Multiresolution Sampling Procedure for Analysis and Synthesis of Texture Images

Jeremy S. De Bonet *
Learning & Vision Group
Artificial Intelligence Laboratory
Massachusetts Institute of Technology

EMAIL: jsd@ai.mit.edu
HOMEPAGE: http://www.ai.mit.edu/∼jsd

Abstract

This paper outlines a technique for treating input texture images as probability density estimators from which new textures, with similar appearance and structural properties, can be sampled. In a two-phase process, the input texture is first analyzed by measuring the joint occurrence of texture discrimination features at multiple resolutions. In the second phase, a new texture is synthesized by sampling successive spatial frequency bands from the input texture, conditioned on the similar joint occurrence of features at lower spatial frequencies. Textures synthesized with this method more successfully capture the characteristics of input textures than do previous techniques.

1 Introduction

Synthetic texture generation has been an increasingly active research area in computer graphics. The primary approach has been to develop specialized procedural models which emulate the generative process of the texture they are trying to mimic. For example, models based on reaction-diffusion interactions have been developed to simulate seashells [15] or animal skins [14]. More recently work has been done which considers textures as samples from probabilistic distributions. By determining the form of these distributions and sampling from them, new textures that are similar to the originals can, in principle, be generated. The success of these methods is dependent upon the structure of the probability density estimator used in the sampling procedure. Recently several attempts at developing such estimators have been successful in limited domains. Most notably Heeger and Bergen [10] iteratively resample random noise to coerce it into having particular multiresolution oriented energy histograms. Using a similar distribution, and a more rigorous resampling method Zhu and Mumford [16] have also achieved some success. In work by Luettgen, *et al* [12] multiresolution Markov random fields are used to model relationships between spatial frequencies within texture images.

In human visual psychophysics research, the focus of texture perception studies has been on developing physiologically plausible models of texture discrimination. These models involve determining to which measurements of textural variations humans are most sensitive. Typically based on the responses of oriented filter banks, such models are capable of detecting variations across some patches perceived by humans to be different textures ([1, 2, 3, 4, 6, 9, 11], for example.) The approach presented here uses these resulting psychophysical models to provide constraints on a statistical sampling procedure.

In a two-phase process, the input texture is first analyzed by computing the joint occurrence, across multiple resolutions, of several of the features used in psychophysical models. In the second phase, a new texture is synthesized by sampling successive spatial frequency bands from the input texture, conditioned on the similar joint occurrence of features at all lower spatial frequencies.

The sampling methodology is based on the hypothesis that texture images differ from typical images in that that there are regions within the image which, to some set of feature detectors, are less discriminable at certain resolutions than at others. By rearranging textural components at locations and resolutions where the discriminability is below threshold, new texture samples are generated which have similar visual characteristics.

2 Motivation

The goal of probabilistic texture synthesis can be stated as follows: to generate a new image, from an example texture, such that the new image is sufficiently different from the original yet still appears as though it was generated by the same underlying stochastic process as was the original texture.

If successful, the new image will differ from the original, yet have perceptually identical texture characteristics. This can be measured psychophysically in texture discrimination tests. To satisfy both criteria, a synthesized image should differ from the original in the same way as the original differs from itself.

From an input texture patch, such as that shown in Figure 1, there are infinitely many possible distributions which could be inferred as the generative process. Sampling from such distributions results in different synthesized textures, depending on the priors assumed. Depending on the accuracy of these assumptions, the resulting textures may, or may not, satisfy the above criteria for "good" synthesis.

One possible prior over the distribution of pixels is that the original texture is the only sample in the distribution, and that no other images are texturally similar. From this assumption, simple tiling results, as shown in Figure 2. Clearly this fails the "sufficiently different" criteria stated above.

Another feasible – though also clearly inadequate – prior is to assume that the pixels in the input texture are independently sampled from some distribution. Textures generated with this model do not capture the non-random structure within the original. The result of such an operation is shown in Figure 3. As expected it fails

*Research supported in part by DARPA under ONR contract No. N00014-95-1-0600 and by the Office of Naval Research under contract No. N00014-96-1-0311.

Figure 1: An example texture image for input to a texture synthesis process.

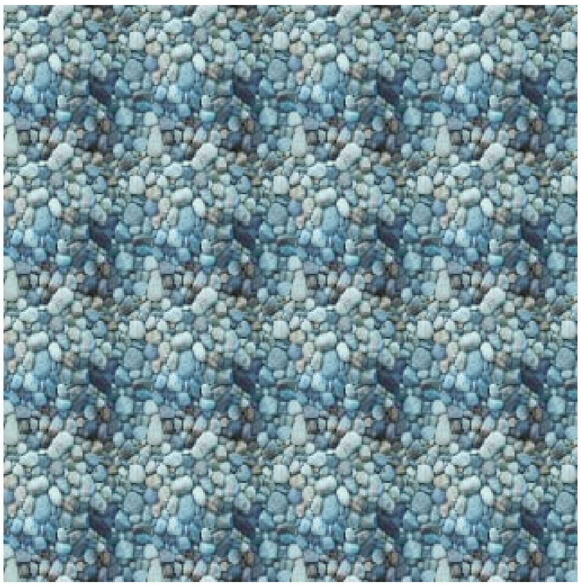

Figure 2: Simple repetition of the image does not result in a texture which appears to have come from the same stochastic distribution as the original.

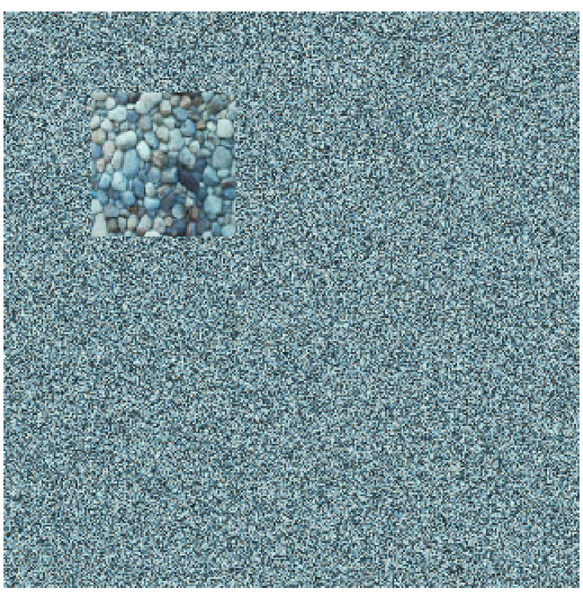

Figure 3: Textures that contain randomness not present in the original are perceptually different textures. This texture was generated by uniformly sampling the pixel values of the original. The original texture superimposed on the synthetic one is easily identified.

Figure 4: Sampling each spatial frequency band from the corresponding band in the original does not capture the detail which is characteristic of the input texture, indicating that relationships between frequencies is critical. The synthesized texture is different from the superimposed original texture, which is clearly discriminable.

Figure 5: The objective is to generate a patch such as the one above which is different from the original yet appears as though it could have been generated by the same underlying stochastic process. This texture, which was synthesized using the technique described in this paper, is perceptually very similar to the original, and the superimposed original is not readily located.

to capture the character of the original and is perceptually different. This is evidenced by the ease with which the original can be located when superimposed on the synthesized texture. This effect, commonly known as "popout" ([3, 9, 11], e.g.), occurs because the textures are perceptually different and do not appear to have been generated by the same process.

The goal of texture synthesis is to generate a texture, such as that shown in Figure 5, which is both random, and indiscriminable from the original texture. Figure 5 satisfies these criteria in that it differs significantly from the original yet appears to have been generated by the same physical process. Because of the perceptual similarity between this texture, which was synthesized by the procedure in this paper, and the input texture (generated by some other process) it is difficult to locate the region which contains the superimposed original.

3 Functional synthesis framework

Mathematically, the goal of texture synthesis is to develop a function, F, which takes a texture image, I_{input}, to a new texture sample, I_{synth}, such that the difference between I_{input} and I_{synth} is above some measure of visual difference from the original, yet is texturally similar. Formally,

$$F(I_{\text{input}}) = I_{\text{synth}} \tag{1}$$

subject to the constraints that

$$D^* \left(I_{\text{input}}, I_{\text{synth}}\right) < T_{\text{max disc}} \tag{2}$$

and

$$V^* \left(I_{\text{input}}, I_{\text{synth}}\right) > T_{\text{min diff}} \tag{3}$$

where D^* is a perceptual measure of the perceived difference of textural characteristics, and V^* a measure of the perceived visual difference between the input and synthesized images. To be acceptable, the perceived difference in textural characteristics must fall below a maximum texture discriminability threshold $T_{\text{max disc}}$, and the perceived visual difference must be above a minimum visual difference threshold, $T_{\text{min diff}}$.

The success of a synthesis technique is measured by its ability to minimize $T_{\text{max disc}}$ while maximizing $T_{\text{min diff}}$.

Human perception of texture differences, indicated by the hypothetical function D^*, depends on our prior beliefs about how textures should vary. These beliefs incorporate much of human visual experience; therefore, determining a computable metric, D, to approximate D^*, is a complex and often ill-defined task. Devising a good approximation for V^* is an even more difficult task. For texture synthesis purposes however, a poor approximation such as direct correlation, is sufficient.

The difficulty of determining a function D, to approximate D^*, depends on the structure and textual complexity of the two images. Many psychophysically based approximations have been proposed (e.g. [4, 6].)

Clearly, more complex textures can be represented in larger images; therefore, determining a discrimination function, say D_{small}, between images which have few pixels is less difficult than determining a similar function D_{large} over larger images.

Using a multiresolution approach, this work approximates D^* with a process which begins from low resolution – small – images. By decomposing the function F into a set of functions F_i which each generate a single spatial frequency band of the new texture, I_{synth}.

The domain of the each function F_i is a subset of the domain of F, as F_i's need only be a function of the information contained in the low spatial frequency bands of I_{input}. An intuitive proof of this is given by the following induction. Consider a new image, I'_{input}, which is generated from an image I_{input} by removing its high frequencies by low pass filtering with a Gaussian kernel. With just I'_{input}, and without knowledge of the additional information in I_{input}, one could still consider generating a new image I'_{synth} which is similar in textural appearance to I'_{input}. Thus, the process of generating I'_{synth} from I'_{input} is independent of highest frequency band of I_{input}. This argument can be repeated to show that I''_{synth} can be generated from I''_{input} without knowledge of I'_{input}, and so on. F_i is then given by:

$$\begin{aligned} F_i\left(I^{\prime\ldots\prime}_{\text{input}}\right) &= \\ &= F_i\left[L_i\left(I_{\text{input}}\right), L_{i+1}\left(I_{\text{input}}\right), \cdots, L_n\left(I_{\text{input}}\right)\right] \\ &= L_i\left(I_{\text{synth}}\right) \end{aligned} \tag{4}$$

where $L_i\left(I_{\text{synth}}\right)$ is the i^{th} spatial frequency octave (or equivalently the i^{th} level of the Laplacian pyramid decomposition.) The original function, F, in equation (1) is then constructed by combining the spatial frequency bands generated by F_0 through F_N The method presented here simplifies the difficulty of minimizing (approximate) D^* difference by initially synthesizing textures which are similar at low spatial frequencies, and then maintaining that similarity as it progresses to higher frequencies. A new texture is synthesized by generating each of its spatial frequency bands so that as higher frequency information is added textural similarity is preserved.

4 Texture generation procedure

4.1 Hypothesis of texture structure

The sampling procedure used by this method is dependent upon the accuracy of the following hypothesis. Images perceived as textures differ from other images in that below some resolution they contain regions which differ by less than some discrimination threshold. Further, if the threshold is strict enough, randomization of these regions does not change the perceived characteristics of the texture. In other words, at some low resolution texture images contain regions whose difference measured by D^* is small, and reorganizing these low frequency regions, while retaining their high frequency detail will not change its textural (D^*) characteristics yet will increase its visual (V^*) difference.

In Figure 6, at each resolution examples of potentially interchangeable regions are highlighted. Rearranging the image at these resolutions and locations, while retaining their high resolution structure, corresponds to moving whole textural units (which in Figure 6 are individual pebbles.)

4.2 Analysis and Synthesis Pyramids

A new texture is synthesized by generating each of its spatial frequency bands so that as higher frequency information is added textural similarity is preserved. Each synthesized band is generated by sampling from the corresponding band in the input texture, constrained by the presence of local features. The general flow of this process is outlined in Figure 7.

In a first phase the input image is decomposed into multiple resolutions. This is done using the standard Laplacian pyramid formulation where band pass information at the point (x, y) at level i, in the image I, is given by:

$$L_i\left(I, x, y\right) = \left(G_i\left(I\right) - 2\!\uparrow\![G_{i+1}\left(I\right)]\right)(x, y) \tag{5}$$

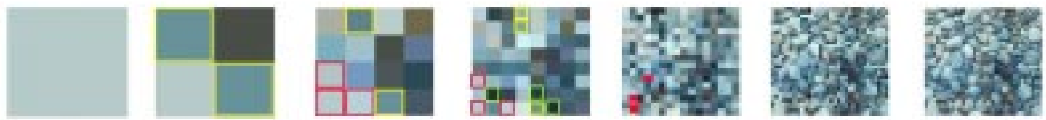

Figure 6: The synthesis procedure is based upon the hypothesis that at lower resolutions there are regions which are below some threshold of discriminability and that the randomness within a texture is in the locations of these regions.

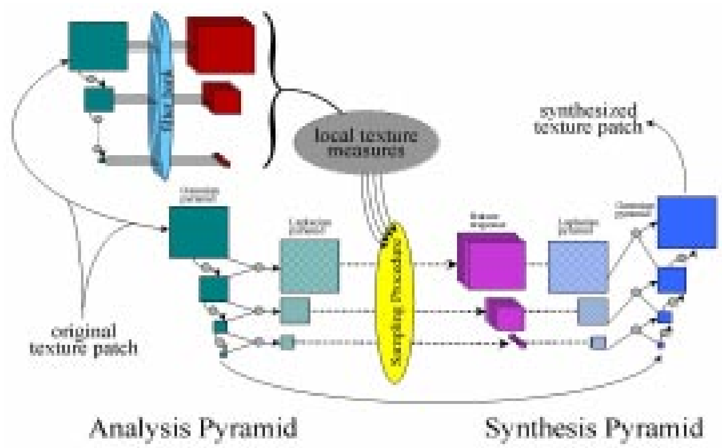

Figure 7: Multiple regions in the analysis pyramid can be candidate values for a location in the synthesis pyramid (as shown in Figure 8).

where $G_i\left(I\right)$ is a low-pass down-sampling operation:

$$G_i\left(I\right) = 2{\downarrow}[G_{i-1}\left(I\right) \otimes g] \tag{6}$$

where $2{\uparrow}[\cdot]$ and $2{\downarrow}[\cdot]$ are the $2\times$ up- and down-sampling operations respectively; g is a two dimensional Gaussian kernel; and $G_0\left(I\right) = I$.

Each level of the Laplacian pyramid contains the information from a one octave spatial frequency band of the input For a complete discussion of Laplacian and Gaussian pyramids, the reader is referred to [5].

From each level of this Laplacian pyramid a corresponding level of a new pyramid is sampled. If this sampling is done independently at each resolution, as shown in Figure 4, the synthesized image fails to capture the visual organization characteristic of the original, indicating that the values chosen for a particular spatial frequency should depend on the values chosen at other spatial frequencies. From the iterative proof, above, we can also infer that these values only depend values at that and at lower spatial frequencies.

However, using only the Laplacian information in the lower frequency bands to constrain selection is also insufficient. Such a procedure which samples from a distribution conditioned exclusively on lower resolutions only loosely constrains the relationship between the 'child' nodes of different 'parents.' Sampling from such a distribution can result in high frequency artifacts which are not present in the intended distribution. To prevent this, constraints must be propagated across children of different parents; however, constraint propagation on a two dimensional network results in dependency cycles, from which sampling requires iterative procedures, and which is not, in general, guaranteed to converge in finite time. This technique constrains the selection process within a spatial frequency band without creating cycles by using image features to constrain sampling.

Because the objective is to synthesize textures that contain the same textural characteristics as the original, yet vary from it in global form, it is assumed that global structure within the input texture is coincidental and should not constrain synthesis. Given this assumption it is sufficient to use the responses of a set of *local* texture measures as features which provide the basis for an approximation to the human perceptual texture-discriminability function D^*. A filter bank of oriented first and second Gaussian derivatives – simple edge and line filters – were used in addition to Laplacian response. At each location (x, y) in the analysis pyramid level i, the response of each feature j, is computed for use in constraining the sampling procedure. When, at the lowest resolutions, the pyramid layers are too small, the features cannot be computed, and a constant value is used.

$$F_i^j\left(I, x, y\right) = \begin{cases} \left(G_i\left(I\right) \otimes f_j\right)(x, y) & \text{if size of } G_i\left(I\right) \geq f_j \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

The constraints provided by these features are stronger than just the "parent" value, because they capture some of the relationships between pixels within a local neighborhood. This "analysis pyramid" which contains the multiresolution band-pass and feature response information, is directly computed from the input image.

4.3 Sampling procedure

A "synthesis pyramid" is generated by sampling from the analysis pyramid conditioned on the joint occurrence of similar feature response values at multiple resolutions. When the synthesized pyramid has been completely generated, the band-pass information is combined to form the final synthesized texture.

Initially the top level – lowest resolution – of the analysis pyramid, which is a single pixel, is copied directly into the synthesis pyramid. When synthesizing a texture larger than the original, the top level of the synthesis pyramid is larger that in the analysis pyramid; in this case the analysis level is simply repeated to fill the synthesis level.

Subsequent levels of the synthesis pyramid are sampled from the corresponding level of the analysis pyramid. At each location in the synthesis pyramid, the local "parent structure" is used to constrain sampling. The parent structure, $\vec{S_i}$, of a location, (x, y), in image I, at resolution i, is a vector which contains the local response for features 1 through M, at every lower resolution from $i+1$ to N:

$$\begin{aligned} \vec{S_i}\left(I, x, y\right) = \\ \Big[\; & F_{i+1}^0\left(\tfrac{x}{2}, \tfrac{y}{2}\right), F_{i+1}^1\left(\tfrac{x}{2}, \tfrac{y}{2}\right), \cdots, F_{i+1}^M\left(\tfrac{x}{2}, \tfrac{y}{2}\right), \\ & F_{i+2}^0\left(\tfrac{x}{4}, \tfrac{y}{4}\right), F_{i+2}^1\left(\tfrac{x}{4}, \tfrac{y}{4}\right), \cdots, F_{i+2}^M\left(\tfrac{x}{4}, \tfrac{y}{4}\right), \\ & \cdots, \\ & F_N^0\left(\tfrac{x}{2^N}, \tfrac{y}{2^N}\right), F_N^1\left(\tfrac{x}{2^N}, \tfrac{y}{2^N}\right), \cdots, \\ & F_N^M\left(\tfrac{x}{2^N}, \tfrac{y}{2^N}\right) \Big]^T \end{aligned} \tag{8}$$

The parent structure of a location in a synthesis pyramid is depicted in Figure 8; in this schematic, each cell represents the set of local feature responses.

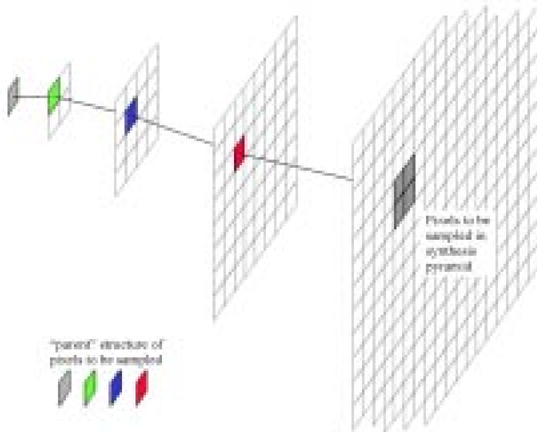

Figure 8: The distribution from which pixels in the synthesis pyramid are sampled is conditioned on the "parent" structure of those pixels. Each element of the parent structure contains a vector of the feature measurements at that location and scale.

Two locations are considered indistinguishable if the square difference between every component of their parent structures is below some threshold. For a given location (x', y') in the synthesis image, I_{synth}, the set of all such locations in the input image can be computed:

$$\mathcal{C}_i\left(x', y'\right) = \left\{ (x,y) \,\middle|\, D\left(\begin{array}{l} \vec{S}_i\left(I_{\text{synth}}, x', y'\right), \\ \vec{S}_i\left(I_{\text{input}}, x, y\right) \end{array} \right) \leq \vec{T}_i \right\} \quad (9)$$

Where the distance function D, between two parent structures u and v, is given by:

$$D\left[u, v\right] = \frac{(u-v)^T (u-v)}{Z} \quad (10)$$

where Z is a normalization constant which eliminates the effect of contrast, equal to $\sum_{x,y} \vec{S}_i\left(I_{\text{input}}, x, y\right)$.

To be a member of set $C_i\left(x', y'\right)$ the distance between each component of the parent structures must be less than the corresponding component in a vector of thresholds for each resolution and feature:

$$\vec{T}_i = \left[\begin{array}{l} T^0_{i+1}\ T^1_{i+1}\ \cdots\ T^M_{i+1} \\ T^0_{i+2}\ T^1_{i+2}\ \cdots\ T^M_{i+2} \\ \cdots \\ T^0_N\ T^1_N\ \cdots\ T^M_N \end{array} \right]^T \quad (11)$$

Where each element T^j_i is a threshold for the j^{th} filter response at the i^{th} resolution.

The values for new locations in the synthesis pyramid are sampled uniformly from among all regions in the analysis pyramid that have a parent structure which satisfies equation (**??**). This yields a probability distribution over spatial frequency band values conditioned on the joint occurrence of features at lower spatial frequencies:

$$P\left(L_i\left(I_{\text{synth}}, x', y'\right) \Rightarrow L_i\left(I_{\text{input}}, x, y\right) \,\middle|\, (x,y) \in \mathcal{C}_i\left(x', y'\right)\right)$$
$$= \quad 1 \,/\, \|\mathcal{C}_i\left(x', y'\right)\| \quad (12)$$

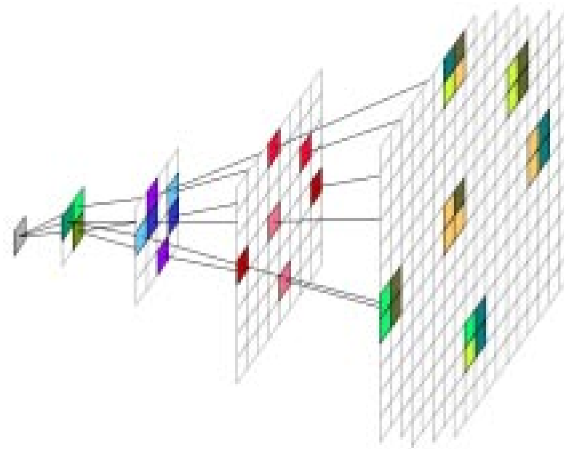

Figure 9: An input texture is decomposed to form an analysis pyramid, from which a new synthesis pyramid is sampled, conditioned on local features within the pyramids. A filter bank of local texture measures, based on psychophysical models, are used as features.

Variations between the analysis and synthesis pyramids occur when multiple regions in the analysis pyramid satisfy the above criterion. The parent structure of such a group of candidate locations is depicted in Figure 9. As the thresholds increase, the number of candidates from which the values in the synthesis pyramid will be sampled, increases. The levels of the thresholds, T^j_i, mediate the rearrangement of spatial frequency information within the synthesized texture, and encapsulate a prior belief about the degree of randomness in the true distribution from which the input texture was generated.

Algorithmically, this sampling procedure can be described with the pseudo-code:

```
SynthesizePyramid
Loop i from top_level-1 downto 0
  Loop (x',y') over Pyr_synth[level i]
    C = ∅
    Loop (x,y) over Pyr_analysis[level i]
      C = C ⋃ {(x,y)}
      Loop v from top_level downto i+1
        Loop j for each feature
          if D ( Pyr_analysis[v][j](x/2^(v-i), y/2^(v-i)) ,
                 Pyr_synth[v][j](x'/2^(v-i), y'/2^(v-i)) )
             < threshold[level v][feature j]
          then
             C = C - {(x,y)}
             break to next (x,y)

    selection = UniformRandom[0, ||C||]
    (x,y) = C[selection]
    Pyr_synth[v](x',y') = Pyr_analysis[v](x,y)
```

With more complex code, additional efficiency can be obtained by skipping whole regions which share a parent structure element that is above threshold difference.

Upon the completion of this sampling process for each level of the synthesis pyramid the synthesized band-pass information is combined to form the new texture using a standard **CollapsePyramid** procedure.

Though each band is sampled directly from the input image, the image which results from the recombination of each of these synthesized layers contains pixel values (i.e. RGB colors) not present

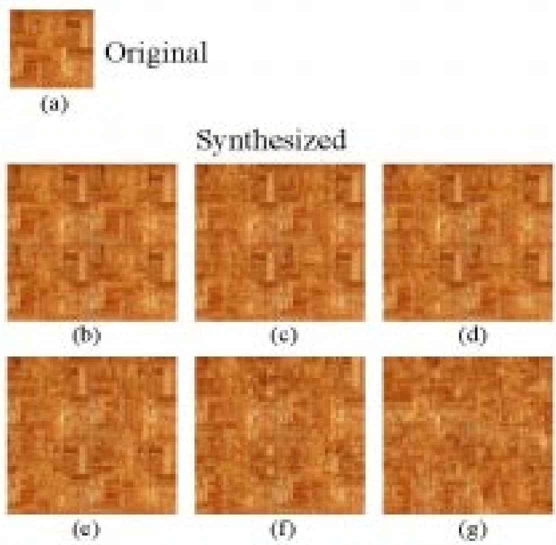

Figure 11: This series of 6 images (b-g) was generated from the original (a). For each a single threshold is used for all features and resolutions. Thresholds increase from 0.05 to 0.3 from (b) to (g).

in the original, because non-zero thresholds allow synthesized spatial frequency hierarchies which differ from those in the original.

Because the Laplacian pyramid representation is over-complete, i.e. the space spanned by Laplacian pyramids is $4/3$ larger than that spanned by images, it is possible to synthesize pyramids that are off of the manifold of real-images. When this occurs, the pyramid is projected onto the closest point on this manifold before reconstruction. This is done by collapsing the pyramid using full precision images, then replacing values above or below the range of legal pixel values with the closest legal value.

5 Examples of texture synthesis

For 800 full color input textures, we synthesized new textures, each four times larger than the original. Some typical results are shown in Figure 10. The results from these examples are indicative of the synthesis performance on the entire set and were chosen only because they reproduce well on paper. The results of all 800 textures are available on the world wide web via the URL:

http://www.ai.mit.edu/∼jsd/Research/TextureSynthesis

In the synthesis examples through out this paper thresholds of the form:

$$T_i^j = \alpha / i^\beta \tag{13}$$

were used with $\alpha \in [0, 0.4]$ and $\beta \in \{0, 1\}$. The parameter α establishes the prior belief about the sensitivity of D^*, the threshold $T_{\text{max disc}}$ in equation (2); larger β incorporates the belief that the 'true' distribution which generated the input texture is spatially homogeneous, and that the low frequency structure within the input image should not be an influential factor in region discrimination.

Shown in Figure 11 are a series of synthesized textures for $\beta = 0$ and $\alpha = \{0.05, 0.10, 0.15, 0.20, 0.25, 0.30\}$. As the threshold increases, progressively more locations in the original become indistinguishable, and the amount of variation from the original increases. For this texture, the synthesized image which balances sufficient difference from the original with perceptual similarity, lies somewhere between $\alpha = .15$ and $\alpha = .20$ (images d-e.) For different images, the ideal threshold is different, reflecting our prior

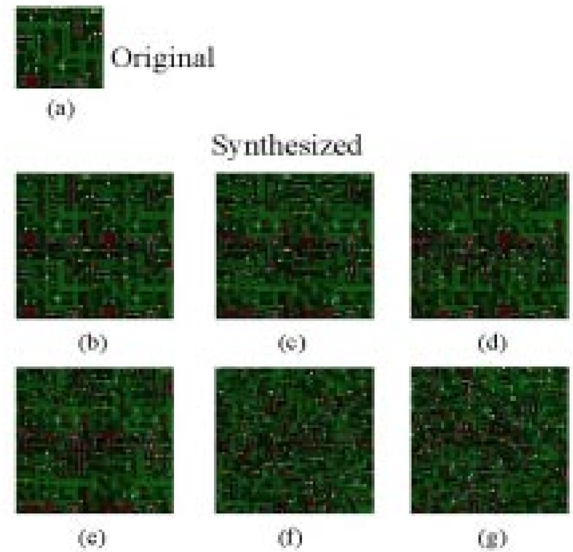

Figure 12: A series of synthesized textures for which the thresholds are inversely proportional to the spatial frequency and proportional to 0.05 in (b) to 0.3 in (g).

belief about the randomness implied by the original. Another synthesis series for a different input image is shown in Figure 12. In this case $\beta = 1$, a varies over the same range, and the ideal threshold is somewhere around $\alpha = 0.25$ (image f.)

6 Discussion

Because it uses only local constraints, the estimator presented here cannot model, texture images with complex visual structures. Such structures include: reflective and rotational symmetry; progressive variations in size, color, orientation, etc.; and visual elements with internal semantic meaning (such as symbols) or which have meaning in their relative positions (such as letters.)

Simply adding additional complex features to attempt to capture these sorts of visual structures over conditions the sampling procedure, and simple tiling results. If appropriate thresholds could be determined through additional analysis of the input image, the effects of complex features could be mediated, and they might provide useful constraints.

Because it samples exclusively from the input image, this model assumes that the 'true' distributions from which each spatial frequency band in the input was generated, can be accurately approximated by only those values present in that image. If there were a model for the probability of values not present in the original, synthesized textures could possibly be generated which contain additional variation from the original which does not increase texture (D^*) difference yet increases the visual (V^*) difference.

7 Conclusion

We have presented a method for synthesis of a novel image from an input texture by generating and sampling from a distribution. This multiresolution technique is capable of capturing much of the important visual structure in the perceptual characteristics of many texture images; including artificial (man-made) textures and more natural ones, as shown in Figure 13. The input texture is treated as probability density estimator by using the joint occurrence of fea-

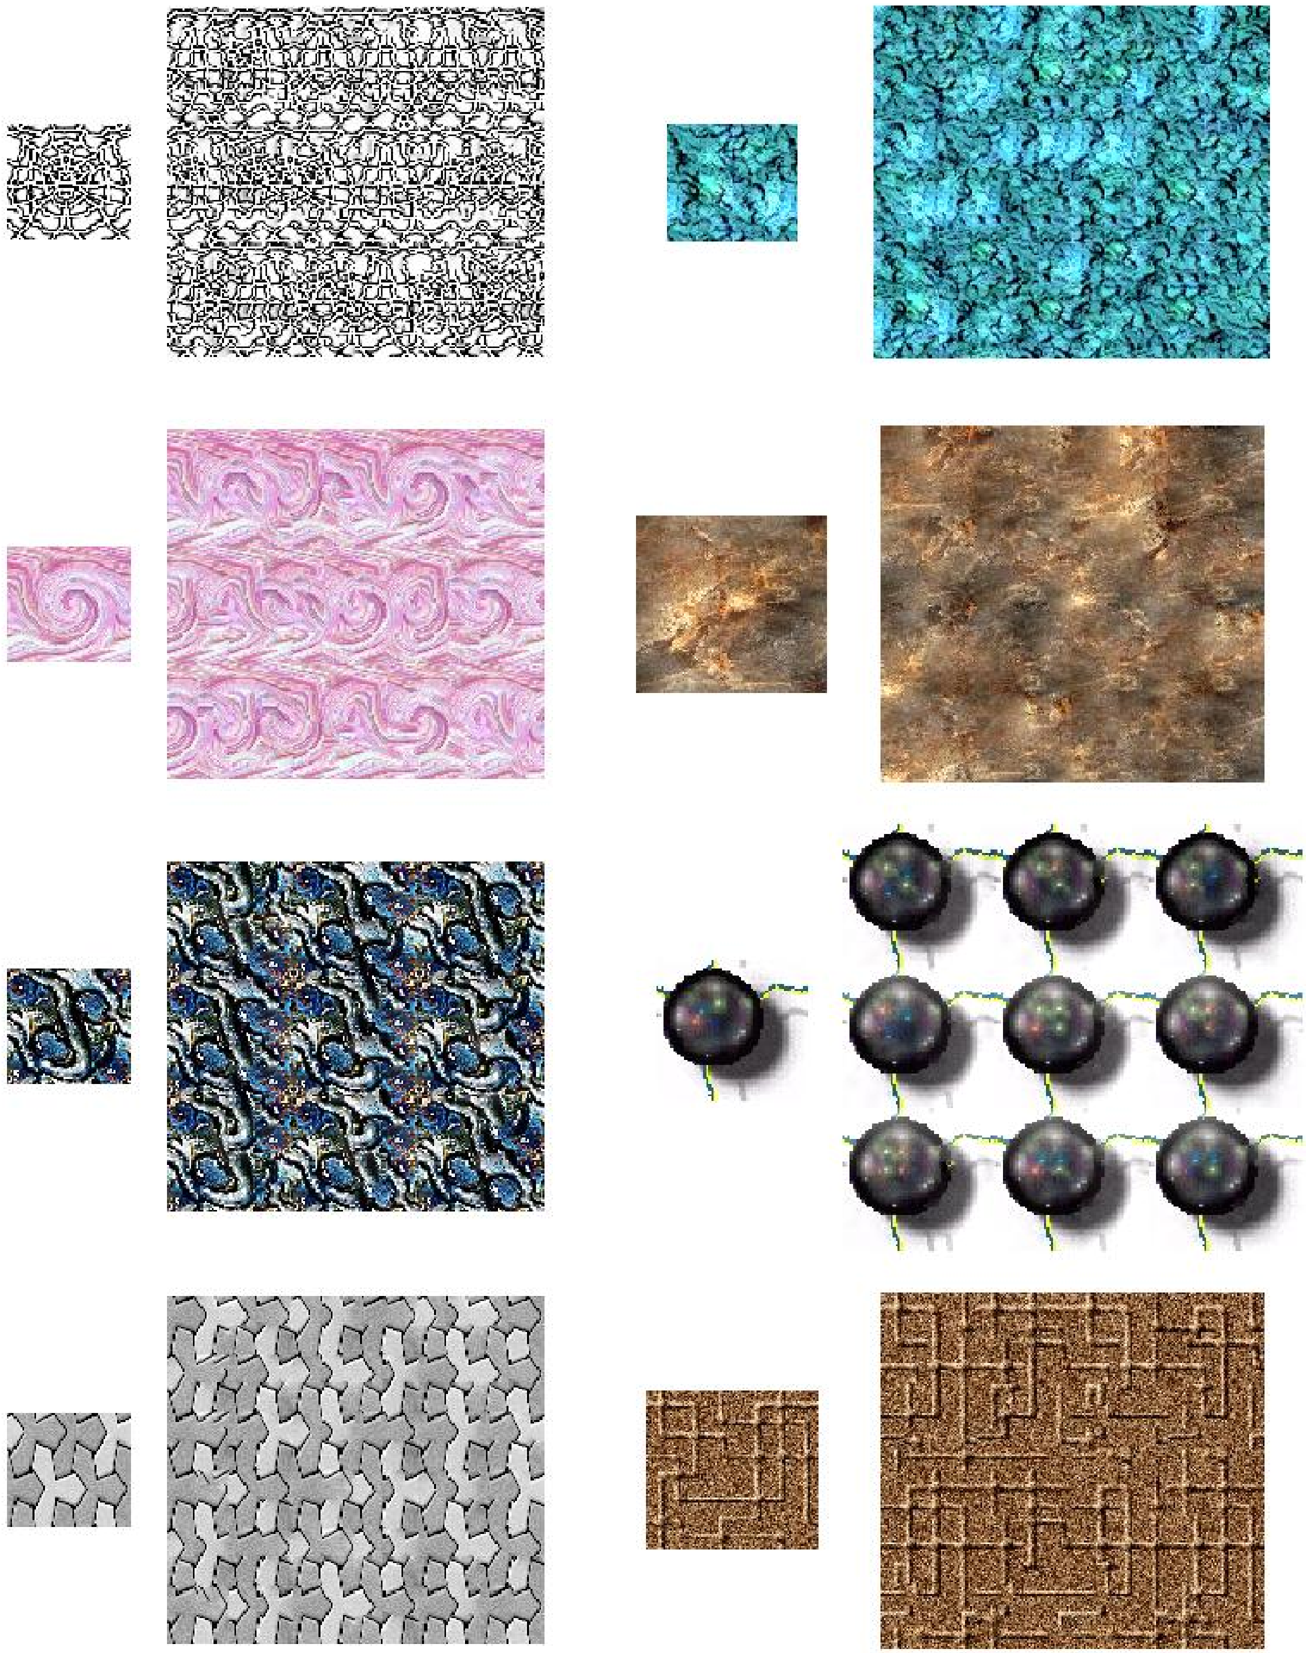

Figure 10: Texture synthesis results. The smaller patches are the input textures, and to their right are synthesized images which are 4 or 9 times larger.

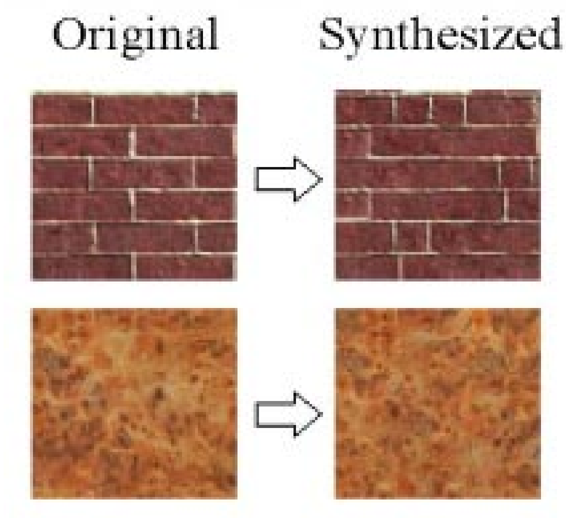

Figure 13: The characteristics of both artificial / man-made and natural textures can be captured and replicated with this process.

tures across multiple resolutions to constrain sampling. Prior beliefs about the 'true' randomness in the input are incorporated into the model through the settings of thresholds which control the level of constraint provided by each feature. Many of the textures generated by sampling from this estimator can simultaneously satisfy two the two criteria of successful texture synthesis: the synthesized textures are sufficiently different from the original, and appear to have been created by the same underlying generative process. These textures can be synthesized from more intricate input examples, and produce textures which appear more akin to the originals, than those produced by earlier techniques (Figure 14.)

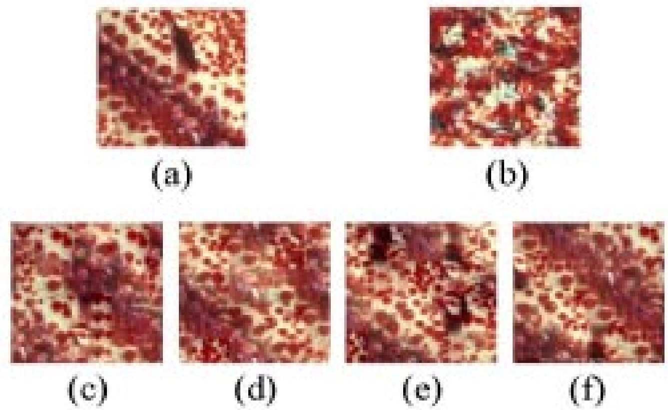

Figure 14: An input texture (a) which is beyond the limitations of Heeger and Bergen (1995) model (b), can be used successfully by this techniques to synthesize many new images. Four such synthesized images, using the same set of thresholds, are shown in (c) - (f).

References

[1] J. R. Bergen. Theories of visual texture perception. In D. Regan, editor, *Vision and Visual Dysfunction*, volume 10B, pages 114–134. Macmillian, New York, 1991.

[2] J. R. Bergen and E. H. Adelson. Early vision and texture perception. *Nature*, 333(6171):363–364, 1988.

[3] J. R. Bergen and B. Julesz. Rapid discrimination of visual patterns. *IEEE Transactions on Systems Man and Cybernetics*, 13:857–863, 1993.

[4] J. R. Bergen and M. S. Landy. Computational modeling of visual texture segregation. In M. S. Landy and J. A. Movshon, editors, *Computational Models of Visual Perception*, pages 253–271. MIT Press, Cambridge MA, 1991.

[5] P. J. Burt and E. H. Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31:532–540, 1983.

[6] C. Chubb and M. S. Landy. Orthogonal distribution analysis: A new approach to the study of texture perception. In M. S. Landy and J. A. Movshon, editors, *Computational Models of Visual Perception*, pages 291–301. MIT Press, Cambridge MA, 1991.

[7] A. Gagalowicz. Texture modelling applications. *The Visual Computer*, 3:186–200, 1987.

[8] A. Gagalowicz and S. D. Ma. Model driven synthesis of natural textures for 3–D scenes. *Computers and Graphics*, 10:161–170, 1986.

[9] N. Graham, A. Sutter, and C. Venkatesan. Spatial-frequency and orientation-selectivity of simple and complex channels in region segregation. *Vision Research*, 33:1893–1911, 1993.

[10] D. J. Heeger and J. R. Bergen. Pyramid based texture analysis/synthesis. In *Computer Graphics*, pages 229–238. ACM SIGGRAPH, 1995.

[11] B. Julesz. Visual pattern discrimination. *IRE Transactions on Information Theory*, IT–8:84–92, 1962.

[12] M. R. Luettgen, W. C. Karl, A. S. Willsky, and R. R. Tenney. Multiscale representations of markov random fields. *IEEE Trans. on Signal Processing*, 41(12):3377–3396, 1995.

[13] S. D. Ma and A. Gagalowicz. Determination of local coordinate systems for texture synthesis on 3-D surfaces. *Computers and Graphics*, 10:171–176, 1986.

[14] G. Turk. Genereating textures on arbitrary surfaces using reaction-diffusion. In *Computer Graphics*, volume 25, pages 289–298. ACM SIGGRAPH, 1991.

[15] A. Witkin and M. Kass. Reaction–diffusion textures. In *Computer Graphics*, volume 25, pages 299–308. ACM SIGGRAPH, 1991.

[16] S. C. Zhu, Y. Wu, and D. Mumford. Filters random fields and maximum entropy(frame): To a unified theory for texture modeling. *To appear in Int'l Journal of Computer Vision*, 1996.