POLITECNICO DI TORINO
Repository ISTITUZIONALE

The Intrinsic Dimensionality of Attractiveness: A Study in Face Profiles

Original
The Intrinsic Dimensionality of Attractiveness: A Study in Face Profiles / Bottino, A.G., Laurentini, A.. - STAMPA. - 7441:(2012), pp. 59-66. (17th Iberoamerican Congress on Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications, CIARP 2012 Buenos Aires (Argentina) 3-6 September 2012) [10.1007/978-3-642-33275-3_7].

Availability:
This version is available at: 11583/2502048 since:

Publisher:
Springer

Published
DOI:10.1007/978-3-642-33275-3_7

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

06 September 2026

Author's version

Published in: Lecture Notes in Computer Science, Volume 7441 LNCS, 2012, Pages 59-66

The Intrinsic Dimensionality of Attractiveness: a Study in Face Profiles

Andrea Bottino, Aldo Laurentini

Politecnico di Torino, Dipartimento di Automatica e Informatica, Italy
{andrea.bottino, aldo.laurentini}@polito.it

Abstract. The study of human attractiveness with pattern analysis techniques is an emerging research field. One still largely unresolved problem is which are the facial features relevant to attractiveness, how they combine together, and the number of independent parameters required for describing and identifying harmonious faces. In this paper, we present a first study about this problem, applied to face profiles. First, according to several empirical results, we hypothesize the existence of two well separated manifolds of attractive and unattractive face profiles. Then, we analyze with manifold learning techniques their intrinsic dimensionality. Finally, we show that the profile data can be reduced, with various techniques, to the intrinsic dimensions, largely without loosing their ability to discriminate between attractive and unattractive faces.

Keywords: manifold learning, intrinsic dimensionality, dimensionality reduction, profiles, facial attractiveness.

1 Introduction

In recent years the scientific analysis of facial attractiveness has been a major research issue both in medical areas such as plastic surgery and orthodontics, and in human science fields such as psychology, psychobiology, anthropology, evolutionary biology, behavioral and cognitive sciences. Many thousands of relevant papers have been presented in these areas. Several results point to the objective nature of the human perception of attractiveness, suggesting that beauty is not, or not only, "in the eye of the beholder". Empirical rating studies have demonstrated high beauty rating congruence over ethnicity, social class, age, and sex ([1], [2], [3], [4]). Recent studies in psychophysiology and neuropsychology lead to the detection of the brain areas where the assessment of facial beauty is processed. Activity patterns related to explicit attractiveness judgement of face images, showed a non-linear response profile, with a greater response to highly attractive and unattractive faces. Finally,

babies as young as three/six months, which are not affected by cultural standards about beauty, were found to be able to distinguish between faces previously rated as attractive or unattractive by adult raters ([5]). These results show that the human perception of attractiveness is essentially data-driven, and largely irrespective of the perceiver. They are the rationale of the use of pattern analysis/image processing techniques for objective attractiveness analysis. Computer analysis of attractiveness has several practical applications such as supporting studies in human science, planning plastic surgery and orthodontic treatments, suggesting the make-up and hairstyle more fitting to a particular face, selecting images for social networks or curricola. Using pattern analysis techniques for analyzing facial attractiveness is an emerging research area, and a number of paper on this subject have been recently published (see [5]). Although many interesting results have been obtained, they have not yet been combined in an overall framework. In particular, the main problems, that is: which are the objective elements of facial beauty, how they combine together and whether they can be expressed in some simple form, are far from being solved.

Using the face space paradigm ([6]), according to which faces represent a d-dimension manifold in the D-dimension space used to describe them, with $d<<D$, most of these unsolved problems can be expressed as the problem of learning the manifolds of faces rated for attractiveness. Manifold learning is an active area of research, aimed at discovering hidden relations between multidimensional data ([7]). Learning a manifold means first understanding its intrinsic dimensionality (ID), that is the number of independent parameters required for describing the manifold. The next step is reducing the high dimensionality of the original data into a space with dimensions near to ID, maintaining, as far as possible, the relations between data points relevant to the problem considered. Up to now, no such research has been performed in the face space with relation to attractiveness. Manifold learning techniques have been found useful for other face analysis, as human age estimation ([8]) and gender classification problems ([9]). Observe that an important requirement for manifold learning is a sufficiently dense sampling. Unfortunately, we have no clear idea of the meaning of “sufficiently dense” in the case of manifolds of faces rated for beauty. In [10] it has been observed that classification accuracy, that is coherence with human rating, increased with the number of samples without showing sign of saturation using around hundred 2D frontal expressionless samples. This and other facts point to a clear undersampling of the face space, in particular for very beautiful faces, even for monochromatic images.

In this paper, we present what to our knowledge is the first study that applies manifold learning techniques to the problem of facial attractiveness. In particular, we will deal with face profiles, in order to reduce possible undersampling problems (w.r.t. frontal images). Actually, profiles are very characterizing face features. In recent studies, they have been found to convey several information, sufficient, for instance, for identity recognition ([13], [14], [15]), for identifying gender and ethnicity ([12]), for planning plastic surgery ([11]), and for recognizing facial expressions ([16]). In addition, it has been demonstrated that beauty ratings of frontal and profile images are strongly correlated ([17]).

The aim of our work is the following. First of all, the research previously quoted that supports the objective nature of human attractiveness, also points to the existence in face space of two well separate manifolds, related to attractive and unattractive faces. This is also strongly supported by the fact that several approaches aimed at automatically rating face attractiveness report great accuracy for the higher and lower beauty levels, while average attractiveness judgments are much more uncertain both for automatic and human ratings ([22], [23], [24], [25]). Therefore, we first analyze the ID of the manifolds of attractive and unattractive face profiles. Then, we show that discriminating the two manifolds can be effectively performed with data reduced, with various techniques, to dimensions near to their ID. This has been done collecting a training set of face profile images rated for attractiveness by a human panel, and constructing, on the basis of the reduced image data, an automatic rater, to be compared for a test set with human ratings, assumed to be ground truth. Human raters are asked to score faces attractiveness with some integer numbers, from which the two classes of attractive and unattractive profiles can be separated. Therefore, attractiveness estimation is considered as a classification problem and its accuracy is evaluated as the percentage of test samples classified into the right class.

The rest of the paper is organized as follows. In section 2, we present the database used. In section 3, we briefly discuss the technique for estimating the ID related to the attractiveness classification. Section 4 is devoted to present and discuss the experimental results obtained.

2 Sampling the manifolds of pleasant and unpleasant faces

The first problem to face for this work, as well as for other 2D or 3D beauty research, is the lack of databases containing faces rated for attractiveness, and in particular beautiful faces. Therefore, we decided to build such a database, collecting an initial set of profile images, with different resolutions, selected from several sources (Bernard Achermann DB, Color FERET DB, CVL Face DB, Flickr and color photographs of volunteers participating to this research). Some examples can be seen in Fig. 1. The reference database contains 510 profile images with neutral expression, different age and ethnicity (45 Africans, 68 Asians and 397 Caucasians) and equally divided between the two genders. In order to identify samples belonging to the manifolds of attractive and unattractive faces for our investigation, we asked a panel of human raters to evaluate their attractiveness. The obtained scores have then been used, to separate these two sets from that of attractively average faces.

Fig. 1. profile images in the DB

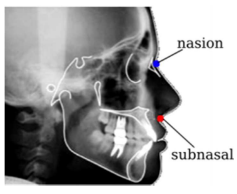

Fig. 2. Nasion ans subnasal points

The samples in the DB were rated through a public website by a panel of students and colleagues of our University, who were asked to express a vote for each subject on a 10 levels scale, ranging from 1 (attractive) to 10 (unattractive). Prior to web evaluation images were properly cropped and scaled to focus raters on the profiles. The raters were almost equally divided between genders (53.3% males and 46.7% females). Since the scores of the human raters are not coincident, the attractiveness value for a profile is considered as the mean of the raters' votes. A total of 82,102 votes, with an average of 160 votes per image, were collected, showing a substantial rating congruence between male and female raters (Pearson correlation of 0.94), consistent with previously reported findings [5]. The final mean ratings were in the range [1.99, 7.91], with a 41% reduction of the initial available rating interval. As we expected, selecting faces from the available face databases strongly reduces the number of samples (very attractive and very unattractive) useful for our study. We underline the fact that this is a problem that seems to affect most of the data sets used in the literature for attractiveness related research.

In order to perform meaningful comparisons, the heterogeneous profiles in the DB have been normalized. This process was first aimed at aligning them and delimiting the same section for all profiles, including the most significant facial features (forehead, nose, mouth and chin) and then at reducing the effects of varying lighting conditions. Geometric normalization is based on the position of two landmarks in the profile contour (Fig. 2): nasion (the point in the skull where the nasal and frontal bone unite) and subnasal (the point, above the upper lip, where the nasal septum begins). These two landmarks are identified using the algorithm described in [18], which first extract the face silhouette by background subtraction and then processes its outline; landmarks are then aligned with two predefined points within a fixed area of interest, whose size is 200x100 pixels. Finally, normalized images are converted to grayscale and their histograms are equalized. Each profile is then represented as a one-dimensional vector of size 20.000, obtained concatenating grayscale image rows.

3 Intrinsic dimensionality and dimensionality reduction

The intrinsic dimensionality ID of a data set with dimension D can be defined as the number of independent parameters that can be used to describe the data set without significant loss of information relative to the problem considered. In other terms, it means that the data points lie on a manifold of dimension ID, where $0 < ID \leq D$. Several methods have been reported in literature for ID estimation. In this work, we use a fractal-based estimator, called Correlation Dimension ([19]). The basic idea for this and other estimators is that the number of points enclosed into a hypersphere of radius r centred on a point of the manifold grows proportionally to r^{ID}. The Correlation Integral $C(r)$, defined as:

$$C(r) = \frac{2}{n(n-1)} \sum_{i=1}^{n} \sum_{j=i+1}^{n} c, \quad \text{where } c = \begin{cases} 1, \text{ if } \|x_i - x_j\| \leq r \\ 0, \text{ if } \|x_i - x_j\| > r \end{cases}$$

where x_i and x_j are points of the dataset, provides the relative amount of pairs of points lying into an hypersphere of radius r. $C(r)$ can be used to estimate ID of the dataset, by computing the limit: $\lim_{r \to 0} \log(C(r)) / \log(r)$.

For a finite set of samples, this limit can be estimated considering the slope of the linear part of the curve $\log(C(r))/\log(r)$. As we already stated in the introduction, the reliability of the ID estimate has been tested by checking if attractiveness can be adequately discriminated in two classes by using ID dimension for each face sample. In other words, the human panel attractiveness scores are used to extract two subsets of attractive and unattractive profiles and discrimination is considered as a classification problem. Given the uncertainty about the adequateness of the density of sampling, we assume the ID computed with this technique as a rough estimation, and for classification we will experiment several other dimensions near ID. For conducting our tests, we have

selected three different linear and non-linear dimensionality reduction methods: PCA, Isomap and Laplacian Eigenmaps ([7]).

4 Experimental results

Our purpose is to estimate how many parameters are required for discriminating profiles belonging, according to attractiveness scores, to the manifolds of attractive and unattractive profiles. Hence, we first estimate the ID of the manifolds containing some of the best and worst classified profiles. The separation of the two classes of attractive and unattractive faces from that of attractively average faces, is given by the lower and upper percentile of all attractiveness scores. Then, to validate these estimates, we reduce dimensionality to various values near to the estimated ID using various techniques and attempt to discriminate classes with different attractiveness using these reduced dimensions. Validation has been done with different datasets to investigate the relevance to profile attractiveness classification of several factors: sex, number of samples and separation in attractiveness of the two manifolds.

	Female			**Male**			**Mixed**		
	1st	***25th***	***50th***	***1st***	***25th***	***50th***	***1st***	***25th***	***50th***
Attractive	7.91	5.72	5.11	7.46	5.07	4.48	7.91	6.07	5.52
Unattractive	1.99	3.06	3.35	2.23	2.92	3.14	1.99	2.71	3.00
Diff	***5.92***	***2.66***	***1.76***	***5.23***	***2.15***	***1.34***	***5.92***	***3.36***	***2.52***

Table 1. attractiveness ratings of the samples in the reference database

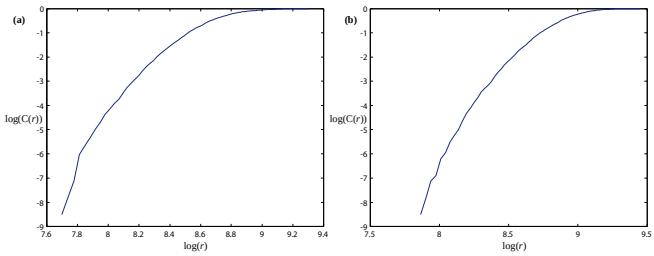

Fig. 3. The plot of Correlation Dimension for attractive (a) and unattractive (b) profile images

It is clear that more samples in each dataset provide a more dense coverage of the manifolds and a better training of the classifiers. It is also clear that the largest the distance, in terms of attractiveness between the classes of attractive and unattractive samples, the better the two classes are separated and, therefore, better classification results can be expected. Unfortunately, these requirements conflict, since increasing the dataset size reduces distances between classes, and vice-versa. This can be seen in **Table 1**, where the attractiveness ratings of the 1st, 25th and 50th best and worst samples of each class for the two distinct sexes and mixed sexes are listed. In order to keep a reasonable interval between the two classes, two datasets for each gender were created. The firsts comprise the 25 best and 25 worst rated profiles, the seconds the 50 best and 50 worst. Finally, we created two other sets of 50 and 100 samples combining the best and worst rated profiles, without regard of their sex.

An estimate of the ID of the manifolds of attractive and unattractive faces has been obtained applying the Correlation Dimension technique to two datasets, combining the 100 best and the 100 worst male and female profiles. The plots of the Correlation Dimension for these datasets are shown in Fig. 4. Since the plots are non linear, we selected three different intervals on that curve and evaluated the mean slope of the lines of best fit. We

estimated 12 as ID for the attractive silhouettes (the results in the various intervals were d1=14, d2=12 and d3=9) and 11 for the unattractive ones (d1=13, d2=11 and d3=9). Since the ID evaluated with this technique can be considered only as a rough estimate, we performed classification experiments with several values near to the estimated IDs. After dimensionality reduction, three different classifiers were used: Support Vector Machines with radial basis kernel (SVM), whose parameters were optimized with a grid approach, Multi Layer Perceptron (MLP), with 10 training epochs and 5 hidden units, and k-Nearest Neighbours (kNN), with k=4. For assessing the classification results, in all experiments we applied a stratified 10-fold cross validation technique. In Table 2 we show a summary of the best classification results for different data sets (Female 50, Female 100, Male 50, Male 100, Mixed 50, Mixed 100) and reduced dimensionality spaces of size 3, 5, 10, 15, 20, 25 and 30, dropping the reference to dimensionality reduction and classification technique. The last column reports the highest classification accuracy obtained for each data set. We recall that ground truth values are given by human panel scores.

	3	*5*	*10*	*15*	*20*	*25*	*30*	*max*
Female 50	0,84	0,90	0,86	0,86	0,86	0,86	0,84	*0,90*
Female 100	0,82	0,83	0,86	0,87	0,85	0,84	0,87	*0,87*
Male 50	0,78	0,72	0,80	0,78	0,80	0,88	0,76	*0,88*
Male 100	0,67	0,71	0,77	0,77	0,77	0,84	0,81	*0,84*
Mixed 50	0,86	0,94	0,90	0,88	0,90	0,86	0,88	*0,94*
Mixed 100	0,76	0,80	0,87	0,85	0,86	0,88	0,88	*0,88*

Table 2. Best classification accuracies

The following main observations stems from the above table.

1. The main result is the effective profile attractiveness discrimination in low dimensionality spaces. Although the beauty rating separations between test datasets is rather low, the classification results in spaces with dimensionality near to the estimated ID in general can be considered in good agreement with the human ratings (94% accuracy for Mixed 50 and 90% for Female 50, both in a 5 dimension space). When the separation is lower, better results are achieved with a dimension somewhat higher (15 for Female 100, 25 for Male 50, Male 100 and Mixed 100), but still close to the estimated IDs.
2. Although not shown in the table, the classification results are not much affected from the data reduction techniques (linear, PCA, or non-linear, Isomap and Laplacian Eigenmaps,). This fact points to a good intrinsic separation of the manifolds of attractive and unattractive face profiles in the face space, which appears to be an interesting result. As for different classifiers, SVM performed consistently better.
3. As expected, more effective classification is obtained for better separated datasets. As can be seen in the table: i) results achieved by 50 element datasets are better in all the cases than those obtained by 100 element datasets; ii) mixed datasets are better than female ones, which are in turn better than male datasets (according to their rating distances in Table 1).
4. Female datasets achieved better classification results than male datasets. One reason is that the average ratings of the attractive males was lower than that of attractive females. Another reason could be that attractive male faces have in general stronger features than attractive female features ([20], [21]), which hints at a worst sampling of the attractive male manifold. In general, according to various results presented in human sciences, as those stating that qualities as averageness and symmetry are much more related to female than male beauty ([5]), computer analysis of female beauty is likely to be easier than male beauty.

5 Conclusions and future work

In this paper, we presented what to our knowledge is the first study that applies manifold learning techniques to the analysis of facial attractiveness. Understanding the intrinsic dimensionality of the manifolds of attractive and unattractive faces is a first step toward understanding which facial elements are relevant to attractiveness, and how they must combine together. In order to reduce possible under-sampling problems, we analyzed the ID and dimensionality reduction techniques for face profiles. The analysis of data sets of attractive and unattractive faces has provided an intrinsic dimensionality ID not much far from 10. Several dimensionality reduction techniques have been experimented, and the discrimination of attractive and unattractive profiles in low dimensionality spaces has been compared with human ratings. The tests show that a number of independent parameters near to the estimated ID are sufficient for attractiveness ratings in good agreement with human judgement. Although we believe that these first results are interesting, much further work is needed to approach a full understanding of the elements of facial beauty and their relations. While the manifolds of attractive and unattractive faces have been shown to be well separated in low dimensionality spaces, the shape of these manifolds is still to determine, as well as the best data reduction techniques. A basic requirement of this research would be a dense sampling, in 2D, or better in 3D, of the manifolds of faces with high beauty ratings differences, and in particular of faces rated for high attractiveness. Since currently no such data set is available, we plan to construct it, starting from that of frontal 2D images.

References

[1] Bashour, M., "History and Current Concepts in the Analysis of Facial Attractiveness", Plastic and Reconstructive Surgery, Vol.118, No.3, pp.741-756, 2006

[2] M.R Cunningham, A.R. Roberts, A.P. Barbee, P.B. Druen, "Their ideas of beauty are, on the whole, the same as ours: Consistency and variability in the crosscultural perceptions of female attractiveness", Journ. Pers. and Social Psyc., Vol.68, pp.261-279, 1995

[3] Etcoff, N., "Beauty and the beholder", Nature, Vol.368, pp.186-187, 1994

[4] Jones, D., "Physical Attractiveness and the Theory of Sexual Selection: Results from Five Populations", University of Michigan Press, Ann Harbour, 1996

[5] A. Bottino, A. Laurentini, "The analysis of facial beauty: an emerging area of research in pattern analysis," ICIAR 2010, June 21-23, 2010, pp.425-435

[6] L. Sirovich, M. Kirby, "Low-dimensional procedure for the characterization of human faces," J. Opt. Soc. Am. Vol.4, pp. 519-524, 1987

[7] L.J.P. van der Maaten. "An Introduction to Dimensionality Reduction Using Matlab". Int. Report MICC 07-07, Universiteit Maastricht, The Netherlands, 2007

[8] G. Guo, Y.Fu, C.R. Dyer, and T. Huang, "Image-based human age estimation by manifold learning and locally adjusted robust regression," IEEE Trans. on Image Processing, Vol. 17 , no.7, 2008, pp. 1178 -1188

[9] S. Buchala, N. Davey, R.J. Frank, and T.M. Gale, "Dimensionality reduction of face images for gender classification," Proc. IEEE Conf. on Intell. Syst. pp. 88-93, 2004

[10] Y. Eisenthal, G. Dror, and E. Ruppin," Facial Attractiveness: beauty and the machine," Neural Computation, Vol. 18, pp. 119-142, 2006

[11] T. Ozkul, M.H. Ozkul, "Computer simulation tool for rhinoplasty planning", Comput. In Biol. And Med., Vol.34, pp.697-718, 2004

[12] U.Tariq, Y.Hu and T.S.Huang, "Gender and Ethnicity identification from silhouetted face profiles," IEEE Proc. ICIP, pp. 2441-2444, Nov. 2009

[13] I.A.Kakadiaris, H. Abdelmunim, W.Yang, and T. Theoharis, “Profile-based face recognition,” IEEE Proc. FG’08, pp.1-8, Sept. 2008
[14] C.J. Wu and J.S. Huang, “Human face profile recognitionby computer,” Pattern Recognition, vol. 23, pp.255–259, 1990
[15] X. Zhou and B. Bhanu, “Human recognition based onface profiles in video,” in Proc. IEEE CVPR, Washington, DC, Jun. 2005
[16] M. Pantic, I. Patras, and L. Rothkrantz, “ Facial action recognition in face profiles image sequences,” Proc. IEEE ICME’02, pp. 37-40, 2002
[17] N. Davidenko, “Silhouetted face profiles: A new methodology for face perception research”. Journal of Vision, vol 7(4):6, pp. 1-17, 2007
[18] A. Bottino, S. Cumani, “A fast and robust method for the identification of face landmarks in profile images”. W. Trans. on Comp. 7, 8 , 2008, 1250-1259
[19] P. Grassberger and I. Procaccia. "Measuring the Strangeness of Strange Attractors". Physica D: Nonlinear Phenomena 9 (1-2), 1983, pp. 189-208
[20] C.F. Keating, “Gender and the physiognomy of dominance and attractiveness,” Social Psychology Quarterly, Vol. 48, pp. 61-70, 1985
[21] M.Zuk, “The role of parasites in sexual selection: current evidence and future directions,” Advances in the Study of Behavior, Vol.21, pp. 39-67, 1992
[22] P.Aarabi, D.Hughes, K. Mohajer, M. Emami, “The automatic measurement of facial beauty,” IEEE SMC 2004, pp. 2168-2174
[23] H. Mao, L. Jin, and M.Du, “ Automatic classification of Chinese female facial beauty using Support Vector Machine, “ IEEE SMC 2009, pp.4842- 4846
[24] D. Sutic, I. Breskovic, R. Huic, and I. Jukic, “Automatic evaluation of facial attractiveness,” MIPRO 2010, May. 24-28, Opatija, pp. 1339 -1342
[25] Y. Eisenthal, G. Dror, and E. Ruppin,“ Facial Attractiveness: beauty and the machine,“ Neural Computation, Vol. 18, pp. 119-142, 2006