Quality-Preserving Image Downsizing

Matthew Trentacoste Rafal Mantiuk Wolfgang Heidrich

University of British Columbia

a place of mind

Abstract

High-quality digital camera designs without through-the-lens optical viewfinders are becoming increasingly common. Historically relegated to low-end point-and-shoot models, digital-only viewfinders have been incorporated into compact models such as the Micro Four-Thirds system.

However, the image quality of a digital viewfinder is considerably lower than that of a through-the-lens optical system. While the sensor may be capable of capturing 10 or 20 megapixels, the screen of the viewfinder is typically constrained to resolutions under 1 megapixel. The limited resolution makes it impossible to discern all the small details of the captured image. Small blurs and noise that are present in the full-size image can render the image unusable for certain tasks, yet these artifacts may be too small to be discernible in the downsampled version shown on the camera viewfinder.

We preserve spatial detail such as blur and noise while creating the thumbnail image for presentation in a digital viewfinder. We present an efficient method for enhancing the artifacts that may be present in both still images and live preview video streams, making them more identifiable, and thus decreasing the chance of capturing low-quality images. The result more accurately reflects the original image, allowing faster and more accurate assessments of quality in the field or when browsing thumbnails.

We analyze the image to detect artifacts present but too small to be visible on a digital viewfinder, and modify the downsampled thumbnail image such that they are large enough to be discernible, as seen in the example below. The three images show a comparison of our quality-preserving downsampling algorithm with a regular thumbnail. Note that the cat's face is the only part of the image that is in focus, a feature that is preserved by our algorithm.

Conventional thumbnail

Quality preserving thumbnail

Related Work

Significant work exists regarding the estimation of blur and noise present in images. Other methods exist for resizing images in manners more representative than simple Gaussian filtered resampling.

Work on estimating the amount of blur present exists in numerous forms. Examples include classification of defocus and motion blur by Liu et al. [2008] and scale-space methods such as Elder and Zucker [1998] and Bae and Durand [2007]. Like Bae and Durand, our blur enhancement produces an image with lower resolution approximation at significantly reduced computational cost.

Similar to our work, Samadani et al. [2010] developed a method for amplifying specific artifacts present in full-size images to be visible in thumbnail images. However, they do not specifically ensure that all pertinent artifacts are amplified enough to be visible. Additionally, their method does not extend to video and has trouble with some textured areas.

Research also exists on the generation of more informative thumbnails. Suh et al. [2003] and Santella et al. [2006] have proposed automatic cropping methods for selecting the most salient region of the image to highlight in a thumbnail.

Approach

In order to make the blur and noise artifacts more apparent in a downsampled thumbnail image, we first must analyze the full-size image to determine how much of each attribute is present. Given an estimate of the significance of both blur and noise, we amplify these characteristics so they remain visible in the generated thumbnail. This process is particularly important when the strength of an artifact is large enough to be of visual significance, but small enough to be invisible in the thumbnail.

The outline of our algorithm is as follows:

1. We identify the strength of both blur and noise present in the full-size image.
2. Based on our blur estimate, we determine how much blur needs to be added to each region of the thumbnail.
3. We downsample the image, and add the required amount of blur.
4. For each image region, we model how much of the original noise is lost due to both the downsampling operation and the subsequent blur enhancement.
5. Finally, we reintroduce the required amount of noise for each pixel.

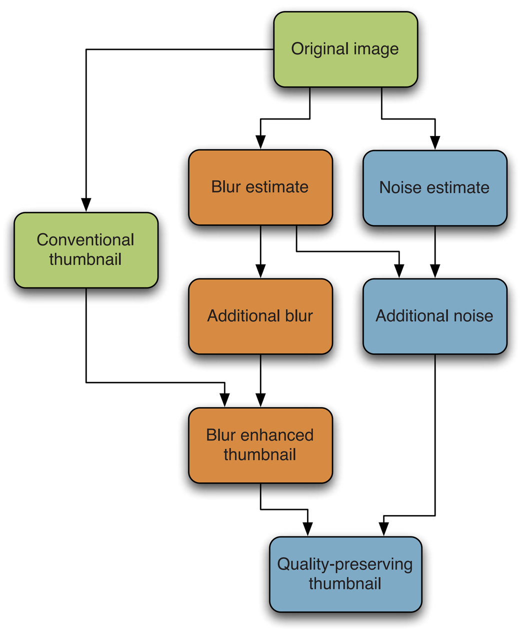

Blur Enhancement

In order to retain the detail of in-focus regions while further removing detail in out-of-focus regions, the blur estimation routine determines a spatially-variant estimate of blur across an image. We employ a technique derived from the work of Hu and de Haan [2006], which uses the convolution formula for Gaussian distributions:

$$g(x,\sigma_1) \otimes g(x,\sigma_2) = g\left(x, \sqrt{\sigma_1^2+\sigma_2^2}\right)$$

The image already has some unknown amount of blur σ present at each pixel. Blurring that image by some σ_a results in an image with blur of magnitude $\sqrt{\sigma^2+\sigma_a^2}$ at each pixel. This process is repeated for another $\sigma_b > \sigma_a$, and the ratio r of differences between the original image I and two blurred versions is used to compute the resulting blur map:

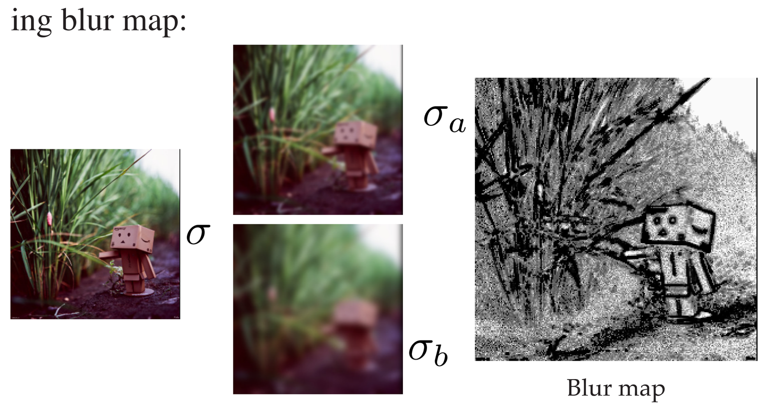

Blur map

Since there is blur already present in the downsampled image, we solve for the amount of additional blur required to achieve the increased amount of blur we desire.

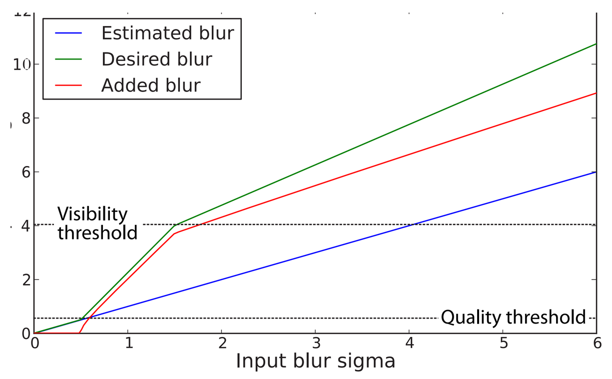

We downsample the original image to produce the normal thumbnail and create the blur enhanced thumbnail from that image. The final thumbnail pixel is the result of linearly interpolating between the two filtered images closest to the desired increased blur radius.

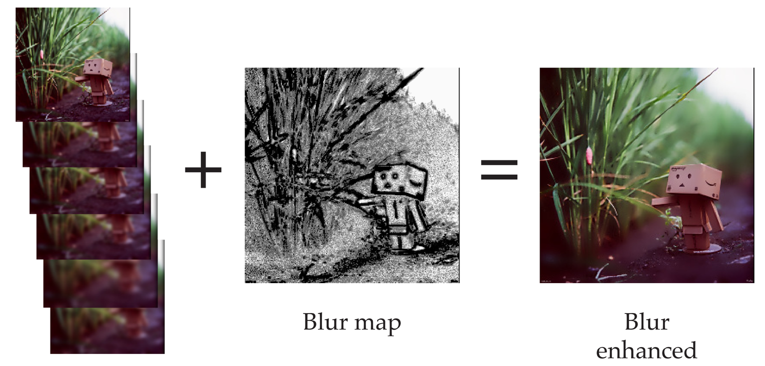

Blurred thumbnails Blur map Blur enhanced

In order to avoid contrast inversion around edge boundaries, we use the bilateral filter to produce the set of blurred images. Consider the case of two areas of uniform color on either side of a sharp edge. Pixels near the edge will not be blurred and will retain their original contrast.

However, if the blur function does not respect boundaries, pixels further from the edge will be blurred with pixels on the opposite side of the edge, reducing their contrast and giving the appearance of ringing artifacts around the edge, as seen below:

Gaussian filter Normal thumbnail Bilateral filter

Noise Enhancement

To estimate the noise, we use the block-based approach of Ibenthal [2007]. Each channel c of the image is divided into N × N blocks and the mean μ_b and standard deviation σ_b are computed. For each channel, the value of f_{NL} for each intensity level I is taken to be the minimum σ_b across all blocks with $\mu_b = I$.

While the noise in the full-size image only varies with intensity, the combination of downsampling and bilateral filter have altered the distribution of the noise in a spatially variant manner. The down-sample averages intensity values over an N × N block, while the bilateral filter computes a weighted sum over a local neighborhood of pixels which varies with the amount of blur added.

These two operations yield the following noise model:

$$\sigma_{It} \propto \left(\frac{1}{N^2}\right)\left(\frac{1}{P}\right)\sqrt{\sum_{j\in F} g_j \sum_{k}^{N^2} I_k},$$

where N is the block downsample size and P = sum(g_j) for Gaussian weights g_j in some neighborhood F.

We make two approximations to simplify this equation. First, since the mean value of the N x N block is I, we can approximate the N^2 samples present in sum(I_k) with N^2 samples of I. Second, the spatial sigma of the bilateral filter causes I only to be averaged with adjacent pixels of similar intensity.

We can approximate the respective I_j for each g_j with I:

$$\sigma_{It} \propto \left(\frac{1}{N^2P}\right)\sqrt{\sum_{j\in F} g_j N^2 I} = \left(\frac{N}{N^2P}\right)\sqrt{PI}.$$

Finally, we simplify and substitute the measured noise estimate f_{NL} for the approximate response:

$$\sigma_{It} = \frac{f_{NL}}{N\sqrt{P}}.$$

The exact value of P depends on the size and weights of the neighborhood F, which we determine from the value of the standard deviation of the additional blur corresponding to that pixel. Finally, we scale Gaussian noise by the blur map and add it to the blurred thumbnail.

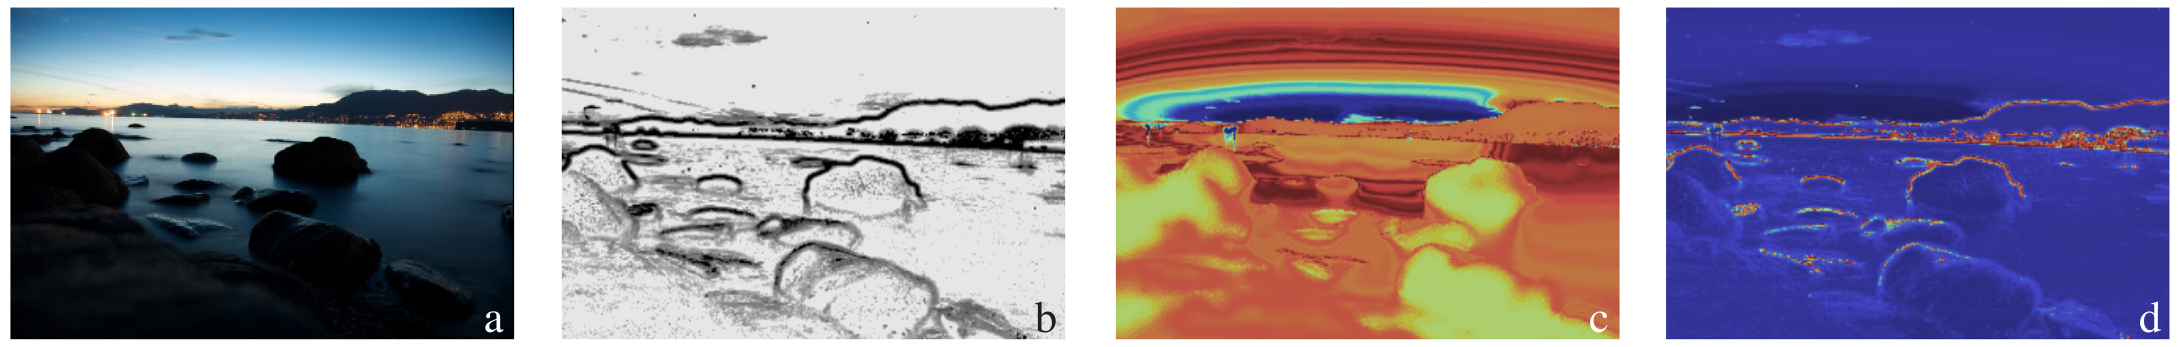

These four images depict the original image (a), estimated blur map (b), estimated noise map (c) and final adjusted noise map (d). Note the noise variance decreases more in regions with larger radii blurs applied to them.

Results

The following images show the results of our algorithm for a number of images. The 1st image demonstrates blur enhancement, the 2nd image demonstrates noise enhancement and the 3rd and 4th images demonstrate both blur and noise enhancement.

Conventional thumbnail

Quality preserving thumbnail

Conventional thumbnail

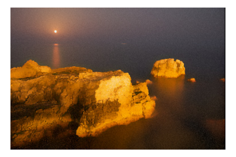

Quality preserving thumbnail

Conventional thumbnail

Quality preserving thumbnail

Conventional thumbnail

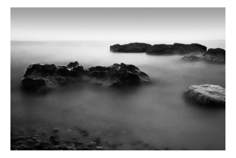

Quality preserving thumbnail

References

1. BAE, S., AND DURAND, F. 2007. Defocus magnification. Computer Graphics Forum 26, 3, 571–579.
2. ELDER, J., AND ZUCKER, S. 1998. Local scale control for edge detection and blur estimation. IEEE PAMI 20, 7, 699–716.
3. HU, H., AND DE HAAN, G. 2006. Low cost robust blur estimator. In Proc. IEEE Image Processing, 617–620.
4. IBENTHAL, A. 2007. Image sensor noise estimation and reduction. Tech. Rep. ITG 3.2, University of Applied Sciences and Arts Hildesheim/Holzminden/Gottingen.
5. SAMADANI, R., MAUER, T., BERFANGER, D., CLARK, J. 2010. Image Thumbnails That Represent Blur and Noise. IEEE Transactions on Image Processing 19(2): 363-373
6. SANTELLA, A., AGRAWALA, M., DECARLO, D., SALESIN, D., AND COHEN, M. 2006. Gaze-based interaction for semi- automatic photo cropping. In Proc. CHI, 771–780.
7. SUH, B., LING, H., BEDERSON, B. B., AND JACOBS, D. W. 2003. Automatic thumbnail cropping and its effectiveness. In Proc. UIST, 95–104.

DOLBY

ICICS CONNECTING KNOWLEDGE