised Cardioidal Transformation model [Tod80]. Unlike the original technique, the proposed polynomial works on the 3D facial structure, allows different local facial areas to grow at different rates, and can be used for age-regression as well as age-progression. After describing the technique, we show a range of possible results and then map the parameters of the polynomial to perceived age (and realism) through a perceptual experiment. Finally, we then compare the anatomical proportions of age-regressed faces to those found in real child pictures of the regressed individual.

2 AGE SYNTHESIS

Age synthesis describes a re-rendering of a face image or a face mesh with an approximation of natural aging or rejuvenating effects.

2.1 Natural aging

The process of natural aging causes significant, idiosyncratic changes in a person's face. The appearance of a person's face at any given age is highly dependent on both inner and outer factors [Enl89, Alb11]. Inner factors are mostly biological. For example, during aging soft tissue looses its elasticity and volume, fatty tissue gets lost, and the face increasingly resembles the bony structure of the skull [Par08, Alb11]. Other inner factors include biological sex and ethnic group [Alb11]. Biological sex determines the intensity and duration of the actual aging process and influences the growth of facial bones, the height of cheekbones, the width of the nose and the existence of facial hair [Enl89]. Outer factors, on the other hand, generally describe changes arising from environmental causes or lifestyle-based behaviors and refer to things such as weight, scars, wrinkles due to frequently used facial expressions, effects of the physical environment (such as frequent exposure to the sun), and the general life experiences of the person [Alb11]. The particular combination of inner and outer factors that occur for a given individual contribute strongly to the individuality of each human face and must be considered when age synthesis is preformed.

Aging is not a linear process. The changes that occur to adults are fundamentally different from those that happen during childhood [Suo07]. During childhood, skin changes are minor while the cranium undergoes various and rapid growth modifications. In adult aging, in contrast, the shape of the cranium changes very little while skin changes can be considerable, including changes due to sunlight and gravity [Pit75, Suo07, Alb11].

2.2 Age synthesis techniques

Age synthesis can be classified into the three groups: explicit mechanical, implicit statistical, and explicit data-driven [Fu10].

Explicit mechanical synthesis: Explicit mechanical synthesis techniques describe changes of skin and other facial tissue during growth, including the synthesis of wrinkles [Sha04, Fu10]. Wu et al. presented a plastic-viscoelastic model, which simulates dynamic wrinkles [Wu94]. In their technique, the face is represented with three different layers: muscles, connective tissue, and skin. During the contraction of the muscles, the connective tissue – which is modeled with hookian springs – simulates the elastic process and thereby pulls the skin in a specific direction. The main use of this model is to create wrinkles for facial expressions, but Wu et al. [Wu94] also state that changing certain parameters such as the stiffness of the spring and the thickness of the connective tissue will result in wrinkles caused by aging.

Implicit statistical synthesis: By far the most common class of technique is implicit statistical synthesis. This approach does not use physics, but instead tries to capture age-related changes by performing an automated, statistical evaluation of the appearance variations of facial images or 3D facial scans. These techniques require the acquisition of a large training dataset at a wide range of different ages. Each image or mesh in this training set is considered to be a high-dimensional point. Scherbaum et al. [Sch07] used a Morphable Model to extract shape and texture variations in faces in such a dataset. With the help of a non-linear Support Vector Regression performed on several shape and texture coefficients, they are able to learn a function which maps every face in their database to a scalar age value. Calculating the gradient of the age function allows Scherbaum et al. [Sch07] to extract aging trajectories for shape and texture and thereby alter the apparent age of 3D facial scans.

 ISBN 978-80-86943-41-1

Another representative approach is the Multi-Resolution Dynamic Model presented by Sou et al. [Suo07]. Given an image as input, they first build a graph structure and then sample it over various age groups according to a beforehand learned dynamic model through a Markov process. This allows them not only to capture and afterwards simulate aging effects relatively well, but also to simulate *variations* in the aging process [Suo07].

Golovinskiy et al. [Gol06] proposed a statistical face model to extract and transfer facial details, like wrinkles or pores. They use high resolution facial scans and split them into a smooth base mesh and a detailed displacement image. Statistics are performed to capture the local orientations and the amount of detail. These statistics can now be used in combination with a displacement image of another facial scan. Thereby they are able to transfer wrinkles of one mesh to another, to make the mesh seem older or smooth the wrinkles, to make the mesh seem younger.

Explicit, data-driven synthesis: Explicit, data-driven synthesis techniques use empirical, mathematical models of human growth to describe the shape changes of the head. Such techniques represent a human face as a 3D-mesh and then alter the mesh with geometrical scaling, translation, and rotation operations. These techniques use explicit models of anatomical changes of the aging process. One of the most well-known models of age progression is the **Craniofacial Growth Model** [Tho17], which simulates the aging process by modeling shape changes caused by craniofacial growth. The model is based on the observational work of D'Arcy Thompson, who believed that the morphological changes of the aging process can be reduced to simple geometrical transformations [Tho17, Tod80]. Pittenger and Shaw asserted that growth was a series of visco-elastic events [Pit75]. They observed that during aging the face grows more rapidly than the rest of the cranium which results in changes in the face shape profiles, the so called the facial angle [Tod80]. Todd et al. mathematically modeled those events with affine shear and cardioidal strain transformations (CST) [Tod80, Fu10]. With the help of experiments, they found that affine shear transformations have less effect on the perceived age of the facial profiles and produce unidentifiable distortions. They subsequently focused on the cardioidal strain transformation, which had a much larger effect on the perceived age [Tod80]. Specifically, they conceptualized the facial profile as a cross-section of a head that is modeled as a sphere filled with fluid. A transformation on a certain point depends on the force which is exerted on that point, which is derived from a function considering gravity, the density of the fluid and the radius of the sphere [Tod80, Ram06]. They assumed that the pressure is directed radially outwards, distributed continuously and symmetrically along the vertical axis [Tod80, Ram06]. Todd et al. revised that approach with the observation that the amount of pressure at any point is depending on the amount of fluid above it and therefore added a term of position to the original CST [Tod80]. The revised cardioidal strain transformation (rCST) was defined mathematically – for 2D facial profiles – in polar coordinates for (R,θ) with the following equation [Tod80]:

$$R' = R(1 + k \cdot (1 - cos(\theta)) \tag{1}$$

with R being the distance of a given point to the origin before transformation and R' being the distance of that point to the origin after transformation. The polar angle θ is the angle between the line segment R and the vertical or polar axis. After the transformation is applied the polar angle stays the same $\theta' = \theta$. The constant k is a growth-related constant which increases with age [Tod80, Ram06], thereby this equation is defined as an age progression. Mark et al. later extended this equation to work with 3D faces with the help of spherical coordinates [Mar83].

3 PROPOSED AGE REGRESSION

The proposed algorithm attempts to rejuvenate three-dimensional facial models using a 3D trigonometric polynomial (Equation 2) inspired by Todd et al.'s revised CST model (Equation 1). Specifically, we converted the equation to spherical coordinates, made it additive, altered the terms that involve the angle θ so that they are always positive and between [0,1], and added a number of terms to allow changes to specific facial areas (i.e., the nose bridge *nb*, the tip of the nose *nt*, and the cheeks *ch*). We also allow changes along the latitude (azimuth or φ) as well as the longitude (inclination or θ) of facial scans. Finally, all terms in the equation now use exponential functions to focus their changes on different areas. By having different exponents for the lateral and the longitudinal cosines, the new terms can – for example – be tightly focused horizontally but loosely focused vertically (such as would be needed to capture nose-related changes). The new equation is:

$$\begin{aligned} R_{yng} = R_{old} \cdot (1 &+ k_h \cdot (\frac{\cos(\theta)+1}{2})^{n_h} \\ &+ k_{nb} \cdot (\frac{\cos(\theta-\alpha_{nb})+1}{2})^{n_{nb}} \cdot |\cos(\frac{\varphi}{2})^{m_{nb}}| \\ &+ k_{nt} \cdot (\frac{\cos(\theta-\alpha_{nt})+1}{2})^{n_{nt}} \cdot |\cos(\frac{\varphi}{2})^{m_{nt}}| \\ &+ k_{ch} \cdot (\frac{\cos(\theta-\alpha_{ch})+1}{2})^{n_{ch}} \cdot |\cos(\frac{\varphi}{2})^{m_{ch}}|) \end{aligned} \tag{2}$$

with as is standard for spherical coordinates $\theta \in [0,\pi]$ and $\varphi \in [0,2\pi]$. The angle θ is between R_{old} and the vertical polar axis while φ defines changes along the orthogonal plane. R_{old} is the distance of one point in the adult facial mesh to its local origin (center of mass).

 ISBN 978-80-86943-41-1

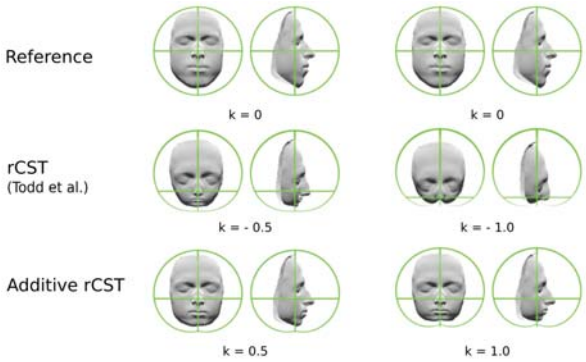

Figure 1: Age Regression equation corresponding to Todd et al. and additive rCST. For illustration purposes the same deformation is applied to the unit circle shown in green.

The age coefficients k_i control how much deformation is applied to distinct areas i. The exponents n_i, m_i control the tightness of the transformation for the distinct areas i. The angles α_i define a shift in the starting point θ of the transformation. The full equation is applied to every vertex of the 3D scan of an adult face.

It should be noted that Todd et al.'s rCST (Equation 1) can already be used to perform an age regression [Mar83]. While small amounts of rejuvenation (using negative values of k close to zero) can produce plausible results, higher rejuvenation effects produce increasing distortions (Figure 1, second row). Note that merely altering the sign of the first term (by adding it and ensuring that the θ terms are in the range [0,1]) already improves the results (Figure 1: third row, only considering the term for the head h in Equation 2). Since the effect of such an additive, positive rCST is still global, it cannot account for the different growth rates of different facial areas (as can be seen, e.g., with the overly-large nose).

To simulate the fact that different facial regions grow at different rates, we added three new terms (affecting the nose bridge, nose tip, and cheeks) based upon anatomical considerations of the differential growth rates (see, e.g., [Ram06]). Initial experiments (see Section 5) confirm that these four terms (the three new ones and the cranium) represent a decent basis model. The angles α_i which assure the focus of the transformation on a certain area are unique for different 3D scans. Note that the angle φ was used in all areas except for the head term h, since its inclusion there leads to unnatural cranial forms. Also note that we use the half angle $0.5 \cdot \varphi$ to ensure that the changes for nb, nt, ch happen only to the face (and not to the back of the head). The aging coefficients $k_h, k_{nb}, k_{nt}, k_{ch}$ are chosen to match biological constraints. During childhood the cranium (the first term in Equation 2) undergoes the most growth modifications [Suo07], thereby the age coefficient k_h should be altered more than, for example, the age coefficient

Figure 2: Transformation for different facial areas with varying parameters in absolute value. Only one area (columns) is transformed at a time. For illustration purposes, the same deformation is also applied to the unit circle shown in green.

for the nose bridge k_{nb} or nose tip k_{nt}. We also defined different exponents n, m for each of the four areas to control the tightness of focus of the cosine term. A rejuvenation for the nose tip nt, for example, needs to affect a much smaller area than the transformation of the cheeks ch. For our reference face, the effect of the different parameters for the four terms are illustrated in Figure 2.

4 RESULTS

The proposed technique can produce a very wide range of facial deformations. Careful choice of the parameter values can reduce the apparent age of a 3D adult facial scan to any desired age. Other parameter values, however, can produce unnatural facial meshes. A small sample of possible results considering a subjectively chosen parameter space can be seen in Figure 4. We chose to first linearly interpolate the parameters k_i, n_i and m_i in a proposed interval to rejuvenate the 3D scan of an adult. The faces in the middle of the parameter range clearly resemble children of different ages. Extreme values for k_i, n_i and m_i decrease the degree to which the face appears human-like. Figure 4 also makes it clear that to control the apparent age of

ISBN 978-80-86943-41-1

a face, more than the age coefficient k needs to be considered. In particular, the size of the exponents can play an important role, consistent with Ramanathan et al.'s [Ram06] claim that different areas of the face grow at different rates. A closer glance at the figure shows that, for example, for coefficients $k_h = 0.6$ through $k_h = 0.8$ (sixth through eighth row) with low exponents of n_i and m_i (left side of the spectrum) will produce almost baby-like faces. Higher values for n_i and m_i (right side of the spectrum) will produce older children. It is important to note that the coefficients k_i can also be used for reducing or expanding a region, e.g. using a negative coefficient can help to regulate the nose tip better.

We also applied our method to commercially purchased head scans [Ten24] of different genders and different ethnic groups. Some example results can be seen in Figure 3 for the evaluated age groups and their corresponding parameters from the proportional study (see Section 5.2). Figure 3 shows that we obtain realistic, child-like faces. It is important to note that even for very different facial meshes, the exact same parameter values for k_i,n_i, m_i (see Figure 4) produced similar amounts of age regression. The local origin of both scans is placed in between both eyes in frontal view and in between the eye and the ear from side view (approximately the center of mass of the previous scan). We chose to also use the same angles α_{nb}, α_{nt} and α_{ch} defined for our previous scan to specify the areas in our new scans. Figure 3 shows that we produce with the same angles reasonable child-like faces. Obviously, a manual adjustment of these angles for the new meshes or even an automatic detection would produce even better rejuvenation results. Note here that the inclusion of the neck in the 3D-scan is important and makes a difference in the perception of age (especially the Adam's apple in the male scan). For further results of the full suggested parameter space from Figure 4 for both of the models please see supplemental material.

The rejuvenated meshes can be animated in a number of ways. The simplest method is to age regress a neutral model and use cluster animation based on motion capture data to move the face [Par08]. One can also use blend-shape animation, where a number scans of the same person in different peak expressions are placed into correspondence [Par08]. Novel expressions and animations are then produced by a weighted sum of the different scans. Note that if the correspondence is established before age-regression, then only the neutral expression needs to be regressed since the other scans in a blend-shape animation system are stored as a deformation of the neutral expression. Initial tests of both cluster-based and blend-shape-based animation produce realistic results. Note that in these animations, the motion of an adult was projected on the face of a child. This technique, then will allow future work to determine the relative contribution of dynamic (e.g., motion) and static (e.g., geometry) information to the perception of age.

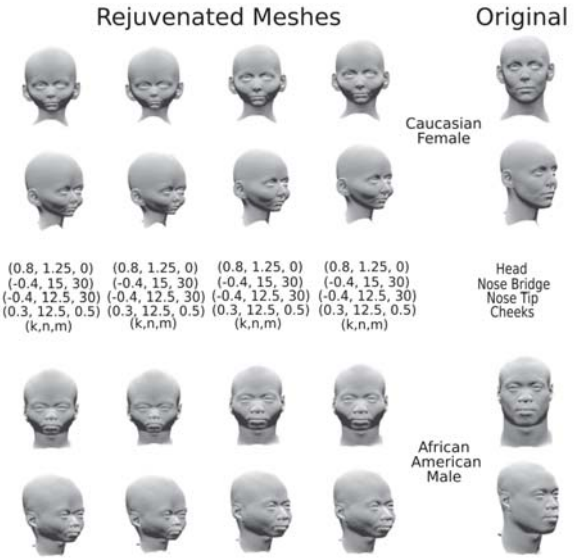

Figure 3: Example results of the proposed technique considering different head models. Values for k_i, n_i, m_i taken from suggested parameter space (see Figure 4). Chosen subset is consistent with proportional study (see Section 5.2).

Additionally, our proposed technique can also be used to make faces older again by simply dividing R by the aging terms instead of multiplying it (see Equation 2). To test this, we first age-regressed an adult face and then progressed it back to the adult face without changing the parameters for k_i, n_i, m_i. We were able to recover reasonable adult faces including the original scan. To use the technique to age-progress an face from young adult through middle-age into old age, an adjustment of the parameter space is needed. During the later stages of aging, the cranium grows less than for example the nose [Ram09], this fact can be modeled by altering the k_i for these areas. Also an adjustment of the k_{ch}, n_{ch}, m_{ch} could be useful to synthesize the loss of fatty tissue.

5 EVALUATION

To evaluate the effectiveness of the technique, and to provide a first approximation of which parameter combinations can yield specific ages, we performed a perceptual experiment and a proportional study using all of the rejuvenated models shown in Figure 4.

5.1 Perceptual Experiment

The perceptual experiment presented a set of stimuli generated by our equation and asked participants to rate the apparent age of the head and how natural or "human-like" the head was.

Methods: Any attempt to systemically test the effect of all 11 parameters of the equation and their combinations would require a prohibitively large number of

ISBN 978-80-86943-41-1

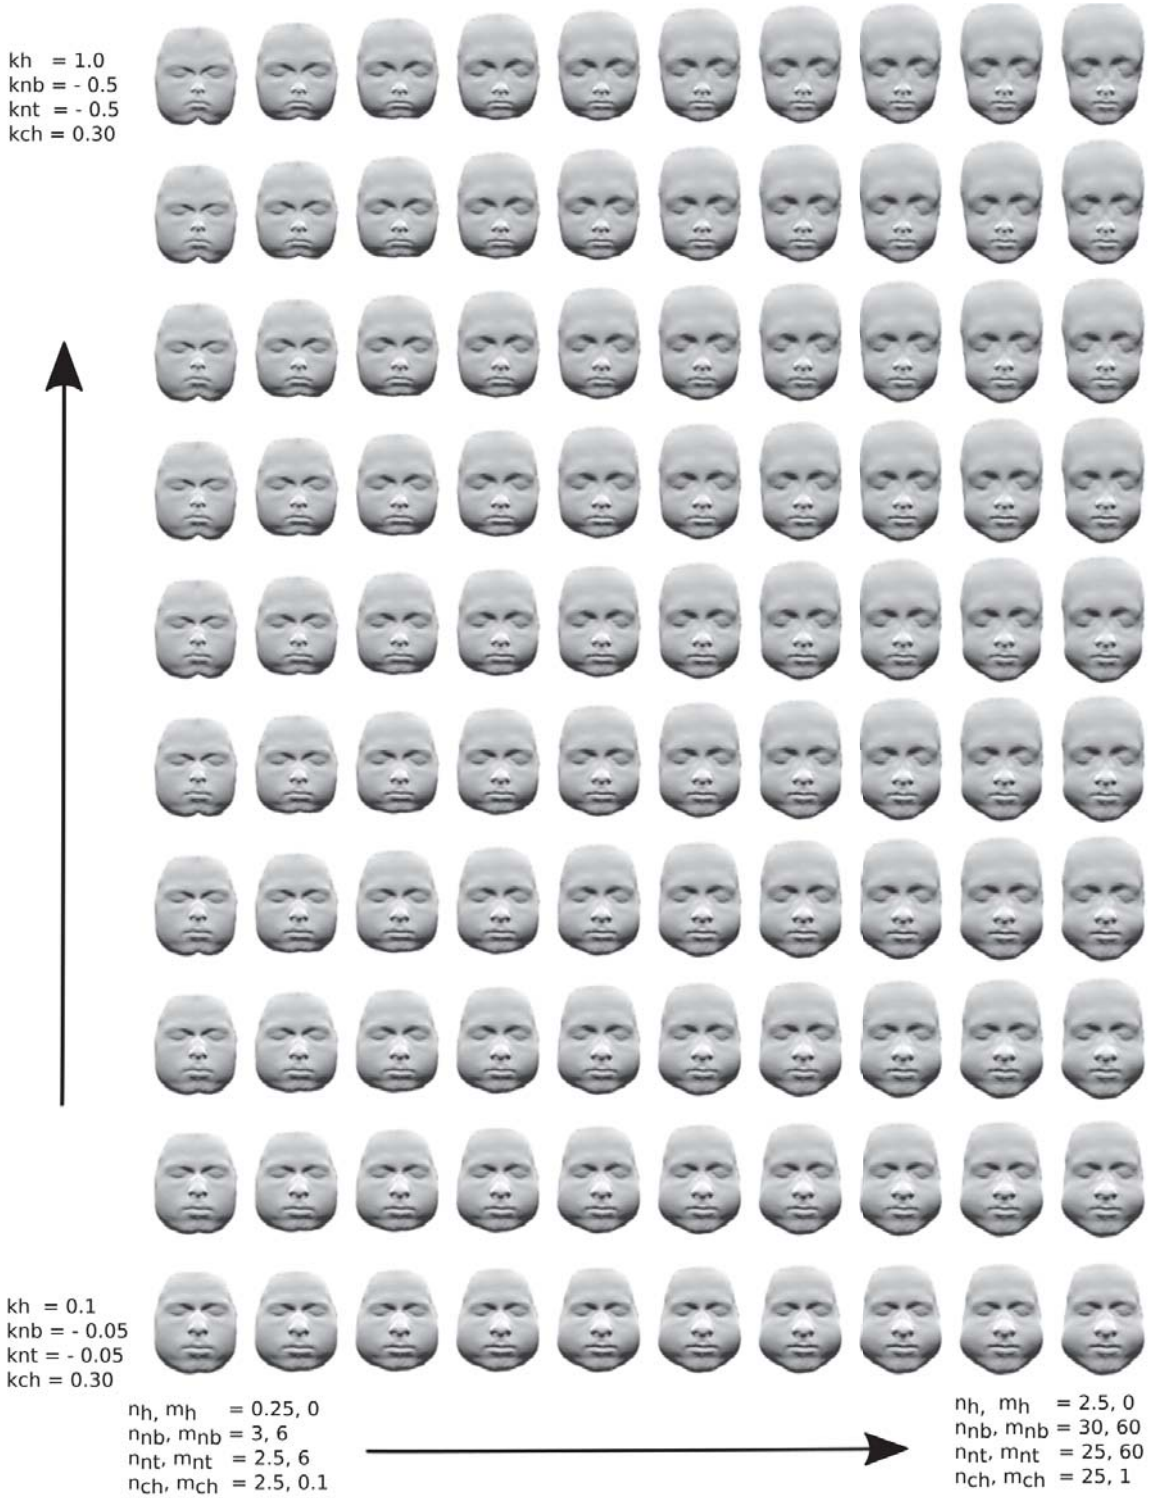

Figure 4: A sample of results of the proposed age regression technique. Parameter values for k_h, k_{nb} and k_{nt} as well as the exponents n, m for the specific areas are given in the picture. For initial testing k_{ch} was held constant with a value of 0.30.

ISBN 978-80-86943-41-1

trials. Therefore, we decided to use only a subset of the possible combinations. Specifically, we set the age coefficients k_h,k_{nb} and k_{nt} to different initial values and then linearly modified these values. Note the age coefficient k_{ch} was held constant for the experiment. All age coefficients co-varied perfectly (any change in one was accompanied by an equal change in the other). Each of the exponents n and m also had a different initial value, but again all co-varied perfectly. The combination of 10 head/nose-bridge/nose-tip coefficients with 10 exponents yielded 100 reconstructed faces (see Figure 4). The stimulus set also included the original facial scan. All stimuli were rendered without texture and under the same lighting conditions. On any given trial, one 3D face was shown both a frontal view and with a 60 degree angle to the camera, rendered with orthographic projection in front of a black background.

Ten people (4 women and 6 men; between 24 and 38 years old) participated in the experiment for financial compensation at standard rates. They were naïve to the purpose of the experiment. All participants provided informed consent. Each participant sat alone in front of a computer monitor in a darkened room and was given the experimental instructions. They were then shown all 101 stimuli in random order, with each participant receiving a different order. First, each participant had to rate the apparent age using a set of different age groups: $<1, 1-2, 2-4, 4-6, 6-8, 8-10, >10$ years old. Second, they had to rate how natural or human-like the head was using a 7-point Likert scale (with 7 being extremely likely to be human and 1 being extremely unlikely to be human). To help anchor the naturalness scale, the real 3D scan of the reference person was shown and the participants were told that this was a real person and should be given a value of 7. The average time for one experiment was 30 min.

Results: Overall, the age regression technique significantly and systematically modified the apparent age of a facial scans. Moreover, nearly all of the results were seen as extremely human-like (see Figure 5c). Since the two head coefficients co-varied perfectly as did all of the exponents, there were effectively two factors: age coefficient and exponent. We performed a two-way, repeated measures ANOVA with both of these as within-participants factors. Both main effects and the interaction are significant ($F(9,81) = 70.23, p < 0.001$; $F(9,81) = 6.229, p < 0.001$; $F(81,729) = 1.63, p < 0.001$ for the age coefficient, exponent, and interaction, respectively). Further analysis showed that the age coefficient and the perceived age are significantly (and negatively) correlated ($r = -0.6918, p < 0.001$; see Figure 5a). If we take a look at the correlation of perceived naturalness and the age coefficient k (see Figure 5b) we see that we produce very natural looking rejuvenated faces (rated above 4 on a 7-point Likert scale) especially in the middle of our spectrum (third till eighth row). As mentioned above, extreme parameter values (bottom and top row) tends to produce somewhat unnatural faces.

Modifying the exponents helped to produce younger faces but also had a distinct effect on the perceived naturalness of the shown faces (see Figure 5d and e). Closer examination shows a small but significant correlation of exponent with the perceived age ($r = -0.1554893, p < 0.001$). As can be seen in Figure 5d, the low correlation is probably due to the fact that their relationship is not linear. Small exponents helped a little to reconstruct younger faces. Larger exponents, however, did not increase the rejuvenation. This effect might be explained with the nose in the profile view of our stimuli set. As the age coefficient of the head k_h mainly stretches the full face, the nose also gets longer and thereby less child-like. As can be seen in Figure 5e, most of the exponents produced natural-looking faces, except for extreme values of our parameter space. In sum, the initial parameter ranges used for the experiment were surprisingly effective, although clearly some fine-tuning on the specific values is needed.

The perceived age was also slightly but significantly correlated with the human-ness of the reconstructed faces (correlation of $r = 0.125098, p < 0.001$). The median (represented as rot dots) and the density of the violins emphasize the fact that almost all of our reconstructed faces were seen as very likely to be human. Note that the distribution of results is non-Gaussian, leading to the assumption that the extreme parameter values are outside the ideal range.

The participants rated the original facial scan older than 10 and on average with 6.4 on the 7-point Likert scale of naturalness.

5.2 Proportional Study

To validate how close we come to the real shape of a child's head we compared pictures of our reference face (who was 28 years old) taken when he was a child (at 2, 4, 7 and 10 years) with our rejuvenated faces considering different anatomical proportion measures. In order to decide which rejuvenated faces best resemble a given real image, we selected the highest rated rejuvenated faces (rated ≥ 5 on a 7-point Likert scale) in terms of naturalness which were most frequently voted for the relevant age group.

Following Chellappa et al.'s and Farkas et al.'s approach [Ram06, Far94] we specify special landmarks in both of the faces, compare their ratios and thereby evaluate if our technique is able to establish the proportions of a specific child's face. The landmarks we chose are a subset of 57 landmarks mentioned by Farkas et al. [Far94] and can be seen in Figure 6 (left-hand side).

The original and the transformed child faces are in point-to-point correspondence. So, we were able to

 ISBN 978-80-86943-41-1

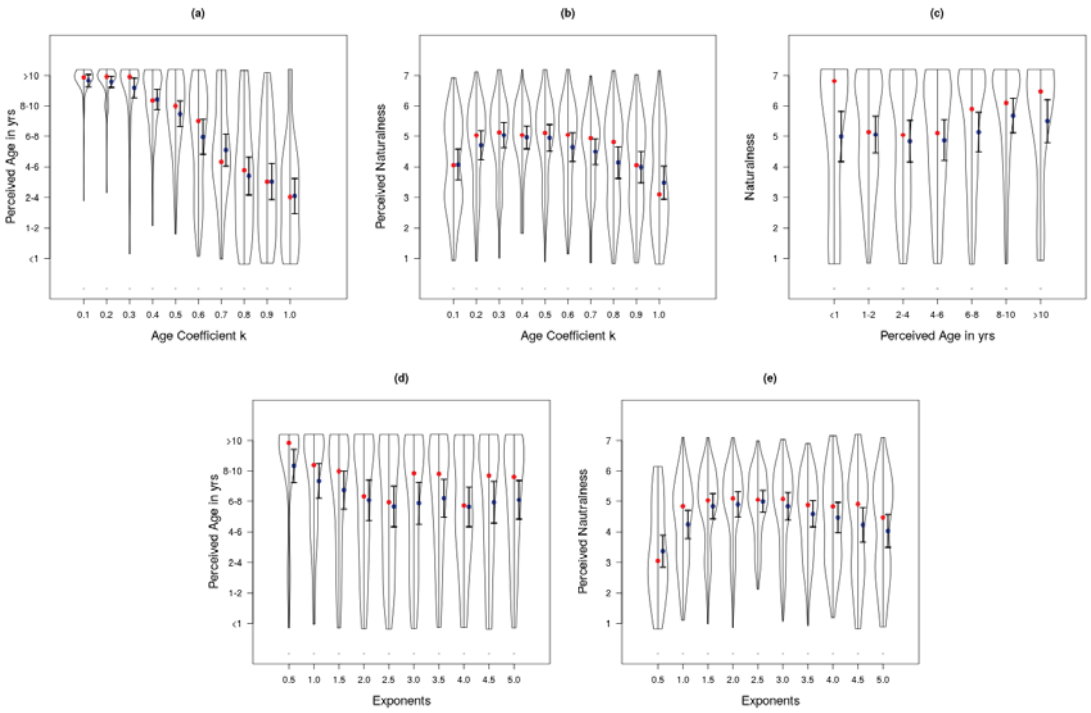

Figure 5: Results of the Experiment: Correlation of (a) perceived age and k, (b) perceived naturalness and k, (c) perceived age and naturalness. (d) perceived age and exponent n, m, (e) naturalness and exponents n, m. Red dots represent the median, blue dots the mean, error bars represent the standard error of the mean, the violins indicate the relative density of the distribution.

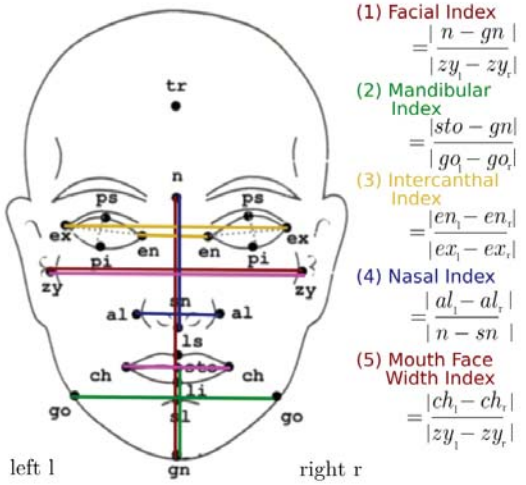

Figure 6: Subset of considered proportions of the face, see [Far94] for the full measurements.

choose exactly the same vertices as landmarks in different rejuvenated models. We then measured the distances among the horizontal and/or vertical axis according to the ratios. To extract the landmarks in the real photos, we manually measured them with the help of a graphics software. Note that low resolution pictures and any out-of-plane rotation (such as turning the head to the side) will artificially alter these "real" values.

In general, our proposed age-regression technique can reconstruct the correct proportions of a specific child. Table 1 shows that facial index (1) values obtained for the rejuvenated faces matched the real images in three out of the four considered age groups. Likewise, the intercanthal index (3) values match very well for every age group. That is a bit surprising since we only modified the eyes as a part of the overall head term and never considered them specifically. The nasal index (4) reveals that our technique works a bit better for older (age ≥ 4) children than younger children. The values for the rejuvenated 2 year old were a bit too low, which might reflect the nose over-stretching mentioned above. The mouth face width index (5) and the mandibular index (2) show a good match for two out of the four reference ages. In particular for index (2) and (5), our values are a little off when it comes to older children.

In Figure 7, the real photos and their corresponding rejuvenated faces with the highest matching ratios among the rated age group are shown. Even though the model and the picture are not always identical, they do have very similar facial proportions (see Table 1) and are clearly perceptually realistic. A finer tuning of the parameters would most likely produce more physically accurate rejuvenated faces.

It is worth mentioning that the age coefficients for 7 and 10 years in Figure 7 are exactly the same. The only difference between these two faces is a small change in

ISBN 978-80-86943-41-1

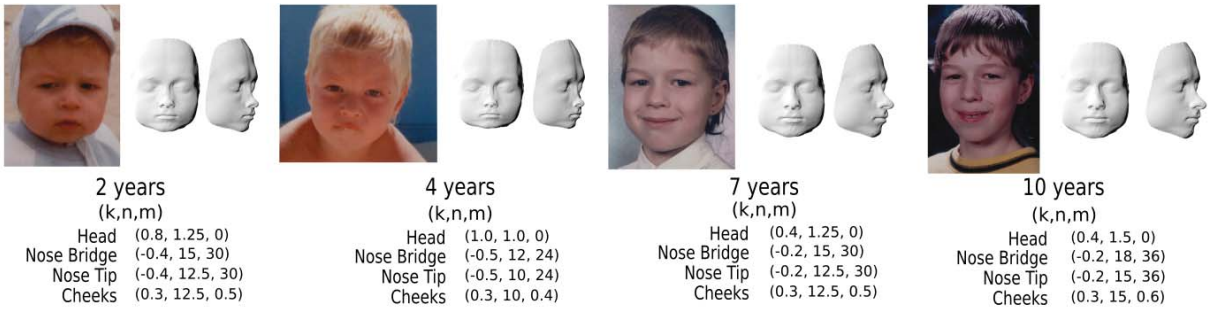

Figure 7: Results for the proportional study: Comparison of the rejuvenated model with original child photographs.

	2 yrs		4yrs		7 yrs		10 yrs	
	real	model	real	model	real	model	real	model
(1)	0.80	0.91	**0.82**	**0.83**	**0.97**	**0.97**	0.96	1.01
(2)	**0.29**	**0.29**	0.38	0.26	0.41	0.31	0.43	0.32
(3)	0.38	0.32	0.36	0.32	0.37	0.31	**0.33**	**0.32**
(4)	0.71	0.59	0.65	0.59	**0.63**	**0.64**	**0.64**	**0.63**
(5)	0.32	0.37	**0.36**	**0.36**	0.52	0.38	0.47	0.38

Table 1: Results of proportional study. Bold numbers show identical ratios. The underlined values show values which are approximately the same (tolerance of 0.06**).**

the exponents n_i and m_i (see Figure 5c). Such a small change in the exponent (e.g. for the k_h a change of 0.5) can be perceived as three years of age difference. We can also see the complex interaction between the age coefficient k_i and the exponents n_i, m_i. Interestingly, in this case the age coefficient was higher (considerably lower for the negatively correlated k_{nb}, k_{nt}) for the four year old than for the two year old. The reduction in age here was caused by the changes in the exponents.

6 CONCLUSIONS AND FUTURE WORK

This paper presents an age-regression algorithm for rejuvenating facial scans. The algorithm takes any 3D polygonal mesh of an adult and applies localized, 3D transformations using a trigonometrical polynomial in order to generate realistic, younger versions of that person. A perceptual experiment varying only a small subset of the parameters showed that the technique can produce young faces that still seem natural. This was confirmed in the proportional study, which showed that it is possible to create faces with similar proportions to that of a specific person when they were a child. We also showed that the proposed parameter space can be used to rejuvenate head models of different genders and ethnic groups. The initial experiments as well as a casual examination of the algorithm have shown that the effect of the different parameters co-varies in some cases and interacts in non-linear ways in other cases. Thus, future work will focus on determining the effect of changes in the cheek area and will also evaluate the individual parameters independently by e.g. showing the participants the modification of the parameters k_i, n_i, m_i separately per facial area and let them decide which parameter value for the head shape, nose bridge, nose tip or cheeks resemble best a child at a certain age group. Further future experimental design could also allow the participants to interactively modify the parameter values for k_i, n_i and m_i for each area (e.g. in form of sliders) to make the modified 3D model match their impression of a child at a certain age. We will also focus on re-formulating the equations to allow a more intuitive selection of a desired age. Furthermore, future work should also incorporate an automatic detection of the location of the core facial regions in the initial 3D-scan to avoid a manual choice of angles. Additionally, an age-regression of the texture should be performed to fully enable a transformation into a child. Finally, since the technique captures the difference between two age groups well, it would be interesting to combine this parameter space with machine learning techniques such Scherbaum et al.'s technique [Sch07]. This would allow us to automatically learn the proper parameter values needed to regress a 3D scan of an adult into a 3D scan of a child.

7 REFERENCES

[Alb11] M. Albert, A. Sethuram and K. Ricanek. Implications of adult facial aging on biometrics. INTECH Open Access Publisher, 2011.

[Cah90] J. Cahn. The Generation of Affect in Synthesized Speech. Journal of the American Voice I/O Society., 8, 1–19, 1990.

[Enl89] D. Enlow. Handbuch des Gesichtswachstums. Quintessenz Bibliothek, 1989.

[Far94] L. Farkas Anthropometry of the head and face Raven Press., 1994.

[Fis16] K. Fisher, J. Towler and M. Eimer. Facial identity and facial expression are initially integrated at visual perceptual stages of face processing. Neuropsychologia., 80, 115–125, 2016.

[Fu10] Y. Fuo, G. Guo and T. Huang. Age Synthesis and Estimation via Faces. IEEE Transactions on Pattern Analysis and Machine Intelligence., 32, 1955 – 1976, 2010.

[Gle75] J. Gleason. Fathers and other strangers: Men's speech to young children. Developmental psy-

ISBN 978-80-86943-41-1

cholinguistics: Theory and applications., 1, 289–297, 1975.

[Gol06] A. Golovinskiy, W. Matusik, H. Pfister, S. Rusinkiewicz and T. Funkhouser. A Statistical Model for Synthesis of Detailed Facial Geometry. ACM Trans. Graph. , 25(3), 1025–1034, 2006.

[Kop05] S. Kopp and L. Gesellensetter and N. Krämer and I. Wachsmuth A conversational agent as museum guide–design and evaluation of a real-world application Springer., 2005.

[Kro17] F. Kron and M. Fetter and M. Scerbo and C. White and D. Becker and others Using a computer simulation for teaching communication skills: A blinded multisite mixed methods randomized controlled trial Elsevier., 2017.

[Kru12] E. Krumhuber, M. Hall, J. Hodgson and A. Kappas. Designing interface agents: Beyond realism, resolution, and the uncanny valley. Proceedings of the 6th Workshop on Emotion and Computing–Current Research and Future Impact., 18 – 25,2012.

[Mar83] L. Mark and J. Todd. The perception of growth in three dimensions. Springer. , 33(2), 193–196, 1983.

[McD12] R. McDonnell, M. Breidt and Heinrich H. Bülthoff. Render Me Real?: Investigating the Effect of Render Style on the Perception of Animated Virtual Humans. ACM Trans. Graph., 31(4), 91:1–91:11, 2012.

[Mor15] L. Morency and G. Stratou and D. DeVault and A. Herholt and M. Lhommet and G. Lucas and F. Morbini and K. Georgila and S. Scherer and J. Gratch SimSensei Demonstration: A Perceptive Virtual Human Interviewer for Healthcare Applications. Proceedings of AAAI., 2015

[Nie09] R. Niewiadomski and E. Bevacqua and M. Mancini and C. Pelachaud Greta: An Interactive Expressive ECA System International Foundation for Autonomous Agents and Multiagent Systems., 2009.

[Par08] F. Parke and K. Waters. Computer facial animation. A.K. Peters Ltd., 2008.

[Pit75] J. Pittenger and R. Shaw. Aging faces as viscal-elastic events: implications for a theory of nonrigid shape perception. Journal of Experimental Psychology: Human perception and performance. , 1(4), 374, 1975.

[Ram06] N. Ramanathan and R. Chellappa. Modeling age progression in young faces. IEEE Computer Vision and Pattern Recognition. , 1, 387–394, 2006.

[Ram09] N. Ramanathan, R. Chellappa and S.Biswas. Age progression in human faces: A survey. Journal of Visual Languages and Computing., 15, 3349–3361, 2009.

[Rya86] E. Ryan, H. Giles, G. Bartolucci and K. Henwood. Psycholinguistic and social psychological components of communication by and with the elderly. Language & Communication., 6(1-2), 1–24, 1986.

[Sch07] K. Scherbaum, M. Sunkel, H.-P- Seidel and V. Blanz. Prediction of Individual Non-Linear Aging Trajectories of Faces. Computer Graphics Forum. , 26(3), 285–294, 2007.

[Sha04] Z. Liu, Z. Zhang and Y. Shan. Image-based surface detail transfer. IEEE Computer Graphics and Applications., 24(3), 30–35, 2004.

[Suo07] J. Suo, F. Min, Z. Songchun, S. Shan and X. Chen. A multi-resolution dynamic model for face aging simulation. IEEE Computer Vision and Pattern Recognition, 2007. CVPR'07. , 1–8, 2007.

[Ten24] http://www.3dscanstore.com/

[Tho17] D.A.W. Thompson. On Growth and Form. Cambridge University Press. , 1917.

[Tod80] J. Todd, L. Mark, R. Shaw and J. Pittenger. The perception of human growth. Scientific American., 242, 132 – 144, 1980.

[Vin06] V. Vinayagamoorthy, M. Gillies, A. Steed, E. Tanguy, X. Pan, C. Loscos and M. Slater. Building expression into virtual characters. Proc. Eurographics Conf. State of the Art Report., 2006.

[Wu94] W. Yin, NM. Thalmann and D. Thalmann. A plastic-visco-elastic model for wrinkles in facial animation and skin aging. Proceedings of the Second Pacific Conference on Computer Graphics and Applications, Pacific Graphics '94. Fundamentals of Computer Graphics., 24(3), 201–214, 1994.

[Yar09] S. Yarosh and G. Abowd. Embodied Interaction for Mediated Communication between Children and Parents. , 2009.

 ISBN 978-80-86943-41-1