Scale and the Shape/Texture Continuum[1]

Eric Saund
Xerox Palo Alto Research Center

Abstract

This paper is about the way that spatial scale mediates the relationship between visual shape and visual texture. We seek ultimately representations for significant image events that will support multiple later visual processes, including refinement and exploitation of figure/ground relations (segmentation), indexing and matching with object and scene models, and directing visual attention for the selection and application of further processing steps. These representations should support description of a scene in abundant detail and multiple levels of abstraction, yet favor omission of information that is unlikely to be useful. We introduce the notion of a texture scale-space *making explicit the relationship between two scales of interest, the characteristic grain size of image elements, and the size of a frame of view. The analysis entails consideration of several interrelated concepts, including the notion of an image feature, frame of view, spatial coherence, scale-dependent representation of shape, feature uniformity in a region, and odd-man-out phenomena. We describe experiments with two algorithmic approaches, one based in spatial filtering, the other in fine-to-coarse spatial aggregation of discrete events.*

1 Introduction

"Texture" occurs when there are too many or too complex signal changes within a contiguous region to describe in detail, so aggregate properties must be used instead. Figure 1 illustrates. When focused on by a frame of view, the figure in the image patch is regarded in terms of its shape. As the frame of view

[1]This is not a completed research paper describing a body of work in relation to other research in the field. There are, for example, no citations. Instead, this is a summarization of thinking and experiments dating to 1991 and 1993 and is intended to put forth a viewpoint in support discussion and collaborative efforts in this area.

zooms back to encompass surrounding image material, this same image patch can assume any of several roles. Among them, it can blend in as an anonymous part of a texture (Figure 1b), it can pop out as special event (Figure 1c), or it can retain its role in defining the shape of the figure (1d). Which of these occurs has to do with the ways and degrees to which the focal figure is "like" its surroundings. This simple observation opens the door to a complex interplay between several concepts widely believed to be of central importance in computational vision. The notions of shape, features, texture, pattern, scale, frame of view, salience, and the purpose of a visual representation itself, all matter. This working paper is a discussion of the relationships among these and a description of two experiments exploring new formulations of image texture analysis. The topic is vast and this treatment is necessarily sketchy. Our main intent is to begin to expose broad outlines of the interrelationships between a spectrum of vital concepts.

1.1 Representing Image Structure by Labeling

The purposes for which a representation for image texture may be used are many and varied. An exhaustive taxonomy would include the following.

- **Classifying materials and their properties.**
- **Detecting boundaries between materials.**
- **Detecting boundaries between surfaces which may or may not be made from the same material.**
- **Estimating surface shape and pose in space.**
- **Detecting significant collections of markings on marked surfaces.**
- **Identifying, classifying, and evaluating spatially distributed objects or collections of objects in scenes.** For example, when viewed at a distance, the leaves of the tree become elements of a texture permitting analysis of the tree's species and spatial extent.
- **Picking out objects or other events differing from their surroundings.**
- **Indexing object models.** Based on the *identities* of texture parts, a subset of object models in a database can be selected for possible matches. See Figure 2.

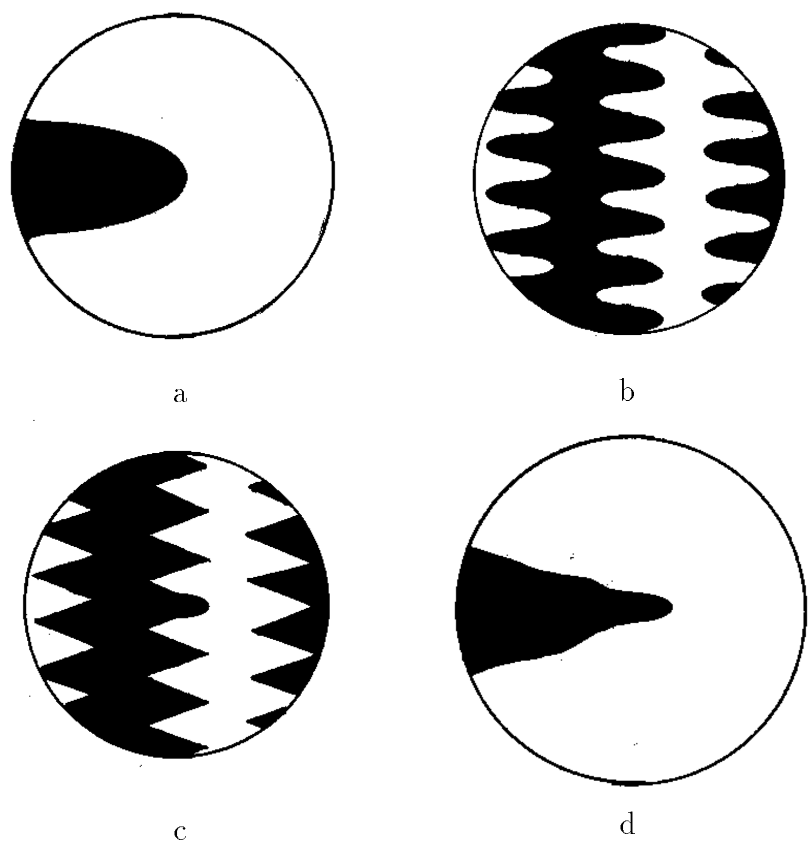

Figure 1:

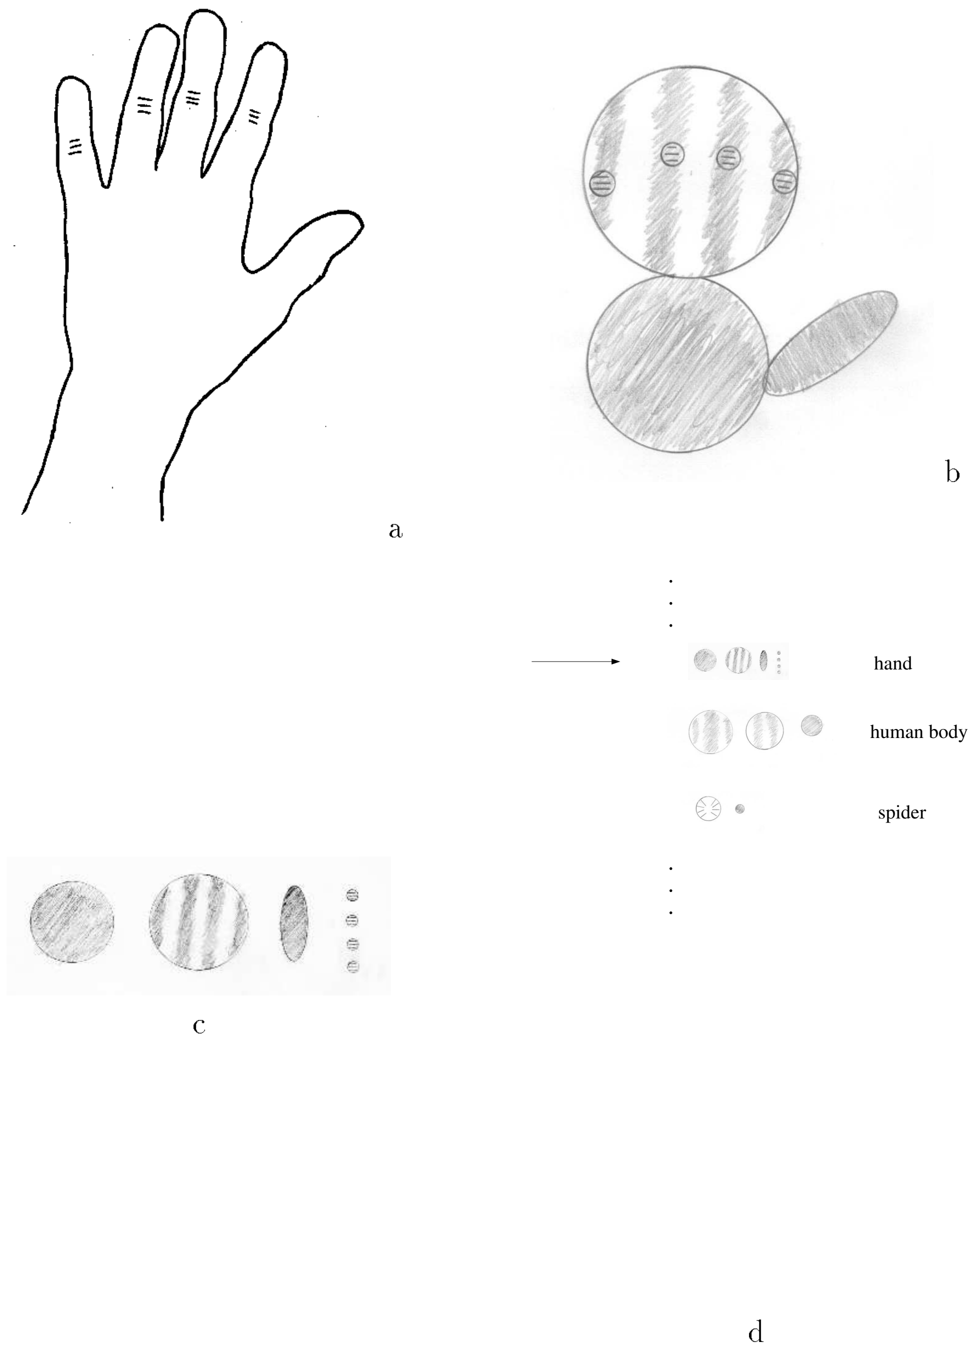

Figure 2: Texture label identities can be used for indexing models in a database. a. Input image. b. Texture label description. Solid circle and ellipse indicate "figure" regions, striped circles denote the assertion of alternating figure/ground with region scale, orientation, and spatial frequency drawn in the style of Gabor filters. c. Texture label identities (no spatial configuration specified. d. Indexing into an object database based on texture label identities.

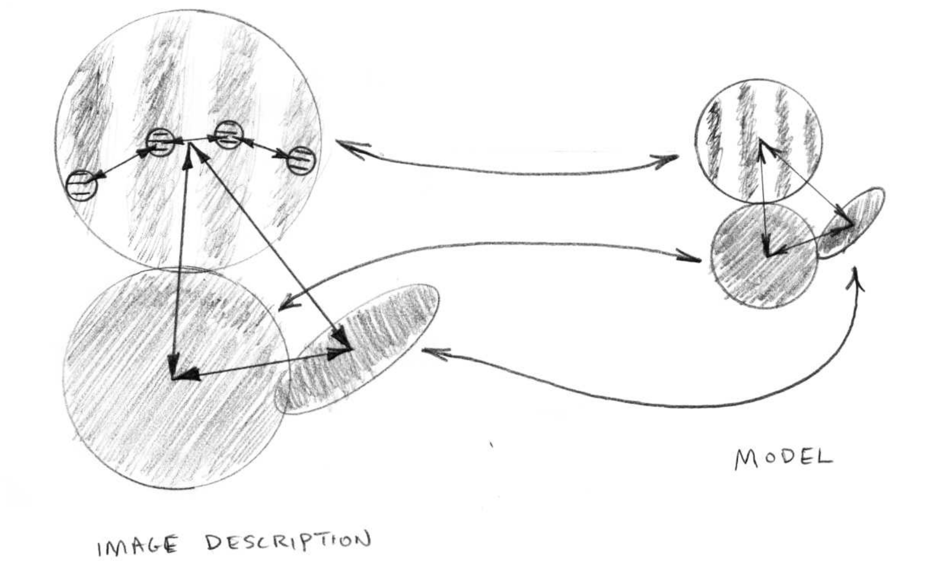

Figure 3: Matching an observed configuration of texture labels with the configuration specified by an object model.

- **Matching to object models.** Based on the *configuration* of texture parts, an object can be matched to an entry in an object model database. See Figure 3.

- **Indexing and initializing visual routines.** In Figure 4, the task is to count the number of "fingers" on the object. The fingers are first identified coarsely as an oriented texture patch. The position and orientation of this patch is used to initialize the parameters of a visual routine that sequentially counts bar items.

Classically, work in image texture analysis has focused on the first four items of this list. The task of classifying material properties calls for local statistical measurement techniques; the task of identifying material and surface boundaries has driven work in detection of texture edges; estimation of surface geometry calls for analysis of of gradual changes in properties over regions. It is worth turning attention, however, to the remaining purposes for texture analysis.

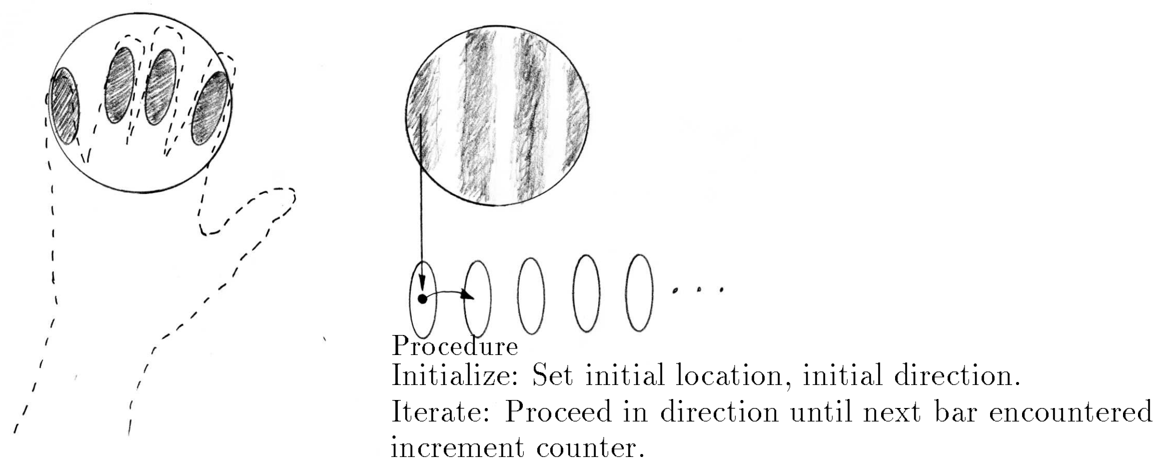

Figure 4: Texture labels can be used to index, launch, and initialize visual routines. Here, the routine performs a counting task by sequentially indexing "finger" objects whose orientation and starting point are indicated by a texture label.

For these tasks, it is appropriate to consider representations for texture in terms of concise *labels* declaring the presence and summarizing the properties of texture regions. The difference between a label and, say, a broad "energy" map of responses over an image, is that a label concentrates information in one place, as it were, so that, following Marr's Principle of Explicit Naming, it may be operated on as a unit. For example, in Figure 5 information useful both to recognizing a broom and judging its pose in space is contained in the spatial configuration of the handle part and the brush part, which in turn encapsulates the distribution and qualities of the bristles. We may certainly gain access to more detailed information about subregions of the brush through further inspection, but, as demonstrated by "Hidden Pictures" illustrations, this dissection is not part of normal processing. Did you notice the pencil in the broom drawing?

2 The Shape Realm and the Texture Realm

Spatially organized labels are useful because they make explicit the *presence* and *attributes* of potentially important structure in images, and they provide a substrate for exploiting spatial configurations of image events. The vocabulary

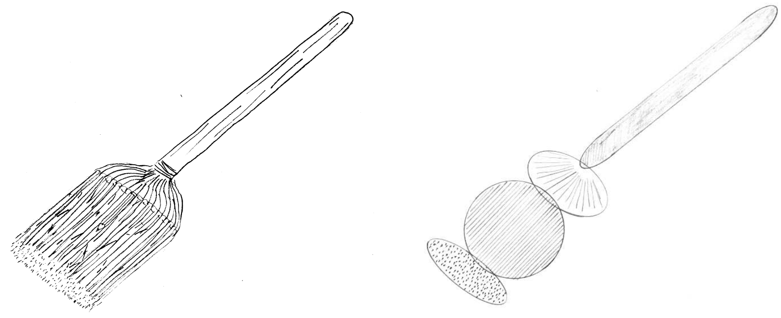

Figure 5: Broom, and its description in terms of texture labels.

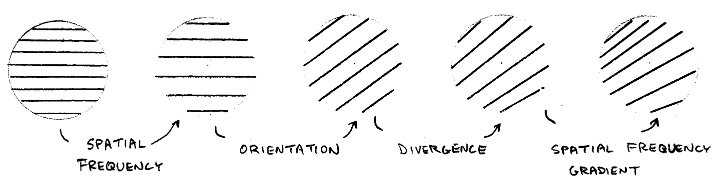

Figure 6: Variations in image texture that could be encoded either as distinct label types or as parameterization of a single descriptor.

of label types depends of course on the purpose for which representation will be used.

Distinctions between similar scene fragments can be described in terms either of distinct labels, or different parameterizations of a common label. See Figure 6. The choice of representations employing a proliferation of label type, versus parameters within a fewer number of types, is perhaps not a natural one, but more an artifact of the way computer scientists organize data structures. Of greater significance is the differences in image structure a given representation is intended to and able to distinguish, and the variations it is intended to and able to generalize over.

2.1 Complexity and Comprehensibility

To illustrate, consider a hypothetical visual subsystem of a bird, designed to analyze clusters of berries. The input to this subsystem, after some sort of feature detection, can be portrayed as dots indicating the positions of individual berries. Figure 7 presents clusters of dots due to various species of berry bushes, some safe, some poisonous, in a taxonomy of six different hypothetical environments. Berry clusters in the sparse environments contain three berries, clusters in the moderate environments contain 6 berries, while clusters in the cluttered environments contain more than a dozen. "Comprehensible" environments differ from the "perplexing" environments by the complexity of the algorithm required to distinguished safe from poisonous berry clusters, where we assume some given tolerance for spatial localization.

Clearly, clusters with more berries can display more possible configurations than clusters with fewer berries. In sparse environments it may be perfectly possible to memorize individual dot patterns signifying safe versus poisonous clusters. This is even possible even in a perplexing-sparse environment. To do this in moderate and cluttered clusters however requires potentially an exponentially increasing amount of memorization. Instead, as more data elements are dealt with, the problem becomes manageable only when some simplifying strategies can be brought into play. Informationally complex environments become comprehensible when they afford rules about overall shape, pattern, uniformity, or other measures of *aggregate* element distribution. Sometimes these generalize statistically over the details of the placement of individual data elements, as in the uniform/clumpy distinction in the complex-comprehensible environment, while sometimes they rely on very precise placement of data elements forming regular patterns such as rows or circular fields.

Even when the problem is not to simply categorize patterns, the overall point is that there exists a continuum between a *shape realm* where analysis can proceed in terms of the detailed configuration of data elements, and a *texture realm* where aggregate measures become appropriate. The texture realm arises amidst the alliance of two factors. First, it becomes computationally intractable to maintain highly detailed information about a large number of highly complex objects. Second, in our world, most of the information present in complex visual objects doesn't matter. Safe clusters of berries and poisonous clusters of berries fortunately do not differ only in the details of their configurations; complex-perplexing environments occur rarely enough in nature that we can usually avoid them.

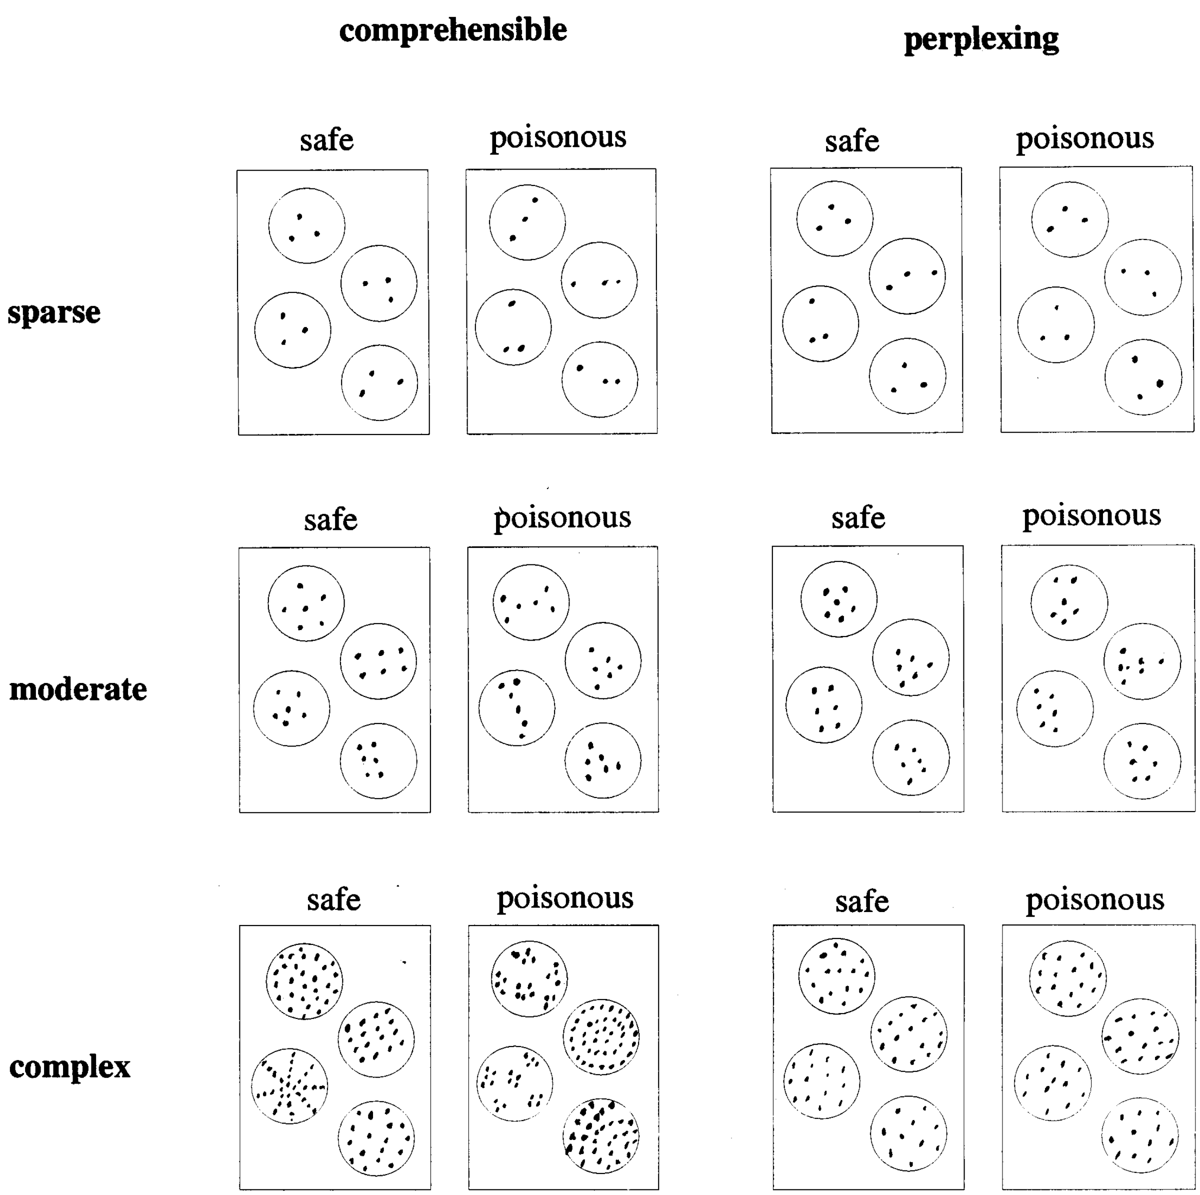

Figure 7: Configurations of safe and poisonous berries in six different environments.

2.2 Charting Attributes Across Shape and Texture

By regarding shape and texture as poles of a continuum, we direct attention to the abstractions performed by a visual representation. In figure 8 we present a *Shape/Texture Attribute Chart* that explores the relationship between the shape realm and the texture realm of local analysis and characterization of visual properties. This chart organizes the relationships among many of the most significant and recognized components of the vocabularies used in the field for describing both shape and texture. The framework for this view is that the purpose of a descriptive language is to characterize some aspect of the attributes and configuration of primitive elements within a field of view. Thus these descriptors begin to form a basis for labeling significant structure in images. In the chart, the visual elements themselves are considered only with respect to the attributes of location and perhaps orientation. A comprehensive chart would include a larger set of attritubes including color, motion, aspect ratio and other aspects of element shape, and so on. Ways of dealing with combinations of elemental attributes are addressed in Section 3.2.

Several interesting features of this chart serve as points of departure for considering the central issues in the relationship between shape and texture.

- A description of element locations in terms of rote templates is expressed in the top row. As discussed in the berry cluster example, the tractability and usefulness of employing rote templates diminishes as the number of data elements increases toward the texture realm.

- The second row shows how proximity abstractions evolve between the shape and texture realms. Distance between two elements becomes average distance among a small number of elements, then finally becomes density in the texture realm.

Some interesting things happen when individual elements are given local attributes. For purposes of illustration, *orientation* will be taken as a representative local attribute. Attributes can be considered alone or in combination with others, and also in combination with elements' locations.

- The third row explores elements' attributes alone. In the shape realm, elements' attribute values themselves can be listed and compared. Moving toward the texture realm, listings, or else histograms, of attributes and their differences with respect to neighboring elements capture significant information. In the texture realm, it becomes sensible to characterize attribute values using parametric statistical distributions such as number of, mean, and variance of distribution modes.

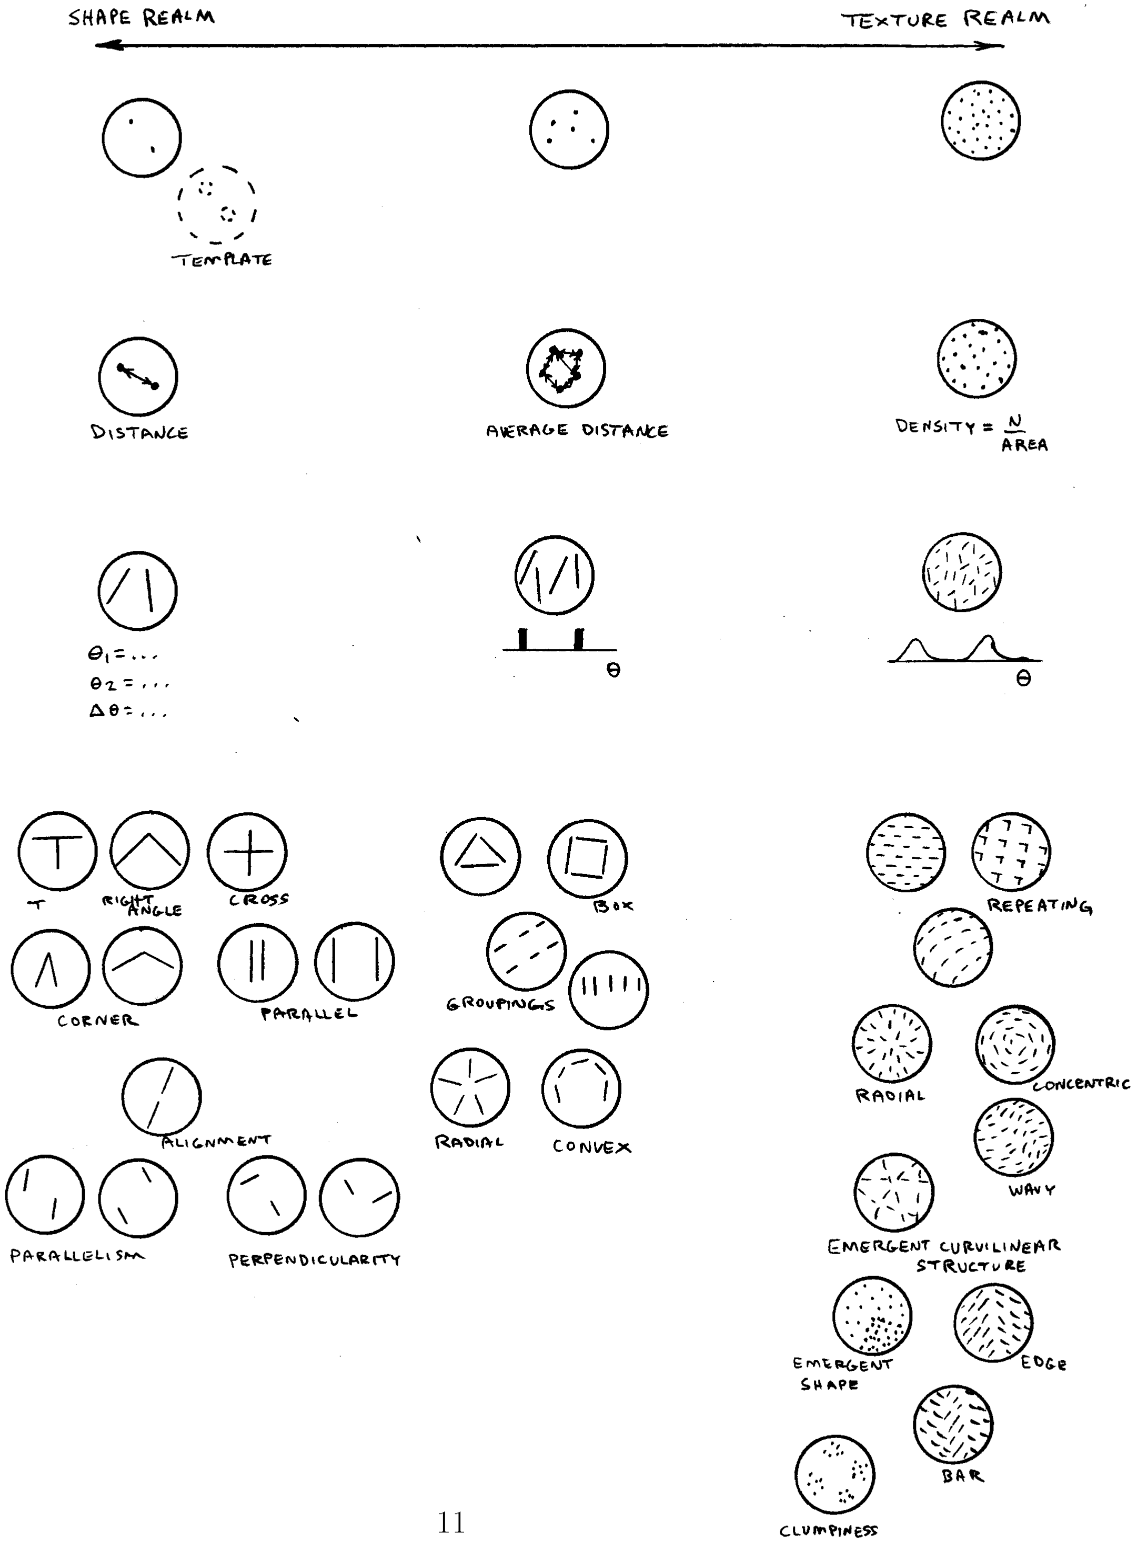

Figure 8: Chart characterizing shape/texture attributes across the shape and texture realms.

- The fourth row considers abstract properties of the conjunction of element attributes and location.

  Within this row, we identify different degrees of specificity in characterizing configurations or patterns, ranging from tight constraint to broadly general regularity. In the shape realm where the attribute is element orientation, some well-known shape figures conforming to rigid templates include the right-angle corner, T-junction, and cross. If templates are allowed to deform, then parallel, and general-angle corner emerge as recognizable types. Finally, the abstractions known as *alignment*, *parallelism*, and *perpendicularity* are weaker characterizations of local shape-like configurations.

  Moving to the intermediate realm, named figure categories become more complex, and hence sparser in relation to the possible configurations that could be depicted. Triangle, box, and star are objects for which names exist. More general abstractions that are sensible in this region include radial structure, convexity, and alignment groupings.

  In the texture realm, high constraint is reflected in so-called repeating textures whose elements form fixed local patterns. Lower constraint gives rise to more varied patterns whose overall structure may still have names like radial, concentric, and so forth. In some cases local groupings of image elements contribute to emergent structure. For example, cuvilinear grouping gives rise to objects that become primitives for analysis in terms of more shape-like descriptors. A uniform field differs from a raked field in two regards. First, grouping of aligned elements produces emergent curvilinear line or bar primitives, and second, the distribution of image elements becomes clumpy.

Some descriptive properties are meaningful only in the shape realm or texture realm, but not on the extreme other end of the continuum.

- Clumpiness in spatial distribution is a property that only makes sense in regard to image texture. In the crossover region between shape and texture realms where a small number of image elements are considered, the relatively small number of modes of clumpiness can be enumerated. Clumpiness reflects the presence of structure at more than one scale.

- Variation in density of elements or attribute values (e.g. orientation) across the frame of view is a phenomenon that occurs only in the texture realm. Shallow gradients and center-surround phenomena are characterized by various moments in the distribution. Sharp gradients are known as "edges"; double edges are known as "bars".

In the next section we explore several central issues raised by the Shape/Texture Attribute Chart: the role of scale, placement of the frame of view, the notion of region coherence, and the phenomena of odd-man-out.

2.3 Significant Scale

The prototypical examples depicted in the Shape/Texture Attribute Chart consider configurations of image elements having attributes of location, and possibly orientation. Let us introduce now the attribute of image element *size*. In one sense, element size is just another continuous-valued attribute that may be uniform or vary among image elements in a region, as shown in Figure 9. But something more interesting happens when we consider image element size *in relation* to the size of the frame of view. Figure 10 poses the question of what aspects of spatial scale or size are most significant in characterizing the visual appearance of a region. Consider that under most viewing conditions, the absolute size of a visible object forming a shape or texture element is not stable, but depends on distance to and focal length of the sensor. Similarly, the dimensions in the image of the frame of view just encompassing a delineated object depends also on distance to and focal length of the sensor. What *is* stable is their *relative* sizes. This suggests that the image shape/texture characteristic worthy of measuring is the ratio of element size to region size, as this is self-similar with respect to magnification.

Figure 11 thus depicts a *Texture Scale-Space* showing the relationship between two significant scales, region scale and characteristic grain size of image elements. Traversing vertically in this diagram is a "pure" scale change that simply magnifies both scales uniformly while preserving their ratio.

The diagonal direction corresponds to varying the frame of view to encompass more or fewer image elements. The horizontal axis corresponds to changing the characteristic grain size of image elements. This is depicted in Texture Scale-Space diagram by changing both dot size and dot density (because characteristic grain size is determined by spans of whitespace as well as dot size). The horizontal axis in Texture Scale-Space thus characterizes the image region along the shape/texture dimension: Simple coarse-grained forms whose configurations can be conveyed with relatively small amounts of information lie in the shape realm, while spatterings of fine-grained elements

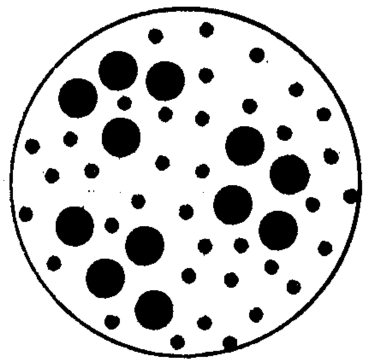

Figure 9: Texture elements in a region may be uniform or vary in their attributes.

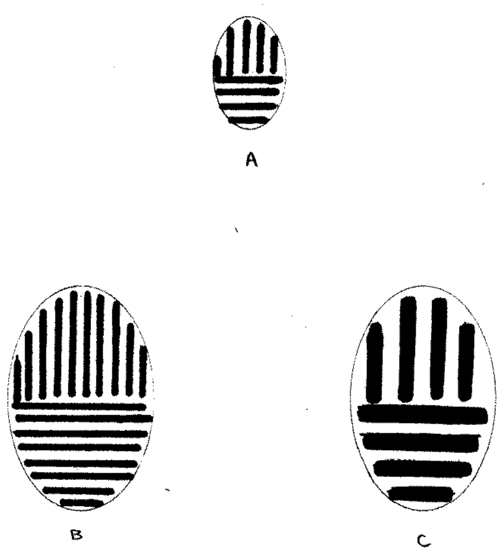

Figure 10: Which is more similar to a, b or c?

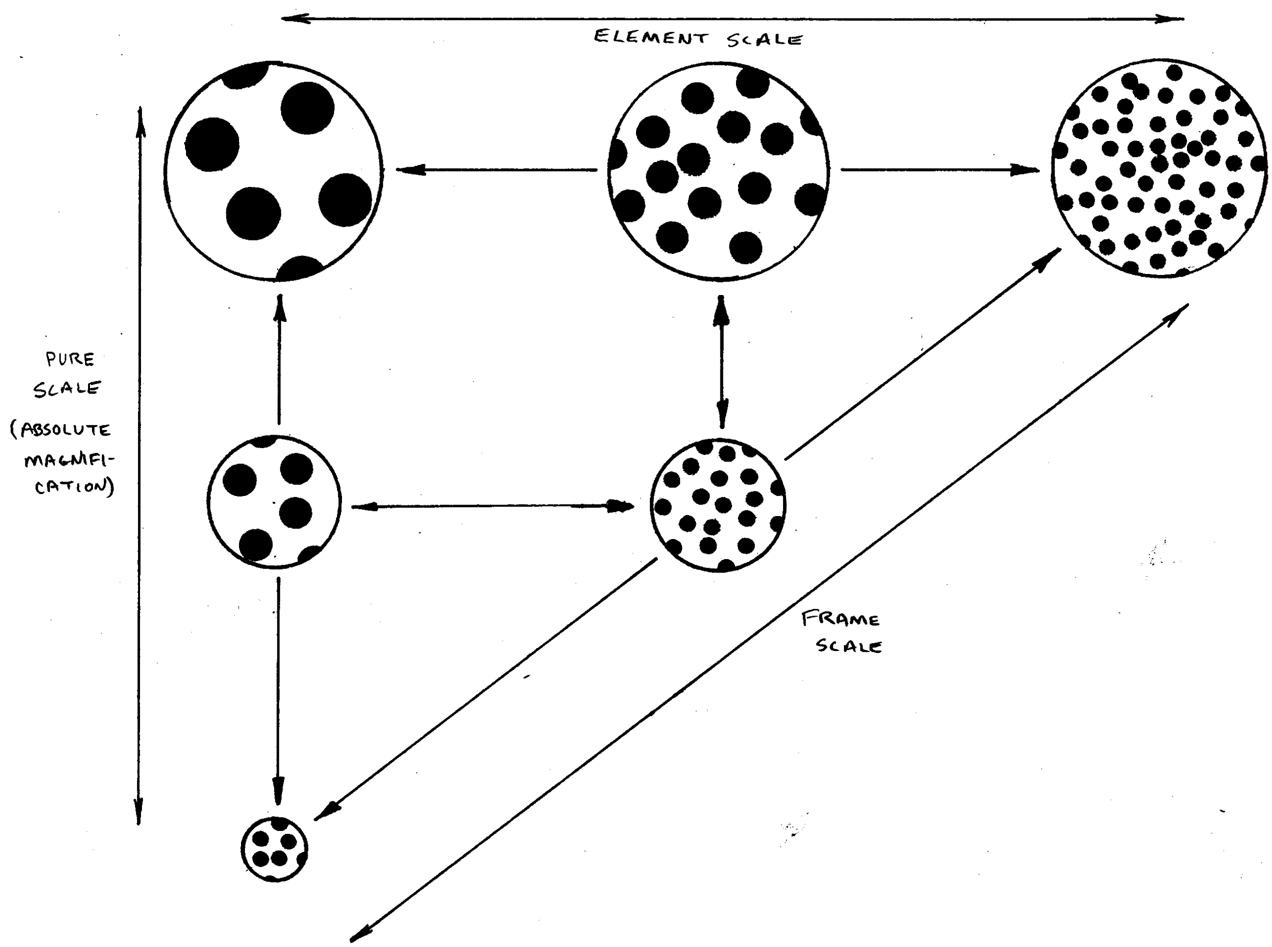

Figure 11: Texture Scale-Space

whose useful characterization is in terms of aggregate properties instead of a detailed accounting, fall in the texture realm.

2.4 Frames of View

Visual Neurophysiology entertains the notion of a neuron's *receptive field*, the region of visual space over which image stimuli influence the response of a cell. By analogy, in the Shape/Texture Attribute Chart, circles depict a "frame of view" on the input within which are found the primitive image elements whose attributes and configuration are to be described. The receptive field model in Neurophysiology helps to frame the question of what information a cell computes, by way of delineating its input domain. It is well recognized, however, that the abstraction requires deliberation in its use, as when the apparent receptive field changes with an animal's task. Similarly, the image domain of view of an effective shape or texture descriptor need not be a fixed circle as depicted in the chart, but may more profitably conform itself in accordance with input data or task demands. For example, Figure 12 shows a curvilinear line that can be identified explicitly, i.e. labeled, as a coherent and salient entity amongst cluttered and undifferentiated surroundings. The circular frame model should be viewed not as a prescription for some sort of receptive field, but as an indication that certain input data is encompassed by the descriptive label, and other is excluded.

A texture label is likely to be most useful when the information contained in its parameters and by virtue of its very assertion offers maximal conformance between observed data and the idealized model implied by the label. By explicitly demarking a frame of view for an image shape or texture descriptor, one could mean either of two things. The frame of view could indicate a *region of support*, a portion of the image from which information is drawn to compute the label's value. Alternatively, the frame of view could indicate *the region about which the label makes an assertion.* The distinction is subtle but important. The former meaning is about how the label is computed, the latter is more concerned with how it is to be used later.

The tradition of "feature detectors" in computational models of perception further accustom our thinking to the idea that at least at some stages information is conveyed by relatively independent channels; the meaning of a label does not vary with the context in which it occurs. While we do not know to what degree natural perceptual systems obey this principle, representation becomes so much more difficult to contemplate when it is violated, that we often consent to accept the premise for purposes of study.

Labels should accurately reflect the region location, size, and shape, plus

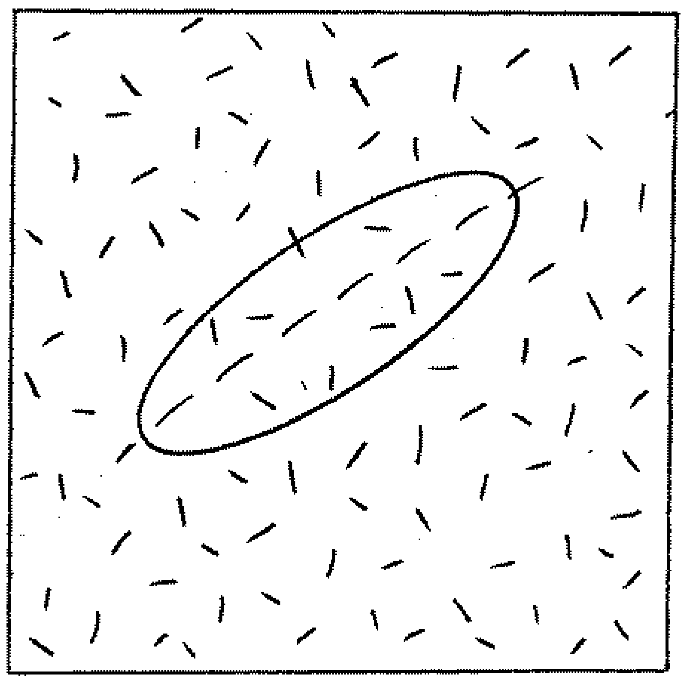

Figure 12: Sensible frames of view for texture labels need not be circular.

auxiliary properties and parameters, corresponding to causal phenomena in the world. But balanced against this ideal, the representation should summarize concisely the information most likely to be significant elsewhere in the system by limiting complexity to generalize over unneeded variation and detail.

2.5 Region Coherence

The assertion that an image region has this or that texture quality becomes most meaningful when the region presents something of a coherent appearance, and likewise, a texture boundary emerges only when locations compared across the boundary are somehow more different than locations compared within regions on either side of the boundary.

A default starting point for considering placement of labels for shape-like and texture-like image structure is geometrically compact, i.e. roughly circular, regions. Figures 13 and 14 suggests that certain placements of frames of view seem to be more natural or meaningful than others. In the shape realm, a label that "centers" the image event of, say, a T-junction, with respect to the label's location coordinates will support more accurate computation with respect to the junction's spatial relations to other image events. Note that scale of frame size is a degree of freedom in this regard as well. Natural frame placements are zoomed in to include just the image data fitting the descriptor's shape or texture event model.

In the texture realm, a key feature of sensible region labels is uniformity of the label property over the frame of view. If a descriptor declares event density, then a frame placement over a region of uniform density provides a meaningful result where a placement over an area of heterogeneous density does not. Similarly, a label identifying a texture boundary seems most aptly placed when the edge is centered in the frame of view. A descriptor whose job it is to label lopsided events such as Figure 14f is intuitively strange, probably because of our valid implicit beliefs that explicit characterization of this "mixed" area is less ecologically informative than either a patch containing only one type of texture exclusively, or the true border between two textures.

2.6 Odd Men Out

Sometimes a texture region is uniform in an attribute *except* for one local subregion. Often the exception subregion is ecologically significant in reflecting an interesting object or feature distinguished in semantic value from the surroundings. An appropriate representation might therefore make explicit such

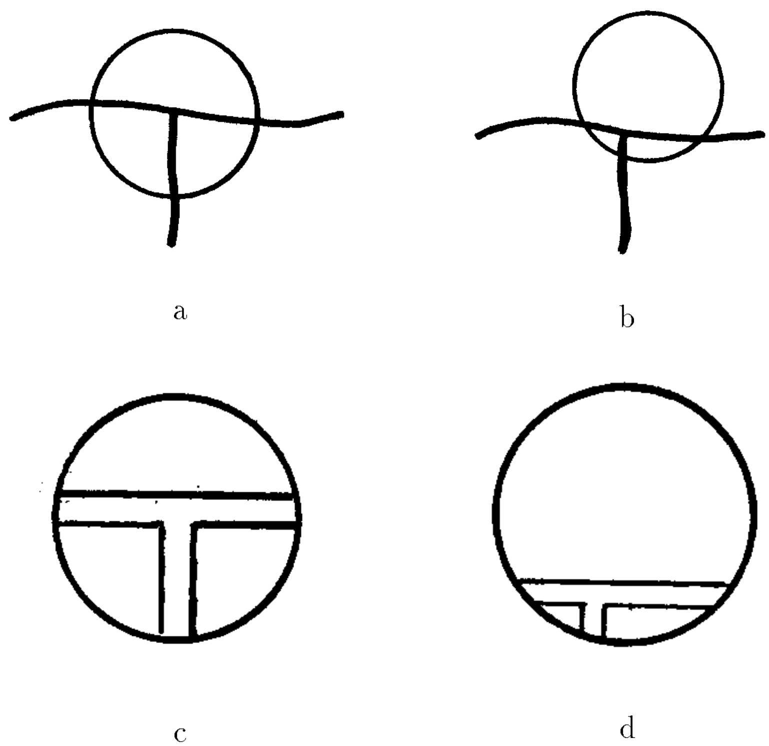

Figure 13: Natural (a) and unnatural (b) placements of a frame of view with respect to a T-junction event occurring in image data. c and d reflect the ideal event model that would justify each of these placements.

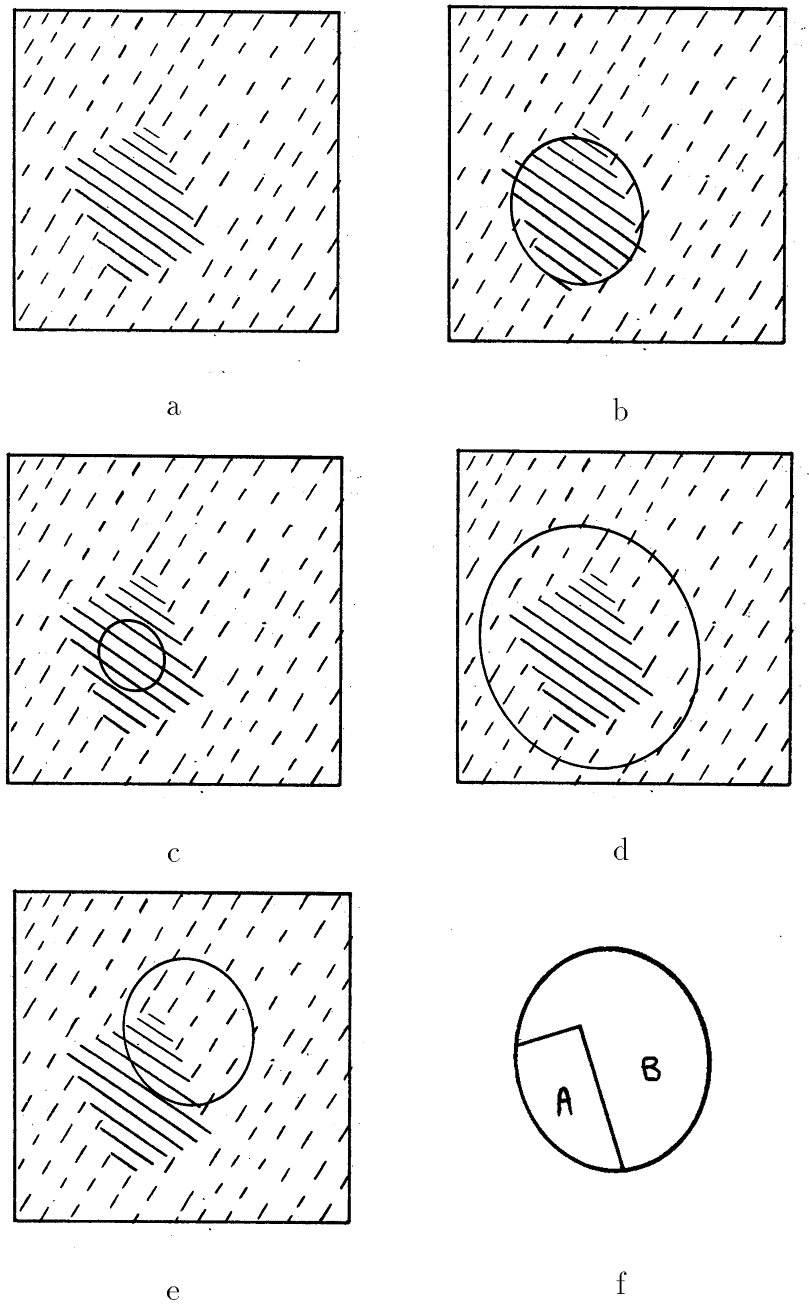

Figure 14: Image data (a) and a sensible frame of view (b) for labeling the differentiated patch. c and d reflect inappropriate scaling of a uniform region model frame. e. Frame placement reflecting nonuniformity in texture element attributes. f. Required texture model that would justify frame placement in e; note the strangely shaped boundary between subregions A and B associated with different element attribute values.

an *odd-man-out* event with its own label, in effect subtracting it away, yet preserving the ability to concisely describe the background as a whole by labeling the remaining region as if the intruder weren't there.

Odd-man-out detection and visual interest operators are intimately tied with dynamic processes of visual attention, which falls beyond the scope of this discussion. Of particular relevance here however are questions about the spatial scales at which odd-men-out are detected and reported.

Figure 15 shows how odd-man-out status interacts with region uniformity across scales. A small region of interest frame focused only the odd-orientation region possess high uniformity in its orientation attribute, and is thus well-suited to support a descriptive label. Zooming back a ways, the region of interest contains a mixture of orientations not easily summarizable. Zooming back more, the region of interest again becomes uniform with respect to orientation, with the exception of the odd-man-out subregion.

Odd-man-out phenomena are not limited to spatially compact outliers. Figure 12 showed a curve that has odd-man-out status with respect to a relatively uniform texture background.

The next two stages in this discussion would address the following. First, we are now equipped to examine characterization of texture boundaries across scales, for example when a texture boundary is blurry or wiggly. Second, the issue of *grouping* processes cuts across the shape and texture realms, and are intimately concerned with label assertions.

The next two sections briefly describe two experiments exploring these ideas through computational models. The first is based on the filter/energy computing paradigm, the second on symbolic linking and value passing.

3 Experiments: Filter-Based Uniformity in Texture Scale-Space

3.1 Filter-Based Uniformity in Texture Scale-Space

This first experiment demonstrates the interplay between two kinds of scale whose relationship reflects the continuum between the shape and texture realms of image events. Using a traditional filter-based paradigm for computing texture properties, we build a two-dimensional space of image maps (texture scale-space) where the degree to which primitive image events in a certain size range occur uniformly over a frame of view is indicated by response strength or "energy" in an image map.

An experiment was carried out using as image data the cereal box of Figure

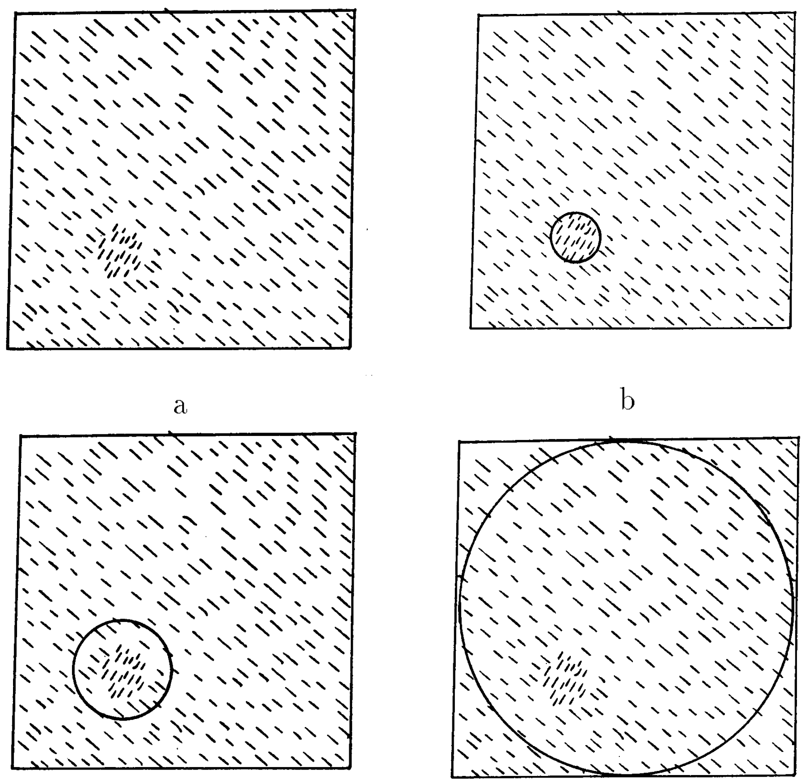

a b c d

Figure 15: a. Region containing an odd-man-out subregion. b. Appropriate frame of view for labeling odd-man subregion properties. c. Inappropriate frame of view. c. Appropriate frame of view for labeling large region, especially if the odd-man subregion is signalled.

16a, in two stages.[2] The first stage was to compute the amount of image energy occurring at various grain sizes throughout the image. Since the geometric property of image element size is independent of image contrast, it would be inappropriate to use a purely linear filter for this measurement. Instead, a contrast normalization technique was used. Figures 16b-d shows image energy at three sample grain sizes. The second stage was to consider frames of view of different sizes. This stage proceeded under the assumption that a texture label is meaningful only when its support uniformly reflects the image structure the label claims to model. Thus, at the second stage a measure of uniformity of grain size energy was applied over each frame of view. For efficiency, pyramid methods were used involving successive layers of blurring and subsampling.

Results are shown in Figure 17. Figure 17a shows energy deep in the texture realm, where response strength reflects small image elements over a large region. Figures 17b through d show responses in the shape realm, where frame size is only somewhat larger than elemental grain size. Figure 17b shows large scale shape realm structure, Figure 17c shows medium scale shape realm structure, while Figure 17d shows small scale shape realm structure. Note that all of these highlight regions of printed text commensurate with the scale of the label.

Because the results of this processing are a set of energy maps, this experiment does not demonstrate computation of concise image texture labels as such. An additional stage would be required to analyze the energy maps to arrive a symbolic tokens or other some other more concise representation.

3.2 Fine-to-Coarse Aggregation of Texture Information

This second experiment aims to explore a strategy of building texture descriptions across multiple scales by aggregating local neighborhood information at successively larger scales. In common with the previous experiment, we develop a notion of grading the salience or aptness of a texture descriptor by virtue of attribute uniformity over a support region. Unlike the previous experiment, the style of computing is not to perform linear and nonlinear transformations of continuous-valued energy maps, but to propagate symbolic information among a sparse array of sample points. Also addressed in this experiment are notions of dynamically defined feature channels, and odd-man-out detection.

As we have discussed, image descriptors in the texture realm summarize attributes of small scale events over relatively larger regions. The simplest and

[2]This experiment was performed with the aid of a CM-2 Connection Machine.

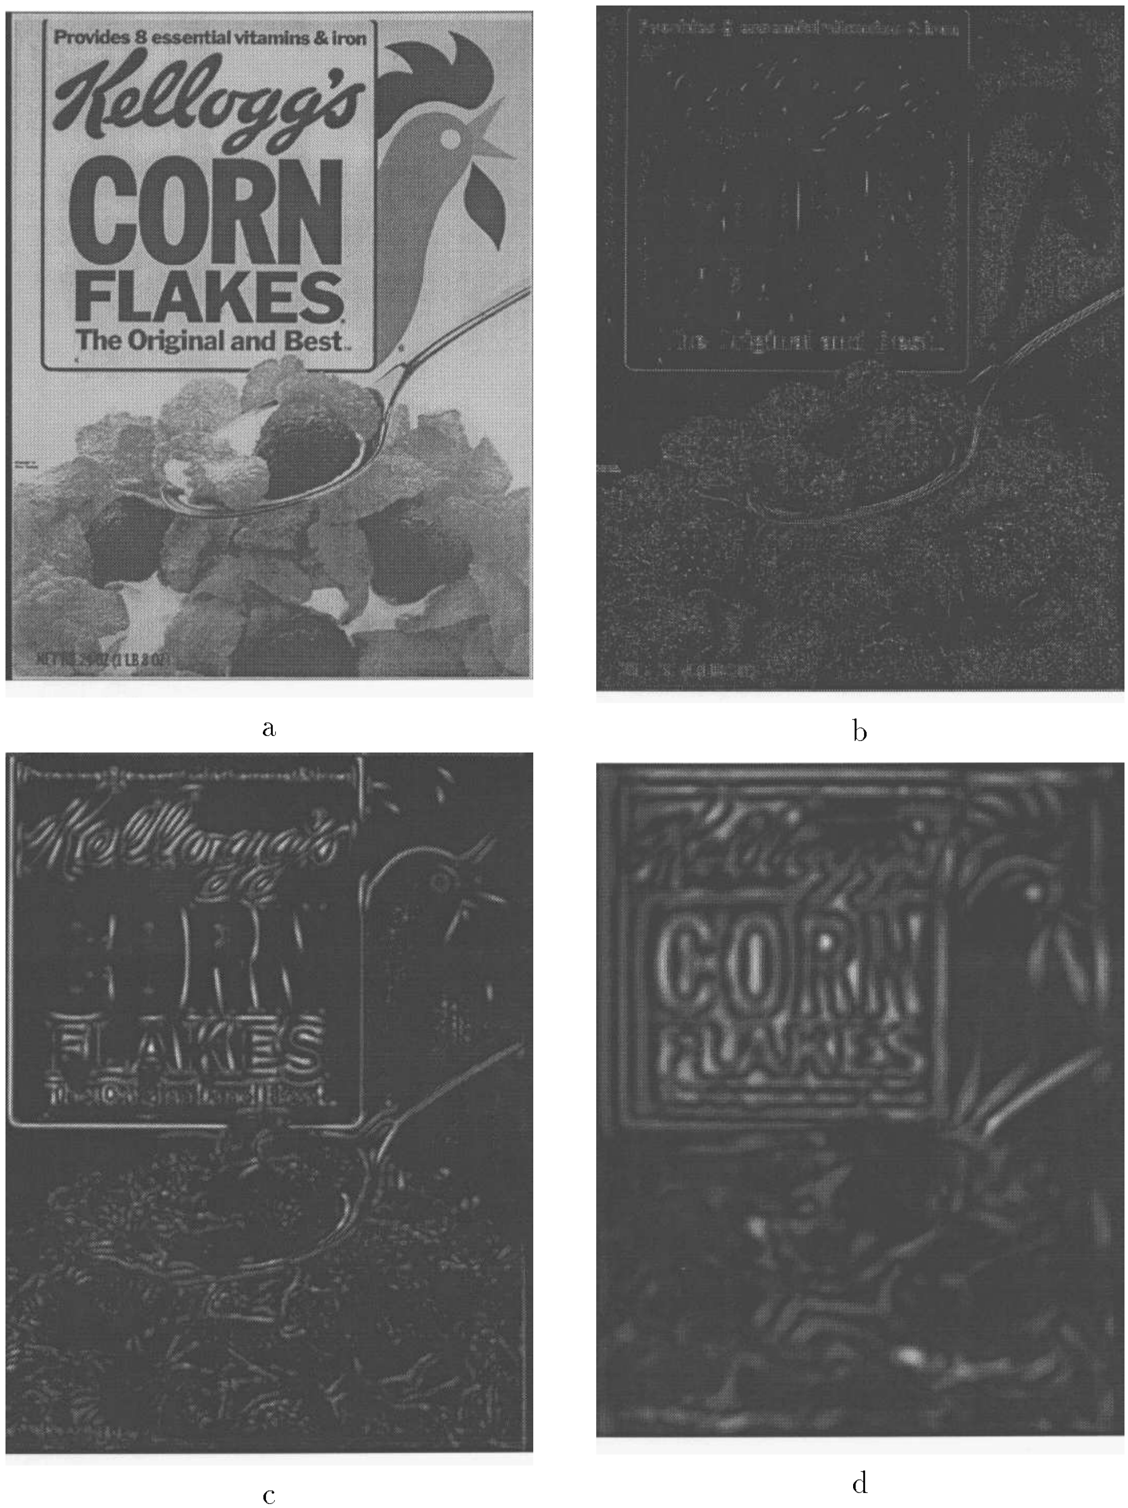

Figure 16: a. Image used in Experiment 1. b, c, and d. Texture element grain size “energy” at small, medium and large scales.

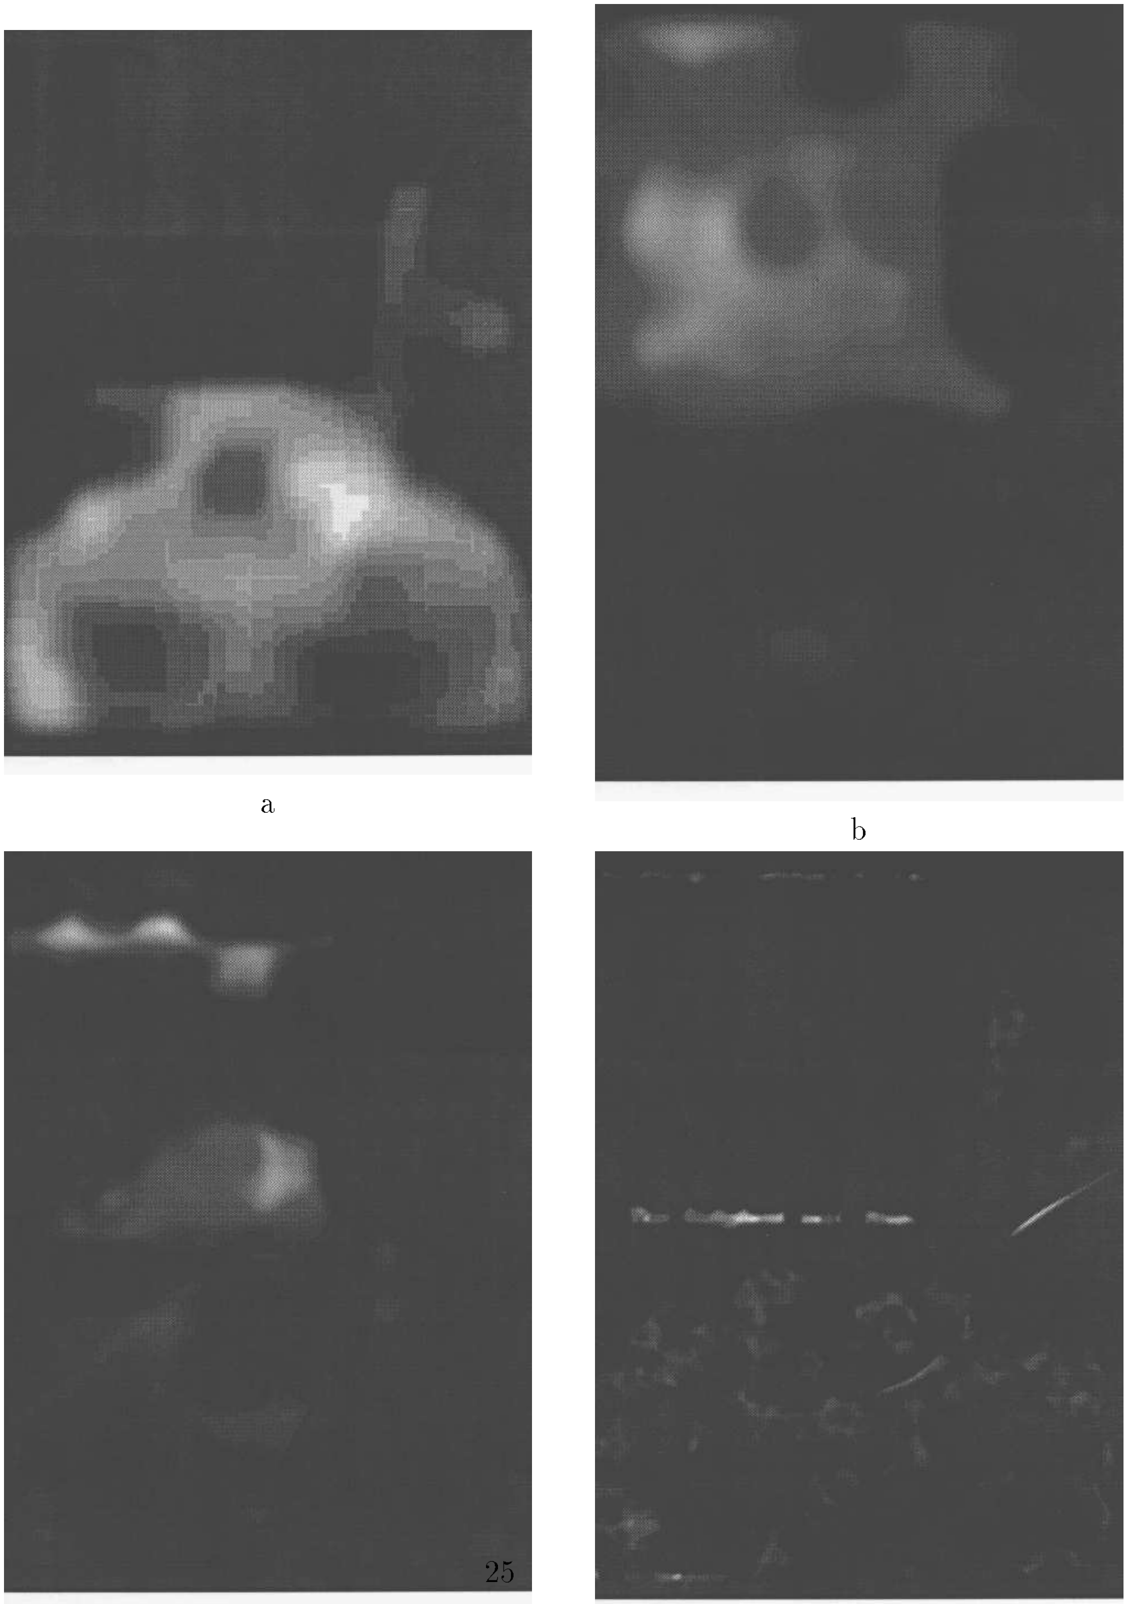

a b c d

Figure 17: Region uniformity "energy" at various points in texture-scale-space. a. Texture realm (large frame, small grain size). b. Shape realm (large frame, large grain size). c. Shape realm (medium frame, medium grain size.) d. Shape realm (small frame, small grain size).

most familiar form of this operation is generalized image blurring, where the primitive event is image intensity and the summarization method is weighted averaging (e.g. Gaussian, Smoothed Gradient, Laplacian of Gaussian). Any blurring kernel whose weights vary slowly over the support region delivers output values which can vary only slowly across the image, and subsampling in the output space is typically used for sake of efficiency. Pyramid methods are used to build hierarchies of successively coarser scale samplings of kernels with successively larger image support.

One objective of Experiment 2 is to design efficient layered methods for aggregative computations other than weighted averaging. The information conveyed by averaging methods is the *amount* of something–the amount of lightness, the amount of energy at $45°$ and $6cpd$, the amount of agreement between image data and a given convolutional template, etc. An alternative representation deals explicitly in the *values* and *distributions* of parameters. Figure 18 illustrates. The representation of texture orientation is often modeled as distribution or histogram of responses to tuned filters, mimicking the response tuning curves of certain neurons. The phenomenon of texture metamers suggests that at some point in the human visual system at least this detailed information gets reduced into a much smaller number of independent "channels".

One form of this transformation is the expression of region aggregate attributes in terms of mean, variance, and possibly higher order moments of the distribution of parameter values. Thus a channel doesn't carry "energy," but numerical values. A representation of attribute distributions in terms of mean and variance of one or more modes of the distribution is particularly convenient for local-to-global scale propagation of information because these can be computed for a large scale distribution from the values for component smaller scale distributions. Each mode is then analogous to an independent channel for that attribute, up to some maximum distinguishable number of modes. This sort of representation would reflect the fact, for example, that a bipartite oriented texture is perceived as either a single region or two disparate regions depending not only on the mean orientations, but on the variances as well.

A local to global scheme for propagating aggregate texture information demands a layout of support cells covering the 2d image space at multiple scales, and satisfying several constraints. First, it must be repeatable across multiple spatial scales. Second, it should provide appropriate spatial sampling. Third, it should support analysis of the uniformity of a property over a region purely by examining summary information obtained at subregions. A layout satisfying these constraints is illustrated in figure 19. The basic idea is that at

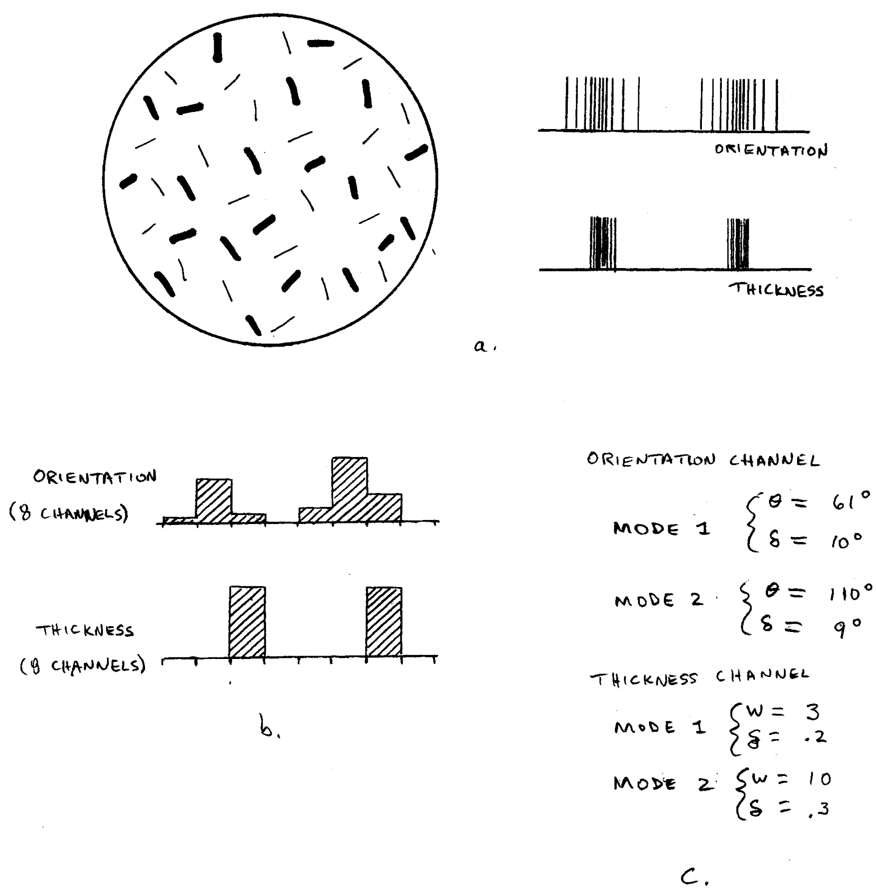

Figure 18: a. Image data in a frame of view, along with detailed accounting of orientations and thicknesses of texture elements. b. "Energy" channel representation in terms of the number of events within each bin of a histogram. c. Numerical channel representation in terms of mean and variance of principal modes in each distribution.

each scale region cells occur in three overlapping fields. The support of a cell is imagined to be a circular region. Circular regions in a field are packed with centers on a hexagonal grid. At a given scale, scale S_i, three such fields overlap one another in staggered fashion (three phases). An array of larger scale cells at scale S_{i+1}, computed from these employs regions three times in diameter. The support for aggregating parameter means and variances consists of seven regions as shown. Subsampling occurs by choosing one S_{i+1} cell for every nine S_i cells in a field, staggering the three fields appropriately.

Although parameter mean and variances for a cell at scale S_{i+1} are computed from means and variances from seven cells occurring in a single field at scale S_i, the other fields are valuable in computing other information. In particular, a total of nineteen S_i cells falls within the support region of a single S_{i+1} cell as shown in Figure 19c, providing a wealth of sample locations for computing uniformity and gradient characteristics of a parameter over the region. I have experimented with a number of uniformity measures; one of the most effective is based on computing an edge measure in three different directions, and asserting uniformity when no edge is found.

Input data in this experiment is discrete events, black bars with variable orientation and thickness. In Figures 20 and 21 test images were synthesized by placing objects varying in orientation or size in a field. Figures 20a and b and 21a and b show uniformity measures at two scales. Darker stippling indicates greater uniformity. Values of means and variances in orientation or element size are also maintained at each cell, along with event counts used in the formula for combining means and variances. The uniformity measure can be viewed as an indicator of the validity of parameter mean and variance as providing a meaningful aggregate description of a cell's support region with respect to that parameter. Odd men out are handled naturally, because at an intermediate scale (corresponding to Figure 15c) uniformity is low, while a zoomed-back field of view contains high uniformity except for one subregion signifying the outlier.

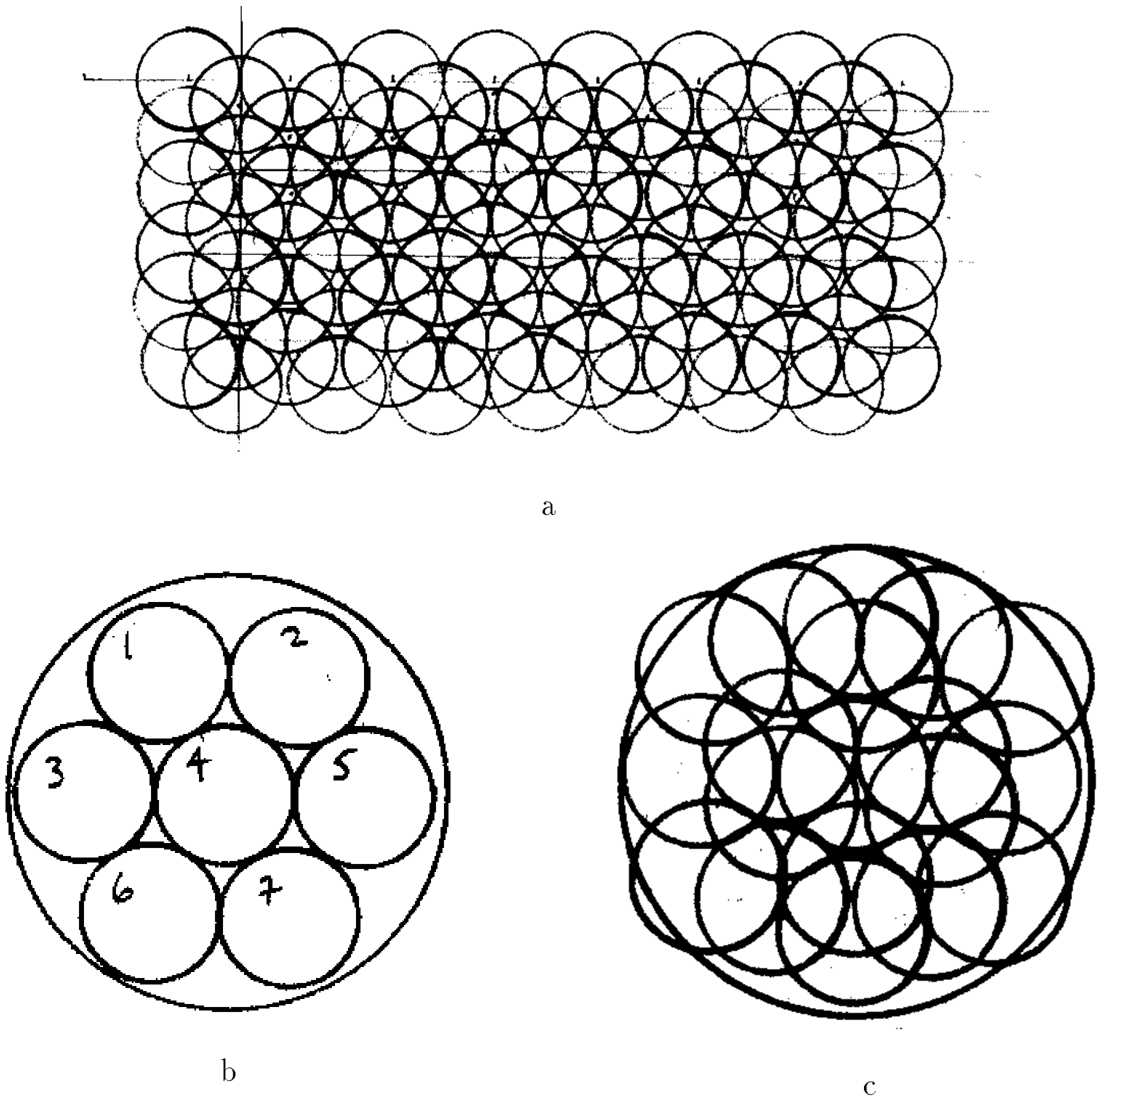

Figure 19: a. Field generated by overlapping three phases of circular frames of view. b. Centering of S+1 scale region over one phase. c. The nineteen frames offering support for texture analysis of an S+1 scale frame.

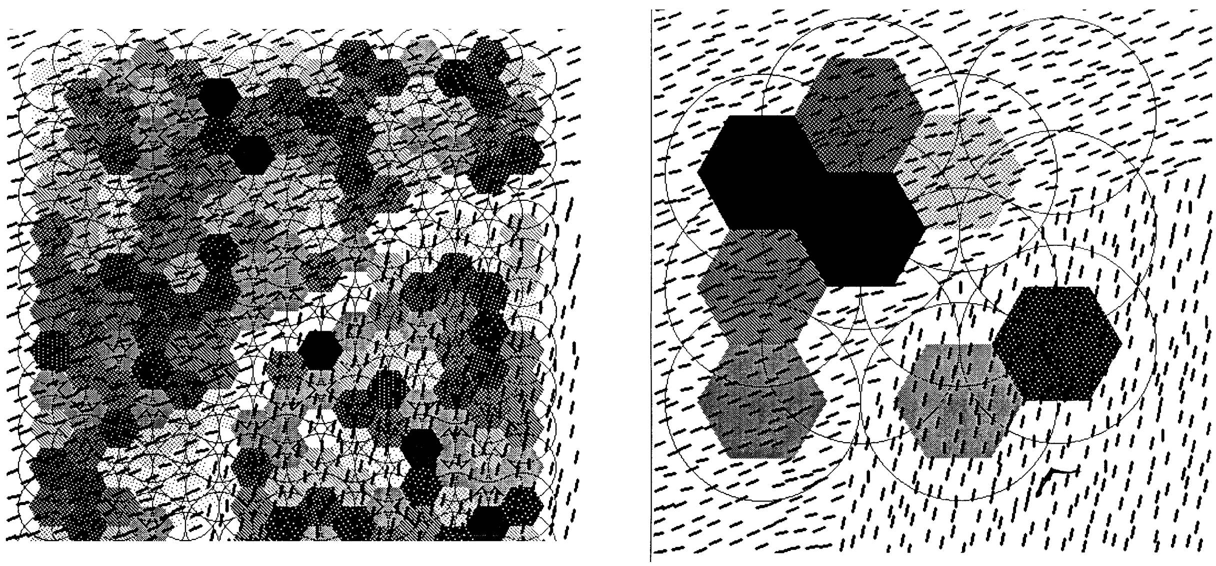

a

b

Figure 20: Uniformity measures for two scales of analysis of element an orientation channel. Darkness of hexagons indicates uniformity of the circular frame centered on that hexagon. Result is overlain on original image data. In addition to uniformity, each frame maintains orientation mean and variance for its support region (not shown).

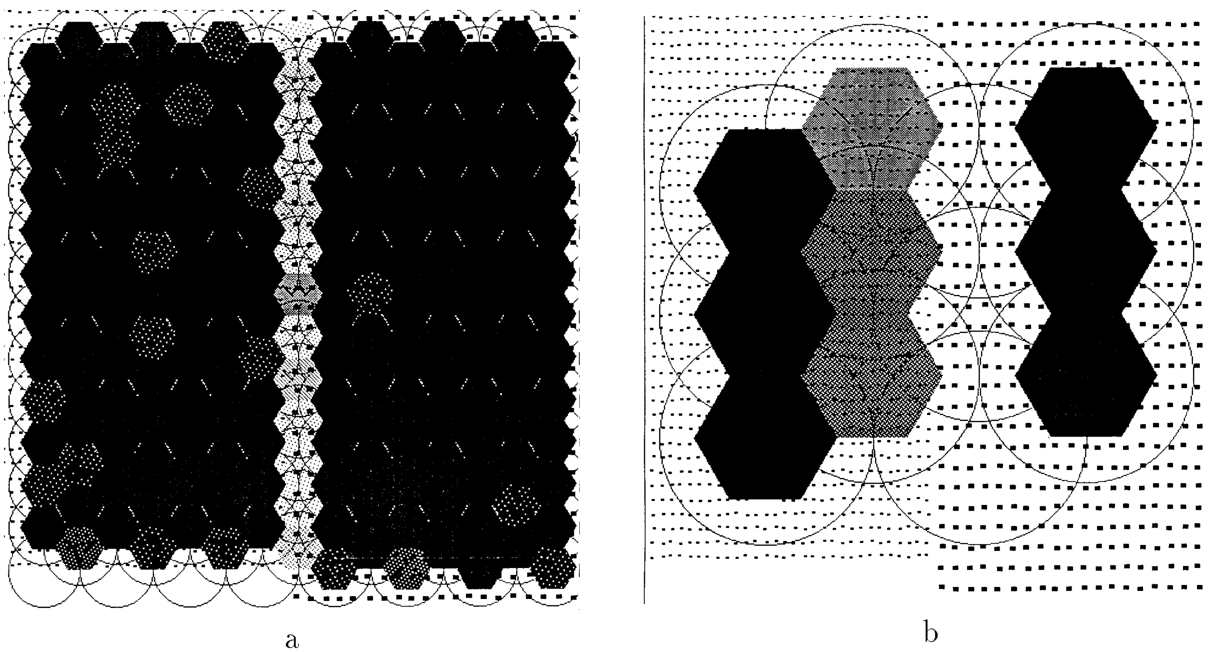

Figure 21: Uniformity measures for two scales of analysis of element a size channel. Darkness of hexagons indicates uniformity of the circular frame centered on that hexagon. Result is overlain on original image data. In addition to uniformity, each frame maintains size mean and variance for its support region (not shown).