Editor's Choice

Combining Deep Learning With Optical Coherence Tomography Imaging to Determine Scalp Hair and Follicle Counts

Gregor Urban, MS,[1] Nate Feil, BS,[2] Ella Csuka, BS,[2] Kiana Hashemi,[2] Chloe Ekelem, MD MPH,[2,4] Franchesca Choi, RPh, MD,[1] Natasha A. Mesinkovska, MD, PhD,[2,3] and Pierre Baldi, PhD[1]*

[1]*Department of Computer Science, University of California, Irvine, California 92697*
[2]*Department of Dermatology, School of Medicine, University of California, Irvine, California 92697*
[3]*Beckman Laser Institute, Irvine, California 92697*
[4]*Department of Dermatology, School of Medicine, University of Utah, 30 North 1900 East, 4A330, Salt Lake City, Utah 84132*

Background and Objectives: One of the challenges in developing effective hair loss therapies is the lack of reliable methods to monitor treatment response or alopecia progression. In this study, we propose the use of optical coherence tomography (OCT) and automated deep learning to non-invasively evaluate hair and follicle counts that may be used to monitor the success of hair growth therapy more accurately and efficiently.

Study Design/Materials and Methods: We collected 70 OCT scans from 14 patients with alopecia and trained a convolutional neural network (CNN) to automatically count all follicles present in the scans. The model is based on a dual approach of both detecting hair follicles and estimating the local hair density in order to give accurate counts even for cases where two or more adjacent hairs are in close proximity to each other.

Results: We evaluate our system on 70 OCT manually labeled scans taken at different scalp locations from 14 patients, with 20 of those redundantly labeled by two human expert OCT operators. When comparing the individual human predictions and considering the exact locations of hair and follicle predictions, we find that the two human raters disagree with each other on approximately 22% of hairs and follicles. Overall, the deep learning (DL) system predicts the number of follicles with an error rate of 11.8% and the number of hairs with an error rate of 18.7% on average on the 70 scans. The OCT system can capture one scalp location in three seconds, and the DL model can make all predictions in less than a second after processing the scan, which takes half a minute using an unoptimized implementation.

Conclusion: This approach is well-positioned to become the standard for non-invasive evaluation of hair growth treatment progress in patients, saving significant amounts of time and effort compared with manual evaluation. Lasers Surg. Med. © 2020 Wiley Periodicals, Inc.

Key words: alopecia; convolutional neural network; deep learning; hair loss; machine learning; OCT; optical coherence tomography

INTRODUCTION

Hair plays a substantial role in defining one's identity, representing age, social status, and even wisdom. Effluvium refers to a pathophysiological process that leads to either reversible or permanent hair loss. Not surprisingly, alopecia patients often encounter great amounts of psychological and social morbidity directly related to their illness [1]. These patients' grief and frustration have, in part, motivated the development of a wide range of hair loss treatment products; however, techniques to determine treatment efficacy are still lacking. Current methods of monitoring hair loss include hair counting and diagnostic identification of scalp disease with histology or trichoscopy [2–4]. Hair counting can be done on horizontal biopsy, which is the current gold standard [5,6]; however, scalp biopsies carry inherent procedural risks, such as bleeding, pain, and infection. Trichoscopy is non-invasive, but is biased by manual interpretation [7]. Services such as TrichoScan trichoscopy provide an automated calculation of metrics [4,8,9], including hair number, density, and phase ratio [10], but the mechanism is unable to capture hair-bearing and non-hair-bearing follicles. TrichoScan requires shaving and coloration of an area of hair, something that many hair loss patients – many of which are women – are unwilling to readily undergo, limiting its clinical practicality.

Optical coherence tomography, or OCT, is a non-invasive light-based imaging technique that uses the interference properties of skin tissue and infrared light to achieve high-resolution cross-section morphology *in vivo*

Conflict of Interest Disclosures: All authors have completed and submitted the ICMJE Form for Disclosure of Potential Conflicts of Interest and none were reported.
Contract grant sponsor: NSF; Contract grant number: NRT 1633631.
*Correspondence to: Pierre Baldi, PhD, Department of Computer Science, University of California, Irvine, CA 92697. E-mail: pfbaldi@ics.uci.edu
Accepted 9 September 2020
Published online in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/lsm.23324

that can be evaluated in multiple dimensions [11], and has the added benefit of not requiring any hair colorization or shaving, allowing for a traceless procedure. The principles and instrumentation components of OCT have been described previously in detail [12]. The relatively high-resolution images generated by OCT spanning depths of 1–3 mm make it a useful tool for evaluation and diagnosis for many cutaneous conditions [13,14]. At these depths, the epidermis, upper dermis, skin appendages, and blood vessels can be visualized, and changes from skin tumors or inflammation can be appreciated for aid in narrowing differential diagnoses [15]. One feature of the image analysis using optical coherence tomography is viewing 1000 generated vertical image slices in a compiled "en-face" view of the *z*-plane, a two-dimensional image that is parallel to the surface of the skin. In this view, hair follicles can be evaluated as dark spots on the skin surface. This view can further be enhanced with a topographical color-coding, or "height"-map, developed by the authors for ease of viewing (Fig. 1). Thus, follicular density is quantified by automated identification of hair-bearing and non-hair-bearing follicles seen in these constructed images.

Machine-learning methods, in particular deep learning methods, have been successfully applied to biomedical data, beginning with the early biometric application in Baldi and Chauvin [16]. Recently, the field has witnessed a rapid expansion [17] with many applications, for instance, for evaluation and diagnosis of skin cancer [18], lung cancer [19], endoscopy [20], analysis of mammograms [21,22], drug design [23], and spinal metastasis detection [24].

The field of ophthalmology specifically has used OCT extensively to devise algorithms to accurately diagnose a multitude of vision-threatening conditions, such as diabetic retinopathy [25,26]. Here, we describe a machine-learning algorithm built to automatically count hairs, hair-bearing follicles, and non-hair-bearing follicles on the scalp from en-face OCT images.

MATERIALS AND METHODS

Data Collection—In Vivo OCT

All images presented were obtained using an in-line fiber-based swept-source OCT system (Thorlabs Inc., Newton, NJ) using a 1310 nm center wavelength, sweep rate of 200,000 a-line/second, an average power of 35 mW, axial resolutions of 16 μm, lateral resolutions of 8 μm, and NA of 0.2665. Using a spacer, the OCT probe was placed 4 cm above the skin. Each image captured a $5 \times 7\ \text{mm}^2$ area with 1000 vertical frames with a penetration depth of 1.3 mm (Fig. 2). The 1000 frames were then integrated into an en-face view that allows two-dimensional (2D) interpretation of the surface of the skin (see Section Preserving Depth Information in 2D Projections for details).

We collected and processed a total of 70 scans of unique locations of the scalp from 14 patients. Two research experts labeled all hairs, hair-bearing follicles, and non-hair-bearing follicles on the 2D en-face view images, and their annotations will be treated as the gold standard in the following. A subset of 20 images was labeled by both investigators to compare and estimate human versus machine-learning model predictions. This study was granted approval in 2017 by the University of California Irvine Institutional Review Board (IRB).

Possible Machine Learning Approaches to the Problem

The raw OCT data of one single scan exists in the form of a three-dimensional grayscale image, with dimensions of $1000 \times 700 \times 1000$ pixels. The goal was to automatically find and count the follicles and hairs in this volume using machine learning (ML). Currently, the type of ML model

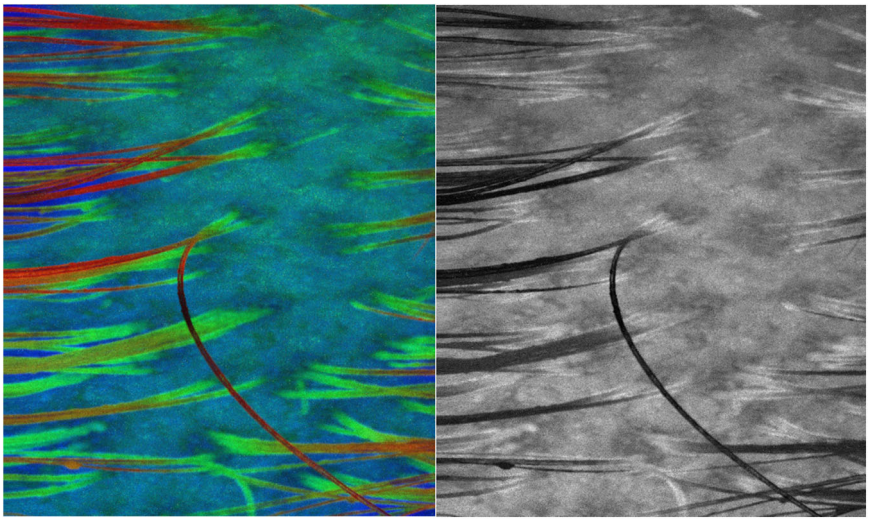

Fig. 1. Hair-bearing scalp: example of a topographically enhanced color-coded OCT scan projection compared with a traditional 2D projection. 2D, two-dimensional; OCT, optical coherence tomography.

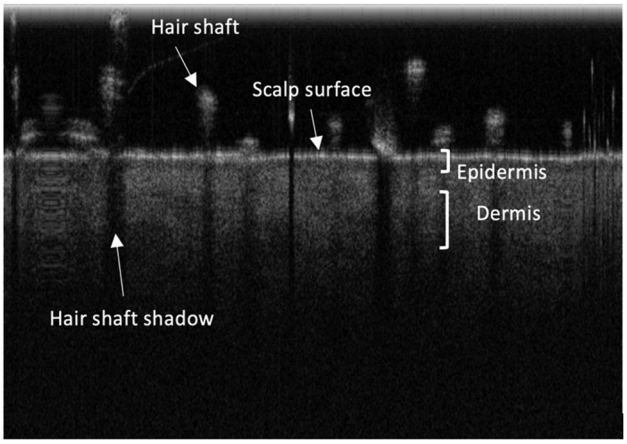

Fig. 2. Example of a cross-section of 2D OCT data obtained on the human scalp. The scalp surface, dermis, epidermis, hair shafts, and hair shaft shadows are labeled. 2D, two-dimensional; OCT, optical coherence tomography.

best-suited for computer vision problems such as this is a convolutional neural network (CNN) [16]. There are several methods by which a CNN can be applied, and the following were considered:

1) Applying a 2D CNN in a scanning fashion that operates on individual 2D slices of the three-dimensional (3D) data. The model would predict the locations of follicles and hairs in each slice. If scanning along the first dimension, this entails processing 1000 images of size 700×1000 pixels and then combining the 1000 predictions of the CNN into a 3D "image" of predictions. A post-processing step is necessary, in which adjacent positive predictions are analyzed to determine if they are assigned to the same follicle or to a different follicle that is merely close by.
2) Applying a 3D CNN that processes the entire scan in one pass, or in smaller chunks if computer memory is a limiting factor. The predictions will be a 3D "image" as in (1) and require similar post-processing to identify unique follicles. The main drawback compared with (1) is that 3D CNNs are extremely resource-intensive, which would complicate widespread use.
3) Converting the 3D data into a 2D projection, that is, a "top-down-view" on the skin, or "en-face" view, such that a 2D CNN can process the entire scan in one pass. This approach has a significant advantage over (1) and (2) of being roughly 700 times faster, assuming the 3D to 2D projection is given, which also makes training and optimizing the CNN model much easier. Furthermore, correctly labeling data for training the model becomes easier. The downside is that a naïve projection compiled simply by averaging across one dimension loses all height information present in that dimension and might produce a blurry image.

Initial attempts were focused on approach (1) but yielded insufficiently accurate predictions to obtain correct counts. This was, in part, due to the difficulty in creating correct annotations in the 3D data for training the model. Therefore, approach (3) is described in this study, as it provides the added benefit of being useable on computers without a dedicated graphics processing unit, which greatly facilitates clinical use.

Preserving Depth Information in 2D Projections

The intermediate goal was to "compress" the 3D data into a single 2D en-face view while preserving as much information as possible, which was accomplished by taking advantage of the grayscale nature of the 3D images and encoding different heights in the 3D data with different colors in the 2D en-face view. For example, anything far above the scalp will be colored red, will become gradually and fractionally greener as the scalp is approached, and then shift to more blue hues at and below the level of the scalp. The final 2D color image is then the sum of the pixel intensities of the 3D image weighted by a height-dependent fraction of red/green/blue colors along the dimension that is perpendicular to the scalp surface (Fig. 3).

Two issues remain: the scalp may not always be at the same height if the OCT probe was held at different heights or angles while the scan was being obtained, and the scalp might be heavily curved and/or angled despite excellent instrumentation by the OCT operator. This was solved by designing a simple scalp height estimator using the following image processing functions: filtering, thresholding, computing the average height of the columns in the 3D data (where one column will later correspond to one pixel in the 2D projection), minimum filtering to remove the height values of hairs, and finally, Gaussian smoothing. The result is a sufficiently accurate "height"-map of the scalp, which was then used as a reference when computing the color projection. By shifting the color weightings used to compute any given pixel in the 2D projection with the heightmap, the goal of "unwarping" the scalp is achieved, and every single color corresponds to a specific height relative to the scalp (Fig. 4). This color consistency should also facilitate the follicle-detection problem for the machine-learning model. A similar technique of 3D reconstruction to accommodate curved surfaces using OCT for vascular visualization has been reported by Ulrich et al. [27].

Deep-Learning Models

The use of CNNs has become commonplace for tasks such as object classification and localization. Generally speaking, there are two approaches for object detection/localization. The first involves training two CNNs, where one model identifies candidate objects regions in the image, and the other model is evaluated on each such region to confirm if an object is there; examples for this are [28] and [29]. The second class of models is single-shot detectors that classify and locate objects with a single model in a single pass, such as YOLOv3 [30] (You only look once), and SSD [31] (single-shot detector).

As opposed to the general object detection framework, the task at hand did not require elaborate classification involving many classes, neither would it benefit from the

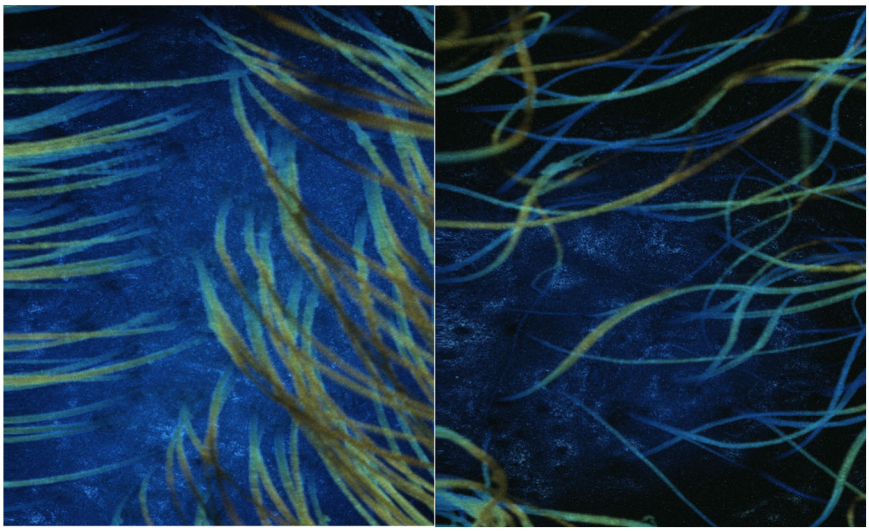

Fig. 3. 2D color projections of 3D grayscale data without accounting for scalp curvature and position. Some follicles are effectively not visible, especially in the image on the right. 2D, two-dimensional; 3D, three-dimensional.

prediction of bounding boxes, as all follicles are of similar size and labeling their exact outline with a bounding box would be error-prone and ambiguous in many cases. Furthermore, all follicles look alike—if they are even clearly visible at all—and can be located very close to each other, even to the point of obscurity. For these reasons, a multi-task fully convolutional CNN was trained (similar to YOLOv3 or SSD), but with a few major differences. The CNN predicts the locations of follicles and its "confidence" that a hair or follicle is present, but not a bounding box width or height, and thus does not require anchor boxes. The same CNN was also trained to predict the local hair density at each output grid cell to ensure that correct hair counts would be obtained in cases where multiple hairs are in close adjacency, which could otherwise be indistinguishable from multiple candidate predictions for a single hair.

Deep-Learning Model Architecture

A CNN was designed with a total of 21 convolutional layers ("same" padding in all but the 15th layer) with a filter size of 3×3, four 2×2-max pooling layers. Furthermore, a layer is included that "inverts" a 2×2 pooling operation and is combined with a skip-connection, which facilitates more precise location predictions of hairs and follicles. Each convolutional layer is followed by a leaky

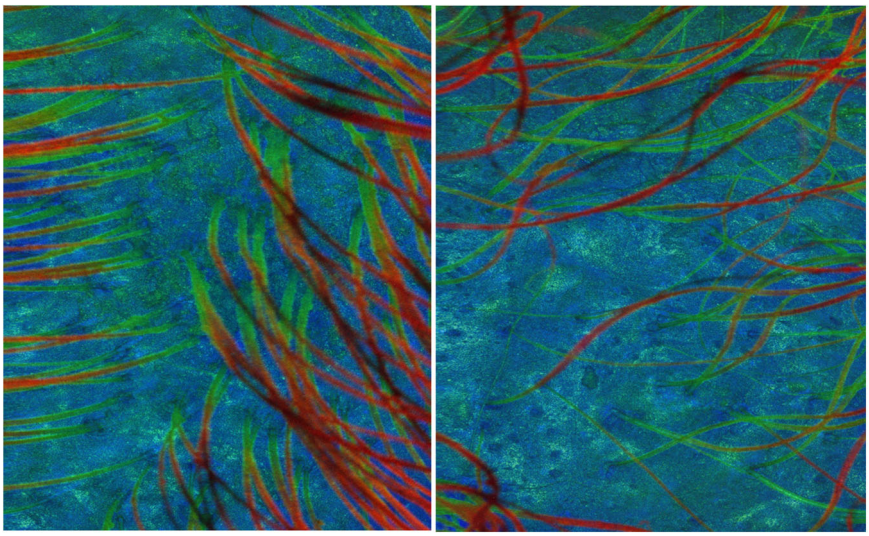

Fig. 4. 2D color projections of 3D grayscale data after automatically detecting and accounting for scalp curvature and position. Identical source data used, as in Figure 2. 2D, two-dimensional; 3D, three-dimensional.

rectified linear activation function [32] with a slope of 0.09, except for the final output layer, which is linear. The architecture is depicted in Figure 5.

The inputs to the CNN are en-face views (2D, colored, projections) with a resolution of 1000 by 1000 pixels. The model output is a grid of 122 by 122, with six output units per grid cell. The model outputs per cell are the: (1) confidence of seeing an object nearby, (2) classification of hair versus follicle, (3, 4) relative X/Y offsets from the grid cell's center to the predicted hair or follicle, and (5, 6) the number of hairs within a radius of 25 and 50 pixels of the grid cell center.

For every labeled follicle and hair in the training set, the model is tasked to make predictions using the closest four grid cells, their importance inversely weighted by distance, instead of just using the closest one, to obtain more robust predictions through a voting and averaging procedure. The local hair count predictions (5) and (6) are used in cases where multiple hairs are very close to each other (e.g., in the relatively frequent case of pili multigemini), which could otherwise be mistaken for multiple model predictions of the same hair in the post-processing phase.

Training and Post-processing

All OCT scans were converted to the color-coded en-face view and then passed on to the CNN to either train the model or obtain predictions.

To improve the model's accuracy and to significantly mitigate overfitting, data augmentation was used during training, a widely employed technique where the training data is perturbed in small, random ways to force the model to make correct predictions under additional and changing conditions. Augmentation techniques, such as random horizontal and vertical flipping of the images, translations, and scaling of up to 5% with appropriate padding, and rotations and shearing by up to $\pm 7^{\circ}$ were used to further enhance the model.

Model training was completed with the Adam optimizer [33] for 500 epochs, with an initial learning rate of 10^{-4}, which is subsequently exponentially decayed by a factor of 100 over the course of training.

The OCT data set contains 70 scans from 14 different patients. Of the 14 hair study patients, 12 were female and 2 male, with an age range of 21–76 and Fitzpatrick skin types I–IV. Seven of the 14 patients had non-scarring alopecia (androgenic alopecia; AGA); the other seven presented with scarring alopecia. The latter group includes two cases of chemotherapy-induced alopecia (CIA), two cases of frontal fibrosing alopecia (FFA), two cases of lichen planopilaris (LPP), and one case of dermatomyositis. Across the 70 scans, a total of 3412 hairs and 3141 follicles were found, and their locations were labeled. The model was evaluated using stratified 10-fold cross-validation with splits across patients. That is, in each fold, all scans of one or several patients were set aside as test

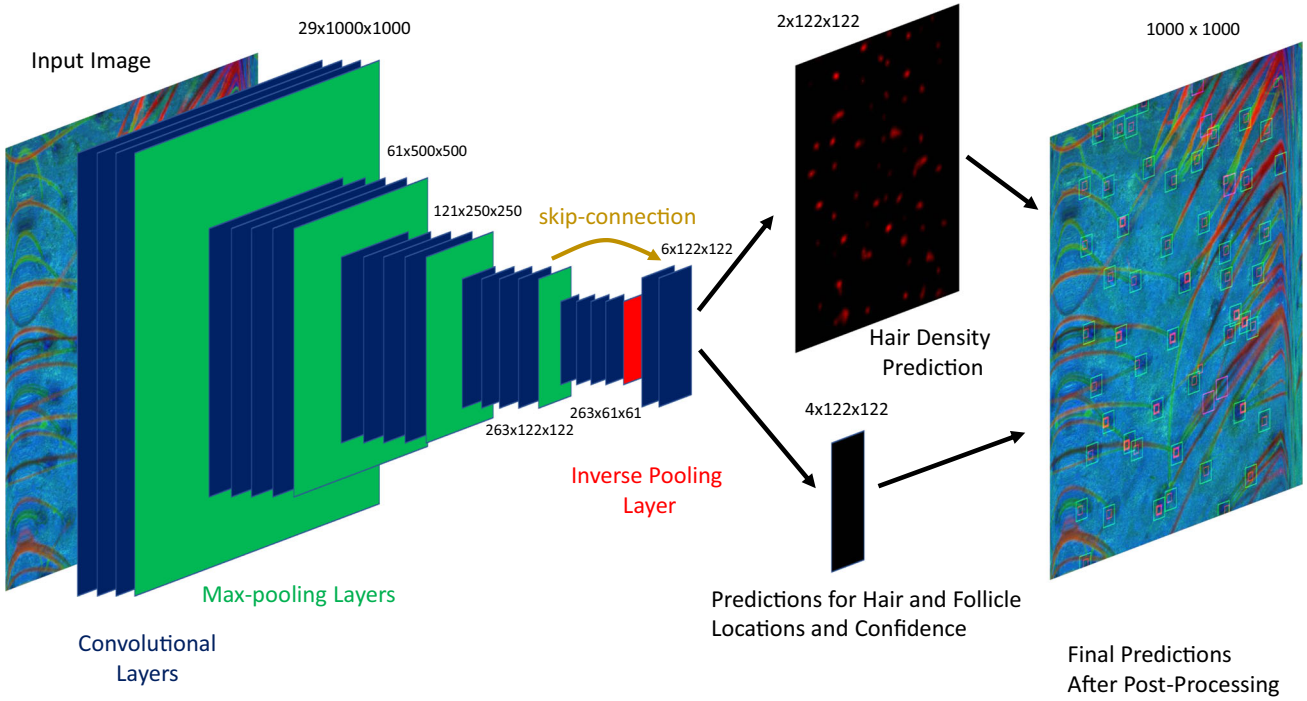

Fig. 5. The CNN network contains a total of 21 convolutional layers (shown in blue) and four max-pooling layers (shown in green). The final layer predicts six values per location in a grid of size 122 by 122 that is spaced evenly over the original 1000 by 1000 pixels image (the spacing is eight pixels and is due to the pooling layers). The first three convolutional layers each use 29 filters, the following four use 61 filters, and more for later layers, which keeps the computational requirements of this model at manageable levels. CNN, convolutional neural network.

data, and the scans of the remaining patients were treated as training data for CNN. This process was then repeated nine more times such that every patient's scans were selected as test data exactly once, and the splits were organized such that the number of test images was as similar as possible across each of them. This theoretically guarantees that the test results were not influenced by any patient-specific features that the model could have picked up during training.

For each en-face view, the model output is a tensor with a shape of $122 \times 122 \times 6$. In a post-processing step, the continuous-valued predictions for all 14884 grid cells were converted to useable discrete predictions of hairs, follicles, and their locations. As a first step, each prediction with confidence below 0.4 was discarded (the first out of the six predicted numbers per grid cell). All hair and follicle predictions were then placed into the original 1000×1000-pixel image as a Gaussian with a fixed standard deviation but weighted by confidence. Extraction of all local maxima of these combined Gaussians was then done to find all candidates for hair and follicle positions, the largest of which were kept in cases where multiple were located within a short distance. For hair predictions, the weighted average of the closest hair-density predictions at all predicted hair locations was collected, which allows the algorithm to predict the correct number of hairs even if there are multiple hairs within a short distance and an earlier step erroneously filtered too many locations out.

Evaluation Methods

The final evaluation scheme consisted of comparing the total number of human-labeled objects (i.e., hairs, hair-bearing follicles, and non-hair-bearing follicles) in each scan to the number of CNN-predicted objects. While this is the ultimate metric of interest, numerous errors made by the model could potentially be hidden. For example, if an equal number of false positive and false negative predictions are made, a perfect score would be inadvertently obtained.

To put the model under higher scrutiny, every individual prediction of the model was further compared with the closest labeled object. If no label was closer than 120 pixels (or approximately 0.8 mm) then the prediction was considered a false positive. Conversely, every label that was not matched to any close-by prediction was counted as a false negative. This process was done separately for hairs, hair-bearing follicles, and non-hair-bearing follicles.

RESULTS

The CNN model was trained on all 70 images in a stratified cross-validation procedure, as discussed earlier. Test predictions for the 20 redundantly labeled images were collected (Table 1). The data obtained indicate that CNN-predicted the total correct number of hairs to within 2–6% and the total number of follicles to within 0.1–15% when compared with the two human labelers. The relative differences between the human raters are 8% for hairs and 18% for follicles, so overall, the CNN predictions are "better" than the human inter-rater variability. However, on the level of individual predictions on individual scans, the CNN model makes a total of 633–692 false-positive and -negative predictions compared with only 432 discrepancies between the two human raters. Expressed in terms of discrepancies per total number of hair and follicle predictions, there are 33.5–36.7% mismatches for the CNN model compared with only 21–23.6% mismatches for the human raters.

The results on the 20 redundantly labeled images are close to those of the larger set of 70 OCT scans, with a 32% CNN hair and follicle prediction mismatch rate when compared with human labels (Table 2). The overall number of predicted hairs and follicles match the human numbers very well, with a difference of only 1.6% for hairs and 8.7% for follicles.

Results for the individual 70 scans can be found in the Supplementary Section (see Supplemental Table S1). Analyzing these results, we infer the average absolute difference in counts between CNN predictions to human annotations to be 9.1 per image for hairs, and 5.3 for follicles. The average number of hairs and follicles are 48.7 and 44.9 per image, respectively. This yields an estimated 18.7% error in hair counts per image and 11.8% error in follicle counts per image, when taking the human annotations as ground truth.

DISCUSSION

The purpose of this study was to construct and evaluate a machine-learning algorithm that could use data obtained from OCT images to accurately count hairs, hair-bearing follicles, and non-hair-bearing follicles. Such a program would add significant functionality to an already

TABLE 1. Comparison of CNN Test Set Predictions and Independent Annotations of Two Human Raters (H1 and H2) on a Set of 20 Processed OCT Scalp Images

	False pos	False neg	Class errors	Total # hairs	Total # follicles
CNN vs. H1	314	362	16	1044 vs. 1067 (−2.2%)	843 vs. 994 (−15.2%)
CNN vs. H2	403	228	2	1044 vs. 987 (+5.8%)	843 vs. 842 (+0.1%)
H1 vs. H2	98	330	4	987 vs. 1067 (−7.5%)	842 vs. 994 (−15.3%)

False positives are predictions with no counterpart nearby, while false negatives are missing predictions at locations where the counter-side had either a follicle or hair marked. "Class errors" are matches of hairs to follicles or vice versa.
CNN, convolutional neural network; OCT, optical coherence tomography.

TABLE 2. Comparison of CNN Test Set Predictions and Human Annotation (H1 and H2) on the Entire Set of 70 Images Using 10-Fold Cross-Validation Stratified by Patient

	False pos	False neg	Class errors	Total # hairs	Total # follicles
CNN vs. labels	1104 (16.8%)	948 (14.5%)	43	3465 vs. 3412 (+1.6%)	2869 vs. 3141 (−8.7%)

widely used technology, including the ability to automatically and quantitatively compute existing differences in scalp structures or changes over time with serial imaging.

For example, the follicular density that could imply follicular "dropout" has traditionally been used to discern scarring versus non-scarring alopecia, and this evaluation has routinely been completed by means of horizontal scalp biopsy and trichoscopy [23]. With the ability to quickly count hair-bearing and non-hair-bearing follicles with relative accuracy and precision, a 2D image could offer a reasonable idea of the active follicles in a certain area of the scalp and might give insight to the progression of hair loss under investigation or to the efficacy of various treatments.

Additionally, scalp biopsies are not practical for long-term evaluation of alopecia, should the follicular density or various hair-to-follicle ratios be desired to track any treatment over time. The algorithm presented by the authors offers an efficient way not only to obtain accurate follicular and hair density at one point in time but to do so longitudinally. An additional benefit includes the opportunity to evaluate multiple sites on the scalp, which may exhibit intrinsic differences, such as variable responses to treatment, without the need to subject a patient to multiple biopsies.

Assuming human evaluators can accurately identify and count the hairs and follicles by means of OCT images, the results obtained show that this model makes relatively good predictions when compared with investigators (accurate within approximately 10%), especially when taking factors such as time of operation into account. Like all machines, humans are also prone to error, and the process of counting follicles and hairs manually from an OCT image is both time-consuming and tedious. The OCT system takes only three seconds to complete one scan for one scalp location, and the CNN takes less than a second to predict all hairs and follicles in a processed en-face view while generating the en-face view from the OCT scan takes roughly half a minute using unoptimized python code.

Current limitations of this technology include the inability of the algorithm to discern different stages of hair growth, which would require a more in-depth evaluation of hair diameter and/or adnexal structures, neither of which would be easily discernable from the 2D en-face view used in this study. This is where histological evaluation reigns supreme. The most challenging aspect of developing the algorithm was to discern multiple hairs in close proximity to their associated follicles in order to predict whether the hairs came from the same, or different, follicles. The only way to affirm predictions with OCT image evaluation would be through comparison of the 2D en-face image with its associated vertical section of the overall 3D image, which together are considered the "orthogonal view," as vertical sections allow the user to visualize deeper structures that aid in this differentiation. Other limitations include the use of only two human evaluators in the inter-rater comparisons used against each other and against the algorithm; having additional human evaluators to add to the comparison validation would provide further insight into how accurately the model performs when compared with manual counting by an investigator. Creating a state-of-the-art non-invasive tool could be achieved by significantly increasing the size of the training data set and including additional (redundant) human annotators.

In future work, the system could be extended to estimate the depth of the bulbi to predict the stage of the hair cycle in individual follicles (i.e., deep bulbi in the anagen phase and high-set bulbi in the telogen phase), something that would not be possible with other non-invasive methods. This could be done by automatically analyzing the 3D data at locations that are identified as follicles in the 2D en-face view.

To the authors' knowledge, there is an obvious lack of methods to quantify follicular density by means of non-invasive methodology, and it is hoped that this technology may be of benefit in studies where these parameters will be useful in tracking treatment and management of various dermatologic conditions.

CONCLUSION

We present a process of non-invasively obtaining follicles and hair counts using OCT and deep learning. The process automatically returns all counts within seconds of completing the OCT scan of a patch of the scalp and with an accuracy that is within the discrepancy range of human raters.

ACKNOWLEDGMENTS

The work of GU and PB was, in part, supported by NSF grant NRT 1633631 to PB.

REFERENCES

1. Nardi AE, Brewster MF, Amor DM. Psychological impact of alopecia [9] (multiple letters). Br Med J 2005;331(7524):1084.
2. Whiting DA, Waldstreicher J, Sanchez M, Kaufman KD. Measuring reversal of hair miniaturization in androgenetic alopecia by follicular counts in horizontal sections of serial scalp biopsies: Results of finasteride 1 mg treatment of men and postmenopausal women. J Investig Dermatol Symp Proc 1999;4(3):282–24.
3. Whiting DA. Diagnostic and predictive value of horizontal sections of scalp biopsy specimens in male

pattern androgenetic alopecia. J Am Acad Dermatol 1993;28(5):755–763.
4. Gentile P, Garcovich S, Bielli A, Scioli MG, Orlandi A, Cervelli V. The effect of platelet-rich plasma in hair regrowth: A randomized Placebo-controlled trial. Stem Cells Transl Med 2015;4(11):1317–1323.
5. Jackson AJ, Price VH. How to diagnose hair loss. Dermatol Clin 2013;31(1):21–28.
6. Dhurat R, Saraogi P. Hair evaluation methods: Merits and demerits. Int J Trichol 2009;1(2):108–119.
7. Birch MP, Messenger JF, Messenger AG. Hair density, hair diameter and the prevalence of female pattern hair loss. Br J Dermatol 2001;144(2):297–304.
8. Gassmueller J, Rowold E, Frase T, Hughes-Formella B. Validation of TrichoScan® technology as a fully-automated tool for evaluation of hair growth parameters. Eur J Dermatol 2009;19(3):224–231.
9. Hoffmann R. TrichoScan: Combining epiluminescence microscopy with digital image analysis for the measurement of hair growth in vivo. Eur J Dermatol 2001;11(4):362–368.
10. Ramos PM, Miot HA. Female pattern hair loss: A clinical and pathophysiological review. An Bras Dermatol 2015;90(4):529–543.
11. Gambichler T, Jaedicke V, Terras S. Optical coherence tomography in dermatology: Technical and clinical aspects. Arch Dermatol Res 2011;303:457–473.
12. Gambichler T, Moussa G, Sand M, Sand D, Altmeyer P, Hoffmann K. Applications of optical coherence tomography in dermatology. J Dermatol Sci 2005;40(2):85–94.
13. Adabi S, Turani Z, Fatemizadeh E, Clayton A, Nasiriavanaki M. Optical coherence tomography technology and quality improvement methods for optical coherence tomography images of skin: A short review. Biomed Eng Comput Biol 2017;8:117959721771347
14. Ekelem C, Feil N, Csuka E. Optical coherence tomography in the evaluation of the scalp and hair: Common features and clinical utility. Skin Res Technol 2001;7(1):1–9.
15. Welzel J. Optical coherence tomography in dermatology: A review. Skin Res Technol 2001;7(1):1–9.
16. Baldi P, Chauvin Y. Neural networks for fingerprint recognition. Neural Comput 1993;5(3):402–418.
17. Baldi P. Deep learning in biomedical data science. Ann Rev Biomed Data Sci 2018;1:181–205.
18. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017;542(7639):115–118.
19. Zhu W, Liu C, Fan W, Xie X. Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification. IEEE Winter Conference on Applications of Computer Vision. 2018.
20. Urban G, Tripathi P, Alkayali T, et al. Deep learning localizes and identifies polyps in real time with 96% accuracy in screening colonoscopy. Gastroenterology 2018;155(4):1069–1078.
21. Wang J, Ding H, Bidgoli FA, et al. Detecting cardiovascular disease from mammograms with deep learning. IEEE Trans Biomed Imaging 2017;36(5):1172–1181.
22. Geras K, Wolfson S, Shen Y, et al. High-resolution breast cancer screening with multi-view deep convolutional neural networks. arXiv preprint 2017;1703:07047.
23. Urban G, Bache K, Phan DTT, et al. Deep learning for drug discovery and cancer research: Automated analysis of vascularization images. IEEE/ACM Trans Comput Biol Bioinf 2018;16(3):1029–1035.
24. Wang J, Fang Z, Lang N, Yuan H, Su M, Baldi P. A multiresolution approach for spinal metastasis detection using deep siamese neural networks. Comput Biol Med 2017;84:137–146.
25. Balyen L, Peto T. Promising artificial intelligence–machine learning–deep learning algorithms in ophthalmology. Asia-Pacific J Ophthalmol 2019;8(3):264–272.
26. Grewal PS, Oloumi F, Rubin U, Tennant MTS. Deep learning in ophthalmology: A review. Can J Ophthalmol 2018;53(4): 309–313.
27. Ulrich M, Themstrup L, De Carvalho N, et al. Dynamic optical coherence tomography in dermatology. Dermatology 2016;232(3):298–311.
28. Ren S, He K, Girshick R, Sun J. Faster R-CNN: Towards real-time object detection with region proposal networks. Paper presented at: Advances in Neural Information Processing Systems, 2015.
29. He K, Gkioxari G, Dollár P, Girshick R. Mask r-cnn. Paper presented at: Proceedings of the IEEE International Conference on Computer Vision, 2017.
30. Redmon J, Farhadi A. Yolov3: An incremental improvement. arXiv preprint 2018;1804:02767.
31. Liu W, Anguelov D, Erhan D, et al. Ssd: Single shot multibox detector. Paper presented at: European Conference on Computer Vision, 2016.
32. Xu B, Wang N, Chen T, Li M. Empirical evaluation of rectified activations in convolutional network. arXiv preprint 2015;1505:00853.
33. Kingma DP, Ba J. Adam: A method for stochastic optimization. ArXiv Preprint 2014;1412:6980.