

THE PERCEPTION AND EVALUATION OF VISUAL BEAUTY

Teresa K. Pegors

A DISSERTATION

in

Psychology

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

2014

Supervisor of Dissertation

Russell A. Epstein, Professor, Psychology

Graduate Group Chairperson

John C. Trueswell, Professor, Psychology

Dissertation Committee

Geoffrey K. Aguirre, Associate Professor, Neurology (Committee Chair)

Joseph W. Kable, Assistant Professor, Psychology

Anjan Chatterjee, Professor, Neurology

THE PERCEPTION AND EVALUATION OF VISUAL BEAUTY

COPYRIGHT

2014

Teresa Kristine Pegors

DEDICATION

This dissertation is dedicated to mother, Karen Pegors, who passed away just before I began the doctoral program. It was because of her unconditional love and her willingness to allow me great freedom in following my interests that I have made it this far in my academic pursuits. Even though I was not able to share the trials and joys of this most recent journey along with her, she was the one who taught me the quiet strength and perseverance that I needed to travel on without her.

ACKNOWLEDGMENTS

I am deeply indebted to so many people who have helped me through this difficult but rewarding journey. First of all, I want to acknowledge and thank my advisor Russell Epstein, who encouraged and supported me even when my research questions began to fall outside the scope of the lab's "standard" research areas. I am profoundly grateful for his mentorship and for his ability to give extremely thoughtful and constructive feedback on ideas, writing, and presentations. Not only did I learn how to be a better scientist, but I also became a much better writer and thinker under his guidance. I also want to thank the rest of my committee for showing great interest in my work, letting me sit in on their lab meetings, and always making the time to meet with me when I had questions or ideas. I cherished the collaborative relationship I had with my committee, and I am thankful to all of them for their willingness to make that possible. Finally, I want to thank all of my labmates and the Center for Cognitive Neuroscience community. I couldn't have asked for a better group of people with which to work.

ABSTRACT

THE PERCEPTION AND EVALUATION OF VISUAL BEAUTY

Teresa K. Pegors

Russell A. Epstein

What are the perceptual and cognitive processes that underlie our experiences of beauty? In this dissertation, I describe a series of experiments where we used functional magnetic resonance imaging (fMRI) and behavioral methods to explore the mechanisms of perception, reward representation, and decision-making during evaluations of face and place beauty. In our first study, we used fMRI to ask whether evaluative signals in frontal cortex contain category-specific information or whether these signals are encoded as a “common currency” across reward types. By comparing neural activity correlated with subjective ratings of face and place beauty, we showed overlapping activity in dorsal ventromedial prefrontal cortex (vmPFC), consistency with the common currency hypothesis. At the same time, our results revealed category-specific patterns of activity in ventral vmPFC and in lateral orbitofrontal cortex (latOFC), suggesting at least a partial distinction in the frontal networks recruited during the processing of different types of rewards. In a follow-up study, we used fMRI to further examine face-responsive “patches” of activity in latOFC by measuring response in these patches while subjects evaluated but did not explicitly rate face beauty. Our results demonstrated a similar pattern of response to that observed during explicit ratings, suggesting that reward-related activity in this region is not dependent on a decision-making task. Lastly, in a series of behavioral studies, we developed a novel experimental design to measure the influence of recent trial history on current judgments of face attractiveness. We found that

attractiveness judgments are simultaneously contrasted away from the attractiveness of the previous face but assimilated towards the previous numerical rating given. Our results also suggested that these influences are not specific to attractiveness judgments but may be linked to more general properties of perception and decision-making. Collectively, this work furthers our understanding of the neural mechanisms underlying evaluations of face and place beauty, and illuminates some of the specific contextual influences on these evaluations.

TABLE OF CONTENTS

DEDICATION	III
ACKNOWLEDGMENTS	IV
ABSTRACT	V
LIST OF TABLES	IX
LIST OF FIGURES	X
CHAPTER 1 – INTRODUCTION	1
FACE BEAUTY	2
<i>Visual features of attractive faces</i>	2
<i>Reward value of attractive faces</i>	4
<i>Neural correlates of face attractiveness</i>	6
PLACE BEAUTY	7
<i>Visual features of beautiful places</i>	7
<i>Reward value of beautiful places</i>	8
<i>Place beauty in the brain</i>	10
RESEARCH AIMS	10
CHAPTER 2 - COMMON AND UNIQUE REPRESENTATIONS IN PREFRONTAL CORTEX FOR FACE AND PLACE ATTRACTIVENESS	13
ABSTRACT	13
INTRODUCTION	14
METHODS	16
<i>Subjects</i>	16
<i>fMRI Acquisition</i>	16
<i>Stimuli</i>	17
<i>Design and Procedure</i>	18
<i>fMRI Data Analysis</i>	19
RESULTS	24
<i>Behavioral Results</i>	24
<i>Activity for face and place attractiveness in frontal cortex</i>	25
<i>Distributed category-specific encoding in vmPFC</i>	28
<i>Whole brain analysis</i>	29
DISCUSSION	31
<i>Common response to face and place attractiveness in vmPFC</i>	31
<i>Category-specific attractiveness coding in vmPFC</i>	33
<i>Face-specific activity in right latOFC</i>	34
<i>Visual region differences</i>	36
<i>Conclusions</i>	37
CHAPTER 3 – FACE-RESPONSIVE ACTIVITY IN LATERAL ORBITOFRONTAL CORTEX	46
ABSTRACT	46
INTRODUCTION	47
METHODS	49
<i>Subjects</i>	49
<i>fMRI Acquisition</i>	49

<i>Stimuli</i>	49
<i>Design and Procedure</i>	51
<i>fMRI Data Analysis</i>	51
RESULTS.....	54
<i>Face attractiveness response in frontal cortex</i>	54
<i>Mean response to faces > places in frontal cortex</i>	54
<i>Anatomical specificity of face-related responses in latOFC</i>	55
DISCUSSION.....	56
<i>Face attractiveness in latOFC and vmPFC</i>	57
<i>Faces > places activity in latOFC and vmPFC</i>	59
<i>Face Specificity</i>	61
<i>Conclusions</i>	61
CHAPTER FOUR - SIMULTANEOUS PERCEPTUAL AND RESPONSE BIASES ON SEQUENTIAL FACE ATTRACTIVENESS JUDGMENTS.....	66
ABSTRACT.....	66
INTRODUCTION.....	67
<i>Contrastive Sequential Biases</i>	68
<i>Assimilative Sequential Biases</i>	69
<i>Simultaneous Stimulus and Response Biases</i>	70
EXPERIMENT 1.....	71
Methods.....	72
Results and Discussion.....	75
EXPERIMENT 2.....	79
Methods.....	80
Results and Discussion.....	81
EXPERIMENT 3.....	83
Methods.....	84
Results and Discussion.....	88
GENERAL DISCUSSION.....	89
<i>Sequential Contrast Bias</i>	90
<i>Sequential Assimilation Bias</i>	91
<i>Relevance to Sequential Tasks</i>	92
<i>Conclusion</i>	93
CHAPTER 5 – FUTURE DIRECTIONS.....	97
QUESTIONS FOR FURTHER EXPLORATION.....	97
THE PERCEPTUAL VERSUS COGNITIVE NATURE OF THE SEQUENTIAL CONTRAST EFFECT.....	99
BIBLIOGRAPHY.....	102

LIST OF TABLES

2.1 Clusters of response for face and place attractiveness in vmPFC and latOFC	40
2.2 Wholebrain results for face and place main effects, attractiveness effects	43

LIST OF FIGURES

2.1 Face and place attractiveness experimental design	38
2.2 Histogram of ratings distributions and RT across ratings scales.....	39
2.3 Cross-subject validation results for univariate group analysis	41
2.4 Multivertex pattern analysis of attractiveness response in vmPFC.....	42
2.5 Wholebrain maps for face and place attractiveness in visual cortex	44
2.6 Relationship between functionally distinct face patches in latOFC	45
3.1 Face attractiveness response during passive attractiveness evaluation.....	63
3.2 Face main effect in passive evaluation and 1-back matching task.....	64
3.3 Anatomical ROI analysis for faces responses in latOFC.....	65
4.1 Face attractiveness/hair darkness alternating task design	94
4.2 Regression results for face/place alternating design	95
4.3 Wholebrain maps for face attractiveness and face attractiveness adaptation.....	96

CHAPTER 1 – Introduction

Questions about the nature of beauty have been around since the beginning of philosophy itself: *What is it that makes a thing beautiful and why? Is beauty universal or cultural? What is the relationship between beauty and goodness, truth, and desire?* While some questions are purely philosophical, researchers in psychology and neuroscience have sought to address empirical questions related to the cognitive and neural underpinnings of our experiences of beauty. Much progress has been made, but there is still tremendous work to be done in understanding how mechanisms such as perception, memory, and evaluation all work together to make up these “aesthetic” experiences (Chatterjee 2004; Conway and Rehding 2013).

In this dissertation, I describe a series of experiments that used functional magnetic resonance imaging (fMRI) and behavioral methods to better understand how we perceive beauty, specifically the beauty of faces and places. We chose these categories of beauty for two main reasons: First, there is no research on the neural correlates of beautiful landscapes. We therefore wanted to establish the set of regions across the brain that specifically responded to this kind of place beauty. Second, while it is true that both faces and places are often given the label of “beautiful” or “attractive,” beautiful faces and beautiful places have very different visual properties and are associated with different kinds of motivation and reward. Comparing neural and behavioral responses between both of these categories, therefore, would allow us to observe which beauty-related processes were recruited across the two categories, and which processes were tied to the visual or reward properties of one particular category.

As background, the subsequent sections present an overview of the current research on face and place beauty. I describe what is known about the visual features and reward properties thought to be associated with both, and I also review what is known about face and place beauty in the brain. Following these sections, I then outline the specific research aims of this dissertation.

Face beauty

Visual features of attractive¹ faces

Research on face attractiveness is a relatively new field, because until recently, it was thought that the perceived attractiveness of a face was completely dependent on culture, upbringing, and individual idiosyncrasies. But studies in the last few decades have highlighted at least three facial characteristics that seem to be universally tied to attractiveness: averageness, symmetry, and sexual dimorphism (Thornhill and Gangestad 1999; Rhodes 2006).

Average faces are defined as those that are close to the mathematical center of the space of all face features in a given population. In a first demonstration that average faces are judged to be attractive, Langlois and Roggman had subjects rate the attractiveness of individual faces and of composite faces that were created by averaging together a large number of individual faces (1990). Their result, that composites were rated as significantly more attractive than individual faces, was surprising, but follow-up studies where researchers using more advanced means of creating composites showed that these average faces were not rated as more attractive solely because composites had smoother skin textures (Rhodes and Tremewan 1996; O'Toole et al. 1999) or because they were more symmetrical (Rhodes, Sumich, and Byatt 1999; Valentine, Darling, and Donnelly 2004). One

¹ Throughout this dissertation, I use the words "attractiveness" and "beauty" interchangeably.

study also provided evidence that the that the averageness of a face is not based on a fixed norm, but instead, people seem to dynamically update what they consider “average” based on the recent set of faces observed in the world: even over the course of a few minutes of exposure to expanded or contracted faces, people will shift their judgments of what they considers attractive to match the average face in the current distribution (Rhodes et al. 2003).

The left-right symmetry of a face, independent of averageness, has also been shown to robustly correlate with facial attractiveness, both for naturally symmetric faces (Grammer and Thornhill 1994), and in comparisons of normal faces to faces in which the left and right halves were mirror reflections. (Rhodes et al. 1998; Perrett et al. 1999). Additionally, one study used monozygotic twins (genetically but not developmentally identical) to show that the level of perceived difference in facial symmetry between the twins was correlated with the level of perceived attractiveness. The twin who had the more symmetrical face was judged to have the more attractive face (Mealey, Bridgestock, and Townsend 1999).

While averageness and symmetry are attractive traits for both male and female faces, researchers have also shown some evidence that dimorphic features (secondary sex characteristics) correlate with facial attractiveness. The degree of femininity in female faces (e.g. higher and wider eyes, small nose, and small chin) is positively correlated with attractiveness ratings given by males, and this was true regardless of whether researchers directly quantified feminine features for each face or simply asked subjects to rate the femininity of each face (Cunningham 1986; Cunningham et al. 1995; O’Toole et al. 1998; Perrett, Lee, and Penton-Voak 1998; Rhodes, Hickford, and Jeffery 2000). The relationship between masculinity in male faces and attractiveness is more complex. A number of studies have suggested that

the degree of masculinity in a male face is actually negatively correlated with attractiveness (Perrett, Lee, and Penton-Voak 1998; Rhodes, Hickford, and Jeffery 2000; Little and Hancock 2002). On the other hand, one insightful study revealed that masculinity is preferred in females who were in the high-risk (for ovulation) phase of their menstrual cycle, whereas at all other times, females preferred more feminized male faces (Johnston et al. 2001).

Reward value of attractive faces

The previous section outlined a number of facial features associated with attractiveness across cultures, but why do humans have a preference at all for one type of face over another? What are the true or perceived rewards associated with attractive faces?

One line of inquiry has used an evolutionary framework to ask whether face attractiveness functions to signal reproductive fitness. For example, facial attractiveness may reflect genetic stability, typical development, and lack of sickness or parasites. Evidence, though, for a direct link between attractiveness and health is weak: one of the most robust studies of this question looked at health records and attractiveness across individuals' lifespans, but they found no relationship between health and attractiveness for either females or males (Kalick and Zebrowitz 1998). Only by examining those individuals below the median in attractiveness, was there found a moderate relationship between attractiveness and health (Zebrowitz and Rhodes 2004). The authors hypothesized that facial attractiveness functions to signal low fitness, but it is not helpful in identifying high fitness. Other evidence suggests that the degree of masculine features in male faces does have a low correlation with health and reproductive fitness (Soler et al. 2003; Rhodes et al. 2003; Rhodes,

Simmons, and Peters 2005). Conversely, feminine features of female faces do not seem to be correlated with a variety of health measures (Rhodes et al. 2003).

Many researchers concede that actual reproductive fitness may no longer be strongly associated with attractiveness due to overriding factors such as modern medicine and nutrition (Thornhill and Gangestad 1999; Rhodes 2006). Nevertheless, it is very well known that people *perceive* attractive individuals to have a whole host of positive traits such as health, intelligence, sexual responsiveness, and sociability (for reviews, see Eagly et al. 1991; Feingold 1992; Langlois et al. 2000). Zebrowitz formalized this idea into what she calls the “anomalous face overgeneralization hypothesis,” in which humans have adapted to recognize facial features that signal low fitness, but these responses are also applied to normal individuals whose “faces resemble those who are unfit” (Zebrowitz and Montepare 2008). By way of contrast, then, attractive people are perceived to have more positive traits than unattractive people (Griffin and Langlois 2006).

Another line of inquiry has looked into whether preferences for attractive faces are by-products of more general information processing mechanisms. For example, it has been argued that symmetry aids recognition (Enquist and Arak 1994; Johnstone 1994), and might therefore be a generally positive trait (Rhodes 2006). Also, quite a number of studies have shown that average or “prototypical” items are preferred over non-prototypes, demonstrating similar effects for average fish, watches, cars, musical compositions, and voices (Smith and Melara 1990; Repp 1997; Halberstadt and Rhodes 2000, 2003; Bruckert et al. 2010). This phenomenon may reflect a preference for familiarity or a more general preference for efficient information processing, or processing “fluency” (Reber, Scharz, and Winkielman 2004).

Neural correlates of face attractiveness

Face perception as a whole has been a topic of great interest to cognitive neuroscientists, and much research has gone into understanding the “face-processing” network (Haxby, Hoffman, and Gobbini 2002; Haxby and Gobbini 2011; Collins and Olson 2014). Many studies have also directly tested for the neural correlates of facial attractiveness, showing both visual and reward-related regions to positively track this trait.

A number of studies showed that response in fusiform cortex correlated with subjects’ individual attractiveness judgments (Kranz and Ishai 2006; Cloutier 2008; Winston et al. 2007). Additionally, at least two studies have shown that activity in visual cortex persists even when subjects are making orthogonal judgments to attractiveness: Chatterjee et al. tested visual cortex response to facial attractiveness while subjects made either attractiveness judgments or identity judgments, and showed a positive response to attractiveness in an extended region of visual cortex, including the fusiform face area (FFA) and lateral occipital cortex (LOC) (2009). Winston et al. also showed similar positive response in occipito-temporal and fusiform cortex, during both attractiveness and gender judgments (2007). It remains unclear from these studies whether these effects in visual cortex are driven by certain visual features or whether these effects are driven by attention or saliency signals.

A meta-analysis of the neural correlates of attractiveness and trustworthiness showed that the most commonly recruited regions are the ventromedial prefrontal cortex (vmPFC), nucleus accumbens (NAcc) and anterior cingulate cortex (ACC) (Mende-Siedlecki, Said, and Todorov 2012). Individual studies of face attractiveness have also consistently reported attractiveness responses in these regions (e.g.

O'Doherty 2003; Cloutier et al. 2008; Kranz and Ishai 2006; Kim et al. 2007). NAcc and vmPFC are interconnected with each other (Ongur and Price 2000) and are considered be part of a larger valuation circuit (Kable and Glimcher 2009). Kim et al. showed that the NAcc was activated earlier than vmPFC for preferred vs. nonpreferred faces, but this differential activity for preferred faces in NAcc only occurred for novel faces (2007). The authors suggest that, based on these results, the NAcc is responsible for forming an automatic, initial affective evaluation of faces, whereas the vmPFC represents information related to the subjects' preference decision.

Place Beauty

Visual features of beautiful places

Places span a much wider range of visual characteristics than faces, and can easily be subdivided into visually distinct categories such as indoor/outdoor and urban/natural. In the following sections, we emphasize landscapes in particular, which we define as places without obvious manmade influence or built structures.

Landscapes as a whole are considered more beautiful than urban environments (Kaplan, Kaplan, and Wendt 1972; Ulrich 1984; Purcell and Lamp 1994). This preference is largely tied to the "naturalness" of an environment, the degree to which a place has natural vs. manmade/built elements. For example, in a review of vegetation in urban environments, Swardon concludes that across studies, urban vegetation has been clearly shown to have both functional and aesthetic value (1988). Additionally, White et al. showed that water elements combined with greenery are considered more beautiful than images with only greenery or only water, and that "built" environments without either are least preferable (2010).

Other research on broad landscape characteristics is sparser. There is some evidence that the subjective complexity of landscapes correlates with beauty ratings (Han 2009), but the relationship between beauty and most landscape qualities such as subjective familiarity, "coherence", "mystery", and spaciousness seems to depend on the type of landscape (Herzog 1985). For instance, Herzog showed that the spaciousness of mountains, canyons, and deserts is positively correlated with beauty (1981), but varied between waterscapes (e.g. seascapes, lakes, swamps) (1985).

Reward value of beautiful places

Evolutionary theorists suggest that place preferences may be due to a natural affinity for the environment or "biome" in which humans evolved, though there is disagreement about whether this may have been the savannah, forest, or grasslands/woodlands (Han 2007). Appleton's prospect/refuge theory is based on the savannah hypothesis of human evolution, and argues that humans have adapted to be attracted to environments that have broad, open vistas, access to refuge (e.g. trees) and resources (water and vegetation) (1975). Balling and Falk showed a preference for savannahs in young American children but no preference for savannahs in American adults (1982). In a more recent study, the same authors tested environmental preferences in Nigerian children from the age of 12-18. Like the American children, Nigerian children also rated savannah images as significantly more preferable than the images of forests and deserts (Falk and Balling 2009). The authors use both of these studies to suggest that children across cultures have an in-born instinct for savannahs which is then overshadowed by experience and cultural norms over the course of development.

Another line of inquiry into place preference has focused on the information-theoretical aspects of beautiful landscapes. A number of studies have shown that

natural scenes in particular have fractal-like, or scale-invariant, properties (Hagerhall, Purcell, Taylor 2004; Redies, Hasentstein, and Denzler 2007; Graham and Field 2007). These fractal properties found in natural scenes are surprisingly similar to the statistical properties found in visual artwork (Taylor and Micolich 1999; Redies, Hasenstein, and Denzler 2007; Graham and Field 2007). This has led to some speculation that these visual statistical properties themselves are factors that contribute to aesthetic preferences (Hagerhall et al. 2004; Redies 2007), and that such preferences may occur due to the fact that scale-invariance of the input leads to efficient encoding within the visual system (e.g. Simoncelli and Olshausen 2001).

The most explored theory of place preference has been to show that exposure to natural landscapes leads to quantifiable physical and mental health benefits (Ulrich 1984; Kaplan 1995; Kweon et al. 2007; Berman, Jonides, and Kaplan 2012; White et al. 2013). Ulrich compared hospital patients with window views looking out onto trees had window views looking out at a brick wall. Those who had views of trees were discharged from the hospital faster and requested less pain medication than patients who had views onto a brick wall (1984). A large-scale study using data from household and health surveys showed a positive relationship between green space and well-being, and a negative relationship between green space and mental distress, after controlling for other possible factors (White et al. 2013). The ideas that natural environments are beneficial to our health and well-being are generally described in evolutionary terms (e.g. Wilson 1984), but these effects may also be due a cultural phenomenon in which a large portion of the population associates urban/suburban settings with work, stress, and normal life, whereas natural settings become associated with rest and "getting away."

Place beauty in the brain

Place images recruit a separate core set of regions in the brain than those for face images (Epstein and Kanwisher 1998). While this general “place-processing” network is well studied, unlike the plethora of studies on face beauty, only two studies have explored the neural correlates of place beauty in the brain, and these both focused on built/indoor settings rather than landscapes.

Kirk et al. had experts (architects) and non-experts give aesthetic ratings in the scanner to a range of buildings and faces (2009). While the behavioral ratings of the images between experts and non-experts were not significantly different, there was greater beauty response in anterior cingulate cortex (ACC) and ventromedial prefrontal cortex (vmPFC) in the experts. Interestingly, there were no regions that responded similarly to beauty across groups in a conjunction analysis. In a study to determine the impact of interior contours on preference and neural response, Vartanian et al. scanned subjects while they judged the beauty of room interiors (2013). The authors showed that visual cortex, specifically the middle occipital gyrus, co-varied with place beauty, but there was no significant response to place beauty in ACC or vmPFC (though ACC did response to place pleasantness).

Across the two studies, regions were found that were similar to those associated with face beauty: the ACC and vmPFC in frontal cortex as well as regions of visual cortex. While this is suggestive of neural similarities between face and place beauty processing, the lack of overlap between the two place studies warrants further exploration and comparative study.

Research Aims

The work described in this dissertation had two overarching goals: first, to explore in more detail the neural and behavioral processes related to face and place

beauty, and second, to use our results to inform more general questions related to perception, reward processing, and decision-making. By using the basic task of evaluating face and place beauty across a variety of designs and methodologies, we attempted to address both of these goals within each of our studies.

In chapter 2, I describe an fMRI experiment in which we compared neural response to face and place beauty in perceptual and reward-related regions of the brain. A number of studies have suggested that the ventromedial prefrontal cortex (vmPFC) encodes a wide range of subjective values (e.g. money, food, trinkets) into a common value signal or “common currency.” Because it is thought that this signal is used to allow for choice comparisons between different kinds of goods, we tested whether such a signal exists even in the case of values that are not typically compared: face and place beauty. Our results supported the common currency hypothesis by showing overlapping response in vmPFC to face and place beauty. At the same time, we also found evidence for category-specific encoding of beauty in both reward and visual regions of the brain. These results extended our understanding of value representation in the reward system, and also gave us insight into the similarities and differences in how face and place beauty is represented in the brain.

In chapter 3, I describe another fMRI experiment that explored the task conditions under which face-related signals exist in lateral orbitofrontal cortex (latOFC). In a previous study (chapter 2), we showed that a region of activity in right latOFC responded to faces but not places, and a separate region in right latOFC responded to face beauty but not place beauty. By scanning subjects while they passively viewed face and place images, we tested whether these signals persisted even when people were only passively evaluating beauty without making explicit

ratings. Our results showed that both these regions showed the same response profile as when subjects made explicit ratings. In contrast, ventromedial prefrontal cortex (vmPFC), which had responded to face beauty during explicit judgments, did not respond during passive evaluation, suggesting a functional distinction between value-related processes in latOFC and vmPFC.

In chapter 4, I describe a series of behavioral studies in which we tested how context affects judgments of face attractiveness. While previous research has shown that ratings of attractiveness are affected by previous trial history, there is contradictory evidence over whether this affect is “contrastive” or “assimilative” in nature. We created a novel behavioral design that allowed us to independently measure the influence of the previous image and the previous rating on current attractiveness judgments. We demonstrated the existence of simultaneous and opposing influences of the previous stimulus and the previous rating on judgments of attractiveness, though these influences were restricted within visual category (place trials did not influence face trials). We also provide evidence that these influences are not specific to attractiveness judgments, but may be a general property of sequential judgments.

CHAPTER 2 - Common and unique representations in prefrontal cortex for face and place attractiveness

Pegors TK; Kable JW; Chatterjee A; Epstein RA. *under review*

Abstract

Although previous neuroimaging research has identified overlapping correlates of subjective value across different reward types in the ventromedial prefrontal cortex (vmPFC), it is not clear whether this “common currency” evaluative signal extends to the aesthetic domain. To examine this issue, we scanned human subjects with fMRI while they made attractiveness judgments of faces and places, two stimulus categories that are associated with different underlying rewards, have very different visual properties, and are rarely compared to each other. We found overlapping signals for face and place attractiveness in the vmPFC, consistent with the idea that this region codes a signal for value that applies across disparate reward types and across both economic and aesthetic judgments. However, we also identified a subregion of vmPFC within which activity patterns for face and place attractiveness were distinguishable, suggesting that some category-specific attractiveness information is retained in this region. Finally, we observed two separate functional regions in lateral orbitofrontal cortex (latOFC), one region that exhibited a category-unique response to face attractiveness and another region that responded strongly to faces but was insensitive to their value. Our results suggest that vmPFC supports a common mechanism for reward evaluation while also retaining a degree of category-specific information, whereas latOFC may be involved in basic reward processing that is specific to only some stimulus categories.

Introduction

Aesthetic evaluations, such as judging beauty and attractiveness, play an important role in our lives, affecting (for example) whom we choose to marry and where we would prefer to live. We can perform aesthetic evaluations on a wide variety of stimuli – a face can be beautiful or ugly, as can a house, room, or landscape. This suggests the possibility that a common evaluative signal might underlie different aesthetic judgments, analogous to the common signal that is believed to underlie different kinds of economic judgments. However, a counterargument is the fact that aesthetic judgments are rarely made across stimulus category (“Is this face more beautiful than this landscape?”). This makes them different from economic judgments, which usually involve comparison across disparate categories of goods, and suggests that a “common currency” for value might not be useful—or calculated—in the aesthetic domain.

Here we test whether human prefrontal cortex computes a common currency for value that is used for aesthetic judgments. Previous studies have shown that fMRI response in ventromedial prefrontal cortex (vmPFC) during economic decision-making corresponds to the “utility” or “subjective value” of a stimulus (Bartra, McGuire, and Kable 2013). This response exhibits some degree of domain-general, insofar as it can be elicited by both money and consumer goods (Chib et al. 2009; Kim et al. 2011; Levy et al. 2011). However, the few studies that have attempted to compare non-economic (e.g. social) to monetary rewards have produced contradictory findings, sometimes showing overlapping activity in vmPFC and sometimes revealing non-overlapping activity or no activity at all in this region (Smith et al. 2010; Lin et al. 2012, Sescousse et al. 2010, Izuma et al. 2008). It

remains unclear whether a common evaluative signal exists that could apply to judgments very far removed from economic exchange such as aesthetic judgments.

To examine this issue, we scanned subjects with functional magnetic resonance imaging (fMRI) while they made attractiveness judgments of faces and places. Our aim was to identify a common neural signal of value that operates across these two very different non-economic stimulus classes. By using faces and places as stimuli, we intended to give our two categories the “best advantage” in terms of possible neural differences. Beautiful faces and beautiful places are associated with distinct reward outcomes: beautiful faces offer the promise of reproductive success and social advancement (see Rhodes 2006 for a review), whereas beautiful landscapes offer the promise of prospect/refuge, physical resources, and rest (Appleton 1975; White et al. 2010; Berman & Kaplan 2008). Faces and places also differ substantially in their visual properties, to the extent that they are processed by distinct regions of visual cortex (Kanwisher et al. 1997; Epstein and Kanwisher 1999). Face and place attractiveness are not typically assigned monetary values (although it is not impossible to assign monetary value to something without a market value - see Smith et al. 2010, Mitchell and Carson 1989). Therefore, a common response for face and place attractiveness is unlikely in the absence of a common currency-type, domain-general evaluative neural signal.

Previous work has shown that activity in vmPFC correlates with face attractiveness (O’Doherty et al. 2003; Ishai 2007; Cloutier et al. 2008), though this is not always the case (Chatterjee et al. 2009). A single study on the neural correlates of indoor place attractiveness showed no activity in vmPFC that correlated with parametric beauty ratings (Vartanian et al. 2013). Consequently, the question of whether face and place attractiveness are represented in the same or different

brain regions remains open. To anticipate, our results suggest that there are both category-general and category-specific representations of attractiveness in vmPFC, and a signal specific to face attractiveness in lateral orbitofrontal cortex (latOFC).

Methods

Subjects

34 healthy, right-handed volunteers with normal or corrected-to-normal vision were recruited to participate in the fMRI study. All subjects gave written informed consent according to procedures approved by the University of Pennsylvania institutional review board. One subject was excluded due to an incidental finding, two subjects were excluded due to technical issues with their anatomical images, and three subjects were excluded due to excessive motion in the scanner. This left us with 28 subjects whose functional data were analyzed (14 females, mean age 22.5).

fMRI Acquisition

Scans were performed at the Hospital of the University of Pennsylvania on a 3T Siemens Trio scanner equipped with a Siemens body coil and a 32-channel head coil. High resolution T1-weighted anatomical images were acquired using a 3D MPRAGE pulse sequence (TR = 1620 ms, TE = 3 ms, TI = 950 ms, voxel size = 0.9766 x 0.9766 x 1 mm, matrix size = 192 x 256 x 160). T2*-weighted images sensitive to blood oxygenation level-dependent (BOLD) contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR = 3000 ms, TE = 30 ms, voxel size = 3x3x3mm, matrix size = 64 x 64, 46 axial slices). The slices were tilted -30 degrees from the AC-PC plane to reduce signal dropout (Deichmann et al. 2003).

Stimuli

Stimuli were digitized 400x400 pixel color photographs of 144 places and 144 faces chosen to span a wide range of attractiveness. The places were natural environments with no man-made elements, spanning scene types such as swamps, mountain ranges, beaches, fields, and forests. The face set had equal numbers of males and females, and all faces were Caucasian, upright and forward-facing, with neutral to pleasant expressions, selected from the Glasgow Unfamiliar Face Database (<http://homepages.abdn.ac.uk/m.burton/pages/gfmt/Glasgow%20Face%20Recognition%20Group.html>), Radboud Database (Langner et al. 2010), the Center for Vital Longevity Face Database (Minear & Park 2004), CVL Face Database (Peter Peer, <http://www.lrv.fri.uni-lj.si/facedb.html>), Diana Theater Face Database (courtesy of Dr. Robert Schultz at the Center for Autism Research), and online searches. Face images were extracted from their original background, blurred slightly along the edges, cropped so that hair did not extend below the chin, and resized to span a height of 400 pixels. They were then placed on an abstract colored background created by phase-scrambling a single place image. This ensured that they subtended the same visual angle as the places while retaining a similar background color. A unique scrambling of the background image was used for each face (code used from: <http://visionscience.com/pipermail/visionlist/2007/002181.html>). These 288 stimuli were chosen from a larger set of 573 face and place images based on pilot testing intended to ensure that they covered a wide range of attractiveness, thus maximizing our ability to see neural activity related to this variability. In these pilot tests, 10 subjects made Likert-scale ratings (1-8) of the visual attractiveness of each place and face. Images were blocked by category (face/place) and subjects used the keyboard to make ratings at their own pace, with instructions to spread

their judgments across the whole 1-8 range. Each subjects' ratings were then z-scored across all images and these z-scores were averaged across subjects for each image. The images were then divided into "low", "average", and "high" attractiveness bins according to whether they had a z-score below -0.5 (low), between -0.5 and 0.5 (average), or above 0.5 (high). Images were then chosen in equal numbers from these three attractiveness bins to make the final stimulus set.

Design and Procedure

The fMRI experiment consisted of six 4 min 57 sec scan runs, each of which was divided into two 36-s face blocks, two 36-s place blocks, and two 36-s fixation (or "rest") blocks in which subjects passively fixated on a central crosshair (Figure 2.1). Between each block was an additional 9 seconds of passive fixation. 18 seconds of fixation were added at the beginning and end of each run to allow the T2* signal to reach a steady state and to model the final HRF, respectively. During each face and place block, subjects used a button box to give "low", "average", or "high" attractiveness ratings to 12 faces or places, each presented for 1 s followed by a 2 s interstimulus interval (ISI) during which only a crosshair was on the screen. Subjects made their ratings any time within the 3-second trial; button assignment was counterbalanced across subjects. Blocks were ordered such that no block type repeated twice in a row (including fixation blocks), and the block orders were counterbalanced across runs.

To acclimate subjects to the distribution of attractiveness in the stimulus set, they were asked to rate prior to the scan the attractiveness of 24 faces and 24 places not used in the main experiment on a 1-8 scale. These 48 images were chosen to span the entire range of attractiveness found in the main stimulus set. In addition, immediately after the scan session, subjects were presented with the

images from the main experiment again in a random order and asked to rate them on the same 1-8 scale. These post-scan ratings were used in the subsequent fMRI analysis because they provided a finer-grained measure of face and place attractiveness than the high/average/low ratings made in the scanner.

Experiment trials within blocks were ordered such that the place and face attractiveness regressors used in the subsequent analyses would have maximal power to detect variability in fMRI response. Although these regressors were ultimately based on the attractiveness judgments made by the subjects in the scanner, which could not be known in advance, we were able to obtain a rough estimate of their shape by using the high/average/low attractiveness ratings provided by the pilot subjects. We generated 10,000 random orders of face and place stimuli that fit our experimental design, convolved the corresponding attractiveness ratings with a canonical hemodynamic response function (HRF), and determined the power by calculating the ratio of variance before and after convolution. A “best” sequence of attractiveness levels was chosen (the one with the highest power), and images corresponding to these bins were drawn randomly from the stimulus set to create a unique stimulus sequence for each subject.

fMRI Data Analysis

Pre-processing and data analysis for individual subjects was performed using the FMRIB Software Library (FSL v.4.1.6) (Jenkinson et al. 2012; Woolrich et al. 2009; Smith et al. 2004). Functional images were corrected for differences in slice time acquisition and then de-obliqued to correct for the 30 degree tilt slice acquisition. For each run, the first six volumes were removed to account for the fMRI signal not yet reaching steady-state, and data were then motion corrected by spatially realigning each image with the central image in the run, registered to the

subject-specific T1-weighted image using 6 degrees of freedom rigid-body transformations, and high-pass filtered to remove temporal frequencies below 0.0074 Hertz.

General linear modeling was used to estimate neural activity correlated with subject-specific ratings of face and place attractiveness. Each subjects' post-scan 1-8 ratings were used to create parametric regressors for face and place attractiveness which extended across all six experimental scans. The model also included regressors corresponding to face and place in-scan reaction times (RTs), as our behavioral data revealed that reaction time was significantly negatively correlated with place attractiveness (but not face attractiveness). We did not include quadratic regressors for face and place attractiveness, in contrast to previous studies (e.g. Winston et al. 2007), because these quadratic regressors were strongly negatively correlated with RT. Finally, categorical regressors were added for face trials, place trials, and instruction screens, and nuisance regressors were added to account for between-scan variability and outliers (outliers calculated with the Gabrieli Lab's Artifact Detection Tools: http://www.nitrc.org/projects/artifact_detect/). All regressors, except scan indicators, were convolved with a canonical HRF.

We used this model to perform two sets of analyses. The first was a set of targeted analyses focusing on regions in the frontal lobe that have been previously implicated in the processing of information about stimulus value. The second was a whole-brain random effects analysis intended to find areas responding to face and place attractiveness without any *a priori* hypotheses about where these areas would be. For the targeted frontal lobe analyses, unsmoothed parameter estimates from the contrasts of interest were registered to the cortical surface using surface templates derived from each subject's T1-weighted anatomical image using

Freesurfer's segmentation function (recon-all). The data were then averaged across subjects by spherically registering these subject-specific surfaces to the group brain (an average of the subject-specific surfaces, where vertex coordinates are calculated as the average Talairach coordinates of that vertex across subjects). During this registration process the data was smoothed on the group-average surface at 3mm FWHM. Using the unfolded cortical surface for inter-subject registration in this manner reduces the variability when averaging across volumetric data associated with regions containing highly variable cortical folding patterns, which was especially important in this case because orbitofrontal cortex is known to have a high degree of anatomical variability between subjects (Chiavaras & Petrides 2000; Chiavaras et al. 2001). Random-effects analyses were then performed on the contrasts of interest to identify regions within vmPFC and latOFC that responded to face and place attractiveness. Output was cluster-corrected for small-volumes in vmPFC and latOFC and Bonferroni corrected to account for observations across 2 hemispheres (clusters defined at $p < 0.05$ uncorrected and then permutation corrected to $p < 0.05$ using Freesurfer's simulation function to estimate the distribution of maximum cluster sizes under the null hypothesis). The medial surface *a priori* small volume was defined using a functional mask for vmPFC (Bartra, McGuire, & Kable 2013) which we translated into surface space. As there is some evidence of category-specific signals in latOFC (e.g. Sescousse et al. 2010), we also used this area as an *a priori* small volume, defined by using the ventral surface of the "lateralorbitofrontal" ROI from Freesurfer's APARC library (taken from the Desikan-Killiany atlas).

We performed a leave-one-subject-out iterative cross-validation analysis (Kriegeskorte et al. 2009, supplementary discussion) on the significant clusters resulting from the targeted frontal lobe analysis to test whether activity within any

clusters responding to face attractiveness or place attractiveness responded differentially to face versus place attractiveness. On each iteration, data from a single subject was held out as the test set. Clusters within vmPFC and latOFC responding to face or place attractiveness were then defined based on a group analysis of the n-1 remaining subjects, using the procedures described above. The response in these clusters in the nth subject was then analyzed, and the procedure repeated. This method gave us an estimate of the response in each cluster to face and place attractiveness, using independent data sets to define the boundary of the cluster and the strength of the effect.

For the whole-brain analyses, pre-processed data were smoothed with a 5mm FWHM kernel and parameter estimates for regressors of interest were obtained for each voxel for each subject. These were normalized to standard volumetric MNI152 space using linear 12 degree of freedom transformations, resampled to 2x2x2 mm voxels in this standard space, and submitted to group level random effects analyses for contrasts of interest. The true Type 1 error rate for each contrast was calculated from FSL's randomise function using Monte-Carlo simulations that permuted the signs of wholebrain data from individual subjects (10,000 relabelings; method based on Freedman & Lane 1983). The resulting reported voxels are significant at $p < 0.05$ corrected for multiple comparisons across the whole brain.

We also performed a follow-up analysis which examined response in functional ROIs in visual cortex. Because we did not conduct independent localizer scans for all subjects, these ROIs were defined by using a set of 40 subjects' localizer contrast files (19 which came from our current study). These group-defined "parcels" were created using an algorithmic method which is fully described in Julian et al. (2012) (We diverged from the Julian et al. description in that we chose a more

liberal threshold of $p < 0.001$ for the Face>Objects contrast maps.) Parcels were defined using the contrasts Faces>Objects (fusiform face area, FFA), Scenes>Objects (parahippocampal place area, PPA), Objects>Scrambled Objects (posterior lateral occipital cortex, LOC) Scrambled Objects > Objects (early visual cortex, EVC). For the ROI analysis, we then translated individual subject contrast maps for face and place attractiveness into standard space and extracted parameter estimates for each subject within these group-defined ROIs.

Our univariate analyses revealed clusters of activity in vmPFC that responded similarly to both face and place attractiveness. To test whether *patterns* of response rather than overall *mean* response within these clusters might reflect category-specific attractiveness information, we performed a pattern classification analysis. Using well-established methods (Haxby et al., 2001), we split the data into independent halves (each consisting of 3 of the 6 scan runs), identified activity patterns for face and place attractiveness in each half, and then compared these patterns across halves to establish their reliability and distinguishability. Activity patterns in each half were calculated based on the same general linear model described above, except that the regressors spanned 3 scans instead of 6. The resulting unsmoothed maps were then registered to the group-defined surface-space, which allowed us to perform the classification analysis within the same independently defined, leave-one-subject-out vmPFC clusters that were used in our univariate analysis. Parameter estimates (beta values) were extracted for each subject's 4 independently defined vmPFC clusters, and classification was run iteratively over all possible split halves of the data (10 ways in which 6 scans can be split into two groups of 3). Classification was considered successful if the average Pearson correlation between the face attractiveness patterns (or place attractiveness

patterns) in opposite halves of the data (within-category comparison) was greater than the correlation between face and place attractiveness patterns in opposite halves of the data (cross-category comparison) (Haxby et al. 2001). Raw correlation difference scores on which classification accuracy was based (e.g. face to face – face to place) were also calculated. Both the accuracy and correlation difference scores were then compared to chance (50%) and zero, respectively, to determine if category-specific information was present in any of the vmPFC clusters.

Results

Behavioral Results

Within-scan ratings of face and place attractiveness were strongly correlated with post-scan ratings (Pearson's r averaged across subjects for faces = 0.74, $t(27)=45.04$, $p<0.0001$ and for places = 0.71, $t(27)=35.17$, $p<0.0001$; p -values reflect repeated-measures t -tests on correlation scores), and there was no significant difference between categories regarding the degree of correlation between these within-scan and post-scan ratings (repeated-measures t -test on the difference between correlation scores: $t(27)=-1.4$, $p=0.17$). Post-scan ratings for face and place attractiveness showed extremely high levels of consistency across subjects (Cronbach's alpha for faces: 0.958; places: 0.956). The within- and between-subjects consistency of attractiveness ratings confirms the validity of using the finer-scale post-scan ratings to analyze the fMRI response.

Subject-specific means of face and place attractiveness ratings were not significantly different, although there was a trend for places to have higher mean ratings than faces ($t(27)=1.94$, $p=0.06$). Figure 2.2a displays histograms of post-scan face and place attractiveness ratings. Post-scan ratings for places were significantly negatively correlated with in-scan response times (Pearson's r averaged

across subjects = -0.24, repeated-measures t-test on correlation scores: $t(27) = -9.07$, $p < 0.0001$). In contrast, there was no significant correlation between post-scan face ratings and in-scan response times ($r = -0.01$, $t(27) = -0.27$, $p = 0.79$). By plotting response time as a function of attractiveness, visual inspection revealed that rather than a linear function, RT exhibited an inverted-U shaped relationship with both face and place attractiveness, peaking in the center range of attractiveness (Figure 2.2b). RTs for face judgments were highly symmetrical around the mean, and therefore were not correlated with face attractiveness ratings, whereas RTs for place judgments had an elongated linear slope for the upper half of the attractiveness scale. For each item, we also calculated the mean RT and the standard deviation of rating judgments across subjects. These measures were moderately correlated for both face and place stimuli (Pearson's r for face stimuli: 0.43, and place stimuli: 0.57). In other words, RTs were slower when there was greater disagreement among subjects about the rating of an image, suggesting that RT may in part reflect the degree to which a subject is uncertain about their rating.

Activity for face and place attractiveness in frontal cortex

We then turned to the primary question of interest: whether there was overlap between regions responding to face and place attractiveness in the frontal lobes. To answer this, we looked for effects of place and face attractiveness within frontal regions known *a priori* to exhibit subjective value signals: ventromedial prefrontal cortex (vmPFC) and lateral orbitofrontal cortex (latOFC). We conducted the analyses in surface-space to better account for the large variability in cortical folding patterns along the ventral surface of frontal cortex. We focused on two (non-overlapping) regions.

Our vmPFC ROI was functionally-defined based on Barta and colleagues' meta-analysis of subjective value responses in the brain (Bartra, McGuire, & Kable 2013). Within this ROI, we observed clusters that responded linearly to face attractiveness at $p < 0.05$ (corrected for small volumes) in both hemispheres ("LvmPFC-face" and "RvmPFC-face"; see Table 2.1 and Figure 2.3) and clusters that responded linearly to place attractiveness in both hemispheres ("LvmPFC-place" and "RvmPFC-place"). Visual inspection makes clear that there is a great degree of overlap between the face and place clusters in each hemisphere. The face attractiveness clusters were found in similar locations to activity correlated with face attractiveness in previous studies (e.g. O'Doherty et al. 2003, Ishai et al. 2007; Cloutier et al. 2008) (Figure 2.3).

To assess whether any of these clusters responded selectively to attractiveness for a single stimulus category, we performed a cross-validation analysis in which we iteratively defined the clusters on $n-1$ subjects and then extracted parameter estimates for the "left-out" subject. This provides a stringent test of whether a cluster responded to both kinds of attractiveness: although each cluster was defined based on its response to either face or place attractiveness, it was tested for its response to the other category in an independent data set. Repeated-measures t-tests comparing face to place attractiveness revealed that all vmPFC clusters showed a significant response to both face and place attractiveness (LvmPFC-place response to face att. $t(27)=2.11$, $p=0.04$, response to place att. $t(27)=4.09$, $p=0.0004$; LvmPFC-face response to face att. $t(27)=4.19$, $p=0.0003$, response to place att. $t(27)=3.70$, $p=0.001$; RvmPFC-place response to face att. $t(27)=3.36$, $p=0.002$, response to place att. $t(27)=3.60$, $p=0.001$; RvmPFC-face response to face att. $t(27)=4.17$, $p=0.0003$, response to place att. $t(27)=2.18$,

$p=0.04$; statistics obtained by t-tests comparing extracted mean parameter estimates against zero). Importantly, all clusters were equally sensitive to both face and place attractiveness regardless of how they were initially defined (Face attractiveness > place attractiveness: LvmPFC-face $t(27)=0.09$, $p=0.93$, RvmPFC-face $t(27)=1.04$, $p=0.30$, LvmPFC-place $t(27)=-1.39$, $p=0.18$, RvmPFC-place $t(27)=-0.34$, $p=0.73$, all n.s., see Figure 2.3; all statistics were obtained by repeated-measures t-tests comparing extracted mean parameter estimates across categories.) These results suggest that vmPFC responds similarly to variation in attractiveness for these two categories; that is, clusters in vmPFC sensitive to face attractiveness are also sensitive to place attractiveness, and vice-versa. Note that this is the case even though the overall response in these regions was greater to faces than to places (Faces > places in LvmPFC-face $t(27)=3.51$, $p=0.002$, RvmPFC-face $t(27)=4.52$, $p=0.0001$; LvmPFC-place $t(27)=4.44$, $p=0.0001$, RvmPFC-place $t(27)=3.18$, $p=0.004$).

Our second region of interest was the lateral orbital-frontal cortex (latOFC). This region has been implicated in subjective value representation of distinct categories of reward (e.g. Sescousse et al. 2010), though value-based response is observed here less frequently than in vmPFC. Our analysis revealed a cluster in the right hemisphere whose response correlated with face attractiveness ($p<0.05$ cluster-wise permutation corrected for small-volumes, see Table 2.1 and Figure 2.3). We again used the cross-validation analysis to determine whether the attractiveness response in this cluster was category-specific. A repeated-measures t-test found significantly higher response for face attractiveness compared to place attractiveness ($t(25)=2.64$, $p=0.01$). Indeed, place attractiveness response in the cluster was not

significantly different from zero (t-test, $t(25)=-0.46$, $p=0.65$). Of note, the overall response to faces and places did not differ in this region ($t(25)=0.08$, $p=0.94$).

Distributed category-specific encoding in vmPFC

Our univariate analyses revealed that clusters in vmPFC that have significant mean response to one category of attractiveness also have a significant mean response to the other category, and that the strength of these responses are not significantly different from each other. While these results are in line with the “common currency” hypothesis for evaluative signals in vmPFC, it remains possible that these brain regions contain separate but intermixed valuation mechanisms for faces and places, which were not discriminable when responses were averaged over all voxel in the cluster. We tested for this possibility within each vmPFC cluster by measuring vertex-wise activation patterns for face and place attractiveness in separate halves of the data and examining whether we could classify face vs. place attractiveness across the split. Classification accuracy for the right hemisphere cluster defined by place attractiveness (RvmPFC-place) was above chance (accuracy = 58%, $t(27)=2.63$, $p=0.014$; see Figure 2.4). Accuracy in this cluster was similar for both categories, though within-place accuracy was just above the threshold for significance (face-face vs. face-place = 58%, $t(27)=2.5$, $p=0.02$; place-place vs. face-place = 58%, $t(27)=2.0$, $p=0.055$). RvmPFC-face also showed a trend toward significant classification (accuracy = 55.8%, $t(27)=1.9$, $p=0.068$) though neither individual category was significant on its own (face-face vs. face-place = 55.9%, $t(27)=1.66$, $p=0.11$; face-face vs. face-place = 55.7%, $t(27)=1.55$, $p=0.13$). Given the fact that RvmPFC-face overlaps to a large degree with RvmPFC-place (see Figure 2.3) but it also extends further anterior, these results suggest that the locus of category-specific information is centered within the RvmPFC-place. Accuracy was

not significant for LvmPFC-place (accuracy = 50.3%, $t(27)=0.1$, $p=0.92$) or LvmPFC-face (accuracy = 46.1%, $t(27)=-1.33$, $p=0.20$). We also examined the raw correlation scores themselves (upon which the accuracy measures were made) by calculating the difference scores between within- and cross-category Pearson correlations. These difference scores were significantly different from zero only within the same RvmPFC-place cluster that showed significant place vs. face attractiveness classification (face/face – face/place difference score=0.08, $t(27)=2.08$, $p=0.05$; place/place – face/place difference score=0.08, $t(27)=2.25$, $p=0.03$).

Whole brain analysis

We next looked for regions whose response correlated with face attractiveness and place attractiveness outside of our *a priori* regions in frontal cortex. We observed bilateral response in fusiform gyrus that was positively correlated with face attractiveness, as well as a response in right intraparietal sulcus (See Table 2.2). In contrast, no attractiveness-related activity for places survived wholebrain corrections, though we observed sub-threshold activity in posterior cingulate, ventral striatum, vmPFC, and in the region of parahippocampus gyrus/collateral sulcus/hippocampus ($p<0.001$ uncorrected). In a direct contrast of face attractiveness vs. place attractiveness, no voxels survived wholebrain corrections.

To explore whether the areas that responded to face and place attractiveness overlapped with face- and place-selective visual regions that have been previously identified in occipitotemporal cortex, we conducted an ROI analysis using independently defined ROIs for fusiform face area (FFA), parahippocampal place area (PPA), a posterior object-selective region (LOC), and early visual cortex (EVC). Somewhat surprisingly, face attractiveness was positively correlated with activity in

all higher level regions (right FFA: $t(27)=2.9$, $p=0.007$, left FFA: $t(27)=2.2$, $p=0.037$; right PPA: $t(27)=3.8$, $p=0.0008$, left PPA: $t(27)=2.5$, $p=0.017$; right LOC: $t(27)=3.37$, $p=0.002$, left LOC: $t(27)=3.05$, $p=0.005$) while place attractiveness only showed positively correlated activity within right LOC ($t(27)=2.1$, $p=0.04$). Correlations between place attractiveness and fMRI response were nonsignificant for right FFA ($t(27)=0.8$, $p=0.41$), left FFA ($t(27)=0.7$, $p=0.46$), right PPA ($t(27)=0.35$, $p=0.73$), and left PPA ($t(27)=0.00$, $p=0.997$). Neither face nor place attractiveness was significantly correlated with activity in EVC. Figure 2.5 shows the location of the functional ROIs and activity related to face and place attractiveness in visual regions.

Finally, for completeness, we compared categorical differences in activity between face and place trials (irrespective of attractiveness). We observed significantly greater response during place compared to face trials in regions previously reported to respond preferentially to places and scenes (bilateral PPA, RSC, OPA/TOS). Conversely, we observed significantly greater response during face compared to place trials in visual regions previously reported to respond preferentially to faces (bilateral FFA, OFA), and also in posterior cingulate, bilateral amygdala, vmPFC, and, surprisingly, a region of right latOFC. (Table 2.2 reports the MNI coordinates of all significant clusters for this comparison.) Because we had also observed activation in right latOFC for face attractiveness, we sought to determine whether the right latOFC region defined by our face > place contrast also responded to face attractiveness. We defined this region as an ROI, thresholded at $t>3.5$ on the group map, and extracted the betas values from each subjects' response to face attractiveness. Unexpectedly, this region's response was not significantly correlated with face attractiveness ($t(27) t=0.75$ $p=0.46$). Figure 2.6 shows an overlay of both the face > place map and the face attractiveness map, demonstrating that the face

> place peak response is located on the posterior orbital gyrus, whereas peak activity for face attractiveness is more medially located within the sulcus.

Discussion

Attractive faces and attractive places promise very different rewards to a person, and comparisons are not often made between these rewards. Despite this, our results demonstrate that the vmPFC exhibits both category-general and category-specific responses to attractiveness. Clusters sensitive to face attractiveness in vmPFC were also sensitive to place attractiveness; however, a multivertex pattern analysis found that place and face attractiveness were distinguishable in one of these clusters located in the posterior and ventral portion in the right hemisphere. These results suggest that some parts of vmPFC might encode category-general reward signals even when the stimuli are not exchangeable goods, while other parts might retain information about category-specific rewards. In right lateral orbitofrontal cortex (latOFC) we observed two distinct face-specific regions: one sensitive to the categorical difference between faces and places but insensitive to face attractiveness, and one sensitive to face (but not place) attractiveness but insensitive to categorical differences. As we found only face attractiveness and not place attractiveness signals in latOFC, this region may be more involved in basic reward processing that is specific to some but not all stimulus categories.

Common response to face and place attractiveness in vmPFC

In vmPFC, we identified clusters that were sensitive to face attractiveness and clusters that were sensitive to place attractiveness, and these face and place clusters were highly overlapping. Further analysis revealed a common response to face and place attractiveness: clusters in vmPFC identified based on their response to face attractiveness responded equally strongly to place attractiveness, while clusters

identified based on their response to place attractiveness responded equally strongly to face attractiveness.

These findings are consistent with a recent meta-analysis that demonstrated a region of vmPFC that encodes a common evaluative signal in studies where two or more categories were directly compared (Levy and Glimcher 2012). Although most of the studies in the meta-analysis involved economic and consumer goods, a few compared various social rewards to monetary rewards. Smith et al. (2010) showed overlapping activity in vmPFC for face attractiveness and monetary value when subjects passively viewed intermixed images of faces and money. Lin et al. (2012) showed overlapping activity in vmPFC between monetary value and another type of social reward, pictures of smiling or angry people (paired with audio of emotionally matched words). Our results extend these findings by showing that even in the case where both judgments are entirely outside the economic domain, in the realm of aesthetics, an overlapping evaluative signal exists in vmPFC.

Indeed, previous findings in the neuroaesthetics literature have suggested that vmPFC is generally involved in aesthetic evaluations (Ishizu and Zeki 2013; Brown et al. 2011). For example, Ishizu and Zeki (2011) found overlapping response in vmPFC for beautiful paintings and music, demonstrating that a common evaluative mechanism in vmPFC is recruited by stimuli of different modalities (visual, auditory). However, it is possible that their subjects evaluated all of the items within the same conceptual/reward framework of "artwork" because the items were presented as such.

Our current results demonstrate an important extension to the results of Ishizu and Zeki, because we found a common signal across reward categories even in an independent ROI test of response to both categories. Furthermore, they did not

explore the possibility of separable distributed responses to visual and auditory artwork, whereas our pattern analysis revealed distinct face and place attractiveness patterns in at least one region that showed equal mean response between categories.

Category-specific attractiveness coding in vmPFC

Our pattern classification analysis revealed that a subregion in right vmPFC (RvmPFC-place) contains separable distributed response patterns to face and place attractiveness, even while showing similar *mean* response. In contrast, the other three vmPFC clusters, which also showed significant mean response to both place and face attractiveness did not display evidence of category-specific encoding, suggesting that category information in vmPFC may be restricted to a posterior and ventral subregion (see Figure 2.3) (though null results in the other clusters cannot be used to ultimately reject the category-specificity hypothesis in those clusters).

Our classification results provides an intriguing parallel with results from a recent study by McNamee, Rangel, and O'Doherty (2013) who also used pattern classification to demonstrate category-specific encoding for food and trinket value in ventral vmPFC. They also demonstrated that dorsal regions of vmPFC showed both mean and distributed response patterns indicative of *category-general* encoding, a result that is not inconsistent with our observation of overlapping mean activity for face and place attractiveness in vmPFC, even in clusters where we did not observe category-specific patterns of response. The main difference between our results and those of McNamee et al. is that the region where we observed category decoding was, in fact, defined by significant univariate response to both categories; in contrast, McNamee et al. showed a ventral to dorsal gradient of multivariate to univariate signal strength, suggesting that the two types of signals were to some

degree spatially separable. Nevertheless, our results provide independent support for the claim that vmPFC contains category-specific as well as category-general reward information. Moreover, we show that these components exist for non-economic rewards.

Face-specific activity in right latOFC

We observed two separate regions in right latOFC that contained face-specific responses, one that exhibited a categorical preference for faces over places and one whose response scaled with face but not place attractiveness.

The previous literature on face attractiveness has not shown consistent results in latOFC. Two studies have found activity that positively correlated with face attractiveness in latOFC (left latOFC: Winston et. al 2007; right latOFC: Tsukiura & Cabeza 2011). In contrast, in a passive viewing task, Liang, Zebrowitz, & Zhang (2010) showed activity negatively correlated with face attractiveness in bilateral latOFC, though this study included disfigured faces, possibly causing an emotional/saliency signal to override a positive attractiveness response. O'Doherty et al. (2003) reported activity negatively correlated with attractiveness in right latOFC, but in their study, subjects were making gender judgments rather than attractiveness judgments. When combined with these previous results, the current data suggest that positive latOFC activity for face attractiveness may only arise when subjects are explicitly evaluating face attractiveness.

There is very little evidence in the literature for a categorical response to faces in latOFC, though one study by Rajimehr, Young, & Tootell showed a very similar region of right latOFC that showed greater activation for faces than places (2009). Interestingly, they also used surface-based group registration, as we did in

the current study, which reduces the noise from high inter-subject anatomical variability.

No human neuroimaging study to our knowledge has shown a disjunction between regions exhibiting face-specific categorical response and regions exhibiting face-specific attractiveness response in latOFC. Our finding of this functional dissociation provides an important link to findings from the macaque, where multiple kinds of face-specific responses have been observed in orbitofrontal cortex (O'Scalaidhe et al. 1997; Rolls et al. 2006; Tsao et al. 2008; Watson & Platt 2012). Using fMRI, Tsao et al. showed that, in macaque OFC, a patch on the orbital surface (lateral orbital sulcus) responded more strongly to faces with emotional expressions than to neutral faces, whereas a more lateral face-specific patch (inferior convexity) showed a categorical face response that did not vary with facial emotions. These anatomical locations are congruent with our own results: the categorical face patch was located on the posterior/lateral gyrus, and the face attractiveness patch was found in the lateral orbital sulcus. While these similarities are suggestive of possible functional homologues, more research is needed to test the robustness and clarify the roles of these regions in humans, especially since the emotion-patch in macaques responded to both positive and negative faces, whereas we were only able to test for linear responses to positive faces.

It is notable that we did not observe place-specific activity in latOFC. Places, unlike faces, may not act as a "basic" reward category in the same way as faces and food. The calculation of place beauty might instead be highly reliant on a dynamic process of integrating "component parts," such as spatial envelope or contrast/lighting, which may be associated with rewards only over time. Barron et al. have shown evidence that online construction of novel reward categories happens

in vmPFC and hippocampus (2013). In line with this theory, place attractiveness was found in our data to correlate with activity in vmPFC and at sub-threshold levels in the hippocampal region.

Visual region differences

Consistent with previous findings (Chatterjee et al. 2009), we observed activity correlated with face attractiveness across a large area of visual cortex, including face-responsive (FFA) and object-responsive (LOC) regions. Importantly, this activity was not simply due to time-on-task, as RTs instead showed non-linear patterns of response to attractiveness. Chatterjee et al. theorized that response to face attractiveness in the FFA reflects processing of face beauty per se, while response in LOC reflects processing of visual aesthetics regardless of the category. They hypothesized that place beauty might activate place-specific mechanisms in the PPA in addition to general visual aesthetic mechanisms in LOC. While we did observe activity in right LOC correlating with place attractiveness, consistent with this prediction, we did not observe attractiveness-related activity in PPA for places. Rather, we observed an unexpected response to face attractiveness in this region.

Why might there be a weaker signal in visual cortex for place attractiveness? As discussed above, it may be the case that faces signify more immediate/basic rewards, and attractive faces may therefore recruit visual cortex as a part of an automatic approach response, whereas places would not. It could also be the case that we saw less activity for place attractiveness because the places spanned a greater variety of environment types (e.g. forests, beaches, deserts, fields) than did faces (males, females). For example, a recent study using a narrower range of place stimuli (indoor built environments) showed that place beauty was correlated with activity in the middle occipital gyrus, although there was no correlated activity within

the ventral visual cortices (Vartanian et al. 2013). One study has reported greater activity in PPA for preferred scenes versus non-preferred scenes (Yue, Vessel, & Biederman 2007), but it is important to note that this study markedly differs from our own in both task and stimuli, in that subjects were asked to make ratings based on the content of the scenes, many of which were not places but images containing highly salient foreground objects, people, and animals. While we did observe response to place attractiveness in the parahippocampal/hippocampal region anterior to the PPA, this did not survive the stringent threshold for significance in the whole-brain analysis. Taken as a whole with these previous findings, our data suggest that there are regions in visual cortex that respond reliably to face attractiveness, but response to place attractiveness may depend on the nuances of the judgment task or the stimulus set.

Conclusions

Our data demonstrate a bridge between aesthetic and economic neural signals, in that a functionally similar evaluative mechanism in vmPFC is recruited for these disparate types of judgments. We expect that further research will continue to illuminate our understanding of this mechanism as well as the unique neural representations underlying specific reward categories.

Figure 2.1

Places and faces were presented to subjects in blocks of 12 images. Two 36s fixation blocks were also included in each scan run. Each face appeared for 1 second with a 2 second ISI. Subjects made coarse attractiveness ratings for each image in the scanner ("low", "medium", "high" attractiveness), and then rated the images again outside of the scanner using a Likert scale, 1-8. The post-rated images were presented to subjects in one randomized block of faces and one randomized block of places (order counterbalanced across subjects).

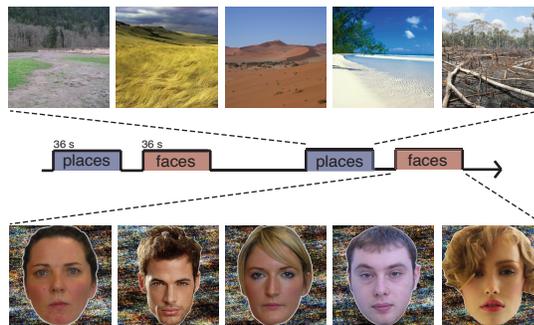


Figure 2.2

(A) Average of subject-specific histograms of post-scan attractiveness ratings. Error bars measure the standard error across subjects. (B) Response time plotted as a function of average attractiveness. Both face and place attractiveness exhibit an inverted-U shaped function, with the longest response times for mid-range images.

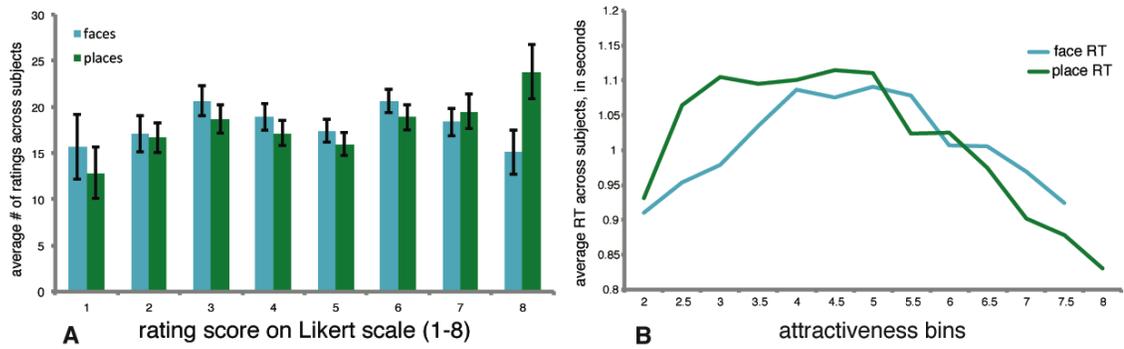


Table 2.1

Summary table of significant clusters of response for face or place attractiveness within the vmPFC and latOFC regions of interest (contrast shown in italics). Coordinates of peak values within the clusters are reported in MNI305 space. Cluster-wise p-values are permutation corrected for multiple comparisons within the vmPFC or latOFC and additionally Bonferonni corrected to account for observations across 2 hemispheres.

		mm²	x	y	z	cluster-wise p
<i>face attractiveness</i>						
vmPFC	L	267.97	-11.1	49	-22	0.016
vmPFC	R	407.91	8.4	36.1	-26.5	0.000
latOFC	R	131.09	33.9	41.1	-21.2	0.044
<i>place attractiveness</i>						
vmPFC	L	588.35	-6.1	42.3	-13.1	0.002
vmPFC	R	145.59	8.9	36.6	-23.5	0.032

Figure 2.3

Cross-subject validation results for univariate cluster-corrected group analysis. Clusters responding to face attractiveness (blue) and place attractiveness (green) are shown on the cortical surface. Bar graphs show mean parameter estimates for face and place attractiveness within these clusters. These parameter values were extracted using a leave-one-subject-out cross-validation procedure so that data used to define the clusters were independent of data used to estimate the size of the effects and response patterns. The black outlines on the medial surface indicate boundaries of vmPFC while black outlines on the orbital (i.e. ventral) surface indicate boundaries of latOFC. All vmPFC clusters show significant response to both face and place attractiveness. The face attractiveness cluster in right latOFC only responds to face attractiveness.

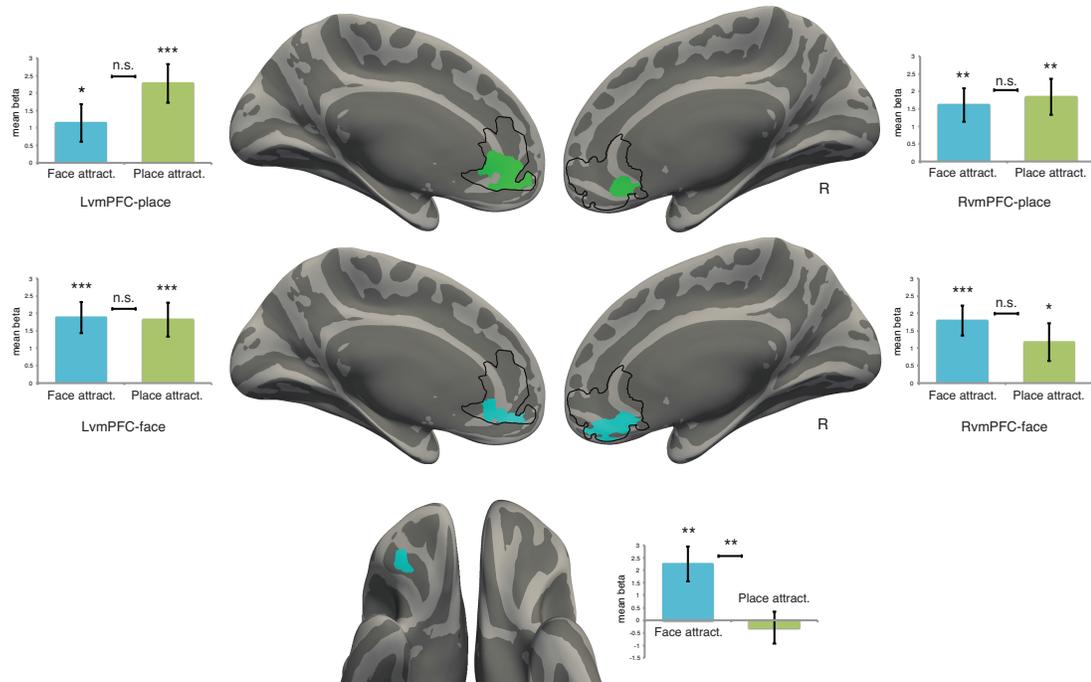


Figure 2.4

Multivertex pattern analysis of responses to face and place attractiveness in vmPFC. Bar graphs show classification accuracy when comparing response patterns across independent halves of the data. Accuracy (orange bars) was determined by comparing pattern similarity for attractiveness within a category to pattern similarity for attractiveness across categories. The breakdown by category (face-face vs. face-place and place-place vs. face-place) is also shown (blue and green bars). Only one cluster (RvmPFC-place) shows greater classification accuracy for same vs. different attractiveness categories, though RvmPFC-face also shows a trend towards significance.

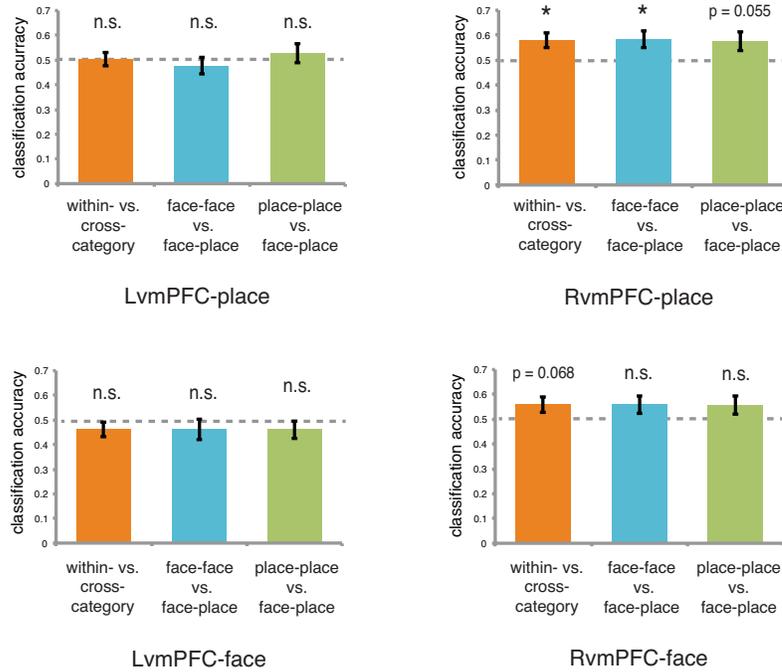


Table 2.2

Coordinates, voxel count, and peak t-values for significant clusters of voxels. Wholebrain maps for each contrast were thresholded at $p < 0.05$, permutation corrected for multiple comparisons. Coordinates are reported in MNI152 space.

		voxel count	X	Y	Z	Max t
<i>face > place</i>						
fusiform gyrus	R	1135	42	-74	-12	11.03
fusiform face area (FFA)	R	*	42	-52	-18	9.8
occipital face area (OFA)	R	*	42	-74	-12	11.03
cingulate sulcus	medial	960	2	52	18	9.43
amygdala	L	921	-20	-8	-16	11.03
	R	710	20	-8	-14	11.7
anterior cingulate gyrus	medial	603	2	28	-12	7.63
posterior cingulate gyrus	medial	404	4	-52	20	8.62
inferior occipital gyrus	L	341	-42	-82	-16	9.62
inferior frontal sulcus	R	251	48	8	20	7.06
middle temporal gyrus	R	127	58	-4	-26	7.46
cerebellum	L	67	-32	-82	-38	6.36
superior frontal gyrus	R	66	20	38	50	6.7
supramarginal gyrus	L	51	-58	-64	26	5.98
	R	37	56	-60	24	7.19
fusiform gyrus (FFA)	L	46	-40	-50	-24	6.93
posterior orbital gyrus		27	38	32	-16	6.04
<i>place > face</i>						
temporal-occipital cortex	L	4415	-26	-50	-10	21.4
	R	4372	28	-48	-12	15.3
parahippocampal gyrus/collateral sulcus (PPA)	L	*	-26	-50	-10	21.37
	R	*	28	-48	-12	15.26
retrosplenial complex (RSC)	L	*	-14	-56	10	10.01
	R	*	14	-52	8	10.35
middle occipital gyrus (occipital place area)	L	*	-34	-84	18	11.89
	R	*	34	-86	18	12.06
lingual gyrus	R	142	8	-98	-6	7.52
cingulate sulcus	R	81	4	-44	50	8.19
	L	66	-10	-36	44	6.55
cerebellum	L	25	-14	-46	-52	6.67
	R	17	16	-46	-52	6.3
<i>face attractiveness</i>						
intraparietal sulcus (posterior)	R	15	30	-80	24	6.38
isiform gyrus	L	9	-42	-54	-16	6.09
	R	4	22	-54	-14	5.8
intraparietal sulcus (anterior)	R	4	30	-72	24	5.83

*peaks are part of a larger cluster

Figure 2.5

Wholebrain maps for face (blue) and place (green) attractiveness displayed at $p < .0005$, uncorrected. The FFA is outlined in light blue, and the PPA is outlined in light green. Note that at this more liberal threshold, vmPFC and ventral striatal activity is visible for place attractiveness, though these regions did not survive volumetric wholebrain corrections.

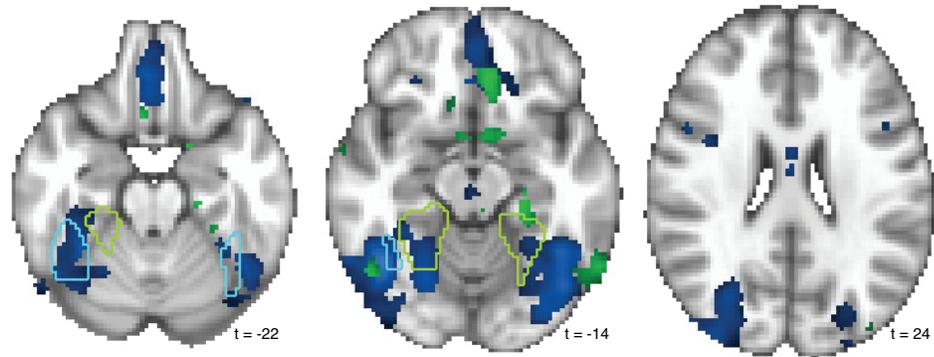
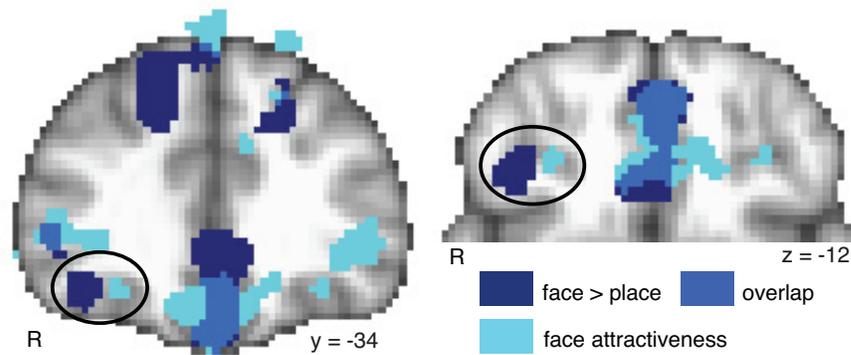


Figure 2.6

Relationship between face category effect and face attractiveness effect in PFC. The face > place effect (dark blue) was thresholded at $t > 4.0$, whereas the face attractiveness effect (light blue) was thresholded at $t > 2.0$. Both contrasts were binarized, and the overlapping activity is displayed in pale blue. Peak response for the categorical effect is located lateral to peak response for face attractiveness in right OFC (circled).



Chapter 3 – Face-responsive activity in lateral orbitofrontal cortex

Pegors, TK; Kable, JW; Epstein, RA. *in preparation*

Abstract

Perceiving and evaluating faces is an important part of making sense of our social world. While much has been discovered about the face-processing network in the brain, less is known about face-selective information in lateral orbitofrontal cortex (latOFC). A recent study out of our lab revealed two anatomically and functionally distinct face-selective “patches” in right latOFC that were active while subjects rated face attractiveness: a lateral region that showed an overall greater response for faces than places and a more central region that showed correlated response to face attractiveness. In the current study, we tested whether these patches showed a similar response during passive evaluations of face attractiveness where subjects were not making explicit ratings. By using ROIs from our previous study, we found the same pattern of response in right latOFC during passive evaluation. Furthermore, we showed that the face attractiveness response is restricted to the central sulcal region of latOFC, whereas the main effect for faces is more distributed across latOFC. Conversely, activity in ventromedial prefrontal cortex (vmPFC) was not significantly correlated with face attractiveness during the passive task, even though it had shown a robust response to face attractiveness during explicit ratings of face attractiveness. In sum, our results suggest that the lateral orbitofrontal cortex contains functionally dissociable regions that may relate to separate aspects of reward-related processing, whereas the vmPFC may be primarily recruited during decision-making.

Introduction

Faces are some of the most information-rich elements of our visual environment, providing cues to a person's identity, gender and age, as well as playing a crucial role in social communication. It is not surprising, therefore, that a tremendous amount of research has gone into identifying regions across the brain that are implicated in processes related to face perception (for reviews, see Haxby and Boggini 2011; Collins and Olson 2014). Even still, there remain regions whose face-related functions have remained largely unexplored, despite growing evidence that they too may be crucially involved in face-selective processing.

In particular, the right lateral orbitofrontal cortex (latOFC) is a reward-related region of the brain where at least two recent studies have observed face-selective activity when measuring neural responses to faces and places (Rajimehr, Yong, and Tootell 2009; Von Der Heide, Skipper, and Olson 2013). Face-selective response in this region is rarely reported in the literature, but it is important to note that these two studies took explicit steps to enhance detection in OFC, a region known to be prone to signal loss and dropout (Kringelbach and Rolls 2004; Deichmann et al. 2003). Additionally, a fair amount of evidence for face-selective encoding in latOFC exists in the monkey literature, where face-selective neurons have been reported in macaque orbitofrontal cortex using both single-neuron and fMRI methods (O Scailidhe, Wilson, and Goldman-Rakic 1997; Rolls et al. 2006; Tsao et al. 2008; Watson & Platt 2012). Interestingly, Tsao et al. showed that this face-selectivity activity appears to be divided into functionally distinct patches: a central region of latOFC showed a greater response to emotionally expressive faces when compared to neutral faces and a more lateral patch of latOFC showed a categorical response to faces over baseline (2008).

Our own lab recently conducted an fMRI study comparing the neural correlates of face and place attractiveness while using a number of techniques to enhance detection in the OFC (Pegors, Kable, Chatterjee, and Epstein, under review). Not only did we replicate the recent human findings for a categorical response to faces in right latOFC (Rajimehr et al. 2009; Von Der Heide et al. 2013), but we also showed evidence for a second, functionally distinct region, similar to that reported by Tsao et al. (2008): we observed a centrally-located region of right latOFC which responded to face attractiveness but which was topographically distinct from the region which showed a categorical response to faces.

In the current study we sought to measure the robustness of these face “patches” under different task conditions. To do this, we first tested whether the face attractiveness patch in central latOFC was also active when subjects only passively evaluated face attractiveness. Secondly, we tested whether the face-category patch in lateral latOFC was active during the passive evaluation task and during a standard 1-back matching task. Finally, we performed an anatomical analysis to determine whether these functional patches were aligned within specific orbital/sulcal boundaries.

To anticipate, our results suggest that activity for face attractiveness and for the main effect of faces persists in latOFC across task types. We furthermore showed that face attractiveness activity appears limited to the orbital sulcus region, while a main effect for faces is more distributed across posterior and medial latOFC.

Methods

Subjects

34 healthy, right-handed volunteers with normal or corrected-to-normal vision were recruited to participate in the fMRI study. All subjects gave written informed consent according to procedures approved by the University of Pennsylvania institutional review board. One subject was excluded for excessive motion, and one subject was excluded for falling asleep, leaving us with 32 subjects whose functional data was analyzed (16 females, average age=21.6).

fMRI Acquisition

Scans were performed at the Hospital of the University of Pennsylvania on a 3T Siemens Trio scanner equipped with a Siemens body coil and a 32-channel head coil. T1-weighted anatomical images were acquired using an MPRAGE iPAT2 pulse sequence (mode=GRAPPA, TR=1630ms, TE=3.11ms, TI=1100ms, voxel size=0.9x0.9x1.0mm, matrix size=350x263x350). T2*-weighted images sensitive to blood oxygenation level-dependent (BOLD) contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR=3000ms, TE=25ms, voxel size=2.5x2.5x2.5mm, matrix size=192x192, 49 axial slices). The slices were tilted -30° from the anterior and posterior commissure plane to reduce signal dropout (Deichmann et al. 2003). The Siemens standard field-mapping scan for B0 correction was also acquired (voxel size=2.5x2.5x2.5mm).

Stimuli

Stimuli were digitized 400x400 pixel color photographs of 288 places and 288 faces chosen to span a wide range of attractiveness. The places were natural environments with no man-made elements, spanning scene types such as swamps,

mountain ranges, beaches, fields, and forests. The face set had equal numbers of males and females, and all faces were Caucasian, upright and forward-facing, with neutral to pleasant expressions, selected from the Glasgow Unfamiliar Face Database (<http://homepages.abdn.ac.uk/m.burton/pages/gfmt/Glasgow%20Face%20Recognition%20Group.html>), Radboud Database (Langner et al. 2010), the Center for Vital Longevity Face Database (Minear & Park 2004), CVL Face Database (Peter Peer, <http://www.lrv.fri.uni-lj.si/facedb.html>), Diana Theater Face Database (courtesy of Dr. Robert Schultz at the Center for Autism Research), and online searches. Face images were extracted from their original background, blurred slightly along the edges, cropped so that hair did not extend below the chin, and resized to span a height of 400 pixels. They were then placed on an abstract colored background created by phase-scrambling a single place image. This ensured that they subtended the same visual angle as the places while retaining a similar background color. A unique scrambling of the background image was used for each face (code used from: <http://visionscience.com/pipermail/visionlist/2007/002181.html>).

Each fMRI subject was assigned a unique (though overlapping) set of 576 face and place images chosen to span the full range of attractiveness scores, which were based on attractiveness ratings given to a larger set of 916 images by a set of 28 independent raters. These raters were instructed to span their ratings across the full 1-8 Likert scale, and they made ratings at their own pace. Faces and place images were presented in separate blocks, and block order was counterbalanced across subjects. Ratings were z-scored across all images within each subject, and a mean attractiveness rating for each image was determined by averaged these z-scores across subjects. The images were then divided into "low", "average", and "high" attractiveness bins according to whether they had a z-score below -0.5 (low),

between -0.5 and 0.5 (average), or above 0.5 (high). Images were then chosen in equal numbers from these three attractiveness bins to create the final stimulus set.

Design and Procedure

The fMRI experiment consisted of six 7 minute 6 sec scan runs, each of which had one sequence of face trials and one sequence of places trials. Each sequence was constructed of 48 images interspersed with 16 null trials in a continuous carryover design (Aguirre 2007), which ensured that “low”, “average”, and “high” attractiveness images were presented in equal numbers and were preceded equally by each other type of trial. 2 “warmup” trials were also included at the beginning of each sequence as a way to allow the first trial of the counterbalanced sequence to have a preceding trial. Each image was presented for 1 sec with a 2 sec interstimulus interval (ISI) during which only the crosshair was on the screen.

The order of the face and place sequences was counterbalanced across runs. Between the face and place sequences were 9 seconds of passive fixation, and 9 seconds of fixation were added at the beginning and 12 seconds at the end of each run to allow the T2* signal to reach a steady state and to model the final HRF, respectively. Subjects were instructed to actively think about the attractiveness of each image, and they were given no other explicit task. To acclimate subjects to the distribution of attractiveness in the stimulus set, they were asked to view prior to the scan 16 faces and 16 places not used in the main experiment.

fMRI Data Analysis

Pre-processing and data analysis for individual subjects was performed using the FMRIB FEAT software libraries (Jenkinson et al. 2012; Woolrich et al. 2009; Smith et al. 2004). For each run, the first three volumes were removed to account for the fMRI signal not yet reaching steady-state, and functional images were then

motion- and slice-time corrected. The data was also high-pass filtered to remove temporal frequencies below 0.02 hertz (50 sec).

General linear modeling was used to estimate neural activity correlated with subject-specific ratings of face and place attractiveness. Mean ratings from an independent set of subjects (described in the *Stimulus* section) were used to create parametric regressors for face and place attractiveness that extended across all six experimental scans. The model also included quadratic (“U-shaped”) regressors for face and place attractiveness (e.g. Winston et al. 2007). Finally, categorical regressors were added for face trials, place trials, “warmup” trials, and nuisance regressors were included to account for between-scan variability and outliers (outliers calculated with the Gabrieli Lab’s Artifact Detection Tools: http://www.nitrc.org/projects/artifact_detect/). All regressors, except those for nuisance spikes and scan indicators, were convolved with a canonical HRF.

Unsmoothed parameter estimates from the first-level models were registered to the cortical surface using surface templates derived from each subject’s T1-weighted anatomical image by Freesurfer’s segmentation function (recon-all). These surface maps were then smoothed at 3mm FWHM and registered to the “fsaverage” brain (Freesurfer’s average brain, derived from 40 subjects).

First, we tested for neural response to face attractiveness in frontal cortex while subjects passively evaluated the visual attractiveness of face and place images. Mean parameter estimates from the face attractiveness contrast were extracted from ROIs which were defined from the same contrast (face attractiveness) in a previous study where subjects explicitly rated the visual attractiveness of faces and places (Pegors, Kable, Chatterjee, & Epstein, under review). Specifically, ROIs were created from this independent dataset by thresholding the face attractiveness contrast at

$p < .05$, uncorrected, and selecting contiguous regions on the orbital and medial surface of frontal cortex (see Figure 3.1).

In a second analysis, we measured response to the main effect of faces > places, again by extracting mean parameter estimates from independently defined ROIs using the same contrast (faces > places) from a previous study where subjects explicitly rated the visual attractiveness of faces and places (Pegors, Kable, Chatterjee, & Epstein, under review) (see Figure 3.2). The ROIs were defined by thresholding the faces > places contrast at $p < .001$, uncorrected, in surface-space and selecting contiguous regions on the orbital and medial surface of frontal cortex (See Figure 2). Additionally, we also measured faces > place activity from an independent localizer task within the same functionally defined ROIs. 20 subjects from our previously cited study (Pegors, Kable, Chatterjee, & Epstein, under review) performed two runs of a standard localizer task (these runs were independent from the main experiment from which we acquired the functional ROIs). During the localizer, each subject performed a 1-back matching task in which they pressed a key every time they noticed an image that repeated twice in a row. Each of the two runs were 6 min 12 secs, and both runs contained 20 blocks of faces, places, objects, and scrambled objects. Each image was presented for 490 ms with a 490 ms interstimulus interval. Regressors signifying each of the four block types were concatenated across runs and convolved with a standard hemodynamic response function and input into a general linear model. Unsmoothed parameter estimates were registered to subject-specific surface space, smoothed by 3mm FWHM, and then registered to the "fsaverage" brain. Mean parameter estimates were extracted from the faces > places contrast for each ROI.

To determine the anatomical specificity of the face-responsive effects in orbitofrontal cortex during the passive task, we conducted an ROI analysis by extracting mean parameter estimates for the face attractiveness and faces > places contrast from two predefined anatomical ROIs in the Destriux atlas (Freesurfer's aparc 2009 library). One ROI was defined as the combination of all gyri on the ventral/orbital surface ("G_orbital"), and the other was defined as the central sulcal region on the orbital surface ("S_orbital_H-shaped") (see Figure 3.3). Here, our estimates were spatially unwarped during preprocessing by a standard B0 map (We did not include B0 correction in our other ROI analyses because spatial unwarping was not available for the data from which the functional ROIs were defined).

Results

Face attractiveness response in frontal cortex

In our first analysis, we asked whether passive evaluation of face attractiveness (without explicit numerical ratings) evoked activity correlated with face attractiveness in lateral orbitofrontal cortex (latOFC) and/or ventromedial prefrontal cortex (vmPFC). In bilateral latOFC, there was a robust response in all three ROIs to face attractiveness (RlatOFC_att $t(30)=3.25$, $p=0.003$; LlatOFC_att-1 $t(30)=3.42$, $p=0.002$; LlatOFC_att-2 $t(30)=3.67$, $p=0.001$) (see Figure 3.1). On the other hand, bilateral vmPFC did not show a significant response to face attractiveness (RvmPFC_att $t(30)=1.08$, $p=0.29$; LvmPFC_att $t(30)=1.09$, $p=0.28$). (All statistics were obtained by comparing mean parameter estimates against zero.)

Mean response to faces > places in frontal cortex

In our second analysis, we asked whether a mean response to faces > places was present in latOFC and/or vmPFC when subjects were not making explicit

judgments of attractiveness. For this analysis we had two datasets: our main dataset in which subjects passively evaluated image attractiveness and a functional localizer in which subjects performed a 1-back matched task on blocked images of faces, places, objects, and scrambled objects. We again defined our ROIs from an experiment where subjects made explicit ratings of attractiveness, using the contrast of face trials > place trials.

The pattern of response was very similar across both tasks (Figure 3.2). In the right latOFC, we observed activity for face attractiveness in both tasks (RlatOFC_face response during passive evaluation $t(30)=6.29$, $p<0.0001$, during 1-back matching $t(19)=4.56$, $p=0.0002$). In bilateral vmPFC, there was a significant response to face > place in the ventral clusters (LvmPFC_face-1 response during passive evaluation $t(30)=6.07$, $p<0.0001$, during 1-back matching $t(19)=4.45$, $p=0.0003$; RvmPFC_face-1 response during passive evaluation $t(30)=6.29$, $p<0.0001$, during 1-back matching $t(19)=4.56$, $p=0.0002$). In the more dorsal clusters there was no significant response during either task (LvmPFC_face-2 response during passive evaluation $t(30)=0.23$, $p=0.82$, during 1-back matching $t(19)=0.68$, $p=0.5$; RvmPFC_face-2 response during passive evaluation $t(30)=0.77$, $p=0.45$, during 1-back matching $t(19)=0.83$, $p=0.42$).

Anatomical specificity of face-related responses in latOFC

Both of our analyses revealed that both face patches in latOFC, the lateral patch for faces > places and the central patch for face attractiveness, continued to show these responses across tasks. The next question we probed was the anatomical specificity of both of these effects. In our previous study where subjects made explicit ratings of attractiveness, we showed that activity correlated with face attractiveness was focused in the sulcal region of latOFC, whereas the main effect of

face > place was focused in the posterior gyrus. It should be pointed out that in our previous study, we did not apply B0 unwarping to the data, meaning that the spatial specificity of the effects may have been affected by signal distortion. For the anatomical analysis in our current study, we were able to apply B0 correction to our functional data, with the goal of obtaining more precise spatial estimates. Furthermore, we used anatomically-defined gyral and sulcal regions of latOFC to extract parameter estimates (see Figure 3.3).

Our results revealed significant bilateral response in both the orbital gyri and sulci for the main effect of faces > places (left orbital gyri: $t(30)=6.0$, $p<0.0001$; right orbital gyri: $t(30)=4.4$, $p=0.0001$; left orbital sulci: $t(30)=3.66$, $p<0.0001$; right orbital sulci: $t(30)=3.17$, $p=0.003$; See Figure 3.3). The face attractiveness response showed a significant bilateral response within the orbital sulcal ROI but not within the gyral ROI (left orbital gyri: $t(30)=1.63$, $p=0.11$; right orbital gyri: $t(30)=1.05$, $p=0.3$; left orbital sulci: $t(30)=4.26$, $p=0.0002$; right orbital sulci: $t(30)=3.02$, $p=0.005$).

We also tested for a “saliency” effect of attractiveness (a “U-shaped” response). No significant activity was observed in any of the ROIs for this effect (left orbital gyri: $t(30)=1.26$, $p=0.21$; right orbital gyri: $t(30)=1.51$, $p=0.14$; left orbital sulci: $t(30)=-0.89$, $p=0.38$; right orbital sulci: $t(30)=1.0$, $p=0.32$).

Discussion

The face-processing network consists of a widely distributed collection of regions across the brain, but while the lateral orbitofrontal cortex (latOFC) has not typically been included in such descriptions, recent evidence from our own lab and others has suggested that face-related processing occurs in this region as well. In the current study, we explored the robustness of this face-related activity in latOFC.

Across multiple task paradigms, we observed a consistent main effect for faces > places in a lateral region of latOFC. We also demonstrated that the central sulcal region of latOFC responded to face attractiveness even during a passive evaluation task. These combined results suggest that the latOFC may play an active role in reward-related aspects of face-processing, and that these processes occur even in the absence of decision-making.

Face attractiveness in latOFC and vmPFC

We looked for frontal activity correlated with face attractiveness during a task where subjects only passively evaluated image attractiveness. In OFC, we found activity significantly correlated with attractiveness in the same regions where we had previously observed this activity during a task where subjects explicitly rated face attractiveness. In contrast, there was no attractiveness-related response in vmPFC during the passive task, even though vmPFC did show such a response during the explicit ratings task.

Evidence that value-related signals are encoded in latOFC even outside of decision-making events comes from both the human and monkey literatures. Single-cell recordings in macaque latOFC show attenuated response to gustatory and olfactory stimuli after satiation on the associated item, suggesting that those neurons encoded the relative value of the stimulus even during passive receipt of the reward (Rolls 1989; Critchley and Rolls 1996). In humans also, two studies have shown attenuated response in latOFC after subjects passively tasted or smelled a stimulus associated with a satiated food item (O'Doherty et al 2000; Kringelbach et al. 2003).

On the other hand, many studies looking for value-related response to face attractiveness do not observe consistent response in latOFC: some studies show

positively correlated activity in this region with attractiveness (Winston et al. 2007; Tsukiura and Cabeza 2011), some show negative correlated activity with attractiveness (Liang, Zebrowitz, and Zhang 2010; O’Doherty et al. 2003) and some show no correlated activity with attractiveness (Chatterjee et al. 2009). It may be the case that tastes and smells naturally evoke a stronger neural response with associated food rewards than the more subtle quality of face attractiveness evokes for associated rewards. In this light, methodological details may be crucial in detecting responses in latOFC to face attractiveness, a region which is already known to suffer from signal loss and dropout due to the nearby sinus cavity (Diechmann 2003) and which has a great deal of sulcal variability between subjects (Chiavaras and Petrides 2000; Chiavaras et al. 2001). To boost detection power in latOFC, signal dropout can be reduced by tilting the image acquisition (Diechmann et al. 2003), and signal distortion can be at least partially corrected by “unwarping” the data with acquired B0 field maps (Stenger 2006). Furthermore, to reduce misalignment due to inter-subject sulcal variability, group analyses can be performed in surface space, meaning, on “unfolded” cortical maps. This procedure allows individual maps to be aligned to a common space *after* being unfolded, removing more of the variance due to individual differences in cortical folds. The robustness of latOFC response to face attractiveness in our own two experiments may be due to the fact that we made use of a number of these methodological techniques.

Our null result in vmPFC is unlikely to be due to less power to detect the effect than in our previous study, as our passive evaluation task had double the number of trials used in our previous explicit evaluation task. Furthermore, we continued to observe significant response in OFC during the passive task, suggesting that our passive task did not simply reduce the overall activity across the brain. A

more likely explanation for this dissociation in vmPFC across tasks is that the vmPFC may calculate value only in those cases where comparisons are being made. Plassman, O'Doherty, and Rangel showed vmPFC activity correlated with willingness to pay values in a bidding task, but they showed that this activity in vmPFC dropped out when subjects were forced to bid a particular amount (the subject had no choice) (2007). Noonan et al. showed that after vmPFC lesions, macaques retained reward-value learning but had impaired performance when making decisions that involved reward comparisons (2010). While our own attractiveness ratings task did not have subjects make comparisons within individual trials, in other work we have shown that attractiveness ratings given to faces in sequence are made in comparison to ratings given in previous trials (Pegors, Mattar, Bryan, and Epstein, under review). Curiously, there are a number of fMRI studies in humans demonstrating vmPFC activity correlated with face attractiveness/preference when subjects are making orthogonal judgments to attractiveness (e.g. O'Doherty et al. 2003; Lebreton et al. 2009). Why vmPFC would show correlated activity with item attractiveness during orthogonal judgments but not during passive evaluation of that attractiveness is unclear.

Faces > places activity in latOFC and vmPFC

In a basic contrast of face attractiveness trials > place attractiveness trials, we observed a significant positive response in latOFC and ventral vmPFC, replicating the effect we showed previously in these regions when subjects made explicit attractiveness ratings of face and place attractiveness. Moreover, we also observed this face selectivity during a standard localizer task in which subjects performed a 1-back matching task.

Why might there be regions of OFC that shows a mean response to faces but are relatively insensitive to face attractiveness? It is likely that the posterior and medial regions are functionally distinct from each other, as there is evidence that medial and lateral OFC receive separate inputs and are part of separate networks of connections (Ongur and Price 2000). The more posterior response may reflect incoming high-level visual information before being combined with associated reward information. Projections from the anterior temporal lobe, known to contain face identity information (Tsao, Moeller, and Freiwald 2008; Von Der Heide, Skipper, and Olson 2008; Collins and Olson 2014), terminate in orbital frontal cortex by way of the uncinate fasciculus (Croxson et al. 2005). It also could be the case that this region reflects a separate type of reward, one that is not correlated with attractiveness but that might be associated with all face stimuli in general. The medial OFC / ventral vmPFC response that we observed in the region of gyrus rectus is actually very commonly observed in face contrasts, robust even in visual localizer scans (e.g. Julian et al. 2012). It has been suggested that this face-selective activity in medial OFC reflects the processing of emotion-related signals, such as social knowledge and self-awareness (Adolphs 2009; Haxby and Gobbini 2011).

Only a dorsal region of vmPFC, in the general region of anterior cingulate cortex (ACC), did not show an explicit response during passive evaluation or matching even though it had done so during explicit attractiveness ratings. It could be the case that this bilateral region is similar to the vmPFC ROI that did not respond to face attractiveness during the passive task, though a majority of the vertices between these functional ROIs do not overlap. Even if this dorsal vmPFC / ACC region is functionally distinct from the face attractiveness ROI, ACC is also known to be recruited during explicit decision-making: activity here has been associated with

post-decisions variables during subjective value choices (e.g. choice outcomes) (Cai and Padoa-Schioppa 2012). The reason why we did not see activity here for the passive or matching task may therefore result from the fact that these tasks required no value-related decisions to be made.

Face Specificity

To what extent are these effects in latOFC specific to faces? In our previous study, we showed that the central patch in latOFC correlated with face attractiveness but not place attractiveness, suggesting at least some degree category selectivity in this area (Pegors, Kable, Chatterjee, and Epstein, under review). At the same time, many gustatory and olfactory studies have shown increased response in latOFC to pleasurable tastes or smells (See Gottfried and Zald 2005 and Small et al. 1999 for meta-analyses of olfactory and gustatory activity in human OFC), suggesting that latOFC may be sensitive to a range of basic/primary rewards. Within this more general region, though, one study using single-cell recordings in monkeys showed evidence that many neurons within latOFC distinguished between social and juice values, and that these populations were intermixed (Watson and Platt 2012). These results would predict that human fMRI data would show similar mean patterns of response for separate basic rewards in latOFC. Techniques such as fMRI adaptation and multi-voxel pattern analysis will likely be more suitable for further exploration of categorical specificity within this central latOFC patch.

Conclusions

Our data suggest that multiple patches of face-specific information in lateral orbitofrontal cortex respond regardless of subject task. Activity in these patches may play an integral role within the larger face-processing network, and we expect that methodological advancements and careful experimentation will further our

understanding of the ways in which these regions uniquely contribute to the perception and evaluation of face stimuli.

Figure 3.1

Response correlated with face attractiveness during passive evaluation of attractiveness. Functional ROIs (displayed in purple) were defined in an independent dataset, from regions positively correlated with face attractiveness while subjects made explicit ratings of attractiveness. During passive evaluation, latOFC clusters show significant correlated response to face attractiveness, but there was no significant correlated response in vmPFC (see bar graph).

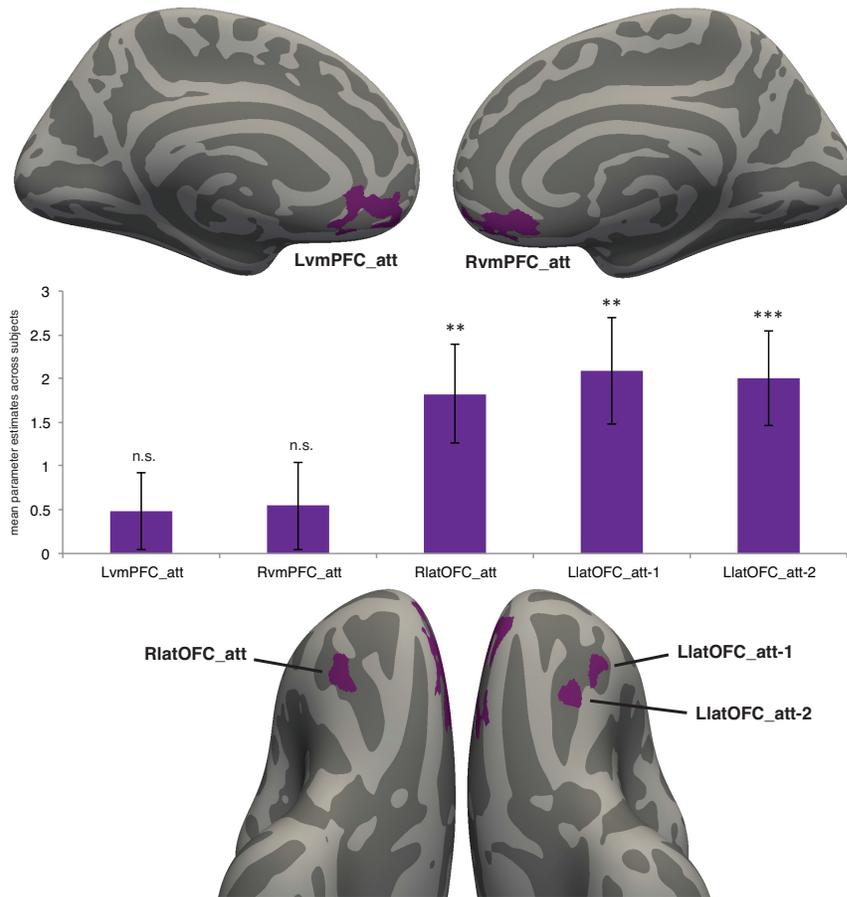


Figure 3.2

Activity for face > place in functional ROIs defined from the face > place contrast in a dataset where subjects performed an explicit attractiveness rating task (ROIs displayed in blue). Face > place activity in orbital and medial clusters was measured during passive evaluation of attractiveness and during a 1-back matching task. In both cases, right latOFC showed a significant response, and ventral medial-frontal cortex also showed a significant response during both tasks. However, there was no response in anterior cingulate, a region that showed face > place response during explicit ratings of face attractiveness.

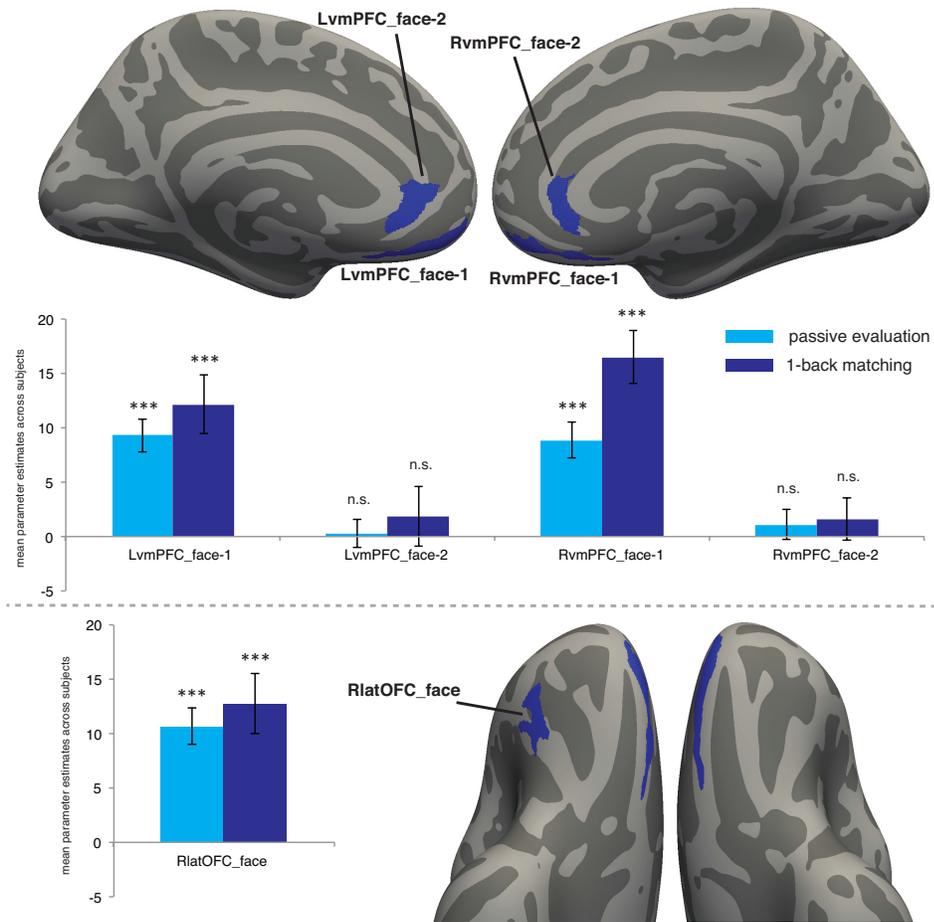
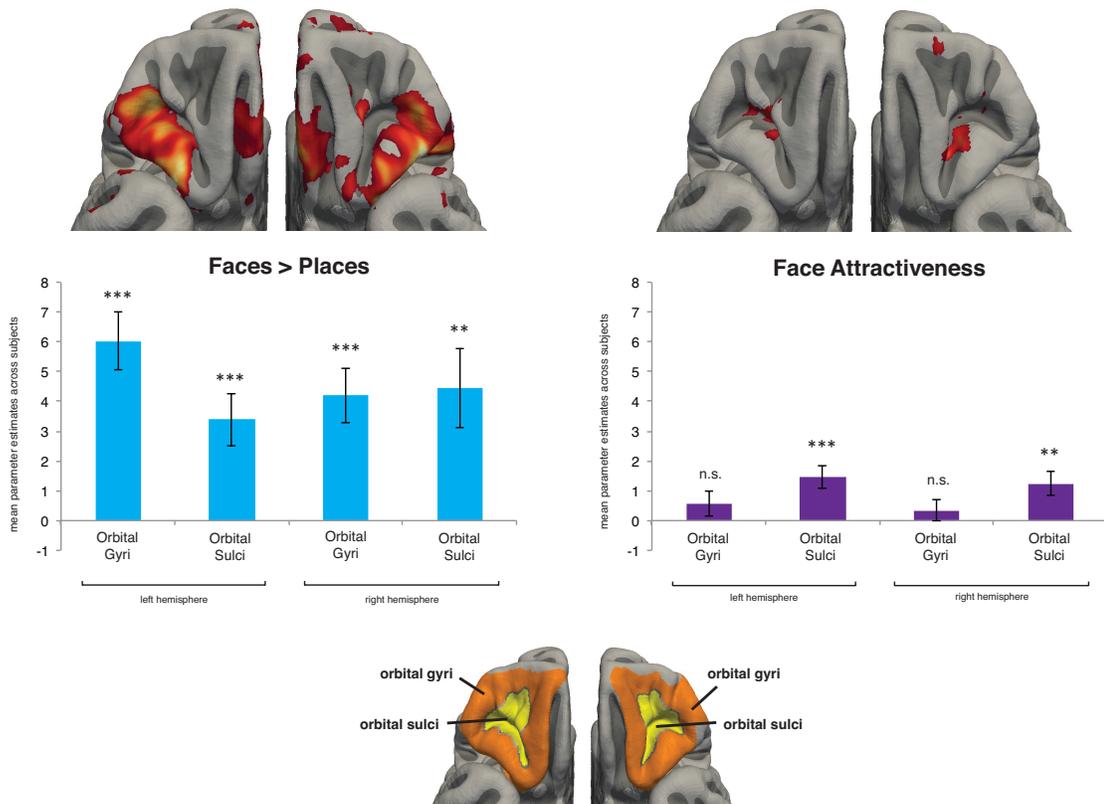


Figure 3.3

Face-responsive activity during passive evaluation of attractiveness. An ROI analysis was performed using anatomical ROIs (Freesurfer's aparc 2009 atlas) that divided the orbital surface into the gyri and central sulcal region. While the main effect of faces > places show more widespread activity through the posterior orbital surface, activity correlated with face attractiveness was specifically located in the orbital sulcal region. Contrast maps of face > place and face attractiveness are shown at $p < .01$, uncorrected. Un-inflated surfaces were used to visually accentuate the gyri and sulci.



CHAPTER FOUR - Simultaneous perceptual and response biases on sequential face attractiveness judgments

Pegors, TK; Mattar, MG; Bryan, PB; Epstein, RA. *in preparation*

Abstract

Face attractiveness is a social characteristic that we often use to make first-pass judgments about the people around us. However, these judgments are highly influenced by our surrounding social world, and researchers still understand little about the cognitive mechanisms underlying these influences. In a series of experiments, we investigated the existence of two opposing influences on face attractiveness ratings that arise from our past experience of faces. By implementing a unique sequential rating design, we showed simultaneous and opposing influences on attractiveness judgments which can be attributed to separate sources: First, we observed a response bias, in which attractiveness ratings shift *towards* a previously given rating, and a stimulus bias, in which attractiveness ratings shift *away* from the actual attractiveness of the previous face. Furthermore, we provide evidence that the contrastive stimulus bias is due to a perceptual “aftereffect,” by showing that this bias is limited to images with perceptual similarities and by localizing this effect to the fusiform face area and nearby visual cortical regions. Overall, our results suggest that even abstract judgments of face attractiveness are influenced by information from our evaluative and perceptual history and that these influences have measurable behavioral effects over the course of just a few seconds.

Introduction

Human faces are rich sources of information that help us to navigate our social world. Face attractiveness, in particular, is a holistic visual trait that we often use to make first-pass assessments of people, as we associate this feature with romantic viability, sociability, and health (for reviews, see Rhodes 2006 and Zebrowitz & Montepare 2008). Interestingly, our judgments of the attractiveness of an individual face are not based solely on that face alone: they are highly influenced by other faces observed in the surrounding context. For example, a person is considered more attractive if seen with an unattractive stranger (Kernis & Wheeler 1981), a very attractive partner or friend (the “radiation” effect: Strane & Watts 1977; Kernis & Wheeler 1981), or by merely appearing within a larger group of people (the “cheerleader” effect: Walker & Vul 2014). Moreover, even faces viewed in isolation are still often judged to be more or less attractive based on faces that have been viewed in the recent past (Wedell, Parducci, & Geiselman 1987; Cogan, Parker, & Zellner 2013; Kondo, Takahashi, & Watanabe 2012). Surprisingly, the nature of this “sequential” attractiveness bias remains unclear, because the results in the literature up to this point have been, at first glance, contradictory. Whereas some studies report a contrastive effect (i.e. if the previous face was very attractive, the current face will be rated as less attractive than usual) (Wedell et al. 1987; Cogan et al. 2013), other studies report an assimilative effect (i.e. if the previous face was very attractive, the current face will be rated as more attractive than usual) (Kondo et al. 2012; Kondo, Takahashi, & Watanabe 2013). The current study attempts to resolve this apparent contradiction.

Contrastive Sequential Biases

One of the first studies to show the influence of recent visual history on current ratings of attractiveness had experimenter “confederates” interrupt undergraduate males who were watching *Charlie’s Angels* to ask them to rate the attractiveness of a girl in a photograph (who was described as a potential date). Males who were watching *Charlie’s Angels*, a show with 3 beautiful women as the main characters, rated the girl in the photograph as less attractive than did other males who were watching another TV show (Kenrick & Gutierrez 1980). Follow-up studies in laboratory-controlled settings provided further evidence for this sequential contrast effect: faces tended to be rated as less attractive when a beautiful face had been previously viewed, and vice versa (Kenrick & Gutierrez 1980; Wedell, Parducci, & Geiselman 1987; Cogan, Parker, & Zellner 2013).

Interestingly, this sequential contrast bias occurs for other kinds of judgments as well, including both hedonic and non-hedonic judgments (Kamenetzky 1959; Shifferstein & Frijters 1992; Shifferstein & Kuiper 1997; Zellner et al. 2003; Parker et al. 2008). For example, a study originally conducted for military taste testing showed that foods were rated as tasting worse when sampled after a good quality food than when sampled after a poor quality food (Kamenetzky 1959). In another study, musical excerpts were given higher ratings if played after a low-rated excerpt than if played after a high-rated excerpt (Parker et al. 2008). Studies on magnitude estimates from the psychophysics literature have even demonstrated sequential contrast biases for estimates of loudness, light intensity, or size (e.g. Jesteadt, Luce, & Green 1977; Ward 1990; DeCarlo & Cross 1990). The fact that contrastive biasing occurs for such a variety of stimulus types raises the question of whether the same

mechanism that underlies other hedonic and psychophysical sequential biases might also be responsible for the sequential bias in face attractiveness judgments.

One general mechanism that may tie together the face attractiveness sequential bias with other sequential biases comes from the perception literature, where the phenomenon of visual “aftereffects” has been used for over a century to describe contrastive sequential perception in motion, color, and shape (Wohlgemuth 1911, McCollough 1965; Gibson 1933). These aftereffects are thought to occur as a result of our visual system constantly adapting to incoming stimulus information and influencing our perception of subsequent input. One striking example from the real world is that looking at an unmoving rock face after staring at a flowing waterfall makes it appear as if the rock face is flowing *upward* (Wohlgemuth 1911). Importantly, aftereffects have more recently been shown to occur even for complex facial features, such as identity, gender, ethnicity, and emotion (Leopold et al. 2001, Webster et al. 2004). In a demonstration of the gender aftereffect, Webster et al. first had subjects view a male or female face for 3 minutes and then judge the gender of a series of ambiguously gendered male/female morphs. Their results showed that if a subject adapts to a male face, they are more likely to judge an ambiguous face as female, and vice versa (Webster et al. 2004). No study has used this paradigm to directly test for aftereffects during face attractiveness judgments (though see Rhodes et al. 2003 for indirect evidence), but these results suggest that the attractiveness contrast effects seen in the social psychology literature may also be the result of a perceptual aftereffect.

Assimilative Sequential Biases

Instead of a contrastive bias, Kondo et al. have reported an *assimilative* bias for sequential face attractiveness judgments (2012). In this study, subjects made

sequential attractiveness judgments of faces using a 1-7 Likert scale. Their results showed a significant assimilative sequential bias: if the previous face was rated as very attractive, the current face would be rated as a little more attractive than usual, and vice versa. The authors attributed this bias to the previous *response*, not to the perception of the previous stimulus itself. Like the contrast effect, this response bias has been reported to occur for more than just attractiveness judgments of faces: this type of bias is broadly known in the decision-making literature as the “anchoring” effect. Tversky and Kahneman originally described this effect as one in which a person’s current decision will be biased towards a previously given value onto which they “anchor and adjust” (1974). Studies in psychophysics have also reported similar assimilative biases in magnitude judgments, though varying theories have been put forward as to the nature of such a bias (Ward & Lockhead 1971; Decarlo & Cross 1990). If it is the case that face attractiveness judgments are influenced by a type of assimilative response bias as the Kondo results suggest, then the contrastive and assimilative results reported in the literature are not necessarily in conflict but may arise from two separate sources, the *stimulus* and the *response*.

Simultaneous Stimulus and Response Biases

No study on face judgments has attempted to observe whether sequential biases simultaneously arise from both the previous stimulus and response. The difficulty in asking such a question is that these two effects are extremely hard to separate because the perception of attractiveness is highly correlated with the actual judgment of attractiveness. While modeling solutions have been proposed in the psychophysics literature to determine the presence of biases arising from the previous stimulus and response (e.g. Jesteadt, Luce, & Green 1977; Ward & Lockhead 1971; DeCarlo & Cross 1990; Matthews & Stewart 2011), these solutions

are necessarily limited in their ability to accurately detect and estimate effects in the presence of multicollinearity (Neter, Wasserman, & Kutner 1989). A more effective method is to decorrelate these possible biases in the experimental design itself.

Given that current designs of sequential attractiveness judgment tasks have highly correlated stimulus and response characteristics, we took two experimental approaches to “de-coupling” the potential biases attributable to the previous stimulus and response. In Experiments 1 and 2, we used a sequential rating design that alternated the task in such a way as to allow us to independently measure the effects of the previous stimulus and response on face attractiveness judgments. This design also allowed us to explore the generalizability of these effects across different judgment and stimulus types. In Experiment 3, we analyzed the neural activity of subjects making sequential face attractiveness judgments in the scanner. We hypothesized that the potential effects would be neuroanatomically “de-coupled,” and that the anatomical loci of these effects might provide further insight into their nature.

To anticipate, our results show that opposing biases due to both the previous stimulus and response are indeed simultaneously present during sequential face attractiveness judgments, and we additionally showed that these effects are not unique to attractiveness judgments. Furthermore, we observed activity correlated with a contrastive sequential attractiveness bias (but not assimilative sequential bias) in visual cortex, providing additional evidence for the perceptual nature of the contrastive bias and its resemblance to perceptual aftereffects.

Experiment 1

The goal of our first study was to test whether face attractiveness judgments made in sequence are biased by the attractiveness of the preceding face, the

response to the preceding face, or both. To answer this question, we asked subjects to make attractiveness judgments and hair darkness judgments of faces on alternating trials. Because all attractiveness trials were preceded by hair darkness trials (a quality which was not correlated with attractiveness in our stimulus set), this design allowed us to separately measure the effect of the attractiveness of the preceding stimulus and the response to the preceding stimulus on face attractiveness judgments. To determine the generality of the effects, we also investigated whether the preceding stimulus and/or response during attractiveness trials affected subsequent hair darkness judgments.

Methods

Stimuli

242 female face images were selected to span a wide range of attractiveness and hair darkness. These came from the Glasgow Unfamiliar Face Database, Radboud Database (Langner et al. 2010), the Center for Vital Longevity Face Database (Minear & Park 2004), CVL Face Database (Peter Peer, <http://www.lrv.fri.uni-lj.si/facedb.html>), Diana Theater Face Database (courtesy of Dr. Robert Schultz at the Center for Autism Research), and online searches. Faces were all Caucasian, had a neutral to pleasant expression, and were forward-facing. They were cropped such that the hair did not extend well below the chin, resized to a height of 400 pixels, and placed on 400x400 pixel backgrounds consisting of phase-scrambled variations of a single scene image (See Figure 4.1 for example stimuli). From this set of 242 images, 10 of these composed a practice set of trial images used across all subjects, and the experimental trial (212 images) and memory task foil images (10 images) were randomly drawn for each subject from the remaining 232 images.

We acquired attractiveness ratings from 28 subjects not participating in our main experiments to calculate an attractiveness score for each face. Each rater made 1-8 Likert scale ratings of 543 male and female faces (244 females). In a separate block, ratings were given to place images (see Experiment 2). Within each block, image order was randomized, and attractiveness ratings were averaged across raters for each item to determine its attractiveness score. In the current experiment, these female face attractiveness scores served as our stimulus values, which were considered to be independent of the stimulus history or the task.

Subjects

Our *a priori* sample size was set at 30, which was based on the number of subjects used in an earlier experiment that implemented a similar (but non-alternating) design (Kondo et al. 2012). 32 total Penn undergraduates were recruited and given class credit for their participation. 2 subjects were excluded for not following instructions, leaving us with a total of 30 participants (21 female).

Procedure

Subjects made a total of 106 hair darkness judgments and 106 attractiveness judgments in an alternating fashion on a 1-8 Likert scale. Importantly, these judgments alternated such that all attractiveness judgments were preceded by hair darkness judgments, and vice versa. Faces were presented on the screen for 4 seconds each, and between face presentations, a fixation cross appeared on the screen for a randomized interstimulus-interval length of 0-0.5 seconds. Faces were displayed in the center of the screen, and buttons indicating the numbers between 1 and 8 were displayed at the bottom of the screen (see Figure 4.1). Subjects were instructed to place 8 fingers on the keyboard row of numerical keys, so that ratings

could be made easily and quickly. To make the task easier, the current judgment type (attractiveness or hair darkness rating) was cued on the screen by the color of an outline around the face and buttons, as well as by the button labels at the anchors of the scale. When the subject made a judgment, the corresponding outline of the button turned white to reinforce their selection.

No face was repeated over the course of experiment. Faces were randomly ordered and randomly assigned to one of the two judgment types for each subject. The first judgment type that subjects made was counterbalanced between subjects. In an attempt to ensure that participants attended to the entire face (and did not just focus on the hair, for example) we asked participants to remember each face for a post-experiment memory test.

To acclimate participants to the range of attractiveness in the experiment, participants were trained beforehand on the alternating task with 10 faces that were not used in the main experiment. Faces for the practice were chosen to span the range of attractiveness and hair darkness. Participants were instructed to spread their ratings during the main experiment across the full scale based on the range of faces they had seen during the practice.

After the main experiment, subjects were shown a random subset of 20 images from the experiment (10 from the hair darkness trials, 10 from the attractiveness trials), and 20 novel images. These images were randomly intermixed, and subjects used a mouse to click a "Y" button on the screen if they had previously seen the image and the "N" button if they had not. Subjects completed the memory task at their own pace.

To acquire hair darkness ratings outside of the context of the alternating task, subjects rated, at their own pace, hair darkness on the full set of female faces (242

images). Subjects also made hair darkness ratings on a separate block of male faces (not used in our subsequent analyses). Faces were presented in a different randomized order for each subject. Whether subjects rated male faces or female faces first was counterbalanced across subjects. The resulting hair darkness ratings were compiled from 28 of the subjects (2 subjects' ratings were not acquired due to technical errors) and averaged across subjects to create a mean hair darkness score for each face.

Results and Discussion

For all analyses, any trial was excluded in which the reaction time (RT) was less than or equal to 0.2 seconds, as a short RT might indicate an anticipation error or a rating attributable to the previous trial. There was no correlation between attractiveness ratings and RT (Pearson's $r = 0.01$, $t(29)=0.42$ $p=0.68$), and an extremely small but trending negative correlation between hair darkness ratings and RT (Pearson's $r = -0.04$, $t(29)= -1.93$ $p=0.06$), indicating that that hair darkness was rated slightly more quickly for faces with darker hair.

Subjects performed above chance in the memory test for correctly identifying faces seen during attractiveness judgments (mean=7.1 out of 10, $t(29)=5.69$, $p<0.001$) and for correctly identifying novel faces (mean=14.6 out of 20, $t(29)=6.76$, $p<0.001$). On the other hand, subjects did not perform above chance for recognition of faces seen during hair darkness judgments (mean=5.23 out of 10, $t(29)=0.68$, $p=0.5$), suggesting that subjects paid less attention to faces during hair darkness judgments.

In our first analysis, we used a time-series regression analysis to determine whether attractiveness ratings were significantly influenced by either the attractiveness of the previous face or the previous response. We created a separate

model for each subject by regressing individual attractiveness ratings against the mean attractiveness of the previous face and that subject's hair darkness response given to the previous face. We also included the mean attractiveness of the current face as a predictor, to account for attractiveness variance not due to sequential biasing. The model used is summarized by the following equation:

$$R_t = \beta_0 + \beta_1 S_t + \beta_2 R_{t-1} + \beta_3 S_{t-1} + \varepsilon \quad (1)$$

where R is the response, S is the average attractiveness of the face (the stimulus value), t is the trial index and ε is the error term. (Note that R_t in this first model is a judgment of attractiveness and R_{t-1} is a judgment of hair darkness.) The dependent variable and all predictors were standardized (z-scored) for each subject in order for the resulting beta estimates to be comparable across subjects. The mean Pearson's r between the previous response R_{t-1} (hair darkness judgment) and the previous stimulus value S_{t-1} (attractiveness) was -0.135. We formally tested for multicollinearity by examining the variance inflation factor (VIF) of each of the independent variables. This number gives us an estimated severity of multicollinearity – the higher the number the more severe, with a lower bound of 1. Each of our independent variables had a very low VIF ($S_t = 1.02$, $R_{t-1} = 1.04$, $S_{t-1} = 1.04$), suggesting that multicollinearity was not a concern. (A VIF of 1.02 means that the variance of the coefficient is 0.02% larger than it would be if that predictor were uncorrelated with all other predictors).

Beta estimates of the previous stimulus and response predictors were extracted for each subject-specific regression model. Results from testing these betas against zero revealed that the response given during hair darkness trials had a

significant and positive effect on subsequent attractiveness ratings (β_2 : $t(29)=2.73$, $p=0.011$), whereas the attractiveness of the preceding stimulus had a significant but negative influence on current judgments of face attractiveness (β_3 : $t(29)=-4.92$, $p<0.001$). That is, the effect of the preceding response was assimilative, while the effect of the preceding stimulus was contrastive.

The first result (β_2) parallels the assimilative bias seen by Kondo et al. (2012; 2013), but extends it by showing that this bias can be linked to the previous response rather than to the attractiveness of the previous face. Notably, this response bias occurs across judgment types: hair darkness ratings influenced attractiveness ratings. This cross-judgment influence echoes results from the decision-making literature, in which seemingly unrelated numerical values influence subsequent decisions (Tversky & Kahneman 1974; Critcher & Gilovich 2008). The second result (β_3) parallels other study results that have shown a contrastive effect for sequential ratings of face attractiveness and other stimulus qualities (Wedell et al. 1987; Cogan et al. 2013; Parker et al. 2008). Moreover, our design directly links this contrastive effect to the attractiveness of the previously viewed face rather than the previous response, even though the subject was attending primarily to the hair rather than the face (as evidenced by the memory results, in which subjects did not remember significantly above chance the faces in the hair darkness trials).

In our second regression analysis, we sought to determine whether hair darkness ratings also showed the same sensitivity to stimulus and response biases. We used the same model as for the attractiveness ratings, but now regressed hair darkness ratings on the hair darkness of the previous face and the attractiveness rating of the previous face (Pearson's r between previous stimulus and response: -0.11, VIF for $S_t = 1.02$, $R_{t-1} = 1.04$, $S_{t-1} = 1.03$). Here, we saw a similar pattern of

stimulus results in that there was a significant contrastive influence due to the previous stimulus (β_3 : $t(29)=-5.5$, $p<0.001$). That is, faces were judged as having darker hair if they were preceded by faces with lighter hair, and vice-versa. We also observed a marginal trend towards a assimilative influence from the previous response (β_2 : $t(29)=1.84$, $p=0.08$).

To get an estimate of the size of these effects in terms of raw ratings scores, we re-ran the regressions using non-z-scored regressors. For the attractiveness ratings model, the averaged beta weight across subjects was -0.08 for the stimulus effect (β_3 range: -0.30 to 0.13) and 0.03 for the response effect (β_2 range: -0.09 to 0.19). This means that, for the stimulus effect, holding all other variables constant, a face that is 1 "rating unit" (on the Likert scale) more attractive than the mean will tend to be rated as 0.08 rating units *less* attractive than it would have been on average. For the response effect, on the other hand, a face that is 1 rating unit more attractive than the mean will tend to be rated as 0.03 rating units *more* attractive than it would have been on average. In this case, the overall effect on a rating score is contrastive, as the contrastive effect of the preceding stimulus is larger on average than the assimilative effect of the preceding response. Given that these values estimate the shift that would occur with only a distance of 1 rating unit from the mean, and the fact that even greater variations of face attractiveness occur in the natural world, it is likely that the true effect of one face on another would be even larger, and possibly quite noticeable. The hair darkness model showed effects on the same order of magnitude: the averaged beta weight across subjects was -0.05 for the stimulus effect (β_3 range: -0.20 to 0.04) and 0.01 for the response effect (β_2 range: -0.08 to 0.08).

In summary, in study 1, we created an experimental design that effectively decorrelated possible biases due to the previous stimulus and the previous response. Our results from this design revealed that there are indeed significant biases on attractiveness judgments that occur simultaneously and in opposite directions. Furthermore, as would be predicted by a perceptual aftereffect account, the contrastive bias due to the previous stimulus does not appear to be unique to attractiveness judgments, as hair darkness ratings were biased in the same manner. Additionally, face attractiveness ratings were assimilated towards hair darkness ratings, suggesting that the response effect as well is not unique to attractiveness judgments (though the fact that we observed only a trend for the opposite case renders the interpretation of the effect less clear).

Experiment 2

In our second study, we explored the boundary conditions of the stimulus and response effects, by testing whether these effects could be obtained between items drawn from different stimulus categories (faces and places). In order to keep the previous stimulus and response decorrelated, subjects viewed face and place images on alternating trials and rated the attractiveness of each face and the “perceived temperature” of each place. If the contrastive stimulus bias observed in Experiment 1 was due to processing of attractiveness per se, then we would expect judgments of face attractiveness to be affected by the attractiveness of the preceding place. On the other hand, if the contrastive bias were attributable to a high-level perceptual aftereffect, then a stimulus from a perceptually-dissimilar category should have no influence. This design also allowed us to examine whether the assimilative response effect operated across judgments made on different stimulus categories. Because we observed both significant and trending cross-task assimilation in our previous

experiment, we expected to observe cross-category response assimilation, regardless of the presence or absence of a stimulus bias.

Methods

Stimuli

The same 242 female face images were used from experiment 1. Additionally, 373 natural scene images were selected from online sources to span a range of scene types (e.g. forests, beaches, mountains). These were cropped to 400x400 pixels to match the size of the face images. Place attractiveness ratings were acquired from the same 28 independent raters used to acquire face attractiveness ratings (see experiment 1). 7 face and 7 place images were used for all subjects as practice images, and the experimental trial images (106 female faces, 106 places) and memory task foil images (20 faces, 20 places) were randomly drawn for each subject from the remaining images. Each subject, therefore, saw a unique (though overlapping) set of images.

Procedure

To match the number of participants used in study 1, we again set our *a priori* N to 30. We ran a total of 31 Penn undergraduate participants, and excluded one participant due to a technical error, leaving us with 30 participants (18 females). Participants received course credit for their participation.

During the experiment, face and place trials were alternated, with participants rating the temperature of the place images on a scale of 1-8, and the attractiveness of face images on a scale of 1-8. The design and procedure were similar to that used in experiment 1, with a few key changes. The place trials were cued with the word “temperature” above the image, and the words “cold” and “hot” at the scale anchors.

In a practice session, subjects completed the alternating task on 14 images (7 place, 7 face) that were not shown in the main experiment and which were chosen to span the range of the attractiveness and temperature. After the practice, subjects were instructed to spread their ratings across the scale based on the images they had just seen. They were also instructed to remember all of the images for a post-experiment memory test. The memory test included 20 place and 20 face images seen during the main experiment, and 20 place and 20 face foils.

Results and Discussion

Any trial where the reaction time (RT) was less than or equal to 0.2 seconds was excluded. There was no correlation between RT and face attractiveness (Pearson's $r = 0.01$, $t(29)=0.45$ $p=0.66$) nor between RT and place temperature (Pearson's $r = 0.01$, $t(29)=0.56$ $p=0.58$).

Subjects performed above chance in the memory test for correctly identifying images seen during the experiment (places: mean=14.9 out of 20, $t(29)=29.45$, $p<0.001$; faces: mean=14.93 out of 20, $t(29)=32.66$, $p<0.001$) and for correctly identifying novel images (mean=29.9, $t(29)=45.93$, $p<0.001$). There was no significant difference between the number of faces and places remembered ($t(29)=-0.06$, $p=0.95$).

To test whether face attractiveness judgments were influenced by either the attractiveness of the preceding place stimulus or the previous response, we regressed subject-specific face attractiveness ratings against the mean attractiveness of the preceding place and the subjects' previous response to place temperature. There was a very low correlation between the temperature judgments and place attractiveness (Pearson's $r = -0.03$, averaged across subjects), suggesting that our design successfully decoupled the potential effects from the previous

stimulus and response. We also included a predictor for the mean attractiveness of the face (see Equation 1). Our test for multicollinearity using the variance inflation factor (VIF) on each of the predictors showed low numbers similar to experiment 1, indicating that multicollinearity was not a concern (VIF for $S_t = 1.02$, $R_{t-1} = 1.04$, $S_{t-1} = 1.04$).

Our regression analysis did not reveal biases on face attractiveness ratings due to either the preceding place stimulus or the previous place temperature rating (beta weights across subjects in t-test against zero: previous temperature rating (β_2): $t(29)=1.47$, $p=0.15$; previous place attractiveness (β_3): $t(29)=1.64$, $p=0.11$). To show that the absence of effects was due to the change of stimulus category rather than to other factors, we ran another regression analysis modeling the previous 4 trials, rather than just the previous trial. This allowed us to look for any significant effects from previous face trials (trials which were 2-back and 4-back) on the current face trial. (It is important to note, of course, that the stimulus and response predictors for these trials were highly correlated, since they both measured attractiveness. While significant results are meaningful, the true strength of the effect cannot be characterized.) Using this model, we showed a significant assimilative influence of both the 2-back and 4-back face attractiveness responses on the participants' current face attractiveness judgment (2-back rating: $t(29)=3.67$, $p<0.001$; 4-back rating: $t(29)=2.8$, $p=0.009$), and a significant contrastive influence due to both the 2-back and 4-back faces (2-back face attractiveness effect: $t(29)=-6.94$, $p<0.001$; 4-back face attractiveness effect: $t(29)=-4.0$, $p<0.001$) (See Figure 4.2).

These results have a number of implications. First, because we saw significant biases arising from the 2- and 4-back trials, our lack of significant weighting on the

1-back place trials cannot be explained by a lack of power to find a sequential effect. Second, these results make clear that these sequential bias effects are modulated by factors other than time, as the strength of the influence is modulated in an alternating fashion by the trial type. Third, there seems to be a limit to the influence of a previous numerical response on the current response, as there was no effect of temperature judgments on subsequent attractiveness judgments. This contrasts with the results of Experiment 1, where hair darkness judgments did have an effect on attractiveness judgments. Fourth, no general hedonic or conceptual representation of attractiveness led to a contrastive bias across trials, providing additional evidence for the perceptual nature of the contrastive bias. In general, these results suggest that both the stimulus and response effects apply across stimuli drawn from the same category (faces preceded by faces), but not across stimuli drawn from different categories (faces preceded by places).

Experiment 3

In our previous two behavioral experiments, we used an alternating sequential design to reveal biases in sequential attractiveness judgments due to both the previous stimulus and the previous response. In this third study, we attempted to locate the source of these biases in the brain (Aguirre 2007; the analyses performed here were conducted on data reported in Pegors, Kable, Chatterjee, & Epstein, under review). If the influence of previous faces on face attractiveness judgments is perceptual in nature and part of a more general phenomenon of neural adaptation, we would expect to see neural activity correlated with this effect in visual cortex. Specifically, we predicted that this effect would arise within fusiform face area (FFA), a region of visual cortex defined by its preferential response to faces

(Kanwisher, McDermott, & Chun 1997). We had no preexisting hypothesis about where we might find neural activity correlated with an assimilative response effect.

Methods

Subjects

34 healthy, right-handed volunteers with normal or corrected-to-normal vision were recruited to participate in the fMRI study. All subjects gave written informed consent according to procedures approved by the University of Pennsylvania institutional review board. One subject was excluded due to an incidental finding, two subjects were excluded due to technical issues with their anatomical images, and three subjects were excluded due to excessive motion in the scanner. This left us with 28 subjects whose functional data were analyzed (14 females, mean age 22.5).

fMRI acquisition

Scans were performed at the Hospital of the University of Pennsylvania on a 3T Siemens Trio scanner equipped with a Siemens body coil and a 32-channel head coil. High resolution T1-weighted anatomical images were acquired using a 3D MPRAGE pulse sequence (TR = 1620 ms, TE = 3 ms, TI = 950 ms, voxel size = 0.9766 x 0.9766 x 1 mm, matrix size = 192 x 256 x 160). T2*-weighted images sensitive to blood oxygenation level-dependent (BOLD) contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR = 3000 ms, TE = 30 ms, voxel size = 3x3x3mm, matrix size = 64 x 64, 46 axial slices). The slices were tilted -30 degrees from the AC-PC plane to reduce signal dropout (Deichmann et al. 2003).

Procedure

Subjects in the scanner made “low”, “average”, and “high” attractiveness ratings on a total of 144 images of male and female faces (72 female). These face images were collected from the same sources as reported in experiment 1 and were cropped and edited in the same manner. Nearly all of the female faces were images from the set used in experiments 1 and 2. Each face was presented for 1 second, with a 2 second interstimulus interval during which only a crosshair was on the screen. Face ratings were made in sequence blocks of 12 images each. These face blocks were interspersed with place attractiveness blocks, data which are not reported here. Blocks were ordered such that no block type repeated twice in a row (including fixation blocks), and the block orders were counterbalanced across runs. Between each block was an additional 9 seconds of passive fixation. 18 seconds of fixation were added at the beginning and end of each run to allow the T2* signal to reach a steady state and to model the final HRF, respectively. In total, the experiment consisted of six 4 min 57 sec scan runs, each of which was divided into two 36-s face blocks, two 36-s place blocks, and two 36-s fixation (or “rest”) blocks in which subjects passively fixated on a central crosshair.

Prior to the scan, subjects rated 24 additional faces and 24 additional places on a 1-8 Likert scale. These were meant to acclimate the subject to the full attractiveness range that they would encounter in the main experiment. Immediately after the scan session, subjects rated all of the images again from the main experiment in a randomized order on the same 1-8 scale. (Both before and after the scan, faces and places were rated in separate blocks.)

Data Analysis

Pre-processing and data analysis for individual subjects was performed using the FMRIB Software Library (FSL v.4.1.6) (Jenkinson et al. 2012; Woolrich et al. 2009; Smith et al. 2004). Functional images were corrected for differences in slice time acquisition and spatially corrected for the 30 degree tilt slice acquisition. For each run, the first six volumes were removed to account for the fMRI signal not yet reaching steady-state, and data were then motion corrected by spatially realigning each image with the central image in the run, registered to the subject-specific T1-weighted image using 6 degrees of freedom rigid-body transformations, and high-pass filtered to remove temporal frequencies below 0.0074 Hertz. Data were then smoothed with a 5mm FWHM kernel.

We used two separate general linear models: one to look for activity correlated with a contrast effect due to the attractiveness of the previous face, the other to look for an assimilative effect due to the previous response. In the "stimulus" model, we constructed the predictor of interest to be the signed difference of the mean attractiveness of the current face minus the mean attractiveness of the preceding face (mean attractiveness scores were calculated from averaging across subjects the 1-8 Likert scale rating given each face in the post-scan phase of the experiment). In the "response" model, the predictor of interest was the signed difference between the in-scan rating of attractiveness on the current trial and the in-scan rating of attractiveness on the preceding trial ("low", "average", and "high" ratings were re-coded as the numerical values 1-3). Also included in these models was a predictor for the current mean or in-scan attractiveness rating, respectively. Finally, we added predictors for in-scan reaction time, categorical predictors for face trials, place trials, and instruction screens, and nuisance regressors to account for

between-scan variability and outliers (outliers calculated with the Gabrieli Lab's Artifact Detection Tools: http://www.nitrc.org/projects/artifact_detect/). All predictors, except scan indicators, were mean-centered and convolved with a canonical HRF. Individual contrasts were registered to the standard MNI152 brain and then submitted to a random-effects group-level analysis.

Using the "stimulus" model, we first conducted an ROI analysis based on our *a priori* prediction about the source of a contrastive bias. Because our hypothesis was that the attractiveness contrast effect is perceptual in nature, we predicted that our stimulus model would reveal activity correlated with this effect in FFA, based on a number of studies that have demonstrated neural adaptation to holistic face features in this region (Andrews & Ewbank, 2004; Eger Shyns, & Kleinschmidt, 2004; Eger et al., 2005; and Winston et al., 2004). Because we did not have localizer scans for all subjects, we defined the left and right FFA by using the intersection of the subject-specific faces > places contrast from the model (thresholded at $t > 3.5$) with the FFA parcels that had been derived from a set of 40 subjects' functional localizer contrast files for faces > places (Parcels were derived using the method described in Julian et al. 2012, the only difference is that we chose a more liberal threshold of $p < 0.001$ for the contrast maps)(19 of these subjects came from our current study, during which we conducted independent localizer scans on these subjects). Subject-specific parameter estimates were extracted and averaged across each ROI.

We also conduct a wholebrain analyses on both the "stimulus" and "response" models to search for areas of activity correlated with sequential biases outside of our *a priori* defined regions. Contrasts from the previously described random-effects group-level analyses were corrected for multiple comparisons using FSL's randomize function to perform Monte-Carlo simulations which permuted the signs of wholebrain

data from individual subjects (10,000 relabelings; method based on Freedman & Lane 1983). Significance at the wholebrain level, therefore, was calculated as $p < 0.05$, corrected for multiple comparisons across the whole brain.

Results and Discussion

We used the “stimulus” model to first ask whether face-selective visual cortex showed activity consistent with a contrastive sequential effect on face attractiveness judgments. By extracting subject-specific beta estimates for the model predictor measuring the difference between the attractiveness of the previous and current face, we observed significant negative activity in left FFA ($t(25) = -2.12$, $p = 0.04$), indicating that the FFA responded more strongly to faces when they were preceded by a less-attractive face, and less strongly to faces when they were preceded by a more-attractive face. We did not see a significant effect in the right FFA ($t(27) = -1.23$, $p = 0.23$). In a wholebrain analysis, no regions survived wholebrain corrections, but at an exploratory threshold of $p < 0.001$, uncorrected, we observed bilateral clusters of activity just posterior and medial to FFA (right hemisphere: $x = 32$, $y = -70$, $z = -6$; left hemisphere: $x = 57$, $y = 28$, $z = 34$ and $x = 63$, $y = 28$, $z = 33$, all coordinates reported in MNI space). We also overlaid this “stimulus bias” map with another map from our model showing positively correlated activity with face attractiveness (See Figure 4.3). (A similar analysis of face attractiveness using subject-specific ratings is reported in Pegors, Kable, Chatterjee, & Epstein, under review; face attractiveness response in visual cortex has also been reported by Chatterjee et al. 2009). It is clear from visual inspection that the major loci of the visual clusters for face attractiveness and sequential contrast are not centered in the same regions of visual cortex.

Together, these results provide further evidence for a perceptual explanation to the contrastive attractiveness bias, and they also open up further questions about the functional relationship between the seemingly distinct populations that respond directly to the attractiveness of the current stimulus and those populations that are subtly influenced by the previous stimulus attractiveness.

A wholebrain analysis using our “response” model did not reveal any significant regions of activity after correction for multiple comparisons, or even at the more liberal threshold of $p < 0.001$ uncorrected. It is very possible that our lack of results is due to the coarseness of the response model. Since subjects were only making “low” “average” and “high” ratings in the scanner, there was less of a range over which ratings were able to be modulated. Further studies using designs in which subjects make finer-grained ratings in the scanner might be better suited to determine the source of the response bias.

General Discussion

To navigate the social world, it is important to be able to evaluate face attractiveness, but these judgments are always made in relation to a larger social and environmental context. In this paper, we provide evidence for the source of at least two contextual influences on face attractiveness judgments. First, we show that a face viewed previously, even if just for a few seconds, will create a perceptual contrast with the next face, and cause, for example, a face to appear slightly less attractive if we have just seen an extremely beautiful face. Second, we also show that this contrast effect on attractiveness ratings is attenuated if we have rated a previous face, because we are biased to give a rating that is more similar to the rating we have just given. Furthermore, our results provide evidence that the

contrast and stimulus effects are not unique to face attractiveness judgments but rather indicative of more general perceptual and decision-making mechanisms.

Sequential Contrast Bias

In both of our behavioral studies, we observed that the attractiveness of previous faces negatively predicted subsequent attractiveness ratings. This effect parallels attractiveness contrast effects seen previously in the social psychology literature (Kenrick & Gutierrez 1980; Wedell, Parducci, & Geiselman 1987; Cogan, Parker, & Zellner 2013). We also observed a strong contrast effect for hair darkness in our first study; that is, faces were judged to have darker hair when preceded by faces with lighter hair, and vice versa. The fact that we showed contrastive biases for two separate perceptual characteristics (attractiveness and hair darkness), that we did not show a contrastive bias across different perceptual categories, and that we observed activity correlated with the contrastive bias in visual cortex, suggests that this bias may be due to a general perceptual “aftereffect” phenomenon that is driven by neural adaptation in visual cortex. Neural adaptation has been seen for face characteristics like gender (Podrebarac et al. 2013), identity and expression (Winston et al. 2004), but this is the first time that neural adaptation has been observed for face attractiveness. Because the main loci of the effect was seen outside of FFA, it is possible that, rather than revealing a population that adapts to perceptual attractiveness per se, this contrast effect arises due to adaptation for lower-level visual features of faces that may correlate with attractiveness, such as skin smoothness or shape. This interpretation would also be extremely interesting, as it would suggest that a behavioral bias on attractiveness ratings is due to a relatively elemental perceptual bias. On the other hand, the fact that our fMRI design intermixed male and female faces likely decreased the holistic perceptual similarity

between images; a design which uses only male or female faces might increase adaptation in FFA. Regardless, the implications of our results are far-reaching: every exposure to a face, even briefly (whether on TV, within our group of friends, etc.), serves to dynamically change our basic perception (and therefore, evaluation) of face attractiveness.

By using the word 'bias', we do not mean to imply that this contrast effect is a maladaptive mechanism. Specifically, neural adaptation has been proposed to improve coding efficiency (Clifford et al., 2007; Wainwright, 1999; Wark et al., 2007) by shifting the neural tuning in a way to prevent response saturation and improve discriminability around previously observed stimuli. In other words, our neural system is constantly adapting to incoming input so that we can best process and discriminate the stimuli in our current environment. In fact, neural adaptation seems to be a ubiquitous mechanism in the brain. For example, recent work in reward processing has shown that many of the neurons in macaque frontal cortex constantly adjust their firing rates according to the range of reward (e.g. amounts of juice) available, even on a trial-by-trial basis (Padoa-Schioppa 2009, Kennerley et al. 2011).

Sequential Assimilation Bias

Both of our behavioral studies also reveal that previous ratings given to faces positively predict current attractiveness ratings. These results replicate the assimilative effect on face attractiveness seen by Kondo et al. (2012, 2013), but we extend their findings by linking the effect directly to the previous rating. It may be the case, then, that Kondo et al. observed an overall assimilative effect due to the fact that their brief image presentations created a weaker stimulus bias relative to the response bias.

In the psychophysics literature, one interpretation of the assimilative relationship between past and current judgments is that it is a reflection of the previous judgment acting as a reference point for comparison. Decarlo and Cross provide evidence for this “relative judgment” model by showing that the assimilation effect on loudness estimates was decreased when subjects were instructed to make their judgments relative to a single reference loudness, presumably meaning that subjects shifted their reference away from the previous trial (1990).

Our results showing no assimilative effect of responses to place temperature on face attractiveness ratings differ from the “anchor and adjust” account in the decision-making literature, in which previous values can be completely unrelated to the current judgment yet still have an assimilative influence (Tversky and Kahneman 1974). On the other hand, we did show *cross-task* assimilation within the same stimulus category, suggesting some level of generality to the effect. Future studies should address, therefore, exactly the set of conditions under which this bias is present.

Relevance to Sequential Tasks

Our results reveal at least two bias-inducing mechanisms that reinforce researchers’ motivation to randomize trial order for each subject when acquiring mean estimates of stimuli. Since randomization is already common practice, our results in no way invalidate the many studies that use sequential rating designs. Rather, having an awareness of these potential biases may help researchers when considering other appropriate experimental designs and analyses, by taking into account the fact that both previous subject responses and stimulus presentations may affect behavior in a measurable way on subsequent trials.

Conclusion

To what extent does our environment and past experience influence the way we evaluate people around us? Our data suggest that our evaluations are constantly shifting and adapting to our world, even within the course of seconds. Rather than these phenomena being specific to the social or hedonic realm, we instead provide evidence that they are subserved by more general mechanisms, provided a possible link between these attractiveness biases and a wide range of biases described in social psychology, psychophysics, and judgment and decision-making literature.

Figure 4.1

“Alternating” experimental design. Subjects rated either the attractiveness or hair darkness of each female face on a Likert scale of 1-8.

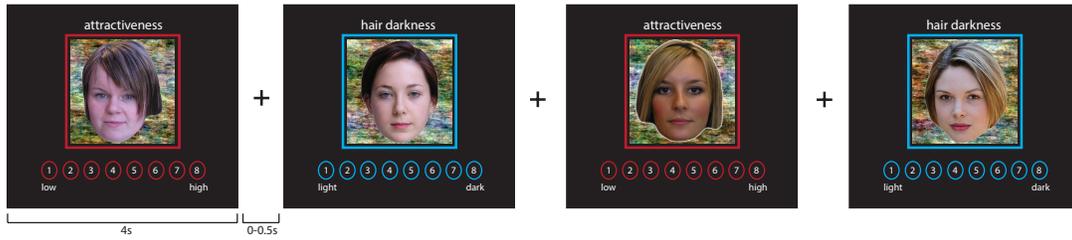


Figure 4.2

Regression results for face/place alternating design. Face attractiveness ratings were regressed against the previous rating and the previous response of the four preceding trials. Because of the alternating design, trials 1-back and 3-back were always place trials in which subjects judged "place temperature", and trials 2-back and 4-back were always face trials in which subjects judged attractiveness. Neither the response to place temperature nor the underlying attractiveness of places significantly predicted current face attractiveness ratings, but face trials even 4 trials back showed predictive power related both to the subjects' response and the mean attractiveness of the face.

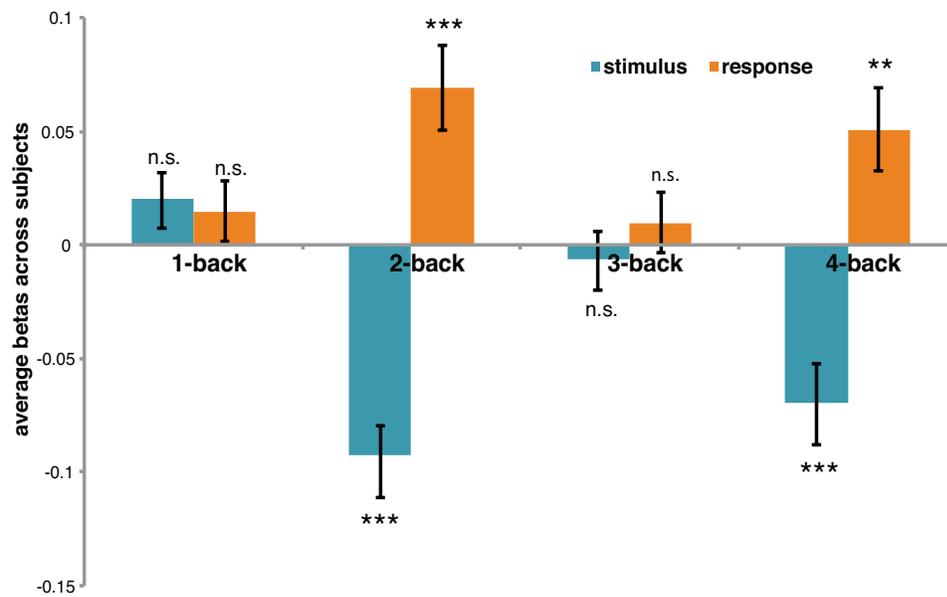
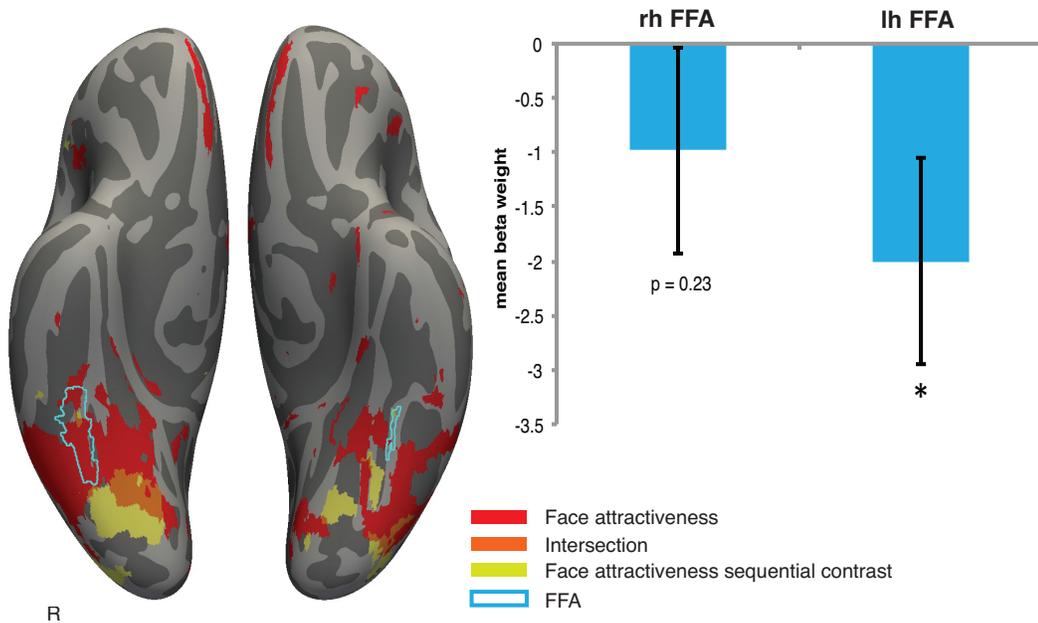


Figure 4.3

Wholebrain image displaying correlated activity with the sequential contrast effect (yellow), face attractiveness (red), and the overlap between the two contrasts (orange). The FFA ROI is outlined in light blue. Beta maps were thresholded at $p < .02$ uncorrected to show activity that did not reach significance at the wholebrain level after corrections for multiple comparisons. While there is some activity correlated with the contrast effect in FFA, it is clear that the locus of the effect is more posterior and medial. Activity correlated with face attractiveness is somewhat overlapping with the contrast effect, but large regions are also unique to one or the other effect.



CHAPTER 5 – Future Directions

In this dissertation, I have described research in which we explored the neural and behavioral underpinnings of how we perceive and evaluate visual beauty. Our results accomplished both major goals of this work, to advance understanding of the processes underlying face and place beauty, and to address broader questions related to perception, reward processing, and decision-making. These results also opened up many further questions for exploration. Below, I will outline a number of questions that arose from the studies described in each of the three chapters, and I will then discuss one of these questions in greater detail, outlining potential experimental designs and hypotheses.

Questions for further exploration

In chapter 2, we compared activity in the brain while subjects rated face and place attractiveness. We found overlapping activity in a region of ventromedial prefrontal cortex (vmPFC), but we also showed that only faces elicited responses in lateral orbitofrontal cortex (latOFC) and in visual cortex. Why did places also not elicit responses in these regions? Might it be the case that latOFC only responds to stimuli associated with more primary rewards? Were the faces simply more rewarding than the places? In visual cortex, all higher-level visual regions responded to face attractiveness but only one region (lateral occipital cortex) showed a slight response to place attractiveness. Do these differences reflect the fact that face attractiveness is more closely associated with a common set of visual features, and that these visual features in turn are modulating the response? Our landscape images spanned a much wider feature range, and so place attractiveness may have been less correlated with any particular set of features across the images. On the

other hand, these differences may be driven by top-down attentional affects. Does face attractiveness modulate attention in a way that place attractiveness does not?

In chapter 3, we showed two distinct regions of activity in latOFC even when subjects were not making explicit ratings of face attractiveness: a region that showed a bulk response to faces (when compared to places), and a region that was correlated with face attractiveness. Other kinds of gustatory and olfactory stimuli have been shown to elicit response in latOFC. Do these rewards also elicit these two distinct types of responses – one that exhibits a bulk response to the category of reward and one that tracks the specific value of the reward? What are the separate functions of these regions in relation to the larger reward and/or emotion network? During the passive evaluation task, we also observed that vmPFC no longer showed a correlated response with face attractiveness. Is this because vmPFC only represents value in choice-making contexts? Or is vmPFC encoding value in a way that is non-linearly related to explicit attractiveness ratings?

In chapter 4, we used a novel experimental design to show that sequential face attractiveness judgments contrast away from the previous stimulus but assimilate towards the previous response. Interestingly, even un-related ratings modulated subsequent judgments (e.g. attractiveness judgments were assimilated towards previous hair darkness ratings). At the same time, ratings given to places did not influence face judgments. Does the strength of this assimilative effect vary continuously based on similarity to the current judgment? Also, is the contrast effect truly due to perceptual aftereffects from the previous face (i.e. the next face actually looks slightly different), or is this effect better explained by some type of cognitive rescaling, in which the mapping between face features and the ratings scale subtly shifts from trial to trial?

The perceptual versus cognitive nature of the sequential contrast effect

In chapter 4, we demonstrated that sequential attractiveness judgments are susceptible to influences from the previous trials. Judgments are assimilated towards the previous response and contrasted away from the previous stimulus. We interpreted the stimulus contrast effect to be caused by perceptual aftereffects, a phenomenon in the literature that has been shown to influence even high-level stimuli such as faces (Webster et al. 2004). Another interpretation of the contrast effect is that this effect is not perceptual but cognitive: the subject may not *perceive* the next face differently, but rather, the subject may be subtly re-mapping facial features to the ratings scale on a trial-by-trial basis. This cognitive interpretation would also suggest that the reason place attractiveness does not influence face attractiveness is that the subject uses separate scales for places and faces. Therefore, viewing a place will not invoke remapping of the subjects' face scale.

Two possible experiments may begin to tease apart whether the contrast effect is due to perceptual adaptation or cognitive remapping. Both experimental designs would use the base alternating design in which face attractiveness and hair darkness judgments would alternate after every trial, but an additional condition manipulation would be pseudo-randomized within the trials. This would allow us to independently measure the strength of the contrast and assimilation effect between two conditions.

In a first experiment, the alternating design would introduce the condition of short vs. long face exposures. While all trials lengths will be held constant, the presentation time of the face will either be short (e.g. 500 ms) or long (e.g. 4 s). Previous research has shown that face adaptation effects strength logarithmically

with the length of the adaption period and decay exponentially over the test duration (Leopold et al. 2005; Rhodes et al. 2007). These facts would suggest that, if the contrast effect is perceptual in nature, short face presentation times in our design would weaken the contrast effect in comparison to the long presentation times, since the adaptation period is shorter and the test duration (length between the face exposure and next face) is longer.

Would the cognitive remapping account predict a different the strength of the contrast effect for short and long exposures? Face attractiveness judgments with exposures of only 100 ms are highly correlated with attractiveness judgments without time constraint (Willis and Todorov 2006). These data suggest that face features that drive attractiveness ratings are perceived within a very short period of time, and so in our own design, the length of exposure should not influence the contrast effect in the case of cognitive remapping. One could argue, though, that short-exposure faces may still correlate less well than long-exposure faces with the average ratings that we use to model attractiveness for hair darkness trials. If this were true, the contrast effect would be weaker for the short duration trials, not because of a weaker perceptual aftereffect but because we had less precise predictors. We could pre-empt this concern by acquiring two independent sets of attractiveness ratings to match the short and long exposure times during the experiment itself.

In a second experiment, rather than varying the exposure time, we would pseudo-randomize the facing-angle of the face stimuli (e.g. looking 45 degrees to the left and 45 degrees to the right). Benton, Jennings, and Chatting varied face viewpoint angle across two face identities to show that identity adaptation decreased as the angle between the adapting face and the test face increased (2006). These

results would lead to the prediction that, if the contrast effect is perceptual in nature, trials preceded by different-viewpoint trials would have a smaller contrast effects if preceded by same-viewpoint trials. Alternatively, if the contrast effect is due to cognitive remapping, it is unlikely that the contrast effect would be affected by changes in facing-direction, as the major face features would still be visible in either condition.

Bibliography

- Adolphs, R. (2009). The Social Brain: Neural Basis of Social Knowledge. *Annual Review of Psychology*, 693–716. doi:10.1146/annurev.psych.60.110707.163514.The
- Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *NeuroImage*, 35(4), 1480–94. doi:10.1016/j.neuroimage.2007.02.005
- Andrews, T. J., & Ewbank, M. P. (2004). Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *NeuroImage*, 23(3), 905–13. doi:10.1016/j.neuroimage.2004.07.060
- Appleton, J. (1975). *The Experience of Landscape*. New York: John Wiley.
- Aron, A., Fisher, H., Mashek, D. J., Strong, G., Li, H., & Brown, L. L. (2005). Reward, motivation, and emotion systems associated with early-stage intense romantic love. *Journal of Neurophysiology*, 94(1), 327–37. doi:10.1152/jn.00838.2004
- Balling, J. D., & Falk, J. H. (1982). Development of Visual Preference for Natural Environments. *Environment and Behavior*, 14(1), 5–28. doi:10.1177/0013916582141001
- Barron, H. C., Dolan, R. J., & Behrens, T. E. J. (2013). Online evaluation of novel choices by simultaneous representation of multiple memories. *Nature Neuroscience*, 16(10), 1492–1498. doi:10.1038/nn.3515
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–27. doi:10.1016/j.neuroimage.2013.02.063
- Benton, C. P., Jennings, S. J., & Chatting, D. J. (2006). Viewpoint dependence in adaptation to facial identity. *Vision Research*, 46(20), 3313–25. doi:10.1016/j.visres.2006.06.002
- Berman, M. G., Jonides, J., & Kaplan, S. (2008). The cognitive benefits of interacting with nature. *Psychological Science*, 19(12), 1207–12. doi:10.1111/j.1467-9280.2008.02225.x
- Bruckert, L., Bestelmeyer, P., Latinus, M., Rouger, J., Charest, I., Rousselet, G. a, ... Belin, P. (2010). Vocal attractiveness increases by averaging. *Current Biology : CB*, 20(2), 116–20. doi:10.1016/j.cub.2009.11.034
- Cai, X., & Padoa-Schioppa, C. (2012). Neuronal Encoding of Subjective Value in Dorsal and Ventral Anterior Cingulate Cortex. *Journal of Neuroscience*, 32(11), 3791–3808. doi:10.1523/JNEUROSCI.3864-11.2012

- Chatterjee, A. (2004). Prospects for a cognitive neuroscience of visual aesthetics. *Bulletin of Psychology and the Arts*, 4(2), 55–60.
- Chatterjee, A., Thomas, A., Smith, S. E., & Aguirre, G. K. (2009a). The neural response to facial attractiveness. *Neuropsychology*, 23(2), 135–43. doi:10.1037/a0014430
- Chatterjee, A., Thomas, A., Smith, S. E., & Aguirre, G. K. (2009b). The neural response to facial attractiveness. *Neuropsychology*, 23(2), 135–43. doi:10.1037/a0014430
- Chatterjee, A., & Vartanian, O. (2014). Neuroaesthetics. *Trends in Cognitive Sciences*, 1–6. doi:10.1016/j.tics.2014.03.003
- Chiavaras, M. M., LeGoualher, G., Evans, A., & Petrides, M. (2001). Three-dimensional probabilistic atlas of the human orbitofrontal sulci in standardized stereotaxic space. *NeuroImage*, 13(3), 479–96. doi:10.1006/nimg.2000.0641
- Chiavaras, M. M., & Petrides, M. (2000). Orbitofrontal sulci of the human and macaque monkey brain. *The Journal of Comparative Neurology*, 422(1), 35–54.
- Chib, V. S., Rangel, A., Shimojo, S., & O’Doherty, J. P. (2009). Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *The Journal of Neuroscience*, 29(39), 12315–20. doi:10.1523/JNEUROSCI.2575-09.2009
- Clifford, C. W. G., Webster, M. a, Stanley, G. B., Stocker, A. a, Kohn, A., Sharpee, T. O., & Schwartz, O. (2007). Visual adaptation: neural, psychological and computational aspects. *Vision Research*, 47(25), 3125–31. doi:10.1016/j.visres.2007.08.023
- Cloutier, J., Heatherton, T. F., Whalen, P. J., & Kelley, W. M. (2008). Are attractive people rewarding? Sex differences in the neural substrates of facial attractiveness. *Journal of Cognitive Neuroscience*, 20(6), 941–51. doi:10.1162/jocn.2008.20062
- Cogan, E., Parker, S., & Zellner, D. a. (2013). Beauty beyond compare: Effects of context extremity and categorization on hedonic contrast. *Journal of Experimental Psychology. Human Perception and Performance*, 39(1), 16–22. doi:10.1037/a0031020
- Collins, J. a., & Olson, I. R. (2014). Beyond the FFA: The role of the ventral anterior temporal Lobes in face processing. *Neuropsychologia*, 61, 65–79. doi:10.1016/j.neuropsychologia.2014.06.005
- Conway, B. R., & Rehding, A. (2013). Neuroaesthetics and the trouble with beauty. *PLoS Biology*, 11(3), e1001504. doi:10.1371/journal.pbio.1001504

- Critcher, C., & Gilovich, T. (2008). Incidental environmental anchors. *Journal of Behavioral Decision Making*, 251(October 2007), 241–251. doi:10.1002/bdm
- Critchley, H. (1996). Olfactory neuronal responses in the primate orbitofrontal cortex: analysis in an olfactory discrimination task. *Journal of Neurophysiology*, 75(4).
- Croxson, P. L., Johansen-Berg, H., Behrens, T. E. J., Robson, M. D., Pinski, M. a, Gross, C. G., ... Rushworth, M. F. S. (2005). Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. *The Journal of Neuroscience*, 25(39), 8854–66. doi:10.1523/JNEUROSCI.1311-05.2005
- Cunningham, M. (1986). Measuring the Physical in Physical Attractiveness : Quasi-Experiments on the Sociobiology of Female Facial Beauty. *Journal of Personality and Social Psychology*, 50(5), 925–935.
- Cunningham, M. R., Roberts, A. R., Barbee, A. P., & Druen, P. B. (1995). "Their ideas of beauty are, on the whole, the same as ours": Consistency and variability in the cross-cultural perception of female physical attractiveness. *Journal of Personality and Social Psychology*, 68(2), 261–279. doi:10.1037//0022-3514.68.2.261
- Decarlo, L. T., & Cross, D. V. (1990). Sequential Effects in Magnitude Scaling : Models and Theory. *Journal of Experimental Psychology: General*, 119(4), 375–396.
- Deichmann, R., Gottfried, J. ., Hutton, C., & Turner, R. (2003). Optimized EPI for fMRI studies of the orbitofrontal cortex. *NeuroImage*, 19(2), 430–441. doi:10.1016/S1053-8119(03)00073-9
- Eagly, A. H., Ashmore, R. D., Makhijani, M. G., & Longo, L. C. (1991). What is Beautiful is good, But. . . : A Meta-Analytic Review of Research on the Physical Attractiveness Stereotype. *Psychological Bulletin*, 110(1), 109–128.
- Eger, E., Schweinberger, S. R., Dolan, R. J., & Henson, R. N. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *NeuroImage*, 26(4), 1128–39. doi:10.1016/j.neuroimage.2005.03.010
- Eger, E., Schyns, P. G., & Kleinschmidt, A. (2004). Scale invariant adaptation in fusiform face-responsive regions. *NeuroImage*, 22(1), 232–42. doi:10.1016/j.neuroimage.2003.12.028
- Enquist, M., & Arak, A. (1994). Symmetry, Beauty, and Evolution. *Nature*, 372, 169–172.
- Epstein, R., Harris, a, Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: recognition, navigation, or encoding? *Neuron*, 23(1), 115–25.

- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601. doi:10.1038/33402
- Falk, J. H., & Balling, J. D. (2009). Evolutionary Influence on Human Landscape Preference. *Environment and Behavior*, *42*(4), 479–493. doi:10.1177/0013916509341244
- Feingold, A. (1992). Good-Looking People Are Not What We Think. *Psychological Bulletin*, *111*(2), 304–341.
- Freedman, D., & Lane, D. (1983). A nonstochastic interpretation of reported significance levels. *Journal of Business & Economic Statistics*, *1*(4), 292–298.
- Gibson, J. J. (1933). Adaptation, after-effect and contrast in the perception of curved lines. *Journal of Experimental Psychology*, *16*(1), 1–31.
- Gottfried, J. a, & Zald, D. H. (2005). On the scent of human olfactory orbitofrontal cortex: meta-analysis and comparison to non-human primates. *Brain Research. Brain Research Reviews*, *50*(2), 287–304. doi:10.1016/j.brainresrev.2005.08.004
- Graham, D. J., & Field, D. J. (2007). Statistical regularities of art images and natural scenes: spectra, sparseness and nonlinearities. *Spatial Vision*, *21*(1-2), 149–64. doi:10.1163/156856807782753877
- Grammer, K., & Thornhill, R. (1994). Human (*Homo sapiens*) facial attractiveness and sexual selection: the role of symmetry and averageness. *Journal of Comparative Psychology*, *108*(3), 233–42.
- Griffin, A., & Langlois, J. (2006). Stereotype directionality and attractiveness stereotyping: is beauty good or is ugly bad? *Social Cognition*, *24*(2), 187–206.
- Hagerhall, C., Purcell, T., & Taylor, R. (2004). Fractal dimension of landscape silhouette outlines as a predictor of landscape preference. *Journal of Environmental Psychology*, *24*(2), 247–255. doi:10.1016/j.jenvp.2003.12.004
- Halberstadt, J., & Rhodes, G. (2000). The Attractiveness of Nonface Averages: Implications for an Evolutionary Explanation of the Attractiveness of Average Faces. *Psychological Science*, *11*(4), 285–289. doi:10.1111/1467-9280.00257
- Halberstadt, J., & Rhodes, G. (2003). It's not just average faces that are attractive: computer-manipulated averageness makes birds, fish, and automobiles attractive. *Psychonomic Bulletin & Review*, *10*(1), 149–56.
- Han, K.-T. (2007). Responses to Six Major Terrestrial Biomes in Terms of Scenic Beauty, Preference, and Restorativeness. *Environment and Behavior*, *39*(4), 529–556. doi:10.1177/0013916506292016

- Han, K.-T. (2009). An Exploration of Relationships Among the Responses to Natural Scenes: Scenic Beauty, Preference, and Restoration. *Environment and Behavior*, 42(2), 243–270. doi:10.1177/0013916509333875
- Haxby, J., & Gobbini, M. (2011). Distributed neural systems for face perception. *The Oxford Handbook of Face Perception*, 93–110.
- Haxby, J. V, Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539), 2425–30. doi:10.1126/science.1063736
- Haxby, J. V, Hoffman, E. a, & Gobbini, M. I. (2002). Human neural systems for face recognition and social communication. *Biological Psychiatry*, 51(1), 59–67.
- Herzog, T. R. (1981). A Cognitive Analysis of Preference for Natural Environments: Mountains, Canyons, and Deserts. *Landscape Journal*, 6(2), 140–152.
- Herzog, T. R. (1985). A cognitive analysis of preference for waterscapes. *Journal of Environmental Psychology*, 5(3), 225–241.
- Ishai, A. (2007). Sex, beauty and the orbitofrontal cortex. *International Journal of Psychophysiology*, 63(2), 181–5. doi:10.1016/j.ijpsycho.2006.03.010
- Ishizu, T., & Zeki, S. (2011). Toward a brain-based theory of beauty. *PloS One*, 6(7), e21852. doi:10.1371/journal.pone.0021852
- Ishizu, T., & Zeki, S. (2013). The brain's specialized systems for aesthetic and perceptual judgment. *European Journal of Neuroscience*, (July 2012), 1–8. doi:10.1111/ejn.12135
- Izuma, K., Saito, D. N., & Sadato, N. (2008). Processing of social and monetary rewards in the human striatum. *Neuron*, 58(2), 284–94. doi:10.1016/j.neuron.2008.03.020
- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, 62(2), 782–90. doi:10.1016/j.neuroimage.2011.09.015
- Jesteadt, W., Luce, R. D., & Green, D. M. (1977). Sequential effects in judgments of loudness. *Journal of Experimental Psychology. Human Perception and Performance*, 3(1), 92–104.
- Johnston, V. S., Hagel, R., Franklin, M., Fink, B., & Grammer, K. (2001). Male facial attractiveness: evidence for hormone-mediated adaptive design. *Evolution and Human Behavior*, 22(4), 251–267. doi:10.1016/S1090-5138(01)00066-6

- Johnstone, R. (1994). Female preference for symmetrical males as a by-product of selection for mate recognition. *Nature*, *372*, 172–175.
- Julian, J. B., Fedorenko, E., Webster, J., & Kanwisher, N. (2012). An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *NeuroImage*, *60*(4), 2357–64. doi:10.1016/j.neuroimage.2012.02.055
- Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron*, *63*(6), 733–45. doi:10.1016/j.neuron.2009.09.003
- Kalick, S., & Zebrowitz, L. (1998). Does human facial attractiveness honestly advertise health? Longitudinal Data on an Evolutionary Question. *Psychological Science*, *9*(1), 8–13.
- Kamenetzky, J. (1959). Contrast and convergence effects in ratings of foods. *Journal of Applied Psychology*, *1*(1), 47–52.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of ...*, *17*(11), 4302–4311.
- Kaplan, S. (1995). The restorative benefits of nature: Toward an integrative framework. *Journal of Environmental Psychology*, *15*(3), 169–182. doi:10.1016/0272-4944(95)90001-2
- Kaplan, S., Kaplan, R., & Wendt, J. S. (1972). Rated preference and complexity for natural and urban visual material. *Perception & Psychophysics*, *12*(1), 354–356.
- Kennerley, S. W., Behrens, T. E. J., & Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, *14*(12), 1581–9. doi:10.1038/nn.2961
- Kenrick, D. T., & Gutierrez, S. E. (1980). Contrast Effects and Judgments of Physical Attractiveness: When Beauty Becomes a Social Problem. *Journal Personality and Social Psychology*, *38*(1), 131–140.
- Kernis, M. H., & Wheeler, L. (1981). Beautiful friends and ugly strangers: Radiation and contrast effects in perception of same-sex partners. *Personality and Social Psychology Bulletin*, *7*(4), 617–620.
- Kim, H., Adolphs, R., O'Doherty, J., & Shimojo, S. (2007). Temporal isolation of neural processes underlying face preference decisions. *PNAS*, *104*(46), 18253–18258.
- Kim, H., Shimojo, S., & O'Doherty, J. P. (2011). Overlapping responses for the expectation of juice and money rewards in human ventromedial prefrontal cortex. *Cerebral Cortex*, *21*(4), 769–76. doi:10.1093/cercor/bhq145

- Kirk, U., Skov, M., Christensen, M. S., & Nygaard, N. (2009). Brain correlates of aesthetic expertise: a parametric fMRI study. *Brain and Cognition*, *69*(2), 306–15. doi:10.1016/j.bandc.2008.08.004
- Kondo, A., Takahashi, K., & Watanabe, K. (2012). Sequential effects in face-attractiveness judgment. *Perception*, *41*, 43–49. doi:10.1068/p7116
- Kondo, A., Takahashi, K., & Watanabe, K. (2013). Influence of gender membership on sequential decisions of face attractiveness. *Attention, Perception & Psychophysics*. doi:10.3758/s13414-013-0533-y
- Kranz, F., & Ishai, A. (2006). Face perception is modulated by sexual preference. *Current Biology: CB*, *16*(1), 63–8. doi:10.1016/j.cub.2005.10.070
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, *12*(5), 535–40. doi:10.1038/nn.2303
- Kringelbach, M. L., & Rolls, E. T. (2003). Neural correlates of rapid reversal learning in a simple model of human social interaction. *NeuroImage*, *20*(2), 1371–83. doi:10.1016/S1053-8119(03)00393-8
- Kringelbach, M. L., & Rolls, E. T. (2004). The functional neuroanatomy of the human orbitofrontal cortex: evidence from neuroimaging and neuropsychology. *Progress in Neurobiology*, *72*(5), 341–72. doi:10.1016/j.pneurobio.2004.03.006
- Kweon, B.-S., Ulrich, R. S., Walker, V. D., & Tassinari, L. G. (2007). Anger and Stress: The Role of Landscape Posters in an Office Setting. *Environment and Behavior*, *40*(3), 355–381. doi:10.1177/0013916506298797
- Langlois, J. H., Kalakanis, L., Rubenstein, A. J., Larson, A., Hallam, M., & Smoot, M. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin*, *126*(3), 390–423. doi:10.1037//0033-2909.126.3.390
- Langlois, J., & Roggman, L. (1990). attractive faces are only average. *Psychological Science*, *1*(2), 115–121.
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S. T., & van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition & Emotion*, *24*(8), 1377–1388. doi:10.1080/02699930903485076
- Laumann, K., Gärling, T., & Stormark, K. M. (2003). Selective attention and heart rate responses to natural and urban environments. *Journal of Environmental Psychology*, *23*(2), 125–134. doi:10.1016/S0272-4944(02)00110-X

- Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An automatic valuation system in the human brain: evidence from functional neuroimaging. *Neuron*, *64*(3), 431–9. doi:10.1016/j.neuron.2009.09.040
- Leopold, D. a, O'Toole, a J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, *4*(1), 89–94. doi:10.1038/82947
- Leopold, D. a, Rhodes, G., Müller, K.-M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings. Biological Sciences / The Royal Society*, *272*(1566), 897–904. doi:10.1098/rspb.2004.3022
- Levy, D. J., & Glimcher, P. W. (2011). Comparing Apples and Oranges: Using Reward-Specific and Reward-General Subjective Value Representation in the Brain. *Journal of Neuroscience*, *31*(41), 14693–14707. doi:10.1523/JNEUROSCI.2218-11.2011
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology*, *22*(6), 1027–38. doi:10.1016/j.conb.2012.06.001
- Liang, X., Zebrowitz, L., & Zhang, Y. (2010). Neural activation in the “reward circuit” shows a nonlinear response to facial attractiveness. *Social Neuroscience*, *5*(3), 320–334. doi:10.1080/17470911003619916.Neural
- Lin, A., Adolphs, R., & Rangel, A. (2012). Social and monetary reward learning engage overlapping neural substrates. *Social Cognitive and Affective Neuroscience*, 274–281. doi:10.1093/scan/nsr006
- Little, A. C., & Hancock, P. J. B. (2002). The role of masculinity and distinctiveness in judgments of human male facial attractiveness. *British Journal of Psychology (London, England : 1953)*, *93*(Pt 4), 451–64.
- Matthews, W. J., & Stewart, N. (2009). Psychophysics and the judgment of price : Judging complex objects on a non-physical dimension elicits sequential effects like those in perceptual tasks. *Judgment and Decision Making*, *4*(1), 64–81.
- McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science*, *149*(3688), 1115–1116.
- Mcnamee, D., Rangel, A., & Doherty, J. P. O. (2013). Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex. *Nature Neuroscience*, (February), 1–9. doi:10.1038/nn.3337
- Mealey, L., Bridgstock, R., & Townsend, G. C. (1999). Symmetry and perceived facial attractiveness: a monozygotic co-twin comparison. *Journal of Personality and Social Psychology*, *76*(1), 151–8.

- Mende-Siedlecki, P., Said, C. P., & Todorov, A. (2012). The social evaluation of faces: a meta-analysis of functional neuroimaging studies. *Social Cognitive and Affective Neuroscience*. doi:10.1093/scan/nsr090
- Minear, M., & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, & Computers: A Journal of the Psychonomic Society, Inc*, 36(4), 630–3.
- Mitchell, R., & Carson, R. (1989). *Using surveys to value public good: the contingent valuation method*. New York: Resources for the Future.
- Neter, J., Wasserman, W., & Kutner, M. H. (1989). Multicollinearity and Its Effects. In *Applied Linear Regression Models* (2nd ed., pp. 295–305). Homewood, IL: Irwin.
- Noonan, M. P., Walton, M. E., Behrens, T. E. J., Sallet, J., Buckley, M. J., & Rushworth, M. F. S. (2010). Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 107(47), 20547–52. doi:10.1073/pnas.1012246107
- O Scalaidhe, S. P., Wilson, F. A. W., & Goldman-Rakic, P. S. (1997). Areal Segregation of Face-Processing Neurons in Prefrontal Cortex. *Science*, 278(5340), 1135–1138. doi:10.1126/science.278.5340.1135
- O’Doherty, J., Rolls, E. T., Francis, S., Bowtell, R., McGlone, F., Kobal, G., ... Ahne, G. (2000). Sensory-specific satiety-related olfactory activation of the human orbitofrontal cortex. *Neuroreport*, 11(4), 893–7.
- O’Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia*, 41(2), 147–55.
- O’Toole, a J., Deffenbacher, K. a, Valentin, D., McKee, K., Huff, D., & Abdi, H. (1998). The perception of face gender: the role of stimulus structure in recognition and classification. *Memory & Cognition*, 26(1), 146–60.
- O’Toole, a. ., Price, T., Vetter, T., Bartlett, J. ., & Blanz, V. (1999). 3D shape and 2D surface textures of human faces: the role of “averages” in attractiveness and age. *Image and Vision Computing*, 18(1), 9–19. doi:10.1016/S0262-8856(99)00012-8
- Ongür, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex*, 10(3), 206–19.
- Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *The Journal of Neuroscience: The Official Journal of the*

- Society for Neuroscience*, 29(44), 14004–14. doi:10.1523/JNEUROSCI.3751-09.2009
- Parker, S., Bascom, J., Rabinovitz, B., & Zellner, D. (2008). Positive and negative hedonic contrast with musical stimuli. *Psychology of Aesthetics, Creativity, and the Arts*, 2(3), 171–174. doi:10.1037/1931-3896.2.3.171
- Perrett, D. I., Burt, D. M., Penton-Voak, I. S., Lee, K. J., Rowland, D. a, & Edwards, R. (1999). Symmetry and Human Facial Attractiveness. *Evolution and Human Behavior*, 20(5), 295–307. doi:10.1016/S1090-5138(99)00014-8
- Perrett, D., Lee, K., Penton-Voak, I., Rowland, D., Yoshikawa, S., Burt, D., ... Akamatsu, S. (1998). Effects of sexual dimorphism on facial attractiveness. *Nature*, 394(August), 884–887.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(37), 9984–8. doi:10.1523/JNEUROSCI.2131-07.2007
- Podrebarac, S. K., Goodale, M. a, van der Zwan, R., & Snow, J. C. (2013). Gender-selective neural populations: evidence from event-related fMRI repetition suppression. *Experimental Brain Research*, 226(2), 241–52. doi:10.1007/s00221-013-3429-0
- Purcell, A., & Lamb, R. (1994). preference or preferences for landscape? *Journal of Environmental Psychology*, 14, 195–209.
- Rajimehr, R., Young, J. C., & Tootell, R. B. H. (2009). An anterior temporal face patch in human cortex, predicted by macaque maps. *Proceedings of the National Academy of Sciences of the United States of America*, 106(6), 1995–2000. doi:10.1073/pnas.0807304106
- Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: is beauty in the perceiver's processing experience? *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc*, 8(4), 364–82. doi:10.1207/s15327957pspr0804_3
- Redies, C. (2007). A universal model of esthetic perception based on the sensory coding of natural stimuli. *Spatial Vision*, 21(1-2), 97–117. doi:10.1163/156856807782753886
- Redies, C., Hasenstein, J., & Denzler, J. (2007). Fractal-like image statistics in visual art: similarity to natural scenes. *Spatial Vision*, 21(1-2), 137–48. doi:10.1163/156856807782753921
- Repp, B. (1997). The Aesthetic Quality of a Quantitatively Average Music Performance: Two Preliminary Experiments. *Music Perception*, 14(4), 419–444.

- Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Review of Psychology*, *57*, 199–226. doi:10.1146/annurev.psych.57.102904.190208
- Rhodes, G., Chan, J., Zebrowitz, L. a, & Simmons, L. W. (2003). Does sexual dimorphism in human faces signal health? *Proceedings. Biological Sciences / The Royal Society*, *270 Suppl*, S93–5. doi:10.1098/rsbl.2003.0023
- Rhodes, G., Hickford, C., & Jeffery, L. (2000). Sex-typicality and attractiveness: are supermale and superfemale faces super-attractive. *British Journal of Psychology*, *91*, 125–140.
- Rhodes, G., Jeffery, L., Clifford, C. W. G., & Leopold, D. a. (2007). The timecourse of higher-level face aftereffects. *Vision Research*, *47*(17), 2291–6. doi:10.1016/j.visres.2007.05.012
- Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G., & Nakayama, K. (2003). Fitting the Mind to the World: Face Adaptation and Attractiveness Aftereffects. *Psychological Science*, *14*(6), 558–566. doi:10.1046/j.0956-7976.2003.psci
- Rhodes, G., Proffitt, F., Grady, J. M., & Sumich, A. (1998). Facial symmetry and the perception of beauty. *Psychonomic Bulletin & Review*, *5*(4), 659–669. doi:10.3758/BF03208842
- Rhodes, G., Simmons, L. W., & Peters, M. (2005). Attractiveness and sexual behavior: Does attractiveness enhance mating success? *Evolution and Human Behavior*, *26*(2), 186–201. doi:10.1016/j.evolhumbehav.2004.08.014
- Rhodes, G., Sumich, A., & Byatt, G. (1999). Are average facial configurations attractive only because of their symmetry. *Psychological Science*, *10*(1), 52–58.
- Rhodes, G., & Tremewan, T. (1996). Averageness, exaggeration, and facial attractiveness. *Psychological Science*, *7*(2), 105–110.
- Rolls, E. T., Critchley, H. D., Browning, A. S., & Inoue, K. (2006). Face-selective and auditory neurons in the primate orbitofrontal cortex. *Experimental Brain Research*, *170*(1), 74–87. doi:10.1007/s00221-005-0191-y
- Rolls, E. T., Sienkiewicz, Z. J., & Yaxley, S. (1989). Hunger Modulates the Responses to Gustatory Stimuli of Single Neurons in the Caudolateral Orbitofrontal Cortex of the Macaque Monkey. *The European Journal of Neuroscience*, *1*(1), 53–60.
- Schifferstein, H. N. J., & Frijters, J. A. N. E. R. (1992). Contextual and sequential effects on judgments of sweetness intensity. *Perception & Psychophysics*, *52*(3), 243–255.
- Schifferstein, H. N. J., & Kuiper, W. E. (1997). Sequence effects in hedonic judgments of taste stimuli. *Perception & Psychophysics*, *59*(6), 900–912.

- Sescousse, G., Redouté, J., & Dreher, J.-C. (2010). The architecture of reward value coding in the human orbitofrontal cortex. *The Journal of Neuroscience*, *30*(39), 13095–104. doi:10.1523/JNEUROSCI.3501-10.2010
- Simoncelli, E., & Olshausen, B. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, *24*, 1193–216.
- Small, D. M., Zald, D. H., Jones-Gotman, M., Zatorre, R. J., Pardo, J. V, Frey, S., & Petrides, M. (1999). Human cortical gustatory areas: a review of functional neuroimaging data. *Neuroreport*, *10*(1), 7–14.
- Smardon, R. C. (1988). Perception and aesthetics of the urban environment: Review of the role of vegetation. *Landscape and Urban Planning*, *15*(1-2), 85–106. doi:10.1016/0169-2046(88)90018-7
- Smith, J., & Melara, R. (1990). Aesthetic preference and syntactic prototypicality in music: 'Tis the gift to be simple. *Cognition*, *34*, 279–288.
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., ... Matthews, P. M. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, *23 Suppl 1*, S208–19. doi:10.1016/j.neuroimage.2004.07.051
- Smith, D. V, Hayden, B. Y., Truong, T.-K., Song, A. W., Platt, M. L., & Huettel, S. a. (2010). Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(7), 2490–5. doi:10.1523/JNEUROSCI.3319-09.2010
- Soler, C., Núñez, M., Gutiérrez, R., Núñez, J., Medina, P., Sancho, M., ... Núñez, a. (2003). Facial attractiveness in men provides clues to semen quality. *Evolution and Human Behavior*, *24*(3), 199–207. doi:10.1016/S1090-5138(03)00013-8
- Stenger, V. A. (n.d.). Technical considerations for BOLD fMRI of the orbitofrontal cortex. In D. Zald & S. Rauch (Eds.), *The Orbitofrontal Cortex* (pp. 423–446). New York: Oxford University Press.
- Strane, K., & Watts, C. (1977). Females judged by attractiveness of partner. *Perceptual and Motor Skills*, *45*, 225–226.
- Taylor, R. P., & Micolich, A. P. (1999). Fractal analysis of Pollock's drip paintings. *Nature*, *399*(June), 422–423.
- Thornhill, R., & Gangestad, S. (1999). Facial attractiveness. *Trends in Cognitive Sciences*, *3*(12), 452–460.
- Tsao, D. Y., Moeller, S., & Freiwald, W. a. (2008). Comparing face patch systems in macaques and humans. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(49), 19514–9. doi:10.1073/pnas.0809662105

- Tsao, D. Y., Schweers, N., Moeller, S., & Freiwald, W. a. (2008). Patches of face-selective cortex in the macaque frontal lobe. *Nature Neuroscience*, *11*(8), 877–9. doi:10.1038/nn.2158
- Tsukiura, T., & Cabeza, R. (2011). Remembering beauty: roles of orbitofrontal and hippocampal regions in successful memory encoding of attractive faces. *NeuroImage*, *54*(1), 653–60. doi:10.1016/j.neuroimage.2010.07.046
- Tversky, a, & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, *185*(4157), 1124–31. doi:10.1126/science.185.4157.1124
- Ulrich, R. (1984). View through a window may influence recovery from surgery. *Science*, *224*(4647), 420–421. doi:10.1126/science.6143402
- Valentine, T., Darling, S., & Donnelly, M. (2004). Why are average faces attractive? The effect of view and averageness on the attractiveness of female faces. *Psychonomic Bulletin & Review*, *11*(3), 482–7.
- Vartanian, O., Navarrete, G., Chatterjee, A., Fich, L. B., Leder, H., Modroño, C., ... Skov, M. (2013). Impact of contour on aesthetic judgments and approach-avoidance decisions in architecture. *Proceedings of the National Academy of Sciences of the United States of America*, *110* Suppl, 10446–53. doi:10.1073/pnas.1301227110
- Von Der Heide, R. J., Skipper, L. M., & Olson, I. R. (2013). Anterior temporal face patches: a meta-analysis and empirical study. *Frontiers in Human Neuroscience*, *7*(February), 17. doi:10.3389/fnhum.2013.00017
- Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Research*, *39*(23), 3960–74.
- Walker, D., & Vul, E. (2014). Hierarchical encoding makes individuals in a group seem more attractive. *Psychological Science*, *25*(1), 230–5. doi:10.1177/0956797613497969
- Ward, L. (1990). Mixed-method mixed-modality psychophysical scaling. *Perception & Psychophysics*, *48*(6), 571–582.
- Ward, L., & Lockhead, G. (1971). Response system processes in absolute judgment. *Perception & Psychophysics*, *9*(18).
- Wark, B., Lundstrom, B. N., & Fairhall, A. (2007). Sensory adaptation. *Current Opinion in Neurobiology*, *17*(4), 423–9. doi:10.1016/j.conb.2007.07.001
- Watson, K. K., & Platt, M. L. (2012). Social signals in primate orbitofrontal cortex. *Current Biology*, *22*(23), 2268–73. doi:10.1016/j.cub.2012.10.016

- Webster, M., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*(April), 357–360. doi:10.1038/nature02361.1.
- Wedell, D., Parducci, A., & Geiselman, E. (1987). A Formal analysis of ratings of physical attractiveness: Successive contrast and simultaneous assimilation. *Journal of Experimental Social Psychology*, *23*, 230–249.
- White, M. P., Alcock, I., Wheeler, B. W., & Depledge, M. H. (2013). Would You Be Happier Living in a Greener Urban Area? A Fixed-Effects Analysis of Panel Data. *Psychological Science*, (April). doi:10.1177/0956797612464659
- White, M., Smith, A., Humphryes, K., Pahl, S., Snelling, D., & Depledge, M. (2010). Blue space: The importance of water for preference, affect, and restorativeness ratings of natural and built scenes. *Journal of Environmental Psychology*, *30*(4), 482–493. doi:10.1016/j.jenvp.2010.04.004
- Willis, J., & Todorov, A. (2006). First Impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science*, *17*(7), 592–598.
- Wilson, E. O. (1984). *Biophilia*. Cambridge: Harvard University Press.
- Winston, J. S., Henson, R. N. a, Fine-Goulden, M. R., & Dolan, R. J. (2004). fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *Journal of Neurophysiology*, *92*(3), 1830–9. doi:10.1152/jn.00155.2004
- Winston, J. S., O’Doherty, J., Kilner, J. M., Perrett, D. I., & Dolan, R. J. (2007). Brain systems for assessing facial attractiveness. *Neuropsychologia*, *45*(1), 195–206. doi:10.1016/j.neuropsychologia.2006.05.009
- Wohlgemuth, A. (1911). *On the after-effect of seen movement* (The British.). Cambridge: Cambridge University Press.
- Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., ... Smith, S. M. (2009). Bayesian analysis of neuroimaging data in FSL. *NeuroImage*, *45*(1 Suppl), S173–86. doi:10.1016/j.neuroimage.2008.10.055
- Yarkoni, T., Barch, D. M., Gray, J. R., Conturo, T. E., & Braver, T. S. (2009). BOLD correlates of trial-by-trial reaction time variability in gray and white matter: a multi-study fMRI analysis. *PloS One*, *4*(1), e4257. doi:10.1371/journal.pone.0004257
- Yue, X., Vessel, E. a, & Biederman, I. (2007). The neural basis of scene preferences. *Neuroreport*, *18*(6), 525–9. doi:10.1097/WNR.0b013e328091c1f9
- Zebrowitz, L. a, & Montepare, J. M. (2008). Social Psychological Face Perception: Why Appearance Matters. *Social and Personality Psychology Compass*, *2*(3), 1497. doi:10.1111/j.1751-9004.2008.00109.x

Zebrowitz, L. a., & Rhodes, G. (2004). Sensitivity to “Bad Genes” and the Anomalous Face Overgeneralization Effect: Cue Validity, Cue Utilization, and Accuracy in Judging Intelligence and Health. *Journal of Nonverbal Behavior*, 28(3), 167–185. doi:10.1023/B:JONB.0000039648.30935.1b

Zellner, D. a, Rohm, E. a, Bassetti, T. L., & Parker, S. (2003). Compared to what? Effects of categorization on hedonic contrast. *Psychonomic Bulletin & Review*, 10(2), 468–73.