Chapter 11: Real and Apparent Motion

Brief summary of chapter

A sequence of rapidly presented, independently generated Glass patterns appears to move consistently, one way or the other, along the path of the transformation used in generating the Glass patterns. In this inquiry, rotation Glass patterns are employed. The apparent motion of rotation Glass patterns is examined along with the situation in which a Glass pattern is really rotated from frame to frame. A nearest neighbour explanation illuminates results from a published study set up to investigate the phenomenon. However, this leads on to consideration of the role of further Delaunay neighbours.

The work outlined in this chapter, with the exception of that indicated of other researchers, is original. Particularly, the interpretation of outcomes from Ross, Badcock, and Hayes (2000), in terms of proximity measures, is original.

Background

In a recent paper, Ross, Badcock, and Hayes (2000) draw attention to a phenomenon in which a sequence of rapidly presented, independently generated Glass patterns gives rise to a perception of coherent motion that, they claim, is 'indistinguishable from real motion' (p. 679), except that its direction is ambiguous. According to Ross et al. (2000), because there is no organization in the motion vectors, any coherence in the motion perceived 'can only come from their static internal structure' (p. 679).

This conclusion is arresting because it is widely believed that form and motion are analysed separately in mammalian visual systems (De Yoe, Fellerman, Van Essen, & McClendon, 1994; Goodale, Milner, Jakobson, & Carey, 1991; Merigan, 1999; Ungerleider & Mishkin, 1982). Instead, Ross et al. (2000) propose that the analysis of form and motion may be 'tightly coupled at all stages of visual analysis' (p. 681), as earlier suggested by Lennie (1998). In accordance with this view, the generative transformation model uses the same routines to detect both form (static or spatial structure) and motion (dynamic or temporal structure). And on this view, 'information about form could contribute to the analysis of motion, just as information about motion can to the analysis of form' (Ross et al., 2000, p. 681, citing Burr, Ross, & Morrone, 1986).

While we agree that form and motion may be closely linked at all stages, we consider the argument for it, presented by Ross et al. in this particular instance, is incomplete and may be misleading. Although Badcock, Ross, and Hayes (2000) recognise that such sequences 'contain motion signals of random velocity' (Suppl. 25), their focus is on the contribution of *static* structure to the organization of this motion into a coherent percept. The analysis of successive, independently generated Glass patterns presented below, however, shows that such sequences provide an abundance of information for motion, while there are also spatio-temporal constraints that make for coherence. This result means that it may not be useful to conceive of the experience elicited by these sequences as a paradoxical process, in which static structure alone is responsible for a perception of coherent and continuous motion.

In their first experiment, Ross et al. (2000) presented observers with a sequence of 10 independently generated Glass patterns. Each pattern in a sequence was generated by starting with an array of 50 dots, distributed uniformly and randomly within a circular area, imposing a rotational transformation about the centre, and then superposing the transformed array onto the original array. According to Ross et al. (2000), '[t]here is no corresponding organization in the motion vectors from a sequence of independent Glass patterns…and thus a percept of coherent motion is not predicted' (p. 679).

In a second stage of their investigation, Ross et al. (2000) found that observers were unable to distinguish the apparent spin of independently generated Glass patterns from that of

a succession of Glass patterns linked by 'real bidirectional motion signals' (p. 680) consistent with a speed of 22º per sec. The independent Glass patterns had a dot pair separation of 8 arc min and a frame duration of 32 msec.[1] In the case of the motion-consistent Glass patterns, the bidirectional signals were intended to mimic the ambiguity of the direction of spin seen with the independent Glass patterns and took the form of alternating clockwise and anticlockwise directions of the displacements in four successive annular regions of the circular display area. According to Badcock (personal communication to D. Vickers, 2002), the motion-consistent Glass patterns also had a frame duration of 32 msec and a dot pair separation of 8 arc min, but had a between-frame separation of 44 arc min. The reason given for this separation is that '[t]he speed was needed to match the apparent speed of rotation of the independent Glass pattern sequence' (Ross et al., 2000, p. 680).

What is puzzling about this last finding is that the succession of independent Glass patterns, with dot pair separations of 8 arc min and frame durations of 32 msec, is consistent with a maximum total separation of just over 4º per sec, provided all individual separations are added without loss. Nonetheless, this succession of static patterns gives rise to an impression of spin that is equated with 'true' motion-consistent displacements with a speed of 22º per sec (i.e., over 5 times as fast). Moreover, this occurs despite the fact that the motion-consistent Glass patterns also have the same static structure as the independent Glass patterns.

Although Ross et al. (2000) do not comment on this finding, it seems paradoxical in terms of the interpretation they propose. What it suggests is that, not only does static structure provide a source of motion information, but it also appears to be several times as powerful as accepted sources of motion information. Moreover, although static structure is effective when successive Glass patterns are independent, it is much less so when successive rotated Glass patterns (endowed with the same static structure) are indeed consistent with motion. The analysis we advocate for a range of dot densities of the same order as those used by Ross et al., suggests a plausible resolution of these apparent paradoxes.

Since Ross et al. (2000) employ rotation Glass patterns in developing their thesis—but state generally equivalent effects for other Glass transformations—in developing our argument we generally restrict ourselves to rotation Glass patterns.

Interpretation in terms of mean nearest neighbour distance

In this interpretation we use distances measured *between* patterns, *not* distances measured *within* patterns. Mean nearest neighbour distance can provide an indication for what occurs with the apparent rotation perceived when independently generated Glass patterns are displayed in quick succession. The model calculates mean nearest neighbour distances from one independently generated Glass pattern to the next in a sequence, and all patterns in the sequence have the same number of pairs.

For real rotation sequences, the model calculates the mean nearest neighbour distance from another independently generated Glass pattern to the selfsame Glass pattern rotated by some angle increment, and this pattern has the same number of pairs as the other independently generated patterns. After generation it is rotated clockwise by 0^o increments, from which mean nearest neighbour distance is calculated from an increment to the next. (This is the null instance.) Then another independently generated Glass pattern is rotated clockwise by 1^o increments, from which mean nearest neighbour distance is calculated from an increment to the next. And then another independently generated Glass pattern is rotated clockwise by 2^o increments, from which mean nearest neighbour distance is calculated from

[1] The term 'pair separation' used here is inaccurate. Strictly it means separation between one pair and another, when generally throughout it is meant to indicate displacement between transformational partners. To be consistent with the connotation of Ross et al. (2000), we here use the same terminology to indicate displacement between transformational partners.

an increment to the next, and so on, up to 180° increments. A new, different, Glass pattern is generated for each real rotation sequence. The same number of increments is applied to each real rotation sequence for the purpose of observing effect, but any two consecutive increments in each sequence suffice for calculating mean nearest neighbour distance.[2]

At real rotation increments corresponding to a mean nearest neighbour distance below a mean nearest neighbour distance range for apparent rotation, real rotation can be discriminated, either when presented on the same display with apparent rotation or when presented by itself. Within and about the mean nearest neighbour distance range for apparent rotation, real rotation becomes indistinguishable from apparent rotation in apparent speed as well as ambiguity of direction. This appears to be the case over a considerable density range of Glass pairs. Hence apparent rotation speeds are in the regions of the upper limits for discriminable real rotation speeds. See Figures 11.1 to 11.4 for graphs of mean nearest neighbour distances for 50, 100, 200, and 300 apparent and real rotation Glass pairs. Note the reduction in angle increments for real rotation required to match the speed of apparent rotation for increasing number of Glass pairs.

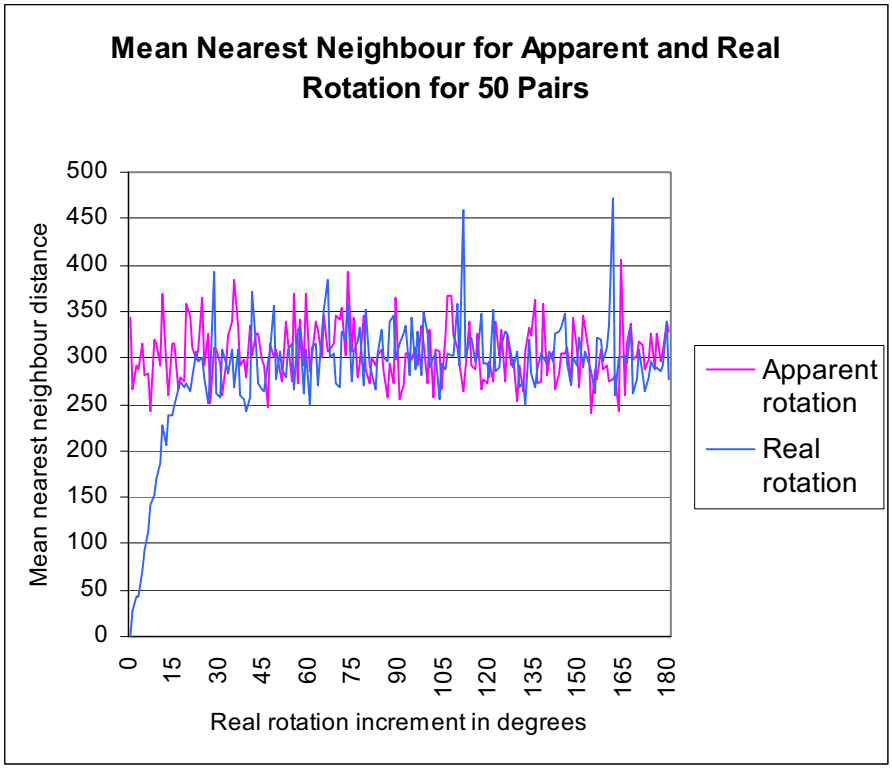

Figure 11.1: Mean nearest neighbour distances for 50 apparent and real rotation Glass pairs. The boundary between unambiguous and ambiguous real rotation is here shown to be about 23°.

[2] Generating a new Glass pattern for each real rotation sequence ensures palliation of any pattern atypicality.

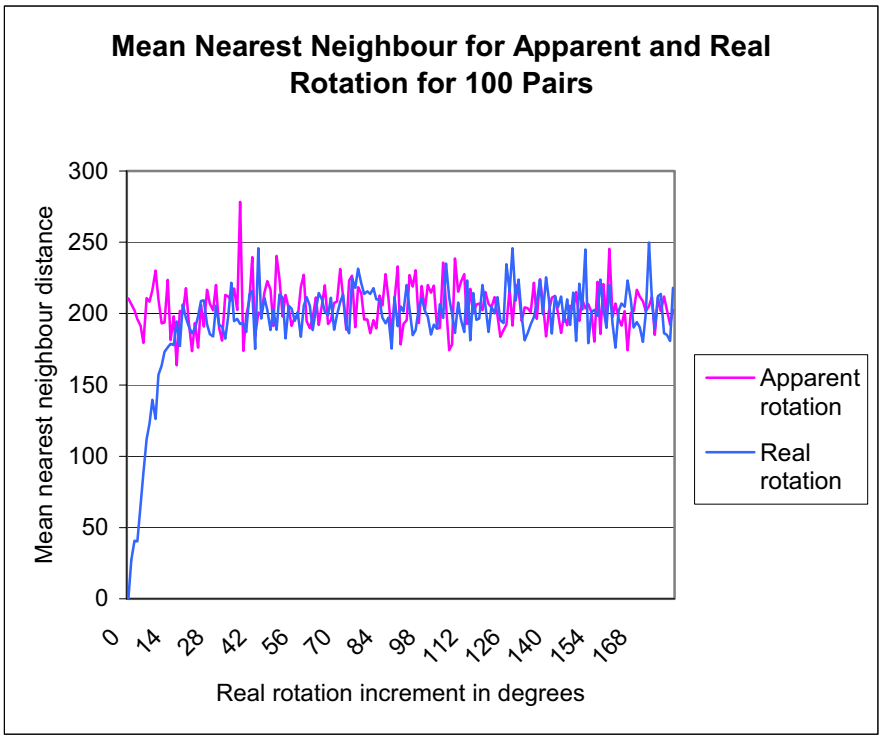

Figure 11.2: Mean nearest neighbour distances for 100 apparent and real rotation Glass pairs; the same number of pairs as used by Ross, Badcock, and Hayes (2000). The boundary between unambiguous and ambiguous real rotation is here shown to be about 19°.

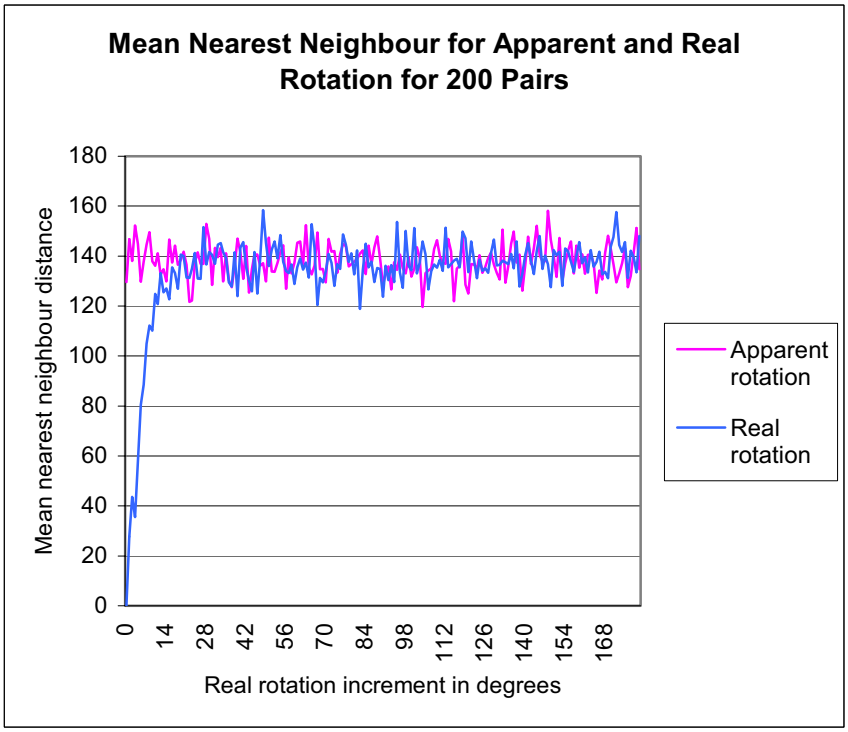

Figure 11.3: Mean nearest neighbour distances for 200 apparent and real rotation Glass pairs. The boundary between unambiguous and ambiguous real rotation is here shown to be about 16°.

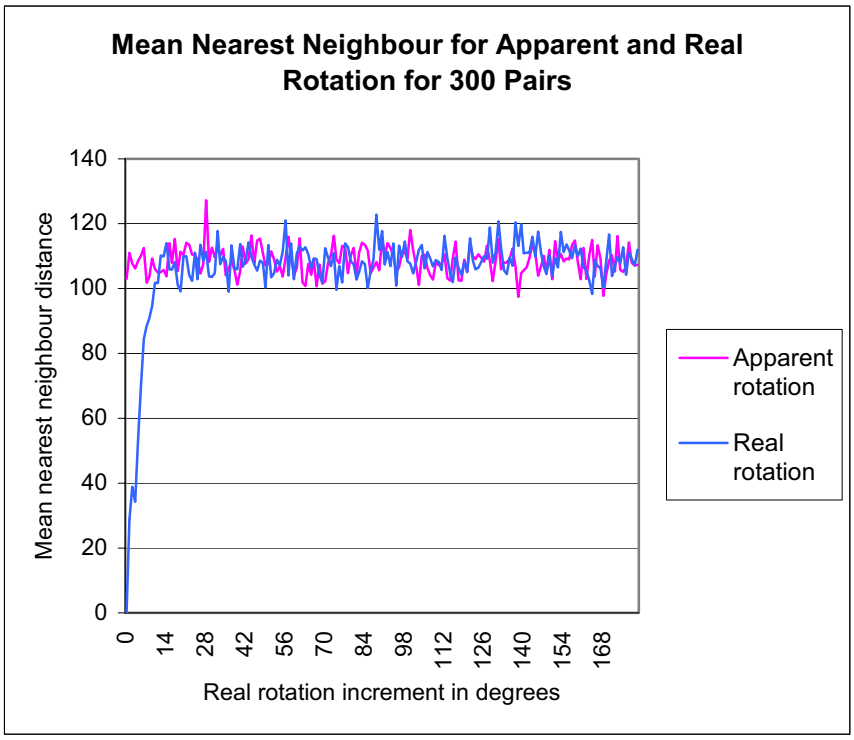

Figure 11.4: Mean nearest neighbour distances for 300 apparent and real rotation Glass pairs. The boundary between unambiguous and ambiguous real rotation is here shown to be about 14°.

Figure 11.5 shows the graph of mean nearest neighbour distances for 100 apparent and real rotation Glass pairs, each averaged over five runs, which provides a better indication of the boundary between unambiguous and ambiguous real rotation: about 24°.

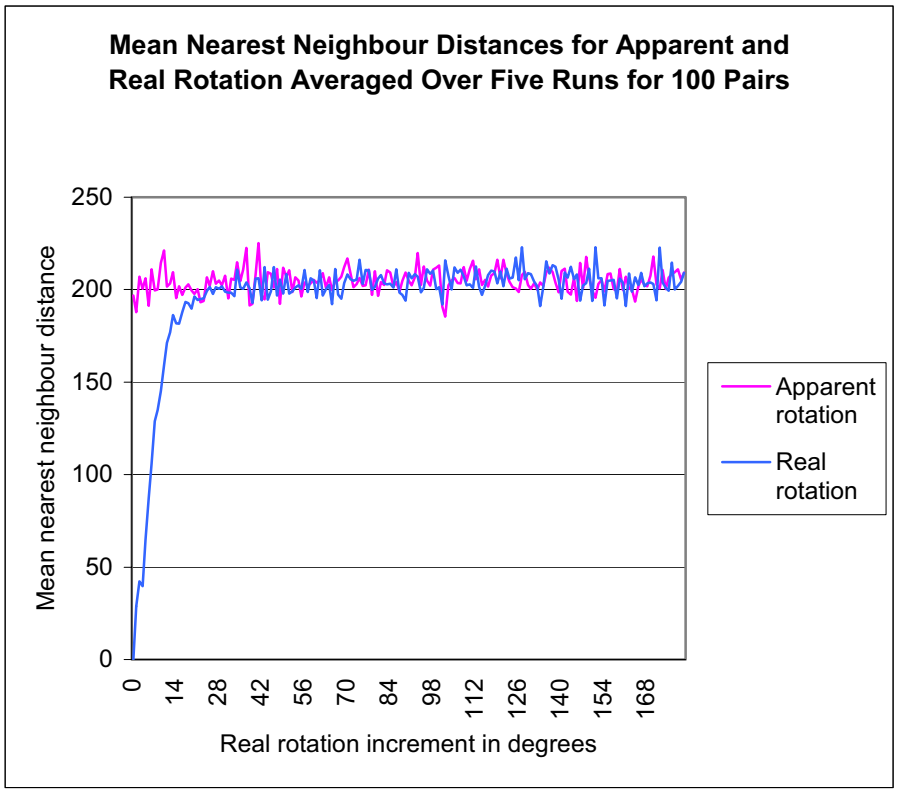

Figure 11.5: Mean nearest neighbour distances for 100 apparent and real rotation Glass pairs, each averaged over five runs. The boundary between unambiguous and ambiguous real rotation is here shown to be about 24°.

Figure 11.6 shows mean nearest neighbour distances for clockwise real rotation of 100 Glass pairs, independently generated for each sequence of increments ranging from 0° to 360°. (That is, a new, different, Glass pattern is generated for each real rotation sequence.) For the relatively smooth initial part of the graph, rotation appears slower than it does for apparent rotation, and unambiguously clockwise. The closer to zero, the slower it appears. For the fluctuating part of the graph, rotation appears indistinguishable from that of apparent rotation; i.e. about the same speed and ambiguous in direction. For the relatively smooth final part of the graph, rotation appears slower than it does for apparent rotation, and unambiguously anticlockwise. The closer to 360°, the slower it appears. Of course, rotation appears unambiguously anticlockwise this time because the pattern is rotating clockwise but pulling up short of a full circle by some, diminishing, few degrees, and the shortest transformation path is reversed. And when the pattern rotates by increments of 360°, then no rotation is observed.

Generating a new Glass pattern for each real rotation sequence ensures palliation of any pattern atypicality. Figure 11.7 shows a typical graph for which the same Glass pattern is involved in all rotation sequences. However, had the pattern been atypical—an unrepresentative number of pairs in one half of the display for instance—then mean nearest neighbour distances would have shown bias.

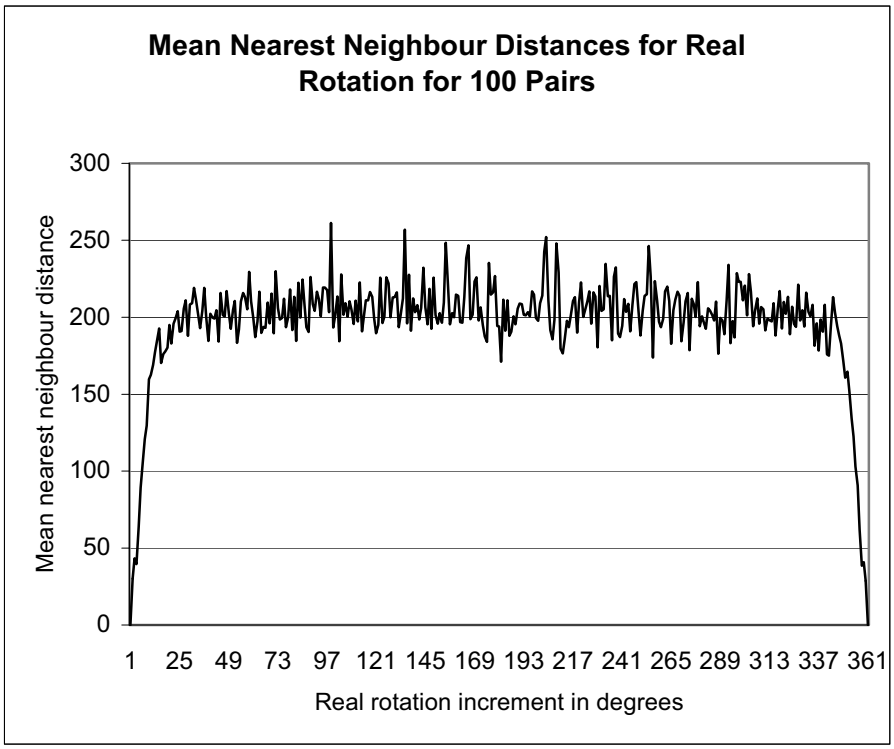

Figure 11.6: For the relatively smooth initial and final parts of the graph, rotation appears slower than for the fluctuating part of the graph, and appears unambiguously clockwise and anticlockwise respectively. For the fluctuating part of the graph, rotation appears about the same speed as for apparent rotation and appears ambiguous in direction.

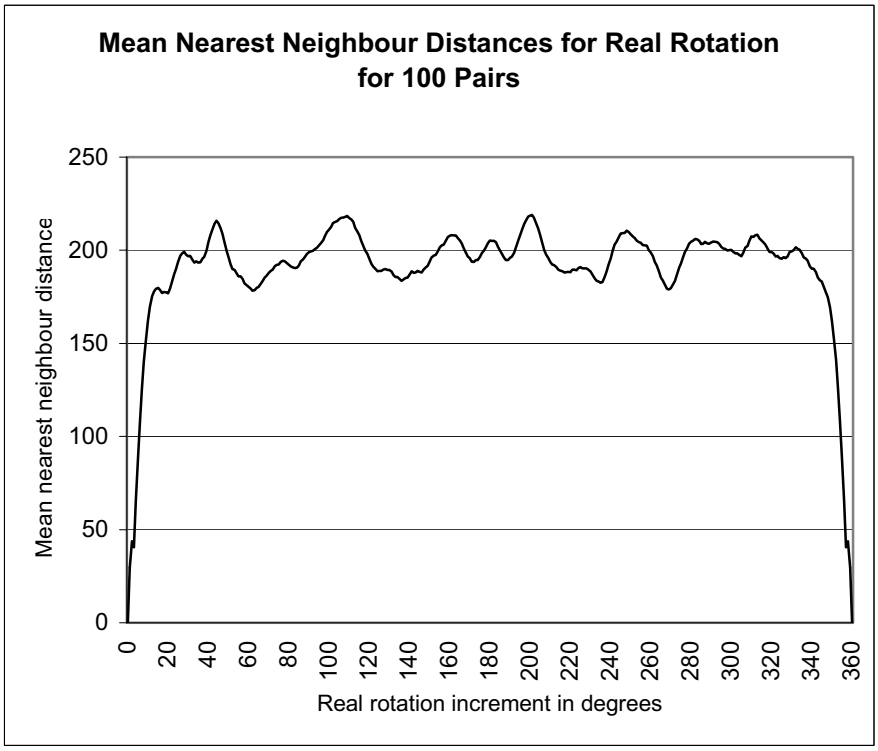

Figure 11.7: Typical graph for which the same Glass pattern is involved in all rotation sequences. Generating a new Glass pattern for each real rotation sequence, as per Figure 11.6, ensures palliation of any pattern atypicality.

Why does speed appear to stabilize at an upper limit, and ambiguity of direction appear to occur for angle increments of real rotation that take its mean nearest neighbour distances into the range of mean nearest neighbour distances for apparent rotation? It concerns the distance at which the regions of shortest path for coherent displacement become populated by interposing pairs. Below mean nearest neighbour distance, coherent displacement is evident because of reduced interposition by other pairs, but beyond mean nearest neighbour distance, displacements are commonly interposed by other pairs.

In the event that both apparent and real rotation patterns are displayed together, interpositions can come from pairs belonging to both patterns. In the event that only real rotation patterns are displayed, then interpositions come entirely from within. See Figure 11.8, which, for sake of simplicity, shows real rotation pairs only. When A is rotated to A' then B is rotated to B', which obscures the path from A to A'.

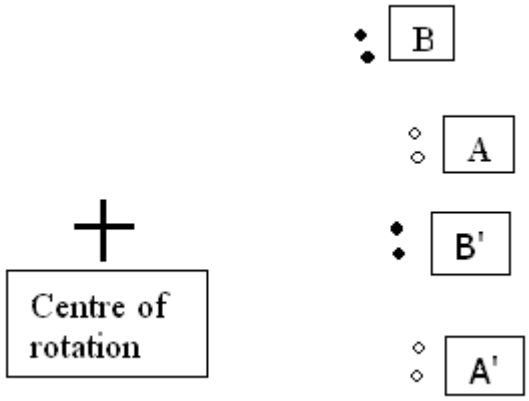

Figure 11.8: A typical situation in which Glass pairs are rotated in angle increments above the range for mean nearest neighbour distance for apparent rotation.

And why, with more dot pairs, is there a reduction of the real rotation angle increment at the boundary between unambiguous and ambiguous real rotation, which coincides with the mean nearest neighbour distance region for apparent rotation? Because for more dot pairs the mean nearest neighbour distance reduces; hence perceived rotation speed slows. This slowing is reflected in a reduction of the real rotation angle increment, as shown in Figure 11.9. The reduction of the real rotation angle increment coincides with the boundary between unambiguous and ambiguous real rotation, and is proportional to decrease in mean nearest neighbour distance with increase in number of dot pairs.

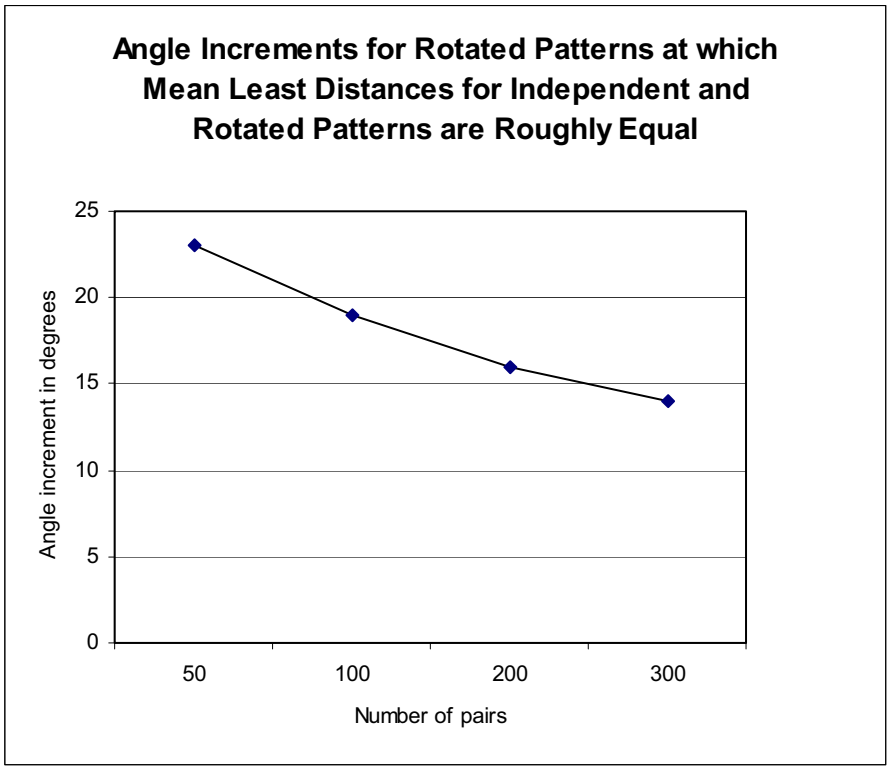

Figure 11.9: Reduction, with increasing number of dot pairs, of the real rotation angle increment, which coincides with the boundary between unambiguous and ambiguous real rotation indicated by the mean nearest neighbour distance region for apparent rotation. Slowing of perceived rotation is reflected in this reduction.

Motion vectors and Delaunay neighbours

Ross et al. (2000), conclude that there is no organization of motion vectors in sequences of Glass patterns. In other words, motion vectors in sequences of Glass patterns are random, as they are for sequences of random dots. However, point-to-point distances for Glass patterns have a grouping due to transformation consistencies, and nearest neighbour distances also can have a concomitant grouping. This is the case whether we are talking about *within* a Glass pattern or *between* Glass patterns, but you are reminded that here we deal with the latter.

If you consider the source of the motion illusion to lie within any one Glass pattern, (which I do not), then you still have your concomitant grouping. However, the source of any such illusion is nothing static. Two adjacent light bulbs permanently on or off, for example, do not account for the motion illusion that materializes when they are alternated on and off. The *source* of the illusion is in the dynamics (energy difference), not the statics! The *character* of the illusion is defined by the statics, which means unchanging arrangement.

Figures 11.10 to 11.15 show graphs of distances derived from configurations having just 10 dots for the sake of clarity. The first triplet of figures shows graphs of point-to-point distances sorted in ascending order, and the second triplet shows corresponding graphs of nearest neighbour distances sorted in ascending order. The first graph of each triplet is for sorted distances between two sets of random dots, the second is for sorted distances between two Glass translation patterns, and the third is for sorted distances between two Glass rotation patterns.

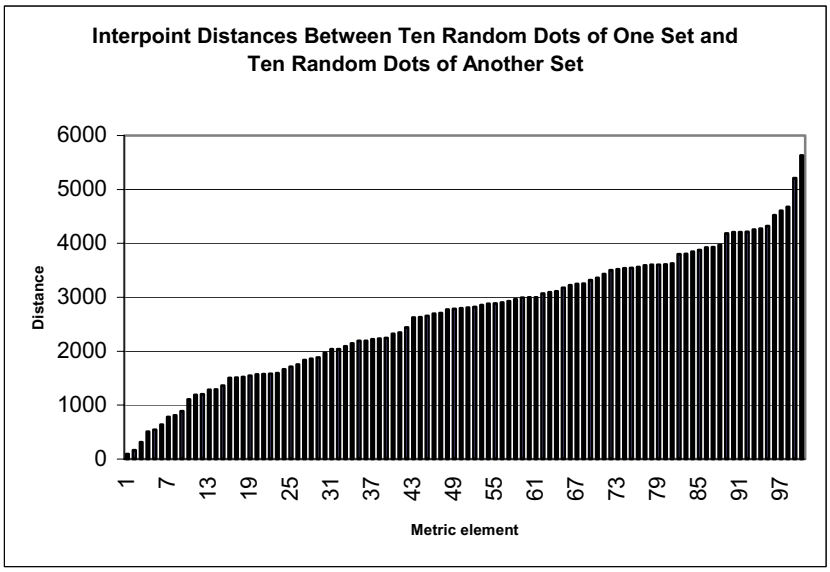

Figure 11.10: Sorted point-to-point distances from one random dot set to the next set in a sequence. No grouping is evident.

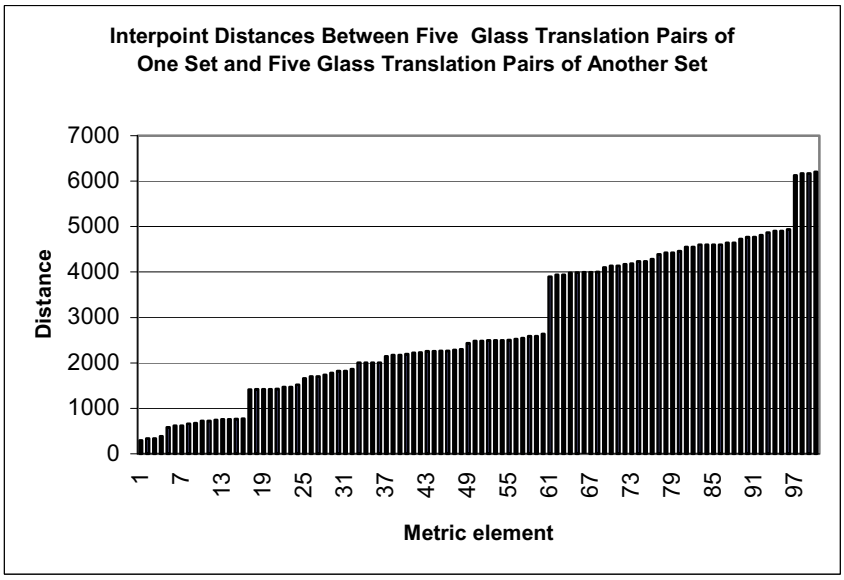

Figure 11.11: Sorted point-to-point distances from one translation Glass set to the next set in a sequence. Quad grouping can be distinguished.

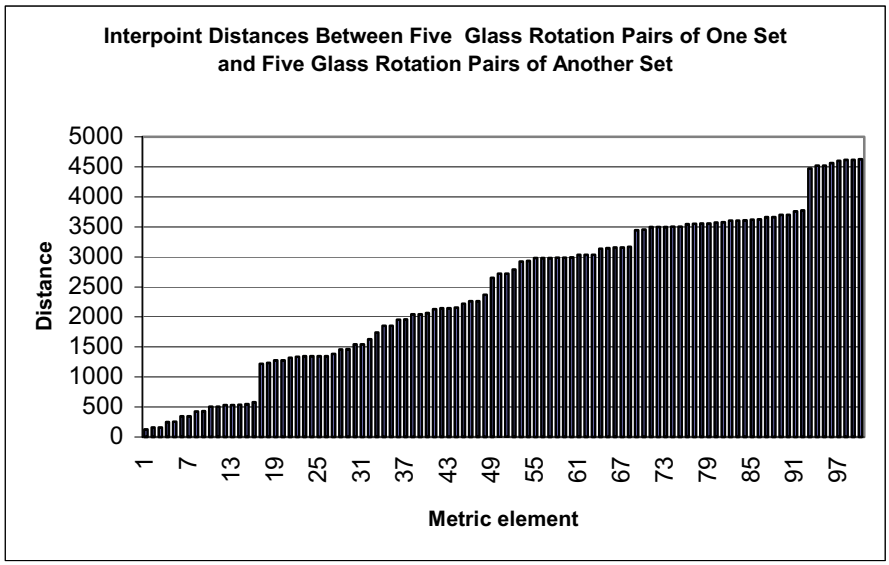

Figure 11.12: Sorted point-to-point distances from one rotation Glass set to the next set in a sequence. Quad grouping can be distinguished.

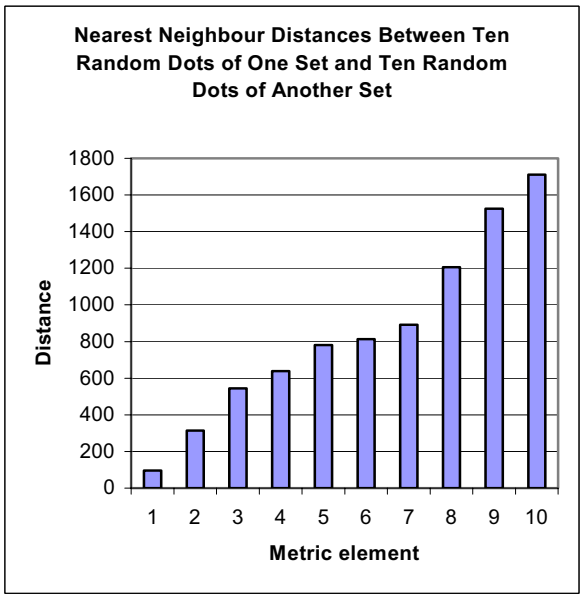

Figure 11.13: Sorted nearest neighbour distances from one random dot set to the next set in a sequence. No pairing is evident.

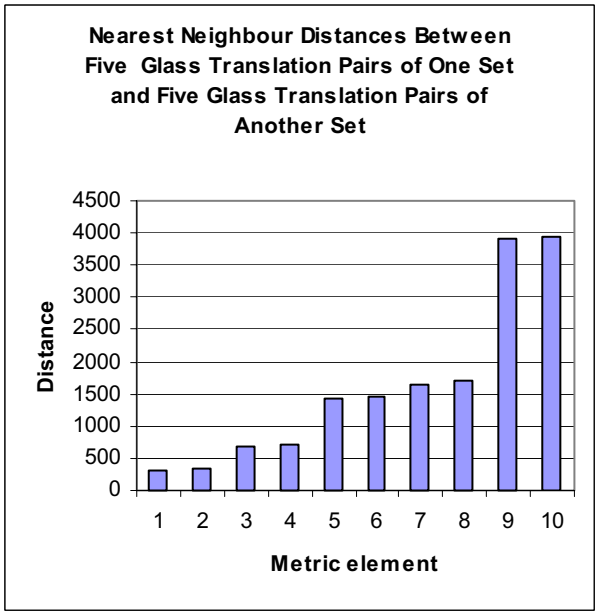

Figure 11.14: Sorted nearest neighbour distances from one translation Glass set to the next in a sequence. Pairing within a narrow order of magnitude is evident.

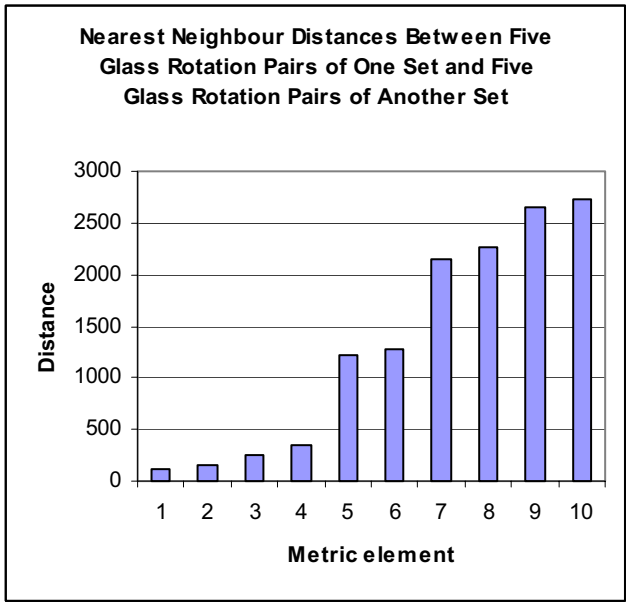

Figure 11.15: Sorted nearest neighbour distances from one rotation Glass set to the next in a sequence. Pairing within a narrow order of magnitude is evident.

For point-to-point distances from a set *A* to a set *B,* Glass patterns exhibit quad groupings, i.e. distance *a* to *b, a* to $b_{transform}$, $a_{transform}$ to *b,* and $a_{transform}$ to $b_{transform}$. For nearest neighbour distances from a set *A* to a set *B,* Glass patterns can exhibit pair-wise grouping because if *b* is nearest *a* then $b_{transform}$ has some chance of being nearest *a* as well.

The number of dots used to generate these particular figures is limited for the sake of graphical clarity. And it is not disputed that as mean nearest neighbour distance becomes less than Glass translation distance or mean Glass rotation distance (or any other Glass transformation distance), then nearest neighbour grouping degrades.

Motion vectors appear to be determined by nearest neighbour distances, at least, and by all Delaunay distances at most. Certainly, the unambiguous real rotation motion vectors are described by a preponderance of transformational paths as shortest distances, and while mean nearest neighbour distance from successive real rotation instances is less than mean nearest neighbour distance from successive apparent rotation instances, there is no ambiguity of real rotation direction. Figure 11.16 shows nearest neighbour edges as vectors from a Glass pattern (set A, hollow dots) to the selfsame Glass pattern rotated by 5^{o} (set B, solid dots).

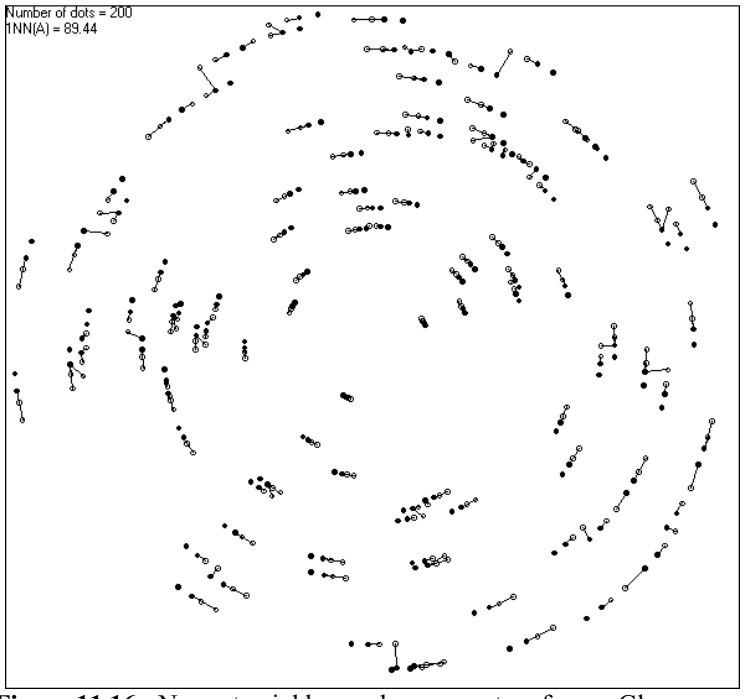

Figure 11.16: Nearest neighbour edges as vectors from a Glass pattern (set A, hollow dots) to the selfsame Glass pattern rotated by 5^{o} (set B, solid dots). (The numbers in the upper left corner are somewhat incidental, per favour of the computer routine. They are number of dots, and mean nearest neighbour distance in screen units from set A to set B.)

This unambiguous direction is lost as mean nearest neighbour distance from successive real rotation instances gets close to, equals, or exceeds, mean nearest neighbour distance from successive apparent rotation instances, but the rotation effect along with apparent speed is not lost. Hence vectors become quasi-vectors: direction is purely in the eye of the beholder. Figure 11.17 shows nearest neighbour edges as quasi-vectors from a Glass pattern (set A, hollow dots) to the selfsame Glass pattern rotated by 30^{o} (set B, solid dots). The quasi-vectors are similar to those for the two independently generated Glass patterns of Figure 11.18.

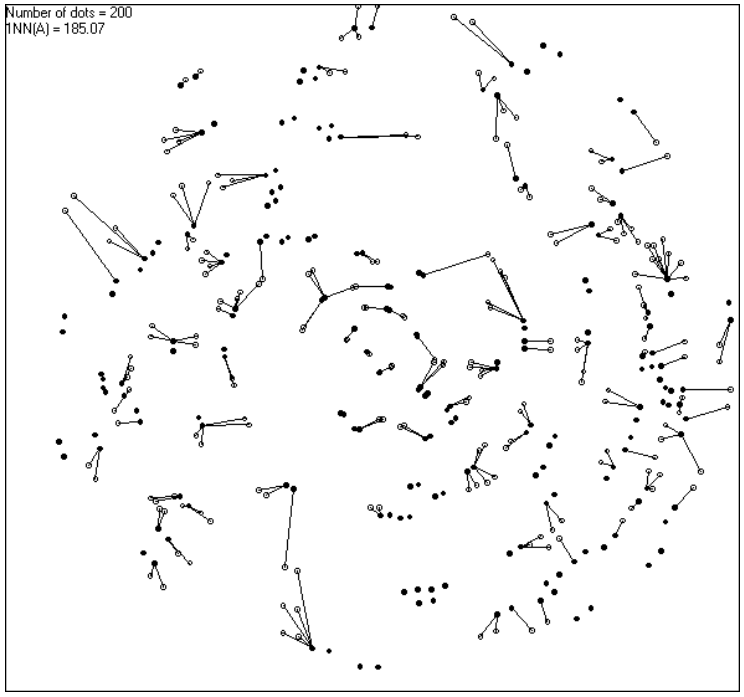

Figure 11.17: Nearest neighbour edges as quasi-vectors from a Glass pattern (set A, hollow dots) to the selfsame Glass pattern rotated by 30^{o} (set B, solid dots). The quasi-vectors are similar to those for the two independently generated Glass patterns of Figure 11.18.

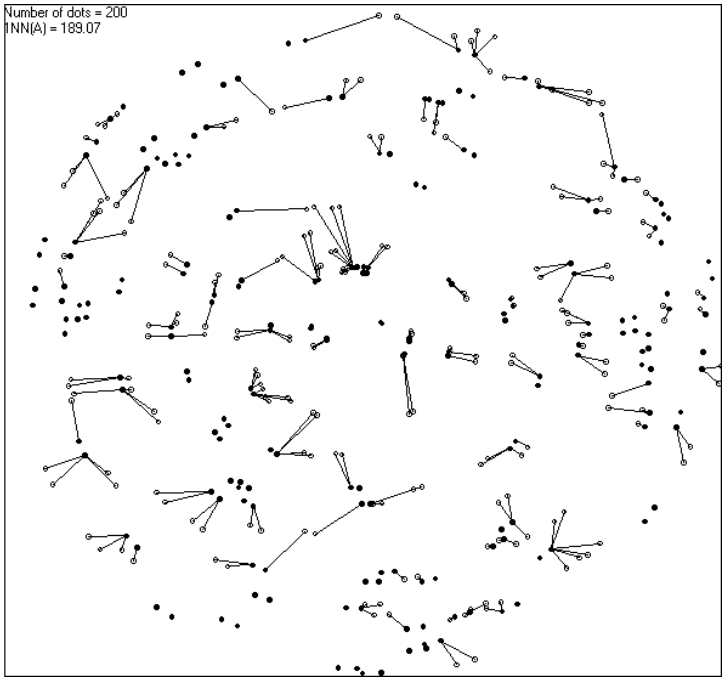

Figure 11.18: Nearest neighbour edges as quasi-vectors from a Glass pattern (set A, hollow dots) to a subsequent independently generated Glass pattern (set B, solid dots).

Mean nearest neighbour distance is germane to directional ambiguity and apparent speed (the larger the mean the faster apparent speed), but what of the rotation effect? When mean nearest neighbour distance is roughly the same for the real and apparent rotation

situations, the rotation effect appears essentially the same for each. In the case of real rotation, transformation paths no longer have a preponderance of shortest distances.

A ploy that teases apart what is happening involves displaying only nearest neighbour edges; not Glass patterns themselves. Figures 11.19 and 11.20 show nearest neighbour edges resulting from markedly different dot densities. Each set of edges are quasi-vectors from a Glass pattern (set A) to a subsequent independently generated Glass pattern (set B). The Glass patterns themselves are not shown.

Figure 11.19: Nearest neighbour edges as quasi-vectors from a Glass pattern (set A) to a subsequent independently generated Glass pattern (set B). The Glass patterns themselves are not shown.

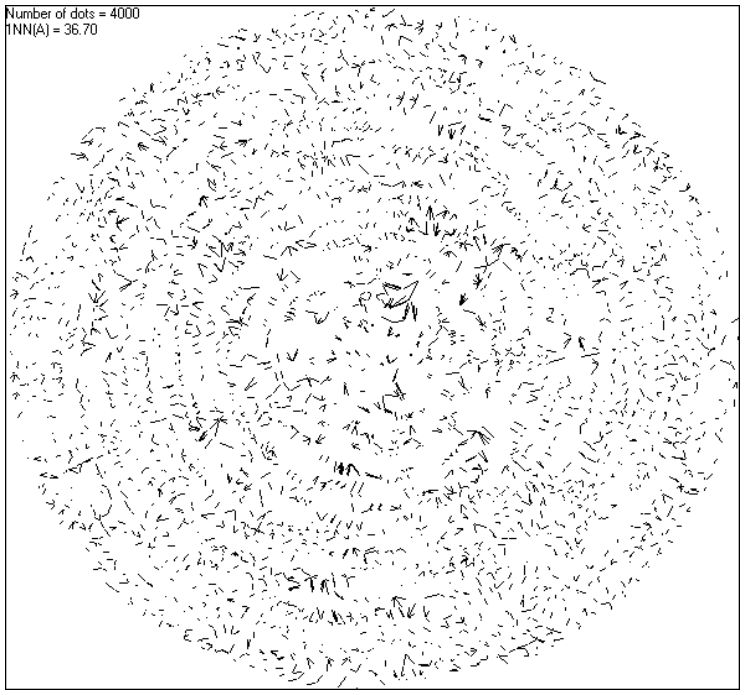

Figure 11.20: Nearest neighbour edges as quasi-vectors from a densely populated Glass pattern (set A) to a subsequent independently generated, densely populated Glass pattern (set B). The Glass patterns themselves are not shown.

This is done for both real rotation and apparent rotation, and, of course, motion vectors essentially develop along transformation paths for the smaller increments of real rotation; for which unambiguous rotation is clear despite the absence of Glass patterns. For the larger increments of real rotation, and for apparent rotation, motion vectors only sometimes develop along transformation paths, but, rotation, albeit ambiguous in direction, is obvious despite the absence of Glass patterns.

This can be explained simply by two considerations. The first is that for ambiguous rotation, vectors are bi-directional. They are not the true vectors perceived for unambiguous real rotation. For the true vectors there is little doubt about direction owing to a preponderance of consistent increases in length that extends all the same ends in the direction of rotation consistent with the smallest transformation. The second consideration involves observation of the quasi-vectors as a sequence of hidden Glass patterns proceeds.

The quasi-vectors extend, break and reconnect from elements of the hidden patterns of set A to set B, and then from set B to set C, and so on. Connections from A to B are taken in apparent fashion from B to C in such a way that the kind of patterns from which they arise is evident. All Glass patterns are anisotropic in some way, which favours increased density orientated along transformation directions, and it is particularly along such aspects that these 'probing feeler' quasi-vectors appear to begin or end as they converge on, or diverge from, transformation paths from display to display.[3] The quasi-vectors themselves do not necessarily align with transformation paths, but point them out in a way that makes them totally obvious. At frame rates in quick succession, the quasi-vectors on their own give the same impression of rapid rotation as observed for Glass patterns on their own. And last,

[3] This is not that evident from looking at the figures: there is no substitute for a dynamic display with ability to control frame rate. However, Figure 11.20, with its more densely inspired rotational form, hints at what the effect might be.

quasi-vectors for a sequence comprising independently generated, hidden random dots give no perception of rotation whatsoever; and why should they? But perception of rotation is obvious enough for real rotation involving hidden random dots.

Motion paths are *not* perceived to cross in any sequence, random or otherwise. This is interesting because Delaunay edges do not cross either: they simply intersect at dots. This means that none of the hierarchy of edges within Delaunay can cross, which augers well for a nearest neighbour interpretation. Nearest neighbour appears to be implicated over a large range of Glass pattern densities; but with some help, especially as densities become extreme.

A subset of point-to-point distances constituted of the range of Delaunay neighbour distances has the same grouping or lack thereof, as that for a whole set of point-to-point distances. All Delaunay neighbours might be profitably recruited in an argument for the perception of apparent rotation. It is clear what 'unobstructed' means by way of what 'obstructed' means, in the sense of nearest neighbour: if a transformation path is not a nearest neighbour path, then it is obstructed. (In spite of this, though, rotation is still adequately *pointed* out.) Enlisting all Delaunay neighbours, with their hierarchy of meanings for 'unobstructed', might help resolve issues for those who would like motion vectors to lie along transformation paths (and, again, it is not necessary for ambiguous motion).

The relevant question is, 'Are transformational paths included in the hierarchy of edges that constitute Delaunay triangulation?' This can be decided by checking that the mean for the farthest Delaunay neighbour distances for each dot from one set to the next is greater than transformation distance. If so, it would then be a relatively simple matter to identify which rank Delaunay neighbours best fit.

Figure 11.21 shows a colour coded Delaunay triangulation as quasi-vectors from a very densely populated Glass pattern to a subsequent independently generated, very densely populated Glass pattern. The Glass patterns themselves are not shown. Delaunay neighbours as quasi-vectors are colour coded with black for nearest neighbours, blue for second rank Delaunay neighbours, green for third rank Delaunay neighbours, cyan for fourth rank Delaunay neighbours, red for fifth rank Delaunay neighbours, magenta for sixth rank Delaunay neighbours, yellow for seventh rank Delaunay neighbours, and so on with other colours, through to farthest rank Delaunay neighbours. The unspecified colours do not matter here because they so rarely present.

Figure 11.21 shows some evidence of rotation, but if it is submitted to colour enhancement for one colour at a time, the eye readily discerns which colour quasi-vectors best load along transformation paths. However, they still look spindly and generally lack impact owing to diminution; hence 'diffusing' is expedient. Diffusing the enhanced colour about the spindles highlights the rotational effect, as shown in Figure 11.22.

Unlike the other figures representative of Glass patterns, which are dispersed over circular forms because that is the way Ross et al. did it, these last two figures are dispersed over square forms. This is to obviate the possibility that circular perimeters might somehow heighten, or at least influence, perception of any pattern.

Figure 11.21: Delaunay triangulation as quasi-vectors from a very densely populated Glass pattern to a subsequent independently generated, very densely populated Glass pattern. The Glass patterns themselves are not shown. Delaunay neighbours as quasi-vectors are colour coded with one colour for nearest neighbours, another colour for second rank Delaunay neighbours, and so on, through to farthest rank Delaunay neighbours.

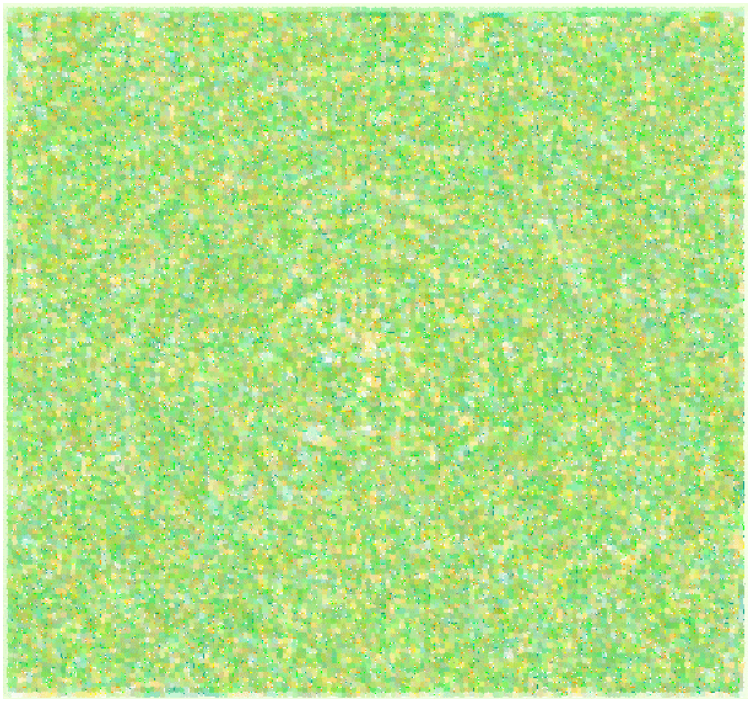

Figure 11.22: Delaunay triangulation depicted in Figure 11.21, with enhancement and diffusing of the colour coded neighbours as quasi-vectors that reveal most spin. At this density, most spin loads on third rank Delaunay neighbours.

Here, green is dominant, so, generally, third rank Delaunay neighbour edges show an appreciable element of rotation. That is not to say that there are no useful rotation indicators among other edges. In any event the significant message regarding stimuli of the density employed by Ross et al. (2000) is that nearest neighbour explains perceived rotation adequately. And nearest neighbour continues to do so adequately through to a considerable density.

Before leaving this topic, one more point might be profitably raised in connection with visual perception generally. It is related to the theory of those who advocate frequency filtering from a computational point of view: Adelson and Bergen, (1985); Heeger, (1988); Jones and Malick, (1992); Koendernink and Van Doorn, (1976a); Malick and Perona, (1990), for example. And it is also related to ecological theory like that of Gibson, (1950, 1966, 1979), which shifts the problem to the environment. The former uses multi-orientation, multi-scale filters as spatial primitives upon which higher level processes operate, and the latter contends that visual perception is simply a process of 'resonating' to the rich variety of information available in the stimulus array.

It seems that just a hint of some consistency, or pattern, primes the perceptual search for more. If no more is forthcoming, then an array is deemed random under the current approach. However, more seems to prime more, in an exponential way, such that if there *is* more then it looms rapidly. In other words, it would not take more than just the odd valid transformational trajectory here and there to tune the perceptual filter, and the resonance of the tuned system would ensure a rapid escalation to global perception. Since this involves self-similarity across scales, such a mental process might be described in fractal terms.

Chapter 12: Neighbours and Beyond

Brief summary of chapter

Neighbour connectivity patterns of Glass patterns without dot displays are examined for a reason as to why they reveal the structure of underlying patterns. Ultimately, Glass patterns with millions of dots are employed in the effort to go beyond Delaunay neighbours or n^{th} nearest neighbours for an explanation. Here, it seems, the concept of anisotropy can help.

All the work outlined in this chapter is original. Particularly, the neighbour connectivity investigation and measures of anisotropy are original, along with the investigation in which millions of dots are employed.

Relationship of neighbours to pattern detection

Figure 12.1 is composed of four computer screenshots. The upper left panel shows a Glass screw pattern comprising 10,000 dots, which was generated from a Glass rotation pattern by rotating the pairs at midpoints by 20^{o}. The upper right panel shows corresponding nearest neighbour edges; the display of stimulus dots has been omitted. The lower two panels show Delaunay neighbour edges: one in colour and the other in grey. The colours, of course, have the same interpretation as that for Figure 11.21, page 225. Displays of stimulus dots for the Delaunay diagrams have been also omitted.

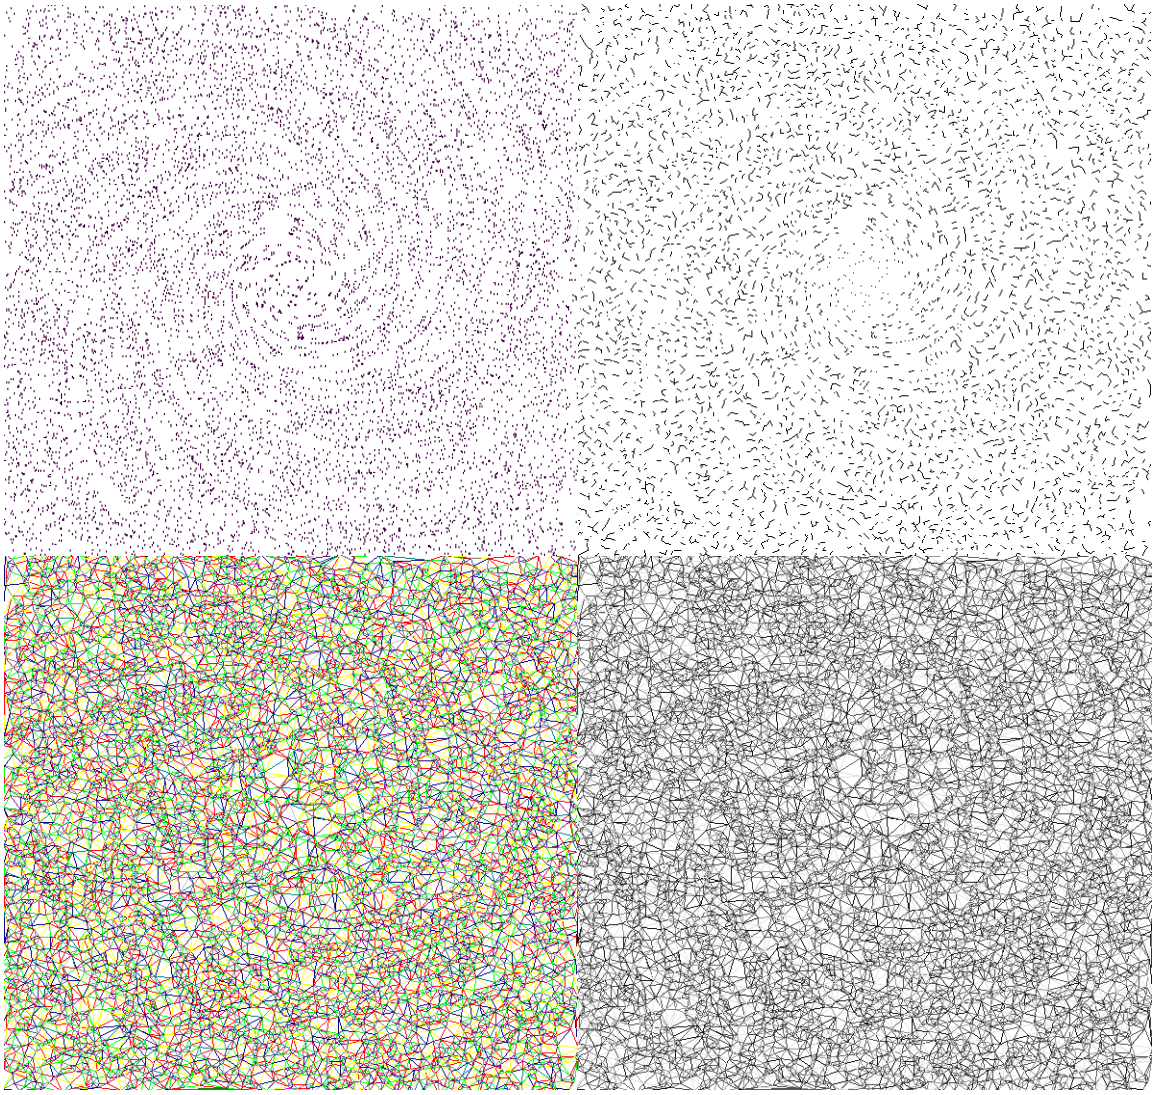

Figure 12.1: Upper Left: Glass screw pattern comprising 5,000 pairs rotated at midpoints by 20°. Mean transformation distance is 100.15 screen units and mean nearest neighbour distance is 25.04 screen units. **Upper Right:** Nearest neighbour edges for Glass screw pattern. **Lower Left and Right:** Delaunay neighbour edges for Glass screw pattern. Transformational structure is clearly seen in edges alone, and even though mean transformation distance well exceeds mean nearest neighbour distance, transformational structure is clearly seen in the nearest neighbour edges.

Figure 12.2 provides a similar example to Figure 12.1, but the number of dots involved is 28,000 and the pairs are rotated at midpoints by 330°. In both cases, mean nearest neighbour distances are much less than mean transformation distances. According to some authors (e.g., Dakin, 1997; Glass, 1979; Maloney, Michinson, & Barlow, 1987; Stevens, 1978), nearest neighbours should be rendered ineffective under such conditions. Since pertinent displays show intended transformational structure with edges alone, particularly with nearest neighbour edges, then nearest neighbour is surely a basic structural device. All Delaunay neighbours make an obvious contribution, but what of the contribution of neighbours beyond those of Delaunay?

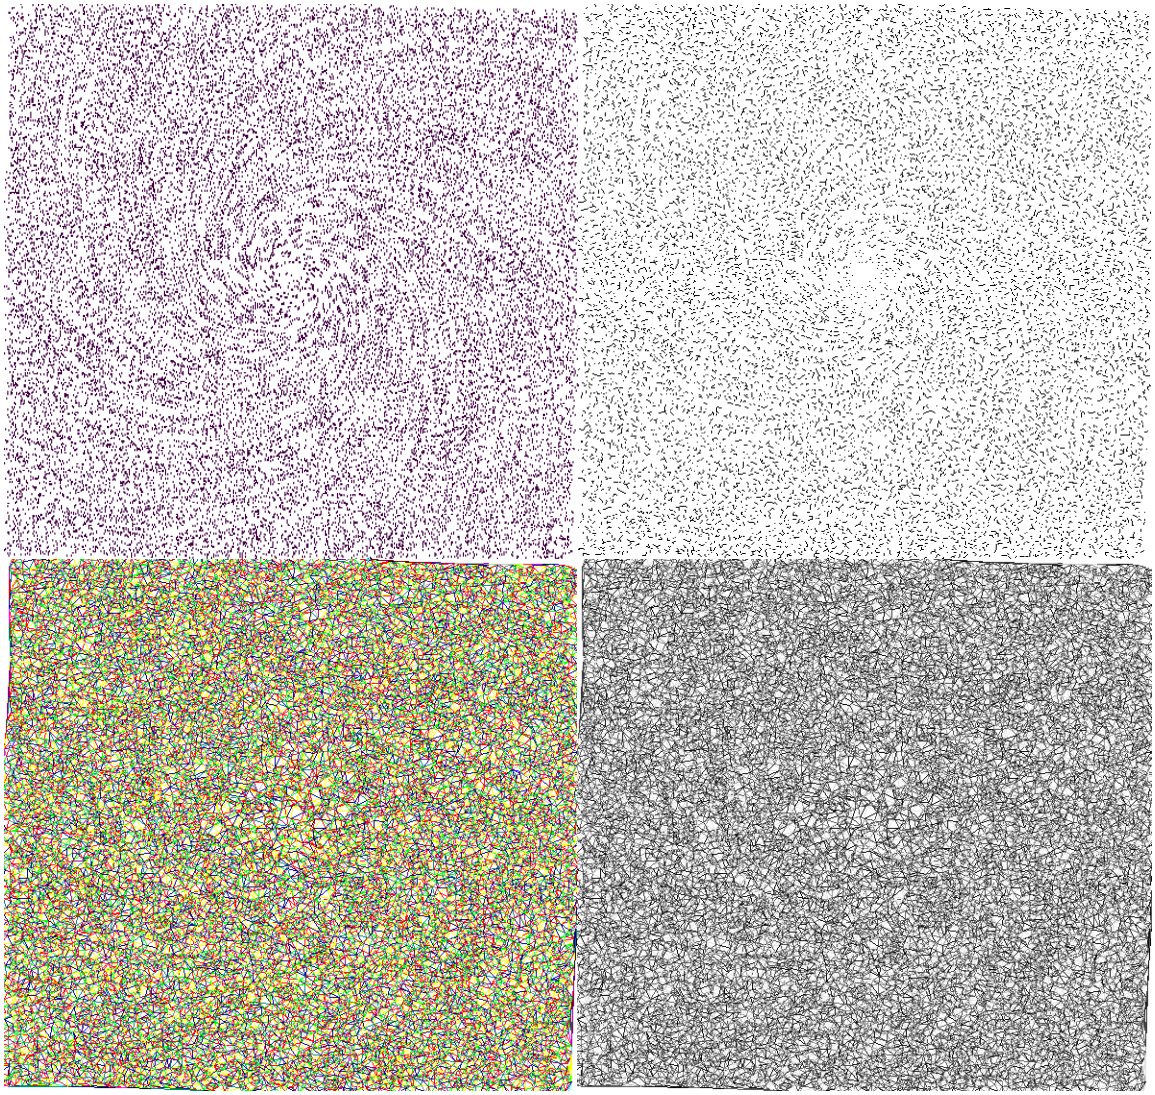

Figure 12.2: Upper Left: Glass screw pattern comprising 14,000 pairs rotated at midpoints by 330°. Mean transformation distance is 100.00 screen units and mean nearest neighbour distance is 15.02 screen units. **Upper Right:** Nearest neighbour edges for Glass screw pattern. **Lower Left and Right:** Delaunay neighbour edges for Glass screw pattern. Transformational structure is clearly seen in edges alone, and even though mean transformation distance exceeds mean nearest neighbour distance substantially more than for the previous figure, transformational structure is seen just as clearly in the nearest neighbour edges.

Figure 12.3 is composed of seven computer screenshots. The upper left panel shows a Glass screw pattern comprising 5,000 dots. The upper right panel shows corresponding nearest neighbour edges. Again, the mean nearest neighbour distance is well less than the mean transformation distance. The lower panels show 3rd nearest neighbour edges, 6th nearest neighbour edges 9th nearest neighbour edges, 12th nearest neighbour edges, and 15th nearest neighbour edges from left to right down the page.

The distribution for number of Delaunay neighbours of each dot for a random display peaks at six, and neighbours are effectively exhausted at around eleven. With reference to Figure 2.18, page 27, the probability of getting 14 neighbours is estimated at just .000005. Hence it may be no coincidence that structure is substantially exhausted at about 12th nearest neighbours.

I hypothesize that neighbours beyond those of Delaunay do not contribute in any substantial way to structure in Glass pattens and some other patterns with similar low-level

attributes. However, the exhaustion of Delaunay neighbour links between transformation partners at high pattern densities should not be regarded as an indication of depletion of the transformation effect. The hypothesis does *not* say that Delaunay neighbour links are required to bridge transformation partners in order to perceive a transformation effect.

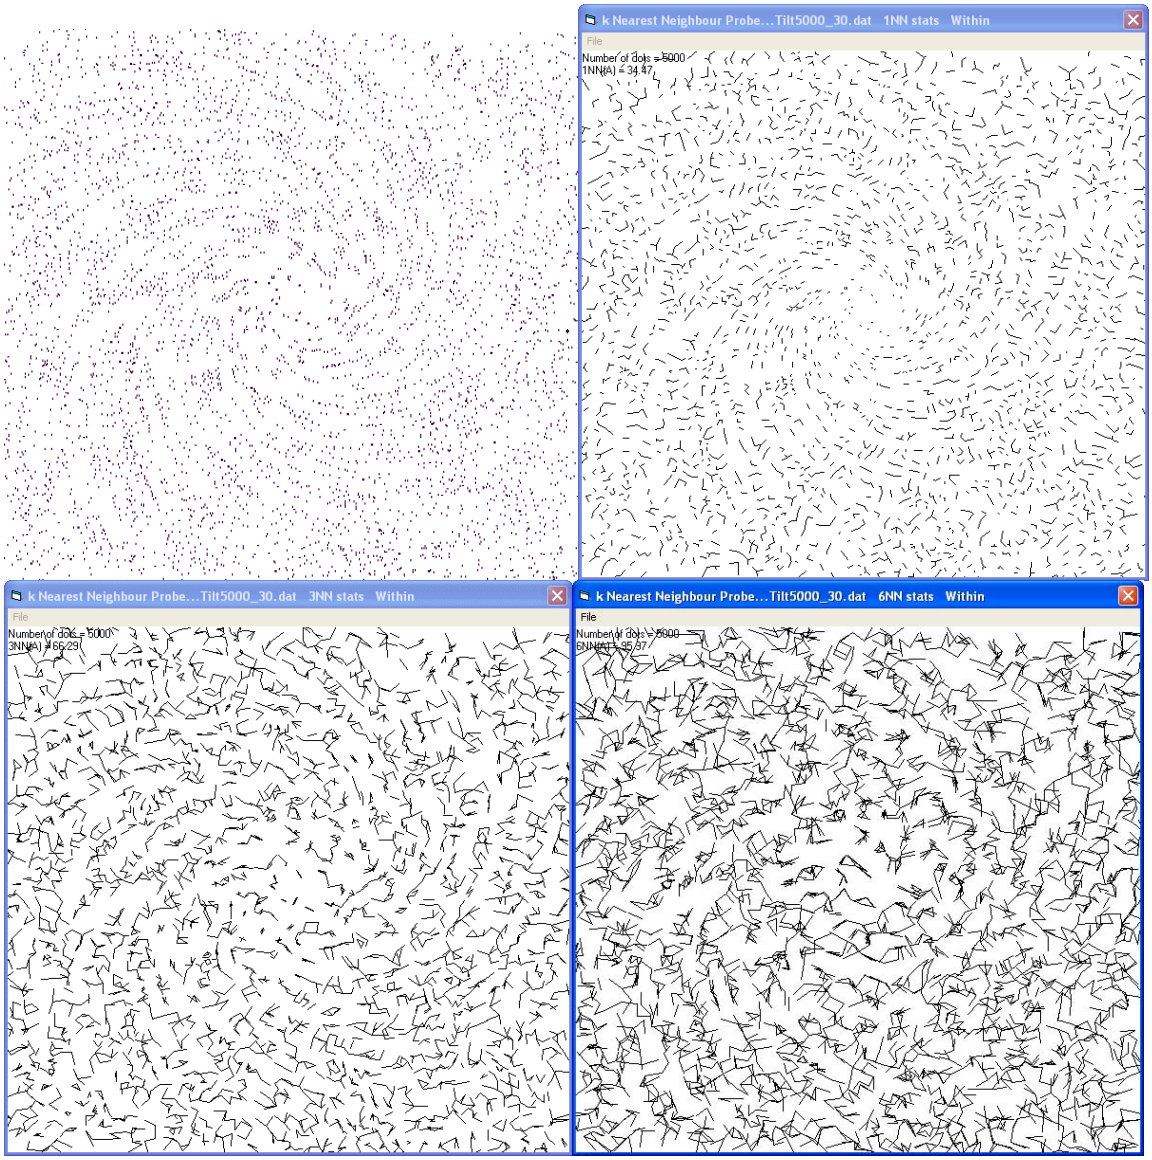

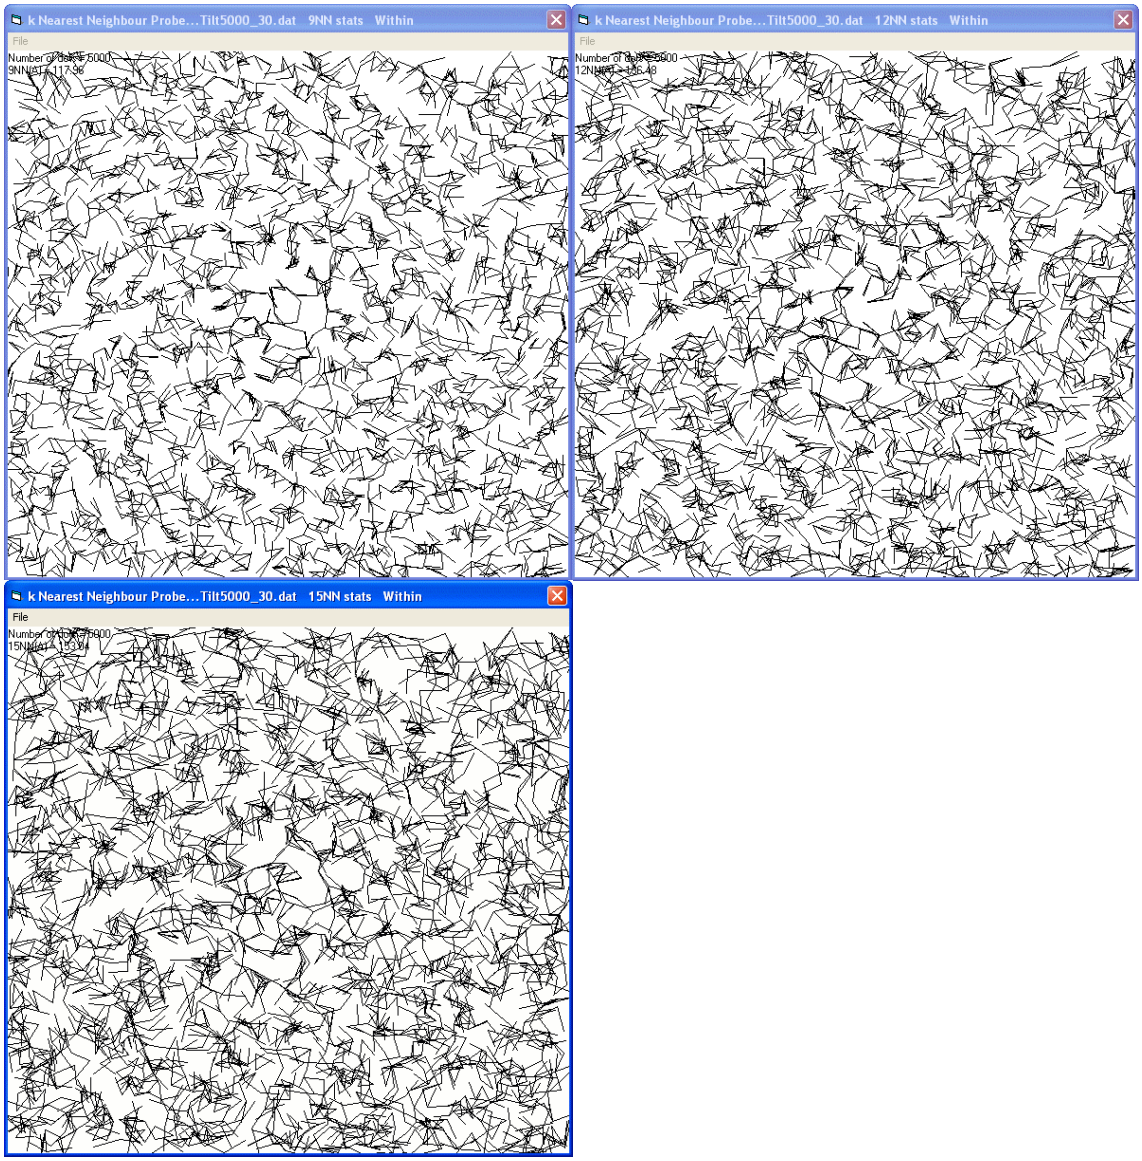

Figure 12.3: Upper Left (previous page)**:** Glass screw pattern comprising 2,500 pairs rotated at midpoints by 30°. Mean transformation distance is 99.41 screen units and mean nearest neighbour distance is 34.47 screen units. **Upper Right** (previous page)**:** Nearest neighbour edges for Glass screw pattern. **Others** (including this page)**:** Progressively larger ranked neighbours, corresponding to progressively longer edges, for Glass screw pattern (see preceding text). Structure is substantially exhausted at about 12th nearest neighbours, shown in the second to last panel.

Caelli (1981) suggested that the mean and variance of nearest neighbour distances for Glass patterns, including those for which mean transformation distance is substantially greater than mean nearest neighbour distance, are different to those for a random display with the same number of dots. Nonetheless if Caelli's proposal is pushed to explain the streakiness in suitably dense Glass patterns, it fails. Differences in means and variances for nearest neighbour distances become negligible. Indeed when means and variances for other measures—nearest neighbour angular offset from transformation direction(s), for example—are used in suitably dense Glass patterns, there is a streakiness that is no longer differentiated.

However, similar analyses of larger ranked neighbours confirm the flavour of Caelli's proposal, but not without a twist. Whereas nearest neighbour analysis indicates clustering for Glass patterns ranging to reasonably dense, and eventually randomness for suitably dense Glass patterns, larger ranked neighbour analysis indicates regularity, and eventually

randomness; while still showing randomness for noise of course. And this can be readily explained.

Figures 12.4 and 12.5 derive from Glass translation patterns of the same density as the Glass screw pattern of Figure 12.3. Figure 12.4 shows a plot of nearest neighbour distances for a Glass translation pattern (magenta) displayed over a plot of nearest neighbour distances for the same number of noise points (black); which provides a better view in this case. For the Glass translation pattern, mean nearest neighbour distance is well less than Glass translation distance. Nearest neighbour distances for the Glass pattern are clipped at translation distance: enough that they stand out against the smaller nearest neighbour distances, and that they are readily distinguished from the larger nearest neighbour distances for noise. Some nearest neighbours for the Glass pattern are not translation counterparts, in which case they have to be closer. It is clear that mean and variance for Glass nearest neighbours are somewhat less than those for noise; hence the verdict of clustering. However, if the translation distance is not less than the larger noise distances, nearest neighbour mean and variance fail to discriminate.

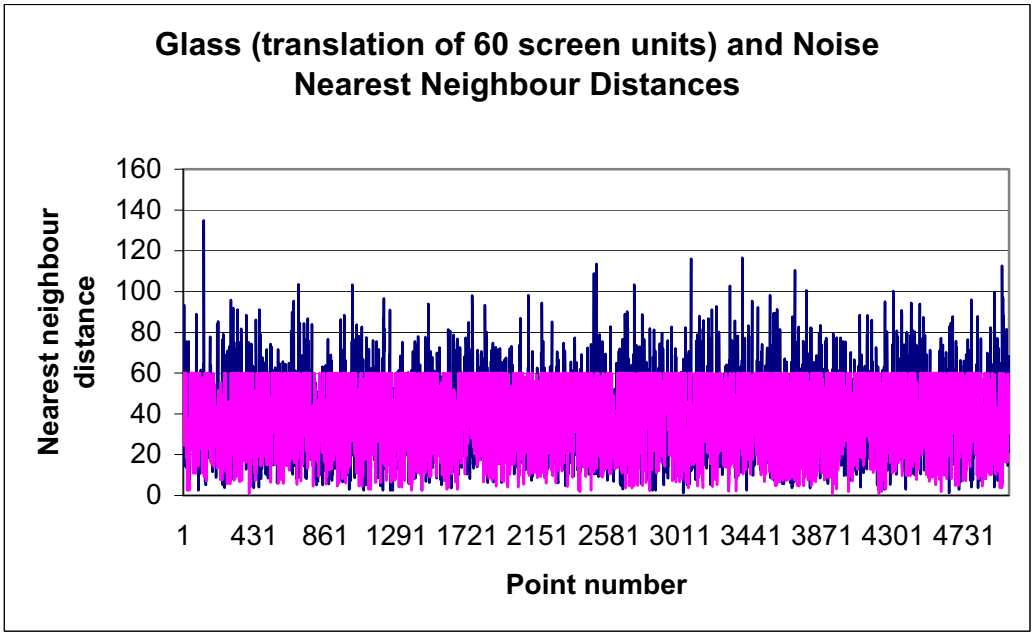

Figure 12.4: Nearest neighbour distances (magenta) derived from a Glass translation pattern comprising 2,500 pairs, compared to nearest neighbour distances (black) derived from 5,000 noise points. Glass transformation distance is 60 screen units, or about 1.71 times the mean nearest neighbour distance for the Glass pattern.

Figure 12.5 shows a plot of second nearest neighbour distances for a Glass translation pattern (magenta). This time, the plot of second nearest neighbour distances for the noise points (black) is displayed over that for the Glass pattern; which provides a better view in this case. Translation distance is now a little greater than the larger nearest neighbour distances for noise, as shown in Figure 12.4. Second nearest neighbour distances for the Glass pattern are clipped at translation distance, but not as many reach the translation distance as do nearest neighbours for the smaller translation distance of the previous example. More second nearest neighbours are closer than translation distance, and third and maybe fourth nearest neighbours are needed to increase clipping. However, some of these extend beyond translation distance as well, which diminishes the measure, and this is where Delaunay ranked neighbours prove

useful.[1] Yet, even without the advantage of Delaunay it is clear that mean and variance for Glass second nearest neighbours are greater than those for noise; hence the verdict of regularity.

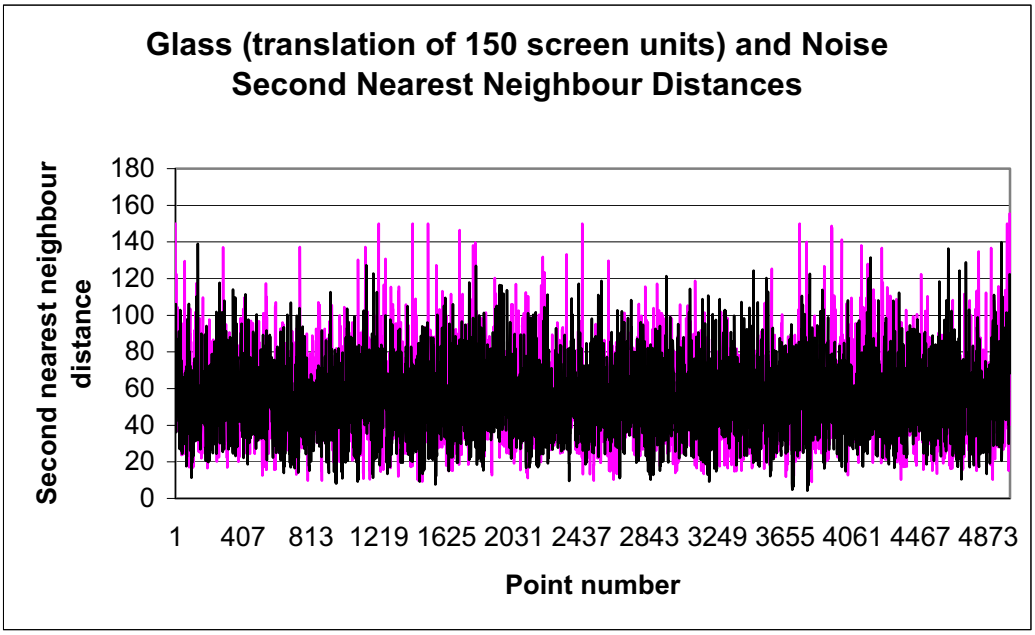

Figure 12.5: **S**econd nearest neighbour distances (magenta) derived from a Glass translation pattern comprising 2,500 pairs, compared to second nearest neighbour distances (black) derived from 5,000 noise points. Glass transformation distance is 150 screen units, or about 2.78 times the mean second nearest neighbour distance for the Glass pattern.

Given some transformation magnitude, there is an optimal loading spread for neighbours as outlined on page 23 of Chapter 2, and by the graphs of Figures 2.11 to 2.20. (For the purpose outlined, increasing the transformation magnitude is equivalent to increasing the density; as long as transformation counterparts taken beyond the display form are included in calculations.) And as consistent distances become *larger* relative to mean distances belonging to respective neighbour rankings for equivalent noise, then any such consistency has to be judged as regular.

To see this more clearly, consider a Glass pattern with a relatively large transformation magnitude alongside an equal number of noise dots. The overall effect of transformation counterparts is to increase mean neighbour distance beyond that for the noise. This equates to some degree of regularity (as outlined in Chapter 3) for the Glass pattern. At a relatively small transformation magnitude, the overall effect of transformation counterparts is to reduce mean neighbour distance below that for the noise. This equates to some degree of clustering (as outlined in Chapter 3) for the Glass pattern.

Nonintersecting paths

Kolers (1972) has extensively investigated apparent motion between successive frames of sparse arrays. Among his findings, Kolers determined that apparent motion paths avoid crossing. Nearest neighbour paths do not cross, and neither do any of the Delaunay

[1] Delaunay neighbour ranking is not the same as near neighbour ranking; even though they are affiliated. This can be appreciated by observing that Delaunay neighbour links do not cross, but near neighbour links, other than nearest neighbour links, can cross; statistically more so with larger ranked neighbours and greater dot density. Moreover Delaunay neighbours are essentially more limited in number of ranks than near neighbours.

paths, whether dealing with neighbours *within* a pattern set or *between* (successive) pattern sets. Nonetheless, regarding near neighbours, Figure 12.3 indicates that even third nearest neighbours, at the relatively high dot density involved, resist crossed paths to a fair extent. The importance of nonintersecting paths with regard to apparent motion, and also travelling salesperson solutions, is addressed in the unpublished study contained in Appendix B (Vickers, Preiss, & Hughes, 2003).

A point of technical interest

In Chapter 8, page 152, it was shown how Glass screw patterns form a continuum between rotation and dilation patterns by simply rotating the dot pairs. The same segment of program code that forms a Glass pattern is used to rotate the pairs by re-referencing from the centre of rotation of the Glass pattern to midpoints of pairs.

Furthermore by referencing the centre of rotation of a Glass pattern to a location far removed from the pattern itself, any Glass translation pattern can effectively be generated. Of course the transformation angle needs to be suitably reduced, in proportion to radial distance. The point is that all the common Glass patterns can effectively be produced with the same few lines of program code by changing a couple of parameters. For those interested in computer programming, the segment of Visual Basic 6 code for a universal Glass pattern generator shown in Appendix D serves to provide some idea of the procedure. (It can be copied into a Visual Basic 6 form, and directly executed for the purpose of manipulating parameters.) The methodology essentially reduces Glass patterns to one type, in which angles subtending Glass pairs form elements of consistency.

A point of practical interest

Undoubtedly, transformation regularity, which is consistency or continuity of orientation along with consistency of magnitude, is important to Glass pattern detection. (Consistency of a quantity can mean consistent change, or consistent change of change, as well as constancy.) Here, such consistency highlights inherent anisotropy, or cardinal direction, in Glass patterns. Rotating pairs of a Glass rotation pattern in either direction by one or another appropriate angle results in a Glass screw pattern. Independently generated sequences of these—as per the procedure described for Glass rotation patterns in the last chapter—produce a strong spiralling effect over a suitable range of angles. The spiralling effect is inwards for one direction of pair rotation and outwards for the other direction.

Direction and distance are well known quantities in neurophysiology concerned with the visual system. An important conjecture of this thesis is that these induce a most likely transformation, which acts to filter out a plethora of less likely possibilities.

Voronoi, Delaunay overture

Voronoi partitioning provides a basis for Glass phenomena to be conceived in terms of MacKay patterns. For dot patterns, members of pairs connected by Delaunay neighbour links generally have Voronoi cell boundaries normal to, and midway between, their links. This is the case whether or not a link is reflexive. (Sometimes, of course, a Delaunay neighbour link has only the extension of a Voronoi cell boundary normal to, and midway between, its link.)

If a Glass pattern was partitioned by the visual system in terms of Voronoi, then there could be a conspicuous proportion of cell boundaries normal to transformation orientation that might produce a MacKay-like preference for transformation direction. Moreover Glass patterns have a range of transformation displacements, not too small nor large, over which Glass pattern structure is more prominent (see Dakin, 1997; Dry, Vickers, Lee, & Hughes,

2004).[2] Hence, as a Glass pattern becomes more densely populated, effect due to decreasing neighbour distances earlier in ranks of Delaunay neighbours abates, while it transfers to later ranked neighbours, and a conspicuous proportion involves transformation pairs. This, while boundaries normal to notional links between members of such pairs could invoke the Mackay effect.

Figures 12.6 to 12.10 depict the theme with a mix of translation and rotation Glass patterns. Transformation partners that are nearest neighbours are encircled in magenta and transformation partners that are other ranked neighbours are encircled in green. Transformation normals (Voronoi boundaries perpendicular to transformation direction) are shown in red. Note that as patterns become more dense the ratio of transformation normals to dots reduces; and this concurs with the description outlined in Chapter 2, pages 22 to 30. For a sparse Glass pattern the ratio of transformation normals to dots is generally .5, i.e. there is a transformation normal for every Glass pair of dots.[3] And for an extremely dense Glass pattern, in which Delaunay neighbours as transformation links are substantially exhausted, the ratio of transformation normals to dots is expected to approach that for Voronoi boundaries that have the same sense of orientation in an equivalent density noise display. However, the situation is not quite that straightforward. Noise displays are provided for comparison with patterned situations.

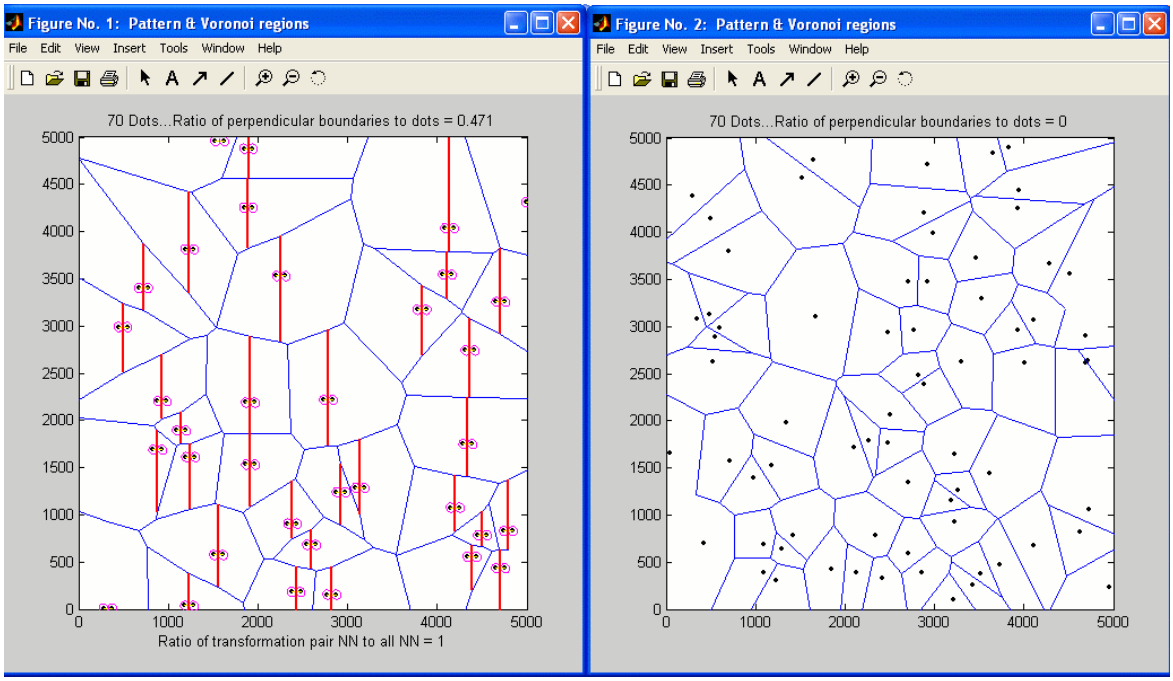

Figure 12.6: Left: Horizontal translation Glass pattern of 70 dots. All nearest neighbour distances are translation distances. Hence the ratio of nearest neighbours for transformation pairs to nearest neighbours for all pairs is 1. Translation normals (vertical Voronoi boundaries in this case) are shown in red, and other Voronoi boundaries are shown in blue. Ratio of translation normals to dots should be .5, instead of .471, but for a foible of MATLAB (see footnote 3). **Right:** Same treatment for the same number of noise dots. No Voronoi boundary has the same sense of orientation as the red translation normals of the left panel, hence the corresponding ratio is 0.

---

[2] Since patterns scale perfectly well over a wide range of comfortable viewing distances, transformation displacements in terms of absolute units, particularly angular units for dipole lengths as subtended at the eye at some viewing distance, are not particularly relevant to this thesis.

[3] Owing to a foible of MATLAB, in which it plots only the *bounded* Voronoi cells, the ratio is a little short in the examples, with less error accompanying increasing density.

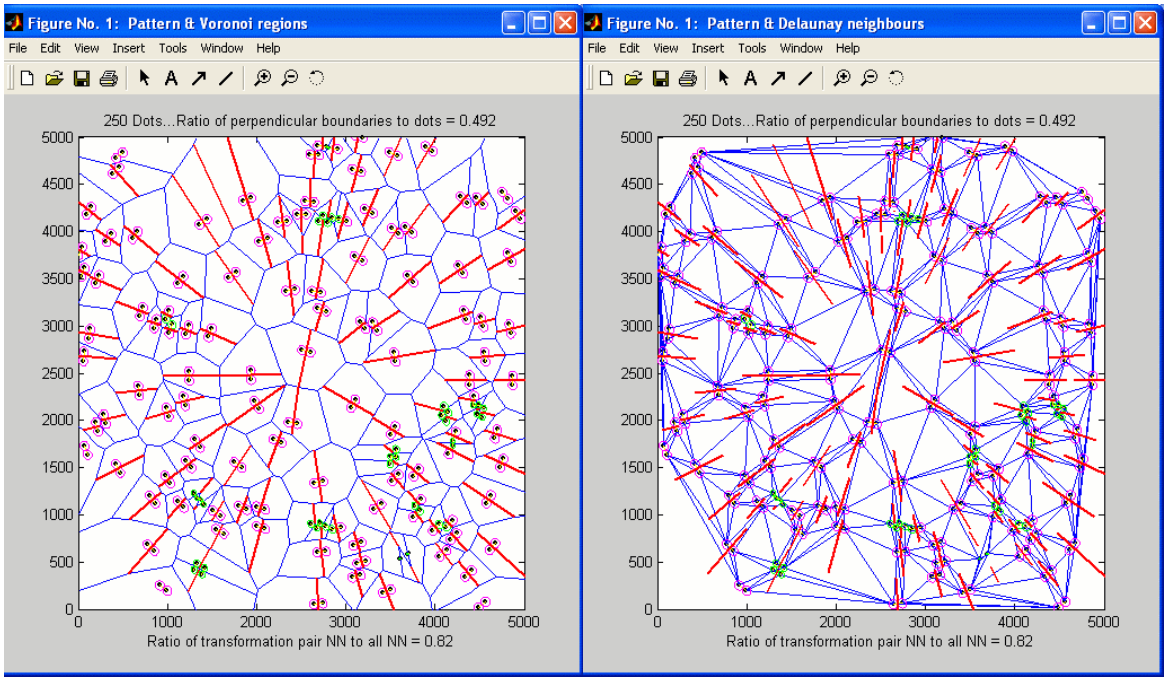

Figure 12.7: Left: Rotation Glass pattern of 250 dots, for which not all nearest neighbour distances are rotation distances. The ratio of nearest neighbours for transformation pairs to nearest neighbours for all pairs is .82. Rotation normals (radial Voronoi boundaries in this case) are shown in red. Ratio of rotation normals to dots is .492. **Right:** Delaunay display along with rotation normals for the same Glass pattern. Strong connections (low rank Delaunay links) exist across the rotation normals for all transformation regularities.

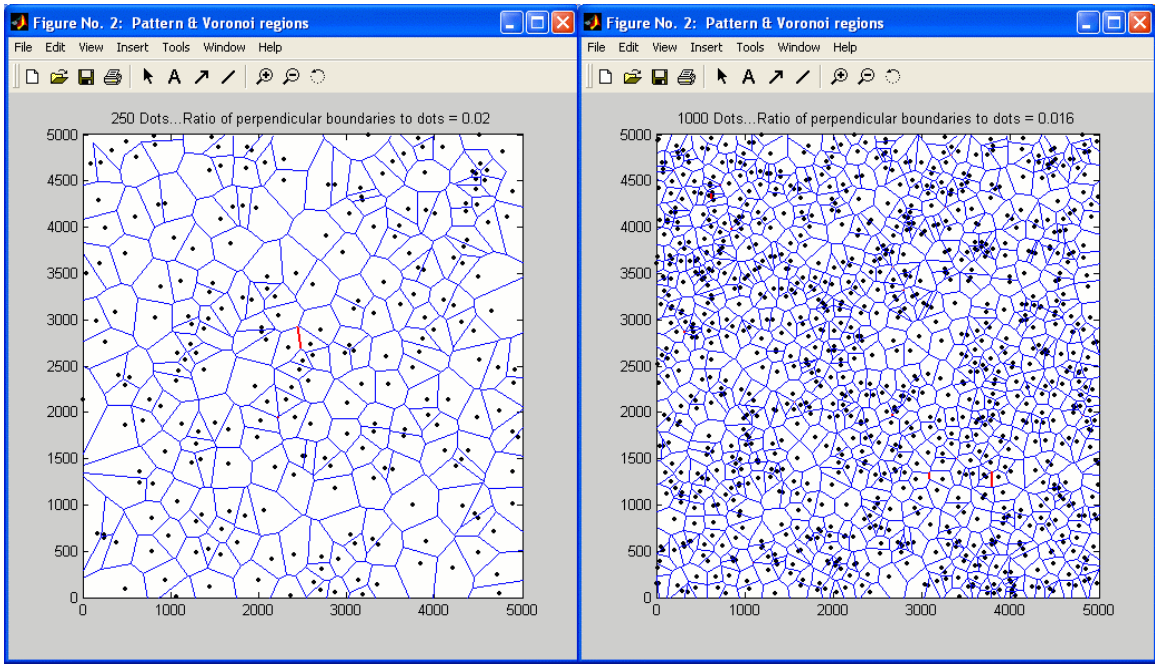

Figure 12.8: Left: Noise display for same number of noise dots as per Figure 12.7. Ratio of radial Voronoi boundaries to dots is .02. **Right:** Display for same number of noise dots as per the left panel of Figure 12.9. Ratio of vertical Voronoi boundaries to dots is .016.

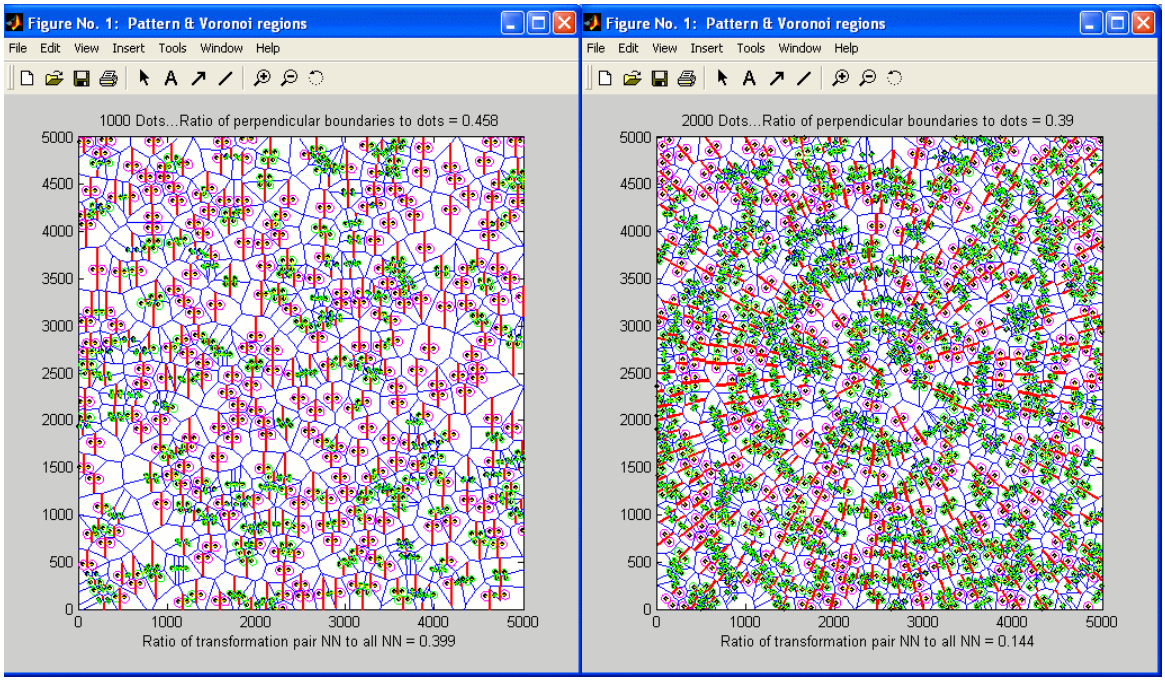

Figure 12.9: Left: Horizontal translation Glass pattern of 1,000 dots. The ratio of nearest neighbours for transformation pairs to nearest neighbours for all pairs is .399. Ratio of translation normals to dots is .458. **Right:** Rotation Glass pattern of 2,000 dots. The ratio of nearest neighbours for transformation pairs to nearest neighbours for all pairs is .144. Ratio of rotation normals to dots is .390.

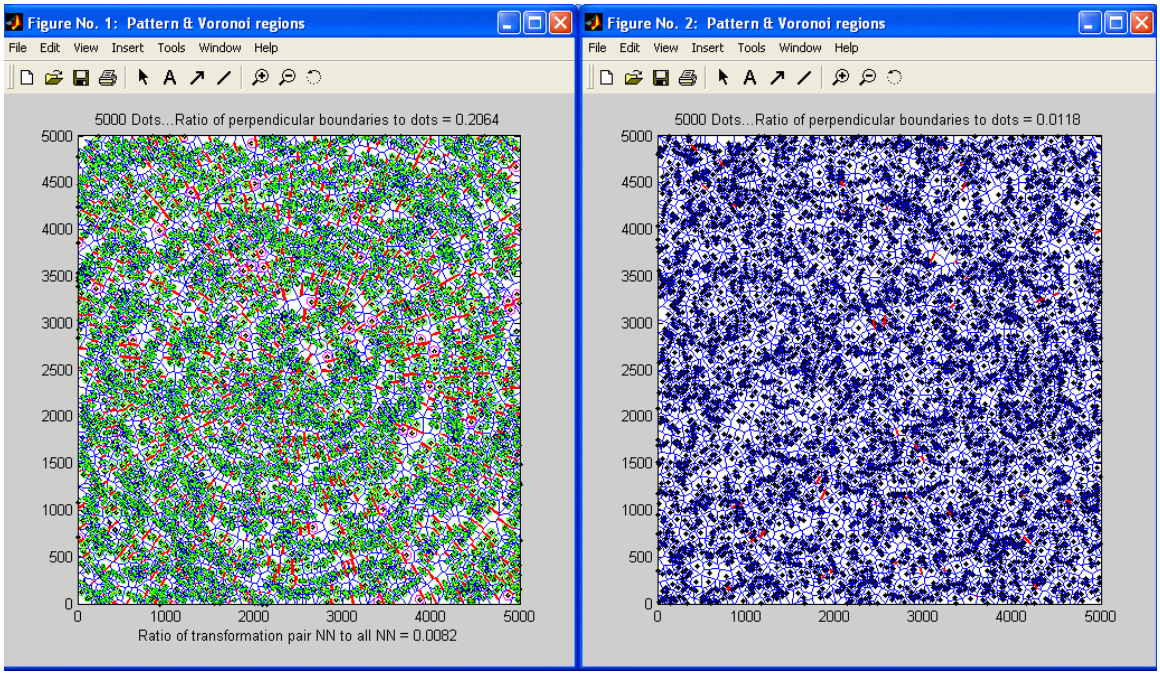

Figure 12.10: Left: Rotation Glass pattern of 5,000 dots. The ratio of nearest neighbours for transformation pairs to nearest neighbours for all pairs is .0082. Ratio of rotation normals to dots is .2064. **Right:** Same treatment for same number of noise dots. Ratio of radial Voronoi boundaries to dots is .0118.

The Delaunay interpretation shown in the right panel of Figure 12.7 indicates strong connections (i.e., relatively close neighbours) across transformation normals for all transformation regularities. This remains the case for even dense Glass patterns, and only

reduces when Delaunay neighbours for transformation distances expire under extreme density.

Figure 12.11 is a rotation Glass pattern of intermediate density. Its density is ten times greater than that for the right panel of Figure 12.7. There is a Delaunay neighbour connection, in blue, shielded by every rotation normal, shown in red. Comparatively strong connections prevail across the rotation normals for transformation regularities.

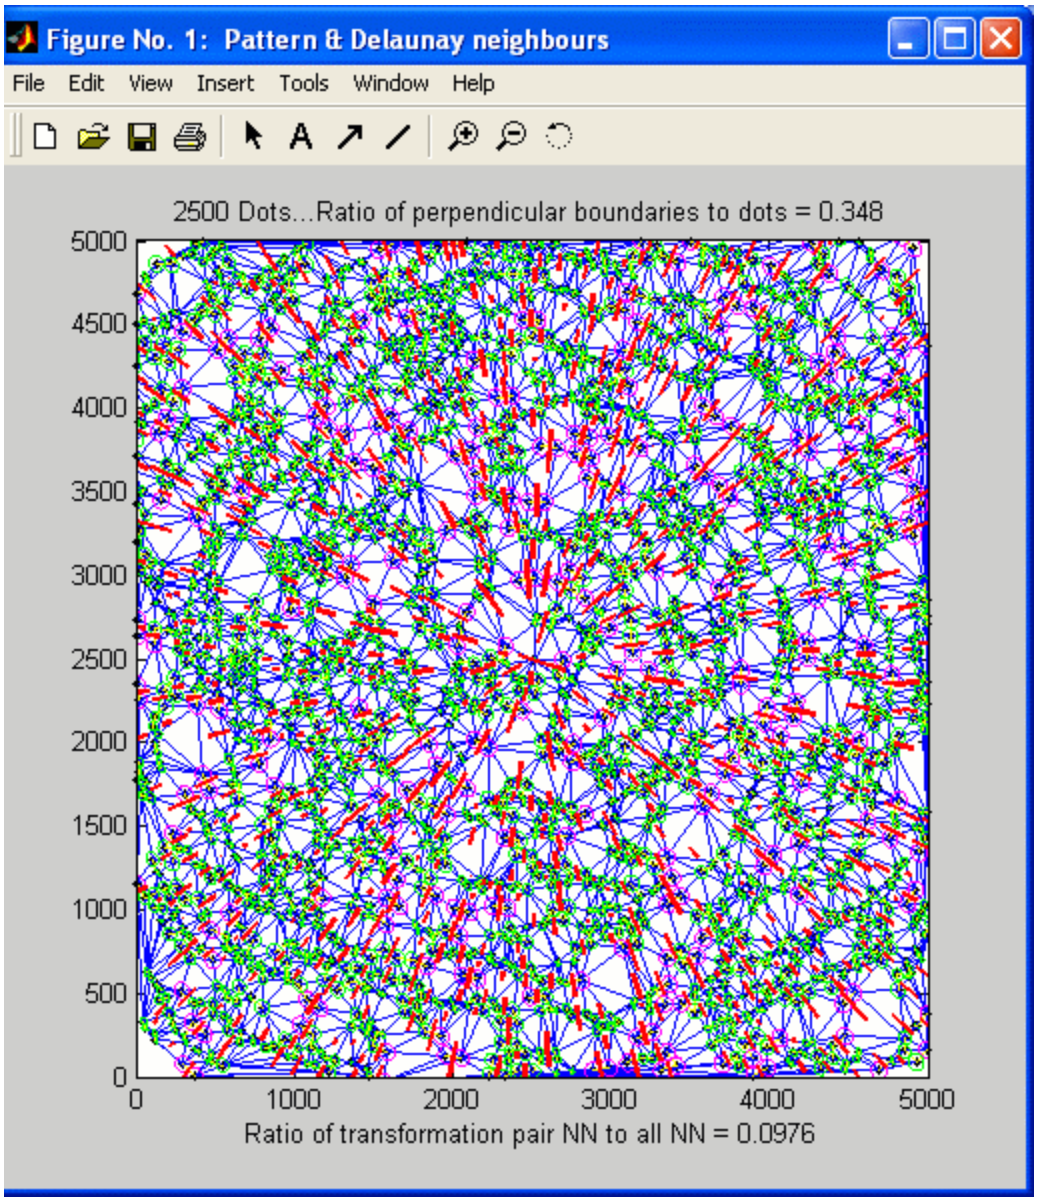

Figure 12.11: Delaunay display along with rotation normals for a rotation Glass pattern of 2,500 dots. Comparatively strong connections exist across the rotation normals for transformation regularities.

Table 12.1 summarises the trend for the Voronoi, Delaunay figures as density increases. The middle column shows percentages of nearest neighbours for transformation pairs relative to nearest neighbours for all pairs (number of transformation nearest neighbours multiplied by 100, then divided by number of all nearest neighbours). The right-hand column shows

corresponding percentages of transformation normals relative to numbers of dots (number of transformation normals multiplied by 100, then divided by number of dots).

Table 12.1: Quantitative changes with increasing Glass pattern density (see preceding text).

Number of dots	Transform NN / all NN	Transform normals / dots
70	100.00%	50.00%
250	82.00%	49.20%
1,000	39.90%	45.80%
2,000	14.40%	39.00%
2,500	9.76%	34.80%
5,000	0.82%	20.64%

Table 12.1, along with the likes of the left panel of Figure 12.10, indicates that connections for which nearest neighbour distances are transformation distances can be a small proportion of all nearest neighbour distances, while the Glass effect remains strong. Further investigation might inquire as to whether or not there is a qualitative difference in the Glass effect upon depletion of such indicators. Even so, the next step is to see if the ratios of transformation normals to dots for extremely dense Glass patterns approaches the regions of those for Voronoi boundaries that have the same orientation(s) in equivalent density noise displays. If so, then few such Glass pairs are connected by Delaunay neighbours.

Figure 12.12 shows a suitably dense rotation Glass pattern. There are relatively few transformation normals, especially about regions at three o'clock and eight o'clock. Despite the fact that the ratio is almost as low as that of equivalent noise (see the right panel of Figure 12.14), the rotation effect is still obvious. And it is just as obvious in the two regions with almost no transformation normals.

Figures 12.13 and 12.14 show the rotation effect, via Voronoi tessellation, for 20,000 and 30,000 dots respectively. As before, displays of stimulus dots for the Voronoi diagrams have been omitted. The rotation effect is conspicuous among neighbours for which distances are obviously not transformation distances. Omitting the display of dots renders them literally dimensionless: they are simply sites for which Voronoi boundaries are loci. And as discussed in Chapter 2, pages 27 to 30, dots of a suitable size in a pattern of suitable density would cause the Glass effect to fall away.[4] Nonetheless invoking such a device does not help with the transformation effect as ultimately indicated by Voronoi, Delaunay with no display of stimulus dots.

It is clear that given a random distribution of dots, corresponding other dots offset from these under some transformation provide a guarantee of anisotropy, or greater concentration along the direction(s) of transformation. Hence interspersed and adjacent Glass pairs, which can upset neighbour considerations between transformation partners, provide connections between partners of different Glass pairs that outline the transformation scheme at work. In Chapter 2, page 23, a range of directions $\pm 10^{o}$ about orientation along transformation direction were considered salient, meaning that salient normals would have the same range. This is shown in Figure 12.15, by allowing what were transformation normals to have $\pm 10^{o}$ slack. And it is clear that allowing a slack around orientation along transformation direction

[4] It seems that this is reached at greater densities for rotation than for translation Glass patterns.

does not confer a device that differentiates the Glass effect from noise. There are about as many salient normals for each display.

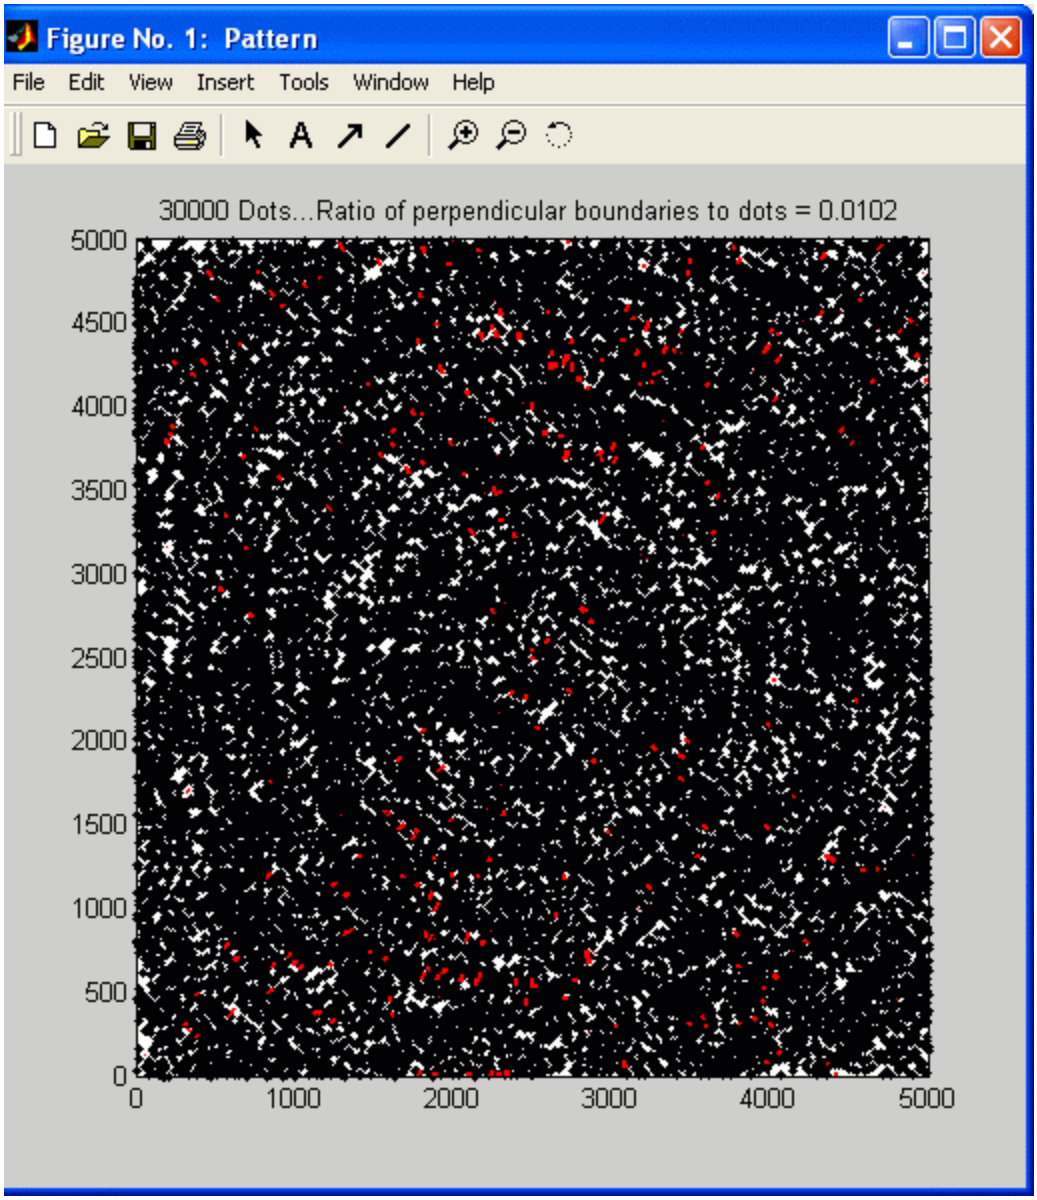

Figure 12.12: Dense rotation Glass pattern. Transformation normals are shown in red. There are relatively few transformation normals, especially about regions at three o'clock and eight o'clock. The ratio of transformation normals to dots is .0102. Despite the fact that the ratio is almost as low as that for equivalent noise (see the right panel of Figure 12.14), the rotation effect is still obvious. And it is just as obvious in the two regions with almost no transformation normals.

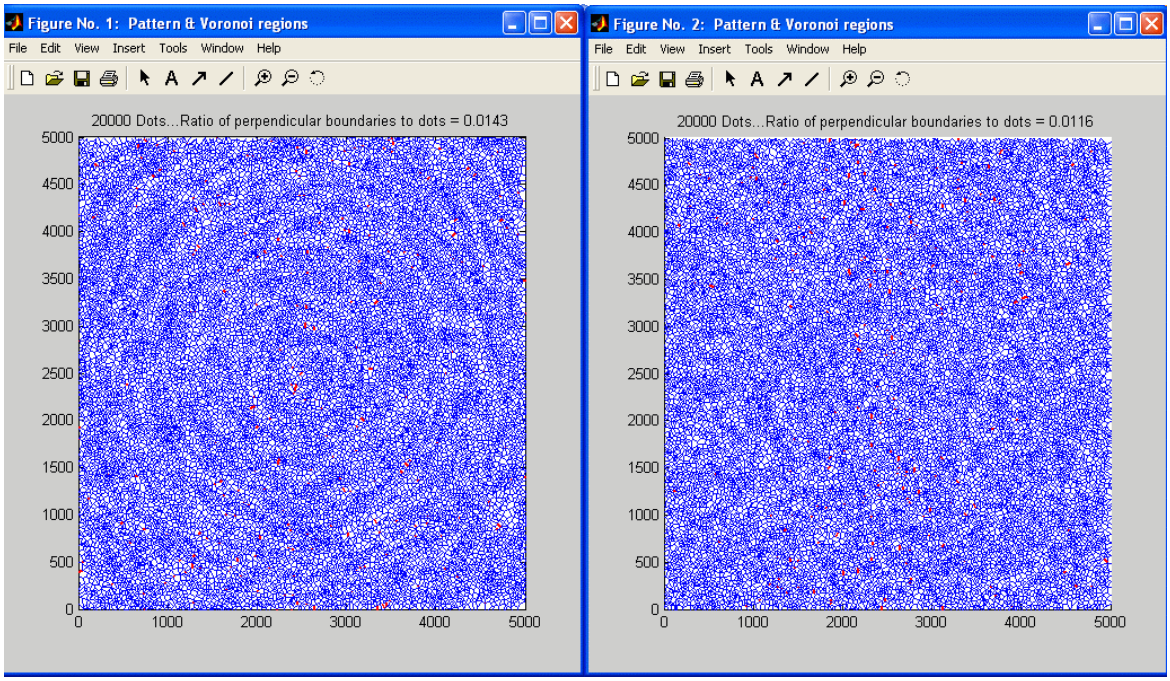

Figure 12.13: Left: Rotation effect, via Voronoi tessellation, for 20,000 dots. As before, displays of stimulus dots for the Voronoi diagrams have been omitted. Transformation normals are shown in red, and the ratio of transformation normals to dots is .0143. The rotation effect is conspicuous among neighbours for which distances are obviously not transformation distances. **Right:** Equivalent density noise display. Voronoi boundaries that form radials from the centre of the display are shown in red, and the ratio of these to dots is .0116.

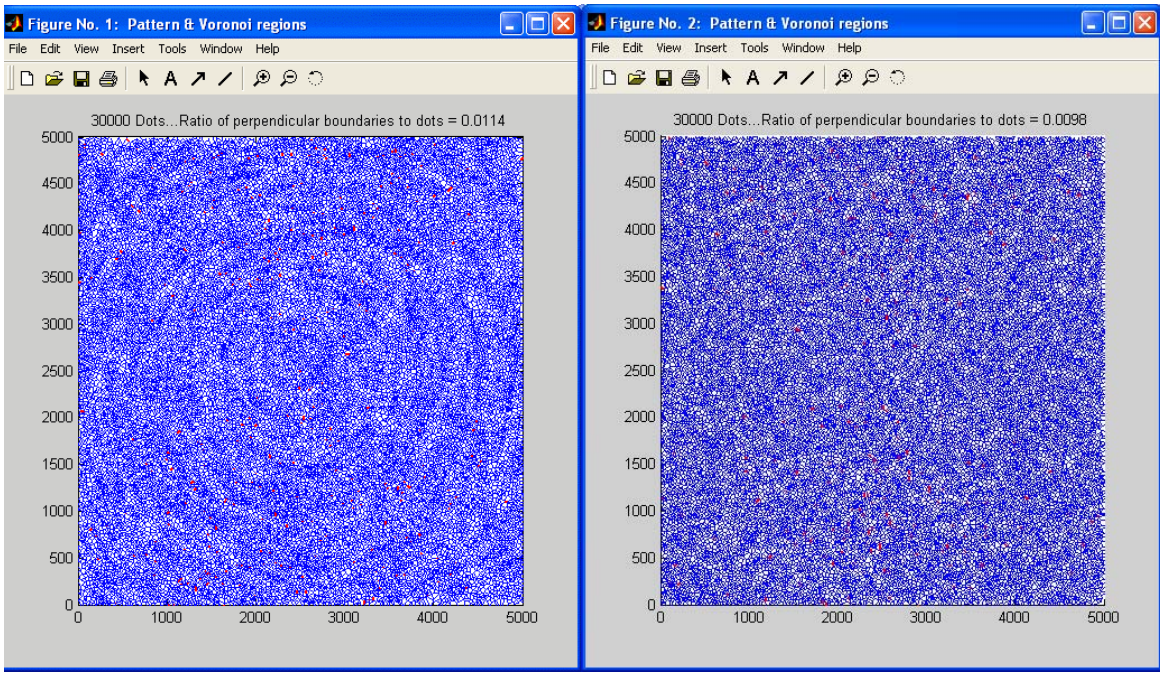

Figure 12.14: Rotation effect, via Voronoi tessellation, for 30,000 dots. As before, displays of stimulus dots for the Voronoi diagrams have been omitted. Transformation normals are shown in red, and the ratio of transformation normals to dots is .0114. The rotation effect is conspicuous among neighbours for which distances are obviously not transformation distances. **Right:** Equivalent density noise display. Voronoi boundaries that form radials from the centre of the display are shown in red, and the ratio of these to dots is .0098.

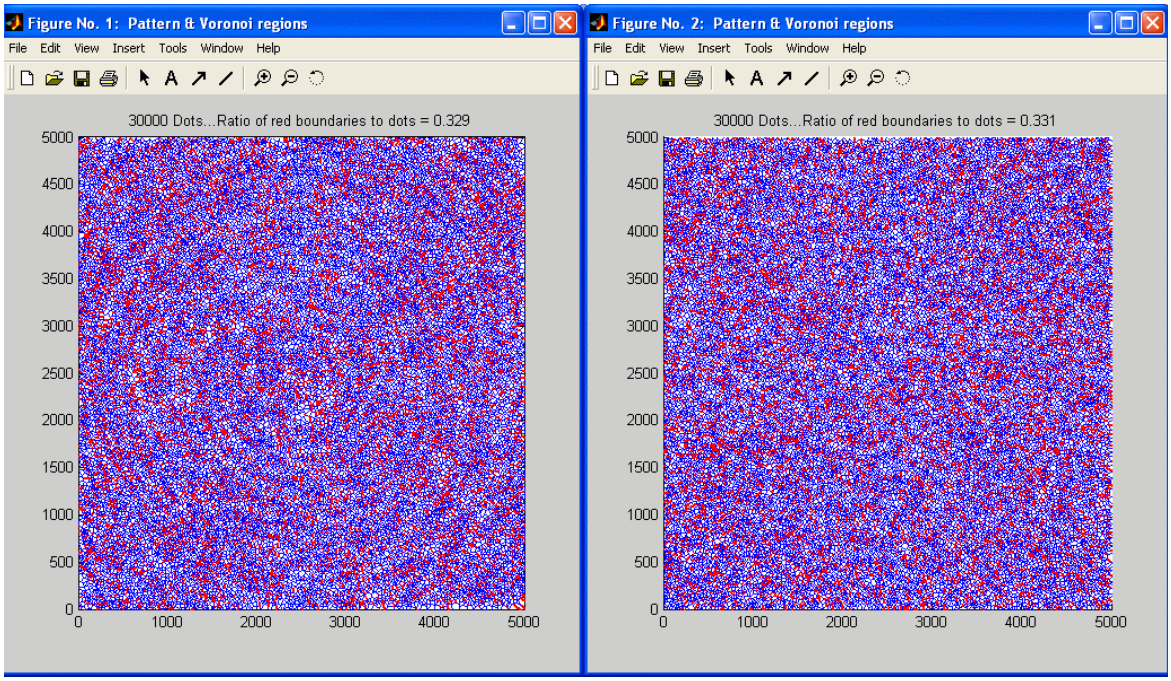

Figure 12.15: Left: Rotation Glass pattern of 30,000 dots. Salient normals (Voronoi boundaries in the range $\pm 10^{o}$ of rotation normals) are shown in red. The ratio of salient normals to dots is .329. **Right:** Same treatment for same number of noise dots. Ratio of salient normals (Voronoi boundaries in the range $\pm 10^{o}$ of radials) to dots is .331. Allowing a slack around orientation along transformation direction does not help to differentiate the Glass effect from noise. There are about as many salient normals for each display.

In summary, at the density associated with 5,000 dots any putative effect due to nearest neighbour links is lost and at 30,000 dots any putative effect due to remaining Delaunay neighbour links is lost. Nonetheless the transformation scheme is made obvious by Voronoi or Delaunay links. When presented with an ultimate consideration that refutes neighbour links between transformation partners, then emphasis on anisotropy seems imperative. Investigation shows that anisotropy is not demonstrated by examination of the difference between average lengths of sets of Voronoi boundaries: one set within a salient range about transformation normals and the other set within the corresponding perpendicular range. The calculation 'average length for the set of Voronoi boundaries in the range $\pm 10^{o}$ about transformation normals' divided by 'average length for the set of Voronoi boundaries in the range $\pm 10^{o}$ about the perpendicular' shows no substantial difference between Glass patterns and equivalent noise.

Nonetheless I hypothesize that with increasing density, notional nearest neighbour links, such as the real ones of Figures 12.1 and 12.2, pages 228 and 229, form notional elements, which are notionally linked by subsequent notional neighbour links and so on. At low to intermediate densities these provide connections between transformation partners, as well as connections between partners of different transformation pairs, that outline the transformation scheme at work. At high densities they predominantly provide connections between partners of different transformation pairs that outline the transformation scheme at work.

Voronoi, Delaunay based connectivity preserves spatial arrangements of underlying dot distributions, while practically all other point-to-point connectivity does not. Hence there must be a measure on some aspect of Voronoi, Delaunay that differentiates Glass patterns from equivalent noise.

Measures of Anisotropy

A measure of anisotropy might involve calculating the mean and standard deviation for each of the *x* and *y* components of nearest, or near, neighbour distances. The ratio (*x* mean) / (*y* mean) and-or ratio (*x* standard deviation) / (*y* standard deviation) could be then taken as the ratio of the axes of a plane elliptical form. Hence a circular form would indicate noise. And the more elliptical a form, the greater would be the degree of anisotropy.

Indeed, this works quite well for horizontal and vertical transformed Glass patterns, even when transformation distance well exceeds mean nearest neighbour distance. And by using components of polar coordinates, it works equally as well for all types of Glass pattern. However, its utility becomes exhausted for suitably dense Glass patterns.

A different procedure that appears to work for Glass patterns of any density that can be reasonably tested involves Hausdorff distance methodology applied to components. In terms of rectangular coordinates, the method involves first finding the set of least distances from each point to every other point in the *x* sense only. The mean and standard deviation are then calculated for the set of least *x* distances. An equivalent procedure yields the mean and standard deviation for the set of least *y* distances. Then the ratio (*x* mean) / (*y* mean) and-or ratio (*x* standard deviation) / (*y* standard deviation) is taken as the ratio of the axes of a plane elliptical form. Note that a point nearest to another point in either the *x* or *y* sense could be well separated from that other point in the resultant sense.

Again, the more elliptical a form the greater is the degree of anisotropy. The method works well for horizontal and vertical translation Glass patterns of any density that can be reasonably tested. For an α^{o} oblique translation Glass pattern, the coordinate system can be rotated by α^{o}. Only least distances are required, even though the method works with second least distances, third least distances, and so on. And by using the ratio of normalized polar coordinate components—proportion of radial distance to proportion of angular distance—the method works equally as well for all types of Glass pattern. Lastly, the reason for the elliptical interpretation, as opposed to simply stating ratios, is that an ellipse, overlying a pattern segment, can confirm degree and orientation of anisotropy by its shape and alignment.

Means and standard deviations for components of second least distance, third least distance, and so on (allied to second nearest neighbour, third nearest neighbour, and so on) work effectively, but the mean and standard deviation for components of least distance (allied to nearest neighbour) clearly differentiates any Glass pattern from noise and pair orientation perturbations without help from the subsequent measures. The measure is shown for the patterns of Figures 12.16 to 12.20. For the horizontal and vertical translation Glass patterns the measure is, respectively, about half and twice that of noise. For the translation Glass pattern with pairs rotated by random amounts, the measure is close to that of noise.

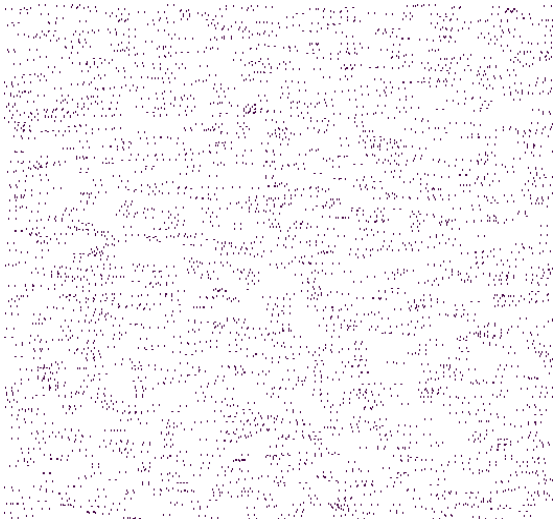

Figure 12.16: Horizontal, translation Glass pattern comprising 2,500 pairs. Transformation distance = 60.00 screen units. Mean nearest neighbour distance = 35.00 screen units and standard deviation = 17.55 screen units. (*x* mean) / (*y* mean) = 0.49 and (*x* standard deviation) / (*y* standard deviation) = 0.50. The measure is about half that for equivalent noise.

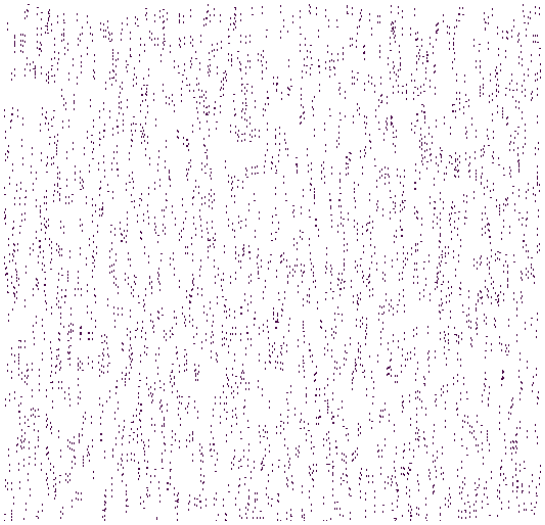

Figure 12.17: Vertical, translation Glass pattern comprising 2,500 pairs. Transformation distance = 60.00 screen units. Mean nearest neighbour distance = 35.01 screen units and standard deviation = 17.09 screen units. (*x* mean) / (*y* mean) = 1.92 and (*x* standard deviation) / (*y* standard deviation) = 1.95. The measure is about twice that for equivalent noise.

Figure 12.18: Horizontal, translation Glass pattern comprising 2,500 pairs. Transformation distance = 200.00 screen units. Mean nearest neighbour distance = 36.65 screen units and standard deviation = 19.19 screen units. (*x* mean) / (*y* mean) = 0.53 and (*x* standard deviation) / (*y* standard deviation) = 0.54. The measure is a little more than half that for equivalent noise because of the larger transformation distance.

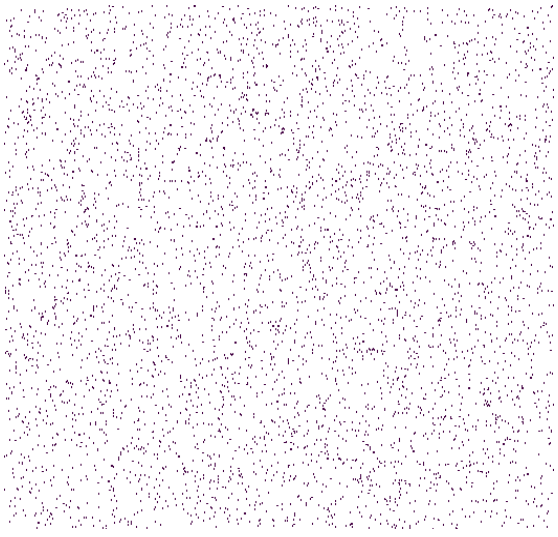

Figure 12.19: Noise comprising 5,000 dots. Mean nearest neighbour distance = 35.62 screen units and standard deviation = 19.05 screen units. (*x* mean) / (*y* mean) = 0.99 and (*x* standard deviation) / (*y* standard deviation) = 0.97.

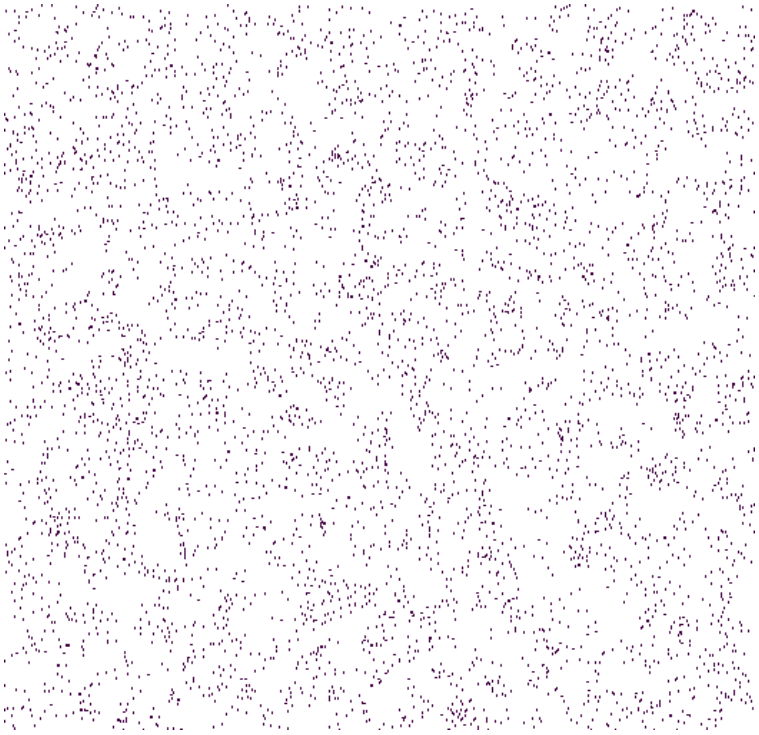

Figure 12.20: Perturbed translation Glass pattern comprising 2,500 pairs that have been rotated by random amounts. Transformation distance = 60.00 screen units. Mean nearest neighbour distance = 36.66 screen units and standard deviation = 17.23 screen units. (*x* mean) / (*y* mean) = 1.01 and (*x* standard deviation) / (*y* standard deviation) = 1.08. The measure is about that for equivalent noise.

Back to Delaunay

The top left panel of Figure 12.21 shows the Delaunay triangulation of 10,000 Glass rotation dots, with constant separation between transformation partners. The bottom left panel shows the Delaunay triangulation of 20,000 such dots. And the two right-hand panels show Delaunay triangulations of respective noise. Again, the display of stimulus dots has been omitted.

The number of triangles fails to differentiate Glass patterns from noise. The mean and standard deviation of least angles of triangles ultimately fails to differentiate Glass patterns from noise. And the mean and standard deviation of lengths of triangle sides ultimately fails to differentiate Glass patterns from noise. That is, distributions of these measures, graphed from associated data, fail to show any differentiation of Glass patterns from noise at middle to higher densities. However, the mean area of triangles is of some help in differentiating Glass patterns from noise, and the standard deviation of areas of triangles proves useful. Table 12.2 shows the means and standard deviations of Delaunay triangulations belonging to respective Glass and noise patterns for the range of measures just outlined.

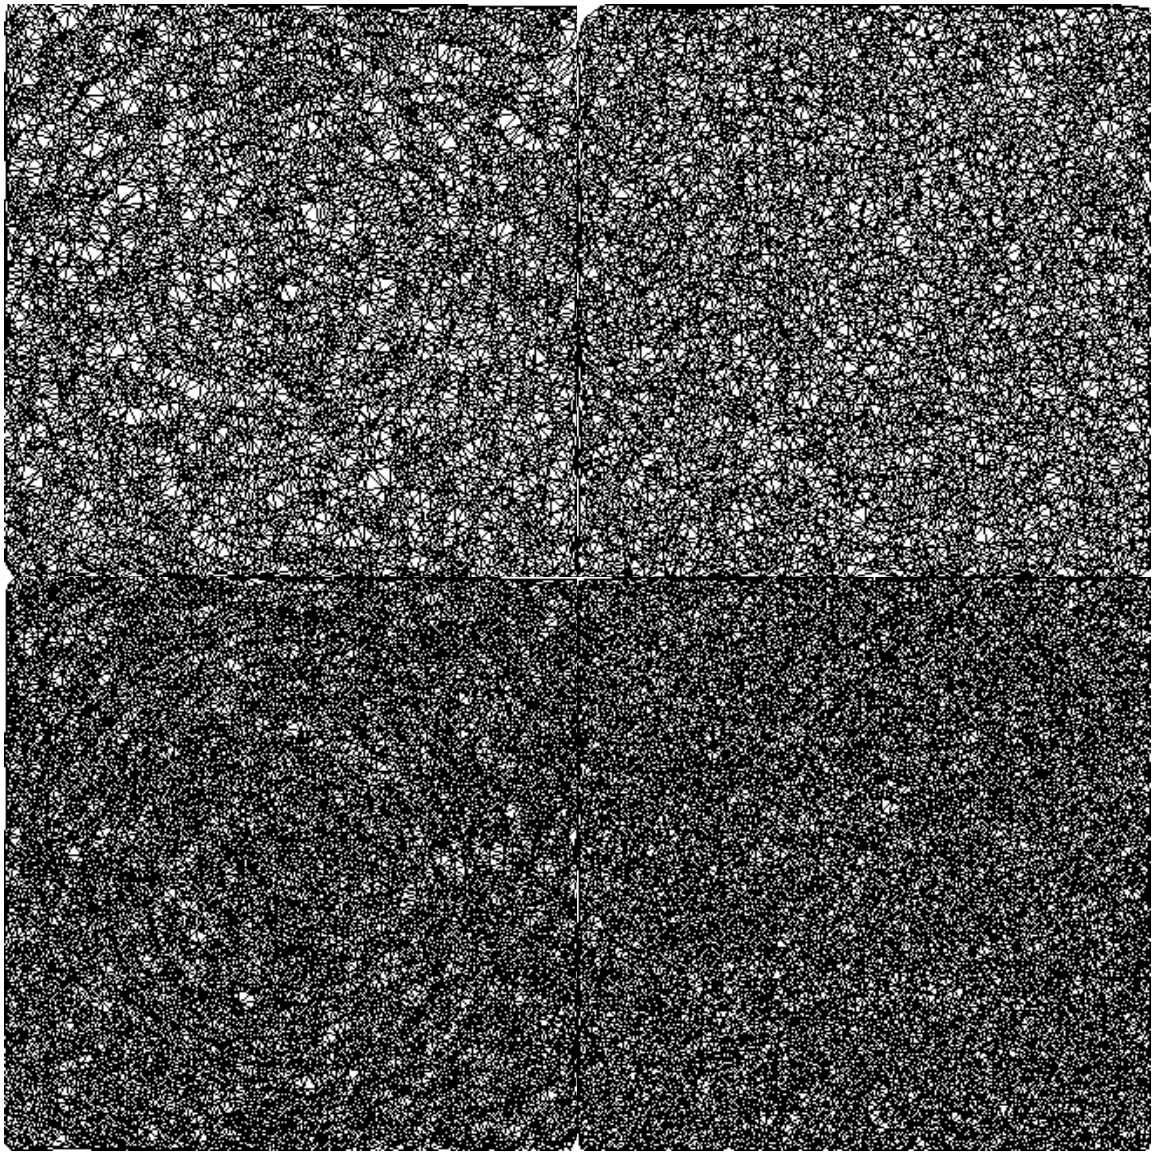

Figure 12.21: Upper Left: Delaunay triangulation of 10,000 Glass rotation dots. **Upper Right:** Delaunay triangulation of 10,000 noise dots. **Lower Left:** Delaunay triangulation of 20,000 Glass rotation dots. **Lower Right:** Delaunay triangulation of 20,000 noise dots. (See text.)

Table 12.2: Means and standard deviations for various measures on Delaunay triangulations for respective Glass and noise patterns. The standard deviation of areas of triangles proves useful.

	10,000 dots		**20,000 dots**	
	Glass	**Noise**	**Glass**	**Noise**
Number of triangles	19,931	19,932	39,928	39,900
Average least angle of triangles	31.17°	30.12°	30.75°	30.51°
Standard deviation	13.16°	13.30°	13.08°	13.09°
Average side length of triangles	57	58	41	41
Standard deviation	34	34	23	22
Average area of triangles	1259	1247	629	625
Standard deviation	**1215**	**1120**	**577**	**549**

The ability of standard deviation—in particular standard deviation squared, or variance—to differentiate areas of triangles via appropriate inferential statistics; the *F* test for equality of variances (σ^2 Glass/ σ^2 noise) for example, holds for Glass patterns of any density that can be reasonably tested.

This concurs with the anisotropic interpretation. For every dot in a Glass *translation* pattern there is at least one other dot that lines up, which, at least, forms a pair with an orientation that is consistent across the pattern. For every dot in a Glass *rotation* pattern there is at least one other dot with the same radial distance from the rotation centre, which, at least, forms a pair with an orientation normal to that of the radial extent, and with an orientation that is consistent across the pattern. Dots are constrained in pairs, with a lane-like tendency, orientated in preferred directions. This also affects the homogeneity of dots normal to transformation direction.

Compared to a random dot pattern with the same number of dots, more dots lie along preferred directions and therefore fewer dots lie across them. Compared to the distribution of Delaunay triangle areas for the random dot pattern, there is a greater frequency of smaller triangle areas for Glass patterns and also a longer tail of larger triangle areas.

Hence, in the anisotropic situation there is a greater spread of triangle areas, which can be overtly appreciated by reference to the upper two panels of Figure 12.21; these being Glass rotation and noise, each involving 10,000 underlying stimulus dots.

The greater spread of triangle areas can be also appreciated by reference to Figure 12.22, which was derived from a circular form rotation Glass pattern and equivalent noise, each involving 10,000 underlying stimulus dots.

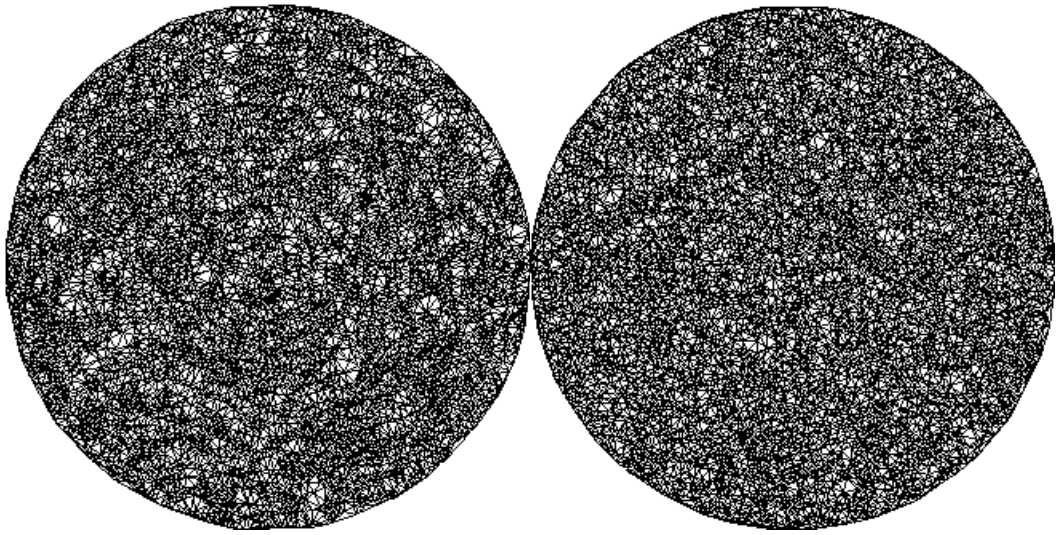

Figure 12.22: Left: Delaunay triangulation of 10,000 Glass rotation dots on a circular form. **Right:** Delaunay triangulation of equivalent noise; i.e., 10,000 noise dots on a circular form. (See text.)

A circular form obviates the overlap problem with square form rotation Glass patterns, which exists at the corners and increasingly along the edges with greater rotation magnitude. Figure 12.23 indicates the general idea. Dot density is half as much at the overlaps owing to transformation counterpart dots in these regions being rotated out of the common area. Just the dots inside a common sub-area should be used, but it is expedient in the present application to use a circular form.

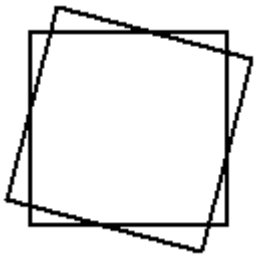

Figure 12.23: Dot density is half as much at the overlaps. Just the dots inside a sub-area should be used.

Figure 12.24 shows *F* ratios for variances of Delaunay triangle areas (circular form) for rotation Glass patterns against equivalent noise ranging from 500 to 50,000 dots at 60 screen units transformation distance, shown in dark blue, and 200 screen units transformation distance, shown in magenta. It also shows corresponding critical values for the *F* tests in cyan and yellow.

The transformation effect is near optimal at 60 screen units transformation separation, for which the plot shows it always significant, but more so at lower densities. For 200 screen units transformation separation, the plot rightly indicates that the transformation effect is not as strong. As close to *F* criticals as such plots might come, there appears to remain some difference out to very high densities. Even so, given the power to detect meaninglessly small differences in large numbers of dots, statistical significance is not the issue. Rather the issue is that the differences are principled. They are all in the right direction, reduce as expected with increasing density and transformation magnitude, and there is no case in which an *F* ratio is less than or equal to 1. And so it appears to be for Glass patterns of any density that can be reasonably tested (of which there is a practical limitation owing to current computing capacity). Moreover, the *t* test on means of triangle areas yields principled differences.

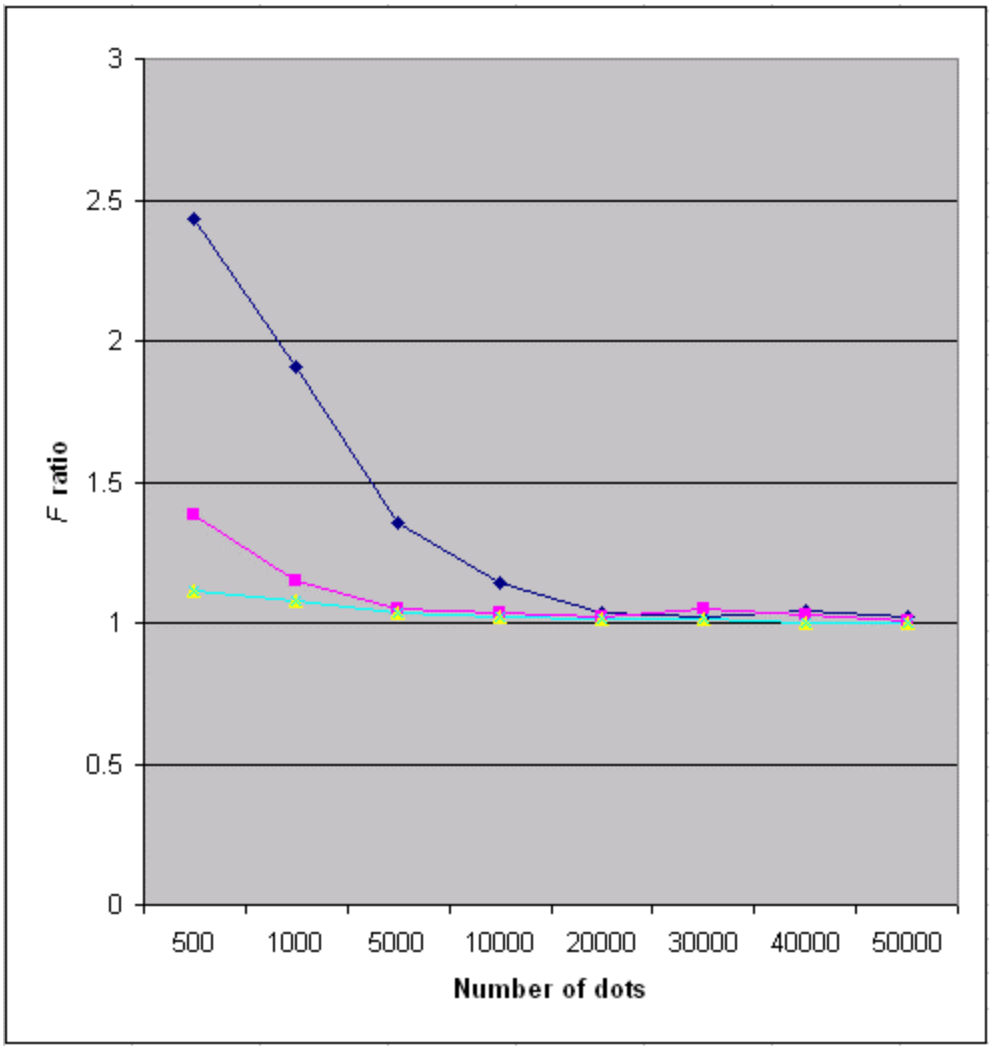

Figure 12.24: *F* ratios for variances of Delaunay triangle areas for rotation Glass patterns (circular form) against equivalent noise, ranging from 500 to 50,000 dots. Transformation distance at 60 screen units is shown in dark blue, and transformation distance at 200 screen units is shown in magenta. Corresponding critical values for the *F* tests are shown in cyan and yellow, which are essentially superposed.

Results from such measures on Glass patterns indicate the transformation effect does not disappear with increasing density. Results also indicate that the transformation effect is less with greater separation between transformation partners.

Upon the evidence it is likely that proximity and orientation measures have some influence at lower densities. There can be quite a difference indicated for such measures between lower density Glass patterns and equivalent noise, which aught to operate in concert with the variance measure for Delaunay triangle areas. Moreover first and second moments for Delaunay triangle edge lengths and least angles, as well as average area for Delaunay triangles, are quite different to those of equivalent noise at lower densities.[5]

For a clearer indication of differences between distributions of Delaunay triangle areas, those for a circular form rotation Glass pattern of 40,000 dots and equivalent noise are shown in Figure 12.25.

[5] Maybe less dense Glass patterns have some extra qualitative difference in transformation effect, with respect to more dense Glass patterns, to match this quantitative difference.

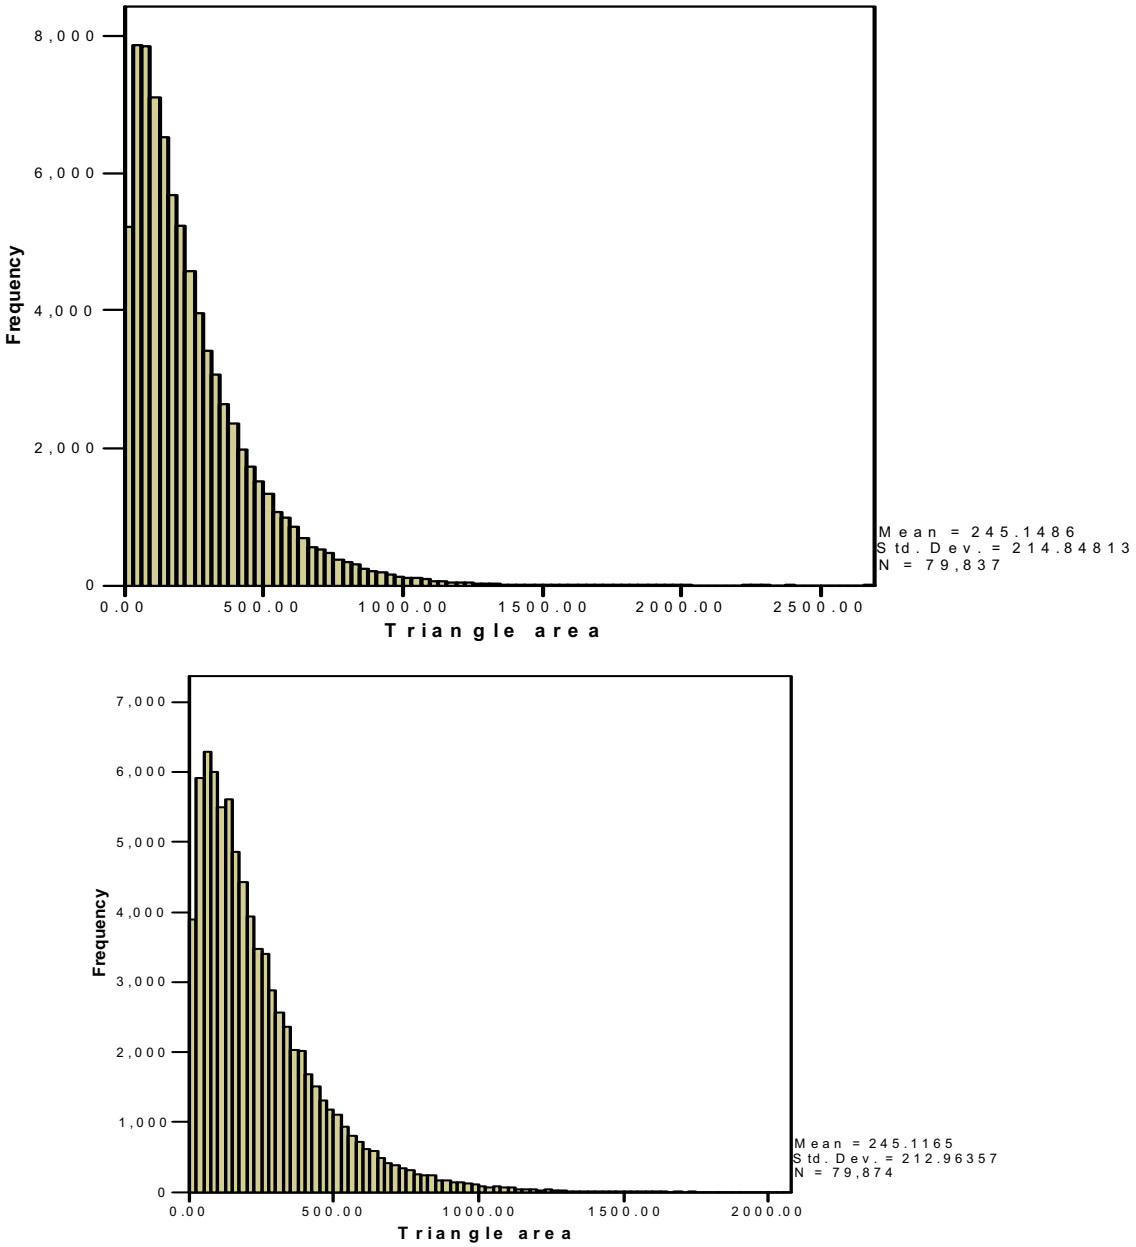

Figure 12.25: Distributions for a circular form rotation Glass pattern of 40,000 dots (upper panel) and equivalent noise (lower panel). The greater frequency of smaller triangle areas for the Glass pattern is obvious. However, the tail of larger triangle areas for the Glass pattern is not obvious owing to the scale of the graph, but the larger scale of the abscissa connotes these areas; albeit at low frequency.

As far as density is concerned, limits appear to be practical ones for which there is an obliterating fusion of dots at some density related to dot size. The transformation effect that characterizes Glass patterns endures at extremely large densities. With the variance measure on Delaunay triangle areas, there is evidently saturation of dot colour for even tiny dots before a Glass pattern measures like noise of the same density.

Upon magnifying rotation or dilation Glass patterns containing up to 10 million or so speck-like dots, by differentially scaling selected areas with respect to dot size, the transformation effect can be still patently evident. Before such scaling, Glass patterns, typical in size of those shown in figures throughout, are characteristically saturated with dot colour.

In the differential scaling employed herein, a form is scaled up more than its dots, and such that the dots tend strongly to individuation in the event of an initially saturated situation.

Where a Glass pattern is initially discernable, then uniform scaling (i.e., form and dots scaled up equally) preserves the effect. To discern a strong transformation effect after differential scaling at increasing magnifications, a Glass pattern requires relatively smaller transformations. Upon differential scaling, a pattern with a larger transformation can be lost though the pattern is obvious beforehand. However, this requires some qualification. If, after such scaling, enough of the pattern is included then it can still be seen; albeit with some reduction of effect. Notwithstanding the global contribution of larger pattern segments, all this suggests an interaction of transformation distance and dot interposition. Moreover the transformation effect is not nearly as persistent for translation Glass patterns as it is for Glass patterns in which transformation is about a proximal central pivot; which seems ecologically plausible.

Figure 12.26 shows a small segment with a few thousand dots belonging to a constant displacement rotation Glass pattern of 5 million dots in all. The segment is rescaled by a form factor of 388. The *F* test for variance of triangle areas with respect to equivalent noise for segments proves principled in the way that the *F* tests were for the previous examples. And if a reasonable part of any constant displacement Glass pattern shows such comport against equivalent noise, then the whole must do so as well.[6]

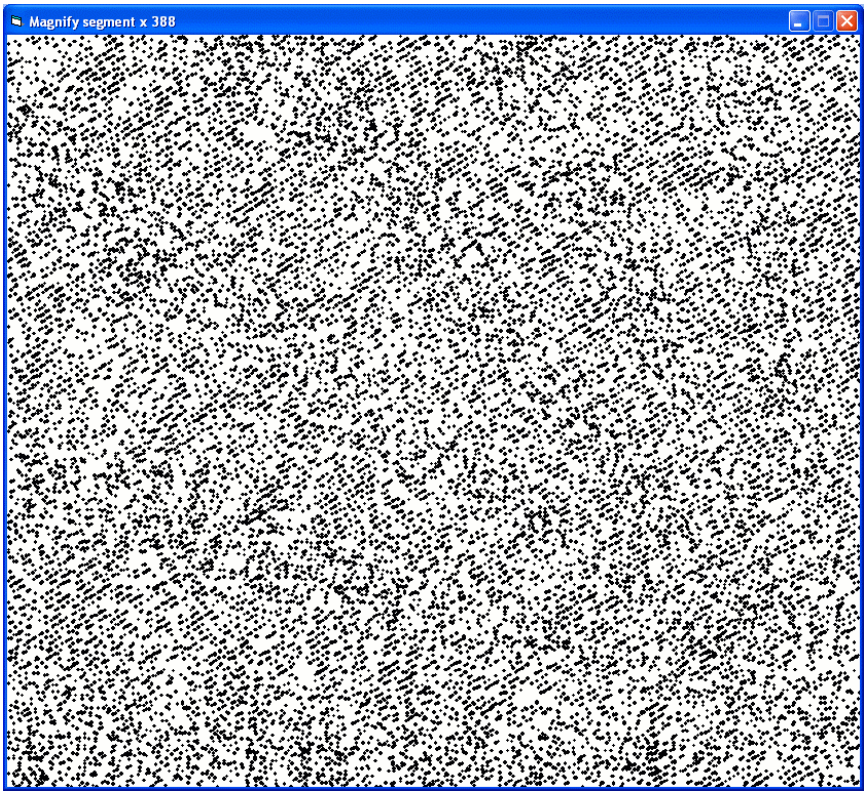

Figure 12.26: Small segment of a constant displacement rotation Glass pattern, rescaled by a form factor of 388. The whole Glass pattern of 5 million dots was on a form not much smaller than that derived for the magnified segment; the latter of which was maximally sized for screen fit while maintaining aspect ratio.

By iterating through arrays preloaded from Glass pattern picture files, millions of dots can be shown at once for each presentation. The method allows for rapid iteration of displays

[6] A result for the whole remains beyond accessible computing resources.

at any practical number of dots; as opposed to generating each display immediately prior to presentation. The latter procedure becomes slow and ineffective for successive displays of large numbers of dots

As an aid to description, Figure 12.27(a) shows a relatively sparsely populated example of a progressive displacement rotation Glass pattern; as opposed to an equivalent constant displacement rotation Glass pattern shown in Figure 12.27(b). Taking a progressive displacement rotation Glass pattern as an exemplar, rotation motion is obvious at hundreds of thousands to millions of dots. At the same time, dot density can be linearly increased throughout a presentation sequence, upon which the rotation effect becomes progressively more localized around the centre of the display. The same 'spin in' effect can be seen by *generating* one such Glass pattern, with each dot pair output to the display in generation sequence, up to millions of pairs.

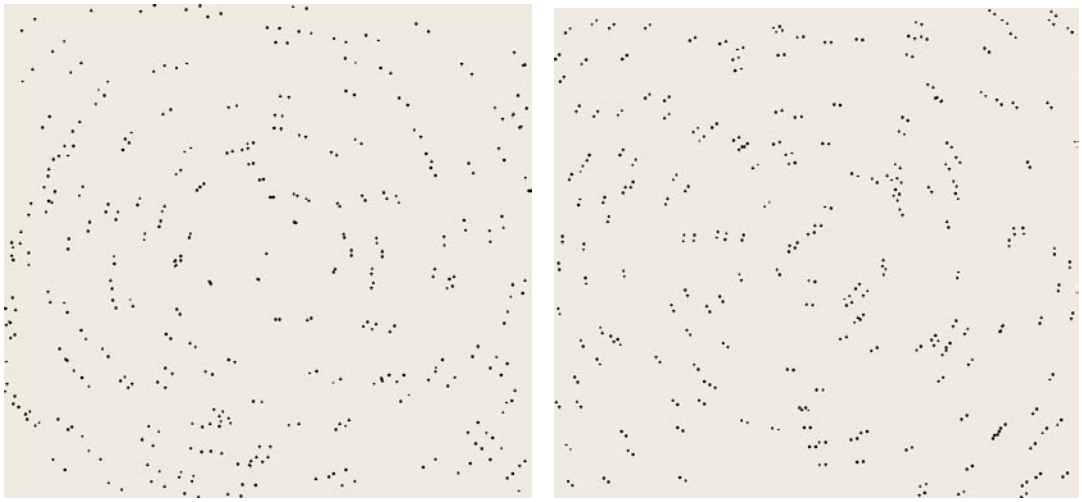

a **b**

Figure 12.27: Progressive displacement rotation Glass pattern (a). Constant displacement rotation Glass pattern (b).

The Glass pattern exemplar has a progressive increase in distance between transformation partners, related to radial distance from transformation centre. In other words, more outwardly located pairs are constituted of larger partner displacements than more inwardly located pairs. With increasing density, the transformation effect is lost *not* because of increasing interposition, which is uniform across any Glass pattern, but because of interaction of increasing interposition with progressively greater transformation distances.

The relation between density and transformation effect is further clarified by iterating the same Glass pattern, or at least the same number of Glass pairs, such that it spans increasing size forms. Initial density is chosen so that the transformation effect is visible just to the middle part of the edges, say, of the small, initial form. It turns out that the same effect persists across the different size forms: the same visibility of effect to the edges of all the forms remains with decreasing density.

Starting with the same Glass pattern, or at least the same number of Glass pairs, on the same small, initial form as before, the number of Glass pairs is now increased proportionally to increasing form area so that density remains constant. And there is no doubt that the transformation effect moves inwards, in a principled manner, as form size is increased. Transformation effect appears to be exactly compensated by reducing density in the first instance, and reduced by increasing transformation distances in the second instance.

Additionally, MacKay patterns with hundreds of thousands of dots can be presented by iterating through arrays preloaded from random dot picture files with barred backgrounds, and the MacKay effect is still patently obvious. In all these kinds of observation there is no movement other than that generated by the visual system. This suggests that an active or *generative transformational* system is at work in the visual system. It appears to be the common element across many apparently different phenomena related to visual perception. Distance between elements is relevant, and structure is certainly elucidated by the Delaney neighbourhood.

A summary hypothesis

Barlow (1999) noted that many researchers with an interest in cognitive function recognised the importance of statistical regularities of the environment. He and others have begun to show how regular statistical properties of images are exploited by neural mechanisms. Chun (2000) showed that statistical regularities in the structure of images were picked up unconsciously by observers. Inasmuch as this thesis is concerned, statistical regularities are inducing consistencies.

Inducing consistencies may produce likely transformations. Voronoi tessellation and the Delaunay neighbour hierarchy capture structure in patterns. However, investigation has shown that the transformation effect, both within Glass patterns and between Glass patterns treated in sequence, persists beyond exhaustion of Delaunay neighbour distances equal to transformation distances. Hence a claim that neighbours of any kind in the Delaunay hierarchy need to link transformation partners to sustain the transformation effect is incorrect. Nonetheless they may be implicated at low to moderate densities in Glass patterns and Mackay-like patterns. An important point is that even when Delaunay neighbour distances equal to transformation distances are exhausted, members of the Delaunay neighbour hierarchy capture structure in Glass patterns.

For unstructured displays—static random dot displays and random dot kinematograms, for example—sporadic, or chance, inducing consistencies may produce likely transformations. Nearest neighbours and-or other Delaunay neighbours appear to pay a key role in the perception of such displays (see Appendix B). Moreover nearest neighbours and-or other Delaunay neighbours prove important for the perception of clustered and regular patterns (see Appendix A), and for the perception of Marroquin patterns as another form of regular pattern. Figure 12.28 (a) shows a piece of Marroquin pattern formed with a rotation of 12.5^{o} and Figure 12.28 (b) shows the pattern outlined by nearest neighbours, for example.

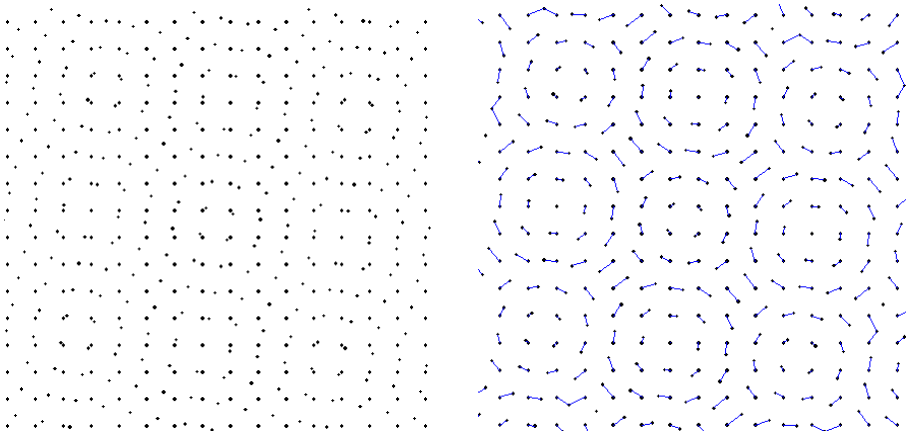

a **b**

Figure 12.28: Structure in the piece of Marroquin pattern (a) is captured by nearest neighbours, shown linked in (b).

Gestalt principles as inducing consistencies

Polygons can be generated using translations and circles can be generated using rotations. Geometric patterns such as these consist of special arrangements of Glass pattern elements. Hence the Glass pattern detection routine can detect these on the same basis as it detects Glass patterns. In fact the Glass pattern routine can serve as a basis for detection of a variety of low-level structures.

Voronoi or Delaunay based connectivity preserves spatial arrangements of dot distributions. With increasing density of Glass patterns, notional neighbour links form notional elements. At low to intermediate densities, these provide connections between transformation partners, as well as connections between partners of different transformation pairs, that outline the transformation scheme at work. This then may trigger a salient transformation in the visual system. At high densities, elements formed by notional neighbour links predominantly provide connections between partners of different transformation pairs that outline the transformation scheme at work. Again, this may trigger a salient transformation in the visual system.

Voronoi, Delaunay based connectivity is a kind of consistency that relates to just one Gestalt principle, namely proximity. The other Gestalt principles, similarity, continuity, and closure, along with proximity, constitute a group in which the elements can interact. All are kinds of consistency, or statistical regularity, that may trigger salient transformations in the visual system.

Links substitute for dots

For a relatively sparse Glass pattern, most nearest neighbours link transformation partners and few nearest neighbours link partners of different transformation pairs. For an intermediate density Glass pattern some nearest neighbours link transformation partners, but these are intermingled with other nearest neighbours that link partners of different transformation pairs. And for a high density Glass pattern there are few nearest neighbours that link transformation partners. Nearest neighbours predominantly link partners of different transformation pairs. Figure 12.29 shows this transition in panels from top left to bottom right.

Whether or not separate nearest neighbour links substitute for two underlying dots or connected nearest neighbour links substitute for three, or more, underlying dots, and whether or not they link transformation partners appears to be irrelevant. What is relevant is that the links substitute for underlying dots and we see the Glass effect in the links, which have a lower density than the dots in the range half density upwards. And if these links, as elements, are linked in similar fashion again, between their centroids for example, and the original links erased, then a decreasing number of larger elements continues to outline structure. The important consideration is that *proximal distances matter.* They outline structure. Moreover they outline transformation schemes, whether or not their corresponding links join transformation partners.

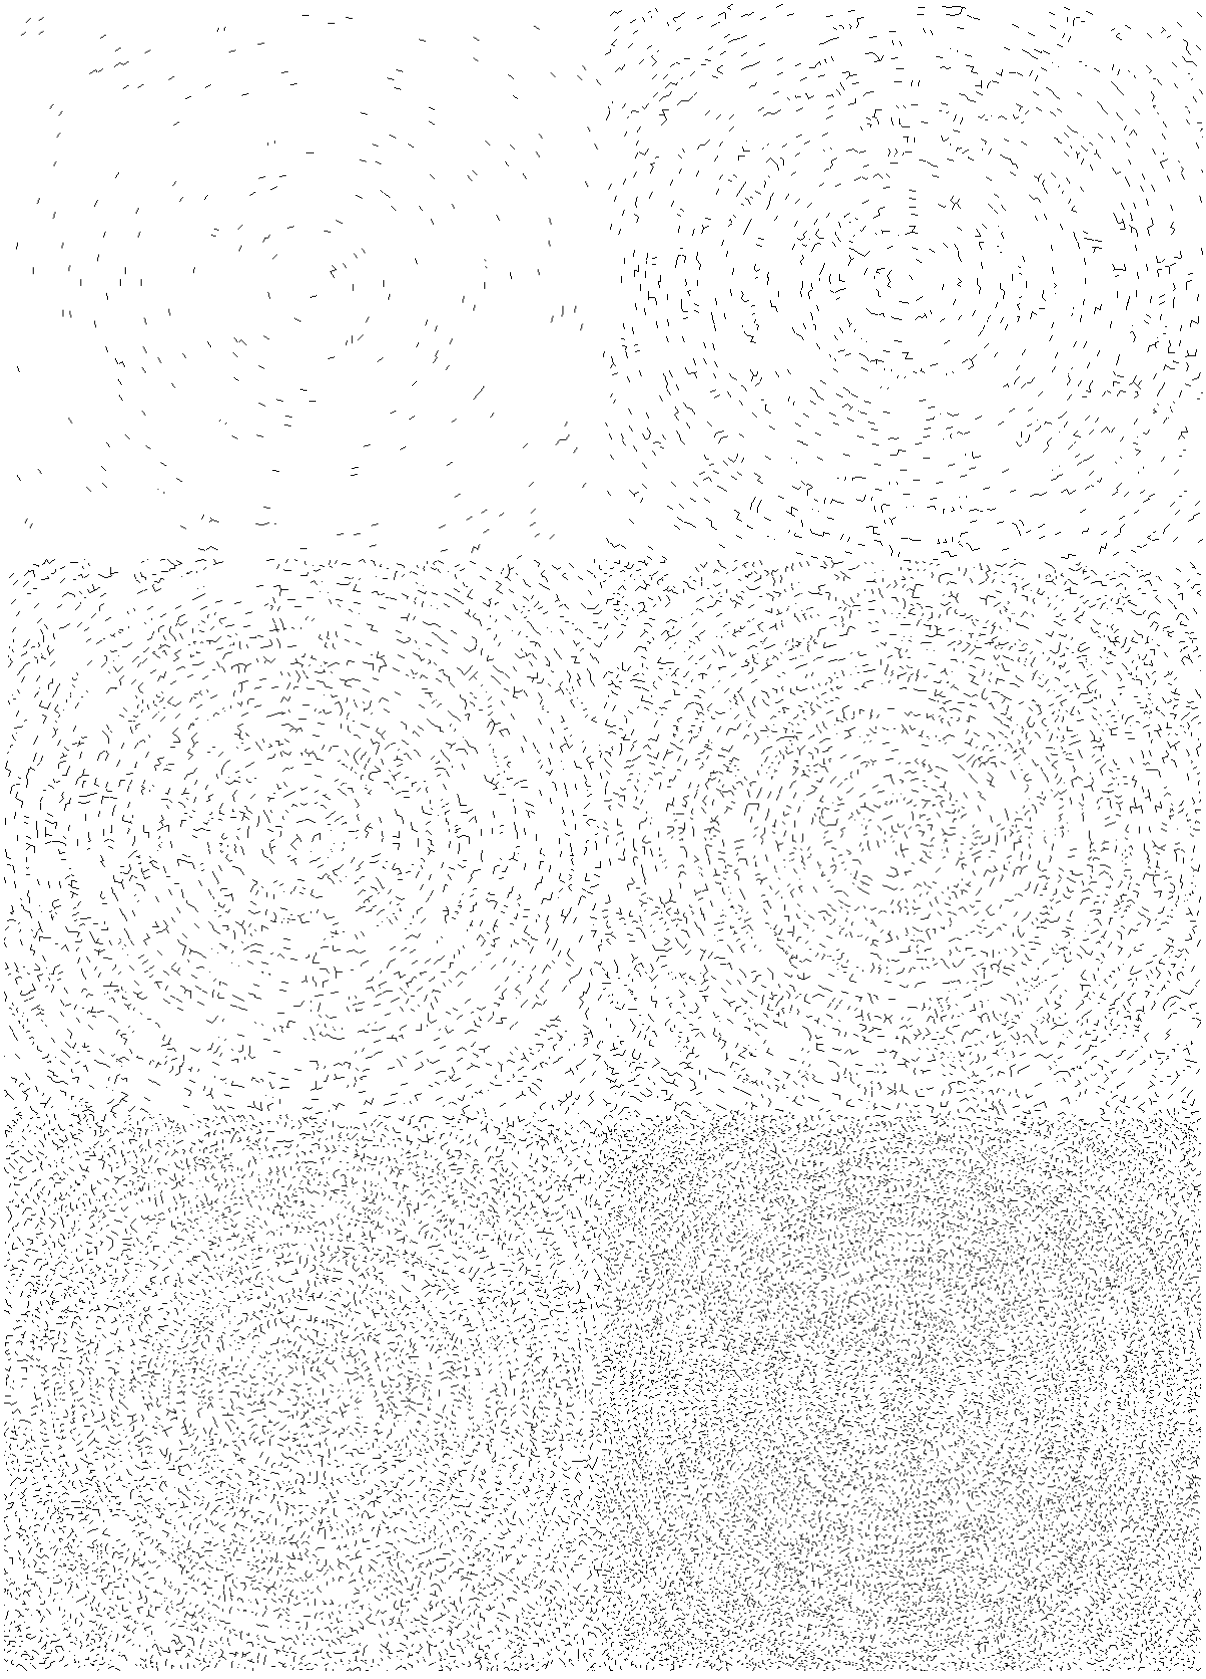

Figure 12.29: Nearest neighbour links for underlying Glass patterns of increasing density. At low density they link transformation partners. At intermediate density they link transformation partners, but these are intermingled with other nearest neighbours that link partners of different transformation pairs. At high density they link partners of different transformation pairs.

Neighbours and spatial frequency

A commitment of this thesis has been to thoroughly explore nearest neighbour and its encompassing family, but these are obviously not the only terms in which explanations can or should be couched. Spatial frequency theory provides an appealing account involving proximity, and is here briefly mentioned.

It is generally known that Neurophysiological studies (Frisby, 1980; Hubel, 1979a; Hubel & Wiesel, 1959, 1968; Mountcastle, 1978) have revealed cortical columns of specialized feature detector cells that are aggregated in hypercolumns with varying degrees of a feature along one dimension and left-eye, right-eye alternating columns for the respective degrees along the complementary dimension.

In spatial frequency theory, the phase and amplitude of a spatial frequency wave are respectively commensurate with the location and degree of brightness of elements belonging to an image. Higher spatial frequencies carry finer detail formed by local contrasts and lower spatial frequencies carry coarser detail formed by the overall pattern of shading. Complex spatial frequency waves can be constructed from the summation of pure sinusoidal waves having a variety of phases and amplitudes. These typically correspond to real scenes, in contrast to the pure sinusoidal waves that correspond to a class of contrived displays used in psychophysical studies.

Much evidence exists to show that we have specialized cells, called spatial frequency filters, in the retina and further levels of the visual system, and that these respond to specific spatial frequencies (De Valois & De Valois, 1980). According to spatial frequency filter theory, the specialized feature detector cells revealed in neurophysiological studies detect information about spatial frequencies. Many bits of spatial frequency information from small areas of the retina are integrated further along in the visual system.

For MacKay patterns, there is an impression of dots streaming faster in directions normal to barred backgrounds as the barred backgrounds are made more dense, and this equates simply to increasing the spatial frequency of the backgrounds as inducing elements. Anisotropy in Glass patterns equates to a kind of textural graininess, which causes Glass patterns to have higher spatial frequencies along transformation orientations than across them. For both MacKay patterns and Glass patterns, impression of orientation is the same as the orientation of highest spatial frequency. That is, impression of orientation is the same as the orientation of *closest proximity* of inducing elements.

One of the earliest cortical cell features discovered was sensitivity to orientation, with progressive changes in orientation across the columns of cells (Hubel & Wiesel, 1959, 1968). Further studies have shown that cortical cells tuned to different spatial frequencies appear to be ordered along the dimension complementary to that of orientation selectivity within each hypercolumn (De Valois & De Valois, 1988).

Scale integration in spatial frequency theory has presented a challenge. However, scale integration that uses a broad and continuous range of scales has been undertaken somewhat successfully by Witkin (1983). Notwithstanding that there may be constraints on distances associated with orientations and that these would need to scale, then, on the above, we *should* see some kind of Glass effect. Moreover orthogonality may bracket the argument. Effects from cells with the most excitation may be compared with effects from cells with the least excitation; which highlights the difference along a grain as opposed to across the grain. At the same time, effects from cells excited a little less than the most excited may be compared with effects from cells excited a little more than the least excited, and so on in trigonometric complementarity.

However, for all but pure translation Glass patterns orientation is not constant, hence there remains a problem for the global percept of Glass patterns. It is the same as that outlined in Chapter 8 for the difficulty with the neural approach in terms of hierarchies of filtering

units. How are the many bits of spatial frequency information at different orientations, each from small areas of the retina, integrated further along in the visual system? And, again, it is hypothesized that such information triggers an appropriate transformation. More generally, however, some or all of the Gestalt principles in interactive combination may trigger an appropriate transformation.

Chapter 13: Discussion and Further Research

General

The generative transformational approach stems from a general hypothesis in response to theoretical problems gleaned from the literature. Investigation into the approach indicates that it can provide a rudimentary, workable model of certain elements of human visual perception (Vickers, Navarro, & Lee, 2000; Vickers & Preiss, 2000). This is further supported by the investigation undertaken herein.

The generative transformation model provides a reasonable account of perception of static and changing visual stimuli. It recognizes transformed stimuli, and can be utilized to draw conclusions about a number of illusions. The model is parsimonious and illuminating, both qualitatively and quantitatively. The discipline associated with the need to specify every process in detail enforced an unavoidable, if not laborious, advantage. Many dimly anticipated, and some unforeseen, elements were pointed up by the operational demands of the model under development. (Needless to say, there is a good deal of difference between a hypothesis and actually being called upon to implement it in any reasonable detail.)

The generative transformation model uses information about relative positions of image elements. In effect each image element is 'seen' by all the others, and relationships of economy that prove most suitable are utilized to select transformations that minimize dissimilarity between arrays. The trajectories of the selected transformations then correspond to perceived structure.

The model decomposes and regenerates fractal objects. It detects and differentiates non-fractal objects, including Glass and Marroquin patterns, and regular and semi-regular geometric objects; any embedded in noise as appropriate. In the same fashion, it detects regular and semi-regular motion. It also provides an explanation for temporal and spatial context effects, such as the 'representational momentum' that allows us to infer the future state of an object (Freyd & Finke, 1984).

The generative transformational approach accords with Leyton (1992), who argues that visual perception consists of recovering the process history undergone by an object. According to Leyton, such recovery proceeds by progressively removing asymmetries or 'distinguishabilities', so as to infer an original object that is maximally symmetric.
All this implies that perceptual information may be remembered in a more active form than is generally reckoned; a form that is connected to the way in which such information can be regenerated (Vickers & Lee, 1997). Moreover the approach includes some interpretation of the development or construction of objects, hence it shifts the orthodox boundary between perception and cognition.

Although detection of structure, as examined in Section 2, rests heavily on the Hausdorff measuring procedure, it should be emphasized that it is not helpful to regard the model as a static, reflex-like process. It is generative in the sense that the response of the model is capable of mimicking the process whereby an image was generated in the first place. This has been shown by the way it can regenerate fractal objects from parameters gathered during decomposition. (Decomposition is akin to seeing, and regeneration is akin to realization.) More generally, because relationships of economy are utilized to select transformations that minimize dissimilarity between arrays, the model is implicitly generative. Whatever is seen to be done in the first place, is 'undone' by the model in the deciphering process.

The generative transformational approach invites many predictions and provides explanations for a multitude of phenomena, just a few of which are mentioned below.

Illusions concerning visual aftereffects

Illusions concerning visual aftereffects are accommodated by the generative transformational approach. MacKay figures, explained opportunely in terms of nearest neighbours at the end of Chapter 2, pages 33 to 35, are a precedent. The inducing, or background, figure has some form of symmetry. A radial background has bidirectional rotation symmetry. A concentric background has dilation symmetry (more on this presently). A barred background has translation symmetry in two directions orthogonal to the bars, and so on.

The generative transformational approach cites these figures in terms of process history, or transformational actions on simplest elements needed to produce them (Leyton, 1992, again; see also Vickers & Preiss, 2000, Appendix C). In the case of the radial MacKay figure for example, a bar would be planar rotated, using one end as a pivot, and stamped in increments to give the barred radii. Another precedent related to MacKay's work (MacKay, 1957a for example) is the waterfall illusion, in which fixating for a while on downward streaming dots, as the inducing figure, results in apparent upward motion of a test figure of stationary dots. Here, the symmetry involves invariance under translation. In terms of process histories, the transformations that reinstate original configurations as invariants are experienced as tendencies by way of aftereffects. In either the MacKay figure or the waterfall illusion, a subsequently viewed bland field of stationary dots appears to move with a tendency to undo the inducing transformation. Note that in some cases inducing transformations result from one of two interpretations related to bi-directionality, as per Ross, Badcock, and Hayes (2000), cited in Chapter 11. Whatever the interpretation, the aftereffect is a complementary effect.

In the case of fixating an oblique line for a while, and then switching to a vertical line, the vertical line appears a little inclined in the opposite sense. Here, again, there is the endeavour by the visual system to undo a transformation, which is triggered in this instance by the mapping from test to inducing line.[1] Like the representational momentum effect described in Chapter 9, pages 187 and 188, compensation occurs in response to a most parsimonious rendering of transformational history.

Directionality of aftereffect is a topic that needs more explanation. Again, MacKay figures serve the purpose. Notwithstanding the predilection of a particular visual system, a MacKay figure with a radial background promotes equal probability for aftereffect streaming in either rotational direction; and similarly with a barred background for either direction orthogonal to the bars. However, for a concentric background the aftereffect is comprehensively experienced as streaming from the centre, outwards; and this could be because people do not have eyes in the backs of their heads!

Concentric circles prime for dilation with greater probability than for contraction because the visual system has internalized the former with greater precedence than the latter. This is because the visual system experiences dilations during forward components of movement, which are ubiquitous in daily experience. The less likely experience of looking to the rear while moving forward, or, indeed, moving backward, produces contractions at the retina. An example that highlights could be the rear view experienced while looking from the rear of a fast moving train. The visual system manages, but the experience seems somewhat atypical. Nonetheless anything that recedes concurs with a contraction, and the visual system accommodates it.

The aftereffect from fixating stationary patterns of alternating black and white stripes for a while and then viewing a uniform field of dots, which appear to show a streaming

---

[1] Of course, a vertical line is chosen as the canonical form because of the sensitivity of the visual system to departure from verticalness. In any event, such a small offset would not be noticed with another oblique line as the test figure.

motion normal to the orientation of the stripes, is the same as the simultaneous effect for *dynamic* MacKay configurations. Careful observation of the dynamic MacKay configurations, described at the end of Chapter 2, show that it is just as easy to urge either of the directions for streaming where a radial or a barred background is concerned, but it is more difficult to urge the inward direction for streaming against the predilection for the outward direction where a concentric background is concerned. Nevertheless it can be forced to some extent.

In keeping with Shepard (1984, 2001), all this argues well for the notion of internalization. In urging one direction or another, one literally wills the direction of effect with the same mental transformation as the apparent streaming transformation that ultimately results. And the biased bi-directionality for the concentric background should be, again, a confirming analogue for internalization.

Müller-Lyer illusion

Müller-Lyer arrowheads, however acute or obtuse, long or short, are simply transformations of one another. It is obvious that the centroids (balance points) of inward pointing arrowheads are further apart than those of outward pointing arrowheads, where each pair cap the same length shafts. See Figure 13.1. In addition to the centroid located at the centre of each shaft, each half of a Müller-Lyer figure has its own centroid constituted of half a shaft and an arrowhead. However, it does not appear to be this sub-centroid that dominates, only the sub-centroid of the arrowhead itself.

A computational analysis indicates that twice the distance from the tip of an arrowhead to its centroid best represents the subjective difference in length between the two shafts. Equivalently, a shaft appears shortened (contracted) or lengthened (dilated) by the distance from the tip of an arrowhead to its centroid. See Figure 13.2. Another equivalent standpoint is that the centroids due to the arrowheads at both ends of a shaft have the effect of shrinking (contracting) or stretching (dilating) the ends halfway toward them. This prediction of the generative transformational approach needs investigation by way of a research study involving subjects, not just a computational analysis.

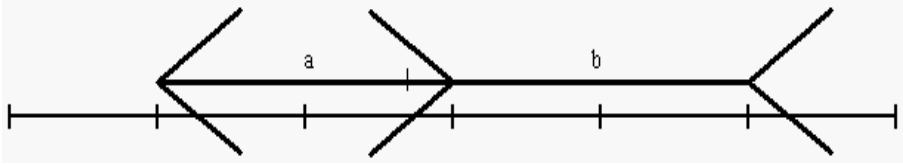

Figure 13.1: Müller-Lyer illusion. Shafts **a** and **b** are the same length. The intersection of the hairline with a shaft indicates the centroid of the corresponding arrowhead.

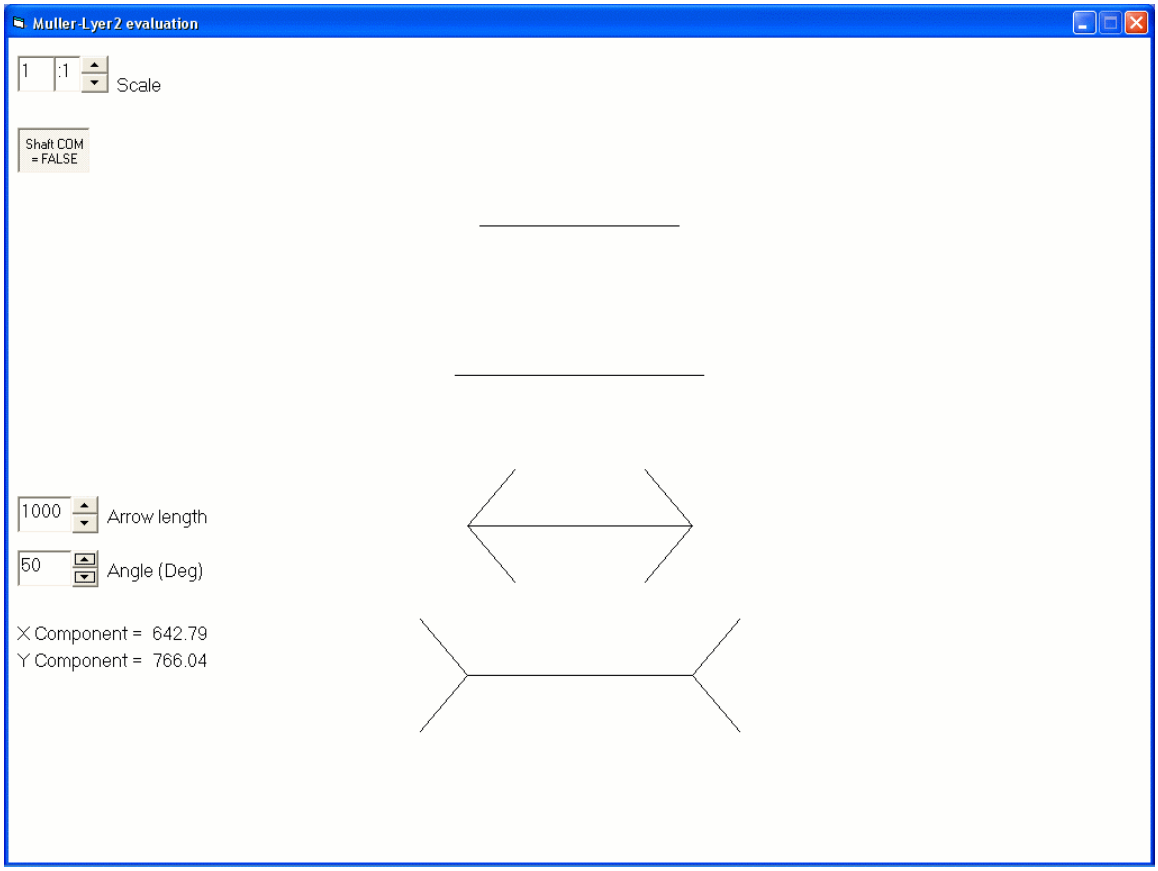

Figure 13.2: Analysis indicates that twice the distance from the tip of an arrowhead to its centroid best represents the subjective difference in length between the two shafts. The top two lines, originally the length of the shafts, are altered in length by the signed distance from the tip of an arrowhead to its centroid. The computer program allows variable length arrowheads, and can spin the angles through acute and obtuse continuously. Any combination appears to ring true.

A related prediction of the generative transformational approach concerns displacements of shafts. In this scenario, one shaft is capped with an inward pointing arrowhead at one end and an outward pointing arrowhead at the other end, and the other shaft is rendered in the opposite sense (an outward pointing arrowhead at the same end as the other shaft has an inward pointing arrowhead and vice versa). As the arrowheads are continuously transformed through acute and obtuse angles, the generative transformational approach predicts that the shafts will appear to translate horizontally in opposite directions while maintaining fixed equal lengths, and this is exactly what happens. The apparent displacement of the shafts with respect to each other is, again, twice the distance from the tip of an arrowhead to its centroid. Equivalently, a shaft appears to translate horizontally about a mean position by the distance from the tip of an arrowhead to its centroid. Whether dealing with opposed or coincident arrowheads, shaft transformations appear seamless with continuous transformation of the arrowheads, and this seamless effect faithfully follows the way the centroids translate.

Distortion illusions

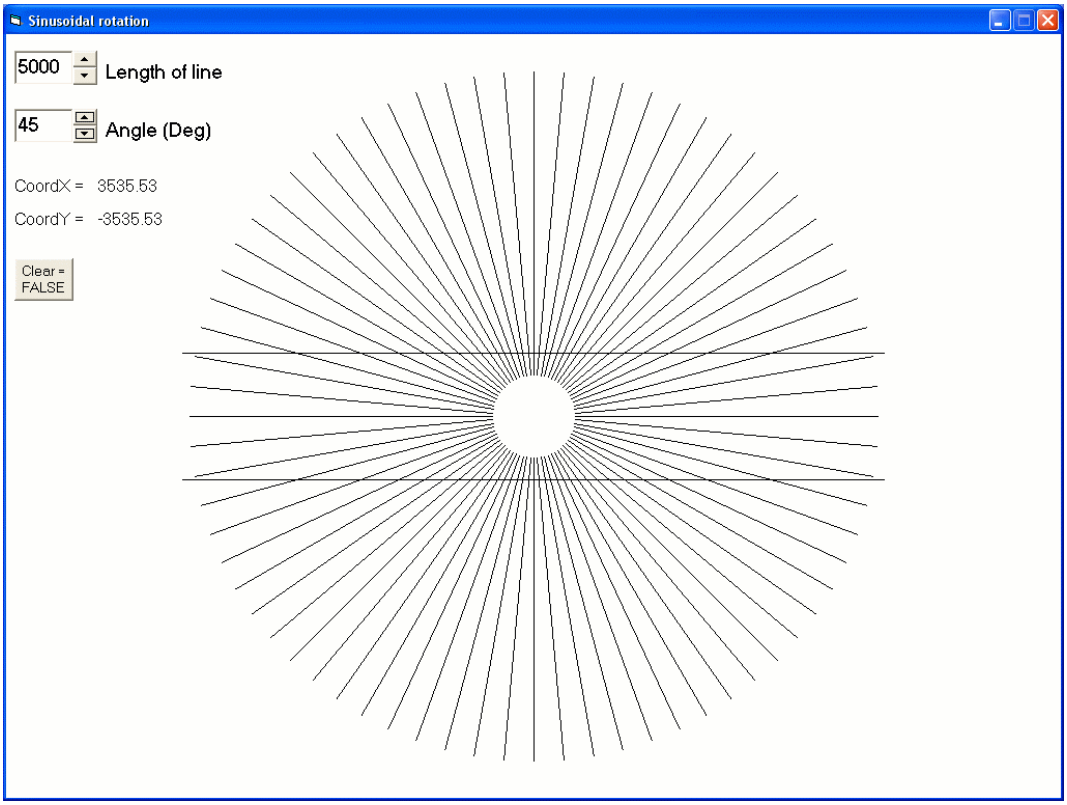

Figure 13.3: Straight lines appear bowed away from the common center of the radial lines.

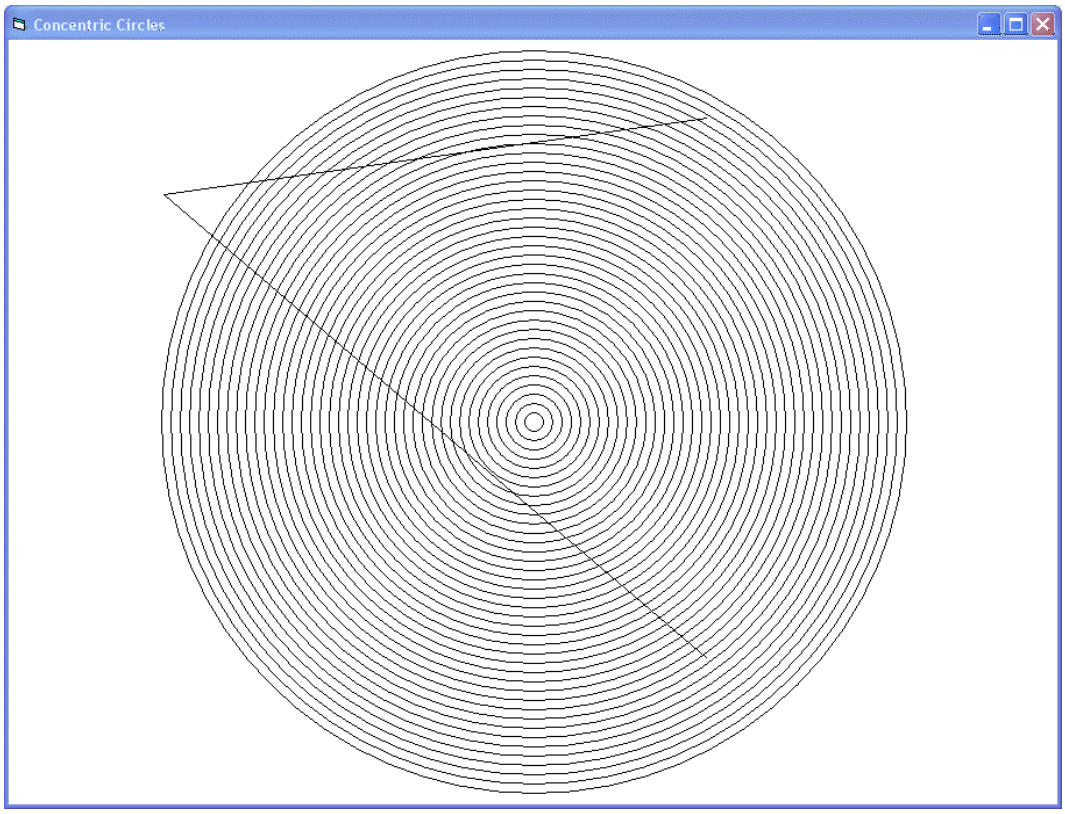

Figure 13.4: Straight lines appear bowed at the tangents, towards the common centre of the circles.

The line distortion illusions shown in Figures 13.3 and 13.4, again, can be explained in terms of bidirectional rotation symmetry for the radial background and dilation symmetry for the concentric background. For Figure 13.3, rotation symmetry in either direction concurs with apparent curving of the lines. For Figure 13.4, more dilation towards the ends of the upper line concurs with them being pushed away in apparent fashion from the centre of the circles. Dilation operates more *along* the lower line, and has more of an apparent stretching effect (which cannot be realized in this figure alone) along with some apparent bowing effect.

Categorization

Feldman (1997) has recently investigated what he terms 'one-shot categorization' (p. 145): the telling ability to infer that a single object, or a small number of objects, belongs to a general class with certain properties. Feldman (1997) proposes that this kind of generalization is equivalent to a generative model for a category: 'a process…that generates legal members by applying a limited set of transformations to some primitive object' (p. 146).

Feldman and Richards (1998) explored the 'natural' transformations undergone by a rectangle and conclude that subjects tend to choose transformations that preserve shape, rather than maintaining constant area, length, or width. This is consistent with the way greyscale charts indicate a shape invariance—via the arrangement of Hausdorff components—of transforming objects (see Chapter 8, pages 174 to 176).

Recognition of different instances of the same letter is a categorization example that could be included under the umbrella of the generative transformational approach. More generally, logo recognition is a categorization example that could be included under the umbrella of the generative transformational approach. Bitmap images, represented by relatively sparse arrays (see Chapter 9, pages 195 and 196), indicate a shape invariance—via the arrangement of Hausdorff components—that could be used to advantage.

Transformed images presented to the retina

Stratton (1896, 1897a, 1897b) and a raft of subsequent researchers, among them Kohler (1962, 1964), report a *plastic* correspondence between the image that impinges upon the retina and what is perceived. Such images, transformed in different ways by various optical devices worn at length by experimenters, are seen as transformed only until adaptation ultimately occurs; usually over some period. A transformation could be a displacement or some rotation, including an inversion, for example. Adaptation occurs via integration of the senses during a rediscovery of veridicality, which then dictates that there is nothing odd about the new world. However, invoking the pre-adaptation mode can interrupt the new experience, such that the world viewed through the device appears odd again. Additionally, adaptation to the new experience causes the world to be seen as transformed in the opposite sense for some period when the device is removed. In any event all adaptations undergone by an individual, including the so-called 'normal' one established during infancy, are available in complementary fashion: only one at a time can be realized. (Invoking an adaptation mode is evidently a bit like urging the alternative perception when looking at a Necker cube, cited in Palmer, 1999, for example, with its two mutually exclusive depth interpretations.)

Stratton, commenting on his 1896 retinal image inversion studies, describes how 'up' is nothing in the visual sensory pattern other than opposite to 'down', and orientation is achieved by the relation of the visual pattern to somothesis and behaviour. He states, 'Any visual field in which relations of seen parts to one another would always correspond to the relations found by touch and muscular movement would give us "upright" vision, whether the optic image lay upright, inverted, or at any intermediate angle whatever on the retina'

(Stratton, 1896, p. 617). The conclusion is that the raft of findings from these studies fit nicely with, and could even compel, a *generative transformational* approach to visual perception.

Perceptual constancy

Adaptation to optically transformed images, including up-down reversal, right-left reversal, deflection with distortion, image discontinuity by shearing at midline, bisection with complementary colours, and the provision of an artificial scotoma, indicates a facility related to wide ranging perceptual constancy. Such adaptation equates with reversion to some perceptual norm, or *correctness,* which services the imperative of successful functioning in the world. J. J. Gibson, interpreting Kohler's findings in the introduction to Kohler (1964), explains that this is achieved by *exploratory,* not *performatory,* activity relating to eye, head, and body orientating systems as a successively inclusive hierarchy. The aim of the exploratory activity is to achieve an optimum of stimulation.

The acquired uniformity of perception, across images transformed in quite different ways, includes size and shape constancy as just two among some number of perceptual constancies. Further, the status of an image need only be such that lawful transformations can validate experience in the world. Any biasing of the way that light delivers information from the environment is not unsatisfactory provided information is not destroyed. Again, this fits nicely with, and could even compel, a *generative transformational* approach to visual perception. In fact, a generative transformational approach accommodates enough that it is difficult not to conclude a substantial part of brain function is given over to transformation operations on representations. Generation of transformations on representations allows the representations to be plastic in any lawful way.

Fractals

By utilizing appropriate transformations, along with minimization of Hausdorff distance mismatch, it is anticipated that any fractal can be decomposed and the representation thereof stored in just a few parameters. This includes modeling of growth and decay, for which a change in size of contracted copies along with one or another convergence factor is needed. An ongoing investigation would develop and optimize methodology to cater for any line, area, or volume fractal with a fixed or variable contraction size.

Model refinement

An immediate need for further research lies in refining the model through comparing predictions with data, especially in the fundamental areas of perception of Glass pattern structure and detection of coherent motion. These are relatively simple, constrained situations in which parameter estimation and the fitting of empirical data seem to be reasonably straightforward.

With respect to Glass patterns, for example, in Chapter 2 natural neighbour loadings were derived. This raises the question of quantification by comparing a formal analysis of natural neighbour loadings with psychophysical measurements of Glass pattern detection. However, in Chapter 12 a significant point was made: proximity links can still show structure without connecting transformation partners. This more inclusive finding subsumes the utility of natural neighbor loadings. New methods relating to anisotropy and F ratios for variances of Delaunay triangle areas were developed. Again, this raises the question of quantification by comparing a formal analysis of these measurements with psychophysical measurements of Glass pattern detection. Such quantification could be accomplished, for example, by comparison of a formal analysis with outcomes in terms of the signal detection theory of Green & Swets, 1966, for at least a 'yes/no' pattern discrimination task, in which Glass patterns in various levels of noise are presented to subjects.

Natural versus person made

Delaunay triangulations for degenerate point sets have different statistical properties than for point sets in general position. But, of course, as point sets in general position approach degeneracy so the statistical properties of their Delaunay triangulations approach those for degenerate point sets. Degenerate point sets are mostly artificial in their 'perfect' geometry, typical, as they are, of Euclidean mathematics. Degenerate point sets are mostly person inspired representations, and are the exception in representations of nature. Nonetheless the striking difference between 'natural' and 'person made' can be marked by the difference in statistical properties of their Delaunay triangulations, and this could be a subject of some research in pattern recognition.

Further statistical development

Pattern detection methodology in psychology could be further developed along some of the following lines. Numbers of Delaunay neighbour points belonging to each point for different point patterns can have different distributions. Statistical operations on these distributions may yield information about point pattern type. Nonetheless Bootes (n.d.) indicates properties of Voronoi—therefore Delaunay—polygons have proved difficult to derive analytically; hence some reliance on Monte Carlo simulation for moment distributions.

Average neighbour distances for Delaunay neighbour sets have different distributions for different point patterns. The F Test for equality of variances—independent samples, $n - 1$ degrees of freedom—could be an appropriate test for differentiating two normalized point patterns based upon such average distances. For differentiating more than two normalized point patterns, one-way analysis of variance could be appropriate.

Average nearest neighbor distance for randomly distributed points appears to approach half the average Delaunay neighbour distance as the number of points becomes large. This could be investigated by way of graphing the two average distances for an increasing number of points (notwithstanding system array size limitations; something which I repeatedly encounter). It may be the case that patterns can deviate from the above-stated observation for random points. However, the procedure is time consuming and cumbersome.

Owing to the lack of significance measures, differences between expected and empirical averages for length of perimeter and area of Voronoi polygons cannot be proposed as tests for complete spatial randomness. However, the distributions for length of perimeter and area are still useful for spatial analysis. Classification of point patterns into clustered, random, or regular, can be performed by homogeneity tests. Pásztor (1994) underscores the point that random and regular point structures have unimodal normal distributions with different variances. Multimodality indicates a hierarchic cluster structure, with the number of modes determined by the number of scales in a sample. And this is indicated by my research.

Pásztor (1994) also underscores the point that the Voronoi model is useful for compiling density maps from point data. Areas of Voronoi polygons gauge inverse values of local intensity. By assigning 1/area to point locations as local densities, a map can be compiled with the aid of an interpolation method; deterministic or stochastic.

The importance of anisotropy in some patterns should not be understated. With regard to Ripley's K-function in Spatial Point Pattern Analysis, anisotropy can be detected by recording directions of points with respect to the point from which they are measured for distance lags. Directions are then processed along with counts for the distance lags. Some patterns show greater counts in a tolerance range for a particular direction over some distance. This methodology needs developing, particularly with regard to Glass patterns and reflection symmetry patterns.

When is a neighbour not a neighbour?

Results from experiments indicate that proximity is not absolute. Rather it is relative, and is shown to be encompassed in its degrees by Delaunay triangulation of a field. Relative proximity, neighbourliness, and adjacency overlap in substance, and are important factors in pattern detection. That elements at larger distances can be relevant while identical elements at lesser distances are not, demands investigation into just what constitutes proximity, neighbourliness, and adjacency, and the way they interact. What changes a more neighbourly situation into a less neighbourly one, and vice versa; and is this consistent with how subjects see a changing field? This could be facilitated by a computer program that examines Delaunay triangulation of a point pattern that continually changes in a principled manner (related to neighbourliness), and from which a list of properties could be formulated.

What constitutes a dot pattern in the context of this enquiry?

Generally, a transformation is a lawful difference from one condition to another. For a Glass pattern an original set of randomly distributed dots is overlaid with a transformed set, and the transformation cannot be too large. A minimum transformation exists that would superpose the two sets of dots. A Marroquin pattern maintains one pattern or another, no matter how large the transformation. This is because of proximal consistencies that are wrought by repetition of the same regularities within each set. For such a pattern, the minimum transformation that would superpose the two sets of dots need not be as large as the generating transformation.

It seems that proximal consistencies are a prerequisite for perceiving many dot patterns. However, reflection symmetry dot patterns do not adhere to this heuristic. What they have in common with other patterns is the element of transformation. (In fact they involve the simplest transform interpretation of all: a contributing factor in the will to develop a transformational model of visual perception.) Proximal consistency is no longer necessary. Transformation partners can be relatively large distances apart!

Owing to the nested arrangement of the aggregate of transformation pairs in reflection symmetry, there is not the masking of more and more separated transformation partners that there is for Glass patterns. However, by nesting Glass pairs generated with a range of transformation magnitudes, the upper limit of which is well beyond typical masking magnitude for a routine Glass pattern, the transformation characteristic appears to remain obvious for an average distance larger than that which obscures the transformation effect in the routine Glass pattern.

This could constitute a study of itself. A simple implementation—to be used advisedly—might involve transformation magnitudes for respective Glass pairs generated within some constrained random range. This can change the number of overlapping pairs in favour of a more nested arrangement. An example of a serious attempt to ‘unmask’ a Glass pattern involves nesting pairs without staggering them. That is to say, no pairs within largest transformation distance overlap and for these pairs more closely spaced transformation partners span the transformation normals of more widely spaced transformation partners. The theory is that what is in between is wholly contained, hence it does not confuse the issue. As long a pattern is not masked—by its own or extraneous elements—then a transformation is seen, even if the transformation is relatively large.

Reflection symmetry involves more than one transformation interpretation, and this could have a bearing on why it is comparatively fundamental. Left/right reflection symmetry, for example, can be generated by overlying an original set of dots with its horizontal distance complement (maximum possible x-coordinate distance minus actual x-coordinate distance, with y-coordinate distance unchanged). Moreover it can be generated by choosing an original set of dots that occupy just the left half of a form, for example, and by filling out the right half

with the horizontal distance complement. Left/right reflection symmetry can be also generated by overlying the original set of dots with the set resulting from a vertical depth rotation of 180^{o} in either direction.

Relative proximity is important to obviate masking interference in many patterns, but it is not necessary in all patterns. *The factor of transformation is the common element across all patterns.* Even for imprecise patterns, transformation to some degree of relaxed precision is relevant. A necessary condition for a pattern is that a transformation exists that at least inclines towards superposing the two sets of dots (although it is not necessarily a sufficient condition).

In any event, outcome may be summarized by the finding that for perception of unstructured visual arrays, distance mechanisms, particularly nearest neighbours, are important. For structured arrays, distance mechanisms along with perceived transformations are important. This might even be the case for reflection symmetry, where reflection counterparts can be widely separated. Once again, this would require that the visual system is sensitive to structure that varies over a wide range of different scales within a single stimulus (Dakin, 1997; Prazdny, 1986).

Constraints

Inverse problems need constraints to enable solutions. The account of 2D projections of views of 3D objects indicates considerabe constraints on matrices of affine transformations, whether pure or compound. (See Chapter 9, pages 196 to 199.) A thorough mathematical investigation is needed to determin what other constraints exist, and then to formulate a minimal representation of the constraints within a set. Constraints may be interrelated, and elements of a minimal set would be necessarily independent.

Appropriate parameters could then be indexed within the set, according to visual demand imposed by a 3D object, and applied to the 2D projection. The subset of parameters could be used to reconsruct the 3D object. In other words, parameters belonging to transformations needed to reconcile 2D projections of different views of a 3D object could be used in conjunction with projections to reconsruct the 3D object. Elaborating this could costitute a line of research.

Minimum Euclidean matching

A further line of investigation based on proximity could involve minimum Euclidean matching (see Edmonds, 1965; Gabow, 1972; Mirzaian, 1993, for example). This requires linking pairs of points to make separate dipoles, such that the sum of the resulting integer(n/2) number of link lengths is minimum.

Minimum Euclidean matching has figured in near optimal solutions to traveling salesperson problems (Christofides, 1976; Preparata & Shamos, 1985). However, it may also reveal structure in dot distributions. Maybe a recursive form of minimum Euclidean matching would be useful. This would involve the already linked dipoles becoming fundamental elements that require linking in pairs to generate new separate dipoles, and so on. In all these connections, the relation of minimum Euclidean matching to Voronoi, Delaunay linking would also need exploring.

Other transformations

It is well recognised that affine transformations can reconcile only part of the story. In categorization of the letters **m** and ***m***, for example, the second letter with respect to the first is contracted and sheared, as well as translated, but it is still affine transformed. However, all manner of homeomorpic transformations that differentially stretch and shrink can be applied

such that the second letter can be still recognised as belonging to the same category as the first.

Just as a distance hierarchy was used in this theseis, so a transformation hierarchy might have been used. Going from least to most inclusive, this is Euclidean, similarity, affine, projective, and topological. Increasing inclusivness allows transformational reconciliation of a wider range of forms. And Kolers (1972) supports the impression that we prefer transformations in order from simplest to most compex in an endeavour to make sense of the world. Indeed, just a shift in position of an object acccords with a Euclidean transformation, at the bottom end of the spectrum, and it is anticipated that handwritten charactes require a topological interpretation, at the top end.

The generative transformation model was conceived as a test project, to be kept relatively simple. The pattern measuring devices introduced in the first section primarily support the model in its pattern detection and discrimination endeavours, which utilize affine transformations as described in the second section. Having shown the utility of the model to date, improvements that include homeomorpic transformations do not seem so ambitious anymore; especially in light of today's computer morphing technology. Moreover treatments of other transformations, such as colour transformations and image discontinuities, included in the range examined by Kohler (1964), also do not seem so ambitious.

Relationship to various approaches

The generative transformational approach adopts elements from various approaches to visual perception, while rejecting others. In sympathy with the ecological approach (as per Gibson, 1979), description at the receptor level is considered to need complementing by description at the level of the optic array. At the receptor level, input for *vision* is a two-dimensional retinal array corresponding to intensity values for separate points of light. However, input for a *perceiver* is a *pattern;* the optic array extended in space and time, which is open to predictable change. Ways of describing the input for a perceiver are important.

The generative transformational approach works with invariants in the optic array, but contrary to the ecological approach they are detected *indirectly;* by mediation in 'early vision', which is 'cognitively impenetrable' as posited by Pylyshyn (1999). Like Marr (1982), the generative transformational approach draws on algorithms to realize properties of the optic array. The role of an algorithm is to turn one representation into another. Contrary to the ecological approach, *representations* are needed. (See Bruce, Green, & Georgeson, 2003.)

The generative transformational approach treats vision in its entirety as having evolved from the internal direction of what were originally overt actions. It posits a *reconstructive* vision, ultimately for the purpose of directing activity. Passive observation is wholly contained as a significant part of this 'active vision'.

The generative transformational approach draws on Leyton (1992). It adopts reconciliation, via inferred process history, of asymmetries with a maximally symmetric exemplar or canonical representation. And the generative transformation model looks to the Gestalt principals of organization in identifying relationships among elements in the optic array. These are economical relationships based on statistical regularities in the structure of the optic array. (See Barlow, 1999; Chun, 2000.) Such relationships are used to select transformations to accomplish reconciliation. Clearly, the generative transformation model could be improved by involving more of the Gestalt principles.

Some approaches to low-level vision posit feature detector cells tuned to respond to fixed attributes, not attributes scaled to pattern configuration. However, the human visual system is demonstrably sensitive to structure that varies over a wide range of scales. Accordingly, the generative transformation model gives exactly the same result for the same pattern proportionally changed by a range of different scaling factors.

Summary of basic assumptions, hypotheses, and implications

Much of the second section of this thesis has its context in a broader persuasion furnished by D. Vickers. The following are assumptions he intended as a guide to subsequent theorizing, a summary of the nature of the perceptual theory he envisaged, and a statement of some of the wider implications of the generative transformational approach. In his words:

> *The guiding assumptions are that:*
>
> 1. The brain has evolved to reflect constraints in the environment to which we are sensitive.
> 2. The most general physical principles should therefore be useful in structuring our explanations of mental processes.
> 3. These mental processes (whether experienced or inferred), the symbol systems (by means of which we represent them), the functional structure of the brain, and the physical events in question can all be understood in terms of a common set of geometric notions operating at a single level.
>
> *According to the perceptual theory that is envisaged:*
>
> 1. The perceptual system codes incoming information in terms of the simplest possible algorithm for generating an output that is "more or less" compatible with the incoming sensory data.
> 2. The simplest possible algorithm is that which requires the simplest set of transformations to generate a replica of the current sensory input.
> 3. This economical coding also corresponds to the objective state of affairs with the highest likelihood of occurring.
> 4. Coding in terms of the simplest possible set of transformations maximizes the symmetry, invariance, or partial self-similarity in the stimulus array.
> 5. This coding is accomplished, not by a sequential, computational algorithm, but by a recurrent neural network.
> 6. The operation of this recurrent network can by described in terms of a nonlinear dynamical system and its output characterized as an attractor of such a system.
>
> *The following, more general hypotheses concerning brain function are consistent with this particular theoretical approach to visual perception:*
>
> 1. The single major function carried out by the brain is to perform multiple geometric transformations on patterns of incoming sensory excitation (i.e., the brain is a massive, parallel transformation engine).
> 2. All significant mental events and processes are determined by symmetries (or invariants, congruencies, or equivalences), associated with these transformations, or to which they give rise through continued iteration.
> 3. Different cognitive activities are differentiated by their different transformation groups and associated symmetries.
>
> *From those interrelated hypotheses four further implications follow as natural extensions:*
>
> 1. The relations between perceptual organization (i.e. perception of the intrinsic structure of an object), the perception of orientation and layout (based on extrinsic relations between the light source, the surface(s) in question and observer), and cognitive and linguistic categorization can all be explicated in terms of the relations between the different invariants associated with groups of progressively less constrained transformations, ranging from similarity transformations, through affine and projective transformations, up to those which may be described as topological;
> 2. The perceptual or psychological distance between two stimuli is determined by the locus of the 'simplest' transformations relating them;
> 3. Natural processes of reasoning and inference are not mediated by adherence to a rule-based formal system of the kind typified by the predicate calculus, but by transformations of a geometric, analogue kind, carried out on the basic elements of a topological system; and
> 4. The response to invariants arising from multiple iterated geometric transformations is made possible by the anatomical structure of the brain and by the nature of its early growth and later development.
>
> (Vickers, n.d., pp. 191-192)

May his *vision* endure in further inquiry.