

Visualization Viewpoints

Editors: Theresa-Marie Rhyne and
Lloyd Treinish

Harnessing Natural Textures for Multivariate Visualization

Victoria
Interrante

University of
Minnesota

In our ongoing quest to convey more information more clearly in a single image, harnessing the full potential of texture for data representation remains an elusive goal. Others have begun excellent work in this area,¹⁻³ and my efforts are inspired by their example. The grail that I seek is a partially ordered multidimensional palette of richly detailed and varying texture patterns that can be used—in conjunction with lightness and hue—to represent multivariate information. The goal is to facilitate the flexible visual appreciation of the correlations of various quantities across the different dimensions. The approach that I outline here departs a bit from the norm, but is motivated by a desire to proceed more directly from my vision of what I want to achieve, unrestrained by the limitations of the tools I have on hand. In the following discussion, I motivate the adoption of rich, natural textures—resembling those from photographic images⁴—as elemental primitives and sketch some of the approaches that we can take to enhance our understanding of how to effectively harness their properties. My intent here is not to present results, but to expound on the issues and conclude with the questions to which we’re still seeking answers.

Why natural textures?

The intricate variety and subtle richness of detail of texture patterns found in nature support possibilities for data representation far more vast and comprehensive than we could ordinarily hope to achieve from standard primitives. Even if we must ultimately rely on synthesized textures for data visualization, by looking to nature for inspiration we have the potential to expand our vision of what to strive for in such a synthesis. The graphic design community has long held that perfectly regular synthetic textures on a flat plane, in particular the infamous hatching patterns that Edward Tufte refers to as “chart junk,”⁵ are discomforting to the eye and annoying to look at. Natural textures are not only more aesthetic, but they also put less extraneous stress on the visual system, leaving our eyes freer to observe and attend to the most intrinsically important texture-pattern characteristics.

Understanding human texture perception

To create a perceptually meaningful multidimensional *texture space* that can be indexed in the same fash-

ion as a color space, we must begin by knowing what we’re looking for. We need to proceed from a rigorous and experimentally supported understanding of how human observers perceive and interpret texture patterns, under the conditions in which we intend for these patterns to ultimately be viewed. This grounding provides a structure for guiding our search through the complex space of possibilities and formalizes the intuition that a good designer calls upon to create a visualization that works.

A number of researchers⁶ have conducted studies to try to elucidate the most significant perceptual dimensions of texture. The results of these experiments will aid us, though some important questions remain. It’s beyond the scope of this article to summarize previous findings further than to say that most of the studies used unaltered images from the Brodatz album, subjects were generally asked to cluster the textures into groups, and there appears to be general agreement that a small number (about three) of characteristic dimensions seem sufficient to describe most of the structure underlying this classification. The interpretation of the dimensions varies from study to study, but most often includes aspects of the following:

- periodic (consisting of repeated discrete elements) <-> nonperiodic
- strongly directional <-> rotationally invariant
- coarse <-> fine (spatial frequency of the dominant detail)
- regular (deterministic) <-> random
- high contrast <-> low contrast
- homogeneous (spatially invariant) <-> heterogeneous

Clearly, there is some overlap in these categorizations. Also, it’s not evident that we can hope to determine an orthogonal basis that encompasses all members of the texture pattern set. However, the apparent low perceptual dimensionality of the space and the strong agreement between the studies bodes well for our application.

Open questions in visual texture perception

Identifying the features according to which people tend to classify texture patterns gives us important insight into how to structure a perceptually meaningful texture space.


But some important issues have been implicitly overlooked in the studies conducted so far. Foremost are the uncontrolled-for influences of higher level processes.

Payne et al.⁷ observed, in studies similar to those referenced above, that material property categories appeared to have a strong influence on a fair number of their observers' clustering decisions. They also noted that their subjects often commented that they felt they were using different criteria to find matches for different texture images. Knowing that people tend to make judgments based on different criteria in different cases, but not knowing who is considering what when or why, weakens the general applicability of experimental results based on these traditional methodologies. I believe that it might be useful to attempt to control for, or at least quantify the impact of, some of these effects by considering alternative paradigms for objectively measuring texture patterns' perceptual similarity.

The question of whether to control for rotation, scale, luminance, and contrast variance among texture samples when seeking insight into the perceptual groupings of texture images is a second important issue that studies using the Brodatz album have frequently overlooked. How do we want to consider apparent texture pattern differences that aren't clearly intrinsic to the pictured material but that could conceivably be attributed to external factors such as viewpoint or lighting?⁸

On the one hand, Ware and Knight³ have said that orientation, size, and contrast are the primary orderable dimensions of texture. But samples of a texture that are differently oriented, scaled, and lit still intrinsically appear to be the same thing and thus remain good candidates for similarity grouping. As pre-attentive features of individual elements, size, contrast, and orientation differences are undisputedly important in facilitating "pop-out". At the same time, our visual system is remarkably adept at maintaining perceptual constancy across changes in illumination or viewpoint. Many computational methods for classifying texture adopt a rotationally invariant texture recognition approach for this reason. A strong argument exists for equalizing characteristics such as scale, luminance, and contrast before classifying or quantifying texture-pattern differences for the purposes of visualizing multiple distributions across a 2D image. Doing so would let us retain the ability to introduce variations in these feature dimensions universally across all the other texture dimensions and use them to encode additional values. To this end, it might be illuminating to try to separately examine the relative effects of rotation, scale, and contrast differences versus other texture-characterizing differences. Our considerations will also differ in the case of visualizing distributions across surfaces through a 3D domain.

The question of how best to factor out the variations in contrast and luminance—when we choose to do so—must also be carefully considered. Although it's relatively straightforward to equalize the intensity histograms of sample texture images before processing them for similarity, it's not clear that histogram equalization adequately preserves the meaningful qualities of the texture patterns.



1 The salience of differences in various properties of texture patterns might be estimated through psychophysical experiments using briefly presented stimuli such as these.

Quantifying the perceptual similarity of texture patterns



In addition to determining which textures tend to cluster, it's important for creating a perceptually linear texture space to quantify the perceptual distances between individual texture patterns. It's also necessary to estimate the magnitude of the perceived distance due to the differences along each of the feature dimensions.

One possible approach is to estimate the magnitude of the change required to enable a "just noticeable difference" between images along individually selected texture dimensions such as scale, contrast, orientation, regularity, and so on using psychophysical methods.

Another possibility, which may be more appropriate for judging the kinds of differences that cannot be easily brought down to threshold levels, is to measure the pre-attentive discriminability or salience of differences in features of individual texture patterns randomly embedded in homogeneous and heterogeneous fields of distracters, as illustrated in Figure 1. The objective in this case then is to determine how large of a difference is required to allow the effortless identification of the "odd man out" in brief, masked stimulus presentations. Studies using individual element arrays have found that salience (or the tendency to "pop-out") tends to increase when the targets are characterized by redundant, unique properties such as luminance and hue or color and orientation.⁹ Similarly, the salience of the target tends to decrease as the heterogeneity of the distracter elements increases, even when the heterogeneity occurs along a different perceptual dimension.


It may additionally be of interest to determine how many different texture types people can simultaneously discriminate, using a methodology similar to what Healey employed for studying color.¹⁰

The long and rich history of research on texture classification algorithms in the image processing community also offers valuable resources for constructing a texture palette. Although the extent and variety of the possible computational approaches for classifying texture patterns is somewhat overwhelming, most of the successful methods work by extracting a finite set of features from the texture patterns (via transforms similar to those shown in Figure 2) and then calculating various statistics across these feature sets. Rubner and



2 (a) Statistics of texture pattern features, such as those evident in this set of power spectra images, form the basis for many computational texture characterization methods. (b) It's easy when looking at the original images to form an intuitive understanding of the types of information carried by the various features of the Fourier coefficients.

3 A small potential texture palette. Scale increases along the horizontal axis, regularity increases along the vertical axis, and intensity increases along the left-to-right descending diagonal.



Tomasi⁸ proposed quantifying texture similarity by using the Earth Mover's Distance as a metric for the goodness of fit between histograms of these features. A key issue for us is ensuring that the computer vision results and the human observer criteria agree.


Creating a texture palette

One approach to creating a texture palette (Figure 3) is to begin with a collection of well-chosen input images. You can then objectively determine where they lie in the best-fitting multidimensional texture feature space that they span and appropriately fill in the remaining open space with intermediate textures that lie at equal perceptual distances along each of the dimensions.

Some of the difficulties in creating a texture palette are that the texture space may not be orthogonal, there may be interaction among certain dimensions (such as contrast and spatial frequency), and some texture type mixtures may not be meaningful. However, my intuition is that the closer we can get to aesthetically filling out the space with real acquired images, the easier it will be to patch the holes.

Texture synthesis methods such as those that Portilla and Simoncelli¹¹ and Zhu et al.¹² have proposed may hold the greatest potential for creating a set of sample textures that fill out a multidimensional palette. Working from these frameworks, we may have the possibility to deterministically create natural-looking intermediate textures that interpolate the characteristic properties of their neighboring swatches in texture space or that

4 A set of four agricultural maps, representing distributions of four different variables across the US counties.



conform to a prespecified set of desired criteria. However, these methods currently invoke too many features (roughly 800 plus) and are not quite mature enough to guarantee the realism that we seek.

Efros and Leung¹³ suggest a very different synthesis method for seamlessly generating highly realistic samples of “more of the same” textures from a given sample. The success of this method is subject to the assumption of stationarity and a reasonable estimate of the extent of the lowest spatial frequency detail that must be preserved. However the feasibility of extending this approach to texture interpolation has yet to be shown.

Another issue we face is how to properly deal with the lighting consistency problems that will inevitably arise when we try to combine acquired textures that exhibit relief. It may be necessary to solve for the surface relief, allowing the material intensity texture to be handled separately. In the 2D case, the most important consideration is simply to maintain consistency. Additionally, it’s desirable to avoid orienting textures so that to preserve convexity, observers must envision the light as coming from below.

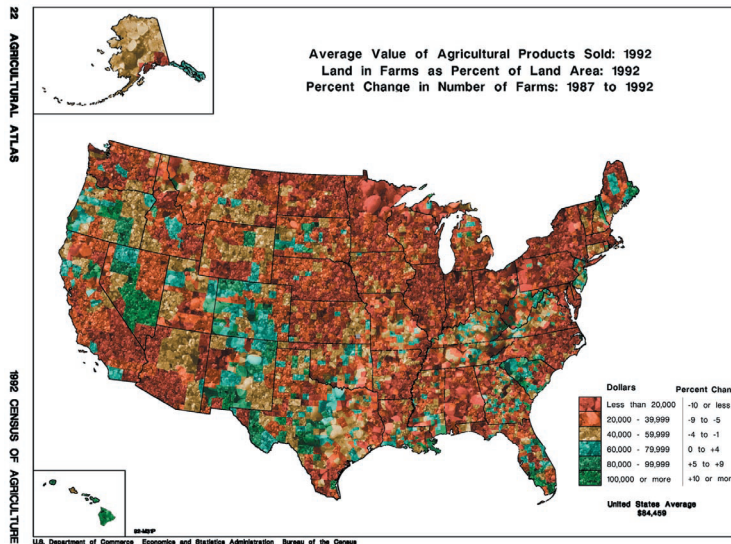


Feasibility issues

Some of the questions that we need to address are

- What does a reasonable partitioning of a natural texture space look like?
- Would it be feasible to try to choose exemplars at the endpoints of each perceptually relevant texture dimension, characterize them statistically, then interpolate to obtain intermediate textures that fill out the space?
- To what extent do we need to guarantee that different textures will meld continuously into each other at the transitions between level set regions?
- How can we most effectively combine color with texture to convey yet more information in a meaningful way?

Clearly, texture has the greatest potential to be effective as a tool for visual data representation when it conveys local values of an underlying function across homogeneous cells of sufficient size to allow the characteristic detail of the resident texture to be discriminated.

The philosophy behind my attempt to more effectively harness the potential of texture for multivariate data visualization is that we should begin from a vision of what we want to achieve and work from there to figure out how



5 A simple hand-crafted demonstration of a display in which three natural textures of increasing scale convey three discrete ranges in the value of a variable along a single

6 A hand-crafted example in which one dimension of color and two dimensions of texture overlap to represent three different variables in each cell.

7 A second hand-crafted example in which one dimension of color and two dimensions of texture represent three different variables in each cell.

to accomplish the desired results. To more freely explore the possibilities before becoming bogged down in implementation issues, I mocked up several tests by hand using Adobe Photoshop and scanned images from the Brodatz texture album⁴ to represent multivariate agricultural data¹⁴ at the county level (Figures 4 through 7).

In Figure 7, texture scale represents one of three

ranges of the average value of the agricultural products produced in each county (with the largest scale corresponding to the lowest value, evoking the metaphor of inhospitable terrain). The texture type represents the direction of change in the amount of land in farms (with rocks representing areas with an overall loss of farmland and weaves representing areas with an overall increase in the amount of land used for farming). Color labels the percentage of the total land area in each county used for farming (with the green tones indicating the higher percentages and the brown tones representing the lower percentages of land occupied by farms).

Representing uncertainty

Both color and texture admit intriguing possibilities for representing uncertainty in data measurements. Texture regularity has particularly good potential as an intuitive marker for certainty, with texture pattern irregularities increasing in prominence where measurement reliability is lower. Color regularity has somewhat weaker potential as a marker for measurement certainty, as textures that contain more balanced or restricted distributions of hues may appear less distinct or vibrant (implying decreased noteworthiness) than textures in which the hues vary more widely across the spectrum. Figures 8 and 9 provide examples of these two different configuration series. It remains to be seen how easily irregularity of the kind shown in Figure 9 can be incorporated into a computational texture synthesis definition for such patterns.

Conclusions

Despite the excellent progress made in recent years,¹⁻³ I believe that there remains great untapped potential for the effective use of texture in multivariate visualization. I have proposed that we might take important steps towards realizing more of this potential by attempting to harness the power of rich natural textures. I envision that a successful approach will begin from a

fundamental understanding of visual texture perception and progress toward an understanding of how to synthesize a multidimensional palette of detailed texture samples whose variations evoke an intrinsic appreciation of the local and global relationships between multiple quantities across a 2D domain. ■


Acknowledgments

In framing this work, I have benefited from informal discussions with a number of different people. I am particularly grateful to Tomás Filsinger for providing valuable criticism and insights into understanding effective uses of texture from an artist/graphic designer's perspective, and to Guillermo Sapiro and Dan Kersten for pointers into the computer vision and texture synthesis literature. My fledgling efforts are partially supported by a 1999-2000 University of Minnesota Grant-in-Aid of Research, Artistry, and Scholarship and an NSF Presidential Early Career Award for Scientists and Engineers (Pecase) award.


References

1. C.G. Healey and J.T. Enns, "Large Datasets at a Glance: Combining Textures and Colors in Scientific Visualization," *IEEE Trans. Visualization and Computer Graphics*, Vol. 5, No. 2, 1999, pp. 145-167.
2. D.H. Laidlaw et al., "Visualizing Diffusion Tensor Images of the Mouse Spinal Cord," *Proc. IEEE Visualization 98*, ACM Press, New York, 1998, pp. 127-134.
3. C. Ware and W. Knight, "Using Visual Texture for Information Display," *ACM Trans. Graphics*, Vol. 14, No. 1, 1995, pp. 3-20.
4. P. Brodatz, *Textures: A Photographic Album for Artists and Designers*, Dover Publications, New York, 1966.
5. E.R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, Conn., 1983, p. 111.
6. A.R. Rao and G.L. Lohse, "Identifying High Level Features of Texture Perception," *CVGIP: Graphical Models and Image Processing*, Vol. 55, No. 3, 1993, pp. 218-233.

8 Variability in hue might be used to convey the range of imprecision or uncertainty in the data value recorded for a region.



9 Increasing irregularity in a texture pattern can intuitively represent increasing amounts of uncertainty.



7. J.S. Payne, L. Hepplewhite, and T.J. Stonham, *Texture Similarity: Using Human Studies to Interpret Retrieval Results*, Tech. Report TR-685, Dept. Electronic and Computer Engineering, Brunel University, UK, 1997.
8. Y. Rubner and C. Tomasi, "Texture Metrics," *Proc. IEEE Int'l Conf. Systems, Man, and Cybernetics*, IEEE Press, Piscataway, N.J., 1998, pp. 4601-4607.
9. H-C. Nothdurft "Salience from Feature Contrast: Additivity Across Dimensions," *Vision Research*, Vol. 40, 2000, pp. 1183-1201.
10. C.G. Healey, "Choosing Effective Colors for Data Visualization," *Proc. IEEE Visualization 96*, ACM Press, New York, 1996, pp. 263-270.
11. J. Portilla and E.P. Simoncelli, "A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients," *Int'l J. of Computer Vision*, 2000, to appear; currently available on the Web at <http://www.cns.nyu.edu/ftp/eero/portilla99.pdf>.
12. S.-C. Zhu et al., "Exploring Texture Ensembles by Efficient Markov Chain Monte Carlo—Toward a 'Trichromacy' Theory of Texture," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 22, No. 6, 2000, pp. 554-569.
13. A.A. Efros and T.K. Leung, "Texture Synthesis by Nonparametric Sampling," *IEEE Int'l Conf. Computer Vision, Volume II*, IEEE Press, Piscataway, N.J., 1999, pp. 1033-1038.
14. *1992 Agricultural Atlas of the United States*, U.S. Census Bureau Geography Div. Map Gallery, <http://www.census.gov/geo/www/map-gallery>.

Readers may contact Interrante at the Dept. of Computer Science and Engineering, University of Minnesota, 4-192 EE/CS Bldg., 200 Union St. SE, Minneapolis, MN 55455, e-mail interrante@cs.umn.edu.

Readers may contact department editors Rhyne and Treinish by e-mail at rhyne.theresa@epamail.epa.gov and lloyd@us.ibm.com

Correction

In the September/October issue, the Visualization Viewpoints article "Visualizing Visualizations: User Interfaces for Managing and Exploring Scientific Visualization Data" by Kwan-Liu Ma incorrectly implied that Fritz Hasler (NASA Goddard) alone worked on the Distributed Image Spreadsheet (p. 17). In fact, the work was done by both Hasler and Kannappan Palaniappan (University of Missouri-Columbia).

your Vision

our Future

Southwest Research Institute™

has exciting opportunities for

Computer Scientists & Engineers

- **Modeling and Simulation**
- **Software Engineering & Design**

Positions require a BS or MS in Computer Science, Computer Engineering, Electrical Engineering, or equivalent; programming experience in C, C++, and Java on PC- or UNIX-based workstations.

Responsibilities include software prototype development, design, testing, implementation, integration, maintenance, program mgt., and traveling to client sites. **Opportunities available in San Antonio, TX; Huntsville, AL; Warner Robins, GA; and Oklahoma City, OK.**

Please visit:

www.swri.org/jobs/

*An Equal Opportunity/Affirmative Action Employer
A Public and Community Services Employer*

- **Competitive salaries**
- **Medical, dental, & vision care**
- **Retirement & benefits**
- **Educational advancement**
- **Professional development**
- **Awards for individual research**
- **Paid relocation**
- **Employee assistance programs**
- **On-site fitness facility**
- **Generous vacation**
- **Ten paid holidays**



**Southwest
Research
Institute**