Keratoconus Classifier for Smartphone-based Corneal Topographer

Siddhartha Gairola[1], Pallavi Joshi[2], Anand Balasubramaniam[2], Kaushik Murali[2], Nipun Kwatra[1], Mohit Jain[1]

Abstract— **Keratoconus is a severe eye disease that leads to deformation of the cornea. It impacts people aged 10-25 years and is the leading cause of blindness in that demography. Corneal topography is the gold standard for keratoconus diagnosis. It is a non-invasive process performed using expensive and bulky medical devices called corneal topographers. This makes it inaccessible to large populations, especially in the Global South. Low-cost smartphone-based corneal topographers, such as *SmartKC*, have been proposed to make keratoconus diagnosis accessible. Similar to medical-grade topographers, *SmartKC* outputs curvature heatmaps and quantitative metrics that need to be evaluated by doctors for keratoconus diagnosis. An automatic scheme for evaluation of these heatmaps and quantitative values can play a crucial role in screening keratoconus in areas where doctors are not available. In this work, we propose a dual-head convolutional neural network (CNN) for classifying keratoconus on the heatmaps generated by *SmartKC*. Since *SmartKC* is a new device and only had a small dataset (114 samples), we developed a 2-stage transfer learning strategy—using historical data collected from a medical-grade topographer and a subset of *SmartKC* data—to satisfactorily train our network. This, combined with our domain-specific data augmentations, achieved a sensitivity of 91.3% and a specificity of 94.2%.**

I. INTRODUCTION

Keratoconus is an eye disease that causes the cornea to become thin, leading to a conical shape. It results in blurred vision, irregular astigmatism, and partial/complete blindness. Studies have shown that keratoconus is highly prevalent in the Global South, *e.g.*, in central India, 1 in 44 individuals were found to suffer from keratoconus [1]. The disease generally affects people between the ages of 10 and 25, and may progress slowly for years. Early diagnosis is crucial to provide timely treatment and prevent worsening of disease. The clinical gold standard to diagnose keratoconus is using a technique called corneal topography. This is done by a sophisticated medical device called corneal topographer. It measures the shape of the corneal surface and outputs curvature heatmaps, helping the doctor to diagnose keratoconus. Although highly accurate, these devices are expensive, bulky (non-portable), and require a trained technician to operate. Such factors limit accessibility and also make frequent mass screening of keratoconus hard in remote areas.

In our prior work, we proposed *SmartKC* [2] (Figure 1), a low-cost smartphone-based corneal topographer that generates topography heatmaps similar to a medical-grade corneal topographer. Through a clinical study comprising of 101 eyes (57 participants) and evaluation by four ophthalmologists, we showed that *SmartKC* achieved a sensitivity of 94.1% and specificity of 100.0%. However, *SmartKC*-generated heatmaps need to be evaluated by doctors for diagnosis, and with only 1 doctor for more than 1000 people in the

¹ Microsoft Research, India {t-sigai, nkwatra, mohja}@microsoft.com
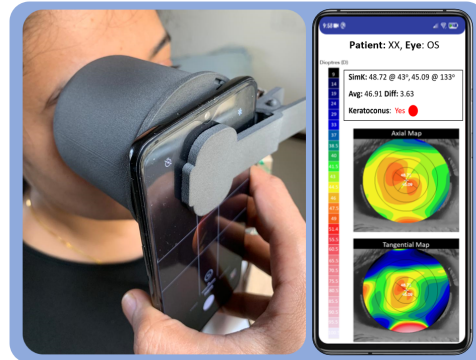² Sankara Eye Hospital, Bengaluru, India

Fig. 1. Left: *SmartKC*, a smartphone-based corneal topographer. Right: Outputs generated by *SmartKC* (axial heatmap, tangential heatmap, sim-K values, and automatic keratoconus diagnosis by our proposed model).

Global South [3], it is hard to obtain a doctor's evaluation for every keratoconus screening test. Additionally, evaluations by doctors suffer from subjectivity. For instance, in *SmartKC* evaluation, we found that for 30.7% of eyes, at least one of the four doctors had a different diagnosis than the rest (this number was 42.6% for Keratron device). Having an accurate automated method to detect keratoconus that works for low-cost topographers can help alleviate such issues and enable mass screening to identify keratoconus patients.

Recent works have demonstrated the efficacy of using deep neural networks for accurately detecting keratoconus using color-coded heatmaps obtained from clinical topographers, such as devices based on Optical Coherence Tomography [4], Scheimpflug imaging [5, 6, 7], and placido disc reflection [8, 9]. However, such methods have been limited to medical-grade topographers. We extend this line of work to adapt it to heatmaps generated from the low-cost *SmartKC* device.

In this work, we propose a deep CNN with two classification heads, that takes *SmartKC*-generated heatmaps as input and provides keratoconus diagnosis as output. Since we had very few samples (114 eyes) from *SmartKC* [2], to effectively train our network we propose a 2-stage transfer learning strategy: (a) pre-training on ImageNet followed by fine-tuning on historical data (2110 samples) collected from Optikon Keratron, a medical-grade topographer, and (b) further fine-tuning on 50% of the *SmartKC*-generated heatmaps. This improved the performance significantly (sensitivity=82.1%, specificity=85.6%), compared to the standard fine-tuning on only *SmartKC* samples (sensitivity=65.2%, specificity=76.5%). On further analysis, we found that Keratron-generated heatmaps had fixed scaling and centering due to the stable head- and chest-rest setup, while *SmartKC* heatmaps had variations owing its handheld nature. Hence, we applied domain-specific augmentations to simulate similar behavior in the Keratron dataset, which helped us achieve a sensitivity of 91.3% and specificity of 94.2%.

The main contributions of our work are:

1) a dual-head CNN-based keratoconus detection algorithm for *SmartKC*, a smartphone-based topographer,
2) a 2-stage transfer learning strategy and domain-specific augmentations for training with very few samples, and
3) evaluation of our method on topography heatmaps from actual patients obtained by *SmartKC* and also a medical-grade topographer (Keratron).

II. RELATED WORK

Prior research has proposed automated methods using quantitative indices [10] (like Percentage Probability of Keratoconus (PPK), Cone Location and Magnitude Index [11]), statistical methods [12, 13], and traditional machine learning algorithms [9, 14, 15] (like logistic regression, K-nearest neighbour, clustering, decision trees, and random forests) to diagnose keratoconus from corneal topography data. However, with the advent of deep learning and its high performance on image classification tasks, deep learning based methods are being widely explored for keratoconus diagnosis and have been shown to be highly accurate [5, 6, 16].

Lavric and Valentin [16] propose to use a CNN-based classifier for keratoconus detection, however they train and test their model only on synthetic eye data. Zéboulon et al. [6] investigate the performance of CNNs to classify corneal topography heatmaps into normal and keratoconus. They evaluate the model on 3000 samples from real patients obtained using Orbscan, a medical-grade topographer, and achieve classification accuracy of 99.3%. Similarly, a ResNet152 based classification model has been proposed [8]. They train and test their model on 354 samples, and achieve sensitivity and specificity of above 90%. Chen et al. [5] use corneal tomography scans, with four type of heatmaps (axial, anterior elevation, posterior elevation, and pachymetry) per sample, and adopt a VGG16 model to learn a keratoconus detection classifier. They train their model on 1115 samples and test it on 279 samples, achieving a sensitivity of 98.5% and specificity of 90.0%.

These systems have achieved high performance for keratoconus detection, but have been limited to training and testing on heatmaps generated by medical-grade topographers. In spite of the progress made in developing low-cost portable topographers [2, 17, 18, 19, 20, 21, 22], there has been no work on adapting a model designed for data obtained from medical-grade topographers that is suitable for such low-cost alternatives. In our work, we aim to bridge this gap by proposing a CNN-based solution that accurately detects keratoconus from corneal topography heatmaps obtained using both, *SmartKC* (a low-cost portable topographer) and Keratron (a medical-grade topographer).

III. METHOD

Dataset: Our dataset is divided into two parts:

SmartKC-data: This an augmented version of the dataset used in [2], collected from 64 patients, with 114 eye samples obtained from both Keratron and *SmartKC* device deployed at the Sankara Eye Hospital in Bengaluru, India. It comprises of 68 non-keratoconus and 46 keratoconus eyes. The ground truth data for each eye was keratoconus classification, as
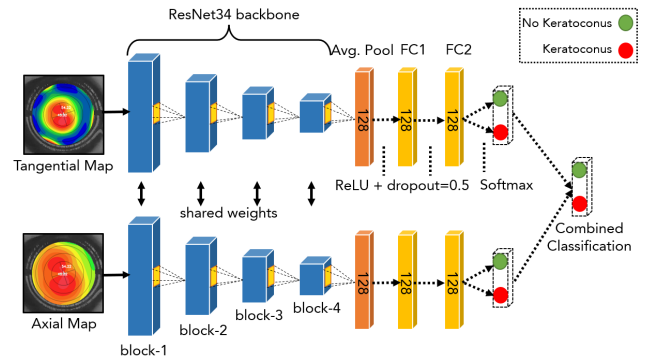

Fig. 2. *Proposed network architecture.*

diagnosed by a senior ophthalmologist at the hospital. This data was collected during Jun-Aug'21.

Keratron-data: Since 114 samples is small for training a deep learning model, we retrospectively downloaded anonymized data from the Keratron device database for all the patients who took the corneal topography examination at the hospital from April'08 to May'10. The dataset consists of 2110 samples (1637 non-keratoconus and 473 keratoconus). Each sample comprises of three images (axial heatmap, tangential heatmap, and placido ring), and three quantitative metrics (sim-K1, sim-K2, and PPK). The ground truth class (keratoconus vs non-keratoconus) was obtained based on the PPK-based classification[1]. All the details of our study were approved by the hospital's Institutional Review Board.

Data Preprocessing: All heatmaps are cropped and resized to a fixed resolution of 512×512. The heatmaps are then standardized by performing Z-normalization on each channel of the RGB image ($\tilde{X} = (X - \mu)/\sigma$, where the mean μ and standard deviation σ are computed for the entire dataset).

Network Architecture: Figure 2 illustrates the architecture of the proposed CNN. The network is organized into two branches—one each for axial and tangential heatmaps—with a shared convolutional backbone, followed by two distinct feed-forward classifiers, one for each branch. The shared backbone comprises of the convolutional layers from a ResNet34 model. The classifiers consist of two *128-d* fully connected (FC) layers with *ReLU* activations. The last layer of each classifier applies *softmax activation* to model class probabilities. *Dropout* is added to each FC layer to prevent overfitting. The network is jointly trained via a *weighted binary cross-entropy*[2] loss to minimize the loss for the two-class (keratoconus vs non-keratoconus) classification task.

Training Procedure: We followed a 2-stage transfer learning strategy to train our network. In stage-1, the network backbone is initialized with weights from a ResNet34 network pre-trained on ImageNet [23], and then *fine-tuned* on the Keratron-data with a 90:10 train:validation split. This achieved good performance on the Keratron heatmaps, however failed on the *SmartKC* heatmaps (see Tables I, II). On deeper analysis of the data, we noticed that as *SmartKC* is a portable hand-held device, the heatmaps generated by it

---

[1]PPK< 20% is non-keratoconus, 20% ≤PPK< 45% is suspect keratoconus, and 45% ≤ PPK is keratoconus, as per Keratron user-manual.
[2]where the weights are inversely proportional to the number of samples in each class, to address class imbalance.

Eval. Data	Model	Se	Sp	Acc
SmartKC-data (57 samples)	SVM	80.4%	100.0%	92.1%
	Dual-head CNN†	65.2%	76.5%	71.9%
	Dual-head CNN*	91.3%	94.2%	93.1%
Keratron-data (114 samples)	PPK	89.4%	94.2%	92.2%
	SVM	78.3%	95.5%	88.6%
	Dual-head CNN‡	94.7%	93.4%	93.9%

TABLE I

*Sensitivity (S_e), Specificity (S_p), Accuracy (Acc) of the proposed method. †: fine-tuned only on SmartKC-data. *: fine-tuned on Keratron-data (stage-1) and 50% of SmartKC-data (stage-2). ‡: fine-tuned only on Keratron-data (stage-1).*

Fig. 3.   *t-sne plots (a) before training, and (b) after training, depicts a good separation of the learned features for keratoconus vs non-keratoconus test samples. (Note: This is shown for the fully-connected features for axial heatmaps; the tangential heatmap features show a similar separation.)*

had variations in terms of zoom level, translation, and aspect ratio. In comparison, the heatmaps obtained from Keratron had minimal variations as they are captured using a stable head- and chin-rest setup. To handle this difference in image characteristics, we added data augmentations (described below) and a stage-2 fine-tuning. In stage-2, we fine-tuned the network on a subset (50%, 57 samples) of the *SmartKC*-data to ensure that the network learns these variations. This helped further improve the performance (see Table II).

Data Augmentation: As the Keratron-data is relatively small (2110 samples) for training a deep neural network, we added *domain-specific data augmentations* like horizontal-flip, rotation ($0^o - 10^o$), scaling ($0.6x - 1.4x$) and translation ($1 - 10\%$). Besides improving generalization of the network, this also helped in making the network robust to variations in the *SmartKC* captured images. Additionally, we also employ *mixup data augmentations* [24]. To combine two randomly selected samples x_i, x_j along with their labels y_i, y_j in the training data, a new sample $\tilde{x}$ and its label $\tilde{y}$ is generated as follows using the mixup scheme: $\tilde{x} = \lambda x_i + (1 - \lambda)x_j$, $\tilde{y} = \lambda y_i + (1 - \lambda)y_j$, where y_i, y_j are one-hot encoded class labels of the feature vectors x_i, x_j, and $\lambda \in [0, 1]$ is a random number generated according to a Beta distribution [24]. Mixup also helps with addressing the data imbalance problem typical in medical datasets.

IV. EXPERIMENTS AND RESULTS

Implementation Details: We train our network on a Tesla v100 GPU on a Microsoft Azure VM. We used the SGD optimizer with momentum of 0.9 and a batch size of 32, with a weighted random sampler to sample mini-batches uniformly from each class. For the stage-1 fine-tuning on Keratron-data, the network is trained for 200 epochs using a fixed learning rate of 1e-3 (took 73.5 minutes). For the stage-2 fine-tuning on *SmartKC*-data, we start with a 1e-4 learning rate with linear decay for 100 epochs (took 3.5 minutes). The inference time for a single image was 0.3 seconds. Note: During both stages of training, *ResNet34 blocks 1-3* of the backbone were kept frozen.

Results: We evaluate the performance of our network on the task of keratoconus detection on heatmaps from the *SmartKC*-data. In stage-2, the network is *fine-tuned* on 50% of *SmartKC*-data and evaluated on the remaining 50%; we perform this evaluation on 5 random splits, and report the average values over the 5 runs. We report sensitivity ($S_e = P_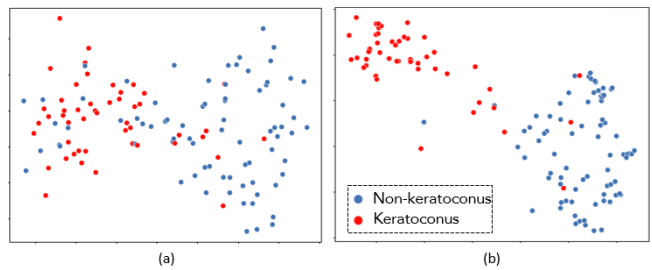k/N_k$), specificity ($S_p = P_n/N_n$), and accuracy ($Acc = \frac{P_k+P_n}{N_k+N_n}$) values (Table I), where P_i, N_i are the number of correctly classified samples and total number of samples in class $i \in \{keratoconus, non\text{-}keratoconus\}$, respectively. The final output of the dual-head CNN model is computed as the worst case possibility from the axial and tangential branch, *i.e.*, if either of the branches predict the heatmap as keratoconus, we consider the final label as keratoconus.

Our network, comprising of dual-head CNN and 2-stage transfer learning, achieved an accuracy of 93.1% (S_e=91.3%, S_p=94.2%) on the *SmartKC*-data (Table I). The 2-stage transfer learning showed significant improvement compared to traditional transfer learning of training only on the *SmartKC*-data, which achieved an accuracy of 71.9% (S_e=65.2%, S_p=76.5%). Analogous to the *SmartKC*-data, we also had access to the heatmaps generated by Keratron for the same set of 64 patients. Our network achieved an accuracy of 93.9% (S_e=94.7%, S_p=93.4%) when evaluated on these Keratron heatmaps (without the stage-2 of fine-tuning). In comparison, the Keratron-generated PPK-value based classification achieved an accuracy of 92.2% (S_e=89.4%, S_p=94.2%). Additionally, we train an SVM on sim-K quantitative values, where the input vector comprises of *[sim-k1, sim-k2, (sim-k1 − sim-k2), (sim-k1 + sim-k2)/2]*. This classifier achieved a high specificity (100.0%) but performed poorly in terms of sensitivity (80.4%) (Table I). This further shows the superiority of the proposed network.

Our proposed network is successful in learning discriminative features for the detection of keratoconus, as visible from the t-sne visualizations [25] (Figure 3) of the learned representations from the second fully-connected layer (FC2 in Figure 2) before and after training it for the axial heatmaps. (Note: the tangential heatmaps show a similar separation).

Furthermore, we analyze the effect of data augmentations and fine-tuning, by conducting an ablation analyses. Since the traditional transfer learning model performed poorly (S_e=65.2%, S_p=76.5%) due to the small size of *SmartKC*-data, for the ablation experiments, we consider the model after stage-1 fine-tuning on Keratron-data as our base model.

Data Augmentations Ablation: Table II demonstrates the utility of adding domain-specific data augmentations. Without adding any augmentations, the network performed poorly with an accuracy of 42.1% (S_e=88.8%, S_p=10.5%). Adding *mixup* or domain-specific (rotation, translation, scaling) augmentations increased the accuracy to 59.9% and 56.9%, respectively. On combining both types of augmentations, the accuracy further improves to 74.7% (S_e=89.1%, S_p=65.0%).

Fine-tuning	Augmentations	Se	Sp	Acc.
Stage-1	None	88.8%	10.5%	42.1%
Stage-1	Mixup	84.7%	43.1%	59.9%
Stage-1	Dom. Spec.	82.1%	39.9%	56.9%
Stage-1	Dom. Spec. + Mixup	89.1%	65.0%	74.7%
Stage-1,2	None	82.1%	85.6%	84.1%
Stage-1,2	Dom. Spec. + Mixup	91.3%	94.2%	93.1%

TABLE II

Ablations for fine-tuning, and domain-specific and mixup data augmentations. Evaluation is performed on the SmartKC test split..
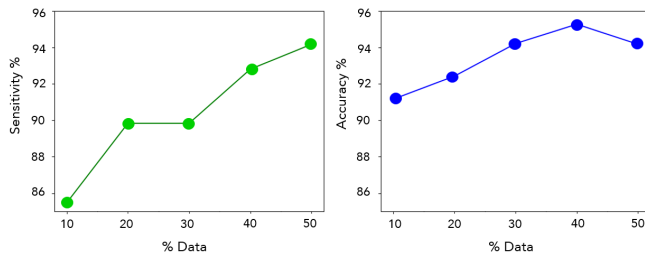

Fig. 4. *Effect of % data used for stage-2 fine-tuning on the SmartKC-data.*

Fine-tuning Ablation: Despite the improvements in performance after adding data augmentations, the performance of the network is not satisfactory. Fine-tuning on a subset of *SmartKC*-data (stage-2) drastically improved performance, with the network achieving an accuracy of 84.1% (S_e=82.1%, S_p=85.6%) (Table II) even without any data augmentations. When combined with data augmentations, the network accuracy improved to 93.1% (S_e=91.3%, S_p=94.2%). The amount of *SmartKC*-data used for fine-tuning impacts performance (Figure 4). Increasing the data for stage-2 fine-tuning from 10% to 50% improves accuracy by 3.1%, thus hinting that more data can further improve the performance of our network. We found the sensitivity to increase uniformly with more data, however the accuracy decreased slightly on increasing stage-2 fine-tuning data from 40% to 50% (Figure 4). This can be attributed to our usage of a weighted loss function, which penalizes the network more for misclassifying keratoconus samples, and also owing to the fact that the number of non-keratoconus samples (68) is much larger than keratoconus samples (46).

V. CONCLUSION

In this work, we demonstrated the efficacy of a deep neural network to detect keratoconus using the curvature heatmaps generated by *SmartKC*, a low-cost smartphone-based corneal topographer. We propose a dual-head CNN (one each for the axial heatmap and the tangential heatmap), and develop a 2-stage transfer learning strategy to train our network on a small-sized dataset (114 samples). We achieve results that are at par with doctors [2], indicating that our proposed system can be used for diagnosis of keratoconus in practice. Although this work focuses on classifying corneal topography heatmaps, the proposed 2-stage transfer learning strategy can be applied generally to similar scarce data settings. In future, to improve the performance of the network and increase robustness of our system, we plan to collect more data for training the network. Finally, we aim to deploy the *SmartKC* system with the proposed novel CNN-based keratoconus classifier to aid in mass screening of keratoconus

in schools, villages, and remote locations. This can help in early diagnosis of keratoconus and timely treatments, thus playing a crucial role in preventing needless blindness.

REFERENCES

[1] Jost B. Jonas et al. "Prevalence and Associations of Keratoconus in Rural Maharashtra in Central India: The Central India Eye and Medical Study". In: *American Journal of Ophthalmology* 148.5 (Nov. 2009), pp. 760–765.

[2] Siddhartha Gairola et al. "SmartKC: Smartphone-Based Corneal Topographer for Keratoconus Detection". In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5.4 (Dec. 2021).

[3] *World Health Organization, The Global Health Repository, Medical Doctors*. 2019. URL: https://bit.ly/3KRfTnT.

[4] Kazutaka Kamiya et al. "Keratoconus detection using deep learning of colour-coded maps with anterior segment optical coherence tomography: A diagnostic accuracy study". In: *BMJ Open* 9 (2019).

[5] Xu Chen et al. "Keratoconus detection of changes using deep learning of colour-coded maps". In: *BMJ Open Ophthal* 6.1 (2021).

[6] Pierre Zéboulon et al. "Corneal Topography Raw Data Classification Using a Convolutional Neural Network". In: *American Journal of Ophthalmology* 219 (2020), pp. 33–39. ISSN: 0002-9394.

[7] Rohit Shetty et al. "Artificial Intelligence Efficiently Identifies Regional Differences in the Progression of Tomographic Parameters of Keratoconic Corneas". In: *Journal of Refractive Surgery* 37.4 (2021).

[8] Bo-I Kuo et al. "Keratoconus Screening Based on Deep Learning Approach of Corneal Topography". In: *Translational Vision Science Technology* 9.2 (Sept. 2020), pp. 53–53. ISSN: 2164-2591.

[9] Ayse Yildiz Tas et al. "Automated Diagnosis of Keratoconus from Corneal Topography". In: *Investigative Ophthalmology & Visual Science* 62.8 (June 2021), pp. 2021–2021. ISSN: 1552-5783.

[10] Eric S. Hwang et al. "Distinguishing Highly Asymmetric Keratoconus Eyes Using Combined Scheimpflug and Spectral-Domain OCT Analysis." In: *Ophthalmology* 125 12 (2018), pp. 1862–1871.

[11] Ashraf M. Mahmoud et al. "CLMI: The Cone Location and Magnitude Index". In: *Cornea* 27 (2008), pp. 480–487.

[12] Philippe Chastang et al. "Automated keratoconus detection using the EyeSys videokeratoscope." In: *Journal of cataract and refractive surgery* 26 5 (2000), pp. 675–83.

[13] Alain Saad and Damien Gatinel. "Combining Placido and Corneal Wavefront Data for the Detection of Forme Fruste Keratoconus". In: *Journal of Refractive Surgery* 32.8 (2016), pp. 510–516.

[14] Illés Kovács et al. "Accuracy of machine learning classifiers using bilateral data from a Scheimpflug camera for identifying eyes with preclinical signs of keratoconus". In: *Journal of Cataract and Refractive Surgery* 42 (2016), pp. 275–283.

[15] Pierre Zéboulon, Guillaume Debellemanière, and Damien Gatinel. "Unsupervised learning for large-scale corneal topography clustering". In: *Scientific Reports* 10 (Oct. 2020).

[16] Alexandru Lavric and Popa Valentin. "KeratoDetect: Keratoconus Detection Algorithm Using Convolutional Neural Networks". In: *Computational Intelligence and Neuroscience* (Jan. 2019), pp. 1–9.

[17] Manuel Campos Garcıa et al. "Design of a null-screen for a compact corneal topographer". In: *Modeling Aspects in Optical Metrology VII*. SPIE, June 2019.

[18] Behnam Askarian et al. "An affordable and easy-to-use diagnostic method for keratoconus detection using a smartphone". In: *Medical Imaging 2018: Computer-Aided Diagnosis*. Ed. by Kensaku Mori and Nicholas Petrick. SPIE, Feb. 2018.

[19] F. I. Pinheiro et al. "Design And Development Of An Ultraportable Corneal Topographer For Smartphones As A Low Cost New Tool For Preventing Blindness Caused By Keratoconus". In: *International Journal of Latest Research in Science and Technology* 4 (3 2015).

[20] B. Askarian et al. "Novel Keratoconus Detection Method Using Smartphone". In: *2019 IEEE Healthcare Innovations and Point of Care Technologies, (HI-POCT)*. 2019, pp. 60–62.

[21] LVPEI. *Students Screened for Keratoconus*. 2018. URL: https://bit.ly/3E3zeiZ.

[22] André Beling da Rosa. "An accessible approach for corneal topography". In: 2013.

[23] Olga Russakovsky et al. "ImageNet Large Scale Visual Recognition Challenge". In: *International Journal of Computer Vision (IJCV)* 115.3 (2015), pp. 211–252.

[24] Hongyi Zhang et al. "mixup: Beyond Empirical Risk Minimization". In: *ArXiv* abs/1710.09412 (2018).

[25] Laurens van der Maaten and Geoffrey Hinton. "Visualizing Data using t-SNE". In: *Journal of Machine Learning Research* 9.86 (2008), pp. 2579–2605.