The Dynamic Interactive Nature of Person Construal

A dissertation submitted by

Jonathan B. Freeman

In partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Psychology

COMMITTEE:

Nalini Ambady, Ph.D. (Chair)

Stanford University

Phillip J. Holcomb, Ph.D.

Keith B. Maddox, Ph.D.

Tufts University

Kerri L. Johnson, Ph.D.

University of California, Los Angeles

TUFTS UNIVERSITY

May 2012

ABSTRACT

A dynamic interactive framework for person construal is proposed. It argues that the perception of other people is accomplished by a dynamical system involving continuous interaction between social categories, stereotypes, high-level cognitive states, and the low-level processing of facial, vocal, and bodily cues. This system permits lower-level sensory perception and higher-order social cognition to dynamically coordinate across multiple interactive levels of processing to give rise to stable person construals. A recurrent connectionist model of this system is described, which accounts for a wide range of experimental findings from a computer mouse-tracking technique that examines social categorization in real time. These include evidence for a) continuously evolving category representations and a dynamic competition process underlying categorization (Studies 1–4); b) continuous face–voice interaction during categorization (Studies 5–6), and c) the continuous top-down influence of stereotype activations on categorization (Studies 7–8). Together, across 8 studies, mouse-tracking and computational simulations provide converging evidence for the dynamic and interactive nature of person construal. Implications and new predictions arising from this framework are discussed.

ACKNOWLEDGEMENTS

Words cannot do it justice to say how indebted I am to Nalini Ambady for the incredible support she has given me over the years. This work could have never developed without the immense freedom to explore my intellectual interests that Nalini gave me in graduate school. She put an unusual level of faith in me, allowing me to delve into what at first blush seemed like distant, unrelated research areas and incorporate them with our primary interests in person perception and our home field of social psychology. It was clear this research program was going to be a high-risk one, and yet she went out on a limb to support my interests. I feel that Nalini's disciplinary breadth and her willingness to let her students travel in sometimes dramatically new directions is truly remarkable. It is one of the many reasons that I am forever grateful for having the opportunity to work with her and to learn from her brilliance and insightfulness. There is no doubt that I was extremely fortunate to train with her, and my only hope is that I have learned from her abilities and am able to carry them forward with me in the future to support my own students in the decades to come.

I would have never gotten to graduate school without the wonderfully supportive mentors I was fortunate to work with during my undergraduate years at NYU. Kerri Johnson initially drew my attention to person perception research. She gave me an enormous amount of independence and flexibility. She would continually tell me to "think big," in her words, and to never stop working at the broadest level. Her theoretical approach then was very exciting to me, and it served to jumpstart my own work in graduate school. She has been an outstanding mentor and friend to me. Liz Phelps and her postdoc at the time, Daniela Schiller, also gave me an unusual degree of independence in

exploring my interests in social neuroscience and delving into neuroimaging work. Finally, although far outside our field, Lisa Duggan was very inspirational to me and indirectly led me to the present work. What I learned from her is that very automatic, routine processes occurring within an individual may in part be the result of large-scale cultural and historical processes. Lisa, together with my mentors in person perception and social neuroscience, paved the way for my quest to understand mental phenomena at multiple levels of analysis: between neuron and society, between biological substrate and psychological experience, or between lower-level sensory perception and higher-order social cognition. I am very grateful for the time and support my undergraduate mentors gave me, which played a fundamental role in bringing about this research program.

With this developing perspective brewing away, it was only when in my first year of graduate school that I read Michael Spivey's *Continuity of Mind* that I experienced a rapid shift in research trajectory, a phase transition, if you will. It was then that my emergent research interests could be grounded and take on a more formal nature through the insights of dynamical systems. Massive connections started to become apparent to me between the dynamical cognitive science Spivey's work drew my attention to, classic social psychological work on person perception, and more current work on ecologically valid, sensory-based person construal. At this juncture is where the present work lies, and I have a great deal of gratitude for Spivey for helping set this in motion.

I am also very thankful for Phil Holcomb. Although coming from the distant area of language processing, Phil was willing to entertain and support my ideas using event-related potentials to study social categorization. I hope he has found this work as fruitful as I have, and I am appreciative of the latitude he gave me in exploring my social

psychological interests. I am also appreciative of all the positive support and feedback he has given me over the years and for his ongoing guidance. In addition, I am grateful to Ken Nakayama who allowed me to visit his lab at Harvard during the latter half of graduate school. The vision science approach there opened my eyes to a wildly different and rich perspective on, ironically, some of the same phenomena studied in recent social psychological research on person construal. I learned an immense amount from working with Ken and Phil, and I am in their debt for their support and for entertaining my ideas.

Starting as far back as my undergraduate years, Rick Dale has been an invaluable “mentor-from-afar” and collaborator, offering a tremendous amount of guidance and dedicating an enormous amount of time to me. From mouse-tracking to computational modeling to explanatory pluralism, he always made important and razor-sharp insights, and I thank him for the many enjoyable conversations. None of this work could have ever taken off without him. I am also indebted to Nick Rule, Anne Krendl, and Max Weisbuch for helping mentor me early on in graduate school and for being fantastic colleagues. In addition, I am thankful for my committee members and for the collaborators who helped bring these specific studies into being, namely, Kristin Pauker, Evan Apfelbaum, Andrew Penner, Aliya Saperstein, and Matthias Scheutz. Finally, I am grateful to the National Institutes of Health and National Science Foundation for financially supporting this work.

Last, but hardly least, I am deeply grateful to the wonderful friends I have had the greatest pleasure of knowing over these past years and to my family. Somehow they have put up with me, which I find genuinely remarkable in and of itself. They keep me fresh, happy, and sane, and I would be nowhere without them.

Thank you!

TABLE OF CONTENTS

THE DYNAMIC INTERACTIVE NATURE OF PERSON CONSTRUAL

INTRODUCTION

With only a fleeting glimpse, a constellation of near-instant judgments are often made about another person. Although frequently warned not to “judge a book by its cover,” our tendency to make meaning out of the sensory information availed by others is typically beyond our conscious control. From minimal cues afforded by the face, voice, and body, we unwittingly infer the intentions, thoughts, personalities, emotions, and category memberships (e.g., sex, race, age) of those around us. While some of these judgments may be expectancy-driven and biased by our stereotypes (Brewer, 1988; Fiske & Neuberg, 1990; Macrae & Bodenhausen, 2000), others may be surprisingly accurate and expose humans’ exquisite ability to perceive other people from only the briefest of observations (Ambady, Bernieri, & Richeson, 2000).

This unique ability to perceive other people, however, is plagued by a basic contradiction. As readily and rapidly as we may dispel our judgments of others, each judgment requires an astonishing complexity of mental processing; and despite their complexity, however, they occur with remarkable ease. From a single face, for example, any number of perceptions (e.g., sex, emotion) are immediately available, but each requires the integration of an enormous amount of information. Unlike objects, other people are highly complex stimuli, embedded in a rich set of contexts and grounded in multiple sensory modalities. All the features and configural properties of a person’s face must be bound together, along with that person’s hair and array of bodily cues. Auditory cues of a person’s voice are available as well, and these must be bound together with the person’s visual cues to form a coherent social percept. Such a complexity of bottom-up sensory information is matched, however, by a similar complexity in top-down information sources that are uniquely present in person perception. For example, people bring a great deal of prior knowledge, stereotypic expectations, and affective and motivational states to the process of perceiving others. The influences of these top-down factors may often, I will argue, seep down into the perceptual process itself. In person perception, therefore, there is a vast array of information—bottom-up and top-down—that must rapidly conspire to drive perceptions in the impressively short time it takes to arrive at a simple judgment of another person.

If person perception is characterized by, on the one hand, being highly complex, and on the other being highly efficient, research has historically placed a great deal of focus on the latter. Seminal work in social psychology by Allport (1954), Sherif (1967), and Tajfel (1969), for example, argued that individuals perceive others via spontaneous, perhaps inevitable, category-based impressions that are highly efficient and designed to economize on mental resources. Since then, a vast array of studies have demonstrated that such category-based impressions bring about a host of cognitive, affective, and behavioral outcomes. Mere exposure to another person has long been known to automatically trigger a relevant social category (e.g., sex, race, age), and along with that category, its corresponding knowledge structure. Activating category knowledge, it has been shown, then spontaneously changes how individuals think about others, feel about them, and behave towards them, often in ways that may occur nonconsciously (e.g., Bargh, 1994; Bargh, 1999; Brewer, 1988; Devine, 1989; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995; Fiske & Neuberg, 1990; Gilbert & Hixon, 1991; Sinclair & Kunda, 1999). For example, activated category representations shape subsequent encoding and representation of any information relevant to the target (Bodenhausen, 1988). Their associated knowledge structures (e.g., stereotypes) become a lens that molds the judgments perceivers make and impressions they form (Brewer, 1988; Fiske & Neuberg, 1990) and distorts perceivers' memories of a target (Hamilton & Sherman, 1994). Perceivers' behavior is even subject to these influences as well (Bargh, 1997). For instance, activation of the category, elder, can lead perceivers to walk more slowly (Bargh, Chen, & Burrows, 1996) and activation of the category, professor, can boost performance on general knowledge tests (Dijksterhuis & Van Knippenberg, 1998). Unquestionably, social category activation is highly consequential.

It became clear with such findings that social categorization had a powerful role in shaping interpersonal interactions. Naturally, therefore, social psychological research placed a great deal of focus on the downstream implications of categorization. Given this focus, research on person perception by and large investigated how perceivers make judgments from written behavioral descriptions, often in ways that influence downstream processing (but see McArthur & Baron, 1983). Real-world social targets, however, are not generally encountered through behavioral descriptions. Rather, in real life perceivers encounter other people first through sensory cues of the face, voice, and body. The theoretical and empirical work examining the links between lower-level perceptual processing and higher-order social cognition began only recently (see Bodenhausen & Macrae, 2006; Zebrowitz, 2006). Although it was long understood that perceivers frequently categorize other people along a variety of dimensions (e.g., sex, race, age) from mere exposure to their face (Brewer, 1988; Fiske & Neuberg, 1990; Strangor, Lynch, Duan, & Glas, 1992), the mechanisms and perceptual determinants underlying these categorizations received considerably less attention.

While social psychologists were documenting the downstream implications of perceiving others, cognitive psychologists and neuroscientists were examining person perception from a different perspective. They were concentrating their efforts on investigating the perceptual mechanisms of face processing (Bruce & Young, 1986; Burton, Bruce, & Johnston, 1990; Calder & Young, 2005; Farah, Wilson, Drain, & Tanaka, 1998; Haxby, Hoffman, & Gobbini, 2000). Recently, by integrating the social cognitive framework of person perception with insights from the cognitive literature on face processing, a growing body of research has begun to link lower-level perceptual processing with higher-order social cognition. This emerging body of work has come to be referred to as "person construal" research. Traditional social cognition research focused on the relatively high-level cognitive processes involved in person categorization and individuation, especially how these shape downstream phenomena (e.g., stereotyping, behavior). Person construal research, on the other hand, seeks to understand the lower-level perceptual mechanisms that produce these social cognitive phenomena in the first place.

A Dynamic Interactive Framework for Person Construal

Over the past two decades, researchers have developed a number of models of person perception, including models that explain how we reason about other people and infer their personality traits, how we categorize and individuate, and how explicit knowledge and memory of other people is learned, stored, and accessed (Bodenhausen & Macrae, 1998; Brewer, 1988; Chaiken & Trope, 1999; Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Fiske, Cuddy, Glick, & Xu, 2002; Fiske & Neuberg, 1990; Higgins, 1996; Kunda & Thagard, 1996; Read & Miller, 1998b; Smith & DeCoster, 1998; Srull & Wyer, 1989; van Overwalle & Labiouse, 2004). These models tend to place categorization as a starting point, after which subsequent interpersonal phenomena are richly explained (e.g., impressions, memory, behavior). Thus, the focus of these models is not to explain the categorization process; it is to explain the higher-order social cognitive processing that comes after.

Person construal research seeks to examine the lower-level perceptual mechanisms and determinants of categorization, including how categories and stereotypes are activated from cues of the face, voice, and body. Here, I will propose a framework that details how such lower-level perceptual processing contributes to higher-order social cognitive phenomena. This framework utilizes popular approaches to cognition, namely connectionism and dynamical systems theory (Kelso, 1995; Port & van Gelder, 1995; Rogers & McClelland, 2004; Rumelhart, Hinton, & McClelland, 1986; Smolensky, 1989; Spivey, 2007). Recently, researchers have applied connectionist models to understand social cognitive phenomena as well (e.g., Kunda & Thagard, 1996; Read & Miller, 1993, 1998a; Read, Vanman, & Miller, 1997; Smith & DeCoster, 1998, 1999; van Overwalle, 2007; van Overwalle & Labiouse, 2004; Zebrowitz, Fellous, Mignault, & Andreoletti, 2003). In the present work, I will apply connectionism and dynamical systems theory to explain the process of person construal. As such, I aim to provide a framework that explains social categorization processes at a perceptual level and links these processes to the higher-order social cognitive phenomena emphasized in prior models of person perception.

Top-Down and Bottom-Up Interactivity

In perceiving the world, we are continually extracting sensory information to guide our attempts in discerning what it is that lies before us. Even with the most mundane kinds of construal, such as perceiving objects or environments, we bring a great deal of knowledge to the perceptual process. This is only truer in the case of perceiving other people. Our rich set of prior experiences with another person or the regularities in our experience with whole groups of people (e.g., sex, race, age) undoubtedly provide a lens through which we construe others. Beyond the prior knowledge that might contextualize perception, our everyday encounters with others are also replete with complex affective and motivational states. Though there is much prior knowledge about the objects or environments we might encounter, this only pales in comparison to what is brought to the table when perceiving other people. We may have stereotypic beliefs about people of a certain sex, we may feel disdain for someone who has made us cry, or we may be motivated to make a good impression in order to land the job. In short, there is an enormity of prior knowledge and high-level states that may be brought to bear on the perception of our social world. Although traditionally it was long assumed that perception is primarily a bottom-up phenomenon and insulated from any top-down influence of higher-order processes (e.g., Fodor, 1983; Marr, 1982), it is becoming increasingly clear that perception arises instead from both bottom-up and top-down influences, likely mediated by large-scale neural oscillations (e.g., Engel, Fries, & Singer, 2001; Gilbert & Sigman, 2007). Even the earliest of responses in primary visual cortex, for example, are altered by top-down factors (Li, Piëch, & Gilbert, 2004). I argue, therefore, that our prior knowledge and expectations about people, our stereotypes, and our affective and motivational states may all interact with incoming sensory information in the perceptual process to shape person construal.

The person construal process invites another form of interactivity as well, one that is driven directly by the incoming sensory information itself. Whereas the perception of an object, for example, generally affords only one focal type of construal (e.g., "that's a table"), multiple construals are simultaneously availed to person perceivers, including sex, race, age, emotion, or inferences of personality characteristics, to name a few. Given how many construals are available, sometimes the perceptual cues supporting certain construals will, by chance, overlap. For instance, the cues specifying another person's sex and emotional state can overlap (Becker, Kenrick, Neuberg, Blackwell, & Smith, 2007). An adult's facial features might by chance happen to overlap with the facial features more common in babies or with the facial features of another person we know, in turn shaping our inferences of his or her personality characteristics (Zebrowitz & Montepare, 2008). Thus, certain person construals may be thrown into interaction with one another because they are directly confounded in the bottom-up sensory information itself.

Time Dependence and Continuous Temporal Dynamics

I argue that the process of person construal is dynamic, in the sense that it takes time and fluctuates over time, and that representations triggered during this process are inherently time-dependent. Accordingly, after catching sight of another person, representations of social categories and stereotypes would be dynamically evolving across hundreds of milliseconds until stabilizing over time. Thus, at each moment during the categorization process, representations would be varying as a function of time, making time-dependent transitions between, for instance, ~0% activation and ~100% activation.

This is not particularly surprising when considering how a social categorization would be implemented in an actual human brain.

For instance, there is now a great deal of evidence suggesting that mental representations, as realized in the brain, are neuronal populations that convey information (e.g., “he’s a man!”) through patterns of activity distributed across many neurons (Rogers & McClelland, 2004; Spivey, 2007; Spivey & Dale, 2004). This was confirmed with regard to representations of the face by studies that recorded populations of temporal cortex neurons in nonhuman primates (Rolls & Tovee, 1995; Sugase, Yamane, Ueno, & Kawano, 1999). Thus, most modern-day accounts assume that mental representations, such as a representation of a social category, involve continuous changes in a pattern of neuronal activity (e.g., Rogers & McClelland, 2004; Smith & Ratcliff, 2004; Spivey, 2007; Spivey & Dale, 2004; Usher & McClelland, 2003). For instance, about 50% of a face’s identity is transiently represented in macaque temporal cortex as early as only 80 ms after a face’s presentation, but the remaining 50% of its representation gradually accumulates over the following hundreds of milliseconds (Rolls & Tovee, 1995). Thus, in early moments of processing representations of a face’s category memberships would reflect a rough “gist,” because the initial rough sketch of the face is partially consistent with multiple interpretations (e.g., both male and female). As the ongoing accrual of more and more information continues, however, the pattern of neuronal activity would gradually sharpen into an increasingly confident representation (e.g., male) while other competing, partially-active representations (e.g., female) are pushed out (Freeman, Ambady, Midgley, & Holcomb, 2011; Smith & Ratcliff, 2004; Spivey & Dale, 2004; Usher & McClelland, 2003).

This approach proposes therefore that a single category representation (e.g., male) would not discretely activate at an instantaneous moment after a target's presentation, nor would a single category representation transition from zero activation to full activation across time. Instead, this approach suggests that person construal would involve alternative, competing categories that are simultaneously and partially active, and these would evolve over time until stabilizing onto ultimate construals. Given such proposed continuous dynamics, I argue that person construal is a temporally dynamic process and that person construal phenomena (e.g., a social categorization; activation of a stereotype) would be best understood as gradual time-dependent transitions between mental states (e.g., from state A, the initial sight of another person, transitioning to state B, the ~100% confident recognition that the person is a White man). Further, I argue that during this time-dependent process, representations of a person's category memberships (e.g., male, White) as well as other candidate category memberships (e.g., female, Black, Asian) would be rapidly fluctuating over time until achieving a stable, steady state.

Complex Integration

Person construal routinely involves complex integration. Even the simplest of construals, such as categorizing a person's sex, requires simultaneous integration of an enormous amount of information. For instance, all the various cues of the internal face in addition to peripheral cues such as hair must be integrated into a coherent interpretation of a target's sex. In many person construal tasks of the laboratory, this may be the only information available to perceivers—and even these simple tasks require already a substantial integration among cues. In everyday person construal and more complicated laboratory tasks, however, the integration is even more complex. For instance, perceivers receive information from multiple sensory modalities at the same time. Thus, to perceive the sex of real-world social targets, all the sex-specifying cues of the face and body arriving in the visual system must be integrated together with the vocal cues arriving in the auditory system. Moreover, not only does bottom-up sensory information need to be integrated, so too do top-down information sources, as described earlier. For instance, high-level motivational states influence the perception of a face's race (Pauker et al., 2009). Moreover, priming context, expectations, stereotypes, cultural knowledge, among many other top-down factors, shape basic perceptions (e.g., Balcetis & Dunning, 2006; Eberhardt, Dasgupta, & Banaszynski, 2003; Hugenberg & Bodenhausen, 2004; Johnson, Pollick, & McKay, 2010; MacLin & Malpass, 2001; Pauker, Rule, & Ambady, 2010). Thus, there is a complexity of information involving many sources—some bottom-up, some top-down—that must be integrated together in a very short amount of time to perceive others.

Summary of Theoretical Claims

In consideration of the above, I propose that perceptions of other people are accomplished by a dynamical system in which they gradually emerge through the ongoing interaction between categories, stereotypes, high-level cognitive states, and the low-level processing of facial, vocal, and bodily cues. As such, this system permits lower-level sensory perception and higher-order social cognition to continuously coordinate across multiple interactive levels of processing to give rise to stable person construals. A computational model, introduced below, will capture this theory of person construal.

Introduction of the Model

A general diagram of the proposed dynamic interactive model of person construal appears in Figure 1. It is a recurrent connectionist network with stochastic interactive activation (McClelland, 1991). The figure depicts a number of pools; in specific instantiations of the model, each pool will contain a variety of nodes (e.g., MALE, BLACK, AGGRESSIVE, FEMALE CUES). Specific details on the model's structure as well as relevant background on connectionist networks of this type are found in Appendix A. The model provides an approximation of the kind of processing that might take place in a human brain (Rogers & McClelland, 2004; Rumelhart et al., 1986; Smolensky, 1989; Spivey, 2007), specifically in the context of perceiving other people. An important aim of the present work is to demonstrate how this model can capture a wide range of person construal phenomena, including those explored in the current experimental studies.

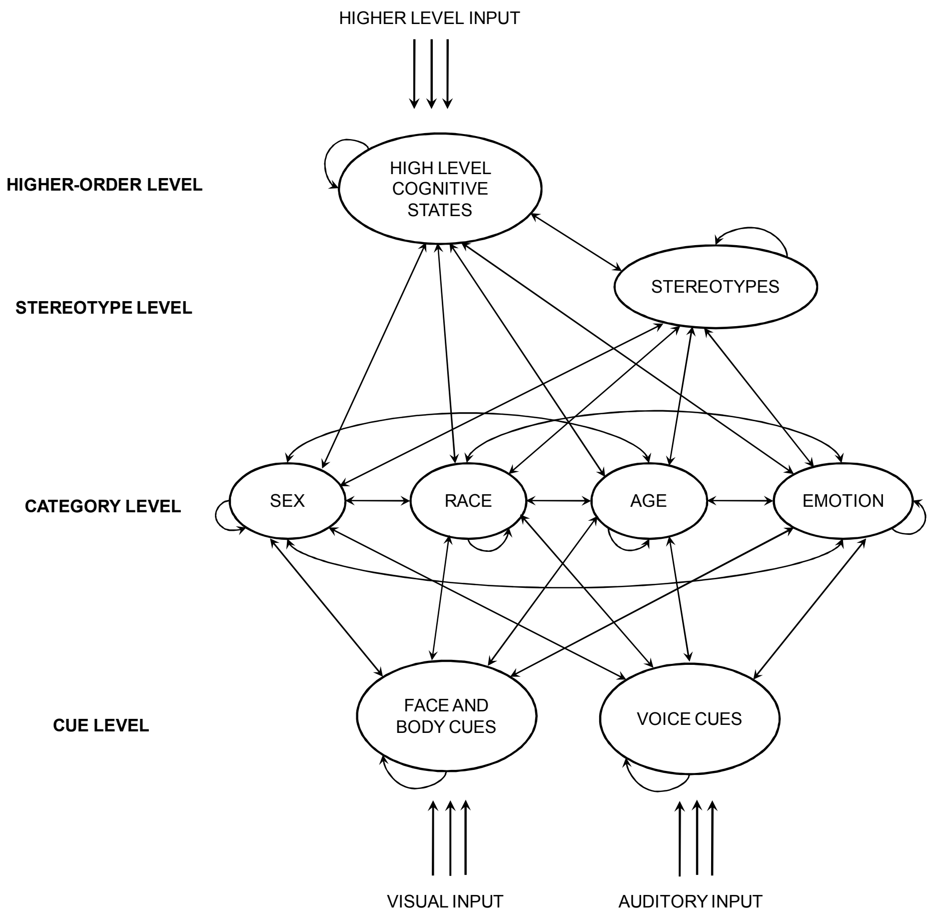

Figure 1. A general diagram of the dynamic interactive model of person construal. The Cue Level contains two pools: a Face/Body Cues pool, which contains detectors for visual features, and a Voice Cues pool, which contains detectors for auditory features. Cue nodes are directly stimulated by bottom-up input from visual and auditory systems. The Category Level contains a number of competitive pools that correspond with social category dimensions, such as Sex, Race, Age, and Emotion (although any dimensions may be used, e.g., Occupation). Each of these pools contain category nodes (e.g., Female, White, Janitor, Happy). The Stereotype Level contains one pool that includes nodes for all category-related stereotypes (e.g., Aggressive). The Higher-Order Level contains nodes corresponding with high-level cognitive states, such as task demands, motivations, prejudice, goals, among others. Higher-order nodes are directly stimulated by input from higher levels of mental processing (e.g., motivational systems or top-down attentional systems).

Initially, the network is stimulated simultaneously by both bottom-up and top-down inputs (see Figure 1). For model instantiations in the present work, this will include inputs such as visual input of another's face, auditory input of another's voice, or higher-level input from a top-down attentional system that directs attention toward a given category dimension (e.g., sex or race) for a particular task. In other simulations, however, the network could be stimulated by other kinds of bottom-up (e.g., vocal cues) and top-down (e.g., motivations, prejudice) inputs. Each model instantiation contains a variety of nodes that are organized into, at most, four interactive levels of processing (one level representing each of the following: cues, categories, stereotypes, and high-level cognitive states). Every node has a transient level of activation at each moment in time. This activation corresponds with the strength of a tentative hypothesis that the node is represented in the input. Once the network is initially stimulated, activation flows among all nodes simultaneously as a function of their connection weights. Activation is also altered by a small amount of random noise, making the system's states inherently probabilistic. Because many connections between nodes are bidirectional, this flow results in a continual back-and-forth of activation between many nodes in the system. As such, nodes in the system continually readjust each other's activation and mutually constrain one another to find an overall pattern of activation that best fits the inputs. Gradually, the flows of activation lead the network to converge on a stable, steady state, where the activation of each node reaches an asymptote. This final steady state, it is argued, corresponds to an ultimate perception of another person. Through this ongoing mutual constraint-satisfaction process, multiple sources of information—both bottom-up cues and top-down factors—are interacting over time toward integrating into a stable person construal. Thus, this model grounds the framework's theoretical claims on the dynamic and interactive nature of person construal. See Appendix A for specific details on the model's structure.

Overview of the Present Research

Across 8 studies, I aim to provide converging evidence for the theory that person construal is a dynamic interactive process, and to show how the proposed computational model can capture this process.

Part I focuses more on the dynamic aspects of person construal. First, in Studies 1–3, a computer mouse-tracking technique is used to test the existence of a dynamic competition process that is argued to underlie our ability to slot others into social categories, such as sex and race. In this process, partially-active categories (e.g., male and female; White and Black) continuously compete to stabilize onto single categorical outcomes over time. In Study 4, simulations are used to demonstrate how the computational model can account for such dynamic competition effects. Together, the studies of Part I aim to provide evidence for the dynamic nature of person construal.

Part II focuses more on interactive aspects of person construal. First, in Studies 5–6, mouse-tracking and computational modeling are used to explore how the competition process underlying social categorization may be dynamically weighed in on not only by face processing, but also by voice processing in parallel. As such, face and voice processing may be thrown into interaction and may be able to influence one another across the categorization process. These studies examine the interaction between two bottom-up sensory modalities. Then, Studies 7–8 extend the work beyond two bottom-up modalities, and examine the interaction between bottom-up face processing and a top-down information source: activated stereotypes. Specifically, these studies use mouse-tracking and computational modeling to investigate how both face processing and activated stereotypes may simultaneously weigh in on categorization, thereby allowing stereotypes to exert a top-down influence on the categorization process. As such, a perceiver's prior stereotypic expectations would be free to alter basic categorizations. Thus, the studies of Part II explore the ability for multiple bottom-up and top-down information sources to potentially interact across the person construal process. They also aim to show that the proposed model can account for this real-time interactivity.

Together, across 8 studies, mouse-tracking and computational simulations are used to provide converging evidence for the dynamic and interactive nature of person construal.

PART I: DYNAMIC NATURE

There is an ongoing line of inquiry in social psychology as to when and how social categories come to be activated because, as discussed earlier, the mere activation of a social category shapes a number of cognitive, affective, and behavioral outcomes. Investigations into the perceptual determinants that lead to social category activation are thus quite crucial, acknowledging the consequences that follow this activation. Recent work in the emerging area of person construal research has examined the role of perceptual features in determining both overt person categorization and category activation itself. Overall, such work has tended to focus on examining how low-level processing of stimulus features maps onto higher-level stages of the person processing pipeline. One series of studies, for example, showed that perceivers can more rapidly and efficiently extract category-cueing information as compared with identity-triggering information, and that the extraction of category cues is uniquely impervious to stimulus manipulation and degradation. The special ease with which perceivers can decode category cues has thus been interpreted as an important determinant of the predominance of categorical thinking at all later stages of person processing (Cloutier, Mason, & Macrae, 2005). Important consequences of perceiving category cues is reaffirmed by findings showing that such cues can function orthogonal to category membership itself in automatic evaluations (Livingston & Brewer, 2002) and stereotypic attributions (Blair, Judd, Sadler, & Jenkins, 2002). Moreover, category cues can automatically trigger category activation itself (Macrae & Martin, 2007). Thus, the extraction of a mere perceptual cue is sufficient to activate a social category representation per se. Thus, this emerging work has provided an important start in opening up the process of social categorization, showing how perceptual cues and their bottom-up operations ultimately lead to the triggering of a social category.

More recently, some researchers have moved beyond examining the perceptual conditions that determine whether a social category will simply be activated or not activated. Instead, they have begun to examine how features can affect the strength of category representations. Using racial morphing, for example, Locke, Macrae, and Eaton (2005) showed that exemplar typicality can modulate the strength with which perceivers activate social category representations in graded fashion. A handful of similar findings have recently been reported in the social psychological literature elsewhere (Blair, Judd, & Fallman, 2004; Livingston & Brewer, 2002; Macrae, Mitchell, & Pendry, 2002; Maddox & Gray, 2002). Indeed, as Locke et al. have noted, such findings raise problems for the dominant "all or nothing" account of social categorization, in which a category is limited to two dichotomous states of either on or off, active or inactive.

The framework proposed here agrees with Locke et al's (2005) view that social categories have graded states. However, it emphasizes that social categorization involves more than just one single graded category representation. It argues, instead, that there are always multiple category representations partially active in parallel, and each has a graded strength. Further it is through a dynamic competition process that such partially-active categories can resolve into a stable categorical outcome. The main rationale behind this argument is that it is consistent with current understandings of the human brain. Specifically , it is consistent with the competitive dynamics of a neuronal population dynamically settling into one of multiple potential patterns (e.g., a "male" pattern vs. a "female" pattern) in response to perceptual input such as a face, as described earlier

(Grossberg, 1980; Rogers & McClelland, 2004; Spivey & Dale, 2004). In the present work, I aim to provide evidence for such a dynamic competition process, which is argued to underlie our ability to categorize other people.

Hand in Motion Reveals Mind in Motion

One way to provide evidence for such dynamic competition would be to examine perceivers' reaching hand movements as they make their way toward settling into one of multiple response options. Although motor responses are classically considered to be the end-result of a one-way route from perception → cognition → action (temporal cortex → "association cortex" → premotor areas), there is now a great deal of behavioral evidence showing that that cognition does not discretely collapse its processing onto movement execution; rather, movement is continually updated by cognitive processing over time (Goodale, Pelisson, & Prablanc, 1986). This is buttressed by neurophysiological evidence. Studies in humans, for example, suggest that the process of categorizing a face immediately shares its ongoing results with the motor cortex to continuously guide a hand-movement response over time (e.g., Freeman, Ambady, et al., 2011). Further, human reaching movements suggest that multiple motor plans are prepared in parallel, and that these cascade over time into visually-guided action (Song & Nakayama, 2006, 2008). Monkey studies show that the hand's position and velocity are tightly yoked to ongoing changes in the firing of population codes within motor cortex (Paninski, Fellows, Hatsopoulos, & Donoghue, 2004), and when a monkey must generate a hand movement based on a perceptual decision, these motor-cortical population codes are yoked to the evolution of the decision process. In short, the dynamics of action do not simply reside in the aftermath of cognition; rather, they are part and parcel with it, and show systematic covariation (Dale, Roche, Snyder, & McCall, 2008; Song & Nakayama, 2009; Spivey, 2007). Fortunately, then, continuous motor responses such as a hand movement may be informative as to what the perceptual system is doing across those fractions of a second between catching sight of another's face and recognizing that person's category memberships (see Freeman, Dale, & Farmer, 2011).

Study 1: Real-Time Dynamics of Sex Categorization

In the present study, a computer mouse-tracking technique is used to examine the real-time social categorization process. This technique measures participants' hand trajectories via mouse movement as they make their way into settling into one of two social category alternatives. Consider a task where two categories appear in the top-left and top-right corners of a computer screen, and participants are asked to move the mouse from the bottom-center to either response. If social categorization indeed results from a dynamic competition process, then when participants are presented with faces containing some degree of cues tied to the opposite category (e.g., male face with slight feminine features), their hand trajectories should be continuously attracted toward the opposite category response (e.g., female) before settling into the correct one (e.g., male). This would be evidence for simultaneously and partially active representations of both categories, which dynamically compete over time until stabilizing onto a single categorical outcome. To easily measure these hand movements, the streaming *x, y* coordinates of the computer mouse may be used (Freeman & Ambady, 2010). The present study first focuses on the categorization of sex using sex-typical faces and sex-atypical faces, which bear some perceptual overlap with the opposite category.

Method

Participants. Twenty-five undergraduates participated for partial course credit or monetary compensation.

Stimuli. To generate highly realistic faces and morph along sex, FaceGen Modeler was used. This software uses a 3D morphing algorithm based on anthropometric parameters of human population (Blanz & Vetter, 1999), in which a continuum from male to female (among other dimensions) can be manipulated while holding other extraneous cues constant. The typical condition included 10 male and 10 female face stimuli whose respective sex was generated at the anthropometric mean. The atypical condition comprised these same male and female face stimuli, except that their sex was generated at a level systematically closer to the anthropometric mean of the opposite sex. This resulted in a total of 40 target faces for the task. See Figure 2 for sample stimuli.

Procedure. To begin each trial, participants were asked to click on a "Start" button located at the bottom-center of the screen. Once clicking this, the target face appeared in its place. Targets were presented in a randomized order and were categorized by clicking either the "male" or "female" response option, located in the top-left and top-right corners of the screen (which category was on the left/right was randomized across participants). During this process, the streaming *x, y* coordinates of the mouse were recorded (sampling rate ≈ 70 Hz). To record, process, and analyze mouse movements, the freely available MouseTracker software package was used (http://mousetracker.jbfreeman.net). Details about the software and a discussion of analytic techniques for mouse trajectory data can be found in Freeman and Ambady (2010).

Results

Data preprocessing. All trajectories were rescaled into a standard coordinate space (top-left: "-1, 1.5"; bottom-right: "1, 0") and normalized into 100 time bins (101 time steps) using linear interpolation to permit averaging of their full length across multiple trials. For comparison, all trajectories will be remapped rightward. To obtain a by-trial index of the degree to which the mouse was attracted toward the opposite sex category (indexing the simultaneous activation of that category), Area Under the Curve (AUC) was compute: The area between the observed trajectory and an idealized response trajectory (a straight line between the trajectory's start and endpoints).

Spatial attraction. Mean trajectories were computed for typical and atypical male targets and for typical and atypical female targets. Plotted in Figure 2, trajectories for atypical targets revealed distinct curvature towards the opposite category. Trajectories for atypical male targets showed a continuous attraction towards "female" relative to trajectories for typical male targets, and trajectories for atypical female targets showed a continuous attraction towards "male" relative to trajectories for typical female targets.

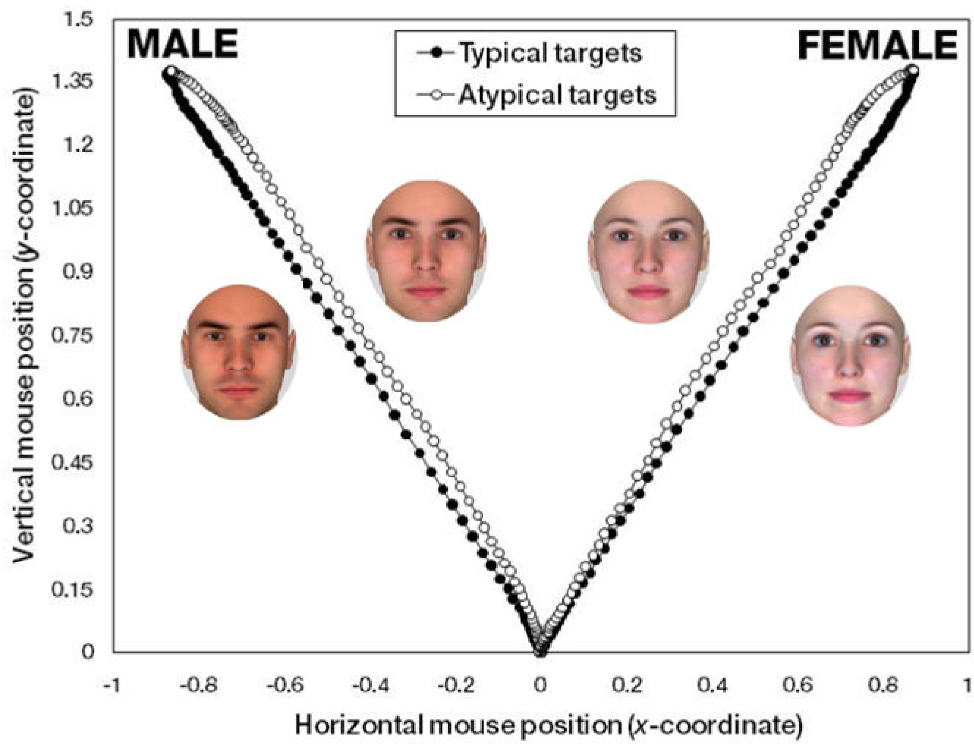

Figure 2. Mean mouse trajectories for sex-atypical faces exhibit a continuous attraction toward the opposite sex-category response (Study 1). Trajectories for male targets were remapped leftward, and trajectories for female targets were remapped rightward. Sample face stimuli are also depicted next to their respective mean trajectories.

To assess these attraction effects statistically, AUC values were submitted to a 2 (target sex) × 2 (typicality) repeated-measures ANOVA. A main effect for sex did not reach significance, $F(1, 24) = .07$, $p = .79$, nor did the interaction, $F(1, 24) < .01$, $p = 1.00$. More critically, however, a main effect for typicality was significant, $F(1, 24) = 15.82$, $p = .001$, such that trajectories for atypical targets were more attracted towards the opposite category relative to trajectories for typical targets. Specifically, trajectories for male targets were more attracted towards "female" when targets were atypical ($M = .70$, $SE = .04$) relative to typical ($M = .54$, $SE = .04$), $t(24) = 3.24$, $p < .01$, and trajectories for female targets were more attracted towards “male” when targets were atypical (*M* = .69, *SE* = .03) relative to typical (*M* = .53, *SE* = .03), $t(24) = 4.16$, $p < .001$.

Distributional analysis. It is possible that such continuous attraction effects in reality reflected the averaging together of some trials involving movement straight to the correct category, with other trials involving movement initially directed straight at the incorrect category, followed by a discrete reanalysis and corrective movement redirecting the trajectory straight to the correct category. If true, the continuous attraction effects would be the product of several discrete-like errors biasing the results. This spurious pattern can be detected by inspecting the distribution of trial-by-trial AUC values for bimodality (see Freeman & Ambady, 2010). A bimodal AUC distribution would suggest that the mean graded attraction effect spuriously reflects one population of trajectories with zero attraction and a separate population of trajectories with extreme attraction (i.e., discrete-like errors). A unimodal AUC distribution would suggest, instead, a single population of trajectories all exhibiting some degree of graded attraction, some high, some medium, and some low, reflecting dynamic competition as predicted.

AUC values for typical male targets and atypical male targets were together converted into *z*-scores within a participant and then pooled across participants. For both distributions, the bimodality coefficient *b* (SAS Institute, 1989) was computed, which has a standard cutoff value of $b = .555$. Values that exceed .555 are considered evidence to reject unimodality in favor of bimodality. Neither distribution had any indication of bimodality (atypical male, $b = .397$; typical male, $b = .425$). Concerns that the distribution for atypical male targets might host underlying bimodality can be alleviated by obtaining evidence that the shapes of the distribution for typical male targets and distribution for atypical male targets are statistically identical. To this end, AUC values were *z*-scored within each participant, separately for typical male and atypical male targets, and pooled across participants. The Kolmogorov-Smirnov test was used to evaluate any reliable departure between the respective shapes of these two distributions. This analysis confirmed that the distribution for typical male targets and distribution for atypical male targets were statistically indistinguishable, $D = .04$, $p = .99$, eliminating the possibility that the distribution for atypical male targets may be hosting latent bimodal features.

These distributional analyses were applied to the distributions for typical and atypical female targets as well. Neither distribution had any indication of bimodality (atypical female, $b = .420$; typical female, $b = .486$), and the Kolmogorov-Smirnov test confirmed that the respective shapes of the distribution for typical female targets and distribution for atypical female targets were statistically indistinguishable, $D = .06$, $p = .86$, eliminating the possibility that the distribution for atypical female targets was hosting latent bimodal features.

Discussion

The present results suggest that sex categorization is a dynamic process and never discretely “occurs” in any single moment. Rather, a target’s multiple perceptual cues trigger multiple partially-active categories, which continuously evolve into a stable categorical outcome over time. In Study 2, the evidence for a dynamic competition across sex categorization is extended to another social category: race. There are a number of differences between these category dimensions that have theoretical importance for understanding the real-time dynamics of person construal more broadly. One difference is the symmetric nature of sex and asymmetric nature of race. Specifically, perceivers are likely to routinely encounter similar amounts of men and women, and thus the functional significance of perceiving cues diagnostic of men and women should be similar. However, in most cases, perceivers are likely to routinely encounter disproportionate amounts of White and Black individuals. Because the White category represents a cultural "default," and White individuals are generally encountered more often than Black individuals in the U.S., the White category tends to be assumed if no other dimension-relevant information is provided (Smith & Zarate, 1992; Zarate & Smith, 1990). Thus, in a White-majority environment, the perceptual system may be biased towards Black-cueing features, which are ecologically salient relative to White-cueing features (Levin, 1996). Indeed, event-related potential evidence has shown that attention is preferentially directed to Black targets early in processing (Ito & Urland, 2003, 2005). Thus, it may be that more salient Black-cueing features and less salient White-cueing features drive the time-course of race categorization in potentially divergent ways. This was explored in Study 2.

Study 2: Real-Time Dynamics of Race Categorization

As in Study 1, participants were asked to categorize faces in a mouse-tracking paradigm. The typical condition included White and Black face stimuli whose level of race was generated at the anthropometric White and Black means. The atypical condition comprised these same White and Black stimuli, except their race was generated at a level closer to the other race. If perceiving race indeed results from a dynamic competition process, the mixture of White-specifying and Black-specifying cues on atypical targets would trigger partially-active race categories (White and Black) that simultaneously compete over time to stabilize onto ultimate categorizations. This would be evidenced by a continuous attraction in participants' hand movements toward the opposite race category before settling into eventual responses for race-atypical faces.

Method

Participants. Twenty-six undergraduates participated for partial course credit or monetary compensation.

Stimuli. As in Study 1, faces were generated using FaceGen Modeler, allowing race-related cues to be manipulated while holding other extraneous cues constant. The algorithm does not make assumptions about what differs between White and Black faces; rather, by averaging across many faces, parameters that emerge as reliably different between the races become incorporated into the morphing (Blanz & Vetter, 1999). Ten unique White faces (5 male) were generated at the White mean and 10 unique Black faces (5 male) were generated at the Black mean, together composing the typical condition. The 10 typical White faces were then morphed 25% toward the Black mean, and the 10 typical Black faces were then morphed 25% toward the White mean, together composing the atypical condition.

Procedure. The procedures were similar to those of Study 1, except here with "White" and "Black" category responses. One methodological limitation of Study 1 was that participants were not encouraged to initiate movement early. Thus, several trials' movements were initiated relatively late in the categorization process, rendering the measure off-line with respect to a large portion of the process. To remedy this, in the present study and all subsequent mouse-tracking studies, participants were encouraged to begin initiating movement early. If initiation time exceeded 400 ms, a message appeared after participants made their response, encouraging them to start moving earlier on future trials even if not fully certain about their response (see Freeman & Ambady, 2010).

Results

Data preprocessing. The preprocessing was identical to that of Study 1.

Spatial attraction. The mean trajectories for typical and atypical White targets appear in Figure 3A and those for typical and atypical Black targets appear in Figure 3B. The mean trajectory for atypical White targets showed a continuous attraction toward the "Black" response, and the mean trajectory for atypical Black targets showed a continuous attraction toward the "White" response. AUC values were submitted to a repeated-measures ANOVA using factors of typicality and target race. The main effect of typicality was significant, $F(1, 25) = 22.87$, $p < .0001$. Trajectories for atypical White targets (M = 0.39, SE = 0.05) curved more toward the "Black" response relative to those for typical White targets (M = 0.28, SE = 0.04), $t(25) = 4.34$, $p < .001$, and trajectories for atypical Black targets (M = 0.43, SE = 0.06) curved more toward the "White" response relative to those for typical Black targets (M = 0.33, SE = 0.04), $t(25) = 2.74$, $p = .01$. Neither the main effect of race [$F(1, 25) = 0.63$, $p = 0.43$] nor the interaction [$F(1, 25) = .02$, $p = .89$] reached significance. This is evidence that, on the way to arriving at categorizations of atypical targets, both race categories (White and Black) were simultaneously and partially active in continuous competition across construal.[1]

[1] Overall, the mean trajectories of the present study and in subsequent studies are considerably more curved than those of Study 1. This is likely due to the present study's methodological improvement of encouraging participants to initiate movement earlier in the categorization process.

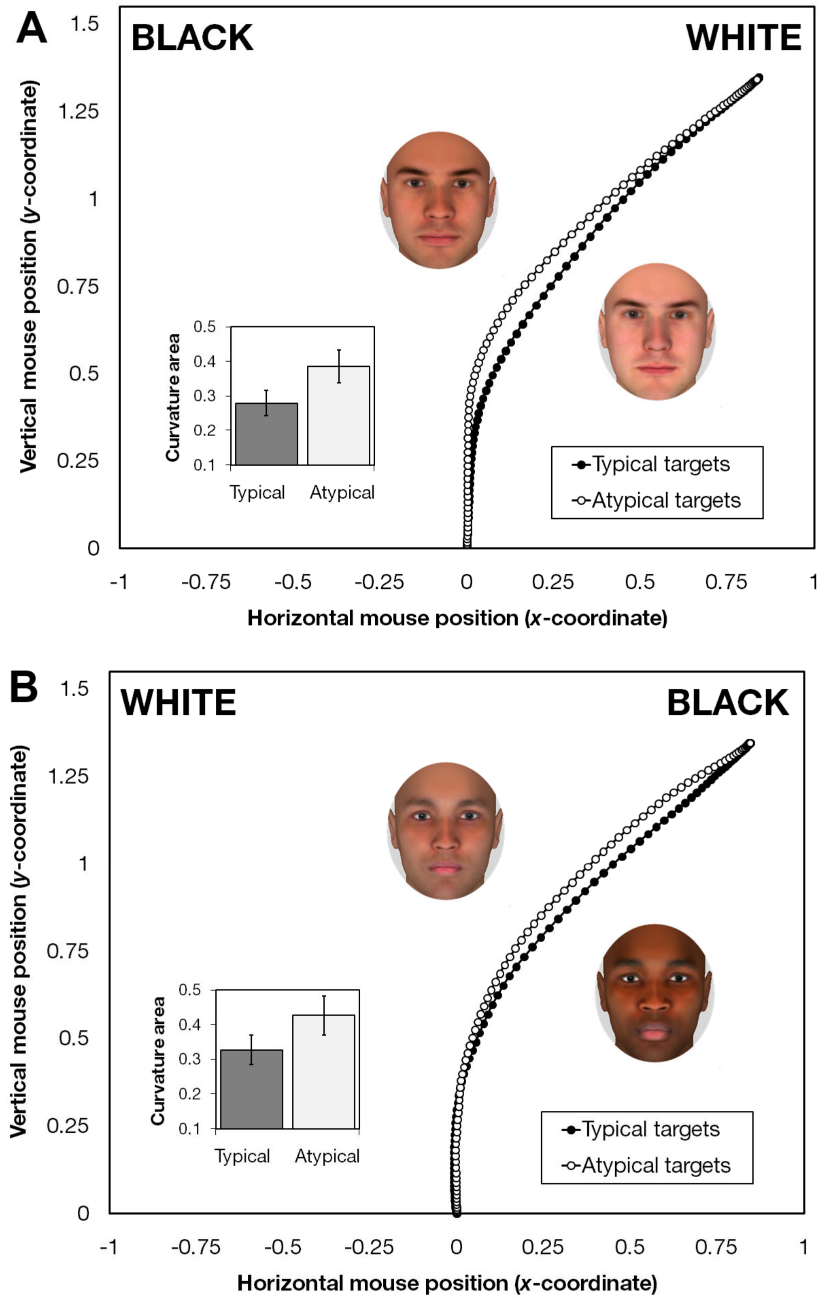

Figure 3. Mean mouse trajectories for race-atypical faces exhibit a continuous attraction toward the opposite race-category response (Study 2). All trajectories are remapped rightward. Sample face stimuli are also depicted next to their respective mean trajectories. Bar graphs shows trajectories' curvature toward the opposite race category, separately for race-typical and race-atypical trials (error bars denote standard error of the mean). (A) Trajectories for White targets. (B) Trajectories for Black targets.

Distributional analysis. As in Study 1, it is possible that the continuous-attraction effects reported here were spuriously produced by a bimodal population where some trajectories exhibit zero attraction and others exhibit extreme attraction (e.g., discrete-like errors). The distribution of AUC values for atypical White targets (b = .444) and for typical White targets (b = .378), however, were within the b < .555 bimodality-free zone. The Kolmogorov-Smirnov test also verified that the shapes of the distribution for typical White targets and for atypical White targets, once standardized, were statistically indistinguishable (D = .06, p = .86). Similarly, neither the distribution for typical or atypical Black targets showed evidence of bimodality: atypical Black (b = .476), and typical Black (b = .372), and the Kolmogorov-Smirnov test confirmed that their shapes were statistically indistinguishable (D = .08, p = .38). These analyses confirm that the continuous-attraction effects were not spuriously produced by a combination of discrete-like movements.

Time-course analysis. Given the a priori hypothesis of a timing difference in the processing of White-specifying and Black-specifying cues, the time course of the attraction effects was examined. To that end, the Euclidean proximity of the mouse position to the opposite race category was calculated at each time step. Greater proximity to the opposite-race category in the atypical condition at any given time step would indicate that category's partial and simultaneous activation as categorization unfolded over time. Separately for White and Black targets, difference scores were computed at each time step by subtracting the proximity of the typical condition from that of the atypical condition. The proximity differences are plotted in Figure 4, indexing across time the degree to which the hand traveled closer to the opposite race category. A casual inspection of this figure indicates that while the attraction effects for both atypical White and Black targets were manifest continuously across the course of categorization, atypical White targets elicited an attraction effect more consistent throughout the time-course whereas atypical Black targets elicited an effect that was relatively small early on, but then rapidly grew in the later portion of the categorization process.

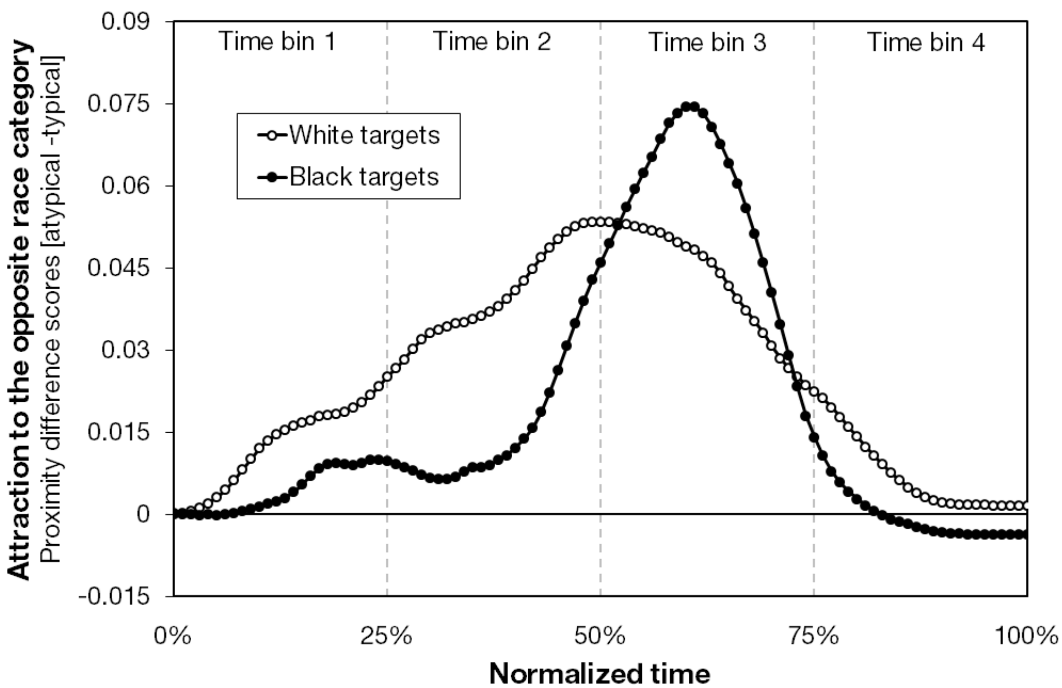

Figure 4. Separately for the White and Black faces of Study 2, difference scores between the atypical and typical conditions [atypical – typical] in proportional Euclidean proximity (1 – distance/max(distance)) to the opposite race category are plotted as a function of normalized time. This indexes across the time the degree to which the hand traveled closer to the opposite race category for atypical targets relative to typical targets. The atypical White targets induced a relatively consistent continuous-attraction effect, with gradual increase and decrease of partial activation. The atypical Black targets, however, induced an effect with weak attraction at the beginning, which then rapidly rose in the later portion of the trial (and settled). These two different patterns of temporal dynamics were statistically distinguishable.

To more rigorously assess these two patterns of temporal dynamics, four time bins were computed for proximity distances (time-steps: 1-25, 26-50, 51-75, 76-101). These were then submitted to a 2 (typicality) × 2 (race) × 4 (time bin) repeated-measures ANOVA. Only effects involving the typicality factor are reported. Beyond the main effect of typicality [$F(1, 25) = 36.62$, $p < .0001$], this analysis revealed a marginally significant three-way interaction [$F(3, 75) = 2.50$, $p = .07$] and a significant typicality × time bin interaction [$F(3, 75) = 18.06$, $p < .0001$]. These interaction effects arose because, for White targets, trajectories in the atypical condition showed strong attraction to the "Black" response (relative to those in the typical condition) consistently throughout time bins 1 [$t(25) = 2.71$, $p = .01$], 2 [$t(25) = 4.44$, $p < .001$], and 3 [$t(25) = 4.56$, $p = .0001$], whereas, for Black targets, trajectories in the atypical condition showed relatively weak attraction at the beginning, with non-significant attraction at time-bin 1 and marginally significant attraction at time bin 2 [$t(25) = 1.83$, $p = .08$], but then a substantial rise in attraction later in time-bin 3 [$t(25) = 4.36$, $p < .001$]. These different patterns of temporal dynamics are best illustrated in Figure 4. Thus, although the atypical White and Black faces induced equivalent amounts of continuous attraction toward the opposite race category [as there was no typicality × race interaction: $F(1, 25) = 0.80$, $p = .38$], these effects were temporally distributed in different ways. This indicates that the partial and simultaneous activation of the opposite race category during categorization of atypical White versus atypical Black faces fluctuated across the course of construal in divergent ways.

Discussion

When categorizing atypical White and Black faces, hand movements exhibited a continuous attraction toward the opposite race-category response. This is evidence that atypical faces triggered simultaneously and partially active race-category representations that dynamically competed across the course of categorization. This therefore extends the results of Study 1, showing that dynamic competition is a generalized process underlying perceivers' ability to categorize along many social category dimensions, such as sex or race.

Importantly, White-cueing and Black-cueing features had different temporal signatures during this competition process. Specifically, White-cueing and Black-cueing features on atypical targets induced an equivalent amount of partial and simultaneous activation of their respective race category, but differed in how this activation fluctuated across the course of construal. Whereas partial Black cues led the hand to travel closer to the "Black" response relatively consistently throughout the categorization process, partial White cues did not bear a substantial effect until approximately 40% of categorization had elapsed, at which point a partial activation of the White category rapidly grew and settled (see Figure 4). These different patterns of temporal dynamics are fitting given the asymmetric nature of White and Black categories, such that Black cues are more likely to capture attention and spontaneously activate category representations than White cues (Smith & Zarate, 1992). These results suggest that, when participants encountered an atypical White face, partial Black cues began biasing the competition early—as Black cues are highly salient—and this competition persisted across the course of categorization until participants finally settled onto a "White" response. However, when encountering an atypical Black target, the majority of Black cues were dominant during the early portion of the competition—as Black cues are highly salient—while the White cues were not well represented. However, after a slight delay now-better-represented partial White cues exerted a strong bias that then yielded to the predominance of accruing perceptual evidence for the Black category. Thus, although prior work has investigated similar asymmetries on the outcomes of social categorization (e.g., Zarate & Smith, 1990), well before these outcomes are even realized, perceptual cues belonging to different sorts of social categories may exert different dynamic patterns of influence over the perceptual–cognitive processing that creates those outcomes.

Thus far, category-atypical targets have been exploited to examine a dynamic competition process underlying social categorization. When targets bear cues that partially overlap with an alternate category (e.g., slight Black-specifying cues on a White face), a competing representation of that alternate category is activated in parallel. But is this competition process truly sensitive to the natural perceptual gradations underlying social category dimensions? For example, an interesting difference between the dimensions of sex and race is the inherently fuzzy nature of race relative to the substantially less fuzzy nature of sex. While it is generally rare to encounter faces that are truly sex-ambiguous—an unlikely situation usually evoking anxiety or a few laughs (e.g., Saturday Night Live's androgynous "Pat" skits)—person perceivers often encounter faces that do not fit squarely into any race category at all. Interactions with mixed-race individuals, for instance, involve the perception of faces that tend to contain major physiognomic overlap between multiple traditionally distinguished race categories (e.g., White and Black). Prior research indicates that, even in instances of extreme racial ambiguity (e.g., mixed-race faces), perceivers readily extinguish this ambiguity by slotting faces into traditionally distinguished race categories (Pauker et al., 2009), particularly during rapid categorization (Peery & Bodenhausen, 2008). The following study examines how this resolution of racial ambiguity is accomplished in real time. Because perceptions of race can be fuzzy and can involve different levels of ambiguity, this afforded the opportunity to examine how graded increases in the ambiguity of a social category may have corresponding graded effects on the real-time evolution of social categorical responses. As such, it permits a test of whether the competition process underlying social categorization is sensitive to the full perceptual continuum underlying category dimensions such as race.

Study 3: Sensitivity to the Perceptual Continuum: Racial Ambiguity

As in Studies 1 and 2, participants engaged in a mouse-tracking task in which they were asked to decide whether a face was White or Black. Rather than using computer-generated stimuli, however, real faces that varied along a continuum of racial ambiguity were used. It was hypothesized that as ambiguity increases, mouse trajectories would show an increasingly stronger attraction to whichever race category is not ultimately selected, indicating competition between alternate race categories settling over time onto a single categorical outcome.

Method

Participants. Thirty-two undergraduates participated for partial course credit or monetary compensation.

Stimuli. Photos were obtained (40 male and 40 female images) of self-identifying White, Black, and mixed-race (half Black, half White) individuals who were participating in a separate study on face memory. These individuals varied widely on how racially typical or ambiguous they appeared. These faces were pre-tested ($N = 16$) on racial appearance using a Likert scale ("Does this person look African American or Caucasian?"), ranging from 1 – "Very African American" to 8 – "Very Caucasian." These ratings were subtracted by a constant of 4.5, converted into absolute values, and rescaled to vary between –0.5 (typicality) and 0.5 (ambiguity).

Procedure. Analogous procedures to Study 2 were used.

Results

Trajectories were preprocessed using the same procedures as the previous studies. Trajectories were remapped (inverted along the x-axis) so that whichever race category was ultimately selected was located at the top-right. Regardless of whether participants perceived a face to be White or Black, of interest was whether the amount of racial ambiguity affected mouse trajectories en route to indicating that perception. To determine this, AUC values were regressed onto pre-test ambiguity scores using a generalizing estimating equations (GEE) regression analysis. This allows for the incorporation of trial-by-trial data while accounting for the intracorrelations in a repeated measures design, thereby permitting more statistically efficient parameter estimates (Zeger & Liang, 1986). In this analysis and all subsequent GEE regression analyses, unstandardized regression coefficients are reported. As predicted, as racial ambiguity increased (i.e., as a face depicted the unselected race category more strongly), mouse trajectories' attraction toward the opposite race category linearly increased, $B = 0.44$, $SE = 0.07$, $p < .0001$.

Discussion

As targets became more racially ambiguous, the competition between race categories correspondingly increased. This was evidenced by increases in racial ambiguity leading to increases in trajectories' attraction toward the unselected race category. This extends the results of Studies 1 and 2, showing that social category competition increases and decreases with the amount of ambiguity challenging the perceptual system. Thus, the competition is sensitive to the perceptual continuum underlying a category dimension, such as the natural diversity in others' race-specifying facial cues. The present study also demonstrates that this category competition process generalizes beyond computer-generated faces to more ecologically valid, real faces. In the following study, the processing underlying this competition process is more rigorously explored at the computational level.

Study 4: Computational Simulation of Social Categorization Dynamics

The present study aims to demonstrate how the proposed model naturally captures the dynamic competition process underlying social categorization explored in Studies 1–3. It focuses specifically on the effects of Study 1 with sex categorization, although the results would be generalizable to race categorization as well. Figure 5 depicts a new instantiation of the general model (see Figure 1) developed for this purpose. Solid-line connections with arrows are excitatory (positive weight) and dashed-line connections with dots are inhibitory (negative weight). Connection weights appear in Appendix B. This instantiation of the model is intended to capture the experience of a perceiver categorizing either sex or race for a particular task.[2]

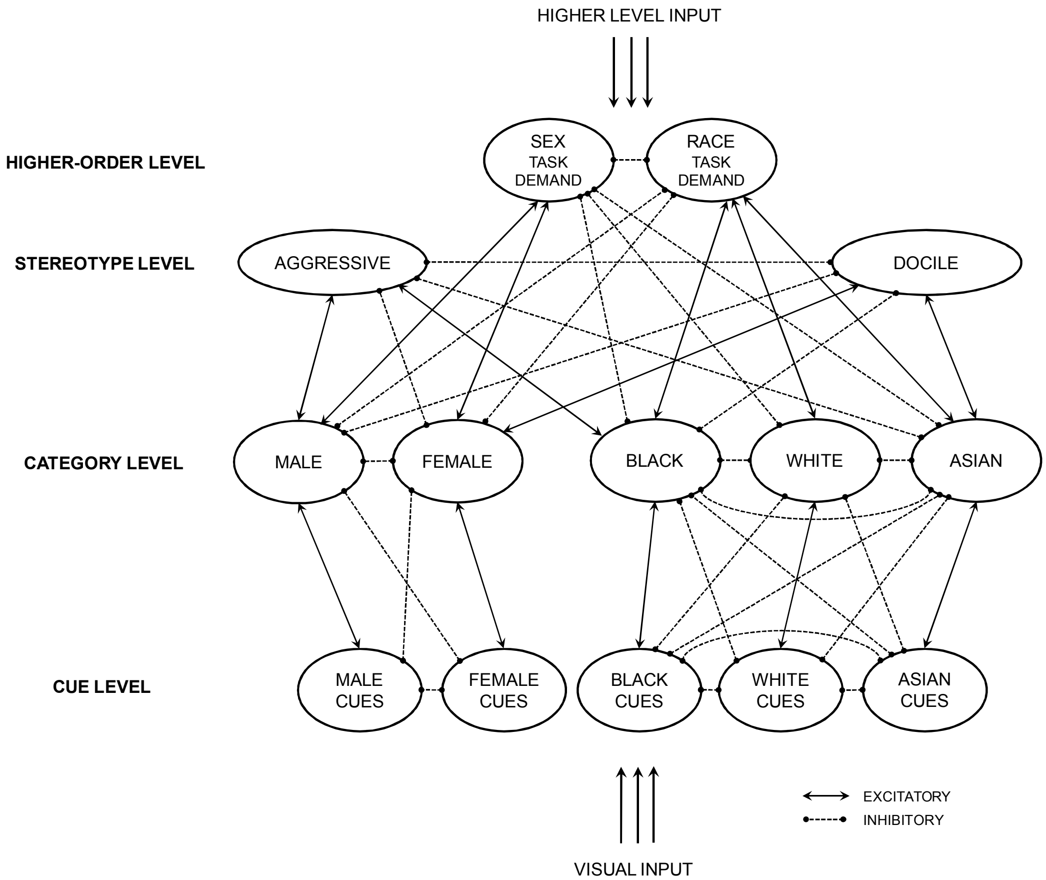

Figure 5. An instantiation of the model presented in Study 4 to account for the results of Study 1.

---

[2] It should be noted that representations in the model are interactive rather than independent. The activation of one representation influences other representations in the system, as one node's activation influences the activation of all other nodes. Consider the presentation of a male face. Stimulation of the MALE CUES node will facilitate the MALE category, which will inhibit the FEMALE category, which will inhibit the stereotype, DOCILE, in turn inhibiting the category, ASIAN, and so on and so forth. These influences gradually taper off as all nodes in the system come to settle into a steady state. Before the system stabilizes, however, representations are in continuous interaction over time. They are dynamically and probabilistically reconstructed in every new instance, and not static. They develop in continuous interaction with other activations across the system, both influenced by these activations and a source of influence over them. For example, there is no stand-still, discrete symbol-like representation of the "male" category. Rather, the system will gravitate toward a stable state involving strong activation of MALE, but this state is not a discrete symbol identically activated every time the system encounters a male target. Thus, the system may frequently visit a similar stable state involving strong activation of MALE every time it encounters a male target. But this is a dynamically reconstructed state of activation that could only approximate an idealized, linguistically identifiable representation of the "male" category (Spivey & Dale, 2006). See Appendix A.

Method

Two simulations are conducted, both of which simulate a sex categorization task. One involves the presentation of a sex-typical face the other involves the presentation of a sex-atypical face. Because this is a sex categorization task and the demands of the task compel attention to sex, higher-level input would directly activate the SEX TASK DEMAND node. Accordingly, higher-level input into the SEX TASK DEMAND node was set at .9 and higher-level input into the RACE TASK DEMAND node at .1.[3] This simulates the task context of sex categorization, where perceivers would be focusing on targets' sex over their race. This thus facilitates activation of MALE and FEMALE category nodes, and inhibits activation of BLACK, WHITE, and ASIAN nodes. At the same time, nodes in the cue level receive direct input from visual processing of the face. To simulate the presentation of a sex-typical White male face, visual input into the MALE CUES node was set at .95 and visual input into the FEMALE CUES node at .05. Thus, this face is inherently 95% masculine and 5% feminine. Because the face is White, visual input into the WHITE CUES node was set at .95 and visual input into the BLACK CUES and ASIAN CUES nodes at .025 each. The simulation was run 100 times, each time for 150 iterations, and the average activation level of each category node was examined over time, appearing in Figure 6.

[3] In all simulations, connection weights and input values were set according to intuitions regarding stimulus and task features. It may be possible in future work to derive these values empirically. However, I am confident given previous studies that the parameters are in accord with participant judgments and task features in these contexts, and parameters that best reflect these intuitions were chosen. In this sense, the current simulations serve as existence proofs for the kind of dynamic interactive processing that may take place during construal, though it is acknowledged that future work may advance these simulations by deriving network parameters empirically.

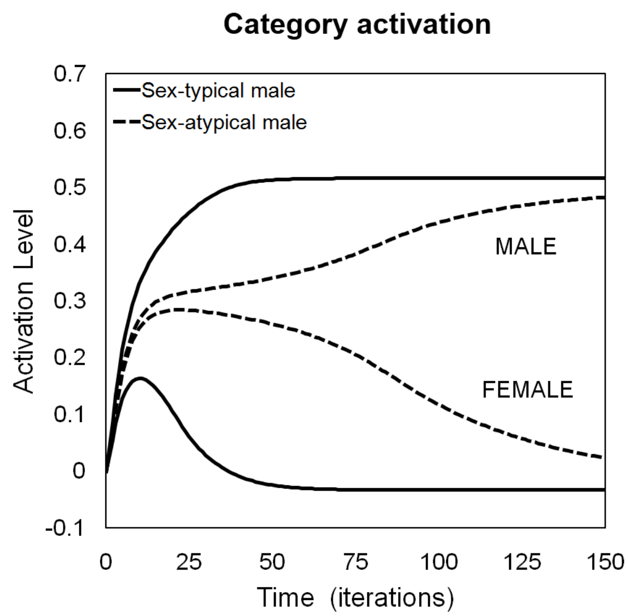

Figure 6. The activation level of the MALE and FEMALE category nodes as a function of time (iterations) following the presentation of a sex-typical male face (solid lines) and sex-atypical male face (dashed lines) in Study 4.

Results

The presentation of a sex-typical White male face sets a process into motion, in which visual processing of the face directly activates cue nodes. Cue nodes inconsistent with one another, such as the MALE CUES and FEMALE CUES nodes, compete for the visual input. The activation of cue nodes, in turn, places excitatory and inhibitory pressures on category nodes (see Figure 5). In this case, the highly activated MALE CUES node places strong excitatory pressure on the MALE category node and inhibitory pressure on the FEMALE category node. The highly activated WHITE CUES node places strong excitatory pressure on the WHITE category node and inhibitory pressure on the BLACK and ASIAN category nodes. At the same time, the higher-order SEX TASK DEMAND node places excitatory pressures on the MALE and FEMALE category nodes and inhibitory pressures on the BLACK, WHITE, and ASIAN category nodes. These simultaneous pressures cause the activation levels of some category nodes to be pushed above their resting levels, whereas others are inhibited and pushed below their resting levels. Excitatory pressure from both the MALE CUES node and the higher-order SEX TASK DEMAND node leads the MALE category node to rise above its resting level. Positive feedback is then produced between these nodes, which causes the MALE category node to rapidly gain activation until gradually settling into a stable state. Because a small amount of feminine features were presented to the network (the FEMALE CUES node was initialized with .05 visual input), the FEMALE category node also becomes slightly active for a very brief moment early on, and then succumbs to strong inhibition from the MALE CUES node and the MALE category node, resulting in it being pushed below its resting level. Excitatory pressure from the WHITE CUES node leads the WHITE category node to rise above its resting level, but the WHITE category node is also inhibited by the SEX TASK DEMAND node. This leads the WHITE category node to gain a meager amount of activation until eventually settling into a stable state (and thus MALE is more strongly active than WHITE). Finally, inhibitory pressures from the WHITE CUES node and the SEX TASK DEMAND node lead the BLACK and ASIAN category nodes to be rapidly pushed below their resting levels. These dynamics are apparent in Figure 6.

Note how each category node gradually works over time to settle into a stable state, such that its activation reaches some asymptotic level and tapers off. This stable state would correspond with the fully confident categorization of the target as male. However, before that 100% confident categorization is achieved, bear in mind that partial, tentative evidence for that categorization actually accumulates gradually over time (the dynamics of the MALE category activation).

Consider, on the other hand, how the person construal system settles into a stable state when presented with a sex-atypical White male face in a sex categorization task. To simulate this, visual input into the MALE CUES node was set at .55, input into the FEMALE CUES node at.45, input into the WHITE CUES node at .95, and input into the BLACK CUES and ASIAN CUES nodes at .025 each. As done previously, higher-level input into the SEX TASK DEMAND node was set at .9 and input into the RACE TASK DEMAND node at .1. This simulates attention on sex induced by the task context of sex categorization. The simulation was run 100 times, and the averaged activation level of each category node over 150 iterations appears in Figure 6.

The activated MALE CUES node begins exciting the MALE category and inhibiting the FEMALE category, while the FEMALE CUES node begins exciting the FEMALE category and inhibiting the MALE category (see Figure 5). The MALE CUES and FEMALE CUES nodes also begin inhibiting one another as well. The highly activated WHITE CUES node excites the WHITE category and inhibits the BLACK and ASIAN categories. At the same time, the higher-order SEX TASK DEMAND node excites the MALE and FEMALE categories and inhibits the BLACK, WHITE, and ASIAN categories. The excitatory pressure from both the MALE CUES node and the higher-order SEX TASK DEMAND node leads the MALE category to rise above its resting level. The excitatory pressure from the FEMALE CUES node and the SEX TASK DEMAND node also leads the FEMALE category to rise above its resting level. Pressures from the cue nodes and higher-order nodes cause the WHITE category to gain a meager amount of activation and the BLACK and ASIAN categories to be rapidly pushed below their resting levels. With the MALE and FEMALE categories now simultaneously activated, they begin competing with one another through mutual inhibition. In this case, the system is simultaneously attracted to be in two different states: one state involving ~100% MALE/~0% FEMALE and another involving ~0% MALE/~100% FEMALE. These are highly stable (leading the system to be attracted to them), whereas an earlier state such as ~55% MALE/~45% FEMALE is highly unstable. Over time, the mutual inhibition between competing MALE and FEMALE categories, in addition to feedback with the cue nodes, leads the FEMALE category to gradually decay while the MALE category gradually rises in activation until a stable state is achieved. This results in the MALE category winning the competition, while the FEMALE category dies off and is cleared from the processing landscape. Thus, simultaneously and partially active sex categories dynamically compete over time to settle onto a single categorical outcome (in this case, a male categorization).

Discussion

The present results show how the proposed computational model replicates effects of temporally dynamic competition, as found with human perceivers. In Study 1, it was found that participants' mouse trajectories exhibited a continuous attraction toward the "female" response before settling into the "male" response when categorizing a sex-atypical male face. Thus, a female category representation was simultaneously and partially active across construal, which led the mouse to partially curve toward the "female" response before clicking the "male" response. This is precisely what is reflected in Figure 6, which shows that the FEMALE category node was partially active, simultaneously with the MALE category's activation, until the system settled into a stable state involving strong MALE category activation and the FEMALE category below resting level. Thus, the mouse-tracking results and computational model converge on the finding that an ongoing competition process underlies our ability to categorize other people. Simultaneously and partially active category representations continuously evolve into single categorical outcomes over time, highlighting the dynamic nature of person construal.

PART II: INTERACTIVE NATURE

One of the most remarkable features of perceiving other people, as compared with everyday objects, is that perceptions of people are frequently grounded in multiple sensory modalities and embedded in a rich set of contexts. The human voice, for example, always contextualizes the human face, continuously over time. The body's motion, for instance, contextualizes the perception of its shape. A growing number of studies have shown that these prevalent contextual and cross-modal cues powerfully constrain the perception of the social percepts under the focus of perceivers' attention. The studies of Part I showed that social categorizations continuously evolve over time, and they demonstrated a dynamic competition process underlying these categorizations that is weighed in on by facial cues. Here, the studies of Part II examine how this competition process is weighed in on not only by one bottom-up sensory modality (e.g., facial cues), but also in parallel by other extraneous information sources, such as other bottom-up sensory modalities (e.g., vocal cues) or top-down sources (e.g., stereotypes). By driving the competition process in parallel, these multiple information sources may interact with one another over time. As such, according to the proposed framework, extraneous factors well beyond the face, such as vocal cues or activated stereotypes, could potentially exert a real-time influence over the process of categorizing a face.

Study 5: Face–Voice Interaction in Social Categorization

The majority of person construal research has focused on visual features, such as facial cues, with little attention paid to auditory features, such as vocal cues. Recently, however, vocal cues were shown to give rise to categorical judgments and stereotypic inferences of others, and these inferences are sensitive to within-category variation (Ko, Judd, & Blair, 2006), as is seen with facial cues (Blair et al., 2002). Thus, like the face, the voice plays an important role in social categorization. How the face and voice are combined into a coherent social percept when categorizing others, however, remains poorly understood. Previous work provided clear evidence that perceivers do combine facial and vocal input. For instance, when a face appears sad but is accompanied by a voice that sounds happy, perceivers consistently report seeing the face as more happy than it really is. This remains true even when participants are instructed to disregard the voice (de Gelder & Vroomen, 2000). Furthermore, congruency between facial and vocal features tends to make the perception of identity or emotion more accurate and efficient (for review, Campanella & Belin, 2007). Very few studies, however have examined face–voice interaction in social categorization.

The present framework proposes that the biases of another person's sensory information (e.g., facial and vocal cues) converge the moment they become available in the input to weigh in on multiple partially-active social category representations. These parallel representations then settle onto categorical outcomes across a process of continuous competition. Ongoing voice-processing results, therefore, should integrate with ongoing face-processing results over time. If true, processing of sex-specifying vocal cues should exert a temporally dynamic influence on face processing across the construal process. Indeed, such continuous cross-modal interactivity would be consistent with evidence for recurrent interactions between the visual and auditory cortices and top-down feedback from higher-order multimodal cortices (e.g., Ghazanfar, Chandrasekaran, & Logothetis, 2008; Kreifelts, Ethofer, Grodd, Erb, & Wildgruber, 2007).

In the present study, mouse-tracking is used to track in real time how voice processing weighs in on resolving facial ambiguities, specifically in the context of sex categorization. By looking at instances in a mouse-tracking paradigm where there is a conflict between a category triggered by facial cues and an opposing category triggered by vocal cues—and how these might compete with one another over time—the process by which information from the face and voice interact and are integrated can be measured. The predicted continuous interaction and gradual integration of the face and voice would be evidenced by the hand's continuous attraction toward the opposite sex-category response in instances where, although the face is categorized as the correct sex, vocal cues partly suggest the opposite sex.

Method

Participants. Forty-one individuals participated for partial course credit.

Stimuli. Face stimuli were generated using FaceGen Modeler to appear somewhat sex-ambiguous. Ten male faces were generated at 60%-Male/40%-Female and 10 female faces at 60%-Female/40%-Male. For voice stimuli, 10 male and 10 female speech samples of American dialect and neutral tone were obtained from the International Dialects of English Archive (http://web.ku.edu/~idea). Clips of 2000 ms were extracted from each sample, the content of which was selected to be naturalistic for a first-impression encounter (e.g., "My family's origins are pretty interesting."). Praat software (http://www.fon.hum.uva.nl/praat) was used to morph each male voice into a sex-typical (masculine) version (formant shift ratio:1/1.1) and a sex-atypical (feminine) version (formant shift ratio:1.1), and to morph each female voice into a sex-typical (feminine) version (formant shift ratio:1.1) and a sex-atypical (masculine) version (formant shift ratio:1/1.1), consistent with prior work (e.g., Groen et al., 2008). Median pitch was not manipulated because its alteration tends to sound computer-like and artificially synthesized.

Procedure. Each face was presented twice in the experiment, accompanied by a same-sex voice (once sex-typical and once sex-atypical). Voice stimuli were randomly paired with face stimuli (without replacement). Participants were instructed to categorize the face's sex (and only use the voice if it could help resolve the face's sex, as correct responses were based on the face). The mouse-tracking task was identical to that of Study 1, except that the additional voice stimulus began playing once the face stimulus appeared.

Results

Preprocessing was identical to that of Study 1. To index trajectories' attraction toward the opposite response, maximum deviation (MD) was computed: the largest perpendicular deviation from an idealized straight line between the trajectory's start and endpoints. This measure is highly correlated with the AUC measure, and they show negligible differences (Freeman & Ambady, 2010). The mean trajectories are plotted in Figure 7.

Participants were more likely to misinterpret the face as the opposite sex when faces were accompanied by sex-atypical (M = 11.5%, SE = 1.0%) relative to sex-typical (M = 5.1%, SE = 0.8%) voices, $t(40) = 5.38$, $p < .0001$, a finding often cited as evidence of face-voice integration (e.g., Hietanen, Leppänen, Illi, & Surakka, 2004). To examine the temporal dynamics of this integration, trials that were correctly categorized were examined.

As indicated by MD, before participants settled into their correct categorizations, the hand was continuously attracted to the opposite sex category while categorizing faces accompanied by sex-atypical voices ($M = 0.33$, $SE = 0.03$) relative to sex-typical voices ($M = 0.26$, $SE = 0.03$), $t(40) = 3.81$, $p < .001$.

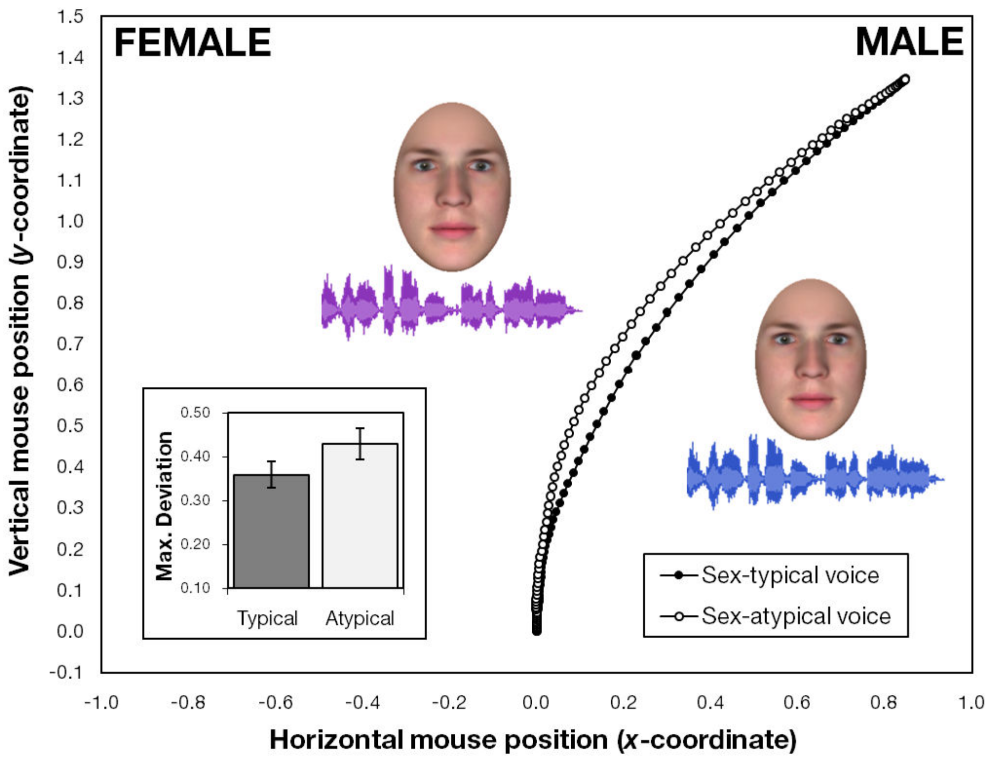

Figure 7. Mean mouse trajectories of Study 5 (aggregated across male and female targets). Trajectories for all targets were remapped rightward, with the opposite sex category on the left and the sex category consistent with the face's sex on the right. A sample male face stimulus is displayed. A voice stimulus typical for the face's sex (masculine) is shown on the right (audio waveform depicted in blue), next to the mean trajectory for sex-typical trials. Its atypical (feminine) counterpart is shown on the left, next to the mean trajectory for sex-atypical trials (audio waveform depicted in purple). During an actual trial, a single face was centered at the bottom of the screen while the voice stimulus played. The bar graph shows trajectories' maximum deviation toward the opposite sex category, separately for sex-typical and sex-atypical trials (error bars denote standard error of the mean).

Distributional analysis indicated that the MD distribution for sex-atypical trials was within the bimodality-free zone ($b < .555$; SAS Institute, 1989), $b = .407$, as was the distribution for sex-typical trials, $b = .428$. Furthermore, the Kolmogorov-Smirnov test confirmed that the shapes of these two distributions were statistically indistinguishable ($D = .02$, $p = .99$), ruling out the possibility of latent bimodality. This ensures that the continuous-attraction effect was not the product of a subpopulation of discrete-like errors.

Discussion

While participants' hands were moving en route to making a sex categorization of the face, the simultaneous processing of a sex-atypical voice led the hand to travel closer to the opposite sex-category, continuously across the course of construal. This suggests that processing of the voice continuously interacted with processing of the face. At each moment during the categorization of sex-atypical pairs, mouse movements were neither in a discrete pursuit straight to the male response, nor in a discrete pursuit straight to the female response. Rather, as seen with the conspicuous curving of the trajectory toward the opposite sex category in Figure 7, at each moment the location of the mouse was in a weighted combination of one pursuit consistent with face processing (e.g., male) and a simultaneous pursuit consistent with voice processing (e.g., female), while the mouse progressively stabilized onto ultimate interpretations of the face. Thus, mouse trajectories reflected an ongoing interaction between category information from the face and voice as they gradually integrated over time. In the following study, simulations are used to show the proposed model converges on this effect.

Study 6: Computational Account of Face–Voice Interaction

The previous study found that, when perceivers correctly categorized the face's sex, auditory processing of sex-specifying vocal cues exerted a temporally dynamic influence on the face-based categorization. Specifically, the simultaneous processing of sex-specifying facial and vocal cues triggered partially-active representations of both sex categories (male and female) that simultaneously competed over time to settle into ultimate categorizations. To account for this continuous face–voice interactivity in sex categorization, another instantiation (Figure 8) of the general model (Figure 1) was developed. Connection weights are provided in Appendix C. In this network, the cue level receives input from both visual processing and auditory processing and has separate nodes for facial and vocal cues.

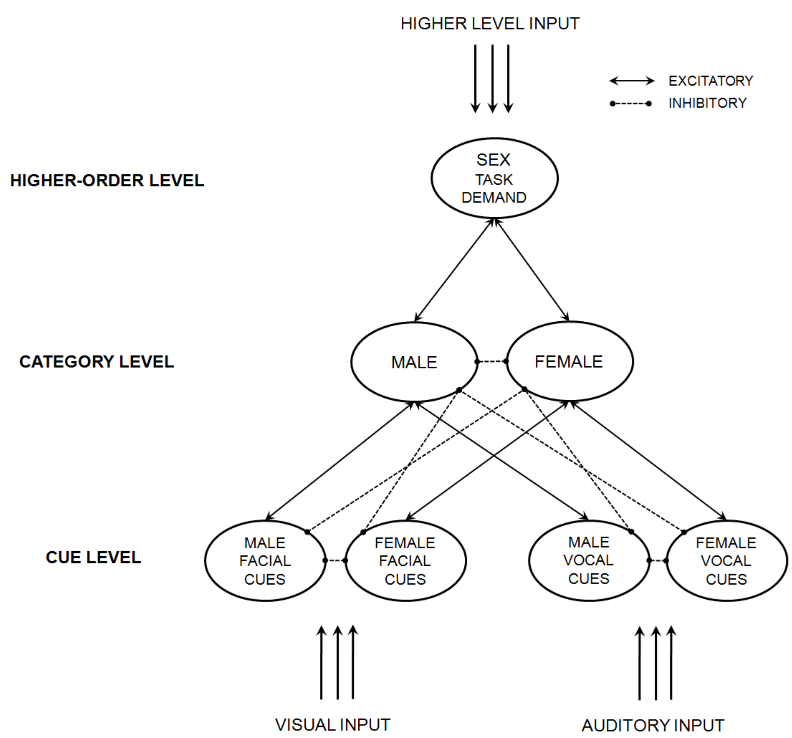

Figure 8. An instantiation of the model used in Study 6 to account for the results of Study 5.

Method

To simulate the presentation of a slightly ambiguous male face, visual input was set at .55 for MALE FACIAL CUES and at .45 for FEMALE FACIAL CUES. To simulate the simultaneous presentation of a sex-typical voice, auditory input was set at .95 for MALE VOCAL CUES and at .05 for FEMALE VOCAL CUES (see Footnote 3). Higher-level input was set at .9 for the SEX TASK DEMAND node to simulate a strong attentional state on targets' sex required by the task. The simulation was run 100 times, each time over 75 iterations, and the averaged level of activation of the category nodes was plotted over time (Figure 9).

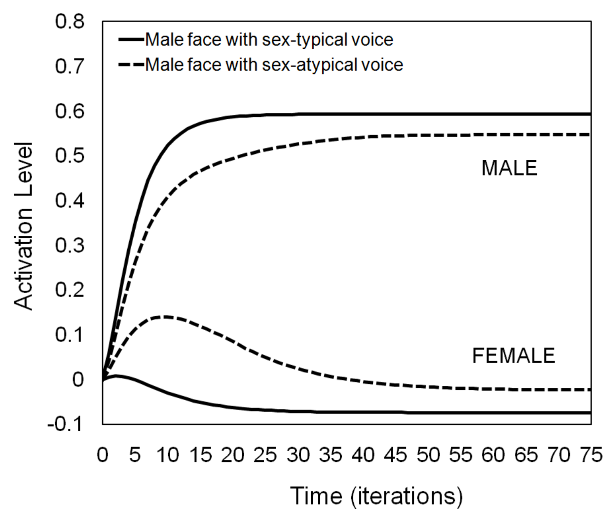

Figure 9. The activation level of the MALE and FEMALE category nodes as a function of time (iterations) following the presentation of the same male face with a sex-typical male voice (solid lines) or sex-atypical male voice (dashed lines).

Results and Discussion

The slightly ambiguous activation of facial cues nodes fed forward activation onto the MALE and FEMALE category nodes. Simultaneously, the activation of the vocal cues nodes also fed forward activation onto the category nodes. In doing so, the simultaneous processing of vocal cues placed an immediate constraint on the face-triggered activation of sex categories. This permitted ongoing updates from voice processing to immediately interact with ongoing updates from face processing, continuously over time. The strong activation of MALE VOCAL CUES was therefore immediately brought to bear on resolving the category competition induced by ambiguous facial input. Strong excitation of the MALE category and inhibition of the FEMALE category, due primarily to the unambiguous vocal cues nodes, led the system to rapidly converge on a stable state involving strong activation of MALE category, with FEMALE category pushed below resting level.

When the voice is more atypical, however, the face-triggered category competition did not resolve so quickly. To simulate the presentation of a slightly ambiguous male face coupled with a sex-atypical voice, the input activation was kept the same except, this time, input into the MALE VOCAL CUES node was set at .6 and input into the FEMALE VOCAL CUES at .4. The simulation was run 100 times, each time over 75 iterations, and the averaged level of activation of the category nodes was plotted over time (Figure 9). The slightly ambiguous activation of facial cues nodes and slightly ambiguous activation of vocal cues nodes simultaneously fed forward activation onto the MALE and FEMALE category nodes. This induced a strong competition between the category nodes. Although the system eventually resolved the competition by arriving at a stable state involving strong activation of MALE and weak activation of FEMALE (i.e., a male categorization), the FEMALE category was partially-active in parallel strongly throughout the process. This partial activation of the FEMALE category was considerably stronger when the voice was sex-atypical rather than sex-typical (see Figure 9).

This pattern of results is consistent with the findings of Study 5. The stronger partial activation of the FEMALE category, which continuously competes with the MALE category, is clearly seen in the human mouse-tracking data of Figure 7. When sex-categorizing a male face, the simultaneous processing of a sex-atypical voice led participants' hands to be continuously attracted toward the "female" response before ultimately arriving at the "male" response. This reflects a stronger partially-active representation of the female category (induced by voice processing) that simultaneously competed over time with the male category during face-based categorization. Thus, in sex categorization, the model predicts (as experimental data show) that voice processing interacts with face processing by simultaneously weighing in on the dynamic competition inherent to the categorization process. As such, the simultaneous processing of facial and vocal cues place parallel constraints on sex categorization (which are dynamically satisfied over time), permitting the ongoing processing of vocal cues to continuously interact with the ongoing processing of facial cues. In short, the present results demonstrate how the proposed model naturally accounts for continuous face–voice interactivity in person construal. The following study examines how face categorization may be shaped by other kinds of extraneous information sources emanating from the top-down rather than the bottom-up, such as stereotype activations.

Study 7: Influence of Top-Down Stereotypes on Race Categorization

In the present study, mouse-tracking is used to assess how stereotype activations may exert a top-down influence on the ongoing race categorization process. To trigger stereotypes during the categorization process, contextual cues were exploited. In many instances, stereotypes may be triggered by the contextual cues that often surround a face in the real world, such as attire. Once activated via contextual cues, stereotypes could potentially alter the processing of a face's category memberships.

For example, businesspeople are stereotypically associated as high-status, whereas janitors are associated as low-status. However, White individuals, too, are associated as high-status, whereas Black individuals are associated as low-status (Devine, 1989). Due to this overlap in the stereotypes associated with both race and occupation categories, contextual cues to occupation (e.g., business attire) might potentially activate stereotypes (e.g., high-status) that then exert top-down pressure on the race categorization process, swaying it toward the associated category (e.g., White). For example, business attire could activate high-status stereotypes that then gradually push the race-category competition—primarily being driven by visual processing of the face—more toward the White category. Conversely, janitor attire could activate low-status stereotypes that then gradually push the race-category competition more toward the Black category. According to this approach, therefore, race categorization could be driven by both the bottom-up processing of facial features, and top-down stereotypes activated by contextual cues, which mutually constrain one another before a stable categorization is achieved.

One implication of this tight exchange between bottom-up and top-down forces theorized here is that, as one force gets weaker, the other force is given sway to exert an increasingly stronger influence on categorization. Thus, as race-specifying facial cues become increasingly ambiguous, the bottom-up ambiguity opens the door wider and wider to stereotypes' top-down influences. This is important because perceivers in the real-world regularly encounter racially ambiguous faces (e.g., multiracial individuals). Despite their ambiguity, however, perceivers rapidly resolve such faces into monoracial categories, such as White or Black (Peery & Bodenhausen, 2008). The present study thus explored whether racial ambiguity might affect the degree to which contextual status cues are able to shift race categorization. Such a finding would provide a compelling illustration of the interactive nature of person construal, showing how bottom-up and top-down forces work in conjunction with one another to drive the construal process.

It is hypothesized that high-status business attire would tend to elicit White categorizations and low-status janitor attire tend to elicit Black categorizations. Further, these effects would be more pronounced as a face's race increases in ambiguity. Even when a status cue would not influence an ultimate categorization response, however, it was predicted that it would still lead perceivers to partially, simultaneously activate the other race category with which it is associated. Such a partial parallel activation of the other race category—due to status cues tied to that category—would be evidenced by a partial attraction in participants' hand movements toward the other category response (e.g., Black) before clicking their final response (e.g., White). Further, this attraction effect would grow stronger as a face's race increases in ambiguity. Thus, this study aims to test how top-down stereotypes interact with the processing of bottom-up facial cues to shape race perception.

Method

Participants. Twenty-two undergraduates participated for partial course credit or monetary compensation. One participant did not follow instructions correctly, leaving 21 participants for analysis.

Stimuli. Face stimuli were comprised of 16 computer-generated face identities (8 male) that were morphed along a 13-point race continuum, from White (morph −6) to

Black (morph +6), using FaceGen Modeler. Using various images of clothing obtained from public domain websites, each face was affixed to a high-status (business) and low-status (janitor) attire (see Figure 10A). Half of the 16 face identities (each containing 13 levels of race) were affixed to a high-status cue, whereas the other half were affixed to a low-status cue (which identities were affixed to which cue was counterbalanced across participants).

Procedure. Participants categorized faces as White or Black as quickly and accurately as possible in a mouse-tracking paradigm. The procedure was identical to that of Study 2.

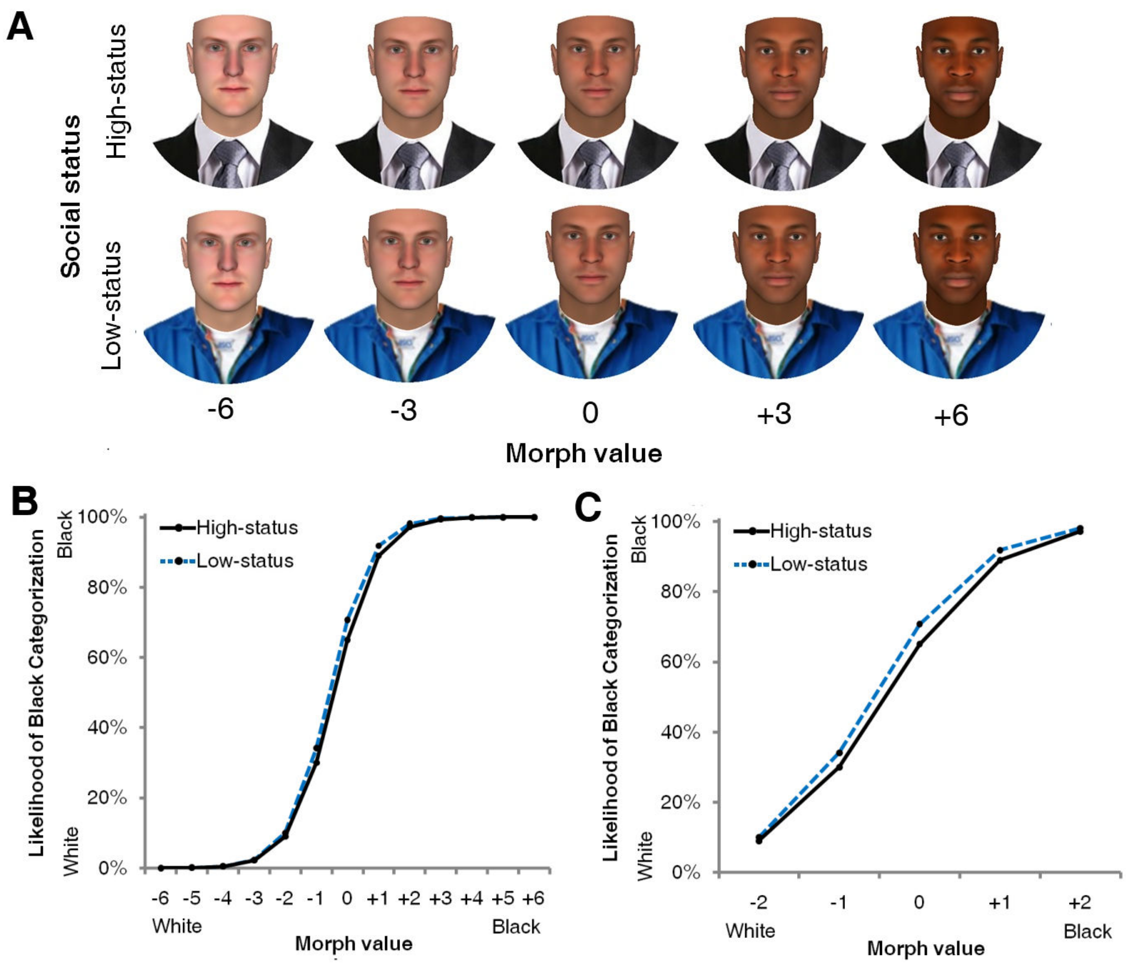

Figure 10. (A) Sample stimuli of Study 7. A high-status or low-status cue was affixed to 13-point morph continua, where race was varied from White (−6) to Black (+6). (B) The likelihood of Black categorization is plotted as a function of morph values, separately for high-status and low-status faces. Note the canonical sigmoidal shape of the curves, consistent with the categorical perception of race (Levin & Angelone, 2002). Also note that the strongest influences of the status cue are in the middle of the continuum (most clearly shown in panel C). (C) The same plot as in panel B, except here zooming in on the middle of the morph continuum, where race is most ambiguous.

Results

Data preprocessing. Preprocessing was identical to the previous mouse-tracking studies. For comparison, all trajectories were remapped rightward, such that the selected response was at the top-right and the unselected response at the top-left.

Categorization responses. First, perceived race (0 = White, 1 = Black) was regressed onto morph values (−0.5 = most prototypically White [morph −6], 0.5 = most prototypically Black [morph +6]), status cue (−0.5 = high-status, 0.5 = low-status), and the interaction (using logistic GEE regression). Expectedly, as morph values rose from White to Black, the likelihood of Black categorization increased, $B = 18.05$, $p < .0001$, $z = 15.22$, confirming the morphing manipulation. Status cues, however, also influenced categorization. A low-status cue raised the likelihood of Black categorization relative to a high-status cue, which raised the likelihood of White categorization, $B = 0.26$, $p < .05$, $z = 2.42$ (Figures 10B and 10C). The interaction was not significant, $B = 0.85$, $p = .41$, $z = 0.82$.

To directly examine whether racial ambiguity may have moderated the influence of status cues on categorization, an index of racial ambiguity was generated by converting morph values into absolute values, multiplying by −1, and centering around 0: −0.5 = most prototypical (morph ±6) to 0.5 = most ambiguous (morph 0). Perceived race was regressed onto racial ambiguity, status cue, and the interaction (using logistic regression). Increases in racial ambiguity overall led to increases in the likelihood of Black categorization, $B = 0.64$, $p = .0001$, $z = 3.88$. This was due to an overall bias of categorizing racially ambiguous faces as Black rather than White, as the most ambiguous face (morph 0) had a likelihood of Black categorization in the 60−70% range, rather than 50%. This is consistent with prior work on hypodescent (the tendency to assign individuals of mixed heritage to the social group of lowest status) in race categorization (e.g., Peery & Bodenhausen, 2008). As in the previous analysis, status cues influenced categorization as well, with a low-status cue raising the likelihood of Black categorization and vice-versa for a high-status cue, $B = 0.06$, $p < .05$, $z = 2.29$. More importantly, a significant interaction indicated that the influences of status cues were exacerbated as racial ambiguity increased, $B = 0.17$, $p = .05$, $z = 1.92$ (Figure 11A). Thus, contextual status cues shaped race perception, and ambiguity moderated their ability to exert an influence.

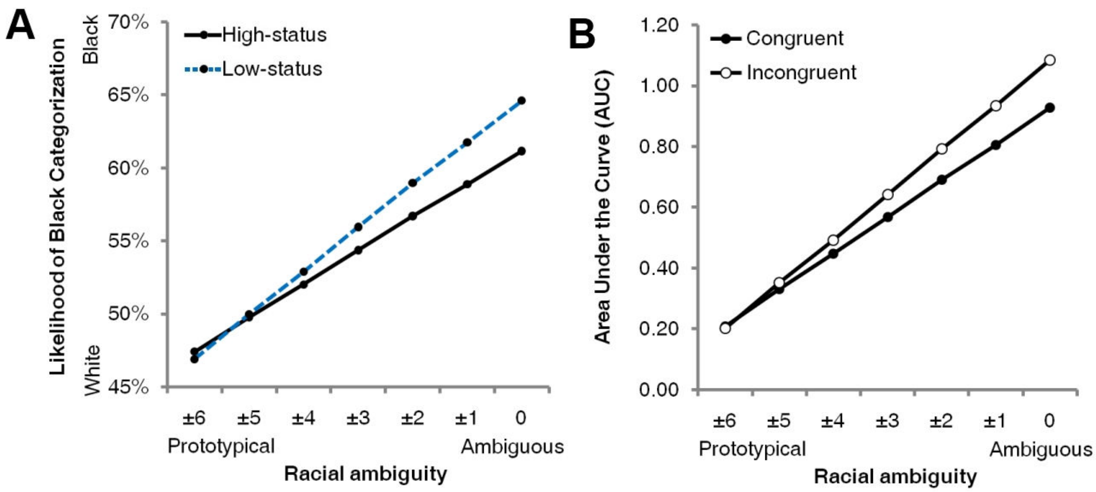

Figure 11. Racial ambiguity's moderation of status cues' influence on race categorization (Study 7). (A) The likelihood of Black categorization is plotted as a function of racial ambiguity, separately for faces surrounded by high-status versus low-status attire. (B) The degree of attraction toward the opposite race-category response (indexed by AUC) is plotted as a function of racial ambiguity, separately for trials where the categorization response was stereotypically congruent versus incongruent with the status cue.

Although the above analysis shows that status cues affected a considerable number of categorizations, there were also many categorizations that remained unaffected by status cues. The account proposed here, however, argues that such seemingly unaffected categorizations are, in fact, still subtly influenced by those cues. This is because the processing of status cues would always partially weigh in on the categorization process, as described above. Thus, even when a face with low-status attire is categorized as White or a face with high-status attire is categorized as Black, the status cue would still trigger the partial parallel activation of the other race category with which it is associated, thereby temporarily altering race perception. This was addressed using the mouse-tracking data. For mouse-trajectory analyses, trials were coded as congruent or incongruent based on whether the categorization response was stereotypically congruent vs. incongruent with the status cue. Thus, trials where a face with high-status attire was categorized as White or a face with low-status attire was categorized as Black were coded as congruent; trials where a face with low-status attire was categorized as White or a face with high-status attire was categorized as Black were coded as incongruent.

Spatial attraction. Trajectories' AUC values were regressed onto racial ambiguity, congruency (−0.5 = congruent, 0.5 = incongruent), and the interaction (using normal GEE regression). As expected given the results of Study 3, there was a significant effect of racial ambiguity. Increases in ambiguity overall led to increases in the attraction toward the opposite side of the screen [$B = 0.80$, $p < .0001$, $z = 9.03$], suggesting that perceivers were tentatively considering the other race category. More importantly, there was a significant effect of congruency. When categorization responses were incongruent (i.e., not influenced by the status cue), hand trajectories nevertheless showed an attraction toward the other race category stereotypically associated with the status cue, relative to hand trajectories for congruent responses, $B = 0.07$, $p < .01$, $z = 2.73$ (Figure 12). Moreover, a significant interaction indicated that the hand's attraction toward the other race category, due to the presence of a status cue tied to that category, became increasingly strong as racial ambiguity increased, $B = 0.16$, $p < .05$, $z = 2.07$ (Figure 11B).

Thus, en route to settling into the White response for a face with low-status attire, the hand showed an attraction to select the Black response; and en route to settling into the Black response for a face with high-status attire, the hand showed an attraction to select the White response. Further, this attraction effect grew stronger as racial ambiguity increased.

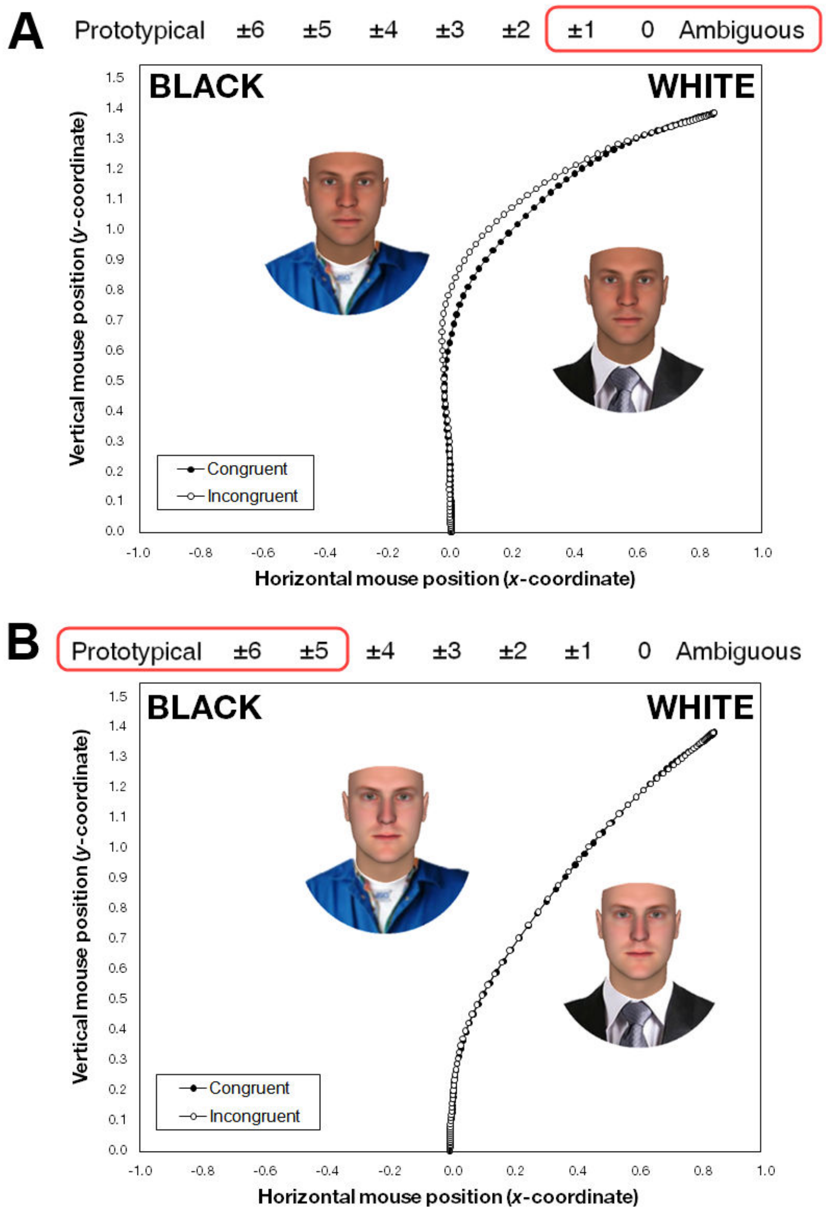

Figure 12. Mean computer mouse trajectories of Study 7. Trajectories for all targets were remapped rightward, with the opposite race category on the left and the selected race category on the right. A sample face stimulus, surrounded by a status cue associated with the selected race category, is shown on the right, next to the mean trajectory for congruent trials (when faces with high-status cues were categorized as White or faces with low-status cues were categorized as Black). On the left is shown that same face stimulus, but with a status cue associated with the opposite race category, next to the mean trajectory for incongruent trials. Panel A shows trajectories averaged across trials for the most ambiguous faces (morphs 0 and ±1), along with a sample ambiguous face stimulus. Panel B shows trajectories averaged across trials for the least ambiguous faces (morphs ±5 and ±6), along with a sample unambiguously White face stimulus.

Discussion

Low-status cues presented with a face increased the likelihood of Black categorization, and high-status cues presented with a face increased the likelihood of White categorization. Further, such influences grew stronger as a face's race became more ambiguous, as the bottom-up ambiguity opened the door to top-down pressures from stereotypes triggered by contextual status cues. Often these influences affected categorization wholesale and drove ultimate responses. In cases where they did not, however, they nevertheless influenced categorization. Even in many cases in which faces with low-status attire were categorized as White or faces with high-status attire were categorized as Black, the processing of a status cue still triggered the partial parallel activation of the other race category with which it was stereotypically associated. This was evidenced by participants' hands temporarily gravitating toward the other race-category response before arriving at their ultimate decision. When status cues do not shape an ultimate categorization, therefore, they nevertheless exert a subtle influence by activating the other, associated race category. These results suggest that extraneous top-down information sources, such as activated stereotypes, may interact with bottom-up face processing across the construal process. Further, the degree of bottom-up ambiguity moderates the ability for top-down information to enter into the categorization process. In the following study, such top-down effects are explored with the computational model.

Study 8: Computational Account of Top-Down Interactivity

To account for the pattern of categorization responses and mouse-tracking data of the previous study, a new instantiation (Figure 13) of the general model (Figure 1) was developed.

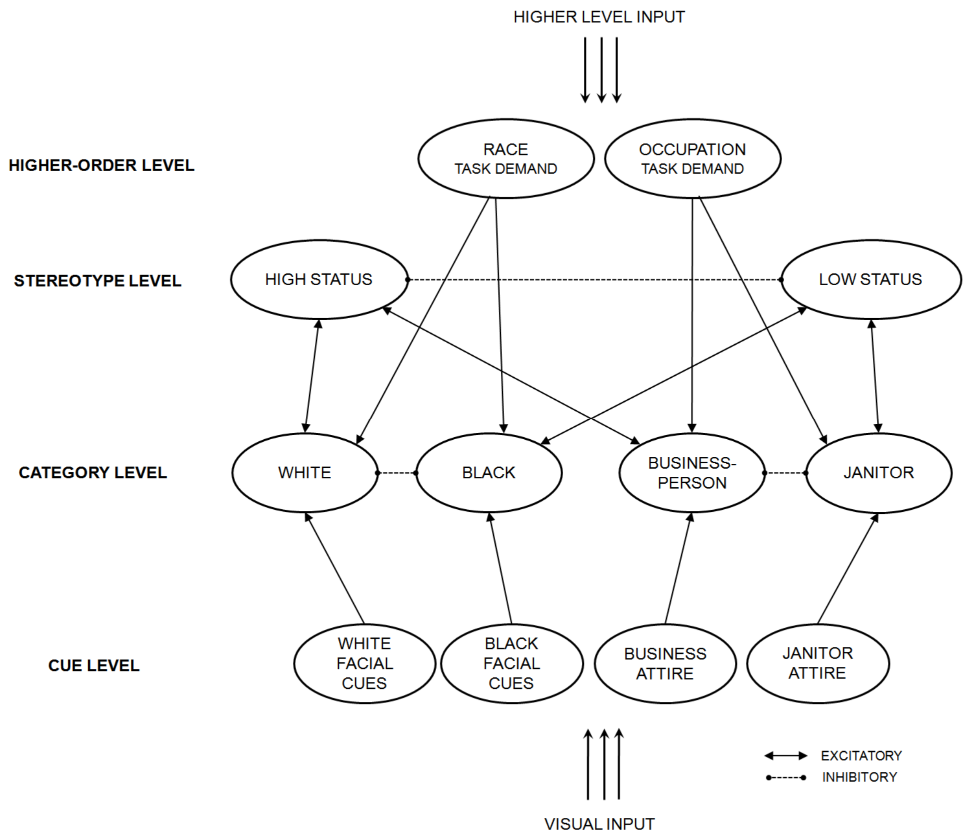

Figure 13. An instantiation of the model used in Study 8 to account for the results of Study 7.

Method

In this instantiation, all excitatory connections have a weight of 0.2 and inhibitory connections a weight of –0.1. Because non-normative categories (e.g., Black) and the stereotypes tied to those categories (e.g., low-status) tend to be more readily activated (Smith & Zarate, 1992), and because race categorization tends to be more strongly swayed by the Black category rather than the White category (e.g., Peery & Bodenhausen, 2008), the bidirectional excitatory BLACK–HIGH-STATUS connection was given a slightly stronger weight of 0.203, thereby capturing this asymmetry. The level of noise also slightly differed from previous instantiations, with σ = 0.007 (see Appendix A). The present instantiation of the model is also a simpler variant than the previous instantiations, in that there are considerably less between-node connections. A more complex instantiation was also used, but because the simpler variant was able to account for the empirical data well, it was adopted for parsimony.

A total of 26 simulations were conducted: 13 morph values × 2 status cues. In each simulation, input into the RACE TASK DEMAND node was set at .9 and into the OCCUPATION TASK DEMAND node at .1, simulating the task demand that requires attention on race rather than occupation. For the high-status condition (where targets had business attire), input into the BUSINESS ATTIRE NODE was set at 1 and into the JANITOR ATTIRE node at 0, and vice-versa for the low-status condition (where targets had janitor attire). Based on a face's morph value, input into the WHITE FACIAL CUES node was set at [1 − (morph + 6)/12] and input into the BLACK FACIAL CUES node at [(morph + 6)/12]. For example, for the most prototypically White face (morph −6), the WHITE FACIAL CUES node was initialized with 1 and BLACK FACIAL CUES node with 0, and vice-versa for the most prototypically Black face (morph +6). For a slightly less White face (morph −5), the WHITE FACIAL CUES node was initialized with 0.92 and the BLACK FACIAL CUES node with 0.08. For the most racially ambiguous face (morph 0), both nodes were initialized with 0.5. See Footnote 3. Each of the simulations was run 100 times. After 200 iterations, the race-category node with the highest activation was selected as the network's categorization response.

Results and Discussion

When the network was presented with the task demand of race categorization and the face stimuli of Study 7, its categorization responses closely mirrored that of human perceivers (R^2 = 0.99, root mean-square-error = 0.03). As shown in Figure 14A, low-status cues made a Black categorization more probable, whereas high-status cues made a White categorization more probable. Further, these influences of status cue grew stronger as racial ambiguity increased. For those categorization responses that were not affected by the status cue (incongruent responses), the processing of the status cue nevertheless triggered the partial parallel activation of the other race category with which it was associated. This is reflected in Figure 14B, showing the maximum activation level of the selected and unselected race-category nodes. When a status cue stereotypically tied to the other race category was present (i.e., incongruent trials), that other, unselected category was partially active in parallel. Further, this partial activation of the unselected race category became increasingly strong as racial ambiguity increased. Such partial activation accounts for why participants' hand movements were simultaneously attracted toward the other race-category response (Figure 12), and why that attraction grew increasingly strong as racial ambiguity increased (Figure 11B).

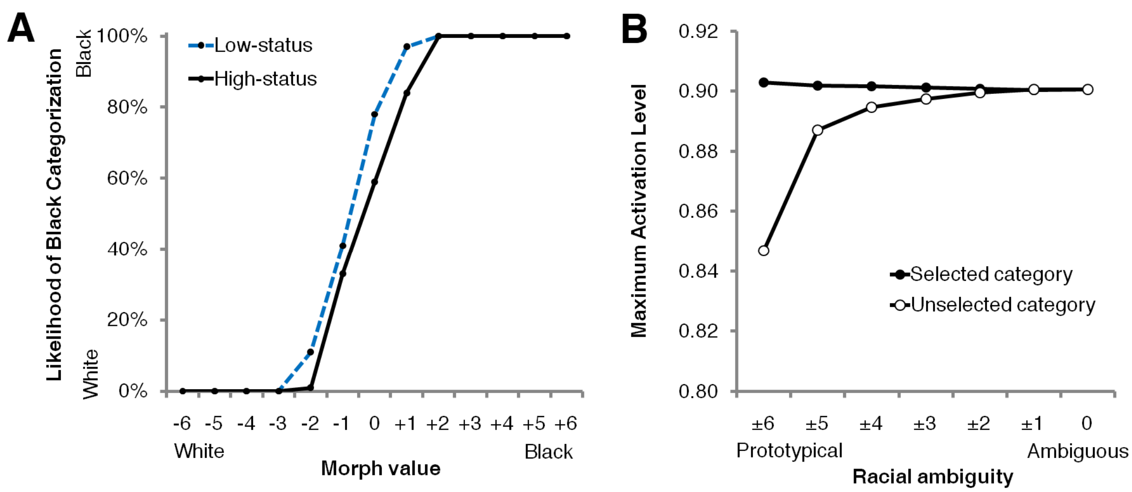

Figure 14. Results of the computational simulations in Study 8. (A) The proportion of times that the network settled into a Black category (the BLACK category node won the competition) after presented with faces that varied along a continuum from White (−6) to Black (+6). Note the close correspondence with the data from human perceivers (Figure 10B). (B) Maximum activation level of the winning WHITE or BLACK category node (selected category) and the losing WHITE or BLACK category node (snselected category) is plotted for incongruent trials (when the BLACK category won the competition for a face with a high-status cue, or the WHITE category won the competition for a face with a low-status cue), across a continuum of racial ambiguity. Note that the unselected category is partially activated as well, and that as racial ambiguity increases the maximum activation level of the unselected category (the opposite race category) increases as well. This accounts for why participants' mouse movements increase in attraction toward that unselected race category, especially as race becomes more ambiguous (Figure 11B).

Consider, for example, the presentation of a relatively unambiguous White face with janitor attire. A process is set into motion where visual input of the face activates cue nodes and higher-level input of the task demand activates higher-order nodes (see Figure 13). Activation of the RACE TASK DEMAND node starts exciting the WHITE and BLACK categories and inhibiting the BUSINESSPERSON and JANITOR categories, leading the race categories to become partially active for the task. Strong activation of the WHITE FACIAL CUES node places strong excitatory pressure on the WHITE category. With both race categories simultaneously active, they begin competing with one another through mutual inhibition to stabilize onto just one. As the competition unfolds, the WHITE category excites the HIGH STATUS stereotype and the BLACK category excites the LOW STATUS stereotype. Now with the conflicting LOW STATUS and HIGH STATUS stereotypes simultaneously active as well, they too begin competing with one another through mutual inhibition to stabilize onto just one. Meanwhile, activation of the JANITOR ATTIRE node excites the JANITOR category and inhibits the BUSINESSPERSON category (but the JANITOR category only gains a meager amount of activation because it is inhibited by the RACE TASK DEMAND node). Ongoing activation of the JANITOR category then excites the LOW STATUS stereotype. At this point, the stereotype nodes are being continually fed activation by both race and occupation categories. Because activation in the network is mutually interactive, however, while the competition is still resolving itself the stereotype nodes also feed activation back to the category nodes. This leads the JANITOR category's excitation of the LOW STATUS stereotype, in turn, to place excitatory pressure on the BLACK category and help it win against the WHITE category.

In some cases, such pressures would be strong enough to make the BLACK category more likely to win the competition, driving ultimate categorization responses. In other cases, such pressures would not be strong enough to drive responses and would only lead to a stronger partial parallel activation of the BLACK category (until it gradually decays, succumbing to the WHITE category). Moreover, these top-down pressures from stereotypes would be given increasingly more room to shape race-category activation as a face's race increases in ambiguity. The lack of bottom-up bias toward either the WHITE or BLACK category would open the door wider and wider to top-down influences, as the race-category competition is increasingly swayed by feedback from stereotype nodes. By activating stereotype nodes, contextual attire cues readily influence race categorization. As such, the network incorporated status cues to categorize a face's race, particularly when race was ambiguous.

Converging with the results of Study 7, the present results show how top-down stereotypes continuously interact in real time with bottom-up face processing. As such, one's stereotypic expectations may enter in on the categorization process. In some cases, this can alter categorizations wholesale; in other cases, it only temporarily alters them by simultaneously activating an alternate race category. Moreover, such effects are more pronounced as a face increases in racial ambiguity, and these effects are evidenced in both the mouse-tracking experimental data and in computational simulations. Thus, although it is often thought that prejudice is a consequence of initially categorizing others (Allport, 1954), the present results suggest that our prejudices also affect even initial categorization itself. These results also extend the findings of the previous studies by documenting yet another source of bottom-up information at play in driving the categorization process: the contextual cues that often surround a face in the real world.

GENERAL DISCUSSION

Across 8 studies, converging evidence provided support for the dynamic and interactive nature of person construal. Specifically, the studies support the proposed theory that the perception of other people is accomplished by a dynamical system in which lower-level sensory perception and higher-order social cognition coordinate across multiple interactive levels of processing. The computational model presented here was intended to capture these theoretical claims, and simulations demonstrated that it can account for many of the experimental results. These include the finding of continuously evolving category representations and a dynamic competition process that underlies social categorization (Studies 1–4), the finding of continuous face–voice interaction during categorization (Studies 5–6), and the finding of activated stereotypes' top-down influences on categorization (Studies 7–8).

Studies 1-2 were used to test the existence of a dynamic competition process argued to underlie our ability to slot others into social categories, such as sex and race. This was evidenced by mouse movements' partial, simultaneous attraction to the opposite category when categorizing a face bearing some overlapping cues tied to that category. These studies show that there is a competition process inherent to social categorization, which allows the natural diversity in others' sensory cues to be translated into a rigid categorization. The was best evidenced in Study 3, where the strength of the competition, indexed by trajectory curvature, sensitively increased and decreased with the amount of racial ambiguity perceivers were presented with. Thus, this competition process is initially sensitive to the full perceptual continuum underlying a social category dimension, but then eventually slots it into a rigid category (e.g., White or Black).

Not only was the person construal process found to be highly sensitive to the bottom-up cues in one sensory modality, but also those originating from multiple modalities at the same time. Study 5 found that, when categorizing a face's sex, the simultaneous processing of vocal cues exerted an ongoing influence on the categorization process (as evidenced in mouse movements). Study 7 then found that other extraneous information sources beyond cues from bottom-up modalities, such as an individual's top-down stereotypes, can also constrain categorization in dynamic fashion. For example, if a face surrounded by low-status attire was categorized as "White", mouse trajectories exhibited a continuous attraction to the "Black" response (stereotypically tied to the surrounding cue). Thus, stereotypes triggered by contextual cues exerted top-down pressure on the categorization process as it was being driven by bottom-up face processing in parallel. As such, prior stereotypic expectations temporarily altered categorization. Together, these studies show how multiple bottom-up and top-down information sources interact and constrain one another in driving perceptions.

Across the various mouse-tracking studies, the reader may have been tempted to inspect the mean trajectories (e.g., Figure 4) and attempt to pinpoint "when" a category decision occurred. For instance, perhaps the moment when the mouse first deviated from the horizontal center ($x = 0$) towards the correct response is when a decision was made. The present work argues, however, that there is no single moment at which a decision occurs; rather, all processing leading up to the mouse-click (including motor execution of the mouse-click itself) is argued to be temporally dynamic. As described earlier, during early moments of face processing a rough sketch of the face rapidly accumulates into neuronal populations, which would afford a quick-and-dirty interpretation of another's face that continuously sharpens across processing (Rolls & Tovee, 1995). Thus, early on, mouse movements could already begin heading in the direction of the correct response, but this need not indicate that a decision has already been made; it would simply indicate that, during these moments, the competition was predominantly swung toward the correct category. But the competition is likely not over. Although the mouse may start travelling in the correct direction, the other category may still be simultaneously active. Imagine a moment where the transient interpretation of a face's race, for example, is 75% in support of the White category and 25% in support of the Black category. The mouse should already be heading in the direction of a "White" response (as the White category is predominantly activated), thus deviating from $x = 0$, but it should nonetheless still be partially attracted to the "Black" response (as the Black category is still 25% activated). Thus, I argue that there is no instantaneous moment to pinpoint at which one social category is discretely selected and other categories vanish from the processing landscape. Rather, an ultimate categorization could simply be the end-result of continuously fluctuating category representations that gradually settle into a stable, steady state (also see Dale, Kehoe, & Spivey, 2007; Spivey & Dale, 2004).

Summary

According to the proposed model, perceptions of other people gradually emerge through the ongoing interaction between social categories, stereotypes, high-level cognitive states, and the low-level processing of facial, vocal, and bodily cues. Internal representations of categories and stereotypes are dynamically and probabilistically reconstructed, rather than behaving like static, symbol-like structures that wait around inertly until discretely accessed (see Footnote 2). The real-time evolution of these probabilistic representations is in continuous interaction with other activations across the system, both influenced by these other activations and a source of influence over them. The entire system's prior history, its visual inputs (e.g., facial cues), auditory inputs (e.g., vocal cues), and top-down inputs (e.g., stereotypes, task demands), its internal constraints, and some random noise jointly determine the construal of other people.

Taken together with the experimental mouse-tracking results, the present work suggests that perceptions of other people continuously evolve over fractions of a second and emerge from the mutually constraining interaction of multiple bottom-up sensory cues (e.g., facial, vocal, and contextual cues) and top-down social factors (e.g., stereotypes). Moreover, multiple conflicting perceptions may often be triggered during this process, and these gradually stabilize over time onto single categorical outcomes through dynamic competition. As such, person construal readily makes compromises between the variety of sensory cues inherent to another person and the baggage an individual perceiver brings to the construal process. Higher-order social cognition, according to this framework, is thus free to constrain and alter lower-level perceptual processes. This therefore furthers the emerging perspective that human perception is driven by an intimate interplay between both sensory and social phenomena (Adams, Ambady, Nakayama, & Shimojo, 2010; Balcetis & Lassiter, 2010).

Comparison with Extant Models

Extant social psychological models have described how perceivers form high-level impressions of other people, whether they utilize category-based or individual-based information, and how knowledge about individuals and groups is learned, stored, and accessed (Bodenhausen & Macrae, 1998; Brewer, 1988; Chaiken & Trope, 1999; Conrey et al., 2005; Fiske et al., 2002; Fiske & Neuberg, 1990; Higgins, 1996; Kunda & Thagard, 1996; Read & Miller, 1998b; Smith & DeCoster, 1998; Srull & Wyer, 1989; van Overwalle & Labiouse, 2004). Models in the cognitive face-processing literature, on the other hand, have described the visual and perceptual mechanisms that permit face recognition (Bruce & Young, 1986; Burton et al., 1990; Valentin, Abdi, O'Toole, & Cottrell, 1994). The proposed dynamic interactive model helps unify these two literatures by describing how the lower-level perceptual processing modeled in the cognitive literature works together with the higher-order social cognitive processes modeled in the social literature to give rise to person construal.

Social psychological models have tended to use categorization as a starting point, with relatively little focus on the perceptual processing that gives rise to it. Thus, in Fiske and Neuberg’s (1990) influential model of impression formation it is argued that the utilization of stereotypes, which is derived from a dominant categorization, is prioritized over more individual-based information in forming impressions, unless the perceiver is motivated to move further and individuate the target. This model, like Brewer’s (1988) and Kunda and Thagard’s (1996) models of impression formation, provide comprehensive accounts of how top-down processes, such as stereotypic expectations, motivation, and attention, interact with the bottom-up process of learning explicit individuated characteristics about a target. In these models, therefore, a target’s category memberships are given, and their influence on subsequent interpersonal phenomena are richly described (e.g., impressions, behavior). This is also the case for other models of person perception, such as Bodenhausen and Macrae’s (1998) stereotype activation and inhibition model. As such, categorization (and corresponding stereotype activation) is the initial input into these models. The focus of these models is not to explain the categorization process itself; it is to explain the higher-order social cognitive processing that comes after.

The present framework builds on these important models by fleshing out the initial category and stereotype activation process and explaining how this process is dynamically driven by both bottom-up sensory information as well as high-level top-down factors. Notably, this expands on extant models by explaining how initial category and stereotype activation may be influenced, sometimes considerably, by top-down factors. Although models of person perception have always emphasized the role of top-down factors (e.g., expectations, motivation, and attention), these factors have not been readily acknowledged to seep down into lower levels of processing, into the initial category and stereotype activation process itself. For example, such top-down factors had an important role in Studies 7–8, where status stereotypes activated by contextual attire cues altered the perception of a face's race. The modeling of the reach of top-down influences into even lower levels of person perception, such as basic category activation, thus builds on extant models that have generally described only the reach of top-down influences into higher levels of processing.

Beyond the importance of accounting for how perceptual processing brings about social cognitive phenomena in general, the modeling of perceptual processing is also important because it can bear a variety of downstream effects. For example, as shown in the present studies, within-category facial or vocal variation affects the dynamic competition inherent to categorization. This can in turn affect the eventual stable category representations that perceivers settle into (Locke et al., 2005). Thus, more prototypically masculine facial or vocal features (relative to less), for instance, affects the competition between male and female categories, which results in a stronger stable representation of the male category and weaker stable representation of the female category (see Figure 6). This can bear a variety of downstream effects, shaping trait attributions (Blair, Chapleau, & Judd, 2005; Blair, Judd, & Fallman, 2004; Blair et al., 2002; Ko et al., 2006; Maddox & Gray, 2002) as well as behavior (Blair, Judd, & Chapleau, 2004; Johnson, Eberhardt, Davies, & Purdie-Vaughns, 2006). Thus, the present framework builds on extant models by shedding new insights into the relationship between the higher-order processes these models have described and the lower-level perceptual processing that has received less attention.

Implications

The present findings and framework have several implications for present understandings of person construal, which are discussed here.

Re-thinking the "Multiple Category Problem"

Individuals naturally vary along any number of category dimensions (e.g., sex, race, age). Extant models have often emphasized that one category and the stereotypes tied to that category come to dominate the processing landscape, whereas others are actively suppressed, making the perceiver's job easier and thereby solving the "multiple category problem" (e.g., Bodenhausen & Macrae, 1998; Macrae, Bodenhausen, & Milne, 1995; Sinclair & Kunda, 1999).

According to the present model, the selection of one category and winnowing of other categories is accomplished by top-down pressure from higher-order nodes. For instance, the task demands of sex categorization, expressed by higher-order nodes, exerts excitatory pressure on to-be-attended categories (male, female) and inhibitory pressure on task-irrelevant categories (e.g., Black, White, Asian). However, the model introduces several nuances to an understanding of how the person construal system comes to arrive at focal categorizations of others.

The present model assumes that these top-down task-demand pressures exert their differential influence on categories dynamically over time. Thus, although for the purposes of sex categorization an applicable sex category rises in activation (thus becoming focally attended) whereas an applicable race category falls (thus becoming ignored), this pattern of excitation and inhibition is not instantaneous. Rather, higher-order task-demand nodes gradually exert excitatory pressure on certain categories while exerting inhibitory pressure on others. Thus, while these pressures are still at work the model predicts that multiple applicable category memberships (e.g., sex, race) are actually flexibly active in parallel. This places the model in line with neural dynamic models of visual attention (Desimone & Duncan, 1995), which assume a similar parallel activation of multiple representations.

Because multiple applicable category memberships (e.g., Black, janitor) may be active in parallel while the system works toward stabilizing on a focal category (e.g., Black), non-focal categories also have the ability to influence perception. This is because their partial parallel activation can powerfully affect the system's trajectory and the stable states it achieves. A clear demonstration of this, for example, is found in the stereotype-mediated race–occupation interactive effects of Studies 7–8. Due to the context of a race categorization task, higher-order nodes exerted excitatory pressure on race-category nodes and inhibitory pressure on occupation-category nodes. While these top-down task-demand pressures were at work, for a great deal of processing time occupation categories were still partially active in parallel. This occupation-category activation had a powerful effect on the trajectory of the system and on race categorization in particular. The competition between race categories, which sometimes was initiated by a completely race-ambiguous face (and thus initially equibiased with respect to bottom-up visual input) was powerfully swayed one way (White) or the other (Black) based on the partial parallel activation of presumably task-irrelevant occupation categories. Specifically, when targets were surrounded by janitor attire, the partial activation of the JANITOR category biased the race-category competition toward a Black categorization. Thus, non-focal, presumably task-irrelevant categories (e.g., occupation in a race categorization task) can bear powerful influences on focal person construals.

The model also implies that, in the absence of strong top-down factors that require all but one category to be inhibited (e.g., task demands, goals), the person construal system could settle into stable states that are quite flexible. For instance, without higher-order nodes exerting inhibitory pressures on particular category nodes, the stable states that the person construal system settles into could easily involve multiple categories (e.g., White, male) that are flexibly active in parallel. Indeed, the quality of having multiple person characteristics (e.g., lazy, friend, lives nearby) partially active in parallel is a central feature of the content-addressable memory modeled in connectionist networks of person memory (Smith, 1996, 2000; Smith & DeCoster, 1998). Just as multiple categories have often been shown to simultaneously constrain high-level impressions and social reasoning (Kunda & Thagard, 1996; Read & Miller, 1998b), a dynamic interactive theory proposes that they also simultaneously constrain lower-level person construals.

Thus, the "multiple category problem" might best be characterized not so much as a "problem" that must be eliminated to keep cognitive efficiency (Allport, 1954), but rather as a reflection of the flexibility of interactive, parallel category representations.

The Dynamic Coextension of Category and Stereotype Activation

Recent research has found that variation in facial features may bear effects on stereotype activation that are independent of a target's category membership. For instance, the presence of Black-specifying cues on a person who is not Black (e.g., a White face) increases Black-related stereotypic attributions (Blair et al., 2005; Blair et al., 2002). These effects may thereafter influence behavior as well. For example, in court trials, targets with more Black-specifying features are punished more severely and more likely to be sentenced to death (Blair, Judd, & Chapleau, 2004; Johnson et al., 2006). Based on such findings, some accounts have argued that these independent feature-based effects on stereotype activation are accomplished by a special feature-based processing route, where features become associated with stereotypes unmediated by any category representation at all (Blair et al., 2002; Livingston & Brewer, 2002). This direct *features* → *stereotypes* route is theorized to be separate from a more typical *categories* → *stereotypes* route.

The present model agrees with these previous accounts that facial features can influence stereotype activation without a discrete categorization. However, because the model permits categorizations to be partially active in parallel, independent feature-based effects on stereotype activation could be mediated by the tentative, partially-active categorization of an alternate category. Specifically, the model suggests that independent feature-based effects on stereotyping are a product of the dynamic processing cascade inherent to the system. Cues of an alternate category (e.g., Black-specifying cues on a White target) trigger partially-active, competing category representations (e.g., "he's [tentatively] White" versus "he's [tentatively] Black"). Both category representations (e.g., White, Black) then immediately pass activation onward to their respective stereotypes before the competition in the category level has resolved and settled into just one alternative. This is reflected in Figure 6, where feminine cues on a man's face triggered the partial and parallel activation of the FEMALE category, which continuously cascaded into the partial and parallel activation of the female-related stereotype, DOCILE, as was shown with human perceivers (Freeman & Ambady, 2009). Thus, the dynamic coextension of category and stereotype activation permits independent feature-based effects on stereotype activation. As such, the present model parsimoniously accounts for independent feature-based effects on stereotyping by one single route involving a dynamic processing cascade.

A Rapidly Adaptive, Ecologically Valid Person Construal System

Like the present model, the ecological approach to social perception (McArthur & Baron, 1983) emphasized the need to study directly the stimulus information that avails perceivers with functionally significant characteristics about other people. It also emphasized the inherently dynamic and multimodal nature of social stimuli. The dynamic interactive framework presented here is in the best spirit of this approach and builds on it in several ways.

This framework brings new and helpful ways of thinking about ecologically-valid person construal. Specifically, it assumes that the person construal system's processing is fully continuous and highly interactive, and that its representations are probabilistic, active in parallel, and changing over time. This is exactly the sort of system required for the ecologically-valid person perceiver—the kind of perceiver that must make sense of others in real-time, on-the-fly, and in a rapidly changing social environment. In real-world social encounters the sensory stimulation of another person is almost always in continuous flux (Gibson, 1979). The most obvious example might be the perception of a face's emotion, which continuously fluctuates over time. Rarely do perceivers encounter a static emotional expression. Rather, for just a few fleeting moments, another's face displays slight anger, which then rapidly transitions into some other expression. By the time perceivers are finished processing that subtle anger, however, there are already hundreds of milliseconds of new visual information that needs to be accrued and dealt with. In real-world person construal, therefore, another's face tends not to fit squarely into any one expression (e.g., angry), but is usually in some in-between state amidst one interpretable expression and the next, and rarely standing still.

For simplicity, in the presented instantiations of the model external input was supplied to the network discretely (at iteration 1). However, the model is flexible to support the more ecologically-valid situation in which external stimulation to the network dynamically changes across time based on changing cues in the social environment. As a face's emotion, a body's subtle nonverbal behavior, or the ongoing stream of vocal cues fluctuate over time, the visual and auditory inputs into cue nodes would continually change across iterations accordingly. This would thereby continually change, iteration to iteration, the amount of excitatory and inhibitory pressures on category and stereotype nodes. As such, at any given moment while the system is trying to settle into one stable state, new sensory information bombarding the system would already start changing the various stable states to which the system will start gravitating (Spivey, 2007). This leaves little time for the system to actually rest in any given stable state, since by the time it starts to stabilize it is already being pushed out of its stability by new constraints (e.g., changes of a face's emotion, of the body's behavior, of the voice stream). Thus, the network outlined here is a rapidly adaptive and dynamic person construal system. Its continually evolving states are able to be tightly yoked to the ongoing sensory stimulation of the social environment.

This adaptive, dynamic person construal system is potentially stimulated by continuous top-down input as well. For instance, ecologically-valid, moment-to-moment changes in one's goals or attentional states, among other top-down factors, would continually stimulate higher-order nodes, which thereafter continually change the amount of excitatory and inhibitory pressures on category and stereotype nodes. Thus, although for the sake of simplicity external inputs into the network were modeled as discrete occurrences, the system is inherently capable of supporting stimulation by a dynamically changing social environment as well as dynamically changing internal cognitive states.

An ecologically-valid person construal system also needs to permit ongoing perceptions of other people to guide action continuously over time. In social interaction, something apparent on individual A's face and gesturing elicits a reaction on individual B's face and gesturing, which then elicits a reaction on individual A's face and gesturing, and so on and so forth. Thus, there is no staccato series of static images and sounds that elicit particular reactions. Instead, ecologically-valid person construal would likely need to involve continuous millisecond-by-millisecond updates of facial, vocal, and bodily information, and these updates need to make their way onto the motor system immediately, not once the system has 100% finalized the processing of each transient image or sound in a social interaction. Indeed, recent neurophysiological evidence suggests that this dynamic person processing is a likely possibility. In a series of event-related potential studies, it was shown that the process of social categorization immediately shares its ongoing results with the motor cortex to guide action continuously over time (Freeman, Ambady, et al., 2011). This is consistent with multi-cell recordings in nonhuman primates (Cisek & Kalaska, 2005, 2010). Thus, person construal is characterized by continuous perceptual–cognitive–motor dynamics, such that perceptual, cognitive, and motor processing are coextensive. Cognitive representations of a face's category memberships develop over hundreds of milliseconds while perceptual processing is ongoing, and these representations evolve alongside accruing perceptual evidence for category alternatives. Further, ongoing results of this social category processing are immediately cascaded into the motor cortex to guide relevant actions continuously over time. Thus, person construal is continuously coextensive with action. This is exactly the kind of processing required by the ecologically-valid person perceiver.

In short, described here is a person construal system that is rapidly adaptive and dynamic. It is able to perceive others in an ecologically-valid, real-time social environment, while also able to coordinate with the motor system to act on ongoing perceptions.

New Predictions and Future Directions

Beyond the present model's ability to explain a variety of experimental results, it also gives rise to a number of new and distinctive predictions, which future work could directly examine. Below are a few examples of important predictions derived from the model that could serve as testable hypotheses in the future.

Category Interactions due to Incidental Stereotypic and Phenotypic Overlap

The present model makes the novel prediction that any incidental overlap in the stereotype or phenotype content of two category memberships would lead the system to throw those categories into interaction. As shown in Studies 7–8, overlapping stereotype content between the Black and janitor categories (e.g., low status) and between the White and businessperson categories (e.g., high status) created top-down pressure that gave rise to race–occupation interactions. However, any number of category interactions are possible and, in fact, quite likely. Many stereotypes are likely to be incidentally shared by multiple categories. In fact, the very existence of some categories may be predicated on the stereotypes of other categories, such as sexual orientation categories and sex-category stereotypes (Kite & Deaux, 1987), and this is evident in perceptual construals (Freeman, Johnson, Ambady, & Rule, 2010; Johnson, Gill, Reichman, & Tassinary, 2007). Future work could empirically estimate the degree of stereotype overlap between categories using explicit or implicit measures, and implement the estimated overlap into the stereotype and category levels. A variety of category interactions could arise in network simulations, and these could then be experimentally tested in the laboratory.

Similarly, categories could also be thrown into interaction through bottom-up processes as well. For example, the perceptual cues contained in the face, voice, and body are likely to, by chance, partly covary between categories. This has been shown with sex and emotion categories (Becker et al., 2007), where angry male and happy female faces are characterized by more efficient processing. Future work could empirically estimate the degree of phenotype overlap between categories and then implement this estimated overlap into the cue and category levels. For example, face-modeling software can derive precise estimates of hundreds of facial cues from a facial photograph (e.g., Blanz & Vetter, 1999). Thus, researchers could derive estimates of cue overlap using representative samples of faces for specific category memberships, and then implement these estimates into instantiations of the model. If category interactions arose in network simulations, these could then be experimentally investigated in the laboratory.

Category interactions could also potentially be driven by both top-down and bottom-up overlap at the same time (see Johnson, Freeman, & Pauker, 2011). For example, not only do male and angry cues and female and happy cues overlap (Becker et al., 2007), but also men are stereotyped as angry and women are stereotyped as happy (Fabes & Martin, 1991). Simulations with the model are uniquely poised to assess the relative contribution of potentially coexistent top-down and bottom-up forces in driving category interactions. Such simulations could also be used to tease apart the time-courses of these two forces' influence on perceptions.

Social Category Blending and Multiracial States

In the mouse-tracking studies reported here, participants were constrained to making dichotomous categorical judgments. Such dichotomous judgments may have a great deal of real-world plausibility for some category dimensions, such as sex, but less so perhaps for other dimensions, such as race. Although recent work suggests that people effortlessly slot racially ambiguous faces into monoracial categories (Bodenhausen & Peery, 2009; Pauker et al., 2009; Peery & Bodenhausen, 2008), an increasingly heterogeneous racial landscape is leading to the recognition of more multiracial identities at both an individual and institutional level (Jones & S., 2001; Lee & Bean, 2004; Pauker & Ambady, 2009; Renn, 2009; Rockquemore, Brusma, & Delgado, 2009). The present research provides perhaps a promising look at the future of race categorization in this heterogeneous landscape. Even when perceivers are constrained to making monoracial categorizations, the present work suggests that before a race-category decision is settled into, the story is a lot fuzzier. For the vast majority of race-category processing, it was found that perceivers entertain many dynamically changing in-between states amidst traditional race categories—even when targets are slotted into traditional categories if demanded by the task (or society). Thus, the present studies show that race categorization is inherently capable of supporting all sorts of fuzzy, graded mixtures of multiracial interpretations. Before perceivers ultimately fit another's face into a traditional race category, it goes through an ongoing process of fluctuating interpretation that is, in a sense, inherently "mixed-race" (e.g., 60% White, 40% Black). The computational simulations were consistent with this dynamic process as well. That race categorization flexibly supports these sorts of graded categorical blends—even for fractions of a second—perhaps provides promise for the future of social categorization in a more heterogeneous cultural milieu, ripe with people that blur traditional categorical lines. It will be interesting for future research to apply a mouse-tracking paradigm and this dynamic interactive framework to understanding the resolution of racially ambiguous targets into explicitly multiracial, rather than monoracial, categories.

Future research should also consider whether such graded categorical blends may have important downstream implications. Since the writings of Allport (1954), for example, a thorny issue that social psychology has had to tackle is the relationship between race categorization and prejudice. Categorization was long thought to be an inevitable consequence of the perception of others (Allport, 1954), with recent work seeking evidence for the avoidable nature of categorization in the hopes that this could mitigate prejudice (e.g., Blair, 2002; Kurzban, Tooby, & Cosmides, 2001; Macrae & Bodenhausen, 2000). The present findings suggest, however, that although others are eventually slotted into monoracial categories, the fuzzy overlap with other tentative categories (e.g., White, for a mixed-race face seen as Black) is—although not reflected in the categorization *outcome*—dynamically retained in the *process*, in those fleeting moments between catching sight of another's face and settling into an eventual category. Knowing that the race categorization process can support these transitory, categorical in-between states, future research might examine whether these could ever manifest in ultimate categorization outcomes, potentially bearing implications for the tendency to prejudge others by oversimplified monoracial categories.

Downstream Consequences of "Hidden" Parallel Activations

Anderson (2002) argued for the importance of bridging psychological phenomena across multiple orders of temporal magnitude. Here a model of person construal was proposed, which fleshes out the process by which an ultimate perception crystallizes on the order of hundreds of milliseconds. But how do these relatively low-level, fine-grained dynamics relate to higher-order phenomena on the order of hundreds of seconds or hours, such as aspects of social interaction or other behavioral outcomes? There are likely many relationships to be uncovered. For example, the model predicts that several unforeseen category and stereotype representations may be simultaneously and partially active before perceivers arrive at an ultimate construal. Subtle bottom-up overlap with an alternate category (e.g., slight feminine facial features on a man) can lead to partial parallel activation of that alternate category (e.g., female). Or, high-level cognitive states or stereotypes can exert top-down influences on category-level processing, in turn triggering partially-active representations of other candidate categories. The proposed model therefore predicts that, for a great many of our construals of others, a variety of "hidden" category and stereotype activations may be partially triggered in parallel—activations that are not reflected in an ultimate perceptual outcome.

Such subtle activations triggered during real-time construal could likely give rise to a variety of unforeseen downstream consequences. The lasting effects of category and corresponding stereotype activation on higher-order social phenomena—even the briefest of kinds (e.g., priming)—have long been documented. Activated stereotypes change how we think about others, judge, and remember them (Bodenhausen, 1988; Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990). They also activate related attitudes and behavioral tendencies, in turn changing how we feel about others and evaluate them (Fazio, Sanbonmatsu, Powell, & Kardes, 1986) and how we interact with others and treat them (Bargh et al., 1996; Chen & Bargh, 1999). Thus, future work could investigate how "hidden" parallel activations of alternate categories and stereotypes computed during the construal process, or other aspects of this real-time process, relate to important downstream phenomena. Moreover, such work could test how variation in the presence of these parallel activations relates to measures of individual differences (e.g., levels of prejudice or motivation) or other behavioral outcomes.

Future Advances to the Model

Future work could advance the model and simulations presented here in several ways. First, the simulations presented here were limited to focusing on how sensory information and high-level cognitive states temporally conspire to shape category and stereotype activations. However, any given change in one node of the system will lead to changes in all other nodes, as the system works over time to maximally satisfy all of its constraints in parallel. Thus, the model is highly interactive and inherently bidirectional. It therefore assumes that, beyond high-level cognitive states shaping lower levels of processing, lower levels of processing also shape high-level cognitive states. As such, the model predicts that sensory information and category and stereotype activations should all lead to a variety of changes in high-level cognitive states. However, in the present work our focus was on category and stereotype activations as the dependent measures of interest. Future work could develop the model further by testing the reverse relationship, making high-level states the dependent measure of interest (e.g., motivation, prejudice, top-down attention, affect) and examining how these states are shaped by a rich interaction with lower levels of processing, as the model predicts.

The model could also be advanced by deriving network parameters empirically (see Footnote 3), and experimental studies could be used to refine and expand the model. For example, data could be collected for estimating the connection weights between category nodes and potentially hundreds of stereotype nodes (e.g., via explicit or implicit measures) and hundreds of cue nodes (e.g., via face-modeling algorithms), and all these nodes and their weighting could be implemented in future versions. This would bring the model closer to the empirical rigor and level of quantification common to connectionist models of speech perception (e.g., McClelland & Elman, 1986). Moreover, future work could opt to replace the cue level with more sophisticated approaches to modeling the uptake of sensory information, such as a pixel-based image processor (e.g., Burton, Bruce, & Hancock, 1999). This would make fewer assumptions about the role of specific features and instead rely more on the emergent properties inherent in other people's sensory information. Together, such advances would allow the model to better reflect the real-world interrelatedness among cues, categories, stereotypes, and high-level states.

CONCLUSION

A new approach to the study of person perception is on the rise, as evidenced by the two recent volumes, *The Science of Social Vision* (Adams et al., 2010) and *The Social Psychology of Visual Perception* (Balcetis & Lassiter, 2010). Social psychologists are working alongside researchers in the cognitive, neural, and vision sciences to provide a unified and more complete understanding of person perception. In the present work, I sought to open up the temporally extended, real-time process of person construal. In this real-time process, person construal is dynamic and interactive, and the connection between the "sensory" and the "social" is an intimate one. The theory, model, and studies presented here together show that many person construal phenomena may be accounted for by a dynamical system that permits lower-level sensory perception and higher-order social cognition to continually collaborate across multiple interactive levels of processing. Low-level sensory information and high-level social factors fluidly work together to give rise to stable and integrated perceptions of other people. Probabilistic and parallel construals gradually emerge through the ongoing interaction between categories, stereotypes, high-level cognitive states, and the low-level processing of facial, vocal, and bodily cues. My hope is that a dynamic interactive framework for person construal will provide a helpful guiding force in the burgeoning interdisciplinary effort to understand the perception of our social worlds.

REFERENCES

Adams, R. B., Ambady, N., Nakayama, K., & Shimojo, S. (2010). *The Science of Social Vision*. New York: Oxford University Press.

Allport, G. W. (1954). *The nature of prejudice*. Oxford: Addison-Wesley.

Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgmental accuracy from thin slices of the behavioral stream *Advances in experimental social psychology, Vol 32* (pp. 201-271). San Diego, CA: Academic Press.

Anderson, J. R. (2002). Spanning seven orders of magnitude: a challenge for cognitive modeling. *Cognitive Science, 26*, 85-112.

Balcetis, E., & Dunning, D. (2006). See what you want to see: Motivational influences on visual perception. *Journal of Personality and Social Psychology, 91*, 612-625.

Balcetis, E., & Lassiter, D. (2010). *The Social Psychology of Visual Perception*. New York: Psychology Press.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition *(1994) Handbook of social cognition, Vol 1: Basic processes; Vol 2: Applications (2nd ed)* (pp. 1-40). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.

Bargh, J. A. (1997). The automaticity of everyday life. In J. A. Bargh & R. S. Wyer (Eds.), *The automaticity of everyday life* (Vol. 10, pp. 1-61). Mahwah, NJ: Erlbaum.

Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361-382). New York: Guilford Press.

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71*(2), 230-244.

Becker, D. V., Kenrick, D. T., Neuberg, S. L., Blackwell, K. C., & Smith, D. M. (2007). The confounded nature of angry men and happy women. *Journal of Personality and Social Psychology, 92*, 179-190.

Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review, 6*(3), 242-261.

Blair, I. V., Chapleau, K. M., & Judd, C. M. (2005). The use of Afrocentric Features as Cues for Judgment in the Presence of Diagnostic Information. *European Journal of Social Psychology, 35*(1), 59-68.

Blair, I. V., Judd, C. M., & Chapleau, K. M. (2004). The influence of Afrocentric facial features in criminal sentencing. *Psychological Science, 15*(10), 674-679.

Blair, I. V., Judd, C. M., & Fallman, J. L. (2004). The Automaticity of Race and Afrocentric Facial Features in Social Judgments. *Journal of Personality and Social Psychology, 87*(6), 763-778.

Blair, I. V., Judd, C. M., Sadler, M. S., & Jenkins, C. (2002). The role of Afrocentric features in person perception: Judging by features and categories. *Journal of Personality and Social Psychology, 83*(1), 5-25.

Blanz, V., & Vetter, T. (1999). *A morphable model for the synthesis of 3D faces*. Paper presented at the SIGGRAPH'99, Los Angeles.

Bodenhausen, G. V. (1988). Stereotypic biases in social decision making and memory: Testing process models of stereotype use. *Journal of Personality and Social Psychology, 55*(5), 726-737.

Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition *Stereotype activation and inhibition* (pp. 1-52). Mahwah, NJ: Lawrence Erlbaum Associates Publishers.

Bodenhausen, G. V., & Macrae, C. N. (2006). Putting a face on person perception. *Social Cognition, 24*(5), 511-515.

Bodenhausen, G. V., & Peery, D. (2009). Social categorization and stereotyping In vivo: The VUCA challenge. *Social and Personality Psychology Compass, 3*, 133-151.

Brefczynski, J. A., & DeYoe, E. A. (1999). A physiological correlate of the 'spotlight' of visual attention. *Nature Neuroscience, 2*, 370-374.

Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer (Eds.), *A Dual-Process Model of Impression Formation: Advances in Social Cognition* (Vol. 1, pp. 1-36). Hillsdale, NJ: Erlbaum.

Bruce, V., & Young, A. W. (1986). A theoretical perspective for understanding face recognition. *British Journal of Psychology, 77*, 305-327.

Burton, A. M., Bruce, V., & Hancock, P. J. B. (1999). From Pixels to People: A Model of Familiar Face Recognition. *Cognitive Science, 23*, 1-31.

Burton, A. M., Bruce, V., & Johnston, R. A. (1990). Understanding race recognition with an interactive activation model. *British Journal of Psychology, 81*, 361-380.

Calder, A. J., & Young, A. W. (2005). Understanding the recognition of facial identity and facial expression. *Nature Reviews Neuroscience, 6*, 641-651.

Campanella, S., & Belin, P. (2007). Integrating face and voice in person perception. *Trends in Cognitive Sciences, 11*, 535-543.

Chaiken, S., & Trope, Y. (1999). *Dual-Process Theories in Social Psychology*. New York: Guilford.

Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin, 25*(2), 215-224.

Cisek, P., & Kalaska, J. F. (2005). Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. *Neuron, 45*(5), 801-814.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience, 33*, 269-298.

Cloutier, J., Mason, M. F., & Macrae, C. N. (2005). The Perceptual Determinants of Person Construal: Reopening the Social-Cognitive Toolbox. *Journal of Personality and Social Psychology, 88*(6), 885-894.

Cohen, J. D., & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In C. Umilta (Ed.), *Attention and Performance XV* (pp. 453-476). Cambridge, MA: MIT Press.

Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The Quad Model of implicit task performance. *Journal of Personality and Social Psychology, 89*, 469-487.

Dale, R., Kehoe, C., & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition, 35*(1), 15-28.

Dale, R., Roche, J., Snyder, K., & McCall, R. (2008). Exploring action dynamics as an index of paired-associate learning. *PloS ONE, 3*, e1728.

de Gelder, B., & Vroomen, J. (2000). The perception of emotions by ear and by eye. *Cognition and Emotion, 14*, 289-311.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18*, 193-222.

Devine, P. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5-18.

Dijksterhuis, A., & Van Knippenberg, A. (1998). The relation between perception and behavior, or how to win a game of trivial pursuit. *Journal of Personality and Social Psychology, 74*(4), 865-877.

Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). The nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology, 33*, 510-540.

Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron, 28*, 287-298.

Eberhardt, J. L., Dasgupta, N., & Banaszynski, T. L. (2003). Believing is seeing: The effects of racial labels and implicit beliefs on face perception. *Personality and Social Psychology Bulletin, 29*(3), 360-370.

Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*

2, 704-716.

Fabes, R. A., & Martin, C. L. (1991). Gender and age stereotypes of emotionality. *Personality and Social Psychology Bulletin, 17*, 532-540.

Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review, 105*, 482-498.

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: a bona fide pipeline? *Journal of Personality and Social Psychology, 69*(6), 1013-1027.

Fazio, R. H., Sanbonmatsu, D., Powell, M., & Kardes, F. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229-238.

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology, 82*(6), 878-902.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum model of impression formation from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology, 23*, 1–74.

Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological Science, 20*, 1183-1188. doi: 10.1111/j.1467-9280.2009.02422.x

Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior Research Methods, 42*, 226-241.

Freeman, J. B., Ambady, N., & Holcomb, P. J. (2010). The face-sensitive N170 encodes social category information. *NeuroReport, 21*, 24-28. doi: 10.1097/WNR.0b013e3283320d54

Freeman, J. B., Ambady, N., Midgley, K. J., & Holcomb, P. J. (2011). The real-time link between person perception and action: Brain potential evidence for dynamic continuity. *Social Neuroscience, 6*, 139-155.

Freeman, J. B., Dale, R., & Farmer, T. A. (2011). Hand in motion reveals mind in motion. *Frontiers in Psychology, 2*, 59.

Freeman, J. B., Johnson, K. L., Ambady, N., & Rule, N. O. (2010). Sexual orientation perception involves gendered facial cues. *Personality and Social Psychology Bulletin, 36*, 1318-1331.

Ghazanfar, A. A., Chandrasekaran, C., & Logothetis, N. K. (2008). Interactions between the superior temporal sulcus and auditory cortex mediate dynamic face/voice integration in rhesus monkeys. *Journal of Neuroscience, 28*, 4457-4469.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Gilbert, C. D., & Sigman, M. (2007). Brain states: Top-down influences in sensory processing. *Neuron, 54*, 677-696.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60*, 509-517.

Goodale, M. A., Pelisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature, 320*, 748-750.

Groen, W. B., van Orsouw, L., Zwiers, M., Swinkels, S., van der Gaag, R. J., & Buitelaar, J. K. (2008). Gender in voice perception in autism. *Journal of Autism and Developmental Disorders, 38*, 1819-1826.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review, 87*, 1-51.

Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of Social Cognition* (2nd ed., Vol. 2, pp. 1-68). Hillsdale, NJ: Erlbaum.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*, 223-233.

Hietanen, J., Leppänen, J., Illi, M., & Surakka, V. (2004). Evidence for the integration of audiovisual emotional information at the perceptual level of processing. *European Journal of Cognitive Psychology, 16*, 769-790.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruganski (Eds.), *Social Psychology, Handbook of basic principles* (pp. 133-168). New York: Guilford Press.

Hugenberg, K., & Bodenhausen, G. V. (2004). Ambiguity in Social Categorization: The role of prejudice and facial affect in race categorization. *Psychological Science, 15*(5), 342-345.

Ito, T. A., & Urland, G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals. *Journal of Personality and Social Psychology, 85*, 616-626.

Ito, T. A., & Urland, G. R. (2005). The influence of processing objectives on the perception of faces: An ERP study of race and gender perception. *Cognitive, Affective, and Behavioral Neuroscience, 5*, 21-36.

Johnson, K. L., Freeman, J. B., & Pauker, K. (2011). Race is gendered: How Covarying Phenotypes and Stereotypes Bias Sex Categorization. *Journal of Personality and Social Psychology*, doi: 10.1037/a0025335.

Johnson, K. L., Gill, S., Reichman, V., & Tassinary, L. G. (2007). Swagger, sway, and sexuality: Judging sexual orientation from body motion and morphology. *Journal of Personality and Social Psychology, 93*(3), 321-334.

Johnson, K. L., Pollick, F., & McKay, L. (2010). Social constraints on the visual perception of biological motion. In R. B. Adams Jr., N. Ambady, K. Nakayama & S. Shimojo (Eds.), *The Science of Social Vision*. New York: Oxford University Press.

Johnson, S. L., Eberhardt, J. L., Davies, P. G., & Purdie-Vaughns, V. J. (2006). Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science, 17*, 383-386.

Jones, N. A., & S., S. A. (2001). *The two or more races population: 2000 (Census 2000 Brief No. C2KBR/01-6)*. Washington, DC: U.S. Census Bureau.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge: MIT Press.

Kite, M. E., & Deaux, K. (1987). Gender belief systems: Homosexuality and implicit inversion theory. *Psychology of Women Quarterly, 11*, 83-96.

Ko, S. J., Judd, C. M., & Blair, I. V. (2006). What the Voice Reveals: Within- and Between-Category Stereotyping on the Basis of Voice. *Personality and Social Psychology Bulletin, 32*(6), 806-819.

Kreifelts, B., Ethofer, T., Grodd, W., Erb, M., & Wildgruber, D. (2007). Audiovisual integration of emotional signals in voice and face: An event-related fMRI study. *Neuroimage, 37*, 1445-1456.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review, 103*, 284-308.

Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences, 98*(26), 15387-15392.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23*, 571-579.

Lee, J., & Bean, F. D. (2004). America's changing color lines: Immigration, race/ethnicity, and multiracial identification. *Annual Review of Sociology, 30*, 221-242.

Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 1364-1382.

Levin, D. T., & Angelone, B. L. (2002). Categorical perception of race. *Perception, 31*, 567-578.

Li, W., Piëch, V., & Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience, 7*, 651-657.

Livingston, R. W., & Brewer, M. B. (2002). What are we really priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology, 82*(1), 5-18.

Locke, V., Macrae, C. N., & Eaton, J. L. (2005). Is person categorization modulated by exemplar typicality? *Social Cognition, 23*(5), 417-428.

MacLin, O. H., & Malpass, R. S. (2001). Racial Categorization of Faces: The Ambiguous Race Face Effect. *Psychology, Public Policy and Law, 7*(1), 98-118.

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology, 51*, 93-120.

Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology, 69*(3), 397-407.

Macrae, C. N., & Martin, D. (2007). A boy primed Sue: Feature-based processing and person construal. *European Journal of Social Psychology, 37*(5), 793-805.

Macrae, C. N., Mitchell, J. P., & Pendry, L. F. (2002). What's in a forename? Cue familiarity and stereotypical thinking. *Journal of Experimental Social Psychology, 38*(2), 186-193.

Maddox, K. B., & Gray, S. A. (2002). Cognitive Representations of Black Americans: Reexploring the Role of Skin Tone. *Personality and Social Psychology Bulletin, 28*(2), 250-259.

Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.

McArthur, L. Z., & Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychological Review, 90*, 215-238.

McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology, 23*, 1-44.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*, 1-86.

Paninski, L., Fellows, M. R., Hatsopoulos, N. G., & Donoghue, J. P. (2004). Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology, 91*, 515-532.

Pauker, K., & Ambady, N. (2009). Multiracial faces: How categorization affects memory at the boundaries of race. *Journal of Social Issues, 65*, 69-86.

Pauker, K., Rule, N. O., & Ambady, N. (2010). Ambiguity and social perception. In E. Balcetis & D. Lassiter (Eds.), *The Social Psychology of Visual Perception*. New York: Psychology Press.

Pauker, K., Weisbuch, M., Ambady, N., Sommers, S. R., Adams Jr., R. B., & Ivcevic, Z. (2009). Not so Black and White: Memory for ambiguous group members. *Journal of Personality and Social Psychology, 96*, 795-810.

Peery, D., & Bodenhausen, G. V. (2008). Black + White = Black: Hypodescent in reflexive categorization of racially ambiguous faces. *Psychological Science, 19*, 973-977.

Port, R. F., & van Gelder, T. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.

Read, S. J., & Miller, L. C. (1993). Rapist or "Regular Guy": Explanatory Coherence in the Construction of Mental Models of Others. *Personality and Social Psychology Bulletin, 19*, 526-540.

Read, S. J., & Miller, L. C. (1998a). *Connectionist models of social reasoning and social behavior*. Mahwah, NJ: Erlbaum.

Read, S. J., & Miller, L. C. (1998b). On the dynamic construction of meaning: An interactive activation and competition model of social perception. In S. J. Read & L. C. Miller (Eds.), *Connectionist models of social reasoning and social behavior*. Mahwah, N. J.: Erlbaum.

Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionism, parallel constraint satisfaction processes, and gestalt principles: (Re)introducing cognitive dynamics to social psychology. *Personality and Social Psychology Review, 1*, 26-53.

Renn, K. A. (2009). Educational policy, politics, and mixed heritage students in the United States. *Journal of Social Issues, 65*, 163-181.

Rockquemore, K. A., Brusma, D. L., & Delgado, D. J. (2009). Racing to theory or re-theorizing race? Understanding the struggle to build a multiracial identity theory. *Journal of Social Issues, 65*, 13-34.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. Boston: Bradford Books.

Rolls, E. T., & Tovee, M. J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology, 73*, 713-726.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). *A general framework for parallel distributed processing*. Cambridge, MA: MIT Press.

SAS Institute. (1989). *SAS/STAT user's guide*. Cary, NC: Author.

Sherif, M. (1967). *Group conflict and co-operation: Their social psychology*. London: Routledge & K. Paul.

Sinclair, L., & Kunda, Z. (1999). Reactions to a black professional: motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77*, 885-904.

Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology, 70*(5), 893-912.

Smith, E. R. (2000). Subjective experience of familiarity: Functional basis in connectionist memory *The message within: The role of subjective experience in social cognition and behavior* (pp. 109-124). New York, NY: Psychology Press.

Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology, 74*(1), 21-35.

Smith, E. R., & DeCoster, J. (1999). Associative and rule-based processing: A connectionist interpretation of dual-process models *Dual-process theories in social psychology* (pp. 323-336). New York, NY: Guilford Press.

Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99*(1), 3-21.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences, 27*, 161-168.

Smolensky, P. (1989). Connectionist modeling: Neural computation/mental connections. In L. Nadel, A. Cooper, P. Culicover & R. M. Harnish (Eds.), *Neural connections, mental computations*. Cambridge, MA: MIT Press.

Song, J. H., & Nakayama, K. (2006). Role of focal attention on latencies and trajectories of visually guided manual pointing. *Journal of Vision, 6*, 982-995.

Song, J. H., & Nakayama, K. (2008). Target selection in visual search as revealed by movement trajectories. *Vision Research, 48*, 853-861.

Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences, 13*, 360-366.

Spivey, M. J. (2007). *The continuity of mind*. New York: Oxford University Press.

Spivey, M. J., & Dale, R. (2004). The continuity of mind: Toward a dynamical account of cognition *Psychology of Learning and Motivation* (Vol. 45, pp. 87-142). San Diego: Elsevier.

Spivey, M. J., & Dale, R. (2006). Continuous dynamics in real-time cognition. *Current Directions in Psychological Science, 15*(5), 207-211.

Srull, T. K., & Wyer, R. S. (1989). Person memory and judgment. *Psychological Review, 96*(1), 58-83.

Strangor, C., Lynch, L., Duan, C., & Glas, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology, 62*, 207-218.

Sugase, Y., Yamane, S., Ueno, S., & Kawano, K. (1999). Global and fine information coded by single neurons in the temporal visual cortex. *Nature, 400*, 869-873.

Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Social Issues, 25*, 79-97.

Usher, M., & McClelland, J. L. (2003). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108*, 550-592.

Valentin, D., Abdi, H., O'Toole, A. J., & Cottrell, G. W. (1994). Connectionist models of face processing: A survey. *Pattern Recognition, 27*, 1209-1230.

van Overwalle, F. (2007). *Social connectionism: A reader and handbook for simulations*. New York: Psychology Press.

van Overwalle, F., & Labiouse, C. (2004). A recurrent connectionist model of person impression formation. *Personality and Social Psychology Review, 8*, 28-61.

Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social Cognition, 8*(2), 161-185.

Zebrowitz, L. A. (2006). Finally faces find favor. *Social Cognition, 24*, 657-701.

Zebrowitz, L. A., Fellous, J.-M., Mignault, A., & Andreoletti, C. (2003). Trait impressions as overgeneralized responses to adaptively significant facial qualities: Evidence from connectionist modeling. *Personality and Social Psychology Review, 7*, 194-215.

Zebrowitz, L. A., & Montepare, J. M. (2008). Social psychological face perception: Why appearance matters. *Social and Personality Psychology Compass, 2*, 1497-1517.

Zeger, S. L., & Liang, K. Y. (1986). Longitudinal Data Analysis for Discrete and Continuous Outcomes. *Biometrics, 42*(1), 121-130.

APPENDIX A

Details on Model Structure

The model has a recurrent connectionist architecture that may be classified as a stochastic interactive activation network (McClelland, 1991; Rumelhart et al., 1986). In such networks, there are a number of nodes with connections that can be positive (excitatory) or negative (inhibitory). These nodes are not intended to represent individual neurons, but the overall structure of a network is often intended to be approximately neurally plausible (Smolensky, 1989). Most of these connections are bidirectional. Thus, as one node's activation tends to excite the nodes connected to it, the excitation of the nodes connected to it send feedback to the original node. Indeed, feedback is highly pervasive across the human brain (e.g., Brefczynski & DeYoe, 1999; Dragoi, Sharma, & Sur, 2000; Lamme & Roelfsema, 2000, also see Spivey, 2007), and thus recurrent connectionist networks have relatively high neural plausibility (Smolensky, 1989). Initially, the network is stimulated by external input. This input could come from bottom-up sources (e.g., facial or vocal cues) as well as top-down ones (e.g., task demands, prejudice, motivation). Activation then spreads among all nodes simultaneously as a function of their connection weights. Because many of the nodes receive feedback, complex feedback loops are produced within the system. This causes the system to gradually converge on an overall stable pattern of activation that best fits the input.

Because a node's activation is a function of all the positive and negative connections to other nodes that are activated in parallel (due to the feedback loops across the system), the final activation of a node (when the system stabilizes) can be thought of as the satisfaction of multiple constraints. The steady states that a recurrent network eventually stabilizes on are end-solutions that maximally satisfy all the constraints in the network, including between-node connections (e.g., MALE–AGGRESSIVE) and the input (e.g., facial cues, vocal cues, task demands, prejudice). As such, nodes in a recurrent network constrain each other in finding a best overall pattern that fits the input.

With respect to the present model, how the activation of a node changes over time is determined by three factors: the node's prior activation, how quickly this activation decays, and the net input of activation into the node from other nodes. It is assumed that excitation and inhibition summate algebraically, and that the influence of input on a node is dependent on the node's prior history of activation. It is also assumed that processing is stochastic rather than deterministic (see McClelland, 1991). On each iteration, therefore, the input to every node is altered by normally distributed random noise. Thus, the system's activation states are inherently probabilistic.

Before the presentation of each stimulus, activations of all nodes in the network are set equal to a resting activation value (zero), and external inputs are presented to certain nodes for processing. Processing occurs over a number of iterations. On each iteration, each node computes its net input from the nodes connected to it based on their latest activation. Specifically, the net input to node *i* is:

$$net_i = \sum_j w_{ij} o_j + ext_i + \varepsilon_\sigma$$

where w_{ij} is the connection weight to node *i* from node *j*, o_j is the greater of 0 and the activation of node *j*, ext_i is any external input to node *i*, and ε_σ is a small amount of normally distributed random noise with mean 0 and standard deviation σ. Once the net input into all nodes has been computed, the activation of node *i* is updated as:

$$\text{If } net_i > 0\text{:}$$

$$\Delta a_i = I(M - a_i)net_i - D(a_i - r).$$

$$\text{If } net_i \leq 0\text{:}$$

$$\Delta a_i = I(a_i - m)net_i - D(a_i - r),$$

such that M is the maximum activation, m is the minimum activation, r is the resting activation level, I is a constant that scales the influence of external inputs on a node, and D is a constant that scales a node's tendency to decay back to rest. Unless otherwise noted, in instantiations of the model the parameters are as follows: $M = 1$, $m = –0.2$, $r = 0$, $I = 0.4$, $D = 0.1$, and $\sigma = 0.01$. These are standard values used in connectionist networks of this type (McClelland, 1991; Rumelhart et al., 1986). Connection weights are specified for each instantiation of the model separately. In simulations, the network's ultimate response is given by the response alternative associated with the node with the largest activation in a pool after a given amount of iterations (once the network has stabilized).

It is assumed that the person construal system is organized into four interactive levels of processing: cue level, category level, stereotype level, and a higher-order level. Within each of these levels are one or several pools of nodes (Figure 1). Most nodes represent some feature or micro-hypothesis. For instance, the RACE pool would include a node for WHITE category and another node for BLACK category. Most of these pools are competitive in the sense that all the nodes are mutually exclusive and related by inhibitory connections. However, this is not necessarily the case for all pools. For instance, in the STEREOTYPES pool would be many nodes for different stereotypes. Some of these may inhibit one another (and thus be competitive), such as AGGRESSIVE and NICE, whereas others might have no relationship with one another, and some others might excite one another, such as AGGRESSIVE and DANGEROUS. Nodes that excite another node have a positively weighted connection, nodes that do not influence another node have no connection (zero weight), and nodes that inhibit another node have a negatively weighted connection.

Each node has a transient level of activation at every moment in time. This level of activation corresponds with the strength of a tentative interpretation or hypothesis that the node is represented in the input (e.g., a face). Thus, in situations where a face is presented, the activation level of the MALE category node could be said to represent, at every moment in time, the strength of the hypothesis that the face is male. A node whose activation level exceeds a threshold excites other nodes with which it has an excitatory connection and inhibits other nodes with which it has an inhibitory connection. Importantly, most of the connections in our model are bidirectional, producing feedback and making the network highly interactive.

Cue Level

The cue level contains a set of detectors for visual features (facial and bodily cues) and auditory features (vocal cues), which are directly stimulated by bottom-up sensory information of another person. The cue level contains two pools: a FACE/BODY CUES pool and a VOICE CUES pool. Sensory information of another person arriving in the visual system (facial and bodily cues) directly activates nodes in the FACE/BODY CUES pool. Sensory information arriving in the auditory system (vocal cues) directly activates nodes in the VOICE CUES pool. Depending on specific modeling interests, these pools have the flexibility to contain different arrangements of nodes. For instance, the FACE/BODY CUES pool could contain one node corresponding with all male facial features and another node corresponding with all female facial features. However, different strategies could be used. For instance, one node could describe a specific feature (e.g., LONG HAIR or DARK SKIN). Similarly, the VOICE CUES pool could contain a node corresponding with all male vocal features or it could contain a node corresponding with something specific such as FORMANT RATIO.

Nodes for cues that are along the same dimension (e.g., MALE CUES and FEMALE CUES) are related by mutually inhibitory connections because they compete for the same visual/auditory input. Thus, excitation of the MALE CUES node will inhibit the FEMALE CUES node, and vice-versa. Nodes that have no direct relationship with one another (e.g., LONG HAIR and DARK SKIN) have no connection between them. Cue nodes excite all category nodes consistent with them and inhibit all of those inconsistent with them. For instance, the cue node for male facial features would activate the MALE category node and inhibit the FEMALE category node. Similarly, the cue node for female facial features would activate the FEMALE category node and inhibit the MALE category node. Note that the connections between cue nodes and category nodes are bidirectional. Thus, cue nodes both influence and are influenced by category nodes. This produces feedback and a recurrent flow of activation, as discussed earlier.

Category Level

The category level contains a number of competitive pools that correspond with social category dimensions. For instance, in Figure 1, there are 4 pools: SEX, RACE, AGE, and EMOTION. Any number of different categories could be used, however (e.g., SOCIAL CLASS, SEXUAL ORIENTATION, OCCUPATION, ETHNICITY). These could include categories that are relatively static (e.g., sex) as well as categories that are dynamic (e.g., emotion). Each of these pools contain category nodes. The pool for SEX would include a MALE node and a FEMALE node; the pool for RACE would include, for example, a WHITE node, a BLACK node, and an ASIAN node. Nodes within a pool compete with one another through mutual inhibition. In the broad model depicted in Figure 1, bidirectional connections exist between all 4 of the category pools. This is not required for all instances of the model, but they are depicted because in some instances category nodes may be directly related to one another. For instance, if perceivers have learned in their lifetime that women tend to be happy and men tend to be angry (see Fabes & Martin, 1991), then the node for MALE (in the SEX pool) may have a bidirectional excitatory connection with ANGRY (in the EMOTION pool). Similarly, the node for FEMALE may have a bidirectional excitatory connection with HAPPY.

Category nodes receive input from cue nodes (which directly receive bottom-up sensory information) and they also send feedback to cue nodes. Category nodes activate stereotype nodes (e.g., MALE excites AGGRESSIVE and FEMALE excites DOCILE), and they also receive feedback from these nodes as well. Thus, not only will the category node, MALE, tend to activate the stereotype node, AGGRESSIVE, but activation of AGGRESSIVE will tend to activate the MALE category. Finally, category nodes may activate and be activated by higher-order nodes.

Stereotype Level

The stereotype level contains one pool including nodes for all category-related stereotypes (e.g., AGGRESSIVE or DOCILE). Within this, nodes could mutually inhibit or mutually excite one another. For instance, AGGRESSIVE and DANGEROUS would mutually excite one another, but AGGRESSIVE and DOCILE may mutually inhibit one another. Stereotype nodes receive input from category nodes and send feedback to them. Stereotype nodes also receive input from higher-order nodes and send feedback to them as well.

Higher-order Level

Nodes in this level may correspond with any number of high-level cognitive states, depending on what is being modeled. They could include factors such as prejudice, motivations, processing goals, task demands, among others. It is assumed that these nodes receive direct input from higher levels of mental processing (e.g., motivational systems or top-down attentional systems). Higher-order nodes may influence category nodes or stereotype nodes, or both. Moreover, they may have a bidirectional connection with these nodes or simply a unidirectional top-down connection only.

For instance, higher-order nodes could be used to model high-level task demands in a particular context. One higher-order node could denote SEX TASK DEMAND and another node could denote RACE TASK DEMAND. During a sex categorization task, the higher-order SEX TASK DEMAND node would be directly activated by higher-level input (e.g., top-down attentional systems, driven by memory of task instructions). Activation of this higher-order node would then have top-down excitatory connections with sex-related category nodes (MALE and FEMALE), but have top-down inhibitory connections with race-related category nodes (WHITE, BLACK, ASIAN), since the task demand compels attention to sex and away from race. As such, attentional effects due to task demands (e.g., placing attention on sex and away from race in a sex categorization task) emerge out of the flows of activation between these higher-order task-demand nodes and the category nodes, consistent with other computational models accounting for task demands (e.g., Cohen & Huston, 1994). This is one example of how the higher-order level could be used to model top-down effects from internal cognitive states, such as task demands, memory, affect, motivations, expectations, situational context, among others.

APPENDIX B

Connection weights in simulations of Study 4 (see Figure 5)

Connection	*Weight*
Category to Category inhibition	-1
Category to Cue excitation	.75
Category to Cue inhibition	-.25
Category to Higher-order excitation	.25
Category to Higher-order inhibition	-.25
Category to Stereotype excitation	.8
Category to Stereotype inhibition	-.3
Cue to Category excitation	.25
Cue to Category inhibition	-.1
Cue to Cue inhibition	-.1
Higher-order to Category excitation	.8
Higher-order to Category inhibition	-.3
Higher-order to Higher-order inhibition	-.5
Stereotype to Category excitation	.8
Stereotype to Category inhibition	-.8
Stereotype to Stereotype inhibition	-.3

APPENDIX C

Connection weights in simulations of Study 6 (see Figure 8)

Connection	*Weight*
Category to Category inhibition	-1
Category to Cue excitation	.5
Category to Cue inhibition	-.7
Category to Higher-order excitation	.33
Cue to Category excitation	.9
Cue to Category inhibition	-.4
Cue to Cue inhibition	-.8
Higher-order to Category excitation	.33