DEPARTMENT OF PSYCHOLOGY

Effects of allocation of attention on initial racial face discrimination

Johanna Lycke
Patrik Setterberg

Bachelor's thesis 2011

Supervisor: Mikael Johansson

Abstract

Utilizing a modified version of the affect misattribution procedure, event-related brain potentials of 32 Caucasian participants were examined in order to study the effects of allocation of attention on early racial face perception and discrimination. When attention was directed away from faces, towards ambiguous non-face stimuli, larger N1s, peaking at approximately 110 ms after stimulus onset, were observed for faces from participants' outgroup as compared to N1s for faces from participants' ingroup, peaking approximately 114 ms after stimulus onset. When attention was directed towards face stimuli no difference in N1 amplitude was observed. Conversely, N2s observed were larger for participants' ingroup when attention was directed towards faces as well as non-faces. Results suggest that allocation of attention appears to have an effect on the early time-course of racial face discrimination.

Keywords: Event-related potentials (ERP); race discrimination; face perception; attention allocation; affect misattribution procedure (AMP)

Table of Contents

Effects of allocation of attention on initial racial face discrimination

Making sense of reality is a surprisingly demanding endeavor. To ease our journey through the physical world our brain helps us by filing places, tasks, emotions and objects into categories. Very quickly, and more often than not outside of our awareness, a novel stimulus is perceived, analyzed, compared with earlier experiences of similar qualities, and subsequently placed in a fitting *category* for that type of stimulus. By no means bound by modality, categorization of visual information is vitally important to our daily lives. Further, we are as a species, with good reason particularly skillful at detecting and discerning between different members of our own kind, specifically their faces. In fact, the ventrally located fusiform gyri in the temporal lobes contain an area called the fusiform face area which is particularly devoted to this very task.

Quickly being able to identify members of one's family or group has intuitively adaptive benefits. Being social creatures, our repositories for different *people categories* quickly grow quite large. We have friends, family, co-workers, teachers, lovers, perhaps even adversaries. Socially or culturally linked groups of people form bonds, and, the more we can identify with the group, the stronger the bonds. This process of identification and subsequent grouping of individuals can be based on e.g. similar moral values, geographical proximity or similarity in appearance. Categorization of people inadvertently leads to the creation of stereotypes[1]. While often associated with prejudice, stereotypes in fact represent frequently occurring characteristics of a representative of a category, and thus facilitate processing of the stimulus in question. Prejudice[2] arises when negatively charged characteristics and preconceived notions intermix with a person or group's categorized information and influences their attitudes regarding another group of people. A readily available way to place a person in a category, perhaps unfortunately, is by the color of their skin, or, race. This, depending on one's point of view, unites or divides people: A sense of unity with one's *ingroup* (be it members of one's own gender, race, religion or

---

[1] The A.P.A. dictionary of psychology (2007) defines *stereotype* as "a set of cognitive generalizations (e.g., beliefs, expectations) about the qualities and characteristics of the members of a particular group or social category." (p. 893).

[2] The A.P.A. dictionary of psychology (2007) defines *prejudice* as "a negative attitude towards another person formed in advance of any experience with that person or group. Prejudices include an affective component […], a cognitive component […], and a behavioral component […]." (p. 723). In line with Breckler (1984), prejudice consists of these three components as it is, in fact, an attitude.

someone sharing the same hobby) sometimes breeds division towards one's *outgroup*. Effectively, members of one's outgroup can start being perceived as outsiders.

Abilities such as detection and identification of a possible threat must necessarily be tuned to ensure survival. Belonging to a group brings about a sense of security; being able to decide friend from foe in a very short amount of time is therefore of great importance. What then, is *a very short amount of time*? How quickly can our brain perform these extractions of categorized information of potential use? In order to answer these questions one must delve into the inner workings of the human brain. As we will see, categorization and category information extraction are remarkably speedy processes that require a number of cognitive functions to complete. Within milliseconds, our brains can execute an almost dizzyingly large number of commands. Research in the fields of neuroscience and neuropsychology has come a long way in mapping and exploring the different regions of the brain: their functions, interactions, influence on behavior and so forth. Yet there is much left to learn.

Stereotypes, while sometimes useful, do as previously noted give rise to prejudiced attitudes[3] and consequently prejudiced behavior. Research in social cognition concerning attitudes and prejudice increasingly relies on so called implicit measures and evaluations (a large body of literature exists regarding explicit and implicit cognition: see, e.g. Nosek (2005) for a distinction regarding evaluations). *Explicit evaluations* are conscious and controlled processes while *implicit evaluations* refers to automatic evaluations over which one cannot exert conscious control, correct and sometimes is not even aware of (Greenwald & Banaji, 1995; Nosek, 2005).

Indeed, a lot of research is done in this field: see, e.g. Fazio, Jackson, Dunton and Williams (1995), who examined racial attitudes as automatically activated memories; Brigham (1993), who investigated college students' racial attitudes; McConahay, Hardee and Batts (1981), who developed a measure of racism in America, and discussed the continuing conflict between ethnicities; Bartholow, Dickter and Sestir (2006), who examined electrophysiological correlates of stereotype activation under the influence of alcohol. Although research has provided new insight and a greater understanding of the problem we have reason to believe that currently available research is not without its flaws.

---

[3] An attitude is an evaluation of and a response to an object in a positive or negative way often described as consisting of three elements: Affect, cognition and behavior (Breckler, 1984).

It is our belief that much can be learned regarding attitude activation towards people if more research is conducted on early stages of perceptual face processing. Our point is brilliantly described by Ito and Urland (2003) who argue that:

> Although much has been learned about the ease with which stereotypes and prejudice are activated, less is known about the categorization process itself and the degree to which it occurs automatically. This may be attributable in part to the difficulty in measuring early attentional processes. (p.616)

Ito and Urland's (2003) experiments provided evidence for remarkably rapid face categorization. Indeed, we will discuss and expand on their reasoning in sections below. Another group of researchers (He, Johnson, Dovidio, & McCarthy, 2009), interested in the perception of faces of different races, specifically the neural correlates thereof, recently conducted a series of experiments utilizing a well-known implicit attitude test (*implicit association test*; Greenwald, McGhee, & Schwartz, 1998). In order to fully appreciate the results from the aforementioned studies and their implications for the present study, however, description of relevant experimental methodology and terminology is needed.

In short, this thesis will focus on automatic encoding of social category information (photographs of ingroup versus outgroup members, where ingroup refers to members of participants' own race and outgroup refers to individuals of other races) and effects of shifted allocation of attention. Neural correlates will be investigated in an attempt to identify the time-course of early face perception and the very first neural categorization of race utilizing a method commonly used in implicit attitude research that allows for manipulation of direction of attention. Additionally, we would like to propose amendments to aforementioned method in order to make it better suited for our purpose.

The complexity of social cognition and attitude research

Many cognitive processes are involved in making an evaluation of a stimulus. The mere perception and possible recognition of what is being presented, retrieval of relevant category information from memory about the object occurs before cognitive evaluation is initiated and a possible motor response (such as the pressing of a key in a laboratory setting) can be initiated. The great cognitive journey which takes place in our minds during discrimination and categorization certainly involves more processes than can be accounted for today. Ultimately, research on implicit attitudes wishes to gain an understanding of what cognitive processes are involved, and how they interact. Models in cognitive psychology can be elaborate and sophisticated but become far more interesting when they can be linked to corresponding neural activity, or, what actually goes on in the brain during testing. Behavioral experimental procedures used in attitude research, including those designed to measure implicit attitudes, run the risk of being figured out. This presents participants with an opportunity to modify their responses (e.g. conceal true attitudes), ultimately changing the outcome of the experiment. Neural activity on the other hand is something one cannot easily modify. To study the neural mechanisms underlying face perception, a prerequisite for the ability to categorize faces, is likely to be of importance for the understanding of stereotype activation and prejudiced behavior.

An advantageous approach, then, is to couple behavioral procedures with different brain imaging methods, providing researchers with additional data and insight into the relationships between neural activity and observed behavior. Functional magnetic resonance imaging (fMRI), positron emission tomography (PET) and electroencephalogram (EEG) are examples of commonly used brain imaging methods at present. While fMRI and PET have a high spatial resolution compared to EEG, the latter is superior to the other two in its high temporal resolution (Fabiani, Gratton, & Coles, 2000). Since face perception involves rapidly occurring cognitive processes, a high temporal resolution is crucial for discovering differences in neural activity. Thus, for the present study EEG was the most appropriate brain imaging method for us to utilize.

Electrophysiological recordings & event-related brain potentials

Using an electroencephalograph, one can measure changes in rhythmic oscillations in voltage (post-synaptic changes in activity, synchronized over a network of neurons). As the EEG has a high temporal resolution, accurate measures can be made of electrical potentials associated with specific events in time, so-called *event-related brain potentials*, or *ERPs*. Extracted from recorded EEG from a large number of trials, ERPs are attained by averaging signals from multiple electrodes (individually or collectively) at particular points in time (repeated occurrences of events, or *epochs*), providing information regarding latency as well as topography of electrical activity. In addition to waveform latency and polarity, the amplitude of the potential is of critical interest and is chiefly where differences between conditions are observed. The following section will focus on how different ERP patterns can be identified, named, and how and what they can teach us about corresponding, underlying cognitive processes.

ERP components

ERPs can be divided into a multitude of components, typically named based on polarity (positive or negative), temporal features as well as topography, or location on the scalp where they are recorded. Thus, a negative-going waveform trough peaking around 200 ms following e.g. the onset of a presented stimulus is usually called *N200*. Similarly, a positive-going waveform peak occurring around 300 ms following stimulus onset is usually called *P300*. Additionally, components can be identified based on their ordinal position in the waveform signature: The first occurring positive peak can be called *P1*, the second negative trough *N2* and so forth.

Different components signify sometimes very different underlying cognitive processes depending on where on the scalp they are located, e.g. whether they are produced in frontal (anterior) or occipital (posterior) cortical regions. Cognitive functions (or compounds thereof) can also produce differing ERP components at different sites, as is the case with the VPP / N170 component. The vertex positive potential (VPP; a positive component peaking roughly at around 170 ms after stimulus onset, and most strongly around the vertex area), also called N170 when topographically located temporally and negative in polarity, is a component that is believed to reflect an early stage in the encoding process where faces are differentiated from non-faces (Ito, Thompson, & Cacioppo, 2004).

Differentiation of race has been observed extremely early in processing. In the study by Ito and Urland (2003), White participants viewed photographs of White as well as Black individuals and were asked to categorize the persons in the photographs by race. The experimenters found significantly larger negativity peaking at 122 ms (N1, or N100) post-stimulus presentation for photographs of Black faces compared to photographs of White faces. Crucially, this increased negativity for Black faces persisted in a follow up study where photographs had been de-saturated, or converted to gray-scale. People can, it seems, categorize individuals by race remarkably rapidly, even in the absence of color as a cue.

The N2 component, sometimes referred to as the N200 (Crites, Berntson, Cacioppo, & Gardner, 1995) as it peaks roughly between 200 and 250 ms after stimulus onset holds information that is both interesting and complex. Ito et al. (2004) found greater (negative going) responses for White faces compared to Black faces, more negative at frontal and central regions than more posterior regions. This, according to the authors, indicates a difference in ERP response as a function of group. The effect described above has more recently been replicated by He et al. (2009). The researchers attribute these findings (extended same-race attention) to expertise in processing faces of the race to which a participant belongs.

Face perception and the role of attention

Having briefly discussed ERPs and their value in studying processing of perceptual information, we can better appreciate the significance of the abovementioned findings of Ito and Urland (2003) and He et al. (2009). The former duo's participants were instructed to categorize photographs of people along the dimensions of race and gender in two separate experiments. They found evidence of racial categorization occurring as early as 122 ms following face onset, while gender categorization was delayed about 50 ms. Explicit instructions to categorize individuals by race, however, obviously requires participants to consciously allocate attention to racially distinguishing facial features. As the authors agree, examining early attentional processes is not unproblematic. In an attempt to examine race processing and the effects of shifted attentional focus, the researchers conducted a subsequent series of experiments in which participants were required to process facial information at a deeper level, making personality judgments of people (whether stimulus-individuals liked or disliked certain vegetables & whether they were introverted or extroverted). They found that Black faces elicited larger N1s than did White faces, again very early: around 120 ms, even when attention was focused away from social category (Ito & Urland, 2005).

He et al. (2009) similarly studied neural correlates of racial face perception. Interestingly, however, they wished to investigate whether and how ERP results correlated with scores on a test evaluating implicit attitudes, the implicit association test (IAT; Greenwald et al., 1998). Although the present study did not utilize the IAT, a brief description of its procedure will serve to explain our reasons for choosing a different implicit attitude test.

IAT participants are asked to categorize stimuli that appear on a screen by pressing one of two buttons. There are four categories – presented in pairs – of stimuli with each button assigned to two categories. Typically in a racial IAT, as utilized by He et al., photographs of Black faces, White faces, words positive in valence and words negative in valence make up the four categories mentioned above. ”The core idea underlying the IAT”, say De Houwer, Teige-Mocigemba, Spruyt and Moors (2009, p.347) ”is that categorization performance should be a function of the degree to which categories that are assigned to the same key are associated in memory.” Consequently, good categorization performance for a certain combination of stimuli, as measured by response speed, should indicate that those two categories of stimuli are more closely associated in memory than are combinations requiring a longer response time. For example, longer response time for a photograph of a Black person and a positive word as compared to a photograph of a White person and a positive word would, according to the IAT, indicate a more positive attitude towards the White person.

He et al. found Black vs. White N2 amplitude difference to correlate significantly with IAT scores, where a high score signifies ingroup (White) preference, at scalp sites FC4 (fronto-central, right hemisphere) and T6 (posterior region of the right temporal area). Participants with a low IAT score did not produce differing N2 amplitudes between Black and White faces whereas participants with a high IAT score produced larger N2s for White faces. These results are profoundly intriguing as they, if replicable; offer a means by which to predict prejudiced attitudes based on neural activity[4].

Ito and Urland (2003, 2005) and He et al. (2009) all managed to study early perceptive processes and racial categorization and, using different approaches, manipulate direction of

---

[4] The later occurring and, thus aptly named, late positive potentials (LPPs) are often studied in attitude research. LPPs are positive ERP deflections that include distinct components such as the P300 (believed to vary as a function of categorization; e.g. Crites, Berntson, Cacioppo, & Gardner, 1995) as well as less distinct yet readily observed increased positivity up to 600 ms after stimulus onset. Larger LPP activity has been observed (chiefly in parietal regions) when participants view people compared to non-people, and when participants view stimuli negative in valence compared to stimuli positive in valence, even when participants are not explicitly instructed to categorize the stimuli based on valence. This people and negativity bias has been suggested to indicate a stronger allocation of processing resources with adaptive benefits, i.e. identifying possible threats in a person's environment, and that this is an automatic response (Ito & Cacioppo, 2000). Moreover, Crites et al. (1995) found that when participants were instructed to report their attitudes towards target trait words, the LPP was larger for negative traits than for neutral traits, and larger still when compared to positive traits. Interestingly, this effect persisted even when participants were instructed to intentionally misreport their attitudes, leading the researchers to suggest that the study of LPPs can be useful when trying to identify attitudes that people would rather not report.

attention. One thing they all have in common, however, is that participants in all tasks are explicitly instructed to attend pictures of faces (albeit in combination with words in the IAT) and then perform the categorization. It would be very interesting to examine whether the processing of race remained similarly expeditious if participants' attention was intentionally diverted from faces, i.e. if the faces were task-irrelevant. As previously hinted at, another test designed to measure implicit attitudes exists which, we believe, can be utilized to achieve this diversion of attention manipulation.

Affect misattribution procedure

In 2005, intrigued by the concept of misattribution, Payne, Cheng, Govorun and Stewart developed an elegant method of attaining implicit attitudes. Their operationalization of the term misattribution reads "mistaking an effect of one source for the effect of another" (p.278) and this encapsulates the core logic behind their affect misattribution procedure (AMP).

The AMP is performed in two blocks referred to as the indirect version (implicit measure) and the direct version (explicit measure). In both versions, participants are presented to stimuli in the same order and for the same amount of time: a target stimulus, i.e. a photograph of a face, followed by an ambiguous stimulus, i.e. a Chinese pictograph. Participants are then asked to rate the *pleasantness* of one of the stimuli and try to ignore the presence of the other stimuli. In the direct version, participants are asked to rate the photograph and in the indirect version, participants are asked to rate the pictographs; thus manipulating whether or not the face-stimuli are task-relevant.

Swift presentation of a target stimulus followed by an ambiguous stimulus appears to have an effect on the affective evaluation of the ambiguous stimulus; i.e. participants misattribute their affect towards the target stimulus (the photograph), in their evaluation of the ambiguous stimulus (the Chinese pictograph). Avoiding deeper contemplation of presented stimuli, i.e. making sure priming has occurred, is of critical importance in the AMP, which is why stimuli are presented rapidly (typically 100 ms) and the response window following stimulus presentation is typically set to 1500 ms. Responses recorded outside this response window are usually excluded from analysis (Payne, Burkley, & Stokes, 2008).

When studying implicit attitudes, the measure of interest is often response-latency, as measured from the time of stimulus presentation onset to a given response (e.g. the pressing of a button) is recorded. Affective priming tasks (developed by Fazio, Sanbonmatsu, Powell, & Kardes, 1986; see also e.g. Fazio et al., 1995, for an example of affective priming tasks used as a measure of racial attitudes), while in many ways dissimilar to the IAT, also rely on response-latency to determine attitudes (Payne et al., 2008, remark that response-latency tasks, in fact, “make up the bulk of implicit measurement”, p.17). Behavioral models and methods, however well-controlled, refined and useful, do possess inherent limitations. When solely focusing on response-latency, a lot of important information regarding processes of potential interest preceding the final motor response could be lost (White, Crites, Taylor, & Corral, 2009).

Another possible risk is that the participants figure out the critical significance of the response speed and therefore consciously regulate their response time. This could potentially, and seriously, confound IAT results. In fact, it calls into question the validity of the IAT; specifically, the very *implicitness* of the implicit association test. If participants are able to understand the structure of the test, then can it really be said to measure something implicit at all? Thus, an advantage of using the AMP is that the implicit measure is based on evaluation (and not response speed) where intentional delaying of responses by participants becomes meaningless.

Any implicit measure becomes meaningful only once it is held in contrast to a corresponding explicit measure. Typically, this is achieved by complementing an IAT or affective priming task with one or more self-report questionnaires – for studies of racial attitudes e.g. the Modern Racist Scale (MRS; McConahay et al., 1981; McConahay, 1983) or the Attitudes Towards Blacks Scale (ATB; Brigham, 1993). However, results from studies of explicit and implicit attitudes often correlate only weakly. Measuring explicit and implicit attitudes in differing settings and using differing methods is a cause for concern, and could, at least in part, account for the weak correlations between explicit and implicit attitudes often found. This problem has been addressed by e.g. Payne et al. (2008) who refer to it as a problem of structural fit and hypothesized that the distinction between implicit and explicit cognition could be exaggerated. For example, there is a significant difference in test structure between a self-report questionnaire and an IAT, which relies on response-latency. Generally, when conducting experiments, the aim is to keep conditions as similar as possible to exclude the possibility of any extraneous variables affecting the dependent variable, thus ensuring that recorded differences over trials or conditions are due to the manipulation of the independent variable or variables studied.

The AMP allows researchers to direct participants' attention between focusing on a target stimulus (face) and an ambiguous stimulus (pictograph). Since one of the aims of this thesis is to investigate the effects of diverted attention in early stages of race categorization, AMP seems an appropriate method to utilize. A key strength of the AMP is that it embodies an implicit version as well as an explicit version, and that these two versions of the test are structurally very similar, making it well suited to combine with ERP-recordings. The only difference between the versions in the experiment, and subsequent explicit and implicit measures – is the instructions regarding what to rate, thus elegantly avoiding possible influence of extraneous variables.

Since ERP-recordings are very sensitive, especially to muscular activity, reducing impact of possible confounds is of great importance. The procedure of an AMP is quite simple and monotonous, and the participant is kept motionless except for the pressing of keys, keeping possible artefacts as effects of muscular activity low.

Modification of AMP. If people's perception of rapidly presented stimuli is at all dependent upon factors such as hue, chromaticity or luminance, a problem emerges in implicit ethnic prejudice research since a difference in these aspects will always exist between ethnic groups. If, for example, strength of luminance can have an influence on people's perception of a stimulus, this phenomenon could make it more difficult to ensure that early racial face discrimination is due to social features. Since one of the aims of this thesis was to identify the time-course of early racial face perception and not early perception of luminance and/or hue we had to find a solution, or workaround to this problem.

Whereas one trial of the racial AMP traditionally consists of photographs representing participants' ingroup or outgroup preceding an ambiguous stimulus (a Chinese pictograph), a third type of stimulus material was implemented to precede the ambiguous stimulus in half of the trials. This was done to achieve a baseline condition for comparison that could make it possible to deduct possible influence from hue, chromaticity or luminance. Scrambled versions of the photographs, containing the same amounts of colored pixels as the photographs, only arranged in such a way that they appeared as seemingly randomly colored horizontal and vertical lines (see

2nd box in Figure 1) were thus included. We will refer to these scrambled photographs as the control stimuli: ingroup and outgroup controls respectively.

Using this approach, essentially all the visual information present in the photographs is retained. However, when presented with a control stimulus, participants are unable to distinguish any socially significant features representing either the participants’ ingroup or outgroup or indeed any facial features at all. The recorded ERPs during presentation of controls can be compared with recorded ERPs during presentation of faces to ensure that a possible difference in neural activation between groups will not be due to luminance or hue.

Summary

Our aim with this thesis was to examine the effects of allocation of attention on early racial face perception and race discrimination. While participants were presented to photographs of ingroup and outgroup members as well as scrambled versions of the same photographs, event-related potentials were examined to assess racial face categorization processes on-line. As Ito et al. (2003, 2005) found evidence for categorization of race around 120 ms following presentation of photographs of faces, we expected the present study to produce congruous results (larger outgroup N1s within a similar time window). Further, since photographs of faces contain socially distinguishable features and are more complex than scrambled ones we expected these to elicit larger N1 amplitudes. As He et al. (2009) and Ito et al. (2004) found greater N2 amplitude for ingroup faces compared to outgroup faces we expected the present study to replicate these findings. The stimuli were presented within the framework of a modified AMP structure. Participants’ allocation of attention shifted as a function of version of the AMP which made it possible to investigate whether directed attention toward non-faces can have an influence on racial face perception and discrimination. Adding the scrambled photographs to the AMP allowed for control of luminance and/or hue, and ensured that possible observed differences in the ERP waveform could be attributed to socially important facial features. Additionally, in order to test the validity of the AMP, we wished to replicate previous AMP results. Particularly, we expected outgroup members to be rated as less pleasant than ingroup members in the direct as well as in the indirect task (Payne et al., 2005; Payne et al., 2008). Social desirability, or a wish to appear unprejudiced in the eyes of others should, if the AMP proves a valid test of implicit and explicit attitudes, lead to more pleasant ratings of outgroup members in the direct version as compared to in the indirect version. In conclusion, our hypotheses were as follows:

Hypothesis 1a: N1 amplitude for outgroup faces will be larger than the N1 amplitude for ingroup faces.

Hypothesis 1b: N1 amplitude for faces will be larger than N1 amplitude for controls.

Hypothesis 1c: Allocation of attention to the face-stimulus will affect early race discrimination reflected in the N1 component.

Hypothesis 2: N2 amplitude for ingroup faces will be larger than the N2 amplitude for outgroup faces.

Hypothesis 3a: Ingroup members will be rated as more pleasant than outgroup members.

Hypothesis 3b: Outgroup members will be rated as more pleasant in the direct version than in the indirect version.

Method

Participants

32 Caucasian volunteers in Lund, primarily students at Lund University (13 male, 19 female; M_{age} = 22.72 years, *SD* = 2.63, range = 18-28 years) were healthy and had normal or corrected to normal vision. All were right-handed; this was a requirement since some left-handed individuals can be less lateralized (specific cognitive functions such as language can be diffused across both hemispheres). Six additional participants were excluded from analysis due to deficient behavioral or ERP data. Participants were randomly assigned to one of four groups: Two main groups (direct evaluation followed by indirect evaluation and indirect evaluation followed by direct evaluation, with a brief pause with new instructions in between versions) were further divided into two sub-groups each. The order in which the stimuli were presented to the participants was reversed in the two sub-groups respectively in order to avoid possible order effects. Initial analysis in the somewhat similar study of attitudes by Crites et al. (1995) found no effects involving gender. For the present study we thus, placing our trust in the validity of Crites et al.'s findings, collapsed results across gender.

Material

For the AMP, 40 photographs of faces representing participants' ingroup (i.e. Caucasian; 20 male, 20 female) and 40 photographs of faces representing participants' outgroup (i.e. other skin colors, e.g. Blacks, Hispanics, Asians, individuals of middle eastern descent; 20 male, 20 female) were used (*N*=80). The individuals in the photographs showed similar facial expressions, were cut out, showing only the face, and on a white background (Face-Place database; http://www.tarrlab.org). A thin, black, square frame surrounded the photographs. Additionally, scrambled versions of the same photographs (discussed above) were included (*N*=80). Chinese pictographs were randomly selected and were presented clearly in black on a white background (*N*=160) with a thin, black, square frame. Stimulus size measured 93x93 mm on screen. The experiment was programmed and shown using E-Prime 2.0 software.

Electroencephalographic recording equipment

Electroencephalographic activity was continuously recorded using the Neuroscan Scan 4.5 Acquire software. 38 Ag/AgCl electrodes installed in a Compumedics Quick Cap were placed according to the extended 10-20 system. EEG data was amplified using a Compumedics NuAmp amplifier, digitized with 32 bit resolution at a sampling rate of 500-Hz. Impedance was kept below 5.7 kΩ for all electrodes included in the analysis. All electrodes were referenced online to the left mastoid electrode, and offline to the average of the left and right mastoids. Additional off-line processing included a digital high pass filter set to .01 Hz (6 dB cut-off) as well as a digital low pass filter set to 15 Hz (12 dB cut-off). Eye movements, blinking and brief tension in facial musculature often produce artefacts in the EEG. To counteract this, horizontal as well as vertical electro-oculographic activity was recorded (HEOG & VEOG respectively). Horizontal electrodes were placed at participants' outer canthi, while vertical electrodes slightly above as well as below participants' left eye (roughly 2 cm). Facial recording sites were cleaned with alcohol and abraded with peeling prior to electrode application.

Procedure

Upon arriving at the laboratory, participants read and signed an informed consent form and were explained procedures for electroencephalographic recording. A laminated hand-out was given to participants, depicting visually the order in which stimuli would be presented on screen for one trial. During electrode application participants were offered modest amounts of marshmallows, grapes and non-caffeinated soft drinks.

After each participant was prepared for EEG recording he or she was placed in front of a computer in a brightly lit room (approximately 60 cm from a 17” CRT-monitor; the experiment was run at 1024x768 resolution at 85 Hz) and was presented to instructions given on the screen. Participants were instructed that the experiment was about making quick judgments while avoiding being distracted. The experiment consisted of two randomly ordered versions, where they were to rate the pleasantness of the stimuli corresponding to the current version while simply ignoring the presence of the other set of stimuli (photographs/scrambled photographs and Chinese pictographs respectively). The instructions emphasized that there were no right or wrong answers, and that the measure of interest was the participant’s spontaneous affective reaction to the stimuli. Participants were encouraged to relax, sit as still as possible and try to blink exclusively when a fixation cross was visible on screen. After the participant had explained the procedure of the experiment to the experimenters – a measure taken to ensure that he or she had understood the instructions – the experiment could begin.

A total of 160 trials were performed (direct evaluation, *N*=80 and indirect evaluation, *N*=80, counterbalanced). Participants were presented to the stimulus material: a modified version of the classic AMP design which included our scrambled photographs. First, a white screen with a fixation-cross was shown (2000 ms). After the fixation-cross and a 200 ms blank screen a photograph or a scrambled photograph was shown for 100 ms. Following a 100 ms blank screen, a Chinese pictograph was shown for 100 ms. Finally, a mask (black and white noise) appeared on the screen for 1500 ms at which time the participants were to evaluate the pleasantness of the stimuli and respond to the photograph or scrambled photograph in the direct version and the pictograph in the indirect version. Responses given at times other than when the mask was present were excluded from analysis. The total time for one trial was 4000 ms. Immediately following the response window (mask), the program automatically initiated the next trial (for a visual depiction of one trial, see Figure 1).

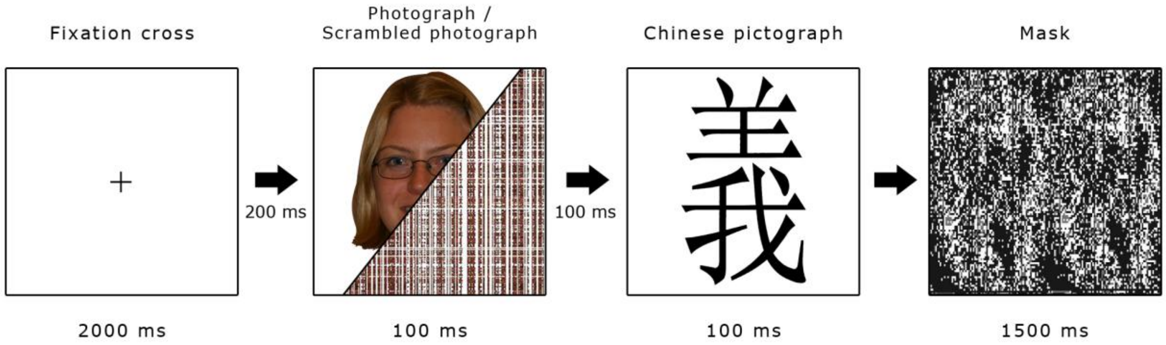

Figure 1: Order of items presented on screen for one trial. The second box illustrates that either a photograph or a scrambled photograph was presented. Numbers below items indicate the duration the corresponding items were present on screen; numbers below arrows indicate blank screen duration between item presentations.

Ratings were made using a four item scale on a computer keyboard with four keys highlighted with stickers labeled with the numbers 1 through 4 on them where 1 = *not at all pleasant*, 2 = *not pleasant*, 3 = *somewhat pleasant* and 4 = *very pleasant* (the Swedish word “tilltalande”, which translates into “pleasant” or “appealing” was used). Stickers were placed on the keys *D*, *F*, *J* and *K*, where *D* = 1 and ascending for half the participants and *D* = 4 and descending for the other half of the participants. This reversing of order served to counterbalance for people’s possible preferences for pressing a certain key over others and possible increased response speed due to their right-handedness.

Having completed the AMP, participants were then asked to fill out a form regarding demographics, habits of caffeine, nicotine and alcohol use (heavy usage, or being under influence during testing can affect electroencephalographic activity). The experiment was now over. The experimenters thanked the participants for their participation, who were offered a shower and were subsequently debriefed. Each session lasted for circa 90 minutes.

Results

ERP results

Results will be related to corresponding hypotheses in order. The section begins with a presentation of the results obtained when analyzing the ERP data and ends with presenting the behavioral data. The ERP-curves measured from electrode site Cz (located in the vertex area) was used in the analysis. The selection of this recording site was based on previous literature as well as an inspection of the ERP waveforms and components of interest.

Hypothesis 1

In order to test our hypotheses 1a, b & c, we investigated the 100-150 ms time window by conducting a three-way 2 (Task: direct, indirect) × 2 (Type: face, control) × 2 (Group: ingroup, outgroup) repeated measures ANOVA. A significant main effect of Group was discovered. Outgroup N1 amplitude was more negative (M = -1.12 µV, SD = 3.18) than ingroup amplitude (M = -.5 µV, SD = 3.29), F(1,31) = 5.19, p = .03. Analysis also revealed a significant main effect of Type, where the negativity was larger for faces (M = -1.61 µV, SD = 3.81) compared to controls (M = -.004 µV, SD = 2.86), F(1,31) = 13.82, p = .001. Ingroup faces' N1 peaked (reached its largest amplitude) at 114 ms after stimulus onset in both tasks. Outgroup faces' N1 peaked at 112 ms in the direct task and at 110 ms in the indirect task.

Hypothesis 1a. We hypothesized that N1 amplitude for outgroup faces would be larger than the N1 amplitude for ingroup faces. The main effect of Group provides evidence for this claim. Our data suggests that racial discrimination of faces occurs between approximately 110-114 ms after stimulus presentation.

Hypothesis 1b. Our hypothesis regarding increased N1 amplitude for faces compared to N1 amplitude for controls is supported by the main effect of Type. Faces elicit a larger N1 than do controls within the same time window.

Hypothesis 1c. We believed the larger N1 amplitude for outgroup faces compared to ingroup faces to be affected by allocation of attention. Examining the ERP waveform, the difference in negativity between ingroup and outgroup faces appeared more prominent in the indirect version of the test (Figure 2, top; dashed lines). To study the differences between groups in each task we conducted a paired samples t-test for direct task and indirect task respectively. In the indirect task N1 amplitude was significantly larger for outgroup faces (M = -2.50 µV, SD = 4.41) than for ingroup faces (M = -1.04 µV, SD = 3.78), t(31) = 2.869, p = .007, while no significant effect of Group was present in the direct task. These results are in line with our hypothesis in that a very early discrimination of race is dependent upon allocation of attention.

Our data suggests that discrimination of race occurs after 110 ms and directed attention toward faces does not appear to be required.

Hypothesis 2

Upon examining the ERP waveform, Task appeared to affect N2 amplitude (Figure 2, top). To verify this possible effect, and our second hypothesis: N2 amplitude for ingroup faces will be larger than the N2 amplitude for outgroup faces, we conducted a two-way 2 (Task: direct, indirect) x 2 (Group: ingroup, outgroup) repeated measures ANOVA on ERP data in the 200-300 ms time-window. Analysis revealed one main effect of Task, $F(1,31) = 55.20$, $p < .000$, where N2 amplitude was larger in the indirect task ($M = -7.91$ µV, $SD = 3.94$) compared to in the direct task ($M = -4.57$ µV, $SD = 4.40$). It also revealed a main effect of Group, $F(1,31) = 26.90$, $p < .000$, where N2 amplitude was larger for ingroup ($M = -7.17$ µV, $SD = 4.25$) than for outgroup ($M = -5.31$ µV, $SD = 3.96$). These results are in line with our hypothesis, as ingroup amplitude was more negative than outgroup amplitude in both tasks. Ingroup faces' N2 peaked at 232 ms in the direct task and 226 ms in the indirect task. Outgroup faces' N2 peaked at 234 ms in both tasks.

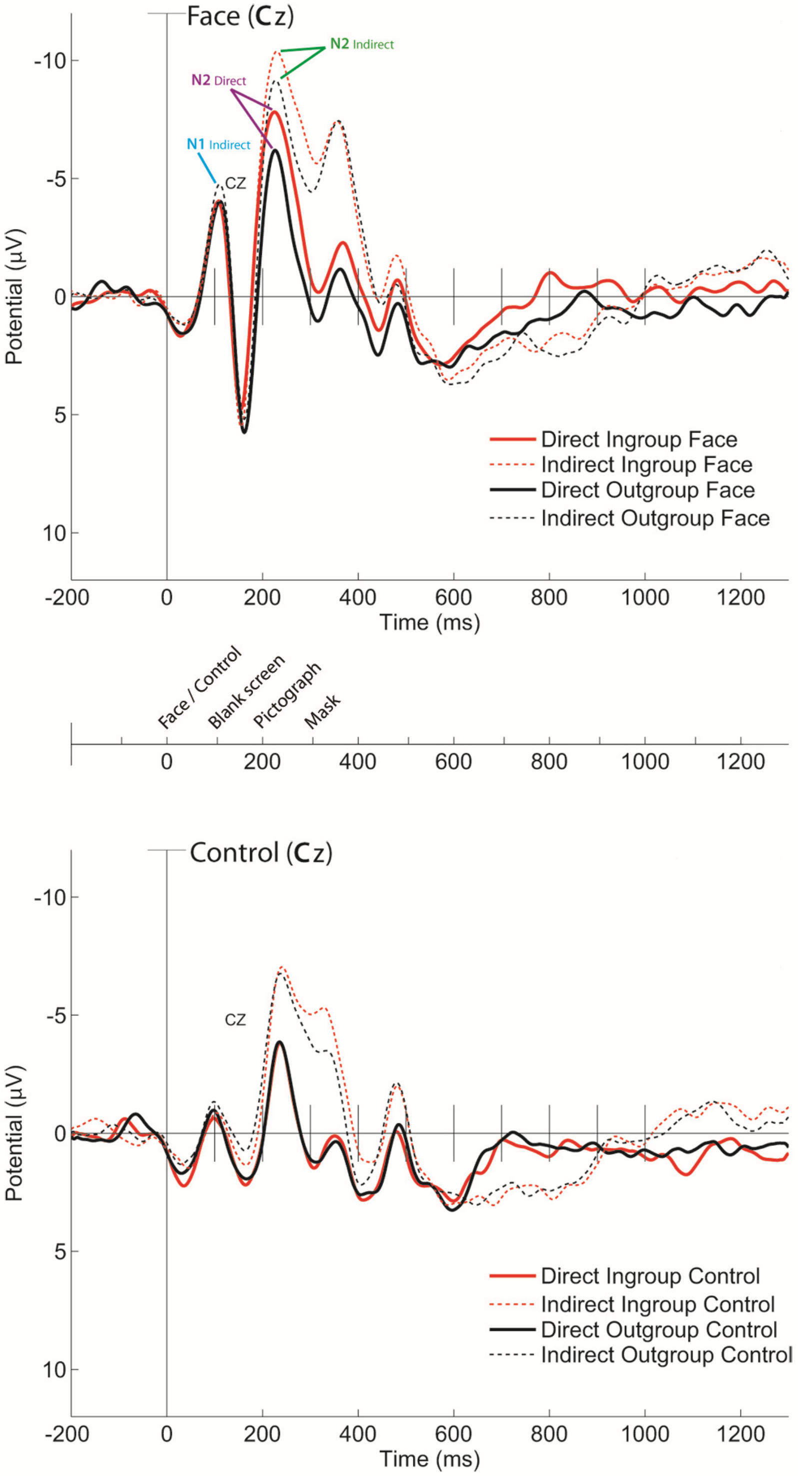

Figure 2: Waveforms at Cz electrode for faces (top) and scrambled photographs, or controls (bottom). For details and significance for *N1 Indirect*, *N2 Direct* and *N2 Indirect*, see Figure 3A, 3B & 3C respectively. The timeline between waveforms indicate times at which stimuli, blank screen (interstimulus interval) and mask appear on screen.

A: N1 Indirect (Ingroup - Outgroup)

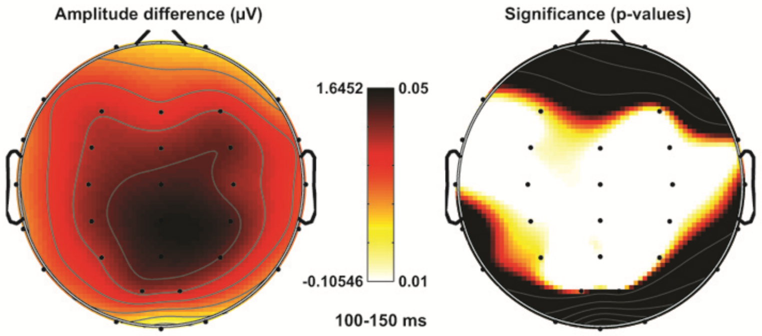

B: N2 Direct (Ingroup - Outgroup)

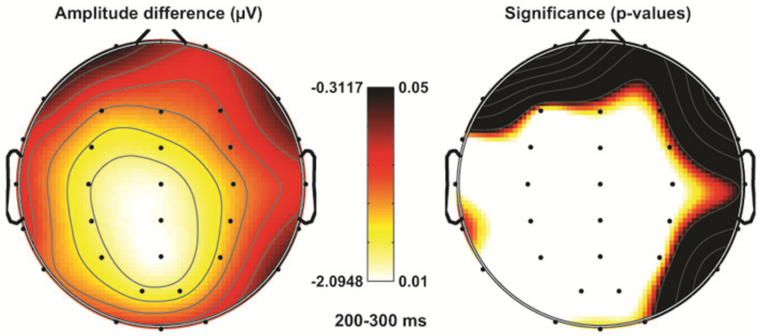

C: N2 Indirect (Ingroup - Outgroup)

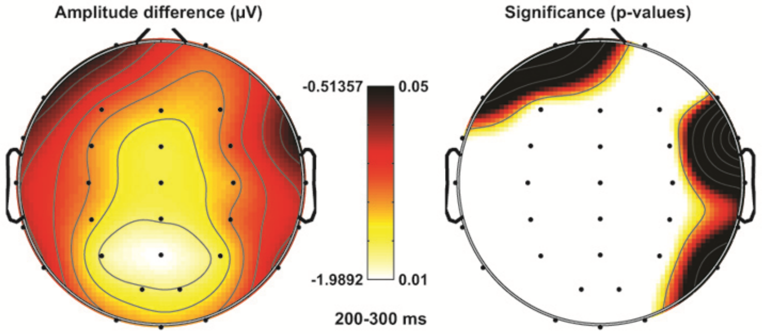

Figure 3: Amplitude difference and significance between ingroup and outgroup at electrode Cz for Faces, A) N1 Indirect, B) N2 Direct & C) N2 Indirect.

Behavioral data

Ratings of pleasantness were transformed into probabilities in the following manner. The positive ratings (somewhat pleasant & very pleasant) were collapsed and divided by total ratings for that condition and participant. Probabilities of positive responses were averaged across participants for all eight conditions. In order to get an index of probabilities of positive ratings for each group unclouded by extraneous variables such as luminance or hue, ratings of controls (for ingroup and outgroup respectively) were subtracted from corresponding ratings of faces. The following analyses will be based on differences between faces and corresponding controls. ‘Ratings’ will refer to *probabilities of positive ratings*.

Hypothesis 3

Main effects and interaction effects were analyzed using a two-way 2 (Task: direct, indirect) × 2 (Group: ingroup, outgroup) repeated measures ANOVA. Analysis revealed a main effect of Group: Outgroup ratings ($M = .05$, $SD = .13$) were overall higher than ingroup ratings ($M = -.02$, $SD = .17$), $F(1, 31) = 6.59$, $p = .015$.

Hypothesis 3a. No support was found for our hypothesis, that ingroup members would be rated as more pleasant than outgroup members. We found a main effect of Group, however contradictory to our hypothesis, outgroup members were rated as more pleasant than ingroup members.

Hypothesis 3b. We hypothesized that outgroup members would be rated as more pleasant in the direct version compared to in the indirect version. The results were not in line with this hypothesis as no interaction was found between Group and Task.

Discussion

The aim of the present study was to investigate the role of attention in initial racial face discrimination using ERPs and a modified version of the AMP. The study produced evidence suggesting that direction of attention plays a role in racial face discrimination. Differences between groups were observed in the N1 component as well as in the N2 component. Previous AMP results reported by Payne et al. (2005) and Payne et al. (2009) were not replicated.

ERP results

N1. Previous studies (Ito & Urland, 2003, 2005) have provided neurophysiological evidence for racial discrimination of faces after approximately 122 ms when participants are explicitly instructed to categorize stimuli by race. In the present study, however, discrimination of race was never task-relevant. When instructed to rate pleasantness of faces, N1 amplitude did not differ across groups. However, instructing participants to allocate attention towards an ambiguous stimulus and ignore photographs of faces, as in our indirect task, resulted in outgroup faces eliciting a larger N1 than did ingroup faces. In conjunction with the absence of a difference in amplitude in the direct task, we argue, these results could indicate that very early decoding of social features is dependent upon the direction of attention. Put differently, the larger outgroup N1 found in the indirect version only suggests that very early race discrimination can occur without consciously attending face-stimuli. To the best of our knowledge, this has not previously been reported.

In addition, outgroup N1 peaked approximately 110 ms after face onset in the indirect task. While an approximate difference of 12 ms is minute at best, it appears as though diverting attention from faces altogether could increase temporal sensitivity to racially distinguishing facial information. It is entirely possible that the difference in N1 latency between Ito and Urland (2003, 2005) and the present study is negligible. Additional studies would have to be performed in order to confirm this claim. A study examining early racial face discrimination including one condition in which participants are instructed to ignore faces (as in the present study) and one condition in which participants are instructed to categorize faces by race (as in Ito & Urland, 2003) would likely produce interesting results.

Analysis of data yielded a main effect for Type which indicates that the N1 component is sensitive to faces. Analysis also revealed a main effect for Group. These two main effects in conjunction suggest that racial discrimination is not affected by hue and luminance. Hence, our modification of the AMP, aimed at ensuring that early racial discrimination be attributable to social features appears to have fulfilled its purpose.

N2. He et al. (2009) reported a larger N2 for ingroup faces as compared to outgroup faces. The present study replicated these findings. Our study, unlike He et al. (2009), included measurements under two separate tasks where allocation of attention was shifted between faces and non-faces. We found a difference in amplitude between tasks. N2s were larger for both groups in the indirect task compared to the amplitude in the direct task (Figure 2, top). It appears as though allocation of attention has an impact on subsequent processing of the faces as well. This is also an interesting finding that, as far as we know, has not previously been reported. It would be interesting to further investigate the underlying processes causing this observed difference in amplitude between tasks. That, however, is unfortunately beyond the scope of this thesis.

Behavioral results

The behavioral results obtained from our modified AMP were not consistent with our hypotheses. We were initially struck by the absence of a preference for any of the groups. As the population of Lund is highly multicultural, we did not expect our participants to be particularly prejudiced. However, because AMP is a known and widely used test of implicit attitudes and the logic of misattribution is inviting, we expected results that at least resembled those previously reported by e.g. Payne et al. (2005) and Payne et al. (2008). Several explanations, not all mutually exclusive, could account for these discrepant findings.

The most convenient approach is to conclude that the AMP is a valid tool for obtaining implicit attitudes. This approach requires that our sample consisted of wholly unprejudiced participants.

Another line of thought rests on the assumption that the AMP is flawed in that it can be – and incidentally was – figured out by some or most participants. Were this the case, it would not be far-fetched to assume that our participants sought to appear unprejudiced (towards outgroup members) and consequently intentionally misreported their attitudes to the stimuli, thus distorting results. This would indicate that the AMP suffers from the same problem as the IAT, of a possibility of being figured out. In hindsight, it is possible that the underlying structure of the AMP could be even easier to figure out than the structure of the IAT. In the AMP participants are instructed to rate the pleasantness of faces, i.e. explicitly report their preferences regarding people’s appearances. It is not implausible to think that participants can understand that some kind of comparison will be made based on their ratings.

General discussion

While the present study was unable to replicate earlier AMP results, that was not our main purpose. Based on obtained results, we are reluctant to assess its validity as a measure of implicit attitudes. The structural similarity between its two versions allowed us to direct participants' attention away from faces while still studying perception thereof, which was of critical importance for the principal purpose of the thesis. The observed difference in N1 amplitude as a function of shifted allocation of attention is intriguing, and it is our hope that the results produced by the present study will be of use to future researchers seeking to investigate early mechanisms involved in processing of perceptual information of people and discrimination of facial features. It is our firm belief that research on these early processes is of great importance in order to gain better understanding of people categorization, stereotype formation and ultimately prejudiced attitudes.

Problems with the present study

Reliability of acquired data. Steps were taken to create an optimal experimental environment. However, future research would possibly profit from acknowledging and amending the following shortcomings of the present study. An easily implementable measure to create cleaner EEG-data as well as more reliable behavioral evaluations from the AMP would be to include a short set of practice trials at the start of the experiment consisting of unrelated stimuli to help participants *get a feel* for the real experiment. Additionally, systematic post-test probing for suspicion of prejudice being part of the test could validate the behavioral results and exclude possible influence of social desirability.

Accuracy of electrophysiological recordings. While all mobile phones and similar possible sources of radio & electromagnetic signals were turned off during the experiment, the laboratory was unfortunately not electromagnetically shielded and weak signals from adjacent rooms might possibly have interfered with recordings, however slightly. Furthermore, the laboratory was not satisfactorily sound-attenuated. Consequently, background noise, however quiet, might have distracted participants, subconsciously diverting attention from the task at hand, affecting electrophysiological as well as behavioral results.

Implications for future research

Our demonstration of larger N2 amplitude for both groups in the indirect task compared to the direct task is enthralling in by itself since it has not previously been reported. To further investigate the underlying processes causing this difference in amplitude between tasks can be of great importance for a deeper understanding of early racial face perception.

As mentioned, the AMP appears to work well in combination with ERP recordings. By performing simple changes (such as replacing photographs) in the existing AMP structure one can create opportunities to study a variety of phenomena involved in the early time-course of perception. Using a different set of stimuli would allow for testing of perception and discrimination along a multitude of dimensions other than race.

The control condition utilized in the present study appears to have corrected for influence from hue or luminance. Even so, it would be intriguing to investigate whether the differences in N1 as well as N2 amplitudes would persist (or differ) with a sample consisting of individuals with different ethnic backgrounds, e.g. Blacks, Hispanics or Asians. Investigating possible effects of attention allocation later in the time-course of racial face perception (e.g. late positive potentials, LPPs) would also be of great interest.

Finally, should the AMP gain further support as being a valid and reliable tool for attaining implicit measures of attitudes, correlating AMP results with neural activity (similar to He et al., 2009) could potentially unlock doors, revealing key pieces to the understanding of stereotyping and prejudice.

Conclusions

Results indicate that discrimination of race occurs very early in processing of facial information. The present study provides evidence suggesting that shifted allocation of attention away from faces has an effect on N1 as well as N2 amplitude; specifically, attending and evaluating ambiguous stimuli (and trying to ignore faces) increases N1 amplitude for outgroup faces, and N2 amplitude for ingroup faces as well as outgroup faces.

The answer to the question stated in the introduction of this thesis (”what, then, is *a very short amount of time?*”) might well therefore be *somewhere around 110 milliseconds.*

References

Bartholow, B. D., Dickter, C. L., & Sestir, M. A. (2006). Stereotype activation and control of race bias: Cognitive control of inhibition and its impairment by alcohol. *Journal of Personality and Social Psychology*, *90,* 272-287.

Breckler, S. J. (1984). Empirical validation of affect, behavior, and cognition as distinct components of attitude. *Journal of Personality and Social Psychology*, *47*, 1191-1205.

Brigham, J. C. (1993). College students' racial attitudes. *Journal of Applied Social Psychology*, *23*, 1933-1967.

Crites, S. L., Berntson, G. G., Cacioppo, J. T., & Gardner, W. L. (1995). Bioelectrical echoes from evaluative categorization: II. A late positive brain potential that varies as a function of attitude registration rather than attitude report. *Journal of Personality and Social Psychology*, *68*, 997-1013.

De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, *135,* 347-368.

Fabiani, M., Gratton, G., & Coles, M. G. H. (2000). Event-related brain potentials: Methods, theory, and applications. In J. Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds), *Handbook of psychophysiology* (2nd ed.) (pp. 53-76). Cambridge: Cambridge University Press.

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,* 1013-1027.

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

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem and stereotypes. *Psychological Review*, *102*, 4-27.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464-1480.

He, Y., Johnson, M. K., Dovidio, J. F., & McCarthy, G. (2009). The relation between race-related implicit associations and scalp-recorded neural activity evoked by faces from different races. *Social Neuroscience*, *4*, 426-442.

Ito, T. A., & Cacioppo, J. T. (2000). Electrophysiological evidence of implicit and explicit categorization processes. *Journal of Experimental Social Psychology*, *36*, 660-676.

Ito, T. A., Cacioppo, J. T., & Thompson, E. (2004). Tracking the timecourse of social perception: the effects of racial cues on event-related brain potentials. *Personality and Social Psychology Bulletin*, *30*, 1267-1280.

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, & Behavioral Neuroscience*, *5*, 21-36.

McConahay, J. B. (1983). Modern racism and modern discrimination: The effects of race, racial attitudes, and context on simulated hiring decisions. *Personality and Social Psychology Bulletin*, *9*, 551-558.

McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined in America? It depends on who is asking and what is asked. *Journal of Conflict Resolutions*, *25*, 563-579.

Monteith, M. J., Voils, C. I., & Ashburn-Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial bias. *Social Cognition*, *19*, 395-417.

Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology*, *134*, 565-584.

Payne, B. K., Burkley, M. A., & Stokes, M. B. (2008). Why do implicit and explicit attitude tests diverge? The role of structural fit. *Journal of Personality and Social Psychology*, *94*, 16-31.

Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, *89*, 277-293.

Prejudice. (2007). In G. R. VandenBos (Ed.), *A.P.A. dictionary of psychology* (p. 723, 1st ed., vol. 2). Washington, DC: American Psychological Association.

Stereotype. (2007). In G. R. VandenBos (Ed.), *A.P.A. dictionary of psychology* (p. 893, 1st ed., vol. 2). Washington, DC: American Psychological Association.

White, K. R., Crites, S. L., Taylor, J. H., & Corral, G. (2009). Wait, what? Assessing stereotype incongruities using the N400 ERP component. *Social Cognitive and Affective Neuroscience, 4*, 191-198.

Acknowledgements

We would like to extend our deepest thanks to our supervisor Mikael Johansson, Professor, department of psychology at Lund University for his invaluable guidance and support, and without whose help this thesis would not have been possible. Mikael has taught us a great deal about psychology, specifically the field of neuroscience and has been a constant source of inspiration. Additionally, we greatly appreciate assistance received from Jean-Christophe Rohner, Associate Professor, department of psychology at Lund University, who is responsible for the majority of our knowledge in experimental methodology. We would also like to thank Robin Hellerstedt, PhD student at Lund University for wholly without obligation patiently teaching us the procedures involved with, and logic behind EEG-recordings.