Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 44 (2004) 2179–2186
Vision Research
www.elsevier.com/locate/visres

Theory of texture discrimination of based on higher-order perturbations in individual texture samples

Christopher W. Tyler *

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

Received 21 July 2003; received in revised form 29 March 2004

Abstract

This analysis addresses the issue that texture properties are defined on ensembles of possible textures, while psychophysical judgments of texture properties must be made on individual texture samples, or regions of uniform texture within a larger texture field. Since the basic discrimination task requires comparison of two sample images (or regions) specified by different ensemble rules, the viewer is thus required to compare the estimates of their ensemble statistics of single textures. This paper develops a theory of texture discrimination incorporating a roving local sampling window that allows the visual system to derive an estimate of the ensemble statistics over the window from any particular texture image, without the need to present multiple samples for evaluation. This approach to texture explains how we can have a clear sense that two patterns derive from different statistical generation rules even though we see only one example of each type. In providing the theoretical basis for texture discrimination of individual samples, this analysis goes beyond previous work to account for our intuitive sense that we can estimate the generation rule underlying particular textures. It also analyzes the decision process for discriminating texture boundaries in extended images, defining a novel "Gregorian attractor" that replaces and extends standard Bayesian decision rules.

Keywords: Texture; Ensemble; Random; Statistics generation rule

1. Relevance of the complete reconstruction analysis to human texture discrimination

A general approach to the classification and analysis of textures is offered by their statistical properties. As an early exponent of the parallel processing approach to texture perception, Julesz (1962) developed the analysis of statistical constraints in random-dot fields. Julesz defined his statistics to enumerate the mean frequency of all occurrences of colorings of pairs, triplets and, generally, *k*-gram sets of points at all spacings throughout a texture ensemble. This analysis precipitated an exploration of the order of texture specification that could be discriminated by human observers (see Klein & Tyler, 1986; Tyler, 2004). To provide a specific example of higher-order textures, Fig. 1 shows a random plaid texture accompanied by a binary random pattern with the same mean, mean digram and mean trigram statistics (from Julesz, Gilbert, & Victor, 1978). It is only at fourth-order and beyond that the mean statistics of these two types of texture differ.

Yellott (1993), on the other hand, published the demonstration that any pattern may be *completely reconstructed* from the full specification of its third-order statistics, which appears to invalidate the notion that there is any value in analyzing texture statistics of fourth and higher-order. This demonstration requires that the pattern has finite 'support', i.e., is of finite extent with a known value (e.g., zero) outside the region where the pattern is specified. As long as the third-order statistics are available to the perceiving system, Yellott argued that any pattern with finite support is discriminable in principle from any other. Indeed, an even stricter constraint has been specified by Chubb and Yellott (2000), that any finite pattern is uniquely defined by its digram or second-order histogram statistics. It should be made clear that both theorems are based on the discrete histogram statistics, which incorporate the particular numbers of all combinations of all luminance levels at all spacings in all directions. As such, Chubb and Yellott

* Tel.: +1-415-345-2065; fax: +1-415-345-8455.
E-mail address: cwt@ski.org (C.W. Tyler).

0042-6989/$ - see front matter
doi:10.1016/j.visres.2004.03.029

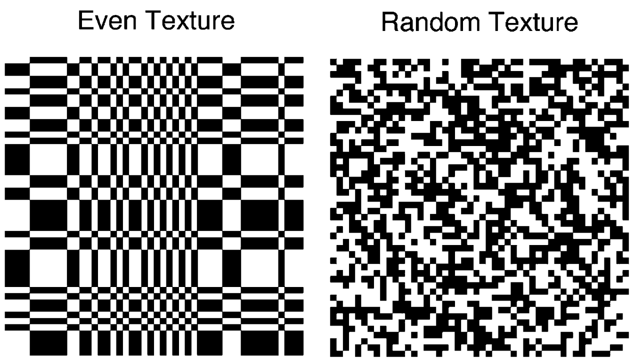

Fig. 1. Examples of a fourth-order random plaid (left) and a random binary texture (right). One sample is sufficient for us to infer that the two patterns has different generation rules.

make clear that the *n*th-order histogram statistics are generally a much larger number set that that of the image itself. The histograms do not provide any compression of the information in the usual sense of statistical cumulation.

As pointed out by Victor (1994), however, this property of reconstructability from the second-order statistics is valid only for the statistics of *particular* images, while the statistical parameters usually are defined on the *ensemble* of all possible images that could be generated under a particular set of statistical constraints. Interpretation in terms of ensembles and their generator rules means that it is only the *probability* of each *k*-gram coloring that is specified by the statistics, not the actual frequency of occurrence in a particular image. Julesz (1962) originally defined texture statistics in this probabilistic sense. The statistic has a defined value whereas the probability has only an expected value that does not mandate a particular value at any point in the image. Chubb and Yellott's (2000) demonstrations therefore do not apply to textures defined by ensemble statistics, so that one may legitimately ask the question of whether a texture drawn from one ensemble is discriminable from a texture drawn from another ensemble. (For maximum clarity, the term 'pattern' will be reserved for the isolated images that are subject to Yellott's (1993) constraints, whereas the term 'texture' will be used for an exemplar of an ensemble generation rule. The ensembles will be assumed to be effectively infinite for the purposes of this theoretical development, because the ensemble of even a binary texture of 100×100 elements has 2^10,000 distinct patterns, far too many for the visual system or even a computer to reasonably enumerate. It will therefore be assumed that all ensemble statistics are estimated from limited random samples of the ensemble.)

This ensemble-based view of texture statistics raises the problem, however, that pattern statistics do not provide a basis for discriminating between an image generated from a particular set of texture statistics and a random image. The reason is that *any* pattern could have arisen from a random generation rule (with equal probability to any other image), and thus any pattern from a given set of image statistics could, in principle, have arisen from the random ensemble. Therefore, the statistics for any particular image cannot definitively distinguish the ensemble from which it was drawn from a random ensemble. In principle, according to this view, one can test the discriminability of ensemble statistics only on *ensembles* of images, which provide an inductive indication of their generation rule by averaging over many samples (Victor, 1994). (In finite ensembles one may be able to distinguish among particular generation rules, but one can never be sure that the pattern did not arise from a random generation rule.)

Adoption of the ensemble approach means that the question of reconstruction of a pattern from its particular second-order histogram frequencies is avoided; ensembles are infinite (or extremely large) sets, making their reconstruction impractical from their statistics at second, third, or any order because the sets of samples do not have finite support. In practice, however, perceptual discriminability must be established for sets of patterns generated from the statistics with the appropriate uncertainty on each parameter, so that each one is a sample from the infinite ensemble. The ensemble is represented by sets of such particular patterns specified to any desired order, treating such sets as a way of approximating the ensemble as a whole. In this way, the ensemble approach advocated by Victor (1994) avoids the second-order limitation propounded by Chubb and Yellott (2000).

2. Problems with the ensemble approach to texture

In one sense, however, the ensemble approach is subject to the same objection that Yellott (1993) made against discriminability of a particular pattern. Because the ensemble as a whole is an infinite (or extremely large) set, only a finite sample of the ensemble of images is accessible to experimental test. The finite set that is actually used in an experiment provides only a sample of the overall ensemble parameters and one that, in principle, could again have been generated by a random rule. Thus, the ensemble approach does not avoid the Yellott objection in principle, it merely enlarges the scope of the test from a single pattern discrimination to an extended experiment (for anything but the smallest patterns).

However, the advantage of the ensemble approach is that the ensemble statistics of all orders up to the pattern size may be estimated from the sample set, in terms of both their mean values and their variances (together with higher-order moments). The important property of these estimates is that they now assign not only a particular probability of occurrence to the individual *k*-

gram figures (which were all equal under the random generation rule), but also estimates of their standard deviation and other statistical moments about the mean. One can then evaluate whether these *ensemble* statistics are significantly different from random or from those of a comparison set of images generated by a different generation rule. By providing statistical estimators rather than simply the defined frequencies of occurrence, the ensemble approach allows meaningful evaluation of statistical parameters beyond the second-order moments.

3. Definition of the statistics of an individual texture

Nevertheless, Victor's (1994) ensemble approach, which is acknowledged by Chubb and Yellott (2000), is unsatisfying as a way to avoid the third-order ceiling. A glance at Fig. 1 reveals that we do not need an ensemble of textures to discriminate a texture with statistical structure from a purely random texture; single examples suffice in practice. An alternative approach to the issue is to base the analysis on the specific definition of the *texture* as opposed to the broader class of all possible patterns. As in Klein and Tyler (1986), a texture may be defined as a pattern that is self-similar over translation up to some scale of analysis (viz, size of the window within which the k-gram statistics are evaluated). The designation "self-similar" is here intended in the sense of ergodic, where an ergodic pattern is one in which the statistical parameters of samples of the pattern are similar to those of the pattern as whole (and of the infinite ensemble from which the pattern was drawn). It is not intended to allude to the fractal quality of self-similarity over scale. Indeed, fractal patterns may or may not exhibit self-similarity over translation, depending on their generation rule.

The concept of self-similarity over translation leads us to a concept of two types of texture (see Fig. 2). A *regular* texture (such as a grating) is one that is exactly self-similar on some discrete translation matrix. A *statistical* texture (such as a random-dot field) is one whose *local* statistics are identical (or not statistically different) throughout the image. Many of the textures developed by Julesz (1975) and co-workers are hybrids of these two types, with random perturbation from a regular matrix. The definitions of regular and statistical textures may be expressed mathematically in terms of a windowed version of the autocorrelation definition of texture statistics, or windowed autocorrelation function, $\mathrm{WACF}_k(\Delta R)$ at order k,

$$\mathrm{WACF}_k(\Delta R) = \int W(\Delta \mathbf{r}) \prod_{i=1}^{k} L(\mathbf{r} + \Delta_i \mathbf{r})\, d\mathbf{r}, \tag{1}$$

where $\mathbf{r} = (x, y)$, $L()$ is the two-dimensional image in $\mathbf{r}$ and $\Delta_k \mathbf{r} = (\Delta_k x, \Delta_k y)$ for a 2D texture, and $W(\Delta \mathbf{r})$ is a windowing function defining the range of validity of the statistics.

(This specification is a reduced representation relative to the histogram statistics preferred by Chubb, Econopouly, and Landy (1994) and Chubb and Yellott (2002). It assumes that the relevant property for visual processing is the average over the intensity histogram at each location, and is insensitive to special combinations of intensity levels over space.)

Thus, a statistical texture is a pattern in which the expected value of the ACF is spatially homogeneous:

$$\overline{\mathrm{WACF}_k}(\Delta \mathbf{r}) \approx \overline{\mathrm{WACF}_k}(\mathbf{r} + \Delta \mathbf{r}) \tag{2}$$

for all $\mathbf{r}$, within the accuracy of statistical estimation. Specification of this latter restriction is the burden of the rest of this paper.

For a regular texture, Eq. (2) is restricted to some discrete subset of points in the space of $\Delta \mathbf{r}$. The textures used by Julesz and co-workers, however, were typically restricted to a discrete grid of neighborhoods in $\mathbf{r}$. In the

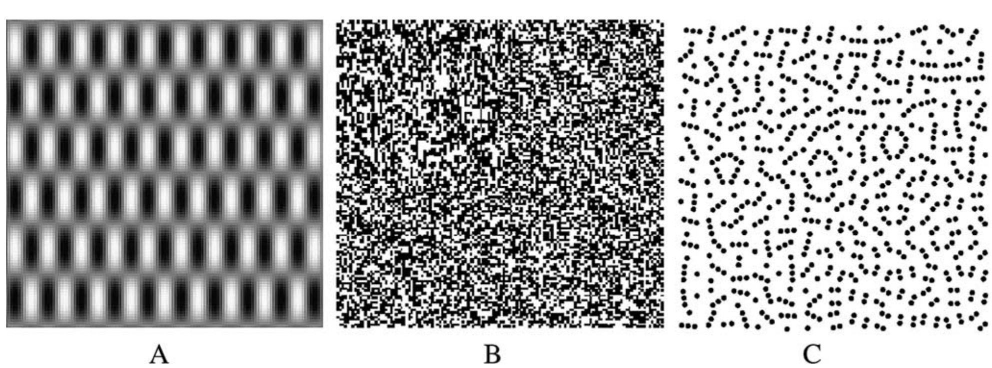

Fig. 2. Texture types. (A) Example of a regular texture. (B) Example of a random binary texture; the upper right quadrant is derived from a different generation rule where elements are vertically paired. (C) Semi-regular texture (modified from Julesz, 1975, designed by the present author as the first violation of Julesz' second-order conjecture) consisting of dot quartets randomly rotated about point defined on a regular grid. The lower right quadrant has quartets in a different configuration that is equated for second-order spatial statistics, but differs at the third order.

Fig. 3. Statistics of the sampling window (exemplified by the gray window). (A–D) show textures with sampling density increasing in factors of 2. Selected pixels (gray) can exhibit extreme statistics (such as uniform coloration) at low density (A), but are tightly constrained close to the mean level at high density (D).

sense defined, a statistical texture may be generated by tiling of the pattern space with an ensemble of patterns of the size of the analysis window (see Fig. 3), each drawn according to the specified ensemble statistics $\overline{\mathrm{WACF}_k}(\Delta\mathbf{r})$. Of course, the texture approach breaks down when the size of the statistical window $W(x,y)$ becomes equal to the size of the full pattern, which is also the point at which the pattern no longer conforms to the definition of a texture. Theoretically, however, a pattern will conform to the definition of Eqs. (1) and (2) as long as the statistical window is smaller than the pattern by a finite amount. Say, for example, that we permit statistical windows up to within one pixel of the edge of the texture (represented as a raster of pixels). For k-grams up to this scale, there are repeated examples in the image from which the ensemble statistics may be estimated (although with a reduced precision as the number of occurrences decreases with increasing size of the k-gram) in the Chubb/Yellott theorem. Any window that contains multiple representations of all k-grams limits the reconstruction to statistical estimators and does not permit exact (or even approximate) reconstruction. Thus, it is only the inclusion of the k-grams having a *single* representation in the texture that allows reconstruction of the texture from its second-order histograms.

This analysis of the second-order reconstruction theorem reveals that it is not, in fact, a theorem of statistical estimation but a theorem of occurrence frequencies (as acknowledged by Chubb and Yellott (2000)). Without access to the full histogram structure, one cannot, in principle, reconstruct an image from the parameters of its statistics (e.g., its means and higher-order moments) because they are inherently probabilistic concepts. Every instance of the reconstruction from such order parameters will be different even if the parameters are known to the full order of the texture size. Only if the precise frequencies of occurrence are determined for the full texture, or the full set of moments for that particular texture, is exact reconstruction possible. The frequencies of occurrence may be statistics in the baseball sense but they are not statistical parameters in the usual probabilistic sense.

4. Ideal observer for the texture generation rule

If one can make a valid inference of the generation rule underlying particular textures, it follows that there must be an ideal observer for such a estimation, one that can extract all the information in the particular texture relating to the generation rule (assuming ergodicity of the generator rule over space). A plausible candidate for the ideal observer would be the mean value of each point in the k-gram histogram H at every order k.

$$\widehat{I}(l_k, \Delta\mathbf{r}) = \overline{H}_k(l, \Delta\mathbf{r}), \tag{3}$$

where l_k is the dimensionality of the color combinations at the kth order.

Thus, for a particular pattern, the observer's best estimate of the mean density in the generation rule (in the absence of other information) is the mean density of that pattern; the best estimate of the luminance histogram is the luminance values available in the histogram of particular pattern, and so on.

Of course, there is the opportunity to go beyond the data if prior information is available on the likelihoods of particular values in the generator histograms. For example, if the mean density in the sample texture is close to 0.5, prior experience with experimenter behavior tells us that the generator density was probably 0.5, rather than the actual value of the mean. However, if the mean value of the prior is well outside the confidence interval for the mean from this sample, we would be justified in assuming that the generator density was some value other than the prior, and is consequently best estimated by the mean itself. In related fashion, one can extract higher-order 'priors' from the sample itself. The null hypothesis for the mean frequency of all histogram values is that derived from the mean frequencies of the constituents through the binomial probabilities of k-grams of their combinations. Thus, for a 0.5 density black/white texture, we should expect the second-order histogram to be 0.25 black/black, 0.5 black/white and 0.25 white/white at all displacements. Any empirical value in the histogram (within the binomial confidence interval for this sample) may be supposed to indicate that the generator value was derived from this expectation. Only values that deviate *significantly* from these expectations (in terms of the statistical confidence interval) should be taken as valid estimators of a non-random component in the generator rule.

5. Incorporating the priors through a Gregorian attractor

Interestingly, this analysis leads us to a post-Bayesian rule for statistical estimation. Let us call this a 'Gregorian attractor'. Instead of Bayesian rule of the posterior probability being the weighted product of the prior probability $P_k(l, \Delta\mathbf{r})$ and the empirically observed probabilities $E_k(l, \Delta\mathbf{r})$, the Gregorian attractor is a decision rule that accepts the prior completely if it is compatible with the observed probability, but switches to the observed probability completely if the latter falls outside the confidence interval of the prior. In symbolic form, we would have:

$$\hat{G}(l_i, \Delta\mathbf{r}) = \begin{cases} P_k(l, \Delta\mathbf{r}), & \text{if } |E_i(l, \Delta\mathbf{r}) - P_i(l, \Delta\mathbf{r})| < z\sigma_i, \\ E_k(l, \Delta\mathbf{r}), & \text{otherwise}, \end{cases} \tag{4}$$

where σ_k is the standard deviation of the occurrence frequency and z is the criterion significance level in units of σ_k (taking a typical value of $2 < z < 4$).

In the context of higher-order texture perception, we may envisage that the P_k are derived from the combinatorics of the lower-order histograms, such as the binomial probabilities in the case of a binary texture. Thus, we have a ready means to fill in the priors at each order from the information available in the texture. This procedure may be used to define the ideal observer for the decision rule as to whether to accept or reject the prior at each order of texture definition.

The Gregorian attractor corresponds to a probabilistic instantiation of Gregory's (1980) concept of perceptual hypothesis testing. In reduced-cue situations, the perceptual system (in visual and other sensory modalities) develops hypotheses as to the external reality underlying the sensory input. The characteristic of these hypotheses is that they are discrete interpretations, each incompatible with the next. The perceptual system evaluates the validity of one hypothesis (such as, for example, that a set of local motions derives from a rigid 3D transformation). If this hypothesis fails (because no rigid transformation can be found that fits the full set of motion data), a decision is made to reject the hypothesis and a new hypothesis is generated (such as that the local motions fit a *pair* of rigid transformations). This process of hypothesis testing obviously must have some error tolerance, even if that tolerance is determined by the intrinsic noise of the sensory signal.

As made explicit by Gregory (1980), the perceptual process of hypothesis testing is fully analogous with the corresponding process of scientific investigation as a whole. The hypothetico-deductive system of science is based on the concept that hypotheses are generated by deduction from one set of data and tested against another set. Further, acceptance or rejection of a hypothesis is based on some level of error tolerance. Strictly, the deviation of the new data from the predictions of the hypothesis has to be outside the statistical error range to be considered a significant violation. Practically, the possibility of experimenter error (either operationally or conceptually) is also factored in by the wider community, so the violations must be large and consistent before a plausible scientific hypothesis is taken to be rejected, particularly if it is well supported by a wide array of prior studies.

Thus, the basic concept of a Gregorian attractor is well developed in statistical analysis, making a binary choice between the experimental and the null hypotheses. There is no case in standard statistical practice where the Bayesian concept is applied—by taking the product of the distributions for experimental hypothesis and the null hypothesis. Statistical estimators always either accept or reject the null hypothesis based on a criterion level derived from the probability distribution.

In perception, however, either the Bayesian or the Gregorian rule may apply, according to circumstance. When the two perceptual hypotheses are very close in form, the Bayesian rule may make sense (for example in estimating surface slant in a visual scene from stereoscopic and texture gradient cues, each with their own perceptual limitations). This slant estimation task is simulated in Fig. 4, in which the statistical decision is made between a prior theoretical slope (Fig. 4A) or a best-fitting empirical slope (Fig. 4B). However, the decision to split from this compromise may be taken when the quantitative deviations become too extreme. This splitting point amounts to a Bayesian rule for abandoning the Bayesian combination principle in favor of the dichotomous decision of Eq. (4), combining the two hypotheses into a joint solution in which *both* are true (Fig. 4C). In perception, the decision to split amounts to a version of perceptual transparency. Instead of a single plane, the scene may be perceived as two transparent planes with the scene coloration split between them. Instead of a moving texture, the motion aftereffect may be split into a stationary object with a flow moving over it. And so on. Thus the statistical priors of the three kinds of hypothesis $p(P)$, $p(E)$ and $p(S)$, may define the choice between the interpretation in terms of (i) the prior hypothesis, P, (ii) the empirical

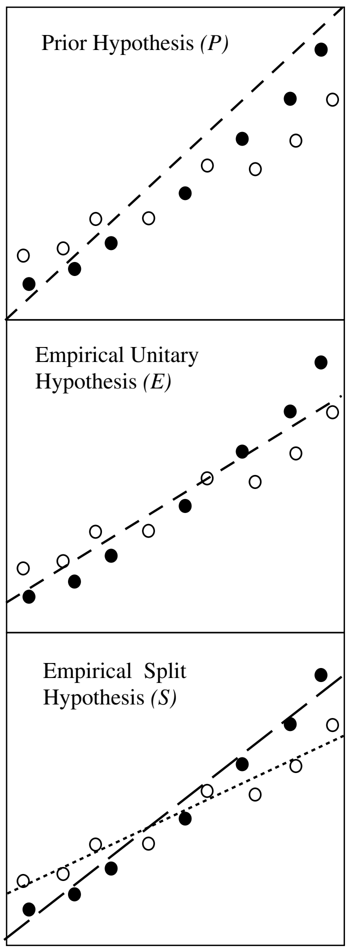

Fig. 4. Illustrative data for a sloped function defined by two separate sensory cues (filled and open circles). The three panels illustrate the Gregorian choice between the Bayesian prior, P (standard slope, panel A), the empirical hypothesis, E (best-fitting slope, panel B) and the transparent split, S (dual slopes, panel C).

hypothesis, E, and (iii) the split hypothesis, S, in which both are simultaneously valid.

From these considerations, we may derive the final version of the decision rule, which may be termed the 'generalized Gregorian attractor'. Here the deviation from the null hypothesis, or Bayesian prior, results in a tripartite choice of options, C_H based on their own Bayesian priors $p(E)$, $p(E)$ and $p(S)$ in this situation.

$$\hat{G}(l_i,\Delta\mathbf{r}) = \begin{cases} P_k(l,\Delta\mathbf{r}), & \text{if } |E_i(l,\Delta\mathbf{r}) - P_i(l,\Delta\mathbf{r})| < z\sigma_i, \\ \max \begin{bmatrix} p(P) * P_i(l,\Delta\mathbf{r}), \\ p(E) * E_i(l,\Delta\mathbf{r}), \\ p(S) * S_i(l,\Delta\mathbf{r}), \end{bmatrix} & \text{otherwise.} \end{cases} \tag{5}$$

If the observation falls within the specified confidence interval $z\sigma$ of the prior, the prior is accepted. If not, the outcome is probabilistically assigned among the prior, the novel stimulus-based hypothesis and the simultaneous application of both hypotheses, according to their own

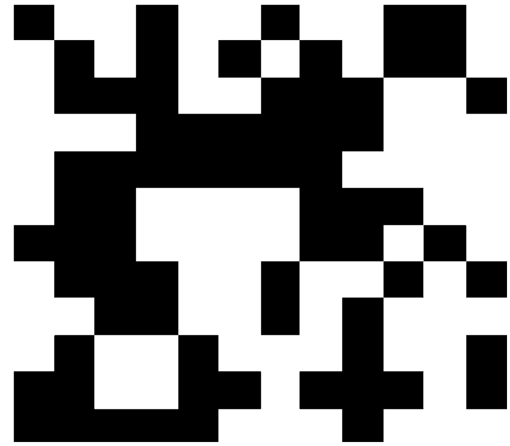

Fig. 5. Biased texture. Is this an extreme example from a uniform texture rule, a texture gradient from high to low density, or a segmented texture consisting of two squares with different density?

relative prior weightings of plausibility. For example, if we looked at a texture that was white on one side with about 50% black dots on the other side (Fig. 5), it would be well outside the distribution of a balanced black/white texture. We could decide that it was an extreme example from an equal black-white distribution (P, with a very high z), a biased texture from a different generation rule (E) or a segmented image with the two sides exhibiting two rules independently (S). In this case, the decision concerning the first-order density rule has no priors within the image to aid the decision. Which we decide will depend on our own prior experience with such textures in the laboratory and in the world in general, and may be subject to large individual differences. The generalized Gregorian attractor may be further expanded to include any other form of competing hypotheses among which perception alternates probabilistically.

6. Conclusions

In conclusion, the local roving window definition of a texture allows one to derive an estimate of the ensemble statistics from any particular texture image. For two sample images drawn from different ensemble rules, the viewer is thus in a position to compare the estimates of their ensemble statistics. This definition of texture allows the ensemble statistics to be estimated from an *individual* sample texture, without the need to present multiple samples for evaluation. The properties of the ensemble are implicit in each individual sample of the texture as long as the sample is substantially larger than the window of analysis. This approach to texture explains how we can have a clear sense that two patterns derive from different statistical generation rules even though we see only one example of each type. In providing the theoretical basis for texture discrimination of in individual samples, this analysis goes beyond those provided either by Chubb and Yellott (2000)—for whom there is no generation rule, only specific patterns—and Victor (1994)—for whom the generation rule is accessible only through multiple samples—to account for our intuitive sense that we have some estimate of the generation rule underlying particular textures.

Analysis of the decision structure required to account for the texture perception leads to a post-Bayesian rule for statistical estimation appropriately termed the 'generalized Gregorian attractor'. Instead of Bayesian rule of the posterior probability being the weighted product of the prior probability and the observed probabilities, the decision rule accepts the prior completely if it is compatible with the observed probability, but switches to an alternative strategy if the latter falls outside the confidence interval of the prior. In perception, various rules may apply, according to circumstance (Gregory, 1980). When the two hypotheses are very close in form, the Bayesian rule of the product of the two competing distributions may make sense, and has been validated by considerable experimental support. A relevant case is that of cue combination, where two or more cues may carry qualitatively similar information that is quantitatively different. In such cases, it may make sense to combine the two estimates in Bayesian compromise. The decision to split from this compromise, on the other hand, may be taken when the quantitative deviations become too extreme, leading to the dichotomous decision strategy of Eq. (4). This splitting point amounts to a Bayesian rule for abandoning the Bayesian combination principle. This splitting decision is the statistical instantiation of Gregory's 'hypothesis-testing' view of perceptual verification.

Perceptually, however, it may make sense to abandon the need for either a compromise or a split, and instead combine the two hypotheses into a joint solution in which both are true (Eq. (5)). For example, the local motion of dots can be rigid to both the left and the right if there is motion within a pair of transparent planes. In the texture context, one could imagine regions of texture in which one rule applied with a different rule in other regions, like a patchwork quilt. In everyday life, such regions are typically marked by changes in the mean color of each type of texture, but it is obviously possible to apply different rules without such color distinction. This brings us to the realm of texture segmentation, in which it is assumed that the image is composed of regions of different textures each of which is homogeneous in its properties. Determination of the boundary of such textures corresponds to an application of Gregory's hypothesis-testing in the texture domain. Having established the hypothesis of one set of texture parameters in one region of the image, the perceptual system must pursue that hypothesis to the boundary of the texture, then decide to terminate the zone of application of that hypothesis as the new texture begins to dominate the statistical estimator. Such processes are commonly incorporated into computer algorithms for texture segmentation, but it is clear that the visual system must perform the same analysis for its interpretation of the structure of the world.

Finally, the discussion of texture segmentation is not meant to exclude the issue of texture gradients and continuous variation of texture parameters for shape or of structural reasons. It should be clear that the Gregorian attractor of Eqs. (4) and (5) allows the decision of a continuous texture when the local change falls within the error tolerance. Although there may be discrete priors on the parameters to be expected in the laboratory situation, these are unlikely to be found in real-world textures. Thus, we may readily allow the perceptual hypothesis being tested to be a function of space rather than to be a fixed value. The visual system may develop the function of a continuous variation in the texture parameters corresponding to a texture gradient of some shape in space (linear receding, spherical, or any other shape). Such continuous variation is embedded within the equations when expressed as functions of space (or time), rather than discrete values.

Acknowledgements

Thanks to Jonathan Victor, Charles Chubb and James Coughlan for comments on an earlier version of this paper. Supported by NIH EY13025.

References

Chubb, C., Econopouly, J., & Landy, M. S. (1994). Histogram contrast analysis and the visual segregation of IID textures. *Journal of the Optical Society of America A, 11*, 2350–2374.

Chubb, C., & Yellott, J. I. (2000). Every discrete finite image is uniquely determined by its dipole histogram. *Vision Research, 40*, 485–492.

Chubb, C., & Yellott, J. I. (2002). Dipole statistics of discrete finite images: Two visually motivated representation theorems. *Journal of the Optical Society of America A, 19*, 825–832.

Gregory, R. L. (1980). Perceptions as hypotheses. *Proceedings of the Royal Society of London, 290*, 181–197.

Julesz, B. (1962). Visual pattern discrimination. *IRE Transactions on Information Theory, IT-8*, 84–92.

Julesz, B. (1975). Experiments in the visual perception of texture. *Scientific American, 232*, 34–43.

Julesz, B., Gilbert, E. N., & Victor, J. D. (1978). Visual discrimination of textures with identical third-order statistics. *Biological Cybernetics, 31*, 137–140.

Klein, S. A., & Tyler, C. W. (1986). Phase discrimination of compound gratings: Generalized autocorrelation analysis. *Journal of the Optical Society of America A, 3*, 868–879.

Tyler, C. W. (2004). Beyond fourth-order texture discrimination: Generation of extreme-order and statistically-balanced textures. *Vision Research*. Submitted for publication. doi:10.1016/j.visres.2004.03.032.

Victor, J. D. (1994). Images, statistics, and textures: A comment on implications of triple correlation uniqueness for texture statistics and the Julesz conjecture. *Journal of the Optical Society of America A, 11*, 1680–1684.

Yellott, J. (1993). Implications of triple correlation uniqueness for texture statistics and the Julesz conjecture. *Journal of the Optical Society of America A, 10*, 777–793.