The use of morphed images in the quantitative study of human spatial vision

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

Journal of Physiology, Volume 506, page 11P--12P - 1997

Most neurophysiological and psychophysical study of spatial vision has relied on the use of simple stimuli such as spots and sinusoidal gratings. We should now ask whether knowledge of low-level visual mechanisms gained with these stimuli allows an understanding of a person's abilities to perform biologically meaningful visual tasks, e.g. discrimination between the faces of different people. We aim to develop a rigorous and quantitative psychophysics that uses visual stimuli made from photographs of real visual scenes for the study of, say, amblyopic anomalies or the differences between foveal and peripheral vision. This is feasible since ordinary computers now have the power to handle the large data structures needed for digitized images, and their graphics systems are capable of high resolution, precise image display.

The basic requirement for quantitative psychophysics is a series of well-defined complex images that differ, one from the next, by only a small amount in some domain. We can then ask how large a difference in that domain is needed for a human observer to be able to distinguish two images in the series reliably. Previous attempts to make visual stimuli from digitized photographs of natural scenes (e.g. Tadmor & Tolhurst, 1994) have not been fully appropriate, since the observer had to discriminate unnatural distortions in visual scenes rather than to discriminate between two different natural scenes. We now describe a paradigm in which each stimulus in the series could potentially be a natural scene.

Morphing is a special-effects technique used in video and cinema, in which a photograph of one object is gradually changed into that of another. The individual spatial features of one photograph are changed in shape step-by-step until they have the shape found in the second photograph. We have used custom-written (Benson, 1994) or commercial software (Morph 2.5, Gryphon Software) to make morphed series for psychophysical experiments. The methods for morphing and some of their pitfalls are given by Benson (1994). In a two-stage operation, both original photographs are first delineated using a set of feature landmarks; joining these points in appropriate order creates an accurate line-drawing of each. The co-ordinates of corresponding pairs of points are compared and the difference computed. The morph transform moves the landmark points of the first original into their corresponding positions in the second original by successively reducing the difference between the two pictures in, say, 5 % steps. Tonal or textural information (for a grey-level image) is transformed in the second stage. A triangular mesh is constructed to tessellate each of the two original images. Individual triangular patches are then affine-warped into the intermediate shape of the morphed image, taking with them the relative contribution of tonal pixel intensity appropriate to that stage in the morphing sequence.

The results of morphing may look unrealistic if there are too few control points, or if the objects in the two originals are very different in size or in tone (e.g. if the originals are of two different faces and if only one of them has strong shadows). This may be a limitation on the possible scenes that can be morphed, but it does tends to ensure that the overall luminance, contrast and spatial-frequency content of all the images in the sequence stay much the same. For experiments, these variables must be controlled otherwise the observer may be able to use spurious cues for discrimination, such as overall brightness or visibility.

Experiments are performed using a modified two-interval forced choice paradigm. There are three time intervals in each trial. In the middle interval, a reference image is shown. In either the first or the third interval, a second copy of that reference is shown; and in the remaining interval, a morphed test image is shown. The observer must identify whether the morphed image appeared in the first or third interval. Depending upon whether the observer's choice is correct, a staircase procedure makes the test stimulus for the next trial more or less different from the reference. We have found that observers can detect when a photograph has been morphed 2Â·5-20 % of the way towards a second photograph, and we have found it necessary to make the morphed series in steps of 1Â·0 % to 2Â·5 % each in order for the staircase procedure to work effectively.

Supported by Fight for Sight (D.J.T. and C.A.P.) and the MRC (P.J.B.).