

The Role of Facial Regions in Evaluating Social Dimensions

David Masip Rodo¹, Alexander Todorov², and Jordi Vitrià Marca³

¹ Universitat Oberta de Catalunya

² Princeton University

³ Universitat de Barcelona

dmasipr@uoc.edu, atodorov@princeton.edu, jordi.vitria@ub.edu

Abstract. Facial trait judgments are an important information cue for people. Recent works in the Psychology field have stated the basis of face evaluation, defining a set of traits that we evaluate from faces (e.g. dominance, trustworthiness, aggressiveness, attractiveness, threatening or intelligence among others). We rapidly infer information from others faces, usually after a short period of time ($< 1000ms$) we perceive a certain degree of dominance or trustworthiness of another person from the face. Although these perceptions are not necessarily accurate, they influence many important social outcomes (such as the results of the elections or the court decisions). This topic has also attracted the attention of Computer Vision scientists, and recently a computational model to automatically predict trait evaluations from faces has been proposed. These systems try to mimic the human perception by means of applying machine learning classifiers to a set of labeled data. In this paper we perform an experimental study on the specific facial features that trigger the social inferences. Using previous results from the literature, we propose to use simple similarity maps to evaluate which regions of the face influence the most the trait inferences. The correlation analysis is performed using only appearance, and the results from the experiments suggest that each trait is correlated with specific facial characteristics.

1 Introduction

Humans continuously infer social judgments from the people surrounding us. One of the most important sources of social information is the face, specially in the first impression, and these impressions have been shown to be pervasive in many situations in our daily life [1, 2]. In spite of its questionability [3], facial trait judgments influence important social outcomes, such as the success in elections [4, 5] or court room decisions [6]. Surprisingly, this social evaluations are performed rapidly [7], and do not involve a deep thought. During the last decade, facial trait judgments have been an important research topic in Social Psychology. Currently, there exist psychological models of facial trait evaluation [8, 9] and these social evaluations have been successfully related to the amygdala, a specific brain region located deep in the medial temporal lobes [10, 11].

This field of research has also attracted scientists from computer science, and computational models to automatically predict facial trait judgments have been developed. Facial inferences can be computationally modeled from the image stimuli, and computer vision techniques have been applied to the problem of facial trait judgments prediction. Specifically, two different descriptors have been developed: a structural modeling of the facial geometry using a set of predefined landmarks [12] and an appearance descriptor of the holistic internal region of the face [13]. The comparative studies favor the appearance over the structure, and some trait judgments can be predicted with high accuracy (80 – 90%).

In this paper we provide more insights about the specific zones of the face that influence the facial trait judgments. We present a study that aims to answer the questions: Which parts of the face play a more active role in the face evaluation? What makes a person look more dominant, aggressive or trustworthy? Using a labeled dataset, we propose the use of simple statistical measures to build correlation maps among pixels and behavioral data. Some traits are strongly correlated with specific regions of the face, and facial trait judgments have also important dependencies on the gender.

The paper is organized as follows: section two describes the basic psychological modeling of facial evaluation, the main trait judgments are enumerated and a brief review of the main previous works is performed. Section three presents the methodology used in this study. Section four shows the experimental results, and section five discusses relevant insights from these results. Finally section six concludes this paper.

2 Basis of Face Evaluation

We constantly evaluate faces in multiple trait dimensions. In [8], Oosterhof and Todorov defined a 2D model of the facial trait judgments we perform in our daily life. The data collected in this seminal work constitutes the base corpus of our study. In this section, we will briefly review the basic notions underlying the trait judgments and the labeled facial dataset used in this paper.

In a first step (see supporting information in [8] for more details), 55 participants in the experiment were asked to write unconstrained descriptions from a set of faces. Participants were instructed to write “everything that came to their mind about the person shown in the screen”. This textual information (1134 descriptions) was analyzed and classified by two researchers in fourteen trait categories (accounting for 68% of the textual statements). These 14 trait categories were: *Attractive, caring, aggressive, mean, intelligent, confident, emotionally-stable, trustworthy, responsible, sociable, weird, unhappy, dominant and threatening*.

In this paper, we will use the same list of trait categories. In [13], some of these traits were computationally learnt using both a structural and appearance descriptors. Besides the trait judgments learnability, authors also suggested that some regions of the face are more significant for facial trait classification. A set of predefined landmark locations were used to correlate facial geometry and facial evaluations. In this paper we propose a computational study that aims to



Fig. 1. Examples from male and female images used in the rating experiments from the Karolinska data set

analyze which regions or parts of the facial appearance are more useful for trait inferences. This study complements the initial findings in [13], given that the holistic appearance is used instead of the simple geometry of the landmarks.

2.1 Annotated Database

In the second experiment from [8], 327 undergraduate students rated a set of 66 images from the Karolinska data set [14] (half of them from males and half from females). Students were instructed to answer relying in their first impressions, and the ratings were performed in a 1 to 9 scale (where 1 meant “not at all [trait]” and 9 meant “extreme [trait]”). Each face was exposed for about 1000 ms.

Using these ratings, the researchers submitted the data to a Principal Component Analysis (PCA). The first component accounted for 63% of the data variance, and the second for 18%. Positive judgments (intelligent, attractive, responsible,...) had positive loading in the first PC and negative judgments (aggressive, mean, threatening,...) had negative loadings in the same first PC. Authors named this first axis as **valence**. The second PC was named **dominance** given the high loadings of the dominant, aggressiveness and confidence traits on this axis. In the light of these results, they suggested that the underlying dimensionality of face evaluation is 2.

In this paper we used the same 66 face images from the Karolinska data base, and we also used the normalized ratings from the previously described paper [8]¹. Figure 1 shows some samples from the database. Images have been previously converted from RGB to gray scale by averaging the 3 color planes, and each image has been centered and aligned using an Ad hoc coded GUI² and a manually annotated set of landmarks.

¹ The evaluations of the 14 traits are available for research purposes in <http://webscript.princeton.edu/~t1lab/databases/> from the authors of [8] upon request.

² A Matlab R2010b script has been used for the registration and normalization purposes.

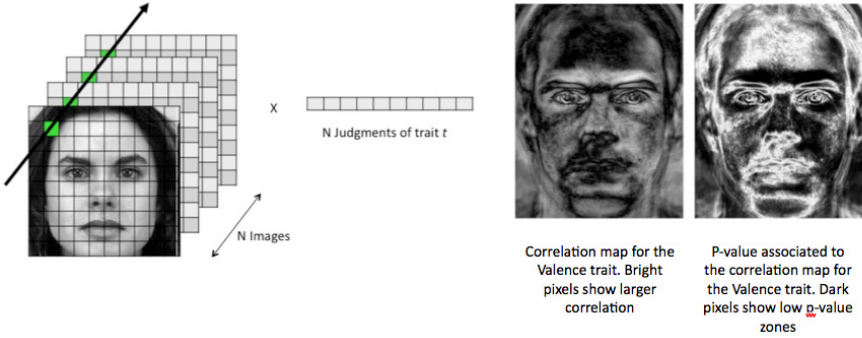


Fig. 2. Scheme of the correlation computation at pixel (i, j) and the resulting maps from the valence trait. The image shows the correlation maps, where light pixels show high correlation with the trait judgment, and the p-value map, where dark values show low p-values in the analyzed zone.

3 Facial Zones in Social Face Evaluation

The human evaluation of the facial trait judgments is not performed uniformly on the face images. Different facial zones play crucial roles in the ratings depending on the trait analyzed. In this study, we propose a methodology to evaluate the visual importance of each region of the face in the human trait judgments inferences. The problem of evaluating the saliency in images has been previously tackled in the object recognition literature [15–17]. Several approaches have been followed, usually filtering the image using: Laplacian, Gaussian, Gabor filters and steerable wavelets [18]. The appropriate scale (coarse or fine) has also been studied in object recognition and perception publications [19, 20]. Moreover, the appropriate scale has also been shown to influence face classification tasks such as face recognition [21]. In this study, a fine strategy has been followed. The trait inferences are studied in a pixel based approach correlating the image information with the behavioral data.

3.1 Similarity Maps Using Correlation

Given a data set consisting of N images represented as $\{\mathbf{X}\}_{1 \dots N}$, where the pixel (i, j) of the image a is denoted as $\{X_{ij}\}_a$; and a set of trait judgments \mathbf{T} , where the value T_{tk} denotes the mean value of t^{th} trait for image k , we define the similarity operator between the image zone and the trait judgments as the correlation:

$$\rho(\mathbf{X}_{ij}, \mathbf{T}_t) = \frac{cov(\mathbf{X}_{ij}, \mathbf{T}_t)}{\sigma_{\mathbf{X}_{ij}} \sigma_{\mathbf{T}_t}} \tag{1}$$

We define the similarity map between the trait judgment t and the pixel images as a new image \mathbf{M} that has at each pixel \mathbf{M}_{ij} the result of $\rho(\mathbf{X}_{ij}, \mathbf{T}_t)$. i.e. the

similarity in terms of correlation between the N values of the pixel at the position (i,j) and the N mean judgments of the trait analyzed. Figure 2 summarizes the algorithm.

Pseudocode of similarity maps construction

Given the set of face images X , and the mean trait judgments T :

```
begin
  for i:=1 to image height H do
    for j:=1 to image width W do
       $M_{ij} := \rho(\mathbf{X}_{ij}, \mathbf{T}_t)$ 
    endfor
  endfor
end
```

Facial images usually suffer from noise artifacts related to the acquisition process or the posterior alignment step. In addition, pixel locations far from the landmarked regions are more prone to small errors. Although visually, the resulting correlation maps convey interesting information about the facial zones of interest according to the trait judgments, we used a filtered version of the face images in the experiment. Briefly, a Gaussian filter with parameter $\sigma = 5$ has been convolved with each original image, obtaining a filtered image where each pixel is the average of its environment. The resulting correlation maps are used as a similarity scores to evaluate the significance of each region of the face.

4 Experiments

The experiments use the Karolinska data set, and the ratings obtained in [8]. Figure 3 graphically shows the results for each one of the 14 facial trait judgments.

Dominance, aggressiveness, and threatening traits show positive correlation (> 0.3) specifically in the chin region (with p-value < 0.05) and forehead (p-value < 0.1). This result is consistent with the intuitive idea of dominance and the prototypical synthetic images of these traits developed in [8]³.

On the other hand, attractiveness, caring, sociable and trustworthiness show positive correlation in the cheek region (> 0.3 , p-value < 0.1). These traits are correlated to the first PC (Valence axis), and the results are consistent with [13] where the upper half of the facial geometric information was found to be strongly related to these traits (specially eyes and eyebrows).

Caring, emotionally-stable, responsible and sociable show high correlations also in the eyebrow regions (with p-value < 0.1), while unhappy and weird traits do not show relevant correlations in any specific zones.

The same experiment has been repeated but using as a labels set the Valence and Dominance. To obtain these artificially generated dimensions, the same procedure as in [8] has been followed. The trait ratings have been submitted to a

³ See visual animations in

<http://webscript.princeton.edu/~tlab/demonstrations/> for more details.



Fig. 3. Similarity maps obtained and their associated p-values. From from left to right and up to down , the traits evaluated are: attractive, trustworthy, caring, responsible, aggressive, sociable, mean, weird, intelligent, confident, unhappy, dominant, emotionally-stable, and threatening. Bright values denote high correlation (and p-value), and darker values low correlation (and p-value)



Fig. 4. Similarity maps for the dominance (2nd Principal Component) and Valence (1st Principal Component)

Principal Component Analysis, and the first two components scores (accounting for 63 and 18% of the variance respectively) have been used as ratings for the proposed 2-Dimensional model of face evaluation. Figure 4 shows the resulting similarity maps obtained from the correlation of these scores (resulting from the dimensionality reduction) and the pixel values of the images. Consistently with the previous findings, dominance shows significant correlation in the chin region ($p\text{-value} < 0.1$), and valence shows positive correlations in the eye-brows ($p\text{-value} < 0.03$) and cheek ($p\text{-value} < 0.1$), which is consistent with the results of the structural descriptor presented in [13].




Fig. 5. Correlation among the 14 facial traits ratings representing the behavioral data and the 14 correlation maps representing the most active facial zones of each trait

4.1 Relative Correlation among Traits

In the second experiment, we explore the relationship between the different facial traits using the behavioral data and the image-based representation. It has been shown that the 14 traits selected are highly correlated, and only two orthogonal dimensions may be sufficient to model face evaluation. This redundancy in the evaluations can be analyzed by computing the correlation among the different trait judgments. We conjecture that similar correlation measures should be found using the behavioral data and the information extracted from the similarity maps.

The experimental evaluation has been performed in two phases, using different data sets at each one. The original database has been randomly split in two independent sets (D_1 and D_2), balancing the presence of male and females at each subset. In the first stage, we used the set D_1 to compute the 14×14 similarity matrix C_b , where each cell (i,j) encodes the correlation between the trait judgments i and j from the behavioral data (mean ratings).

In the second stage, we used the set D_2 to compute the 14×14 similarity matrix C_m , where each cell (i,j) encodes the correlation between the similarity maps M obtained from the traits i and j . This similarity matrix contains the relationships between the different regions of the faces that convey useful information for each trait judgment. Notice that the splitting protocol assures independence in the construction of both similarity matrices, no behavioral data from the set D_1 is used in the computation of the similarity maps (pixel space) from D_2 .

The figure 5 illustrates the resulting 14×14 correlation matrices from the behavioral and pixel domains at the two independent stages. For the sake of clarity, the trait judgments have been sorted in increasing weight order in the PCA decomposition of the mean ratings of the whole image set. Notice that although the similarity between judgments is not exactly the same in magnitude, they are similarly related even when the data sets have been independently extracted.

5 Discussion

The similarity maps developed show significant correlations between facial trait judgments and specific regions of facial images. Although this work is preliminary, and the availability of labeled data is scarce, some of the conclusions are consistent with previous works and also with the intuitive idea of the social appearance we infer for some of these traits.

An important remark must be performed regarding to the gender topic. Some of the correlation maps depicted in figure 3 show strong similarities with female images, and others look more similar to male images (usually related to mean, threatening and aggressive traits). Some of this relations can be found in the behavioral data collected in the facial ratings. Figure 6 shows the mean judgments of each trait using the male images (red) and the female images (blue). Notice that prior to the similarity computation there is a strong bias in the initial ground truth behavioral data. Participants in the study rated faces taking into account the gender topic.

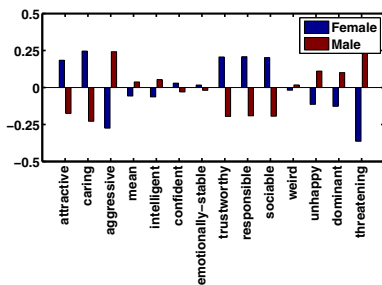


Fig. 6. Mean trait judgements (mean score of the students ratings) of each one of the 14 traits in the male (red) and female (blue) samples

6 Conclusions and Future Work

In this paper, an experimental evaluation of social values inferred from face images is performed. Taking previous works from Psychology as a starting point, and using previous published results on a specific dataset, we used simple similarity measures to evaluate which facial zones convey more information regarding to the facial traits judgments inferred. We propose the use of correlation maps to perform this experimental evaluation. The main contribution of the paper is therefore methodological. Nevertheless, the results of these trait evaluations could be used in multiple applications. The appropriate evaluation of the specific features that make look a face less trustworthy, or more aggressive, could be used in domains where static images should be corrected to improve the inferred first impression. Typical examples could be the pictures in magazines publicity, or improving candidates image in election campaigns.

The paper performs a first exploratory study on the topic. The availability of labeled data determines the statistical significance of the results. Larger image sets (> 500 images) and larger trait evaluations (using > 1000 graduate students for rating) would increase the significance. Nevertheless, we conjecture that the best improvement in the facial zones evaluation could be achieved by means of using eye tracker data [22, 23]. We plan as a future work to acquire a new data set of trait judgements. In addition to the trait ratings, information from the eye tracker should be recorded from each participant and for each trait. Previous works in visual attention and object recognition [24, 25] define specific criteria to evaluate image zones in terms of visual attention. Specifically, the notion of *fixation* is defined as specific zones where eyes are fixed between two saccades, and its position is obtained by weighted sums of pixel locations within fixations (averaged by their duration) [26].

We also plan as a future work to explore more robust hybrid approximations (using both appearance and eye tracking information). Object recognition literature shows successful attempts in relating visual attention and image saliency in the context of object detection and classification [27]. A similar Bayesian approach to take into account both sources of information could be used for the facial trait inference problem.

Acknowledgements. This research was partially supported by MEC grants TIN2009-14404-C02-01 and CONSOLIDER-INGENIO 2010 (CSD2007-00018).

References

1. Zebrowitz, L.: Reading faces: Window to the soul? Westview Press (1997)
2. Montepare, J., Zebrowitz, L.: Person perception comes of age: The salience and significance of age in social judgments. *Advances in Experimental Social Psychology* 30, 93–161 (1998)
3. Eagly, A., Ashmore, R., Makhijani, M., Longo, L.: What is beautiful is good, but: A meta-analytic review of research on the physical attractiveness stereotype. *Psychological Bulletin*; *Psychological Bulletin* 110(1), 109 (1991)
4. Ballew, C., Todorov, A.: Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences* 104(46), 17948 (2007)
5. Little, A., Burriss, R., Jones, B., Roberts, S.: Facial appearance affects voting decisions. *Evolution and Human Behavior* 28(1), 18–27 (2007)
6. Blair, I., Judd, C., Chapleau, K.: The influence of afrocentric facial features in criminal sentencing. *Psychological Science* 15(10), 674 (2004)
7. Willis, J., Todorov, A.: First impressions making up your mind after a 100-ms exposure to a face. *Psychological Science* 17(7), 592–598 (2006)
8. Oosterhof, N., Todorov, A.: The functional basis of face evaluation. *Proceedings of the National Academy of Sciences* 105(32), 11087 (2008)
9. Todorov, A., Said, C., Engell, A., Oosterhof, N.: Understanding evaluation of faces on social dimensions. *Trends in Cognitive Sciences* 12(12), 455–460 (2008)
10. Phelps, E.A., LeDoux, J.E.: Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* 48(2), 175–187 (2005)

11. Adolphs, R., Tranel, D., Damasio, A., et al.: The human amygdala in social judgment. *Nature*, 470–473 (1998)
12. Rojas, Q., Masip, D., Todorov, A., Vitria, J., et al.: Automatic point-based facial trait judgments evaluation. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2715–2720. IEEE (2010)
13. Rojas, M., Masip, D., Todorov, A., Vitria, J.: Automatic prediction of facial trait judgments: Appearance vs. structural models. *PloS one* 6(8), e23323 (2011)
14. Lundqvist, D., Flykt, A., Ohman, A.: Karolinska Directed Emotional Faces: Database of standardized facial images. Psychology Section, Department of Clinical Neuroscience, Karolinska Hospital, S-171 76 Stockholm, Sweden (1998)
15. Torralba, A.: Contextual priming for object detection. *International Journal of Computer Vision* 53(2), 169–191 (2003)
16. Torralba, A.: Modeling global scene factors in attention. *JOSA A* 20(7), 1407–1418 (2003)
17. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(11), 1254–1259 (1998)
18. Simoncelli, E., Freeman, W.: The steerable pyramid: A flexible architecture for multi-scale derivative computation. In: *Proceedings of the International Conference on Image Processing*, vol. 3, pp. 444–447. IEEE (1995)
19. Schyns, P., Oliva, A.: Flexible, diagnosticity-driven, rather than fixed, perceptually determined scale selection in scene and face recognition. *Perception* 26, 1027–1038 (1997)
20. Oliva, A., Schyns, P.: Coarse blobs or fine edges? evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology* 34, 72–107 (1997)
21. Keil, M.S., Lapedriza, A., Masip, D., Vitria, J.: Preferred spatial frequencies for human face processing are associated with optimal class discrimination in the machine. *PloS one* 3(7), e2590 (2008)
22. Rayner, K.: Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin* 124(3), 372 (1998)
23. Rao, R., Zelinsky, G., Hayhoe, M., Ballard, D.: Eye movements in iconic visual search. *Vision Research* 42(11), 1447–1463 (2002)
24. Henderson, J., Pollatsek, A., Rayner, K.: Covert visual attention and extrafoveal information use during object identification. *Attention, Perception, & Psychophysics* 45(3), 196–208 (1989)
25. Kowler, E., Anderson, E., Doshier, B., Blaser, E.: The role of attention in the programming of saccades. *Vision Research* 35(13), 1897–1916 (1995)
26. Henderson, J., McClure, K., Pierce, S., Schrock, G.: Object identification without foveal vision: Evidence from an artificial scotoma paradigm. *Attention, Perception, & Psychophysics* 59(3), 323–346 (1997)
27. Torralba, A., Oliva, A., Castelhan, M.S., Henderson, J.M.: Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychological Review* 113(4), 766 (2006)