



Honest Image Thumbnails: Algorithm and Subjective Evaluation

Ramin Samadani, Tim Mauer, David Berfanger, Jim Clark, Suk Hwan Lim, Dan Tretter
Media Technologies Laboratory
HP Laboratories Palo Alto
HPL-2007-88
June 1, 2007*

image thumbnails,
image quality,
image browsing,
blur, noise

Image thumbnails are commonly used for selecting images for display, sharing or printing. Current, standard thumbnails do not distinguish between high and low quality originals. Both sharp and blurry originals appear sharp in the thumbnails, and both clean and noisy originals appear clean in the thumbnails. This leads to errors and inefficiencies during image selection. In this paper, thumbnails generated using image analysis better represent the local blur and the noise of the originals. The new thumbnails provide a quick, natural way for users to identify images of good quality, while allowing the viewer's knowledge to select desired subject matter. A subjective evaluation using twenty subjects shows the new thumbnails are more representative of their originals for *blurry* images. In addition, there are no significant differences between the results of the new thumbnails and the standard thumbnails for *clean* images. The noise component improves the results for *noisy* images, but degrades the results for *textured* images. The blur component of the new thumbnails may always be used. The decision to use the noise component of the new thumbnails should be based on testing with the particular image mix expected for the application.

Honest Image Thumbnails: Algorithm and Subjective Evaluation

Ramin Samadani¹, Tim Mauer², David Berfanger², Jim Clark², Suk Hwan Lim¹ and Dan Tretter¹

1) HP Labs, Media Technologies Lab and 2) Rainbow Image Science Team

Ramin.Samadani@hp.com

Abstract

Image thumbnails are commonly used for selecting images for display, sharing or printing. Current, standard thumbnails do not distinguish between high and low quality originals. Both sharp and blurry originals appear sharp in the thumbnails, and both clean and noisy originals appear clean in the thumbnails. This leads to errors and inefficiencies during image selection. In this paper, thumbnails generated using image analysis better represent the local blur and the noise of the originals. The new thumbnails provide a quick, natural way for users to identify images of good quality, while allowing the viewer's knowledge to select desired subject matter. A subjective evaluation using twenty subjects shows the new thumbnails are more representative of their originals for *blurry* images. In addition, there are no significant differences between the results of the new thumbnails and the standard thumbnails for *clean* images. The noise component improves the results for *noisy* images, but degrades the results for *textured* images. The blur component of the new thumbnails may always be used. The decision to use the noise component of the new thumbnails should be based on testing with the particular image mix expected for the application.

1 Introduction

Image thumbnails are pervasive. Computer operating systems and applications display image thumbnails of folders or albums. Photo kiosks let users review thumbnails, touch the screen at the thumbnails, and then print the selected photos. Image sharing sites display thumbnails of photo albums. Small displays on printers, cameras, cell phones, and video players let users preview images before taking actions such as viewing, mailing, printing, or deleting.

As the examples show, image thumbnails are very useful for selecting images since one can inspect more images at once. In addition, one can recognize personally meaningful images from their thumbnails. Unfortunately, standard thumbnails are unreliable for selection since they lose information about image quality. Our new thumbnails improve image selection by better representing their originals.

Standard thumbnail generation involves lowpass filtering and downsampling. This process results in thumbnails that do not represent the quality of the high resolution originals. None of the many sources of image blur, including unintentional misfocus and camera shake, as well as intentional depth of field local blurs are represented. Image noise, particularly prevalent in night or indoor scenes, is also not preserved.

Browsing with standard thumbnails leads to errors and inefficiencies. While browsing, one can easily select a normally appearing thumbnail to find that the original is blurred, noisy, or both. The same problem on printer or camera LCDs leads to erroneous selections for printing or deleting. Browsing with standard thumbnails is inefficient since they ambiguously represent the quality of their originals. This ambiguity means that it takes extra time to inspect the originals by trial and error to ensure they are of high quality.

This paper describes new thumbnails that alleviate these problems by representing original image quality in addition to image composition. Figure 1 shows examples of the results for the cropped originals shown in Figure 2. It is best to view these images in the PDF document, since they are not tuned for the print process.

A benefit of the new thumbnails is that they are natural to interpret; there is no learning necessary to understand the blur and noise shown in the new thumbnails. The alternate approach of automatic image ranking by quality [1] is extremely difficult because the knowledge about the subjects of interest resides with the user, not with the image. For example, with the new thumbnails, the user can quickly check whether the subject of interest is in focus.

Section 2 describes the algorithm for the new thumbnails and Section 3 describes the subjective comparison of the new and standard thumbnails. In Section 2.1, we formulate the general thumbnail extraction problem. Section 2.2 describes the limitations of currently used standard thumbnails. These limitations are overcome with the new thumbnails described in Section 2.3. Section 3.1 describes the softcopy method used for the subjective user study. The findings of the study are summarized in Section 3.2 and the implications are discussed in Section 3.3.

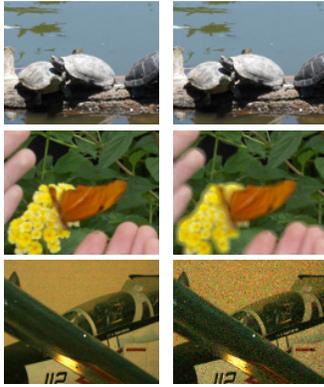


Figure 1: Standard thumbnails (left column) and new thumbnails (right column) for three different images. The Figures are best viewed using the original PDF file.

2 Algorithm

2.1 Image Model and Solution Formulation

For simplicity of notation, images are considered column stacked vectors [2]. The image model we use is

$$\mathbf{d} = B\mathbf{c} + \mathbf{n}. \quad (1)$$

In this equation, the vector \mathbf{c} represents an ideal image captured with infinite depth of field. The matrix B represents, in general, a space-varying blur that may correspond to unintended distortions such as camera shake, motion blur or misfocus, and \mathbf{n} represents an additive, perhaps correlated, noise. Well taken digital photographs will not have unintended distortions. In this case, the noise $\mathbf{n} = 0$. But the matrix B may not be the identity, still representing the space-varying depth of field blur. In the special case of infinite depth of field, $B = I$, and therefore $\mathbf{d} = \mathbf{c}$.

Our work takes advantage of prior work on the very difficult problems of image denoising [3] and blind deconvolution [4], where the goal is to recover \mathbf{c} in Equation 1 from \mathbf{d} . The goal for our work, however, is to generate a low resolution thumbnail \mathbf{t}_n , not the exact reconstruction of high resolution \mathbf{c} . This changes the requirements of our component algorithms. For example, our solution works well with both shake and defocus blurs, by applying an appropriate space-varying Gaussian blur. The details of the applied blur kernel is not critical to our results. Similarly for noise, we do not need an extremely accurate noise estimate, but rather a rough, fast one may be sufficient.

2.2 Standard thumbnails

We first consider the limitations of standard thumbnail generation. Commonly used thumbnails differ in the details of the filters applied [2], but they consist of a linear process, first applying an antialiasing lowpass filter, A , followed by subsampling, S . The thumbnails are thus given by

$$\mathbf{t}_s = T_s \mathbf{d} = S A \mathbf{d}, \quad (2)$$

where T_s represents the combination of filtering and subsampling.

Expanding Equation 2 using the image model in Equation 1 results in,

$$\mathbf{t}_s = S(AB\mathbf{c} + A\mathbf{n}). \quad (3)$$

Analysis of the quantity in parenthesis explains why standard thumbnails appear sharp and clean, even if the input image \mathbf{d} has blur and noise added. First, the bandwidth of a typical blur filter B is broader than the bandwidth of the antialiasing filter A for typical subsampling factors (the ones used in our tests were between 10 and 23). Thus

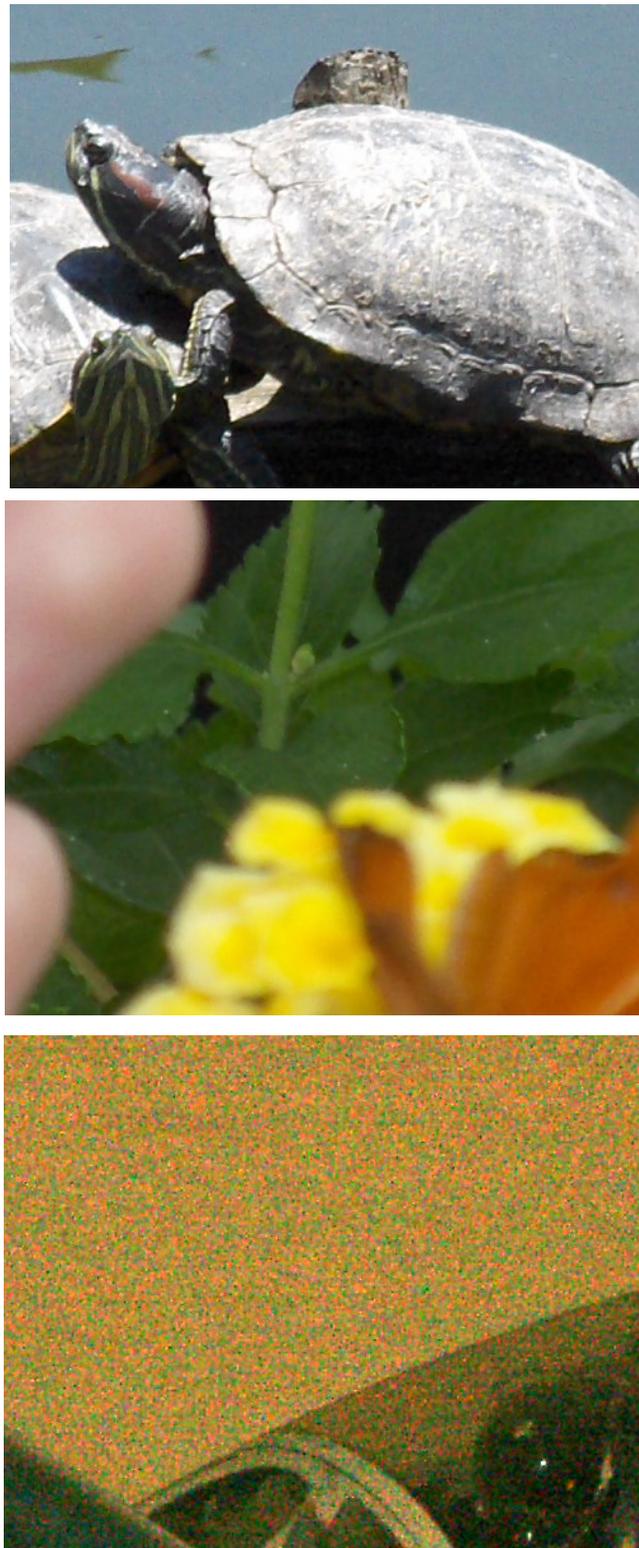


Figure 2: Cropped portions of the three originals corresponding to the thumbnails seen in Figure 1

$AB \approx A$, in Equation 3. Considering noise next, antialiasing filter A applied to n will result in output filtered noise variance much lower than the input variance, so that $An \approx 0$. This is true for typical noise levels and for any practical antialiasing filter. The case of a $k \times k$ boxcar filter, for a subsampling factor k is particularly easy to analyze. If the input noise is uncorrelated, the output noise variance will be reduced by a factor of $\frac{1}{k^2}$. For antialiasing filter,

A, the simulations in this section used a boxcar filter corresponding to $k = 10$, resulting in noise standard deviation of $\frac{1}{10}$ of the input standard deviation.

With these approximations, $T_s d \approx T_s c$. The standard thumbnail for the distorted image d will be very similar to the thumbnail for the ideal image c for typical levels of blur and noise.

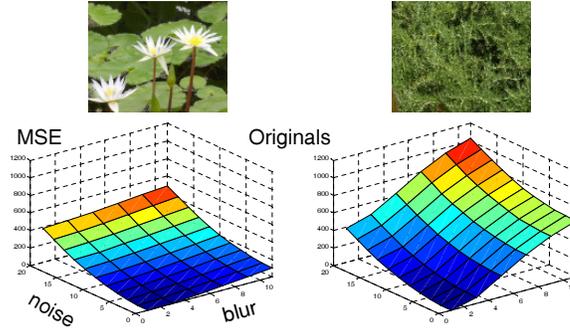


Figure 3: Mean square errors for two different images

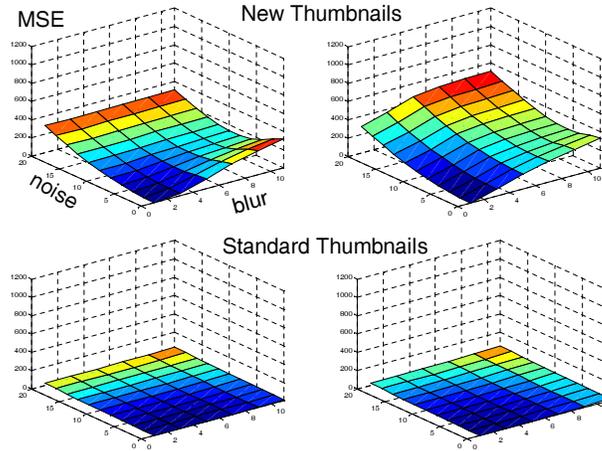


Figure 4: Mean square errors for thumbnails

These approximations are confirmed by simulations that apply differing amounts of blur and noise to input test images to generate different distorted images d in Equation 1. For the matrix B , $m \times m$ boxcar filters with $m \in \{1, 3, 5, 7, 9, 11\}$ were applied. For the noise, a moving average noise generated by filtering white Gaussian noise with a 3×3 boxcar filter (to roughly simulate the observed noise correlation in actual photographs [5]) was applied. The standard deviation of the noise was set to $\sigma_n \in \{0, 2, 4, \dots, 20\}$. Thus, for each test image, this generates a set of images d_{ij} indexed by blur and noise, that include the original and 65 distorted versions with differing amounts of blur and noise.

The mean square error (MSE) between a distorted image d_{ij} , and undistorted image $c \equiv d_{00}$, is proportional to the square norm between the images interpreted as vectors, given by $\frac{1}{N} \|c - d_{ij}\|^2$, where N is the total number of pixels in the image.

Considering the original image c as a vector in a very high dimensional space, a distorted image may be ex-

pressed as the addition of two vectors to the original image, given by,

$$\mathbf{d}_{ij} = \mathbf{c} + (B_i\mathbf{c} - \mathbf{c}) + \mathbf{n}_j = B_i\mathbf{c} + \mathbf{n}_j. \quad (4)$$

This equation shows that the noise component is independent of the input image, but that the blur distortion will depend on image content. Higher spatial frequency content in the input image results in more blur distortion.

Figure 3 plots the MSE between each distorted image and the undistorted version. Two different images, each with 1024×1280 pixels, are used in the simulations, to illustrate the changing nature of the blur, depending on image content. The image of the water plants on the left has typical spatial frequencies, but the image of the ground cover on the right has higher spatial frequencies, that are seen in the figure by observing the faster increase of MSE along the blur axis of the ground cover image. Since the blur and noise are independent, the MSEs add for images with both blur and noise. Viewing the graphs, along the noise direction at the zero blur axis shows the same noise for both images, confirming that the noise MSE in this case is image independent.

Figure 4 plots the MSE of thumbnails, of size 102×128 pixels, (subsampling factor $k = 10$) generated by our new approach as well as standard thumbnail MSEs. The mean square errors shown in each plot are between the reference thumbnail, for the input image without blur and noise, and the thumbnails for the other simulated images. Relevant to the discussion in this section, the bottom plots of the figure show the surfaces for the standard thumbnails are near zero for all the different thumbnails (corresponding to the different input images \mathbf{d}_{ij}). The thumbnails are also observed to be visually very similar, showing neither blur nor noise.

2.3 New thumbnails

The new thumbnails, \mathbf{t}_n , are generated by starting first with the standard thumbnail, which was shown in the previous section to be clean even for distorted input images. To this standard thumbnail, blur and noise are applied to correspond to the blur and noise in the original,

$$\mathbf{t}_n = \mathbf{t}_b + \mathbf{n}_t = B_t\mathbf{t}_s + \mathbf{n}_t. \quad (5)$$

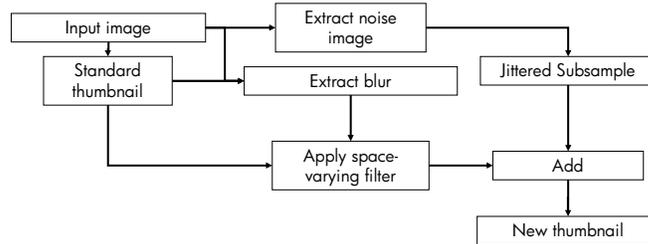


Figure 5: To generate a new thumbnail, we start with a standard thumbnail and use image analysis to estimate and apply local blur and noise.

Figure 5 shows the processing that generates the new thumbnails. The *Extract blur* block results in a two dimensional thumbnail resolution *blur map*, \mathbf{m} , with estimates of the amount of blur at each location and the block *Apply space-varying filter* applies a filter based on this blur map. This local computation accounts for depth of field as well as undesired blurs. The blur map is determined without identification of the type of blur. The assumption is that users will not be able to distinguish between different types of blur in a thumbnail very easily.

The local amount of blur is computed by noting that the image edge profiles differ between sharp and blurry images. At an edge, for example, the profile of a blurry high resolution image will be more gradual than its corresponding low resolution standard thumbnail, \mathbf{t}_s , whose profile will be steeper [6]. Applying successively larger blurs to \mathbf{t}_s will cause its edge profile to become more gradual, and to correspond better to the blurry original. To have the system work with various image features, and not just edges, the computation is based on *pixel range* (difference between maximum and minimum pixel values in a spatial neighborhood) to determine the local image profiles.

During the building of the *blur map*, a set of Gaussian filtered low-pass versions of the standard thumbnail, l_σ are created,

$$l_\sigma = g_\sigma * t_s. \quad (6)$$

Eleven l_σ , with $\sigma_j \in \{0, 0.5000, 0.7111, 0.9222, 1.1333, \dots, 2.4\}$, are generated. The image l_0 represents the unblurred thumbnail t_s . The remaining images correspond to increasing blurs, starting with $\sigma = 0.5$, ending with $\sigma = 2.4$ and with increment 0.2111.

From the original image, d , a low resolution *range map*, is computed. First, the maximum absolute difference for each center pixel and its eight neighbors is calculated. This high resolution range map is reduced to thumbnail size by taking the maximum in a high resolution neighborhood of the same size as the resampling factor k . The low resolution original range map is called r_o .

Similarly, from each of the images l_σ , a low resolution range map, r_σ , is generated. The blur map value at each pixel index, i , is then computed using

$$m(i) = \min_j \{j \mid r_{\sigma_j}(i) \leq \rho r_o(i)\}, \quad (7)$$

where ρ is a constant that sets the amount of blur added. Equation 7 implements the idea, described earlier, of reflecting in the thumbnails the local pixel range found in the high resolution original. The blur map shown in the *Extract Blur* block of Figure 5 is computed by comparing local range of the images in the scale space expansion of the standard thumbnail, r_{σ_j} of Equation 7 with the resampled local range of the original image, r_o . Using this blur map, at each pixel, i , a space-varying blurred thumbnail t_b , the first term in Equation 5, is created by selecting values from the appropriate blurred thumbnail,

$$t_b(i) = l_{\sigma_m}(i). \quad (8)$$

This is shown in the block labeled *Apply space-varying filter* in Figure 5. More accurate blur maps and space-varying filter implementations are possible [7], but this simple approach has worked well with the tested images.

For the noise component, n_t , a simple, modified wavelet based soft thresholding [8] (known as *VisuShrink*), was used. The noise residual is based on a high-pass filtered original, $h = d - g_1 * d$, where g_1 is a Gaussian filter with $\sigma = 1$. Following Donoho, the high resolution noise, \hat{n} at each pixel i , is estimated using

$$\hat{n}(i) = h(i) - \text{sgn}(h(i))(|h(i)| - \lambda)_+, \quad (9)$$

with $x_+ = x$ if $x \geq 0$ and $x_+ = 0$ if $x < 0$.

The threshold λ is determined by first estimating noise standard deviation, $\hat{\sigma}_n = h_m / .6745$, where h_m is the median of absolute values of the pixels $|h(i)|$. Then, threshold $\lambda = \hat{\sigma}_n \log(N)$ is used, where N is the number of pixels in the original. The noise found in Equation 9 is multiplied by the empirically determined gain factor of 1.6. This gain factor adjusts for the proportion of noise that passes through the high-pass filter, given that noise in digital photographs is typically correlated [5]. The same fixed factor was used to process all of the images in our tests.

From \hat{n} , the low resolution thumbnail noise n_t in Equation 5 is generated by subsampling. The subsampling of the noise component is justified by considering the autocorrelation of a *discrete time*, stationary Gaussian noise process after subsampling. In particular, the noise variance will remain unchanged after subsampling on a regular sampling grid. The noise generation process used is not perfect, however, allowing some high spatial frequency signal to appear in the noise image. Jittered sampling [9] is used to reduce potential Moire from any residual image textures that appear in the noise image.

Our research software generated new thumbnails on a 2 Gigahertz Pentium IV laptop in around .14 second per image.

The top plots of Figure 4 show the MSEs for the new thumbnails, comparing the thumbnail of each distorted input image, d_{ij} , with the new thumbnail for the image without blur and noise, c . These plots show that the new thumbnails discriminate much better than the plots for the standard thumbnails shown in the bottom of the figure. The slight dip observed in the plots at high blur levels show that the blur estimation is somewhat sensitive to noise. The MSEs, however, show that significant blur is still present in the thumbnails. Noise resistant blur estimation [6] may provide improvements to the plots. On the other hand, careful visual study of the interaction of blur and noise may show that for noisy images, correct blur estimation is not critical to image quality determination.

3 Subjective Evaluation

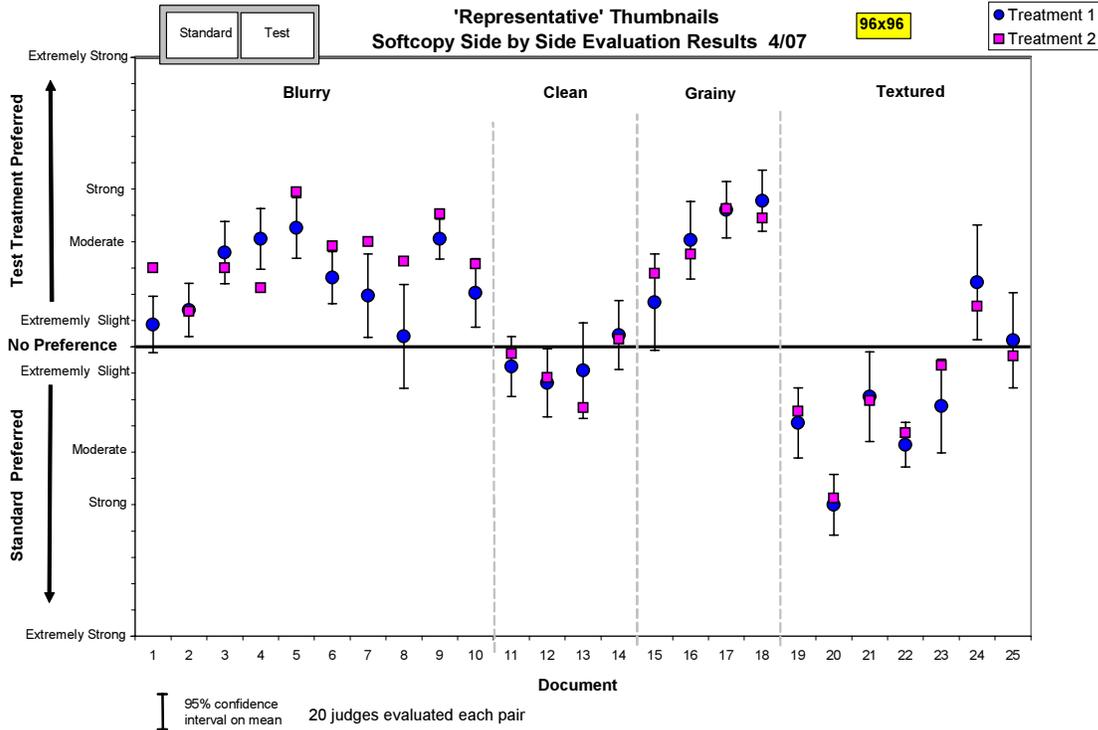


Figure 6: Plots of results for 96x96 pixel thumbnails for 25 images

Computer simulations, shown in Figure 4, show better blur and noise representation using the new thumbnails. In addition, the algorithm was first tested informally with several hundred images. The results found the algorithm to be effective for blurry and noisy images, but also found differences between the standard and new thumbnails for textured images. By turning off the noise processing, corresponding to term n_t in Equation 5, it was determined that the differences for the textured images (as well as the noisy images) were due to the noise term. This is understandable since the currently used noise algorithm does not always distinguish between image noise and texture, both of which contain high spatial frequencies.

These findings identified four categories of input images, *Blurry*, *Clean*, *Grainy* (as used here, the term grainy is equivalent to noisy) and *Textured* for further study. The next section discusses a softcopy subjective evaluation conducted using these four image categories.

3.1 Evaluation Method

The evaluation consisted of comparing a standard thumbnail vs. a candidate treatment thumbnail for best representation of the original image. Two thumbnail sizes were included, one with 96x96 pixel bounding box and one with 300x300 pixel bounding box. Two thumbnail treatments, corresponding to different blur factors ρ in Equation 7, were tested for each size. For the 96x96 thumbnails, the values used were $\rho = .85$ for *treatment 1* and $\rho = 1.25$ for *treatment 2*. For the 300x300 thumbnails, the values used were $\rho = 1.25$ for *treatment 3*, and $\rho = 1.5$ for *treatment 4*. The image suite consisted of a total of twenty five photos, divided into the four categories: blurry, clean, grainy, and textured. A total of one hundred pair were judged in a softcopy evaluation. The judges were asked to determine which thumbnail version of a pair best represented the original full image. Twenty judges participated.

The evaluation was conducted in the HP Rainbow Image Science Test Lab in Vancouver, Washington using the softcopy display monitors. The monitors are HP L2335 Active Matrix TFT's (thin film transistor), which have a 23

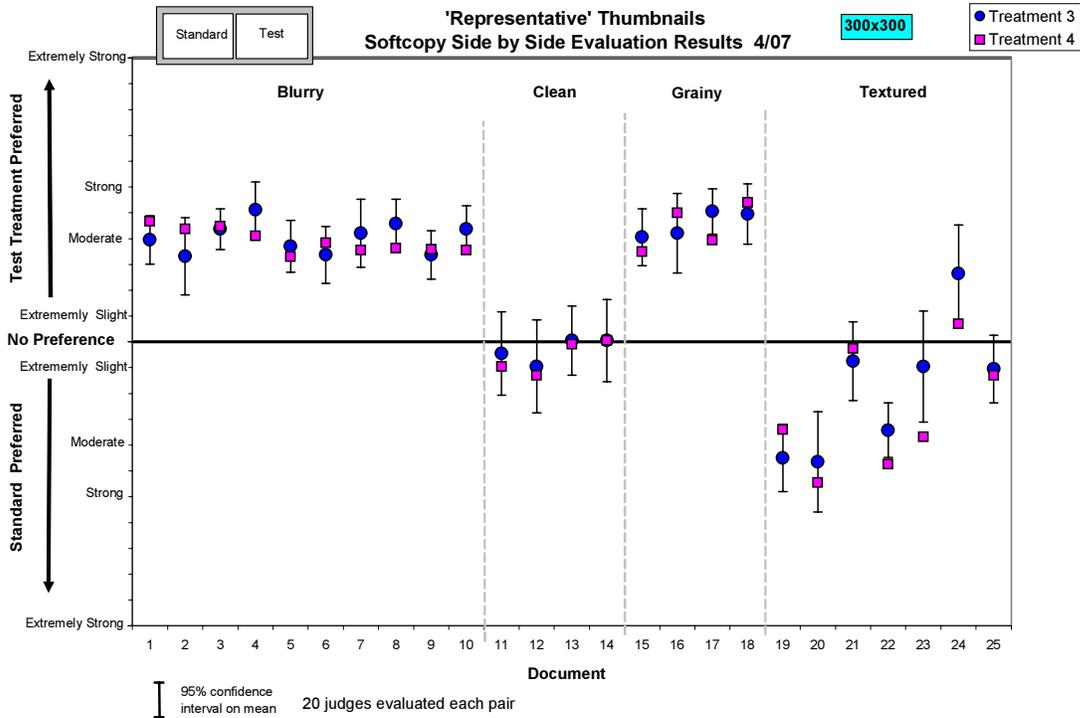


Figure 7: Plots of results for 300x300 thumbnails for 25 images

inch diagonal viewing screen length and a native resolution of 1920 x 1200 pixels and a pixel pitch of 0.258mm. The monitors were calibrated just prior to the testing using the Monaco Optix 2.0 software and sensor produced by X-Rite.

The overhead room lighting was turned off for this evaluation. There were a few task lamps turned on across the room to provide enough light for the judges to see to walk safely to the test cubicle where the softcopy monitors are stationed.

The observers were selected from a pool of HP employees at the Vancouver site that have satisfactorily completed color vision discrimination testing and observer orientation training. They are experienced at evaluating image quality primarily because of involvement in the evaluations that are conducted weekly. They are not considered *experts*, but are believed to predict the general consumer response. Periodic external evaluations have validated the calibration of the HP internal testing with real customers. The observer pool is diverse in gender and age. Special care is taken to keep the observers unwitting with regards to the source of the samples and the objectives of the test.

Thumbnail treatments were positioned to appear on the right and left side randomly.

Observers were presented with a full version view of an image in softcopy and then two candidate thumbnail versions side by side. They were asked to indicate which thumbnail version was the most representative of the full-size original and indicate the degree of their preference using the scale provided. They were able to toggle the screen back and forth between the full image and the thumbnails. After recording their response, the next image set would load and the observer would proceed until all 100 samples were evaluated.

Samples were presented to each observer in a different random order. This technique distributes any *start-up* or *fatigue* effects over different samples.

3.2 Results

Graphs of the results for the 96x96 thumbnails is shown in Figure 6 and the results for the 300x300 thumbnails is shown in Figure 7. The data on each graph are grouped by image category. The results are now described by image

category:

- *Blurry images* - for the small thumbnails (96x96) treatment 2 produced slightly better representations of the original than did treatment 1, and both treatments were better than the standard. For the large thumbnails (300x300) there was no difference between treatment 3 and 4 performance, both were significantly better than the standard, and the result was consistent across the document suite and on average better than the 96x96 results.
- *Clean images* - there was no difference in the treatment performance for either size, and no significant preference between standard and treated thumbnails (that is to say, the treatments didn't *break* the clean images).
- *Grainy images* - no difference in performance between the thumbnail treatments for either size, and for both sizes the treated versions were more representative than the standard thumbnails.
- *Textured images* - the treatment applied to the thumbnails in most cases was not as good a representation of the original as the standard thumbnail. This was true for both sizes of thumbnails. In general, the treatment added *speckle*. In the case of the bird image with a screen door in the background, shown in Figure 9, and also shown as Document 20, the second textured plot in Figure 6, significant and non-representative distortion was visible. Treatment 4 on the 300x300 size was slightly worse than Treatment 3.

Overall, the treatments appear to work well for blurry and grainy images, do not harm clean images, and do harm the textured images. For the blurry and grainy images, testing new values of ρ may provide for *better* representation of the original than the treatments tested provided.

3.3 Discussion

The main findings of the subjective study, even without careful algorithm parameter tuning, are encouraging. At both resolutions, users clearly find that the new thumbnails better represent the blurry and noisy images. There are not significant differences between the two thumbnails for the clean images. For most of the textured images, however, the users prefer the standard thumbnails.

Figures 1 and 2 show thumbnail comparisons and originals for three examples used in the tests, including ones with noise and blur. The images are not tuned for the print process. They are best viewed in the original PDF document. The top image shows an example where the new thumbnails and standard thumbnails are indistinguishable for a good quality image. In the middle image, the hands, yellow flowers and red butterfly in the middle image are misfocused, as is seen in the new thumbnail, but not the standard thumbnail. The bottom image is noisy, as seen in the new thumbnail, but not the standard thumbnail. The originals shown in Figure 2 are cropped to save space while showing the thumbnails and originals at the correct relative scales. It also seems that the blur and noise are natural for users to interpret.

Figures 8 and 9 show two examples of textured images, with high spatial frequency content. The current noise component of our thumbnail algorithm does not distinguish well between noise and texture. For example, for the top image of Figure 9, the screen door in the background has repetition frequency above Nyquist. The screen door disappears (at 96x96 bounding box), as expected, using the standard thumbnail shown in the top left of Figure 8. With the new thumbnail, however, the screen door appears as a patterned noise. Neither thumbnail, in this case, accurately represents the original but for this image users prefer the standard thumbnail. This image corresponds to document 20, the second textured plot shown in the subjective test results shown in Figure 6.

For most of the textured images, including ones containing spraying water and sand, users preferred the standard thumbnail. The textured image for which users preferred the new thumbnail is shown in the bottom of Figure 8 and 9. In this case, it appears that the new thumbnail better represents the carpet texture whereas the texture is not apparent in the standard thumbnail. The subjective results for this image correspond to document 24, the sixth textured plot shown in Figure 6.

The results of the subjective evaluation show that the blur component of the algorithm may always be turned on with improved effects. The noise component of the algorithm, however, improves the grainy images but degrades

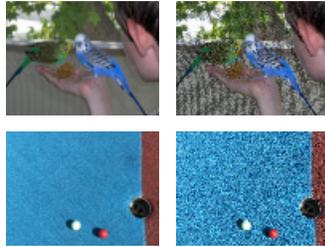


Figure 8: Standard thumbnails (left column) and new thumbnails (right column) for two *textured* images. The Figures are best viewed using the original PDF file.



Figure 9: Cropped portions of the two *textured* originals corresponding to the thumbnails seen in Figure 8

the textured images. The decision to use the noise component requires further testing with the expected image mix for the particular application. For the subjective evaluation, roughly equal numbers of the different image categories were used to better test the algorithm. How often noisy images and textured images occur for a given application setting may help determine whether the noise component should be turned on or turned off. Future work may develop a noise component that better separates between noise and texture, allowing the noise component to always

be turned on without degrading textured images.

4 Conclusions

New thumbnails were created that represent, at thumbnail resolution, the blur and noise found in high resolution images. This is done without recovering actual blur kernels. The thumbnails should be useful in any image selection situation.

Thumbnail generation, with its less exacting requirements, may be well suited for application of techniques developed for more traditional deblurring and denoising applications.

Finally, how human perception of images changes with image size is an area of fundamental study that could be generally useful to image browsing where thumbnails are the primary interface. Results in noise visibility, blur visibility, and the interaction between the two when they are masked by consumer digital photographs at different scales would be very useful to this work.

References

- [1] Yan Ke, Xiaoou Tang, and Feng Jing, "The design of high-level features for photo quality assessment," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2006, vol. 1, pp. 419–426.
- [2] Anil K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Inc., 1989.
- [3] Alan S. Wilsky, "Multiresolution markov models for signal and image processing," *Proceedings of the IEEE*, vol. 90, no. 8, pp. 1396–1458, August 2002.
- [4] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE Signal Processing magazine*, pp. 43–64, May 1996.
- [5] Suk Hwan Lim, "Characterization of noise in digital photographs for image processing," in *Proceedings of SPIE*, January 2006, vol. 6069.
- [6] L.J. Ferzli, R. Karam, "No-reference objective wavelet based noise immune image sharpness metric," in *IEEE ICIP 2005 Proceedings, Vol. 1*, Sept. 2005, pp. 405–408.
- [7] Javier Portilla and Rafael Navarro, "Efficient method for space-variant low-pass filtering," in *VII National Symposium on Pattern Recognition and Image Analysis*, Barcelona, Spain, 1997, vol. 1, pp. 287–292.
- [8] David L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. on Inf. Theory*, pp. 613–627, Dec. 1995.
- [9] Robert L. Cook, "Stochastic sampling in computer graphics," *ACM Transactions on Graphics*, pp. 51–72, Jan. 1986.