PERGAMON

Available online at www.sciencedirect.com
SCIENCE DIRECT®

Vision Research 43 (2003) 2649–2657

Vision Research

www.elsevier.com/locate/visres

Peak localization of sparsely sampled luminance patterns is based on interpolated 3D surface representation

Lora T. Likova *, Christopher W. Tyler

Smith-Kettlewell Eye Research Institute, 2318 Fillmore St., San Francisco, CA 94115, USA

Received 6 May 2002; received in revised form 24 September 2002

Abstract

Objects in the world are typically defined by contours and local features separated by extended featureless regions. Sparsely sampled profiles were therefore used to evaluate the cues involved in localizing objects defined by such separated features (as opposed to typical Vernier acuity or other line-based localization tasks). Objects, in the form of Gaussian blobs, were defined at the sample positions by luminance cues, binocular disparity cues or both together. Remarkably, the luminance information in the sampled profiles was unable to support localization for objects requiring interpolation when the perceived depth from the luminance cue was cancelled by a disparity cue. Disparity cues, on the other hand, improved localization substantially over that for luminance cues alone. These data indicate that it is only through the interpolated depth representation that the position of the sampled object can be recognized. The dominance of a depth representation in the performance of such tasks shows that the depth information is not just an overlay to the 2D sketch of the positional information, but a core process that must be completed before the position of the object can be recognized.

Keywords: Position; Localization; Sampling; Object; 3D; Stereoscopic

1. Introduction

The ability to interact with the environment, would be impossible if the perceptual system did not address the questions of both "what" and "where". The answers to these basic questions intersect in the issue of the visual localization of objects, which are typically defined by contours and local features separated by extended featureless regions. We have only to think of faces, household appliances or cookware to realize that much of the shape of the object is undefined except at edge regions. Localization is particularly difficult under conditions where the objects could be considered as "sampled" by overlapping noise or partial occlusion—the tiger behind the trees, the face behind the window curtain.

Nevertheless, the visual system allows remarkably precise localization even when the stimuli have poorly defined features and edges (Toet & Koenderink, 1988). Furthermore, sample spacing is a critical parameter for an adequate theory of localization. Specifically, interpolation becomes much poorer than would be expected from linear filtering beyond about 3′ sampling for static images (Morgan & Watt, 1982), or about 5′ for moving images (Morgan & Watt, 1983). Thus, low-level filter integration could not account for interpolation behavior with the 16′ sampling used in our study, or the 48–64′ sampling of Kontsevich and Tyler (1998). Conversely, accuracy of localization by humans is almost independent of the sample spacing. For samples spaced from 30 to 3 min apart, localization is not improved by increasing sample density (Kontsevich & Tyler, 1998). This limitation poses an additional challenge with regard to the localization task, raising the 'long-range interpolation problem' that has generated much recent interest in relation to the position coding for extended stimuli, such as Gaussian blobs and Gabor patches (Hess & Holliday, 1992; Kontsevich & Tyler, 1998; Levi, Klein, & Wang, 1994; Morgan & Watt, 1982).

How does the brain solve the task of feature localization within its network of neuronal representation? And how does it overcome the long-range interpolation problem in particular? We addressed these questions in a position task requiring localization of the peak of (i) a

* Corresponding author.

0042-6989/$ - see front matter © 2002 Published by Elsevier Ltd.
doi:10.1016/S0042-6989(02)00575-8

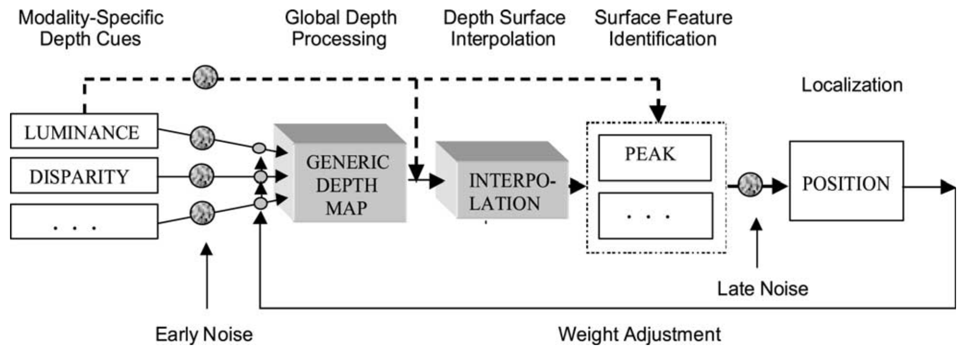

Fig. 1. A model of localization processing based on separate or unitary interpolation processes.

sampled Gaussian luminance profile, (ii) the corresponding Gaussian profile defined by binocular disparity of the samples, or (iii) the sampled Gaussian profile defined by combination of both luminance and disparity (see Fig. 1). In each case, the location of the set of samples was randomized from trial to trial so that the sample locations themselves provided no information as to the location of the profile. The observers' task was to localize the peak of the Gaussian profile regardless of the (randomized) sample positions. Except coincidentally, the peak was never represented in either the luminance or depth information in the samples. We chose this sampled peak-localization task because localization accuracy when the peak position is not present in the stimulus poses a major problem for both visual localization by the "where" mechanism and long-range interpolation mechanisms interpreting "what" stimulus is present.

2. Conceptual framework

Localization information is available from multiple visual cues, as diagrammed in Fig. 1. Position information is available from luminance form, disparity profile, color, texture and other visual cues. Localization in the sampled stimulus might employ interpolation over many such cues.

Three alternative hypotheses of the role of long-range interpolation of a sampled profile can be formulated:

(1) Long-range interpolation plays no role in localization. On this hypothesis, localization of interpolated profiles is determined by the localization of the sample identified as being nearest the peak of the profile. The localization would consequently be limited to the sample positions. Such a strategy would also require a vestige of a long-range process in order to identify which sample was the relevant one for the task, but it will be assumed that this identification could be achieved by attentional comparisons rather than by a continuous field of long-range interactions.

(2) Long-range interpolation is accessible for localization separately in each visual modality (Fig. 1). On this hypothesis, there is some form of interpolation mechanism available within each visual modality to generate information about the form of the profile between the sample positions (Kontsevich & Tyler, 1998). This interpolated information may then be used to define the location of the peak of the profile to greater accuracy than the locations of the samples. In particular, this improved accuracy would be available for both luminance and disparity information.

(3) Long-range interpolation is restricted to one visual modality and feeds into a subsequent unitary localization system (Fig. 1, excluding the dashed line). In this case, only stimuli with access to that modality would support accurate localization, while stimuli providing information that did not reach the site of interpolation would provide poor localization performance.

In our luminance/disparity task, the basic sources of noise determining the localization error are (i) early noise in each visual modality contributing to the position determination, (ii) late noise in the peak localization process (Fig. 1). If localization is performed in separate visual modalities, the position thresholds might be expected to combine according to their absolute signal/noise ratios, assuming that the signals from separate visual modalities have independent noise sources. The observers would be able to interpolate one estimate of the position of the profile from the luminance information alone and a second estimate from the disparity information alone. In this case, signals from the various modalities $(L, D, \ldots, X)$ would combine to improve the localization performance. Adding information about the Gaussian profile from a second modality would always improve detectability and could never degrade it.

On the other hand, the alternative suggested by our observers' reports and by the pattern of the data is that the peak position was judged from interpretation of the input information as a *depth* profile. Even in the absence of disparity information, the luminance was seen as an interpolated Gaussian profile protruding from the background. Thus the luminance information alone was sufficient to evoke an unambiguous depth percept of the brighter bars appearing closer for all observers. These observations suggested the alternative view that the long-range position task was accomplished by the sensory information (both luminance and disparity) being fed into a *unitary depth map*. The local cues in this map of depths would then be subject to a depth surface interpolation process to identify the peak of the Gaussian depth profile.

The important property of the unitary depth map interpretation is that the cues from the separate input modalities could oppose to each other, so that a particular combination of cues would result in a *cancellation* of the depth cues to result in a flat profile with no position information. Thus, the forward bulge of the luminance cues could be opposed by a backward trough from disparity cues. Such cancellation could not happen in a system of independent depth cues, but requires that the cues feed into a common stage with sign-preserving signals. This behaviour provides a strong test for the operation of a unitary depth interpolation mechanism in the task of long-range localization of sampled object profiles.

3. Methods

The key to evaluating the luminance versus the disparity cues to depth was to use a sparsely sampled array. The luminance of the samples carried the luminance profile information while the disparity in their positions in the two eyes carried the disparity profile information. In this way, the two separate depth cues could be combined or segregated as needed. Both luminance and disparity profiles were identical Gaussians, and the two types of profile were always congruent in both peak position and width. The observer's task was to make a left/right judgement on each trial of the position of the jointly specified Gaussian bulge relative to a reference line, using whatever cues were available for the task. Position thresholds were measured for peak disparities of 0, ±0.25, ±0.5, ±1, ±2, ±4 and ±8′. Threshold performance ($d' = 1$) was determined by means of the Ψ staircase procedure (Kontsevich & Tyler, 1999) limited to 60 presentations. Each plotted point derives from the average of at least three staircase measures.

Observers were presented the sampled Gaussian profiles defined either by luminance modulation alone (Fig. 2A), by disparity alone (Fig. 2B), or by combination of luminance and disparity defining a single Gaussian profile (Fig. 2C), It should be noted that the luminance profile evokes a strong sense of depth as the luminance fades into the black background. An unambiguous convex profile was reported by all observers (of a sample of more than 100) in accordance with the diffuse illumination heuristic proposed by Tyler (1997). This depth interpretation should be evident in the printed panels of Fig. 2A and C, and it was certainly seen clearly on the monitor screens. Free fusion of Fig. 2B allows perception of the stereoscopic depth profile (forward for crossed fusion). The third panel shows a combination of both cues at the level that produced cancellation to flat plane under the experimental conditions.

Stimulus samples were vertical 1′ bars separated by uniform gaps of 15′, on a dim background. Bar luminance at the peak of the Gaussian was 20 cd/m^2 and at its minimum was 6 cd/m^2. The Michelson contrast of the luminance profile was 50% (Weber contrast of 185%). The Gaussian depth profile had a width at half-height of 1 deg. The disparity of the reference was offset from the position of the peak of the Gaussian by a randomly varying amount (to avoid depth matching). The overall position of the configuration was randomly jittered by up to 16′ (a full cycle of the sampling interval) on each trial to avoid the judgement being based on position relative to the fixation marker. Each plotted point is an average of at least three staircase threshold estimates.

4. Analysis

Sampled patterns of both luminance and disparity evoke smooth surface interpolation over long distances between the samples (as may be seen by free fusion of the examples in Fig. 2). These observations suggest also the involvement of a unified interpolation process operating at the level of a generic depth representation, and feeding a subsequent unitary localization system. Such unitary depth interpolation would be contrary to hypothesis 2, that interpolation is specific to the processing of separate depth cues.

To evaluate the depth-interpolation hypothesis and the hypothesis that the luminance patterns might be interpreted as an object during localization, we modeled the titration of the depth profile against the perceived depth of luminance profile to reveal the point of cancellation of the depth information. In this model, thresholds are determined by the signal/noise ratio within the luminance (L/σ_L) and disparity (D/σ_D) modalities, together with a source of late noise in the localization of the peak of the Gaussian (σ_P). If localization was performed in separate visual modalities, the noises in each modality would be expected to be statistically independent and therefore to combine according to the Euclidean norm of the signal/noise ratios (Eq. (1)). Thus, the equation for the overall position signal (Θ_P) is simply

$$\Theta_P = [(\sigma_L/L)^2 + (\sigma_D/D)^2 + \sigma_P^2]^{1/2} = 1 \quad (1)$$

Threshold would be defined as a value of unity for the overall position signal. For positions of mixed

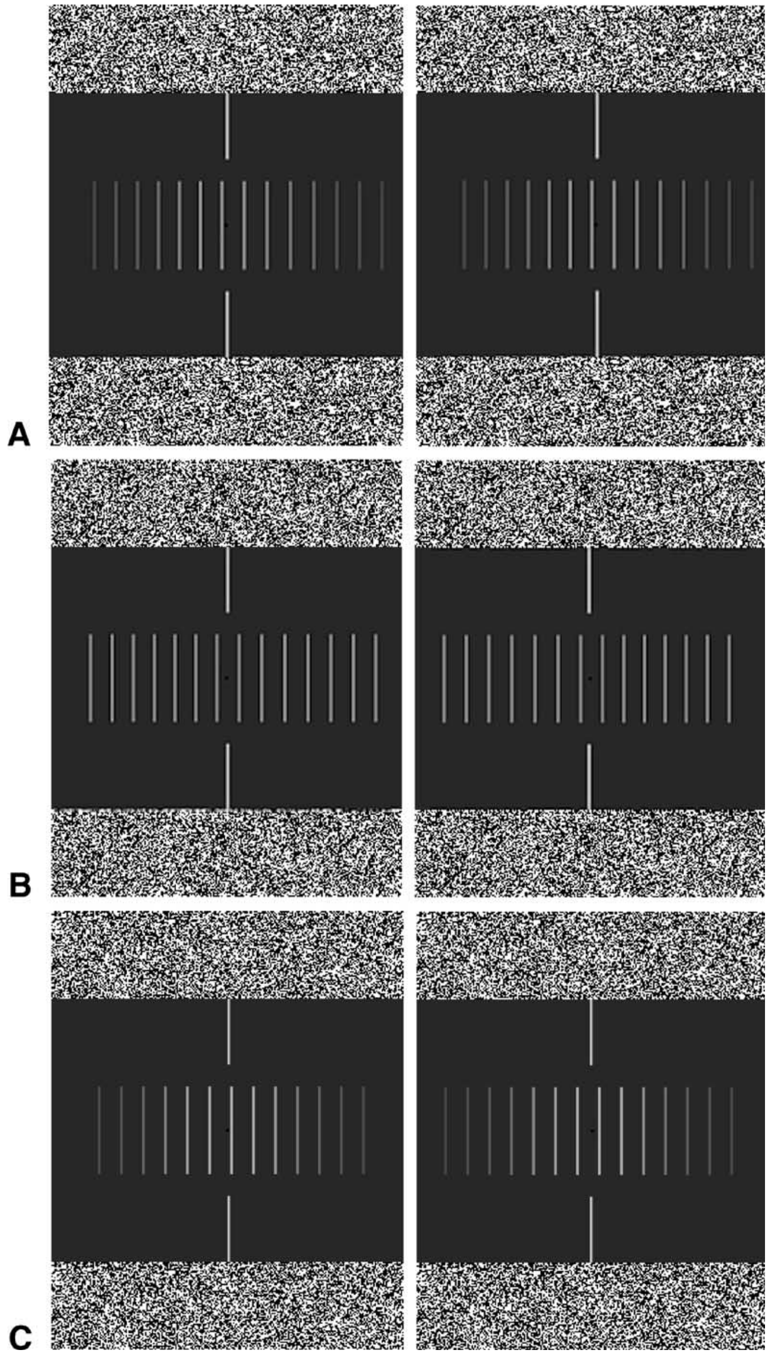

Fig. 2. Examples of the sampled Gaussian profiles used in our experiments, defined by (A) luminance alone, (B) disparity alone, and (C) a combination of luminance and disparity.

luminance and disparity profiles, the position error (as a function of disparity) would take the form shown in Fig. 3A. For a pure disparity profile, the threshold curve would rise with the reciprocal of disparity towards the point of zero disparity (dashed gray curve in Fig. 3A). For a pure luminance profile, the threshold would be independent of disparity (horizontal line in Fig. 3A). For combined disparity and luminance profiles (solid curve in Fig. 3A), the thresholds would always be as good as either cue, improving to better than each cue by $\sqrt{2}$ at the point where the two separate cues have equal effect (point of intersection of the dashed curve and the horizontal line in Fig. 3A).

If interpolated localization is restricted to a pure depth interpolation mechanism, on the other hand, the luminance cue will contribute only by virtue of the *depth* it evokes. In this case, luminance and disparity both feeding the depth module (symbolized by the lower case

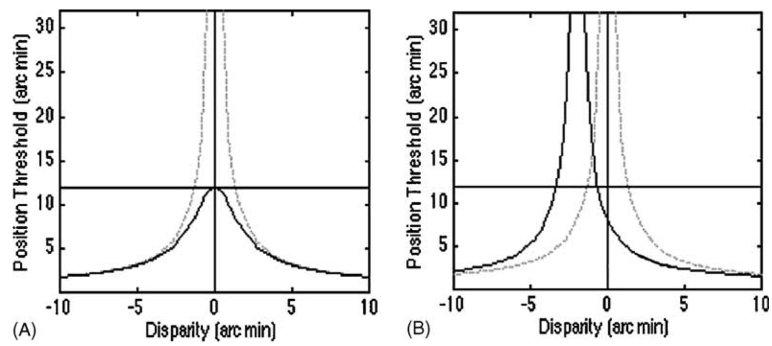

Fig. 3. Predictions for localization performance in the case of a sampled pure disparity profile (dashed grey curve) a pure luminance profile (horizontal line) or a joint disparity/luminance profile (solid curve), as a function of disparity based on the Euclidean norm for the distributed localization model (A) and based on a unified localization model, restricted to a depth interpolation process (B).

d) will be directly summated to form the multiplier of the combined depth noise σ_d (Eq. (2)), and will thus open the opportunity for cancellation between them:

$$\Theta_P = [(\sigma_d/d)^2 + \sigma_P^2]^{1/2} \tag{2}$$

where $d = D + kL$ and $\sigma_d = [\sigma_D^2 + \sigma_L^2]^{1/2}$.

According to Eq. (2), the perceived depth could be canceled when the depth signal d from the disparity profile was equal and opposite to the (protruding) depth signal kL evoked by the luminance profile. Since the position is provided only by the depth signal in this unitary depth interpolation model, localization would become impossible at the cancellation point.

The quantitative prediction of the hypothesis of unitary depth interpolation is shown in Fig. 3B. The disparity only and luminance only curves are naturally identical with that of the first model. However, in the combined-cue case (solid curve) the luminance cue now acts as a fixed depth signal, shifting the entire curve to the left and introducing a cancellation where the depth from disparity is equal and opposite to the depth from luminance information. Note that, at its optimal point (the right intersection between the pure disparity curve and the pure luminance line), the resulting shift of the localization curve allows much greater threshold reduction by the reinforcing combination of the two cues than can occur in the Euclidean norm of Eq. (1), which has a minimum value of 0.71. It should be emphasized that this unitary depth interpolation model (Fig. 3B) incorporates no neural position information from the luminance per se, so when the depth is at cancellation there is no interpolation signal for position discrimination despite the presence of the luminance information in the stimulus.

5. Results

Results of the position localization task for four observers with normal or corrected-to-normal vision are shown in Fig. 4. Their s.e.m. values were 11%–14% of the mean threshold values. The model of Eq. (2) was fit simultaneously to the full set of data by the method of least squares. The average deviation from the model fit was 2.8′, or about 10% of the rage of the measured data. Even this value was heavily increased by one observer, so the other three had much closer fits to the model, of the order of 1′ per point. The parameter values were: $\sigma_D = 3.32 \pm 0.61$, $\sigma_L = 0.59 \pm 0.20$, $k = 1.19 \pm 0.57$, $\sigma_P = 1.98 \pm 0.29$, where the error terms are confidence intervals for the parameter values of the fitted curves at $p < 0.01$, estimated by the Jackknife statistic over the six test conditions (Miller, 1974). Note that the standard errors are very small, of the order of 10% of the mean value, implying that the parameter values of the fits were very similar for the four observers and that each term in the equation was a necessary component of the model.

The results of Fig. 4 show a point of complete cancellation in the combined stimuli, in full support of the model that form interpolation could be performed only at the level of the pure depth interpretation. Specifically:

(A) Localization from disparity alone (Fig. 4, grey curve) was much more accurate than from luminance alone (Fig. 4, horizontal line), immediately suggesting that depth processing plays an important role in the localization of sampled stimuli. Localization accuracy from disparity alone was as fine as 1′–2′, requiring accurate interpolation to localize the peak of the function between the samples spaced 16′ apart. This performance contrasted with the localization accuracy for pure luminance profiles, which was as much as a log unit greater, at about 12′. The fit of the quantitative model indicates that interpolated disparity localization accuracy improved directly with disparity until reaching an almost constant level at higher disparities. This fit is in accord with the idea implied in the equations that localization was limited by the signal/noise ratio in the disparity signal until it became limited by the source of late position noise.

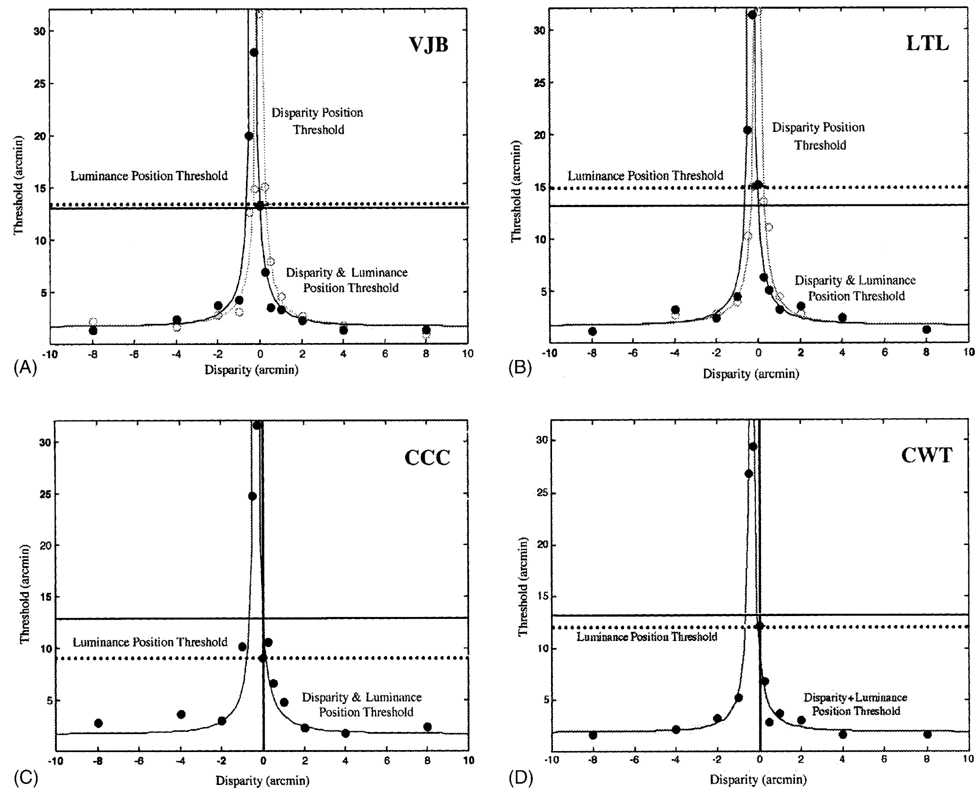

Fig. 4. Results of the position localization task for two principal observers, with the key condition verified in another two observers. The open grey circles are the thresholds for the profile defined only by disparity, the filled black circles are the thresholds defined by disparity and luminance. The grey curve shows the model fit for disparity alone, the black curve that for combined disparity and luminance (fitted with the same parameters to all four observers). The horizontal dashed line shows the threshold for the pure luminance, the horizontal solid line shows the model fit for luminance alone. Note leftward shift of the null point in the combined luminance/disparity function.

(B) Combining identical luminance and disparity Gaussian profiles (Fig. 4, black curve) provides a localization performance that is qualitatively similar to that given by disparity alone (Fig. 4, grey curve). Rather than showing the hump-shaped function predicted by the multiple-cue interpolation hypothesis (Fig. 3A), it exhibits the null behaviour of extreme threshold elevation depicted in Fig. 3B, where localization is impossible within the range measurable in the apparatus. Contrary to the multiple-cue hypothesis, combining cues *degrades* performance by nulling perceived depth when the information is conflicting (similar magnitude, but opposite depth sign). Thus, the stimulus with full luminance information becomes impossible to localize as soon as it is perceived as a flat surface (within the upper limit of measurement, i.e., 32′). Nulling the perceived depth abolished the ability to localize the peak of the profile, despite the fact that the luminance information was clearly visible. The occurrence of this null point must mean that luminance information per se is insufficient to specify the position of the luminance profile in this sampled stimulus.

(C) The model fit to the data reveals that the curve for the combined luminance and disparity cues is shifted leftwards relative to the symmetric curve for localization by pure disparity cues. The fitted curve was shifted to a negative disparity of about 0.4′ (black curves in Fig. 4). The fact that a uniform shift accounts for the full set of data implies that the *only* factor introduced into the position task by the additional luminance information was a positive depth signal of about 0.4′. Once this depth had been cancelled by the corresponding negative disparity, all other aspects of the position task fell on the same curve with no change in parameter values.

(D) To validate this shifted null point, observers were asked to perform an auxiliary experiment of discriminating whether the bars bulged forward or backward as the amount of disparity in the Gaussian profile was varied, again using the Ψ staircase procedure (Kontsevich & Tyler, 1999). The null point for perceived flatness did not occur at zero disparity, but at disparities of −0.4′, −0.3′, −0.4′ and −0.35′ for the four observers in Fig. 4, respectively. This null point validates that the perceived depth from the luminance profile lies in the same qualitative dimension as the perceived depth from disparity cues (i.e., that it is a "true" depth percept rather than just a cognitive inference of some kind). The equivalent disparities are small but reliable. It will be seen below that they are also predictive of the properties of localization performance.

At the null point, the array of bars had the appearance of an approximately flat plane, although small second-order deviations remained. This lack of an accurate null may be expected if there are any nonlinearities in the depth function for either the luminance or the disparity domains. Nonlinearities of perceived brightness in relation to physical luminance are well known (Tyler & Liu, 1996), so a nonlinearity is to be expected here. Defining the exact nonlinearity to achieve a perfect null is beyond the scope of the present project but could be attempted in future studies. For the present work, therefore, the null was defined as the point where the center of the bulge appeared at the same depth as the flanks, even if minor wrinkles may be seen in the transition regions. Note, however, that any residual cue did not allow the observers to resolve the position to better than 32′, beyond the range measurable in our apparatus.

(E) In the region where the luminance and depth cues provided similar position information (at the rightmost intersection of the horizontal line and the grey curve in Fig. 4), the improvement in localization accuracy from the combined cues (black curve) is much greater than the factor of $\sqrt{2}$ predicted from the independent-cue hypothesis, averaging about a factor of 3. This improvement is the obverse of the degradation seen at the point of cancellation, and is a manifestation of the leftward shift of the predicted localization function in the unitary depth-cue model.

(F) Quantitatively, the hypothesis of unitary depth interpolation implies that, at the value of the nulling disparity, the position threshold for disparity should be the same as it is for luminance. To a good approximation, the position thresholds support this prediction: 13.1′ for the nulling disparity alone vs 13.0′ for luminance alone for VJB; and 10.1′ vs 15.3′ for LTL. (Note, the nulling disparity predictions are not shown in Fig. 4.) Thus, when perceived depth is equated for the two cues, the interpolated position thresholds are also equated (within experimental error). This correspondence is one more line of evidence that position thresholds depend on interpolation of the 3D profile of the Gaussian bulge.

6. Discussion

Perhaps the most startling aspect of the results is that position discrimination in sampled profiles can be completely nulled by the addition of a slight disparity profile. It should be emphasized that the position information from disparity was identical to the position information from luminance on each trial, so addition of the second cue would be expected to *reinforce* the ability to discriminate position if the two cues were processed independently. Instead, the *nulling* of the luminance-based position information by the *depth* signal implies that the luminance target is processed exclusively through the depth interpretation. Once the depth interpretation is nulled by the disparity signal, the luminance information does not support position discrimination at all.

This evidence implies that *depth surface reconstruction* is the key process in the accuracy of the interpolated localization profile from both depth and luminance signals. This interpretation is consistent with previous work on interpolation processes within the disparity domain (Würger & Landy, 1989; Yang & Blake, 1995), although they did not generalize their conclusions to other depth cues. It appears that visual patterns defined by different depth cues are interpreted as objects in the process of determining their location. Buckley, Frisby, and Mayhew (1989) also showed that diverse depth cues were integrated in the global depth interpretation of object shape, but they did not measure the role of the separate cues in the process of feature localization. For our data, only an interpolation mechanism operating at the level of generic *depth* representation can account for the results.

Specifically, only depth interpolation accounts for the impossibility of position discrimination at the cancellation point and the asymmetric shift of the cancellation point by the luminance cue. The sampled stimuli were designed so that the peak, which is missing in the sampled physical stimulus, could be identified only from information in an interpolated representation. Inability to interpolate the profile of the Gaussian waveform would restrict observers to localizing on the basis of the 16′ sample points alone, resulting in a localization accuracy of about half that value, or 8′. The fine resolution of the performance when disparity information is present (1′–2′) clearly implies that an interpolation process is involved in the performance. The six properties of the data described in Section 5 clearly indicate that such long-range localization is restricted to the interpolated depth representation of visual image information. Understanding these processes thus provides new insight into the operation of the human visual system, but may also suggest important directions for the development of artificial intelligence systems.

In terms of the noise sources limiting performance in the long-range localization task, localization accuracy can be explained with only *additive* sources of noise in the disparity cue and a small contribution of interpolation noise. The model fit indicates that improvement in performance was linear with disparity until approaching the absolute limit, whereas if there had been some source of disparity-dependent noise, performance would have been expected to improve more gradually. We therefore conclude that disparity processing is limited by local noise in the monocular signals forming its input, which would be independent of the relative position, or disparity, between the two eyes, whereas noise arising in the pathway after disparity had been computed would be likely to depend on disparity.

Objects in the world are typically defined by contours and local features separated by featureless regions (such as the design printed on a beach ball, or the smooth skin between facial features). Leonardo's (1509) depiction of a dodecahedron (Fig. 5) illustrates the point. The facets extending between the edges are perceptually vivid, and yet their locations are not defined by any features in the image. The shading does not define these perceived surfaces, because it is not homogeneous although each surface is perceived as flat. (The inhomogeneity of the shading is interpreted as the painter's brush strokes lying in the surface defined by the edges alone.) The mean differences among the shading fields on different facets are interpreted as consistent with the angles of the surfaces, helping to support the 3D interpretation, but the surfaces themselves are *interpolated* from the locations of the edges without regard to the details of the shading. Thus, the facets are treated as featureless regions extending between the edge locations, and any features that are actually found there are attributed to a separate source, the paint on the page.

These observations do not address the issue of the locations of the edges themselves. Our experiments focus on the long-range interpolation of shape and its position and have no bearing on the localization of the *local* features from which the interpolation takes place. There are two possibilities, although we have not been able to design a paradigm to distinguish between them

1. Short-range localization is defined by luminance information. Given the Gabor structure of typical cortical receptive fields, luminance position coding would need to be a combination of envelope position and phase information. That is, the coding mechanism would need to know whether the receptive-field profile had odd or even symmetry, and of which sign, in order to correctly localize the elements making up a continuous line, for example. These structures illustrate that luminance position coding is not as simple a matter as might initially be supposed.
2. Short-range localization is not processed in the luminance domain but requires the same process of encoding through the depth map as long-range coding. On this account, the location of the edges is not defined by luminance information. Instead, the visual system is assumed to be blind to the locations of its individual receptive fields (avoiding the problem of integrating over the varieties of phase symmetry in the fields). The receptive field outputs are considered as inputs into the computation of the depth map, which, in the case of Leonardo's figure, has sharp depth edges. It would be the localization of these depth edges in space that defines the locations of the edges, not the localization of the luminance boundaries that happen to coincide with them.

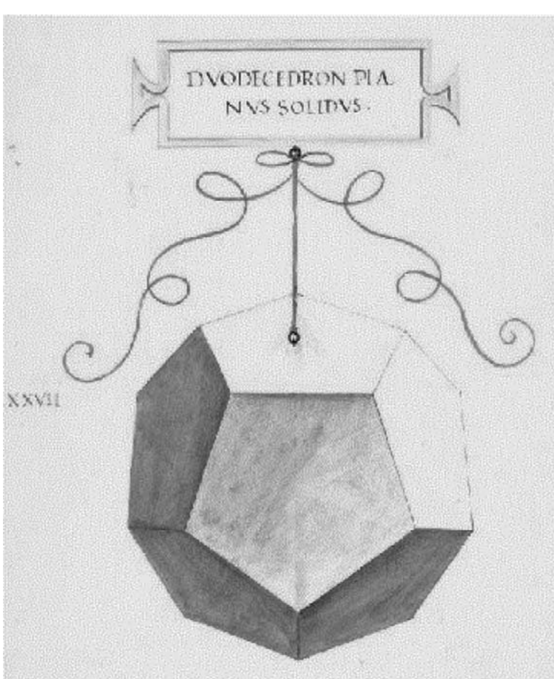

Fig. 5. A dodecahedron drawn by Leonardo da Vinci (1509).

7. Conclusion

It seems clear that the full specification of objects in general requires extensive interpolation to take place, even though some textured objects may be well defined by local information alone. We therefore regard the interpolated position task as more representative of real-world localization of objects than the typical Vernier acuity or other line-based localization tasks of the classic literature. It consequently seems remarkable that luminance information per se is unable to support localization for objects requiring interpolation. The data indicate that it is only through the interpolated depth representation that the position of the interpolated object can be recognized. One might have expected that positional localization would be a spatial form task depending on the primary form processes (Marr, 1982). The dominance of a depth representation in the performance of such tasks indicates that the depth information is not just an overlay to the 2D sketch of the positional information, but a core process that must be completed before the position of the object is known.

Acknowledgement

Supported by NIH grant EY 7890.

References

Buckley, D., Frisby, J. B., & Mayhew, J. E. (1989). Integration of stereo and texture cues in the formation of discontinuities during three-dimensional surface interpolation. *Perception, 18*, 563–588.

Hess, R. F., & Holliday, I. E. (1992). The coding of spatial position by the human visual system: effects of spatial scale and contrast. *Vision Research, 32*, 1085–1097.

Levi, D. M., Klein, S. A., & Wang, H. (1994). Discrimination of position and contrast in amblyopic and peripheral vision. *Vision Research, 34*, 3293–3313.

Kontsevich, L. L., & Tyler, C. W. (1998). How much of the visual object is used in estimating its position? *Vision Research, 38*, 3025–3029.

Kontsevich, L. L., & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. *Vision Research, 39*, 2729–2737.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

Morgan, M. J., & Watt, R. J. (1982). Mechanisms of interpolation in human spatial vision. *Vision Research, 25*, 1661–1674.

Morgan, M. J., & Watt, R. J. (1983). On the failure of spatiotemporal interpolation: a filtering model. *Vision Research, 23*, 997–1004.

Toet, A., & Koenderink, J. J. (1988). Differential spatial displacement discrimination thresholds for Gabor patches. *Vision Research, 28*, 133–143.

Tyler, C. W., & Liu, L. (1996). Saturation revealed by clamping the gain of the retinal light response. *Vision Research, 36*, 2553–2562.

Tyler, C. W. (1997). Diffuse illumination as a default assumption for shape-from-shading in the absence of shadows. *Journal of Imaging Science and Technology, 42*, 319–325.

Würger, S. M., & Landy, M. S. (1989). Depth interpolation with sparse disparity cues. *Perception, 18*, 39–54.

Yang, Y., & Blake, R. (1995). On the interpolation of surface reconstruction from disparity interpolation. *Vision Research, 35*, 949–960.