

Perceptual image preview

Liang Wan · Wei Feng · Zhouchen Lin ·
Tien-Tsin Wong · Zhi-Qiang Liu

Received: 21 March 2007 / Accepted: 27 February 2008 / Published online: 5 April 2008
© Springer-Verlag 2008

Abstract Image preview is a convenient way to browse large or multiple images on small displays. However, current signal-level image resampling algorithms may remove many features of interest in the preview image. In this paper, we propose perceptual image preview which retains more perceptual features such that users can inspect features of interest by viewing the preview image only and without zooming in. This technology has two components, structure enhancement and perceptual feature visualization. Structure enhancement enhances the image structure while suppressing subtle details using a gradient modulation method, thus making the succedent perceptual features more apparent. For perceptual feature visualization, features of interest detected in the picture is visualized on the structure enhanced preview image. We demonstrate with two *examples* of most commonly used image quality features, image blur and noise. The effectiveness of the proposed method is validated by experimental results.

Keywords Image preview · Perceptual features · Structure enhancement · Perceptual feature visualization · Blur measurement

Communicated by Changsheng Xu.

L. Wan (✉) · T.-T. Wong
The Chinese University of Hong Kong, Shatin,
Hong Kong SAR, People's Republic of China
e-mail: lwan@cse.cuhk.edu.hk

W. Feng · Z.-Q. Liu
City University of Hong Kong, Kowloon,
Hong Kong SAR, People's Republic of China

Z. Lin
Microsoft Research Asia, Beijing,
People's Republic of China

1 Introduction

Image preview is a technology that shows a small version of a large image on a small display of a portable device like a digital camera or a mobile phone, or shows a large collection of image thumbnails on an ordinary display like a computer monitor. It provides a convenient way to quickly access the content of a large image. Image preview is usually achieved by applying image resampling techniques [21]. Although different resampling methods [9, 11, 15, 16, 22] have been reported, they do not suffice all situations of image browsing. For example, when taking photos with a digital camera, people are often not sure whether they just had a good shot, because features of interest may be lost in the preview image due to great reduction in image size (e.g. a 5 megapixel camera takes images with $2,592 \times 1,944$ pixels and these images are displayed in the preview window at a size of 400×300 pixels). The similar situation exists when using a digital photo printer which supports image preview on the printer LCD. People may want to identify those good-quality photos before printing in order not to waste the photo paper and the ink.

On the other hand, scholars have proposed thumbnail generation techniques to display (important portions of) pictures in a small size [2, 4, 19, 20]. Many thumbnail generation methods do not retain the global picture view or may impair the image structure. They may also remove perceptual features in the resultant thumbnail. In order to inspect features of interest, zooming in and shifting the large image across the preview window is a common practice. However, by zooming in people often lose the global appearance of the large image and observe only a small portion. Furthermore, zooming and shifting is time consuming.

In this paper, we propose a new image preview methodology to retain the lost information when images are shrunk.

Fig. 1 **a** Shows a blurry picture with a resolution of $2,592 \times 1,944$ reproduced in 3R photo size. **b** Is the naïve preview image in a resolution of 400×300 on a 2 in. display. It looks without image blur. **c** Shows our perceptual preview image after structure enhancement and blur visualization



The basic idea is to investigate the large picture and superimpose the features of interest on the preview image. With our technology, users can easily identify more features of the original large image by inspecting the preview image alone and without zooming in. Figure 1 gives an example to show our idea. Figure 1a is an out-of-focus picture of a pixel size $2,592 \times 1,944$. Its 400×300 naïve preview image (Fig. 1b), i.e., one computed by simple down-sampling, looks rather sharp. In comparison, our perceptual preview image (Fig. 1c) reveals that the original image is not well focused. To our knowledge, our work is the first attempt focusing on perceptual image preview. Although graphical representation can also indicate the existence and the extent of perceptual features, as one graphical representation usually accounts for only one feature, it may be too crowded on the small display if several perceptual features are to be considered. Moreover, such graphical representation has difficulty in visualizing spatially variant features, e.g., spatially variant blur or the position of red eyes. In short, graphical representation may be useful for a few globally homogeneous features only, while our perceptual preview image is *straightforward information visualization*, especially when the features are spatially variant.

The proposed framework has two parts, image structure enhancement and perceptual feature visualization (see Fig. 2). The first part highlights image structure and reduces subtle details in the preview image so that the succedent features of interest can be easier detected. More specifically, we perform non-linear modulation in the gradient domain and reconstruct the structure enhanced preview image by solving a Poisson equation [18]. The computation is fast since the preview image usually has a small size. The structure enhanced preview image itself can serve as a better preview image in many applications. For example, to organize an image album, leaving out subtle details may speed up classification. In ordinary photographing or medical imaging, people may also want a clear image structure in the preview image at the beginning.




Fig. 2 Our framework of perceptual image preview. The structure enhanced preview image is first generated by adaptive gradient modulation. Then the perceptual features detected in the original image is visualized into the structure enhanced preview image to produce the final perceptual preview image. In our current implementation, we take noise and blur as the examples of perceptual features. Most of the traditional image preview techniques output the naïve preview image only

The second part visualizes features of interest detected in the original large image. Note that it is unnecessary to faithfully reproduce the perceptual features, since our purpose is to *notify* the user of some information that exists in the original image but is lost in the naïve preview image. In this paper, we take two most common features of image quality, image blur and noise, as *examples* to demonstrate our methodology. In addition, we propose an efficient blur measurement which depends on the investigation of the gradient field. We also develop a simple yet functional method to retain noise shown in the original image.

Our technology is validated effective through several experiments, including subjective user evaluation.

2 Related work

Research on perception-based image processing has been reported in recent years. The philosophy is that signal-level fidelity is not the exclusive objective of processing. Perception-based image processing also cares about what people are interested in and how the resultant image can help people in specific tasks. For example, people often use the pseudo-color to highlight tumors in medical images so that

abnormal regions are easily detectable for human, or visualize the temperature fields so that such abstract and complex quantities are easily comprehensible to human. Another example is that Gooch et al. [8] converted color images into grayscale ones not by naively computing the luminance of every pixel. The information of color difference between adjacent regions is also kept by adjusting the contrast between them.

Traditional image preview usually employs image resampling techniques [9–12, 14–17, 22]. The simplest downsampling method is the decimation, which generates noticeable artifacts. The aliasing can be reduced by using anti-aliasing techniques, i.e., low-pass filtering the images before decimation, where the low-pass filters can be box functions, bilinear, bicubic [11], higher order splines [15, 22], or Gaussian [9, 10, 14]. However, downsampling is not necessarily related to low-pass filtering explicitly [16]. All these image downsampling methods are more or less signal-level processing. They are not designed to retain perceptual features. As a result, features of interest may be lost in the shrunk image.

Thumbnail generation also displays large images in a small size. Most existing algorithms [2, 4, 19, 20] focus on displaying objects of interest in the thumbnail. Chen et al. [4] dynamically selected important regions out of the large image for various display sizes based on the attention model. Suh et al. [20] detected an optimal cropping rectangle in the image so that it contains the most important portions. With the cropping-based methods, people are able to catch the important objects at the first sight, but may fail to understand the image because the global picture is lost. It is actually not natural for people to preview on imaging devices in this manner.

The retargeting approach [19] minimizes less important spaces between salient objects. It shows the overall structure of the image while maximizing the objects that receive attention. But this method may require user intervention and may impair the image structure. Although the seaming carving approach [1] is able to generate thumbnails which maintain the global view, it is slow to resize a big picture to a small preview size. In addition, these existing thumbnail generation methods may also remove perceptual features in the picture.

There has work to quantitatively measure perceptual features of an image. For example, Caviedes and Oberti [3] proposed a no-reference sharpness metric based on the local frequency spectrum around the image edges. Marziliano et al. [13] measured the blur metric by analyzing the spatial spread of the edges in an image. These quantities can be displayed as texts beside (or be overlaid upon) the naïve preview image. However, as not all people are experts in photography, it would be more straightforward and intuitive if we could visualize these quantities in the preview image directly. Furthermore, preview image can convey more information than a single quantitative value. For instance, in the preview image we could retain different amount of blur in different regions.

3 Perceptual image preview

There are various perception-related features in a picture, for example, image blur, blooming effects, red eyes, half-/closed eyes, etc. They are useful for users to master some characteristics of the picture, e.g., to determine whether the picture is of high quality. Many perceptual features, however, become less detectable in the small preview image. One possible way to retain perceptual features is to merge them into the image resampling process, as done in [12, 17] which preserve the strong edges in image magnification. But this scheme would make the resampling process rather complex and is inflexible to support new perceptual features.

In this paper, a general framework is proposed (Fig. 2). The basic idea is to investigate the original image and superimpose the features of interest with the preview image. Each perceptual feature can be processed independently. Clearly, our framework is a general methodology which is flexible to include a number of different perceptual features. In the following, we discuss respective steps.

3.1 Image structure enhancement

Structure enhancement, the first step in our framework (Fig. 2), provides a clear image structure by strengthening salient edges and flattening weak details. Since salient edges often separate objects, this step actually increases the inter-object contrast and reduces the intra-object contrast. Therefore, the image structure becomes more visually apparent to users, and then the perceptual features added in subsequent steps are more detectable for users.

We perform structure enhancement in the image gradient field. As salient edges have large gradient magnitudes and weak details have small values, the problem becomes to increase large gradients and reduce small gradients. Denote the gradient field of the naïve preview image I_n to be G , where $G = \nabla I_n$ and ∇ is the gradient operator. Let I_s be the structure enhanced preview image, and G' is the modulated gradient field. By solving a Poisson equation [18]:

$$\nabla I_s = \text{div}(G'), \quad \text{with } I_s|_{\partial\Omega} = I_n|_{\partial\Omega}, \quad (1)$$

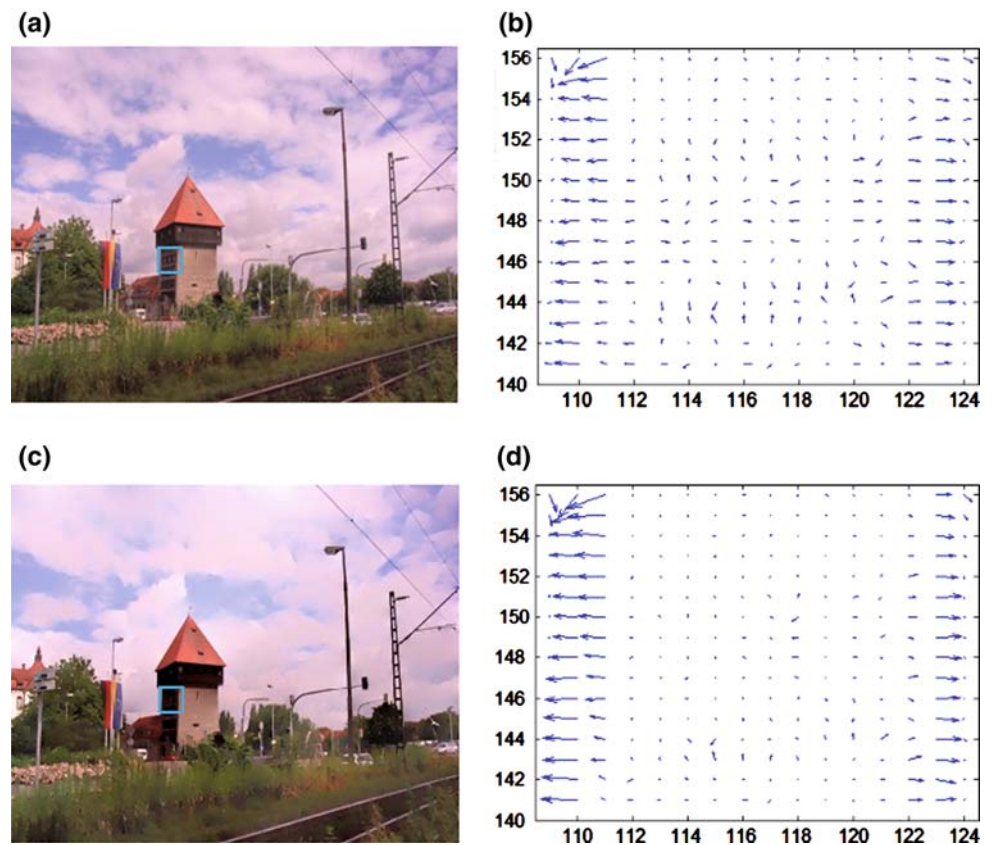
where $\partial\Omega$ is the boundary of image space Ω , we can reconstruct I_s from G' . By this way, the manipulation in the gradient field is reflected in the resultant image.

We use a sigmoid function to manipulate the gradient magnitude,

$$|G'_i| = \frac{\alpha}{1 + \exp(-k(|G_i| - \beta))}, \quad (2)$$

where G_i is the gradient for pixel i in I_n and G'_i is the adjusted gradient in I_s . Note that we maintain the original gradient direction, since local direction manipulation may tamper the consistency of gradient field and may result in unexpected

Fig. 3 **a** is the naïve preview image and **b** is the gradient field inside the region indicated by a *box* at the side of the tower; **c** is the structure enhanced preview image and **d** the gradient field inside the same region. The gradient directions are unchanged, while the gradient magnitudes are increased for salient edges and reduced for trivial details. The structure enhanced image presents a clearer image structure



artifacts. The coefficient α controls the maximum value of gradient magnitudes. The parameter k affects the modulation slope. The larger k is, the more large gradient magnitudes are magnified and the more small gradient magnitudes are suppressed. The parameter β defines the threshold to differentiate large and small gradient magnitudes. In our experiments, we set $\alpha = 2$ and $k = 1.5$ empirically. In order to preserve locally salient tiny details, we choose β adaptively:

$$\beta = \begin{cases} \beta_1, & \text{for } \beta_1 < \beta_g, \\ \beta_g, & \text{otherwise,} \end{cases} \quad (3)$$

where β_g is a global threshold, and β_1 a local threshold that varies in the image space. We evaluate β_1 as the average gradient magnitude in the pixel neighborhood (weighted by a Gaussian kernel) minus a constant (this offset affects the output gradient magnitude, empirically set as 3.5). β_g is computed in the same way over the whole image. In fact, Eq. (3) respects strong edges and favors weak yet locally salient edges. Due to the small size of the reduced image, the computation is quite fast.

For a color image, we convert the naïve preview image to YUV color space and perform the modulation in the luminance channel. The chrominance components keep intact which guarantees that the image color does not change significantly. As an example, Fig. 3c and d show the structure enhanced image and the gradient field inside the box region.

Compared to their counterparts, the increase of large gradient magnitude strengthens the important edges, and the reduction on small gradient magnitudes suppresses the trivial details. Figure 4 presents more results.

3.2 Perceptual feature visualization

As discussed before, there are various perceptual features that can be of interest in different applications. Currently we can only choose some to realize our framework. Here we use image blur and image noise as two examples, which are the most common features of image quality when using digital camera.

3.2.1 Image blur visualization

Image blur occurs when a relative motion happens between the camera and the subjects of the shot, for example, a shaky camera or moving objects, or when the subjects are out of focus. Blur may appear across the whole image or just around some regions. However, unless the blur is significant, the naïvely reduced image at a low resolution may still look rather sharp (e.g., Fig. 1b), and it is hard to tell whether the picture is blurred by inspecting the preview image only.

Our basic idea for blur visualization is to detect the blur locally in the large image, and blur the corresponding region

Fig. 4 Examples of structure enhancement. The *first row* shows two naïvely reduced images, with many details of trivial interest. The *second row* shows the structure enhanced images. Our method flattens subtle edge regions, while boosts strong edges




Fig. 5 The flowchart of image blur visualization

in the preview image according to the detected amount of blur (Fig. 5). Because the blur in edge regions is more obvious for users, even when the whole image are blurred, we only perform blur manipulation around edge regions in the preview image. As shown in Fig. 5, we first apply edge detection on the reduced image. Here the reduced image is the result after structure enhancement. Then for each edge pixel i we estimate its blur degree. Finally the operation, blur synthesis, renders the blur in the neighborhood of pixel i according to the amount of the pixel blur degree.

In our approach, the most important is blur estimation which quantitatively measures the blur degree. Most of the existing quantitative blur measurements [3, 13] give a global evaluation. Our work, on the other hand, requires a local measurement, since different portions in the image may undergo different blurriness. Although there are intensive work on spatially variant blur estimation [5, 7], many algorithms involve optimization which is time consuming for image preview, and some algorithms require multiple source images.

Since image blur is more detectable in salient edge-regions¹ than intensity-continuous regions, we investigate statistics of gradient field in edge regions. We observe that blurry edge-regions usually have smaller variances among gradient angles. Based on this, we propose the following blur metric:

$$B_i = \eta \cdot \exp\{-\text{var}(A_i)^\tau\}, \tag{4}$$

where $\text{var}(A_i)$ is the variance of gradient angles in the edge-region A_i . The parameters η and τ control the estimated amount of image blur. From our experiment, we empirically set $\eta = 4.5$ and $\tau = 2.5$.

Given the estimated blur degree B_i , we then blur the reduced image using a Gaussian kernel with its sigma set as B_i . A large value of B_i corresponds to a large Gaussian kernel, and a small value means to use a small kernel. Figure 1c shows the blur visualized preview image. Compared to the naïve preview image (Fig. 1b), which appears rather sharp, the blur visualized preview image intuitively notifies the viewer that the original image is not focused.

We next validate the blur estimation through a simple experiment. A blurry image is generated by smoothing a clear image with a Gaussian kernel $G(\sigma)$. By increasing σ , a series of blurry images are generated. We then compute the average of blur degrees for each blurry image. Figure 6 plots the curve of the average blur degree with respect to σ . As shown, the

¹ Each pixel i in the preview image corresponds to a small region A_i in the original image. If i is an edge pixel, A_i is an edge-region.




Fig. 6 Estimated blur degree vs. Gaussian kernel size σ

average blur degree B increases rapidly for small σ values, and becomes saturated for big σ values. The proposed blur metric is compatible with blur visualization. For instance, when the picture is subject to significant blur, blur is already kept in the naïve preview image to certain extent, and then blur visualization is to strengthen the blurriness. On the other hand, our blur metric is sensitive to medium or slight blur such that medium or slight blur can also be visualized in the preview image.

We further shoot the same scene with two different camera settings. One produces a blurred image (Fig. 7a). The other results in a sharp image (Fig. 7d). As shown, the naïvely reduced images of the sharp image (Fig. 7e) and the blurred image (Fig. 7b) both look quite sharp. In contrast, using our method, the reduced image of the original sharp image (Fig. 7f) still looks sharp, while the one for the blurred image (Fig. 7c) exhibits obvious blur.

3.2.2 Image noise visualization

Image noise usually comes from image sensors during exposure or A/D conversion, and looks like color grains and scatters across the whole image. It becomes less visible after image reduction to a small size. Noise visualization helps users to be aware of severe image noise when previewing the image, and then users could adjust camera settings to get better shots.

Digital camera usually produces three types of noise, random noise, fixed noise and banding noise. Random noise is characterized by intensity and color fluctuations and it is most influenced by a high ISO value. Fixed noise appears in unique patterns and generally occurs in long exposure. Banding noise is camera-dependent which is introduced by the camera when reading data from image sensors. It is the most visible at high ISO speed and may increase at certain

white balance modes. For simplicity, we assume that noise has additive property and is evenly distributed over the image. Although this assumption may not count all situations, it is sufficient for noise visualization.

Figure 8 illustrates the process for noise visualization. First, the noise is crudely estimated from the large image using a noise detection method and stored as a source noise map N_s . Based on the assumption of even distribution, we prepare the destination noise image N_d , which is in the same size as the preview image, by uniform sampling from N_s . The noise visualized preview image is then obtained by adding the destination noise image to the reduced image. Note that the reduced image can be the naïve preview image, the structure enhanced preview image, or other preview images.

Since we just want the noise on the reduced image look similar to that on the large image, it is unnecessary to precisely estimate noise over the whole large image which is often difficult and time-consuming. Furthermore, image noise is easier to detect in a uniform region which contains few salient edges so high frequency details in the region can be taken as noise. We only perform noise detection in a small color-uniform region Ω_s in the large image. As a speed-up trick, we search color-uniform region Ω in the reduced image and locate its position Ω_s in the large image. The search for color-uniform regions is efficient and reliable as the reduced image is less noisy and in a much smaller size. More specifically, we divide the reduced image into non-overlapped blocks (of a size 12×12), and select the block Ω with the smallest color variance. Then we apply discrete stationary wavelet transform [6] on the corresponding Ω_s (which is in the large image) to estimate noise map N_s . The computation is fast as the region size is rather small.

Given the detected noise map N_s , we can apply texture synthesis techniques to generate the destination noise map N_d from N_s . In practice, we generate a noise map N_d by uniform sampling from N_s . Although N_d may not exactly match the noise distribution in the large image, it offers users a similar visual experience as in the original one (see Fig. 9). Then the final image I_f is,

$$I_f = I_d + \gamma \cdot N_d, \quad (5)$$

where I_d is the reduced image; γ is a parameter to control how salient the noise is visualized.

Figure 9 shows noise visualized results for a noisy image and a clean image respectively. The middle column are the naïve preview images, and the right column are the noise visualized preview images. The blow-ups on the left column show the noise levels in the two original images. Though one original image is noisy and the other is almost noise-free, both naïvely reduced images look clean and noise-free. In comparison, our noise visualized preview images display the difference in the noise levels. For the clean image, our method does not introduce obvious forgery noise.

Fig. 7 Blur visualization on a blurred image (a) and a sharp image (d) is shown. Both images are taken from the same scene with two sets of camera parameters. b and c are the naïve and the perceptual preview image of a, while e and f are the naïve and the perceptual preview image of (d), respectively. Note that with our method, the preview image remains sharpness for the original sharp image, while becomes blurry for the blurred image. Additionally, our method is able to detect local blur (ellipse regions) in the large image (d) and retain it in the preview image (f)





Fig. 8 The flowchart of noise visualization

4 Experimental results

We now show an example of how a final perceptual preview image looks like in a small display if multiple perceptual features are considered simultaneously. To do so, the steps introduced separately in the previous sections must be combined. The structure enhancement is done first. Then the blur is visualized. Finally, the noise is superimposed.

Figure 10a is a large picture shot at a resolution of $2,592 \times 1,944$. The blow-ups Fig. 10b and c show that the picture has obvious image noise and is subjected to image blur (e.g., near


candles). Figure 10f is the final perceptual preview image. We also show the intermediate results after structure enhancement in Fig. 10d and image blur visualization in Fig. 10e. The naïvely reduced image can be found in Fig. 10g. It is obvious that the naïvely reduced image loses clues of both image blur and image noise. Our perceptual preview image, on the other hand, has image blur and noise directly visible and conveys a similar impression as the original picture does.

We also conduct subjective experiment to validate the effectiveness of our method. In the subjective experiment we use a set of 30 color pictures of resolution $2,592 \times 1,944$. The preview images are in a size of 400×300 . In the data set, 10 pictures are well-shot, while another 20 pictures suffer from image blur or severe image noise, as the example in Fig. 10a. For each picture, both naïve preview image and perceptual preview image are provided. To reduce bias, preview images are randomized in showing for different pictures. Ten volunteers (six males and four females) with normal sights were asked to score the images on the scale of 5. Here 1 means that the preview image does not give a similar visual

Fig. 9 Examples of image noise visualization. The *first row* shows the results for an original noisy image. The *upper blowup* is one uniform region in the noisy image, while the *lower blowup* shows the detected noise map in this region. The *central image* is the result after naïve image reduction. The *right image* adds the noise to the structure enhanced image according to the detected noise map. The *second row* shows the results for an originally clean image. The result after noise visualization still looks clean as expected



Fig. 10 An example of perceptual preview image that incorporates all features of interest. The picture in (a) contains both blur (see the blow-up in (b)) and noise (see the blow-up in (c)). (d) is the preview image after structure enhancement. The weak details have been removed. We then feed (d) for blur visualization, and get the result as (e). Notice the blur around the candles. Finally we perform noise visualization and produce the final perceptual preview image shown in (f). Compared to the naïvely reduced image (g), the final perceptual preview image (f) displays more obvious noise and blur



impression as the large picture, while 5 means the preview image looks like the large picture, regarding image blur and noise. During the test, the subjects were required to view two preview images at first and then look at the corresponding large picture before scoring the perceptual similarity.

Let the subjective rating per person be his/her average score for a given type of preview images. Figure 11 shows the rating curves for naïve preview images (bottom) and perceptual preview images (top). Except for one subject, all subjects

think perceptual preview images convey more perceptual information than naïve preview images do.

We then count the subjective rating per picture as the average of ten scores assigned to one picture for the same type of preview images. As shown in Fig. 12, the scores of perceptual preview images are always above three for both well-shot pictures and degraded pictures. Naïve preview images give better perception in the case of well-shot pictures. However, they hardly convey the information of image




Fig. 11 Subjective rating per person is shown for both naïve preview images and perceptual preview images. Among ten volunteers, nine give higher scores to perceptual preview images




Fig. 12 Subjective rating per picture for the test images is shown for both naïve preview images and perceptual preview images. The first 10 pictures are well-shot pictures, and the rest 20 pictures are either blurred or noisy

blur and noise in the degraded images. We also examine why subjects gave higher rates to naïve preview images for well-shot pictures. Some viewers think perceptual preview images have higher contrast and are clearer, but sometimes the high contrast makes the preview image look a bit artificial in bright regions.

5 Conclusion

In this paper, we address a novel problem of perceptual image preview. It aims to convey more perceptual information in the preview image such that users can inspect features of interest like the image quality by viewing the preview image only. For this purpose, a general framework is proposed which consists of image structure enhancement and perceptual feature visualization.

In image structure enhancement, we present an adaptive gradient modulation algorithm such that the image structure becomes clearer. Hence, the perceptual features added later are more apparent to users. Perceptual feature visualization detects perceptual features in the original large image and visualizes them in the preview image. We demonstrate this idea using two features of image quality, namely blur and noise, as examples. Besides visual quality comparison, subjective test is conducted to validate the effectiveness of our technology.

As a basic framework, our perceptual image preview does not restrict the individual algorithms for respective features. For example, we may evaluate different function for gradient modulation; other blur or noise estimation techniques can be employed in the framework. Furthermore, our framework can be extended to support more perceptual features, for example, red eyes or eye blinking, in the preview image.

Acknowledgments This project is supported by the Research Grants Council of the Hong Kong Special Administrative Region, under RGC Earmarked Grant (Project No. CUHK 4170107).

References

1. Avidan, S., Shamir, A.: Seam carving for content-aware image resizing. *ACM Trans. Graph.* **26**(3), 10 (2007)
2. Burton, C., Johnston, J., Sonenberg, E.: Case study: an empirical investigation of thumbnail image recognition. In: *Proceedings of Information Visualization*, pp. 115–121 (1995)
3. Caviedes, J., Oberti, F.: A new sharpness metric based on local kurtosis, edge and energy information. *Signal Process. Image Commun.* **19**(2), 147–161 (2004)
4. Chen, L.Q., Xie, X., Fan, X., Ma, W.Y., Zhang, H.J., Zhou, H.Q.: A visual attention model for adapting images on small displays. *ACM Multimed. Syst. J.* **9**(4), 353–364 (2003)
5. Chiang, M.C., Boulton, T.E.: Local blur estimation and super-resolution. In: *CVPR '97: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, pp. 821–826 (1997)
6. Donoho, D.L.: De-noising by soft-thresholding. *IEEE Trans. Infor. Theory* **41**(3), 613–627 (1995)
7. Elder, J.H., Zucker, S.W.: Local scale control for edge detection and blur estimation. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(7), 699–717 (1998)
8. Gooch, A.A., Olsen, S.C., Tumblin, J., Gooch, B.: Color2gray: salience-preserving color removal. *ACM Trans. Graph.* **24**(3), 634–639 (2005)
9. Greene, N., Heckbert, P.: Creating raster omnimax images from multiple perspective views using the elliptical weighted average filter. *IEEE Comput. Graph. Appl.* **6**(6), 21–27 (1986)
10. Heckbert, P.: *Fundamentals of texture mapping and image warping*. Master’s thesis, CS Division, UCB/CSD 89/516, U.C. Berkeley (1989)
11. Keys, R.G.: Cubic convolution interpolation for digital image processing. *IEEE Trans. Acoust. Speech Signal Process* **29**, 1153–1160 (1981)
12. Li, X.: *Edge directed statistical inference with applications to image processing*. Ph.D. thesis, Princeton University, Princeton (2000)

13. Marziliano, P., Dufaux, F., Winkler, S., Ebrahimi, T.: A no-reference perceptual blur metric. In: Proceedings of the International Conference on Image Processing, vol. 3, pp. 57–60 (2002)
14. McCormack, J., Perry, R., Farkas, K.I., Jouppi, N.P.: Feline: fast elliptical lines for anisotropic texture mapping. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, pp. 243–250 (1999)
15. Meijering, E.H.W., Niessen, W.J., Viergever, M.A.: Piecewise polynomial kernels for image interpolation: a generalization of cubic convolution. In: Proceedings of IEEE International Conference of Image Processing, pp. 25C28 (1999)
16. Muñoz, A., Blu, T., Unser, M.: Least-squares image resizing using finite differences. *IEEE Trans. Image Process.* **10**(9), 1365–1378 (2001)
17. Muresan, D., Parks, T.: Adaptively quadratic (aqua) image interpolation. *IEEE Trans. Image Process.* **13**(5), 690–698 (2004)
18. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Trans. Graph.* **22**(3), 313–318 (2003)
19. Setlur, V.R.: Optimizing computer imagery for more effective visual communication. Ph.D. thesis, Northwestern University, Evanston (2005)
20. Suh, B., Ling, H., Bederson, B.B., Jacobs, D.W.: Automatic thumbnail cropping and its effectiveness. In: *UIST '03: Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology*, pp. 95–104 (2003)
21. Thévenaz, P., Blu, T., Unser, M.: *Image Interpolation and Resampling*. Academic, New York (2000)
22. Unser, M.: Splines: a perfect fit for signal and image processing. *IEEE Trans. Signal Process.* **16**, 22–38 (1999)