



OPEN

Corneal pachymetry by AS-OCT after Descemet's membrane endothelial keratoplasty

Friso G. Heslinga¹✉, Ruben T. Lucassen¹, Myrthe A. van den Berg¹, Luuk van der Hoek¹, Josien P. W. Pluim¹, Javier Cabrerizo^{2,3}, Mark Alberti² & Mitko Veta¹

Corneal thickness (pachymetry) maps can be used to monitor restoration of corneal endothelial function, for example after Descemet's membrane endothelial keratoplasty (DMEK). Automated delineation of the corneal interfaces in anterior segment optical coherence tomography (AS-OCT) can be challenging for corneas that are irregularly shaped due to pathology, or as a consequence of surgery, leading to incorrect thickness measurements. In this research, deep learning is used to automatically delineate the corneal interfaces and measure corneal thickness with high accuracy in post-DMEK AS-OCT B-scans. Three different deep learning strategies were developed based on 960 B-scans from 50 patients. On an independent test set of 320 B-scans, corneal thickness could be measured with an error of 13.98 to 15.50 μm for the central 9 mm range, which is less than 3% of the average corneal thickness. The accurate thickness measurements were used to construct detailed pachymetry maps. Moreover, follow-up scans could be registered based on anatomical landmarks to obtain differential pachymetry maps. These maps may enable a more comprehensive understanding of the restoration of the endothelial function after DMEK, where thickness often varies throughout different regions of the cornea, and subsequently contribute to a standardized postoperative regime.

Corneal thickness is a key biomarker for corneal disorders, including Fuchs' endothelial dystrophy^{1,2}, keratoconus^{3,4}, and keratitis^{5,6}. Measurements on the corneal thickness, called pachymetry, enable detection of thickness changes that are indicative of restoration of corneal endothelial function after surgical treatment. For visualization of the restoring cornea, anterior segment optical coherence tomography (AS-OCT) has become the preferred imaging modality due to its high resolution and reproducibility^{7,8} (Fig. 1). While current OCT software works well for delineating the boundaries of healthy corneas, it often fails for corneas that are irregularly shaped due to pathology, or as a consequence of surgery (Fig. 2). Manual correction of the delineation mistakes is time consuming and not practical for a clinical setting.

In recent years, automated image analysis using deep learning⁹ has shown to be promising for ophthalmic applications¹⁰, including the analysis of AS-OCT images¹¹⁻¹⁴. Deep learning is a subset of machine learning techniques with models that contain many (typically millions of) trainable weights. These weights are iteratively updated with respect to some loss function which compares model predictions to ground truth labels. In contrast with classical machine learning techniques, no handcrafted features have to be selected as the relevant features are automatically learned from the (image) data. Recent work already showed the potential of deep learning for corneal pachymetry by AS-OCT, specifically for keratoconus¹⁵. We hypothesized that a similar approach could be used for corneal pachymetry for cases with irregular inner corneal curvature and/or structures that look similar to the corneal boundaries, both of which can lead to delineation failures by standard AS-OCT software.

In this study we focus on OCT scans acquired after Descemet's Membrane Endothelial Keratoplasty (DMEK)¹⁶. During DMEK, the diseased corneal endothelium and Descemet's membrane are replaced with a donor graft. After placement of the graft, a gas bubble is injected into the anterior chamber to support graft attachment to the host cornea. Both the procedural gas bubble and donor graft can mimic the appearance of the corneal interface and result in incorrect delineation.

We evaluate three different deep learning techniques that were developed or used for ophthalmology applications and shown to be highly effective. We validate our thickness measurements for the central 9 mm diameter (Fig. 1), whereas previous work only did so for 3.1 mm¹⁵. This is essential to assess corneal regeneration after DMEK surgery which uses a graft of ~ 8.5 mm. In addition, we present an automatic approach for reconstructing

¹Department of Biomedical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands. ²Ophthalmology Department, Rigshospitalet - Glostrup, Copenhagen, Denmark. ³Copenhagen Eye Foundation, Copenhagen, Denmark. ✉email: fgheslinga@gmail.com

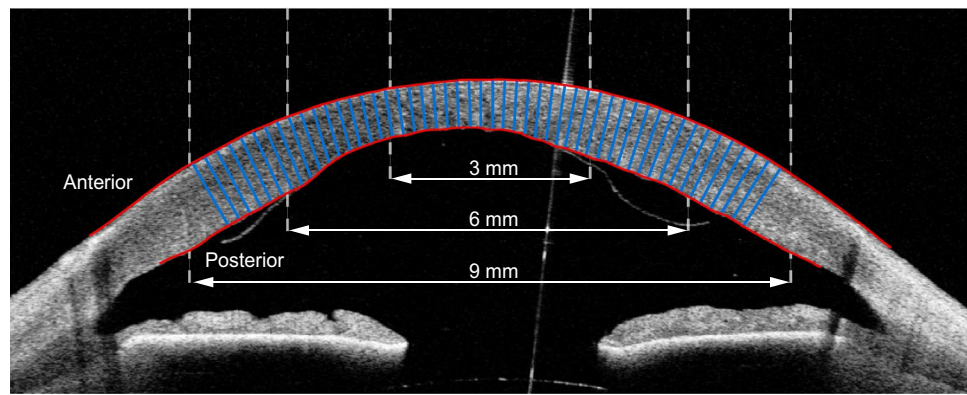


Figure 1. Single image (B-scan) from an AS-OCT scan, showing the cornea and the anterior chamber. This B-scan was cropped centrally and horizontally aligned as reported by Heslinga and Alberti¹⁴. Manual delineations of the corneal interfaces are shown in red. Corneal thickness is measured as the distance between the anterior and posterior interface, perpendicular to the anterior interface. The blue lines illustrate a subset of these thickness measurements. For evaluation of the thickness measurements, we distinguish the central 3 mm, 6 mm, and 9 mm diameter with respect to the corneal apex.

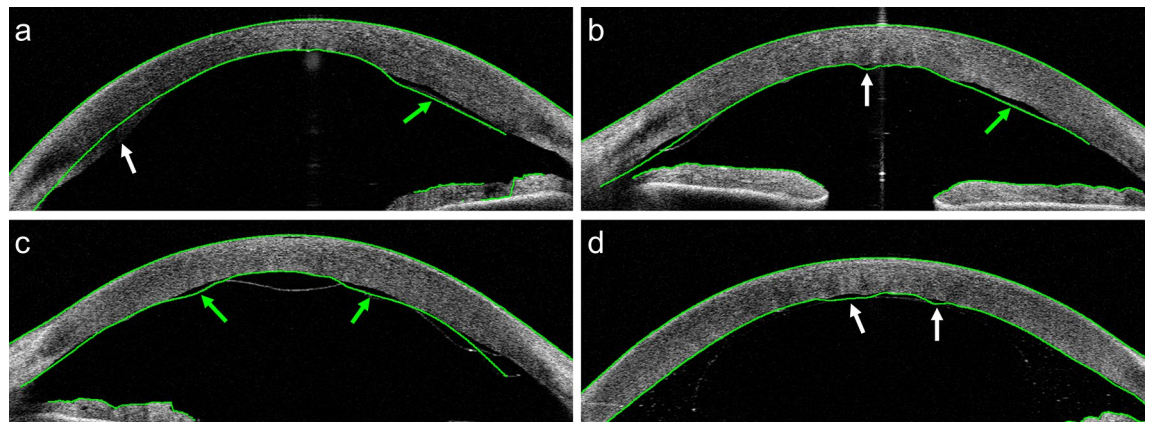


Figure 2. AS-OCT B-scans collected from patients after DMEK surgery. The green lines represent delineations of the (corneal) interfaces by the built-in software of the OCT system. These examples were selected to show the types of delineation errors encountered. In (a)–(c) the delineation partly follows the DMEK graft (green arrows) instead of the posterior interface. Other types of mistakes are indicated by white arrows: (a) Some of the posterior part of the cornea is missed. (b) The delineation does not follow the irregularly shaped interface in the center. (d) The system confuses the boundaries of the gas bubble used in DMEK with the posterior corneal interface.

differential pachymetry maps that locates the center of the cornea in subsequent images and visualizes thickness differences over time.

Results

AS-OCT data and annotations. The AS-OCT scans used in this study were collected as part of a randomized controlled trial conducted at the Department of Ophthalmology, Rigshospitalet - Glostrup, Denmark. The trial was designed to compare air and sulfur hexafluoride (SF6) DMEK surgery in patients with Fuchs' endothelial dystrophy or pseudophakic bullous keratopathy¹⁷. Repeat DMEK procedures and patients with prior keratoplasty were excluded. A total of 80 swept-source AS-OCT scans (CASIA2; Tomey Corp. Nagoya, Japan) from 68 participants were acquired either immediately after surgery, 1 week after surgery, or both. Each scan consists of 16 images (B-scans) acquired in a radial pattern, corresponding to 1280 B-scans in total.

AS-OCT scans were preprocessed similar as reported by Heslinga and Alberti¹⁴. In brief, a deep learning-based localization model was applied to each B-scan to identify the scleral spur, a landmark in the anterior chamber of the eye¹⁸. B-scans were horizontally aligned and cropped based on the scleral spur locations, centering around the corneal apex (Fig. 1). Final crop sizes were 960 by 384 pixels (width by height) with a pixel size of 15.0 μm . For a detailed description of the participant characteristics and processing methodology, we refer the reader to¹⁴.

Diameter (mm)	Models versus annotations				Inter-observer comparison		
	Patch-based CNN	U-Net	CNN with dim. red.	<i>p</i> value	1 versus 2	1 versus 3	2 versus 3
3	14.40 ± 0.69	13.94 ± 0.38	13.94 ± 0.25	0.26	23.49	14.69	17.95
6	14.80 ± 0.51	13.84 ± 0.22	14.17 ± 0.22	< 0.01	23.71	13.91	18.80
9	15.50 ± 0.59	13.98 ± 0.15	14.40 ± 0.15	< 0.01	23.39	13.66	19.26

Table 1. Mean absolute error in μm of corneal thickness predictions on test set. Comparisons represent deep learning models versus annotations (left) and annotator versus annotator (right). Mean \pm SD of 5 training repetitions. *p* values were calculated by one-way ANOVA and represent the chance that model performances are similar.

For each B-scan, the anterior corneal interface was annotated inside a 12 mm diameter from the radial center. A diameter of 10 mm was used for the posterior interface. Partial DMEK graft detachments were excluded from posterior interface annotations. The data set was randomly split on a participant level in a training set of 752 images, a validation set of 208 images and a test set of 320 images. B-scans of the training and validation set were annotated by one of three observers under supervision of a cornea specialist. The test set was annotated by all three observers to assess inter-observer variability.

Thickness measurements. Three deep learning-based models were trained to locate the anterior and posterior corneal boundaries: (1) a patch-based convolutional neural network (CNN), (2) a U-Net¹⁹ based model, and (3) a CNN with dimension reduction. Details about the model architectures and training process are provided in the methods section.

Corneal thickness was measured perpendicularly to the anterior interface (see Fig. 1), similar to²⁰. For each B-scan of the test set, we evaluated thickness for every pixel on the anterior interface inside a 3 mm, 6 mm, and 9 mm diameter. Corneal thickness estimates by the deep learning models were compared with all three sets of annotations (960 in total). Mean absolute errors (MAE) (shown in Table 1) are very similar for the three deep learning models and across different diameters. The smallest error was found for the U-Net model for the 6 mm diameter (13.84 μm), while the largest error was found for the patch-based CNN for the 9 mm diameter (15.50 μm). Apart from the latter, all mean absolute errors are smaller than one pixel (15.0 μm). The small standard deviations shown in Table 1 show the high repeatability over multiple training runs. In addition, we calculated the standard deviation over the mean absolute errors across the B-scans. Averaged over five training runs, these standard deviations for the central 9 mm diameter were 4.35 μm (patch-based), 4.77 μm (U-Net), and 4.90 μm (CNN with dimension reduction).

We investigated inter-observer variability by comparing the corneal thickness annotations between observers. The results of this comparison are shown in Table 1. Only for the combination of observer 1 versus 3, the mean absolute error is similar to that obtained with the deep learning models (13.66–14.69 μm), while the differences between the other combination of observers are substantially larger (17.95–23.71 μm). In addition, a cornea specialist manually assessed all delineations of the test set by the built-in software (version 3H.1) of the OCT system. In 90 out of 320 (28%) B-scans a delineation mistake occurred that resulted in a clinically relevant thickness inaccuracy within the central 9 mm. Out of these 90 B-scans, 74 (82%) included an inaccuracy that overlapped with the location of a partial DMEK graft detachment. For 37 (12%) B-scans the thickness errors were considered severe. Sixteen out of 20 (80%) AS-OCT scans contained at least one B-scan with a clinically relevant thickness inaccuracy, resulting in an incorrect pachymetry map.

Outlier analysis. We further inspected the origin of the deviations in thickness measurements between the deep learning models and the manual annotations by investigating the cases with the largest deviations. Figure 3 shows two example outliers with annotations and delineations by the CNN with dimension reduction. In Fig. 3a the remnant tissue from the host or donor prevents the graft from completely attaching at the right posterior side of the cornea. The tissue was correctly excluded from the annotation, but included in the delineation by the network.

Another example is presented in Fig. 3b, where the enlarged region shows a shortfall in detail of the annotation compared to the network delineation. Note that the graft is not entirely attached at the right side of the posterior interface, which was correctly recognized by the network and mistakenly included in the annotation.

Pachymetry mapping. Corneal thickness measurements from 16 radial B-scans were combined to construct pachymetry maps as shown in Fig. 4. The pachymetry map was divided into three circular regions with diameters of 3 mm, 6 mm, and 9 mm. The outer two rings were divided into octants where the average thickness is displayed. The inner circle displays the average of the four quarters, as well as the average apex thickness inside the central 1 mm diameter. Thickness values were mapped to corresponding colors of a discrete colormap. Similar to conventional pachymetry colormaps the corneal thickness at 600 μm is displayed in green, thinner regions in red, and thicker regions in blue^{2,21}.

Effect of training set size. The precise manual annotation of the corneal boundaries is a time-consuming process. For future projects it is useful to know whether a smaller set of annotated data can be used to obtain similar results. We therefore investigated the effect of the size of the annotated training set on the quality of the

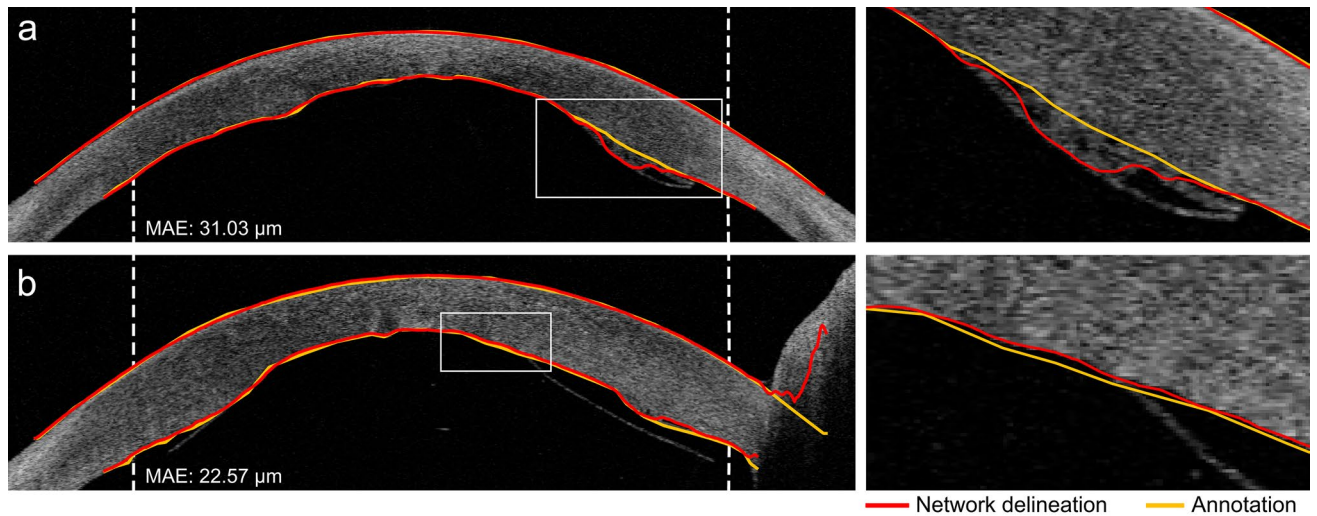


Figure 3. Two examples of B-scans including the annotations and delineations by the CNN with dimension reduction with substantial deviations in predicted thickness. The rectangular areas are enlarged and displayed to the right of the B-scan. Vertical dashed lines indicate the 9 mm diameter. Note that the thickness was not evaluated outside of the 9 mm diameter.

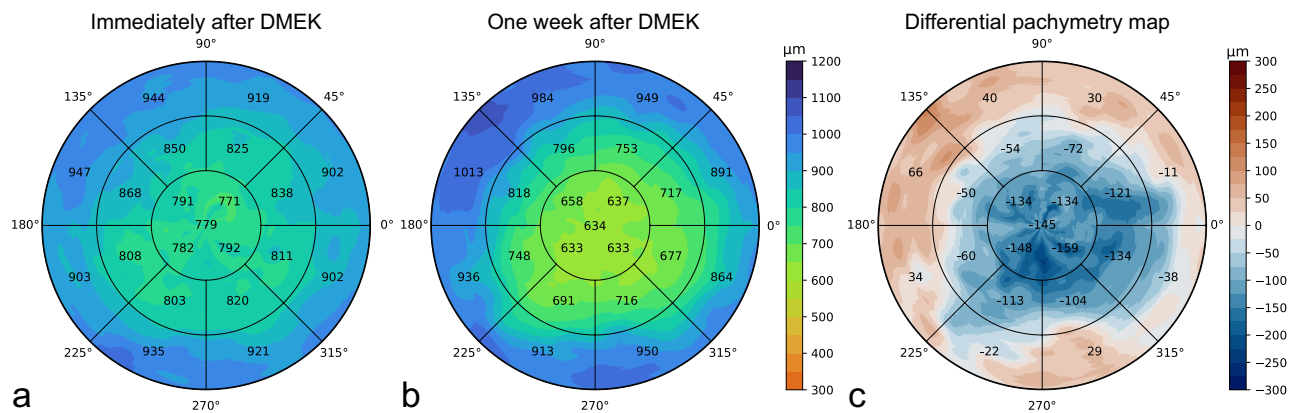


Figure 4. Example of pachymetry maps from one participant in the test set. (a) Pachymetry map of AS-OCT scan acquired immediately after DMEK; (b) Pachymetry map of AS-OCT scan acquired 1 week after DMEK; (c) Differential pachymetry map of difference in corneal thickness between (b) and (a).

thickness measurements. Table 2 shows thickness measurement errors on the test set B-scans when the deep learning models are trained with 100%, 50%, 25%, and 10% of the original training set. For the patch-based CNN and the CNN with dimension reduction we found only a marginal increase in the mean absolute error (of about 2 μm) for the whole central 9 mm range, when trained with only 10% of the data. In contrast, the U-Net model performance did decrease more substantially for the 9 mm range when trained with 10–50% of the data. Further inspection learned this was caused by some substantial error in the 6–9 mm range for a small portion of the B-scans.

Effect of data split. We studied the effect of heterogeneity within our data set using different splits of the data into training, validation and test sets. Next to the original split of the data, two more unique splits were created. Results of the experiment are shown in Table 3. For all three deep learning models, the mean absolute errors over the three splits were similar or slightly smaller than for the original test set.

Discussion

This research shows the feasibility of automated corneal thickness measurements in post-DMEK AS-OCT scans using deep learning. While our data set contains many examples of irregularly shaped corneas and partly detached DMEK grafts, all models are able to measure corneal thickness with an average error of 13.98 to 15.50 μm for the central 9 mm range. In comparison with a typical central corneal thickness of 540 μm ^{22,23}, this corresponds to an error of less than 3%. The quality of our thickness measurements is at least on par with manual annotations, as indicated by the inter-observer distance of 13.66 to 23.39 μm for the central 9 mm range. Based on the accurate thickness measurements, detailed pachymetry maps can be constructed. The preprocessing

Diameter	Training data (%)	Patch-based CNN	U-Net	CNN with dim. red.	<i>p</i> value
3 mm	100	14.40 ± 0.69	13.94 ± 0.38	13.94 ± 0.25	0.26
	50	14.28 ± 0.33	14.36 ± 0.51	15.04 ± 0.69	0.08
	25	15.04 ± 1.13	16.22 ± 1.45	15.03 ± 0.49	0.19
	10	15.93 ± 1.30	16.19 ± 1.38	16.76 ± 1.18	0.60
6 mm	100	14.80 ± 0.51	13.84 ± 0.22	14.17 ± 0.22	< 0.01
	50	14.65 ± 0.39	14.30 ± 0.46	15.10 ± 0.47	0.04
	25	15.16 ± 0.95	15.25 ± 0.86	14.92 ± 0.25	0.77
	10	16.08 ± 1.39	15.35 ± 0.74	15.88 ± 0.87	0.54
9 mm	100	15.50 ± 0.59	13.98 ± 0.15	14.40 ± 0.15	< 0.01
	50	15.23 ± 0.43	17.25 ± 5.97	15.24 ± 0.36	0.58
	25	15.71 ± 0.89	20.53 ± 7.94	15.12 ± 0.20	0.17
	10	16.76 ± 1.73	21.12 ± 6.24	16.08 ± 0.76	0.11

Table 2. Mean absolute error in μm of corneal thickness predictions on test set for varying partitions of the training set. Mean \pm standard deviation of 5 training repetitions. *p* values were calculated by one-way ANOVA and represent the chance that model performances are similar.

Diameter (mm)	Patch-based CNN	U-Net	CNN with dim. red.	<i>p</i> value
3	13.83 ± 0.84	12.53 ± 1.42	13.36 ± 0.50	0.34
6	14.39 ± 0.71	12.58 ± 1.31	13.34 ± 0.48	0.12
9	15.26 ± 0.88	12.77 ± 1.26	13.68 ± 0.50	0.04

Table 3. Results for multiple splits of the image data. Presented is the mean absolute error \pm standard deviation in μm over 3 unique test sets. *p* values were calculated by one-way ANOVA and represent the chance that model performances are similar.

based on the scleral spur locations largely eliminates spatial translation in the coronal plane and sagittal rotation between follow-up scans¹⁴. Our method does not correct for coronal rotation, but we expect the effect of this type of head tilt or eye rotation to be small. The mean localization error for a single scleral spur point was previously reported to be 0.155 mm. Given that the corneal apex location estimate follows from an average of 32 scleral spur point estimates, the registration error with respect to the corneal apex should be smaller than 0.155 mm.

Since follow-up scans can be registered, differential pachymetry maps can be constructed to monitor thickness changes. This may enable a more comprehensive understanding of the restoration of the endothelial function after DMEK, where thickness often varies throughout different regions of the cornea and the restoration of corneal thickness is associated with success of the procedure²⁴. Typically, only the central corneal thickness (CCT) is reported, while this single parameter does not necessarily reflect restoration of the full cornea after DMEK. The DMEK graft is about 8.5 mm²⁵ and partial graft detachment happens most frequently in the peripheral region^{26,27}. Detachment can sometimes be ambiguous, even in high-quality OCT imaging, as the graft can be close to the inner cornea yet not attached¹⁴. Local changes in corneal thickness could then be indicative of corneal restoration and thus graft attachment. The differential pachymetry maps presented in this research enable both qualitative and quantitative progression tracking options within the central 9 mm range, covering the whole region of the DMEK graft.

The data for this study only included post-DMEK AS-OCT scans of patients with previous Fuchs' endothelial dystrophy or pseudophakic bullous keratopathy. The application of the here presented deep learning models for scans from other pathologies or taken in different centra requires further research.

Three deep learning methods were compared in this research. Despite substantial differences in the approaches, the results for the different models were all at least on par with manual annotations. This could indicate that delineating the corneal boundaries is well-defined and relatively easily solvable using deep learning. This finding is in line with other research aiming at delineating layers in ophthalmic OCT imaging^{15,28–34}. Based on the thickness results for the models trained with 100% of the training data, none of the models seems to outperform the other. However, construction of a pachymetry map with the patch-based CNN takes considerably longer because of the large number of patches involved and the extensive post-processing steps. Kugelman et al.³⁴ compared several patch-based and fully-convolutional deep learning methods for segmentation of retinal layers and the chorio-scleral interface. They found that all models performed comparable for retinal layers, which are considered to have well-defined boundaries. However, for the more ambiguous chorio-scleral interface, the fully-convolutional methods outperformed the patch-based methods. The authors attributed this to the additional context available to the network when the whole image is processed at once.

Results of the deep learning models could not directly be compared with the delineations by the built-in software of the OCT system. Nevertheless, 28% of B-scans were found to contain delineations mistakes by the built-in software leading to a clinically relevant thickness inaccuracy. Moreover, 80% of the AS-OCT scans of

the test set contained at least one B-scan with a thickness inaccuracy of clinical relevance. We observed that these errors often occurred at locations of high clinical relevance, such as an irregularly shaped corneal center, or where the DMEK graft detached.

Training the models with smaller partitions of the data provides insight in the value of adding extra annotated data. Since no improvements were obtained by using 100% of the available training data compared to only 50%, it can be concluded that the performance has saturated, and additional data would not contribute to much further improvement. An exception could be the addition of data from rare cases or examples that led to errors in the current test set (e.g. Fig. 3a). For the U-Net model, training with 10% or 25% of the training data did result in an increase in the thickness error in the 6 to 9 mm region. Further inspection revealed that this was due to a small number of B-scans for which the segmentation failed. Interestingly, for the CNN with dimension reduction and the patch-based CNN, even training with 80 B-scans from 5 patients does not seem to reduce the performance of the thickness measurements substantially (MAE of 16.08 μm and 16.76 μm respectively). These results indicate that future projects on delineation of corneal boundaries could already be developed with less annotated data, yet still obtain reasonable results. In addition, the experiment with different data splits showed that similar performance could be achieved when different splits of the data were used.

The high resolution of the AS-OCT combined with deep learning for automated image processing supports fast and accurate analysis of the corneal thickness after DMEK. The here presented (differential) pachymetry maps enable tracking of local corneal thickness changes indicative of corneal restoration. As such, these tools can contribute to the ongoing research efforts towards further improvement of the DMEK procedure and management of the postoperative regime.

Methods

Models and training. We implemented three different deep learning models based on recent successful applications related to segmentation or interface delineation in ophthalmic OCT imaging. The first model is based on the patch-based approach by Wang et al.²⁹ which was adopted from earlier work on CNNs and graph search methods by Fang et al.²⁸. The patch-based model was previously used for the identification of five different retinal interfaces in spectral domain OCT images²⁹. Using patches of 33×33 pixels and a relatively shallow network of 5 trainable layers, Wang et al. achieved localization accuracies of 89–98%. The second model is based on the U-Net architecture which has become the de facto standard method for medical image segmentation¹⁹. Multiple adaptations of U-Net were evaluated by dos Santos et al.¹⁵ to segment three corneal layers in images captured with a custom-built ultrahigh-resolution OCT system. Using cross-validation, a mean segmentation-accuracy of 99.56% was achieved. The third model was also inspired by U-Net, but modified by Liefers et al.³⁰ to reduce the dimensionality and output one-dimensional arrays with y -locations of three retinal layers in OCT images. For the localization of these retinal layers, the authors obtained a mean absolute difference between the predictions and annotations of 1.31 pixels.

For our application of corneal interface localization, both the patch-based approach and the CNN with dimension reduction allow for direct delineation of the corneal interfaces. In contrast, a U-Net approach is used to segment the cornea and requires a post-processing step to obtain the interface delineations from the segmented mask. Details about the model adaptations, implementations, and training procedures are described below for each model. Depending on the model requirements we also adapted the preprocessing and post-processing steps. All models were implemented in Keras³⁵ with TensorFlow³⁶ backend and optimized using Adam³⁷.

Patch-based CNN. The architecture of the patch-based model was similar to that of Wang et al.²⁹ with 2 modifications: (1) 5×5 convolutional layers were replaced by two 3×3 convolutional layers as factorization is considered more efficient³⁸; (2) 3×3 average pooling operations in the final two layers were replaced with 2×2 max pooling operations. From the full images of training and validation sets, patches of 33×33 pixels were extracted for each x -coordinate where the interfaces had been annotated (center 12 mm for the anterior interface and center 10 mm for the posterior interface). Anterior and posterior patches were sampled using the respective annotations as center pixel locations. Similarly, for each x -coordinate within the central 12 mm, a non-interface patch was constructed for one random pixel not part of the interface annotations. All patches (1.70 million for the training set and 0.47 million for the validation set) were extracted prior to training to speed up the training process.

The model was optimized by minimizing the categorical cross-entropy between the pixel ground truth and model predictions. Online data augmentation was added by rotating the patches with a maximum of 30° . Based on preliminary experiments, the model was trained for 20 epochs with a variable learning rate: 0.001 for epoch 1 to 12, 0.0001 for epoch 13 to 16, and 0.00001 for epoch 17 to 20.

For evaluations on the test set, patches were extracted for all pixels within the center 12 mm (width) and processed by the trained model, predicting either anterior interface, posterior interface, or background for the center pixel of the patch. Based on preliminary results on the validation set, the following post-processing steps were performed for the pixels identified as interface: (1) small connected regions (0–250 pixels) were removed; (2) the largest connected region was considered to be true; (3) other regions were considered to be part of the true prediction only when those would be at the same height as the largest connected region; (4) per interface, y -values of positive pixels were averaged to obtain a single value per x -coordinate; (5) any remaining gaps were filled using linear interpolation of adjacent interface locations.

U-Net. As an alternative to directly delineating the interfaces, a U-Net¹⁹ was implemented to segment the whole cornea. The U-Net consisted of the standard 4 downsampling (and upsampling) segments and we included batch normalization and residual layers to accelerate training. Binary masks of the cornea were created using the

interface annotations. As a preprocessing step, the original images and masks were cropped to 800 by 256 pixels (width by height) and split into a superior and inferior half. This step was included to reduce the size of the input to the U-Net while doubling the number of training examples.

For optimization of the U-Net we experimented with different loss functions (Dice, binary cross-entropy, and weighted binary cross-entropy) and learning rate schedules. We found only minor differences in the results on the validation set and used binary cross-entropy for the final model. We trained for 30 epochs with an initial learning rate of 0.001 that was divided by two at every 3 epochs. We also experimented with data augmentation (brightness adaptation and addition of Gaussian noise) but did not identify any improvements on the validation results.

For evaluations on the test set, the inferior anterior crops were processed by the trained U-Net and combined. From the predicted mask, the maximum and minimum y -values were used to reconstruct the anterior and posterior interface respectively. In contrast to the patch-based model, the predicted y -values of the interfaces only consisted of integers.

CNN with dimension reduction. The architecture of the third model was designed by Liefers et al.³⁰ to return a one-dimensional array of y -coordinates for a two-dimensional image as network input. The model consists of a downsampling and upsampling path to incorporate a large contextual region. While U-Net uses direct shortcut connections to provide local context, this architecture resorts to so-called funneling subnetworks between the downsampling and upsampling path to resolve the mismatch in activation map height. The original network architecture was designed for images of 512 by 512 pixels. To avoid unnecessary computations, we adapted the architecture to work for images of 512 by 256 pixels. The downsampling path and all funneling subnetworks therefore contain one less downsampling operation and the upsampling path one less upsampling operation. 1×1 residual blocks at the lowest level were replaced by 1×3 residual blocks in our architecture. Furthermore, we experimented with the addition of batch normalization layers, which did not result in improved performance.

As a preprocessing step, the original images were cropped to 960 by 256 pixels (width by height) and split into a superior and inferior half with some overlap along the horizontal axis, resulting in model inputs of 512 by 256 pixels. Since the annotations did not span the entire width of 512 pixels, the network output layer was cropped to only include the annotated positions. The mean squared error between the annotation and predicted delineation was used as loss function. Training was done for 200 epochs with a learning rate set to 0.0002 at the start and divided by two after every 50 epochs. Based on preliminary experiments, we used the following data augmentation: B-scans were translated (≤ 10 pixels) or rotated ($\leq 12^\circ$) before inferior and superior crops were made. We also added uniform noise (≤ 0.05), Gaussian blurring with $\sigma \leq 1$, and sigmoidal contrast changes with a gain between [4, 5].

For evaluations on the test set, the superior and inferior crops were processed and combined by averaging the central overlapping section.

Thickness measurements and evaluation. Outputs of the deep learning models were y -values, describing the height of the anterior and posterior interface for each x -coordinate within the central 9 mm. The posterior interface of the cornea generally contains more irregular shapes after DMEK. We therefore measured corneal thickness perpendicularly to the anterior interface (see Fig. 1), similar to²⁰. First, the coefficient of proportionality was determined for the anterior interface. A 71 point moving average filter was used to reduce the effect of small deviations from the general curvature. We found that the proportionality coefficient was still affected by local irregularities after filtering with smaller filters, whereas larger filters introduced inaccuracies near the sides. The distance was then measured perpendicularly to a tangent with the corresponding coefficient of proportionality for every pixel on the anterior interface inside a 9 mm diameter. Performance of the models was measured by comparing the thicknesses predictions with the thicknesses following from the three sets of manual annotations. The mean absolute error was then calculated for a 3 mm, 6 mm, and 9 mm diameter. To obtain a measure of variation, we trained all models five times from scratch using different random seeds. Mean absolute errors and sample standard deviations shown in Tables 1 and 2 were based on these five repetitions. Comparison of the model performances was done for each region (3 mm, 6 mm, and 9 mm) to test for statistically significant differences. p values were calculated based on a one-way ANOVA and represent the chance that model performances were similar.

For assessment of the interface delineations errors by the built-in software of the OCT system, all B-scans of the test set were semi-quantitatively assessed by a cornea specialist. With the built-in drawing tool, missed parts of the cornea or areas mistakenly classified as cornea were selected. This was done approximately for the central 9 mm diameter although these B-scans were not centered or horizontally aligned. Sometimes the posterior delineation was not complete for the peripheral cornea. In such cases no extra misclassified area error was added; these regions were simply ignored. B-scans with a total incorrectly classified area of more than 0.1 mm^2 were considered to contain a clinically relevant inaccuracy. When this area was larger than 0.25 mm^2 the inaccuracy was considered severe.

Pachymetry mapping. Corneal pachymetry maps were constructed with the thickness profiles of the 16 radial B-scans (cross-sections) of a single AS-OCT scan. For the maps shown in Fig. 4 we used the CNN with dimension reduction to obtain one-dimensional thickness profiles of length 600 pixels corresponding to cross-sections within the central 9 mm. These arrays were plotted on a two-dimensional polar coordinate axis based on the spatial configuration of the cross-sections in the original AS-OCT scan. We then used cubic interpolation along the polar coordinates to fill in space between the cross-sections.

Before applying the deep learning models, B-scans were cropped based on the scleral spur locations, similar to¹⁴. First, a deep convolutional neural network was trained to find the scleral spur in the right and left side of

each B-scan. Per full AS-OCT scan, an ellipse was fitted through the scleral spur points of all 16 B-scans to ensure that the points lie in the same plane and to refine point locations. The refined scleral spur locations were used to horizontally align the B-scans by assuming that the left and right scleral spur points should be at the same height in the image. Then, the B-scan was centered by assuming that the corneal apex is located equidistant from both scleral spur points.

Since this preprocessing step automatically centers the (cropped) B-scans, follow-up scans and pachymetry maps are registered with respect to the corneal apex. The differential pachymetry map shown in Fig. 4 was obtained by subtracting the pachymetry map obtained after the procedure from the map obtained postoperative day 7.

Reduced training set and cross-validation. Partitions of the training set were made by randomly selecting all B-scans from a subset of the study participants. For 50% of the training set this equaled 24 participants (384 B-scans). For 25% 12 participants (192 B-scans) and for 10% 5 participants (80 B-scans). Partitions were randomly sampled for each of the five training repetitions, and partitions were the same for each deep learning model. All other training parameters were similar as for the 100% training set size models.

In order to examine the influence of the data split into training, validation, and test sets, two extra splits of the data were created, next to the original split. Splitting the data was done on a participant-level, with 18 participants in each test set and without overlap in participants between the test sets of the three splits. Training was done once per split, using the same hyperparameters as described for the other experiments. For the original split, only the first of five training repetitions was used. Since only a single set of manual annotations was available for the original training and validation data, we opted to also use only a single annotation set for the original test data set for this experiment.

Additional information

The code for constructing the (differential) pachymetry maps is available via https://github.com/tueimage/corneal_pachymetry/.

Received: 21 February 2021; Accepted: 16 June 2021

Published online: 07 July 2021

References

- Kopplin, L. J. *et al.* Relationship of Fuchs' endothelial corneal dystrophy severity to central corneal thickness. *Arch. Ophthalmol.* **130**(4), 433–439 (2012).
- Patel, S. V., Hodge, D. O., Treichel, E. J., Spiegel, M. R. & Baratz, K. H. Predicting the prognosis of Fuchs endothelial corneal dystrophy by using Scheimpflug tomography. *Ophthalmology* **127**(3), 315–323 (2020).
- Ambrósio Jr, R., Alonso, R. S., Luz, A. & Coca Velarde, I G. Corneal-thickness spatial profile and corneal-volume distribution: Tomographic indices to detect keratoconus. *J. Cataract Refract. Surg.* **32**(11), 1851–1859 (2006).
- Li, Y. *et al.* Keratoconus diagnosis with optical coherence tomography pachymetry mapping. *Ophthalmology* **115**(12), 2159–2166 (2008).
- Cook, C. & Langham, M. Corneal thickness in interstitial keratitis. *Br. J. Ophthalmol.* **37**, 301–304 (1953).
- Wilhelmus, K. R., Sugar, J., Hyndiuk, R. A. & Stulting, R. D. Corneal thickness changes during herpes simplex virus disciform keratitis. *Cornea* **23**(2), 154–157 (2006).
- Lim, S. H. Clinical applications of anterior segment optical coherence tomography. *J. Ophthalmol.* **2015**, 1–12 (2015).
- Wang, S. B., Cornish, E. E., Grigg, J. R. & McCluskey, P. J. Anterior segment optical coherence tomography and its clinical applications. *Clin. Exp. Optom.* **102**(3), 195–207 (2019).
- Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- Ting, D. S. W. *et al.* Artificial intelligence and deep learning in ophthalmology. *Br. J. Ophthalmol.* **103**, 167–175 (2019).
- Xu, B. Y., Chiang, M., Pardeshi, A. A., Moghimi, S. & Varma, R. Deep neural network for scleral spur detection in anterior segment OCT images: The Chinese American eye study. *Trans. Vis. Sci. Technol.* **9**(2), 18 (2020).
- Fu, H. *et al.* A deep learning system for automated angle-closure detection in anterior segment optical coherence tomography images. *Am. J. Ophthalmol.* **203**, 37–45 (2019).
- Treder, M., Lauermann, J. L., Alnawaiseh, M. & Eter, N. Using deep learning in automated detection of graft detachment in Descemet membrane endothelial keratoplasty: A pilot study. *Cornea* **38**(2), 157–161 (2019).
- Heslinga, F. G., Alberti, M., Pluim, J. P. W., Cabrerizo, J. & Veta, M. Quantifying graft detachment after Descemet's membrane endothelial keratoplasty with deep convolutional neural networks. *Trans. Vis. Sci. Technol.* **9**(2), 48 (2020).
- Dos Santos, V. A. *et al.* CorneaNet: Fast segmentation of cornea OCT scans of healthy and keratoconic eyes using deep learning. *Biomed. Opt. Express* **10**(2), 622–641 (2019).
- Melles, R. J., Ong, T. S., Ververs, B. & van der Wees, J. Descemet membrane endothelial keratoplasty (DMEK). *Cornea* **25**(8), 987–990 (2006).
- Alberti, M. Air versus SF6 for Descemet's membrane endothelial keratoplasty (DMEK). <https://clinicaltrials.gov/ct2/show/NCT03407755>. Accessed May 9, 2020.
- Ang, M. *et al.* Anterior segment optical coherence tomography. *Prog. Retin. Eye Res.* **66**, 132–156 (2018).
- Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Proc. Int. Conf. Med. Image Comput.-Assist. Intervent* 234–241 (2015).
- Li, Y., Shekhar, R. & Huang, D. Corneal pachymetry mapping with high-speed optical coherence tomography. *Ophthalmology* **113**(5), 792–799 (2006).
- Bourges, J. L. *et al.* Average 3-dimensional models for the comparison of Orbscan II and Pentacam pachymetry maps in normal corneas. *Ophthalmology* **116**(11), 2064–2071 (2009).
- Ma, R. *et al.* Distribution and trends in corneal thickness parameters in a large population-based multicenter study of young Chinese adults. *Investig. Ophthalmol. Vis. Sci.* **59**, 3366–3374 (2018).
- Hashemi, H. *et al.* The distribution of corneal thickness in a 40- to 64-year-old population of Shahroud, Iran. *Cornea* **30**(12), 1409–1413 (2011).
- Vasiliaskaitė, I. *et al.* Descemet membrane endothelial keratoplasty: Ten-year graft survival and clinical outcomes. *Am. J. Ophthalmol.* **217**, 114–120 (2020).

25. Röck, T., Bramkamp, M., Bartz-Schmidt, K. U., Röck, D. & Yörük, E. Causes that influence the detachment rate after Descemet membrane endothelial keratoplasty. *Graefes. Arch. Clin. Exp. Ophthalmol.* **253**, 2217–2222 (2015).
26. Bucher, F. *et al.* Spontaneous long-term course of persistent peripheral graft detachments after Descemet's membrane endothelial keratoplasty. *Br. J. Ophthalmol.* **99**, 768–772 (2015).
27. Deng, S. X., Sanchez, P. J. & Chen, L. Clinical outcomes of Descemet membrane endothelial keratoplasty using eye bank-prepared tissues. *Am. J. Ophthalmol.* **159**, 590–596 (2015).
28. Fang, L. *et al.* Automatic segmentation of nine retinal layer boundaries in OCT images of non-exudative AMD patients using deep learning and graph search. *Biomed. Opt. Express*. **8**(5), 2732–2744 (2017).
29. Wang, Y. Z., Galles, D., Klein, M., Locke, D. G. & Birch, D. G. Application of a deep machine learning model for automatic measurement of EZ Width in SD-OCT Images of RP. *Transl. Vis. Sci. Technol.* **9**(2), 15 (2020).
30. Liefers, B., González-Gonzalo, C., van Ginneken, B. & Sánchez, C. I. Dense segmentation in selected dimensions: application to retinal optical coherence tomography. In *Proc. Int. Conf. Med. Imaging with Deep Learning* 337–346 (2019).
31. Roy, A. G. *et al.* ReLayNet: Retinal layer and fluid segmentation of macular optical coherence tomography using fully convolutional networks. *Biomed. Opt. Express* **8**(8), 3627–3642 (2017).
32. Hamwood, J., Alonso-Caneiro, D., Read, S. A., Vincent, S. J. & Collins, M. J. Effect of patch size and network architecture on a convolutional neural network approach for automatic segmentation of OCT retinal layers. *Biomed. Opt. Express* **9**, 3049–3066 (2018).
33. Kugelman, J., Alonso-Caneiro, D., Read, S. A., Vincent, M. J. & Collins, S. J. Automatic segmentation of OCT retinal boundaries using recurrent neural networks and graph search. *Biomed. Opt. Express* **9**, 5759–5777 (2018).
34. Kugelman, J. *et al.* Automatic choroidal segmentation in OCT images using supervised deep learning methods. *Sci. Rep.* **9**(1), Art. No. 13298 (2019).
35. Keras (2015). Software available from keras.io.
36. TensorFlow: Large-scale machine learning on heterogeneous systems (2015). Software available from tensorflow.org.
37. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
38. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the Inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016).

Acknowledgements

This research is financially supported by the TTW Perspectief program and Philips Research.

Author contributions

F.H., M.A. and M.V. designed the study. M.A. and J.C. provided the material and approved the study. R.L., M.B., L.H. and F.H. implemented the models, conducted the experiments, and analyzed the results. M.A., J.C., J.P. and M.V. reviewed the analysis. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to F.G.H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021