

Lucky averaging: quality improvement of adaptive optics scanning laser ophthalmoscope images

Gang Huang,* Zhangyi Zhong, Weiyao Zou, and Stephen A. Burns

Optometry School, Indiana University Bloomington, Bloomington, Indiana 47405, USA

*Corresponding author: gghuang@indiana.edu

Received August 1, 2011; accepted August 8, 2011;

posted August 31, 2011 (Doc. ID 151683); published September 22, 2011

Adaptive optics (AO) has greatly improved retinal image resolution. However, even with AO, temporal and spatial variations in image quality still occur due to wavefront fluctuations, intraframe focus shifts, and other factors. As a result, aligning and averaging images can produce a mean image that has lower resolution or contrast than the best images within a sequence. To address this, we propose an image postprocessing scheme called “lucky averaging,” analogous to lucky imaging [J. Opt. Soc. Am. **68**, 1651 (1978)] based on computing the best local contrast over time. Results from eye data demonstrate improvements in image quality. © 2011 Optical Society of America

OCIS codes: 100.0100, 100.2000, 100.2980.

Adaptive optics scanning laser ophthalmoscopy (AOSLO) provides high-resolution *in vivo* retinal images and has become a reliable tool for studying the structure and function of the retina [1]. AOSLO systems typically provide real-time imaging and thus record images over time. However, in almost all real-time image sequences, image quality varies slightly over both space and time. That is, local regions have higher contrast in one frame than in another. Based on the point-scanning mechanism of AOSLO, the variance can be attributed to both time-related and space-related factors, most likely including dynamic changes in human eye aberrations such as the tear film, local retinal shape changes due to ocular pulse, accommodative fluctuations, and variations in the retinal shape that can interact with the above factors. In addition, images from subjects who have retinal defects such as detachments, drusen, and edema produce even more pronounced variations in the retinal shape, which interact to enhance these variations in image quality.

Conventionally, AOSLO raw images are postprocessed to decrease the impact of noise. Typically this is achieved by either manually or automatically choosing acceptable frames, omitting those with very large eye movements or blinks, then aligning to minimize the effect of both intraframe and interframe movements [2], and finally averaging the aligned frames. The more frames used for the averaging, the greater the increase in signal-to-noise ratio. On the other hand, since the averaging is frame based, the intraframe contrast variance is not altered. Therefore, the calculation of the mean can degrade the optimal image quality achieved from the image sequence. That is, the frame averaging is not taking the best advantage of the sequential nature of AOSLO imaging.

This problem is also true if recording views of the sky from a ground-based telescope. By choosing only the best frames, images are selected when the atmospheric turbulence along the line of sight is minimal. That is, the process takes advantage of lucky moments when turbulence has a minimal effect and has therefore been called “lucky imaging” [3]. In analogy, AOSLO is measuring the retina point by point, and its high-frequency pixel clock

makes it well suited for this lucky imaging approach, but now within a frame as well as between frames.

Thus, we propose a region-based averaging scheme we call “lucky averaging.” Image acquisition and alignment are performed in the normal manner described above. However, before averaging, the image sequence is analyzed region by region. Small areas in the image represent short exposure times and small spatial scales, greatly eliminating the intraregion variance in optical quality and enabling the process to capture those “lucky moments” when a small region’s image quality is at its best. The region quality is evaluated using an appropriate objective image quality metric and assigned to the central pixel of the region. Each pixel becomes the center of a region, and consequently, pixel quality maps are built up in sequence. The final image is remapped by only including the times when each pixel has the highest image quality as evaluated using the metric. The scheme is shown in Fig. 1.

To test this approach, cone photoreceptor images were acquired from different subjects and processed. The AOSLO system has been described previously [4]. The pixel clock was 30 MHz, and the frame rate was 30 Hz. The adaptive optics control loop operated at 15 Hz. Each image was 800 by 520 pixels, corresponding to a $2.79^\circ \times 1.78^\circ$ imaging field in the retina. Each image series consisted of approximately 100 frames acquired sequentially in ~ 3 s. Once aligned, a window of 9 by 7 pixels was created and swept across the image pixel

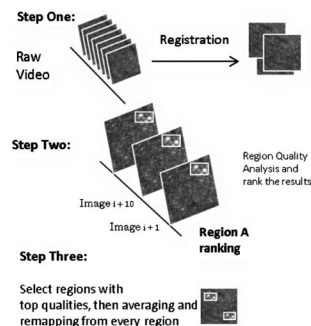


Fig. 1. Schematic information flow for lucky averaging.

by pixel to compute an image quality metric. A window of this size represented $450\ \mu\text{s}$ in time and about $70\ \mu\text{m}^2$ in space.

While any appropriate image quality metric can be used, for this Letter we implemented two examples—a sharpness metric and a gray level co-occurrence matrix (GLCM). The sharpness metric has been used for incoherent imaging in astronomy and synthetic radar image reconstructions [5]. It represents a sum of a nonlinear point transformation of the image and is generally defined as

$$s = \sum_{(x,y)} w(x,y) I(x,y)^\beta, \quad (1)$$

where x, y are pixel coordinates, β is the power of the pixel intensity, $w(x,y)$ is the weight for each pixel, and $I(x,y)$ is the pixel intensity.

In practice, we set β as 2 and $w(x,y)$ as 1. Normalization is also considered, as [6] indicates. The metric is finally defined as

$$S_{(x_0,y_0,\text{index})} = \frac{\sum_{(x_0-\Delta x \leq x \leq x_0+\Delta x, y_0-\Delta y \leq y \leq y_0+\Delta y)} I_{\text{index}}(x,y)^2}{\left(\sum_{(x_0-\Delta x \leq x \leq x_0+\Delta x, y_0-\Delta y \leq y \leq y_0+\Delta y)} I_{\text{index}}(x,y) \right)^2}, \quad (2)$$

where x, y are pixel coordinates and x_0 and y_0 represent the coordinates of central pixels of the sampling window. Δx and Δy represent half the sampling window width and height, respectively. The *index* represents the frame index number of the analyzed image within the sequence.

The second image metric is the GLCM, measuring how often different combinations of gray levels co-occur in an image region [7], computed as in Eq. (3). For a given Δd , which represents pixel separation between a pair of pixels, we compute $\text{GLCM}_{\Delta d}$ by examining pixel pairs along eight different directions from 0° to 315° with 45° increments, where the element $\text{GLCM}_{\Delta d}(i,j)$ represents the number of the pixel pairs with gray level pair (i,j) . The parameter d represents the two-dimensional position vector of the pixel in the image. We typically normalize the GLCM by dividing the whole matrix with the sum of all elements in the matrix. Once a GLCM has been constructed, several types of statistical analysis can be performed. Local contrast is one of these measures, taking the GLCM as contrast weights, computed as in Eq. (4), where α is the power of the gray level difference between a pixel pair, and β is the power of the GLCM. In practice, for cone photoreceptor images, we calculated GLCM_3 or GLCM_4 . α and β are set as 2 and 1, respectively:

$$\text{GLCM}_{(\Delta d)}(i,j) = \sum_{d \in \text{Analyzed Region}} \begin{cases} 1 & \text{if } I(d) = i, I(d + \Delta d) = j \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

$$\text{Contrast}(\alpha, \beta) = \sum_i \sum_j (i-j)^\alpha \text{GLCM}(i,j)^\beta. \quad (4)$$




Fig. 2. Image comparison. Each image is 1.39° by 1.17° cropped from original frames. (a) Single-frame image; (b) image with normal averaging scheme; 20% frames are selected. (c) Image after lucky averaging top 30% pixels from image sequence, with GLCM contrast metric; (d) image after lucky averaging top 30% pixels from image sequence, with sharpness metric.

The calculation result $S(x_0, y_0, \text{frame index})$ from Eq. (2) or Eq. (4) is then assigned to the central pixel (x_0, y_0) . A quality metric $S(x, y, \text{frame index})$ is then generated for every pixel in the image sequence.

For either metric the final step is to set a threshold for including a pixel. That is, for a given pixel position (x, y) , we can include (for instance) the 30% of the instances within the sequence when those pixel instances have the highest value of the image quality metric. The final image is then the average of those pixels. This is computationally fairly simple, taking about 3 min for 100 frames, using a 2.8 GHz PC with 8 Gb of memory and implemented in MATLAB.

Figures 2 and 3 compare results from the different averaging methods. Similar results were obtained on two other subjects. In single-frame images, we observe noise and contrast variance across the image [Fig. 2(a)]. The average image is less noisy [Fig. 2(b)], and in an enlargement more details can be observed due to the decreased noise [Fig. 3(b)]. However, some cones are dim, and some regions are still blurry. For the lucky averaging [Figs. 2(c), 2(d), 3(c), and 3(d)], the contrast reduction is decreased. Cones are somewhat brighter and more




Fig. 3. Enlarged region location indicated in white box as shown in Fig. 2. (a)–(d) correspond to the same region in Figs. 2(a)–2(d).




Fig. 4. (Color online) Spectral power ratio between lucky and conventional averaging. Dotted curve in red, ratio of Fig. 3(c) to Fig. 3(b); dashed curve in blue, spectrum ratio of Fig. 3(d) to Fig. 3(b).

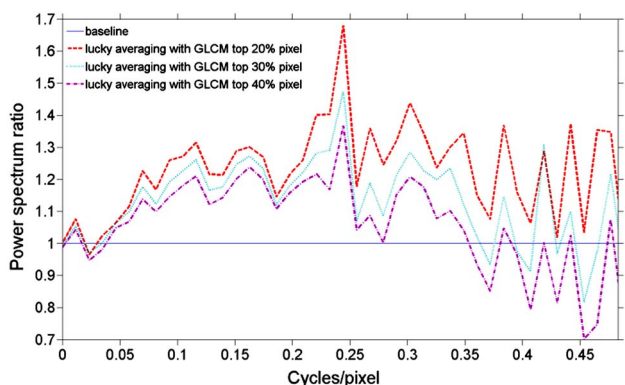


Fig. 5. (Color online) Spectral power ratios for different inclusion criteria. The baseline was based on conventional averaging. The dotted-dashed line in purple, the dotted line in light blue, and the dashed line in red represent inclusion criteria for the best 40%, 30%, and 20% of pixels, respectively.

distinct, suggesting local contrast and resolution improvement. To quantify the difference between lucky averaging and normal averaging, a power spectral analysis was performed. We calculated the power ratio as a function of frequency between lucky averaging and standard averaging (Fig. 4). Thus, a ratio greater than one indicates a higher spectral power for lucky averaging at a given frequency. The average increase in powers between either lucky averaging imaging metric and standard averaging is 20%, and the improvement extends all the way to the optical cutoff frequency, which is about 0.4 cycles/pixels based on the system resolution $2.5\ \mu\text{m}$ [8] for this subject's pupil size of 6 mm. Beyond the cutoff frequency, the ratio fluctuates around one, indicating

that the noise level is close between the two schemes. The two metrics used for lucky averaging produced similar improvements for these images, but because the scheme can be used with most local quality evaluation metrics the approach is general and can be tuned to optimize extraction of the most appropriate information.

To investigate the role of the criterion for including pixels, we processed the same region in Fig. 3 with varied criteria to include the best 20%, 30%, and 40% of pixels using the GLCM method (Fig. 5). As expected, the improvement is largest with the lowest percentage criterion; however, there is a limit since, as the criterion uses fewer samples, the impact of photon noise increases. Thus, the choice of criterion will be a balance between reducing random noise through averaging and improving contrast by picking the best "moments." One simple solution to this trade-off is to collect more data for the same region of interest. It is also possible to base the criterion on the information itself, optimizing the signal-to-noise ratio.

In summary, lucky averaging, analogous to lucky imaging, has been applied to AOSLO image postprocessing. The image sequence is evaluated based on small regions instead of frames. The quality is enhanced by 20%, suggesting a better use of the data by this processing scheme. A larger data set is necessary, but collecting 3 or 4 s worth of data for each retinal region is realistic.

This work was supported by National Institutes of Health (NIH) grants EY14375, EY04395, and P30EY019008.

References

1. A. Roorda, *Optom. Vis. Sci.* **87**, 260 (2010).
2. S. B. Stevenson and A. Roorda, *Proc. SPIE* **5688**, 145 (2005).
3. D. L. Fried, *J. Opt. Soc. Am.* **68**, 1651 (1978).
4. R. D. Ferguson, Z. Zhong, D. X. Hammer, M. Mujat, A. H. Patel, C. Deng, W. Zou, and S. A. Burns, *J. Opt. Soc. Am. A* **27**, A265 (2010).
5. R. A. Muller and A. Buffington, *J. Opt. Soc. Am.* **64**, 1200 (1974).
6. Y. Geng, L. A. Schery, R. Sharma, A. Dubra, K. Ahmad, R. T. Libby, and D. R. Williams, *Biomed. Opt. Express* **2**, 717 (2011).
7. R. M. Haralick, K. Shanmugam, and I. Dinstein, *IEEE Trans. Syst. Man Cybern.* **3**, 610 (1973).
8. W. Zou, X. Qi, and S. A. Burns, *Biomed. Opt. Express* **2**, 1986 (2011).