Available online at www.sciencedirect.com
ScienceDirect
Vision Research 48 (2008) 235–243
Vision Research
www.elsevier.com/locate/visres

A machine learning predictor of facial attractiveness revealing human-like psychophysical biases

Amit Kagian [a], Gideon Dror [b], Tommer Leyvand [a], Isaac Meilijson [c], Daniel Cohen-Or [a], Eytan Ruppin [a,d,*]

[a] School of Computer Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel
[b] School of Computer Sciences, The Academic College of Tel-Aviv-Yaffo, Tel-Aviv 64044, Israel
[c] School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel
[d] School of Medicine, Tel-Aviv University, Tel-Aviv 69978, Israel

Received 16 June 2007; received in revised form 28 September 2007

Abstract

Recent psychological studies have strongly suggested that humans share common visual preferences for facial attractiveness. Here, we present a learning model that automatically extracts measurements of facial features from raw images and obtains human-level performance in predicting facial attractiveness ratings. The machine's ratings are highly correlated with mean human ratings, markedly improving on recent machine learning studies of this task. Simulated psychophysical experiments with virtually manipulated images reveal preferences in the machine's judgments that are remarkably similar to those of humans. Thus, a model trained explicitly to capture a specific operational performance criteria, implicitly captures basic human psychophysical characteristics.

Keywords: Face perception; Facial attractiveness; Machine learning; Aesthetics; Computational neuroscience

1. Introduction

Philosophers, artists and scientists have been trying to capture the nature of beauty since the early days of philosophy. Although in modern days a common layman's notion is that judgments of beauty are a matter of subjective opinion alone, recent findings suggest that people share a common taste for facial attractiveness and that their preferences may be an innate part of our primary constitution. Several experiments have shown that 2–8 months old infants prefer looking at faces rated by adults as more attractive (Langlois et al., 1987). In addition, attractiveness ratings show very high agreement between groups of raters belonging to the same culture and even across cultures (Cunningham, Roberts, Wu, Barbee, & Druen, 1995). Such findings give rise to the quest for common factors which determine human facial attractiveness. Accordingly, various hypotheses, from cognitive, evolutional and social perspectives, have been put forward to describe and interpret the common preferences for facial beauty.

Inspired by Sir Francis Galton's photographic method of composing faces (Galton, 1878), Langlois and Rogmann have created averaged faces by morphing multiple images together. Human judges found these averaged faces to be attractive and rated them with attractiveness ratings higher than the mean rating of the component faces composing them, proposing that averageness is the answer for facial attractiveness (Langlois & Roggman, 1990; Rubenstein, Langlois, & Roggman, 2002). Investigating symmetry and averageness of faces, Grammer and Thornhill concluded that symmetry was more important than averageness in facial attractiveness (Grammer & Thornhill, 1994). Other

* Corresponding author. Address: School of Computer Sciences and School of Medicine, Tel-Aviv University, Tel-Aviv 69978, Israel. Fax: +972 3 640 9357.
E-mail addresses: amit.kagian@gmail.com (A. Kagian), ruppin@post.tau.ac.il (E. Ruppin).

0042-6989/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2007.11.007

studies have agreed that average faces are attractive but claim that faces with certain extreme features, such as extreme sexually dimorphic traits, may be more attractive than average faces (Little, Penton-Voak, Burt, & Perrett, 2002). Yet other researchers have suggested various conditions which may contribute to facial attractiveness such as neonate features, pleasant expressions and familiarity (Zebrowitz & Rhodes, 2002). Finally, Cunningham et al. have suggested a multiple fitness model in which there is no single constructing line that determines attractiveness (e.g. perception of fitness as implying an ideal romantic partner). Instead, different categories of features signal different desirable qualities of the perceived target (Cunningham, Barbee, & Philhower, 2002). Even so, the multiple fitness model agrees that some facial qualities are universally physically attractive to people.

Apart from eliciting the facial characteristics which account for attractiveness, modern researchers have aimed to describe the mechanisms underlying these preferences. Many contributors refer to the evolutionary origins of attractiveness preferences (Andersson, 1994; Møller & Swaddle, 1997; Thornhill & Gangsted, 1999). According to this view, facial traits signal mate quality and imply chances for reproductive success and parasite resistance. Some evolutionary theorists suggest that preferred features might not signal mate quality but that the "good taste" by itself is an evolutionary adaptation (individuals with a preference for attractiveness will have attractive offspring that will be favored as mates) (Thornhill & Gangsted, 1999). Another mechanism explains attractiveness' preferences through a cognitive theory—a preference for attractive faces might be induced as a by-product of general perception or recognition mechanisms (Rubenstein et al., 2002; Zebrowitz & Rhodes, 2002): attractive faces might be pleasant to look at since they are closer to the cognitive representation of the face category in the mind. Halberstadt and Rhodes have further demonstrated that not just average *faces* are attractive but also birds, fish, and automobiles become more attractive after being averaged with computer manipulation (Halberstadt & Rhodes, 2003). Such findings led researchers to propose that as perceivers can process an object more fluently, aesthetic response becomes more positive (Reber, Schwarz, & Winkielman, 2004). A third view suggests that facial attractiveness originates in a social mechanism, where preferences may be dependent on the learning history of the individual and even on his social goals (Zebrowitz & Rhodes, 2002).

Other studies have used computational methods to analyze facial attractiveness. In several cases faces were averaged using morphing tools (e.g. Perrett, May, & Yoshikawa, 1994; Rubenstein et al., 2002). Laser scans of faces were put into complete correspondence with the average face in order to examine the relationship between facial attractiveness, age and averageness (ÓToole, Price, Vetter, Bartlett, & Blanz, 1999). A genetic algorithm, guided by interactive user selections was programmed to evolve a "most beautiful" female face (Johnston & Franklin, 1993). Machine learning methods have been used recently to investigate whether a machine can predict attractiveness ratings by learning a mapping from facial images to their attractiveness scores (Eisenthal, Dror, & Ruppin, 2006). The latter predictor achieved a correlation of 0.6 with average human ratings, demonstrating that facial beauty can be learned by a machine, at least to some moderate extent. However, as human raters significantly outperform the predictor of Eisenthal et al., the challenge of constructing a facial attractiveness machine predictor with human-level accuracy has remained open.

A primary goal of this study is to surpass these results by developing a machine which obtains human-level performance in predicting facial attractiveness and, thus, passes what Kurzweil calls a *subject matter expert turing test* (SME TT) (Kurzweil, 2005). Having accomplished this, our second main goal is to conduct a series of simulated psychophysical experiments and study the resemblance between human and machine judgments. This latter task carries two potential rewards: first, to determine whether the machine can aid in understanding the psychophysics of human facial attractiveness, capitalizing on the ready accessibility of manipulating and studying its performance, and second, to study whether learning an explicit operational ratings prediction task also entails learning implicit human-like biases, at least for the case of facial attractiveness.

In the past decades machines have achieved human-level performance in rule-based systems such as playing games (Schaeffer & Herik, 2002) and in various expert systems (Slezak, 1991). Impressive progress has been displayed in simulating various tasks which involve face perception, such as face detection (Hjelmas & Low, 2001), face recognition (Becker, 1999; Zhao, Chellappa, Rosenfeld, & Phillips, 2000) and tasks of facial category learning such as emotion (Dailey, Cottrell, Padgett, & Adolphs, 2002) and gender (Graf, Wichmann, Bülthoff, & Schölkopf, 2006) recognition. The task of evaluating human attractiveness ratings adds the notion of *judgment of taste* to the previous achievements in machine perception of faces. Learning the concept of facial attractiveness could form an important demonstration of a computer's ability to learn to master a quantitative, basic, human judgment task.

To this end we have collected human scores of facial attractiveness for a given dataset of female facial images. We developed an algorithm for automatic extraction of a very large set of geometric facial features, which, combined with a set of global features, yields a principled representation of each facial image via a set of image-features in an appropriate dimension-reduced space. Using this data of facial representations and their associated rating scores, we have employed standard supervised learning algorithms to construct a facial attractiveness prediction machine. Given a new, unseen face, this machine predicts its human attractiveness score in an accurate manner. We then turned to performing a series of simulated psychophysical experiments, modeled after known experiments in the psycholog-