How Looking at Someone You Don't Know Can Help You to Recognize Someone You Do

Nathan Witthoft (witthoft@mit.edu)
Jonathan Winawer (winawer@mit.edu)
Department of Brain and Cognitive Sciences, MIT, 77 Mass Ave
Cambridge, MA 02139 USA

Lera Boroditsky (lera@psych.stanford.edu)
Department of Psychology, Stanford University, 410 Jordan Hall
Stanford, CA 53706 USA

Abstract

Adaptation to faces has been shown to influence judgments of many different features of subsequently viewed faces. For example, after viewing a face that has had its internal features compressed, subjects report that the features of a normal face seem unnaturally expanded (Webster & MacLin, 1999). Recent work has extended these findings to identity, showing that the judgments of the identity of a neutral face can be biased by adaptation to another face (Leopold, D., O'Toole, A., Vetter, T., & Blanz, V., 2001). These results have been interpreted as supporting face-space models of recognition where faces are coded with respect to the prototype. While fascinating, these results require extensive training and depend on participants' subjective reports. We present an objective method for demonstrating that adaptation can affect identity judgments without extensive training in the lab thus supporting the notion of identity based aftereffects. However, the fact that our stimuli are chosen without consideration of a face space suggests that there may be alternative mechanisms underlying face representation and adaptation that do not rely on a prototype.

Introduction

Perceptual adaptation and aftereffects have long interested psychologists both for the clues they provide about what kinds of features (and where) might be coded in the visual pathway and because they demonstrate the ways in which the visual system adjusts itself in response to experience. While many well known aftereffects have been described for 'simple' visual features like color and the orientation of lines more recent work has found that adaptation and aftereffects are found with more 'complex' stimuli such as shapes and faces (Feng & He 2005; Webster & MacLin 1999). Work on face adaptation in particular has shown that adaptation can bias many different kinds of judgments about faces including gender, normality, race, and emotion (Webster et al, 2004; Rhodes et al, 2004; Rhodes et al, 2003). It has also been shown that these adaptation effects are somewhat robust to changes in position, orientation, and size (Yamashita, J., Hardy, J., De Valois, K, & Webster, M., 2005; Zhao & Chubb, 2001; Leopold, D. A., Rhodes, G., Muller, K. M., & Jeffery, L., 2005). These invariances are thought to be typical of neurons higher up in the visual hierarchy and the persistence of adaptation under these conditions is an indication of a relatively high level locus of adaptation.

Most of these studies induce their effects in the laboratory using similar paradigms. For example, Mike Webster (1999) demonstrated adaptation to faces that had been distorted by either compressing or expanding a region between the eyes. Subjects were first shown faces spanning the continuum from expanded to contracted and asked to judge how normal the faces looked. Then subjects adapted to either a contracted or expanded face by staring at it for some time and then were again presented with the continuum of faces and asked to judge their normality. In between judgments of normality subjects again stared at the distorted face to maintain the adapted state (what is often referred to as top-up adaptation). Webster found that when subjects adapted to an expanded face, subsequent judgments were biased in such a way that slightly expanded faces looked normal and undistorted faces looked compressed. The effect is quite powerful and given sufficient adaptation can be seen by an observer on a single trial which is probably why there are few objections to the fact that the results are based on subjective report. One interpretation of this data is that the viewer's representation of what is normal shifts towards recent experience and that judgments of normality are really comparisons to the normal point. So, when adapting to an expanded face, the neutral point begins to move towards that face and subsequent testing with what used to be normal will appear contracted.

Using a similar paradigm, Leopold et al (2001) demonstrated identity specific aftereffects. They created a stimulus set based on the notion of a face space in which faces are coded with respect to a norm, that is the identity of a face is the distance from the center of the space on however many dimensions there are. Having a model of the space allowed them to do a very sophisticated experiment showing that when adapted to a particular face A, subjects were likely to judge neutral faces as having the identity of anti-A (that is the face opposite A in the face space). They further showed that this shift in judgments following adaptation was specific to the trajectory connecting a face

and its anti-face through the neutral point. They took the evidence as suggesting an opponent coding mechanism for faces with the neutral point (or prototype) playing a special role.

One difficulty with the Leopold et al (2001) study is that subjects required extensive training not only on the endpoints but also on faces along the trajectories. That is subjects spent hundreds of trials learning to identify a face that was only 10% of the way towards a particular identity. Subjects might have been trained to represent these faces with a prototype that did not exist before the experiment. In addition they used a subjective method which could be affected by response bias. Here we present an objective method for showing adaptation can selectively improve peoples' ability to recognize famous faces. At no time before or during the experiment are subjects told the names of the people they must recognize, nor are they shown unaltered pictures of the famous faces until the end of the experiment.

Experiment 1

In experiment 1, subjects tried to identify faces taken from morphs between famous faces and two unknown faces.

Stimuli and Methods

Ten grey-scale pictures of faces (8 famous and 2 unknown) were used to construct the stimuli in the experiment. Each image was scaled to the same size and then masked using an oval to eliminate external features like the ears and hairline. Using Morph 2.1 software, four of the famous faces (Bill Clinton, Tom Cruise, Christian Slater, and Sylvester Stallone) were morphed with an unknown face (hereafter face A). The other four famous faces (Ben Affleck, Bill Cosby, Dustin Hoffman, and Elvis Presley) were morphed with a different unknown face (hereafter face B). Each morph line (the set of faces including the endpoints and the graded transformations) was created by hand positioning corresponding points and lines on each pair of endpoint faces. The software then generated 58 intervening images between the endpoints. Counting the unknown face as 1, the 24th, 30th, 36th, and 40th faces were chosen from each morph for use in the experiment. For each morph, these are referred to as the 40, 50, 60 and 67% famous faces respectively. The unknown faces and the famous faces are considered 0% and 100% famous. Examples of the stimuli can be seen in figure 1.

Stimulus presentation and response collection was done using VisionShell software on Macintosh iMac computers. All subjects were recruited at Stanford and tested in quiet darkened rooms and received either 5 dollars payment or credit in an introductory psychology course.
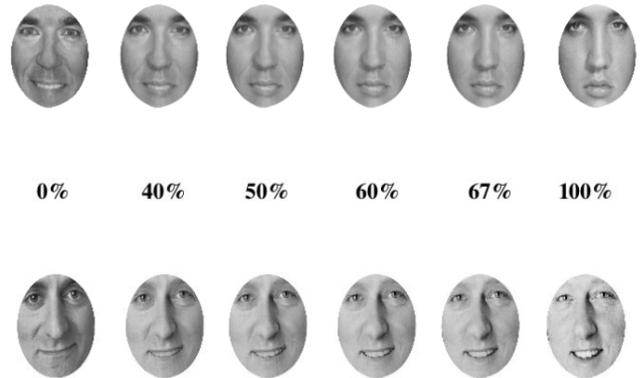

Figure 1: Examples of faces taken from two morphs and used in the experiment. The top row shows faces from Face B to Elvis Presley while the bottom row shows faces in morph from Face A to Bill Clinton.

40 subjects participated in the experiment. 10 subjects adapted to face A during the experiment, another 10 adapted to face B, and the final 20 did not undergo any adaptation. The test stimuli for all the subjects were identical. Each version of the experiment had 6 blocks with the percent famous face increasing across the blocks. In the first block, subjects were told that they would be shown a series of faces each of which resembled a famous person, and that they should try to guess who it was. The subjects were then shown the 8 40% famous faces in random order and asked to guess what famous person they resembled if they could. At no time before or during the experiment were subjects told the identity of the famous faces used to make the morphs. In most cases subjects could not identify the 40% morph.

For the subjects in the adapting conditions, each of the next 4 blocks began with 55 seconds of adaptation to one of the two unknown faces. They were then presented with one of the 40% faces for 1 second and asked to identify it. This was followed by 5 seconds of top-up adaptation and then another 40% face until all the 40% faces had been shown. After the block of 40% faces, subjects did similar adaptation blocks with the 50%, then the 60% and then the 67% faces. In the last block, subjects were shown the 100% famous faces without further adaptation and asked to identify them. The 100% faces remained onscreen until subjects made a response. Subjects in the baseline (no adapt) condition had the identical sequence of test stimuli but no intervening adapting faces.

Our prediction is that subjects should perform better at recognizing famous faces that have been morphed with the stimulus they are adapting to. In other words, less of a famous face needs to be present if it is part of a morph with the adapting stimulus. Since subjects do not know the identities of the famous faces used in the experiment, any information they have must be gotten from the stimuli rather than changes in response strategy (bias). Furthermore, since the test stimuli are the same for all subjects, any differences must be attributable to the effect of adaptation.

Analysis and Results

For each subject, faces were excluded from the analysis either if they were recognized at the first exposure of 40% prior to any adaptation or if they were not recognized in the final block at 100%. Data from each of the subjects who were adapted was divided into two groups dependent on whether the test face had been morphed with the unknown face the subject was adapting to, or the one the subject never saw. Test faces from morphs using the adapting stimulus as an endpoint are referred to as 'same', and those using the other face were labeled different. So if a subject adapted to face A, the morphs made with face A were considered 'same' and the morphs made with face B 'different'. The situation is reversed for those subjects adapting to face B. Within each of these groups, the number of correctly recognized faces from a level of morph was divided by the number of faces the subject recognized at the 100% level. This yielded a normalized percent correct score for each level of morph.
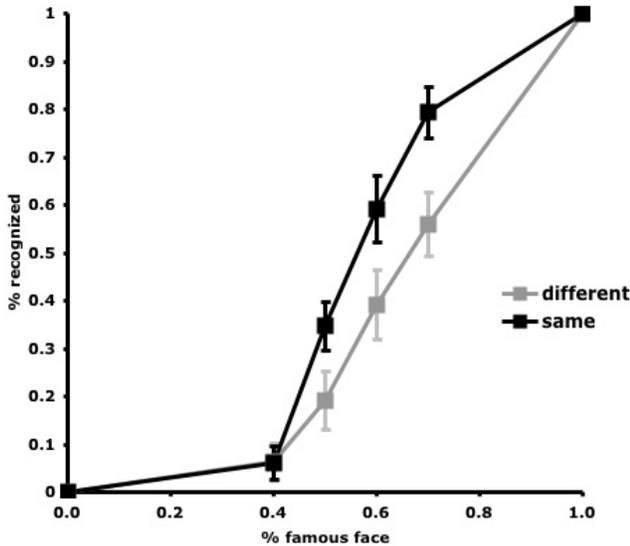

Figure 2. Recognition performance of subjects in experiment 1. Error bars represent standard error of the subject means. Note that stimuli were presented in blocks with each block increasing the percentage of famous face. So the data for the 40% faces on the x axis come from block 2, the 50% from block 3 and so on.

Overall, subjects did fairly well at recognizing the faces, on average correctly identifying 72% at some point during the experiment (mean number recognized = 5.75 faces out of 8, SEM = 0.32). All subjects who did not leave the room during the experiment are included in the analysis, and every subject recognized at least 3 faces. Figure 2 shows the average data from the 20 subjects who adapted to one of the two unknown faces in experiment 1. The graph shows that faces that were morphed with the adapting stimulus (same) were better recognized than those morphed with the unknown face not adapted to (different). Another way of

putting it is that less of the famous face is required to be present for successful recognition in the same condition relative to the different condition.

To quantify the result a logistic function was fit to the each subject's data (Palmer, Huk, & Shadlen, 2005). The logistic had two parameters, one corresponding to the slope, and one corresponding to the center of the curve (the "% celebrity" needed for 50% recognition). Betas representing the center of the curve for each subject were determined independently for the same, different, and baseline conditions. Curves that are (centered) shifted farther to the left represent better recognition performance (see figure 3). These parameter estimates were then entered into a 1x3 (adapting conditions: same, different, baseline) Anova. The results showed a main effect of adapting condition ($F_{(57,2)} = 4.687$, $p < 0.05$). Post hoc t-tests conducted using Tukey's Least Significant Difference procedure showed that a significant shift leftward for the same adaptation relative to the different condition ($p = 0.003$). The baseline condition was intermediate between the two adapting conditions but did not significantly differ from either ($p = 0.1$ compared to same and $p = 0.14$ compared to different).
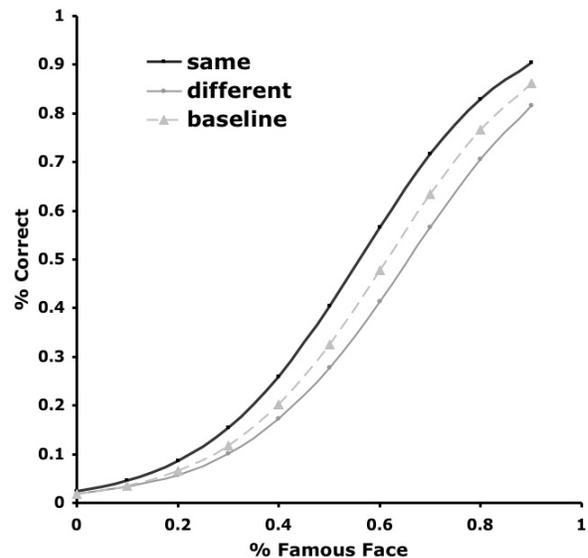

Figure 3: Fitted logistic functions for subjects in experiment 1. Recognition performance is significantly better for faces morphed with the adapting stimulus (black line) than for those not morphed with the adapting stimulus (solid grey line). Performance in the baseline condition (dashed grey line) where subjects did not undergo any adaptation is intermediate.

While the preceding finding demonstrates that adaptation can improve recognition performance, the data can also be analyzed by item to see whether or not the effect occurs for all the faces. To analyze the items two logistic functions were fitted simultaneously the same and different performance on each item and a set of betas representing the difference in horizontal shift between the two curves was

generated (1 beta for each item). These betas were t-tested against zero and showed that better recognition was seen when items were morphed with the adapting stimulus ($t(7)$ = 3.7 , 0.0082, mean shift 13% famous face). The fitted curves are depicted in figure 4. This way of looking at the data highlights the fact that the improvement in recognition performance cannot be due to some items being easier to recognize than others. The item analysis shows that recognition of a particular face is modulated in the predicted fashion by adaptation.
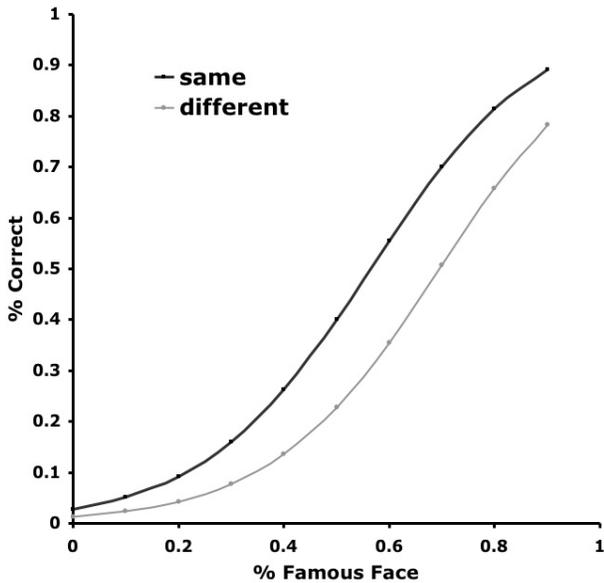

Figure 4. Logistic functions fitted to item data from experiment 1. Faces were recognized better when morphed with the adapting stimulus (same – dark line) than when morphed with the stimulus not seen (different – lighter line). Organizing the data in this way highlights the fact that performance for a particular face is dependent on the adapting stimulus.

Experiment 2

Experiment 1 showed that adaptation to an unknown face decreased the amount of a famous face needed in an image in order to recognize it. This effect was specific to morphs that used the same unknown face as an endpoint as an adapting stimulus. The effect of adaptation depended on the adapting stimulus thus ruling out the possibility that the improved recognition was simply due to some of the faces being more recognizable than others. The effect was also present in an item analysis, with 7 of the 8 items showing recognition performance consistent with the adaptation effect. To see how general this effect is, experiment 2 replicates experiment 1 using 10 new faces (8 famous and 2 unknown).

Methods

The stimuli and design of experiment 2 were exactly like experiment 1, except that 8 new famous faces and 2 new unknown faces were used (hereafter referred to as Faces C and D). To increase the likelihood that subjects would be able to recognize the 100% famous faces a separate group of subjects was asked to identify a number of famous faces that had been scaled and masked in the same fashion as they would appear in the experiment. Arnold Schwarzenegger, Elijah Wood, Nicholas Cage, Matt Damon, John Travolta, Harrison Ford, Brad Pitt, and Johnny Depp were recognized the most often were chosen and used in the experiment. One other minor difference was that instead of the 67% morph the 70% morph was used.

32 subjects recruited from the Stanford undergraduate population participated in return for course credit. 10 subjects adapted to face C during the experiment, 10 adapted to face D, and the remaining 12 did not adapt. As in experiment 1 subjects were shown the test faces in blocks with the percentage of famous face in each image increasing across the blocks.

Analysis and Results

Analysis proceeded exactly as in experiment 1 with faces recognized on the first presentation (40%) prior to adaptation and those faces not recognized at 100% discarded from the analysis. Overall recognition performance in the adapting conditions was again good, with subjects recognizing 74% of the faces on average (mean number recognized = 5.95, SEM = 0.36).
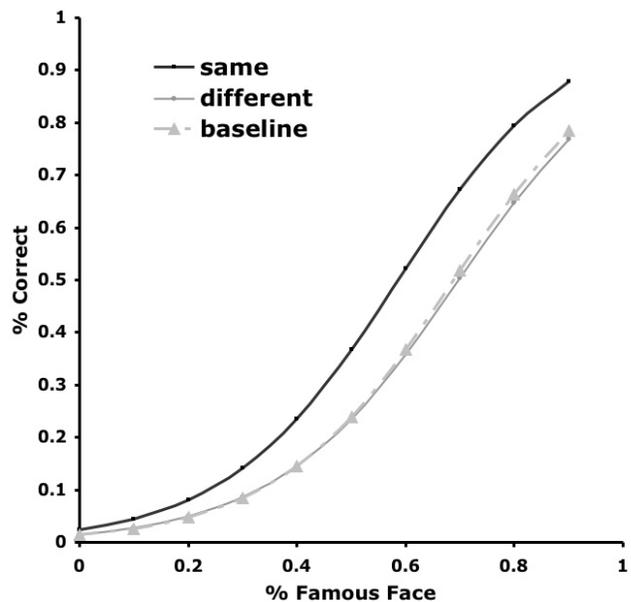

Figure 5. Logistic functions fitted to data from experiment 2. As in experiment 1 subjects recognition performance was better when they guessed the identities of faces from a morph that were generated using the adapting stimulus as an endpoint.. However, in this experiment, performance in the baseline condition is indistinguishable from the different condition.

The results were similar to those of experiment 1. Logistic functions were separately fitted to data taken from the same, different, and baseline conditions. Estimates of the parameters measuring the horizontal shift of the curve were entered into a 1x3 Anova. Results showed a significant main effect of adaptation condition ($F(49,2) = 5.945$, $p = 0.005$). Post-Hoc contrasts showed that again stimuli in the same condition were better recognized than those in the different ($p = 0.002$). However, in this case, same performance was also better than baseline ($p=0.036$) while there was no difference between baseline and different ($p>.4$). This supports the idea that adaptation is boosting performance on the 'same' faces rather than disrupting performance on the 'different' faces.
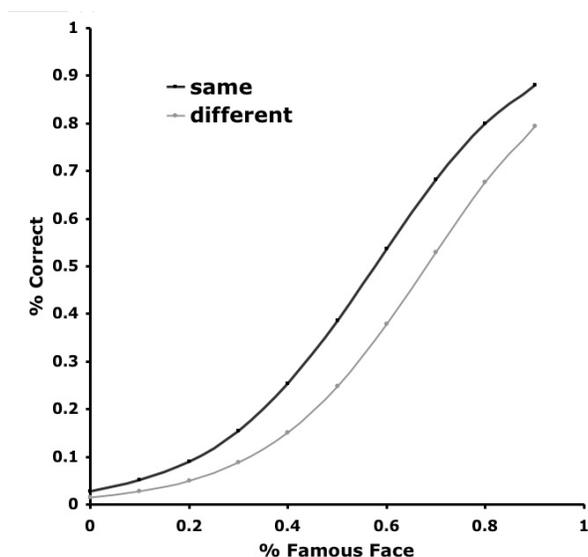

Figure 6: Results from item analysis of experiment 2.

As in experiment 1, data were also grouped by item and analyzed. Logistic functions were fit to each item and a parameter representing the shift between the same and different curves obtained for each item. Testing the collection of parameters representing the horizontal shift against zero showed again that faces were better recognized following adaptation to the unknown face they were morphed with ($t(7) = 2.41$, $p = 0.0468$ mean shift ~10% famous face). Figure 6 shows the curves fit to the aggregate data.

General Discussion

Here we have presented results showing that adaptation can selectively enhance the ability to recognize faces. Since subjects did not know which faces out of all the ones they know they would need to recognize, the results can not be attributed to response bias and must be attributed to a change in the way subjects were processing the test faces. Which morphs showed improvement was dependent on the adapting stimulus and not the particular faces to be recognized. This supports earlier work showing that there are identity specific aftereffects while eliminating concerns that recognition performance is really due to changes in response strategy or a result of training during the experiment.

The results of the experiments are very robust considering that a subject has only 4 stimuli in each condition and that only 10-12 subjects are used for each group. Furthermore, it is worth considering that the faces were chosen without taking into consideration where they might be relative to one another in a face space. This suggests that perhaps models which take face adaptation to reveal the primacy of the prototype for explaining face recognition and aftereffects may be missing other mechanisms. This is particularly true since the advantage for adapting in the 'same' condition can be seen in most of the items. Not all of the trajectories from the unknown faces to the famous faces are likely to pass through the center of the face space. Given that the origin or prototype is usually taken to be the average of the faces a person has seen (Rhodes et al 2003, Leopold et al 2001), having chosen a set that matches this would be a striking coincidence. Furthermore, it is not clear from such accounts whether or not there are multiple prototypes (i.e. separate ones for gender or race) or whether all faces are coded as deviations from the global average.

An alternative possibility is that adaptation causes changes such that differences from the adapting stimulus are emphasized. Note that such an account is consistent with the results found in the Leopold (2001) study. Testing across the center of the face space is just one trajectory among many that could have been made. Further supporting the idea is the fact that Leopold et al found improvement on face recognition following adaptation to the prototype in their face space. If aftereffects were merely the result of shifting the prototype towards recent experience adapting to the prototype should have no effect. However, if adaptation (at least in part) results from emphasizing change from recent experience than this is exactly the predicted effect.

Finally, our adaptation effects were obtained without any training on the to-be-recognized faces. This eliminates the possibility that the desired structure of the categories is learned during the training (though this would be an interesting effect in its own right) and applies just to the stimulus set used in the experiment. That is, it is possible that training people extensively to classify the stimuli may induce a prototype representation for those stimuli. While this may seem unlikely, Leopold has shown that the same effects using human faces as stimuli can be induced in monkeys following training (Leopold & Bondar, 2005). Since the prototype is meant to reflect the average of a person's experiences it is curious that monkeys would share that prototype and is suggestive that they learned the desired structure during training.

Acknowledgments

References

Fang, F., & He, S. (2005). Viewer-centered object representation in the human visual system revealed by viewpoint aftereffects. *Neuron*, 45(5), 793-800.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89-94.

Leopold, D. A., Rhodes, G., Muller, K. M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. Proceedings. Biological Sciences/The Royal Society, 272(1566), 897-904.

Leopold, D. A. & I. Bondar. (2005) Adaptation to complex visual patterns in humans and monkeys. in Clifford, C. W. G., & Rhodes, G. (Eds). *Fitting the mind to the world : adaptation and after-effects in high-level vision* (1st ed.). Oxford ; New York: Oxford University Press.

Palmer, J., Huk, A., & Shadlen, M. (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. Journal of Vision 5(5), 376-404.

Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W., & Nakayama, K. (2003). Fitting the mind to the world: face adaptation and attractiveness aftereffects. *Psychological Science*, 14(6), 558-566.

Rhodes, G., Jeffery, L., Watson, T. L., Jaquet, E., Winkler, C., & Clifford, C. W. (2004). Orientation-contingent face aftereffects and implications for face-coding mechanisms. *Current Biology*, 14(23), 2119-2123.

Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review*, 6(4), 647-653.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557-561.

Yamashita, J. A., Hardy, J. L., De Valois, K. K., & Webster, M. A. (2005). Stimulus selectivity of figural aftereffects for faces. *Journal of Experimental Psychology: Human Perception & Performance*, 31(3), 420-437.

Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, 41(23), 2979-2994.