

Are You Smiling, or Have I Seen You Before? Familiarity Makes Faces Look Happier

Psychological Science
1–16
© The Author(s) 2017
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0956797617702003
www.psychologicalscience.org/PS


Evan W. Carr¹, Timothy F. Brady², and Piotr Winkielman^{2,3,4}

¹Columbia Business School, Columbia University; ²Psychology Department, University of California, San Diego;

³Behavioural Science Group, Warwick Business School, University of Warwick; and ⁴Faculty of Psychology, SWPS University of Social Sciences and Humanities

Abstract

It is clear that unreinforced repetition (familiarization) influences affective responses to social stimuli, but its effects on the perception of facial emotion are unknown. Reporting the results of two experiments, we show for the first time that repeated exposure enhances the perceived happiness of facial expressions. In Experiment 1, using a paradigm in which subjects' responses were orthogonal to happiness in order to avoid response biases, we found that faces of individuals who had previously been shown were deemed happier than novel faces. In Experiment 2, we replicated this effect with a rapid "happy or angry" categorization task. Using psychometric function fitting, we found that for subjects to classify a face as happy, they needed less actual happiness to be present in the face if the target was familiar than if it was novel. Critically, our results suggest that familiar faces appear happier than novel faces because familiarity selectively enhances the impact of positive stimulus features.

Keywords

familiarity, affect, perception, facial expressions, exposure

Received 7/11/16; Revision accepted 3/7/17

Psychologists have been fascinated with familiarity for more than a century (Fechner, 1876). Titchener (1910), for example, characterized the feeling of familiarity as a "warm glow" (p. 408). Dominant explanations of this feeling propose that familiarization (via unreinforced repetition) associates the stimulus with an absence of negative consequences (Zajonc, 2001) and reduces uncertainty (Lee, 2001), or that repetition facilitates processing (Bornstein & D'Agostino, 1994) and such fluency is experienced as positive (Winkielman, Schwarz, Fazendeiro, & Reber, 2003). The preference for previously encountered stimuli has been well documented across many tasks, modalities, and stimuli (e.g., the classic *mere-exposure effect*; Zajonc, 1968). Surprisingly, however, the link between familiarity and actual perception of emotion has been unexplored.

One key question concerns the processing stage at which familiarity creates positivity: Does familiarity affect early perceptual stages of stimulus processing or only late-stage judgments? Another key question concerns the nature of any changes in valence: Does familiarity change reactions to positive features, negative features, or both?

In this article, we explore these key questions and present the results of experiments testing several novel predictions using important social stimuli—emotional facial expressions. We propose that familiarity enhances the perceived happiness of facial expressions, and that this effect involves the selective enhancement of positive stimulus features (rather than the weakening of negative features).

These predictions are grounded in several lines of previous research on familiarity. Past studies have found that stimulus repetition increases a variety of preference judgments. Such experiments have often used ratings of liking or attractiveness as the dependent variable, but preliminary evidence suggests that the effect holds with ratings of happiness, at least for ratings of neutral faces (Claypool, Hugenberg, Housley, & Mackie, 2007). This

Corresponding Author:

Evan W. Carr, Columbia Business School, 3022 Broadway, 7L Uris,
New York, NY 10027

E-mail: ewc2138@columbia.edu

raises the possibility that familiarity actually makes faces “look” happier. This idea fits with evidence that the mere-exposure effect can be detected with physiological measures of affect (Harmon-Jones & Allen, 2001) and in tasks without explicit evaluative ratings (Garcia-Marques, Prada, & Mackie, 2016). However, the extant research is silent regarding whether familiarity affects early perceptual processing of facial expressions and whether any changes involve positive features, negative features, or both.

The Current Research

We tested how familiarity with another individual affects rapid perceptual decisions (Experiment 1) and classification judgments (Experiment 2) concerning that person’s emotional facial expressions. We first manipulated familiarity by systematically exposing subjects to neutral expressions of certain individuals but not others. Next, subjects judged the level of happiness in face blends (morphs on a continuum from angry to happy) from both familiar and unfamiliar individuals. Our paradigms are related to those used in previous work suggesting that there are tight links between observers’ affect and their perception (including their perception of facial expressions; e.g., Phelps, Ling, & Carrasco, 2006). Our studies were also designed to dissociate predictions from four frameworks regarding the connection between familiarity and responses to valenced facial features (see Table 1).

First, the *amplification* account (which is similar to theories of *nonspecific activation*; Albrecht & Carbon, 2014; Mandler, Nakamura, & Van Zandt, 1987) assumes that repetition intensifies the appearance of dominant stimulus features, regardless of their valence, such that the apparent positive features become more positive, and the apparent negative features become more negative. This account suggests that, in our experiments, exposure would simply augment perception of the existing valenced features in the facial expressions—so that happy expressions of familiar targets would appear happier than happy expressions of unfamiliar targets, but angry expressions of familiar targets would also appear angrier than angry expressions of unfamiliar targets.

Second, the *generalized-positivity-shift* account assumes that familiarity elicits broad positive affect that imbues positivity to all stimuli, regardless of their intrinsic valence. This idea is implicit in the notion of a generalized “warm glow” (Titchener, 1910), and some psychologists have proposed that such a glow functions the way positive mood does, making everything better (Monin, 2003; Schwarz & Clore, 2003). This account suggests that, in our experiments, exposure would make both happy and angry expressions of familiar targets appear happier than corresponding expressions of unfamiliar targets.

The third and fourth accounts, unlike the first two, assume that familiarity has separable effects on perceived positive and negative valence. According to the *negative-skew* account, familiarity selectively dampens negative

Table 1. Summary of Four Alternative Frameworks Proposing a Relationship Between Familiarity and Perception of Valenced Facial Features

Framework	Main prediction	Does familiarity differentially affect responses to positive and negative faces?	Key references
Amplification (nonspecific activation)	Familiarity intensifies the appearance of already-existing stimulus features, regardless of their valence (i.e., positive features are perceived as more positive, and negative features are perceived as more negative).	No	Albrecht and Carbon (2014); Mandler, Nakamura, and Van Zandt (1987)
Generalized positivity shift	Familiarity elicits broad positive affect that imbues positivity to all stimuli, regardless of their valence (i.e., everything familiar is perceived as more positive).	No	Monin (2003); Schwarz and Clore (2003); Titchener (1910)
Negative skew	Familiarity dampens negativity and reduces uncertainty, without enhancing positive features (i.e., negative features are perceived as less negative).	Yes	Lee (2001); Zajonc, Markus, and Wilson (1974)
Hedonic skew	Familiarity selectively enhances positive features but not negative features (i.e., positive features are perceived as more positive).	Yes	Garcia-Marques, Mackie, Claypool, and Garcia-Marques (2004); Harmon-Jones and Allen (2001); Winkielman and Cacioppo (2001)

Note: These frameworks do not necessarily account for all familiarity effects (e.g., general mere-exposure effects). We describe them here in the context of our specific focus on how familiarity might affect the perception of facial emotion.

responses, without enhancing positive responses. This view is consistent with results from some of Zajonc's early animal research, which suggested that familiarity reduces initial distress to novelty (Zajonc, Markus, & Wilson, 1974), and also aligns with models in which repetition reduces initially unpleasant uncertainty (Lee, 2001). This framework predicts that familiarity has the greatest benefit for stimuli with negative (and possibly neutral) valence. Thus, it suggests that, in our experiments, angry faces (and possibly neutral faces) of familiar targets would appear less angry than angry faces of unfamiliar targets, but that familiarity would not affect perception of happy faces.

Finally, according to the *hedonic-skew* account, familiarity selectively enhances responses to positive features. This account suggests that, in the current study, happy expressions of familiar individuals would appear happier than happy expressions of unfamiliar individuals, but familiarity would have little effect on perception of angry expressions. Theoretically, these differential effects could involve attribution of target-dependent features; that is, positive affect from familiarity could reasonably be attributed only to positive (and possibly neutral) facial features (Schwarz, 2014). This type of effect would be related to perceptual aftereffects of early visual adaptation that are expressed only when the correct stimulus features are present. For example, aftereffects of color perception can depend on the orientation of lines displayed in a test grating (e.g., the aftereffect will occur only with a grating that has horizontal lines and not vertical lines, or vice versa; McCollough, 1965). If effects of familiarity on emotion perception emerge in a similar way, faces from familiar people might appear happier, but only when positive features are already embedded in those faces. Moreover, if the effects of familiarity are stimulus dependent, this could also indicate that familiarity has dissociable effects on the positive- and negative-affect systems (Cacioppo & Berntson, 1994), which would be consistent with other studies demonstrating that familiarity specifically increases positive affect, and does not reduce negative affect (Garcia-Marques, Mackie, Claypool, & Garcia-Marques, 2004; Harmon-Jones & Allen, 2001; Winkielman & Cacioppo, 2001).

Experiment 1

To test these alternative frameworks in Experiment 1, we adapted a paradigm designed to examine influences on perception independently of decision and response biases. This paradigm was originally developed by Carrasco, Ling, and Read (2004) and was adapted to face stimuli by Störmer and Alvarez (2016). After exposing subjects to neutral expressions of certain individuals but not others (thus systematically manipulating familiarity), we gauged how familiarity affected subjects' perceptual judgments of happiness in facial expressions of these

familiar individuals and novel individuals. A familiar (trained) and a novel (untrained) face with objectively the same expression were presented on each trial. Our main dependent measure was the frequency with which subjects selected the trained face as happier than the untrained face. Figure 1 displays hypothetical data for this experiment that are consistent with the predictions of each of the four frameworks.

Method

Subjects. Fifty undergraduates (mean age = 20.90 years, $SD = 5.02$ years; 35 females) from the University of California, San Diego (UCSD), participated for course credit. All signed consent forms approved by the Human Research Protection Program at UCSD. We planned our sample size on the basis of a priori power calculations and in accordance with previous studies on perceptual judgments for faces (e.g., Störmer & Alvarez, 2016). Using G*Power (Version 3.1.9.2; Faul, Erdfelder, Lang, & Buchner, 2007), we estimated that with a total sample of 41 to 67 subjects, we would have 80% power to detect a small-to-medium effect, $d_z = 0.35$ – 0.45 , given a two-tailed test and α level of .05. We therefore targeted a sample size of 50.

Materials

Stimulus preparation. We created our facial stimuli using still images from the Amsterdam Dynamic Facial Expression Set (ADFES; van der Schalk, Hawk, Fischer, & Doosje, 2011). From the ADFES, we selected 12 different models to use for morphing (6 males and 6 females). Using the 100%-angry, 100%-happy, and neutral images for each model, we generated a morph continuum at five different levels: 50% angry, 25% angry, neutral (the original ADFES image), 25% happy, and 50% happy. This created a set of 60 unique stimuli (12 different models, each displaying five different levels of emotion). Note that all the stimuli were single-person morphs; that is, images of different individuals were never blended together. All the faces were then cropped so that only the facial features were visible.

We then divided the stimuli into two sets (A and B), each of which contained the images of a different half of the models (three males and three females). One set was presented during training (Phase 1); whether this was Set A or Set B varied across subjects. Both sets were presented in a follow-up task in which subjects judged the models' emotional expressions (Phase 2). Thus, in Phase 2, each subject had to respond to each of the models' emotional expressions, but a given model was either trained or untrained, depending on the subject's exposure set. For example, if a subject was assigned to study the models in set A, he or she was exposed to neutral expressions of those models during training and then

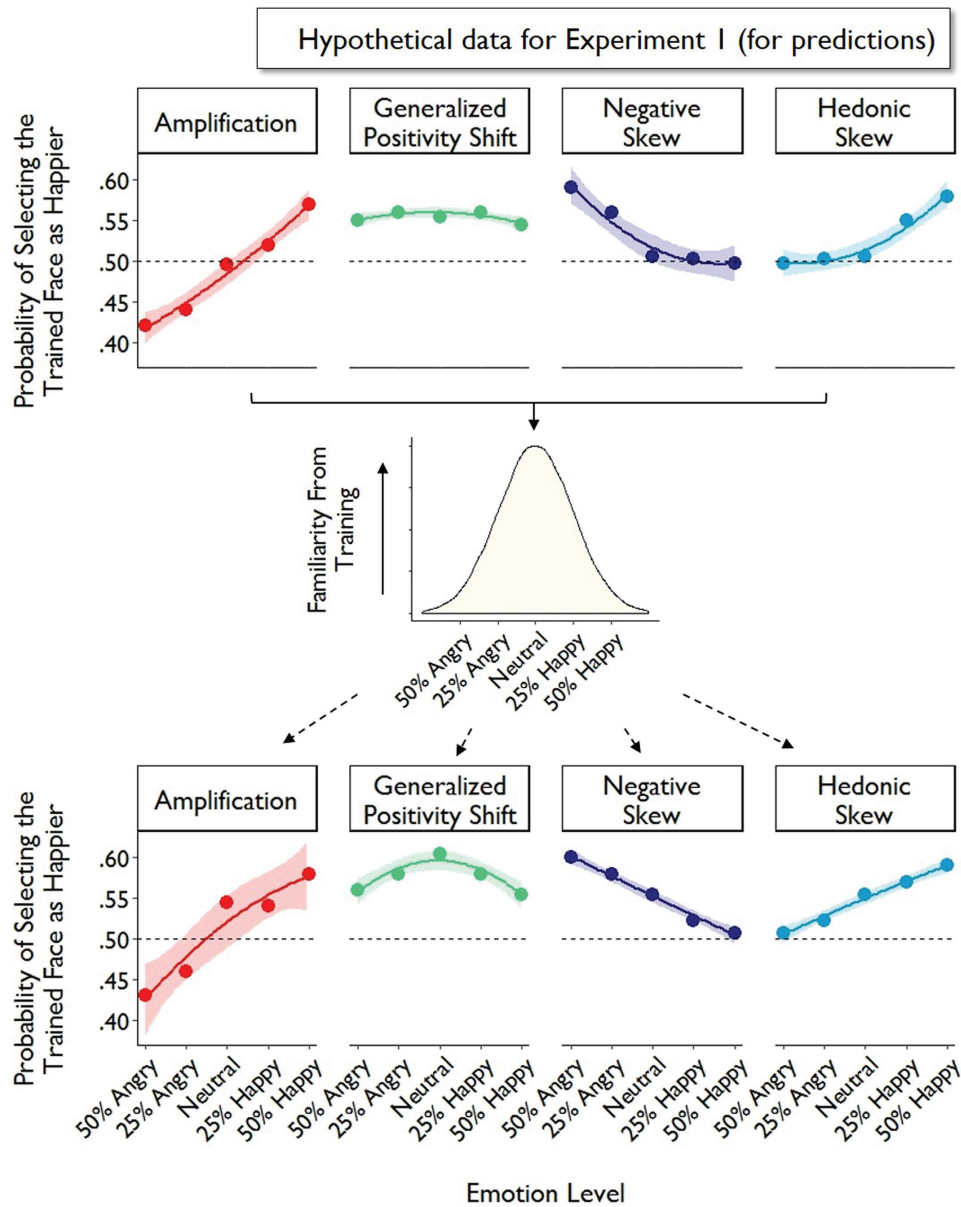


Fig. 1. Illustration of the four frameworks' qualitative predictions for Experiment 1. The graph at the top shows hypothetical data that would support each of these frameworks, before accounting for the fact that subjects were familiarized with neutral expressions only. Because subjects were trained on neutral expressions, each of these functions would be filtered through a familiarity distribution (middle panel). That is, because the neutral expressions were exact replicates of the stimuli shown during training, they would have the greatest familiarity (and hence positivity), and familiarity and positivity would decrease as the expressions deviated farther from neutral (here, we assume a simple similarity gradient based on literature about mere-exposure generalization effects; Gordon & Holyoak, 1983). The integrated predictions of the four frameworks after the original hypothetical data are combined with the familiarity distribution are shown in the bottom panel.

saw the emotion morphs and neutral expressions of both set A's and set B's models in a follow-up task (Phase 2).

Stimulus norming. Our predictions for this experiment were based on the assumption that subjects would perceive the happy and angry morphs as positive and negative (respectively). Indeed, Wingenbach, Ashwin, and

Brosnan (2016) had already mostly verified this assumption with stimuli similar to those we used. They created low-, intermediate-, and high-intensity versions of the ADFES facial expressions, which they termed the Bath Intensity Variations (ADFES-BIV), by extracting from the ADFES videos consecutive frame sequences that started with a neutral frame and continued to the full apex of an

expression. The low-intensity category included subtle expressions (i.e., only limited magnitudes of facial action units). The high-intensity category included the apexes of the emotional expressions within the videos (i.e., maximal values of all the relevant facial action units). The stimuli in the intermediate-intensity category were at the midpoint between the low- and high-intensity stimuli. Therefore, our 25% and 50% morph stimuli roughly corresponded to the low- and intermediate-intensity stimuli in the ADFES-BIV, respectively. Subjects were highly accurate (> 55%)—performing far above chance level (10%)—on emotion-recognition tasks with happy and angry ADFES-BIV expressions at all three intensity levels. Subjects also responded quickly, within approximately 1,000 ms, to all expressions at all three intensities. Note that this high level of performance was achieved even when subjects did not simply indicate the valence of the expressions (i.e., positive or negative), but had the more difficult task of categorizing specific emotions (i.e., whether the faces expressed anger, contempt, disgust, embarrassment, fear, happiness, pride, sadness, or surprise or were instead neutral).

To follow up on the results from Wingenbach et al. (2016), we ran an online norming study on our own stimuli, using the same UCSD subject pool as in the main experiment ($N = 102$; mean age = 20.16 years, $SD = 1.97$ years; 62 females). Subjects randomly cycled through all 60 unique stimuli and used sliding scales (0–100) to indicate the extent to which they thought each face was positive or negative (0 = *completely negative*, 100 = *completely positive*) and how confident they were in their valence rating (0 = *not at all confident*, 100 = *completely confident*). We included confidence ratings so that we could assess whether subjects were more certain about negative ratings of angry faces than about positive ratings of happy faces (or vice versa). We analyzed the valence and confidence ratings for the online norming study using multilevel models (MLMs). Both models were fitted with stimulus emotion (five levels: 50% angry, 25% angry, neutral, 25% happy, 50% happy; within subjects) as the only fixed-effects factor; maximal random effects were fitted on both subjects' and models' identities. We found main effects of stimulus emotion for both valence ratings, $F(4, 16.16) = 124.41, p < .001$, and confidence ratings, $F(4, 20.58) = 30.18, p < .001$. A parallel repeated measures analysis of variance (ANOVA) showed similar main effects for both valence ratings, $F(4, 404) = 654.23, p < .001, \eta^2 = .83$, and confidence ratings, $F(4, 404) = 78.70, p < .001, \eta^2 = .11$.

The results for valence demonstrated that subjects judged 25%-angry expressions ($M = 38.08, SD = 6.36$; 95% confidence interval, $CI = [33.39, 42.78]$) and 50%-angry expressions ($M = 27.32, SD = 9.49, 95\% CI = [22.80, 31.83]$) as significantly more negative than the scale's midpoint, $t(101) = -18.93, p < .001, d_z = 1.87$, and $t(101) = -24.13,$

$p < .001, d_z = 2.39$, respectively. Similarly, subjects judged both the 25%-happy expressions ($M = 59.85, SD = 8.28, 95\% CI = [55.03, 64.66]$) and the 50%-happy expressions ($M = 79.53, SD = 11.31, 95\% CI = [75.97, 83.10]$) as more positive than the scale's midpoint, $t(101) = 12.02, p < .001, d_z = 1.19$, and $t(101) = 26.37, p < .001, d_z = 2.61$, respectively. Further, the valence ratings for the 25%-happy, 50%-happy, 25%-angry, and 50%-angry expressions were also each significantly different from the valence ratings for the neutral expressions ($M = 44.91, SD = 5.13, 95\% CI = [40.61, 49.21]$), $|b|s \geq 6.82, |t|s(12.0-27.3) \geq 4.53, ps < .001, d_zs \geq 0.45$.

The results for confidence showed that subjects were similarly confident in their negative ratings of angry faces and their positive ratings of happy faces. Not surprisingly, they were the least confident in their valence ratings of neutral expressions ($M = 74.30, SD = 13.37, 95\% CI = [71.66, 76.95]$), but their confidence increased when they rated the valence of the 25%-happy ($M = 75.82, SD = 13.32, 95\% CI = [72.55, 79.09]$) and 25%-angry ($M = 75.97, SD = 12.54, 95\% CI = [73.28, 78.66]$) expressions, and increased further when they rated the valence of the 50%-happy ($M = 86.28, SD = 10.73, 95\% CI = [83.53, 89.02]$) and 50%-angry ($M = 80.34, SD = 11.81, 95\% CI = [77.46, 83.22]$) expressions. These results demonstrate that subjects had high confidence in their valence ratings of expressions at all positions along the morph continuum (all $M_s \geq 74.30$ out of 100) and that subjects had comparable levels of certainty when judging the positive and negative valence of our happy and angry stimuli.

Given these results and the work of Wingenbach et al. (2016), we are confident that our subjects perceived the valence of our face stimuli as we intended.

Design and procedure. Experiment 1 consisted of a training task (Phase 1), in which subjects were exposed to some of the face stimuli, and a follow-up task (Phase 2), in which subjects made speeded perceptual judgments on all the face stimuli. The stimuli for these tasks were presented using E-Prime 2.0 software (Psychology Software Tools, 2012) on 17-in. Dell flat-screen PCs running Windows 7 (resolution of 1,280 × 1,024 pixels; 60-Hz refresh rate).

Prior to performing the training task, all subjects were told that they would be completing a memory task in which they would have to track and recall the color and number of blue and green square probes that would appear randomly on various images. These images were the neutral faces for the specific models in the set (A or B) that subjects had been randomly assigned to study (also see Carr et al., 2017).

The task consisted of 20 exposure trials with each of the six models in the assigned training set (total of 120 training trials). Before each trial, subjects saw a prompt that said, "Remember the color and number of squares!"

After the prompt, they pressed a button to trigger the start of the trial. During the trial, a neutral expression for one of the training models appeared in the center of the screen for 5,000 ms. Depending on the trial, some combination of blue squares, green squares, both or neither flashed on the face for 200 ms each at random intervals. Squares of each color could appear anywhere from zero to nine times with equal probabilities (so it was possible for no squares or only one square of one color to be shown on a given trial). After the trial ended, another screen asked subjects, “How many BLUE squares did you see (0-9)?” and “How many GREEN squares did you see (0-9)?” To encourage high attention and effort throughout this task, we told subjects that they would advance to the next phase of the experiment only after they hit a satisfactory level of performance (in reality, all subjects completed the same number of training trials so that exposure to the faces would be consistent). With this task, we were able to give subjects many passive exposures to neutral expressions for only certain models, thus giving them familiarity with some models but not others.

After subjects finished the 120 training trials in Phase 1, they moved on to Phase 2. Our Phase 2 paradigm was a modified version of the attentional-cuing task used by Carrasco et al. (2004). This paradigm is usually used to measure the effects of exogenous attention on perceptual processing; although we were not interested in attention in the current study, a benefit of this type of task is that it controls for decision and response biases, and results obtained with this task have therefore been repeatedly used to support claims that observed effects have a perceptual locus. For example, in a study by Störmer and Alvarez (2016), two faces were presented simultaneously during each trial on the left and right sides of a computer screen; one face was shifted upward and the other was shifted downward along the vertical axis. Subjects’ task was to report whether the face they perceived as more attractive was shifted upward or downward, using the up- or down-arrow key on the keyboard. The fact that the response was orthogonal to the dimension of interest eliminated the possibility that the observed effects were due to a simple response bias and reduced the likelihood that they originated in late decision-making stages. We adapted this paradigm to investigate how our Phase 1 training task affected speeded perceptual judgments of happiness in trained (familiar) relative to untrained (novel) faces.

Figure 2 shows a schematic of our Phase 2 task. Each trial began with a prompt reminding subjects to report whether they thought the happier face was above or below the line (by pressing the up- or down-arrow key, respectively). They then pressed a key to trigger the onset of a fixation cross that appeared for 750 ms. Two lines then appeared on the left and right sides of the fixation cross for 500 ms, to mark the horizontal axis of the

screen. Next, two faces were displayed on the left and right sides of the screen (one face 128 pixels to the left and the other 128 pixels to the right of the center fixation); one face was shifted slightly upward (154 pixels), and the other face was shifted slightly downward (154 pixels) from the center fixation. If subjects thought the upper face appeared happier, they pressed the up-arrow key; if they thought the lower face appeared happier, they pressed the down-arrow key. Subjects were given up to 3,000 ms to respond (any responses not logged within this time were excluded from analysis), and after they gave a response, a 1,000-ms response-confirmation screen was displayed before the next trial.

Critically, we manipulated the types of faces shown on each trial. We always displayed one trained (familiar) model and one untrained (novel) model, and the two faces always displayed the same objective level of emotion (i.e., 50% angry, 25% angry, neutral, 25% happy, or 50% happy). Therefore, on each trial, no response could be considered correct or incorrect, given that both faces displayed the same type and level of emotion. We were interested in how subjects’ training with certain faces would influence their perceptual judgments of happiness. If the training made certain faces appear happier, subjects would choose the trained faces consistently more often than the untrained faces (regardless of their spatial location on the screen). If training made certain faces appear angrier, subjects would choose the trained faces consistently less often (regardless of their spatial location on the screen). The pattern of results would be informative as to how familiarity affects the perception of valence (i.e., the predictions of the four previously described frameworks; see Table 1 and Fig. 1). Our instructions emphasized that there were no correct or incorrect answers, and subjects were told that the tasks in Phases 1 and 2 were unrelated.

This Phase 2 task consisted of six blocks of 60 trials each. Each of the six trained faces was matched twice with each of the six untrained faces at each of the five levels of emotion, appearing once on the left and once on the right (i.e., 6 trained faces \times 6 untrained faces \times 5 emotion levels \times 2 display positions = 360 trials). In order to get accustomed to the Phase 2 task, subjects completed 8 practice trials using ADFES models that were not incorporated in either the Phase 1 or the Phase 2 task.

Results

All repeated measures analyses used MLMs and restricted maximum likelihood, because this method offers numerous analytical advantages over traditional ANOVA methods—including more effective handling of unbalanced data with missing observations, reliance on fewer assumptions regarding covariance structures, and increased parsimony and flexibility in specifying model parameters (Bagiella, Sloan, & Heitjan, 2000). Note that

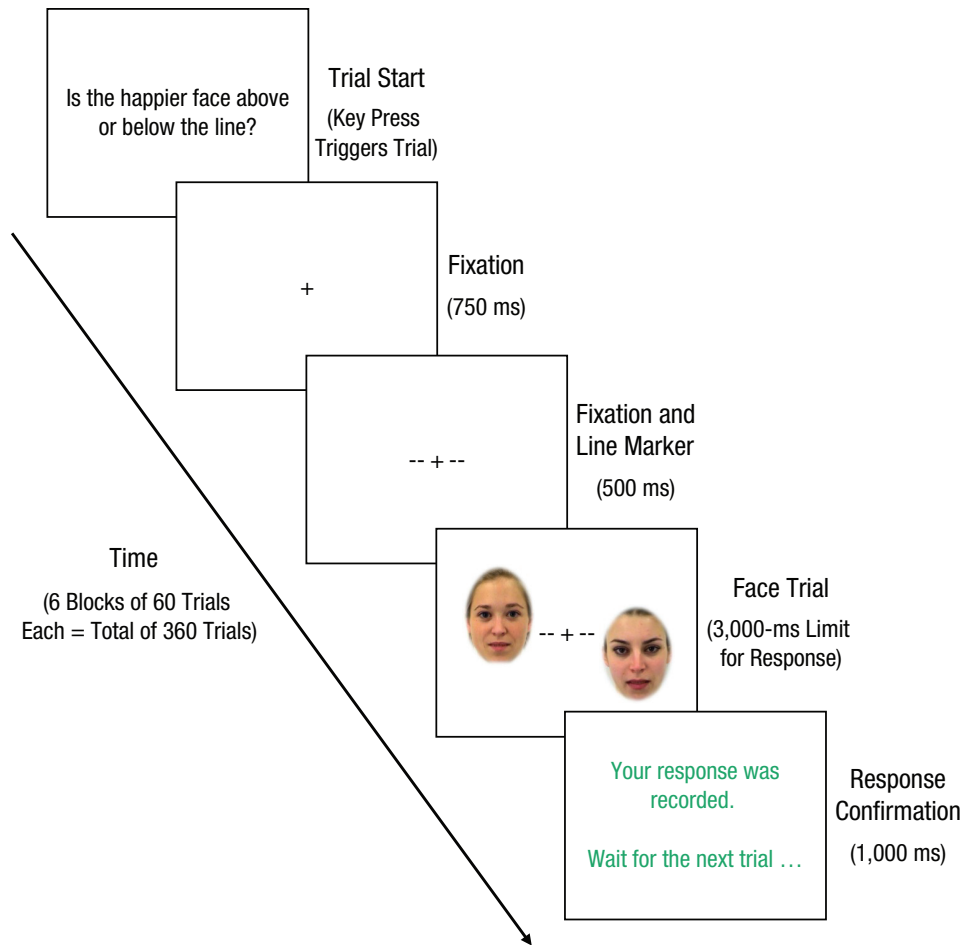


Fig. 2. Design and procedure for the Phase 2 task in Experiment 1. Subjects were reminded that their task was to report which of two upcoming faces was happier, and then a fixation cross appeared. Next, a line marked the horizontal axis of the screen, and then the screen displayed two faces, one shifted upward and the other shifted downward relative to the horizontal axis. After subjects pressed the up- or down-arrow key to record their response, a response-confirmation screen ended the trial.

with our trial-level data, we had missing observations (we excluded all trials with response times, RTs, less than 200 ms) and could model stimuli as random effects (i.e., of the different models in our face stimuli). MLMs are well suited to handle both of these issues because they provide more powerful estimates of the effects in question (Judd, Westfall, & Kenny, 2012). All models were built with the *lmerTest* package in R (Kuznetsova, Brockhoff, & Christensen, 2016), and we used the maximal random-effect structure on subject identities that would allow for model convergence after 10,000 iterations (West, Welch, & Galecki, 2014). Stimulus identities were not modeled as random effects in this experiment because two stimuli were shown on each trial. To obtain p -value estimates for fixed effects, we used Type III Satterthwaite approximations, which can sometimes result in decimal degrees of freedom (Kuznetsova et al., 2016). We also report results from parallel repeated measures ANOVAs.

We analyzed the probability of the trained face being selected as the happier one using an MLM with stimulus emotion (five levels: 50% angry, 25% angry, neutral, 25% happy, 50% happy; within subjects) as a fixed-effects factor. Figure 3 shows the mean probability of the trained face being chosen as happier than the untrained face at each level of emotion across all trials. We fitted a bootstrapped logistic psychometric function to these response probabilities, using the *quickpsy* package in R (Linares & Lopez-Moliner, 2016).

Overall, across emotion levels, subjects selected trained faces as happier more often than expected by chance (.50), 95% CI for the mean = [.51, .57], $t(49) = 2.35$, $p = .02$, $d_z = 0.33$, but the probability of the trained face being selected as happier also varied as a function of the emotion level, $F(4, 84.81) = 3.59$, $p = .009$. This main effect was confirmed by a repeated measures ANOVA, $F(4, 196) = 5.41$, $p < .001$, $\eta^2 = .03$. As the positive features in the faces

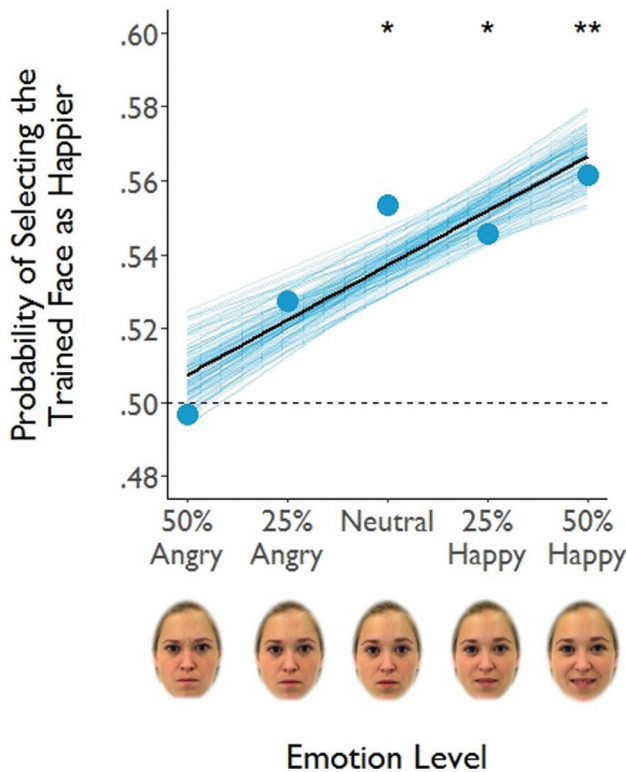


Fig. 3. Results for the Phase 2 task in Experiment 1: mean probability of the trained face being selected as happier than the untrained face at each of the five levels of emotion. The thin lines show the 100 bootstrapped samples from psychometric fitting of a logistic function, and the heavy black line is the final estimate of this curve. Although the final curve appears to be linear, it is indeed a logistic function. Asterisks indicate significant differences from the .50 chance level (* $p < .05$, ** $p < .01$).

increased (from 50% angry to 50% happy), subjects became more likely to judge the trained face as happier than the untrained face. Specifically, the probability of the trained face being selected did not differ from chance for 50%-angry faces, $t(49) = -0.35$, 95% CI for the mean = [0.47, 0.52], $d_z = 0.05$, or 25%-angry faces, $t(49) = 1.46$, 95% CI for the mean = [0.49, 0.57], $d_z = 0.21$, but the trained face was selected as the happier face significantly more often than chance level when the faces were neutral, $t(49) = 2.56$, 95% CI for the mean = [0.51, 0.59], $p = .01$, $d_z = 0.36$; 25% happy, $t(49) = 2.20$, 95% CI for the mean = [0.50, 0.58], $p = .03$, $d_z = 0.31$; or 50% happy, $t(49) = 3.35$, 95% CI for the mean = [0.52, 0.60], $p = .002$, $d_z = 0.47$. These results mirror the prediction of the hedonic-skew framework (see Table 1 and Fig. 1), which posits that the warm glow of familiarity operates via the selective enhancement of positive features.

Experiment 2

Experiment 1 demonstrated that familiarity influences perceptions of happiness, suggesting that positivity is

enhanced at early processing stages. The strength of this association between familiarity and positivity increased as the test expressions became happier (there were no effects for 25%- or 50%-angry morphs). In Experiment 2, we tested whether we could replicate and extend these effects using a rapid-categorization and judgment task. After training, a single face was presented on each test trial, and subjects quickly decided whether that face (from a familiar or novel individual) was “happy or angry.” The key benefit of this single-face design in Experiment 2 over the dual-face design in Experiment 1 is that it allowed for a direct measure of the happiness level of both trained and untrained faces at each level of emotion (via fitting separate psychometric functions for trained and untrained faces), rather than an indirect measure of the happiness level of trained faces relative to simultaneously presented untrained faces. Another consideration was that the perceptual task in Experiment 1 may have given subjects a goal of detecting positive features (because they were repeatedly asked about which face looked happier), and such a goal would likely be reduced (if not eliminated) in a binary “happy or angry” classification paradigm.

In Experiment 2, subjects also gave a percentage estimate (0–100%) for how happy they thought the face on each trial looked. These percentage ratings not only provided a secondary measure of perceived happiness, but also focused on a more deliberative judgment (no time limit), rather than a first impression (rapid categorization). Figure 4 displays hypothetical data for Experiment 2 that would be consistent with the four frameworks described in Table 1.

Method

Subjects, materials, and equipment. Forty undergraduates (mean age = 20.30 years, $SD = 1.22$ years; 27 females) from UCSD participated for course credit. All signed consent forms approved by the Human Research Protection Program at UCSD. We planned our sample size on the basis of a priori power calculations and in accordance with previous studies on rapid responses and classification judgments for faces (e.g., Carr, Korb, Niedenthal, & Winkielman, 2014; Carr, Winkielman, & Oveis, 2014). For Experiment 2 we projected a d_z effect size between 0.40 and 0.50. Using G*Power (Version 3.1.9.2; Faul et al., 2007), we estimated that with a total sample of 34 to 52 subjects, we would have 80% power to detect such an effect, given a two-tailed test and α level of .05. Therefore, we decided to target a sample size of 40. Our face stimuli, equipment hardware, and software were the same as in Experiment 1.

Design and procedure. Our main changes for Experiment 2 concerned the design of the Phase 2 task. The

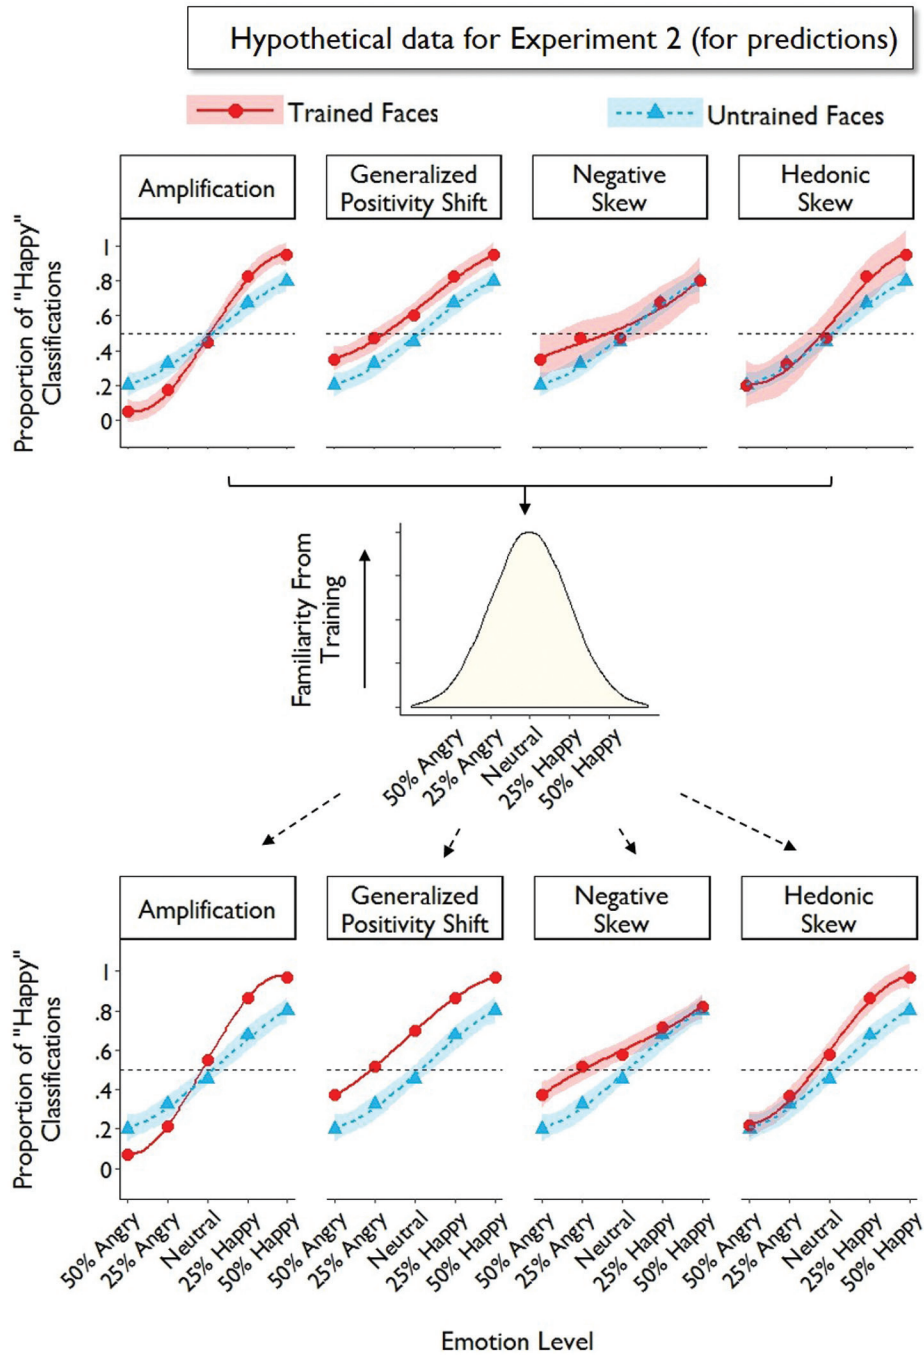


Fig. 4. Illustration of the four frameworks' qualitative predictions for Experiment 2. The graph at the top shows hypothetical data that would support each of these frameworks, before accounting for the fact that subjects were familiarized with neutral expressions only. Because subjects were trained on neutral expressions, each of these functions would be filtered through a familiarity distribution (middle panel; here, we assume a simple similarity gradient based on literature about mere-exposure generalization effects; Gordon & Holyoak, 1983). The integrated predictions of the four frameworks after the original hypothetical data are combined with the familiarity distribution are shown in the bottom panel.

Phase 1 training task was the same as in Experiment 1, but we replaced the dual-face perceptual task in Phase 2 with a speeded forced-choice classification paradigm.

Figure 5 shows a schematic of the Phase 2 task in Experiment 2. There were five blocks of 60 trials (total of

300 trials); in each block, each of the 60 unique stimuli (12 models × 5 emotion levels) was presented once. On each trial, a fixation cross was displayed for 1,000 ms. Next, one face stimulus (a trained or untrained model displaying a 50%-angry, 25%-angry, neutral, 25%-happy,

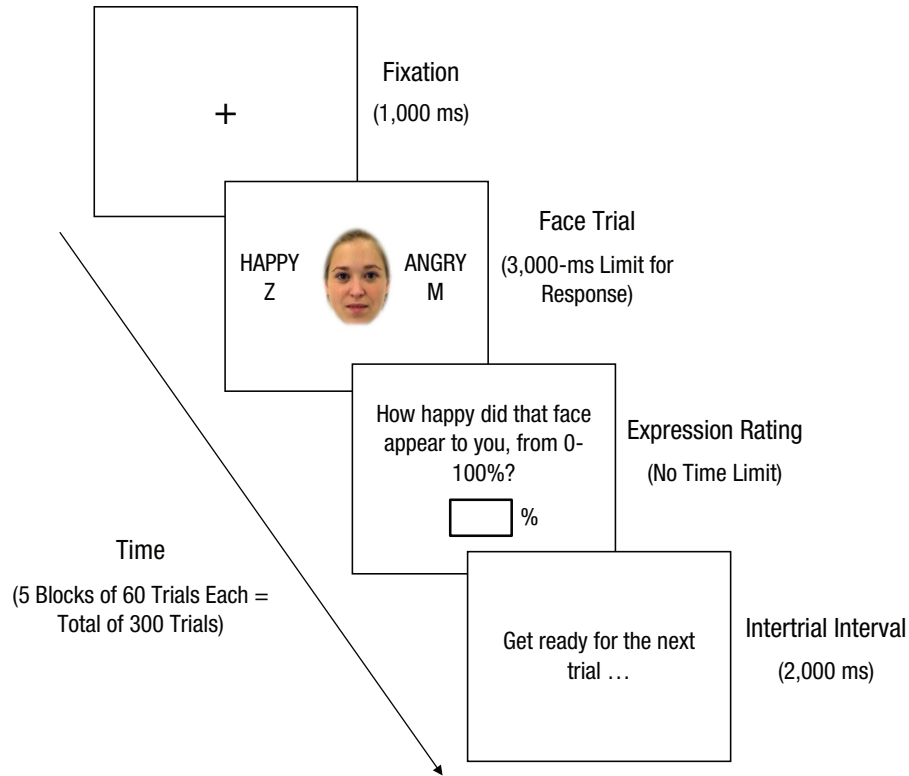


Fig. 5. Design and procedure for the Phase 2 task in Experiment 2. Each trial began with a fixation screen, after which one of the stimulus faces was displayed. Subjects classified the face as happy or angry and then rated how happy they thought the face appeared (scale from 0% to 100%).

or 50%-happy expression) appeared in the center of the screen. Subjects were instructed to categorize the face, as quickly and accurately as possible, as happy or angry, using the “Z” and “M” keys on the keyboard (assignment of responses to keys was randomized across trials). They were told that they would have up to 3,000 ms to respond and that any response after that period would be counted as incorrect. After the classification, another question asked, “How happy did that face appear to you?” Subjects typed their answers, from 0% (*not happy at all*) to 100% (*as happy as possible*), in a response box shown on the screen. There was no time limit for responding to this question. After subjects completed all 300 trials of this task, they were debriefed and given credit for their participation.

Results

As in Experiment 1, we analyzed the data using MLMs, excluding all responses with RTs (for any response) less than 200 ms. In this experiment, maximal random effects were modeled on both subject and stimulus identities (because of the single-face design).

We also fit logistic psychometric functions to subjects’ forced-choice classifications, separately for the trained

and untrained faces, using the *quickpsy* package in R (Linares & Lopez-Moliner, 2016). These psychometric curves were fitted using the direct maximization of the likelihood, according to the following form:

$$\varphi(x) = \gamma + (1 - \gamma - \lambda) * \text{fun}([1 + \exp(-\beta * (x - \alpha))]^{-1}),$$

where γ is the guess rate, λ is the lapse rate, and $\text{fun}(\cdot)$ is the sigmoidal-shape logistic function with asymptotes at 0 and 1 (Linares & Lopez-Moliner, 2016). The fitting of these psychometric functions allowed us to calculate slopes and thresholds at different emotion levels.

Modeling the happiness classification probabilities. First, we analyzed the probability of a face being classified as happy using an MLM with a Training (two levels: trained, untrained; within subjects) \times Stimulus Emotion (five levels: 50% angry, 25% angry, neutral, 25% happy, 50% happy; within subjects) fixed-effects structure. We observed a significant interaction, $F(4, 50.46) = 3.17, p = .02$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 2.58, p = .04, \eta^2 = .01$. Follow-up tests revealed that although subjects were more likely to classify trained than untrained faces as happy when the faces were

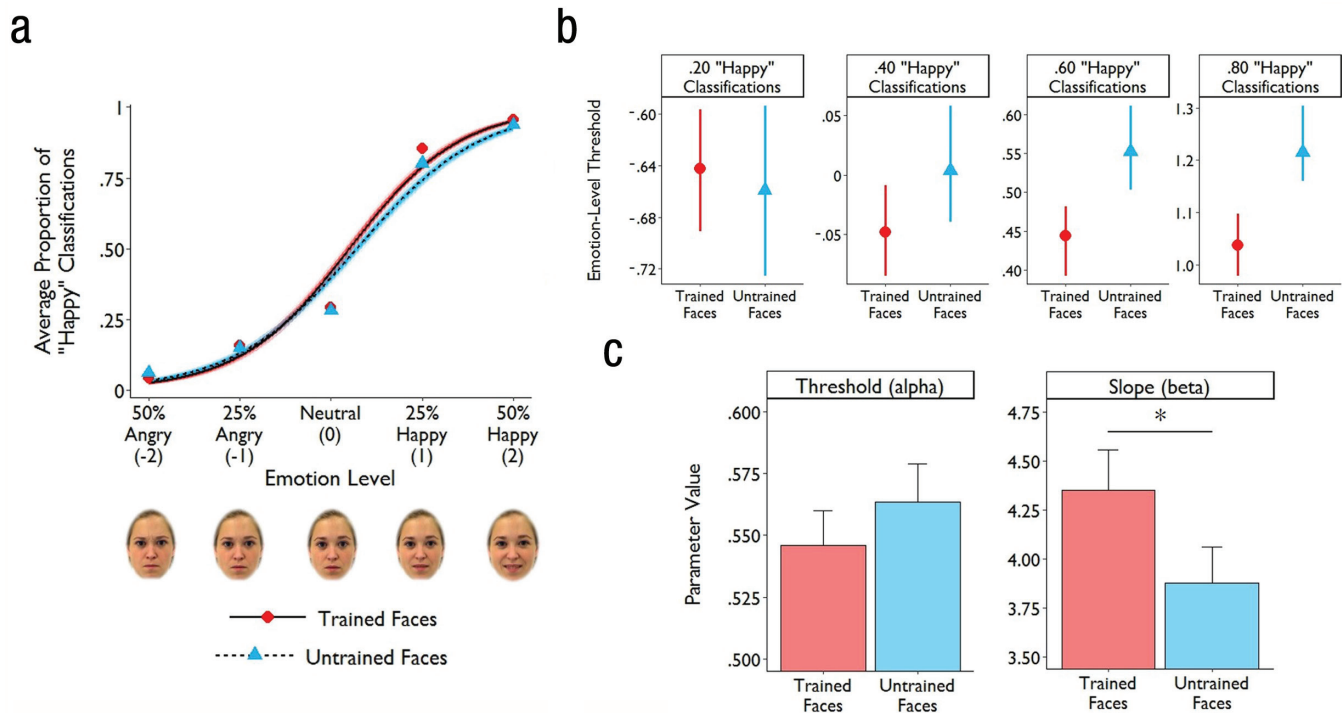


Fig. 6. Results for the Phase 2 task in Experiment 2. The graph in (a) shows the probability of a face being classified as happy as a function of emotion level, separately for trained and untrained faces. The solid and dashed black lines are logistic functions fitted to the data, and the shaded regions around the functions are bootstrapped 95% confidence intervals, which were generated using 100 samples. The graphs in (b) show the emotion-level thresholds necessary for trained and untrained faces to be classified as happy with probabilities of .20, .40, .60, and .80, separately for trained and untrained faces. The error bars represent 95% confidence intervals around the point estimates. The bar graphs (c) show the mean estimates of the alpha parameters (thresholds) and beta parameters (slopes) for subjects' cumulative Gaussian functions for classifying faces as happy, separately for trained and untrained faces. Error bars indicate ± 1 SEM. The asterisk indicates a significant difference between the parameters for trained and untrained faces ($*p < .05$).

neutral, 25% happy, and 50% happy, this difference was greatest at the 25%-happy level, $b = 0.06$, 95% CI = [0.02, 0.09], $t(68.40) = 3.29$, $p = .002$, $d_z = 0.52$. The probability of being classified as happy did not differ between trained and untrained faces at the 50%-angry level, $b = -0.02$, 95% CI = [-0.05, 0.02], $t(67.90) = -1.13$, $d_z = 0.18$, or the 25%-angry level, $b = 0.01$, 95% CI = [-0.03, 0.04], $t(68.70) = 0.48$, $d_z = 0.08$. These results all replicate those of Experiment 1. Note that we also detected a main effect of stimulus emotion, $F(4, 56.52) = 147.99$, $p < .001$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 942.62$, $p < .001$, $\eta^2 = .91$. This result indicated only that the likelihood of faces being classified as happy increased as they became more positive (going from 50% angry to 50% happy; see Fig. 6a).

Psychometric function fitting

Group-level analysis of thresholds. As previously mentioned, we also fit logistic psychometric functions to the classification data, separately for each training condition (i.e., trained vs. untrained faces). We did this to calculate emotion-level thresholds at different response probabilities.

More specifically, using the fitted curve for each training condition, we were able to obtain points on the morph continuum (somewhere between 50% angry and 50% happy) that corresponded to certain proportions of "happy" classifications. By bootstrapping these curves, we were also able to estimate 95% CIs around these thresholds, in order to compare the training conditions (for this analysis, we generated 100 bootstrap samples for each function).

Figures 6a and 6b display the results. To gauge how the thresholds for trained and untrained faces changed across the morph continuum, we assessed thresholds at four different response probabilities: .20, .40, .60, and .80 "happy" classifications. The pattern of results was similar to what we found in Experiment 1: The logistic functions showed that subjects classified trained faces as happy more often than untrained faces, but only when those faces contained neutral to positive features (see Fig. 6a). Consequently, subjects required less actual happiness to be present in the trained faces, compared with the untrained faces, in order to classify them as happy 60% or 80% of the time. Note that the differences between the thresholds for trained and untrained faces were significant

for .60 and .80 probabilities of “happy” classifications, $p < .05$, but not for .20 and .40 probabilities of “happy” classifications.

Subject-level analysis of thresholds. We also fitted subject-level logistic functions to the classification data. Specifically, we followed the same steps as with the group-level data, but fitted the functions to each individual subject’s classification data separately. From these functions, we were able to calculate emotion-level thresholds for each subject at the four “happy” response probabilities (.20, .40, .60, and .80), for both trained and untrained faces. Next, we created an MLM to predict these subject-level thresholds, using a Training (two levels: trained, untrained; within subjects) \times Response Probability (four levels: .20, .40, .60, .80; within subjects) fixed-effects structure (we built the maximal random-effects structure on each subject’s identity that allowed for model convergence).

This analysis yielded results similar to those of the group-level analyses depicted in Figures 6a and 6b. We observed a Training \times Response Probability interaction, $F(3, 195.00) = 6.94$, $p < .001$, which was confirmed by a repeated measures ANOVA, $F(3, 117) = 4.17$, $p = .008$, $\eta^2 = .01$. Thresholds for faces to be classified as happy 60% of the time were marginally lower for trained faces than for untrained faces, $b = -0.11$, 95% CI = $[-0.24, 0.02]$, $t(72.30) = -1.67$, $p = .10$, $d_z = 0.26$, and thresholds for faces to be classified as happy 80% of the time were significantly lower for trained faces than for untrained faces, $b = -0.20$, 95% CI = $[-0.33, -0.07]$, $t(72.30) = -3.00$, $p = .004$, $d_z = 0.47$. There were no differences between trained and untrained faces at the .20 response probability, $b = 0.05$, 95% CI = $[-0.08, 0.18]$, $t(72.30) = 0.76$, $d_z = 0.12$, or the .40 response probability, $b = -0.04$, 95% CI = $[-0.17, 0.09]$, $t(72.30) = -0.57$, $d_z = 0.09$.

Comparative analysis of thresholds and slopes. The previous analyses on thresholds examined the bias toward categorizing trained, rather than untrained, faces as happy at different emotion levels. However, we were also able to analyze the slopes of the psychometric functions to gauge whether familiarity increased discrimination, apart from shifting the curve (Kingdom & Prins, 2016). Analyzing the slopes was also relevant for determining which of the four frameworks (see Table 1) accounted for the data best. For example, if familiarity simply led to a broad bias to classify trained models as happier than untrained models (as would be predicted by the generalized-positivity-shift framework), we would expect a significant upward shift in the overall function for trained faces, but no difference between the slopes of the functions for trained and untrained faces. In contrast, if familiarity selectively enhanced positive features (as would be predicted by the hedonic-skew framework), we would

expect the psychometric function for trained faces to have a steeper slope than the psychometric function for untrained faces, because familiarity would increase the probability of “happy” classifications of expressions in the neutral-to-happy range.

To investigate these possibilities, we used the Palamedes toolbox in MATLAB 2015a (Prins & Kingdom, 2009) to fit cumulative Gaussian psychometric functions to each subject’s classification data for trained and untrained faces. From these fits, we were able to calculate alpha parameters (i.e., indices of overall bias, or shift, in the curve) and beta parameters (i.e., indices of slope of the curve) for each subject for both trained and untrained faces. We then averaged alpha and beta parameters separately across subjects and then compared these averages between training conditions.

Figure 6c displays the results. Critically, the psychometric functions had steeper slopes for trained faces than for untrained faces, $t(39) = 2.07$, 95% CI for the mean difference = $[0.01, 0.94]$, $p < .05$, $d_z = 0.33$. Trained faces also had lower thresholds than untrained faces, but this difference was not significant, $t(39) = -1.25$, 95% CI for the mean difference = $[-0.05, 0.01]$, $d_z = 0.20$. Taken together with the previous analyses on group- and subject-level thresholds across emotion levels, these results demonstrate that the observed effects of familiarity on the probability of a face being classified as happy cannot be entirely explained by an overall shift in the entire curve (alpha). Instead, the differences between classifications of trained and untrained faces were also driven by increased discrimination (slope, or beta) of trained faces, which affected perception of faces mainly in the neutral-to-happy range. These results provide evidence for the hedonic-skew framework (and conflict with the predictions of the other three frameworks; see Fig. 4).

Classification RTs. We also analyzed subjects’ classification RTs using similar methods. Our model had a Training (two levels: trained, untrained; within subjects) \times Stimulus Emotion (five levels: 50% angry, 25% angry, neutral, 25% happy, 50% happy; within subjects) fixed-effects structure. We included only those RTs between 200 and 3,000 ms, and we \log_{10} -transformed the RTs to normalize the response distribution.

Figure 7 (left panel) shows the results. We observed a Training \times Stimulus Emotion interaction, $F(4, 151.01) = 3.28$, $p = .01$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 3.29$, $p = .01$, $\eta^2 = .03$. Post hoc tests revealed that subjects were marginally faster when classifying trained than untrained faces at the 25%-happy level, $b = -0.009$, 95% CI = $[-0.0193, 0.0004]$, $t(183.60) = -1.88$, $p = .06$, $d_z = 0.30$, and the 50%-happy level, $b = -0.009$, 95% CI = $[-0.0191, 0.0005]$, $t(181.10) = -1.87$, $p = .06$, $d_z = 0.30$. However, subjects were also slower to

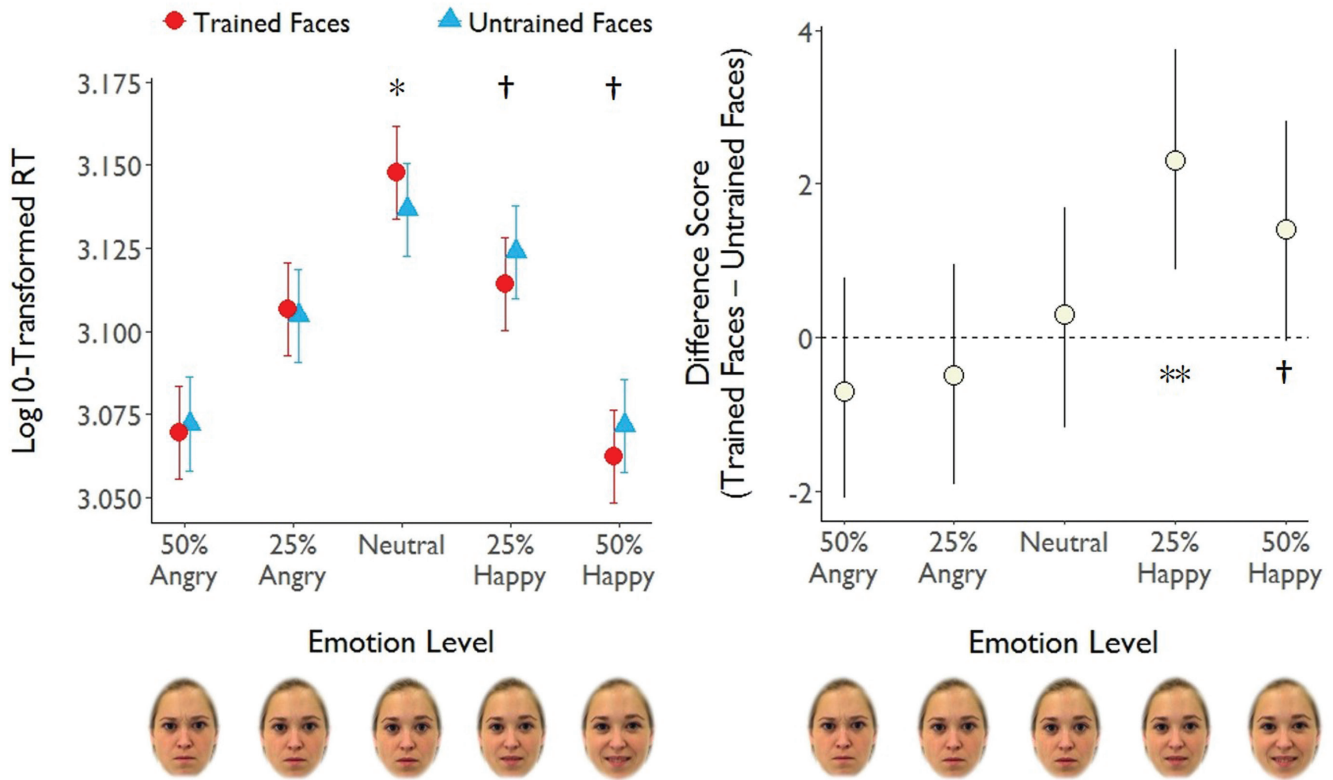


Fig. 7. Response times (RTs) and self-reported estimates of the faces' happiness in Experiment 2. The left panel shows \log_{10} -transformed RTs for the classification decision as a function of emotion level, separately for trained and untrained faces. The daggers and asterisk indicate marginally significant and significant differences between RTs in the two training conditions ($^{\dagger}p < .10$, $^*p < .05$), and the error bars indicate ± 1 SEM. The right panel shows difference scores for the happiness estimates (estimates for trained faces – estimates for untrained faces) as a function of emotion level. The dagger and asterisks indicate difference scores that are marginally significantly different and significantly different from zero ($^{\dagger}p < .10$, $^*p < .05$, $^{**}p < .01$), and the error bars show 95% confidence intervals.

classify neutral expressions of trained models compared with untrained models, $b = 0.011$, 95% CI = [0.001, 0.021], $t(187.80) = 2.24$, $p = .03$, $d_z = 0.35$. There were no RT differences between training conditions for 25%-angry expressions, $b = 0.002$, 95% CI = [-0.008, 0.012], $t(185.20) = 0.43$, $d_z = 0.07$, or 50%-angry expressions, $b = -0.003$, 95% CI = [-0.013, 0.007], $t(181.50) = -0.53$, $d_z = 0.08$. Note that we also observed a main effect of stimulus emotion, $F(4, 68.27) = 18.33$, $p < .001$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 86.87$, $p < .001$, $\eta^2 = .12$; subjects generally had the slowest classification RTs for neutral expressions.

Happiness estimates. Subjects' free-response estimates of the level of happiness they saw in each face gave us an alternative metric of happiness perception, reflecting a more deliberative judgment, rather than a first impression during a rapid classification. To analyze these data, we ran an MLM with a Training (two levels: trained, untrained within subjects) \times Stimulus Emotion (five levels: 50% angry, 25% angry, neutral, 25% happy, 50% happy;

within subjects) fixed-effects structure. Figure 7 (right panel) displays the difference between the estimates for trained and untrained faces at each emotion level. The pattern is similar to that of the perceptual responses in Experiment 1 (see Fig. 3) and the classification responses in Experiment 2 (see Fig. 6a). We observed a Training \times Stimulus Emotion interaction, $F(4, 39.68) = 4.54$, $p = .004$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 5.01$, $p < .001$, $\eta^2 = .001$. Follow-up tests revealed that subjects estimated trained faces as significantly happier than untrained faces at the 25%-happy level, $b = 2.31$, 95% CI = [0.87, 3.75], $t(63.20) = 3.21$, $p = .002$, $d_z = 0.51$, and marginally happier than untrained faces at the 50%-happy level, $b = 1.38$, 95% CI = [-0.06, 2.81], $t(62.60) = 1.92$, $p = .06$, $d_z = 0.30$. There were no differences between the training conditions for faces at the 50%-angry level, $b = -0.66$, 95% CI = [-2.10, 0.77], $t(62.70) = -0.92$, $d_z = 0.15$, or the 25%-angry level, $b = -0.49$, 95% CI = [-1.92, 0.95], $t(63.40) = -0.67$, $d_z = 0.11$. Also, although happiness estimates were higher for trained expressions than for untrained expressions at the neutral

level, this difference did not reach significance, $b = 0.26$, 95% CI = [-1.19, 1.70], $t(63.90) = 0.35$, $d_z = 0.06$. We also observed a main effect of stimulus emotion, $F(4, 152.63) = 116.14$, $p < .001$, which was confirmed by a repeated measures ANOVA, $F(4, 156) = 201.55$, $p < .001$, $\eta^2 = .60$; subjects' happiness estimates increased as the faces became more positive (going from 50% angry to 50% happy).

Discussion

The current results suggest that familiarity influences early stimulus processing, modifying the perception of other individuals' facial expressions. Across different tasks that involved speeded perceptual judgments (Experiment 1), rapid forced-choice classifications (Experiment 2), and deliberative estimates of happiness (Experiment 2), subjects deemed familiar individuals' expressions as happier than unfamiliar individuals' expressions—particularly when the expressions were in the neutral-to-positive range (see Figs. 3, 6, and 7). Psychometric function fitting also revealed that familiarization led to increased discrimination (i.e., steeper slopes in the psychometric functions), which was driven by expressions with more positive features (see Fig. 6). Critically, our findings cannot be explained by simple response biases, given that subjects judged trained faces as happier than untrained faces only at certain levels of emotion. This specific pattern also emerged across multiple tasks and even when responses were orthogonal to the dimension of happiness (Experiment 1). Our results support the hedonic-skew framework, which proposes that familiarity selectively enhances positive features (Garcia-Marques et al., 2004; Harmon-Jones & Allen, 2001; Winkielman & Cacioppo, 2001). Generally, our findings seem inconsistent with models proposing that repetition leads to amplification of preexisting features (nonspecific activation; Albrecht & Carbon, 2014; Mandler et al., 1987), specific decreases in negative affect (negative skew; Lee, 2001; Zajonc et al., 1974), or a generalized positivity shift (Monin, 2003; Titchener, 1910).

Why was the effect of familiarity specific to faces in the neutral-to-positive range? One possible explanation is that familiarity elicits positive (but does not reduce negative) affect (Cacioppo & Berntson, 1994) and that such positive affect is then selectively attributed to stimulus properties that could act as a plausible cause of the affect (i.e., primarily positive features; Schwarz, 2014). Another possibility is that the effects we observed reflect changes in attentional processing. That is, familiarity might free up cognitive resources, allowing for more efficient detection of stimulus features. Depending on the processing goal, these features can be positive (as in the current experiments) but perhaps also negative. Future work should pit this attentional framework against the attributional account to test which provides a better explanation.

In any case, our findings show that the observed effect of familiarity occurs at early stages of processing. Intriguingly, in Experiment 2, subjects not only judged trained faces as happier than untrained faces, but also were faster to classify trained faces as happier (see Fig. 7, left panel). This fits the notion that the observed effect is partially tied to increased fluency (Winkielman et al., 2003), but this interpretation is speculative, as more intense positive affect will also speed up classification. Although the precise mechanistic distinctions between “pure” fluency (ease in stimulus processing) and “pure” familiarity (exposure history and sense of knowing the stimulus) are not essential for our main points, future research should disentangle these constructs, perhaps using neural measures such as electroencephalography (e.g., Nessler, Mecklinger, & Penney, 2005).

To our knowledge, these experiments provide the first evidence that familiarity modulates the perception of facial affect. Our results are consistent with past findings that familiarity changes ratings of neutral expressions (e.g., Claypool et al., 2007), but they go beyond previous findings. First, we used tasks designed to assess early perceptual processes, rather than only scale ratings. Second, our facial expressions varied in the type and intensity of emotion being displayed, which allowed us to estimate psychometric functions, thresholds, and slopes (see Fig. 6). We showed that the effects of familiarity on happiness judgments are dependent on the positive features in the test expressions. As the expressions became happier, subjects were more likely to judge the trained faces as happier than the untrained faces (both when trained and untrained faces were compared directly, in Experiment 1, and when they were presented alone, in Experiment 2). Finally, our findings help to distinguish between prominent theoretical accounts of and predictions for the relationship between familiarity and the perception of valenced features.

Future work should evaluate the boundary conditions of these findings and, in particular, investigate why familiarity did not reduce the perceived negativity of angry expressions. We suggest that familiarity selectively elicits positive affect, which is more easily attributed to positive than to negative stimulus features. However, anger could be special and perhaps gated from familiarity influences (e.g., Pinkham, Griffin, Baron, Sasson, & Gur, 2010). One could determine whether this is the case by using morphs displaying negative emotions other than anger, but note that our data do not offer any real support for the idea that anger expressions are special, given that our subjects were as sensitive to neutral-to-happy transitions as to neutral-to-angry transitions (see Fig. 6a).

Finally, claims about the perceptual nature of any effect are often hotly debated (Firestone & Scholl, 2016), and we hesitate to claim that familiarity influences early

vision. Nevertheless, evidence for affective influences on perception is reasonably strong (Vetter & Newen, 2014), so future research should examine such early influences with tasks that can gauge visual pop-out for trained faces (e.g., continuous flash suppression or visual search paradigms). For now, however, our results suggest that familiar faces do “look” happier.

Action Editor

Alice J. O’Toole served as action editor for this article.

Author Contributions

E. W. Carr developed the study concept. All the authors contributed to the study design. Testing and data collection were performed by E. W. Carr. E. W. Carr performed the data analysis and interpretation under the supervision of T. F. Brady and P. Winkielman. E. W. Carr drafted the manuscript, and T. F. Brady and P. Winkielman provided critical revisions. All the authors approved the final version of the manuscript for submission.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

References

- Albrecht, S., & Carbon, C.-C. (2014). The Fluency Amplification Model: Fluent stimuli show more intense but not evidently more positive evaluations. *Acta Psychologica, 148*, 195–203. doi:10.1016/j.actpsy.2014.02.002
- Bagiella, E., Sloan, R. P., & Heitjan, D. F. (2000). Mixed-effects models in psychophysiology. *Psychophysiology, 37*, 13–20.
- Bornstein, R. F., & D’Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition, 12*, 103–128.
- Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin, 115*, 401–423. doi:10.1037/0033-2909.115.3.401
- Carr, E. W., Huber, D. E., Pecher, D., Zeelenberg, R., Halberstadt, J., & Winkielman, P. (2017). The ugliness-in-averageness effect: Tempering the warm glow of familiarity. *Journal of Personality and Social Psychology, 112*, 787–812.
- Carr, E. W., Korb, S., Niedenthal, P. M., & Winkielman, P. (2014). The two sides of spontaneity: Movement onset asymmetries in facial expressions influence social judgments. *Journal of Experimental Social Psychology, 55*, 31–36. doi:10.1016/j.jesp.2014.05.008
- Carr, E. W., Winkielman, P., & Oveis, C. (2014). Transforming the mirror: Power fundamentally changes facial responding to emotional expressions. *Journal of Experimental Psychology: General, 143*, 997–1003. doi:10.1037/a0034972
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience, 7*, 308–313. doi:10.1038/nn1194
- Claypool, H. M., Hugenberg, K., Housley, M. K., & Mackie, D. M. (2007). Familiar eyes are smiling: On the role of familiarity in the perception of facial affect. *European Journal of Social Psychology, 37*, 856–866. doi:10.1002/ejsp.422
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39*, 175–191. doi:10.3758/BF03193146
- Fechner, G. T. (1876). *Vorschule der Aesthetik* [Preschool of aesthetics]. Leipzig, Germany: Breitkopf & Härtel.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral & Brain Sciences, 39*, Article e229. doi:10.1017/S0140525X15000965
- Garcia-Marques, T., Mackie, D. M., Claypool, H. M., & Garcia-Marques, L. (2004). Positivity can cue familiarity. *Personality and Social Psychology Bulletin, 30*, 585–593. doi:10.1177/0146167203262856
- Garcia-Marques, T., Prada, M., & Mackie, D. M. (2016). Familiarity increases subjective positive affect even in non-affective and non-evaluative contexts. *Motivation and Emotion, 40*, 638–645. doi:10.1007/s11031-016-9555-9
- Gordon, P. C., & Holyoak, K. J. (1983). Implicit learning and generalization of the “mere exposure” effect. *Journal of Personality and Social Psychology, 45*, 492–500. doi:10.1037/0022-3514.45.3.492
- Harmon-Jones, E., & Allen, J. J. B. (2001). The role of affect in the mere exposure effect: Evidence from psychophysiological and individual differences approaches. *Personality and Social Psychology Bulletin, 27*, 889–898. doi:10.1177/0146167201277011
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology, 103*, 54–69. doi:10.1037/a0028347
- Kingdom, F. A. A., & Prins, N. (2016). *Psychophysics: A practice introduction*. Waltham, MA: Academic Press.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). *Package ‘lmerTest’*. Retrieved from <http://cran.uib.no/web/packages/lmerTest/lmerTest.pdf>
- Lee, A. Y. (2001). The mere exposure effect: An uncertainty reduction explanation revisited. *Personality and Social Psychology Bulletin, 27*, 1255–1266. doi:10.1177/01461672012710002
- Linares, D., & Lopez-Moliner, J. (2016). quickpsy: An R package to fit psychometric functions for multiple groups. *The R Journal, 8*(1), 122–131.
- Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 646–648. doi:10.1037/0278-7393.13.4.646
- McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science, 149*, 1115–1116. doi:10.1126/science.149.3688.1115
- Monin, B. (2003). The warm glow heuristic: When liking leads to familiarity. *Journal of Personality and Social Psychology, 85*, 1035–1048. doi:10.1037/0022-3514.85.6.1035
- Nessler, D., Mecklinger, A., & Penney, T. B. (2005). Perceptual fluency, semantic familiarity and recognition-related

- familiarity: An electrophysiological exploration. *Cognitive Brain Research*, 22, 265–288. doi:10.1016/j.cogbrainres.2004.03.023
- Phelps, E. A., Ling, S., & Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological Science*, 17, 292–299. doi:10.1111/j.1467-9280.2006.01701.x
- Pinkham, A. E., Griffin, M., Baron, R., Sasson, N. J., & Gur, R. C. (2010). The face in the crowd effect: Anger superiority when using real faces and multiple identities. *Emotion*, 10, 141–146. doi:10.1037/a0017387
- Prins, N., & Kingdom, F. A. A. (2009). Palamedes: Matlab routines for analyzing psychophysical data [Computer software]. Retrieved from <http://www.palamedestoolbox.org/>
- Psychology Software Tools, Inc. (2012). E-Prime 2.0 [Computer software]. Retrieved from <http://www.pstnet.com>
- Schwarz, N. (2014). *Cognition and communication: Judgmental biases, research methods, and the logic of conversation*. New York, NY: Psychology Press.
- Schwarz, N., & Clore, G. L. (2003). Mood as information: 20 years later. *Psychological Inquiry*, 14, 296–303.
- Störmer, V. S., & Alvarez, G. A. (2016). Attention alters perceived attractiveness. *Psychological Science*, 27, 563–571. doi:10.1177/0956797616630964
- Titchener, E. B. (1910). *A textbook of psychology*. New York, NY: Macmillan.
- van der Schalk, J., Hawk, S. T., Fischer, A. H., & Doosje, B. (2011). Moving faces, looking places: Validation of the Amsterdam Dynamic Facial Expression Set (ADFES). *Emotion*, 11, 907–920. doi:10.1037/a0023853
- Vetter, P., & Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and Cognition*, 27, 62–75. doi:10.1016/j.concog.2014.04.007
- West, B. T., Welch, K. B., & Galecki, A. T. (2014). *Linear mixed models: A practical guide using statistical software* (2nd ed.). Boca Raton, FL: CRC Press.
- Wingenbach, T. S. H., Ashwin, C., & Brosnan, M. (2016). Validation of the Amsterdam Dynamic Facial Expression Set – Bath Intensity Variations (ADFES-BIV): A set of videos expressing low, intermediate, and high intensity emotions. *PLoS ONE*, 11(1), Article e0147112. doi:10.1371/journal.pone.0147112
- Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation elicits positive affect. *Journal of Personality and Social Psychology*, 81, 989–1000. doi:10.1037/0022-3514.81.6.989
- Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 189–217). Mahwah, NJ: Erlbaum.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2), 1–27. doi:10.1037/h0025848
- Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, 10, 224–228. doi:10.1111/1467-8721.00154
- Zajonc, R. B., Markus, H., & Wilson, W. R. (1974). Exposure, object preference, and distress in the domestic chick. *Journal of Comparative and Physiological Psychology*, 86, 581–585. doi:10.1037/h0036173