(appeared in *Proceedings SIGGRAPH '98*, pp. 67-74)

An Anthropometric Face Model using Variational Techniques

Douglas DeCarlo, Dimitris Metaxas and Matthew Stone
Department of Computer and Information Science, University of Pennsylvania
{dmd@gradient | dnm@central | matthew@linc}.cis.upenn.edu

Abstract

We describe a system that automatically generates varied geometric models of human faces. A collection of random measurements of the face is generated according to anthropometric statistics for likely face measurements in a population. These measurements are then treated as constraints on a parameterized surface. Variational modeling is used to find a smooth surface that satisfies these constraints while using a prototype shape as a reference.

Keywords: face modeling, anthropometry, variational modeling, crowd generation

1 Introduction

A hallmark of the diversity and individuality of the people we encounter in daily life is the range of variation in the shape of their faces. A simulation or animation that fails to reproduce this diversity—whether by design or circumstance—deprives its characters of independent identities. To animate a bustling scene realistically or to play out an extended virtual interaction believably requires hundreds of different facial geometries, maybe even a distinct one for each person, as in real life.

It is a monumental challenge to achieve such breadth with existing modeling techniques. One possibility might be to use range scanning technology. This involves all the complexities of casting extras for a film: with scanning, each new face must be found on a living subject. And although scanning permits detailed geometries to be extracted quickly, scanned data frequently includes artifacts that must be touched up by hand. Another alternative is manual construction of face models, by deforming an existing model or having an artist design one from scratch; this tends to be slow and expensive.

This paper describes a new alternative: a system capable of automatically generating distinct, plausible face geometries. This system constructs a face in two steps. The first step is the generation of a random set of *measurements* that characterize the face. The form and values of these measurements are computed according to *face anthropometry*, the science dedicated to the measurement of the human face. Anthropometric studies like [11, 12] report statistics on reliable differences in shape across faces within and across populations. Random measurements generated according to the anthropometric profile of a population characterize the distinctive features of a likely face in that population.

In the second step, our system constructs the best surface that satisfies the geometric constraints that a set of measurements imposes, using *variational modeling* [16, 31, 33]. Variational modeling is a framework for building surfaces by constrained optimization; the output surface minimizes a measure of *fairness*, which in our case formalizes how much the surface bends and stretches away from the kind of shape that faces normally have. Having a fairness measure is necessary, since the anthropometric measurements leave the resulting surface underdetermined. Bookstein [4] uses this same fairness measure as a method of data interpolation for sparse biometric data, supporting its utility for determining the geometry of an underdetermined biological shape. Variational modeling provides a powerful and elegant tool for capturing the commonalities in shape among faces along with the differences. Its use reduces the problem of generating face geometries into the problem of generating sets of anthropometric measurements.

The remainder of the paper describes our techniques in more detail. We begin in Section 1.1 by introducing the problem of representing and specifying face geometry. In Section 2, we summarize the research from face anthropometry that we draw on; Section 3 describes how random measurements are generated from these results. In Section 4, we describe our use of variational techniques to derive natural face geometries that satisfy anthropometric measurements. We finish in Section 5 with illustrations of the output of our system.

1.1 Background and related work

Human face animation is a complex task requiring modeling and rendering not only of face geometry, but also of distinctive facial features (such as skin, hair, and tongue) and their motions. Most research in face modeling in computer graphics has addressed these latter problems [21, 23, 25, 26].

Research on human geometry itself falls into two camps, both crucially dependent (in different ways) on human participation. The first approach is to extract geometry automatically from the measurement of a live subject. Lee, et al. [21] use a range scan of a subject, and produce a physics-based model capable of animation. Akimoto, et al. [1] use front and profile images of a subject to produce a model.

The second approach is to facilitate manual specification of new face geometry by a user. A certain facility is offered already by commercial modelers (though of course their use demands considerable artistic skill); several researchers have sought to provide higher levels of control. Parke [25] provides parameters which can control the face shape; and Magnenat-Thalmann, et al. [23] describe a more comprehensive set of localized deformation parameters. Patel [27] offers an alternative set of parameters similar in scope to [23] but more closely tied to the structure of the head. DiPaola [8] uses a set of localized volumetric deformations, with a similar feel to [23] in their effects. Lewis [22] discusses the use of stochastic noise functions as a means of deforming natural objects (including faces). In this case, the control maintained by the user

is limited to noise generation parameters.

In contrast, we adopt a different approach: generating new face geometries automatically. More so than interactive methods, this approach depends on a precise mathematical description of possible face geometries. Many conventional representations of face shape seem inadequate for this purpose.

For example, the simple scaling parameters used by manual modeling techniques can perform useful effects like changing the width of the mouth or the height of the head; but they are unlikely to provide sufficient generality to describe a wide sampling of face geometries.

Meanwhile, for models based on principal components analysis (PCA)—an alternative representation derived from work in face recognition [32]—the opposite problem is likely. PCA describes a face shape as a weighted sum of an orthogonal basis of 3D shapes (called principal components). This basis is constructed from a large bank of examples that have been placed in mutual correspondence. (This correspondence is very much like that required for image morphing [3]; establishing it is a considerable task, but not one that has evaded automation [32].)

PCA typically allows faces nearly identical to those in the bank to be accurately represented by weighting a truncated basis that only includes a few hundred of the most significant components. However, because components are individually complex and combined simply by addition, alternative weightings could easily encode implausible face shapes. Identifying which basis weights are reasonable is just the original problem (of characterizing possible faces) in a different guise. Bookstein [5] describes this problem in terms of "latent variables," and notes that principal components often bear little resemblance to the underlying interdependent structure of biological forms. (In other words, it is quite difficult to extract non-linear dependencies between different shape aspects using a linear model like PCA.) At the same time, there is no guarantee that faces considerably outside the example set will be approximated well at all.

We therefore adopt a representation of face shape based on constrained optimization. The constraints—generated as described in Section 3—are based on the anthropometric studies of the face of [11, 12, 20] described in the next section; we avoid the difficulty of learning possible geometries since these studies identify the range of variation in real faces. The constraint optimization, as described in Section 4, is accomplished by variational surface modeling.

2 Face Anthropometry

Anthropometry is the biological science of human body measurement. Anthropometric data informs a range of enterprises that depend on knowledge of the distribution of measurements across human populations. For example, in human-factors analysis, a known range for human measurements can help guide the design of products to fit most people [9]; in medicine, quantitative comparison of anthropometric data with patients' measurements before and after surgery furthers planning and assessment of plastic and reconstructive surgery [12]; in forensic anthropology, conjectures about likely measurements, derived from anthropometry, figure in the determination of individuals' appearance from their remains [12, 30]; and in the recovery of missing children, by aging their appearance taken from photographs [12]. This paper describes a similar use of anthropometry in the construction of face models for computer graphics applications.[1]

[1] An alternative source of such information might come from morphometrics [5], the study of the overall shape of biological forms, their development, and the interrelations of different aspects of their geometry. Morphometric analyses also provide detailed characterizations of the variability in the shape of faces.

In order to develop useful statistics from anthropometric measurements, the measurements are made in a strictly defined way [19]. The rest of this section outlines one popular regime of such measurements and the information available from analyses of the resulting data. This provides an overview first of the anthropometric structure that our model embodies and then of the statistical results our model exploits.

Anthropometric evaluation begins with the identification of particular locations on a subject, called *landmark* points, defined in terms of visible or palpable features (skin or bone) on the subject. A series of measurements between these landmarks is then taken using carefully specified procedures and measuring instruments (such as calipers, levels and measuring tape). As a result, repeated measurements of the same individual (taken a few days apart) are very reliable, and measurements of different individuals can be successfully compared.

Farkas [12] describes a widely used set of measurements for describing the human face. A large amount of anthropometric data using this system is available [11, 12]. The system uses a total of 47 landmark points to describe the face; Figure 1 illustrates many of them. The landmarks are typically identified by abbreviations of corresponding anatomical terms. For example, the inner corner of the eye is *en* for *endocanthion*, while the top of the flap of cartilage (the tragus) in front of the ear is *t* for *tragion*.

Two of the landmarks determine a canonical horizontal orientation for the head. The horizontal plane is determined by the two lines (on either side of the head) connecting the landmark *t* to the landmark *or* (for *orbitale*), the lowest point of the eye socket on the skull. In measurement, anthropometrists actually align the head to this horizontal, in what is known as Frankfurt horizontal (FH) position [12, 20], so that measurements can be made easily and accurately with respect to this coordinate system. In addition to this, a vertical mid-line axis is defined by the landmarks *n* (for *nasion*), a skull feature roughly between the eyebrows; *sn* (for *subnasale*) the center point where the nose meets the upper lip; and *gn* (for *gnathion*), the lowest point on the chin.

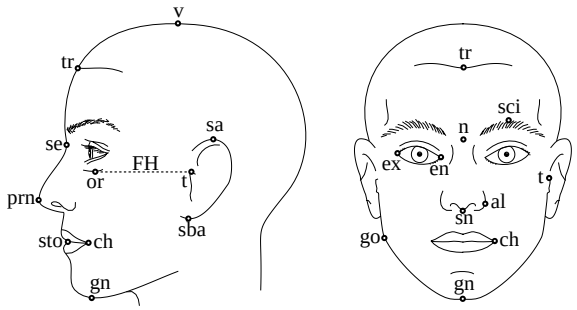

Figure 1: Anthropometric landmarks on the face [12]

Farkas's inventory includes the five types of facial measurements described below and illustrated in Figure 2:

- the *shortest distance* between two landmarks. An example is *en-ex*, the distance between the landmarks at the corners of the eye
- the *axial distance* between two landmarks—the distance measured along one of the axes of the canonical coordinate system, with the head in FH position. An example is *v-tr*, the vertical distance (height difference) between the top of the head (*v* for *vertex*) and hairline (*tr* for *trichion*).
- the *tangential distance* between two landmarks—the distance measured along a prescribed path on the surface of the face. An example is *ch-t*, the surface distance from the corner of the mouth (*ch* for *cheilion*) to the tragus.

- the *angle of inclination* between two landmarks with respect to one of the canonical axes. An example is the inclination of the ear axis with respect to the vertical.
- the *angle between locations*, such as the mentocervical angle (the angle at the chin).

We must represent measurements of each of these types to apply Farkas's anthropometry in creating models for graphics.

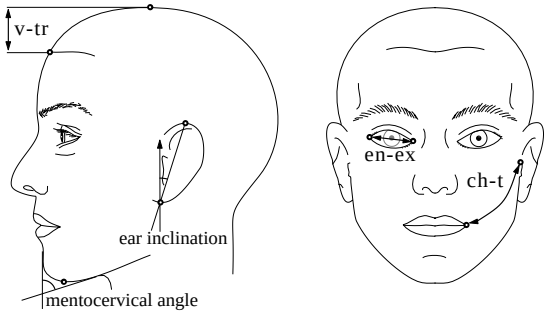

Figure 2: Example anthropometric measurements [12]

Farkas describes a total of 132 measurements on the face and head. Some of the measurements are *paired*, when there is a corresponding measurement on the left and right side of the face. Until recently, the measurement process could only be carried out by experienced anthropometrists by hand. However, recent work has investigated 3-D range scanners as an alternative to manual measurement [6, 12, 20].

Systematic collection of anthropometric measurements has made possible a variety of statistical investigations of groups of subjects. Subjects have been grouped on the basis of gender, race, age, "attractiveness" or the presence of a physical syndrome. Means and variances for the measurements within a group, tabulated in [12, 15], effectively provide a set of measurements which captures virtually all of the variation that can occur in the group.

In addition to statistics on measurements, statistics on the *proportions* between measurements have also been derived. The description of the human form by proportions goes back to Dürer and da Vinci; anthropometrists have found that proportions give useful information about the correlations between features, and can serve as more reliable indicators of group membership than can simple measurements [11]. Many facial proportions have been found to show statistically significant differences across population groups [19]. These proportions are averaged over a particular population group, and means and variances are provided in [11].

3 Generating measurements

The rich descriptions of human geometry developed in anthropometry provide an invaluable resource for human modeling in computer graphics. This goes for artists as well as automatic systems: Parke and Waters [26] describe the importance of having a set of "conformation guidelines" for facial shape, which draw from artistic rules of face design. These guidelines provide qualitative information about the shape and proportion of faces, respecting the quantitative information found in anthropometric measurements.

In using such descriptions, automatic systems immediately confront the problem of bringing a model into correspondence with a desired set of measurements. A widely-used approach is to design a model whose degrees of freedom can be directly specified by anthropometric measurements. For example, in the early visualization frameworks for human factors engineering surveyed in [9]—where anthropometric data first figured in graphics—articulated humans were made to exhibit specified body measurements by rigidly scaling each component of the articulation. Grosso, et al. [17] describe a similar model, but scale physical characteristics (such as mass) as well, to produce a model suitable for dynamic simulation and animation. Azuola [2] builds on Grosso's work, and generates random sets of (axis-aligned distance) measurements using covariance information (but not proportions). The purpose of this generation is to produce a fairly small sampling of differently sized people for human factors analysis.

Our work represents a departure in that we use anthropometric data to constrain the degrees of freedom of the model indirectly (as described in Section 4). This is a must for the diverse, abstract and interrelated measurements of face anthropometry. The flexibility of generating measurements as constraints offers additional benefits. In particular, it allows statistics about proportions to be taken into account as precisely as possible.

This section describes how our system uses published facial measurement and proportion statistics [11, 12] to generate random sets of measurements. The generated measurements both respect a given population distribution, and—thanks to the use of proportions—produce a believable face.

3.1 The need for proportions

Start with a given population, whose anthropometric measurements are tabulated for mean and standard deviation (we later use the measurements from [12]). We can assume that the measurements are given by a Gaussian normal distribution, as corroborated by statistical tests on the raw data [12]. This gives a naive algorithm for deriving a set of measurements—generate each measurement independently as if sampled from the normal distribution with its (estimated) mean and variance. Such random values are easily computed [29]; then, given the constraint-based framework we use, a shape can be generated to fit the resulting suite of measurements as long as the measurements are geometrically consistent.

Mere geometric consistency of measurements is no guarantee of the reasonable appearance of the resulting face shape, however. Anthropometric measurements are not independent. On the face, one striking illustration comes from the inclinations of the profile, which are highly intercorrelated. In the population described in [11], the inclinations to the front of the chin from under the nose (*sn-pg*) and from the lower lip (*li-pg*) take a wide range of values, but, despite the many curves in this part of the face, tend to agree very closely.

Published proportions provide the best available resource to model correlations between measurements such as these. For example, [11] tabulates the mean and variance for statistically significant ratios between anthropometric measurements for a population of young North American Caucasian men and women. Given a calculated value for one measurement, the proportion allows the other measurement to be determined using a random value from the estimated distribution of the proportion. Since the proportion reflects a correlation between these values, the resulting pair of measurements is more representative of the population than the two measurements would be if generated independently.

With many measurements come many useful proportions, but each value will be calculated only once. We must find the proportions that provide the most evidence about the distribution. The next section describes the algorithm we use to do that. It assumes that proportions can be applied in either direction (by approximating the distribution for the inverse proportion) and that we are generating a set of measurements all of which are related by proportions. (We can split the measurements into groups before applying this algorithm.) The algorithm also assumes that we are given a fixed initial measurement (or measurements) in this set from which other measurements could be generated. If we are generating a ran-

dom face, the choice of which initial measurement to use is up in the air. We therefore find the best calculation scheme for each possible initial measurement, and then use the best of those. Random values for this initial measurement are generated by sampling its distribution. Thereafter, randomly generated proportions are used to generate the remaining dependent measurements.

The same algorithm could also be used to fill in measurements specified by a user (as a rough guide of the kind of face needed) or selected to be representative of an extreme in the population (for use in human-factors analysis). In this case, the algorithm gives a way of generating a plausible, random variation on this given information.

3.2 An algorithm for proportions

Given base measurements, our goal is to find the best way to use an inventory of proportions to calculate dependent measurements. We can describe this problem more precisely by viewing measurements as vertices and proportions as edges in a graph. Figure 3(a) shows a portion of this graph, given the measurements and proportions from [11, 12] (some edge labels are omitted for the sake of readability). The presence of cycles in this graph exhibits the need to select proportions. A particular method for calculating measurements using proportions can be represented as a *branching* in this graph—an acyclic directed graph in which each vertex has at most one incident edge. The edge e from s to d in this branching indicates that d is calculated by proportion e from s. By assumption, we will require this branching to span the graph (this means adding dummy edges connecting multiple base measurements). An example branching is illustrated in Figure 3(b), and contains a single base measurement (the vertex marked with a double circle).

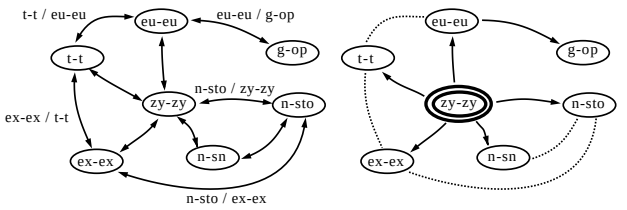

Figure 3: Interpreting measurements and proportions as a graph (a); Example branching used to compute measurements (b)

The algorithm associates each vertex v in the branching with a mean μ_v and variance σ_v^2. The variance is an indication of the precision of the statistical information applied in generating the measurement at v from given information. The smaller σ_v/μ_v, the more constrained the measurement. We take σ_v/μ_v as the weight of d.

For base measurements, σ_v is simply the standard deviation of the measurement. Thereafter, if an edge connects s to d with a proportion with mean μ_e and standard deviation σ_e, and s has mean μ_s and standard deviation σ_s, then the induced distribution at d is characterized by:

$$\mu_d = \mu_s \mu_e$$
$$\sigma_d^2 = \mu_s^2 \sigma_e^2 + \mu_e^2 \sigma_s^2 + \sigma_e^2 \sigma_s^2$$

(This assumes proportions and measurements are independent and Gaussian.) Note that the weight of d is always larger than the weight of s—this means the precision of the information concerning the distribution decreases as we go deeper into the branching.

Our goal in selecting proportions is to derive a branching T_M which assigns a minimum total weight to its vertices. This allows the most constrained features to determine the remaining features via proportionality relationships. We can modify Prim's algorithm for minimum spanning tree to solve this problem. Our algorithm maintains a subtree T of some optimal branching. Initially, the subtree contains just the root for the initial measurement. At subsequent stages, each vertex is associated with the least weight induced by any edge running from the branching to it. The algorithm incorporates the vertex v whose weight is the least into the tree, by the appropriate new edge e.

As with Prim's algorithm (c.f. [13]), the argument that this algorithm works ensures inductively that if T is a subtree of some optimal branching T_M, then so is $T+e$. If e is not an edge in T_M, then T_M contains some other directed path to v, ending with a different edge e'. This path starts at the root of T, so it must at some point leave T. Because e was chosen with minimum weight and weights increase along paths, in fact the path must leave T at e'; since the algorithm chose e, e and e' induce the same weight for v. The inductive property is now established, since $(T_M - e') + e$ is an optimal branching of which T is a subtree.

4 Variational Modeling

Using the method outlined in Section 3, we generate complete sets of anthropometric measurements in Farkas's system. These constraints describe the geometry of the face in great detail, but they by no means specify a unique geometry for the face surface. For example, Farkas's measurements are relatively silent about the distribution of curvature over the face—the particular measurement that specifies the angle formed at the tip of the chin (the mentocervical angle; as in Figure 2), does not actually specify how sharply curved the chin is. What is needed then, intuitively, is a mechanism for generating a shape that shares the important properties of a typical face, as far as possible, but still respects a given set of anthropometric measurements. This intuition allows the problem of building an anthropometric face model to be cast as a constrained optimization problem—anthropometric measurements are treated as constraints, and the remainder of the face is determined by optimizing a surface objective function. This characterization allows us to apply variational modeling techniques [7, 16, 18, 24, 31, 33, 34].

This section briefly introduces variational modeling, and describes how we adapt existing variational modeling techniques to develop the anthropometric face model. Our approach to variational modeling greatly resembles the framework in [33]; a key difference is that we perform most of the variational computation in advance and share results across different face generation runs. This amortization of computational cost makes it feasible to construct larger models subject to many constraints. However, it requires careful formulation of constraints and algorithms to exploit the constancy of the face model and its inventory of constraints.

As described in Section 4.1, we begin by specifying a space of possible face geometries using a parametric surface $\mathbf{s}(u,v)$, and locating the landmark points on the surface. We use a B-spline surface [10] to represent $\mathbf{s}$. This surface is specified by a control mesh, where the mesh degrees of freedom are collected into a vector $\mathbf{p}$. A particular instantiation $\mathbf{p}'$ of $\mathbf{p}$ provides a *prototype shape*, a reference geometry that epitomizes the kind of shape faces have. Both $\mathbf{s}(u,v)$ and $\mathbf{p}'$ are designed by hand, but the same parameterized surface and prototype shape are used to model any set of anthropometric measurements.

Given this shape representation, the task of the face modeling system is to allow a given set of anthropometric measurements $\mathbf{m}$ to be used as degrees of freedom for $\mathbf{s}$, *in place of* $\mathbf{p}$. It does so in two logical steps: (1), expressing $\mathbf{m}$ as constraints on $\mathbf{p}$ in terms of the landmark points as described in Section 4.2; and (2), using variational techniques as described in Section 4.3 through Section 4.5

to find a surface that satisfies the constraints and which minimizes bending and stretching away from the prototype face shape.

4.1 Surface representation

We choose a B-spline surface as a shape representation because of the demands both of anthropometric modeling and variational techniques. Our shape must be smooth, must permit evaluation of our constraints, and must have surface points and tangent vectors that are defined as linear combinations of its control mesh points. This scheme meets all of these requirements.

The specification of $\mathbf{s}(u,v)$ involved the manual construction of a B-spline control mesh for the face, shown in Figure 4. The mesh is a tube with openings at the mouth and neck; the geometry follows an available polygonal face model and (as required for accurate variational modeling) is parameterized to avoid excessive distortion of (u,v) patches.

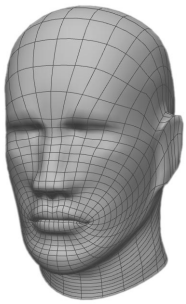

Figure 4: The prototype face model

Anthropometric landmarks are assigned fixed locations on the surface in (u,v) parameter space; some are also associated with constraints that enforce their fixed geometric interpretations. For example, in the case of the *v* landmark, which represents the top of the head, we ensure that the tangent to the surface at the point representing the landmark is in fact horizontal. We likewise add constraints to keep the model in FH position, so that the horizontal axis of the model is consistent with the axis by which landmarks are identified (and measurements taken). These constraints together constitute a set of *base* constraints which must be satisfied to apply any anthropometric measurement. Further constraints are then added to the model—one for each measurement.

4.2 Surface constraints

Our framework derives a shape by applying both linear and nonlinear constraints. The linear constraints are derived from axial distance anthropometric measurements and the base constraints on the model; both can be represented as a linear function of the degrees of freedom of the model, **p**. A matrix **A** describes how the values of all linear constraints are calculated, while a vector **b** encodes the intended values for those measurements. Thus solutions to these constraints satisfy:

$$\mathbf{Ap} = \mathbf{b} \tag{1}$$

Because **A** depends only on the *types* of constraint measurements, **A** can be solved in advance; then values of **p** can be computed directly from **b** given particular measurements **m**.

Many of the constraints are non-linear, however. Each nonlinear constraint is associated with a positive function measuring how far the surface is from the correct measurement. These functions are summed to give an overall penalty function P so that nonlinear constraints impose the equation:

$$P(\mathbf{p}) = 0 \tag{2}$$

($P(\mathbf{p}) \geq 0$ for all **p**). The remainder of this section describes the penalty functions associated with each type of measurement constraint.

The shortest distance measurement constrains the points $\mathbf{x}_i$ and $\mathbf{x}_j$ at a distance r apart using the penalty:

$$P_{\text{dist}}(\mathbf{x}_i,\mathbf{x}_j) = \left(\left\|\mathbf{x}_i - \mathbf{x}_j\right\| - r\right)^2 \tag{3}$$

The tangential distance constraint, which specifies the length of a surface curve to be r, is approximated using the chord-length approximation of a curve [10] using the points $\mathbf{x}_1 \ldots \mathbf{x}_n$:

$$P_{\text{arc-len}}(\mathbf{x}_1,\ldots,\mathbf{x}_n) = \left(\sum_{i=1}^{n-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\| - r\right)^2 \tag{4}$$

The points $\mathbf{x}_i$ all lie on a predetermined curve specified in (u,v)-space (using a B-spline), and are adaptively sampled as to achieve a good estimate of the arc length using the chord-length approximation.

The inclination measurement constraint fixes a vector **v** at an angle θ to a fixed axis **a**:

$$P_{\text{incl}}(\mathbf{v}) = (\hat{\mathbf{v}} - \mathit{Rot}(\mathbf{a},\theta))^2 \tag{5}$$

Using the rotation *Rot*, the axis **a** is aligned with the "goal" direction. **v** can be the direction between two points on the surface, as well as a surface tangent vector.

The angle measurement constraint positions the vectors $\mathbf{v}_1$ and $\mathbf{v}_2$ to be separated by the angle θ. It is treated as two independent inclination constraints:

$$\begin{aligned} P_{\text{angle}_1}(\mathbf{v}_1) &= (\hat{\mathbf{v}}_1 - \mathit{Rot}(\hat{\mathbf{v}}_2,\theta))^2 \\ P_{\text{angle}_2}(\mathbf{v}_2) &= (\hat{\mathbf{v}}_2 - \mathit{Rot}(\hat{\mathbf{v}}_1,-\theta))^2 \end{aligned} \tag{6}$$

4.3 Fairing

A fair surface can be constructed by minimizing an objective function $E(\mathbf{s})$. We will be using the *thin-plate* functional [7, 18, 33] which measures the bending of the surface **s**. It includes the thin-plate term E_{p} to measure bending, and a membrane term E_{m} which ensures the approximation does not become inaccurate:

$$\begin{aligned} E_{\text{p}}(\mathbf{s}) &= \int (\mathbf{s}_{uu}\cdot\mathbf{s}_{uu} + 2\mathbf{s}_{uv}\cdot\mathbf{s}_{uv} + \mathbf{s}_{vv}\cdot\mathbf{s}_{vv})\,du\,dv, \\ E_{\text{m}}(\mathbf{s}) &= \int (\mathbf{s}_u\cdot\mathbf{s}_u + 2\mathbf{s}_u\cdot\mathbf{s}_v + \mathbf{s}_v\cdot\mathbf{s}_v)\,du\,dv \end{aligned} \tag{7}$$

where the subscripts on **s** denote parametric differentiation. The overall fairness of the surface is determined by combining these terms together using weights α and β (where typically α is just large enough to prevent approximation error):

$$E(\mathbf{s}) = \alpha E_{\text{m}}(\mathbf{s}) + \beta E_{\text{p}}(\mathbf{s}) \tag{8}$$

For linear surface representation schemes (including B-splines), the objective function in (8) can be evaluated exactly as a quadratic form $\frac{1}{2}\mathbf{p}^\top\mathbf{Hp}$ [18, 33], where **H** is determined based on the surface representation scheme; the construction for B-splines is given in [33]. Due to the local refinement property of B-splines, **H** is sparse.

The objective function can also be measured with respect to the prototype shape $\mathbf{p}'$ [33], so that the minimization is performed with respect to $(\mathbf{p}-\mathbf{p}')$, resulting in $\frac{1}{2}(\mathbf{p}-\mathbf{p}')^\top\mathbf{H}(\mathbf{p}-\mathbf{p}')$. The use of a prototype shape instructs the fairing process to ignore expected regions of sharp curvature, such as the ears and nose on the face.

Given $\mathbf{H}$, the problem of fairing given purely linear constraints as in (1) is reduced to the following linearly constrained quadratic optimization problem [18, 33]:

$$\min_{\mathbf{p}} \left\| \frac{1}{2}(\mathbf{p}-\mathbf{p}')^\top\mathbf{H}(\mathbf{p}-\mathbf{p}') \right\| \quad \text{subject to } \mathbf{Ap}=\mathbf{b} \tag{9}$$

4.4 Fairing with constraints

There are a number of approaches for solving the constrained minimization problem in (9) including Lagrange multipliers and penalty methods [33] and null-space projection [18], each of which transform the problem to a unconstrained problem.

The Lagrange multiplier $\mathbf{y}$ yields the unconstrained minimization:

$$\min_{\mathbf{p},\mathbf{y}} \left\| \frac{1}{2}(\mathbf{p}-\mathbf{p}')^\top\mathbf{H}(\mathbf{p}-\mathbf{p}') + (\mathbf{Ap}-\mathbf{b})^\top\mathbf{y} \right\| \tag{10}$$

At the minimum, the partial derivatives of the bracketed terms vanish. Differentiation leads to the linear system:

$$\begin{vmatrix} \mathbf{H} & \mathbf{A}^\top \\ \mathbf{A} & 0 \end{vmatrix} \begin{vmatrix} \mathbf{p} \\ \mathbf{y} \end{vmatrix} = \begin{vmatrix} \mathbf{Hp}' \\ \mathbf{b} \end{vmatrix} \tag{11}$$

Solving such a system requires selecting a technique that is mathematically sound and computationally feasible. For example, interactive modeling, with varying constraints and response time demands, requires the use of iterative solution methods, such as the conjugate gradient technique [16, 34]. However, we can solve this system without iteration, using a sparse LU decomposition technique [14]; producing the decomposition takes $O(n^2)$ time given a $O(n)$ sparse $n \times n$ system. This technique is applicable because the set of constraints is hand-constructed, so we can guarantee that the constraint matrix $\mathbf{A}$ contains no dependent rows, and hence that the LU decomposition is well defined. It is feasible because the control mesh topology and the constraint matrix are unchanging, so that only one decomposition ever needs to be generated. Finding solutions is then quite efficient. In general, solving a system given an LU decomposition takes $O(n^2)$ time. However, we have found that the LU decomposition is roughly $O(n)$ sparse given our constraints. (This is not too surprising given that the each constraint involves only a few points on the surface; note that an LU decomposition can be sparse even if the actual inverse is dense.) This means that, in practice, solution steps require roughly linear time.

4.5 Non-linear constraints

As described in Section 4.2, the non-linear constraints are specified using the penalty function $P(\mathbf{p})$. Since this function is positive, it is simply added into the minimization (10) [28, 33]. The extended linear system (11) has $\mathbf{Hp}' - \partial P(\mathbf{p})/\partial\mathbf{p}$ in place of $\mathbf{Hp}'$. Due to the non-linearity of P, this system must be solved iteratively. (By contrast, Section 4.2 described a non-iterative method for solving the linear constraints.)

At iteration i, we determine C_i to be used in place of $-\partial P(\mathbf{p})/\partial\mathbf{p}$ as:

$$C_i = C_{i-1} - \mu_i \frac{\partial P(\mathbf{p}_{i-1})}{\partial \mathbf{p}} \tag{12}$$

with $C_0 = \mathbf{0}$. The scalar value μ is a positive weight (analogous to a time-step in ODE integration), determined using an adaptive method such as step-doubling (for ODE solution) [29]. This results in the iterative linear system:

$$\begin{vmatrix} \mathbf{H} & \mathbf{A}^\top \\ \mathbf{A} & 0 \end{vmatrix} \begin{vmatrix} \mathbf{p}_i \\ \mathbf{y} \end{vmatrix} = \begin{vmatrix} \mathbf{Hp}' + C_i \\ \mathbf{b} \end{vmatrix} \tag{13}$$

where $\mathbf{p}_0$ is the solution corresponding to (11). Note that we still exploit the LU decomposition to allow steps to be solved quickly and exactly; this technique is stabler and faster to converge than the combination of a conjugate gradient technique with the penalty method. We experimented with linearizations of some of the non-linear constraints (and added them into $\mathbf{A}$), but found little gain in efficiency, and decreased stability in solving.

In practice, the simultaneous use of all anthropometric constraints will lead to conflict. For example, some measurements lead to linearly dependent constraints; they are easily identified by inspection, and culled to keep $\mathbf{A}$ invertible. Similarly, when multiple measurements place non-linear constraints on similar features of nearby points on the model (without providing additional variation in shape), including all can introduce a source of geometric inconsistency and prevent the convergence of C. Our constraint set was selected by following a strategy of including only those constraints with the most locally confining definitions (i.e. constraints which affected fewer facial locations or more proximate facial locations were favored).

5 Results and discussion

Sample face models derived using this technique are shown in Figure 5. To produce the measurements for these models, we ran the generation algorithm described in Section 3 on the measurements from [12] and the proportions from [11] for North American Caucasian young adult men and women. Faces for the random measurements were realized by applying the variational framework to a B-spline mesh (a grid 32 by 32) so as to satisfy the base constraints (a total of 15) and 65 measurements that give good coverage both of the shape of the face and of the kinds of measurements used in Farkas's system. There were a total of 120 proportions used as input to the algorithm in Section 3.2.

Producing the LU decomposition used for all these examples involved a one-time cost of roughly 3 minutes on an SGI 175 MHz R10000. Faces typically found their rough shape within 50 iterations; our illustrations were allowed to run for up to 200 iterations to ensure convergence to millimeter accuracy, resulting in runs that took about 1 minute for each face. Models were rendered using RenderMan.

Individual variation across the example males and females in Figure 5 encompass a range of features; for example, clear differences are found in the length and width of the nose and mouth, the inclinations of forehead and nose, as well as the overall shape of the face. At the same time, traits that distinguish men and women—such as the angle at the chin, the slope of the eyes and the height of the lower face (particularly at the jaw)—vary systematically and correctly (based on qualitative comparisons with the anthropometric data). Examining the variation within a population group, the thirty generated males in Figure 6 exhibit the expected range of geometric variation.

In order to quantify this comparison, the proportion-based measurement generation algorithm from Section 3.2 was validated by generating a large number of measurement sets, and comparing the resulting measurement distributions to the published figures from the corresponding population groups. On average, the means differed by about 1% (with a maximum deviation of 4.5%)—well

below the differences in means between population groups. The standard deviations agreed comparably, where the generated measurements had standard deviations that range from being 5% lower to 20% higher than the published values. While this validation guarantees the plausibility of measurements on the generated face models, data is unfortunately not available for comparing the entire geometry (this would require having, for example, a set of measurements of an individual along with a corresponding range scan). One would not expect such a comparison to precisely agree anyway, as the prototype shape has a measurable effect on the resulting geometry. However, this effect decreases with the use of additional measurements, which suggests the need to search out additional data on face geometry (morphometrics [5] seems to be a good starting point).

Despite the many changes, a single prototype shape was used for all examples. This gives the models commonalities in shape where anthropometric data is silent. Further, all the faces use the same texture so as not to exaggerate their differences (having a variety of textures would of course produce nicer results, but would be overlooking the main point of this work). The ears remain coarsely modeled (partly as a result of scarcity of measurements within the ear).

6 Conclusions

This paper has described a two step procedure for generating novel face geometries. The first step produces a plausible set of constraints on the geometry using anthropometric statistics; the second derives a surface that satisfies the constraints using variational modeling. This fruitful combination of techniques offers broader lessons for modeling: in particular, ways to scale up variational modeling—a technique previously restricted to modeling frameworks that have seen limited use to surface fitting tasks—for constrained classes of shapes, and ways to apply anthropometric proportions—long valued by artists and scientists alike—in graphics model generation.

Of course, our models must ultimately be more richly represented. Possible extensions might apply variational techniques to construct the face surface and the interior skull simultaneously; this would form the basis of a face animation model as in [21]. Similarly, landmarks on the face could be used to drive texture synthesis, deriving distinct but plausible patterns of skin and hair.

In the meantime, our work already suggests new computational approaches for tasks that rely on anthropometric results, like forensic anthropology, plastic surgery planning, and child aging. It could also figure in a user interface for editing face models, by allowing features to be edited while related features systematically changed—preserving natural proportions or ensuring that faces respect anthropometric properties common to their population group. Both tasks underscore the importance of continuing to gather and analyze anthropometric data of diverse human populations.

Acknowledgements

We would like to thank Will Welch, Nick Foster, Michael Collins, Max Mintz, Michael Gleicher, Scott King, Nathan Loofbourrow and Charles Loop for their helpful comments and discussion. This research is partially supported by ONR-YIP grant K-5-55043/3916-1552793; ONR DURIP N0001497-1-0396 and N00014-97-1-0385; NSF IRI 95-04372; NSF Career Award grant 9624604; NASA-96-OLMSA-01-147; NIST grant 60NANB7D0058; and ARO grant DAAH-04-96-1-007.

References

[1] T. Akimoto, Y. Suenaga, and R.Wallace. Automatic creation of 3D facial models. *IEEE Computer Graphics and Applications*, 13(5):16–22, September 1993.

[2] F. Azuola. *Error in representation of standard anthropometric data by human figure models*. PhD thesis, University of Pennsylvania, 1996.

[3] T. Beier and S. Neely. Feature-based image metamorphosis. In *Proceedings SIGGRAPH '92*, volume 26, pages 35–42, July 1992.

[4] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.

[5] F. Bookstein. *Morphometric Tools for Landmark Data: Geometry and Biloogy*. Cambridge University Press, 1991.

[6] K. Bush and O. Antonyshyn. 3-dimensional facial anthropometry using a laser-surface scanner–validation of the technique. *Plastic and reconstructive surgery*, 98(2):226–235, August 1996.

[7] G. Celniker and D. Gossard. Deformable curve and surface finite elements for free-form shape design. In *Proceedings SIGGRAPH '91*, volume 25, pages 257–266, 1991.

[8] S. DiPaola. Extending the range of facial types. *Journal of Visualization and Computer Animation*, 2(4):129–131, 1991.

[9] M. Dooley. Anthropometric modeling programs – a survey. *IEEE Computer Graphics and Applications*, 2:17–25, November 1982.

[10] G. Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, 1993.

[11] L. Farkas. *Anthropometric Facial Proportions in Medicine*. Thomas Books, 1987.

[12] L. Farkas. *Anthropometry of the Head and Face*. Raven Press, 1994.

[13] A. Gibbons. *Algorithmic Graph Theory*. Cambridge University Press, 1985.

[14] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1989.

[15] C. Gordon. *1988 anthropometric survey of U.S. Army personnel: methods and summary statistics*. United States Army Natick Research, Development and Engineering Center, 1989.

[16] S. Gortler and M. Cohen. Hierarchical and variational geometric modeling with wavelets. In *1995 Symposium on Interactive 3D Graphics*, pages 35–42, April 1995.

[17] M. Grosso, R. Quach, and N. Badler. Anthropometry for computer animated human figures. In N. Magnenat-Thalmann and D. Thalmann, editors, *State-of-the-art in Computer Animation: Proceedings of Computer Animation '89*, New York, 1989. Springer-Verlag.

[18] M. Halstead, M. Kass, and T. DeRose. Efficient, fair interpolation using Catmull-Clark surfaces. In *Proceedings SIGGRAPH '93*, volume 27, pages 35–44, August 1993.

[19] A. Hrdlicka. *Practical anthropometry*. AMS Press, 1972.

[20] J. Kolar and E. Salter. *Craniofacial Anthropometry: Practical Measurement of the Head and Face for Clinical, Surgical and Research Use*. Charles C. Thomas Publisher, LTD, 1996.

[21] Y. Lee, D. Terzopoulos, and K.Waters. Realistic face modeling for animation. In *Proceedings SIGGRAPH '95*, pages 55–62, 1995.

[22] J. P. Lewis. Algorithms for solid noise synthesis. *Proceedings SIGGRAPH '89*, 23(3):263–270, 1989.

[23] N. Magnenat-Thalmann, H. Minh, M. de Angelis, and D. Thalmann. Design, transformation and animation of human faces. *The Visual Computer*, 5(1/2):32–39, March 1989.

[24] H. Moreton and C. Séquin. Functional optimization for fair surface design. In *Proceedings SIGGRAPH '92*, volume 26, pages 167–176, 1992.

[25] F. Parke. Parameterized models for facial animation. *IEEE Computer Graphics and Applications*, 2(9):61–68, 1982.

[26] F. Parke and K. Waters. *Computer Facial Animation*. A K Peters, 1996.

[27] M. Patel and P. Willis. FACES: The facial animation construction and editing system. In *Eurographics '91*, 1991.

[28] J. Platt and A. Barr. Constraint methods for flexible models. In *Proceedings SIGGRAPH '88*, volume 22, pages 279–288, 1988.

[29] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.

[30] S. Rogers. *Personal Identification from Human Remains*. Charles C. Thomas Publisher, LTD, 1984.

[31] D. Terzopoulos and H. Qin. Dynamic nurbs with geometric constrains for interactive sculpting. *ACM Transactions on Graphics*, 13(2):103–136, 1994.

[32] T. Vetter and T. Poggio. Linear object classes and image synthesis from a single example image. *IEEE Pattern Analysis and Machine Intelligence*, 19(7):733–742, 1997.

[33] W. Welch and A. Witkin. Variational surface modeling. In *Proceedings SIGGRAPH '92*, volume 26, pages 157–166, 1992.

[34] W. Welch and A. Witkin. Free–Form shape design using triangulated surfaces. In *Proceedings SIGGRAPH '94*, volume 28, pages 247–256, July 1994.

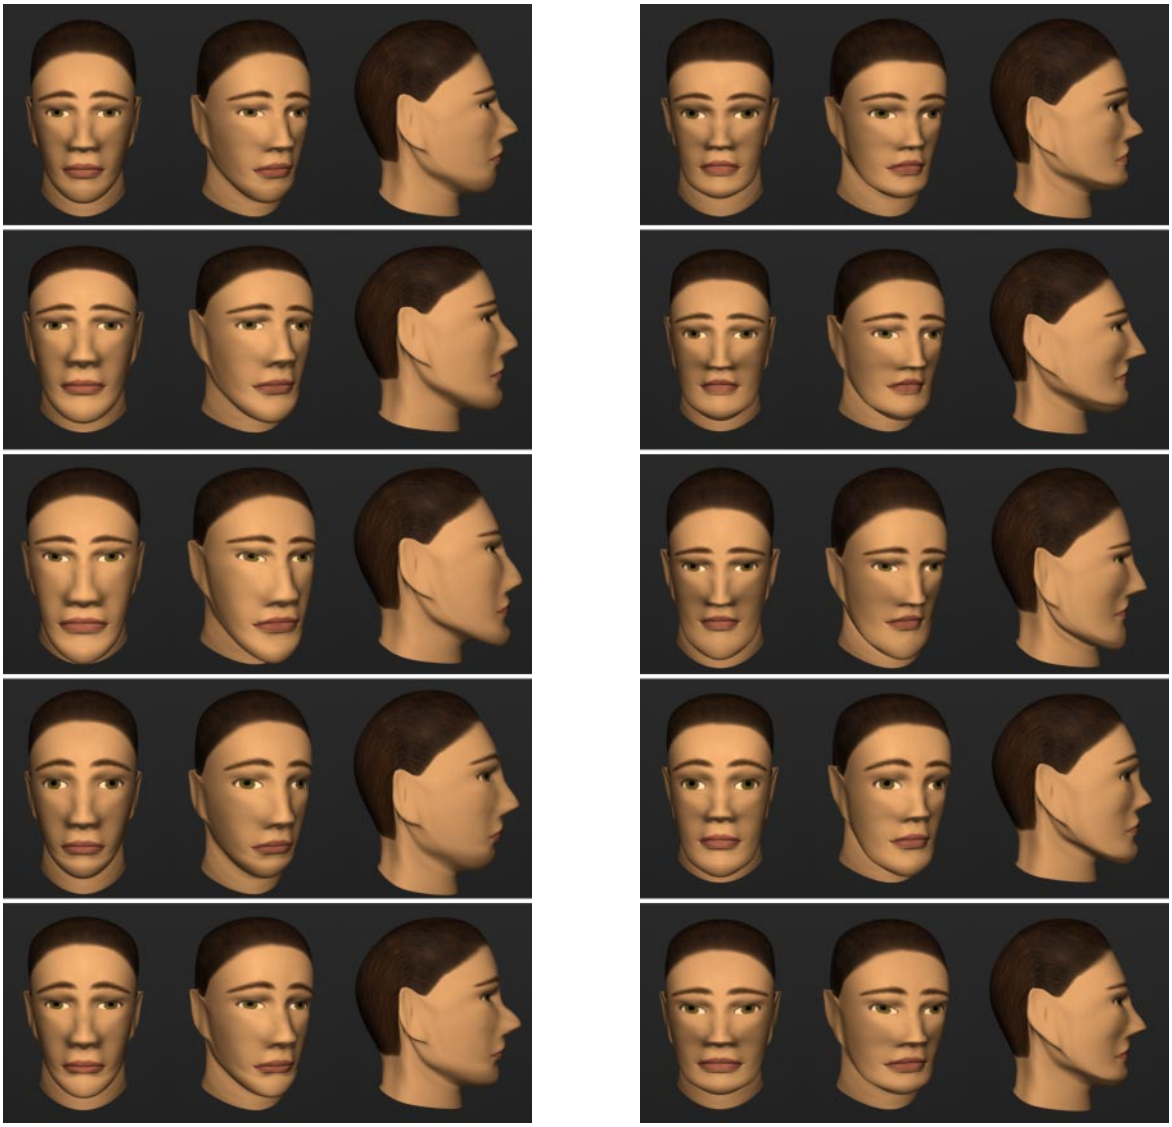

Males Females

Figure 5: Automatically generated face models (3 views of each)

Figure 6: A male a minute