

Priming and Habituation for Faces: Individual Differences and Inversion Effects

Cory A. Rieth and David E. Huber
University of California, San Diego

Immediate repetition priming for faces was examined across a range of prime durations in a threshold identification task. Similar to word repetition priming results, short duration face primes produced positive priming whereas long duration face primes eliminated or reversed this effect. A habituation model of such priming effects predicted that the speed of identification should relate to the prime duration needed to achieve negative priming. We used face priming to test this prediction in two ways. First, we examined the relationship between priming effects and individual differences in the target duration needed for threshold performance. Second, we compared priming of upright and inverted faces. As predicted, the transition from positive to negative priming as a function of prime duration occurred more slowly for inverted faces and for individuals with longer threshold target durations. Additional experiments ruled out alternative explanations.

Keywords: face perception, short-term priming, individual differences, habituation

Searching for a face in a crowd is a seemingly easy task. However, to complete this task, we must identify the subtle differences between different faces while generalizing face recognition across variations in image size, orientation, hairstyle, lighting, and other more salient perceptual dimensions. Because a crowd is in constant motion, and because we generally wish to find a desired face quickly, not much time is afforded to consider each face. Nonetheless, some minimum duration is needed to integrate sufficient information for identification. The constant integration of perceptual input over time implies the possibility that we may incorrectly blend together one face and the next as our eyes scan across positions, or as the crowd moves. However, we do not suffer greatly from such source confusion between successive faces. To explain this and other results involving the effects of successively presented stimuli, we suggest that the perceptual system includes a discounting mechanism that appropriately reduces the response to a previously presented face, and this reduction serves to offset the effect of source confusion.

Huber and O'Reilly (2003) proposed that neural habituation is the basic mechanism behind temporal discounting, which automatically parses the stream of perceptual events. This theory was developed to explain priming effects with words, and the model correctly predicted a variety of lexical effects as well as electrophysiological recordings when viewing words. Because habituation is a general mechanism, similar effects should exist with many

other tasks and stimuli. To explore the generality of this account, we examined face identification because it is a type of perceptual expertise that is thought to differ from word reading in a variety of ways. Beyond generalization to a new stimulus class, face identification allowed us to test a key prediction of this account. Because identification is the driving force behind habituation, situations that allow faster identification should also produce faster habituation and thus faster prime discounting. In the reported experiments we tested this prediction in two ways. First, we compared individuals who could identify a target face with short exposure durations to those who needed a longer duration. Second, we compared upright and inverted faces because inverted faces are more difficult to recognize. For both of these tests, we examined immediate face repetition priming to assess temporal discounting.

To motivate Experiment 1, which examined individual differences in face habituation as indicated by priming, we next (1) present the a theory of temporal segregation through habituation; (2) review previous work with word priming based on average performance; (3) present a reanalysis of individual differences with word priming, which failed to confirm the predicted individual differences; and (4) review previous word on habituation with faces.

Temporal Segregation Through Habituation

The integration of information over time while avoiding temporal source confusion is a general problem in perception and these effects are likely to exist beyond the domains of face perception and word priming. Our perceptual system effortlessly separates the constant stream of input into discrete objects and events. Nonetheless, this temporal segregation is not trivial, requiring neither too little nor too much integration in calculating both where and *when* one object begins and another ends. Temporal integration allows the perceptual system to appropriately combine different features contained within a single object, such as

Cory A. Rieth and David E. Huber, Department of Psychology, University of California, San Diego.

This research was supported by NIMH Grants MH63993-01, MH063993-04, and by NSF IGERT Grant DGE-0333451 to GW Cottrell/VR de Sa. The authors would like to thank Jim Tanaka for his generosity in allowing for the use of his real face stimuli.

Correspondence concerning this article should be addressed to Cory A. Rieth, Department of Psychology, 0109, 9500 Gilman Drive, University of California, San Diego, La Jolla, CA 92093-0109. E-mail: crieth@ucsd.edu

when an object is partially blocked by a moving occluder. However, too much integration can inappropriately blend perception of a previous object with a subsequent object (e.g., combining Ernie's eyes with Bert's nose). This problem of temporal segregation cannot be solved with fixed discrete timing intervals considering that each stimulus in each situation may necessitate a different time scale for proper temporal segregation. Instead, we suggest that there is a constant interplay between the process of integrating the current percept and the process of discounting percepts that were previously identified.

Mechanisms that reduce inappropriate perceptual blending while maintaining appropriate perceptual binding have been proposed both at the cognitive level, such as in Treisman's feature integration theory (Treisman & Gelade, 1980) or Biederman's GEON theory (Hummel & Biederman, 1992), as well as at the neural level, such as with Singer's synchrony theory (Singer, 1999) or Grossberg's adaptive resonance theory (Carpenter & Grossberg, 2003). To bridge these levels of description, we have addressed the problem of temporal segregation by proposing interrelated models in terms of optimal Bayesian 'explaining away' (Huber, Shiffrin, Lyle, & Ruys, 2001; Huber, 2008a) as well as with a specific neural account in the form of habituation through synaptic depression (Huber & O'Reilly, 2003). In line with David Marr's levels of explanation (Marr, 1982), we seek to find a mapping between the computational "why" and the implementational "how" of perceptual segregation through discounting.

For word priming, perceptual discounting through neural habituation has been examined both with behavioral measures as well as with neural measures. Behaviorally, repetition word priming was found to produce a gradual transition from positive priming (more accurate identification of primed targets) to negative priming¹ (less accurate identification of primed targets) as a function of increasing prime duration (Huber, 2008b; Huber, Shiffrin, Lyle, & Quach, 2002; Weidemann, Huber, & Shiffrin, 2005; Weidemann, Huber, & Shiffrin, 2008). In terms of neural behavior, repetition priming produced Event-Related Potential (ERP) and evoked Magnetoencephalography (MEG) responses in early perceptual processing that were modulated by increasing prime duration (Huber, Tian, Curran, O'Reilly, & Woroch, 2008). Furthermore, the time course of these neural effects was predicted by the habituation account as previously fit to behavioral data.

In keeping with the proposal that dynamic discounting is a general mechanism of perception, habituation through synaptic depression (Abbott, Varela, Sen, & Nelson, 1997; Tsodyks & Markram, 1997), is found in the majority of pyramidal cells in many cortical areas (Thomson & West, 1993). Synaptic depression refers to the finding that postsynaptic depolarization (i.e., the message delivered to the receiving cell) is rapidly weakened as a function of recent presynaptic action potentials (i.e., the number of messages recently sent), with this loss of responsiveness lasting anywhere from hundreds of milliseconds to several seconds. This finding is explained by hypothesizing the existence of resources that enable effective signaling across the synapse (e.g., presynaptic neurotransmitter availability, although many other mechanisms produce similar effects), with these resources becoming depleted as a function of recent activity. Therefore, synaptic depression is a natural candidate for automatic temporal segregation because it is the ongoing activation that drives depletion of synaptic resources. Thus, habituation does not occur until after activation is sufficient

both in magnitude and duration (i.e., habituation follows on the heels of identification).

This hypothesis predicts a specific relationship between the speed of identification and the speed of habituation. For instance, individuals who identify a class of object more quickly (i.e., individuals who have greater perceptual expertise) should produce priming data indicative of greater habituation. As reported below in a reanalysis of Experiment 1 from Huber (2008b), this prediction was not confirmed for threshold identification of words. Before presenting this reanalysis, we review results using this particular priming paradigm with words.

Habituation Effects With Immediate Word Priming

Immediate word priming has been used for more than 30 years to measure various aspects of language (e.g., Meyer & Schvaneveldt, 1971). However, in our work, we use word priming to investigate the basic mechanisms behind perceptual identification and lexical retrieval—mechanisms that may well generalize to other stimuli and other tasks. Priming is traditionally measured with speed of lexical decision, speed of naming, or accuracy in the form of threshold identification. For threshold identification, the task is usually to name the briefly flashed word (e.g., Meyer, Schvaneveldt, & Ruddy, 1975), although forced-choice testing provides a variant of threshold identification that allows separate measurement of the costs and benefits of priming (e.g., Ratcliff & McKoon, 1997) by comparing conditions that prime the correct answer (the target primed condition) versus conditions that prime the incorrect answer (the foil primed condition). Huber et al. (2001) adopted this paradigm for the study of immediate priming effects as a function of the duration and degree of processing for primes (see Figure 1 for the equivalent paradigm using faces). With this paradigm, short duration primes produced benefits for the target primed condition and deficits for the foil primed condition (i.e., a preference to choose the repeated choice alternative, or 'positive priming'). In contrast, long duration primes produced the opposite pattern for repetition or orthographic priming, with deficits for the target primed condition and benefits for the foil primed condition (i.e., a preference against the repeated choice alternative, or 'negative priming').

Huber et al. (2001) explained both the positive and negative priming results with a Bayesian model based upon probabilistic feature activation. In the model, source confusion between the prime and target results in positive priming, but negative priming occurs through the discounting of prime features. The appropriate level of discounting is calculated by lowering the evidence for features known to have been primed. Too little discounting of prime-activated features produces positive priming (i.e., these features lend extra evidence in favor of the primed choice) and too much discounting of features activated by primes produces negative priming (i.e., discounted features result in a relative loss of evidence as compared to no priming). This Bayesian model describes these results in terms of too little or too much discounting,

¹ We use the phrase 'negative priming' to refer to a situation in which performance is worse for primed trials as compared to unprimed trials, and this should not be confused with the specific paradigm of *negative priming* (Tipper, 1985), in which priming deficits are found as a function of a prime's status as a to-be-ignored distractor.

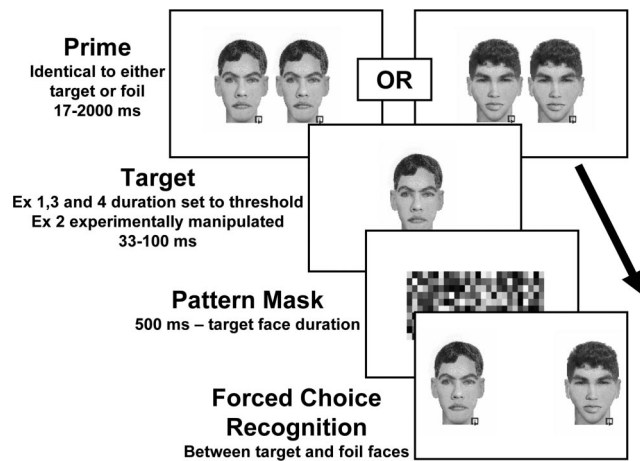


Figure 1. Presentation sequence used in all experiments. Side by side replicated images of either the foil or the target face appeared as the prime display, followed by a single image of the target face. All experiments manipulated prime duration, using 17, 50, 150, 400, and 2,000 ms (Experiments 1, 3, and 4) or 17, 150, and 2,000 ms (Experiment 2). Target duration in Experiments 1, 3, and 4 was set for each participant to a duration that produced 75% accuracy. In Experiment 2, 33, 50, and 100 ms target durations were tested for all participants. Target presentation was followed by a checkerboard pattern mask and then two choice faces (the target and foil) that remained on the screen until a response was made. Choices were presented at the top and bottom of the screen in Experiments 1, 3, and 4 or to the sides of the screen as pictured here in Experiment 2. All experiments discouraged strategic use of the primes by informing participants that the prime was just as likely to indicate the incorrect answer as the correct answer and by providing trial by trial accuracy feedback.

but it does not specify why prime duration results in this change. Therefore, Huber and O'Reilly (2003) developed a model of dynamic discounting through neural habituation in a perceptual cascade of rate-coded neurons (see O'Reilly & Munakata, 2000, for a description of such rate-coded neurons). Huber (2008b) tested this account by parametrically varying prime duration, revealing a gradual nonlinear transition from positive to negative priming. Beyond the form of this transition, other predictions of the habituation account were confirmed for the relationship between priming and forward masking (habituation predicted a relationship between the onset of negative priming and the elimination of forward masking) as well as the difference between repetition priming (deficits because of orthographic habituation) and associative priming (benefits because of lexical-semantic top-down expectations).

In the Huber and O'Reilly (2003) model, the transition from positive to negative priming was captured through synaptic depression. For most neurons, if a presynaptic cell is constantly driven, a receiving postsynaptic cell initially shows a large depolarization, but after a short time (e.g., hundreds of milliseconds), postsynaptic depolarization drops off sizably (Abbott et al., 1997; Tsodyks & Markram, 1997). Including synaptic depression allowed the model to capture the interaction between priming condition (e.g., target primed versus foil primed) and prime duration. For short duration primes, the prime presentation is too brief to cause much depletion of synaptic resources and so lingering activation

from the prime results in enhanced processing of the primed choice word. Longer duration primes cause a build up of synaptic depression for the prime representation, which more than offsets the advantage of lingering activation. In other words, the prime is still active, but because synaptic resources are depleted, the relative advantage of this activation is lost. This depletion results in sluggish responding for the primed representation both in terms of its presentation as a briefly flashed target as well as its presentation as a choice alternative. Not only can this account explain choice accuracy, but it also correctly predicts reaction times; there is a speed up for chosen alternatives (correct response in the target primed condition and error responses in the foil primed condition) after brief primes but a slow down for chosen alternatives after long-duration primes (Huber & Cousineau, 2004).

An Analysis of Individual Differences in Word Priming

Assuming that different individuals differ in their word identification expertise and that the target duration needed for threshold performance is a good indication of these differences, temporal segregation through habituation predicted that individuals with lower threshold durations should reveal faster/stronger habituation as determined by their priming data. As seen in Figure 2, which is a median split of the priming data from Huber (2008b; Experiment 1), this prediction was not confirmed. The median split was based on the target duration needed to achieve 75% accuracy in a condition where the prime was unrelated to the choices (i.e., threshold target duration without priming). The group on the left, labeled the 'fast group', is the 50% of participants with shorter threshold target durations while the group on the right, the 'slow group', was the upper 50%. Even though these groups differed in the duration needed to identify a word, both groups show similar habituation as evidence by the crossover point between the target primed and foil primed conditions. For both the fast and the slow groups of participants, 50 ms of priming produced the maximal positive priming and both groups of participants produced sizable negative priming for the 400 and 2,000 ms prime durations.

This reanalysis failed to support the predicted relationship between the speed of identification and the onset of negative priming. However, it may be that individual differences in the target duration needed for threshold word identification are not an accurate measure of the speed of perceptual word identification. This may be the case if these target duration differences are because of low level visual processing differences (e.g., visual acuity) or perhaps motivational differences (e.g., staying alert on all trials). Alternatively, the observed individual differences might reflect differences in verbal processing (e.g., the ability to maintain words in verbal working memory) rather than differences in perceptual word identification. Regardless of the reason behind this failure, we sought to reexamine the situation with face priming. By using a face identification task, we manipulated the speed of identification within subjects by comparing upright and inverted faces. In addition, by using images of unfamiliar faces, the task was more clearly perceptual—unlike word identification, individual differences in verbal working memory should not matter because there is no associated verbal label for an unfamiliar face. Because the reported experiments concern the speed of habituation with faces, we next review the literature on face habituation.

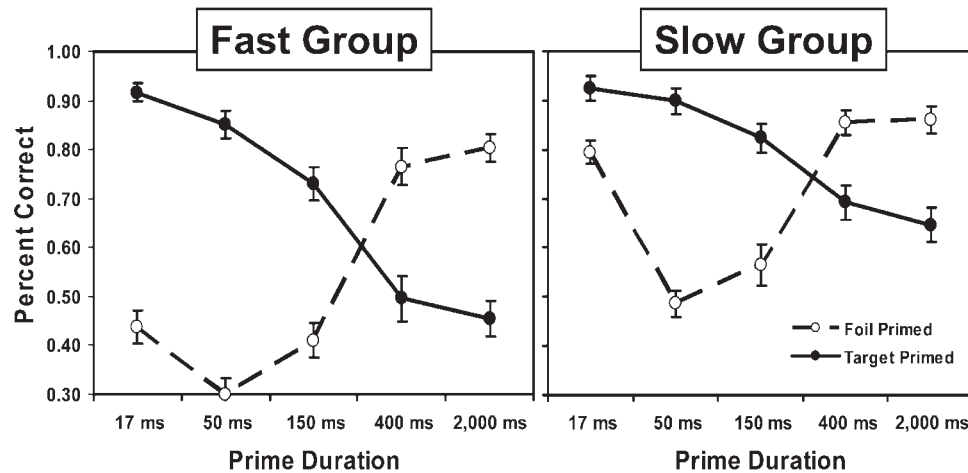


Figure 2. Analysis of the word priming data from Huber (2008b) Experiment 1 by splitting participants into fast and slow groups. Error bars are ± 1 SE.

Habituation Effects With Immediate Face Priming

There is a sizable literature examining habituation² aftereffects with faces, although these experiments do not typically use repetition priming or test different durations of initial exposure. One exception is a study by Bennett, Lleras, Oriet, and Enns (2007) that examined identification for the emotion of a face. In the basic condition there was a priming benefit when the prime emotion matched the target emotion, but presentation of a neutral expression between the prime and target reversed this effect. Our habituation account might make sense of this result if the lingering activation of the prime, which is the source of positive priming, fades during presentation of the intervening neutral face. Habituation operates on slower time scale, and so an intervening stimulus might serve to eliminate lingering activation but not eliminate lingering habituation, which is the source of negative priming.

Demonstrating that these face aftereffects are seen with manipulations of prime-target similarity, other studies find that extended viewing of a face can cause shifts in the categorical boundary along a series of morphed faces (Rothstein, Henson, Treves, Driver, & Dolan, 2005), criterial shifts in a high dimensional 'face space' (Leopold, O'Toole, & Blanz, 2001), and distortions of face configuration (Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003) or emotion expression (Fox & Barton, 2007). In light of these results and others, we take the existence of face habituation as a given. However, instead of testing a single duration that produces strong habituation, we seek to examine the dynamic properties of face habituation by using multiple prime durations.

The Current Investigation

Experiment 1 was a replication of Huber's (2008b) Experiment 1, replacing words with upright views of computer generated faces. In light of the observed individual differences with faces, Experiment 2 tested whether target duration manipulations could produce these differences. Experiment 3 replicated the individual differences using images of real faces and included a baseline

priming condition to rule out a specific strategic interpretation of the results. Besides using faces to test for greater variability in perceptual expertise, another advantage of using faces is that they are viewed in a conical upright orientation more so than words and other objects, and, correspondingly, faces reveal preferentially large inversion effects as compared to other objects (Yin, 1969). With faces, a given individual can be turned from a face expert into a face novice simply by inverting the face. Therefore, Experiment 4 replicated Experiment 1, but also introduced the manipulation of face inversion, revealing an apparent interaction between face inversion and individual differences. Finally, the neural habituation account was quantitatively applied to the data of Experiment 4, which also allowed us to explore possible accounts of the interaction between individual differences and face inversion.

Experiment 1

Experiment 1 replicated the immediate repetition word priming found in Experiment 1 of Huber (2008b) except that upright faces were used rather than words. A range of prime durations were tested with both target primed conditions (the benefits of priming), in which the target repeated the prime face, as well as with foil primed conditions (the costs of priming), in which the incorrect choice repeated the prime face. It was predicted that individuals who needed less time to identify the briefly flashed target would also reveal more rapid face habituation. Time to identify a face was measured by the threshold target duration needed to achieve 75% accuracy. Habituation was measured by the prime duration needed to produce negative priming (foil primed performance better than target primed performance).

² This literature often refers to these effects as face adaptation rather than habituation. However, because there may be a tendency to confuse the term adaptation with long-term learning and memory, and because adaptation does not necessarily imply a negative effect, we instead adopt the term habituation.

Method

Participants. Twenty-eight undergraduate psychology students participated in this experiment for course credit.

Procedure. As seen in Figure 1, the basic task was to identify the briefly flashed target faces that appeared in the center of the screen, which was immediately replaced with a pattern mask. All faces were novel computer generated male faces and no face appeared on more than one trial during the experiment. Immediately before the target face, a prime face appeared for durations ranging from 17 ms (i.e., subliminal face priming) to 2,000 ms (i.e., excessive face priming). Because these experiments investigated immediate repetition priming, it was not possible to present the prime in the exact same location as the target. In other words, a spatial cue was needed to separate primes from targets. Therefore, two identical versions of the prime face were displayed side by side. In this manner, the prime face also served as a forward mask of the target face. Face identification was tested by means of a forced-choice between the target face and a foil face. Performance was not speeded and accuracy feedback was provided on every trial. Participants were explicitly informed that half the time the prime was identical to the correct answer and half the time the prime was identical to the incorrect answer, and that there was no effective response strategy that used knowledge of the prime alone. Thus, the most effective strategy was to focus on the center and wait for the briefly flashed target face.

The experiment began by establishing the target duration each participant needed to achieve accuracy at the threshold value of 75%. Once this target duration was established, it remained fixed for the remainder of the experimental session. In the analyses of individual differences, a split half grouping of these target durations defined slow and fast face identification groups of participants.

Two variables were manipulated within subject: prime duration with five levels (17, 50, 150, 400, or 2,000 ms) and prime type with two levels (target primed or foil primed). There were 32 trials for each condition occurring randomly throughout the experiment. Each trial began with a fixation cross presented for a duration such that the combined duration of the fixation cross and the prime face was 2,500 ms. After the target face was presented, it was immediately masked with a checkerboard pattern mask. The mask was presented for a length of time such that the duration of mask and target combined was 500 ms, thus maintaining a constant duration between onset of the target and onset of choice alternatives. Following presentation of the mask, the target and foil faces of the forced choice decision appeared above and below the center line, separated by ~ 6 degrees of visual angle. Responses were collected by keyboard. The choice faces were presented until a decision was made, after which feedback was presented for 1,500 ms. The position of the target face was randomly counterbalanced across top versus bottom (i.e., in every condition, the correct answer was the top face for half the trials).

During the initial block of trials that determined the appropriate target duration for threshold performance, primes that were different than both the target and the foil were used with a prime duration of 150 ms. In other words, these were 'neither-primed' trials. Unlike Huber's (2008b) Experiment 1, which used a staircase procedure to obtain threshold target duration, this experiment mapped out the psychometric function by testing specific target

durations of 33, 50, 67, and 83 ms equally often as randomly mixed across trials. If 75% identification was not achieved with the 83 ms target, a target duration of 100 ms was adopted for the remainder of the experiment. Otherwise, the target duration that produced performance closest to 75% was selected.

At the beginning of the experiment, participants completed 12 practice trials, with the first 6 using a target duration of 167 ms (which was sufficiently easy for all participants) and the second 6 using a target duration of 83 ms to accommodate participants to threshold presentations. Participants then completed a block of 80 threshold determination trials (20 at each target duration, randomized), followed by four blocks of 80 priming trials (32 at each prime duration for each prime type, randomized). Before the third priming block, participants completed a 10-min word search puzzle to give them a rest. The experiment was self paced and lasted approximately 60 min on average.

Materials. There were 1,000 face images randomly generated using the FACES computer program that creates faces from composites of features including head shape, eyes, nose, mouth, ears, eye brows, chin shape, and facial hair. This was achieved by repeated use of the software's random function followed by elimination of faces that did not appear plausible or sufficiently male. Facial hair and different hair styles were allowed. All features were sampled randomly with replacement by the program, and thus the faces had varying degrees of similarity. All face images were grayscale and 150×137 pixels. As viewed each face was ~ 5 degrees visual angle in height and 3.5 degrees in width. The pattern mask presented after the target face was 150×140 pixels, consisting of a 5×5 grid of blocks. On each trial, a random pattern mask was created by sampling each block of the grid separately from 256 levels of gray.

Equipment. The experiment was conducted on CRT monitors with a 60 HZ vertical refresh rate. Presentation times were synchronized with the display refresh rate. All items were displayed on a gray background and the display resolution was set to 640×480 pixels.

Results and Discussion

Collapsing across individuals, a repeated measures ANOVA revealed a main effect of prime type, $F(1, 27) = 21.46$, $p < .001$, prime duration, $F(4, 108) = 15.33$, $p < .001$, and a prime type \times prime duration interaction, $F(4, 108) = 14.51$, $p < .001$. There was no difference between the target primed and foil primed conditions either at 400 ms, $t(27) = -1.35$, $p = .18$, or at the longest prime duration of 2,000 ms, $t(27) = 0.81$, $p = .43$. However, for the 50 and 150 ms prime durations, there was sizable positive priming, with target primed performance better than foil primed performance (for both tests, $t(27) < -6.9$, $p < .001$). This was also true of the shortest prime duration, however the magnitude of difference was smaller, $t(27) = -2.68$, $p = .013$. Thus, the interaction was a pattern of increasing then decreasing positive priming. This qualitative pattern was the same as previously found with immediate repetition priming of words. However, contrary to words, longer prime durations only served to eliminate positive priming, rather than produce negative priming. Next, individual differences are analyzed based on threshold target duration.

Results from the threshold determination block of trials are presented in Appendix A, which shows results collapsed across all

participants, as well as for the median split according to the target duration required to achieve threshold performance. These median split groupings will hereafter be referred to as the fast group (shorter target durations to achieve threshold performance) versus slow group (longer target durations to achieve threshold performance), although this is not meant to imply that there are literally two distinct groups of people.³ Instead, this is merely a convenient way to parse the data while maintaining sufficient numbers of trials. This median split was created in the following manner for each experiment. Participants were first rank ordered by the target duration needed for threshold performance. Because there are only a few possible target durations because of the refresh rate of the monitors, the median target duration did not cleanly divide the participants into two equal sized groups. Therefore, individuals with the median target duration were further divided according to their average accuracy during the threshold block of trials. This provided a fine grained measure so that these median target duration participants could be assigned to the fast group or the slow group to produce two equal sized groups. If the number of participants was odd, the median participant was included in fast group. For Experiment 1, each group consisted of 14 participants. The average target duration thresholds for the fast and slow groups were 46.43 ms ($SE = 3.11$) and 83.33 ms ($SE = 3.91$), respectively.

Figure 3 shows the effects of prime duration and prime type separately for the fast and slow median split groups of participants. Collapsing across conditions, accuracy was .70 ($SE = .03$) and .77 ($SE = .03$) for each group, respectively. A mixed design ANOVA was conducted by adding the factor of fast/slow group to the analysis of prime type and prime duration. There was a main effect of prime type, $F(1, 26) = 21.90, p < .001$, participant group, $F(1, 26) = 6.62, p < .016$, and prime duration, $F(4, 104) = 17.92, p < .001$. Additionally, there was an interaction between prime duration and group, $F(4, 104) = 5.556, p < .001$, between prime type and prime duration, $F(4, 104) = 17.55, p < .001$, and a three-way interaction between prime type, prime duration and group, $F(4, 70) = 6.65, p < .001$.

As seen in Figure 3, the two groups differed on the pattern of the two-way interaction between priming condition and prime duration. Both groups demonstrated an increase followed by a decrease in positive priming as a function of prime duration. However, for the fast group, the prime duration with the maximal positive priming was 50 ms. In contrast, for the slow group, the prime duration with the maximal positive priming did not occur until 150 ms of prime exposure. Furthermore, at the longest prime duration of 2,000 ms, the fast group produced negative priming, with performance in the target primed condition actually worse than the foil primed condition, $t(12) = 4.38, p = .001$. In contrast, there was no difference between these conditions for the slow group at the 2,000 ms prime duration, $t(12) = -1.40, p = .185$. Thus, only the fast group revealed a full crossover from positive to negative priming.

A main effect of group is visible in Figure 3, with the slow group revealing higher accuracy compared to the fast group. Ideally, the use of individually set threshold target durations should have equated performance across the groups. However, the target threshold was determined based on a limited number of forced choice trials and thus reflects some degree of chance. Because of chance, some people did better during the threshold block of trials

than they would have over more trials and so these individuals were assigned a target duration that was briefer than their true threshold. Chance also worked in the opposite manner, serving to assign some individuals a target duration that was too long. When these misassigned participants subsequently performed the priming trials, they did better or worse than expected because they were assigned a target duration that was correspondingly too long or too short. Because the assignment to the slow or fast group is a selection based on extremes, the role of chance worked against the groupings in a systematic manner. Thus, the failure to fully equate the slow and fast groups was expected because of 'regression to the mean', producing better performance for the slow group than for the fast group.

The distinction between the slow and fast group was based on accuracy during the threshold determination block of trials. However, if there is a speed accuracy tradeoff in this paradigm, there may also be a difference in reaction time between the groups. Such a speed difference may be relevant to the transition from positive to negative priming if it takes longer to strategically correct for the influence of the prime. Thus, perhaps the slow group participants responded too quickly, not allowing time to discount the effect of the prime face. We calculated the median reaction time for each participant's correct and incorrect responses both overall and by priming condition, excluding reaction times recorded as shorter than 100 ms or longer than 6,000 ms. No significant differences between the groups nor interactions with the grouping variable were found. The pattern of reaction times over different priming conditions were similar to those found with words, which suggested a race model between choice alternatives (Huber & Cousineau, 2004). For correct trials, reaction time followed a similar pattern to accuracy data, although mirrored, with shorter reaction times to target primed trials, and longer reaction times to foil primed trials. As with the accuracy data, the reaction time difference between the priming conditions increased than decreased with increasing prime duration. Incorrect reaction times mirrored this pattern (target primed trials were slower than foil primed trials) but otherwise appeared similar. In summary, explanations based on speed accuracy tradeoff are ruled out because there were no reaction time differences between the groups and because correct reaction times were faster for the conditions that produced higher accuracy.

A post hoc median split of participants is merely a convenient method for exploring the behavior of individuals who are better versus worse at identifying the briefly flashed target face. The theory that motivated this work makes no particular predictions for distribution of face processing abilities but it does predict that there should be a continuous relationship between the speed of face identification and the speed of habituation. Figure 4 provides a scatter plot of individuals comparing their mean accuracy over

³ Performing a regression analysis on continuously varying measures of individual differences is statistically preferred to median split analyses (Irwin & McClelland, 2003). However, the current situation only included five possible values for the measure of individual differences. In any case, reliability of these individual differences is additionally assessed through replication across all four experiments, and our use of a median split is used to illustrate the qualitative difference between ends of the face identification spectrum.

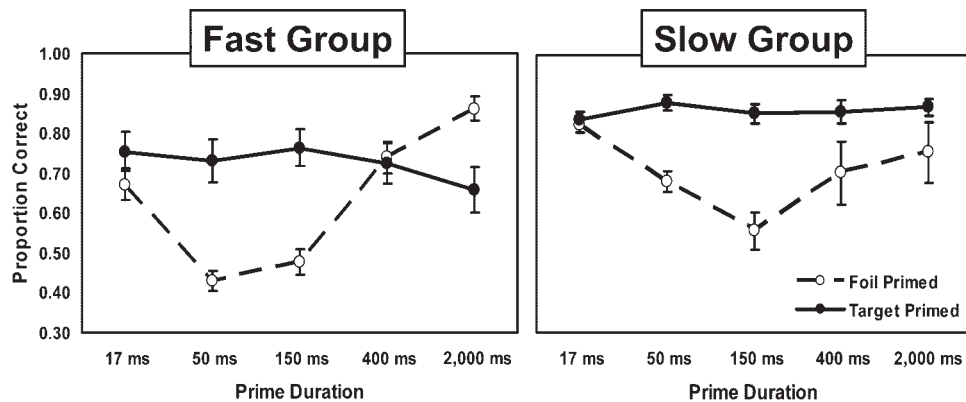


Figure 3. Experiment 1 results broken down by fast versus slow median split of participants based on the target duration necessary for threshold performance. Error bars are ± 1 SE.

the threshold determination block of trials (our proxy for rapid face identification) to the difference between foil and target primed accuracy with a 2,000 ms prime duration (our proxy for the strength of face habituation). The correlation between these measures was $r = .598$, $p = .001$, suggesting that the relationship between rapid face identification and face habituation is continuous. As seen in the figure, the two leftmost individuals were possibly outliers. However, elimination of these individuals did not change any of the statistical conclusions drawn from the median split analyses.

Comparison of Figure 2 and Figure 3 reveals that individual differences are apparent with faces but not words. This comparison

also reveals that face habituation (as indicated by the relationship between the priming conditions as a function of prime duration) appears to occur more slowly than word habituation. However, it is misleading to compare prime durations between face and word priming considering that the displays in the two experiments differed in a variety of ways, such as retinal size (the faces were much larger) and the type of masks (the words were masked with line segments rather than checkerboards). Another key difference is that the faces were unknown, which makes it difficult to assign a verbal label to the face, whereas the words were high frequency words.

In summary, face priming produced increasing and then decreasing positive priming as a function of prime duration. Furthermore, there was a relationship between the speed of face identification (target threshold duration) and the speed of the prime duration effects. If the observed individual differences in the transition from positive to negative priming with face priming reflect true differences in face processing expertise, this confirms the prediction that a higher degree of expertise should correspond to a faster rate of habituation. However, the median split analysis confounds rate of habituation with target duration. This leaves open the possibility that negative priming only occurs for conditions of reduced stimulus energy for the target (i.e., shorter target durations). In other words, a weak target may be more susceptible to the negative effect of a strong prime. The opposite result of this prediction has been found with word priming, in which case reducing target duration actually flipped negative priming to positive priming (Huber et al., 2002; Weidemann et al., 2008), which makes this alternative explanation seem unlikely. Nevertheless, the interaction between target duration and prime duration has not been tested with faces. Therefore, Experiment 2 replicated Experiment 1 while eliminating the setting of different target durations for different individuals. Instead, Experiment 2 fully crossed different prime durations with different target durations for each individual. Thus, the relationship between priming and target duration was examined as a within-subjects manipulation.

Experiment 2

Experiment 2 manipulated target duration within subjects to see if the interaction between prime type and prime duration is af-

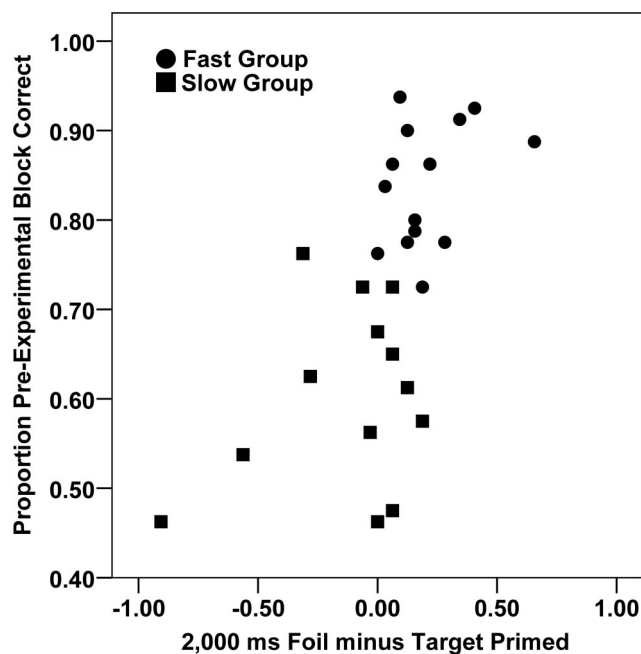


Figure 4. Scatter plot of individual differences comparing average accuracy during the threshold determination block of trials versus the accuracy difference between the target primed and foil primed conditions at the 2,000 ms prime duration. Circles and squares indicate individuals placed into the fast and slow groups, respectively.

affected by target duration. The between subject difference found in Experiment 1 could be explained by the target duration differences between the fast and slow groups. If this explanation is correct, then the within subjects design used in Experiment 2 should produce a priming reversal for short target durations but not for long target durations. Experiment 2 did not use an initial block of trials to set the target duration, but instead contained 'neither primed' trials during the experimental blocks of trials to assess individual differences. These trials allowed separation of the participants into fast and slow groups to test for replication of the individual differences observed in Experiment 1. Including target duration as a manipulation crossed with prime duration greatly increased the number of conditions, and so only three prime durations and three target durations were tested. In addition, unlike Experiment 1, half of the trials reused previously seen faces.

Method

All stimuli, equipment, and procedures were identical to Experiment 1 except as noted.

Participant. Fifty-eight undergraduate psychology students participated in this experiment for course credit.

Procedure. Participants were given 48 practice trials, with progressively decreasing in target durations, followed by five blocks of 96 priming trials with breaks between blocks. The trial procedure was the same as described for Experiment 1 except that target duration was manipulated and the choice faces were offset to the left and right instead of top and bottom. Target durations were tested at 33, 50, and 100 ms and were completely crossed with prime durations of 17, 150, and 2,000 ms. In addition to the target and foil primed trials, there were trials that presented a prime face for 150 ms that was unrelated to either choice (i.e., neither-primed trials). Thus, these trials were identical to threshold trials from Experiment 1 and were included to determine the threshold of each participant. The 96 trials within each block consisted of 72 target or foil primed trials, with four replications at each combination of prime and target duration, and 24 neither primed trials, with eight replications at each target duration. Because this experiment involved a greater number of total trials, half of the trials used entirely new faces while the other half of trials used previously seen faces.

Results and Discussion

A repeated measures analysis determined that there was no main effect of reusing faces, $F(1, 56) < 1$, $p = .586$, but there was an interaction between reuse of faces and prime duration, $F(2, 114) = 4.201$, $p = .02$. However, reuse of faces did not interact with the prime duration by priming condition interaction, $F(2, 56) = 2.273$, $p = .112$, and so subsequent analyses collapsed over trials that used new and reused faces.

To test the relationship between target duration and face priming, we conducted an ANOVA on the factors of prime type, prime duration, and target duration. While there was a three-way interaction of prime type, prime duration and target duration, $F(4, 228) = 5.30$, $p < .001$, it was not in a direction compatible with the proposal that shorter target durations produce a more rapid transition to negative priming. Instead, the nature of this interaction was for the opposite pattern, with a lesser degree of positive

priming for longer target durations. Pairwise tests found that the target primed condition was significantly greater than the foil primed condition at all prime durations (for all tests, $t(56) < -3.04$, $p < .004$) except for 100 ms targets and 2,000 ms primes, $t(57) = -1.92$, $p = .06$. Focusing on this result, we performed an ANOVA using only the longest prime duration, revealing a target duration X prime type interaction, $F(2, 114) = 26.39$, $p < .001$. As seen in Figure 5, longer target durations caused the benefit of target priming versus foil priming to *decrease* (i.e., the data moved in the direction of negative priming with increasing target duration). An explanation of Experiment 1 based on target duration predicted that priming should have instead moved in the direction of positive priming with increasing target duration. Thus, the use of different target durations for different individuals in Experiment 1 was not the cause of the individual differences in the transition from positive to negative priming.

Results from the 150 ms neither primed trials are presented in Appendix A, both for the collapsed results and for the results broken down by median fast/slow split. Based on these trials, threshold target durations were determined for each participant with the same procedure as Experiment 1 (keeping in mind that unlike Experiment 1, Experiment 2 did not set the target durations to these individual threshold durations). The mean threshold target duration for the fast group was 62 ms ($SE = 4.97$). All participants in the slow group had threshold target durations of 100 ms (note that only 33, 50, and 100 ms target durations were tested in this experiment and so this estimate of threshold target duration may be somewhat inflated because of the absence of intermediate target durations).

Figure 6 shows the effects of prime type and prime duration on accuracy for the fast and slow groups. To test for individual differences, we conducted a mixed design ANOVA testing the factors of prime type, prime duration, target duration, and median split of the threshold target durations on accuracy. Accuracy was higher with longer target durations, $F(2, 112) = 243.24$, $p < .001$, for the fast median split group, $F(1, 56) = 61.33$, $p < .001$, and for target primed trials, $F(1, 56) = 127.77$, $p < .001$. There was a main effect of prime duration, $F(2, 112) = 24.33$, $p < .001$, and a prime type X prime duration interaction, $F(2, 112) = 18.88$, $p < .001$.

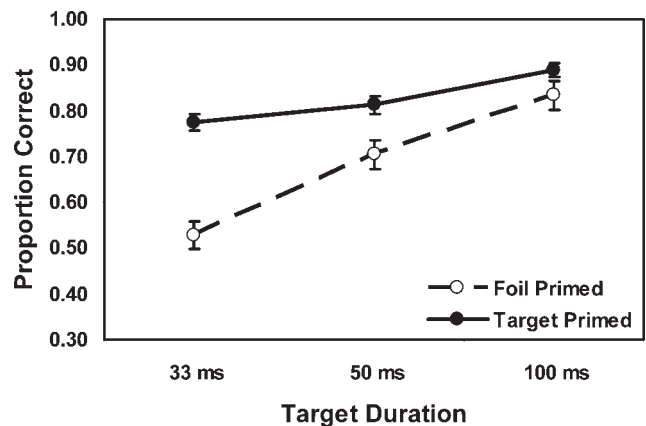


Figure 5. Experiment 2 results collapsed over all participants for the case of strongest prime discounting, which occurred in the 2,000 ms prime duration conditions. Error bars are ± 1 SE.

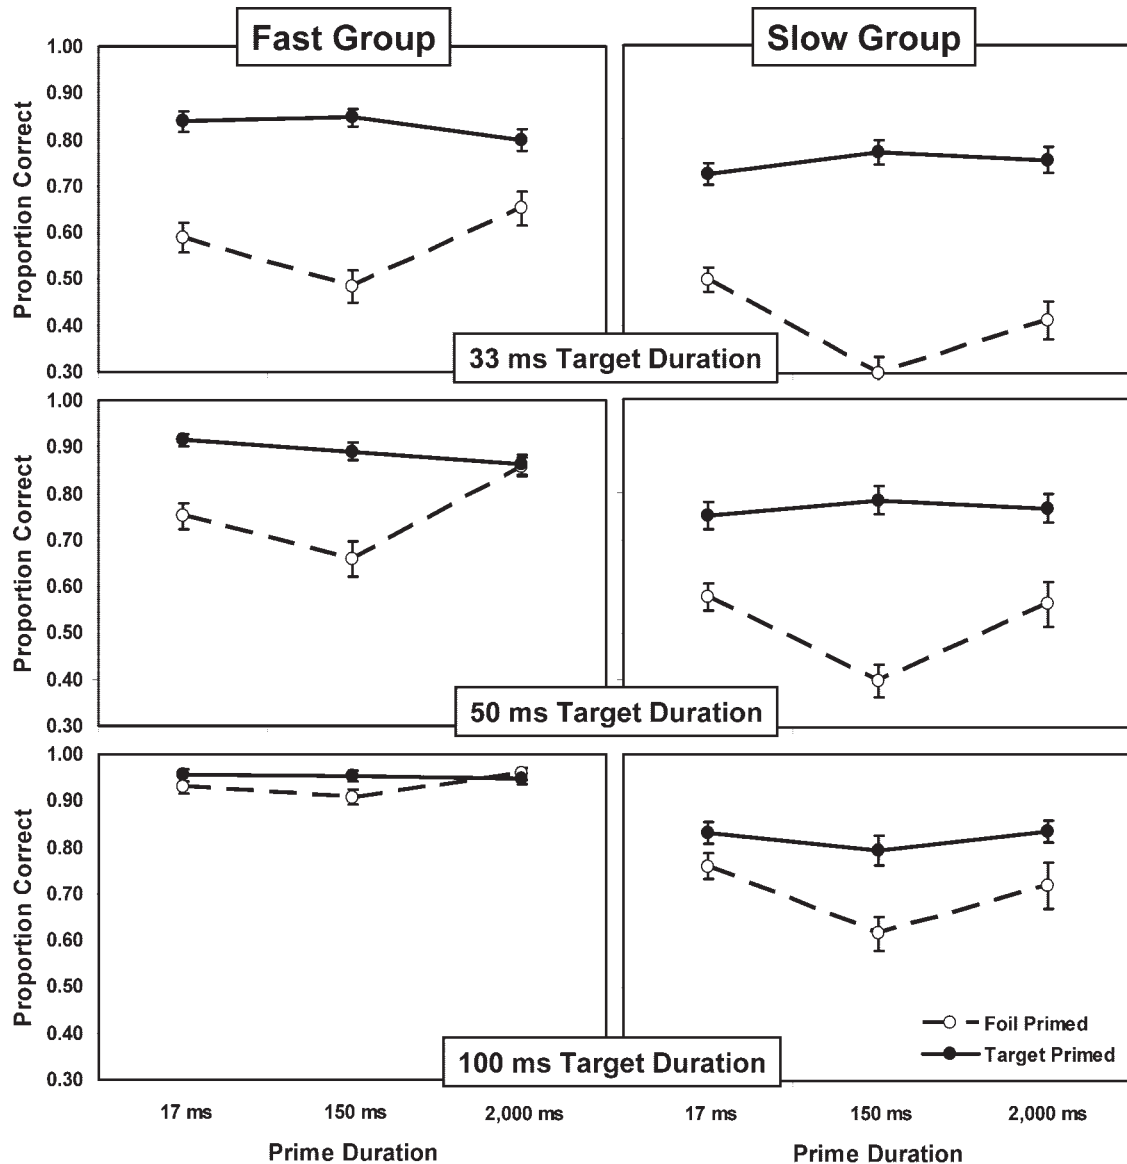


Figure 6. Experiment 2 results split by target duration and by fast/slow median split of participants. Error bars are ± 1 SE.

.001, similar to Experiment 1. There were also interactions of target duration X median split, $F(2, 112) = 5.80$, $p = .004$, prime type X median split, $F(1, 56) = 11.62$, $p = .001$, prime duration X median split, $F(2, 112) = 3.079$, $p = .05$, and target duration X prime type, $F(2, 112) = 86.63$, $p < .001$. Importantly, there was an interaction between prime type, prime duration, and median split, $F(2, 112) = 6.37$, $p = .002$, thus replicating the individual differences of a more rapid elimination of positive priming for the fast participants found in Experiment 1, although without the crossover into negative priming. For the slow group, target primed accuracy was higher than foil primed accuracy at all prime durations and target durations (for all tests, $t(28) < -2.38$, $p < .02$). For the fast group, target primed accuracy was higher than foil primed accuracy at all prime and target durations (for all tests, $t(28) < -2.98$, $p < .006$), except when the prime duration was

2,000 ms and the target duration was either 50 ms, $t(28) = -.21$, $p = .836$, or 100 ms, $t(28) = .9$, $p = .38$, or when prime duration was 17 ms and target duration was 100 ms, $t(28) = -1.87$, $p = .07$.

Emphasizing the importance of setting target duration separately for each participant to avoid ceiling or floor effects, some subjects in the 100 ms target duration achieved 100% accuracy. To compensate for this ceiling effect, we reanalyzed the data using d' (Egan, 1975). Adjustments of $\pm 1/40$ were made to conditions where a subject achieved 0 or 100% accuracy. No changes in the significance of any main effect or interaction were found.

In summary, Experiment 2 demonstrated that use of longer target durations was not the underlying cause of individual differences. While the effect of prime duration was not as striking as in Experiment 1, the same general pattern of more rapid elimination

of positive priming for participants requiring more brief target flashes was obtained both as measured with proportion correct and d' . Thus, these individual differences in the prime duration effect are not caused by different target durations. Instead, these differences must reflect either differences in how participants perform the task or different face identification abilities.

Experiment 3

One difference between Figures 2 and 3 is that there appears to be more symmetry between the target primed and foil primed conditions with word priming; for word priming, prime durations that decreased foil primed performance also increased target primed performance. However, even with word priming, note that the two conditions are not perfect mirror images of each other. These apparent asymmetries between the target primed and foil primed conditions are in truth misleading because neither figure includes the baseline neither-primed condition that presents an unrelated word or face for the appropriate prime duration. Although the above face experiments did not include this condition, the word priming experiment (Huber, 2008b; Experiment 1) did include such a condition, which revealed a u-shaped pattern as a function of prime duration for the baseline neither primed condition. In other words, an intermediate prime duration proved to be the most disruptive to performance even though the prime was unrelated to the choices. In light of this 'forward masking' effect for the baseline condition with words, there was in fact a fair degree of symmetry for the costs and benefits for the target primed and foil primed conditions.

Because of the differences found between words and faces, it is important to evaluate the priming effects we have found relative to a baseline condition, where the identity of the prime face is neither the target nor the foil choice. Without this condition, we cannot make strong conclusions about the symmetry of priming. Furthermore, considering that the target primed condition appears to be unaffected by priming duration, is possible that no priming occurred in this condition (rather than increasing then decreasing priming against a baseline that falls then rises). In Experiment 3 we include this baseline condition to rule out this possibility. In addition, because there is evidence that computer generated faces may be problematic in the

study of face processing (Carlson & Gronlund, 2007), we used photographs of real faces rather than computer generated faces. These photographs were cropped to include just the interiors of faces, which should reduce any reliance on hairstyle in the identification process. Similar to word priming, we hypothesized that the neither primed baseline condition would lie in between the other priming conditions and thus reveal a u-shaped forward masking effect as a function of increasing prime duration.

Method

All stimuli, equipment, and procedures were identical to Experiment 1 except as noted.

Participants. There were 102 undergraduate psychology students participated in this experiment for course credit.

Procedure. In addition to target and foil primed conditions there was also a neither-primed condition where the prime identity was different from both the target and foil choices. Each participant completed 10 priming practice trials. This was followed by two blocks of 40 threshold trials testing target durations of 50, 67, 83, and 100 ms. Participants not achieving 75% accuracy at the 100 ms target duration were assigned a target duration of 117 ms. Participants next completed five blocks of 60 priming trials. Across all blocks, each participant completed 20 trials of each prime condition. Prime durations of 17, 50, 150, 400, and 2,000 ms were tested. Between blocks, participants took breaks of at least 15 s. Because of the smaller number of available faces, each face was repeated four times in either the threshold trials or priming trials.

Materials. Face stimuli were created by cropping the faces of 238 White faces (Tanaka & Pierce, 2009; Tanaka, (2007), "A race face database") to be 149 pixels in width by 155 pixels in height such that they contained only the interior of each face. This was done to place primes as close together as possible, and also to reduce similarity between the faces (all faces in the database have the same exterior features).

Results and Discussion

The priming results split for the fast and slow groups are presented in Figure 7. We conducted a repeated measure ANOVA

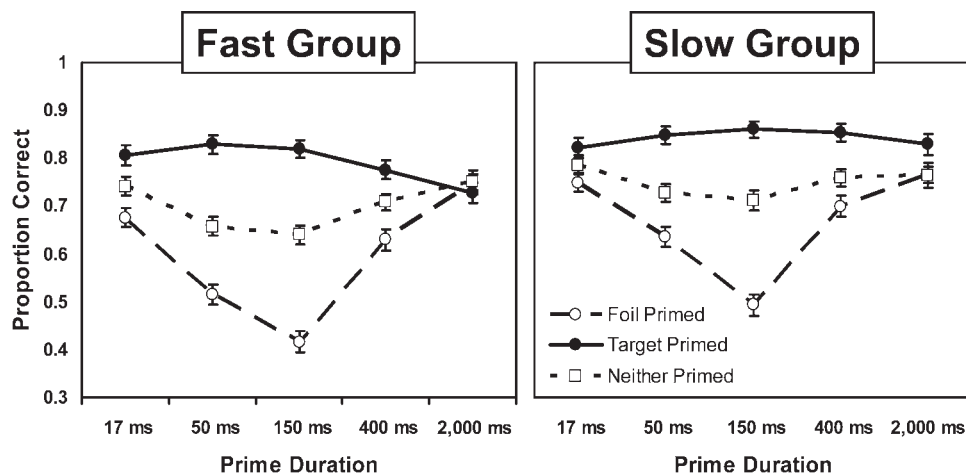


Figure 7. Experiment 3 results broken down by fast versus slow median split of participants based on the target duration necessary for threshold performance. Error bars are ± 1 SE.

on the factors of prime type, prime duration, and fast/slow group. There was a significant three way interaction between prime type, prime duration, and participant group $F(8, 800) = 9.291, p = .002$, as well as a two way prime type and prime duration $F(4, 400) = 45.20, p < .001$, and main effects of participant group, $F(1, 100) = 10.2, p = .002$, prime type, $F(24, 200) = 208.22, p < .001$, and prime duration, $F(4, 400) = 43.64, p < .001$. Planned comparisons of the difference between the prime type conditions for the 2,000 ms prime duration trials revealed positive priming prime (accuracy was lower for foil primed than the target foil primed condition) for the slow participants, $t(50) = 2.14, p = .037$, but for fast participants, positive priming was eliminated, $t(50) = 1.138, p = .26$.

As expected, when a prime unrelated to either choice was presented, accuracy was between that of the target primed and foil primed conditions. Thus, the priming effects are fairly symmetric about this baseline condition and there is indeed a rise and fall in priming for the target primed condition. Furthermore, as expected, the baseline condition revealed u-shaped forward masking. According to the neural habituation model, this is explained through the rise and fall of activation to the prime face combined with inhibition between any concurrently active faces.

Similar to Experiment 1, there was an increase and decrease of positive priming with increasing prime duration. Furthermore, the individual differences in these priming effects were again replicated. This generalizes the effect of Experiments 1 and 2 to the use of real face images. Unlike Experiment 1, there was no crossover to negative priming for the fast group of participants. However, negative priming is not a qualitative prediction of the habituation account. According to this account, habituation is designed to offset the unwanted effect of a previous face blending with a subsequent face. Ideally, this offset through habituation should perfectly balance between integration and separation, producing no difference between the target primed and foil primed conditions. With particularly salient primes, there may be excessive habituation, which produces negative priming. For some reason, the images used in Experiment 1 produced this excessive habituation whereas these real face images did not. There are a number of possible causes for the lack of negative priming compared to Experiment 1, including the change to real faces, repetition of faces, inclusion of exterior face features, and perceived task difficulty. Nevertheless, the predicted differences between groups in terms of the elimination of positive priming as a function of prime duration were confirmed.

Experiment 4

Experiments 1–3 demonstrated the predicted relationship between the speed of face identification and the speed of face habituation by comparing individuals. However, because this involved a subject-variable, a comparison across individuals may involve more than just differences in the speed of face identification. Therefore, Experiment 4 sought to manipulate face expertise in a within-subjects design. This was done by comparing upright versus inverted faces.

While the basic mechanisms that support face processing are hotly debated (Gauthier & Logothetis, 2000; Gauthier & Nelson 2001; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Kanwisher, 2000; McKone, Kanwisher, & Duchaine, 2006) it is un-

contested that we are skilled at identifying upright faces. Investigations into the nature of face expertise suggest that face identification is reliant on information contained in more than just feature identity, such as the precise configurations of features within a face or holistic information over large regions of a face (Maurer, Le Grand, & Mondlock, 2002; Rakover, 2002). Configuration refers to the unique spatial arrangement of features (e.g., eyes, nose, and mouth) on each face. Similarly, holistic face processing occurs without decomposition into separate face parts (Farah, Wilson, Drain, & Tanaka, 1998; although for evidence that holistic face processing can result from decisional factors see Richler, Gauthier, Wenger, & Palmeri, 2008; Wenger & Ingvalson, 2002, 2003). It has been suggested that face inversion disrupts holistic or configuration information (Farah, Tanaka, & Drain, 1995; Freire, Lee, & Symons, 2000; Leder & Bruce 2000; Rossion & Gauthier, 2002), leading to the face inversion effect, which is a disproportionate deficit for inversion of faces as compared to inversion of other objects (Yin, 1969). In the present work, the aspect of face perception that is disrupted by inversion is not under investigation, and so we will refer to this as face configuration for brevity, keeping in mind that identical arguments could be made for holistic information. Regardless of the mechanism, we seek to exploit face inversion to manipulate the strength of response for the information that underlies face identification.

Experiment 4 manipulated face expertise by comparing priming of upright versus inverted faces. The design was identical to Experiment 1, except that half the trials used upright presentations of prime, target, and choice faces, and the other half of trials used inverted faces for all presentations. Because both upright and inverted faces were drawn from the same pool of face stimuli, this allows manipulation of expertise while equating many low-level visual aspects of the images. It was predicted that because recognition ability for inverted faces is diminished, there should be less habituation for inverted faces. According to this prediction, there should be a slower transition from positive to negative priming for immediate repetitions of inverted faces.

Method. All stimuli, equipment, and procedures were identical to Experiment 1 except as noted.

Participants. Forty undergraduate psychology students participated in this experiment for course credit.

Procedure. Half of the trials for each condition, including practice and threshold blocks, used inverted face stimuli for primes, targets, and choice faces. This resulted in 10 trials per condition in threshold trials, and 16 trials per condition in priming trials. The calculation of each subject's threshold duration was taken over both upright and inverted threshold trials and the same target duration was used for both trial types during the experimental blocks of trials.

Materials. Inverted faces were 180 degree rotations from the upright faces. Each face was only presented once, either upright or inverted.

Results and Discussion

Figure 8 shows the interaction between prime type and prime duration separated by fast and slow groups and also by face orientation. These results also appear in Appendix A, along with the threshold results, and results collapsed across groups. It is clear that for upright faces, both the overall prime duration effect and the

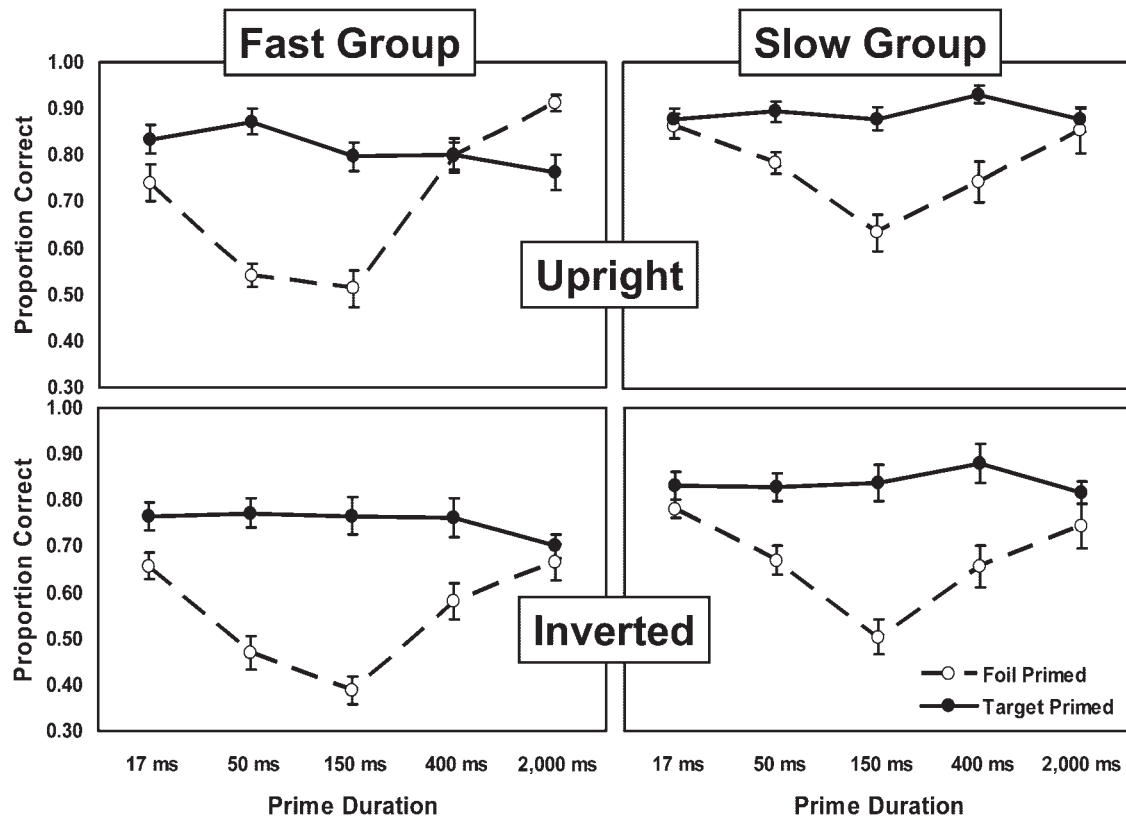


Figure 8. Experiment 4 results separated by upright and inverted conditions and by fast/slow median split of participants. Error bars are ± 1 SE.

median split differences were replicated. First, we consider only the upright face conditions, which replicate the results of Experiment 1, and then we consider inverted faces and include face orientation as a factor in the ANOVA.

For upright faces, as in Experiment 1, there were effects on accuracy of prime type, $F(1, 38) = 80.89$, $p < .001$, prime duration, $F(4, 152) = 31.063$, $p < .001$, and a prime type \times prime duration interaction, $F(4, 152) = 32.36$, $p < .001$. Replicating the median split distinction, there was a three way interaction between prime type, prime duration and participant group, $F(4, 152) = 8.16$, $p < .001$. There was also a main effect of group, $F(1, 38) = 12.87$, $p = .001$, where overall accuracy was higher in the slow group, which as elaborated previously is explained by regression to the mean.

For inverted faces, the time course of the prime effect looks similar for both groups. Accuracy was lower for inverted faces compared to upright faces, $F(1, 38) = 160.04$, $p < .001$, as expected. The effect of prime type was larger with inversion, $F(1, 38) = 30.10$, $p < .001$, presumably because of the lack of any negative priming for inverted faces, even for the fast group. The three-way interaction of inversion, prime type, and prime duration was marginally significant, $F(4, 152) = 2.15$, $p = .08$. The four-way interaction between orientation, prime type, prime duration, and participant group, $F(4, 152) = 2.23$, $p = .07$, was marginally significant, and a planned comparison of prime type, orientation, and participant group for only the 2,000 ms prime

duration highlighted the manner in which face inversion affected the groups differently, $F(1, 38) = 4.16$, $p = .05$. For the fast group, 2,000 ms of foil priming resulted in higher accuracy than target priming (i.e., a pattern of negative priming) when viewing upright faces, $t(19) = 5.02$, $p < .001$, but at this prime duration, there was no effect of prime type when viewing inverted faces, $t(19) = -.85$, $p = .41$. In contrast to this inversion priming effect for the fast group, for the slow group there was neither a priming effect for upright faces, $t(19) = -.498$, $p = .62$, nor for inverted faces, $t(19) = -1.306$, $p = .21$ in the 2,000 ms prime duration conditions.

In summary, the individual differences with upright faces were again replicated, but these differences apparently vanished with inverted faces. For inverted faces, the rate of prime habituation, as revealed by the transition from positive to negative priming, was essentially the same for both participant groups, and similar to the upright slow group. For the fast group, the comparison between upright and inverted faces demonstrates that a within subjects manipulation of expertise can produce the predicted change in the speed of habituation. However, the absence of a habituation face inversion effect for the slow group was unexpected. To explore alternative interpretations of this apparent interaction between individual differences and face orientation, the habituation model was applied to these data, as reported next.

Experiment 5: Simulation Study

As predicted from a neural habituation explanation of face priming, both individual differences and face orientation modulated the transition from positive to negative priming—in both cases, better performance (i.e., upright faces or individuals with low perceptual thresholds) produced more sizable negative priming after long duration primes. These results were predicted under the assumptions that activation drives habituation and that a higher degree of performance corresponds to stronger activation. However, the precise manner in which individuals differ and upright versus inverted faces differ is unclear, and there may be several alternative implementations of this theory that are qualitatively compatible with these results. In this simulation study, we explored one such model to assess whether it was able to quantitatively handle these results, and whether it could shed light on the apparent interaction between face orientation and individual differences. Additionally, by using the same model that explained similar results with word priming (Huber, 2008b), this simulation study might explain the different time course for word versus face priming. The dynamic activation functions of the model and the three layer structure are identical to that presented by Huber and O'Reilly (2003), and are presented here briefly for completeness.

Model Structure

As applied to words, Huber (2008b) assumed that the top layer of this three layer model represented lexical-semantic information, the mid level represented orthographic information, and the bottom visual line segments, thus accounting for different dynamics for masks, orthographic priming, and associative priming. In the current application, we do not have results that constrain the exact nature of the information at each level of face processing beyond inversion effects (i.e., we didn't manipulate face similarity), and so we more generically refer to the bottom layer as visual processing, the mid layer as processing of face parts (which are equally represented when inverted) and the top layer as processing of face wholes (which are disrupted when inverted). We are not making any strong claims about what these face parts or whole are (e.g., we are not claiming that these face parts are necessarily face features such as the eyes, nose, and mouth), just that the higher level of perception is disrupted by inversion. Thus, the goal of this simulation study is to investigate the dynamics of face processing rather the representation of faces.

The model has a three layer structure as shown in Figure 9. Each node in the model can be viewed as describing the average firing rate of a large group of similarly connected neurons. In the bottom layer of the model there exists a simulated node for each possible presented visual object (including pattern masks) in each possible spatial location. The projections of these nodes to the next layer converge so that each node in the second layer codes for the face parts unique to each face, regardless of where the face is presented. Pattern masks, being a simple visual pattern, are only represented at the bottom layer. The model used a localist representation such that each node projects only to higher level nodes that code for the same face; a distributed representation was not necessary because we did not manipulate similarity. Thus, layers 2 and 3 were simply implemented through a single node for each face. Connection weights were set at fixed parameter values, and no learning was

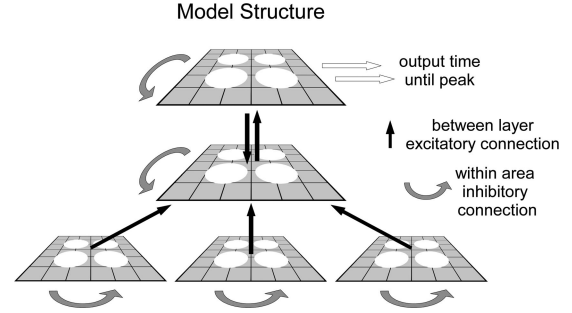


Figure 9. The three layer structure of the artificial neural network applied to the reported results. The bottom layer is sensitive to spatial position, and so the same face presented in different locations (e.g., left, center, or right) activates different nodes. Each possible item in a trial (target, foil, and mask) is represented at this layer. At higher layers, connections from layer 1 converge such that the same face presented in different locations contacts the same nodes. A simple one to one mapping is used in connecting layers 2 and 3 and the use of two separate layers is designed to capture differences in the dynamics of activation, rather than making specific representational assumptions. Activation of each node is achieved through a real value (i.e., rate-coded rather than spiking). Within each area, nodes inhibit each other. The time to achieve maximum response in layer 3 is used to compute accuracy.

used in the simulations. The weights between the first and second layers of the model were fixed at one for connections between the same face at each level and zero otherwise. Connections between the second and third layers were likewise set to zero between different faces, but were set to one of two possible free parameter values for same face connections, one for upright faces and the other for inverted faces. In this manner, the strength of response for upright versus inverted faces was manipulated. Additionally, by allowing these two values to differ between the fast and slow groups, strength of response for individual differences was manipulated.

Neural Dynamics

The dynamics of each node use the same assumptions as found in O'Reilly and Munakata's (2000) LEABRA framework, but with the addition of synaptic depression. Lateral inhibition was used within each layer of the model, simulating the effects of inhibitory interneurons, which dampen excessive excitatory activity. The membrane potential (i.e., activation) of each node is updated by Equation 1. The equations provided here assume updates are performed every millisecond. More model details can be found in Huber and O'Reilly (2003). Membrane potential (v_i^n) for node i in layer n is increased by excitatory input weighted by the connection strength between sending (j) and receiving (i) node ($w_{ij}o_j^{n-1}$) from the previous layer or from input. Membrane potential is decreased by within layer inhibition (I) and constant leak currents (L).

$$\frac{\Delta v_i^n}{S_n} = (1 - v_i^n) \left\{ \sum_j w_{ij} o_j^{n-1} \right\} - v_i^n \left\{ L + I \sum_l o_l^n \right\}$$

$$o = \begin{cases} (v - \theta)a & v > \theta \\ 0 & v \leq \theta \end{cases}$$

$$\frac{\Delta a}{S_n} = R(1 - a) - Do$$

Unlike traditional artificial neural networks, which only have an activation value that varies with time, the nodes of this model contain two time varying parameters that multiply to produce synaptic output as determined by presynaptic activity as well as by the level of synaptic resources. In this manner, the model includes habituation because of recent activity in the specific form of synaptic depression (Abbott et al., 1997; Tsodyks & Markram, 1997). The output (o) of each node is 0.0 until the membrane potential increases past a threshold value (θ), after which it is scaled by a dynamically varying factor (a) representing available resources for that particular connection (see Equation 2). In other words, the effect of a node is the product of two terms, with the first representing the on-average firing rate ($v - \theta$), and the second representing the effect of each action potential in light of the available nondepleted resources (a). One example of synaptic resources is available presynaptic neurotransmitter—if a neuron has recently been highly active and has depleted its neurotransmitter supply, there will be little effect of ongoing action potentials. Equation 3 specifies the dynamic update for synaptic resources, and contains a constant for the recovery rate (R), and a constant for the rate of depression of activity (D). The dynamic update Equations 1 and 3 additionally contain a parameter, S_n , which regulates the speed of information integration for layer n , thus specifying the speed of processing at each layer of the model. The model was run in time steps of one millisecond. Input to a bottom layer node was set to 0.0 in general, but changed to 1.0 for the time steps when a visual object was presented that was the preferred input for that node.

Decision Rule

Because the behavioral task was threshold identification based on partial information, the decision rule in the model was to choose the test face identified more rapidly. As previously assumed in modeling threshold identification of words, we assumed that this perceptual fluency was used to assess the residual activation from the briefly flashed target face. This fluency measure was calculated as the “time to peak” response found separately for each choice face. In other words, if the target choice face reached its peak activation first, it was chosen, resulting in an accurate trial, but if the foil choice face reach peak activation first, an error occurred. Thus, the choice face identified more quickly was chosen. This measure is analogous to a horserace model of forced choice, which was demonstrated to capture correct and error reaction time distributions with the word priming version of this task (Huber & Cousineau, 2004). Although we do not include it in the current model, the calculation of perceptual fluency could be implemented in simulated neurons through the difference between a fast and a slow integrator, with the latter inhibiting the activation of the former (Huber & O'Reilly, 2003).

Simulations were not run with any specific source of noise, and were thus a single deterministic run for each condition. To capture trial by trial variability, the simulated time to peak for the target and foil nodes were assumed to be normally distributed, and the deterministic simulation values provided the average time to peak for the target and foil in that particular condition. Assuming that variance is proportional to the mean, a noise constant (m) was used to calculate the variance of a condition based on the average time to peak. This assumption is appropriate if each millisecond of

processing provides an independent normally distributed error term with mean zero and variance m . Accuracy then becomes a convolution of two normal distributions. More simply, this can be reformulated as a single normal distribution for the difference between the peak time for the target and the peak time for the foil, with the average difference found by subtracting the average peak time of target and foil, and the variance of the difference found by adding the variances of peak times for target and foil. Thus, probability correct is the inverse z-transformation (i.e., cumulative standard normal) based on the value $(F-T)/[m(T^2+F^2)]^{1/2}$ where T is the time to peak for the target choice face and F is the time to peak for the foil choice face in the condition of interest. This is a slight, but not critical change from the method used in Huber and O'Reilly (2003). The logistic function used by Huber and O'Reilly works equally well, although the current use of separate distributions for target and foil is somewhat better motivated on theoretical grounds. Because the parameter m only operates on the generated peak times, it cannot influence the dynamics of the model, and only serves to transform model results into the accuracy scale (see Appendix Figure A1 for example simulations demonstrating the role of the free parameter m). Thus, allowing different values of m in modeling the two groups of participants cannot explain the different habituation rates.

Modeling the Current Experiments

The fundamental test of the theory posed by these experiments was to examine the relationship between the speed of face identification and the speed of habituation, with the speed of face identification manipulated both by individual differences and by face inversion. However, there are in fact two methods for manipulating the speed of face identification to capture different degrees of face expertise. These correspond to changing the connection weights between layers (w) or changing the processing speed within a layer (S). Increasing either the connection strength or processing speed produces faster activation for a briefly presented face (i.e., a shorter target duration is needed for threshold performance) as well as a faster accrual of habituation, thus producing a more rapid change from positive to negative priming. Therefore, we cannot distinguish between these two explanations. However, if the speed of processing for a layer of representation is related to an anatomical constraint, then it not obvious why this parameter would vary between individuals. In keeping with a long tradition of connectionist modeling (Rumelhart & McClelland, 1986a, 1986b; O'Reilly & Munakata, 2000), we hypothesized that connection weights are modulated by experience. Therefore, we parameterized connection weights in different conditions to capture individual differences in face processing and to capture the difference between upright versus inverted faces.

To capture the differences between the dynamics of words and faces, we allowed several parameters to vary, although many were kept the same. Leak ($L = 0.15$), depletion ($D = 0.0324$), recovery ($R = 0.022$), and firing threshold ($\theta = 0.15$), are considered generic properties of all neurons and these were fixed to the same values for all simulated nodes as reported by Huber and O'Reilly (2003). Likewise, the speed of processing ($S_i = 0.054$) and inhibition ($I_i = 0.30$) for the visual input layer of the model is presumably the same type of visual response as for words (e.g., primary visual cortex). Finally, the same degree of feedback ($F = 0.25$) from layer 3 onto layer 2 was used to capture attractor dynamics. The three parameters that were

truly free, and allowed to take on different values as compared to the previously published parameters used for word priming, were the speed of processing for layer 2 (S_2), and layer 3 (S_3), and the inhibition for layers 2 and 3 (I_{23}).

The within subject effect of inversion and the between subject difference between groups were modeled with the connection weight between the second and third layer (four values for w_{23}). Additionally, each group was allowed to have a different variance multiplier (two values for m), corresponding to different levels of variability for different individuals. As discussed above, this parameter does not affect the model dynamics, just the manner in which the model dynamics map into the accuracy scale. The model was fit with chi-squared error as calculated by log likelihoods (e.g., Batchelder & Riefer, 1990), using nine free parameters to capture the 40 conditions, with 360 data points per condition.

Results and Discussion

Figure 10 shows the fit of the model to the data from Experiment 4. The model clearly fits the qualitative pattern of the data, and is quantitatively accurate for most conditions. The median chi-square per conditions was 1.01, and the model was not statistically different from the data in 31 of the 40 conditions (according to a chi-squared test with 31/40 degrees of freedom, considering

the ratio of 9/40 of parameters per condition). The best-fit parameters were as follows: $S_2 = 0.015$ (speed of integration for the second layer), $S_3 = 0.022$ (speed of integration for the third layer), $I_{23} = 0.52$ (inhibition in the second and third layer), $w_{23}\text{-fast-upright} = 2.05$, $w_{23}\text{-slow-upright} = 0.83$, $w_{23}\text{-fast-inverted} = 0.73$, $w_{23}\text{-slow-inverted} = 0.50$, $m_{\text{fast}} = 0.21$, and $m_{\text{slow}} = 0.10$.

A number of interesting observations can be made beyond demonstrating the sufficiency of the model in capturing these data. Indeed, in descriptive modeling, such as an application of signal detection theory (Egan, 1975), the goal is not just a good fit to data, but additionally to transform data into theoretically meaningful parameter values. Looking at the best-fit parameters, it is notable that the processing speed of the third layer is slightly faster (larger numerically) than that of the second layer (the word priming values were $S_2 = 0.046$, $S_3 = 0.015$). The overall differences in speed when comparing faces and words does not necessarily indicate anything important about these classes of visual objects considering that the types of displays (e.g., retinal size) and masks were vastly different. However, that the third layer runs faster than the second layer for faces but not for words suggests that higher level face information is processed quickly relative to lower level face information. This is sensible if face wholes are based on spatial configuration without needing to wait for full identification

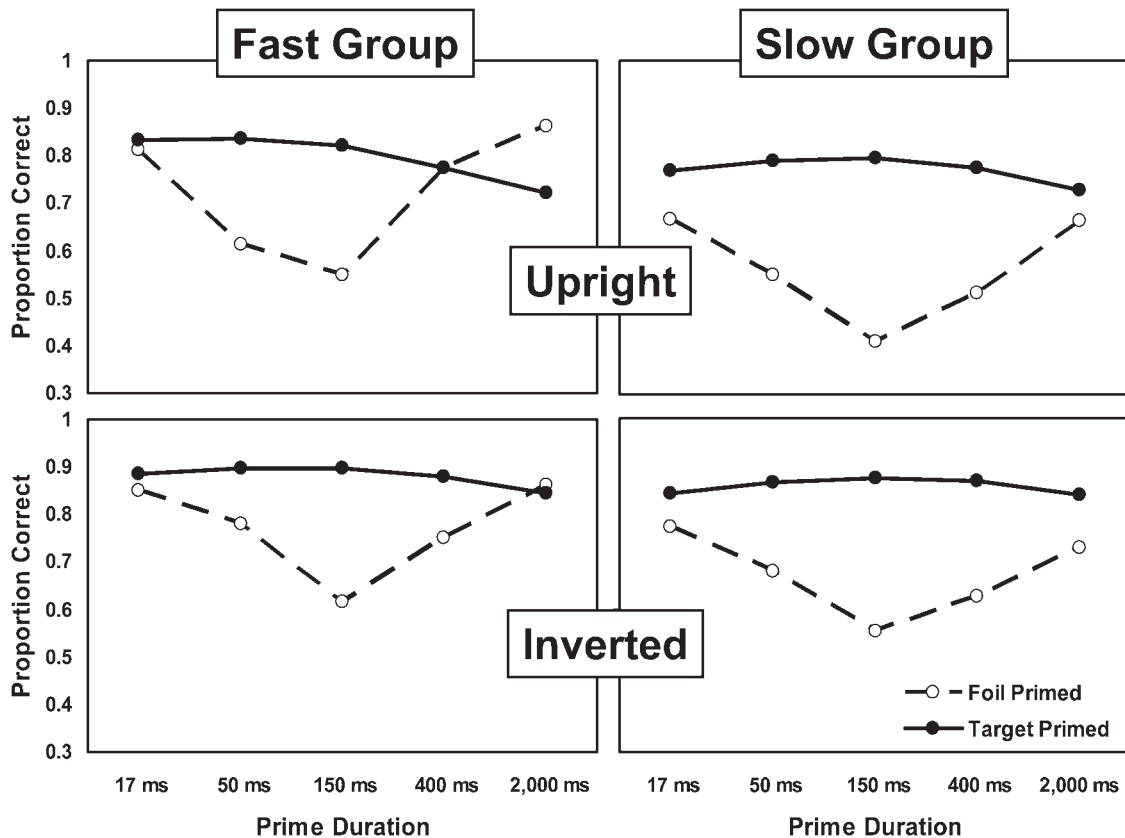


Figure 10. Experiment 4 model results from best-fitting parameters. The connection weights between layers 2 and 3 were allowed to vary for each of the four conditions. The performance gain parameter, m , was allowed to vary between the fast and slow groups. All other parameters were the same in all conditions and many were set to default constants suggested by Huber & O'Reilly (2003).

of the particular face features. In contrast, word identification critically hinges upon the letters that comprise the word. It is also interesting to note that inhibition for face processing is high relative to words (all I values were 0.3 for words). Anecdotally this makes sense. We commonly need to maintain several words at once to extract meaning over a sentence, but it might be confusing to simultaneously identify and maintain several faces (i.e., we usually only attend to one face at a time).

Despite the apparent behavioral trends, both the fast and slow groups reveal a sizable effect of inversion for the connection weights. In other words, application of the model places a cautionary note on the apparent conclusion from the behavioral data that only the fast group produced an inversion effect. The observed pattern of data might result from greater sensitivity in the behavioral measure for the fast group; the very small inversion effects for the slow group may in fact represent sizable changes to the underlying representation. Sensibly, the connection strengths for the fast group were overall greater than the corresponding values for the slow group, corresponding to the fast group's ability to more rapidly identify the target face and the fast group's faster transition from positive to negative priming with upright faces. Interestingly, connection strength differences between groups were actually larger in magnitude than the inversion effects within groups, suggesting that these individual differences in face processing may be something of importance and are perhaps more sizable than face inversion.

General Discussion

Summary of Results

Four behavioral studies examined immediate face repetition priming and confirmed predictions of the claim that habituation temporally segregates streams of perceptual events. According to this theory, temporal integration between recently presented stimuli (i.e., primes) and currently presented stimuli (i.e., targets) produces positive priming, as revealed by accuracy benefits when priming the target but deficits when priming the incorrect choice (i.e., the foil). However, this positive priming is offset and possibly reversed by habituation to primes presented for longer durations. This theory predicted that the speed of identification for a perceptual object should relate to the speed of the transition from positive to negative priming, and this prediction was confirmed both through individual differences in face identification and by manipulating face orientation.

Experiment 1 used a paradigm that previously established the positive to negative priming transition with words, and found similar priming effects with computer generated face images. However, unlike word priming, there were reliable individual differences found in all four experiments. Participants who could more rapidly identify the target face (as determined by threshold target duration) produced a faster transition from positive to negative priming. Experiment 2 demonstrated that this was not because of the use of different target durations for the fast and slow median split groups of participants. Experiment 3 included a condition where neither the target nor the foil face was primed and tested photographs of real faces cropped to the interior of the face. This verified equal costs and benefits of priming relative to a neutral baseline. Experiment 4 found that upright faces produced a

faster transition to negative priming as compared to inverted faces. This inversion effect was particularly prominent for the fast group of participants and seemingly absent for the slow group. Experiment 5 was a simulation study that applied a dynamic neural network model with habituation to the individual difference and inversion effect data from Experiment 4. This provided an accurate account of the data, demonstrating that these qualitative predictions were also quantitatively in accord with neural habituation. The best-fitting connection strength parameters revealed sizable effects of face inversion for both the fast and slow groups, which suggests that the apparent lack of inversion effect with the slow group was a floor effect for the ability of habituation to produce negative priming.

Relationship to Other Paradigms and Measures

This habituation theory may explain a wide variety of paradigms that produce transitory deficits or changes in neural activation after prolonged exposure to a stimulus. For instance, similar phenomena occur with studies of perception, lexical processing, attention, evaluation, and episodic recognition (see for a review Huber, 2008b). Behavioral, ERP, and MEG experiments have confirmed predictions from the claim that habituation between perceptual and lexical-semantic processing underlies immediate repetition priming deficits with words (Huber, Tian, Curran, O'Reilly, & Worocho, 2008). The idea that habituation occurs because of the loss of association between levels of representation has also been applied to semantic satiation, which is the feeling that a word loses meaning when repeated many times. Experiments ruled out explanations based on lexical or semantic fatigue, suggesting that is the inability to access meaning from the repeated word form that explains the phenomenon (Tian & Huber, in revision). Demonstrating the generality of this theory, something similar to our results was found in the phenomenon of 'repetition blindness', in which an observer fails to report the second occurrence of a word in a sentence presented in rapid serial order (Kanwisher, 1987). This has been explained as a failure to separate each word as a separate word token (i.e., the first occurrence vs. the second occurrence) despite adequate identification of the word type (i.e., evidence that a particular word was seen recently, regardless of where and when). Habituation may be the underlying cause of this failure to produce identifiably separate responses to repeated words. Similar to our face inversion results, Coltheart and Langdon (2003) produced the usual repetition blindness effect with known words, but in the same experiment, repetitions of nonwords instead produced a benefit (i.e., more likely to report a second occurrence of a nonword). These results are explained by habituation because nonwords are unfamiliar and habituate more slowly than words; because the first occurrence of an unfamiliar nonword produces little or no habituation, it only serves to provide lingering activation that helps identify the second occurrence (i.e., positive priming).

This theory also may also provide insight into the nature of repetition suppression, which is the attenuation of the neural response to a repeated stimulus as measured with single cell recording, fMRI, or EEG/MEG (Grill-Spector, Henson, & Martin, 2006; Ranganath & Rainer, 2003). An account of these effects because of habituation depends on the time scale over which the repetitions occur. Behavioral facilitation is typically seen for a target that is an

immediate repetition of a brief or subliminal prime when there is no delay between prime and target. In contrast, longer prime durations, multiple prime presentations, or the inclusion of a very short delay (seconds or less) between prime and target often produce a behavioral and neural deficit, such as with repetition blindness. The model used in Experiment 5 naturally produces behavioral facilitation with subliminal no-delay priming because of temporal integration between prime and target. The model also produces reduced neural responses following longer duration primes (e.g., when habituation more than offsets the benefits of temporal integration) or when there is a delay between prime and target. With a short delay between prime and target, residual prime activation fades (i.e., there is no temporal integration and thus no facilitation). However, recovery from habituation operates on a slower time scale, and so the neural activation to the target is weak because of lingering habituation. On a longer time scale, with seconds or longer between prime and target, behavioral benefits are almost always found. These longer time scale benefits are presumably because of learning, which is not included in the current model. Habituation is a temporary effect, and with longer delays, the neural response recovers, thus unmasking the behavioral benefits of learning. Although the neural response is temporarily placed in a habituated state through repetitions, the brain is presumably learning the repeated stimulus, which supports a behavioral benefit after a longer delay. Furthermore, a well learned visual object (e.g., an upright face rather than an inverted face) is expected to produce a faster transition to habituation, and so a measure with poor temporal resolution, such as fMRI, is expected to produce less activation across an averaged time window. The prediction for well learned visual objects as measured with high temporal resolution techniques (e.g., scalp EEG) are more complicated, as discussed next.

The explanation of the reported face inversion effects assumed that upright faces have an initially stronger response than inverted faces, which then falls more rapidly to a habituated state. However, this assumption seems at odds with the finding that inverted faces produce larger N170 ERPs than upright faces (Rossion et al., 1999). However, this finding is only problematic if more negative N170 responses are assumed to indicate more neural activation for the neurons responsible for face identification. However, source localization of ERP responses is complicated not only by anatomical orientation (e.g., there is often a larger P170 as measured at frontal regions concurrent with the reduced N170) but also because multiple cortical sources may be simultaneously active (e.g., non-face identification processes may contribute to the N170). Therefore, an ERP difference between conditions cannot be uniquely attributed to a particular source, and, furthermore, the direction of the difference does not necessarily indicate a particular direction of change in the underlying neural activity. Because there is no unique solution to the so-called 'inverse problem' of electrophysiology (Moshier, Baillet, & Leahy, 1999), all that can be concluded is that the summed neural activity for upright faces is different than that of inverted faces at ~170 ms after presentation. Thus, in the absence of source modeling, the ERP N170 face inversion results are equally compatible with the claim that the neurons most directly responsible for face identification are initially more active for upright faces as compared to inverted faces.

Implications of the Reported Individual Differences

Perhaps the most well studied individual difference with faces is the deficit of prosopagnosia, which is characterized by a selective inability to identify faces as compared to other visual objects (Duchan, Yovel, Butterworth, & Nakayama, 2006). Demonstrating the opposite end of the face expertise continuum, recent investigations have uncovered individuals who are extremely skilled at face identification, to the point of needing to hide this ability in social situations where it might seem inappropriate to easily identify someone seen just fleetingly many years beforehand (Russell, Duchaine, & Nakayama, 2009). In combination with our results with normal individuals, these clinical deficits and enhancements suggest there may be larger variations in face perception ability than has been previously appreciated. For instance, application of the habituation model to our Experiment 4 data revealed a larger effect for the median split of individuals as compared to the effect of face inversion. However, it has been reported that use of computer generated faces may tend to reduce inversion effects as compared to photographs of real faces (Carlson & Gronlund, 2007). Providing some closure on this issue, we recently replicated the inversion and individual difference results of Experiment 4 using face photographs similar to those used in Experiment 3, again finding that individual difference effects were larger than inversion effects. For reasons of space we did not currently report these results.

The reliability of these individual differences demonstrates that this immediate priming paradigm may be useful in clinical diagnoses and treatment of perceptual processing disorders. For instance, this forced choice priming paradigm has been used with words, revealing that the extent of language disorder in schizophrenia relates to the magnitude of semantic priming benefits (Quelen, Grainger, & Raymond, 2005). This was found for the both-primed condition (not currently used), which, according to Huber (2008b), reflects the strength of top-down facilitation as compared to bottom-up support. The current application of the priming paradigm with multiple prime durations demonstrates the predicted link between speed of identification for a stimulus and the rate at which rapid presentations can be handled with minimal source confusion. However, the reported individual differences are correlative in nature, and a training study is needed. Providing an example of such a training study, it has been found that training on sound identification increases reading speed (Merzenich, Jenkins, Johnston, Schreiner, Miller, & Tallal, 1996). Finally, our paradigm for identifying and understanding individual differences may be useful in the study of autism-spectrum disorder (ASD). Although much research on ASD has focused on a lack of social motivation, perceptual processing in general, and face processing in particular, is also disrupted (Dawson, Webb, & McPartland, 2005; Deruelle, Rondan, Gepner, & Tardif, 2004; McCleery, Allman, Carver, & Dobkins, 2007). Variants of our priming paradigm might, for instance, identify whether ASD face processing deficits are in part because of difficulty segregating different views of a face.

Conclusions

Four experiments demonstrated that immediate face priming produces positive priming for short prime durations but negative or no priming for longer target durations. This effect was previously

shown with immediate word priming and this generalization to faces was expected by a theory of temporal segregation through habituation. According to this theory, the temporal integration between short duration prime faces and target faces produces source confusion (see also, Anaki, Boyd, & Moscovitch, 2007). However, habituation after long duration prime faces more than offsets this effect, producing negative priming. More importantly, the reported results tested the prediction that more rapid face identification should produce more rapid habituation and, therefore, more quickly produce negative priming as a function of increasing prime duration. This prediction was confirmed both with individual differences relating face detection threshold to priming and also by comparing priming for upright versus inverted faces. The observed differences in face perception were large and found in all experiments. Beyond verifying key predictions of this habituation account, these results suggest a new technique for examining individual differences in face processing.

References

- Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275, 220–224.
- Anaki, D., Boyd, J., & Moscovitch, M. (2007). Temporal integration in face perception: Evidence of configural processing of temporally separated face parts. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 1–19.
- Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, 97, 548–564.
- Bennett, J. D., Lleras, A., Oriet, C., & Enns, J. T. (2007). A negative compatibility effect in priming of emotional faces. *Psychonomic Bulletin & Review*, 14, 908–912.
- Carlson, C. A., & Gronlund, S. D. (2007, November). *Computer-generated faces are not processed like real faces*. Presented at Psychonomics, Long Beach, CA.
- Carpenter, G. A., & Grossberg, S. (2003). Adaptive resonance theory. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2 ed., pp. 87–90). Cambridge, MA: MIT Press.
- Coltheart, V., & Langdon, R. (2003). Repetition blindness for words yet repetition advantage for nonwords. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 29, 171–185.
- Dawson, G., Webb, S. J., & McPartland, J. (2005). Understanding the nature of face processing impairment in autism: Insights from behavioral and electrophysiological studies. *Developmental Neuropsychology*, 27, 403–424.
- Deruelle, C., Rondan, C., Gepner, B., & Tardif, C. (2004). Spatial frequency and face processing in children with autism and asperger syndrome. *Journal of Autism and Developmental Disorders*, 34, 199–209.
- Duchan, B. C., Yovel, G., Butterworth, E. J., & Nakayama, K. (2006). Prosopagnosia as an impairment to face-specific mechanisms: Elimination of the alternative hypothesis in a developmental case. *Cognitive Neuropsychology*, 23, 714–747.
- Egan, J. P. (1975). *Signal detection theory and ROC analysis*. New York: Academic Press.
- Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception and Performance*, 21, 628–634.
- Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1998). What is 'special' about face perception? *Psychological Review*, 105, 482–498.
- Fox, C. J., & Barton, J. J. S. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research*, 1127, 80–89.
- Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception*, 29, 159–170.
- Gauthier, I., & Logothetis, N. K. (2000). Is face recognition not so unique after all? *Cognitive Neuropsychology*, 17, 125–142.
- Gauthier, I., & Nelson, C. A. (2001). The development of face expertise. *Current Opinion in Neurobiology*, 11, 219–24.
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2, 568–573.
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10, 14–23.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, 108, 149–182.
- Huber, D. E. (2008b). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, 137, 324–347.
- Huber, D. E. (2008a). Causality in time: Explaining away the future and the past. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for rational models of cognition* (pp. 351–376). Oxford: Oxford University Press.
- Huber, D. E., & Cousineau, D. (2004). *A race model of perceptual forced choice reaction time*. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 687–692). Hillsdale, NJ: Erlbaum.
- Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science*, 27, 403–430.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning Memory and Cognition*, 28, 1120–1136.
- Huber, D. E., Tian, X., Curran, T., O'Reilly, C., & Woroch, B. (2008). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 1389–1416.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, 99, 480–517.
- Irwin, J. R., & McClelland, G. H. (2003). Negative consequences of dichotomizing continuous predictor variables. *Journal of Marketing Research (JMR)*, 40, 366–371.
- Kanwisher, N. (2000). Domain specificity in face perception. *Nature Neuroscience*, 3, 759–763.
- Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27, 117–143.
- Leder, H., & Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 53, 513–536.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4, 89–94.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.
- Maurer, D., Le Grand, R., & Mondlock, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6, 255–260.
- McCleery, J. P., Allman, E., Carver, L. J., & Dobkins, K. R. (2007). Abnormal magnocellular pathway visual processing in infants at risk for Autism. *Biological Psychiatry*, 62, 1007–1014.
- McKone, E., Kanwisher, N., & Duchaine, B. C. (2006). Can generic expertise explain special processing for faces? *Trends in Cognitive Sciences*, 11, 8–15.
- Merzenich, M. M., Jenkins, W. M., Johnston, P., Schreiner, C., Miller,

- S. L., & Tallal, P. (1996). Temporal processing deficits of language-learning impaired children ameliorated by training. *Science*, 271, 77–81.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227–234.
- Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1975). In P. M. A. Rabbitt (Ed.), *Attention and performance V* (pp. 98–118). London: Academic Press.
- Mosher, J. C., Baillet, S., & Leahy, R. M. (1999). EEG source localization and imaging using multiple signal classification approaches. *Journal of Clinical Neurophysiology*, 16, 225–238.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience*. Cambridge, MA: MIT Press.
- Quelen, F., Grainger, J., & Raymond, P. (2005). An investigation of semantic priming in schizophrenia using a new priming paradigm. *Schizophrenia Research*, 80, 173–183.
- Rakover, S. S. (2002). Featural vs. configurational information in faces: A conceptual and empirical analysis. *British Journal of Psychology*, 93, 1–30.
- Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience*, 4, 193–202.
- Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, 104, 319–343.
- Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G., & Nakayama, K. (2003). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science*, 14, 558–566.
- Richler, J. J., Tanaka, J. W., Brown, D. D., & Gauthier, I. (2008). Why does selective attention to parts fail in face processing? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 34, 1356–1368.
- Rossion, B., Delvenne, J. F., Debatisse, D., Goffaux, V., Bruyer, R., Crommelinck, M., & Guérin, J. M. (1999). Spatio-temporal localization of the face inversion effect: An event-related potentials study. *Biological Psychology*, 50, 173–189.
- Rossion, B., & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews*, 1, 63–75.
- Rotshtein, P., Henson, R. N. A., Treves, A., Driver, J., & Dolan, R. J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nature Neuroscience*, 8, 107–113.
- Rumelhart, D. E., & McClelland, J. L. (1986a). *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1: Foundations*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., & McClelland, J. L. (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 2: Psychological and biological models*. Cambridge, MA: MIT Press.
- Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16, 252–257.
- Singer, W. (1999). Neuronal synchrony: A versatile code review for the definition of relations? *Neuron*, 24, 49–65.
- Tanaka, J. W., & Pierce, L. J. (2009). The neural plasticity of other-race face recognition. *Cognitive, Affective, & Behavioral Neuroscience*, 9, 122–131.
- Tanaka, J. W. A race face database: Standardized set of over 200 Caucasian, African American, Asian and Hispanic faces. Department of Psychology, University of Victoria, Victoria, BC, Canada.
- Thomson, A. M., & West, D. C. (1993). Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neuroscience*, 54, 329–346.
- Tian, X., & Huber, D. E. (in press). Testing an associative account of semantic satiation. *Cognitive Psychology*.
- Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *The Quarterly Journal of Experimental Psychology Section A*, 37, 571–590.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94, 719–723.
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2005). Confusion and compensation in visual perception: Effects of spatiotemporal proximity and selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 40–61.
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: Is primed evidence discounted, even when it reliably indicates the correct answer? *Journal of Experimental Psychology: Learning Memory and Cognition*, 34, 257–281.
- Wenger, M. J., & Ingvalson, E. M. (2002). A decisional component of holistic encoding. *Journal of Experimental Psychology Learning Memory and Cognition*, 28, 872–892.
- Wenger, M. J., & Ingvalson, E. M. (2003). Preserving informational separability and violating decisional separability in facial perception and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1106–1118.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81, 141–145.

Appendix A

To facilitate future modeling research, we provide mean accuracy and standard errors (in parentheses) for all conditions in all

experiments collapsed across participants, as well as broken down by fast/slow group.

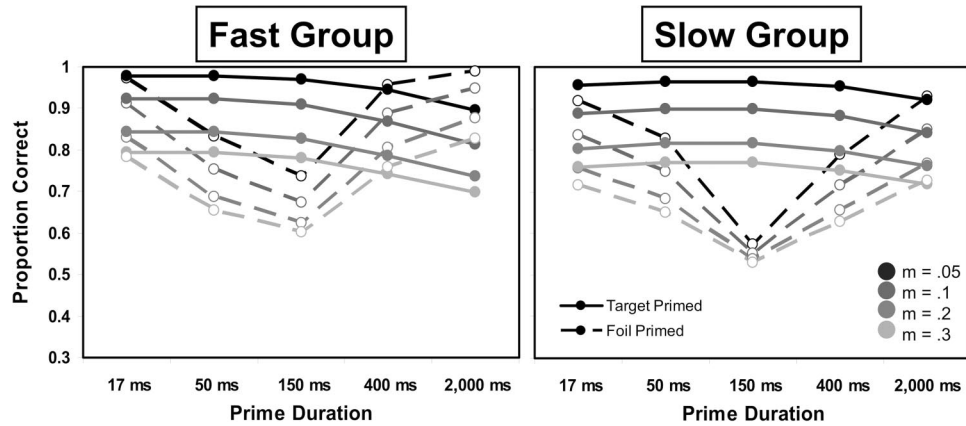


Figure A1. Model accuracy for variations in parameter m while holding other parameters to the best fit values.

Table A1

Experiment 1, Neither Primed Results

Target duration	Fast group	Slow group	All
33 ms	.66 (.03)	.48 (.03)	.57 (.03)
50 ms	.86 (.02)	.53 (.03)	.69 (.04)
67 ms	.90 (.02)	.66 (.05)	.78 (.03)
83 ms	.94 (.02)	.78 (.04)	.86 (.03)
<i>N</i>	14	14	28

Table A2

Experiment 1, Priming Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.67 (.04)	.82 (.02)	.75 (.03)
50 ms	.43 (.03)	.68 (.03)	.56 (.03)
150 ms	.48 (.03)	.56 (.05)	.52 (.03)
400 ms	.74 (.04)	.70 (.08)	.72 (.04)
2,000 ms	.86 (.03)	.75 (.08)	.81 (.04)
Target primed			
17 ms	.75 (.05)	.84 (.02)	.80 (.03)
50 ms	.73 (.05)	.88 (.02)	.80 (.03)
150 ms	.76 (.05)	.85 (.02)	.81 (.03)
400 ms	.73 (.05)	.86 (.03)	.79 (.03)
2,000 ms	.66 (.06)	.87 (.02)	.76 (.04)
<i>N</i>	14	14	28

(Appendix continues)

Table A3
Experiment 2, 33-ms Target Duration Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.59 (.03)	.50 (.02)	.54 (.02)
150 ms	.48 (.04)	.30 (.03)	.39 (.03)
2,000 ms	.65 (.04)	.41 (.04)	.53 (.03)
Target primed			
17 ms	.84 (.07)	.73 (.02)	.78 (.02)
150 ms	.85 (.10)	.77 (.03)	.81 (.02)
2,000 ms	.80 (.12)	.75 (.03)	.78 (.02)
Neither primed			
150 ms	.67 (.02)	.54 (.01)	.60 (.01)
<i>N</i>	28	28	56

Table A4
Experiment 2, 50-ms Target Duration Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.75 (.03)	.58 (.03)	.66 (.02)
150 ms	.66 (.04)	.40 (.03)	.52 (.03)
2,000 ms	.86 (.02)	.56 (.05)	.70 (.03)
Target primed			
17 ms	.91 (.01)	.75 (.03)	.83 (.02)
150 ms	.89 (.02)	.78 (.03)	.84 (.02)
2,000 ms	.86 (.02)	.77 (.03)	.81 (.02)
Neither primed			
150 ms	.80 (.02)	.58 (.02)	.69 (.02)
<i>N</i>	28	28	56

Table A5
Experiment 2, 100-ms Target Duration Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.93 (.01)	.76 (.03)	.84 (.02)
150 ms	.91 (.02)	.62 (.04)	.76 (.03)
2,000 ms	.96 (.01)	.72 (.05)	.83 (.03)
Target primed			
17 ms	.96 (.01)	.83 (.02)	.89 (.02)
150 ms	.95 (.01)	.80 (.03)	.87 (.02)
2,000 ms	.95 (.01)	.84 (.02)	.89 (.01)
Neither primed			
150 ms	.94 (.01)	.74 (.03)	.84 (.02)
<i>N</i>	28	28	56

Table A6
Experiment 3, Neither Primed Results

Target duration	Fast group	Slow group	All
33 ms	.67 (.02)	.54 (.02)	.61 (.01)
50 ms	.81 (.01)	.57 (.01)	.69 (.02)
67 ms	.84 (.02)	.71 (.02)	.78 (.01)
83 ms	.90 (.01)	.81 (.02)	.85 (.01)
<i>N</i>	51	51	102

Table A7
Experiment 3, Priming Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.68 (.02)	.75 (.02)	.75 (.01)
50 ms	.52 (.02)	.64 (.02)	.64 (.02)
150 ms	.42 (.02)	.49 (.02)	.49 (.02)
400 ms	.63 (.02)	.70 (.02)	.70 (.02)
2,000 ms	.75 (.03)	.76 (.03)	.76 (.02)
Target primed			
17 ms	.81 (.02)	.82 (.02)	.82 (.01)
50 ms	.83 (.02)	.85 (.02)	.85 (.01)
150 ms	.82 (.02)	.86 (.02)	.86 (.01)
400 ms	.78 (.02)	.85 (.02)	.85 (.01)
2,000 ms	.73 (.02)	.83 (.02)	.83 (.02)
Neither primed			
17 ms	.74 (.02)	.79 (.02)	.76 (.01)
50 ms	.66 (.02)	.73 (.02)	.69 (.01)
150 ms	.64 (.02)	.71 (.02)	.68 (.01)
400 ms	.71 (.02)	.76 (.02)	.73 (.01)
2,000 ms	.75 (.02)	.76 (.02)	.76 (.01)
<i>N</i>	51	51	102

Table A8
Experiment 4, Neither Primed Results

Target duration	Fast group	Slow group	All
Upright faces			
33 ms	.54 (.03)	.52 (.03)	.53 (.02)
50 ms	.74 (.04)	.61 (.04)	.67 (.03)
67 ms	.88 (.03)	.59 (.04)	.73 (.03)
83 ms	.90 (.03)	.71 (.03)	.80 (.03)
Inverted faces			
33 ms	.56 (.04)	.48 (.04)	.52 (.03)
50 ms	.61 (.04)	.49 (.03)	.55 (.03)
67 ms	.73 (.04)	.51 (.04)	.62 (.03)
83 ms	.83 (.04)	.69 (.04)	.76 (.03)
<i>N</i>	20	20	40

Table A9
Experiment 4, Upright Face Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.74 (.04)	.86 (.03)	.80 (.02)
50 ms	.54 (.02)	.78 (.02)	.66 (.03)
150 ms	.51 (.04)	.63 (.04)	.57 (.03)
400 ms	.80 (.03)	.74 (.04)	.77 (.03)
2,000 ms	.91 (.02)	.85 (.05)	.88 (.03)
Target primed			
17 ms	.83 (.03)	.88 (.02)	.86 (.02)
50 ms	.87 (.03)	.89 (.02)	.88 (.02)
150 ms	.80 (.03)	.88 (.02)	.84 (.02)
400 ms	.80 (.04)	.93 (.02)	.87 (.02)
2,000 ms	.76 (.04)	.88 (.03)	.82 (.02)

(Appendix continues)

Table A10
Experiment 4, Inverted Face Results

Prime duration	Fast group	Slow group	All
Foil primed			
17 ms	.66 (.03)	.78 (.02)	.72 (.02)
50 ms	.47 (.04)	.67 (.03)	.57 (.03)
150 ms	.39 (.03)	.50 (.04)	.45 (.03)
400 ms	.58 (.04)	.66 (.05)	.62 (.03)
2,000 ms	.67 (.04)	.74 (.05)	.70 (.03)
Target primed			
17 ms	.77 (.03)	.83 (.03)	.80 (.02)
50 ms	.77 (.03)	.83 (.03)	.80 (.02)
150 ms	.77 (.04)	.84 (.02)	.80 (.02)
400 ms	.76 (.04)	.88 (.03)	.82 (.03)
2,000 ms	.70 (.03)	.82 (.03)	.76 (.02)
<i>N</i>	20	20	40

Received April 5, 2008
Revision received October 24, 2009
Accepted November 2, 2009 ■