Stimuli based on natural scenes for studying the contributions of luminance and colour to human spatial vision

D.J. Tolhurst, C. Alejandro Parraga, Tom Troscianko

Department of Physiology, Downing Street, Cambridge CB2 3EG and *Department of Experimental Psychology, 8 Woodland Road, Bristol BS8 1TN, UK

We are interested in studying how the human visual system deals with spatial discrimination tasks like those met in everyday life (Parraga et al. 2000). Starting from digitised photographs of natural scenes, we use a morphing technique to construct sets of naturalistic visual stimuli (Tolhurst et al. 1998); the pictures in the sets differ slightly, one from the next, in their spatial organisation (shape, brightness, texture, shading). We measure how well a human observer can discriminate small spatial changes in the pictures (of faces, say). Experiments can be performed with picture sets that have been distorted in various ways, to address the question of whether human vision is ideal for discriminating detail in natural (i.e. undistorted) pictures.

Up until now, we have used stimuli derived only from monochrome photographs. Now, we extend our methods using full-colour photographs of natural scenes in order to address questions about the relative contributions of luminance and chrominance information in spatial discriminations. As specific examples, we can morph pairs of photographs of different fruits, which differ in shape, colour and texture. Such subjects lend themselves readily to morphing and, more importantly, they may be especially appropriate to the study of vision in the natural setting, since it is proposed that primate red-green colour vision may have evolved specifically to solve the task of finding ripe fruit, hidden in green leaves (Mollon, 1989; Osorio & Vorobyev, 1996).

We start by photographing fruits under controlled illumination, using a Nikon digital camera. A pair of such pictures is then used to build a morphed sequence of 40 steps, each step representing a shift of 2Â·5 % between the two parent pictures. Human observers are asked to determine how big a change in one parent picture (i.e. how many of the 2Â·5 % steps) is needed for reliable discrimination. The wavelength sensitivity of the camera's three sensors (R, red; G, green; B, blue) is determined by photographing a white surface viewed through a series of narrow-band interference filters. The camera's output is calibrated against radiometric measurements of the filters' transmittance. The sensitivity curves allow us to transform the three (RGB) planes of the morphed photographs into three new planes (LMS), which represent how well the three human cone types (L, long; M, medium; S, short) would respond to the scene. Before display on a computer monitor, the three LMS planes in the pictures must be transformed to match the spectral emission of the monitor's three (RGB) phosphors.

The pictures in the morphed set may be manipulated in various ways, depending upon the particular psychophysical question. We could, for instance, change the slopes of the power spectra in all three colour planes (LMS) together, giving the effect of blurring or 'sharpening' a coloured picture, much as we have already done with monochrome pictures (Parraga et al. 2000). Or, we could systematically disrupt the phase spectra (Thomson & Foster, 1997) in the three planes together. The more interesting potential of these stimuli is that we can manipulate the spatial properties of different colour planes differently. In fact, we do not deal with the three cone planes themselves (LMS), but with linear transformations of them. A luminance plane (L + M) can be separated from a red-green opponent plane ([L - M]/[L + M]) and a blue-yellow opponent plane ([S - Y]/[S + Y]), where Y is [L + M]/2. We can then manipulate the power spectrum or phase spectrum of the luminance plane differently from those of the two colour-opponent planes taken together, before reconstructing a changed LMS picture. In this way, we can attempt to distinguish the different roles of luminance and colour spatial information in natural visual discrimination tasks.

As a future possibility, we could imagine applying non-linear transformations to LMS pictures, to generate three new planes corresponding to 'value' or 'lightness', 'hue', and 'chroma' (saturation) in the CIE L*C*H or Munsell colour spaces. These perceptual attributes of coloured pictures could be manipulated separately before reciprocal non-linear transformations generate modified LMS pictures.