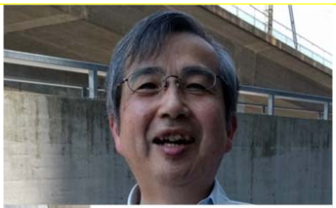

2D Generative Faces for Evolutionary Social Simulation (Installation)

Topic: (Art)

Authors:
Tatsuo Unemi
Japan, Soka University, Department of Information Systems Science
www.intlab.soka.ac.jp/~unemi/
Daniel Bisig
Switzerland, Zurich University of the Arts, ICST
http://bitingbit.org/about/contact.php

Abstract
The simulation of crowds constitutes a useful technique for improving the appearance of CG animations in films and computer games. Since it would look weird if the faces of all the simulated people look exactly the same, there is a need to develop generative algorithms that can automatically generate a wide variety of individual faces. Such algorithms are not only useful for the purpose of interacting with simulated characters such as in computer games but also for observing and understanding activities and events in social simulations.
In our previous work about a simulated evolutionary human society [1], each person was drawn as a simple two dimensional polygon with two colors. In this revised version, the visualization renders each person as a simple face. Similarly to the previous simulation, each person's face is drawn in such a way that a human observer can easily deduce the underlaying hereditary traits coded within the genome of the corresponding person.
While the number of parameters for drawing a visual representation of a human face is potentially enormous, our simulation employs a highly simplified visualization method that requires only five genetically encoded scalar values for drawing a face. Four of these parameters are mapped to the deformation of the face shape, and the fifth parameter affects the tanning of the skin color. The age of the simulated people also affects the shape and color of their faces.
The computational cost for rendering these faces is sufficiently low on a PC that is equipped with a recent GPU to allow the display of more than six thousands agents during each simulation step while preserving a smooth frame rate.

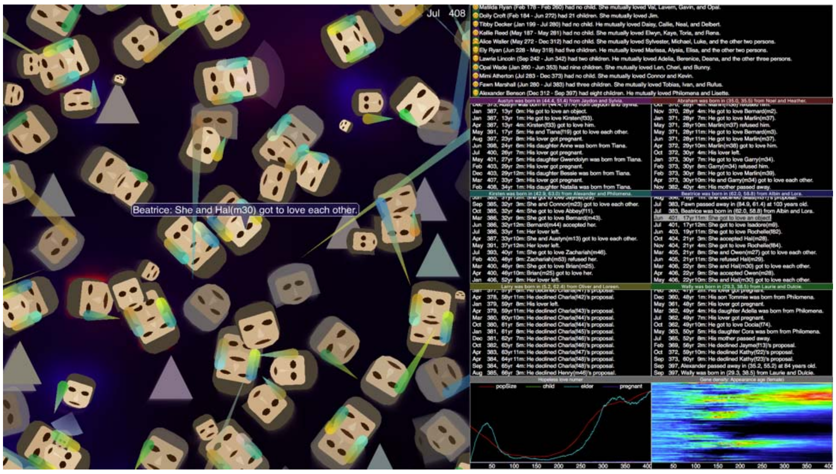

unemi@soka.ac.jp /
Soka University
Tangi-machi 1-236
Hachioji, Tokyo, Japan

Key words: face generation, evolutionary society, 2D CG

Main References:
[1] T. Unemi and D. Bisig, "*Rapid Biography in a Society of Evolutionary Lovers*," in Proceedings of the 20th Generative Art Conference, Ravenna, Italy, 2017.

2D Generative Faces for Evolutionary Social Simulation

Prof. T. Unemi, BEng, MEng, DEng.
Department of Information Systems Science, Soka University, Hachioji, Japan
www.intlab.soka.ac.jp/~unemi/
e-mail: unemi@soka.ac.jp

Dr. D. Bisig, BSc, MSc, PhD.
ICST, Zurich University of the Arts, Zurich, Switzerland.

Premise

The simulation of crowds constitutes a useful technique for improving the appearance of CG animations in films and computer games. Since it would look weird if the faces of all the simulated people look exactly the same, there is a need to develop generative algorithms that can automatically generate a wide variety of individual faces. Such algorithms are not only useful for the purpose of interacting with simulated characters such as in computer games but also for observing and understanding activities and events in social simulations.

In our previous work about a simulated evolutionary human society [1], each person was drawn as a simple two dimensional polygon with two colors. The shape of the polygon represents the sex of the person and the colors form part of a system of hereditary traits for appearance and aesthetic preferences. In this revised version of the social simulation, the visualization renders each person as a simple face. Similarly to the previous simulation, each person's face is drawn in such a way that a human observer can easily deduce the underlaying hereditary traits coded within the genome of the corresponding person.

The following sections describe the mapping between genetic information and rendering parameters, efficiency of rendering process, effects of visualization, and alternative drawings not by polygons but pixel-based images.

1. Parametric face shapes

In a real living organism, its physical appearance emerges through a complex morphological process of growth during which the underlaying genetic code only plays an indirect role. In our simulation, we omit this complicated process and employ a direct mapping between the genetic code and the drawing parameters of the face. While the number of parameters for drawing a visual representation of a human face is potentially enormous, our simulation employs a highly simplified visualization method that requires only five genetically encoded scalar values. Four of these parameters are mapped to the deformation of the face's shape. The left-hand side of figure 1 shows the variation of adult male faces with respect to concentration and spread. A face at the center in this figure is created from mean values of these parameters. Facial elements are represented by polygons that possess a fixed number of vertices. In order to realize the deformation, the following function f is applied to the original x and y coordinates of each vertex to reposition them.

$$f(x) = \frac{\alpha x}{1+(\alpha-1)|x|}, \alpha = \beta^g (g \in [-1,1], \beta = 2) \quad (1)$$

The origin of the coordinate system is assumed to be at the center. The coordinate range of the drawing area is [–1, 1] in each axis. The parameter g controls the amount of concentration. For a g value of –1, the face is mostly spread. For a g value of 1, the face is mostly concentrated. The coefficient β adjusts the maximum deformation rate in the range of $[1, \infty]$.

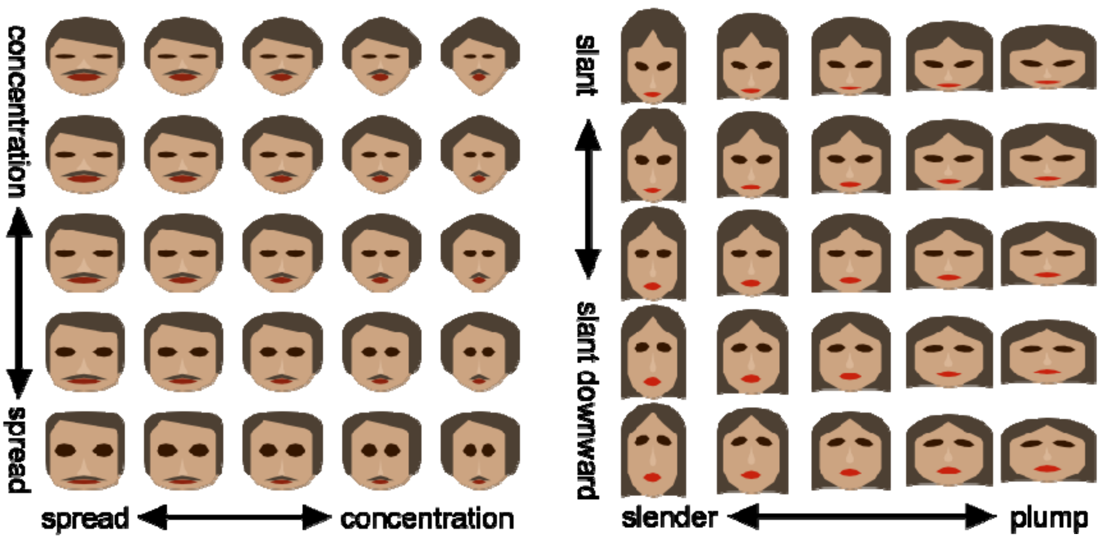

Fig 1. Left: Variation of male adult faces with respect to concentration / spread in the horizontal and vertical direction. Right: Variation of female adult faces with respect to plumpness and slant.

The right-hand side of figure 1 shows the variation of adult female faces with respect to plumpness and slant. The former variation is realized by simply adding a linear offset to each vertex coordinate according to the following equations.

$$\Delta x = \eta k, \Delta y = -\eta k, (k \in [-1,1], \eta = 0.2) \quad (2)$$

The parameter k controls the amount of plumpness. For a k value of -1, the face is mostly slender. For a k value of 1, the face is mostly plump. The coefficient η adjusts the maximum deformation in the range of [0, 1]. The latter deformation is achieved by shifting the vertical coordinate y according to the following equation.

$$\Delta y = \gamma h max\left(0,1-(x^2+y^2)\right), (h \in [-1,1], \gamma = 0.25) \quad (3)$$

The parameter h controls the amount of slant. For a h value of -1, the face is mostly slanted downward. For a h value of 1, the face is mostly slanted upward. The coefficient γ adjusts the maximum deformation rate in the range of [0, 1].

The fifth parameter affects the tanning of the skin color. The age of the simulated people also affects the shape and color of their faces. The position of the eyes and the top vertex of the nose are lower during childhood. With increasing age, the left and right edges of a face shift downwards. The color of the hair gradually becomes gray and finally white once a simulated person reaches a senior age. To increase the visual difference among the sexes, we added a mustache to the face of an adult male older than 16 years. Figure 2 shows a variation of male faces with respect to aging and tanning.

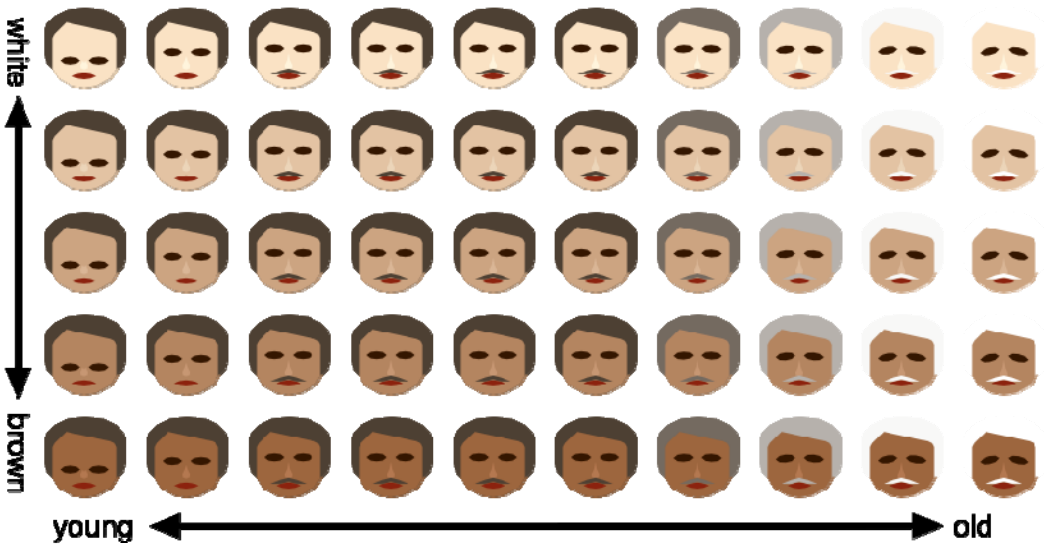

Fig 2. Variation of male faces with respect to aging and tanning.

2. Rendering efficiency

The computational cost for rendering these faces is sufficiently low on a regular personal computer that is equipped with a recent graphic processing unit. Such a computer allows the display of more than six thousands agents at each simulation step while preserving a smooth frame rate. However, it can be useful to further decrease the computational cost for example in order to achieve a smooth animation also on a low cost machine, or to reduce energy consumption, and so on. An effective method to increase drawing efficiency involves a reduction of the number of vertices for facial elements that are drawn at a small scale. For this purpose, we have prepared alternative polygon shapes whose number of vertices is almost half of the original ones. The original shapes contain a total of 146 vertices for a male face and 150 vertices for a female face. For the reduced versions, the number of vertices was decreased to 78 and 74 for male and female faces, respectively. The number of vertices differ between the sexes due to the addition of a mustache for the male face and distinct hair shapes. As described in [1], the visualization possesses the capability to render the 2D animation at an arbitrary zoom factor. When zoomed out, it is effective to show the reduced face versions, whereas the normal versions are shown when zoomed in. Figure 3 illustrates the differences between normal and reduced shapes for male and female faces.

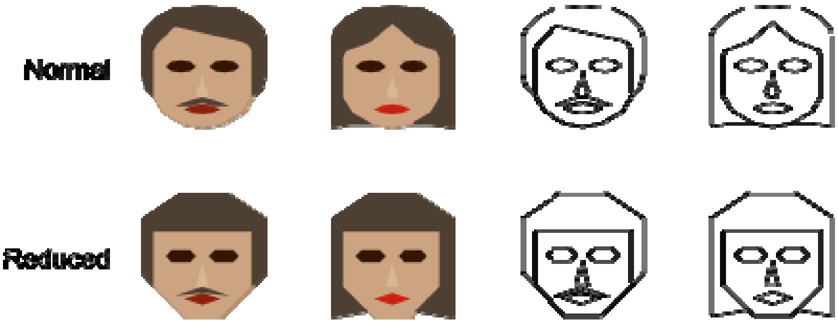

Fig 3. Normal and reduced shapes for male and female faces.

3. Effects of visualization

As has already been demonstrated in [1], the separation of appearances between different sexes is promoted through an evolutionary process that exhibits a selection pressure based on the reproductive advantage of heterosexual couples. This phenomenon becomes apparent in the simulation when observing the differences between facial characteristics of men and women drawn on the screen, even if these characteristics don't appear as particularly natural when compared to real humans. The visualization also reveals a racial separation that becomes apparent as an uneven distribution of colors when observing the entire population. Figure 4 shows an example of a population after 2,000 years of evolution. In this simulation, the parameters that specify face shapes are sex-influenced, but the skin color is not. We can observe a spatial separation of the skin colors. A comparison of two different magnified regions within the simulation (see figure 4) highlights that the different characteristics of male and female faces are race dependent. In the middle image, men appear slender and with slightly slanted eyes whereas women appear plump and vertically concentrated. In the right imaged, men appear with thin eyes and women possess bigger eyes. The size of the eyes is not directly controlled by a genetic parameter but results as a consequence of the vertical spread / concentration of a face.

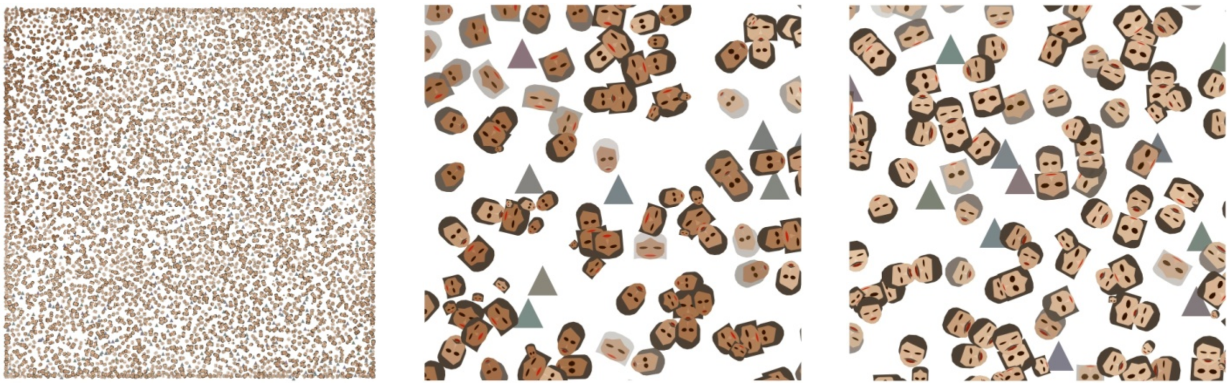

Fig 4. A sample population after 2,000 years of simulation. From left to right: view of the entire population consisting of approx. 6,000 people, magnified view of the top left area, and magnified view of the middle bottom area. Gray triangles represent objects.

4. Utilizing pixel-based images

In general, there exist two distinct methods for drawing 2D visuals by a digital computers. One method employs numerical vectors for representing positions and colors. The other method employs a 2D grid of colored pixels. A vector representation offers the advantage that visuals can be drawn at an arbitrary scale without introducing visual artifacts such as jagged lines. On the other hand, a pixel representation permits the rendering of detailed textures and smooth color gradations both of which are essential for displaying digital photographs and sophisticated paintings.

To examine the possibility of an alternative visualization method, we introduced a pixel-based face representation. The previously described equations for face deformation are still applicable. For this type of visualization, the equations modify the position at which color values are sampled within the pixel array of the original face image. The coordinate values of the sampling positions are calculated using the inverse function of each equation. The higher computational costs for calculating the inverse functions can be alleviated by running the calculations within a shader program on a graphics processing unit. Following this approach, it is possible to render each frame at a frame rate that is sufficient for a smooth animation.

For our first trial, the face images in Emoji characters were examined. Figure 5 shows examples of deformed face images for adult males and females. Apple's Emoji character set includes seven faces for different ages and genders; that is, baby, boy, girl, adult man, adult woman, elderly man, and elderly woman. The Emoji characters also exist in five different skin colors. To realize a continuous variation of skin colors, a face is rendered as a composite image of white and brown skins whose respective appearance is controlled by a blending factor. Unfortunately, there currently exists no method for continuously changing a face with respect to age. As a consequence, a face shape changes suddenly at the boundary of each age span.

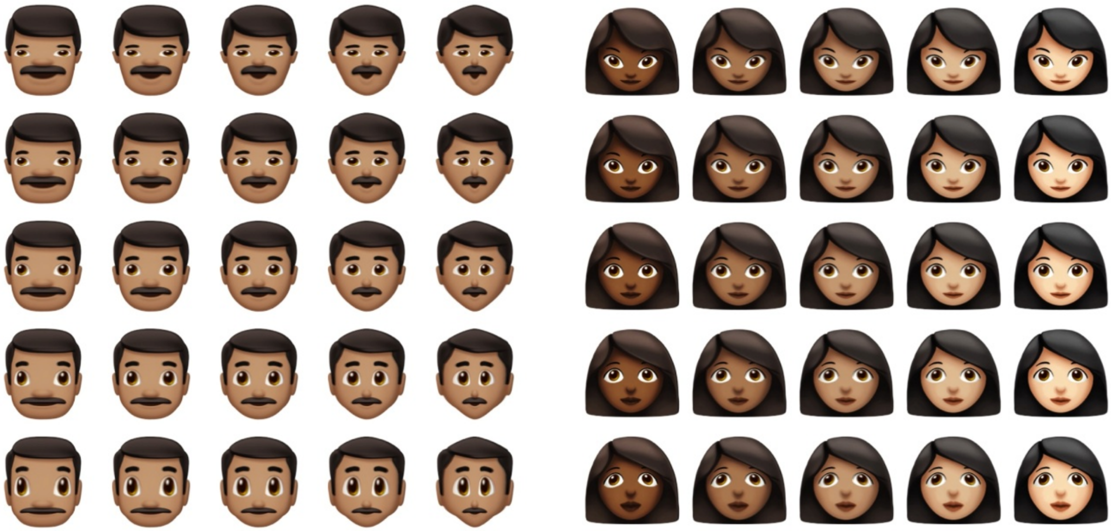

Fig 5. Face variation using pixel-based images extracted from the Emoji character set of Apple's macOS 10.12.

5. Related works

Research on parametric representations of human faces has been conducted within several fields and with different objectives. Chernoff's 2D face [2] provides a useful and intuitive method for clustering by expressing vector data as parametric 2D face shapes. This method can be used to express a maximum of 18 parameters. Typical examples described in [2] include the representation of fossil data (8 parameters), nummulite specimen (6 parameters), and mineral content (12 parameters). The currently most active research field deals with the animation and variation of facial expressions in 3D. One of the most canonical publications that describes this approach can be found in [3]. Due to recent improvements in computational power and machine learning techniques, it has become possible to generate realistic facial animations from a few training examples [4]. While the issue of drawing efficiency is generally considered to be important, the criteria of what constitutes an acceptable computational cost varies greatly between different applications. When a computer that offers higher computational performance becomes available for us, it might become feasible to extend our visualization by rendering each agent as 3D figure with full body and garments.

6. Concluding remarks

In this publication, we described an extension to our simulation that allows to render each agent with a face shape. The current version uses only six parameters to alternate the appearance a face. For future versions, it might be interesting to expand the diversity of face shapes for instance by including parameters for controlling the ratio between a face top and bottom width, and the scales for each individual face elements. By increasing the diversity of face shapes, it would become feasible to display emotional expressions. These expressions could appear in response to the occurrence of important events in the simulated society such as birth, proposal, acceptance, rejection, separation, and death. Though these extension would also make it possible to render the faces in a more realistic manner, an increase in realism is not necessarily our main aim since it likely less effective in provoking the imagination of the audience. Rather, we are hoping to improve the uniqueness and artistic value of our installation.

References

[1] T. Unemi and D. Bisig, Rapid Biography in a Society of Evolutionary Lovers, in *Proceedings of the 20th Generative Art Conference*, pp. 432–441, 2017.

[2] H. Chernoff, The Use of Faces to Represent Points in K-Dimensional Space Graphically, *Journal of the American Statistical Association*, Vol. 68, No. 342, pp. 361–368, June 1973.

[3] F. I. Parke, Computer Generated Animation of Faces, in *Proceedings of the ACM annual Conference*, Vol. 1, pp. 451–457, Boston, MA, 1972.

[4] A. Pumarola, A. Agudo, *et al*: GANimation: Anatomically-aware Facial Animation from a Single Image, *arXiv*:1807.09251, August 2018.