Perception, 2010, volume 39, pages 1230 – 1239

doi:10.1068/p6658

Rapid object category adaptation during unlabelled classification

David Hadas, Nathan Intrator, Galit Yovel#
School of Computer Science (# Department of Psychology), Tel-Aviv University, Ramat Aviv, 69978 Tel Aviv, Israel; e-mail: david.hadas@gmail.com
Received 30 January 2010, in revised form 23 June 2010; published online 1 September 2010

Abstract. Recent reports from electrophysiological and psychophysical experiments provide evidence that repeated exposure to an ordered sequence of morphed stimuli may over time adapt a prelearned object category such that the category may generalise the entire sequence as belonging to the same object. Here, a new protocol that includes a single exposure to a morphing sequence is presented. Subjects exposed to the new protocol replaced a prelearned face with an entirely different face within just 3 days, significantly faster than in previous reports.

1 Introduction

Categorical perception is said to exist where stimuli from a continuum are labelled as belonging to different classes, and where stimuli labelled differently are well discriminated, whereas others, labelled as the same, are less-well discriminated (Ehret 1987). Categorical perception therefore requires that there be a sharp perceptual boundary between stimuli from two different categories. Monitoring changes in the perceptual boundary between two categories sheds light on the internal representation of the category (Harnad 1987). A perceptual boundary appears between general categories (eg faces, chairs, cars, etc), but also between specific examples within a general category (eg face stimuli from Jerry, George, and Elaine). Thus all stimuli sharing the same label (eg Jerry) are perceived as the same object category (Beale and Keil 1995; Newell and Bülthoff 2002). Although in most cases categorical perception results in a stable perception of the observed stimuli, there are known unstable cases where perception fluctuates between multiple categories (see, for example, Leopold et al 2002).

Recent results from electrophysiological and psychophysical experiments provide evidence that exposure to an ordered sequence of morphed stimuli between two prelearned object categories adapts the perceptual boundary between the categories and suggests that exposure to a morphing sequence may expand an object category (Leutgeb et al 2005; Blumenfeld et al 2006; Preminger et al 2007). Leutgeb et al reported that hippocampal place cells are capable of incremental plastic deformation. These researchers trained rats inside a square or circular enclosure in a random sequence. The rats' task was to search for food scattered around the enclosure. Pretraining continued for 16 – 19 days, until the two enclosure conditions reliably activated different subsets of neurons in the associative CA3 network. Cell ensemble activity was then recorded while the square box was morphed into a circle, and vice versa, through a sequence of five intermediate shapes. The researchers showed that during exposure to a gradually morphed enclosure the cell ensemble firing patterns of intermediate shapes depended on the direction of the morphing (stimulus order). The researchers then repeated the sequence in alternating directions on subsequent days and reported that the difference between the firing patterns of the previously established end-points had attenuated. In another hippocampal place cells recording experiment, rats were pretrained on similar stimuli but later exposed to a scrambled order of intermediate shapes. While exposure to the gradually morphed stimuli resulted in attenuated differences in firing patterns, exposure to the scrambled order, instead, resulted in an abrupt switch between firing patterns of square-like and circle-like intermediate shapes (Wills et al 2005).

On the basis of these results, Blumenfeld et al (2006) showed that attractor networks (Hopfield 1982) may also display a stimulus order effect. The researchers modulated the Hebbian learning mechanism (Hebb 1949) of the attractor network using weights determined by the perceived novelty of the stimuli. Using such novelty-facilitated modulation, they showed that, when morph patterns are presented in a gradually increasing order, the network exhibits a stronger tendency for merged representations compared to random order. The model, therefore, predicts that when network inputs are correlated, as in the case of a morphing sequence, different prelearned representations may expand and collapse into a single unified representation. A similar stimulus order effect was also observed during psychophysical experiments. Wallis and Bülthoff (2001) conducted an experiment in which subjects were exposed to rotated heads, artificially constructed through morphing from a face in a frontal view and another in a profile view. Using discrimination tests, they showed that the subjects associated the two faces as belonging to the same object. When subjects were exposed to the morphed head images in a randomised order, no association was discovered. The authors suggested that spatiotemporal correlation was required to associate the face representations. Later, Wallis (2002) showed that the association established is modulated by the spatial correlation, and Wallis (2009) showed associations achieved using spatiotemporal correlation of changing illumination and orientation.

In a recent study, Preminger et al (2007) have shown that association can be obtained with spatial correlation, even when temporal correlation is avoided. In their study, subjects initially learned to classify a group of facial images as 'friends' and others as 'nonfriends'. A sequence of morphed images was generated between a 'friend' and a 'nonfriend'. At the onset, the subjects associated about half of the morphed sequence as a 'friend'. In other words, the object category boundary between a 'friend' and a 'nonfriend' was located around the 50%–50% morphed image. In 10 of 18 cases, exposure to the ordered sequence repeatedly, over many days, resulted in a gradual change in which more of the intermediate morphed images were classified as a 'friend'. Finally, after repetitive sessions, these subjects identified the two morphed faces as 'friends'. The authors called this change in the object boundary a 'morph effect'. In the remaining 8 of the 18 subjects a morph effect was not detected. The researchers indicated that the subjects who did not experience the morph effect had initially been more discriminative in identifying the morphed images as a 'nonfriend'. A control group exposed to a random unordered presentation of the same morphed images did not produce a morph effect. Note that, when preparing the morphing sequence, the researchers ensured that there were no conspicuous features to cause trivial discrimination between faces. The constructed faces were chosen to be neither too similar nor too different from each other.

Taken together, the above reports suggest that unsupervised classification of spatially correlated images may add information to a representation. As a result, the category may gradually drift from its prelearned position in the stimulus continuum and expand. However, all reports correlated between similar objects, ie objects that are not easily distinguishable. In Leutgeb et al (2005), the rat subjects required intensive training to learn to distinguish between the two objects that were later correlated. In Wallis and Bülthoff (2001) and in Wallis (2002) correlation was achieved in human subjects presented with the challenging task of differentiating between frontal and profile face views, and therefore can be easily confused as similarly looking faces. In Preminger et al (2007) the correlated face images were intentionally chosen to be not too different. Note that even with such choice of face images, and following repeated exposure to the morph sequence, only half of the subjects experienced morph effect and expanded the category.

In the current study, a novel protocol which arguably intensifies the morph effect is presented. A morphing sequence between significantly different face images was created to test the protocol. Using the new protocol, category drift was achieved following a single presentation of the morphing sequence, and significantly faster than in previous reports. Former protocols, such as that of Preminger et al (2007), generated drifts of a category both away from and back towards the prelearned image owing to the following reasons. First, since each session starts the morphed sequence from the prelearned object, some of the change in the perceptual boundary already achieved in the previous session is canceled, and, second, since every session presents the entire sequence of morphed images, subjects cross the categorisation boundary between the representation of the prelearned image and an alternative representation on each session. Once the categorisation boundary is crossed, the visual system learns to discriminate between images of the morphed sequence rather than to generalise all images in the morphed sequence as belonging to the same representation.

In the experimental protocol presented here the drift of a category in the stimulus continuum is ensured to always accumulate away from the prelearned object by: (i) starting each new session with the intermediate morphed images that subjects classified as 'targets' in the previous session, rather than going back to the first image in the morphed sequence; (ii) presenting no more than one-third of the sequence of morphed images in every session to avoid crossing the category boundary. Note that earlier protocols repetitively presented the morph sequence, leading subjects to generalise all morphed views. The novel experimental protocol was designed such that subjects are not forced to generalise the entire sequence. This is achieved by presenting the stimulus change more gradually and only once. Using this novel experimental protocol, in just 3 days of exposure to a 10-min session, subjects accumulated substantial drift and eventually classified a completely different image as the target.

2 Experimental procedures

2.1 *Stimuli*

Faces were taken from a 100-face Nottingham scans image set (downloaded from http://pics.psych.stir.ac.uk/index.html). Eight faces were hand-picked from the database providing easy identification between them. Each of the eight faces was prepared by blacking out the background, hair, and ears. Two of the faces were named A and B and were used for the preparation of a sequence of morphed images. The constructed sequence included 99 morphed images (see figure 1). The sequence was prepared with Sqirlz Morph version 1.2e. A separate pairwise similarity-rating test confirmed that A and B are perceived as different from each other by naive observers. Observers evaluated 36 pairs of images, among which the A and B pair appeared three times. On a scale in which 1 indicates that the images are exactly the same and 5 indicates that the images are completely different, the A and B image pair received a mean rating of 4.09 (SD 1.04, $n = 11$).

The remaining six images were picked as alternative images for the construction of the distracting face images. Ten morphed images were prepared from the first alternative image (inclusive) to the second (exclusive), ten more from the second image (inclusive) to the third (exclusive), etc. The resulting fifty alternative images were picked randomly whenever an image from the distracting face images was required.

2.2 *Apparatus*

Images were presented on a 14-inch LCD of a laptop with 1024×768 pixel resolution (refresh rate 60 Hz) in a small, quiet room containing fluorescent lighting and no windows. Presented images were approximately 250×300 pixel Jpeg images on a black background. A fixation point was shown when an image was not presented.

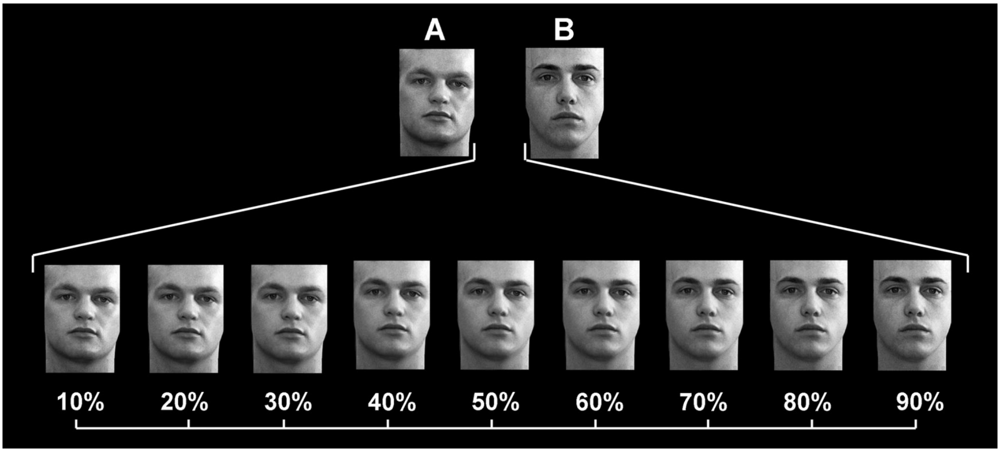

Figure 1. Images A and B and the sequence of morphed images prepared. The experiment presented 99 intermediate images with morphing ratio starting from 1 : 99 and ending at 99 : 1 between images B and A, respectively. Here 10% intervals in the morphing sequence are shown.

2.3 *Subjects*

Nine undergraduate students volunteered to participate in the experiment, to fulfill their introduction to psychology course requirements. They were notified that they would be required to attend 30-min daily sessions on 4 different days. Students filled in a short questionnaire. They were screened to ensure that none of them was taking medications that might affect their memory or attention. They were told that they would be required to learn a face and would be tested on how well they remembered that face during the experiment.

The experiment was performed in accordance with the guidelines and regulations of the Ethics Committee of the Department of Psychology at Tel-Aviv University. Subjects' consent was obtained prior to participation in the experiment.

2.4 *Procedure*

The experimental procedure is summarised in figure 2a. During all phases, subjects performed a target identification task. Subjects were required to press the space bar in order to view a face and in this way controlled the presentation rate. Then they were required to indicate whether they saw the prelearned face by pressing 1 or an alternative face by pressing 2. The stimulus presentation lasted until the subject completed classification or 1000 ms poststimulus onset, whichever occurred first. Instructions were verbally provided by the instructor, on paper, and were also written at the bottom of the screen. During the target-learning phase, the upper screen was used to provide feedback to the subject. No feedback was provided in later phases.

2.4.1 *Preliminary testing.* Preliminary testing to obtain an adequate set of parameters for the protocol phase was employed. The parameters studied included the number of sessions, the interval used between sessions, the number of trials per session, the interval between images of the morphing sequence, and the exposure time. One lesson learned was that subjects became confused if the morphing sequence progressed too quickly. For example, using more than 33% of the morphing sequence per day appeared to confuse many subjects. To avoid such confusion a presentation tool was programmed to limit the progress made per day to a range of 33%.

2.4.2 *Target-learning phase.* In the target-learning phase, subjects learned image A as their target and were asked to remember it. Image A was presented for 3 s during 3 iterations. Subjects were asked to press key 1 on the keyboard after each presentation.

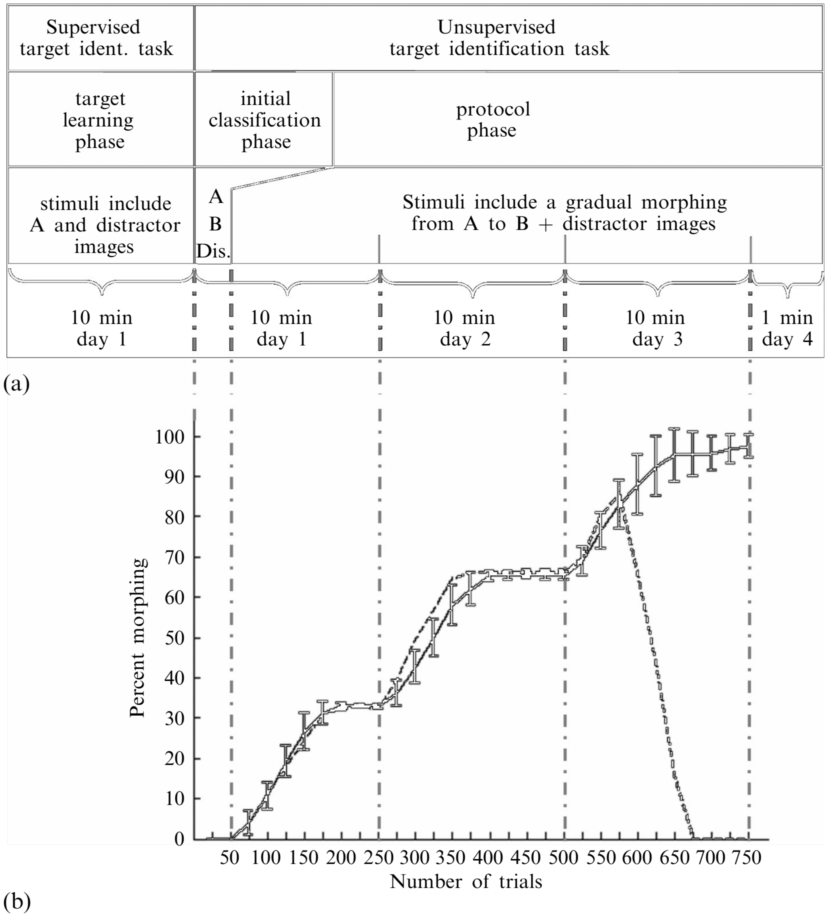

Figure 2. (a) Experiment layout. On day 1, subjects were trained to identify image A among non-target distractor face images. Once successful, subjects performed an initial classification phase in which they discriminated image A from B and distractor faces. Then, unbeknown to the subject, the protocol phase began in which the A-to-B morphing sequence and distractors were presented, during which time, image A gradually transformed to image B in 1% morphing steps. During each day, subjects were limited to a maximum progression of 33% morphing. On the 4th day, subjects completed the classification of the entire morphing sequence. (b) Subject classification performance along the morphing sequence. Classifying a morphed image as belonging to the prelearned target face (face image A) was followed by a positive 1% progress in the morphing sequence. Classifying a morphed image as not belonging to the target was followed by a −3% change in the sequence. Apart from irregular lapses, eight subjects quickly progressed each day until the enforced daily limit of 33% progress (presented as a continuous line). The limit is expressed as a plateau at the end of each day. On average, subjects progressed during the 3 days from 0% morphing to 97.25% (SD 3.4%). All eight subjects reached a morphing of 100% in the 4th day of the experiment (not shown). One subject (shown with a dotted line) performed equally well as other subjects until reaching a critical stage following which she had stopped referring to images from the morphing sequence as the target and her experiment was discontinued.

Next, subjects started a target identification task with feedback using 20 trials. The images were chosen randomly with a 1 : 1 chance between image A and the distracting face images. The subject received feedback after classifying each image. At the end of this phase, the results were summarised and presented. A subject who performed this phase perfectly moved to the next phase. Eight of the nine subjects performed this phase perfectly while one subject made a single mistake and therefore repeated this phase.

2.4.3 *Initial classification phase.* Several minutes after the target-learning phase, subjects started the initial classification phase. In this phase the subjects performed 50 trials of a target identification task without feedback. The images were chosen randomly with a probability of 1 : 2 between image A and an alternative image. The alternative image was chosen with a probability of 1 : 4 between image B and the distracting face images. Note that image B was not presented prior to the initial classification phase, and the test did not provide an indication to allow subjects to single out image B from the distracting face images.

2.4.4 *Protocol phase.* A protocol phase immediately followed the initial classification phase, without a break. Subjects were unaware of the transition between the tests and continued the target identification task without feedback as before. The presentation schedule during the protocol phase was chosen using the presentation tool and included a random 1 : 2 chance between images from the morph sequence and the distracting face images. Images A and B were not presented during this phase. The presentation tool was programmed to progress to the next image of the morphing sequence following a positive classification by a subject and to go back three images of the morphing sequence with every negative classification by a subject (staircase method). The protocol phase started on the 1st day and ended on the 4th day when a subject completed the morphing sequence (reached 100% morphing). In each daily session the experiment ended when the subject completed a continuous target identification task of 250 images (taking about 10 min). The first morphed image presented on each day was based on the last morphed image presented on the previous day. The presentation tool ensured that all morphed images presented during the session did not exceed a progression of 33% morphing, such that a minimum of four sessions was required to complete the morphing sequence.

2.4.5 *Debriefing.* After each session, the subjects were asked to describe their experiences and indicate whether they had faced any difficulties during the task. Subject responses were documented.

3 Results

The results of the 4 experiment days are reported here. Nine subjects, unaware of the purpose of the experiment, learned to identify facial image A as their target and then undertook a target identification task by classifying facial images as the same or different from the prelearned target. On average, when excluding the top and bottom 1% tail, reaction time from stimulus onset was 803 ms (SD 232 ms). The average resulting exposure time was 721 ms (SD 111 ms). In just 3 days, during which a morphing sequence from face image A to a very different face image B was presented, eight of the nine subjects reverted to classifying image B as the prelearned target. This drastic change in representation occurred with just 210 ± 30 exposures to images from the morph sequence. All nine subjects showed similar performance during the 1st and 2nd days of the experiment. On the 3rd day, one subject drastically changed the classification patterns and reverted to classifying images from the morphing sequence as not belonging to the prelearned target. Further, when presented with facial image A, the subject classified 94% of facial image A trials as not belonging to the prelearned target. At this point the experiment was terminated.

During the protocol phase, eight subjects classified 98.7% (SD 1.35%) of the morphed faces as the prelearned face. Accordingly, the staircase method pushed these subjects along the morphing sequence away from image A and towards image B. Figure 2b shows the average progress made by the eight subjects over the first 3 days. Note that all subjects made definite progress and accumulated considerable drift from image A. Apart from irregular lapses, drift accumulated steadily each day until the enforced daily limit of 33% progress. The limit is expressed as a plateau at the end of each day. On average, the eight subjects progressed during the 3 days from 0% morphing to 97.25% (SD 3.4%). The protocol phase ended early on the 4th day when all subjects completed 100% of morphing. The figure also shows the one subject (shown with dotted line) who drastically changed the classification patterns and reverted to classifying one of the distracting face images as the target. Following the critical stage, this subject classified morphed face images as not belonging to the target and, at the same time, classified 93.75% of the images presenting a particular distracting face image as belonging to the target (compared to 0% prior to the critical stage).

Subject classification patterns were affected by the morphing sequence progress as can be seen from figure 3 (averaged over the eight subjects). A correlation of 0.87 was found between the spatial difference increase along the morph sequence and the percentage of morphed images classified as nontargets. The average percentage of morphed images classified as nontargets during the protocol phase was 3% with SD of 3% across the experiment days. The highest percentage of nontarget classifications was recorded during the last 10% of the sequence with 9% (SD 11%) nontargets. In comparison, the average error rate of distracting face images during the protocol phase was relatively constant across the 3 training days (mean: 2.3; SD 0.2%).

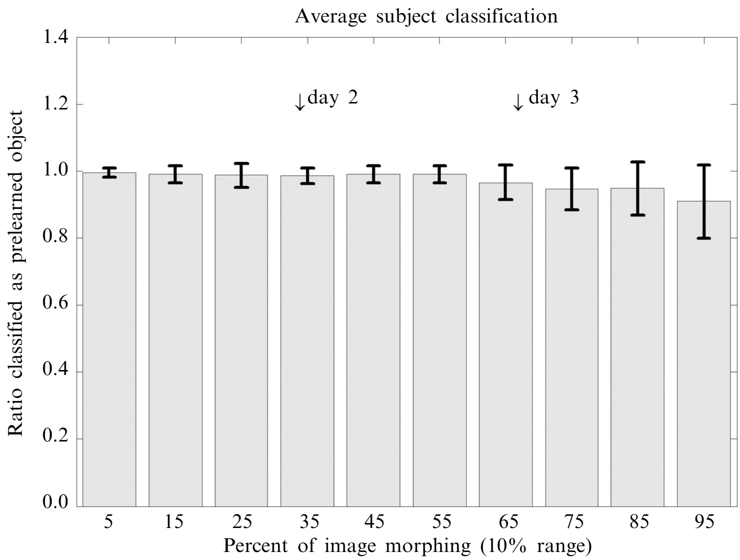

Figure 3. The ratio of images from the morphing sequence classified as the target. The ratio is presented for each 10% morphing range. Strikingly, subjects classified the 90% – 99% morphing range (which includes only 1% – 10% of the prelearned image A), as nontarget only 9% (SD 11%) of the trials.

During the initial classification phase, all nine subjects classified 100% of image A trials as the prelearned face. Image B and the distracting face images were classified as the target only 0.7% (SD 2%) and 2% (SD 2%) of the times respectively. These data suggest that all subjects clearly distinguished image B as not belonging to the prelearned target during that phase.

By day 3 of the experiment, six of the subjects repeatedly classified the last image in the sequence, representing 99% morphing as the prelearned face (100% of the time), one subject classified the image representing 97% morphing as the prelearned face 18 of 22 times, and another subject was indecisive regarding the classification of images above 90% morphing, with just 68% of trials classified as the prelearned face. As day 4 started, all eight subjects classified all remaining images up to 100% morphing as belonging to the target and thus ended the protocol phase.

4 Subjects' experiences during protocol phase

Overall, upon debriefing, the eight subjects did not report confusion and indicated that they were confident in their classifications. Reports were made that the sessions became harder each day. The subject who referred to one of the distracting face images during the 3rd day as the target did not report any difficulty during the first 2 days of the experiment. At the end of the 3rd day, the subject had reported that during the sessions, she had suddenly realised that she is classifying incorrectly and that the image she had learned 2 days earlier was different from the image she was targeting. Following this discovery she had turned to classify the other image as the target.

5 Discussion

Our findings are consistent with earlier reports (Wallis and Bülthoff 2001; Leutgeb et al 2005; Preminger et al 2007), which demonstrated that the visual system may adapt object categories without an external supervisor. All reports collectively indicate that exposure to a morphing sequence between two object categories may result in all sequence images being classified the same. However, contrary to previous reports, it is shown here that the said adaptation is rapid, is found in almost all participants, and can be induced even between very different object categories.

In the current study, subjects learned to identify face image A as their target and discriminate it from other face images including face image B. The representation of the target category therefore included information that allowed them to differentiate A from B. Yet, during the exposure to the morphing sequence, subjects continued to classify the entire sequence of morphed images as their target, even when the morph included less information from the prelearned image A and more information from a clearly different image B. This evidence is consistent with previous findings (Wallis and Bülthoff 2001; Wallis 2002; Leutgeb et al 2005; Preminger et al 2007) and suggests that exposure to images spatially correlated to a prelearned object may change the object category representation. In contrast to previous reports, it is demonstrated here that this process does not require that the morphed sequence be repeated over many sessions. Instead, it is apparent that a single presentation of the ordered morphed sequence suffices to achieve a morph effect.

The findings presented here differ from those of Preminger et al (2007), although in both studies subjects were exposed to a protocol that included a presentation of morphing sequence during an unsupervised classification task. Here, all subjects exposed to the protocol experienced a morph effect and eight of nine completed the sequence within 3 days, whereas in Preminger et al (2007) only 55% were reported to experience a morph effect and completing the sequence required more exposures and more time. The difference in the results may be explained by the different protocols used during the presentation of the morphing sequence. The protocol used here included two significant modifications to the protocol used by Preminger et al. (i) The protocol used by Preminger et al starts from the beginning of the sequence in each daily session. The protocol used here starts each new session with the image that subjects classified as a 'target' in the previous session to avoid cancellation of the drift achieved so far. (ii) In each daily session the protocol used by Preminger et al stops at the end of the sequence after crossing the categorisation boundary. Once the categorisation boundary is crossed, the visual system is trained to discriminate between the morphed images. The protocol used here avoids crossing the category boundary by presenting no more than 33% of the morphing sequence per session. The intent behind the two modifications was to allow drift of a category in the stimulus continuum to accumulate away from the prelearned object and avoid the accumulation of drift in the reverse direction toward the prelearned object.

Two major differences between the results presented here and those of Preminger et al (2007) suggest that the procedures used here intensified the morph effect. First, as mentioned above, all the subjects exposed to the novel protocol experienced a morph effect, whereas only 55% of the subjects exposed to the protocol of Preminger et al did so. Second, the progress made by subjects exposed to the novel protocol appears to be significantly faster than progress made by subjects exposed to the Preminger et al protocol (3 days and 235 exposures versus about 4 – 15 days and 400 – 1500 exposures). The new results add to the results of Preminger et al (2007) as they may suggest that morph effect is a robust phenomenon in the visual system rather than occurring in only about one out of two subjects, as may be suggested by the previous findings.

Another important difference between the experiment described here and the one described by Preminger et al (2007) arises when considering the preparation of the morphing sequence. In Preminger et al (2007), the sequence preparation ensured that there were no conspicuous features that may cause trivial discrimination between faces. The constructed faces were chosen to be neither too similar nor too different from each other. The choice of not too different faces while creating the morphing sequence may leave the reader wondering if a morph effect only occurs between fairly similar faces. Here, the face images were selected to be significantly different when creating the sequence of morphed images. The results show that a morph effect also occurs with significantly different faces. Finally, previous studies employed a protocol that repeatedly presented subjects with a morph sequence (Wallis and Bülthoff 2001; Leutgeb et al 2005; Preminger et al 2007). The repetitive presentation leads subjects to generalise the entire morph sequence. The adaptation observed is therefore coupled with the generalisation process. Here, a new protocol is described in which subjects may adapt without necessarily generalising the entire sequence. The results show that by not forcing subjects to generalise the entire morph sequence, adaptation occurs very quickly. This process is interesting, since it is analogous to many real-life processes occurring in natural objects. Examples include drastic face changes of people getting older, object wear and tear, etc. In such real-life processes, the visual system adapts to sometimes dramatic environmental changes. Yet, as the change is gradual and historical views do not reappear, it is possible for the visual system to achieve such adaptation even without generalising all object historical views.

6 Conclusions

The visual system appears to adapt object categories based on stimulus similarity and without an external supervisor as demonstrated here and by Wallis and Bülthoff 2001; Leutgeb et al 2005; and Preminger et al 2007. Furthermore, the collective evidence suggests the accumulated change is not arbitrary but follows progress made by the stimuli. Subjects classifying stimuli that gradually divert from a prelearned object seem to follow such divergence. Apparently, stimulus classification leads to unsupervised training during which the perceptual boundary of previously learned object categories may adapt. The process that underlies the observed adaptation may be considered as unsupervised retraining. Such a process may offer the visual system the remarkable ability to adapt quickly to ecological changes without a supervisor and with a relatively small number of exposures. Importantly, our study reveals that the accumulation of change may be very rapid and that the observed phenomenon is robust. This visual system ability is not at all trivial when considering how difficult it would be to train a machine to adapt in the same way, as demonstrated in Hadas et al (under review).

Acknowledgments. Thanks to Hani Ron and colleagues from iamba who helped during the pre-experimental stage to build the experimental paradigm.

References

Beale J M, Keil F C, 1995 "Categorical effects in the perception of faces" *Cognition* **57** 217 – 239

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

Ehret G, 1987 "Categorical perception of sound signals: Facts and hypotheses from animal studies", in *Categorical Perception: The Groundwork of Cognition* Ed. S Harnad (Cambridge: Cambridge University Press) pp 301 – 331

Hadas D, Yovel G, Intrator N, under review "Using unsupervised incremental learning to cope with gradual concept drift"

Harnad S (Ed.), 1987 *Categorical Perception: The Groundwork of Cognition* (Cambridge: Cambridge University Press)

Hebb D O, 1949 *The Organization of Behavior: A Neuropsychological Theory* (New York: John Wiley)

Hopfield J J, 1982 "Neural networks and physical systems with emergent collective computational abilities" *Proceedings of the National Academy of Sciences of the USA* **79** 2554 – 2558

Leopold D A, Wilke M, Maier A, Logothetis N K, 2002 "Stable perception of visually ambiguous patterns" *Nature Neuroscience* **5** 605 – 609

Leutgeb J K, Leutgeb S, Treves A, Meyer R, Barnes C A, McNaughton B L, Moser M B, Moser E I, 2005 "Progressive transformation of hippocampal neuronal representations in 'morphed' environments" *Neuron* **48** 345 – 358

Newell F N, Bülthoff H H, 2002 "Categorical perception of familiar objects" *Cognition* **85** 113 – 143

Preminger S, Sagi D, Tsodyks M, 2007 "The effects of perceptual history on memory of visual objects" *Vision Research* **47** 965 – 973

Wallis G, 2002 "The role of object motion in forging long-term representations of objects" *Visual Cognition* **9** 233 – 247

Wallis G, Backus B T, Langer M S, Huebner G, Bülthoff H H, 2009 "Learning illumination- and orientation-invariant representations of objects through temporal association" *Journal of Vision* **9**(7):6, 1 – 8

Wallis G, Bülthoff H H, 2001 "Effects of temporal association on recognition memory" *Proceedings of the National Academy of Sciences of the USA* **98** 4800 – 4804

Willis T J, Lever C, Cacucci F, Burgess N, O'Keefe J, 2005 "Attractor dynamics in the hippocampal representation of the local environment" *Science* **308** 873 – 876

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 39 2010

www.perceptionweb.com