



Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2005-008
AIM-2005-002
CBCL-244

February 7, 2005

Using Computational Models to Study Texture Representations in the Human Visual System

Benjamin Balas



Abstract

Traditionally, human texture perception has been studied using artificial textures made of random-dot patterns or abstract structured elements. At the same time, computer algorithms for the synthesis of natural textures have improved dramatically. The current study seeks to unify these two fields of research through a psychophysical assessment of a particular computational model, thus providing a sense of what image statistics are most vital for representing a range of natural textures. We employ Portilla and Simoncelli's 2000 model of texture synthesis for this task (a parametric model of analysis and synthesis designed to mimic computations carried out by the human visual system). We find an intriguing interaction between texture type (periodic v. structured) and image statistics (autocorrelation function and filter magnitude correlations), suggesting different processing strategies may be employed for these two texture families under pre-attentive viewing.

Acknowledgment of support: BJB is supported by a National Defense Science and Engineering Graduate Fellowship.

Introduction

The visual perception of textures has been an area of interest spanning a wide variety of disciplines from art to computer science. The fields of computer vision, perception, and graphics have each made significant contributions to our overall understanding of texture perception and representation, albeit in quite different ways.

Psychophysical studies of texture perception

Psychophysicists are of course most interested in what representations and rules the human visual system uses to process textures. In this endeavor, Bela Julesz stands out as one of the earliest and arguably most important contributors to the field. The famous “Julesz conjecture” continues to guide research in human texture perception to this day, having originated in a 1962 article (Julesz, 1962) wherein Julesz hypothesized that textures that differed only in 3rd-order or higher pixel statistics would be indiscriminable by human observers. This early version of the conjecture was proved false by Julesz himself years later (Julesz, Gilbert, Shepp, & Frisch, 1973) and the hypothesized “bar” for human discriminability of textures has been pushed past 3rd-order statistics (Julesz, Gilbert, & Victor, 1978) to a possible resting place at 4th order statistics (Klein & Tyler, 1986). However, recent work analyzing the formalism of creating extreme-order textures (Tyler, 2004a) suggests that the global statistics should not be the sole focus of texture research. Local processes that human observers use to compare different texture samples may be of more importance. (Tyler, 2004b).

The majority of studies concerned with the psychophysics of human texture perception make use of random-dot textures, making pixel-level texture analysis a relevant tool. Though useful as a model world for examining texture processing strategies, it must be acknowledged that random textures are hardly representative of the set of natural textures we encounter in everyday experience. Indeed, these textures violate key facts regarding the structure of natural images that have been well-known for some time, specifically the redundancy of natural images (Attneave, 1954; Barlow, 1961). Human observers have implicit knowledge of this redundancy (Kersten, 1987), suggesting that it may be better to study natural textures that match statistical properties of the real world. Natural images have been used to study what higher-level image qualities are used to group textures along salient dimensions (Rao & Lohse, 1996), but little effort has been made to examine low-level representations of photographic textures using psychophysical methods.

Analysis and synthesis of photographic textures

A useful body of work to consider as a means of resolving this difficulty is the growing number of algorithms proposed in the computer vision literature for texture analysis and synthesis. All of these algorithms share the goal of using small samples of some original texture as a starting point for the reconstruction of arbitrarily large amounts of the same texture. The end result should ideally be indistinguishable from the true texture, although no algorithm can truly remove all artifacts of the synthesis process. Rather than random-dot textures, these algorithms are most often applied to natural textures and have been very successful at creating convincing images for graphics applications. Given that these algorithms operate on natural textures, we will consider them as a useful vehicle for studying the perception of such images by human observers.

Clearly the quality of the final reconstruction produced by any of these algorithms informs us as to the utility of both the representation used for the original texture and the process by which that representation is used to generate novel images. However, for us to truly feel confident in relating the computational procedure used for texture synthesis to human perceptual processes it is helpful if the algorithm uses representations employed by the human visual system. For this reason, several texture synthesis strategies that produce strikingly good reproductions of target textures will not be considered here. For example, “image quilting” strategies (Efros & Freeman, 2001) have no true “representation” of a texture, in that patches of the original image are reassembled to make the synthetic version. In a sense, the original image is the only representation of the texture used. Likewise, pixel-growing strategies (Efros & Leung, 1999) are equally problematic in that they represent texture in terms of the distribution of individual pixels in the original image. Synthesis requires a time-consuming search process through the sample provided for analysis. While both of these procedures (and their associated variants) are extremely useful for graphics applications, we shall not consider non-parametric processes at present.

To achieve a deeper insight as to what statistics are important for the visual processing of natural textures, we turn instead to parametric models of texture analysis and synthesis. These models utilize the idea that filters resembling those found in early visual cortex provide information useful for texture segmentation and classification (Bergen & Adelson, 1988; Bergen & Adelson, 1986). Texture analysis by such filters has proven quite successful at modeling pre-attentive segmentation performance (Malik & Perona, 1990). Filter-based analysis has also contributed to a formal definition of Julesz’ “textons” (Julesz, 1981) in terms of clustered filter outputs (Malik, Belongie, Leung, & Shi, 2001).

In terms of texture synthesis, Heeger and Bergen's model (Heeger & Bergen, 1995) demonstrated the utility of "steerable filters" (Simoncelli & Freeman, 1995) for the synthesis of stochastic textures that lacked global structure or distinct textural sub-regions.

The model of Portilla and Simoncelli

Heeger and Bergen's model has been improved upon in many ways since its initial presentation. In particular, to overcome the inability of the original model to reproduce extended contours and other large-scale structures in the target texture, additional constraints across scales and orientations were introduced by Portilla and Simoncelli (Simoncelli, 1997; Simoncelli & Portilla, 1998; Portilla & Simoncelli, 1999; Portilla & Simoncelli, 2000). We opt in the current study to use their model as a basis for exploring the necessary and sufficient statistics for the successful synthesis of various kinds of photographic texture. There are several reasons for this choice. First, Portilla and Simoncelli's model produces very high-quality syntheses of textures. Second, synthesis can be achieved relatively quickly, meaning a library of synthesized textures can be created in a reasonable time frame. This is in contrast to the FRAME model of texture synthesis (Zhu, Wu, & Mumford, 1996; Zhu, Wu, & Mumford, 1997), which is very powerful, but slow. Finally, the implementation of the algorithm allows for "lesioning" of the code to remove certain parameters from the synthesis process. This last aspect of the model makes it particularly attractive for our purposes, as it allows us to synthesize textures lacking certain statistical constraints, and assess how well the final image approximates the target texture.

The Portilla-Simoncelli model utilizes four large sets of parameters to generate novel texture images from a specified target. In all cases, a random-noise image is altered such that its distributions of these parameters match those obtained from the target image. The first of these parameter sets is a series of 1st-order constraints (Marginals) on the pixel intensity distribution derived from the target texture. The mean luminance, variance, kurtosis and skew of the target are imposed on the new image, as well as the range of the pixel values. Second, the local autocorrelation of the target image's low-pass counterparts in the pyramid decomposition is measured (Coeff. Corr), and matched in the new image. Third, the conditional histograms of coefficient amplitude pairs at neighboring scales, spatial positions, and orientations are matched in the new image (Mag. Corr.). Finally, cross-scale phase statistics are matched between the old and new images (Phase). Portilla and Simoncelli report on the utility of each of these parameter subsets in their description of the model, but offer no clear perceptual evidence beyond the visual inspection of a few example images. The current study aims to carry out a true psychophysical assessment, in the hopes that doing so will more clearly demonstrate which statistics are perceptually important for representing natural textures.

We present the results of two experiments, designed to test the aforementioned parameter subsets value in producing textures that are indiscriminable from the target texture under pre-attentive conditions. We note that this is markedly different than analyzing the resulting images under full scrutiny, as the kinds of artifacts and errors that may seem glaring given an attentive analysis of an image may be invisible under pre-attentive conditions. Our strategy is to first produce synthetic textures that are not matched to the target texture for one or more of the parameter families previously mentioned. We then determine how discriminable patches drawn from these images are from patches drawn from the original texture under brief presentation. In so doing, we explicitly assume a local windowing model of texture processing similar to a recently proposal of Tyler's (Tyler, 2004b). We compare discriminability of "lesioned" textures to the discriminability of synthetic textures created using the full set of statistical parameters in the model. This allows us to determine how much each parameter subset contributes to the final synthesis. Further, we break down our target textures into two families (roughly "periodic" and "structured" textures) to see whether or not different statistics are needed to convincingly synthesize specific categories of images.

Methods

Subjects

A total of 16 subjects participated in the two experiments described here, eight in each of our two experiments. Subject age ranged from 19-27 years, and all subjects had normal or corrected-to-normal vision.

Stimuli

Original textures - 12 256x256 texture samples were chosen from a set of textures available via the NYU Laboratory for Computational Vision (<http://www.cns.nyu.edu/~eero/software.html>). Several textures are Brodatz images (Brodatz, 1996) while the remainder are original photographs collected by the NYU laboratory. The images were selected to conform to two pre-conceived visual categories, pseudoperiodic and structured textures. For our purposes, we will consider pseudoperiodic textures to be those textures with repeating middle to large-scale structures. Structured textures are defined as those textures composed of discrete elements that are not repeated in a predictable way across the image (Figure 1).

Nonperiodic structured textures



Pseudoperiodic textures




Figure 1 – The collection of textures used to create synthetic images for Experiments 1 and 2. The top row contains textures that are composed of repeated structural elements but lack strong periodicity or global structure. The bottom row contains textures that have strong periodicity.

"Lesioned" textures - Five synthetic versions of each original texture image were created using Portilla and Simoncelli's algorithm. The first four images were created by choosing to ignore one family of statistical measurements taken from the original image while performing the synthesis procedure. In order, marginal statistics, raw autocorrelation statistics, filter magnitudes, and cross-scale phase measurements were removed from consideration one at a time for each condition. The fifth category of synthesized textures was created by synthesizing each texture using the full set of statistical constraints.

Each synthesized image was 256x256 pixels in size, using parameters extracted from a 192x256 pixel patch taken from the original texture. The original textures were cropped to remove any artifacts introduced by the small text credits at the bottom of each image. Examples of the synthesized textures created from a particular target texture are displayed in Figure 2.




Figure 2 – “Lesioned” texture images created using the Portilla and Simoncelli algorithm to synthesize textures from our original images using either the full set of statistical parameters (far right) or using all but one subset of those parameters. From left to right, the images in this figure were constructed without explicit matching of first-order constraints (mean, range, variance, kurtosis and skew), subband coefficient correlation, subband magnitude correlation, and cross-scale local phase information.

"Pair-wise impoverished" textures - For Experiment 2, we create four new categories of texture images by synthesizing texture patterns using the marginal statistics alone, and also the marginal statistics plus each of the three remaining parameter subsets added in one at a time. While the images in Experiment 1 allow us to discuss the necessity of each subset of parameters for texture synthesis, these images are designed to give us insight as to the sufficiency of these subsets for successful texture reconstruction. The reason for using “pair-wise” images rather than synthesizing textures using each parameter subset in isolation is that in inspecting Figure 2, it is obvious that those images lacking the same first-order statistics as their parent textures are strikingly different from the other lesioned images. This is the case because encapsulated in those first-order measurements are highly salient global image properties like overall contrast and mean luminance of the image. From this, we expect that first-order properties will certainly prove to be necessary for good synthesis in Experiment 1. This means that testing sets of images that lack these properties will be relatively

useless. Instead, we include these parameters in all cases, allowing us to test the first-order properties themselves for sufficiency as well as the remaining parameter subsets (with the caveat that pixel distributions are always matched).




Figure 3 – “Pairwise” lesioned images created by including the 1st order statistics in all synthetic textures with the addition of: (from left to right) nothing additional, subband coefficient correlation, subband magnitude correlation, cross-scale phase information, and all parameters in the Portilla and Simoncelli algorithm.

Texture Quadrants – Finally, each texture was divided into mutually exclusive quarters 128x128 pixels in size and windowed by a circular mask to remove any orientation-specific interactions between the contours of the image frame and those contours contained within the texture itself (Figure 4).



Figure 4 – An original texture (left) and a windowed patch taken from the top left quadrant of the original image (right). This simple circular mask was meant to reduce any unwanted enhancement of horizontal and vertical edges brought on by the rectangular window of the original texture.

Procedure

Subjects were seated approximately 100 cm. from a 17" Dell Ultrasharp monitor. All stimulus display and response recording functions were controlled via the Matlab Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). In all three experiments, subjects were to perform a 3AFC "oddball" task, in which three unique texture patches were presented on each trial, one of which was a patch drawn from a synthesized version of the original texture from which the two distractor patches were drawn. Subjects were not familiarized with the textures previously, and all three texture patches in a given trial were distinct images. These measures were taken to ensure that neither high-level information nor pictorial matching strategies could contribute to subjects' performance.

On each trial, the three images were displayed at the vertices of an equilateral triangle such that the distance between each image and central fixation was approximately 3.5 degrees of visual angle. Each stimulus was approximately 2 degrees of visual angle in diameter, and the entire stimulus triad was onscreen for 250 ms before a response was collected. Breaks were scheduled every 240 trials. Subjects indicated the location of the oddball texture patch via the '1', '2', and '3' keys to indicate left, top, and right respectively. Both accuracy and response time were recorded on each trial.

All 6 pairs of texture quadrants were used as distractors twice for each texture in each condition, and two texture "oddball" patches were used 6 times each in each condition to ensure that overall frequency of synthesized and veridical textures remained balanced. In total, subjects completed 72 trials per "lesion" condition for each of our two texture families for a total of 720 trials.

Results

Experiment 1 – In our first experiment, we are looking for evidence that subsets of statistical constraints collected by the Portilla and Simoncelli algorithm are differentially important for the successful synthesis of our two texture families (periodic and structured). In particular, this experiment assesses the degree to which each subset of parameters is necessary for the synthesis of each type of texture by removing one set of constraints at a time.

A 2-way ANOVA (with repeated measures) was run on both the accuracy and reaction time data. The accuracy data revealed a highly significant effect of texture “lesioning” ($p < 0.0001$) as well as a highly significant interaction between texture category and lesion ($p < 0.0001$). There was no main effect of texture category ($p > 0.15$). The RT data yields little in the way of interesting results, save for a main effect of texture lesioning ($p < 0.001$) which indicates that all subjects were fastest to respond to images lacking first-order statistics.

In Figure 5, we see that as we predicted the first-order statistics of our texture distributions are clearly necessary for successful synthesis. Subjects are at ceiling at detecting the “oddball” texture when these constraints are removed. Further, the interaction between lesion and texture category appears to be driven by the differential importance of raw coefficient correlation and magnitude correlation for our two families of textures. To be more specific, pseudoperiodic textures seem to rely relatively equally (and weakly) on both of these sets of parameters, given that the removal of each results in a relatively low level of correct detections. In contrast, the magnitude correlation statistics are clearly quite necessary for successful synthesis of structured textures, while the coefficient correlations seem to contribute almost nothing to the full synthesis. We note that in neither case do the constraints on cross-scale phase contribute substantially to the performance of subjects on this task, indicating that under pre-attentive conditions these constraints matter very little.

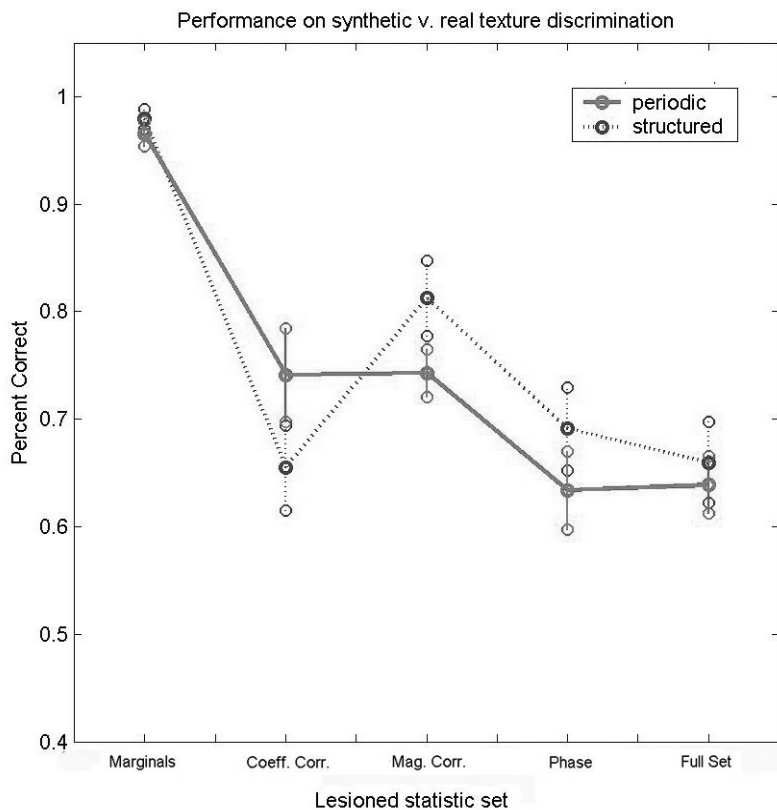


Figure 5 – Plot of the average performance on the oddball detection task as a function of both texture category and texture lesion (mean values \pm standard error). Note both the clear importance of first-order statistics at left, as well as the interaction between the necessity of coefficient and magnitude correlation for periodic and structured textures.

To confirm this assessment of the results, we conducted post-hoc Tukey-Kramer tests within each texture category between each of the 4 “lesion” conditions and the “Full Set” condition. We find that for pseudoperiodic textures, only the removal of the first-order statistics produces a rate of oddball detection significantly greater than the “Full Set” images ($p < 0.05$). However, for the structured textures, we find that both the removal of the first-order statistics and the removal of the magnitude correlation statistics produce rates of oddball detection significantly greater than that of the “Full Set” textures ($p < 0.05$).

Experiment 2 – In this second experiment, we are testing the sufficiency of both 1st order statistics in isolation and pairwise combinations of 1st-order information and the remaining three parameter subsets for producing successful synthetic texture images. In these results we will be looking for cases where the inclusion of parameter subsets gives

rise to low rates of oddball detection. This will indicate that the subsets included may be sufficient for producing synthetic textures viewed under pre-attentive conditions.

As in Experiment 1, we ran a 2-way ANOVA with repeated measures on the accuracy and RT data, with inclusion condition and texture family as factors. As before, we find no effect of texture family ($p = 0.10$) but significant effects of inclusion condition ($p < 0.00001$) and a significant interaction between our two factors ($p < 0.0001$). The RT data again provided nothing more than a main effect of included statistics, with faster RTs in the “Marginals only” condition.

We note in Figure 6 that the inclusion of 1st order statistical constraints alone results in a rate of oddball detection that is at ceiling. This indicates that though these parameters are certainly necessary for synthesis, they are certainly not sufficient. Of interest however, is the relationship between the other three parameter subsets. Specifically, we notice that for structured textures magnitude correlation proves to be quite useful for synthesis, producing rates of oddball detection comparable to the “Full Set” images. For pseudoperiodic textures, the rate of oddball detection is higher, indicating that these measurements are less useful for producing successful textures within this family of images. Raw coefficient correlation produces comparable rates of detection for both texture families, while the inclusion of phase constraints produces quite high rates. For structured textures, this rate is somewhat lower than ceiling (perhaps an indication that cross-scale phase provides some small amount of useful information), but overall demonstrates the insufficiency of phase information for texture synthesis.

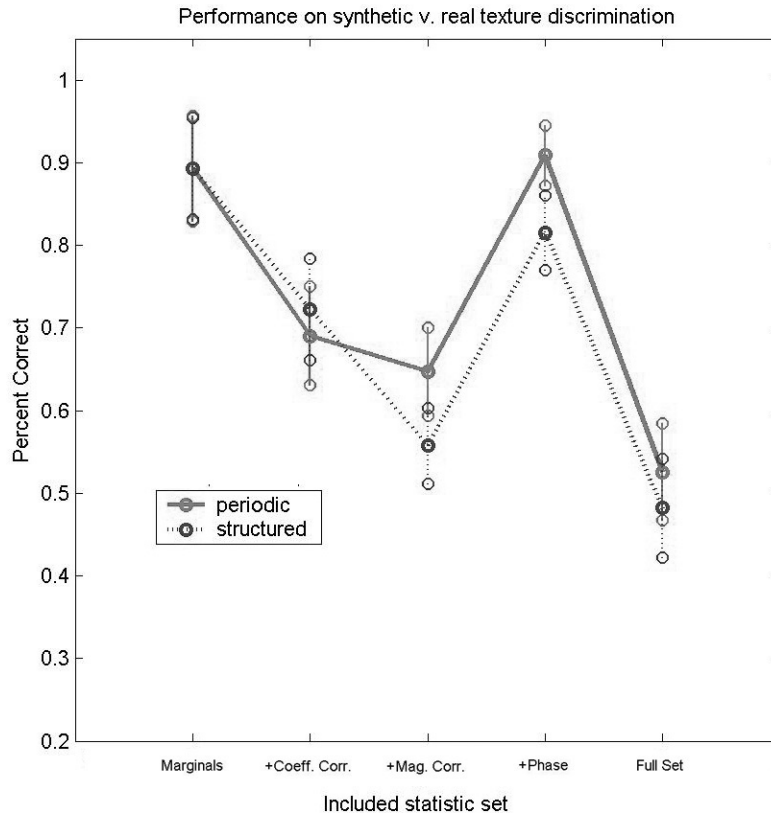


Figure 6 – Rates of oddball detection for both texture families as a function of included statistics. Marginals alone, and the pair-wise inclusion of marginals and cross-scale phase provide poor syntheses. Magnitude correlations and marginals together result in good synthesis of structured textures. For pseudoperiodic textures, the inclusion of both raw coefficient correlations and magnitude correlations result in intermediate quality syntheses, indicating weak sufficiency.

As before, post-hoc Tukey-Kramer tests were run to confirm our intuitions regarding the interaction of inclusion condition and texture family. In comparing the “Full Set” responses to the other conditions within texture families, we find that for structured textures all conditions differ significantly ($p < 0.05$) from the “Full Set” rate of oddball detection, with the sole exception being subband magnitude correlation. For pseudoperiodic textures, we find that only the 1st order-only condition and the cross-scale phase condition differ significantly ($p < 0.05$) from the “Full Set.” Both raw coefficient correlation and magnitude correlation could be considered weakly sufficient given this analysis, although we note that the size of these differences are large enough to warrant skepticism. It is our belief that

the sufficiency of 1st-order statistics and magnitude correlation for structured textures is more clearly indicated by this experiment.

Discussion

We have found in our pre-attentive discrimination task that different sets of statistics are important for the synthesis of pseudoperiodic and structured texture images. In terms of the necessity of each subset of parameters, we find that 1st order pixel statistics such as the mean, variance, and range of luminance values are vitally important for creating perceptually matched textures from a target image. This is hardly surprising given how easily human observers can discriminate between different brightness and contrast levels. Of more interest is the reliance of each texture family on autocorrelation and filter magnitude statistics. Periodic textures demonstrate a weak need for both of these measures, although neither set of statistics alone proved absolutely necessary for the synthesis of these images. Structured textures, by comparison, appeared to rely quite heavily on the magnitude correlation statistics, while demonstrating no need for preservation of the local autocorrelation statistics. Neither texture appeared to rely on cross-scale phase statistics for synthesis, suggesting that these measurements may only be important for texture images that undergo scrutiny.

In terms of the sufficiency of our parameter subsets, we find that preserving only 1st-order measurements of the pixel distribution is clearly not enough to create a convincing synthetic image. Again, this is not surprising given that the human visual system is known to have strong representations of higher-order features (like edges) that will not be preserved through balancing only pixel-based statistics. Also, as expected from the results of Experiment 1, cross-scale phase statistics combined with proper 1st-order measurements result in extremely poor syntheses. Again, it is the imposition of the autocorrelation and coefficient magnitude constraints that prove most useful in this task. Mirroring the data from Experiment 1, we find that pseudoperiodic textures of intermediate quality can be produced by including the autocorrelation constraints or magnitude constraints alone (so long as 1st-order properties are preserved). Structured textures also conform to our expectations given the data from Experiment 1, as the magnitude constraints coupled with 1st-order properties result in synthetic images that are of good quality compared to the images produced with the full set of constraints.

The results of these two experiments align with some previously reported results concerning periodic textures, especially with regard to the role of the autocorrelation function in representing such textures (Fujii, Sugi, & Ando, 2003). The necessity (and insufficiency) of 1st-order image properties is also not a new or surprising contribution. However, what we see as the primary point of interest presented by this work is the perceptual role of cross-scale phase statistics and the magnitude correlations introduced by Portilla and Simoncelli. In the first case, we point out that neither the inclusion or absence of cross-scale phase information affected the synthesis process in any way that indicated this information was of perceptual use under pre-attentive viewing. This is sensible given the parafoveal viewing conditions of the task and the brief presentation time, but still implies that such measurements may not be used to characterize textures at a low level. In the second case, we note that the magnitude correlations imposed on synthetic images are shown by this study to be at least as important as the autocorrelation function for the perceptual similarity of periodic textures under pre-attentive conditions. This is somewhat surprising, as for highly repetitive textures one might have suspected the raw coefficient correlations would capture the majority of important image structures. More strikingly, we note that for structured textures these statistics are extremely important for perceptual similarity. Indeed, through matching only these parameters and 1st-order properties one can create synthetic images that are not of significantly lower quality than those made with the entire set of constraints.

At this point, it is useful to reflect upon whether or not these findings are a result of the specific set of textures chosen for this task. While individual textures can behave idiosyncratically in the synthesis process, we find qualitatively similar results across the individual textures used here and across different subjects. The only exception to this rule are two synthetic textures that almost all subjects found highly discriminable from the target images regardless of what parameter subsets were included: typed text and the coarse stone tiling displayed in the upper left of Figure 1. Given that the deviation in these two cases from the average performance was an overall ceiling effect, we are not concerned about the effects on the group data (especially given the fact that these two textures came from different texture families). We take this as good evidence that the two families of textures we have used are legitimately different in the context of this task. More generally, it may be that different models of synthesis (and therefore perceptual analysis) could be more sensible than a unified model for all texture families.

An important final caveat however concerns the discriminability of the “Full Set” images from the target textures. In our 3AFC task, chance performance was 33%, a rate of oddball detection lower than that displayed by all but a few of our subjects. Overall, this indicates that even in the most difficult condition our synthetic textures were still reliably discriminable from their respective targets. In all cases, we are only able to consider the necessity and sufficiency of the parameters included relative to this baseline. We do not see this as especially problematic either, but

it does indicate that there is still a fair amount of work to be done as far as creating more powerful texture synthesis algorithms. We are limited to testing the statistical constraints imposed by this particular model, and though they seem both reasonable and useful we must remember that there remains an infinite number of image statistics that may prove better in the future. A simple extension of this work would be the inclusion of color images, as the results from the Portilla and Simoncelli algorithm when applied to full-color textures are extremely impressive. The added salience of color information may push discriminability down closer to chance performance, making claims about absolute necessity and sufficiency more reasonable. Creating these images requires a somewhat more complicated synthesis process (Liang, Simoncelli, & Lei, 2000), but would also get us closer to our stated goal of examining the perception of “real world” textures.

Conclusions

We have used a parametric model of texture synthesis as a tool for examining the necessity and sufficiency of different statistical measures for the perceptual similarity of texture images. We have found that different requirements apply for periodic textures as opposed to structured textures, notably in the need for autocorrelation measurements and conditional histograms of edge-like filter magnitudes. Cross-scale phase statistics were found to be of little use under pre-attentive conditions, while 1st-order pixel properties were demonstrated to be vital for capturing global image similarity. These results demonstrate the value of using computational models for texture synthesis to address perceptual questions regarding texture processing. It is hoped that this may help to bridge the gap between the communities of graphics, machine vision, and psychophysical texture research. Moreover, the 3AFC task presented here represents a modest contribution towards the formulation of texture discrimination tasks that make explicit the importance of local texture analysis in the human visual system.

Acknowledgements

The author would like to thank Bernd Heisele, Yuri Ivanov, and Pawan Sinha for their support of this work. Thanks also to Erin Conwell and David Cox, who provided helpful comments regarding earlier versions of this manuscript.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychol. Rev.*, *61*, 183-193.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. Rosenblith (Ed.), *Sensory Communication*. Cambridge, MA: MIT Press.
- Bergen, J., & Adelson, E. (1988). Early vision and texture perception. *Nature*, *333*, 363-364.
- Bergen, J. R., & Adelson, E. H. (1986). Visual texture segmentation based on energy measures. *J. Opt. Soc. Am. A*, *3*, 99.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433-436.
- Brodatz, P. (1996). *Textures: A Photographic Album for Artists and Designers*. New York: Dover.
- Efros, A. A., & Freeman, W. T. (2001). *Image Quilting for Texture Synthesis and Transfer*. Paper presented at the SIGGRAPH '01, Los Angeles, CA.
- Efros, A. A., & Leung, T. K. (1999). *Texture synthesis by non-parametric sampling*. Paper presented at the Proc. Int'l Conference on Computer Vision, Corfu.
- Fujii, K., Sugi, S., & Ando, Y. (2003). Textural properties corresponding to visual perception based on the correlation mechanism in the visual system. *Psychological Research*, *67*, 197-208.
- Heeger, D., & Bergen, J. (1995). *Pyramid-based texture analysis/synthesis*. Paper presented at the ACM SIGGRAPH.
- Julesz, B. (1962). Visual Pattern Discrimination. *IRE Trans. Info. Theory*, *IT*, *8*, 84-92.
- Julesz, B. (1981). Textons, the elements of texture perception and their interactions. *Nature*, *290*, 91-97.
- Julesz, B., Gilbert, E. N., Shepp, L. A., & Frisch, H. L. (1973). Inability of humans to discriminate between visual textures that agree in second-order statistics-revisited. *Perception*, *2*(391-405).
- Julesz, B., Gilbert, E. N., & Victor, J. D. (1978). Visual discrimination of textures with identical third-order statistics. *Biological Cybernetics*, *31*, 137-140.
- Kersten, D. (1987). Predicability and redundancy of natural images. *Journal of the Optical Society of America A*, *4*, 2395-2400.
- 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.
- Liang, Y., Simoncelli, E. P., & Lei, Z. (2000). *Color Channels Decorrelation By ICA Transformation in the Wavelet Domain for Color Texture Analysis and Synthesis*. Paper presented at the 2000 Conference on Computer Vision and Pattern Recognition, Hilton Head, SC.

- Malik, J., Belongie, S., Leung, T., & Shi, J. (2001). Contour and Texture Analysis for Image Segmentation. *International Journal of Computer Vision*, 43(1), 7-27.
- Malik, J., & Perona, J. (1990). Preattentive texture discrimination with early vision mechanisms. *J. Opt. Soc. Am. A*, 7, 923-932.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Portilla, J., & Simoncelli, E. (1999). *Texture modeling and synthesis using joint statistics of complex wavelet coefficients*. Paper presented at the IEEE Workshop on Statistical and Computational Theories of Vision, Fort Collins, CO.
- Portilla, J., & Simoncelli, E. (2000). A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. *International Journal of Computer Vision*, 40(1), 49-71.
- Rao, A. R., & Lohse, G. L. (1996). Towards a texture naming system: Identifying relevant dimensions of texture. *Vision Research*, 36, 1649-1669.
- Simoncelli, E., & Portilla, J. (1998). *Texture characterization via joint statistics of wavelet coefficient magnitudes*. Paper presented at the Fifth IEEE Int'l Conf. on Image Proc., Chicago.
- Simoncelli, E. P. (1997). *Statistical models for images: Compression, restoration and synthesis*. Paper presented at the 31st Asilomar Conf. on Signals, Systems and Computers, Pacific Grove, CA.
- Simoncelli, E. P., & Freeman, W. T. (1995). *The steerable pyramid: A flexible architecture for multi-scale derivative computation*. Paper presented at the Second Int'l Conf. on Image Proc., Washington, D.C.
- Tyler, C. W. (2004a). Beyond fourth-order texture discrimination: generation of extreme-order and statistically-balanced textures. *Vision Research*, 44, 2187-2199.
- Tyler, C. W. (2004b). Theory of texture discrimination based on higher-order perturbations in individual texture samples. *Vision Research*, 44(2179-2186).
- Zhu, S. C., Wu, Y. N., & Mumford, D. (1996). *Filters, random fields and maximum entropy (FRAME) - Towards the unified theory for texture modeling*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition.
- Zhu, S. C., Wu, Y. N., & Mumford, D. (1997). Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9, 1627-1660.