Expert Systems with Applications 37 (2010) 5086–5093

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Predicting trait impressions of faces using local face recognition techniques

Sheryl Brahnam [a], Loris Nanni [b,*]

[a] Computer Information Systems, Missouri State University, 901 S. National, Springfield, MO 65804, USA
[b] DEIS, Università di Bologna, Viale Risorgimento 2, 40136 Bologna, Italy

ARTICLE INFO

Keywords:
Trait impressions
Face recognition
Face classification
Human–computer interaction
Overgeneralization effects
Classifier ensembles

ABSTRACT

The aim of this work is to propose a method for detecting the social meanings that people perceive in facial morphology using local face recognition techniques. Developing a reliable method to model people's trait impressions of faces has theoretical value in psychology and human–computer interaction.

The first step in creating our system was to develop a solid ground truth. For this purpose, we collected a set of faces that exhibit strong human consensus within the bipolar extremes of the following six trait categories: intelligence, maturity, warmth, sociality, dominance, and trustworthiness.

In the studies reported in this paper, we compare the performance of global face recognition techniques with local methods applying different classification systems. We find that the best performance is obtained using local techniques, where support vector machines or Levenberg-Marquardt neural networks are used as stand-alone classifiers. System performance in each trait dimension is compared using the area under the ROC curve. Our results show that not only are our proposed learning methods capable of predicting the social impressions elicited by facial morphology but they are also in some cases able to outperform individual human performances.

1. Introduction

Over the last few decades, research in face recognition has expanded its focus to include some of the social aspects of faces, such as facial expression recognition. Emotional information is transitory in nature and is conveyed by deformations of facial features and by the dynamic movements involved in facial muscle contractions. Originally of theoretical interest, machine recognition of facial expressions is an established area of research, and we have witnessed rapid growth over the last decade in the application of this technology in such diverse areas as interface design, lip reading, synthetic face animation, human emotion analysis, face-image compression, security, and teleconferencing, to mention a few (Fasel & Luettin, 2003).

Another social aspect of faces that has sparked considerable interest in social psychology and other fields, but not so much in the area of machine classification, relates to the personality impressions that the static morphology of the face produces in the typical observer. Evidence indicates that people across cultures, age groups, genders, and social-economic status are consistent and similar to one another in their impressions of various facial forms (Zebrowitz, 1998). Several theoretical frameworks have been advanced that attempt to explain why certain facial characteristics consistently elicit specific impressions. One prominent approach is to examine facial appearances in terms of social affordances. In one variant of this theory, impression formation is based on a number of overgeneralization effects (McArthur & Baron, 1983). As outlined in Section 2, significant overgeneralization effects include facial attractiveness, maturity, gender, and emotion. According to this theory, specific facial characteristics advertize important information that lead people to behave in ways that increase their chances of survival. Recognizing a healthy potential mate, for example, increases the chances of producing robust offspring, and facial features indicative of health (rosy cheeks and unblemished skin textures) are generally sexually attractive. It is theorized that the positive effects of these facial characteristics are then overgeneralized so that attractive people, for instance, are assumed to possess a host of other desirable traits, such as intelligence, optimism, extroversion, and leadership and social skills.

In this paper, we attempt to classify faces according to the personality impressions they elicit. This task presents unique difficulties primarily because the ground truth is not based on ascertainable information about the subjects, such as identity, gender, and age. As noted in Sun, Sebe, and Lew (2004), developing a sound ground truth for detecting social aspects of faces is complicated even in the case of facial expression recognition, where research has identified a basic set of emotions that can be recognized using a well-established facial action coding system

* Corresponding author.
E-mail addresses: sbrahnam@missouristate.edu (S. Brahnam), loris.nanni@unibo.it (L. Nanni).

0957-4174/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.eswa.2009.12.002

(Ekman & Friesen, 1978). Lacking a strong theoretical foundation for building a ground truth for the social impressions of facial morphology, our classes must be rooted in the impressions a set of faces make on the average observer. We are thus asking machines to map faces to a set of subjective social and psychological categories.

Few face classification experiments have focused on matching consensual subjective judgments. One early attempt is Brahnam (2002). In that study, principle component classifiers were trained to match the human classification of faces along the bipolar trait dimensions of adjustment, dominance, warmth, sociality, and trustworthiness. Results varied depending on the trait dimension; the classifiers matched the average observer in dominance (64%), adjustment (71%), sociality (70%), trustworthiness (81%) and warmth (89%). We believe, however, that these classification rates were inflated by the low number of faces (<15) in each trait class.

Related to our work are experiments designed to detect facial attractiveness. In Kanghae and Sornil (2007), for instance, two sets of photographs of 92 Caucasian females, approximately 18 years old, were rated by subjects using a 7-point scale. In their experiments, they classified faces into two classes, *attractive* (highest 25% rated images) and *unattractive* (lowest 25% rated images). Performance, based on percentage of correctly classified images, averaged 76.0% using *K*-nearest neighbor and 70.5% using linear support vector machines (SVMs). In this case, the authors believed that the low number of faces in the two classes ($\sim$24 images each) accounted for the poor performance of the classifiers.

As illustrated in Brahnam (2002) and Kanghae and Sornil (2007), it is important to develop a database that contains a sufficient number of representative faces within each trait class. In Section 3, we describe the databases first used in Brahnam and Nanni (2009) and in this study that attempt to circumvent some of the difficulties mentioned above in developing a solid ground truth for this problem domain. For the studies reported in Brahnam and Nanni (2009), six databases of faces were produced for the following six trait dimensions: intelligence, maturity, dominance, sociality, trustworthiness, and warmth. Each database was divided into two classes: *high* and *low*. Because artists were asked to design the faces, the number of faces verified by human subjects as belonging in each trait class averaged 111 faces, a much higher number of faces in each class then those used above. Single classifiers and ensembles were then trained to match the bipolar extremes of the faces in each of the six trait dimensions. With performance measured by the Area Under the ROC curve (AROC) and averaged across all six dimensions, results showed that single classifiers, especially linear SVM (0.74) and Levenberg-Marquardt neural networks (LMNNs) (0.73), performed as well as human raters (0.77). These single classifiers, however, performed poorly in the trait dimension of maturity. Ensembles of 100 LMNNs, constructed using Bagging (BA) (0.76), Random Subspace (SUB) (0.77), and Class Switching (CS) (0.75) compared equally well to rater performance but were better than the single classifiers at handling all six trait dimensions.

This study proposes using local methods for this classification task. It is well known in the biometric literature that methods based only on the global information of the images are not very effective with differences in facial expressions, illumination conditions, and poses. In the last few years, several face recognition methods based on local information have been proposed (Gottumukkal and Asari, 2004; Nanni and Maio, 2007; Shan et al., 2006; Tan and Chen, 2005). In Shan et al. (2006), the authors proposed using ensembles of classifiers, where the face image is first divided into *M* subimages (with all subimages having the same size without overlapping) according to their spatial locations. *M* feature vectors were then obtained by extracting the Local Binary Pattern (LBP) from the Gabor convolved subimage. A Fisher Discriminant Analysis (FDA) was then used to project the Gabor features onto a lower dimensional space. Finally, a nearest neighbor (NN) was performed in each space, and the results were combined.

LBP is a well known descriptor for textures and has been used for representing faces (Ojala, Pietikainen, & Maeenpaa, 2002; Shan et al., 2006). In Shan et al. (2006) it is shown that if the LBP features are extracted from the Gabor convolved image, it is possible to obtain better performance with respect to that obtained by extracting the LBP features from the gray level image.

The remainder of this paper is organized as follows. In Section 2, we provide background information regarding this problem domain by summarizing the characteristics and traits associated with the overgeneralization effects of facial maturity, attractiveness, emotion, and gender. In Section 3, we present the study design and our method of generating the stimulus faces and of evaluating subject ratings. In Section 4, we describe our system architecture, and in Section 5 we compare the performance of simple classifiers and global classifier ensembles to the performance of different local methods. Our tests show that the local methods are more stable in performance across all six trait dimensions and closely approximate, if not exceed, the performance of individual human raters.

2. Social impressions of faces

In this section we provide some background material regarding the social impressions of facial morphology from the perspective of social psychology. As reviewed in Zebrowitz (2006), several theories have been advanced to explain why people form impressions of facial morphology that are fairly consistent. No one theory provides a complete understanding of this process. Several prominent social psychologists involved in facial impression research take an ecological approach by examining facial appearances in terms of their social affordances. As noted in the introduction, one important elaboration of the ecological approach is based on a number of overgeneralization hypothesis (McArthur & Baron, 1983). Research has shown that the structural characteristics of these effects are associated with specific sets of personality traits. As an introduction to this problem domain, we provide an overview in this section of the morphological characteristics and associated trait impressions of the following overgeneralization effects: facial attractiveness, maturity, gender, and emotion.

2.1. Overgeneralization effect of attractiveness

Attractive people are sexually desirable. Although debated, there is some evidence that attractiveness broadcasts actual indicators of physical, emotional, and intellectual fitness (Jones, 1995; Zebrowitz, Hall, & Murphy, 2002). The facial characteristics associated with attractiveness include proportion (Alley & Hildebrandt, 1988), symmetry (Grammer & Thornhill, 1994), straightness of profile (Carello, Grosofsky, & Shaw, 1989), closeness to the population average (Langlois & Roggman, 1990), and sex specific features (Alley, 1992).

The overgeneralization effect of attractiveness hypothesizes that attractive people, being more sexually desirable, would be attributed more positive character traits. Much evidence supports what has come to be called the attractiveness *halo effect*. Attractive people are generally considered more socially adroit (Kuhlenschmidt & Croger, 1988), better leaders (Cherulnik, 1989), healthier, psychologically more adapted (Langlois, Kalakanis, & Rubenstein, 2000), intellectually more capable, and more moral than those less attractive. As would be expected, unattractive faces are associated with a host of negative traits, and they elicit more negative behaviors from people. Unattractive people are thought to be