

Pattern recognition, attention, and information bottlenecks
in the primate visual system

D. C. Van Essen¹, B. Olshausen¹, C. H. Anderson², and J. L. Gallant¹

¹Biology Division, Caltech, Pasadena CA 91125

²Jet Propulsion Lab, Pasadena CA

ABSTRACT

The primate visual system has evolved impressive capabilities for recognizing complex patterns in natural images. This process involves many stages of analysis and a variety of information processing strategies. Here, we concentrate on the importance of 'information bottlenecks', which restrict the amount of information that can be handled at different stages of analysis. We believe these steps are crucial for reducing the overwhelming computational complexity associated with recognizing countless objects from arbitrary viewing angles, distances, and perspectives. The process of directed visual attention is an especially important information bottleneck because of its flexibility in determining how information is routed to high-level pattern recognition centers.

1. INFORMATION BOTTLENECKS

Our visual system routinely and effortlessly carries out a variety of difficult tasks, one of the most important of which is the process of pattern recognition. How do we attend to and recognize a particular object, such as an individual face, among the multitude of objects in a typical visual scene? This is an immensely challenging problem, and despite much progress over the past decade our understanding remains quite fragmentary. A better understanding of the underlying biological mechanisms may have a major impact on future designs for artificial pattern recognition devices.

Our approach to the problem emphasizes the need to understand how information flows through various stages of visual processing. We see this as a crucial, yet somewhat neglected, aspect of both biological and artificial pattern recognition. In biological vision, inherent limitations are imposed at both early and late stages of processing on the amount of information available to pattern recognition mechanisms and on the nature of the information that can be transmitted. In particular, at early stages there is an initial bottleneck that reflects the large differences in visual acuity between the center of gaze and the visual periphery. At late stages, a very different bottleneck is imposed by the process of visual attention. While both of these bottlenecks have been studied intensively in their own right, only a few studies¹⁻³ have analyzed how they interact to constrain overall performance in pattern recognition tasks. Here, our major emphasis will be to develop a biologically oriented framework for addressing the relationship between acuity-related and attention-related bottlenecks.

Figure 1 illustrates how these two transmission bottlenecks interact in human vision. This display, adapted from Anstis⁴, contains two concentric rings of letters surrounding a central target. When one fixates the letter in the center, all of the letters are equally visible, insofar as they are scaled to compensate for the sharp decline in visual acuity with increasing distance from the center of gaze. However, in a single glance it is not possible to recognize all of the letters in the array simultaneously - there is simply too much information for the pattern recognition system to cope with all at once. In order to identify different letters in the pattern, we must direct our attention to the individual letters or, at most, to small clusters of letters. Normally, we make frequent eye movements while carrying out such tasks, but attention and eye movements can be uncoupled by forcing one's eyes to remain fixed on the central letter while switching attention across the display. If one fixates the central letter and concentrates attention exclusively on that letter, it is easily recognized as an 'A'. If, instead, one consciously switches visual attention to other positions in the display (while remembering to keep the eyes fixed), it is easy to identify each of the other letters serially, i.e., one by one. This represents the operation of directed visual attention, a process that




Figure 1. A visual display to illustrate the capabilities and the limitations of directed visual attention. While keeping the eyes fixed on the center, one's ability to identify different characters depends upon whether attention is directed to the central letter, to specific locations on the peripheral rings, or to the entire display.

facilitates pattern recognition within a restricted 'window of attention'^{5, 6}. The window of attention can be rapidly switched from one location to another in the visual field, independent of eye movements, and it can also be adjusted in its spatial extent^{7, 8}. For example, if one attends to the full display all at once (again without moving the eyes), it is easy to visualize the overall pattern of two concentric rings, each consisting of many discrete elements. However, as long as the window of attention is spread over the entire array, the individual letters are no longer easily identifiable.

This simple example gives an intuitive appreciation of the the highly dynamic nature of visual attention and of the fact that only a tiny fraction of the information actually available in the retinal images passes through the window of attention at any given moment. It also suggests the importance of analyzing information flow in terms of a small number of basic measures, such as the position, spatial extent and spatial frequency bandwidth of the window of attention, and of determining how performance limitations vary across the visual field.

2. ANATOMICAL BOTTLENECKS

The next two figures show some of the major structural components of the primate visual system and illustrate how they are related to the information bottlenecks discussed above. Figure 2 is a lateral view of the brain of the macaque monkey, an intensively studied species whose visual system is very similar to that of humans. The outputs from the two eyes are relayed to the primary visual cortex (area V1, or striate cortex, on the left). From there, information is distributed to a collection of 31 extrastriate (higher) visual areas that have been identified in the occipital, temporal, parietal, and frontal lobes⁹. The region most directly involved in high-level pattern recognition is the inferotemporal cortex (IT), a complex of perhaps a half-dozen visual areas in the inferior part of the temporal lobe^{9, 10}.

Figure 3 shows a more schematic diagram of the major routes of information flow from the retina to the inferotemporal cortex. Information in retinal images (far left) is initially encoded by the photoreceptor layer and is




Figure 2. Lateral view of the cerebral cortex of the macaque monkey. Stippled regions represent cortical areas associated with visual processing. Visual areas specifically referred to in the text are identified in the drawing.

subjected to substantial additional processing and data compression within the retina. Retinal ganglion cells, whose axons in the optic nerve are the sole output of each retina, provide a high-acuity representation of the center of gaze and a low-acuity representation of the visual periphery. Thus, the optic nerves form the early, acuity-related bottleneck in information flow described in the introduction. These signals are relayed by way of the lateral geniculate nucleus (LGN) to area V1, which contains an orderly representation of the visual world mapped across its surface. After an initial round of cortical processing, the outputs of V1 are distributed to the mosaic of extrastriate areas, which are arranged in a hierarchy of processing stages. The main route leading to inferotemporal cortex passes through areas V2 and V4.⁹

Because the neurobiological basis of attention is only poorly understood, we cannot link the attention-related bottleneck described in the introduction simply and unequivocally to a unique anatomical site, as we did in assigning the acuity-related bottleneck to the optic nerve. Nonetheless, it is possible to place some constraints on where the bottleneck might occur. It is unlikely that the attentional bottleneck is a completely abrupt constriction in information flow at a single level of visual processing. Instead, we suspect that it arises through stepwise restrictions in information flow as it is funnelled through several successive stages of the visual cortical hierarchy^{11, 12}. We suspect that the projection from V4 to inferotemporal cortex represents a late stage of this funneling process; indeed, Figure 3 explicitly proposes this pathway as a candidate for the attentional bottleneck. However, this is almost certainly an oversimplification. There are other sources of inputs to inferotemporal cortex besides V4, and it is quite possible that the V4 to IT projection is not at the narrowest part of the attentional funnel. On the other hand, a precise anatomical model is not crucial for our present purposes, and the scheme proposed here gives a useful anatomical framework for our first-order modeling efforts. The critical point is that only a tiny fraction of the information leaving the retina passes through this higher, attention-related bottleneck.

3. VISION UNDER THE SPATIAL FREQUENCY TENT

Images of natural scenes contain useful information at many different spatial scales, from the high spatial frequency fluctuations of the just-discernible ridges on one's fingertip to the low spatial frequency undulations in a pattern of clouds or sand dunes. This reflects the fact that the power spectrum of natural images approximates a 1/f distribution^{13, 14}. Although our visual system is capable of extracting information at many different scales, pattern




Figure 3. Information flow and information bottlenecks in the primate visual system. Each hemisphere receives inputs from both eyes, relayed via the lateral geniculate nucleus (LGN), but only from those portions looking out at the opposite half of the visual field. A hierarchy of processing stages leads through areas V1, V2, V4, and the inferotemporal cortex (IT). In addition to the ascending flow of information in this hierarchy, there are major feedback pathways from higher to lower visual areas.

recognition does not operate with maximal efficiency across all scales and all spatial locations at any particular moment.

Figure 4 illustrates the information capacity available at two early stages of the visual pathway, using a format that plots the spatial frequency bandwidth as a function of retinal eccentricity (distance from the center of gaze). The two curves represent the information capacity at the stages of raw retinal images (before transduction by the receptors) and of the optic nerves, respectively.

3.1 Retinal images

At the level of retinal images, the primary limitations are imposed by the physical optics of the eye. In the fovea, where acuity is the highest, images are blurred by an optical point-spread function with a width of about 1'. This corresponds to a cutoff frequency of about 60 cycles/deg in the spatial frequency domain.¹⁵ Image quality is somewhat worse at peripheral eccentricities¹⁶, but the information bandwidth remains broad across the entire retina. We will refer to the domain under this curve as the 'spatial frequency tent' for retinal images. For simplicity of visualization, our illustrations of this tent will show only 1 spatial dimension, but the underlying mathematical analysis can be generalized to 2 spatial (and 2 spatial frequency) dimensions.

The region underneath the spatial frequency tent can be sliced horizontally into multiple layers that represent the information contained in a set of band-limited filters (such as the Laplacian of a Gaussian). This is the essence of the 'pyramid' representations that have proven increasingly powerful in image-processing applications^{17, 18}. If successive bands are an octave apart, a total of 10 levels suffice to contain the information over the range of natural vision (0.1 - 60 cycles/deg). The highest spatial frequency bands contain the most information, and each octave lower resolution contains progressively less information.

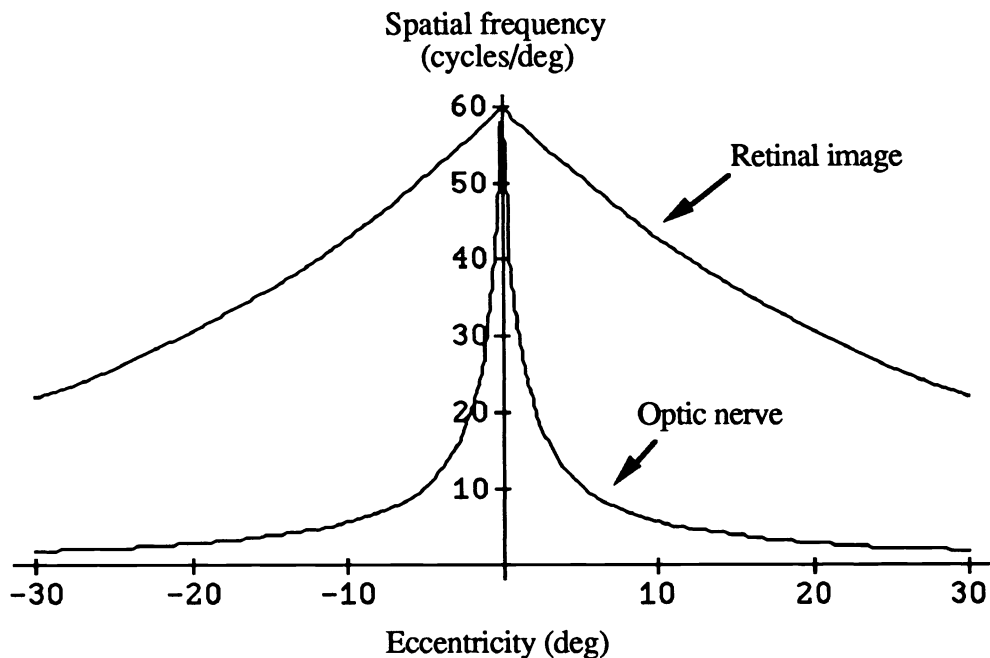


Figure 4. Spatial frequency tents for retinal images and for the optic nerve. The upper and lower curves represent the maximum spatial frequencies at which information is represented in retinal images and in the optic nerve, respectively. The cutoff frequency at a given retinal locus is not completely abrupt at either of these stages, however, as the high frequency roll-off is like that expected from a gaussian blurring function.

3.2 Optic nerve

The corresponding spatial frequency tent for the optic nerve has a very different profile, declining steeply on either side of the center. This reflects a variable-resolution sampling strategy implemented within the retina. It is convenient to express this sampling strategy in terms of an abstract array of sampling nodes, whose density across the visual field reflects the spatial variations in visual resolution¹². We estimate the total number of sampling nodes within the visual field to be approximately 300,000. Their density is very high in the center of gaze ($\sim 10^4$ sampling nodes/deg²) and is vastly lower in the far periphery (~ 2 sampling nodes/deg²). Overall, only about 2% of the spatial information available in the retinal images is transmitted through the optic nerve.

Information from each sampling node is conveyed not by a single retinal ganglion cell, but rather by several (2 - 4) ganglion cells. This allows separate channels to be used for signalling bright vs. dark features and for conveying spectral and temporal as well as brightness information. Individual retinal ganglion cells have a characteristic difference-of-gaussian receptive field profile, which allows them to respond maximally to patterns containing spatial contrast within a range of spatial scales. Interestingly, the center mechanism is usually stronger than the surround mechanism¹⁹. As a consequence of this center-surround imbalance, ganglion cells are neither purely low-pass nor purely band-limited, and they can respond (although not maximally) even to very low spatial frequency patterns. Consequently, visual cortex receives information from all regions under the spatial frequency tent for the optic nerve. In relation to the pyramid representations discussed above, the steep slope of the optic nerve tent implies that progressively higher spatial frequency bands in the pyramid are transmitted over progressively more narrow portions of the visual field. Consequently, the amount of information that can be transmitted over each octave under the optic nerve tent is approximately constant.

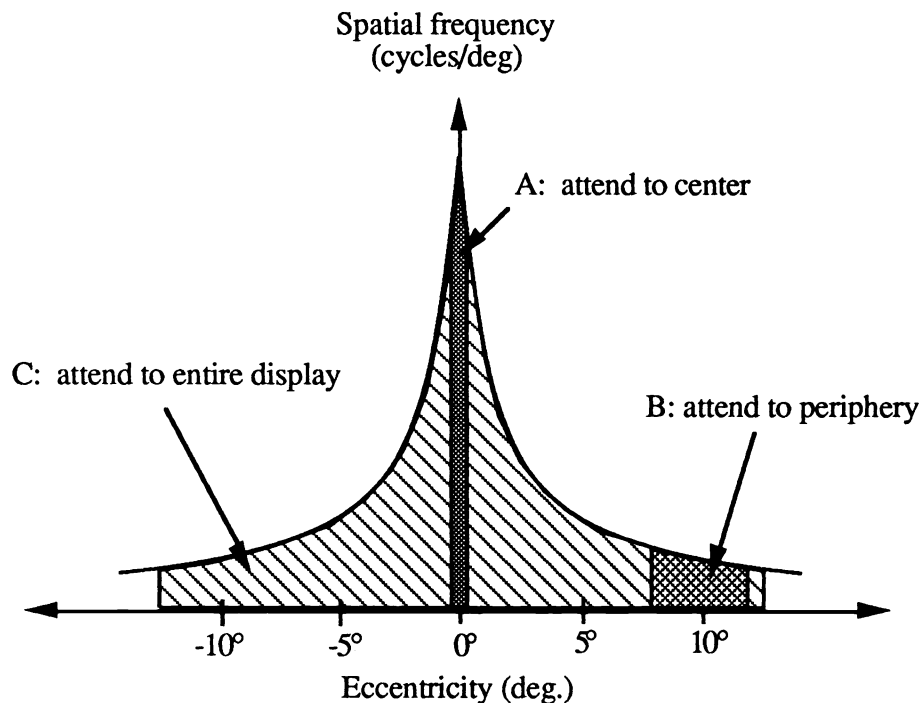


Figure 5. The relationship between the window of attention and the optic nerve spatial frequency tent. Depending on the position, size, and scale of the window of attention, the amount of information potentially available under the spatial frequency tent can vary enormously. However, not all of this information can pass through the attentional bottleneck to reach higher centers. For simplicity, we have drawn the window of attention as though it has sharp edges. This assumption is probably not valid, but neither is it critical to our analysis.

4. THE WINDOW OF ATTENTION

To explore the relationship between the window of attention and the spatial frequency tent for the optic nerve, it is useful to reconsider the example of the concentrically arranged letter patterns shown in Figure 1. When one attends to the single letter in the center of Fig. 1, the window of attention will be restricted to a very small region, probably less than a degree in diameter. This is indicated in Figure 5 by the narrow hatched region under the spatial frequency tent in condition A. Because this is in the region of maximal acuity, a large amount of information is potentially available within this region. To be specific, if we assume that the window is 0.5° in diameter and that it has sharply defined borders, this would correspond to approximately 2,000 potential sampling nodes, or 0.6% of the total number discussed above for the optic nerve. When one attends to a single letter on the outermost ring, while maintaining fixation, the attended region has a much lower acuity (8-fold lower at 10° eccentricity), but the letter occupies a much larger visual angle. If we assume that the window of attention at this peripheral location subtends 4° , this would again correspond to about 2,000 potential sampling nodes (condition B, Fig. 5). Finally, when one attends to the full display ($15^\circ - 20^\circ$ at reading distance), the region inside the window of attention corresponds to 120,000 - 140,000 potential sampling nodes, which is nearly half of the total available information content under the spatial frequency tent (condition C, Fig. 5).

This example illustrates how moment-to-moment shifts in attention can radically change the amount of information from the optic nerve that is available as input to the window of attention. As discussed in the next

section, not all of this information can pass through the window; we estimate the information bandwidth to be roughly about 1% of that available under the entire optic nerve spatial frequency tent.

4.1 Limits on the window of attention

Information concerning limitations on the window of attention comes primarily from psychophysical studies of human performance. These include studies of visual search, in which an observer must locate a predefined target in a field of distractors, and studies of detection and discrimination thresholds for simple signals such as luminous spots or spatial frequency gratings. Two important findings have emerged that are particularly relevant to our analysis: (i) the window of attention is probably unitary and cannot be distributed across discrete locations in the visual field; and (ii) attention is bandpass in the spatial frequency domain.

Unitary nature. In search tasks where complex targets are distributed across widely disparate locations in the visual field, observers cannot attend simultaneously to two distinct locations any more efficiently than they can to the entire area as a whole^{20, 21}. Under some conditions observers may appear to split their attention across locations²². However, these are probably instances of observers using complex decision criteria, a mechanism well known in signal detection theory that could be mediated more centrally in the brain than the window of attention^{21, 23}. Thus the evidence suggests that the window of attention is essentially unitary, and that it cannot be broken apart or distributed arbitrarily across visual space.

Bandpass characteristics. Sperling and Melchner²⁴ measured performance in a search task containing an array of characters placed in two concentric rings; the characters were either all the same size or were small in the inner ring and large in the outer ring (cf. Fig. 1). They found that observers could not attend efficiently to two different letter sizes at a time. Since the spatial frequencies most useful for letter recognition depend on the size of the letter, this suggests that the window of attention has a limited spatial frequency passband. Other evidence suggests that the window of attention is genuinely bandpass rather than lowpass. For example, discrimination and detection efficiency fall off with increasing bandwidth of signals^{25, 26}. When observers are asked to detect the presence of a near-threshold grating, their performance is best if the grating has a known spatial frequency. Performance declines as the bandwidth of likely gratings increases, even if observers know the likelihoods of different frequencies appearing beforehand²⁶ (but see ref. 27). These effects are seen for both low and high frequency areas of the spectrum, suggesting that low frequency signals do not have a privileged status.

Figure 6 illustrates a scheme for the window of attention which is compatible with these constraints. The stippled region indicates what spatial and spatial frequency information is passing through the window of attention at one instant in time. The window cannot be broken into two or more separate regions, and it cannot be bent so as to permit different bandpass filtering at different spatial locations. It can be expanded in either the spatial or spatial frequency dimensions, but only at the cost of decreasing signal quality. This model might seem overly restrictive, but the visual system possesses several means of minimizing these limitations. Nonlinear processes in early vision can serve to overcome certain aspects of the bandpass characteristic (see below). In addition, people are capable of devising sophisticated strategies for positioning the window and adjusting it quickly and efficiently to meet the demands of any particular task.

4.2 The attentional sampling grid.

In an earlier section we invoked the notion of an abstract sampling grid to indicate the spatial grain at which different parts of the visual field are represented at the optic nerve bottleneck. The interval between sampling nodes in this grid is dictated by the Nyquist sampling theorem as applied to visual acuity (two samples needed for each cycle of the highest resolvable spatial frequency). Here, by analogy, we suggest that the notion of an 'attentional sampling grid' is a useful construct for evaluating specific aspects of the aforementioned bandpass model for visual attention. From the Nyquist sampling theorem, we argue that there must be two samples in the attentional sampling grid for each cycle of the highest spatial frequency that is being attended.

What is the size of the attentional sampling grid? Nakayama¹ suggested a grid size of 10 x 10 for the window of attention (the 'iconic bottleneck' in his terms), based mainly on studies of the number of characters that




Figure 6. A bandpass model for the window of attention. This model presumes that the window of attention can be adjusted in its position, size, preferred spatial frequency, and spatial frequency bandwidth, as long as the attended region is spatially uniform and lies within the bounds of the optic nerve spatial frequency tent.

can be simultaneously assimilated when reading. We suspect the grid may have several times better resolution in each linear dimension and thus about 10 times greater informational bandwidth. This is based on several lines of psychophysical evidence, including (i) the pixel density needed for efficient communication of sign language and other complex patterns; (ii) the analysis of contrast threshold vs. grating size; and (iii) spatial displacement thresholds and their dependence on feature size and spacing. None of these studies give unequivocal answers, but together they support an internally consistent model.

Optimal pixel density. Sperling and colleagues²⁸ studied how the intelligibility of American Sign Language (ASL) communication depends upon image pixel density. One of their key findings was that 24 x 16 gray-scale images (spatially smoothed and presented at 15 frames/sec) were 86% as efficient in communication as 96 x 64 images, which have 16-fold greater information rate. If the strategy of the ASL observer in this paradigm is to set his/her window of attention to match the total size of the image viewed, then performance should deteriorate markedly once the image resolution becomes worse than that of the attentional sampling grid. Thus, the Sperling et al. study is consistent with an attentional sampling grid of about 24 x 16. This is only a rough approximation, though, because they did not explore image pixel densities below 24 x 16, nor did they explicitly attempt to control or measure how attention was directed in this paradigm.

Grating size. A second line of evidence comes from studies of how the threshold for detecting a small patch of a spatial frequency grating depends upon the number of cycles within the patch. As the number of cycles increases, the contrast threshold decreases until an asymptote is reached at about 12 - 16 cycles, independent of the

spatial frequency of the grating^{29, 30}. It has been argued that this asymptote might reflect the limitations of visual acuity, because increasing the extent of a grating patch adds extra cycles in progressively more peripheral regions of the visual field, where acuity is reduced³¹. However, by our estimates, performance in this task may not be limited by acuity; we suspect instead that it may be governed by the size of the attentional sampling grid. If so, the Nyquist theorem would suggest that the grid is about 25 - 30 sampling elements in one dimension.

Spatial displacement thresholds. A third line of evidence comes from the analysis of thresholds for detecting spatial displacements of a central stimulus relative to a pair of stimuli on either side. Toet et al.³² used gaussian-blurred 'blobs' for their stimuli, and they measured thresholds as a function of both the separation between blobs and the size of the blobs. For a given separation between the outer blobs, they found that displacement thresholds decreased when the blobs were reduced in size. This makes sense if one assumes that the inherent accuracy with which the position of a blob can be determined is better for smaller blobs. Interestingly, once the blobs were reduced below a critical size, relative to the overall separation, there was no additional improvement in performance. This transition took place at a critical ratio of separation:size of about 25. We hypothesize that this transition reflects a limitation in attentional processing that arises because the alignment task requires the window of attention to cover the entire display pattern. Once the blobs are smaller than that which can be resolved by the attentional sampling grid, the attentional system becomes unable to take advantage of the potentially greater accuracy with which small blobs can be localized, and performance therefore reaches a plateau. If this interpretation is correct, it would again suggest an attentional window size about 20 - 30 nodes across.

We conclude from the above analysis that the size of the attentional sampling grid may be roughly 25 x 25. An interesting but unresolved question is whether the sampling grid can become elongated for tasks such as reading that are more demanding in one dimension than the other. More generally, there is a clear need for careful psychophysical studies to probe the characteristics of the attentional sampling grid in the dimensions of space, spatial frequency, and time.

4.3 Feature dimensions within the attentional sampling grid.

In our earlier discussion of the retinal sampling grid, we argued that there are several retinal ganglion cells subserving each sampling node. This allows for multiple channels to convey diverse types of information pertaining to each small region of space. This concept is even more critical in relation to the attentional sampling grid, where there is likely to be even greater diversity in the types of information that can be transmitted about each small region of the visual field. The different types of information may include texture, depth, motion, and color, as well as simple luminance, as suggested in Figure 7.

The key points underlying this scheme can be illustrated in relation to the following example. Suppose that one stands near a tree and sets the window of attention wide enough to subtend the entire tree, so that the overall shape of the tree and the positions of the major branches are readily perceived. While keeping the window of attention wide, it is possible to discern the leafy texture of the branches, even though individual leaves are much finer than the spacing of the attentional sampling grid. This could be a consequence of non-linear processing at early stages of the visual system, which could shift information about average textural characteristics from the higher frequency bands in which it is initially represented to lower frequency bands in which it can be represented more compactly. This abstracted representation would provide inputs to the attentional window, even when it is operating at a much coarser spatial resolution than the original texture. In all, we estimate that about 10 such featural dimensions may be required to describe the characteristics of each patch within the attentional sampling grid.

Another crucial characteristic of the attentional sampling grid, in our view, is its capability for preserving information about detailed spatial relationships within the attended region of the image. This is necessary in order to provide flexibility for handling the countless variety of possible visual inputs and matching them to the enormous numbers of possible templates that can be stored in our visual memories.

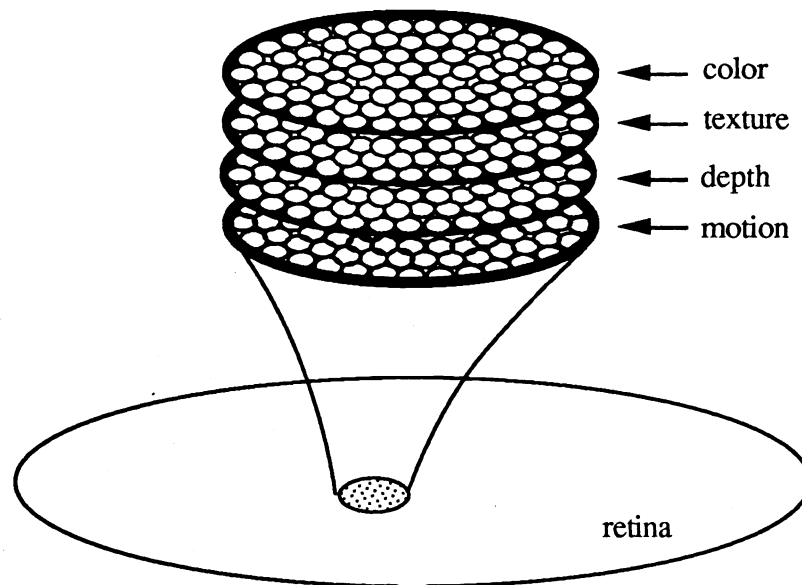


Figure 7. The attentional sampling grid. This scheme presumes that the window of attention is represented by a sampling mosaic covering about 25 x 25 samples in its spatial dimensions and about 10 levels deep at each location in order to convey information about texture, color, depth, and motion.

5. CONCLUDING REMARKS

Biological systems have evolved extremely robust and flexible strategies for carrying out a variety of pattern recognition tasks. These strategies reflect principles of good engineering that are needed to find practical solutions to problems of immense computational complexity. This article has concentrated on what we regard as a critical, but largely neglected, issue of how information is routed and represented through the window of attention in order to provide optimal inputs to high level associative memory processes for pattern recognition. Our analysis provides an attractive model for certain aspects of visual attention. Equally importantly, it raises a number of questions that are amenable to detailed experimental analysis. The problem of understanding the dynamics of information flow is certainly challenging, but is by no means intractable, in our opinion. What has been most notably lacking is an incisive conceptual framework for formulating and attacking these key issues. The framework proposed here addresses only a subset of the critical issues relating to attention and pattern recognition. We have said nothing about how the visual system integrates multiple shifts of attention in order to achieve a coherent percept of a single object³³. Nor have we addressed in any detail what the control system might be for redirecting attention many times per second^{5, 8}. Our hope, though, is that the present formulation is a step in the right direction that can be naturally expanded and refined to encompass these and other important issues.

6. ACKNOWLEDGMENTS

This project was supported by ONR grant N00014-89-1192. We thank Bela Julesz and Milan Jovovic for valuable discussions.

7. REFERENCES

1. K. Nakayama, "The Iconic Bottleneck and the Tenuous Link Between Early Visual Processing and Perception," Vision: coding and efficiency, C. Blakemore (ed.), Cambridge University Press, Cambridge, 1991.

2. C. J. Downing and S. Pinker, "The spatial structure of visual attention," M. I. Posner and O. S. Marin (eds.), Attention and Performance XI, pp. 171-188, Hillsdale, 1985.
3. R.J. Watt, "Scanning from Coarse to Fine Spatial Scales in the Human Visual System after the Onset of a Stimulus." Journal of the Optical Society of America A, vol. 4, 2006-2021, 1986.
4. S.M. Anstis, "A Chart Demonstrating Variations in Acuity with Retinal Position," Vision Research, vol. 14, pp. 589-592, 1974.
5. B. Julesz, "Early Vision and Focal Attention," Review of Modern Physics, in press, 1991.
6. W.A. Johnson and V.J. Dark, "Selective Attention," Ann. Rev. Psychol., vol. 37, pp. 43-75, 1986.
7. G. Sperling and B.A. Doshier, "Strategy and Optimization in Human Information Processing," K.R. Boff, L. Kaufmann, and J.P. Thomas (eds.), Handbook of Perception and Human Performance, vol. I, Wiley, New York, 1986.
8. C.W. Eriksen and T.D. Murphy, "Movement of Attentional Focus Across the Visual Field: A Critical Look at the Evidence," Perception and Psychophysics, vol. 42, 299-305, 1987.
9. D. Felleman and D. Van Essen, "Distributed Hierarchical Processing in Primate Visual Cortex," Cerebral Cortex, in press, 1991.
10. G.C. Bayliss, E.T. Rolls and C.M. Leonard, "Functional Subdivisions of the Temporal Lobe Neocortex," J. Neurosci., vol. 7, 330-342.
11. C.H. Anderson and D. Van Essen, "Shifter Circuits: A Computational Strategy for Dynamic Aspects of Visual Processing," Proc. Natl. Acad. Sci., vol. 84, pp. 6297-6301, 1987.
12. D. Van Essen and C.H. Anderson, "Information Processing Strategies and Pathways in the Primate Retina and Visual Cortex," Introduction to Neural and Electronic Networks, S.F. Zometzer, J.L. Davis, C. Lau (eds.), pp. 43-72, Academic Press, Orlando, 1990.
13. R.F. Voss, "1/f (flicker) Noise: A Brief Review" Proc. 32nd Ann. Symp. on Frequency Control, Atlantic City, pp. 40-46, 1979.
14. D.J. Field, "Relations between the Statistics of Natural Images and the Response Properties of Cortical Cells," Opt. Soc. Am., vol. 4, no. 12, pp. 2379-2394, 1987.
15. W.N. Charman and G. Walsh, "The Optical Phase Transfer Function of the Eye and the Perception of Spatial Phase," Vision Research, vol. 25, pp. 619-623, 1985.
16. N. Drasdo and W. Fowler, "Non-linear Projection of the Retinal Image in a Wide-angle Schematic Eye," Brit. J. Ophthalmol., vol. 58, pp. 709-714, 1974.
17. P.J. Burt and E.H. Adelson, "The Laplacian Pyramid as a Compact Image Code," IEEE Trans. Commun. COM-31, pp. 532-540, 1983.
18. C.H. Anderson, P.J. Burt and G.S. van der Wal, "Change Detection and Tracking Using Pyramid Transform Techniques," Proc. SPIE Conf. on Intell. Robots and Computer Vision, pp. 72-78, 1985.
19. A.M. Derrington and P. Lennie, "Spatial and Temporal Contrast Sensitivities of Neurones in Lateral Geniculate Nucleus of Macaque," J. Physiol., vol. 357, pp. 219-240, 1984.
20. A.H.C. van der Heijden, G. Wolters, J.C. Groep, and R. Hagenaar, "Single letter recognition accuracy benefits from advance cueing of location." Perception and Psychophysics, vol. 42, 503-509, 1987.
21. H.J. Muller and J.M. Findlay. "Sensitivity and criterion effects in the spatial cueing of visual attention." Perception and Psychophysics, vol. 42, 383-399, 1987.
22. M.L. Shaw and P. Shaw, "Optimal Allocation of Cognitive Resources to Spatial Locations." Journal of Experimental Psychology: Human Perception and Performance, vol. 3, 201-211, 1977.
23. M.L. Shaw, "Attending to Multiple Sources of Information: I. The Integration of Information in Decision Making," Cognitive Psychology, vol. 14, pp. 353-409, 1982.
24. G. Sperling and M.J. Melchner, "The Attention Operating Characteristic: Examples from Visual Search," Science, vol. 202, pp. 315-318, 1978.
25. C.J. Downing, "Expectancy and visual-spatial attention: Effects of perceptual quality." Journal of Experimental Psychology: Human Perception and Performance, vol. 14, 188-202, 1988.
26. E.T. Davis and N. Graham. "Spatial Frequency Uncertainty Effects in the Detection of Sinusoidal Gratings." Vision Research, vol. 21, 705-712, 1981.
27. E.T. Davis, "Allocation of attention: Uncertainty effects when monitoring one or two visual gratings of noncontiguous spatial frequencies." Perception and Psychophysics, vol. 29, 618-622, 1981.
28. G. Sperling, M. Landy, Y. Cohen and M. Pavel, "Intelligible Encoding of ASL Image Sequences at Extremely Low Information Rates," Computer Vision, Graphics, and Image Processing, vol. 31, pp. 335-391, 1985.

29. V. Virsu and J. Rovamo, "Visual Resolution, Contrast Sensitivity, and the Cortical Magnification Factor," Exp. Brain Res., vol. 37, pp. 475-494, 1979.
30. J. Jamar and J.H. Koenderink, "Sine-wave Gratings: Scale Invariance and Spatial Integration at Suprathreshold Contrast," Vision Res., vol. 23, pp. 805-810, 1983.
31. N. Graham and J.G. Robson, "Probability Summation and Regional Variation in Contrast Sensitivity Across the Visual Field," Vision Research, vol. 21, pp. 409-418, 1981.
32. A. Toet, M.P. van Eekhout, H.L.J.J. Simons and J.J. Koenderink, "Scale Invariant Features of Differential Spatial Displacement Discrimination," Vision Research, vol. 27, pp. 441-451, 1987.
33. R.J. Baron, The Cerebral Computer, Erlbaum, Hillsdale, 1987.