

Interpolating sampled contours in 3D: perturbation analyses

Paul A. Warren^{a,*}, Laurence T. Maloney^b, Michael S. Landy^b

^a Department of Psychology, Cardiff University, Tower Building, Park Place, P.O. Box 901, Wales, UK

^b Department of Psychology and Center for Neural Science, New York University, New York, NY, USA

Received 6 March 2003; received in revised form 31 October 2003

Abstract

In four experiments, observers interpolated parabolic sampled contours confined to planes in three-dimensional space. Each sampled contour consisted of eight visible points, placed irregularly along the otherwise invisible parabolic contour. Observers adjusted an additional point until it fell on the contour. We sought to determine how each visible point influenced interpolation by measuring the effect of slightly perturbing its location. Influence fell rapidly to zero as distance from the interpolated point increased, indicating that human visual interpolation of parabolic contours is local. We compare the measured influence for human observers to that predicted by three standard interpolation algorithms. The results were inconsistent with a fit of a quadratic to the points, but were reasonably consistent with a cubic spline and most consistent with an algorithm that minimizes the variance of angles between neighboring line segments defined by the sampled points.

© 2003 Elsevier Ltd. All rights reserved.

1. Introduction

The white clothesline in Fig. 1 is readily visible, although the visual evidence signaling its presence is fragmented and sparsely distributed. Assembling this information into a coherent estimate of the location of the clothesline is a formidable computational problem. For convenience, it can be broken down into three subordinate problems. The first is to determine that there are one or more fragmented contours present in the scene (the *Detection Problem*). The second is to decide which parts of the image carry information about the hypothetical fragmented contour (the *Grouping* or *Segmentation Problem*). To solve the third (the *Interpolation Problem*), the visual system must correctly estimate the position of the fragmented contour both where it is and is not obscured by other objects.

Visual segmentation and interpolation of fragmented contours is a research topic of major interest (Field, Hayes, & Hess, 1993; Kanizsa, 1979). Most previous work, though, has concentrated on the Detection Problem. Field et al., for example, required their observers to judge whether a contour, consisting of a



Fig. 1. Contour completion. The white clothesline is readily visible although much of it is obscured by clothing.

small number of Gabor patches lying along a curved path, was present in a visual display containing other, masking Gabor patches. It is important to recognize that successful performance of their task does not entail that the observer has correctly grouped all of the Gabor patches that belong to the contour or can correctly interpolate them.¹ In viewing their stimuli, we may feel that we can both group and interpolate, but successful completion of the task does not require that we do

* Corresponding author. Tel.: +44-29-2087-0078; fax: +44-29-2087-4858.

E-mail address: warrenpa@cardiff.ac.uk (P.A. Warren).

¹ To determine which of two intervals contains a contour, the observer need only detect a subset of the Gabor patches that lie on the contour.




Fig. 2. Schematic of the stimulus. The sampled contour is shown as black points on a white background. The adjustable point p_0 is constrained to move within a *setting plane* that intersects the invisible contour at the invisible true point τ (marked with a circle). The coordinate axes (n, b) in the setting plane are shown. The figure is a stereogram. For crossed fusion, use the left-hand pair of images; for diverged fusion, use the right-hand pair.

either. Similarly, much research on illusory contours emphasizes their presence or absence but not their precise location (e.g., Kanizsa, 1979), although the task used by Ringach and Shapley (1996) does require discrimination of the sign of curvature of an illusory contour.

There is less work directly addressing the grouping problem (for a review, see Warren, Maloney, & Landy, 2002) and very little work concerning the Interpolation Problem. Warren et al. measured how accurately human observers could interpolate linear and parabolic sampled contours in three-dimensional space. One of their stimuli,² a sampled parabolic contour, is illustrated in Fig. 2. On each trial, the observer saw such a collection of points, and adjusted one of them to fall on the contour shared by the remainder. The adjustable point was constrained to lie in an invisible plane that intersected the invisible contour once. The plane was orthogonal to the contour at the point of intersection. The computational theory appropriate for solving this sort of problem is the theory of splines³ (de Boor, 1978).

Warren et al. found that observers' settings were very reliable. The standard deviation in repeated settings was about 1/30th of the gap between points across which the observer interpolated, comparable to human performance in Vernier acuity tasks (Klein & Levi, 1987). Significantly, the settings displayed no patterned biases away from the parabolic contour across conditions and observers. While it is scarcely surprising that human observers interpolated linear contour segments as linear,

it is interesting that observers interpolated parabolic contour segments as parabolic rather than as some other family of contours, slightly flatter or more curved at the point of interpolation. These results suggest that parabolic (and linear) contours play a special role in visual spline interpolation.

Warren et al. (2002) varied the number of points that the observer could see in interpolating a linear contour (2, 4, 6, and 8 points) or a parabolic contour (4, 6, 8, and 10), removing pairs of points furthest from the point of interpolation in succession. Their goal was to examine how the amount of visible information (measured in number of points) affected the variability⁴ of observers' settings. In particular, they sought to determine whether the setting variability of parabolic spline interpolation with 6, 8 or 10 points might compare with the setting variability of linear interpolation with two points. If the two were comparable, then they could conclude that providing additional visual information canceled the increase in uncertainty associated with interpolating a curved contour.

Surprisingly, setting variability did not decrease significantly as the number of visible points increased for either the linear or parabolic contour. One possible interpretation of their results is that observers, in interpolating, are using only the innermost four points for the parabolic segment⁵ or the innermost two for the linear segment. In terms of the theory of splines, the hypothesis can be restated as the claim that the "human visual spline" is local. A local spline with window m is one that bases its interpolation of the gap between points p_i and p_{i+1} on the $2m$ points $p_{i-m+1}, \dots, p_i, p_{i+1}, \dots, p_{i+m}$, ignoring the remainder. The hypothesis that emerges from Warren et al. can be restated as: Human visual interpolation of parabolas is local with window 2 and interpolation of lines is local with window 1. This *local spline hypothesis* is illustrated in Fig. 3.

This local spline hypothesis with $m = 2$ is consistent with the contour grouping models of Feldman (1997) and approximately so with the contour grouping model of Pizlo, Salach-Golyska, and Rosenfeld (1997). We will return to the latter model in the general discussion.

In this article, we test the local spline hypothesis for sampled parabolic contours in three-dimensional space. We also address other hypotheses that we explain once we have described the coordinate systems and dependent measures we employ. We will make use of a perturba-

² The actual stimuli were self-luminous "blobs" suspended in three-dimensional space against a featureless black background. In the figures we will represent the stimuli as black points on a white background for clarity, and refer to the self-luminous blobs as "points".

³ A spline algorithm assigns a unique curve to an ordered series of points p_1, \dots, p_N with the property that the resulting curve passes through each of the given points. Spline algorithms can also include constraints on the slope of the resulting curve at each point. Since each point here is indicated by an unoriented blob, these more complex spline algorithms are not relevant to our task and stimuli.

⁴ The observers adjusted a point confined to a plane and the variability of the settings was characterized by the settings' covariance matrix. We use the term *setting variability* to refer to this covariance matrix.

⁵ Since the parabolic segment is confined to a plane, it is possible that the observer is using only three of the four points displayed. Earlier results for parabolic contours confined to the fronto-parallel plane (Koh & Maloney, 1988) do, however, show a decrease in setting variability with an increase from three to four visible points.




Fig. 3. The local spline hypothesis. The visual system uses a moving $2m$ -point window in computing the spline interpolation of the invisible contour from the visible points. In the figure, m is 2, and the window is shown for interpolation at three different points along the invisible contour. A spline algorithm that can be expressed in this form is a *local spline*.

tion technique first described by Maloney and Landy (1989) (Landy, Maloney, Johnston, & Young, 1995) and first applied to visual interpolation by Hon, Maloney, and Landy (1997). The goal is to measure certain partial derivatives that characterize how each of the visible points in Fig. 2 contributes to the “human visual spline”.

2. Coordinate systems and the influence matrix

We first develop mathematical notation needed to represent the visual splining task just described. The experimenter selects visible points $p_{-N}, \dots, p_{-1}, p_1, \dots, p_N$ with coordinates $p_i = (x_i, y_i, z_i)$ that fall on an invisible parabolic segment constrained to lie on a plane in three-dimensional space referred to as the *contour plane*. A *setting plane* Π_S is chosen that intersects the invisible parabolic segment only once between p_{-1} and p_1 and that is perpendicular to the parabolic contour at the point of intersection (Fig. 2). This point of intersection is referred to as the *true point*, τ . In the notation, we replace $p_{-N}, \dots, p_{-1}, p_1, \dots, p_N$ by \vec{p} for convenience. The observer’s task is to select a point p_0 in Π_S that “falls on the contour”. In our experiments, observers do so by adjusting the position of a point constrained to lie in the plane Π_S .

We denote the point that the subject selects as the interpolation point by $p_0 = S(\Pi_S; \vec{p})$. In general, $S(\Pi_S; \vec{p})$ is a random variable: the observer will not necessarily select the same point given the same stimulus. Based on Warren et al. (2002), we can model this trial-to-trial variability as additive, zero-mean Gaussian noise and write

$$S(\Pi_S; \vec{p}) = s(\Pi_S; \vec{p}) + \varepsilon, \quad (1)$$

where $s(\Pi_S; \vec{p})$, the *visual spline function*, is the expected value of the observer’s setting which we estimate by averaging the observer’s settings across many trials.

Since the observer’s setting is constrained to be in the setting plane, it is convenient to report settings in a coordinate system confined to the setting plane. The origin of this *setting plane coordinate system* is the true point, τ , and the two axes are referred to as n and b (Fig.

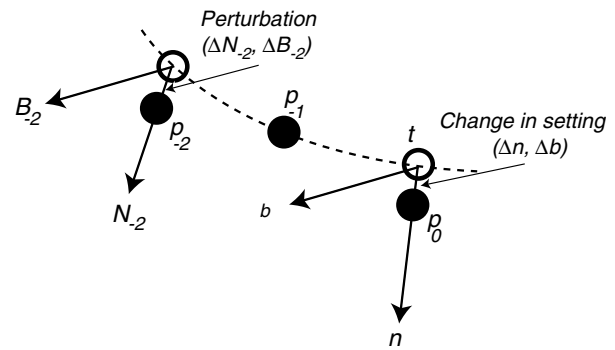


Fig. 4. Setting coordinates and perturbation coordinates. The sampled parabolic contour consisted of eight points in the fronto-parallel plane, three of which are shown here. The observer’s task was to move the adjustable point p_0 , constrained to lie in the setting plane, until it appeared to fall on the contour. The setting plane was orthogonal to the invisible parabolic contour at the point where it intersects it (the true point τ). The coordinate system of the setting plane (n, b) is shown. On some of the trials, one of the contour points (in the figure, p_{-2}) was perturbed away from the invisible contour. The perturbed point was constrained to lie in a *perturbation plane* that intersected the contour at the unperturbed location of the point and was orthogonal to the contour at that point. The coordinate system (N_i, B_i) of one perturbation plane is shown. A possible perturbation ($\Delta N_i, \Delta B_i$) is shown, as is a possible change in setting in response, ($\Delta n, \Delta b$). The magnitudes of perturbation and response shown in this schematic are much larger than those employed in the experiments reported here.

4). Both axes are perpendicular to the contour at the true point (since they are contained in the setting plane Π_S which is perpendicular to the contour at the true point). The axis b (the *binormal*) is orthogonal to the contour plane. The axis n (the *normal*) is in the contour plane, and is orthogonal to both b and to the contour at the true point τ . We represent the observer’s setting in setting plane coordinates, and we define the “output” of the vector-valued function $s(\Pi_S; \vec{p})$ to be in setting plane coordinates as well,

$$\begin{bmatrix} n \\ b \end{bmatrix} = \begin{bmatrix} s_n(\Pi_S; \vec{p}) \\ s_b(\Pi_S; \vec{p}) \end{bmatrix}. \quad (2)$$

Our first goal is to test the local spline hypothesis. The local spline hypothesis is the claim that only a subset of the points $p_{-N}, \dots, p_{-1}, p_1, \dots, p_N$, the ones nearest to the point of interpolation, have any influence on the expected setting $s(\Pi_S; \vec{p})$. Intuitively, we can imagine testing the hypothesis by reaching into Fig. 4 (or Fig. 2) and wiggling any one of the visible points p_i a little bit. If small displacements of the point do not alter the interpolation point then we conclude that the point we “wiggled” had no influence on the visual spline function $s(\Pi_S; \vec{p})$.⁶ We formalize this intuition in terms of a

⁶ We could also move the visible point a large distance, rather than just “wiggling” it. However, large displacements may lead the visual system to re-segment the scene, in effect removing the point from the contour. Hon et al. (1997) investigated the effect of large displacements and found that this was the case.

particular set of partial derivatives of $s(\Pi_S; \vec{p})$, which we define next.

Let p_i denote any one of the visible points $p_{-N}, \dots, p_{-1}, p_1, \dots, p_N$ that we designate as the *perturbed point*. Let Π_i denote the plane through p_i orthogonal to the parabolic contour at p_i , which we refer to as the *perturbation plane*. Just as we did for the setting plane, we establish a *perturbation coordinate system* (N_i, B_i) with its origin at the perturbed point, and N_i and B_i defined analogously (Fig. 4).

We are interested in estimating the effect of small displacements of the perturbed point p_i within the perturbation plane. These effects can be summarized by the matrix of partial derivatives

$$I_i = \begin{bmatrix} \frac{\partial s_n}{\partial N_i} & \frac{\partial s_n}{\partial B_i} \\ \frac{\partial s_b}{\partial N_i} & \frac{\partial s_b}{\partial B_i} \end{bmatrix}, \quad (3)$$

which we refer to as the *influence matrix* corresponding to the point p_i . It is a portion of the Jacobian matrix of the visual spline function s . We will estimate this influence matrix experimentally as explained in the methods section below. The influence matrix is closely related to measures of influence used in the theory of robust statistics (Hampel, Ronchetti, Rousseeuw, & Stahel, 1986).

To see the significance of the influence matrix for our task, consider a perturbation of point p_i . For small enough perturbations we can approximate $s(\Pi_S; \vec{p})$ by a truncated Taylor series

$$\begin{bmatrix} n \\ b \end{bmatrix} \approx \begin{bmatrix} n_0 \\ b_0 \end{bmatrix} + \sum_{i=-N}^N I_i \begin{bmatrix} \Delta N_i \\ \Delta B_i \end{bmatrix}, \quad (4)$$

where $[\Delta N_i, \Delta B_i]'$ represents the perturbation applied to p_i , and $[n_0, b_0]'$ is the expected setting in the absence of a perturbation. Thus, the influence matrices characterize the response of the visual spline function to small perturbations of the visible points; they comprise the linear component of the human visual spline function.

For a given visible point p_i , the influence matrix provides information about how that point enters into the computation of the human visual spline. If, for example, a visible point has no role in interpolation, then the entries of its influence matrix should be zero. The prediction of the local spline hypothesis, then, is that the influence matrices for visible points far from the point of interpolation should be zero. We test this prediction in the experiments below.

The stimuli we use are contours confined to a plane in three-dimensional space. It is natural to ask whether the visual system makes use of this planar constraint. Consider, for example, what might happen if we perturb one of the visible points in the N_i -direction only so that, although perturbed, it remains in the contour plane. Will the resulting displacement of the setting point also

remain in the contour plane? Conversely, if we push the visible point along the B_i -direction, orthogonal to the contour plane, will the resulting change in interpolation be a displacement along the b -direction only? In terms of the influence matrix, we are asking whether the off-diagonal entries, $\partial s_n / \partial B_i$ and $\partial s_b / \partial N_i$ are zero. We refer to this hypothesis as the *dimensional independence hypothesis*. It is interesting to note that previous work concerning interpolation in the fronto-parallel plane effectively assumes this hypothesis without testing it by assuming that the contour defined by visible points (or Gabor patches or fragmentary contours) in the fronto-parallel plane must be contained completely in the fronto-parallel plane. We test this hypothesis for the fronto-parallel plane and for a non-fronto-parallel plane.

Any proposed model of human visual interpolation of parabolic contours must, of course, reproduce the performance of human observers in Warren et al. (2002) (Hon et al., 1997). Unfortunately, most spline algorithms approximate parabolic contours very well and it is difficult to reject models based on interpolation performance alone. However, a valid model of human performance must not only match human performance in interpolation, but must also have the same influence matrices. Intuitively, it must make use of the same points in the same way as characterized by influence. The influence matrices measured in the experiments below provide a potentially useful tool for selecting among models of human visual interpolation.

In the experiments that follow, influence matrix elements will be estimated experimentally. The degree of linearity will be tested, as will dimensional independence. Finally, the measured influence matrices will be compared to those predicted by several possible models of the human visual spline.

3. General methods

3.1. Apparatus

We used two Sony Trinitron Multiscan G500 monitors, positioned on either side of the observer, to display stimuli (Fig. 5). The two monitors formed part of a Wheatstone stereoscope: the image from the left monitor was projected to the observer's left eye by a small half-silvered mirror placed at 45° to the observer's Cyclopean line of sight. A second half-silvered mirror reflected the image of the second monitor to the observer's right eye. The partial transparency of the mirrors facilitated spatial calibration of the monitors (described below) but played no other role in the experimental sessions. The optical distance from each eye to its corresponding monitor was approximately 70 cm. From this distance, the central region of each




Fig. 5. The experimental apparatus. The observer was seated in a Wheatstone stereoscope. The left and right retinal images of a stereo pair were presented on computer monitors and viewed by means of small mirrors. The fused image consisted of a small number of self-luminous points. The apparatus was contained in a large box lined with black, flocked paper to absorb stray light. A combination of anti-aliasing software and calibration procedures permitted sub-millimeter accuracy in positioning points in visual space.

screen, used to display our stimuli, subtended approximately 24×18 deg of visual angle. The screens of the G500 monitors are close to physically flat. All stimuli were generated using MatLab[®] (The Mathworks Inc; Hanselman & Littlefield, 1997) on a Dell Precision 410 workstation running under the Linux operating system.

Observers were positioned in a chin rest and were asked to keep their heads still, although no head restraint was imposed. The apparatus was housed within a large box, the interior of which was covered in black, flocked paper (Edmund Scientific), a highly light-absorbent surface. The observer could see only the points defining the stimulus, apparently floating in front of him or her against a black background. In the experiments described below, the task of the observer was to move a point in space until it appeared to lie on a contour sketched by other points. Each point was a trivariate Gaussian 3-D “blob” of light that could be positioned in space with high resolution. At 70 cm directly in front of the viewer, this resolution was 0.07 mm in the horizontal and vertical directions and 0.14 mm in depth, corresponding to $21''$ visual angle in the vertical and horizontal directions and $42''$ of disparity resolution. This resolution is small compared to observers’ setting variability in these tasks. The computational methods used to present these points in stereo are described by Warren et al. (2002), including anti-aliasing methods that permitted display of points in space with high spatial resolution (Georgeson, Freeman, & Scott-Samuel, 1996).

We calibrated the apparatus spatially before each experimental session. Using only the left eye, the observer first viewed a 4×5 array of points on the left monitor superimposed on a physical reference by one of the half-silvered mirrors. The calibration reference tar-

get was a 4×5 array of points on a rigid, planar surface placed 70 cm in front of the observer. The observer moved each point separately until it appeared to lie on top of the corresponding physical reference dot. This process was then repeated for the right eye.

3.2. Stimulus configurations

We used a parabolic contour segment. The equation of the parabola in cm relative to the bottom left hand corner of the central viewing area of the screen for the fronto-parallel condition was $y = 0.0033x^2 - x + 152.4$ (the parabola was rigidly rotated for the non-fronto-parallel condition). The observer saw only sample points constrained to lie on the segment as illustrated in Fig. 2. The vertical and horizontal extents of the parabolic section in the contour plane were approximately 5.5 and 26 cm (corresponding to 4.5° and 20.4° visual angle for the fronto-parallel plane condition). We first selected nine points: the two endpoints of the segment and seven more points whose positions were computed to be equally spaced in arc-length along the contour. If the observer knew that the distance between successive points along the contour was always the same, he or she might use spacing as a cue in interpolating the contour. Therefore, we jittered the positions of the points by sliding them a random amount along the (invisible) contour. After jittering, spacing between successive points was evidently non-uniform (Fig. 2). The average linear distance between the sampled points was 35.3 mm (corresponding to approximately 3° visual angle in the fronto-parallel condition). The locations of the points on the contour were not varied on a trial-by-trial basis; they remained constant throughout the experiment.

3.3. Procedure

In each trial in all four experiments, observers saw an eight-sample planar parabolic contour in 3-D space. The invisible ninth point τ (the *true point*) was always in the middle of the series of visible sampled points defining the contour (Fig. 2). The observer was instructed to move an additional point (p_0 , the *adjustable point*) until it lay on the perceived contour (method of adjustment). Movement of p_0 was confined to a plane that we will refer to as the *setting plane*. The initial position of p_0 was chosen randomly within the setting plane. The setting plane intersected the contour at τ , and was perpendicular to the contour at τ (Fig. 4).

Observers used six buttons throughout the course of the experiment. Four of these moved p_0 in the setting plane. Prior to the experiment we selected two direction vectors in the setting plane for each angle and curve condition. To each direction vector we assigned two of the four keys. Pressing one key of the pair moved the point one way along the vector, pressing the other

moved it in the opposite direction. We found that observers were best able to learn to move p_0 when one of the direction vectors had zero absolute depth component, lying as close to vertical as possible while still constrained to lie in the setting plane. The other direction vector was perpendicular to the first, and therefore had a depth component. Observers quickly became comfortable with this mapping of keys to movements of p_0 . Defining the actions of the movement keys in this manner also ensured that no inherent significance was attached to the normal and binormal directions to the contour at τ (the directions along which we measured the components of the influence matrix).

At the start of a trial the control program permitted “quick” movement of the point—each key press displaced the point by approximately 0.5 mm in a direction in space defined by the experimental conditions. When the observer judged that the adjustable point p_0 was near the contour, they pressed a fifth key which allowed them to move the point with greater precision (at the limit of resolution of the apparatus) until they were satisfied with their setting. A final press of the sixth key recorded the observer’s setting and triggered the next trial. No feedback was given to observers since the results of Warren et al. (2002) demonstrated that none was necessary to ensure accurate interpolation performance.

3.4. Rotations

In the experiments described below, the stimuli were either presented in the fronto-parallel plane or were rotated by 70° about a vertical axis through the true point τ (Fig. 2). The fronto-parallel and 70° stimuli alternated on successive trials to prevent observers from noticing any changes in position of points associated with changes in perturbation condition, as described below.

3.5. Coordinate systems

Throughout this paper we will refer to *absolute* and *intrinsic* coordinate systems. The *absolute coordinate system*, (X, Y, Z) , is simply the fixed frame of reference of the apparatus centered on the true point, τ . We chose the convention that this frame is left-handed with the X -, Y - and Z -directions corresponding to rightwards, downwards and towards the observer, respectively.

Intrinsic coordinates were introduced previously (Warren et al., 2002). They are the coordinates in the plane orthogonal to the contour at a specified point along the contour. At the true point, τ , this plane is the setting plane and the setting coordinates axes in the plane are denoted b and n (Fig. 4). The axis b is perpendicular to the contour plane; the axis n is in the contour plane. We also define the unit tangent vector, t , to the curve at the true point. We will only make use of

this last coordinate axis in discussing modeling results in the conclusion. The *perturbation coordinate system* is defined as the intrinsic coordinate system at the current perturbed point p_i and is denoted (N_i, B_i) .

3.6. Measuring the influence matrix

In each of the following experiments, we perturb some of the visible points on some of the trials (*perturbed trials*). On other trials, no points are perturbed (*unperturbed trials*). When a point is perturbed, we will alter its position in either the N_i -direction by an amount ΔN_i or in the B_i -direction by an amount ΔB_i , but not in both. We estimate the effect of the perturbation ΔN_i in the setting plane by computing the difference between the mean settings on the unperturbed and the perturbed trials, $(\Delta \hat{n}_{N_i}, \Delta \hat{b}_{N_i})$. Similarly, we estimate the effect of the perturbation ΔB_i in the setting plane by computing the difference between the mean settings on the unperturbed and the perturbed trials, $(\Delta \hat{n}_{B_i}, \Delta \hat{b}_{B_i})$. The resulting *estimated influence matrix* for the perturbed point is

$$\hat{I}_i = \begin{bmatrix} \frac{\Delta \hat{n}_{N_i}}{\Delta N_i} & \frac{\Delta \hat{n}_{B_i}}{\Delta B_i} \\ \frac{\Delta \hat{b}_{N_i}}{\Delta N_i} & \frac{\Delta \hat{b}_{B_i}}{\Delta B_i} \end{bmatrix}, \quad (5)$$

which is identical to the influence matrix of Eq. (3) with empirical estimates replacing some of the quantities. We are especially interested in conditions of sufficiently small perturbations for which \hat{I}_i is a valid estimate of the Jacobian of the human visual spline function. In the next two experiments we test whether the perturbations we use are small enough for linearity to hold, allowing us to consider \hat{I}_i as an estimate of the Jacobian.

4. Experiment 1

In this experiment we examine how scaling the magnitude of the size of the perturbations ΔN_i and ΔB_i affects measured influence. Over the range where perturbation is homogeneous, we expect that scaling ΔN_i by one half would scale the corresponding effects of perturbation $(\Delta \hat{n}_{N_i}, \Delta \hat{b}_{N_i})$ by the same factor, and that scaling ΔB_i by one half would halve $(\Delta \hat{n}_{B_i}, \Delta \hat{b}_{B_i})$. Further, we expect that the inverse perturbation $-\Delta N_i$ should result in the inverse effect $(-\Delta \hat{n}_{N_i}, -\Delta \hat{b}_{N_i})$ and similarly for reversing ΔB_i . If these predictions hold, there is no net effect on the estimated influence matrix \hat{I}_i since the factors of 1/2 or -1 appear in both the numerators and denominators of each entry in Eq. (5) and cancel. We test homogeneity separately for stimuli in the fronto-parallel plane and for stimuli rotated 70° about a vertical axis.

4.1. Methods

On each trial, observers viewed an eight-point sampled contour, and adjusted the position of p_0 within the setting plane so that it appeared to lie on the perceived contour. On some trials a single point (p_1 , the point immediately to the right of p_0 along the contour) was perturbed in the normal direction N_1 .

The factors in this experiment were perturbation size (0, 0.8 or 1.6 mm, corresponding to visual angles of 0', 4' and 8' in the 0° condition), contour rotation angle (0° or 70°) and perturbation direction (−1, +1). The factors were fully crossed resulting in 12 experimental conditions. Note, however, that with perturbation size zero (i.e., no perturbation) the direction factor collapses. This naturally gives twice as many settings for any unperturbed condition as any perturbed condition. Observers completed 12 repetitions of each condition over a period of three hour-long sessions, leading to a total of 144 settings per subject.

The perturbation sizes were chosen based on the results of Warren et al. (2002). There it was found that variability was within 8' in the fronto-parallel plane. Thus, it is unlikely that subjects could detect even the largest perturbation, which was approximately the same size as the upper bound on setting variability, and examination of the stimuli confirmed this claim. The angle conditions were interleaved so that observers could not complete the task by remembering the location of the point between trials. Each trial took approximately 1 min to complete so it is unlikely that observers were able to use location memory to complete the task. All other conditions were randomized.

4.2. Observers

Three observers participated in the experiment. Two of the observers, PAW and LTM, were authors. The third observer, JT, was an experienced psychophysical observer unaware of the purpose of the experiments. All observers had normal or corrected-to-normal vision.

4.3. Results

Fig. 6 shows the entire data set for a single observer (PAW) in setting plane coordinates. First, note that the distributions of settings around each mean setting point are similar in all perturbed and unperturbed conditions. The main effect of perturbation is a shift in mean setting. The setting distributions in this and subsequent experiments are similar to those of Warren et al. (2002) who reported uncertainties of ± 1.5 mm ($\pm 7.5'$ visual angle) in the X - and Y -dimensions and ± 2 mm (less than $\pm 1'$ disparity) in depth. In the intrinsic coordinate frame, for the unperturbed condition, the maximum standard deviation across all three observers, two angle conditions and both coordinate directions was 1.1 mm (3% of the distance to the nearest contour point). Similarly the maximum standard deviations for the 0.8 and 1.6 mm perturbation conditions pooled over positive and negative perturbation directions were 1 and 1.2 mm. In all three perturbation conditions the maximum standard deviation was observed for the b -component of the 0° rotation condition. Since this corresponds to the depth dimension the value can be equivalently expressed as approximately 0.5' disparity for all three perturbation

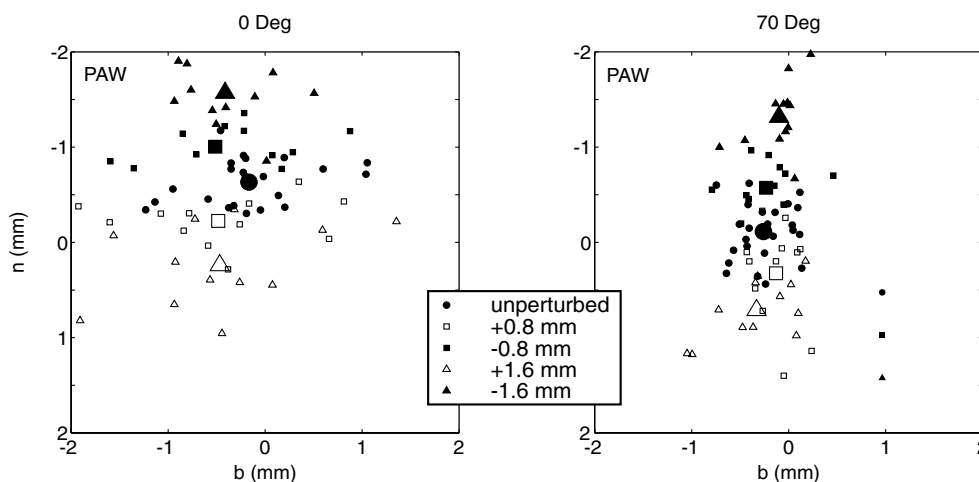


Fig. 6. Experiment 1: Settings in response to perturbations in the N -direction. The settings for observer PAW are plotted in setting plane coordinates (b, n). The origin is the point τ where the invisible parabolic contour intersected the setting plane. For the unperturbed condition and each of four perturbations in the N -direction (± 0.8 and ± 1.6 mm), the distribution of setting points is approximately bivariate Gaussian. Large symbols correspond to the centroids of these distributions. The centroids primarily shift in the n -direction with large displacements for the larger magnitude perturbation. The sign of the shift agrees with the sign of the perturbation. Results are shown for the 0° and 70° conditions. For the 0° condition, note that within a perturbation condition the settings are slightly more variable in the b -direction, which is roughly along the observer's line of sight.

conditions. These quantities are very similar to those reported previously (Warren et al., 2002).

The mean of the setting errors in the n -direction is displaced for the perturbed trials, and perturbations of opposite sign lead to displacements in the opposite direction. The larger the perturbation magnitude, the larger the displacement was in the n -direction. Also, note that perturbation in the N_f -direction has little effect on the mean of the distribution of settings in the b -direction. This is consistent with the dimensional independence hypothesis. We will test this hypothesis more rigorously in Experiments 3 and 4.

Fig. 7 shows $\Delta\hat{n}_{N_f}/\Delta N_1$, and $\Delta\hat{b}_{N_f}/\Delta N_1$ for both the 0° and 70° angle conditions and all four perturbation sizes for both observers. We conducted a variety of t -tests on the data to assess the significance of the influence measures. In total, we performed 28 tests per subject and accordingly applied a Bonferroni correction at the 5% significance level. Thus, the null hypothesis was rejected

only if $p < 0.0018$ ($0.05/28$). The results of some of these tests can be found in Table 1. We first tested whether the components $\Delta\hat{b}_{N_f}/\Delta N_1$ and $\Delta\hat{n}_{N_f}/\Delta N_1$ were significantly different from zero for the positive and negative perturbation conditions separately (a total of 16 tests). For all subjects, in all angle, perturbation size and direction conditions the b -component of influence was not significantly different from zero. Turning now to the n -component, for the 0.8 mm perturbation size condition the n -component of influence was significant for only one subject (PAW) in all angle and direction conditions. The other two subjects showed less consistent results which approached but did not achieve significance. When the perturbation size was 1.6 mm, all three observers displayed a significant influence in all conditions. When the positive and negative perturbation direction data sets were combined (thereby doubling the size of the data set in a single t -test, resulting in eight additional tests) we again found that for both perturbation sizes and angle conditions the b -component was not significantly different from zero for all subjects (not shown). However, the n -component of the influence showed significant differences from zero for two subjects (PAW, LTM) in all conditions (Table 1). For subject JT when the perturbation size was 1.6 mm, the n -component of influence was significantly different from zero in both the 0° and 70° conditions. In the 0.8 mm perturbation size conditions the influence approached significance ($p < 0.01$ and $p < 0.05$ for the 0° and 70° conditions, respectively; Table 1). Thus, we conclude that the perturbations had a measurable effect. In addition, we tested whether the n - and b -components of influence were significantly different for the two perturbation sizes (four additional tests). For all three subjects, no significant differences were found (Table 1). We conclude that homogeneity holds over the range investigated (perturbations of -1.6 to $+1.6$ mm).

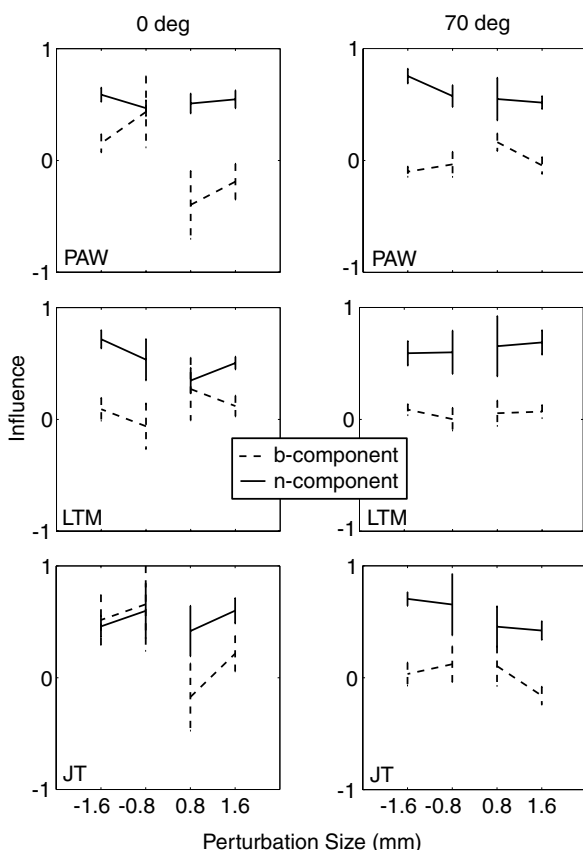


Fig. 7. Experiment 1: Tests of homogeneity. Influence is the ratio of magnitude of effect to magnitude of perturbation. In the linear perturbation region magnitude of effect is proportional to magnitude of perturbation. Thus, influence should be independent of magnitude of perturbation. Measured influence is plotted for three observers for two magnitudes and two signs of perturbation (± 0.8 and ± 1.6 mm). The results are consistent with the claim that these magnitudes of perturbation fall within the linear perturbation region. Data for three observers are shown.

5. Experiment 2

This experiment was designed to test the superposition condition that must be satisfied for local linearity. That is, we test whether the effect on observer settings of perturbing two points equals the sum of the effects resulting from each of the separate perturbations.

5.1. Methods

On each trial observers viewed the eight-point sampled contour and had to adjust p_0 in the setting plane until it lay on the inferred contour. The perturbed points were always p_{-1} and/or p_1 , the two points flanking p_0 . On a given trial, either no points, one point or two points were perturbed by 0.8 mm in the corresponding

Table 1
Tests of homogeneity for Experiment 1

Angle	Observer		
	PAW	LTM	JT
<i>0.8 mm perturbation</i>			
0°	$t_{23} = 8.61, p < 0.001$	$t_{23} = 4.06, p < 0.001$	$t_{23} = 3.03, \text{n.s.}$
70°	$t_{23} = 5.26, p < 0.001$	$t_{23} = 3.78, p < 0.001$	$t_{23} = 3.37, \text{n.s.}$
<i>1.6 mm perturbation</i>			
0°	$t_{23} = 11.16, p < 0.001$	$t_{23} = 12.31, p < 0.001$	$t_{23} = 7.06, p < 0.001$
70°	$t_{23} = 14.56, p < 0.001$	$t_{23} = 8.30, p < 0.001$	$t_{23} = 10.85, p < 0.001$
<i>Difference</i>			
0°	$t_{46} = -1.01, \text{n.s.}$	$t_{46} = -1.21, \text{n.s.}$	$t_{46} = -0.10, \text{n.s.}$
70°	$t_{46} = -0.54, \text{n.s.}$	$t_{46} = -0.06, \text{n.s.}$	$t_{46} = -0.04, \text{n.s.}$

The t -statistics testing whether measured influence on settings in the n -direction of the ± 0.8 and ± 1.6 mm perturbations in the N -direction were significantly different from zero are shown for the three observers in both angle conditions. The degrees of freedom of each t -statistic are shown as a subscript. All were significantly (or close to significantly) non-zero, showing that observers reliably changed their settings in the n -direction in response to perturbations in the N -direction. In the panel labeled “Difference” we compare the magnitude of measured influence in response to the two sizes of perturbation. There were no significant differences and thus we did not reject the homogeneity hypothesis.

normal direction(s) N_{-1} and N_1 . The factors in this experiment were the angle (0° and 70°) and perturbation pair type for the two perturbed points ((0,0), (0,0.8), (0.8,0), (0,-0.8), (0.8,0.8) and (0.8,-0.8)). We doubled the number of (0,0) (unperturbed) conditions to guarantee twice as many unperturbed settings. Thus, there were 14 conditions. Observers completed nine repetitions of each condition over three hour-long sessions leading to a total of 126 settings per subject.

As in the first experiment the angle conditions were interleaved, and all other conditions were randomized.

5.2. Observers

The observers were the same as those used in Experiment 1.

5.3. Results

Once again setting variability was comparable to that reported by Warren et al. (2002). The maximum standard deviation across all three subjects, six perturbation conditions, two rotation conditions and both coordinate directions was 1.39 mm (approximately 4% of the distance to the nearest contour point). As in Experiment 1 this maximum variability was observed in the b -component of the 0° rotation condition. Since this corresponds to the depth dimension the value can be equivalently expressed as approximately 0.6' disparity.

The law of superposition implies that the effect of the perturbation of two points (e.g., perturbation pair type (0.8,0.8)) should be the sum of the effects of the constituent single-point perturbations ((0.8,0) and (0,0.8)). Fig. 8 compares two point perturbations effects with the sum of the constituent single-point effects. This is shown separately for the b - and n -components of the effect (e.g.,

$\Delta \hat{b}_{N_{-1}} + \Delta \hat{b}_{N_1}$) and for the three subjects. Note that we report effect rather than influence as influence is not well defined for two simultaneous perturbations. Effects are calculated using the mean of the perturbed settings in each condition relative to the mean of the unperturbed settings.

We performed a variety of t -tests on the data. In total we performed eight tests per subject—two coordinate directions (n, b) by two angles (0,70) by two double perturbation conditions ((0.8,-0.8) and (0.8,0.8)). We accordingly applied a Bonferroni correction to achieve an overall 5% Type I error rate. Thus, the null hypothesis was rejected only if $p < 0.0063$ (0.05/8).

In Table 2 we show results of these tests of superposition for the n -components of the effect measure. In each condition, we tested whether the effect of the double perturbations was significantly different from the sum of the effects of the corresponding single perturbations. In all conditions for LTM and JT, and all but one condition for PAW (70° rotation with (0.8,-0.8) perturbation), there was no significant difference. Thus, superposition appears to hold in most cases. Taken together with the results of Experiment 1, we conclude that effect is a linear function of perturbation for the range of perturbations employed here.

6. Experiment 3

The purpose of this experiment is to derive estimates $\Delta \hat{n}_{N_i} / \Delta N_i$ and $\Delta \hat{b}_{N_i} / \Delta N_i$ of all components of the influence matrices for all eight points in our stimulus ($i = -4, -3, -2, -1, 1, 2, 3, 4$). In this experiment, we only perturb points in the N_i -direction, for both angle conditions. In Experiment 4, we will estimate the remaining two components of the influence matrices by perturbing

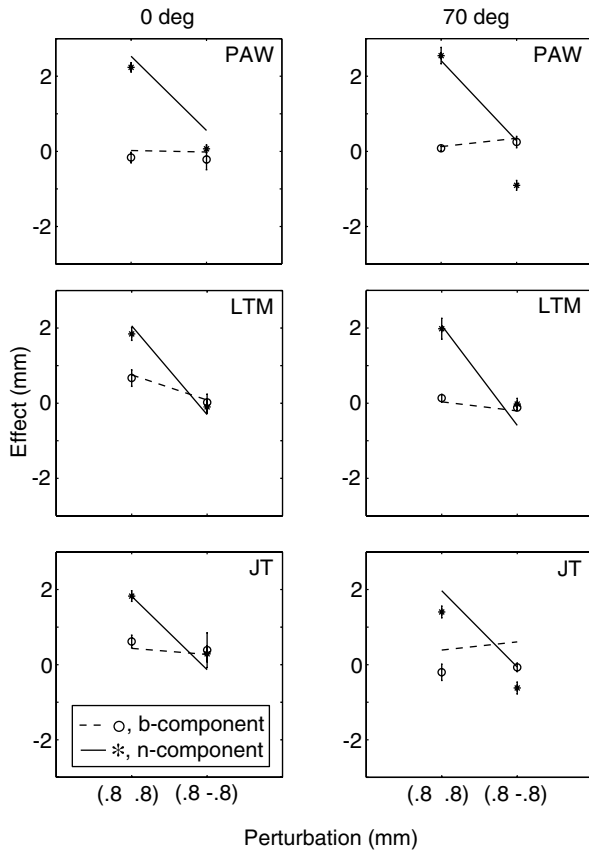


Fig. 8. Experiment 2: Tests of superposition. We have tested the hypothesis that the sum of the effects of two separately performed perturbations is not significantly different from the effect obtained when the perturbations are applied simultaneously. The symbols in each plot correspond to the effects of two perturbations applied simultaneously to the two points immediately flanking the adjustable point ((0.8, 0.8) or (0.8, -0.8)). The end points of the straight lines correspond to the effect obtained when the effects of two corresponding single perturbations are summed (e.g., (0.8, 0) + (0, 0.8)). Results are shown for each of the two angle conditions, the two double perturbations and the two coordinates of the setting plane. If superposition were satisfied exactly then for each of the conditions the symbols should lie on the end points of the corresponding lines. Data for three observers are shown.

in the B_i -direction. If the local spline hypothesis is true, then we expect that both components will decrease rapidly in magnitude with increasing distance from p_0 (i.e., with increasing $|i|$). We also will test dimensional independence, i.e., whether small perturbations in the N_i -direction result in displacements confined to the n -direction, as the results of the last two experiments suggested.

The results of the last two experiments indicate that over the range -1.6 to 1.6 mm the effect of a perturbation is a linear function of the perturbation. We take 0.8 mm as the perturbation magnitude in this and the following experiment, ensuring that we are within this locally linear range.

6.1. Methods

In this experiment, observers saw the eight-point sampled parabolic contour and were asked to adjust p_0 in the setting plane until it lay on the perceived contour. On a given trial any one of the points forming the contour might be perturbed by 0.8 mm in the positive or negative normal direction N_i of the point p_i being perturbed.

The factors in this experiment were angle (0° and 70°), direction of perturbation (+0.8 and -0.8) and perturbed point number (-4, -3, -2, -1, 1, 2, 3, 4). The angle conditions were interleaved and all other conditions were randomized. We ran twice as many unperturbed trials as any single perturbation condition which adds a further four conditions—one for each of the angle and direction conditions. Thus, there were 36 experimental conditions and subjects saw 8 repetitions of each condition over eight 30-min sessions. This leads to a total of 288 settings per subject.

6.2. Observers

The observers were the same as those in Experiments 1 and 2 with the addition of IM who was naive to the purpose of the experiments but was not available to run in Experiments 1 and 2. Thus, we are assuming that the perturbations used here are within the linear range for this observer as they are for the other three.

6.3. Results

As in Experiments 1 and 2 we report maximum setting standard deviations across all eight perturbation conditions, two rotation conditions and both intrinsic coordinate directions (due to the results of Experiment 1 positive and negative perturbations are pooled). However, here we report the value for each observer separately in millimeters and as a percent of the distance to the nearest contour point. For observers PAW, LTM, IM and JT, the maximum standard deviations were: 0.84 mm (2.4%), 0.87 mm (2.5%), 0.78 mm (2.2%), and 1.63 mm (4.7%), respectively. Each of these maximum variability values was observed in the b -component of a 0° rotation condition and can thus be equivalently expressed as a disparity of $0.4'$, $0.4'$, $0.3'$ and $0.7'$, respectively. Consequently, variability was again similar in magnitude to that reported by Warren et al. (2002).

Fig. 9 shows estimates $\Delta \hat{n}_{N_i} / \Delta N_i$ and $\Delta \hat{b}_{N_i} / \Delta N_i$ of the influence measure for each subject. The results of Experiment 1 suggest that we can combine positive and negative perturbation data to estimate influence, and all tests were carried out on these combined data. In total we conducted 32 t -tests on each observer's data set to assess whether the influence of each point was significantly different from zero (i.e., eight points, two influ-

Table 2
Tests of superposition for Experiment 2

Perturbation hypothesis	0°			70°		
	PAW	LTM	JT	PAW	LTM	JT
H0: (0.8, 0.8) = (0.8, 0) + (0, 0.8)	$t_{16} = 1.30$, n.s.	$t_{16} = 0.75$, n.s.	$t_{16} = -0.02$, n.s.	$t_{16} = -0.44$, n.s.	$t_{16} = 0.20$, n.s.	$t_{16} = 2.12$, n.s.
H0: (0.8, -0.8) = (0.8, 0) + (0, -0.8)	$t_{16} = 2.26$, n.s.	$t_{16} = -0.73$, n.s.	$t_{16} = 1.31$, n.s.	$t_{16} = 4.86$, $p < 0.001$	$t_{16} = -1.65$, n.s.	$t_{16} = 2.09$, n.s.

The reported t -statistics test whether the sum of effects (in the n -direction) of two single perturbations is significantly different from the effect of the perturbations when applied simultaneously. Tests were carried out for the three observers in both angle conditions for the two double perturbations tested ((0.8, 0.8) and (0.8, -0.8)). The only significant exception to the Superposition Hypothesis is for observer PAW in the 70°, (0.8, -0.8) condition.

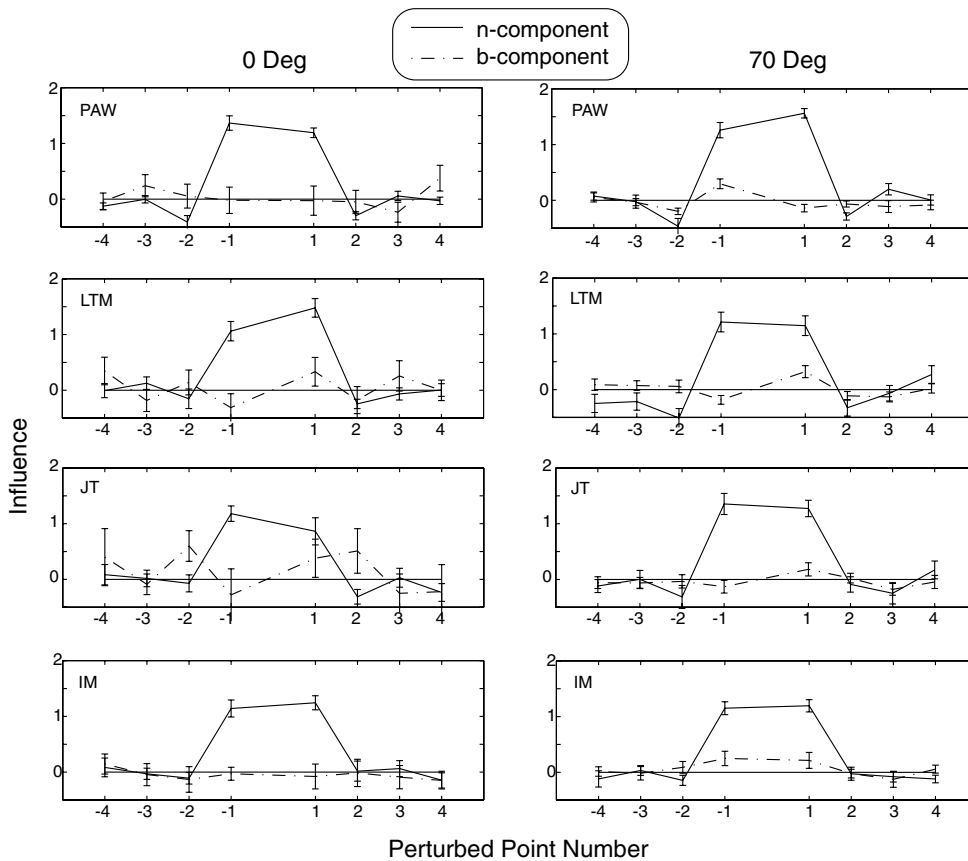


Fig. 9. Experiment 3: Influence functions for perturbations in the N_i -direction. The influence of perturbations of each of the fixed points is shown for four observers. The solid curve is the influence of perturbations of p_i in the N_i -direction on the n -component of the observer's setting, $\Delta n_{N_i}/\Delta N_i$. The dashed curve is the corresponding influence on the b -component of the observer's setting, $\Delta b_{N_i}/\Delta N_i$. Data for four observers are shown.

ence axes, two orientation conditions). A Bonferroni correction was applied and thus significance was achieved at the 5% level only if $p < 0.0016$ ($0.05/32$).

Fig. 9 shows that in both the 0° and 70° conditions the b -component of influence is close to zero for all perturbed points (except perhaps observer JT). This is consistent with dimensional independence. The statistical tests agreed with this observation. For all four subjects there were no significant influences in the

b -direction for any perturbed point in either angle condition.

Turning to the n -component, for all subjects Fig. 9 shows that the influence is close to zero in both angle conditions for all points except p_{-1} , p_1 and perhaps p_{-2} and p_2 . Clearly, we would expect at least three points to exert influence on the setting since it requires at least this many to define a parabola. Statistical tests showed that there was a significant n -component of influence for

points p_{-1} and p_1 in all conditions for all four observers. However, only for observer PAW was there a significant effect of points p_{-2} and p_2 , with the three other observers displaying non-significant influence or influence approaching significance (e.g., point p_{-2} for observer LTM in the 70° condition).

For all but one subject in one condition (point p_2 for observer IM in the 0° condition) the n -component of influence of points p_{-2} and p_2 was negative. This means that, for example, a positive perturbation of point p_2 causes the setting to be displaced away from the average unperturbed setting in the negative direction and vice versa.

These results are similar to those found by Hon et al. (1997) in the fronto-parallel plane. Only the nearest four points contributed to the interpolation setting with the nearest two (p_{-1} , p_1) exerting a positive influence and the next two (p_{-2} , p_2) exerting a negative influence. In addition, our data suggest that influence is not affected by rotating the stimulus into the third dimension. Also, we find evidence of dimensional independence. In the following experiment we look for the opposite pattern of results. We perturb a point in the B_i -direction and expect the setting to be influenced only in the b -direction.

7. Experiment 4

The purpose of this experiment was similar to that of Experiment 3 except that the perturbations were made in the B_i -direction rather than the N_i -direction. Thus we will measure how the estimated 2-D influence vector ($\Delta\hat{n}_{B_i}/\Delta B_i$, $\Delta\hat{b}_{B_i}/\Delta B_i$) varies with position on the contour. We expect to find that the b -component of influence is very small for all but the sample points closest to p_0 (the local spline hypothesis) and to find very small values for the n -component of influence for all perturbed points (the dimensional independence hypothesis).

7.1. Methods

The methods were identical to those of Experiment 3 except that the perturbation was now performed in the B_i -direction at each sample point.

7.2. Observers

The observers were the same as those in Experiment 3.

7.3. Results

As in Experiment 3, for each observer we report the maximum standard deviations of settings across all eight perturbation conditions, two rotation conditions and

both intrinsic coordinate directions (due to the results of Experiment 1, positive and negative perturbations are pooled). We report each value in millimeters and as a percent of the distance to the nearest contour point. For observers PAW, LTM, IM and JT, the maximum standard deviations were: 1.07 mm (3%), 1.04 mm (3%), 0.85 mm (2.4%), and 1.84 mm (5.3%), respectively. Each of these maximum variability values was observed in the b -component of a 0° rotation condition and can thus be equivalently expressed as a disparity of 0.5', 0.5', 0.4' and 0.8', respectively. Consequently, variability was, once again, similar in magnitude to that reported by Warren et al. (2002).

Fig. 10 shows estimates $\Delta\hat{n}_{B_i}/\Delta B_i$ and $\Delta\hat{b}_{B_i}/\Delta B_i$ of the components of the influence matrix for all subjects. Again, 32 t -tests were performed (with Bonferroni correction, and combining positive and negative perturbations) on each observer's data set to assess whether the influence of each point was significantly different from zero. Fig. 10 shows that in both the 0° and 70° conditions the n -component of influence was close to zero for all perturbed points. Again, this is what we would expect in the case of dimensional independence since all perturbations were made in the B -direction. With the exception of point p_1 , for observer PAW in the 70° condition, the statistical tests agreed with this observation; no point had a significant influence on setting in the n -direction. Turning to the b -component, Fig. 10 shows that the influence is close to zero in both angle conditions for all points except p_{-1} , p_1 and perhaps points p_{-2} and p_2 . Again, we would expect at least three points to exert influence on the setting since it requires at least this many to define a parabola. With the exception of observer JT in the 0° condition, statistical tests showed that, for points p_{-1} and p_1 , the b -component of influence was significantly different from zero in all conditions for all observers. No observers displayed, a significant influence of points p_{-2} and p_2 in any condition, however the b -direction influence of observer PAW was nearly significant for points p_{-2} and p_2 in the 70° condition. Finally, as in Experiment 3, the b -component of influence of points p_{-2} and p_2 was always either near zero or negative.

Due to the similarities in the data between the subjects and both angle conditions in Experiments 3 and 4, we collapsed the subject and angle conditions to give a single influence measure in each of the b - and n -directions (Fig. 11, top panels). Tables 3 and 4 shows the results of t -tests to assess whether the influence was significantly different from zero in the b - and n -directions for Experiments 3 and 4, respectively. We conducted 16 tests per experiment and accordingly applied a Bonferroni correction such that the 5% significance level was only achieved if $p < 0.0031$ (0.05/16). For a perturbation in the N -direction, Table 3 shows that no point had a significant influence on the setting in the

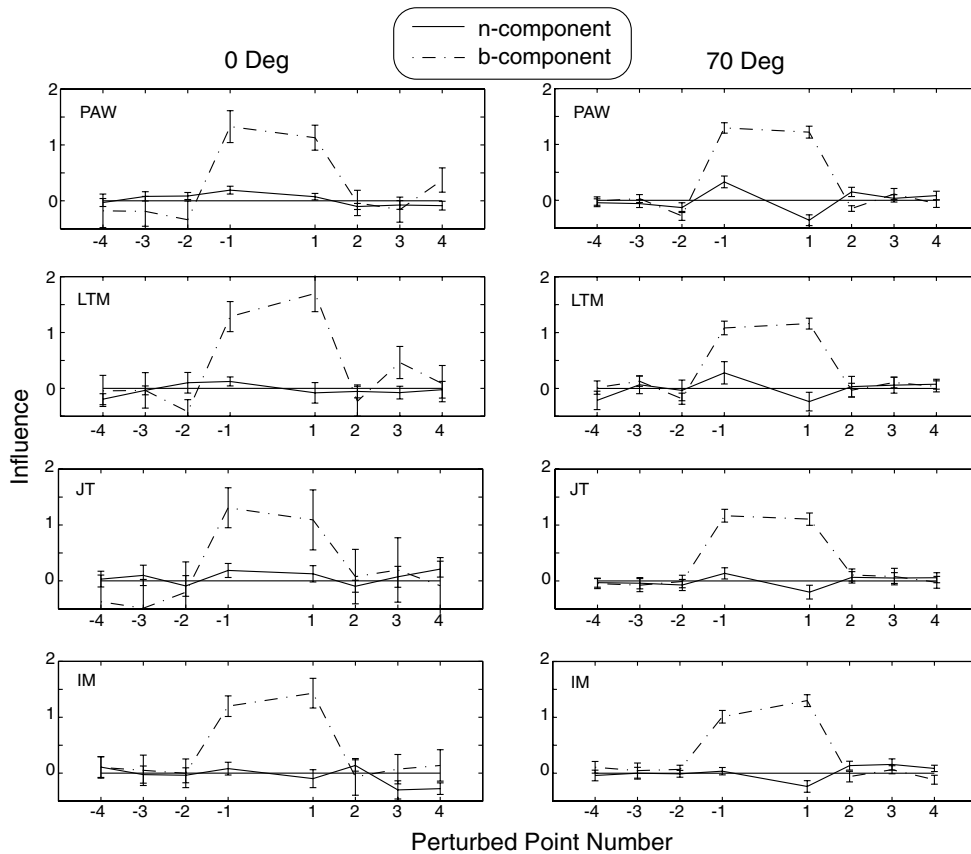


Fig. 10. Experiment 4: Influence functions for perturbations in the B_i -direction. The influence of perturbations of each of the fixed points is shown for four observers. The solid curve is the influence of perturbations of p_i in the B_i -direction on the n -component of the observer's setting, $\Delta n_{B_i}/\Delta B_i$. The dashed curve is the corresponding influence on the b -component of the observer's setting, $\Delta b_{B_i}/\Delta B_i$. Data for four observers are shown.

b -direction, and only perturbation of points p_{-2} , p_{-1} , p_1 and p_2 significantly affected settings in the n -direction. The pattern of results seen in Experiment 4 is less clear. However, there is still evidence of dimensional independence and only the nearest points have an influence on the setting. For a perturbation in the B -direction, Table 4 shows a significant influence on the setting in the n -direction of point p_{-1} and of points p_{-1} and p_1 on settings in the b -direction.

8. General discussion

We have addressed the question of how much of the available information is used in a 3-D interpolation task. Several studies have suggested that when interpolating regions of space, all but the most proximate information is overlooked. We have tested these claims by asking observers to interpolate sampled planar contours in 3-D.

Our experiments replicated some of the conclusions of Hon et al. (1997) while extending those conclusions to non-fronto-parallel contours. The results of Experiment 1 showed that doubling the size of the perturbation

simply doubles the effect on the observer's setting. Experiment 2 demonstrated that the effect of perturbing two points is approximately the sum of the effects of the individual perturbations. Thus, for the range of perturbations tested, interpolation is linear. Experiments 3 and 4 show that influence falls to zero quickly as the perturbed point moves away from the adjustable point p_0 , indicating that limited information is used in carrying out the task. Experiments 3 and 4 extend the work of Hon et al. to 3-D. Dimensional independence held, so that when a point adjacent to p_0 is perturbed, either within or orthogonal to the plane of the contour, the effect on the observer's setting is largely in the same dimension. For example, perturbing point p_{-1} , in the N_{-1} -direction leads to non-zero influence in the n -direction only.

Note that prior psychophysical results suggest a special role for local element orientation in grouping (Field et al., 1993; Geisler, Perry, Super, & Gallogly, 2001). It remains to be seen whether a similar effect is seen for interpolation. A study in which the high contrast positional information found in this study is replaced with oriented elements (e.g., Gabor patches) will form the basis of further research.

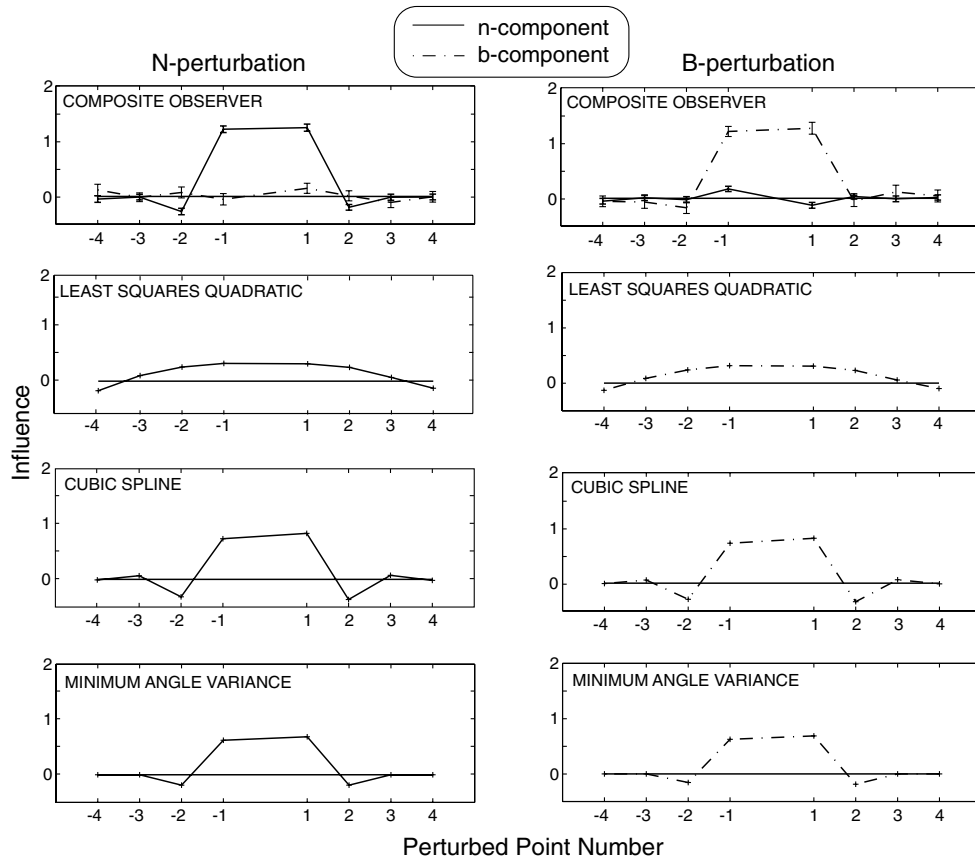


Fig. 11. The composite measured influence matrix and influence matrices for candidate spline algorithms. The influence matrices for three spline algorithms are plotted, together with the averaged human data.

Table 3
Experiment 3: The composite influence function for perturbations in the *N*-direction

Point	<i>n</i> -component	<i>b</i> -component
<i>p</i> ₋₄	$t_{127} = -0.81$, n.s.	$t_{127} = 1.17$, n.s.
<i>p</i> ₋₃	$t_{127} = -0.22$, n.s.	$t_{127} = -0.18$, n.s.
<i>p</i> ₋₂	$t_{127} = -4.76$, $p < 0.001$	$t_{127} = 0.74$, n.s.
<i>p</i> ₋₁	$t_{127} = 20.81$, $p < 0.001$	$t_{127} = -0.50$, n.s.
<i>p</i> ₁	$t_{127} = 19.77$, $p < 0.001$	$t_{127} = 1.65$, n.s.
<i>p</i> ₂	$t_{127} = -3.89$, $p < 0.001$	$t_{127} = 0.13$, n.s.
<i>p</i> ₃	$t_{127} = -0.25$, n.s.	$t_{127} = -1.18$, n.s.
<i>p</i> ₄	$t_{127} = -0.18$, n.s.	$t_{127} = -0.06$, n.s.

The *t*-statistics shown indicate whether influence is significantly non-zero for the *n*- and *b*-directions in response to perturbations in the *N*-direction. These data were obtained by averaging settings over the four observers, two angle conditions and two signs of perturbation. The data are, however, representative of results for the individual observers.

8.1. Piecewise interpolation

The most striking result reported here is that observers use so little of the available information to complete the task. One explanation is that visual interpolation mechanisms are confined to fixed regions of the retina roughly centered on the fovea. This explanation can be tested by changing the size of the stimulus, e.g., by scaling the stimulus size by 0.5. Such a scaling would increase the number of points that fall within the

hypothetical interpolation region. If this explanation were true we would expect to see more points with significant influence. This control experiment was conducted in the fronto-parallel plane by Hon et al. (1997). They found that changing the size of the stimulus did not affect the characteristic shape of the influence functions (i.e., interpolation was *scale-invariant*). Thus, when the stimulus was halved in size, the same number of points had a significant influence on observers' settings, even though some of the points that were not used

Table 4
Experiment 4: The composite influence function for perturbations in the *B*-direction

Point	<i>n</i> -component	<i>b</i> -component
<i>P</i> ₋₄	$t_{127} = -1.03$, n.s.	$t_{127} = -0.57$, n.s.
<i>P</i> ₋₃	$t_{127} = 0.21$, n.s.	$t_{127} = -0.59$, n.s.
<i>P</i> ₋₂	$t_{127} = -0.45$, n.s.	$t_{127} = -1.67$, n.s.
<i>P</i> ₋₁	$t_{127} = 3.59$, $p < 0.001$	$t_{127} = 13.53$, $p < 0.001$
<i>P</i> ₁	$t_{127} = -2.42$, n.s.	$t_{127} = 11.98$, $p < 0.001$
<i>P</i> ₂	$t_{127} = 0.72$, n.s.	$t_{127} = -0.50$, n.s.
<i>P</i> ₃	$t_{127} = -0.18$, n.s.	$t_{127} = 0.98$, n.s.
<i>P</i> ₄	$t_{127} = 0.34$, n.s.	$t_{127} = 0.41$, n.s.

The *t*-statistics shown indicate whether influence is significantly non-zero for the *n*- and *b*-directions in response to perturbations in the *B*-direction. The format is the same as that of Table 3.

in the smaller stimulus were within the interpolation region for the larger stimulus.

In Warren et al. (2002) we replicated this finding using setting variability as a measure of performance. Observers performed a parabolic interpolation task similar to that used here, but the number of points describing the contour was varied. We showed that halving the size of the stimulus led to smaller variability for each number of points, but that the *difference* in variability between these conditions was small. Adding extra points beyond the three or four used for parabolic interpolation did not enhance performance. These control experiments make it unlikely that visual interpolation mechanisms can only operate in a small, central retinal region.

Of course, only three points are actually needed to completely specify the parabola. Thus, an alternative explanation is that the extra points carry little useful information. Care must be taken with this proposal because the neural mechanisms responsible for interpolation do not have access to the exact location of the point, since human estimates of location are unreliable, leading to the variability in the location of the setting point. Thus, one might expect that an optimal strategy for minimizing this setting variability would use location information from all of the points, weighting the contribution of each point by its uncertainty. We tested this hypothesis using a linear sampled contour stimulus in the fronto-parallel plane (Warren et al., 2002). In that study we found evidence that human observers use only two points when interpolating a sampled linear contour. We described an “ideal Gaussian interpolator” model that calculated the contribution of each point relative to a model of 3-D Gaussian uncertainty in its location. The additional information provided by using four points rather than two led to a reduction in setting variability of 30%. We concluded that it is *not* the case that the additional points carry negligible information in the linear task. It is reasonable to assume that additional points in the parabolic case would also contain useful information. Thus, in spite of its great accuracy, human interpolation performance is far from ideal when setting variability is considered.

These considerations suggest that interpolation is a local process in the sense that it relies on only the very few nearest “information samples” from a contour or surface. Note, however, that these nearest samples may be a relatively large distance from the interpolation region.

We have suggested that non-optimal performance may occur as a result of attempting to provide a robust solution to the interpolation problem (Warren et al., 2002). Since there is more uncertainty in the location of more peripheral points, their contribution should be down-weighted. Such a scheme might lead to decreased interpolation reliability unless (a) a precise estimate of internal noise is available and that (b) the information is used optimally. If either (a) or (b) is violated then increased shape estimation errors may result from using more than the most proximate sources of information. Additionally, more distant points are more likely to belong to another object or contour, lending additional impetus to excluding them from the interpolation calculation.

8.2. The human visual spline

Taken together, the results of this paper, Hon et al. (1997) and Warren et al. (2002) provide strong evidence for the proposal of Feldman (1997) that human visual interpolation of contours is very much like a piecewise spline. Thus, the human visual system localizes a sampled contour by computing an estimate of location based on a series of four-point sections. We have so far restricted our attention to parabolic contours, so we cannot conclusively describe the polynomial order of these four-point splines. However, Warren et al. (2002) suggest that any bias away from a parabolic spline is tiny (roughly the width of a few sheets of paper). These results suggests that the family of curves that the human visual system can interpolate without appreciable bias (the “human visual spline”) includes linear and parabolic contours. We do not rule out the possibility that other curve families (circles, cubic polynomials, etc.) could also be part of the “human visual spline”. Tests of such hypotheses will be the goal of future research.

8.3. Comparisons with standard interpolation algorithms

We now compare the measured influence functions of observers with the computed influence functions of three interpolation algorithms: (1) a least-squares fit to a parabola, (2) a standard cubic spline algorithm, and (3) an algorithm that minimizes the variance of angles in the contour (defined over all neighboring triplets of points along the curve). This latter model is an instantiation of the minimum angle variance criterion implicit in the work of Pizlo et al. (1997) and Vos and coworkers (personal communication) (see van Assen & Vos, 1999). For each of these interpolation algorithms, we computed the partial derivatives of the model's setting of the adjustable point p_0 relative to perturbations of the other points, i.e., the influence matrices I_i for perturbations of each point p_i , which constitute each model's predictions for the measured influence functions of human observers.

Since the data for human observers suggests that the "human visual spline" is approximately dimensionally independent, we broke the problem of three-dimensional interpolation into two separate two-dimensional interpolations. For convenience, we let t , n , b denote the intrinsic coordinates at the true interpolation point. Recall that t denotes the unit vector tangent to the curve at the true point, n is perpendicular to t and in the contour plane, and b is perpendicular to both n and t . The first two-dimensional interpolation takes place in the nt -plane which contains the contour and the sampled points. The second two-dimensional interpolation takes place in the bt -plane, orthogonal to the nt -plane. We will interpolate the n -coordinate with respect to the t -coordinate and, separately, the b -coordinate with respect to the t -coordinate. These are precisely the directions in which we measured influence. We will compute the corresponding influence measures for different spline algorithms and compare them to the measured influence functions for human observers.

If dimensional independence were exactly satisfied then the two partial derivatives characterizing $\Delta\hat{n}_{B_i}/\Delta B_i$ and $\Delta\hat{b}_{N_i}/\Delta N_i$ would have expected value 0. We do not plot these data as the algorithms considered all satisfy dimensional independence. Moreover, the three interpolation algorithms are invariant under rotation about the n -axis. Thus, it is legitimate to compare each of them with the results for human observers averaged across angle conditions (Fig. 11, composite observer).

The numerical estimates of the partial derivatives corresponding to $\Delta\hat{n}_{N_i}/\Delta N_i$ and $\Delta\hat{b}_{B_i}/\Delta B_i$ for the least-squares quadratic fit (Davis, 1975, Chapter VIII) are shown in Fig. 11 (least squares quadratic). To compute these values we first fit a quadratic equation to the unperturbed visible points in the 0° condition, and estimated the intersection of the resulting quadratic curve with the setting plane. Then we refit the visible

points with one perturbed in either the N - or B -direction and compute the intersection of the fitted contour with the setting plane. We used this information to compute numerical estimates $\Delta\hat{n}_{N_i}/\Delta N_i$ and $\Delta\hat{b}_{B_i}/\Delta B_i$ of the partial derivatives, just as we did for the human observers.

Warren et al. (2002) found that the interpolation settings of human observers, for unperturbed sampled parabolic contours, fell very close to the parabolic contour. This is also the case for the least-squares quadratic fit and, if we only considered unperturbed interpolation performance, then we could not reject the hypothesis that human observers are simply fitting a parabolic contour to the visible sampled points. When we consider human response to perturbations, however, it is evident that the estimated derivatives for the least-squares quadratic fit are qualitatively different from the human data. Human observers use the available visual information in a different fashion than does the least-squares quadratic algorithm.

The second algorithm considered is a standard cubic spline (de Boor, 1978) computed using the function *csapi* in the MatLab[®] Spline Toolbox (The Mathworks Inc; Hanselman & Littlefield, 1997). This spline routine provides a perfect fit to parabolic data and, like the least-squares algorithm just considered, it can duplicate human performance in unperturbed conditions. The numerical estimates of $\Delta\hat{n}_{N_i}/\Delta N_i$ and $\Delta\hat{b}_{B_i}/\Delta B_i$ are shown in Fig. 11 (cubic spline).

We note first of all that the estimates for the cubic spline data in Fig. 11 are very different from those for the least-squares quadratic fit, and the pattern of results is qualitatively similar to the average human observer data. The effect of perturbing either of the points adjacent to the interpolation point is to move the interpolation point in the same direction in both the n - and b -coordinates. Perturbing the visible points that are two steps removed from the interpolation point leads to a small response in the contrary direction in either n or b . The effect of perturbation rapidly diminishes with separation from the point of interpolation. However, qualitatively, the response of the human observer to perturbation of the adjacent point is about twice as great as the response of the cubic spline algorithm, and the effect of perturbation seems to drop off more rapidly for the human observers.

The last algorithm considered is based on the minimum angle variance criterion of Pizlo et al. (1997). Any three successive points p_{i-1} , p_i , p_{i+1} along the contour define an angle θ_i , $i = -(N-1), \dots, N-1$. The MAV algorithm that we use selects the setting p_0 that minimizes the variance of the $\{\theta_i\}$. For example, if the points p_i are equally spaced, then the angle variance is zero when the points lie on a segment of a circle or a straight line. Intuitively, the MAV criterion favors near-circular contours and the limiting case of a straight line. As Pizlo et al. note, angle variance and, therefore, the setting

selected by the MAV algorithm are size-invariant (as are the other two algorithms considered here). The numerical estimates of $\Delta\hat{n}_{N_i}/\Delta N_i$ and $\Delta\hat{b}_{B_i}/\Delta B_i$ for the MAV algorithm are shown in Fig. 11 (minimum angle variance). They are in qualitative agreement with the average human data but, as was the case with the cubic spline algorithm, the human observer is considerably more responsive to perturbations of the points adjacent to the setting point p_0 than the algorithm.

We conclude by noting the interpolation task requires that the visual system gather discrete items of information sparsely distributed across a large part of the visual field. For the sampled contours considered here and by Warren et al. (2002), human performance is very accurate both in a fronto-parallel plane and in general position. The large separations between adjacent sample points in a contour, and the three-dimensional character of the task both pose challenges to simple models of cortical interpolation that presuppose interactions between retinotopically adjacent neurons (Polat & Sagi, 1993).⁷ We do not know what classes of contours the human observer can interpolate without bias, but this class includes parabolic and linear contours at a minimum. Consequently, any model of human visual interpolation must reproduce this aspect of human performance in three-dimensional space.

An accurate computational model of human interpolation performance is likely to advance our understanding of long-range interactions in cortex (Kapadia, Ito, Gilbert, & Westheimer, 1995; Knierim & van Essen, 1992; Levitt & Lund, 1997; Sillito, Grieve, Jones, Cudeiro, & Davis, 1995), and it is natural to begin with the large mathematical literature on interpolation (see Davis, 1975) in modeling human performance. The perturbation measures advanced here are useful in rejecting classes of models of human visual interpolation.

One of the most obvious candidates considered here (a least-squares quadratic fit) can be rejected as a model of human interpolation performance based on perturbation analyses, even though it reproduces human performance in unperturbed conditions. The other two models considered are in qualitative agreement with our results, but human observers are more affected by perturbation of adjacent points than either of them, and may show faster “damping” of the effect of perturbation with separation from the setting point than the cubic spline algorithm. In searching for the human visual spline, then, we are seeking an algorithm that is less “springy” in its response than the cubic spline and that accurately interpolates parabolas and, of course, lines.

⁷ We do not completely rule out a contribution of the long range interaction model to the interpolation mechanism. However, a first step in testing this would require an experiment linking the tasks of detection and localization.

The approximate dimensional independence of human visual interpolation is also a constraint on possible models.

In this article and in Warren et al. (2002), we have considered only parabolic and linear contours, confined to frontal and slanted planes in space. It would, therefore, be of interest to look at other classes of curves, both polynomial and non-polynomial, and to consider sampled contours in space that have significant torsion, no longer confined to a single plane.

Acknowledgements

This research was supported by Grant EY08266 from the National Institutes of Health and by Human Frontiers Science Program grant RG0109/1999-B.

References

- Davis, P. J. (1975). *Interpolation & approximation*. New York: Dover.
- de Boor, C. (1978). *A practical guide to splines*. New York: Springer-Verlag.
- Feldman, J. (1997). Curvilinearity covariance, and regularity in perceptual groups. *Vision Research*, 37, 2835–2848.
- Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Research*, 33, 173–193.
- Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41, 711–724.
- Georgeson, M. A., Freeman, T. C., & Scott-Samuel, N. E. (1996). Subpixel accuracy: Psychophysical validation of an algorithm for fine positioning and movement of dots on visual displays. *Vision Research*, 36, 605–612.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., & Stahel, W. A. (1986). *Robust statistics: The approach based on influence functions*. New York: Wiley.
- Hanselman, D. C., & Littlefield, B. R. (1997). *Mastering MATLAB 5: A comprehensive tutorial and reference*. Upper Saddle River, NJ: Prentice-Hall.
- Hon, A. K., Maloney, L. T., & Landy, M. S. (1997). The influence function for visual interpolation. In B. E. Rogowitz & T. N. Pappas (Eds.), *Human Vision and Electronic Engineering II, Proceedings of the SPIE*, 3016 (pp. 409–419).
- Kanizsa, G. (1979). *Organization in vision*. New York: Praeger.
- Kapadia, M. K., Ito, M., Gilbert, C. D., & Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in VI of alert monkeys. *Neuron*, 15, 843–856.
- Klein, S. A., & Levi, D. M. (1987). Position sense of the peripheral retina. *Journal of the Optical Society of America A*, 4, 1543–1553.
- Knierim, J. J., & Van Essen, D. C. (1992). Neuronal responses to static texture patterns in area VI of the alert macaque monkey. *Journal of Neurophysiology*, 67, 961–980.
- Koh, K., & Maloney, L. T. (1988). Visual interpolation and extrapolation of linear and quadratic contours through discrete points. *Investigative Ophthalmology and Visual Science*, 29, 408.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, 35, 389–412.

- Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature*, 387, 73–76.
- Maloney, L. T., & Landy, M. S. (1989). A statistical framework for robust fusion of depth information. In W. A. Pearlman (Ed.), *Visual Communications and Image Processing IV, Proceedings of the SPIE*, 1199 (pp. 1154–1163).
- Pizlo, Z., Salach-Golyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research*, 37, 1217–1241.
- Polat, U., & Sagi, D. (1993). Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Research*, 33, 993–999.
- Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research*, 36, 3037–3050.
- Sillito, A. M., Grieve, K. L., Jones, H. E., Cudeiro, J., & Davis, J. (1995). Visual cortical mechanisms detecting focal orientation discontinuities. *Nature*, 378, 492–496.
- van Assen, M. A., & Vos, P. G. (1999). Evidence for curvilinear interpolation from dot alignment judgements. *Vision Research*, 39, 4378–4392.
- Warren, P. A., Maloney, L. T., & Landy, M. S. (2002). Interpolating sampled contours in 3D: Analyses of variability and bias. *Vision Research*, 42, 2431–2446.