

Identity modulates short-term memory for facial emotion

MURRAY GALSTER

Brandeis University, Waltham, Massachusetts

MICHAEL J. KAHANA

University of Pennsylvania, Philadelphia, Pennsylvania

HUGH R. WILSON

York University, Toronto, Ontario, Canada

AND

ROBERT SEKULER

Brandeis University, Waltham, Massachusetts

For some time, the relationship between processing of facial expression and facial identity has been in dispute. Using realistic synthetic faces, we reexamined this relationship for both perception and short-term memory. In Experiment 1, subjects tried to identify whether the emotional expression on a probe stimulus face matched the emotional expression on either of two remembered faces that they had just seen. The results showed that identity strongly influenced recognition short-term memory for emotional expression. In Experiment 2, subjects' similarity/dissimilarity judgments were transformed by multidimensional scaling (MDS) into a 2-D description of the faces' perceptual representations. Distances among stimuli in the MDS representation, which showed a strong linkage of emotional expression and facial identity, were good predictors of correct and false recognitions obtained previously in Experiment 1. The convergence of the results from Experiments 1 and 2 suggests that the overall structure and configuration of faces' perceptual representations may parallel their representation in short-term memory and that facial identity modulates the representation of facial emotion, both in perception and in memory. The stimuli from this study may be downloaded from <http://cabn.psychonomic-journals.org/content/supplemental>.

The international community has instructed its member nations to frown upon smiles. By convention, the photo in any new passport must show the passport's bearer unsmiling, with lips pressed together and an expression whose neutrality is beyond question. This convention is meant to aid automated facial recognition technology that is used to protect against identity fraud or terrorism. Smiles and other emotional expressions are frowned upon because they alter the geometry of key facial features and could subvert an automated biometric system's ability to match the passport photo to the traveler's own face. So, as far as the automated biometric system is concerned, facial expression and identity are inextricably linked.

Although biometric recognition systems link facial identity and facial expression, human observers seem not always to follow suit. In fact, independence of facial identification from other aspects of face processing was a keystone of Bruce and Young's (1986) influential account of face recognition. That account describes the face recognition system as comprising (1) a module specialized for recognizing identity (the individual to whom

the face belongs) and (2) a module specialized for analyzing the face's expression (the emotion that the face is expressing).

The perceptual modularity of these two processes has been supported by a number of studies. For example, prosopagnosia, or face blindness, sometimes impairs a patient's ability to determine somebody's identity from their face alone, while leaving recognition of facial emotion unaffected (Duchaine, Parker, & Nakayama, 2003; Sacks, 1985). The converse has also been observed: Processing of facial emotion is affected, but the processing of identity is spared (G. W. Humphreys, Donnelly, & Riddoch, 1993). Also, adults with social developmental disorders, such as autism and Asperger's disorder, show no systematic relationship between ability to recognize facial identity and ability to recognize facial expression (Hefter, Manach, & Barton, 2005). These dissociations are consistent with the proposition that processing of facial expression and processing of facial identity activate pathways in the brain that are at least partially separable (de Gelder, Pourtois, & Weiskrantz, 2002; de Gelder, Vroomen, Pourtois,

R. Sekuler, sekuler@brandeis.edu

& Weiskrantz, 1999; Hasselmo, Rolls, & Baylis, 1989; Posamentier & Abdi, 2003; Sergent, Ohta, MacDonald, & Zuck, 1994).

Claims about the independence of identity and emotion perception have not gone unchallenged (for a review, see Posamentier & Abdi, 2003). Thus, a pair of recent studies produced conflicting results, using the adaptation after-effect (Ellamil, Susskind, & Anderson, 2008; Fox, Oruc, & Barton, 2008). Using a different experimental paradigm, Ganel and Goshen-Gottstein (2004) asked subjects to categorize photographs either according to the identity they showed or according to the emotion displayed. The stimuli were photographs of two different people (identities) shown in two different emotional expressions and in two different views. In either case, the dimension that subjects attempted to ignore continued to influence the speed with which the items were classified on the other, attended dimension. Baudouin, Martin, Tiberghien, Verlut, and Franck (2002) used a similar experimental design, with schizophrenic as well as normal subjects. Both groups of subjects showed a striking asymmetry of influence between the two dimensions, with differences in identity exerting far less influence on judgments of emotion than vice versa. Moreover, the degree of interattribute interference varied strongly with the severity of negative symptoms in the schizophrenic subjects.

Unfortunately, these studies, and other behavioral studies of interactions between identity and emotion, have based their conclusions on stimuli that were a narrow, unsystematic sample from the task-relevant stimulus space. For example, in each experiment, Ganel and Goshen-Gottstein (2004) and Baudouin et al. (2002) tested subjects with just two different faces, each in combination with just two different emotions. This arrangement makes it difficult to characterize precisely the cross talk between the two dimensions. These studies made no real effort to test the boundary conditions for perceptual independence, which would have required stimulus materials whose similarity relations were carefully chosen and controlled (Sekuler & Kahana, 2007). In the experiments reported here, similarity relations among faces and emotions were manipulated so as to challenge subjects' ability to perceive and remember facial emotion independently of facial identity. Our goals were to determine whether short-term memory for emotion is influenced by identity and whether this potential influence could be explained by perceptual similarity, and to do so with stimuli whose perceptual similarities were well defined.

Divergence between the processing of stimulus attributes at one level, such as at a perceptual level, does not guarantee that their independence will be preserved at a subsequent stage of processing, as Ashby and Maddox (1990) have pointed out. Although their argument focused on decision making, their general point would seem to apply to memory storage and recall as well. Given the distributed, interactive nature of the neural network that may be responsible for processing faces (Haxby, Hoffman, & Gobbini, 2000), Ashby and Maddox's argument might be particularly relevant for face-related tasks. Demonstrations that facial identity and facial emotion are perceptually

independent do not guarantee that this independence is preserved when it comes to memory. Therefore, we set out to assess the relationships between facial identity and facial emotion at the perceptual level and at the level of short-term memory.¹

EXPERIMENT 1

The choice of appropriate stimuli is clearly an important one for any experiment (Stevens, 1951), but, arguably, never more so than for an experiment that seeks to examine connections between memory for facial expression and facial identity. In fact, the similarity of one face's identity to another or one emotional expression to another limits performance in a variety of settings, such as learning name-face associations (Pantelis, van Vugt, Sekuler, Wilson, & Kahana, 2008) and cases of mistaken identity (Wells, 1993), to name just two. Wilson, Loffler, and Wilkinson (2002) developed a class of synthetic faces that offer the qualities needed. These include good realism and recognizability, achieved without sacrificing control of stimulus characteristics. These realistic synthetic faces are known as *Wilson faces*, and full details of their construction have been given elsewhere (Wilson et al., 2002). Figure 1 shows examples of Wilson faces. In brief, these faces are synthesized from several dozen independent measurements made on photographs of a person's face. For example, the nose and mouth of a synthetic face reflect measurements of an individual's mouth width, lip thickness, and nose width and length. Filtering the faces removes textures and hair, leaving the hairline intact. In their initial applications, Wilson faces were neutral with respect to emotional expression. Subsequently, though, Goren and Wilson (2006) developed a method for introducing distinct emotional expressions at graded levels of intensity. With these faces, one can construct stimulus sets with graded and distinct emotional expressions (Goren & Wilson, 2006) and varying similarity relations among faces belonging to different individuals (facial identity). Wilson faces are good approximations to actual faces, producing levels of neural activation in the brain's face regions that are comparable to the activations evoked by high-fidelity photographs of individual faces (Loffler, Yourganov, Wilkinson, & Wilson, 2005) and exhibiting substantial face-inversion effects (Goren & Wilson, 2006), which is taken as a hallmark of face processing (but see R. Russell, Biederman, Nederhouser, & Sinha, 2007). In addition, Goren and Wilson demonstrated that emotional expressions on Wilson faces can be extracted quickly—that is, from presentations lasting 110 msec. This rapidity of processing was important to our choice of stimuli, since conspecifics, including those from our own species, often can interpret an emotional expression that is glimpsed in a single brief fixation (Ekman, 1993). Furthermore, Wilson faces preserve facial identity, which allows a face to be accurately matched to the photograph of the individual from which it was derived (Goren & Wilson, 2006; Loffler, Gordon, Wilkinson, Goren, & Wilson, 2005; Wilson et al., 2002; Wilson & Wilkinson, 2002). Finally, the value of Wilson faces has been demonstrated




Figure 1. Examples of the Wilson faces used as stimuli here. (A–C and E–G) Variants of a single female face. Panels A–C and G show this female with expressions that are neutral (A), maximally sad (B), maximally fearful (C), and maximally angry (G). Panels E–F show this same female exhibiting different degrees of anger, from mild (E) to maximal (G). For comparison with panel C, panel D shows maximal fear exhibited by a different female face. Other examples are available in the supplementary material to this article.

in studies of short-term recognition memory (Sekuler & Kahana, 2007; Yotsumoto, Kahana, Wilson, & Sekuler, 2007) and associative learning of face–name pairs (Pantelis et al., 2008). They have also have proved useful in studying the way that a face’s emotional expression affects an onlooker’s gaze (Fung et al., 2008; Isaacowitz, Toner, Goren, & Wilson, 2008; Isaacowitz, Wadlinger, Goren, & Wilson, 2006a, 2006b).

Guided by Ekman and Friesen’s (1975) facial action coding system, Goren and Wilson (2006) demonstrated that a synthetic face’s emotional expression could be controlled by modifying the relative positions and/or orientations of 10 facial structures that contribute to facial expression of emotion. With a nominally neutral face as a starting point, some desired emotion is produced by adjusting variables such as the amount of visible sclera (Whalen et al., 2004), the positions of upper and lower eyelids, the shape and position of the eyebrows (Sadr, Jarudi, & Sinha, 2003), the width of the nose, and the curvature and size of the mouth. Moreover, the *degree* of each emotion can be systematically manipulated by varying the *amount* by which these structural features are altered. Additional details about the construction of synthetic faces that convey emotions have been given elsewhere (Goren & Wilson, 2006).

While photographs were being taken to provide the raw material for the Wilson face, each person photographed was instructed to assume and hold a neutral expression. However, even perfect compliance with this instruction does not guarantee that the resulting facial expression will consistently be perceived as neutral. Since Darwin’s (1872) pioneering explorations of facial expression, it has been known that some people’s faces when physically at rest—that is, with relaxed facial muscles—are perceived by others as expressing some particular emotion. This point was recently confirmed and generalized in an ex-

tensive, systematic study (Neth & Martinez, 2009). As a result of such observations, we were not surprised that a casual examination showed that some of the 79 “neutral” Wilson faces seemed to display some characteristic emotion, which differed among faces. We therefore decided to identify Wilson faces whose neutral expressions, and other expressions as well, were most likely to be perceived accurately. These faces would then be used as stimuli in our experiments.

The stimulus pool comprised 79 nominally neutral Wilson faces. Each was based on a photograph of a different person, 40 Caucasian females and 39 Caucasian males. Prior to being photographed, the subject of the photograph had been instructed to adopt a neutral, relaxed pose that displayed no emotion. Using an algorithm introduced by Goren and Wilson (2006), we generated nine variants of each neutral face by morphing the neutral expression into three distinct emotional expressions—fear, anger, and sadness—with three degrees of each. We opted against including happy faces in the stimulus set because Goren and Wilson’s algorithm prevented the lips from parting when a smile was called for. This constraint, which kept any teeth from showing, caused many subjects in preliminary tests to describe the smiles as inauthentic. The morphing operation was guided by Ekman and Friesen’s (1975) description of the feature changes that normally characterize each emotion. In particular, changes were applied to 10 features: height and shape of eyebrows, amount of visible sclera, upper and lower eyelid positions, flare of the nostrils, and curvature and size of the mouth and lips. As Table 1 indicates, combinations of changes in these features produced three distinct emotions. Texture information was removed from the faces, forcing emotions to be recognized mainly on the basis of geometric information.

Table 1
Transformations That Generated Expressions
of Emotion From Neutral Faces

Feature Change	Sad	Anger	Fear
Nostril flare	☐	■	☐
Brow distance	■	■	■
Brow upward curve	■	■	■
Brow curvature	■	☐	■
Brow height	■	■	■
Upper eyelid height (visible sclera)	☐	■	☐
Lower lid position	■	■	■
Outer edges of mouth	■	☐	☐
Width of mouth	■	☐	☐
Position of lower lip	■	■	■

Note—Each ■ marks a feature that was changed in order to generate the named emotion; each ☐ entry in the table marks a feature that was not changed in generating the emotion.

Earlier, Goren and Wilson (2006) had determined the stimulus conditions that produced each face's strongest perceived emotion. This was based on naive individuals' judgments of faces whose arrangement of characteristic features were deemed to be most effective in expressing a particular emotion, using criteria that included an expression's believability. More extreme displacements of facial features produced emotions that appeared to be counterfeit. We started with Goren and Wilson's values for the maximum intensity (100%) of each emotion that could be expressed by each face, and generated faces with emotional expressions were at a metric distance of 33%, 67%, or 100% from neutral to maximum. These values can be redescribed in terms of the threshold changes in facial features needed for correct identification of a facial emotion. On the basis of the emotion identification thresholds reported by Goren and Wilson, 33%, 67%, and 100% of the maximum achievable facial emotion correspond to ~1.3, 2.7, and 4.0 times the threshold for identifying the emotion. Note that individual differences in identification thresholds (Goren & Wilson, 2006) make these values only approximate.

Experiment 1 comprised two phases. In the first phase, the large stimulus pool was culled in order to identify those faces whose neutral expression was clearest—that is, with the least tinge of some emotion. This first phase was necessitated by the fact that some faces appear to express a characteristic emotion even when the face's owner does not mean to express an emotion (e.g., Neth & Martinez, 2009). In the second phase of the experiment, the faces identified in the first phase were used as the stimuli in a visual recognition task. Independent groups of subjects served in these two phases. The following describes each phase in turn.

Method: Phase 1

Subjects. Seven undergraduate students naive as to the purpose of the experiment participated. Four served in a single session each; 3 served in two sessions. For the subjects who served in more than one session, data were averaged over sessions for analysis. All the subjects were paid for their participation.

Procedure. The stimuli were presented on a computer display located 114 cm from the subject. Each trial was begun by a keypress that caused a fixation cross to appear at the center of the display. After 900 msec, the fixation cross was replaced by a single Wilson face, which remained visible for 110 msec. The subject then

responded using the keyboard, indicating whether the face appeared to be angry, fearful, happy, neutral, or sad. The relatively brief stimulus presentation was chosen for consistency with conditions that had been used in most previous studies with Wilson faces, and also in order to mimic the limited information that a person might take in during a single brief glance, minimizing reliance on detailed scrutiny.

Stimuli were displayed using MATLAB with functions from the Psychophysics Toolbox (Brainard, 1997). On the display screen, each Wilson face measured approximately $5.5^\circ \times 3.8^\circ$. The spatial position of the face was randomized by the addition of horizontal and vertical displacements drawn from a uniform random distribution with a mean of 12.6 minarc and a range of 2.1–23.1 minarc. No feedback was provided after a response.

In each session, the subjects saw, in random order, nominally neutral faces and faces that were intended to express some emotion. In that random order, each nominally neutral face was presented four times, with each face intended to express some emotion presented just once. Since others had demonstrated that naive judges easily and consistently recognized Wilson faces' emotions (Isaacowitz et al., 2006b), we decided not to give the subjects an opportunity to become familiar with Wilson faces prior to testing.

Results

To select the faces that would later be used in our memory experiment, we identified the faces whose nominally neutral versions were most consistently categorized as neutral. This screening process was stimulated by earlier observations that individual differences, including differences in facial anatomy, influence the success with which individuals' emotional expressions are categorized (Gosselin & Kirouac, 1995). In our experiment, 13 of the 79 nominally neutral faces were correctly classified as neutral at least 85% of the time. From these 13 faces, we identified 6, 3 male and 3 female, whose emotions were most accurately identified when the emotion was presented at its weakest intensity—that is, 33% of its maximum strength. For the 6 faces, the mean accuracy (and range) with which each emotion was correctly categorized was .63 (.50–.90) for anger, .45 (.20–.80) for sadness, and .60 (.50–.70) for fear.

As was expected, strongly expressed emotions were identified more accurately than weaker ones were. In particular, faces meant to express emotions more strongly—that is, either 67% or 100% of maximum—were correctly categorized on virtually every presentation, despite the brief viewing time. As will be explained below, various forms of these 6 individuals' faces were used in Phase 2 as stimuli with which to assess recognition memory and, then, in the subsequent experiment to assess perception of the stimuli.

Method: Phase 2

With the stimuli that had been identified as most representative of the emotional states of interest, we carried out a short-term recognition procedure to examine possible connections between memory for facial identity and memory for facial expression. On each trial, the subjects saw 2 successively presented faces, each representing some combination of facial identity and emotion (including neutral faces). The subjects judged whether the category and degree of emotion on a subsequently presented probe face matched the category and degree of emotion that had been seen on 1 of the 2 study faces. We presented 2 study faces on each trial, rather than just 1, in order to increase the difficulty of the memory task. Stimuli for each trial

were chosen from a set of 60 possible faces, with 10 variants based on each of the 6 faces identified as best exemplifying both neutral expressions and appropriate emotions. Each set of 10 comprised a neutral version of the face, and three different levels of intensity (33%, 67%, and 100% of maximum) for each of three emotion categories (sadness, anger, and fear).

Subjects. Seventeen subjects, 9 of them female, participated in this and the following experiment. None had participated in Phase 1's effort to identify the most suitable face stimuli. The subjects, 18–29 years of age, had normal or corrected-to-normal vision and were paid for their participation. One female subject's data had to be discarded prior to analysis because of chance-level performance.

Procedure. On each trial, the subject first saw two study faces presented one after the other. These two faces were followed by a test or probe face. The subject's task was to determine whether the expression on the probe face exactly matched the expression that had been seen on either of the two list faces, ignoring the identities. The subjects were instructed that a match required that *both* the category of emotion (neutral, fearful, angry, or sad) *and*, for nonneutral faces, the intensity of that emotion (33%, 67%, or 100% of maximum) had to match. This instruction was reinforced on each trial by postresponse feedback, distinctive tones that signaled whether the subject's response had been correct.

At the start of a trial, a fixation cross appeared at the center of the screen for 500 msec. The fixation point disappeared, and the first study item face, s_1 , appeared for 110 msec, after which it disappeared. This was followed by a 200-msec interstimulus interval and, then, presentation of the second study item face, s_2 , for 110 msec. A 1,200-msec delay followed the two study items, during which a tone sounded, indicating that the next stimulus face would be the probe, p , which also was displayed for 110 msec. A subject responded by pressing one key if the probe face's expression (category as well as intensity) matched that of either study item and by pressing a different key if the probe face's emotion and intensity did not match either study item. To minimize the usefulness of vernier cues, the horizontal and vertical location of each was jittered by a random displacement drawn from a uniform distribution with a mean of ± 12.6 minarc (range = 2.1–23.1 minarc). The viewing distance was 114 cm.

The relationship between a trial's study items and its p gave rise to two types of trials, target and lure. A *target* trial was one on which p 's expression (category and intensity) matched the expression of 1 of the 2 previously seen faces; here, the correct response would be *yes*. A *lure* trial was one on which p 's expression matched that of neither study item; here, the correct response would be *no*. Note that lure trials included ones on which the emotional *category*, but not the *intensity* of the probe's expression, matched that seen on a study item. Lure and target trials occurred in random order. With equal probability, the 60 different stimulus faces (combinations of emotion category and intensity) could be the probe on any trial.

Each subject served in four sessions of 320 trials each. In each session, target and lure trials occurred equally often. The first 8 trials of each session (4 lure and 4 target trials) were practice, and their results were discarded before data analysis.

In addition to equalizing the frequencies of target and lure trials, several other constraints controlled the makeup of stimuli on particular subclasses of target and lure trials. In describing these constraints, the term *identity* signifies the person whose face is represented in a stimulus. More specifically, a face's *identity* refers to the particular person of the six different people, three males and three females, whose faces passed the sieve of Phase 1. Target and lure trials were constructed so that the probe's identity matched that of one item, both study items, or no study item. In addition, any combination of emotion category and intensity could appear only once in the study set. When two study items happened to share the same identity, our experimental design forced those two faces to display emotions from different categories. Finally, on some lure trials, the probe's emotion matched the emotion category but not the intensity

of one of the study items, which allowed us to examine the effect of an intensity mismatch. Figure 2 provides a more detailed account of the distribution of target and lure trials.

Analysis. Across subjects and sessions, 17 trials had to be discarded because of technical errors in recording the data. The remaining 21,643 trials were used in the analyses described below. Recognition performance level was quite stable over the four sessions of testing [$F(3,45) = 0.04, p = .989$]. In addition, a t test for two independent samples showed that male and female subjects' performance did not differ from one another [$t(14) = -0.311, p = .76$]. Therefore, the data for the subsequent analyses were averaged over both sessions and subjects' gender.

Results and Discussion

To determine whether memory for emotional expression was influenced by the identity of the expressing face, we first compared recognition of emotion (1) when the probe face's emotion category and intensity matched those of a study face but their identities did not match and (2) when the probe face's emotion category and intensity, as well as the face's identity, matched those of a study face. For every one of the 16 subjects, correct recognitions were higher in the second case, when emotion and identity both matched, than in the first case, when emotion category and intensity matched but identity did not. The values of $P(\text{Yes})$ for the two cases differed significantly [$t(15) = 10.383, p < .0001$]. So, even though judgments were supposed to be based on emotional expression alone, the identity of the probe face clearly influenced memory for emotional expression.

Recall that a correct recognition required that the probe's emotion match not only the category of emotion of a study face, but also its intensity (33%, 67%, or 100% of maximum). The subjects were to respond *no* if the probe shared an emotion *category* with one of the study items but expressed that emotion in a different degree. Recall also that a match or lack of match on facial identity was entirely irrelevant to the judgment. Figure 3 shows the impact of varying the degree of mismatch between the probe's emotional intensity and the emotional intensity of a study item. This effect is shown in the figure for several conditions of match or mismatch between probe identity and study item identity. Note that when the mismatch was zero (0% on the figure's horizontal axis), a *yes* response consisted of a correct recognition. So the trio of open symbols at the left side of the graph all represent correct recognitions—that is, *yes* responses on trials that were consistent with the instructions given to the subjects. The filled symbols indicate false recognitions—that is, *yes* responses to cases in which there was a mismatch between the emotional intensity of the probe face and the emotional intensity of the study items. In these cases, the correct response would have been *no*. As may be evident from the graph, the overall proportion of *yes* responses declines as the emotional intensity of the probe increasingly deviates from that of the probe [$F(2,30) = 52.588, p < .001$].

Data points on the higher of the two curves represent trials on which not only the probe's emotion category, but also its facial identity, matched a study item (\square and \blacksquare). Data points on the lower of the two curves represent tri-

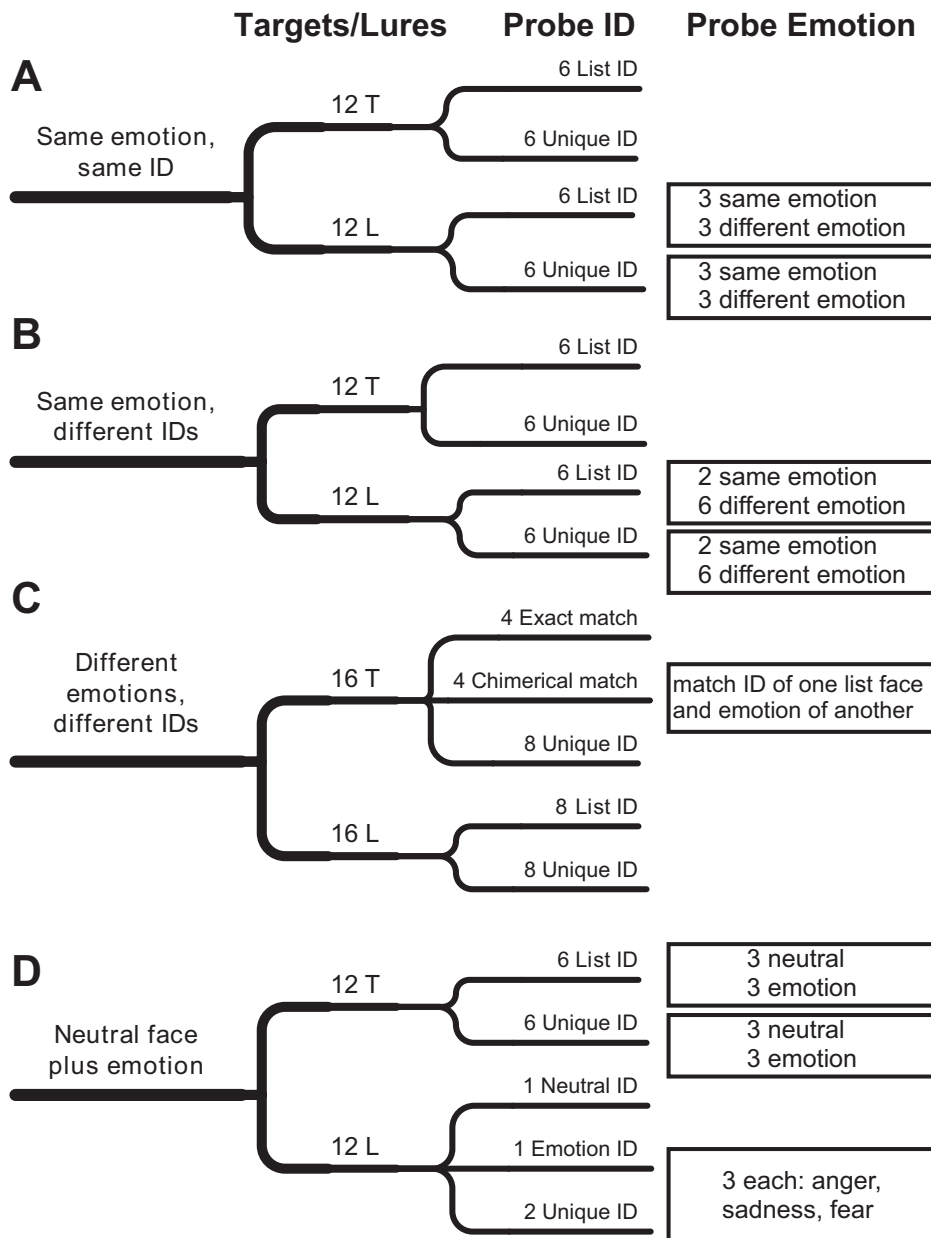


Figure 2. Constitution of trial subtypes: Distribution of trials across different categories of study lists and probes. The broadest categories are enumerated in the leftmost column. (A) Both study items expressed the same emotion, although to different degrees, and both had the same facial identity (ID; 24 trials). (B) Both study items expressed the same emotion, although to different degrees, but the faces represented distinct IDs (24 trials). (C) The study items expressed different emotions and represented distinct IDs (32 trials). (D) One face in the list was neutral, the other expressed some emotion (24 trials). Within each of these four categories, half the trials consisted of target trials (T; the emotion category and intensity displayed by the probe matched those of one of the study items), and the other half consisted of lure trials (L; the emotion category and intensity displayed by the probe matched those of neither study item). The third column in the figure describes the IDs of the probe faces for each type of trial, and the fourth column describes the emotion associated with those trial types.

als on which the match was on emotion category alone, with the identity of the probe face being distinct from the identity of either study item (○ and ●). Although differences between the two sets of data points may seem modest, $P(\text{Yes})$ for a match on both identity and emotion category is consistently and significantly higher than

$P(\text{Yes})$ for a match on emotion category alone [$F(1,15) = 16.554, p < .001$]. The interaction between the two main variables, degree and kind of match, was not significant [$F(2,30) = 1.072, p > .355$]. The difference between corresponding points on the two curves confirms that judgments of emotional expression are influenced by the

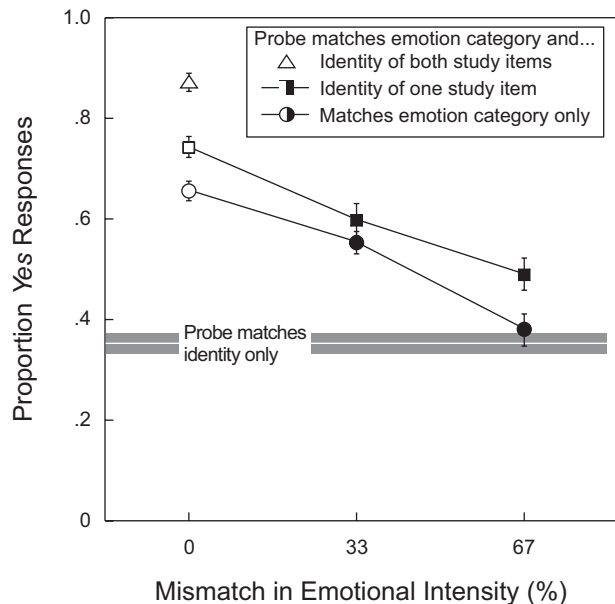


Figure 3. Effect of varying the degree of mismatch between the probe’s emotional intensity and the emotional intensity of a study item. The horizontal axis represents the degree of mismatch, with the leftmost value, 0%, signifying a complete match in intensity. The open symbols, plotted against that one value, represent correct recognitions; the closed symbols, plotted against the two other values (33% and 67%), represent false recognitions—that is, incorrect assertions that the probe’s emotion category and intensity match those of a study item. Data points on the upper curve represent *P*(Yes) when the probe’s emotion category and its facial identity both matched those of a study item; data points on the lower curve represent *P*(Yes) when the probe’s category of emotion matched that of a study face but their facial identities differed. The single open data point at the upper left (Δ) represents *P*(Yes) when the probe’s emotion category and emotion intensity matched those of one study item, while, additionally, the probe’s facial identity matched that of both study items. The vertical center of the gray horizontal bar gives the mean value of *P*(Yes) for trials on which a probe’s identity matched that of one study item, but the probe expressed a category of emotion different from that of either study item; the thickness of that bar corresponds to ± 1 SEM.

presence of match on identity—despite the fact that the subjects were explicitly instructed to ignore that aspect of the stimulus when making recognition judgments. This point is reinforced by the value of *P*(Yes) produced when the probe’s emotion category and intensity matched those of one study item but the probe’s identity matched that of both study items (Δ). Because such trials were relatively rare, as compared with the other trial types represented in the figure, we opted against applying a genuine statistical analysis. However, the figure does suggest an incremental effect when the probe matched the identity of both study faces (Δ) over and above what is seen when it matched the identity of only one study face (\square).

The vertical center of the horizontal gray bar in Figure 3 corresponds to the mean proportion of *yes* responses made when the probe’s identity, but not its emotion category, matched the identity of a study item; the thickness of that same bar corresponds to ± 1 standard error of the mean. This mean value provides a baseline that isolates the in-

fluence of shared facial identity. Note that when probe and study faces share the same emotion category, but the intensity of that expression differs by 67% (rightmost \bullet), *P*(Yes) falls to a level that would be expected from a match in identity alone. So, a face whose emotion category is the same as the probe’s but whose emotional intensity differs by 67% is no more likely to promote a false recognition than is a face that shares the probe’s identity but expresses an entirely different category of emotion.

EXPERIMENT 2

Results from various physiological studies have encouraged the view “that storage and rehearsal in working memory is accomplished by mechanisms that have evolved to work on problems of perception and action” (Jonides, Lacey, & Nee, 2005). More specifically, a number of studies have supported the idea that perception and short-term memory share essential resources (e.g., Agam & Sekuler, 2007; Kosslyn, Thompson, Kim, & Alpert, 1995; Postle, 2006). With behavioral measures alone, we could not directly test this hypothesis of shared resource, but we could test one of its implications—namely, a structural parallel between (1) the memory representation of facial emotion and identity, as revealed in Experiment 1, and (2) the perceptual representation of the same faces.

As a first step in exploring possible parallels between perceptual and memory representations, the subjects made perceptual judgments of faces’ similarity or dissimilarity to one another. These perceptual judgments, made while all the stimuli remained visible to the subjects, were transformed into a description of the perceptual space within which the faces were embedded. Finally, we determined whether the similarity relations within the perceptual description could account for the correct or false recognitions that the subjects made in Experiment 2’s tests of short-term memory. To generate the description of the faces’ perceptual representation, we used multidimensional scaling (MDS). Using only faces with neutral expressions, Yotsumoto et al. (2007) used this same method to characterize the perceptual similarity space for a set of Wilson faces comprising just four different identities, all female. Our first goal was to characterize the similarity space for the faces that had been used in the second phase of Experiment 1. These included not only several different male and female identities, but also expressions of various emotions and intensities of emotions. We then used the pairwise distances in this perceptual space in order to account for the short-term recognition memory results from Experiment 2.

Method

Stimuli. The stimuli were the same 60 combinations of facial identity and facial expression of emotion that were used earlier, in Experiment 1.

Subjects. All the subjects who served in Experiment 1, plus 1 additional female subject, served here. The subjects, 18–29 years of age, had normal or corrected-to-normal vision and were paid for their participation.²

Procedure. To quantify the similarity among faces of various identity–emotion combinations, we used the method of triads (Torg-

erson, 1958), a procedure that worked well previously, in a study using emotionally neutral Wilson faces (Yotsumoto et al., 2007). At trial onset, a fixation cross appeared in the center of the screen for 100 msec. After the cross disappeared, three faces appeared, all equidistant from the center of the screen and from each other. The subjects viewed the three faces and then used the computer mouse to make two responses. The first response selected the two faces that seemed perceptually most similar; the second response selected the two faces that seemed perceptually most dissimilar. In order to remind the subject of the current criterion, “similar” or “different” appeared in the center of the screen. Choices were made by mouse clicks on the desired faces. Once a particular face had been selected, a thin blue frame was inserted around it, which confirmed the selection. If the subjects misclicked on some face, the selection could be erased by clicking again on the same face. The subjects were not instructed to base their judgments on any particular characteristic (i.e., eyes, mouth, expression, hair) but were instructed to go at a relatively rapid pace so as to reduce in-depth scrutiny of each face. Note that all three faces remained visible to the subjects until both responses were made. As a result of the stimuli’s persistence, the subjects did not have to draw upon memory in making their similarity/dissimilarity responses; instead, responses could be based on the perceptual responses elicited by the stimuli.

With 60 different face stimuli, presenting each face with every other pair of faces would have required 205,320 trials ($60 \times 59 \times 58$) for just a single presentation of each possible trio. To reduce that number to a more manageable value, we resorted to a balanced-incomplete block design (BIBD; Weller & Romney, 1988). BIBD balances the number of times each *pair* of stimuli is seen but does not balance the frequency with which each *triad* is seen. Each time that the schedule called for some pair of faces to be presented, the face that accompanied them varied. As a result, repetitions were of pairs and not triads.

Each subject served in one session of 930 trials, which allowed every *pair* of faces to be presented six times. Each subject saw the 930 trials in a different order. In order to satisfy the stringent requirements for BIBD, we added 2 additional neutral faces to our set of 60 faces; the results from those neutral faces were discarded without analysis.³ Trials were self-paced.

The subjects’ responses were entered into a dissimilarity matrix in which the pair judged most similar was assigned a value of three, the pair judged most different was assigned a value of one, and the undesignated pair received an intermediate value, two. The individual subjects’ matrices were summed and then normalized by the frequency with which each pair of faces occurred during the experiment. The resulting matrix was processed by SPSS’s ALSCAL routine using a Euclidean distance model to relate dissimilarities in the data to corresponding distances in the perceptual representation. We evaluated the MDS solutions that were achieved in one through four dimensions.

Results and Discussion

To determine the appropriate number of dimensions for an MDS solution, we compared the goodness of fit achieved in varying numbers of dimensions. The results, expressed as values of stress (Borg & Groenen, 2005), can be seen in Figure 4. The fact that stress declines from a 1-D solution to a 2-D solution and remains relatively constant thereafter supports the idea that two dimensions constitute an appropriate space within which to analyze the MDS results. Using Kruskal’s rules of thumb as a guide (Borg & Groenen, 2005), the MDS fit to the data can be characterized as between *good* and *excellent*. The quality of the 2-D MDS solution is confirmed by Figure 5. The figure shows a scatterplot, known as a Shepard plot, where all 1,770 interpoint distances ($60 \times 59/2$) from the




Figure 4. Stress associated with multidimensional scaling solutions in varying numbers of dimensions. Stress, which is inversely related to goodness of fit, declines substantially from a 1-D to a 2-D solution and is roughly constant thereafter.

MDS solution have been plotted against their corresponding dissimilarities from the original data set. The pattern of scatter around values for which $x = y$ (white line in graph) gives no evidence of an unaccounted, nonlinear component.

Figure 6 is a graphical presentation of the MDS-based distribution of the 60 stimulus faces. As an aid to interpreting the distribution of points, we regionalized the plot in order to link the obvious geometric properties of the distribution to the known features of the objects represented by the plotted points (Borg & Groenen, 2005). On the basis of a visual inspection, we identified compact regions of interest within the MDS graphical solution. Note first that the points corresponding




Figure 5. Shepard plot based on the 2-D multidimensional scaling (MDS) solution. Points represent the interpoint distances from the MDS solution space versus the original dissimilarities among items. Note the relatively narrow scatter of points around the white 1:1 line and the absence of obvious systematic trends to the scatter.

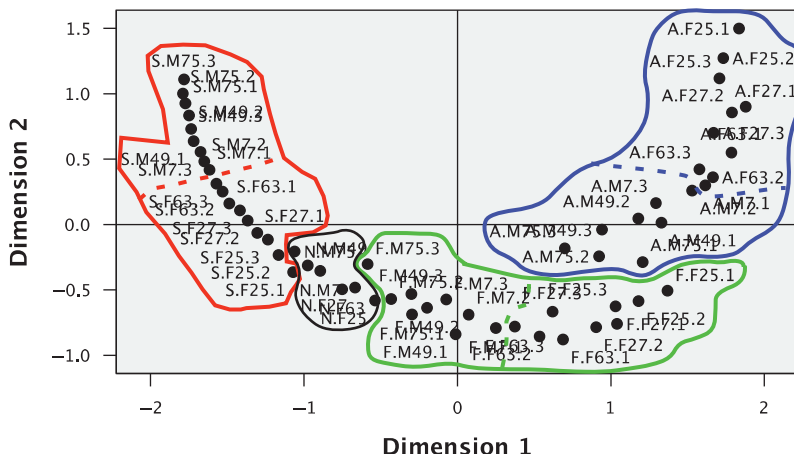


Figure 6. The 2-D multidimensional scaling solution derived from subjects’ triadic similarity judgments. Labels are of the form E.M/F#.1/2/3, where E is the first letter of the emotion’s designation (angry, fearful, neutral, sad), M/F# is the gender and identity number of the Wilson face, and 1/2/3 is the intensity of emotion (33%, 67%, and 100%). Red, black, green, and blue lines enclose points that represent sad, neutral, fearful, and angry faces, respectively. Within each region (except that for neutral faces), a dashed line separates male (M) faces from female (F) faces. Note that faces representing different intensities of the same emotion and facial identity are near each other and that the order of facial identity relations is maintained across expression types.

to each category of emotion form clusters that occupy distinct locations in the 2-D space. These clusters of emotion categories range from sadness (at the extreme left), through neutral, fear, and anger (at the extreme right). Note also that the clusters vary in their internal dispersions, with the points corresponding to neutral faces being more tightly clustered than the points corresponding to any of the other categories. Probably, this reflects the fact that with the neutral category, each face had just one exemplar, but with all other emotional categories, each face had three exemplars, one for each intensity of emotion.

At a finer level of analysis, one sees that within each category of emotion, the points representing the three exemplars of the same face (e.g., a particular male face expressing three levels of anger) are adjacent to one another. This confirms that both emotional intensity and facial identity entered into the subjects’ perceptual judgments of similarity/dissimilarity. Finally, excepting the relatively tight cluster for neutral faces, within each category of emotion, male faces are separated from female faces, as shown by the dashed lines that separate each emotion cluster into gender subsets. If one takes the location of any identity’s neutral exemplar as the starting point, the distance of that same identity’s faces in any of the emotion categories provides a rough index of that face’s ability to express that emotion. Interestingly, on this index, we see no consistent difference between exemplars corresponding to male and female faces. For example, although female faces seem to express fear and anger somewhat more vividly than do our male faces, male faces more vividly express sadness. We will return to this point later, in the General Discussion section.

Previous studies demonstrated that short-term recognition for various stimulus types is strongly dependent on the similarity or dissimilarity among items. This result has been shown for patches of differing hue (Nosofsky & Kantner, 2006), complex sounds (Visscher, Kaplan, Kahana, & Sekuler, 2007), Gabor patches (Kahana, Zhou, Geller, & Sekuler, 2007), and synthetic faces with neutral expressions (Sekuler, McLaughlin, Kahana, Wingfield, & Yotsumoto, 2006). This relationship is strongest when similarity/dissimilarity is expressed as a scalar value derived by summing the similarity or dissimilarity of a *p* item to each of the items in memory (Clark & Gronlund, 1996; Sekuler & Kahana, 2007). To assess this variable’s influence on short-term recognition memory for facial emotions, we regressed *P*(Yes) responses from Experiment 1 against the summed separations in MDS space of the *p* and each study face (Figure 6). Although both correct and false recognitions should be related to summed separation, that relationship ought to differ for the two types of responses. In particular, one expects that the proportion of correct recognitions will exceed the proportion of false recognitions. Therefore, we performed separate linear regressions for *P*(Yes) on target trials, where a *yes* response constituted a correct recognition, and on lure trials, where a *yes* response constituted a false recognition.

To carry out the regressions, we first ordered trials of each type according to the trial’s summed distance between *p* and study items. Then trials were separated into a number of equally populous bins, each representing a different mean level of summed distances. The mean proportions of hits and false recognitions were then calculated for the trials in each bin. Figure 7 shows the result of the

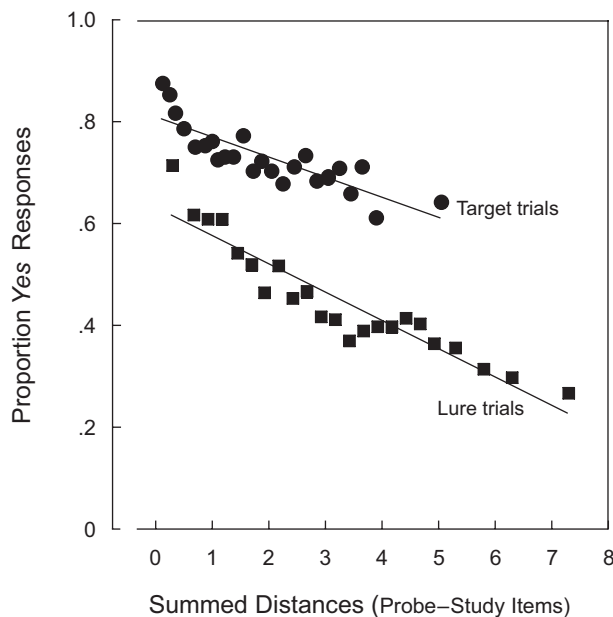


Figure 7. Proportions of *yes* responses on target trials (●, correct recognitions) and on lure trials (■, false recognitions) as a function of the summed distances of the probes to the study items. Data from individual trials were ordered and then put into equally populous bins for the purposes of linear regression.

linear regressions on target trials and on lure trials. Letting D represent the mean value of summed distances, the best-fitting lines were

$$P(\text{Yes}) = -.0398D + 0.812$$

and

$$P(\text{Yes}) = -.0556D + 0.633$$

for targets and lures, respectively. The difference in slopes was not reliable, since the 95% confidence intervals around the two values overlapped. For fits to target and lure trials, the root-mean squared deviations were 0.032 and 0.039; the values of r^2 associated with the two best-fitting lines were .71 and .88. Together, this pair of values indicate that D , the independent variable derived from simple, Euclidean measurements in the 2-D similarity space, does a reasonable job of predicting the proportion of recognitions, correct and false, produced by different combinations of study faces and probes.

To gauge how important it was to include not only facial emotion, but also facial identity, when computing regressor variables, we did two additional, unidimensional regression analyses. For these, regressors were generated by collapsing over the dimension of facial identity or over the dimension of emotion (category and intensity). Specifically, to collapse over emotion, we averaged all the MDS coordinates that were associated with a particular identity, and to collapse over identity, we averaged all the MDS coordinates that were associated with a particular category and intensity of emotion. With each set of unidimensional values as the starting points, a regression was done on the binned values of summed dissimilarity associated with individual trials.

When either dimension of the similarity space was muted as a regressor, the resulting values of r^2 decreased. When facial identity was muted, leaving emotion alone to enter into the computations, regressions for target and lure produced r^2 values of just .48 and .52. The substantial decline when facial identity was omitted from the perceptual space parallels the demonstration that facial identity influenced the subjects' memory-based judgments of facial emotion (see Figure 3). Finally, when facial emotion was muted, leaving only facial identity to enter into the computations, the quality of the model declined further, to just $r^2 = .32$ and .26, for target and lure trials, respectively. Although a precise numerical comparison is impossible, this weak but nonzero value of r^2 is reminiscent of the finding that facial identity alone, with no match on emotion whatever, attracts some small proportion of memory-based false recognitions (Figure 3).

The vertical separation between the two regression lines in Figure 7 suggests something of theoretical importance. As was explained above, some models of short-term memory (e.g., Clark & Gronlund, 1996; Sekuler & Kahana, 2007) assume that recognition responses are based on a probe item's summed similarity to the items in memory. In carrying out summed similarity, the regressions whose results are shown in Figure 7, we substituted summed Euclidean distances as measured in Figure 6 for the summed similarity variable. (This substitution avoided the need to make assumptions about the function relating similarity to distance.) If no variable other than summed distance governed recognition responses, one would expect that, for any given value of summed distance, $P(\text{Yes})$ on target trials would equal $P(\text{Yes})$ on lure trials. In terms of Figure 7, the mostly shared values of summed distance should cause the regression lines for lure and target trials to be superimposed. That the two are not superimposed suggests that some additional variable influenced the recognition responses. Whatever this additional variable might prove to be, its potency can be gauged from the substantial y -intercept difference between the two regression lines in Figure 7. This additional variable might take the form of contextual information—for example, information about the episodic or temporal context in which the probe face most recently appeared (Howard & Kahana, 2002). A complete characterization of this variable, though, will require additional experiments targeted at this hypothesis.

GENERAL DISCUSSION

The memory-based recognition judgments from Experiment 1 showed that emotional expression and facial identity are processed in combination with one another. Experiment 2 produced an analogous result, but for perceptual processing of the same set of faces. The fact that a metric 2-D perceptual representation gave a good account of memory performance suggests at least a rough structural parallel between the psychological space in which perceptual similarity/dissimilarity judgments are made and the psychological space used for short-term recognition memory.

This result extends a previous finding with emotion-neutral synthetic faces in which variation and blending of facial identities produced stimuli that lay along the orthogonal axes and along the diagonal of a 3-D space. In that study, Yotsumoto et al. (2007) showed that the MDS-derived perceptual space for these stimuli preserved much, but not all, of the structure that had been built into the stimuli and could be used in a stimulus-oriented account of recognition memory.

As was mentioned earlier, our application of MDS incorporated a metric, Euclidean distance model. Although the pattern of results in the Shepard diagram (Figure 5) is consistent with such a model, the use of this model rests on two assumptions that bear particular examination. First, the model assumes that different emotional expressions were perceptually equivalent in intensity. We did not think it was necessary to test that assumption ourselves, since others had already carried out such a test. Using a superset of our stimulus faces, Isaacowitz et al. (2006b) showed that the strongest expressions of different emotions were, in fact, comparable to one another in perceived intensity. Second, the model implicitly assumes at least an interval scale of measurement for the three levels of emotional intensity used in our experiments—that is, 33%, 67%, and 100% of the maximum. A direct, formal test of this assumption would require a massive, targeted data collection effort of its own. In place of such a test, we note that our existing results from Experiment 1 bear on the relationships among the three intensities of emotion. The horizontal axis in Figure 3 is calibrated in terms of the percentage of mismatch between the intensity of the probe face's emotion and the intensity of a study item face's emotion. Note that there is an approximately linear change in $P(\text{Yes})$ as the intensity mismatch increases from 0% through 67%. This result alone is not dispositive, but it is consistent with the assumption that the physical changes in emotional intensity produce comparable, corresponding changes in subjective intensity.

Parallels Between Memory and Perception

Experiments 1 and 2 point to possible structural parallels between the space in which faces are represented perceptually and the space in which faces are represented in short-term memory. Although far from dispositive, these parallels are consistent with an emergent-processes account of short-term memory. Such an account postulates that for temporary storage, working memory exploits brain regions that are shared with sensory processing (Jonides et al., 2008; Postle, 2006). There are some obvious limitations to our ability to elucidate possible parallels between perception and short-term memory for facial emotions. One limitation arises from the fact that the data presented here represent an average over subjects, over trials, and, in some cases, over diverse sets of study faces and probes. A second limitation is that Experiments 1 and 2 used a relatively narrow sample of facial identities and emotions. For one thing, as will be explained below, those experiments had to exclude Wilson faces that displayed smiles. In addition, the faces used in those experiments were derived from just the six individuals whose emotional expressions were judged to match best the intended target categories. Clearly, the pos-

sibility of parallels between processing in memory and perception should be tested further, with faces that represent even wider ranges of identities and emotional expressions.

Perceptual Processing

Experiment 1 began by identifying a set of *neutral* Wilson faces that were most reliably judged as expressing no emotion. In that process, the subjects attempted to categorize the emotion, if any, that was shown by each briefly presented face. A response was considered correct if it was consistent with the category of emotion that had been imparted to the face. As was expected, success in categorizing an emotion varied with the intensity of that emotion (Woodworth & Schlosberg, 1954). Our subjects made essentially no errors when an emotion was either 67% or 100% of the maximum for that emotion. However, the subjects did miscategorize displays of emotion when emotional intensity was weak—that is, when an emotion was just 33% of its maximum intensity. We believe that this result is not an artifact of our having used synthetic faces or even of our brief presentations. In fact, this finding reproduces classic results using a variety of face types and display durations (Woodworth & Schlosberg, 1954). More recently, Pollick, Hill, Calder, and Paterson (2003) used point-light displays of moving faces and demonstrated a small but significant relationship between duration and perceived emotional expression, with longer durations generally producing emotions that were judged more intense. In an account of the circuitry responsible for recognition of emotions, Adolphs (2002) described a possible neural basis for the sharpening of emotion judgments with duration—namely, time-varying feedforward and feedback interactions among components in a distributed network of brain regions.

It is interesting to note that in classic studies, when subjects were asked to categorize weak displays of emotion, not only were their judgments frequently wrong, but also those judgments were easily manipulated—for example, by an intentionally misleading label given by the researcher (Fernberger, 1928). These powerful emotion-priming effects are not easily accommodated by simple face recognition models that postulate a strong modularity of the participating components (e.g., Bruce & Young, 1986). However, by adding reciprocal feedforward and feedback connections between components, an interactive version of the basic model could accommodate top-down modulation by nonvisual factors.

Our characterization of perceptual processing was based on similarity/dissimilarity judgments with triads of faces (Experiment 2). The perceptual space produced from subjects' judgments by MDS (Figure 6) suggested interesting differences in the emotional displays between male and female faces. In particular, within each emotion category, male and female faces occupied distinct but neighboring locations in the perceptual space. As an illustration, consider faces' locations within the sadness category. There, male faces occupy more extreme locations, relative to their neutral positions, than do female faces, but the ordering by gender is reversed within the fear and anger categories. We cannot tell whether these results reflect genuine gender-related differences or are somehow a result of our having initially culled faces from the large library of faces. Recall

that we limited our stimuli to faces that most convincingly appeared as neutral when they were intended to be neutral. To determine whether the gender-related differences were genuine would require an extensive, random collection of faces that had not been selected, as in Experiment 1, against some potentially biasing characteristic. It may well be that some sexually dimorphic characteristics of human faces make it differentially easy for male and female faces to express diagnostic signs of particular emotions (Becker, Kenrick, Neuberg, Blackwell, & Smith, 2007).

Short-Term Memory

In Experiment 2, recognitions of an expressed emotion increased when the probe face's identity matched the identity of one or more items in memory. So, even though facial identity was notionally irrelevant to the task, it still exerted a strong influence on subjects' recognition responses. This result, which is illustrated in Figure 3, was confirmed by the regression analysis in which the recognition results were far better fit using a regressor that took both identity and expression into account than using a regressor that treated the two as independent dimensions.

Our study examined faces that were expressing the emotions of anger, sadness, and fear but intentionally omitted faces with happy expressions. Such faces were excluded from our stimulus set because many nominally "happy" Wilson faces struck pilot subjects as peculiar—that is, not genuinely happy. The smiles were described as inauthentic or insincere, what others have dubbed *professional* smiles. Such smiles are generated out of courtesy, rather than as an expression of happiness (Wallace, 1997, p. 289). We believe the perceived lack of sincerity arose from the fact that Wilson faces' lips were not permitted to part. In order to control simple, luminance-based cues that might, on their own, differentiate one emotion from another, Goren and Wilson's (2006) algorithm for expression generation foreclosed a display of teeth, even when the zygomaticus major muscle, a principal participant in smiles, would have parted the lips of an actual face. This design decision was consequential because, when human judges attempt to identify a happy expression, they make important use of the parting of the lips (Smith, Cottrell, Gosselin, & Schyns, 2005).

In Experiment 2, the subjects most accurately recognized a probe's *emotion* when that probe's *identity* matched the identity of both study items (see Figure 3). When two study items happened to share the same identity, our experimental design forced those two faces to display emotions drawn from different categories. Essentially, this constraint meant that as the subjects were encoding what they had seen, they were exposed to two different examples of the same person's face. In principle, these two examples might enable subjects to extract some consistent facial structures from the two study faces, which could be helpful in evaluating the probe face's emotion. In everyday social encounters, we often need to read and make sense of emotions that are expressed on the faces of strangers—that is, people with whom our prior experience is limited or even nonexistent. In such circumstances, we do not have access to the optimum neutral baseline—that same person's own face when it is at rest. Therefore, we must judge and react to the emotional

expression more or less on its own, in an absolute framework that very likely is influenced by long-term memory of previous encounters with other people. Previous experience making emotion judgments might explain Elfenbein and Ambady's (2002) finding that facial emotions expressed by members of one culture are more accurately perceived and interpreted by members of that same culture than they are by people who are not members of that culture. It may be that this difference arises from experience-based learning of the characteristics of that culture's "typical" face(s) when the faces are physically at rest.

As was just mentioned, experience with particular faces or types of faces seems to boost perceptual interpretation of the emotional expressions on those faces. It is important to note, therefore, that both our perceptual and memory measurements were made with faces that were unfamiliar to the subjects, at least upon initial viewing. Although the best-known neuropsychological tests of face recognition memory depend on face stimuli that are not familiar to people being tested (Duchaine & Weidenfeld, 2003), many key results for face perception and memory come from studies in which at least some of the faces were familiar to the subjects (e.g., Bredart & Devue, 2006; Bruce & Young, 1986; Ganel & Goshen-Gottstein, 2004; Ge, Luo, Nishimura, & Lee, 2003; Kaufmann & Schweinberger, 2004; Laeng & Rouw, 2001; Troje & Kersten, 1999). The distinction between familiar and novel faces may bear on the parallels we observed between perception and recognition memory. Arguably, increasing familiarity with a particular face allows recognition memory to approach a limit set by a subject's ability to discriminate visually some change in the face. Thus, with images of Mao Tse-Tung as stimuli, Ge et al. demonstrated that Chinese subjects' ability to recognize small distortions of Mao's face from memory alone closely matched the subjects' ability to discriminate such distortions perceptually. This striking convergence of recognition memory and perception was confirmed using the faces of people who, although not famous, were highly familiar to particular subjects (Bredart & Devue, 2006).

Representing Facial Emotion

Figure 6 shows the 2-D MDS description of the perceived similarities among our stimuli, which varied in facial emotion, facial identity, and gender. This representation of similarities among multidimensional stimuli can be compared with some well-known characterizations of the psychological space representing emotion. For example, on the basis of categorical judgments of facial expressions shown in a series of photographs, Schlosberg (1952) proposed a 2-D representation for emotion. That representation, inspired by Newton's circular color surface, was spanned by two principal, opponent axes (pleasant–unpleasant and attention–rejection). In analogy to the location of desaturated colors in Newton's color system, neutral expressions occupied a position at the center of the surface. Even a casual comparison of Schlosberg's space and our own suggests that the two are not well correlated. For example, in Figure 6, judgments of neutral expressions do not occupy a special position central to all other emotions but are sandwiched between the judgments for two other emotions.

J. A. Russell (1980) derived an alternative representation of affective space, starting from categorical judgments of emotion-related words. Although this space shared the dimensionality of Schlosberg's space, in it emotions were not mapped onto a circular *surface*, but onto a circumplex; that is, different emotions were systematically arranged around the perimeter of a circle (for a rigorous description of circumplex structure, see Revelle & Acton, 2004). Note that in a circumplex model, a neutral emotion would occupy a location on the circumference of the circle, surrounded by other emotions (compare Figure 6). Russell's basic result has been supported by several subsequent studies, including ones in which subjects judged the similarity of photographs of facial expressions (see Posner, Russell, & Peterson, 2005, for a review).

The path of points in Figure 6's MDS solution suggests a roughly circular arc. The absence of smiling faces from Experiment 2's stimulus set might account for some of the obvious gap in a path that otherwise would have been consistent with the circumplex model. Potential data from other emotions that were omitted from our stimulus set could account for another portion of the gap. However, some interemotion proximities in Figure 6 are not consistent with their counterparts in J. A. Russell's (1980) original circumplex account of affective states or with a recent revision (Posner et al., 2005) of that account. Most notably, as compared with the standard circumplex model, in Figure 6, the locations occupied by angry faces and fearful faces are interchanged. Neither this discrepancy nor the obvious gap in what might be a fragment of a circular path bear strongly on the validity of a circumplex model of emotion. After all, the judgments upon which our MDS solution was based reflect variation not only in category of emotion, but also in intensity of emotion, facial identity, and gender. Unavoidably, the outcome of the scaling process reflects not only the underlying structure of psychological or affective processes, but also the way in which those underlying processes happen to be sampled by whatever particular set of stimuli is used to probe that structure (e.g., Bosten, Robinson, Jordan, & Mollon, 2005). That limitation points the way toward future studies that expand the analysis presented here by including a more complete range of emotions and by studying a subject population whose representational spaces might be atypical (e.g., D'Argembeau & van der Linden, 2004; Habak, Wilkinson, & Wilson, 2008; Riby, Doherty-Sneddon, & Bruce, 2008). Such studies might also usefully clarify the nature of the linkage between representations of facial identity and facial emotion. A clarification of this point would have considerable theoretical value, since there are several different mechanisms by which two perceptually distinct aspects of some remembered stimulus might appear to influence one another. For example, a stimulus's distinct attributes could be bound together into a single unified multidimensional representation. This is the arrangement typically used in various computational models of memory (Estes, 1994; M. S. Humphreys, Pike, Bain, & Tehan, 1989; Sekuler & Kahana, 2007). In an alternative formulation, the representations of a stimulus's aspects are kept separate—say, in different, specialized regions of the brain or within differ-

ent populations of neurons within the same region. In a recognition experiment, each separate representation would then be compared against the representation of the probe stimulus's corresponding aspect, and the aggregate of all the comparisons would be used to guide the recognition response. A third and final formulation assigns one stimulus aspect priority over the other aspects. Specifically, the representation of one aspect in the probe provides an index, or *address*, to the representations of the study items. This hypothesis, which owes an obvious debt to computer architecture, resembles a number of previous proposals in the literature on human memory (e.g., Hollingworth & Henderson, 2002; Treisman & Zhang, 2006; Wheeler & Treisman, 2002). Note that in this account, although one aspect has not been encoded as an integral part of study items' representations, that aspect does have an indirect influence on recognition by guiding the retrieval of study items' representations. Selecting among these three competing models, any one of which could produce results like those we presented, requires a behavioral experiment that is designed and targeted for such selection or, possibly, an appropriately designed study with functional neuroimaging.

As was mentioned earlier, our experiments' results are consistent with the proposition that perception of faces and memory for faces draw on a common, shared representation of similarity. Moreover, this shared similarity representation serves both facial identity and emotional expression. This outcome was a surprising one. After all, distinct computational demands are associated with perception and with memory. Similarly, distinct neural computations are needed to recognize a face's identity (i.e., who the face's owner is) and to interpret a face's emotional expression (i.e., what its owner's affective state is; Tsao & Livingstone, 2008). Some behavioral studies are consistent with the hypothesis that face identity and emotion draw upon the same representation (e.g., Ellamil et al., 2008). This shared representation of face similarity corresponds to the gateway stage postulated by recent multistage models of face processing. The initial stage in such models comprises brain regions that are sensitive to changes in facial structure, regardless of whether those changes are generated by differences between the identities of two faces' owners or by changes in the emotional expression of a single face. In such models, this initial, structure-sensitive stage of processing is followed by a stage or stages that produce categorical information about identity or emotional expression. A recent study using fMRI adaptation gives strong support to models of this class. Exploiting the fact that the BOLD signal declines as a stimulus is repeated, Fox, Moon, Iaria, and Barton (2009) examined this adaptation of fMRI signals as subjects viewed stimulus sequences of morphed faces. These morphed faces were designed to simulate varying degrees of changing identity or changing emotional expression. In support of the idea that emotion and identity exploit a shared similarity space, Fox et al. (2009) identified an area in the occipital cortex that was sensitive to structural changes in a face, independently of whether those changes reflected a change in identity or in emotion. Activation patterns in entirely other brain areas were associated with subjects' judgment of which category

of change was being displayed. However valuable such studies are, they represent only a beginning of an attempt to understand how humans process, remember, and assess the faces of their conspecifics. In particular, these studies only begin to shed light on the neural circuits and computations that are required in order to translate facial similarities into full-fledged, memory-dependent judgments, including judgments that incorporate episodic and biographical information about a face's owner (van Vugt et al., 2009). In fact, behavioral as well as functional neuroimaging studies suggest that the initial stage of processing, which is shared by some facial attributes and functions, comprises just the raw material for complex, temporal interactions that are carried out by a spatially distributed network of specialized brain regions (Vuilleumier & Pourtois, 2007).

AUTHOR NOTE

This research was supported by National Institutes of Health Grants MH068404, MH55687, MH61975, and EY002158, and by Canadian Institute for Health Research Grant 172103. Many thanks to Shivakumar Viswanathan and Deborah Goren for contributions to the research reported here. Jason J. S. Barton made valuable comments on an earlier version of the manuscript. This report is based on data collected at Brandeis for M.G.'s master's thesis in Neuroscience. Correspondence concerning this article should be addressed to R. Sekuler, Volen Center for Complex Systems, Brandeis University, Waltham, MA 02454 (e-mail: sekuler@brandeis.edu).

REFERENCES

- ADOLPHS, R. (2002). Neural systems for recognizing emotion. *Current Opinion in Neurobiology*, **12**, 169-177.
- AGAM, Y., & SEKULER, R. (2007). Interactions between working memory and visual perception: An ERP/EEG study. *NeuroImage*, **36**, 933-942.
- ASHBY, F. G., & MADDOX, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 598-612.
- BAUDOIN, J.-Y., MARTIN, F., TIBERGHIE, G., VERLUT, I., & FRANCK, N. (2002). Selective attention to facial emotion and identity in schizophrenia. *Neuropsychologia*, **40**, 503-511.
- BECKER, D. V., KENRICK, D. T., NEUBERG, S. L., BLACKWELL, K. C., & SMITH, D. M. (2007). The confounded nature of angry men and happy women. *Journal of Personality & Social Psychology*, **92**, 179-190.
- BORG, I., & GROENEN, P. J. F. (2005). *Modern multidimensional scaling: Theory and applications* (2nd ed.). New York: Springer.
- BOSTEN, J. M., ROBINSON, J. D., JORDAN, G., & MOLLON, J. D. (2005). Multidimensional scaling reveals a color dimension unique to "color-deficient" observers. *Current Biology*, **15**, R950-R952.
- BRAINARD, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, **10**, 443-446.
- BREDART, S., & DEVUE, C. (2006). The accuracy of memory for faces of personally known individuals. *Perception*, **35**, 101-106.
- BRUCE, V., & YOUNG, A. (1986). Understanding face recognition. *British Journal of Psychology*, **77**, 305-327.
- CLARK, S. E., & GRONLUND, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, **3**, 37-60.
- D'ARSEMBEAU, A., & VAN DER LINDEN, M. (2004). Identity but not expression memory for unfamiliar faces is affected by ageing. *Memory*, **12**, 644-654.
- DARWIN, C. (1872). *The expression of the emotions in man and animals*. London: Murray.
- DE GELDER, B., POURTOIS, G., & WEISKRANTZ, L. (2002). Fear recognition in the voice is modulated by unconsciously recognized facial expressions but not by unconsciously recognized affective pictures. *Proceedings of the National Academy of Sciences*, **99**, 4121-4126.
- DE GELDER, B., VROOMEN, J., POURTOIS, G., & WEISKRANTZ, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *NeuroReport*, **10**, 3759-3763.
- DUCHAIINE, B. C., PARKER, H., & NAKAYAMA, K. (2003). Normal recognition of emotion in a prosopagnosic. *Perception*, **32**, 827-838.
- DUCHAIINE, B. C., & WEIDENFELD, A. (2003). An evaluation of two commonly used tests of unfamiliar face recognition. *Neuropsychologia*, **41**, 713-720.
- EKMAN, P. (1993). Facial expression and emotion. *American Psychologist*, **48**, 384-392.
- EKMAN, P., & FRIESEN, W. V. (1975). *Unmasking the face: A guide to recognizing emotions from facial clues*. Englewood Cliffs, NJ: Prentice Hall.
- ELFENBEIN, H. A., & AMBADY, N. (2002). On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychological Bulletin*, **128**, 203-235.
- ELLAMIL, M., SUSSKIND, J. M., & ANDERSON, A. K. (2008). Examinations of identity invariance in facial expression adaptation. *Cognitive, Affective, & Behavioral Neuroscience*, **8**, 273-281.
- ESTES, W. K. (1994). *Classification and cognition*. Oxford: Oxford University Press.
- FERNBERGER, S. W. (1928). False suggestion and the Piderit model. *American Journal of Psychology*, **40**, 562-568.
- FISHER, R. A., & YATES, F. (1963). *Statistical tables for biological, agricultural, and medical research* (6th ed.). New York: Hafner.
- FOX, C. J., MOON, S. Y., IARIA, G., & BARTON, J. J. S. (2009). The correlates of subjective perception of identity and expression in the face network: An fMRI adaptation study. *NeuroImage*, **44**, 569-580.
- FOX, C. J., ORUC, I., & BARTON, J. J. S. (2008). It doesn't matter how you feel: The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision*, **8**, 11.1-11.13.
- FUNG, H. H., LU, A. Y., GOREN, D., ISAACOWITZ, D. M., WADLINGER, H. A., & WILSON, H. R. (2008). Age-related positivity enhancement is not universal: Older Chinese look away from positive stimuli. *Psychology & Aging*, **23**, 440-446.
- GANEL, T., & GOSHEN-GOTTSTEIN, Y. (2004). Effects of familiarity on the perceptual integrality of the identity and expression of faces: The parallel-route hypothesis revisited. *Journal of Experimental Psychology: Human Perception & Performance*, **30**, 583-597.
- GARNER, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- GE, L., LUO, J., NISHIMURA, M., & LEE, K. (2003). The lasting impression of chairman Mao: Hyperfidelity of familiar-face memory. *Perception*, **32**, 601-614.
- GOREN, D., & WILSON, H. R. (2006). Quantifying facial expression recognition across viewing conditions. *Vision Research*, **46**, 1253-1262.
- GOSSELIN, P., & KIROUAC, G. (1995). Decoding facial emotional prototypes. *Canadian Journal of Experimental Psychology*, **49**, 313-329.
- HABAK, C., WILKINSON, F., & WILSON, H. R. (2008). Aging disrupts the neural transformations that link facial identity across views. *Vision Research*, **48**, 9-15.
- HASSELMO, M. E., ROLLS, E. T., & BAYLIS, G. C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behavioural Brain Research*, **32**, 203-218.
- HAXBY, J. V., HOFFMAN, E. A., & GOBBINI, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, **4**, 223-233.
- HEFTER, R. L., MANOACH, D. S., & BARTON, J. J. S. (2005). Perception of facial expression and facial identity in subjects with social developmental disorders. *Neurology*, **65**, 1620-1625.
- HOLLINGWORTH, A., & HENDERSON, J. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception & Performance*, **28**, 113-136.
- HOWARD, M. W., & KAHANA, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, **46**, 269-299.
- HUMPHREYS, G. W., DONNELLY, N., & RIDDOCH, M. J. (1993). Expression is computed separately from facial identity, and it is computed separately for moving and static faces: Neuropsychological evidence. *Neuropsychologia*, **31**, 173-181.
- HUMPHREYS, M. S., PIKE, R., BAIN, J. D., & TEHAN, G. (1989). Global matching: A comparison of the SAM, Minerva II, Matrix, and TODAM models. *Journal of Mathematical Psychology*, **33**, 36-67.
- ISAACOWITZ, D. M., TONER, K., GOREN, D., & WILSON, H. R. (2008). Looking while unhappy: Mood-congruent gaze in young adults, positive gaze in older adults. *Psychological Science*, **19**, 848-853.

- ISAACOWITZ, D. M., WADLINGER, H. A., GOREN, D., & WILSON, H. R. (2006a). Is there an age-related positivity effect in visual attention? A comparison of two methodologies. *Emotion*, *6*, 511-516.
- ISAACOWITZ, D. M., WADLINGER, H. A., GOREN, D., & WILSON, H. R. (2006b). Selective preference in visual fixation away from negative images in old age? An eye-tracking study. *Psychology & Aging*, *21*, 40-48.
- JONIDES, J., LACEY, S., & NEE, D. (2005). Processes of working memory in mind and brain. *Current Directions in Psychological Science*, *14*, 2-5.
- JONIDES, J., LEWIS, R. L., NEE, D. E., LUSTIG, C. A., BERMAN, M. G., & MOORE, K. S. (2008). The mind and brain of short-term memory. *Annual Review of Psychology*, *59*, 193-224.
- KAHANA, M. J., ZHOU, F., GELLER, A. S., & SEKULER, R. (2007). Lure-similarity affects visual episodic recognition: Detailed tests of a noisy exemplar model. *Memory & Cognition*, *35*, 1222-1232.
- KAUFMANN, J. M., & SCHWEINBERGER, S. R. (2004). Expression influences the recognition of familiar faces. *Perception*, *33*, 399-408.
- KOSSLYN, S. M., THOMPSON, W. L., KIM, I. J., & ALPERT, N. M. (1995). Topographical representations of mental images in primary visual cortex. *Nature*, *378*, 496-498.
- LAENG, B., & ROUW, R. (2001). Canonical views of faces and the cerebral hemispheres. *Laterality*, *6*, 193-224.
- LOFFLER, G., GORDON, G. E., WILKINSON, F., GOREN, D., & WILSON, H. R. (2005). Configural masking of faces: Evidence for high-level interactions in face perception. *Vision Research*, *45*, 2287-2297.
- LOFFLER, G., YOURGANOV, G., WILKINSON, F., & WILSON, H. R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, *8*, 1386-1391.
- NETH, D., & MARTINEZ, A. M. (2009). Emotion perception in emotionless face images suggests a norm-based representation. *Journal of Vision*, *9*, 5.1-5.11.
- NOSOFSKY, R. M., & KANTNER, J. (2006). Exemplar similarity, study list homogeneity, and short-term perceptual recognition. *Memory & Cognition*, *34*, 112-124.
- PANTELIS, P. C., VAN VUGT, M. K., SEKULER, R., WILSON, H. R., & KAHANA, M. J. (2008). Why are some people's names easier to learn than others? The effects of face similarity on memory for face-name associations. *Memory & Cognition*, *36*, 1182-1195.
- POLLICK, F. E., HILL, H., CALDER, A., & PATERSON, H. (2003). Recognising facial expression from spatially and temporally modified movements. *Perception*, *32*, 813-826.
- POSAMENTIER, M. T., & ABDI, H. (2003). Processing faces and facial expressions. *Neuropsychology Review*, *13*, 113-143.
- POSNER, J., RUSSELL, J. A., & PETERSON, B. S. (2005). The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development & Psychopathology*, *17*, 715-734.
- POSTLE, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*, 23-38.
- REVELLE, W., & ACTON, G. S. (2004). Evaluation of ten psychometric criteria for circumplex structure. *Methods of Psychological Research Online*, *9*, 1-27.
- RIBY, D. M., DOHERTY-SNEDDON, G., & BRUCE, V. (2008). Exploring face perception in disorders of development: Evidence from Williams syndrome and autism. *Journal of Neuropsychology*, *2*, 47-64.
- RUSSELL, J. A. (1980). A circumplex model of affect. *Journal of Personality & Social Psychology*, *39*, 1161-1178.
- RUSSELL, R., BIEDERMAN, I., NEDERHOUSER, M., & SINHA, P. (2007). The utility of surface reflectance for the recognition of upright and inverted faces. *Vision Research*, *47*, 157-165.
- SACKS, O. (1985). *The man who mistook his wife for a hat*. New York: Harper-Collins.
- SADR, J., JARUDI, I., & SINHA, P. (2003). The role of eyebrows in face recognition. *Perception*, *32*, 285-293.
- SCHLOSBERG, H. (1952). The description of facial emotion in terms of two dimensions. *Journal of Experimental Psychology*, *44*, 229-237.
- SEKULER, R., & KAHANA, M. J. (2007). A stimulus-oriented approach to memory. *Current Directions in Psychological Science*, *16*, 305-310.
- SEKULER, R., MCLAUGHLIN, C., KAHANA, M. J., WINGFIELD, A., & YOTSUMOTO, Y. (2006). Visual recognition and temporal order information are both well-preserved in aging. *Psychology & Aging*, *21*, 632-637.
- SERGET, J., OHTA, S., MACDONALD, B., & ZUCK, E. (1994). Segregated processing of facial identity and emotion in the human brain: A PET study. *Visual Cognition*, *1*, 349-369.
- SMITH, M. L., COTTRELL, G. W., GOSSELIN, F., & SCHYNS, P. G. (2005). Transmitting and decoding facial expressions. *Psychological Science*, *16*, 184-189.
- STEVENS, S. S. (ED.) (1951). *Handbook of experimental psychology*. New York: Wiley.
- TORGERSON, W. S. (1958). *Theory and methods of scaling*. New York: Wiley.
- TREISMAN, A., & ZHANG, W. (2006). Location and binding in visual working memory. *Memory & Cognition*, *34*, 1704-1719.
- TROJE, N. F., & KERSTEN, D. (1999). Viewpoint-dependent recognition of familiar faces. *Perception*, *28*, 483-487.
- TAO, D. Y., & LIVINGSTONE, M. S. (2008). Mechanisms of face perception. *Annual Review of Neuroscience*, *31*, 411-437.
- VAN VUGT, M. K., SCHULZE-BONHAGE, A., SEKULER, R., LITT, B., BRANDT, A., BALTUCH, G., & KAHANA, M. J. (2009). Intracranial electroencephalography reveals two distinct similarity effects during item recognition. *Brain Research*, *1299*, 33-44.
- VISSCHER, K. M., KAPLAN, E., KAHANA, M. J., & SEKULER, R. (2007). Auditory short-term memory behaves like visual short-term memory. *Public Library of Science-Biology*, *5*, e56.
- VUILLEUMIER, P., & POURTOIS, G. (2007). Distributed and interactive brain mechanisms during emotion face perception: Evidence from functional neuroimaging. *Neuropsychologia*, *45*, 174-194.
- WALLACE, D. F. (1997). *A supposedly fun thing I'll never do again: Essays and arguments*. Boston: Little, Brown.
- WELLER, S. C., & ROMNEY, A. K. (1988). *Systematic data collection* (Vol. 10). Newbury Park, CA: Sage.
- WELLS, G. L. (1993). What do we know about eyewitness identification? *American Psychologist*, *48*, 553-571.
- WHALEN, P. J., KAGAN, J., COOK, R. G., DAVIS, F. C., KIM, H., POLIS, S., ET AL. (2004). Human amygdala responsivity to masked fearful eye whites. *Science*, *306*, 2061.
- WHEELER, M. E., & TREISMAN, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, *131*, 48-64.
- WILSON, H. R., LOFFLER, G., & WILKINSON, F. (2002). Synthetic faces, face cubes, and the geometry of face space. *Vision Research*, *42*, 2909-2923.
- WILSON, H. R., & WILKINSON, F. (2002). Symmetry perception: A novel approach for biological shapes. *Vision Research*, *42*, 589-597.
- WOODWORTH, R. S., & SCHLOSBERG, H. (1954). *Experimental psychology* (2nd ed.). New York: Holt.
- YOTSUMOTO, Y., KAHANA, M. J., WILSON, H. R., & SEKULER, R. (2007). Recognition memory for realistic synthetic faces. *Memory & Cognition*, *35*, 1233-1244.

NOTES

1. Ashby and Maddox (1990) cast their argument in terms of the distinction between separable and integral dimensions (Garner, 1974). To avoid an unnecessary theoretical commitment, we will use the more common term, *independent*, as a substitute for *separable*.
2. Experiment 2 was run before Experiment 1. This was done so that the subjects in Experiment 2 would not have been previously exposed to the synthetic faces.
3. BIBDs cannot be constructed for every possible number of stimulus triplets, including 60. Adding two extra neutral faces allowed us to construct a BIBD that could be used with the number of stimuli we had (Fisher & Yates, 1963).

SUPPLEMENTAL MATERIAL

The stimuli from this article in .tif format, as well as sample stimuli of male and female faces with various emotions and intensity levels, may be downloaded from <http://cabn.psychonomic-journals.org/content/supplemental>.