Article

Fully Automated Estimation of Spacing and Density for Retinal Mosaics

Robert F. Cooper[1,2], Geoffrey K. Aguirre[3], and Jessica I. W. Morgan[1,4]

[1] Scheie Eye Institute, Department of Ophthalmology, University of Pennsylvania, Philadelphia, PA, USA
[2] Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA
[3] Department of Neurology, University of Pennsylvania, Philadelphia, PA, USA
[4] Center for Advanced Retinal and Ocular Therapeutics, University of Pennsylvania, Philadelphia, PA, USA

Correspondence: Jessica I. W. Morgan, University of Pennsylvania, Department of Ophthalmology, 3400 Civic Center Boulevard, Ophthalmology 3rd floor WEST 3-112W, Philadelphia, PA 19104, USA. e-mail: jwmorgan@pennmedicine.upenn.edu

Copyright 2019 The Authors

Purpose: To introduce and validate a novel, fully automated algorithm for determining pointwise intercell distance (ICD) and cone density.

Methods: We obtained images of the photoreceptor mosaic from 14 eyes of nine subjects without retinal pathology at two time points using an adaptive optics scanning laser ophthalmoscope. To automatically determine ICD, the radial average of the discrete Fourier transform (DFT) of the image was analyzed using a multiscale, fit-based algorithm to find the modal spacing. We then converted the modal spacing to ICD by assuming a hexagonally packed mosaic. The reproducibility of the algorithm was assessed between the two datasets, and accuracy was evaluated by comparing the results against those calculated from manually identified cones. Finally, the algorithm was extended to determine pointwise ICD and density in montages by calculating modal spacing over an overlapping grid of regions of interest (ROIs).

Results: The differences of DFT-derived ICD between the test and validation datasets were 3.2% ± 3.5% (mean ± SD), consistent with the differences in directly calculated ICD (1.9% ± 2.9%). The average ICD derived by the automated method was not significantly different between the development and validation datasets and was equivalent to the directly calculated ICD. When applied to a full montage, the automated algorithm produced estimates of cone density across retinal eccentricity that well match prior empiric measurements.

Conclusions: We created an accurate, repeatable, and fully automated algorithm for determining ICD and density in both individual ROIs and across entire montages.

Translational Relevance: The use of fully automated and validated algorithms will enable rapid analysis over the full photoreceptor montage.

Introduction

Adaptive optics (AO) ophthalmoscopy facilitates the routine acquisition of images of the living human retina in both health and disease. After acquisition, image sequences are registered to a reference frame[1,2] and montaged. Historically, these steps were performed semiautomatically, but as the technology has matured, fully automated methods have been developed to automatically select reference frames,[3] remove residual distortion,[4] and montage the registered images.[5,6]

A major remaining analysis step in need of automation is the identification of cones to derive a measure of cone density. Typically, cone density is measured in manually selected regions of interest (ROI), the size and location of which are selected by the investigator. Then, cells within each ROI are identified manually before structural metrics are extracted from the cell locations. Even state-of-the-art cell identification algorithms[7–9] require a reader to correct misidentified cells. Thus, either by ROI selection or cell identification, every quantitative metric of the photoreceptor mosaic is ultimately governed by the subjective assessment of a grader.

To remove the subjective influence and time-consuming need for graders on the structural analysis

of photoreceptors, we present and assess a novel, fully automated algorithm for determining pointwise intercell distance (ICD) and cone density for montages of the photoreceptor mosaic.

Methods

Human Subjects

This research followed the tenants of the Declaration of Helsinki and was approved by the institutional review board at the University of Pennsylvania. All light exposures adhered to the maximum permissible exposure limits set by the American National Standards Institute.[10] Subjects provided informed consent after the nature and possible consequences of the study were explained. Axial length measurements were obtained using an optical biometry device (IOL Master; Carl Zeiss Meditec, Dublin, CA). Image pixels were converted to micrometers on the retina by first acquiring images of a Ronchi ruling positioned at the focal plane of a lens with a 19-mm focal length; from this, we determined the conversion between image pixels and degrees. An adjusted axial length method[11] was used to calculate the scaled retinal magnification factor (micrometers/degree) and ultimately convert the images to a micrometer scale.

Imaging the Photoreceptor Mosaic

Adaptive optics scanning laser ophthalmoscope (AOSLO) images from a previously published work[12] were used for algorithm testing and validation. Briefly, the dataset consisted of nine subjects (14 eyes) without retinal pathology at two time points. We used a previously described AOSLO[13] to obtain a 2×2-mm montage of the photoreceptor mosaic at two time points. Each image within a montage was obtained while the subject maintained fixation on a target and a series of image sequences were collected across the horizontal and vertical meridians of the retina. For each image sequence at each retinal location, intraframe distortion due to the sinusoidal motion of the horizontal scanner was removed by resampling each image using a Ronchi ruling. After desinusoiding the image sequence, a reference frame was manually selected and subsequently used to register the image sequence using a previously described strip-registration method.[2] Fifty registered frames were then averaged to create a single, high signal-to-noise image from each sequence. For each time point, each average image

from each location was stitched together using image-editing software (Photoshop CS6; Adobe, San Jose, CA).

In each eye and at each time point, ROIs were obtained 0.19, 0.35, 0.50, 0.90, and 1.5 mm from the fovea along the temporal meridian. The second time point ROI was aligned to the first at each location using an affine registration. Cone coordinates were identified semiautomatically using custom software (MOSAIC; Translational Imaging Innovations, Inc., Hickory, NC) in each ROI, and ICD was directly calculated from the coordinate locations as previously defined (ICD_D). All ROIs and associated coordinates from the first time point were randomly assigned to either test or validation datasets, with the corresponding ROI from the second time point being assigned to the opposite dataset. All algorithm development was performed solely on the test dataset. All directly calculated densities were finalized prior to development of the modal spacing algorithm described below.

Determining Modal Spacing

This algorithm uses multiscale fitting and residual peak finding to determine modal spacing. First, the discrete Fourier transform (DFT) of the input image was calculated (Fig. 1A, B). We then calculated the log power spectrum of the DFT and obtained a radial average over the upper and lower 90° (Fig. 1C, 1D) of the power spectrum. This radial average was then fit with a first-order exponential (Fig. 1E), which was then subtracted from the radial average (Fig. 1F). The residuals were smoothed using splines, and the location of maximal difference from the fit was used to establish an initial estimate of the modal spacing. A piecewise first-order exponential function was then fit to the adjusted radial average (Fig. 1G), using the initial estimate from Figure 1F as the position of the link between two halves of the function. As before, this fit was subtracted from the original radial average and smoothed using splines. Using a peak crawling algorithm, we found the first peak in the direction of decreasing frequency (Fig. 1H), beginning at the location of the initial estimate, and took this value as the modal spacing for the image. The modal spacing corresponds to the row spacing of the cones in the image. Thus, assuming an asymmetric hexagonally packed mosaic[14] with row spacing of s_r micrometers, we can calculate the modal ICD (ICD_M, Equation 1) or density in cells/mm² (D_M, Equation 2).
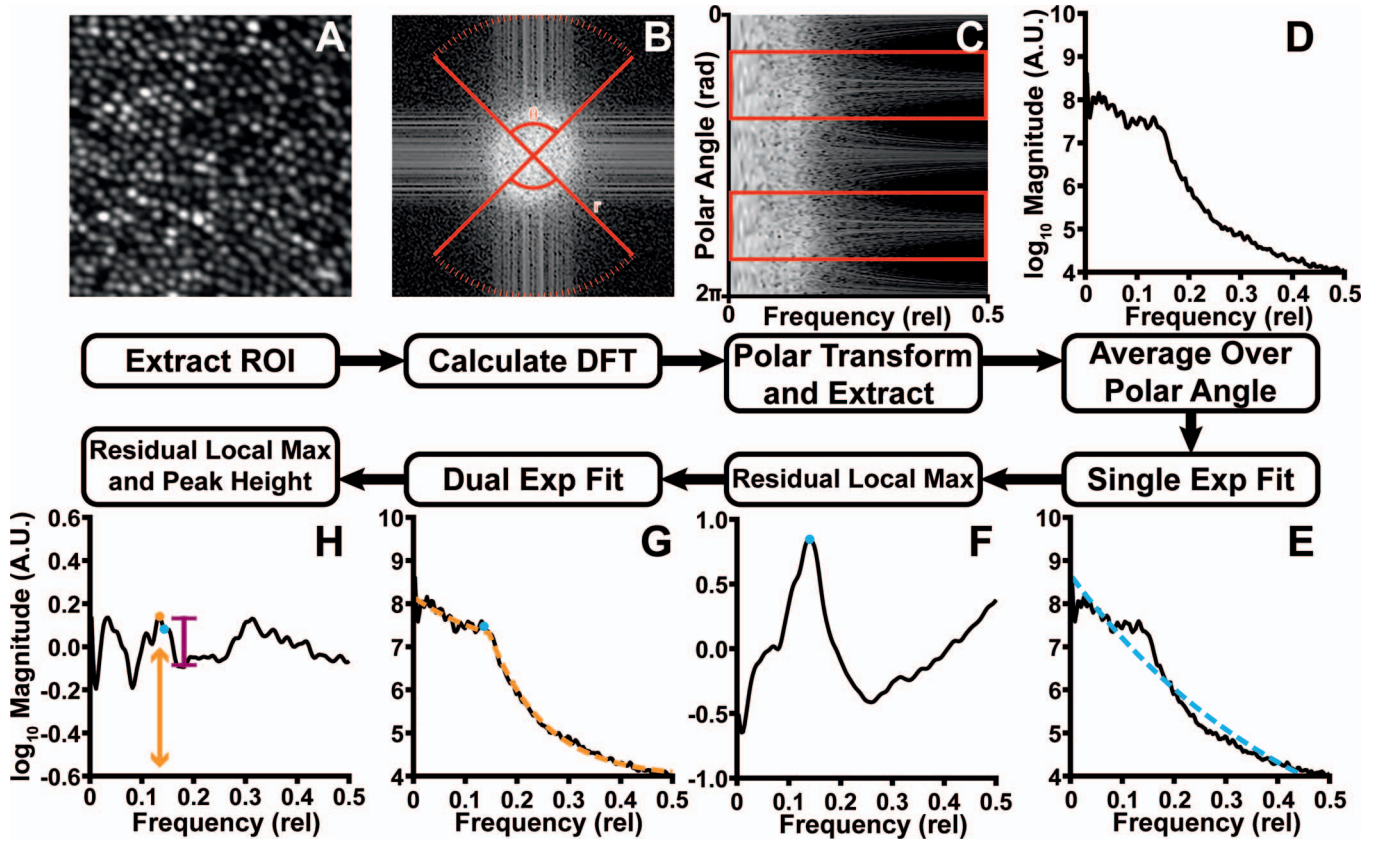
Figure 1. A flowchart of the proposed algorithm. From any given region of interest (A), we calculated the log power spectrum of the DFT of the region (B), converted the DFT to polar space (C), and averaged the upper and lower 90° of the DFT (*red boxes*) to obtain a radial average (D). The radial average was fit with a first-order exponential (E) and subtracted. (F) The residuals were smoothed, and residual maximum location was used to establish a rough estimate of the modal spacing. The rough estimate was used to set the initial location of the piecewise first-order exponential function, which was also fit to the radial average (G). Again, the fit was subtracted from the radial average and smoothed using splines. A peak crawling algorithm was used to find the first peak in the direction of decreasing frequency, starting at the previously determined rough estimate location (H). The final modal spacing estimate corresponds to the row spacing of the cones in the image (*orange arrow*). In addition, we estimated algorithm confidence from the maximum peak prominence of the modal spacing, normalized by the maximum residual amplitude. The maroon caliper shows the side of the peak with the most prominence.

$$ICD_M = \frac{2}{\sqrt{3}} s_r \qquad (1)$$

$$D_M = \frac{1000^2 \sqrt{3}}{2 s_r{}^2} \qquad (2)$$

To provide a confidence rating for the modal spacing estimate, we calculated the modal peak distinctiveness, defined as the maximal modal peak height normalized by the maximum residual amplitude (Fig. 1H, maroon caliper).

### Assessing Algorithm Performance

We first assessed the algorithm's reproducibility by calculating the percent difference of the ICD_M

between the test and validation sets and their limits of agreement (LOA).[15] Statistical significance for differences between the ICD_M were assessed through linear regression analysis using a mixed effects model.[12] Next, we assessed the accuracy of the algorithm by evaluating the agreement between DFT-derived and directly calculated ICD across the entire dataset. We used a Bland-Altman plot[15] to determine the mean bias and 95% LOA between the two methods, where the limits were defined as:

$$LOA = \bar{D} \pm 1.96S \qquad (3)$$

Where $\bar{D}$ is the mean difference between the two methods, and S is the standard deviation of the differences.
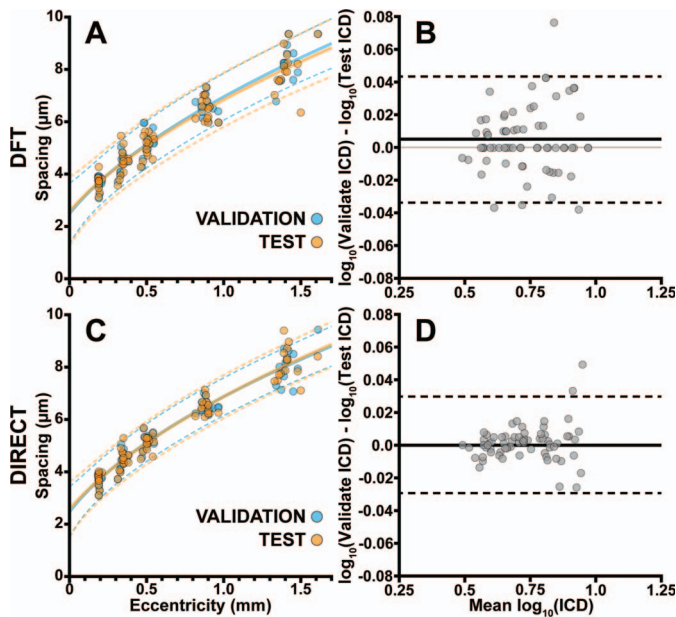
Figure 2. Assessing consistency between the algorithm test and validation dataset. (A) After dividing the dataset into test and validation, the algorithm (Fig. 1) was run on each image, and the results from the test and validation dataset were fit using Equation 4 (Test r^2: 0.91; Validate r^2: 0.93). The test and validation sets were then compared with a Bland-Altman plot (B). On average, the modal ICD of the validation set was 0.005 $\log_{10}$ μm (95% LOA: ±0.040 $\log_{10}$ μm) greater than the test set. On a linear scale, the validation set was on average 1.00 (95% LOA: 0.98–1.03) times the test set. (C) For comparison, we performed the same analysis on directly calculated ICD (Test r^2: 0.94, Validate r^2: 0.95). (D) On average, the directly calculated ICD of the validation set was 0.0001 $\log_{10}$ μm (95% LOA: ±0.030 $\log_{10}$ μm) greater than the test set. On a linear scale, the validation set was on average 1.00 (95% LOA: 0.98–1.02) times the test set. Neither the directly determined nor modal intercell spacing were significantly different between the test and validation datasets ($P > 0.05$).

Results

Examining Algorithm Reproducibility

Modal cone spacing and directly calculated cone spacing increased with eccentricity (Fig. 2), consistent with previously reported measurements.[16–18] We fit each dataset (test, validation, and ICD_D, ICD_M) with the exponential:

$$ICD_\text{fit} = a(x_\text{dist} - b)^c \qquad (4)$$

where x_dist is the distance to the foveal center in micrometers and the fit coefficients a, b, and c were allowed to vary. We then calculated the 95% confidence interval of each fit (Fig. 2A, 2C). All datasets were fit well using this model (ICD_M test r^2:

0.91, validate r^2: 0.93; ICD_D test r^2: 0.94, validate r^2: 0.95) and exhibited similar fit coefficients and confidence interval widths. We determined that the percent difference of ICD_M between the test and validation datasets was 3.2% ± 3.5% (mean ± SD; range, 0%–18%), consistent with ICD_D (1.9% ± 2.9%; range, 0%–21%). We created a Bland-Altman plot between the test and validation datasets (Fig. 2B, 2D). Due to differences between the two datasets violating assumptions of normality, the data were first $\log_{10}$ transformed.[15] After log transformation, all test and validation sets passed normality tests (Shapiro-Wilk, $P > 0.05$), allowing use of the standard parametric Bland-Altman form (Equation 3). On average, the ICD_M of the validation set was 0.005 $\log_{10}$ μm greater than the test set, and the ICD_D of the validation set was 0.0001 $\log_{10}$ μm greater than the test set. The 95% LOA for ICD_M were ±0.040 $\log_{10}$ μm, and ±0.030 $\log_{10}$ μm for ICD_D. On a linear scale, this means that the ICD_M validation set was on average 1.00 (95% LOA: 0.98–1.03) times the test set, and the ICD_D validation set was 1.00 (95% LOA: 0.98–1.02) times the test set. The test and validation sets for both methods were not significantly different ($P > 0.05$), leading us to combine the test and validation sets for the remainder of this work.

Assessing Agreement With Directly Counted Cone Metrics

To quantify intermethod agreement, we compared the modal cone spacing with the directly measured cone spacing. We determined that ICD_M was on average 3.3% ± 2.8% (range: 0.02%–15%) different from the corresponding ICD_D measurement on the same data. We created Bland-Altman plots (Fig. 3). As before, differences between the two datasets violated assumptions of normality, and the ICD data were $\log_{10}$ transformed. After log transformation, all datasets again passed normality tests (Shapiro-Wilk, $P > 0.05$). On average, the difference between the two methods was minimally biased (−0.00014 $\log_{10}$ μm), and the 95% LOA were found to be ±0.019 $\log_{10}$ μm. On a linear scale, this means that the ICD_M was on average equivalent to the ICD_D, where the 95% LOA were 0.97 to 1.03.

### Pointwise Spacing and Density Within Montages of the Photoreceptors

Following validation, the algorithm was extended to create pointwise density maps from AOSLO montages of the photoreceptor layer. To test the
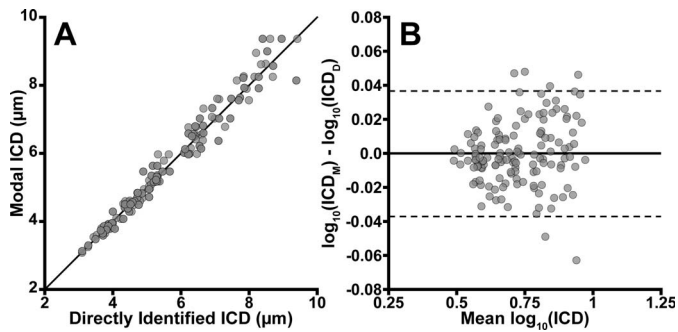
Figure 3. Determining the agreement between the directly determined ICD and the modal ICD. On average, the difference between the two methods was minimally biased (-0.00014 $\log_{10}$ µm), and the 95% LOA were ± 0.019 $\log_{10}$ µm. On a linear scale, the two methods were equivalent, where the 95% LOA were 0.97 to 1.03.

extension of the algorithm to creating pointwise maps of density and spacing, we obtained an additional approximately 4×4-mm montage of the photoreceptor mosaic using the previously described AOSLO[13] and post processing methods.[2,3,5]

Each image within the montage was divided into an overlapping grid of $n \times n$ pixel ROIs, where each ROI was displaced from the previous by p pixels. Here we empirically chose $n = 128$ and $p = 32$ (for the subject in Fig. 4, this corresponds to 58.8 and 14.7 µm, respectively) to provide a balance between pointwise accuracy and map smoothness. Next, each ROI was analyzed using the previously described algorithm. Using the modal spacing estimate confidence (Fig. 1H) as a weight, we then combined the ICD_M (Equation 1) or density (Equation 2) values of each overlapping ROI across the montage by weighted average. Regions with poor confidence (the 5th percentile of the confidence values) or low mean ROI pixel intensity (<10 arbitrary units) were
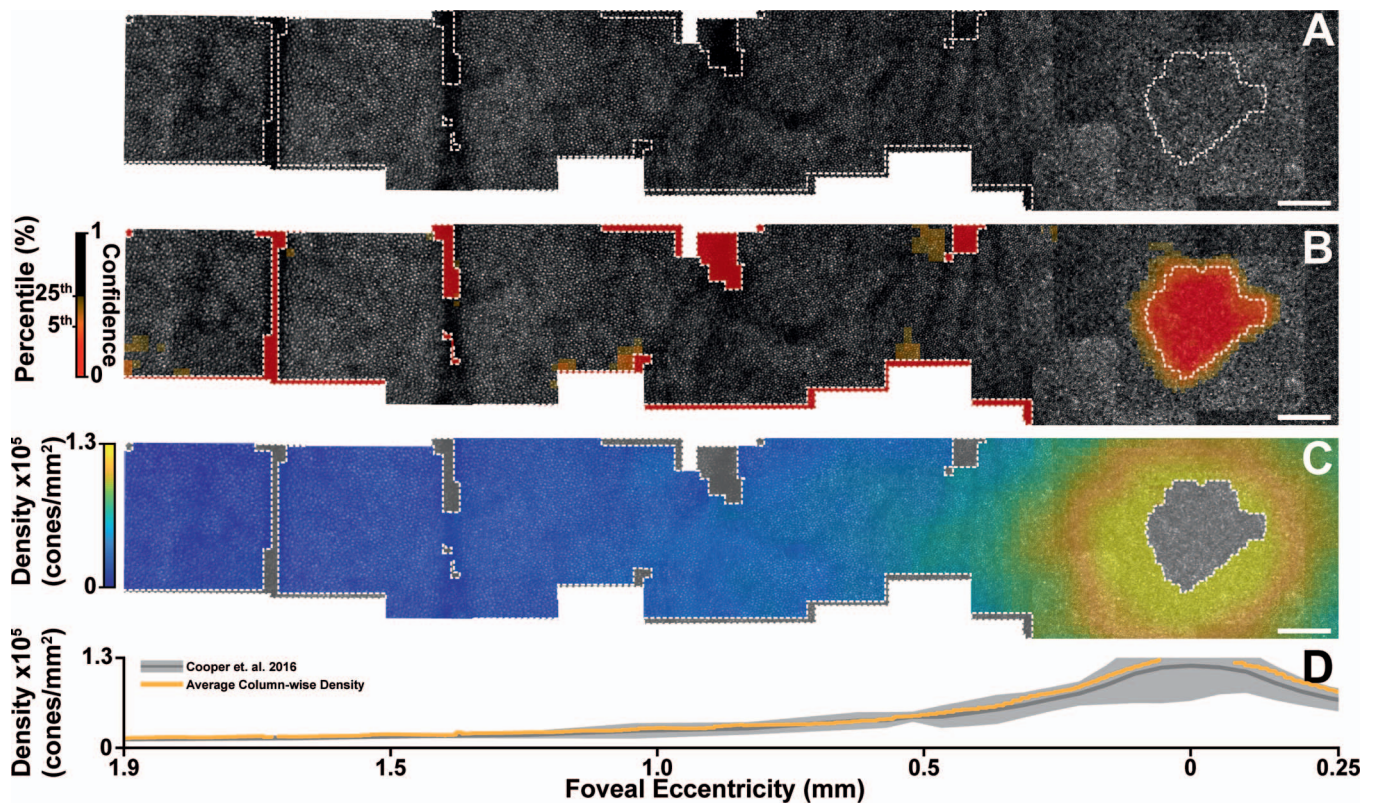
Figure 4. Applying the algorithm to an AOSLO montage. (A) We used a previously obtained montage of the photoreceptor layer to assess the algorithm's performance (2×0.5 mm shown). Using a sliding 128 pixel ROI, we assessed the confidence of our modal spacing estimate (B) at each point in the montage. We then calculated cone density from the modal intercell spacing (Equation 2) at each point in the montage (C). (D) The *orange line* is the column-wise average density obtained from a 200-µm strip along the horizontal meridian from the map in (C). The *gray line* and *shaded region* is the mean directly calculated density and 95% prediction interval obtained from a set of 20 subjects without pathology.[16] The area within *white dashed lines* delineate regions that are excluded from the map due to low image intensity, poor confidence, or foveal artifacts. *Scale bar* is 100 µm.
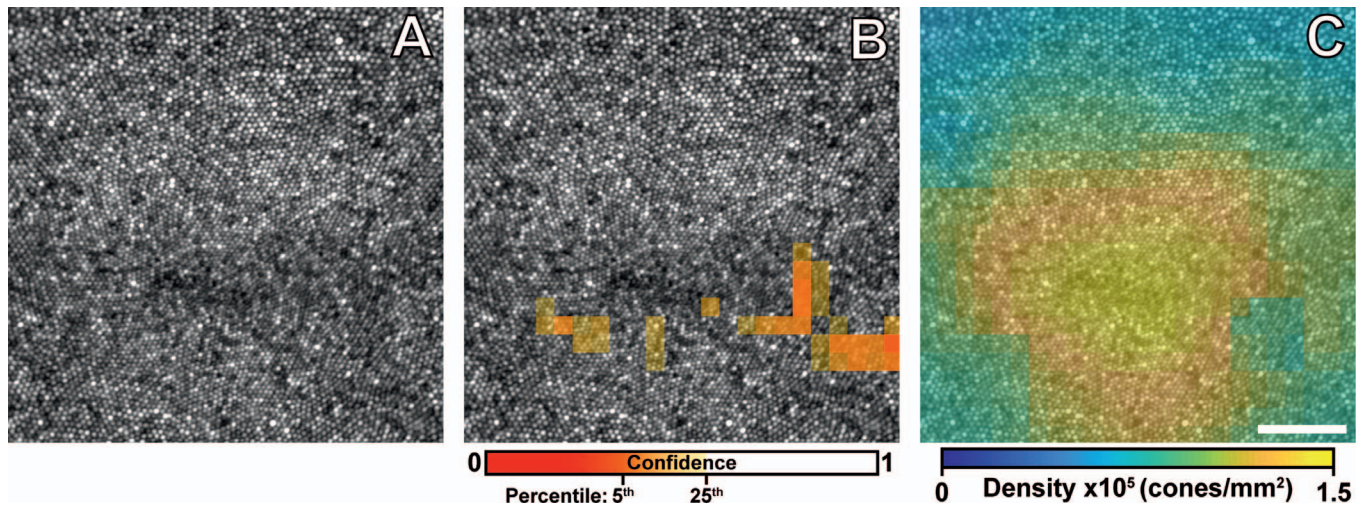
Figure 5. Reduced confidence in ICD_M estimates at regions of high density such as the fovea is caused by local regions of unresolved photoreceptors (Fig. 4B). Indeed, in subjects with resolvable foveal cones (A), algorithm confidence in the fovea (B) is sufficient to provide an estimate of foveal density (C). *Scale bar* is 50 μm.

excluded from the map (Fig. 4B). Due to poor resolution of the foveal cones in some images, density values at the foveal center were erroneously low. To mitigate this artifact, we enhanced the existing exclusion area at the foveal center by creating a mask using active contours that found the beginning of the decreasing density gradient.

Using this technique, we successfully obtained pointwise cone density measurements across a full montage (Fig. 4C). In the montage, we observed the stereotypical exponential falloff of cone density as a function of eccentricity (Fig. 4D), and these data were consistent with previously observed density functions.

Discussion

We have developed a robust, fully automated algorithm that estimates ICD and cone density. The algorithm provides consistent results across two datasets and agrees closely with directly calculated spacing. Moreover, the algorithm is readily extended across entire AO ophthalmoscope montages, enabling montage-sized pointwise maps of cell density and ICD with no grader involvement. Using a desktop processor with 32 GB of RAM (AMD Ryzen 1800×; AMD, Santa Clara, CA) and a programming platform (MATLAB 2018a; MathWorks, Natick, MA), our approximately 4 × 4-mm montage (Fig. 4) took approximately 15 minutes to create a complete map of the density, spacing, and confidence of the montage.

Previously published techniques of ICD and cone density rely on the accurate detection of cone locations within an ROI.[16] Previously published cone detection algorithms can be either heuristic[7,19–21] or inferential,[8] but all algorithms ultimately produce an output of coordinates that must be verified by a human grader. To reduce the time burden of this process, measurements are often made by a limited number of graders, each of whom imparts subjective influence on cone selection.[22,23] Therefore, the reproducibility and accuracy of these measurements are inherently limited by human graders. Additionally, existing techniques for creating spacing and density maps embed subjective decisions regarding ROI placement. Consequently, existing measurements are made over a limited number of ROIs, reducing sensitivity for detecting subtle spatial variations in cell metrics. Overall, the analysis approach presented here substantially reduces subjective input, enabling relatively rapid and accurate characterization of large datasets.

In Figure 4, the algorithm indicated poor confidence in the region near the fovea. This feature, while prominent in the subject we used for the example above, is an imaging system limitation and not an algorithm limitation. Indeed, the algorithm's performance is representative of a common problem in AOSLO imaging: the inability to resolve foveal cones in many individuals. Indeed, in subjects where every foveal cone can be resolved (Fig. 5), the algorithm is reliable in its estimate of cone density (Fig. 5B, 5C).

Our algorithm is still Fourier-domain based. We are therefore unable to determine local regularity or
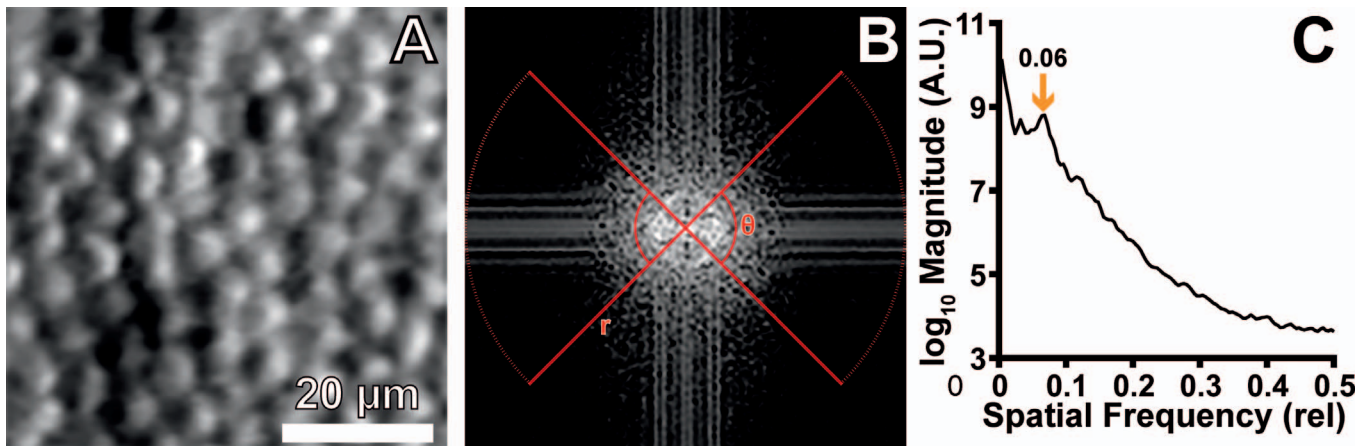
Figure 6. This technique can be used on any image containing periodic structures (e.g., other modalities such as split detection [A] or other cell types). The polar angle range and spacing equation may need to be adjusted accordingly. For example, in split-detection images (A) we can use the left/right 90° of the DFT (B) to extract ICD_M without needing to multiply by the correction factor in Equation 1.

tessellation information in the same way as algorithms operating in the spatial domain do. To achieve an analogous measurement of local regularity, the algorithm could be extended to use pointwise mean and standard deviation; however, this would necessitate a much finer ROI sampling (<8 pixels) and would in turn exponentially increase processing time.

Additionally, without fine ROI sampling, the algorithm is unable to capture pathological changes unless spatial remodeling has also occurred. Indeed, any highly localized loss of photoreceptors such as a small scotoma or spacing gradient smaller than the grid spacing will diminish the accuracy of the method. Moreover, while our confidence measure is effective for detecting poorly defined spatial structure in an image, it can be affected by the infiltration of rod photoreceptors or other structures. Additional structures such as rods would create a corresponding peak in the radial DFT, potentially leading to an incorrect spacing and/or confidence estimate.

In this algorithm, we considered only the upper and lower 90° of the DFT (Fig. 1C) to minimize the effect of the slightly asymmetric row and cone spacing[14] on our estimate of ICD. However, there is nothing implicitly good (or bad) about using the upper/lower versus the left/right for cone photoreceptors. Indeed, the same algorithm could be used with a left/right division, simply omitting the use of Equation 1. Moreover, it may be beneficial to adjust the angle range used for analysis depending on orientation of the feature of interest. For example, for a left/right dominant feature such as the cones in split-detection images,[24] the left and right 90° can be analyzed (Fig. 6).

As the amount of AO ophthalmoscope data grows commensurate with studies of retinal degeneration, it is becoming increasingly impractical to semiautomatically assess the large volume of data available without significant quantization. Algorithms such as these enable large-scale normative studies and, ultimately, comparative retinal degeneration studies, to reach their endpoints in a timely and objective manner.

References

1. Arathorn DW, Yang Q, Vogel CR, Zhang Y, Tiruveedhula P, Roorda A. Retinally stabilized cone-targeted stimulus delivery. *Opt Express*. 2007;15:13731–13744.

2. Dubra A, Harvey Z. Registration of 2D images from fast scanning ophthalmic instruments. In: Fischer B, Dawant B, Lorenz C, eds. *Biomedical Image Registration*. Berlin: Springer-Verlag; 2010: 60–71.

3. Salmon AE, Cooper RF, Langlo CS, Baghaie A, Dubra A, Carroll J. An automated reference frame selection (ARFS) algorithm for cone imaging with adaptive optics scanning light ophthalmoscopy. *Trans Vis Sci Techn*. 2017;6:9.

4. Bedggood P, Metha A. De-warping of images and improved eye tracking for the scanning laser ophthalmoscope. *PLoS One*. 2017;12:e0174617.

5. Chen M, Cooper RF, Han GK, Gee J, Brainard D, Morgan JIW. Multi-modal automatic montaging of adaptive optics retinal images. *Biomed Opt Express*. 2016;7:4899–4918.

6. Davidson B, Kalitzeos A, Carroll J, et al. Fast adaptive optics scanning light opthalmoscope retinal montaging. *Biomed Opt Express*. 2018;9: 4317–4328.

7. Li KY, Roorda A. Automated identification of cone photoreceptors in adaptive optics retinal images. *J Opt Soc Am A*. 2007;24:1358–1363.

8. Cunefare D, Fang L, Cooper RF, Dubra A, Carroll J, Farsiu S. Open source software for automatic detection of cone photoreceptors in adaptive optics ophthalmoscopy using convolutional neural networks. *Sci Rep*. 2017; 7.

9. Davidson B, Kalitzeos A, Carrol J, et al. Automatic cone photoreceptor localisation in healthy and stargardt afflicted retinas using deep learning. *Sci Rep*. 2018; 8.

10. ANSI. *American National Standard for Safe Use of Lasers ANSI Z136.1-2014*; Laser Institute of America; 2014.

11. Bennett AG, Rudnicka AR, Edgar DF. Improvements on Littmann's method of determining the size of retinal features by fundus photography. *Graefes Arch Clin Exp Ophthalmol*. 1994;232:361–367.

12. Jackson K, Vergilio GK, Cooper RF, Ying G, Morgan JIW. A 2-year longitudinal study of normal cone photoreceptor density. *Invest Ophthalmol Vis Sci*. 2019;60:1420–1430.

13. Dubra A, Sulai Y. Reflective afocal broadband adaptive optics scanning ophthalmoscope. *Biomed Opt Express*. 2011;2:1757–1768.

14. Curcio CA, Sloan KR. Packing geometry of human cone photoreceptors: variation with eccentricity and evidence for local anisotropy. *Vis Neurosci*. 1992;9:169–180.

15. Bland JM, Altman DG. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet*. 1986;1:307–310.

16. Cooper RF, Wilk MA, Tarima S, Carroll J. Evaluating descriptive metrics of the human cone mosaic. *Invest Ophthalmol Vis Sci*. 2016;57:2992–3001.

17. Curcio CA, Sloan KR, Kalina RE, Hendrickson AE. Human photoreceptor topography. *J Comp Neurol*. 1990;292:497–523.

18. Song H, Chui TY, Zhong Z, Elsner AE, Burns SA. Variation of cone photoreceptor packing density with retinal eccentricity and age. *Invest Ophthalmol Vis Sci*. 2011;52:7376–7384.

19. Garrioch R, Langlo C, Dubis AM, Cooper RF, Dubra A, Carroll J. Repeatability of in vivo parafoveal cone density and spacing measurements. *Optom Vis Sci*. 2012;89:632–643.

20. Bukowska DM, Chew AL, Huynh E, et al. Semi-automated identification of cones in the human retina using circle Hough transform. *Biomed Opt Express*. 2015;6:4676–4693.

21. Cunefare D, Cooper RF, Higgins B, et al. Automatic detection of cone photoreceptors in split detector adaptive optics scanning light ophthalmoscope images. *Biomed Opt Express*. 2016;7:2036–2050.

22. Liu B, Tarima S, Visotcky A, et al. The reliability of parafoveal cone density measurements. *Br J Ophthalmol*. 2014;98:1126–1131.

23. Gale MJ, Harman GA, Chen J, Pennesi ME. Repeatability of adaptive optics automated cone measurements in subjects with retinitis pigmentosa and novel metrics for assessment of image quality. *Trans Vis Sci Techn*. 2019;8:17.

24. Scoles D, Sulai YN, Langlo CS, et al. In vivo imaging of human cone photoreceptor inner segments. *Invest Ophthalmol Vis Sci*. 2014;55: 4244–4251.

translational vision science & technology