

# ANALYSIS OF HUMAN ATTRACTIVENESS USING MANIFOLD KERNEL REGRESSION

*B. C. Davis and S. Lazebnik*

The University of North Carolina at Chapel Hill  
Department Of Computer Science  
Chapel Hill, North Carolina, USA

## ABSTRACT

This paper uses a recently introduced manifold kernel regression technique to explore the relationship between facial shape and attractiveness on a heterogeneous dataset of over three thousand images gathered from the Web. Using the concept of the Fréchet mean of images under a diffeomorphic transformation model, we evolve the average face as a function of attractiveness ratings. Examining these averages and associated deformation maps enables us to discern aggregate shape change trends for male and female faces.

*Index Terms*— Human attractiveness, Manifold Kernel Regression, Diffeomorphic Registration, Fréchet Mean

## 1. INTRODUCTION

The study of facial attractiveness brings together questions of aesthetics, emotion, biology, and, more recently, computation. The psychology literature contains several theories of what makes a face attractive, including general characteristics like symmetry and averageness, specific facial features, as well as evolutionary, social, and cultural factors [1, 2]. In the last few years, computational approaches have been developed to create attractive composite images [3, 1, 4] or to learn a statistical model for predicting attractiveness [5]. These approaches have been tested on small datasets (not exceeding a hundred images) that have been collected, registered, and annotated under controlled laboratory conditions.

In this paper, we explore a heterogeneous dataset of over three thousand images gathered from a website that allows users to rate pictures based on attractiveness [6]. The range of variability in this dataset is considerable: viewpoint, facial expression, lighting, and image quality all vary widely. Unlike in previous studies, the attractiveness ratings were produced not by experimental subjects in the course of a controlled experiment, but by anonymous computer users with possibly divergent motivations and cultural biases. The variance of ratings under such conditions is much greater than in psychological experiments. Given these issues, it is quite challenging to build a regression model to predict individual attractiveness from this database [6]. But are there any stable patterns that can still be discerned despite all the noise?

Regression is a natural tool for distilling the effect of attractiveness score amid distracting individual variations. In particular, kernel regression is capable of capturing trends with no a priori model [7]. In this paper, we adopt the methodology of Davis et al. [8], where the notion of Fréchet expectation is used to generalize regression to manifold-valued data and applied to the study of anatomical shape change as a function of age in a random design database consisting of 3D MR images of healthy adults. To our knowledge, the present contribution represents the first attempt to regress a shape model of human faces as a function of attractiveness ratings on a large-scale heterogeneous dataset.

The deformation framework underlying our approach is based on Joshi et al. [9], where large-deformation diffeomorphisms are applied to a collection of 3D MR brain scans to compute a *Fréchet mean image*. This is the image that can be registered with the input images using the least amount of “deformation energy,” and it is found by an iterative algorithm of continuous joint alignment of the input images. Related work includes [10], where an affine transformation is used to align 2D face images using a similar method of continuous joint alignment, though no average image is generated. In [11] a diffeomorphic transformation is used to continuously align a collection of 2D face images into a common coordinate system producing a continuously evolving mean. This work uses only ten images taken under controlled poses and lighting conditions. The method of [11] makes a quadratic number of image comparisons at each iteration, which in principle produces a sharper average image. However, because of its computational complexity, this scheme would not scale well to datasets composed of thousands of images, so our approach, which uses only a linear number of comparisons, has an advantage.

## 2. METHODS

This section describes our manifold kernel regression approach to determining average shape trends as a function of attractiveness. It is based on the methodology of [8]. Let  $\{I^i, \alpha^i\}_{i=1}^N$  be the database of images  $I^i$  and associated attractiveness scores  $\alpha^i$ . We formally define the set of images  $\mathcal{I}$  as  $L^2$  functions from  $\Omega \equiv [0, 1] \times [0, 1] \subset \mathbb{R}^2$  to the reals. Attractiveness scores are on a scale of 1 to 10 (the score

associated with each image in the database is the average of multiple users’ ratings, so it is a real number).

The  $L^2$  structure of  $\mathcal{I}$  provides a normed vector space where classical regression techniques, such as the Nadaraya-Watson estimator, can be applied [7]. This estimator can be interpreted as a pixel-wise average of images weighted by kernel functions  $K_h : \mathbb{R} \rightarrow \mathbb{R}$  with bandwidth  $h$ :

$$\hat{I}_h^{L^2}(\alpha) = \frac{\sum_{i=1}^N K_h(\alpha - \alpha^i) I^i}{\sum_{i=1}^N K_h(\alpha - \alpha^i)}. \quad (1)$$

In this work we are interested in average face geometry, not average intensity values. An average that only relies on  $L^2$  structure does not provide a link to geometry: while the images can be added pixel-wise (as in the previous equation), this results in a loss of any identification with the shape configuration of the image contents.

Geometric differences can be accounted for by transformations of underlying image coordinates (see, e.g., [12]). In our work, these transformations  $\phi \in \mathcal{H}$ ,  $\phi : \Omega \rightarrow \Omega$ , are constrained to be diffeomorphic by enforcing laws of continuum mechanics derived from visco-elastic fluid modeling [13]. The action of  $\phi$  on images is defined by  $I_\phi = I \circ \phi^{-1}$ . These transformations form a group where composition is defined but there is no notion of addition, so regression techniques that rely on vector space properties, such as Eq. (1), are not applicable. However, a metric between images that takes geometric change into account can be defined by

$$d^2(I, J) = \min_{\phi \in \mathcal{H}} \left[ d_{\mathcal{H}}^2(\text{ID}_{\mathcal{H}}, \phi) + \frac{1}{\sigma^2} d_{\mathcal{I}}^2(I, J_\phi) \right]. \quad (2)$$

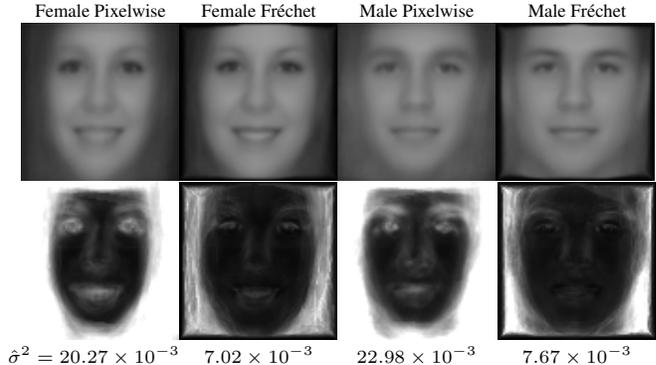
In this equation  $d_{\mathcal{H}}$  measures the amount of geometric change required to match the images [13, 14]. Since there is usually no diffeomorphism that will exactly match two given faces, the second term  $d_{\mathcal{I}}$  measures the residual error between image  $I$  and the deformed image  $J$ . We use  $d_{\mathcal{I}}(I, J_\phi) = \|I - J_\phi\|_{L^2}$ , but other metrics, such as mutual information or cross-correlation, could also be used. The free parameter  $\sigma$  governs the relative importance of these terms. In our experiments, we set  $\sigma = 0.08$ , which gives a high weight to the image match term. Given the metric defined by Eq. (2), the empirical Fréchet mean for our image collection is given by

$$\hat{I} = \operatorname{argmin}_{I \in \mathcal{I}} \frac{1}{N} \sum_{i=1}^N d^2(I, I^i).$$

Figure 1 contrasts the Fréchet mean and the pixel-wise mean. Examples of individual face images deformed to the Fréchet mean are presented in Figure 2.

Starting with the idea of the Fréchet mean, we can define a regression estimator that returns an image  $\hat{I}_h(\alpha)$  as a function of attractiveness value  $\alpha$  and kernel bandwidth  $h$ :

$$\hat{I}_h(\alpha) = \operatorname{argmin}_{I \in \mathcal{I}} \frac{\sum_{i=1}^N K_h(\alpha - \alpha^i) d^2(I, I^i)}{\sum_{i=1}^N K_h(\alpha - \alpha^i)}. \quad (3)$$



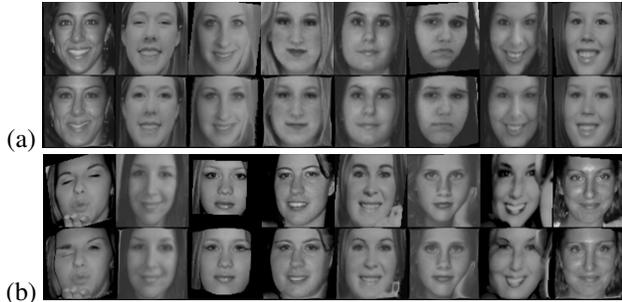
**Fig. 1. Mean face images.** First row: averages of 150 female and male images each. Second row: Maps of residual per-pixel variances. Values above 0.03 (mostly corresponding to the background) are clipped to white. The average variance of each image is shown on the bottom. Note that the faces in the pixel-wise mean images look thinner and their vertical boundaries are softer because of horizontal shifts in the position of the face. The variance maps for the pixel-wise means also reveal bimodalities, e.g., two dominant eyebrow positions for the males. In the Fréchet means, the eyebrows are brought into registration and the bimodality is eliminated.

Eq. (3) expresses the following intuitive idea: For any attractiveness score  $\alpha$ , the average face image is centrally located, according to the metric  $d$ , among the observations that occur near in attractiveness score to  $\alpha$ . As in Eq. (1), the weights are determined by the kernel  $K_h$  (in our experiments, the kernel is Gaussian and  $h = 1$ ). Eq. (3) is solved using an iterative method that results in a continuous joint alignment algorithm. At each iteration, the current estimate for  $\hat{I}$  is fixed and the deformations  $\phi^i$  required to solve (2) for each  $i$  are *incrementally* updated. These estimates of  $\phi^i$  are then used to update  $\hat{I}$  and the process is repeated. This procedure is described in more detail in [8, 9].

Once  $\hat{I}$  is computed, we can quantify its geometric change with the attractiveness-score-indexed transformation  $g_\alpha : [0, 10] \rightarrow \mathcal{H}$ . For each value of  $\alpha$ ,  $g_\alpha$  deforms  $\hat{I}(0)$  so that it matches the regressed image  $\hat{I}_h(\alpha)$ . That is,  $\hat{I}_h(0) \circ g_\alpha^{-1} \approx \hat{I}_h(\alpha)$  (see [8, 12] for details). This deformation is analyzed to determine local spatial patterns of expansion and contraction as a function of attractiveness.

### 3. RESULTS AND DISCUSSION

The images and associated attractiveness scores have been harvested from the website [www.hotornot.com](http://www.hotornot.com) and preprocessed by White et al. [6]. The preprocessing consists of automatic face detection and cropping, followed by an affine rectification that maps automatically detected landmarks (eyes, nose, and corners of the mouth) onto canonical locations. Rectification fails in a small number of cases where one of the landmarks is missed, and we manually removed such images from the database. The resulting database contains 2097 female images and 1523 male images. We converted these images to grayscale and normalized their in-



**Fig. 2. Deformation examples.** (a) Relatively “clean” images. Top row: original images from the database; bottom row: result of deforming the images toward the Fréchet means of Fig. 1. (b) Images with artifacts: closed eyes, background, cropping, head pose, hand in image, shadow on nose, shadows around eye, reflection on forehead. In most cases, our algorithm handles these artifacts relatively robustly, with the exception of the leftmost image, where the closed eyes and their uneven vertical positioning have caused one of the eyebrows to be warped into the eye position.

tensity distributions. The images retain a number of artifacts that make registration extremely challenging. These include resampling artifacts, uncontrolled lighting and pose, overly aggressive cropping, external objects such as eyeglasses or hands, and large variations in background intensities (Fig. 2).

Figure 3 shows results of manifold kernel regression applied to the female and male cohorts separately, and animations of the results can be viewed on the Web at the address given in the caption. Besides having unequal numbers of images, the distributions of the scores for the two cohorts are also different. The average attractiveness score is higher for males than for females (8.4 vs. 7.4) and has lower variance. Both score distributions are skewed toward the high end of the scale, i.e., there are more “attractive” than “unattractive” individuals [6]. There are fewer than 30 males with scores below 6, and the low-rated images tend to be of lower quality and higher variability than the higher-rated ones. This causes registration problems in the regressed male image for the score of 5.5. Apart from this, the results of our method are plausible and allow us to make a number of interesting observations, as explained below. Note that shape change is captured most explicitly not in the mean images, but in the expansion/contraction maps that show the log-determinant of the Jacobian of the derivative of the deformation with respect to attractiveness score. Lighter (resp. darker) values in those maps show areas that are undergoing instantaneous expansion (resp. contraction).

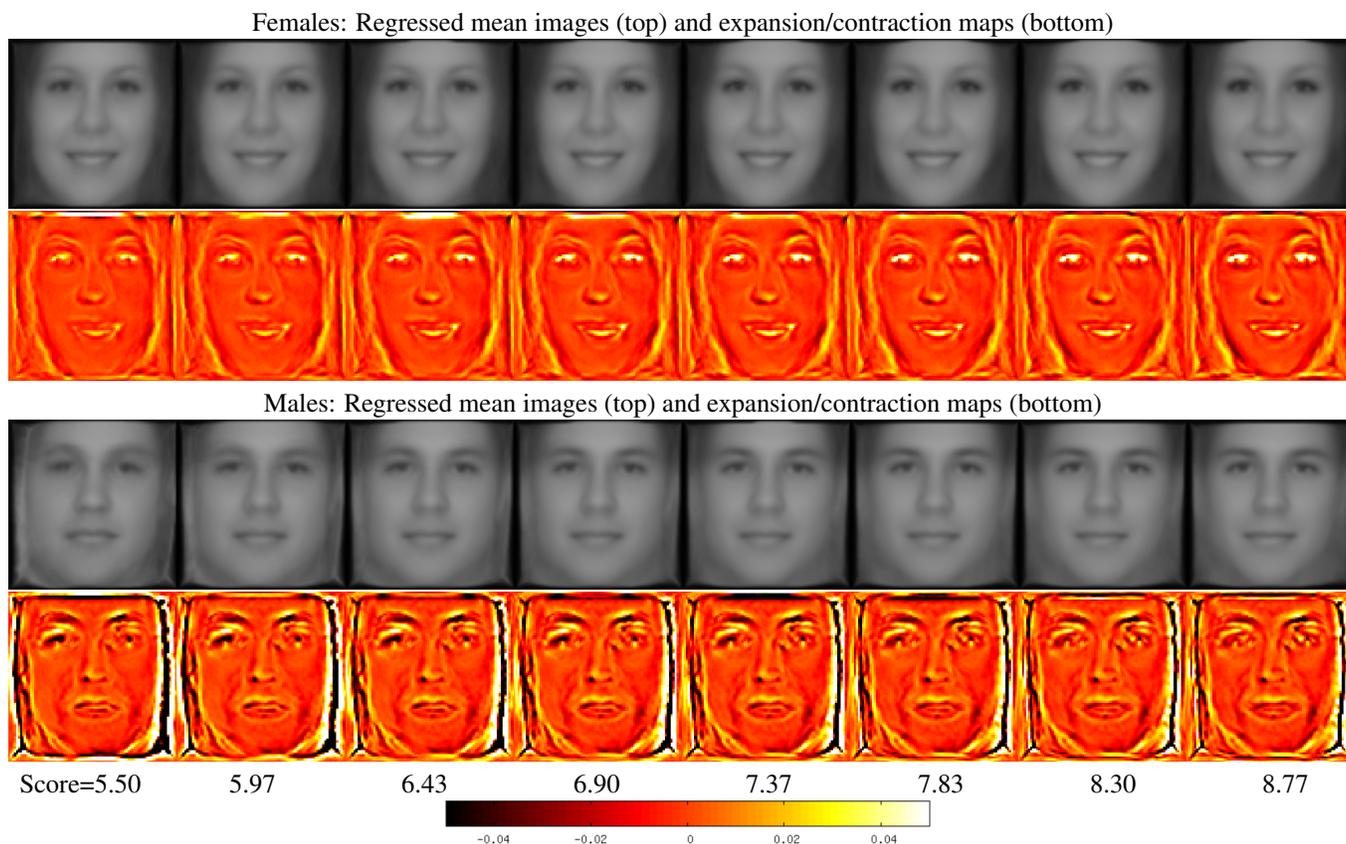
For the females, we can spot several salient trends. As attractiveness score increases, the lower part of the face tends to get thinner and more tapered, the lips expand into an open-mouthed smile, the eyes become notably larger, and the eyebrows are raised. These characteristics of more attractive faces are consistent with the ones reported in the psychology literature [2]. For the males, the shape changes are less salient and more difficult to interpret, but nevertheless we can notice

a few stable trends: the eyebrows are *lowered*, the lower half of the face contracts and the jawline becomes better defined, and the mouth changes from down-turned to more up-turned. The latter change may be caused either by a change in facial pose or expression, and we would like to investigate further whether these two factors become confounded in our regression analysis. As attractiveness rating increases, the Fréchet mean for the male face tends to become better defined. This could imply that more attractive male faces are more “chiseled” and “masculine,” or simply that they are more “average” and thus easier to register. Overall, the male dataset seems to be much more ambiguous than the female one. In the future, we would like to untangle the “spurious” reasons for this, such as biases introduced by the target audience of [www.hotornot.com](http://www.hotornot.com) or by data collection, from the “deep” ones, such as innate or cultural differences that cause males and females to be rated and to rate each other differently.

Averaging plays a central role in the study of attractiveness. Multiple studies have reported that averaging several faces tends to produce a face that seems more attractive than any of the original images [3, 4], though faces of attractive individuals are not necessarily “average” [15]. In the present work, we are not concerned with this phenomenon per se. While the mean faces regressed by our procedure may appear more attractive than the individuals that compose them, there is still a notable trend of shape change from the least to the most attractive faces (at least for the females). Finally, while it is widely reported in the literature that attractive faces are characterized by symmetry, our regressed Fréchet mean faces exhibit a slight skew. This may be an artifact of planar affine rectification applied to non-frontal faces during preprocessing. If there is a preferential turning direction for faces in the database, then the average shape of the face, in particular along the jawline, will not be symmetric. We are currently investigating this issue further.

#### 4. CONCLUSION AND FUTURE WORK

The work presented in this paper is of an exploratory nature. However, our preliminary results are extremely intriguing and our regression methodology is promising. In the future, we would like to carefully study the different factors (pose, lighting, expression) that currently confound our analysis. For example, we plan to detect facial pose (frontal, facing left, facing right) and apply regression separately to each subset. This will give us better insights into facial symmetry. We would also like to investigate alternative deformation models for faces. We currently use a viscous fluid model, which works well for medical images, but can sometimes produce unnatural deformations for faces. An elastic or spline model may be more appropriate. Finally, we would like to solve the *inverse* regression problem: instead of regressing face images as a function of attractiveness, we would like to predict attractiveness as a function of the image. For this, we plan to use



**Fig. 3. Regression results.** We regressed mean images for 16 attractiveness scores linearly spaced from 5.5 to 9, and every other regressed image is shown (the corresponding scores are listed on the bottom). The second and fourth rows show maps of instantaneous local expansion (bright areas) and contraction (dark areas) as a function of attractiveness score. These results are best visualized as animations on <http://www.cs.unc.edu/~davisb/research/Regression-Attractiveness/FaceRegression.html>.

statistical features based on the geometric deformation of the target face to the regressed Fréchet means.

## 5. REFERENCES

- [1] D.I. Perrett, K.A. May, and S. Yoshikawa, "Facial shape and judgments of female attractiveness," *Nature*, vol. 368, pp. 239–242, 1994.
- [2] G. Rhodes and L.A. Zebrowitz, Eds., *Advances in Visual Cognition, Volume 1, Facial attractiveness: Evolutionary, Cognitive and Social Perspectives*, Ablex, Westport, CT, 2002.
- [3] J.H. Langlois and L.A. Roggman, "Attractive faces are only average," *Psychological Science*, vol. 1, pp. 115–121, 1990.
- [4] A.J. OToole, T. Price, T. Vetter, J.C. Bartlett, and V. Blanz, "3d shape and 2d surface textures of human faces: the role of "averages" in attractiveness and age," *Image and Vision Computing*, vol. 18, pp. 9–19, 1999.
- [5] A. Kagian, G. Dror, T. Leyvand, D. Cohen-Or, and E. Ruppiner, "A humanlike predictor of facial attractiveness," *Adv. in Neural Inf. Proc. Systems 19*, pp. 649–656, 2007.
- [6] R. White, A. Eden, and M. Maire, "Automatic prediction of human attractiveness," UC Berkeley Tech. Rep., 2004. [http://www.ryanmwhite.com/research/tr\\_hot.html](http://www.ryanmwhite.com/research/tr_hot.html).
- [7] M. P. Wand and M. C. Jones, *Kernel Smoothing*, Number 60 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, 1995.
- [8] B. C. Davis, P. T. Fletcher, E. Bullitt, and S. C. Joshi, "Population shape regression from random design data," In *Proc. of Int. Conf. on Comp. Vision*, 2007.
- [9] S. C. Joshi, B. C. Davis, M. Jomier, and G. Gerig, "Unbiased diffeomorphic atlas construction for computational anatomy," *NeuroImage (Supplemental issue on Mathematics in Brain Imaging)*, vol. 23, pp. S151–S160, 2004.
- [10] E. G. Learned-Miller, "Data driven image models through continuous joint alignment," *IEEE Trans. on Patt. Anal. and Mach. Intell.*, vol. 28, no. 2, pp. 236–250, 2006.
- [11] G. Charpiat, O. Faugeras, and R. Keriven, "Image statistics based on diffeomorphic matching," *Proc. Int. Conf. on Comp. Vision*, 2005.
- [12] M. I. Miller, "Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms," *NeuroImage*, vol. 23, pp. S19–S33, 2004.
- [13] P. Dupuis, U. Grenander, and M.I. Miller, "Variational problems on flows of diffeomorphisms for image matching," *Quarterly of App. Math.*, vol. 56, no. 3, pp. 587–600, 1998.
- [14] M.F. Beg, M. Miller, A. Trounev, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. of Comp. Vision*, vol. 61, no. 2, 2005.
- [15] T.R. Alley and M.R. Cunningham, "Averaged faces are attractive but very attractive faces are not average," *Psychological Science*, vol. 2, pp. 123–125, 1991.