Provided for non-commercial research and educational use only.
Not for reproduction or distribution or commercial use.

VISION RESEARCH

An International Journal for Functional Aspects of Vision

Biochemistry & Cell Biology • Molecular Biology & Genetics
Anatomy, Physiology, Pathology & Pharmacology • Optics, Accommodation & Refractive Error
Circuitry & Pathways • Psychophysics • Perception • Attention & Cognition
Computational Vision • Eye Movements & Visuomotor Control

ISSN 0042-6989 | Volume 47 | Number 7 | March 2007

This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

http://www.elsevier.com/locate/permissionusematerial

Available online at www.sciencedirect.com

ScienceDirect

Vision Research 47 (2007) 965–973

Vision Research

www.elsevier.com/locate/visres

Authors personal copy

The effects of perceptual history on memory of visual objects

Son Preminger, Dov Sagi *, Misha Tsodyks

Department of Neurobiology, The Weizmann Institute of Science, Rehovot 76100, Israel

Received 12 September 2006; received in revised form 24 November 2006

Abstract

We investigated how the recognition and perception of memory-stored visual objects are influenced by cumulative experience with similar stimuli. The memory of a face was established by training observers to identify a set of faces as either “friends” or “non-friends”. Subsequently, for multiple daily sessions, observers continued to perform this identification task, in which presented faces included a sequence of morphed faces, gradually transforming from a friend face (*source*) to another initially distinguishable non-friend face (*target*), interleaved with other faces. Initially observers identified only the first part of the morph sequence as “friends”. In experimental conditions for which the initial “friends” portion was at least 54% of the sequence, this portion increased along repeated daily practice, until eventually most of the sequence was identified as “friends”. After this practice, perceived similarity between source and target faces was much higher than the average similarity between the other face images. These effects did not occur when the morph images were shown in random order using a similar protocol. In addition, corresponding recognition confusions between source and target faces were found. Our findings suggest that memories of objects can be changed as a result of exposure to similar stimuli and show the dependency of these changes on the order in which stimuli are presented and on their level of similarity.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Long-term memory; Face perception; Learning; Visual object

1. Introduction

The theoretical view of long-term memory as associative neural networks (Amit, 1992; Hebb, 1949; Hopfiled, 1982) suggests that memories emerge through learning by overlapping patterns of synaptic modifications adhering to the Hebbian paradigm (Hebb, 1949). The storage of multiple memories in a common network, encoded by common synaptic connections, implies that the formation or modification of a memory of one object may modify the representation of another object. Thus, interactions between memories stored in the same network are anticipated by associative network models, suggesting effects of perceptual experience on existing memories. In these models, random uncorrelated patterns are usually used to represent the stored memories. It was shown (Amit, Gutfreund, & Sompolinsky, 1985) that storing too many patterns in the same network results in unreliable memory storage due to increasing effect of random correlations between the patterns. In our recent theoretical study we have shown that a sequence of correlated memory patterns, transforming from one random pattern to another, may collapse into a single unified representation (Blumenfeld, Preminger, Sagi, & Tsodyks, 2006). However, although increased interaction between correlated patterns is indicated by theory, interactions between similar (correlated) stimuli stored into memory have not been shown experimentally. In addition to similarity, another factor that might affect representation is the order of exposure to stimuli. Interestingly, while standard theory portrays memory representations which depend on which stimuli are stored into memory but not on the order in which they are stored (Hopfiled, 1982), more recent models suggest dependency of representation on the order of exposure (Blumenfeld et al., 2006; Griniasty, Tsodyks, & Amit, 1993; Mongillo, Amit, & Brunel, 2003; Wallis & Rolls, 1997).

---

* Corresponding author. Fax: +972 8 934 4131.
E-mail address: dov.sagi@weizmann.ac.il (D. Sagi).

0042-6989/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2007.01.003

Author's personal copy

Here we test if the memory of a visual object (face) can be modified by experience with other similar objects. Further, we examine whether such memory changes depend on the order in which the stimuli are presented, and whether these changes affect recognition and perception. In our psychophysical paradigm, we first establish a small number of memories by training observers to recognize four faces as "friends" (enforced with feedback). Subsequently, observers engage in a practice lasting multiple daily sessions, in which they are exposed to morphed faces between one of the pre-learned friend faces (source) and another face, whilst performing a friend/non-friend identification task on these morphed faces and other unrelated faces (see Section 2). To prevent any supervised change of the pre-learned stimulus-response mapping, no feedback is provided during this task. The first faces of the morph sequence, being very similar to the source, are bound to be identified as "friends" while others are distant enough to be identified as "non-friends". The identification transition point (threshold) is considered here to correlate with the boundary of the related memory, that is, the memory corresponding to the source face. We are interested in changes of this threshold with practice and in the dependence of these threshold changes on experimental parameters such as order of presentation of the morphed faces and their similarity. Following this practice, we perform identification and perceptual-similarity tests on the encountered faces, to evaluate the scope of effects due to the exposure to the morph-sequence. While intuition derived from perceptual learning predicts decreasing confusion between the different faces of the morph sequence, associative learning models predict the formation of strong associations between the faces, resulting in increased confusion.

There are some reports in the literature pointing to effects of perceptual history on object perception and memory. Recently, psychophysical studies have shown that adaptation to a face can result in changes to the perceived identity of related faces shortly after removal of the adapting face (Anderson & Wilson, 2005; Leopold, O'Toole, Vetter, & Blanz, 2001; Rhodes & Jeffery, 2006). While these findings demonstrate short-term effects on the perception of objects, long-term effects of perceptual history have also been found, mostly between stimuli presented in immediate temporal adjacency. Recordings of single neuron activity in monkeys performing match-to-sample tasks indicate that images learned in temporal proximity acquire correlated neuronal activities (Ericson & Desimone, 1999; Miyashita, 1988; Sakai & Miyashita, 1991). In humans, it was shown that temporal adjacency is used to group 3D face-views into a common facial identity (Wallis & Bulthoff, 2001). These findings reflect the brain's ability to group together objects that are presented consecutively, forming temporal association between them. The current study is designed to examine the ability of the brain to associate between visual percepts that are similar to each other, even if they are not presented adjacently. To accomplish this, the morphed faces are presented interleaved with other faces, preventing the generation of short-term temporal associations between them.

2. Methods

2.1. Observers

Observers were students between 17 and 27 years old, with normal or corrected-to-normal vision. They were not familiar with the stimuli, were naïve to the objective of the experiment, and were paid for their participation. Twenty-four observers (21 females, 3 males) participated in the Learning phase as a preparation stage for the Gradual Morph experiment, 16 of them (15 females, 1 male) subsequently performed the Gradual Morph practice (see Section 2.3.2 below). Eight observers (8 females) participated in the Learning phase as preparation for the Mixed Morph practice, 5 of them performed the Mixed Morph practice (see Section 2.3.4).

2.2. Stimuli and apparatus

2.2.1. Faces

Stimuli were grayscale, frontal view images of Caucasian male faces with neutral expression, with no facial hair and no glasses. Stimuli were constructed from images of the FERET database (Phillips, Moon, Rizvi, & Rauss, 2000), cropped by a uniform oval to remove hairline. To ensure that there are no conspicuous features causing trivial discrimination between faces, one face of the database was morphed with all the other faces, and the resulting midpoint merged faces were used for the experiment (sample faces, see Fig. 1a). Outlier faces were removed based on pairwise similarity ratings by naïve observers, to make sure that the constructed faces were not too similar or too different from each other. These face stimuli were shown in one of 3 slightly different blur levels, blur level being chosen randomly at each presentation. The 3 blur levels included the constructed face image plus two blurred versions of it which were produced by Adobe Photoshop 6 using the Gaussian blur function with 0.4 and 0.6 pixels radius (see example in Fig. 1b – although it is difficult to notice the difference, one might notice for example the change in the black spot on the left cheek or the white area in the right eye). These blur versions where constructed to subjectively match the blurriness of morph sequence images which appear slightly more blurred towards the middle of the sequence (example in Fig. 1b, M-26%, M-50%). This was done in order to make sure that in the parts of the experiment where the face stimuli were presented interleaved with morph-sequence images (see below), observers could not differentiate between morph-sequence images and regular face stimuli based on blurriness.

2.2.2. Stimuli presentation

In all experiments faces were shown for 200 ms, preceded by a 200 ms mask (randomly shuffled face fragments of 10 × 10 pixels) and a 200 ms white frame. Stimuli were displayed on a Mitsubishi Diamond Pro 2060 22″ monitor, using a PC with Intel processor. Throughout all experiments screen background was of gray color with luminance of 56 cd/m^2. Face stimuli were presented on this background in a white frame (7.87° × 10.87°) with the face oval image (7.12° height × 5.25° width) placed at its center. Average face luminance was 10.5 cd/m^2, luminance of white frame was 99 cd/m^2. Observers used a computer mouse for response in all tasks except for the similarity judgment test (see Section 2.3) in which they used a keyboard.

2.2.3. Morph-sequence

A sequence of face images gradually changing from a *source* face stimulus to another *target* face stimulus (Fig. 1d) generated using MorphMan 4 software. Each morph-sequence contained 50 intermediate faces between source and target, with each intermediate morph face shown twice. Thus, a morph-sequence includes a total of 100 face images. Throughout the experiment, five different *source-target pairs* were used and the morph-sequences between them were generated. Source-target pairs were chosen such that they had close to average perceived similarity based on pairwise

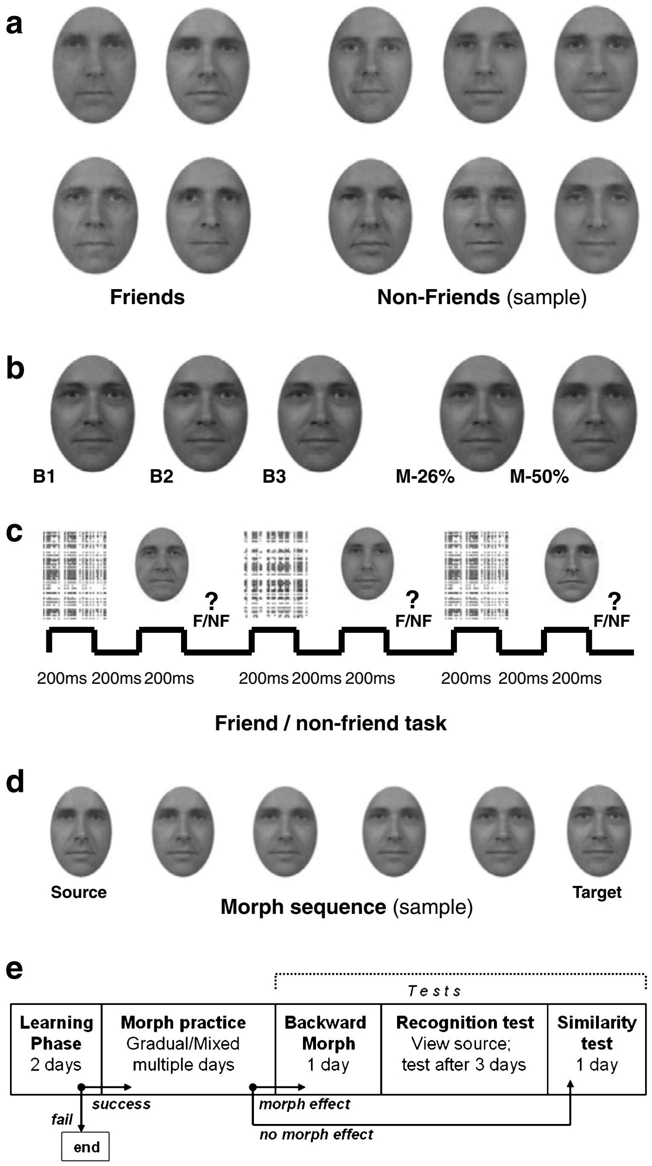

Fig. 1. Learning "friends" and the Gradual Morph protocol. (a) Sample images of "friends" and "non-friends". (b) Image blur levels: B1 is an example of a regular face image, B2 and B3 are slightly blurred versions of it matching blur levels of morph images. M-26% and M-50% are examples of a morph face at 26% and 50% of morph sequence (respectively). (c) Friend/Non-Friend task. (d) Example of a morph-sequence from source to target face. In the experiments morph-sequences were of length 100. (e) Flowchart of experiment. Morph practice could be either Gradual or Mixed.

similarity ratings by naïve observers ($n = 6$; number of faces 44.; rating on a 1–5 scale [1 = "very different" up to 5 = "exactly the same"], mean similarity rate 2.58; STD 0.46; similarity rates of source-target pairs 1–5 is 2.5, 2.33, 2.5, 2.67, 2.5, respectively).

2.3. Procedures

As outlined above, our experimental paradigm included three main parts: Learning phase, morph practice and Post-Morph tests. Fig. 1e shows a flowchart of the experiment.

2.3.1. Learning phase

Twenty-four observers were trained to identify 4 faces as "friends" as opposed to other "non-friend" faces. Initially images of the 4 friend faces were presented and observers were asked to watch the faces and remember them as "friends". Each friend was presented 3 consecutive times, followed by a triplet of images of another friend. Each friend-triplet was shown 4 times, with triplets ordered randomly ([3 images of particular friend] × [4 friends] × [4 times]). Subsequently, observers performed the *Friend/Non-Friend* task (*FNF task*) **with feedback** – images of one of the 4 "friends" or one of 20 other "non-friend" faces were presented at random, observers were asked to make a friend/non-friend decision on each face (Fig. 1a and c). Sound feedback was provided for wrong responses. This learning procedure was repeated in 2 daily sessions of 1 h. At the end of this training, 16 observers (15 females, 1 male) performed the FNF task without any errors during 64 consecutive trials which included randomly ordered 8 presentations of each friend and 32 presentations of various non-friends. Only these 16 error-free observers continued to the next phase, the other 8 observers who had errors, did not continue with the experiment.

2.3.2. Gradual Morph (GM) practice

At the day after completing the Learning phase successfully, observers ($n = 16$) continued to perform the FNF task, **without feedback**, in daily sessions of approximately 700 trials. Three types of stimuli were shown: the images of a morph-sequence (Fig. 1d) between one of the pre-learned friends (*source*) gradually changing to an initially unfamiliar face (*target*) (100 trials), 3 regular friends (the other friends except for source, 250 trials) and non-friend faces (including the non-friends from the Learning phase plus additional 10 non-friend faces, ~350 trials). The type of face to be shown at each trial was chosen randomly. Thus, trials of morphed faces were usually separated by one or more trials of friend or non-friend faces, preventing perceptual comparison between morph images. In the 100 trials of the morph-sequence, morph images were ordered from source to target, each morph image shown twice followed by the next morph image towards target; In the regular friends trials, 2 friends were shown for 100 trials each and another friend was shown for 50 trials to counterbalance the fact that it is perceived as if the source friend disappears in the middle of the session.

This GM practice was repeated until the observer reached stable performance; if no convergence trend was seen for 14–15 days, training was terminated (except for 2 cases in which training was terminated earlier due to technical reasons). In the first day of the GM practice, morph images were shown changing gradually from source to target, and then gradually back from target to source (instead of the regular gradual order shown in all other days of the GM practice, here each intermediate morph-image was shown once on the way up and once on the way down). This first-day procedure was performed for all the observers except for two observers who performed the regular GM practice in the first day as in the other days, and one observer who saw the morph sequence in random order in the first day, followed by the GM practice in the other days (no significant difference in results was found). The experiment was preformed on 5 source-target pairs (8 observers on pair 1, 7 observers on pair 2, with 2 observers performing on both pairs, and one observer for each of the pairs 3, 4, 5, see Stimuli). Observers were told that the purpose of the GM practice was to test how their memory of the friends evolves over time. They were instructed daily to identify each face-image as friend/non-friend by comparing image identity to the identity of the learned "friend" faces. To track how FNF identification along morph-sequence images evolves throughout practice sessions, we defined the *FNF-threshold*: responses to the 100 ordered morph images were binned into 10 bins and normalized to produce fraction of "friend" responses as a function of morph index; the FNF-threshold was estimated as the morph point [0–1] yielding 50% "friend" response.

2.3.3. Post-Morph tests

After completion of the Gradual Morph practice, observers performed several tests to evaluate the effect of the practice on the long-term memory and the perception of faces. Any observer who reached an FNF-threshold close to 1 towards the end of the GM practice (i.e. identified most morph images as "friend", we call this a *morph effect*), initially performed the following 2 tests (see Fig. 1e).

2.3.3.1. Backward Morph test.

To verify that the identification of most of the morph-sequence as a "friend" is not solely due to within-session

Author's personal copy

Author's personal copy

effects, observers performed one session of the FNF-task with stimuli as in the GM practice but with the morph-sequence presented backwards, from target to source. A *staircase-like rule* determined which index of the morph-sequence was presented next: for any "non-friend" response on a morph-sequence image, the index went downwards (towards source); after 6 consecutive "friend" responses on the same index the index went upwards (towards target). The test terminated after 6 "friend" answers on the target or after a maximum of 100 exposures to the morph-sequence.

2.3.3.2. Recognition test. Observers were presented with the source face after being instructed to remember it for later recognition. After 3 days, they were presented with various faces shown one after the other and had to answer for each face, whether this is the person they had memorized (Yes/No). The images were randomly ordered friends, non-friends and images from the morph-sequence, starting from the **target**. A staircase-like rule (as described in the Backward Morph test) determined which index of the morph-sequence will be shown next.

To examine the effect of the GM practice on perception, observers who participated in the GM practice performed the Similarity Judgment test (observers who reached a threshold close to 1, performed the Similarity test after completing the Backward Morph and Recognition tests).

2.3.3.3. Similarity Judgment test. Observers rated the perceived similarity between pairs of faces. For each pair, faces were presented sequentially and then the observer rated their similarity on a 1–5 scale (1 = "very different" up to 5 = "exactly the same"). Pairs of faces included a total of 18 faces: 4 friends (including source), 1 target face, 8 familiar non-friends, 5 unfamiliar non-friends (all possible pairs; each pair presented 4 times, twice in each direction).

2.3.4. Mixed Morph

To examine the importance of the order of presentation of morph images, a separate group of 5 observers was exposed to the morph-sequence in a different order: after successfully completing the Learning phase, these observers performed the Mixed Morph practice, which was identical to the GM practice with the exception that at every daily session the images of the morph-sequence were presented in a random order instead of gradual order. This Mixed Morph practice was followed by a Similarity Judgment test.

3. Results

3.1. Gradual/Mixed Morph practice – Identification of morph sequence images (FNF-thresholds)

Fig. 2a depicts friend/non-friend (FNF) responses on images of the morph-sequence during the GM practice (top 4 rows). Typically, in the first daily session (top line in the response plot of each observer), observers answered "friend" on the first part of the morph-sequence (close to source) and "non-friend" on the second part (closer to target). The observers presented in the two top rows demonstrate the *morph effect* – from day to day, responses on the morph faces gradually changed, until most responses on the morphed faces, including the target face, were "friend" (bottom lines of the response plots). Although occurring in variable paces, a clear morph effect is demonstrated in 10 (1st and 2nd rows) out of 18 cases of the GM practice (3rd and 4th rows show cases where the effect did not occur).

The FNF-threshold (the point on the morph-sequence of transition from "friend" to "non-friend", see Section 2) obtained in the first daily session appears to be a critical parameter determining whether the morph effect will take place. Fig. 2b shows that a morph effect required an initial threshold of 0.54 or above (red circles, obtained with pair 1 by six observers, with pair 2 by two observers, and with pair 3 and pair 4 by one observer each), while 0.54 or lower thresholds resulted in a failure of the morphing procedure (blue squares, five observers with pair 2, two observers with pair 1, one observer with pair 5). Fig. 2c shows average FNF-thresholds along training sessions for *high-threshold* (above 0.54, red) and *low-threshold* (below 0.54, blue) GM observers. In addition, Fig. 2 shows that none of the observers who performed Mixed Morph protocol (see Section 2) exhibited a morph effect, regardless of their initial FNF thresholds ($n = 5$, Fig. 2a, bottom row, and Fig. 2b, green diamonds), even though high-threshold Mixed Morph and Gradual Morph observers have started with similar average thresholds (Fig. 2c, green and red lines, respectively). Thus, the order of exposure to the morph-sequence appears to be a critical factor for the morph effect to occur. Importantly, throughout the GM practice there was no significant change in responses to other friends (false negative rates were usually below 5%, except for two observers (Ob2, Ob13) who starting from the first sessions of the practice identified one of the control friend as a non-friend most of the times), and the average false alarm rate was less than 5% with errors distributed evenly along practice days and within sessions.

3.2. Post-Morph tests

Nine out of the ten observers who had the morph effect performed the Backward Morph and Recognition tests to verify the effect of the practice on long-term memory (one observer was terminated due to technical reasons).

3.2.1. Backward Morph test

Out of the 9 observers with morph effect, 6 have identified the target as a friend six consecutive times (thus, on all their exposures to target) and for another 3 the highest image index identified as friend was close to target (index 0.86, 0.82, and 0.72, where index 0 is source, 1 is the target). The average highest image index identified as a friend was 0.9 ($SE = 0.04$, $n = 9$ observers). For all observers, false positive response rates on any of the other images were between 0% and 1%. Thus, since observers tend to identify the target face as a "friend" even when it is presented at the beginning of the experiment we conclude that the morph effect is not merely due to within-session effects. This result strengthens the long-term nature of the morph effect.

3.2.2. Recognition test

Results of the recognition test closely matched the results of the Backward Morph test. Out of the 9 observers with morph effect, 5 identified the target as the person they had to remember (the source) six consecutive times (thus, on all their exposures to target) and 2 have identified a morph image close to target (index 0.86 and 0.78) as the memorized person. The average highest image index

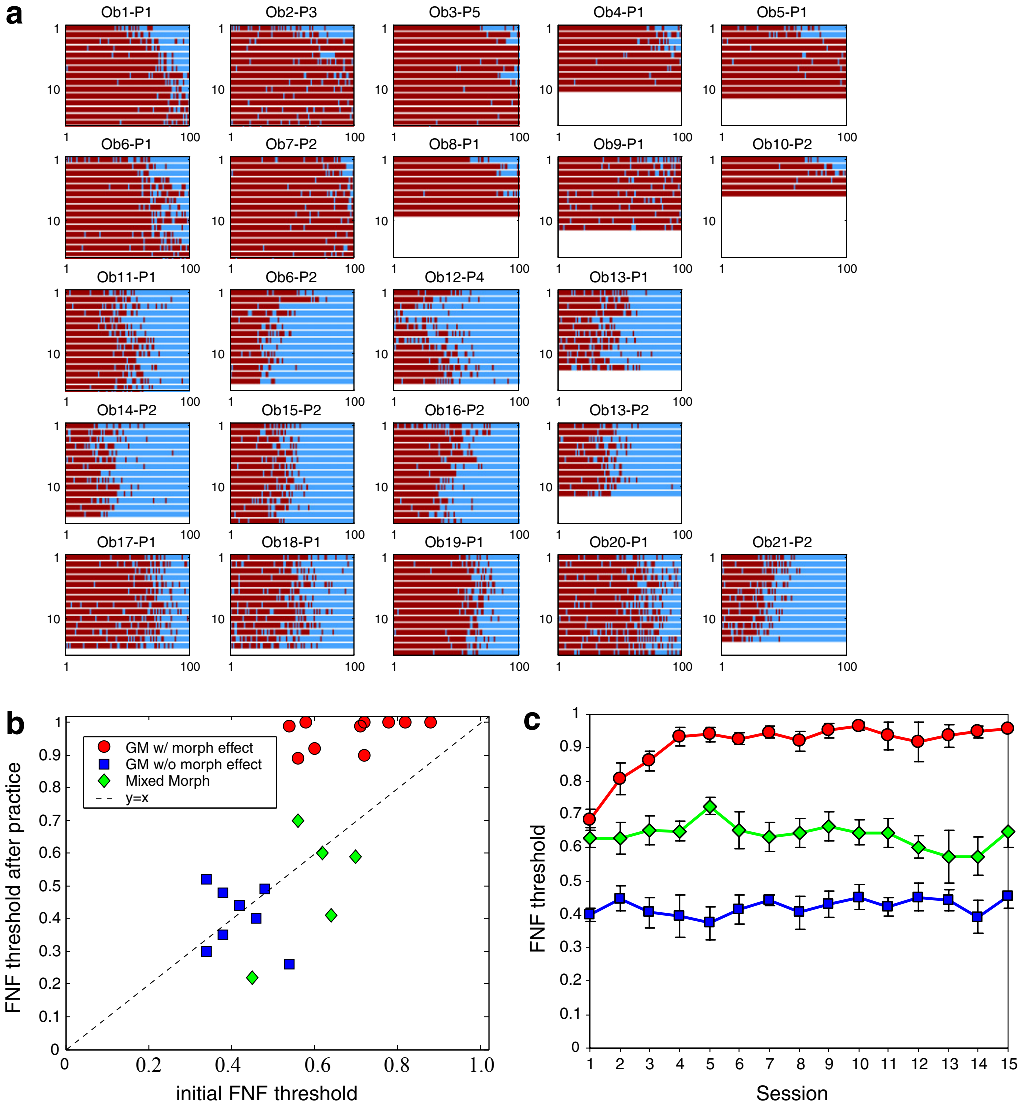

Fig. 2. Response to FNF task during Morph practice. (a) Response plots show FNF responses on morph-sequence images (x-axis) in the first 15 days (y-axis) of GM practice for all observers; dark red – “friend” response, light blue – “non-friend”. Each plot is headed by observer number (Ob) and pair number (P). (b) The relationship between the FNF threshold before and after practice. Each point depicts initial (x-axis) and ending (y-axis) thresholds for one morph practice (1 Ob with 1 source-target pair). Results are shown for GM practice (red circles for positive results (morph effect) and blue squares for negative results) and for the Mixed-Morph (green diamonds) observers. (c) Average thresholds of GM observers starting with high threshold (>0.54; red circles, top line; days 1–6 $n = 9$, days 7–9 $n = 8$, day 10 $n = 7$, day 11 $n = 6$, days 12–15 $n = 4$), low threshold (<0.54; blue squares, bottom line; days 1–11 $n = 7$, day 12 $n = 6$, days 13–14 $n = 5$, day 15 $n = 3$), and Mixed Morph with high threshold (>0.54; green diamonds, middle line; days 1–14 $n = 4$, day 15 $n = 2$). Two GM observers with threshold = 0.54 were excluded (Ob1-P1 had a morph effect, Ob6-P2 did not). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this paper.)

identified as the source was 0.8 ($SE = 0.1$, $n = 9$ observers). For all observers, false positive response rates on any of the other images were between 0% and 1%, except for Ob2 who confused one control friend with the source. These results suggest a change in the long-term memory of the source face, in many cases resulting in recognizing the initially distinguishable source and target faces as the same face.

Observers who performed the Gradual and Mixed Morph practices performed the Similarity Judgment test at the end of the experiment in order to test the effect of practice on perception.

3.2.3. Similarity Judgment test – Gradual vs. Mixed

The results of the Similarity Judgment test are presented in Fig. 3. Statistical significance was assessed using a-parametric tests: Wilcoxon test for within group comparisons, Mann-Witney U test for between groups comparisons. Similarity scores for source-target pair by high-threshold GM observers (see Section 3.1) were significantly higher than the similarity scores they gave to other pairs of faces (3A,C –average scores, $n = 8$; 4.2 ± 0.2 vs. 1.9 ± 0.02, mean ± SE; $p < 0.01$). The average source-target similarity score given by high-threshold GM observers was more than twice as high as the average score given by high-threshold

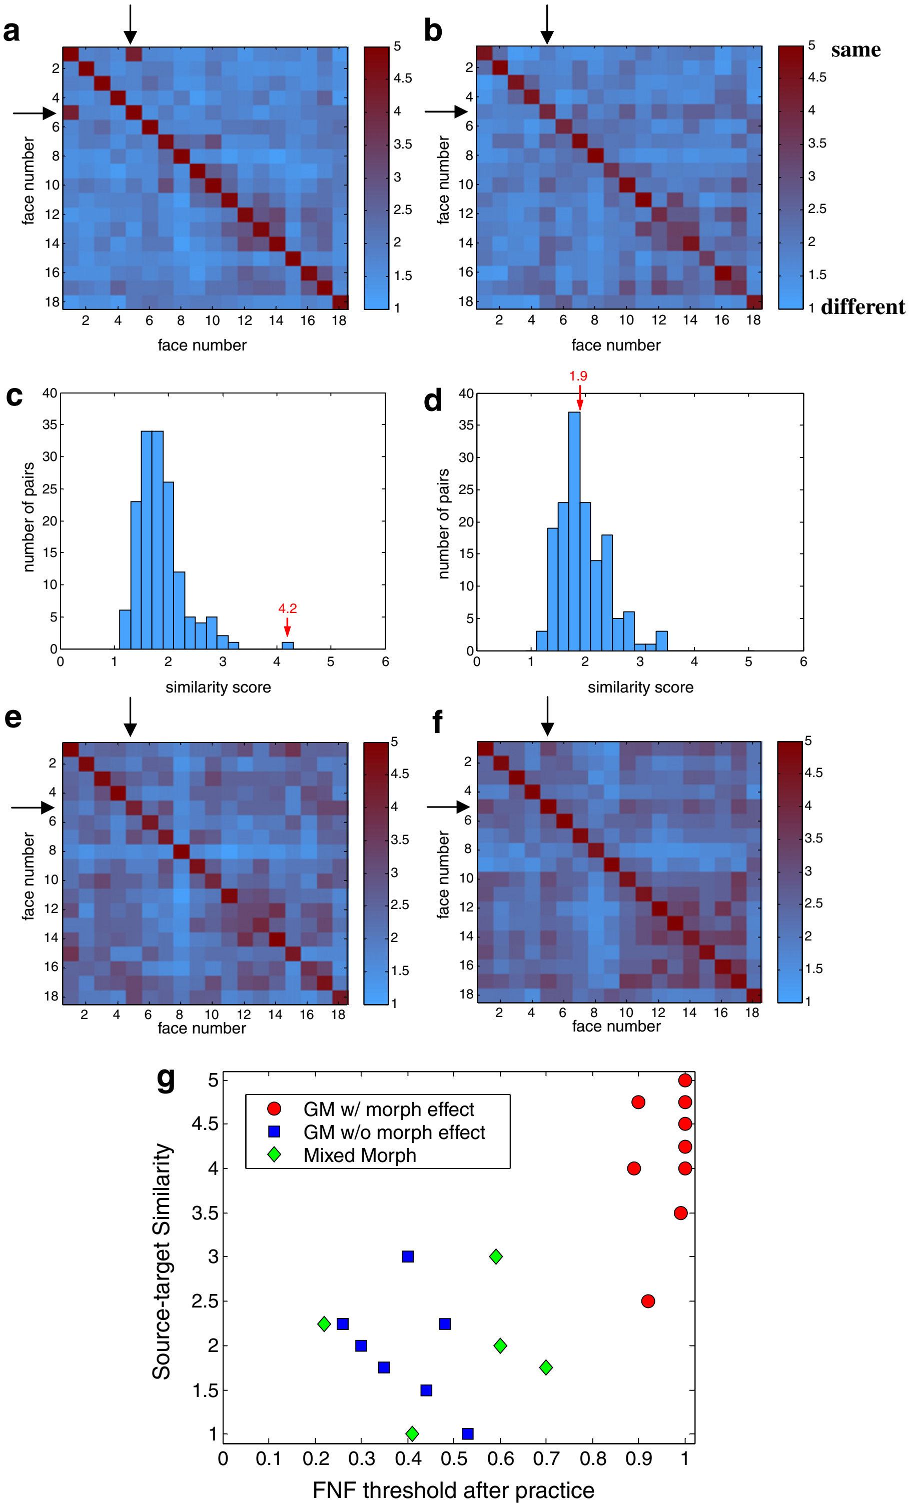

Fig. 3. Results of Similarity Judgment test. (a) and (b) Average similarity scores by high-threshold GM observers (a) and high-threshold Mixed Morph observers (b). Each cell in the matrix shows a similarity score between two faces averaged across subjects. Similarity score for source-target (pair [1,5]) is pointed by arrows. Face indexes: 1–4: friends (1 is source); 5 is target; 7–10: non-friends that appeared in learning and in morph procedure; 11–14: non-friends that appeared only in morph procedure; 6, 15–18: non-familiar faces. (c) and (d) Histogram of similarity scores by high-threshold GM observers (c) and by high-threshold Mixed Morph observers (d) of all pairs of faces that were rated in the similarity test. The arrow marks the score for the source-target pair. (e) Average similarity scores by naïve observers ($n = 5$); source-target score 2.2 ± 0.26 vs. 2.3 ± 0.07 mean of all other pairs 1.9 ± 0.02. (f) Average similarity scores by observers who performed the Learning phase ($n = 5$); source-target score 3.2 ± 0.1 vs. 2.4 ± 0.1 mean of all other pairs 1.9 ± 0.02. (g) The relationship between the similarity score of the source-target pair and the FNF-threshold after practice. Each point depicts similarity score (x-axis) and ending threshold (y-axis) for one morph practice (1 Ob with 1 source-target pair). Results are shown for GM practice (red circles for positive results (morph effect) and blue squares for negative results) and for the Mixed-Morph (green diamonds) observers. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this paper.)

Author's personal copy

Mixed Morph observers ($n = 4$; 1.9 ± 0.02, which is very similar to the average score given by high-threshold Mixed Morph observers for all other face pairs 2.0 ± 0.09; $p > 0.5$, Fig. 3b and d). The similarity gain (score for source-target pair minus average score on other pairs) of high-threshold GM observers was significantly higher than the similarity gain of high-threshold Mixed Morph observers ($p < 0.01$). In addition, the similarity gain by high-threshold GM observers was significantly higher compared with the gain obtained by two other separate control groups – naïve observers who performed only the similarity tests ($n = 5$, $p < 0.005$, Fig. 3e) and observers who performed the similarity tests after completing the Learning phase ($n = 5$, $p < 0.01$, Fig. 3f). Importantly, as can be seen in Fig. 3g, there is a close correspondence between the FNF threshold obtained at the end of the morph practice and the similarity score given to the source-target pair. These results show that the GM practice changed the perceived similarity between the source and target faces, causing them to be perceived as very similar. Note that similarity scores between friends (within regular friends 1.74 ± 0.35 (no. regular friends = 3), between source and regular friends 1.76 ± 0.26;, mean ± SE over observers ($n = 8$)) are not statistically different from similarity scores between friends and non-friends (between regular friends and non-friends (no. non-friends = 8) 1.7 ± 0.21; $p > 0.5$; between source and non-friends 1.63 ± 0.14; $p > 0.5$), pointing to an absence of response category effects on the similarity judgments (Livingston, Andrews, & Harnad, 1998). Thus we conclude that the increased similarity between source and target is due to the morphing process.

4. Discussion

Our findings show that under certain conditions, repeated exposure to a sequence of morphed faces between a pre-learned *source* face and another *target* face results in a long-term change in identification of morphed faces and in increased source-target perceived similarity. Altered performance in both the trained and untrained tasks, suggests that the long-term memory of the source face was modified over the course of the experiment, demonstrating the interaction between memories of visually similar stimuli. Furthermore, since this effect occurred only when morphed faces were presented in order of increasing visual distance from the source image, we conclude that this interaction depends on the order of presentation. Finally, the fact that this so called "morph effect" occurred only when source and target images were not perceived as too different (having an initial FNF threshold 0.54 or larger) hints to the limits of the plasticity of the memory system.

Previous studies have demonstrated interactions between memories of temporally adjacent stimuli (Ericson & Desimone, 1999; Miyashita, 1988; Sakai & Miyashita, 1991; Wallis & Bulthoff, 2001). Such temporal associations can be explained by Hebbian-based mechanisms which strengthen connections between neurons that fire in temporal proximity (Griniasty et al., 1993; Mongillo et al., 2003; Wallis & Rolls, 1997). In the current experiment, the morphed faces are not presented consecutively, rather are separated by other faces. Thus, our results point to the existence of different long-term memory mechanisms, which enable linkage between visually similar stimuli even if the stimuli are presented interleaved with other stimuli, as often happens in natural viewing. Our experimental results also show that this "similarity-association" mechanism depends on the order in which the morphed stimuli are presented, producing different results for gradual vs. random presentation. The dependence on order indicates the existence of history-dependent modifications occurring in the brain at the time of learning. Thus, the magnitude of synaptic update due to visual exposure to a stimulus depends not only on the current stimulus (as in Hebbian learning) but also on previous experience, captured by the existing memory representation. In our recent theoretical work (Blumenfeld et al., 2006) we showed that the merging of a morph-sequence into a single memory representation inevitably results from interactions induced by the correlations between stored patterns in an associative neural network model. We proposed and analyzed a history-dependent learning mechanism that is based on novelty-facilitated memory modifications. This mechanism can counter-balance the merging effect and enable order-dependent learning. Other models of unsupervised learning, such as approaches introduced by statistical machine learning (Pearl, 1988), could also potentially account for the dependence on the order of presentation as long as the update induced by learning in each step takes into account previous learning.

An important result of this study is the strong interaction found here between perception and memory. Similarity measurements showed a strong effect of memory on the comparison between two stimuli presented consecutively. This result suggests that memory affects perceptual qualities and argues against a pure stimulus derived metric within which perceptual similarities are computed. Instead, perceptual judgments on faces apparently involve higher level memory-based analysis. A recent study used backward masking with synthetic faces to show a relatively slow processing time (130 ms) for face discrimination (Loffler, Gordon, Wilkinson, Goren, & Wilson, 2005), suggesting that discrimination (and thus similarity judgments) is affected by higher-level processes that depend on previous experience with the specific stimuli. While experience-dependent effects, traditionally studied within the context of perceptual learning, are thought to improve discrimination performance, here we show the opposite: experience with stimuli increases their perceived similarity and decreases their discriminability. Our finding that these changes in perceived similarity correlate with changes in identification and recognition supports a strong link between recognition and subjective similarity (Ashby & Perrin, 1988).

The dependency of the morph effect on the initial FNF-threshold suggests that although grouping of perceptual stimuli into a single object memory is dynamic and modifiable, initial perceptual similarity constrains this bundling, determining which object instances could be bundled together. The mechanisms that underlie similarity perception and allow the grouping of different visual percepts into a single object (object invariance) have been extensively studied in the literature. Some theories propose that grouping of object instances is based on metric properties (e.g. Bulthoff, Edelman, & Tarr, 1995; Poggio & Edelman, 1990; Ullman, 1996), while others suggest that non-accidental structural properties underlie object invariant representation (Biederman, 1987; Biederman & Bar, 1999; Vogels, Biederman, Bar, & Lorincz, 2001). Our research focuses on the effect of perceptual history and context on object memory, suggesting that in addition to their physical and structural basis, invariant qualities (whether structural 3D attributes or pictorial attributes) can also be modulated by the particular way in which object instances are experienced.

Some studies (Goldstone, 1994; Goldston, Lippa, & Shiffrin, 2001; Kurtz & Gentner, 1998; Livingston et al., 1998) report that learning to categorize objects often results in acquired distinctiveness between categories or acquired equivalence within categories. In our study we have demonstrated that the GM practice can induce increased similarity between source and target, as well as cause morph images to gradually become identified as "friends". One could attempt to interpret our results as pure categorical perception effect, construing the FNF task as a classical categorization task (friends and non-friends). However, our results do not agree with the predictions inferred from this view (with the morph sequence and the friends considered a category). As described in the Results section, our protocol does not induce increased similarity between friends or decreased similarity between friends and non-friends. Furthermore, we show two sets of conditions under which the morph effect does not occur: Mixed Morph and GM with initial low-threshold. The discrepancy between the results found by our protocol and the classical categorization protocols are probably due to major differences in experimental paradigms, including: training duration (many daily sessions in our case vs. usually single or a few sessions in the other studies), no feedback in our morph practice as opposed to feedback training in other studies, one learned category - friends (without any overt commonality) as opposed to categorization into two or more categories (sometimes rule based). Thus, simply interpreting the FNF task as a classical categorization task is insufficient for explaining our results. Despite this, we believe that the mechanisms underlying category learning and the morph effect are of a similar nature. We propose (consistently with our model in Blumenfeld et al., 2006) that at the time of exposure to a morphed face, if it is similar enough to the current memory of the source, the morphed face is associated with the memory of source. We show in our model how this association could be implemented by strengthening the representation of the morphed stimulus proportionally to its novelty relative to the existing memory of the source. If this process is repeated enough times, the original memory of the source is modified until eventually all morphed faces have a unified memory representation. Our experimental paradigm is designed to track this dynamic change of memory representation.

Our results, which shed a light on the process of updating long-term memory of objects, are also consistent with recent theories of memory consolidation positing that upon retrieval, memories become fragile and thus can be modified (Walker, Brakefield, Hobson, & Stickgold, 2003; for review, Dudai, 2004; Nadar, 2003). Our practice procedure and its modifications could be used to study these and other contextual effects on memory.

Short-term effects of perceptual history on identification of faces were recently demonstrated in studies showing aftereffects of adaptation to a face which cause a perceptual bias away from the adapted face right after its removal (Anderson & Wilson, 2005; Leopold et al., 2001; Rhodes & Jeffery, 2006; Webster, Kaping, Mizokami, & Duhamel, 2004). The relation between these short-term adaptation effects and the long-term effects presented in the current study remains to be investigated.

Author's personal copy

Acknowledgments

We thank G. Fuhrman, Y. Adini, R. Raba, H. Gelbard, O. Barak, S. Gilad, Y. Rydzinski and A. Loebel for contribution to the project, B. Blumenfeld for help with statistics, Y. Dudai, R. Malach, S. Ullman, D.A. Levy and S. Gilad for comments on earlier versions of the paper. Supported by the Pearl M. Levine Trust, Anne P. Loderer Research Institute, and Grodetsky Center for Higher Brain Functions. This study used the FERET database of facial images collected under the FERET program.

References

Amit, D. (1992). *Modeling brain function: The world of attractor neural networks.* Cambridge University Press.

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Storing infinite number of patterns in a spin glass model for neural networks. *Physical Review Letters, 55*, 1530–1533.

Anderson, N. D., & Wilson, H. R. (2005). The nature of synthetic face adaptation. *Vision Research, 45*, 1815–1828.

Ashby, F. G., & Perrin, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review, 95*(1), 124–150.

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review, 94*, 115–147.

Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research, 39*, 2885–2889.

Blumenfeld, B., Preminger, S., Sagi, D., & Tsodyks, M. (2006). Dynamics of memory representations in networks with novelty-facilitated synaptic plasticity. *Neuron, 52*, 383–394.

Bulthoff, H. H., Edelman, S., & Tarr, M. J. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex, 5*(3), 247–260.

Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram. *Annual Reviews Psychology, 55*, 51.

Ericson, C. A., & Desimone, R. (1999). Responses of macaque perirhinal neurons during and after visual stimulus association learning. *Journal of Neuroscience, 19*, 10404–10416.

Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123*(2), 178–200.

Goldston, R., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition, 78*, 27–43.

Griniasty, M., Tsodyks, M. V., & Amit, D. J. (1993). Conversion of temporal correlations between stimuli to spatial correlations between attractors. *Neural Computation, 5*, 1–17.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.

Hopfiled, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America, 79*, 2554–2558.

Kurtz, K. J., & Gentner, D. (1998). Category learning and comparison in the evolution of similarity structure. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, 1236.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience, 4*, 89–94.

Livingston, K. R., Andrews, J. K., & Harnad, S. (1998). Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 732–753.

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*(17), 2287–2297.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature, 335*, 817–820.

Mongillo, G., Amit, D. J., & Brunel, N. (2003). Retrospective and prospective persistent activity induced by Hebbian learning in a recurrent cortical network. *European Journal of Neuroscience, 18*, 2011–2024.

Nadar, K. (2003). Memory traces unbound. *Trends in Neuroscience, 26*, 65–72.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Palo Alto: Morgan Kaufmann.

Phillips, P. J., Moon, H., Rizvi, S. A., & Rauss, P. J. (2000). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 22*, 1090–1104.

Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature, 343*, 263–266.

Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research, 46*, 2977–2987.

Sakai, K., & Miyashita, Y. (1991). Neural organization of the long-term memory of paired associates. *Nature, 364*, 152–155.

Ullman, S. (1996). High-level vision. Cambridge, MA: MIT.

Vogels, R., Biederman, I., Bar, M., & Lorincz, A. (2001). Inferior temporal neurons show greater sensitivity to nonaccidental than to metric shape differences. *Journal of Cognitive Neuroscience, 13*(4), 444–453.

Walker, M. P., Brakefield, T., Hobson, J. A., & Stickgold, R. (2003). Dissociable stages of human memory consolidation and reconsolidation. *Nature, 425*, 616–620.

Wallis, G., & Bulthoff, H. H. (2001). Effects of temporal association on recognition memory. *Proceedings of the National Academy of Sciences of the United States of America, 98*, 4800–4804.

Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in Neurobiology, 51*(2), 167–194.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature, 428*, 557–561.

Author's personal copy