Vision Research 49 (2009) 1757–1764

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Adapting to altered image statistics using processed video

Michael Falconbridge [a,*], David Wozny [b], Ladan Shams [c], Stephen A. Engel [d]

[a] *Department of Psychology, UCSD, La Jolla, CA 92093, USA*
[b] *Biomedical Engineering IDP, UCLA, Los Angeles, CA 90095, USA*
[c] *Department of Psychology, UCLA, Los Angeles, CA 90095, USA*
[d] *Department of Psychology, U Minnesota, Minneapolis, MN 55455, USA*

ARTICLE INFO

Article history:
Received 25 September 2008
Received in revised form 11 January 2009

Keywords:
Learning
Natural images
Gabor
Facilitation
Plasticity

ABSTRACT

Perceptual systems can be altered by immersing observers in environments with statistical properties that differ from those naturally encountered. Here we present a novel method for placing observers in naturalistic audio visual environments whose statistics can be manipulated in very targeted ways. We present the results of a case study that used this method. Observers were exposed to an environment where there was a novel statistical relationship between two simple, visual patterns in otherwise natural scenes. Exposure to this altered environment strengthened perceptual interactions between the two patterns.

1. Introduction

Our environment has a strong influence on the properties of our perceptual systems over a range of timescales. Classic experiments like the monocular deprivation experiments of Wiesel and Hubel (1963) and the "stripe-rearing" experiments of Blakemore and Cooper (1970) and Hirsch and Spinelli (1970) revealed critical periods in perceptual development during which exposure to a normal environment is crucial for the development of normal perception (see Barlow, 1975 for an early review). Exposure to unnatural environments shifts neural resources towards the distribution of features within the altered sensory input. For example, depriving developing animals of information at specific orientations reduces the number of cortical neurons that are selective for that orientation (Blakemore & Cooper, 1970; Blasdel, Mitchell, Muir, & Pettigrew, 1977; Sengpiel, Stawinski, & Bonhoeffer, 1999). Clinical studies similarly indicate that there are critical periods for the development of visual capabilities in humans (Fawcett, Wang, & Birch, 2005; Olitsky, Nelson, & Brooks, 2002) – at least for capabilities mediated in early cortical levels of the visual system (see, e.g. Ostrovsky, Andalman, and Sinha (2006) for a case where mid to high levels appeared to recover from long term early blindness).

Input driven plasticity in *adult* humans has received a lot of attention recently, due in part to clinical implications. Information about adult plasticity is helpful in designing recovery programs, for example, following a stroke where some visual capability is lost or following cataract removal where capability is gained (Huxlin, 2008). Knowing about adult plasticity is made more important by the emergence of prosthetic and genetic technologies that may be used to restore low-level sensory capabilities (see, for example, Mancuso et al., 2007; Weiland & Humayun, 2006), but which may require substantial amounts of adaptation on the part of the patient.

Successful research in plasticity will depend on appropriate methods for manipulating the statistics of the environment. Optical methods have featured predominantly in past studies. Early experiments involved full-field shifts of incoming natural signals beginning with the use of the inverting lenses worn by Stratton (1897). Others used similar optical devices to invert, displace, and otherwise distort visual input. These were worn for hours, days, or even weeks (see list of methods in Rock, 1966). There is no doubt that the adaptation effects seen in these cases involved remapping between sensory and motor systems – it is not so clear that there was a change in *perception* (Linden, Kallenbach, Heinecke, Singer, & Goebel, 1999) (although note Gibson, 1933; Kohler, 1962). Another early study used full-field colour shifts. Using bi-coloured lenses, the left visual hemifield was tinted blue whilst the right was tinted yellow (Kohler, 1964). After weeks of adaptation, colour judgments with the lenses in place became as reliable as prior to adaptation. Remarkably, after removal of the coloured lenses the observer experienced a gaze contingent, bi-coloured visual environment – looking left tinted the world yellow and looking right made the world blue. This after effect lasted about a month. Simpler arrangements using lenses of a single colour have

---

* Corresponding author.
E-mail addresses: mfalconbridge@ucsd.edu, mickfalconbridge@hotmail.com (M. Falconbridge).

0042-6989/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2009.03.027

produced results that concord with this one (Neitz, Carroll, Yamauchi, Neitz, & Williams, 2002). Optical methods such as these, as well as those producing spatial shifts, are useful but are clearly limited in their scope for manipulating specific statistical properties of the input.

There has been a recent trend to study environmentally driven adult perceptual plasticity in the lab. In these cases, computer generated stimuli have been used to study, for example, learning of associations between complex unnatural shapes (Fiser & Aslin, 2001, 2002a, 2002b), the ability to recruit cues to disambiguate scenes (Backus & Haijiang, 2007; Haijiang, Saunders, Stone, & Backus, 2006), adaptation to associations between grating patterns (Carandini, Barlow, O'Keefe, Poirson, & Movshon, 1997; Falconbridge & Badcock, 2006), and adaptation to colour and orientation contingencies (McCullough, 1965; Vul & MacLeod, 2006). The advantage of these laboratory manipulations is that the statistics are under tight control, but the use of unnatural inputs in unnatural settings limits the potential for long term exposure, and may make it difficult to translate findings to real world situations.

Here we present a novel method for placing observers in a naturalistic environment where specific statistical features have been altered by the experimenter. A case study, consisting of two experiments on a total of 35 observers, demonstrates the usefulness of the method. Our results reveal the exciting possibility that the adult visual system is capable of adjusting low-level aspects of its perceptual model of the world to match the statistics of the environment.

The case study addresses the specific question of whether the adult visual system is capable of learning relationships between low-level visual features. By low-level features we mean simple visual patterns that match the receptive field profiles of early stage cortical neurons in primates. We chose "Gabor" patterns which are effective stimuli for activating simple cells in primary visual cortex (Jones & Palmer, 1987). We were motivated by two studies showing that the phenomenon of perceptual linking of co-linear Gabor-like patterns (Field, Hayes, & Hess, 1993; Li & Gilbert, 2002) may be *learned* from the environment during development (Hou, Pettet, Sampath, Candy, & Norcia, 2003; Kovacs, Kozma, Feher, & Benedek, 1999). This conclusion is based on three key points. (1) Infants do not exhibit a differential EEG responses to co-linear versus, e.g. parallel elements, and young children do not perform well in contour integration tasks, (2) adults do both, and (3) co-linearity is prevalent in the natural environment (Geisler, Perry, Super, & Gallogly, 2001; Sigman, Cecchi, Gilbert, & Magnasco, 2001). Our aim was to test for the possibility of such environmentally driven learning in adults. We chose a relationship between pairs of Gabors that is less common in natural images, and boosted its occurrence in real world video sequences. Specifically, we exposed adult observers to a strong *parallel* relationship between Gabor features.

We assessed learning by measuring the effect of a Gabor flanker on the apparent contrast of a target Gabor. Exposure to the parallel relationship increased the strength of the flanker effect, when the configuration of target and flanker matched the parallel relationship present in the altered environment. The change was positive so that a high contrast parallel flanker increased the perceived contrast of the target. This result suggests that the adult visual system retains the ability to learn new relationships between low-level features.

2. Methods

2.1. Adaptating stimuli

Observers viewed episodes of a popular television show that were manipulated to boost the prevalence of a parallel relationship between local oriented Gabor features. The use of a popular television show was designed to engage the attention of the observers. Informal reports by observers following exposure indicated high levels of attention. The materials of choice were video episodes of the television program "The Office" (NBC's US version). Seven episodes of "The Office" were converted to gray-scale (resolution: $480(h) \times 720(w)$ pixels). A control group of observers viewed the gray-scale episodes as they were, whilst experimental observers viewed manipulated versions of the video, described below. All videos subtended $16.6^\circ \times 25.3^\circ$ visual angle. The original audio soundtrack was presented to all observers along with the video. A sample manipulated video sequence is presented in Supplementary material A.

Wherever a Gabor pattern (here called "$g1$") occurred in any one of the original video frames, a second Gabor ("$g2$") with the same properties, but offset spatially was added to the frame at the same intensity (see Fig. 1a). $g1$ and $g2$ each subtended 24 minutes of arc viewing angle and were spaced 28 min. arc apart in a parallel arrangement. They were oriented at 135°, had zero phase, and had a peak spatial frequency of 4.3 cycles/deg. A formal description of the manipulation process follows.

The manipulated movie frame M is a weighted sum of the original frame O and an "added image" A:

$$\mathrm{M} = \mathrm{O} + \alpha \mathrm{A},$$

where α is a constant chosen to make the average amplitude of $g2$ equal to that of $g1$ for the first 20 movie frames of a given movie. To calculate A: let o represent a 2D fast fourier transform (MATLAB's 'fft2' function was used) of O, let b represent the fft2 of $g1$, and let c be the fft2 of $g2$. Then the fft2 of A (denoted by a) is

$$a = o \cdot * \, b \cdot * \, c,$$

where $\cdot *$ represents point-wise multiplication. This is equivalent to filtering[1] the original movie frame with $g1$, then convolving the resulting "amplitude map" with $g2$. This produces an image that consists of $g2$ s added at each point in the array according to some constant times the amplitude of $g1$ at that point. As stated, the constant was chosen to equate the amplitude of the added Gabor and the pre-existing Gabor. An example $g1$ amplitude map is shown in Fig. 1b.

Fig. 1c shows the manipulated image corresponding to this map. Overall, the manipulation increased the conditional probability of finding significant $g2$ energy in the image given the presence of $g1$.[2]

2.2. Relationship between added and pre-existing Gabors

In order to understand the effect of the manipulation process on the relationship between our parallel features, 100 movie frames were chosen randomly from an episode of The Office and another one hundred from a manipulated version of the same episode. $g1$ and $g2$ amplitude maps were produced by convolving the frames with $g1$ and $g2$, respectively (note that each $g2$ map was just a translated $g1$ map where the translation reflects the spatial relationship of $g2$ to $g1$). The correlations between corresponding points in the two maps were calculated for all 100 pairs of frames. The conditional probability of $g2$ given $g1$ as well as the $g1/g2$ joint probability were also calculated.

To avoid incorporating responses to clearly non-$g1$-like components (such as the top edge of the paper shredder in Fig. 1b) in our calculations, a cutoff of 15% of the maximum amplitude was applied, and only amplitudes above this value were considered. Fif-

---

[1] Filtering is equivalent to convolving but with the filter matrix rotated 180°. As our filter matrix ($g1$) is the same after rotation by 180°, we can just convolve D with $g1$.

[2] Note that g1 and g2 are identical Gabors but "g1" and "g2" labels are assigned to particular Gabors within a pair to depict their spatial relationship to one another.