CMU-CS-78-152

The Computation of Immediate Texture Discrimination

Bruce R. Schatz

Computer Science Dept.
Carnegie-Mellon University
Pittsburgh, PA 15213 USA

December, 1978

Abstract

One approach to understanding human texture discrimination is to consider what image properties should be examined by a system which computes it. This is investigated in a domain of line drawings. It appears that the orientation and length of lines should be examined. However, there exist textures in which these properties are identical yet discrimination occurs. Consideration of virtual lines (imaginary lines between special points) is introduced to explain such textures. The relation of this method to that of Julesz is discussed with the conclusion that it is strictly less powerful. This lack of sensitivity appears to be psychologically correct; instances are given of indiscriminable textures with different second-order statistics but identical (equivalent) actual and virtual lines. An implementation of the computation is sketched.

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1551.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Introduction

Suppose one wanted to understand how a system, such as the human visual system, makes immediate distinctions between textured regions. One method of accomplishing this is to consider how to design a procedure that computes the region boundaries. In particular, what should go into such a procedure?

The research described below attempts to determine, by examining simple textures, what attributes of an image a computation of human immediate texture discrimination should be sensitive to. It could also be considered as an attempt to determine the psychological properties that characterize texture discrimination, with the ultimate objective of finding a necessary and sufficient set of properties for its computation (or at least properties equivalent to such a set).

In this paper, the problem will be considered in a limited domain. The textures will be those generated by repeating an element regularly (with small perturbations) over a large area. Experiments will consist of examining a texture generated by one element embedded within another texture generated by a different element. The textures will be considered discriminable (and the generating elements called "different in the texture case") if the shape and location of the inner region can be perceptually distinguished within 200 milliseconds. Although it forces a somewhat arbitrary binary decision, this time limit prevents use of directed eye movements. The limit is a useful restriction since allowing the scrutiny provided by such saccades would require a computation which can distinguish between arbitrarily similar but non-identical texture elements. Elements which are the same in the immediate texture region case in fact often look quite different when compared directly (c.f. Julesz[1975],Beck[1966]). Figure 1 is an example of discriminable textures. As in this example, a number of further restrictions will be made.

The elements will be composed of lines and points; the domain thus corresponds to that of line drawings. A single identical element will be used throughout generation of a given texture. Textures will be black/white, no motion will be allowed, and no depth information will be used. Only differences in texture elements will be considered. The spacing between elements will be large to minimize interaction effects.

With the above definition of immediate texture discrimination, the properties which characterize it will now be investigated.

Properties affecting discrimination

What properties of the image affect discrimination?

Orientation differences can cause strong discrimination. See figure 1 for example. However, these differences must be sufficiently great since, empirically, lines with orientations closer than 10 degrees apart (and similar lengths) cannot be distinguished in the texture case (c.f. Campbell[1966]). This is true except that parallel lines are unusually distinguishable. For example, a texture with lines all of one orientation θ can be distinguished from a texture with lines of several orientations varying slightly from θ.

Length should be considered as well, since length differences can cause discrimination even when the orientations are the same. This can be seen in figure 2, where the lines in the inner texture are twice as long as those in the outer. In this domain, the class of length differences encompasses the first-order intensity discriminations, i.e. those due to a difference in the density of points between the texture regions. These first-order distinctions might be considered to consist of two types: (1) the two regions have the same density of lines but differing lengths of lines (as in figure 2), and (2) the two regions have differing density of lines, this could be termed spatial frequency difference. Examples also exist of discriminable textures with identical point densities but differing length distributions (see Schatz[1977,2.2]). In examining the other classes of differences (e.g. orientation, virtual lines), no first-order intensity distinctions will be permitted.

These two properties are easy to measure and compute, and one might thus propose that texture boundaries could be found by a process similar to the following. Tabulate a two-dimensional histogram of orientation and length within windows of appropriate size. Where the histogram differs sufficiently between adjacent windows in the image, assert a boundary. However, this (first-order) strategy is not sufficient to predict discrimination.

For consider figure 3. The orientation and length of the lines is the same in both generating elements, yet the textures are discriminable. This appears to be due to the fact that the lines are arranged differently. As suggested by such textures as figure 4, one possible solution is to introduce <u>virtual lines</u>, imaginary lines filled in between points which behave as though physically present (albeit somewhat more weakly). Thus, in figure 3 the outer generating element has a virtual line at 135 degrees and the inner one at 45; this difference can explain the discrimination.

This proposed use of virtual lines is one technique for describing the relative positioning of lines and points (c.f. Stevens[1977], Attneave[1974]). As the primary properties of lines are orientation, length, and position, something would seem to be needed to measure the effects of position on discrimination. Virtual lines are not the only possible explanation of discriminable textures generated by elements with identical orientation and length of lines, but they seem to provide the simplest consistent hypothesis for the textures considered thus far. In addition they possess other desirable properties. A number of other possibilities, such as using the overall orientation of the texture elements, appear not to be psychologically tenable (see Schatz[1977,2.7]). The references in the literature to the effects of slope and arrangement (e.g. Beck[1972,1967],Olson & Attneave[1970]) could accordingly be replaced by references to (effects of the orientation of) actual and virtual lines. It is the case that, as with slope and arrangement, actual lines produce stronger discrimination than virtual (compare figure 1 to figure 4).

In order to maintain a feasible computation, the virtual lines considered must be restricted to a small finite subset. First, theoretically any two points in the image could be connected. So a locality assumption is made, allowing only <u>local</u> virtual lines (those between points within some small neighborhood of each other) to be inserted. Second, between a line and a point there are an infinite number of possible connections. A restriction to some sort of special points seems reasonable, and the logical choice is the <u>terminators</u>. These delimit the ends of lines and thus include isolated points, endpoints of lines, and corners. Figures 4 and 3 are examples of the first two; corners are included primarily for consistency (the end of a line is a terminator whether or not it is connected to another line) but do provide additional useful explanatory power (see Schatz[1977,2.3]).

One set of properties, then, which seems to be needed to compute immediate texture region discrimination (in this domain) is:

(1) length and orientation of lines, and
(2) length and orientation of local virtual lines between terminators.

Use of only these properties suffices for a large number of instances. These include all of the hundreds of textures examined or considered thus far by the author. However, evidence from other domains of texture indicates that additional properties, such as special consideration of parallel lines, may also be needed.

Relation to the work of Julesz and of Marr

The most comprehensive theory of texture discrimination at present is that of Julesz[1973]. He conjectured that two textures are indiscriminable if their second-order statistics are the same. Informally these statistics can be calculated by considering the "dipoles" (possible lines between points) as follows. Choose an orientation and a length, and find the probability that a dipole (vector) of that orientation and length touches a black point on both ends when randomly placed on one texture. Do the same for the other texture and compare. If the same probability is obtained in both textures for all orientations and all lengths, then the second-order statistics are the same. For example, two generating elements which are 180 degree rotations of each other have identical second-order statistics, and in fact produce indiscriminable textures (see figure 5). Julesz,et.al.[1973,1975] describe how, and Gilbert and Shepp[1974] prove that, this and several other transformations preserve second-order statistics.

The properties proposed here are sensitive to a proper subset of the dipoles: all those corresponding to maximal actual lines (i.e. not including subsegments) and those corresponding to virtual lines locally drawn between terminators. This implies that discrimination based on the properties here is strictly less sensitive than discrimination based on full second-order statistics. Consider figure 6. The actual and virtual lines are identical and the textures are indiscriminable, yet the second-order statistics differ. That this is true can be seen by the following argument.

All the elements consist of an "X" with a horizontal line through the center and two vertical lines joined to this. The actual lines are then clearly the same throughout. In the outer element, the vertical lines are in the lower left and upper right partitions, while in the inner they are in the upper left and lower right. The virtual lines are then the same in both elements since the only difference between the elements is in the vertical lines and these are connected to existing terminators, thus creating no additional points to join virtual lines to. (And two vertical lines are changed from virtual to actual in either case.) However, the dipoles differ, as can be seen by connecting points on one vertical line to those on the other within an element; this produces dipoles slanting to the right in the outer element and ones slanting to the left in the inner. This dipole difference is due to the fact that second-order statistics does not restrict dipoles to connection of special points.

More recent work (Caelli and Julesz[1978]) has uncovered a few examples of texture element pairs which are clearly discriminable despite possessing identical second order statistics. Julesz has accordingly modified his theory to include a small set of detectors which account for discrimination between such textures (termed "class B") as well as the original dipole detectors described previously ("class A"). The approach presented here can thus be seen as a (more precise) mechanism for computing class A texture discrimination, and must be augmented by a few special feature detectors (e.g. colinearity,closure) to provide a complete explanation for texture discrimination in this domain.

Marr[1976] argued for a process-oriented explanation of texture vision. He gave a set of primitives (the "primal sketch"), claimed to represent the first stage in the vision process, and conjectured that first-order distinctions on some preliminary grouping of these would be sufficient to predict discrimination. These primitives from the image include lines, edges, and blobs, described by their orientation, length or size, position, termination points, and local contrast. Thus the only information available from the image consists of these basic primitives and simple groupings on them ("place-tokens"). In the texture computation here, the terminators would be the place-tokens that are grouped together; the groupings include virtual lines but not necessarily texture elements. The work described in this paper could be considered an attempt to make a feasible, finite computation from the dipoles suggested by Julesz (by culling out the "inessential" ones) using as input the primitives suggested by Marr.

Based on the above, the complexity of texture discrimination could be summarized as in figure 7. As shown in the chart, in statistical terms discrimination appears to require more than first-order (actual lines only) but less than full second-order (full dipoles needed). As noted above, a few special feature detectors may also be needed.

Implementation

A straightforward implementation strategy consists of comparing histograms in adjacent windows and declaring a boundary wherever there is a sufficient change. As the windows are calibrated to be roughly the size of an element, this essentially compares adjacent elements. The lines are put into orientation buckets (equivalence classes) every 10 degrees with widths of ±10 degrees. A coarse length histogram is calculated for each bucket. Corresponding histograms are compared and the resulting

$(\sum \text{Difference}) / (\sum \text{Total}) > \text{Threshold}$ calculated to check for a boundary. Here, for each term, "difference" refers to subtraction of the length histograms of the two buckets (summed) and "total" to the total length of lines in the two buckets.

This ratio and threshold is needed since the elements must differ by a sufficient amount to cause discrimination. This is important in such textures as figure 8 where the virtual lines (and hence the full dipoles) differ, but not by enough to cause discrimination. A threshold on the difference required for discrimination is relatively easy to implement in the computational model presented here, but is hard to fit into a statistical model. For example, the implementation correctly smooths out small bumps in lines (the leeway allowed in the orientation and length equivalence classes causes them to be ignored) while the bumps do cause the dipoles to differ (second-order statistics are too sensitive since an exact match is required between dipoles for two textures to be identical). This additional sort of control is one reason why a computational theory such as proposed here may be preferable to a phenomenological one such as Julesz's. It also lends support to the contention that consideration of the essential properties, and their embodiment in a computation, may be more informative than statistical or other descriptive observations.

The current implementation computes discrimination for actual and for virtual lines separately. However, they should probably be considered together, with the actual lines weighted more heavily. This weighting is psychologically realistic (since as previously mentioned actual lines produce stronger discrimination) and is needed (since considering the sums of difference/total as consisting of equally weighted actual and virtual lines would declare textures generated by a "U" and a "C" respectively to be non-discriminable, which is not the case). Thus one could imagine the visual system perceiving a texture boundary whenover there is sufficient difference in the orientation and length of lines between regions, where the lines are a combination of actual and virtual, with the latter perceived more weakly.

To use the implementation on real images, a preprocessor such as that for the primal sketch must be used to find the "input line drawing". By using such other techniques as density calculation, it might be possible to extend the texture computation to other input primitives: edges and blobs, for example. Note also that filling in virtual lines between elements (which happens when the elements are close together) can often predict discrimination due to boundary effects (e.g. subjective contours).

The above computation in the domain of line drawings makes no mention of curved structures. These would raise a particular problem for the placement of virtual lines since the notion of terminator is poorly defined. One possible approach is to approximate smooth arcs by line segments drawn between points of local maximum curvature (as with linear splines). This may be quite effective, given that even crude approximations are perceptually very similar to smooth curves (c.f. Beck[1973], Schatz[1977]).

As discussed above, the proposed properties seem plausible psychologically and computationally. To test this approach to investigating texture discrimination, more detailed work needs to be done. For example, further psychophysical experimentation could determine the precise tuning curves of how much orientation or length causes discrimination, as well as the interrelationships between the effects of orientation versus length and of actual versus virtual lines. The implementation is still in a rudimentary state; problems ranging from the value of the threshold(s) to the composition and comparison of the histograms remain unsettled. Although the discriminations have not at present been tested in a rigorous experimental fashion, this paper has attempted to provide evidence that a fruitful approach to understanding immediate texture discrimination is consideration of methods for its computation, and that the properties of the image that such a computation should examine include orientation and length of actual lines and of local virtual lines between terminators.

Acknowledgements

This paper summarizes an M.S. thesis submitted to the Dept. of Electrical Engineering and Computer Science at M.I.T. More detailed discussions can be found in the revised version of the thesis (Schatz[1977]).

Shimon Ullman, Berthold Horn, David Marr, Kent Stevens, Whitman Richards, John Kender, John Fox, and Joseph Schatz provided helpful criticism and suggestions. Their comments on previous drafts substantially improved the content and presentation of this material.

The author was supported by a graduate fellowship from the National Science Foundation. This research was conducted primarily at the MIT Artificial Intelligence Laboratory supported in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

References

Attneave, F., 1974, "Apparent motion and the what-where connection", Psychologia, 17, 108-120.

Beck, J., 1966, "Effect of orientation and of shape similarity on perceptual grouping", Perception & Psychophysics, 1, 300-302.

Beck, J., 1967, "Perceptual grouping produced by line figures", Perception & Psychophysics, 2, 491-95.

Beck, J., 1972, "Similarity grouping and peripheral discriminability under uncertainty", American J. of Psychology, 85, 1-19.

Beck, J., 1973, "Similarity grouping of curves", Perceptual & Motor Skills, 36, 1331-1341.

Caelli, T. and Julesz, B., 1978, "On perceptual analyzers underlyzing visual texture discrimination", Biological Cybernetics, 28, 167-175, 201-214.

Campbell, F.W. and Kulikowski, J.J., 1966, "Orientational selectivity of the human visual system", J. Physiology, 187, 437-445.

Gilbert, E.N. and Shepp,L.A., 1974, "Textures for discrimination experiments", Bell Laboratories memo TM-74-1218-6.

Julesz, B., 1975, "Experiments in the visual perception of texture", Scientific American, 232, 34-43 (April).

Julesz, B., Gilbert,E.N., Shepp,L.A., and Frisch,H.L., 1973, "Inability to humans to discriminate between visual textures that agree in second-order statistics -- revisited", Perception, 2, 391-405.

Marr, D., 1976, "Early processing of visual information", Philosophical Trans. Royal Society London (series B), 275, 483-524.

Olson, R.K. and Attneave, F., 1970, "What variables produce similarity grouping?", American J. of Psychology, 83, 1-21.

Schatz, B.R., 1977, The Computation of Immediate Texture Discrimination, MIT Artificial Intelligence Lab memo 426.

Stevens, K.A., 1977, "Computation of locally parallel structure", MIT Artificial Intelligence Lab memo 392.

Figures

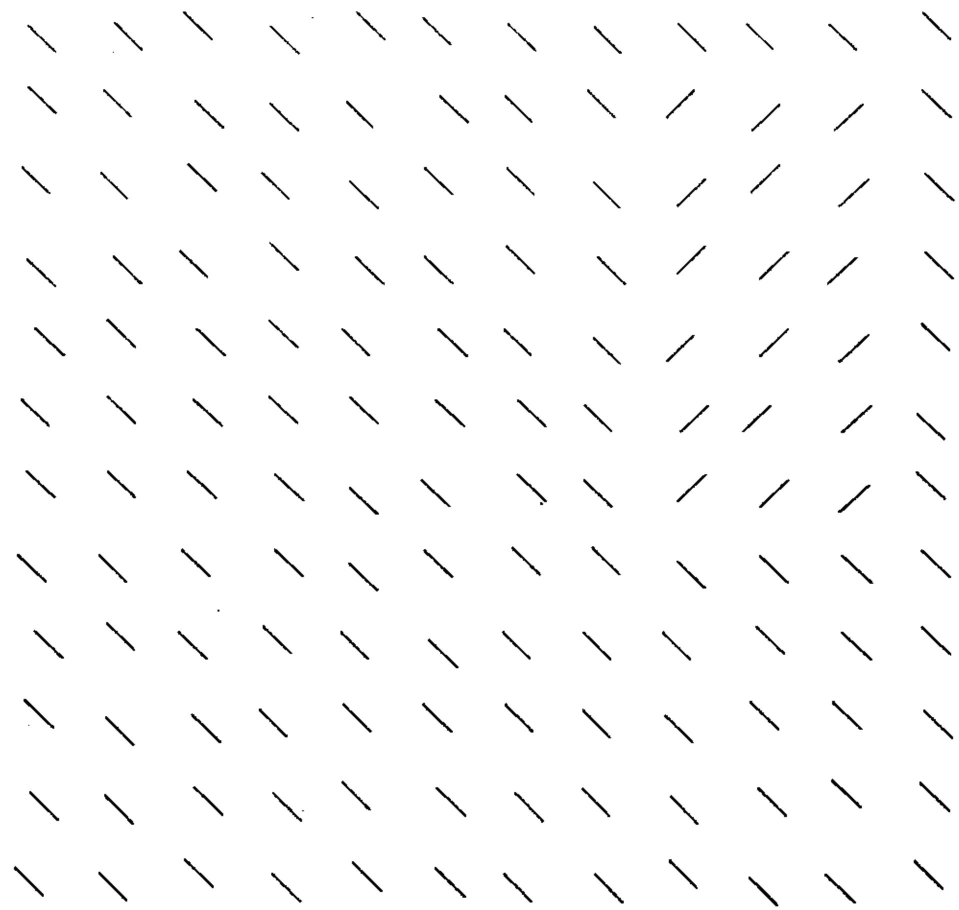

Figure 1.   Discrimination due to orientation difference.

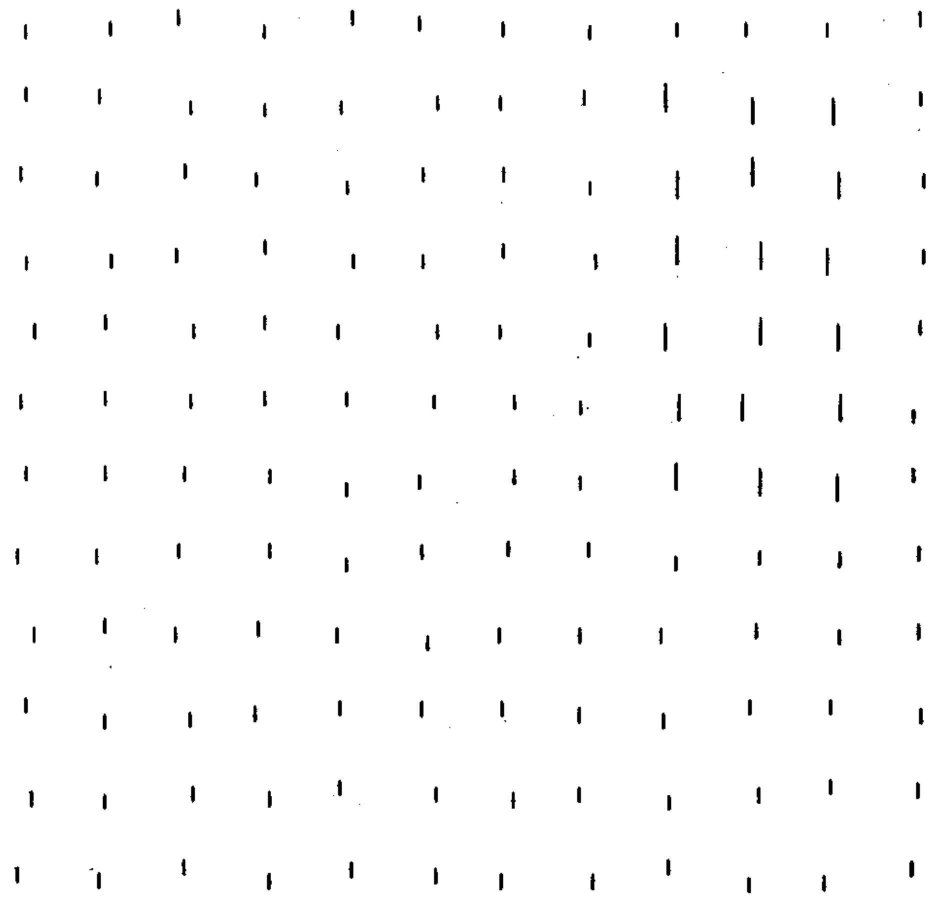

Figure 2. Discrimination due to length difference.

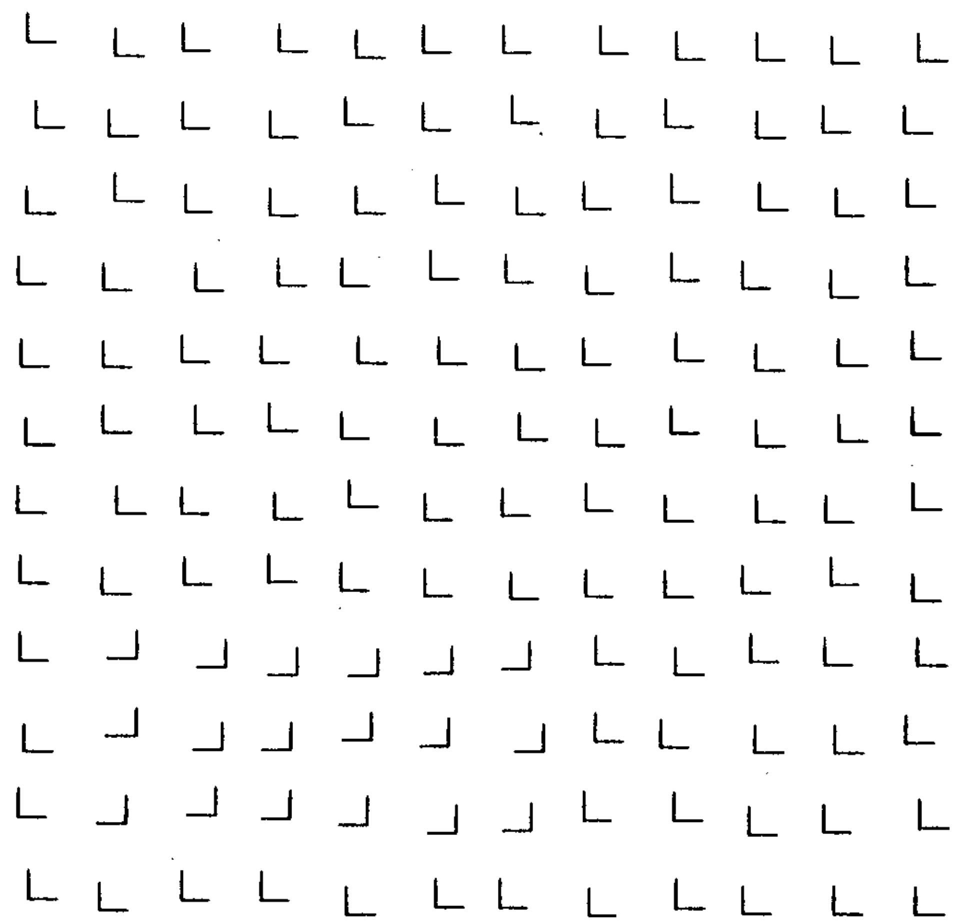

Figure 3. Discrimination due to virtual line difference.

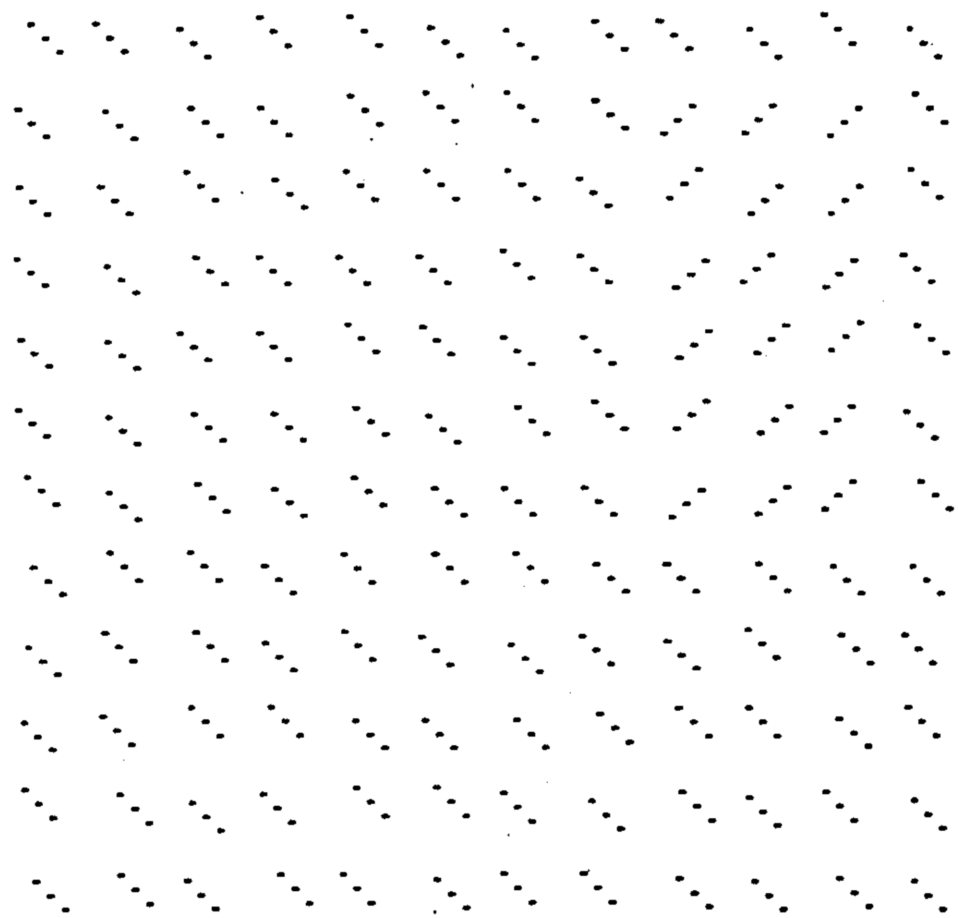

Figure 4. Virtual lines between points.

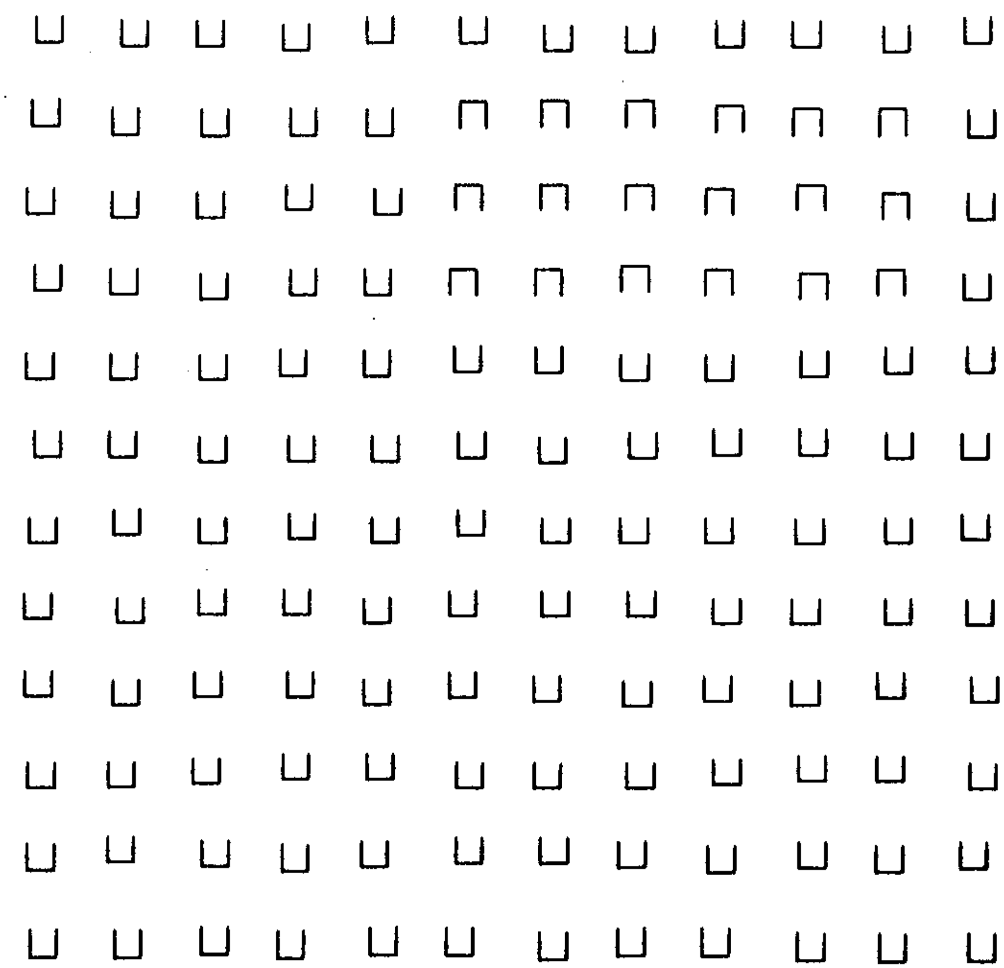

Figure 5. Identical second order statistics (nearly always) implies no discrimination.

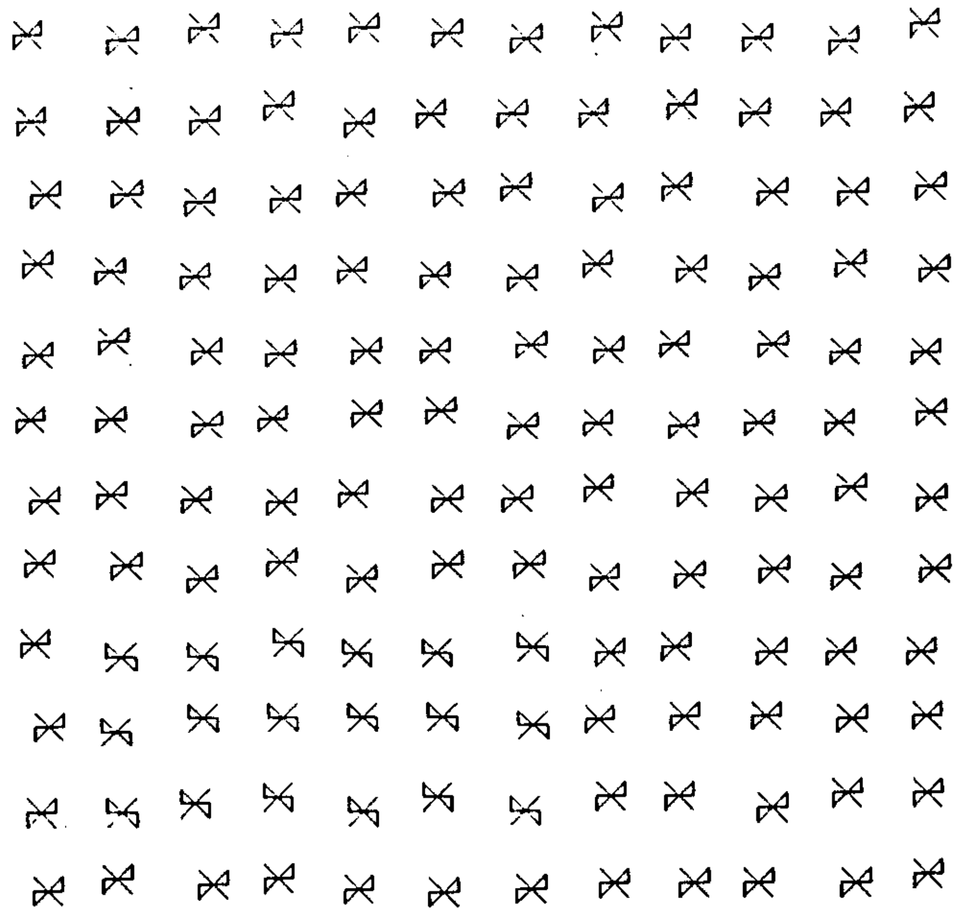

Figure 6. Actual and virtual lines same; second order statistics different yet no discrimination.

sufficient difference in orientation and length of lines?		difference in full second order statistics?	produces immediately perceived texture difference?	example in
actual	virtual	dipole	discrimination	figure
yes	-	yes	yes	1
no	yes	yes	yes	3
no	no	yes	no	6
-	-	no	no	5

Figure 7. Complexity of texture discrimination. ("-" implies property is of no significance).

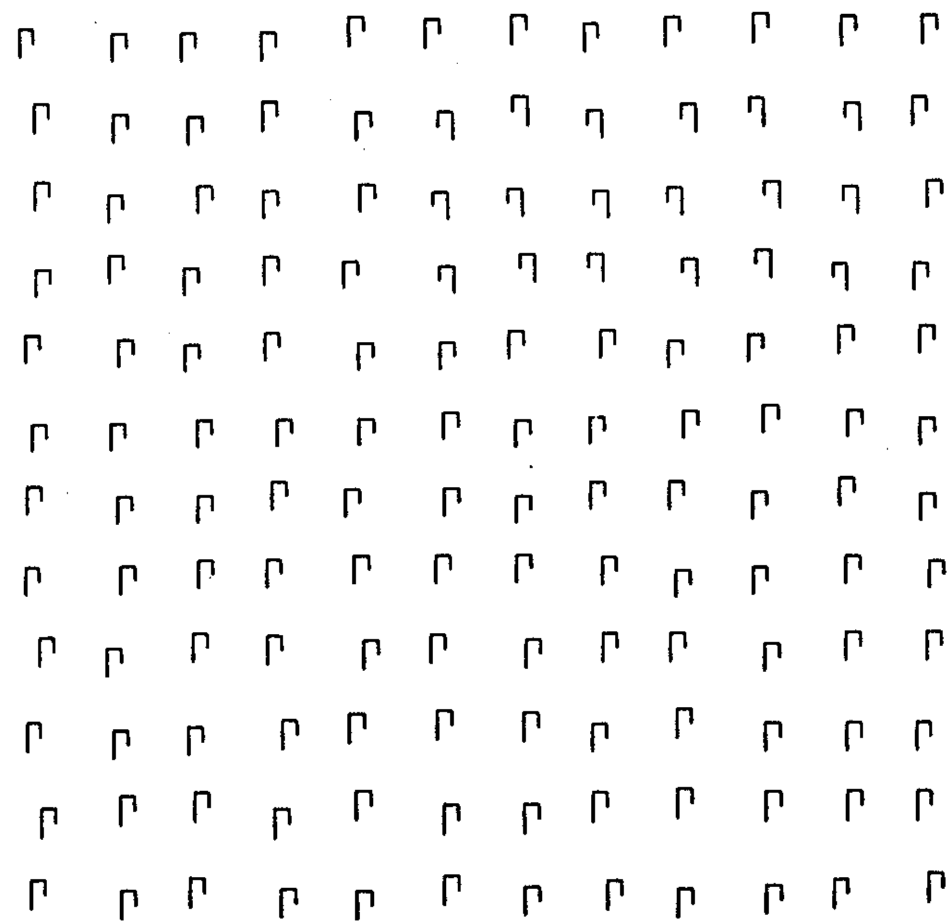

Figure 8. Dipoles different but not by enough to cause discrimination.