Neuropsychologia 48 (2010) 3596–3605

Contents lists available at ScienceDirect

Neuropsychologia

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

The amygdala and FFA track both social and non-social face dimensions

Christopher P. Said[a], Ron Dotsch[b], Alexander Todorov[c,*]

[a] Department of Psychology and Center for Neural Science, New York University, New York, NY, United States
[b] Department of Psychology, Behavioural Science Institute, Radboud University Nijmegen, The Netherlands
[c] Department of Psychology and Princeton Neuroscience Institute, Princeton University, Green Hall, Princeton, NJ 08540, United States

ARTICLE INFO

Article history:
Received 18 May 2010
Received in revised form 5 August 2010
Accepted 9 August 2010
Available online 18 August 2010

Keywords:
Amygdala
Face perception
FFA
OFA
pSTS
Social cognition

ABSTRACT

The amygdala is thought to perform a number of social functions, and has received much attention for its role in processing social properties of faces. In particular, it has been shown to respond more to facial expressions than to neutral faces, and more to positively valenced and negatively valenced faces than faces in the middle of the continuum. However, when these findings are viewed in the context of a multidimensional face space, an important question emerges. Face space is a vector space where every face can be represented as a point in the space. The origin of the space represents the average face. In this context, positively valenced and negatively valenced faces are further away from the average face than faces in the middle of the continuum. It is therefore unclear if the amygdala response to positively valenced and negatively valenced faces is due to their social properties or to their general distance from the average face. Here, we compared the amygdala response to a set of faces that varied along two dimensions centered around the average face but differing in social content. In both the amygdala and much of the posterior face network, we observed a similar response to both dimensions, with stronger responses to the extremes of the dimensions than to faces near the average face. These findings suggest that the responses in these regions to socially relevant faces may be partially due to general distance from the average face.

1. Introduction

The amygdala has been implicated in a number of social functions (Adolphs & Spezio, 2006; Adolphs, Tranel, & Damasio, 1998; LeDoux, 2007; Morris et al., 1996). For example, patients with lesions in the amygdala have problems identifying expressions of fear (Adolphs, Gosselin, Buchanan, Tranel, Schyns, & Damasio, 2005; Adolphs, Tranel, Damasio, & Damasio, 1994), feel comfortable invading the personal space of other people in dyadic social interactions (Kennedy, Glascher, Tyszka, & Adolphs, 2009), and judge faces that appear to most people untrustworthy as trustworthy (Adolphs et al., 1998). Consistent with the human findings, monkeys with experimentally induced amygdala lesions demonstrate uninhibited social interaction (Amaral, 2003). Human functional neuroimaging studies have provided a wealth of data supporting the importance of the amygdala in social perception (Breiter et al., 1996; Canli, Sivers, Whitfield, Gotlib, & Gabrieli, 2002; Costafreda, Brammer, David, & Fu, 2008; Cunningham, Van Bavel, & Johnsen, 2008; Pessoa, Japee, Sturman, & Ungerleider, 2006; Whalen et al., 2004; Whalen, Rauch, Etcoff, Mclnerney, Lee, & Jenike, 1998; Winston, O'Doherty, & Dolan, 2003; Winston, O'Doherty, Kilner, Perrett, & Dolan, 2007).

Following human lesion studies (Adolphs et al., 1998; Todorov & Duchaine, 2008), there have been a number of functional neuroimaging studies implicating the amygdala in social judgments from faces (Engell, Haxby, & Todorov, 2007; Said, Baron, & Todorov, 2009; Todorov, Baron, & Oosterhof, 2008; Todorov & Engell, 2008; Winston, Strange, O'Doherty, & Dolan, 2002). Most of these studies used judgments of trustworthiness. For example, Engell et al. (2007) used such judgments to predict brain responses to faces in a task that did not require an explicit evaluation of the faces. Nevertheless, the amygdala response increased with decreases in the perceived trustworthiness of faces. While the initial studies primarily reported a negative linear response in the amygdala, most recent studies have found a quadratic non-monotonic response (Said, Baron, et al., 2009; Said, Haxby, & Todorov, submitted for publication; Todorov, Said, Oosterhof, & Engell, submitted for publication; Todorov, Baron, & Oosterhof, 2008). Faces that are highly untrustworthy and highly trustworthy elicit the strongest responses, while faces near the middle of the continuum elicit the weakest responses. The same response function was also observed in the inferior temporal cortex. As we show in the present study, this apparent contradiction in the literature may be due to differences in the stimulus properties of the faces used in the respective studies.

* Corresponding author.
E-mail address: atodorov@princeton.edu (A. Todorov).

0028-3932/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neuropsychologia.2010.08.009

It is experimentally useful to measure face trustworthiness, because this trait is an excellent approximation of the valence evaluation of faces (Oosterhof & Todorov, 2008; Todorov, Pakrashi, & Oosterhof, 2009). Therefore, one interpretation of the neuroimaging findings implicating the amygdala in trustworthiness evaluation is that the amygdala is specifically tracking facial properties that define face valence.

However, when these findings are viewed in the context of a multidimensional face space, some important questions emerge. Face space is a vector space where each dimension can be thought of as a physical property of faces, and every face can be represented as a point in the space (Valentine, 1991). The origin of the space represents the average face. According to the model of face trustworthiness proposed by Oosterhof and Todorov (2008), the trustworthiness of a face near the average face can be increased maximally by moving it in one direction in face space, and decreased maximally by moving it in the opposite direction. In this respect, the finding that the amygdala responds more strongly to highly trustworthy and untrustworthy faces than to faces in the middle of the continuum is confounded by the fact that highly trustworthy and highly untrustworthy faces are further away from the average face than faces near the middle of the continuum. Indeed, electrophysiology and fMRI studies have shown that the fusiform response increases with distance from the average face (Leopold, Bondar, & Giese, 2006; Loffler, Yourganov, Wilkinson, & Wilson, 2005). Therefore, it is unknown if observations about the trustworthiness dimension can be attributed specifically to facial properties that convey specific social signals, or if they are instead due to general distance from the average face, regardless of the dimension.

In this fMRI experiment, we compare the valence response profile to the response profile for a control dimension that is perceived to be less socially relevant but is matched on face distance to the valence dimension. As described in the methods section, the valence dimension was obtained from a principal components analysis (PCA) of nine different social judgments of faces (see Table S6 in Oosterhof & Todorov, 2008). The control dimension was selected from a large number of randomly generated dimensions that were orthogonal to all social dimensions. To compare the responses for the valence and control dimensions, we use both a whole brain approach and a region of interest (ROI) approach. Specifically, we targeted the amygdala, the fusiform face area (FFA), the occipital face area (OFA), and the face-selective regions of the posterior superior temporal sulcus (pSTS), as these have all been implicated in trustworthiness judgments or general face processing (Haxby, Hoffman, & Gobbini, 2000; Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Gore, & Allison, 1997). We expected to find a larger quadratic response to valence than to the control dimension in the amygdala and the FFA (Fig. 1).

There are many physical and psychophysical metrics that can be used to measure the distance between faces. It is therefore impossible to find a control dimension in which the range is matched to the valence range on all possible metrics. For metrics in which the ranges are unmatched, it is most conservative to have a smaller range for valence. This provides a stringent test of our hypothesis, as any effect driven by general distance along that metric will be stronger for the other dimension.

A series of preliminary experiments were used to measure the properties of the valence dimension and the control dimension. First, we show that there is less change in perceived trustworthiness, threat, and dominance for the control dimension than the valence dimension. Second, we show that for all the metrics we tested, the range of the control dimension used in the fMRI experiment was either matched to the range of the valence dimension, or unmatched in a conservative direction. Two of the metrics were tested experimentally. A third metric, which could be measured directly on the 3D meshes that defined the shape of our stimuli, was the average vertex displacement of the face mesh.

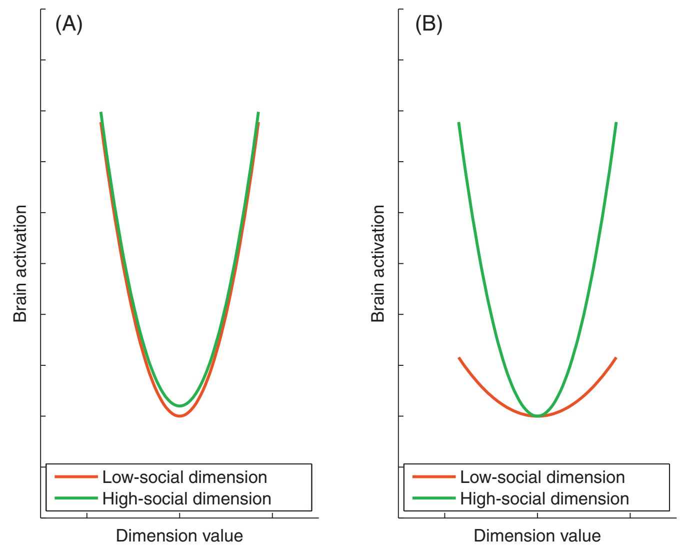

Fig. 1. Possible outcomes. (A) A similar quadratic response to both the high-social dimension (valence) and the low-social dimension (control). (B) A larger quadratic response for the high-social dimension (valence) than the low-social dimension (control).

2. Preliminary experiments

2.1. Stimuli

Faces were generated with FaceGen software and custom code. FaceGen defines face shape using a 50-dimensional face space, where each dimension is a component from a PCA performed on the 3D face vertex positions–defined on a face mesh–of a large sample of laser-scanned human faces. Because the top 50 components account for most of the shape variance, any face can be reasonably approximated as a point in this space. Only face shape was manipulated; the reflectance properties were held fixed. Oosterhof and Todorov (2008) obtained trait ratings for a large number of faces sampled from this space. A separate PCA performed on these trait ratings revealed that more than 54% of the variance could be explained by the first principal component, which can be referred to as valence (see Table S6 in Oosterhof & Todorov, 2008). Trustworthiness judgments were highly correlated with this component (>.90) even when the component was estimated without trustworthiness judgments in the PCA. Next, to build dimensions corresponding to social judgments, Oosterhof and Todorov performed a multiple regression with judgments as the dependent variable and the 50 face shape dimensions as the predictor variables. The same approach was used for building a model of face valence. Specifically, the first principal component derived from the PCA of social judgments was regressed on the shape dimensions. The valence dimension was then defined as the vector of coefficients from this regression. This vector can be added to any face in order to change its predicted value on valence. Under the assumptions of the linear model, a unit change along this dimension is expected to result in a maximal change in the valence of the face.

The control dimension was chosen as a dimension in a face space orthogonal to the valence dimension. First, we randomly generated 100 face dimensions that were orthogonal to valence and 9 other social face dimensions (e.g., threat, competence, extraversion, etc.)