A Generative Model Reveals the Influence of Patient Attributes on Fundus Images

Sarah Müller[1] SAR.MUELLER@UNI-TUEBINGEN.DE
Lisa Koch[1] LISA.KOCH@UNI-TUEBINGEN.DE
Hendrik P. A. Lensch[1] HENDRIK.LENSCH@UNI-TUEBINGEN.DE
Philipp Berens[1] PHILIPP.BERENS@UNI-TUEBINGEN.DE
[1] *University of Tuebingen*

Editors: Under Review for MIDL 2022

Abstract

Screening for ophthalmic diseases routinely relies on retinal fundus images. These images are highly heterogeneous and little is known about how patient attributes such as age and ethnicity contribute to the variability in appearance. As the image variation due to such factors may ultimately confound automated image interpretation using deep learning models, understanding the influence of patient attributes on retinal fundus images is key for reliable AI applications in ophthalmology. Here, we draw on recent advances in generative modeling and present a population model of retinal fundus images which is capable of generating highly realistic images and allows for an analysis of how the patient attributes age and ethnicity are organized in the latent space of the generative model.

Keywords: generative models, fundus images, latent space, patient attributes

1. Introduction

Retinal fundus images are routinely used to screen for major ophthalmic diseases such as Diabetic Retinopathy (DR). For automatic grading of DR, deep learning techniques have been successfully used overall (Sengupta et al., 2020), and even implemented and approved in commercially available medical devices. While there exists a vast body of work on automated analysis of retinal fundus images, many studies rely on publicly available datasets without clinical information beyond the disease label. However, fundus images are highly heterogeneous with a wide range of appearances (see Fig. 1 a) and little is known about how patient attributes such as age and ethnicity contribute to this variation. Interestingly, such attributes are predictable from retinal images (Poplin et al., 2018), indicating a systematic relationship. As the image variation due to such factors may ultimately confound automated image interpretation using deep learning models, understanding the influence of patient attributes on retinal fundus images is key for reliable AI applications in ophthalmology. For example, in the case of ethnicity, not taking potential variations into account could even lead to the reproduction of racial disparities in medical practice.

Here, we use a generative model to (1) generate realistic fundus images and (2) investigate the relationship between patient attributes and image appearance. To this end, we analyze how the patient attributes age and ethnicity are organized in the latent space of the generative model. While the latent space of GANs has been used to study morphological characteristics of tissue in histopathology images (Quiros et al., 2021), to our knowledge this is the first contribution towards a population model of fundus images that takes into account relevant patient attributes.

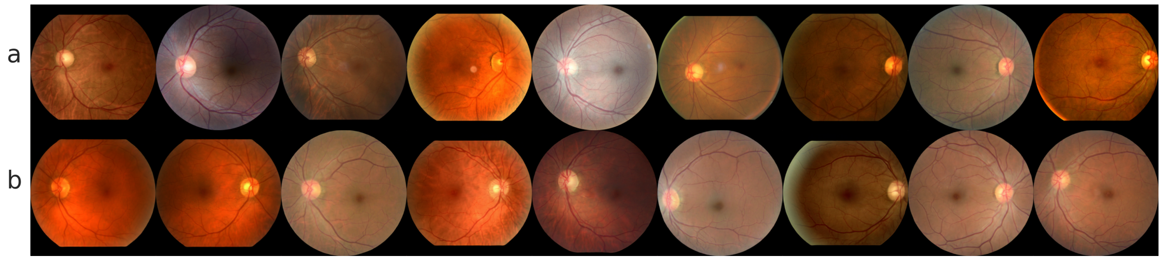

Figure 1: Examples of (a) real and (b) model generated fundus images.

2. Methods

We used 92,691 macula-centered fundus images provided by EyePACS, Inc, which were graded as good image quality. Images were scale normalized and resized to 256×256 pixels. We trained a StyleGAN2-ADA (Karras et al., 2020) which stabilizes GAN training for limited available data with an adaptive discriminator augmentation mechanism. The architecture resembles a classical GAN with a generator and a discriminator, but instead of generating images from a random latent vector, the latent vector is first mapped to an 512-dimensional intermediate latent space. This intermediate latent space can disentangle high-level attributes and stochastic variation. As the GAN approach does not allow direct encoding of real images to the intermediate latent space, we approximated each image's representation by minimizing the distance between the feature activations of the target and the generated image via backpropagation (Karras et al., 2020).

3. Results

Our model generated realistic fundus images (Fig. 1) with optic disc and macula and different fundus colors. Even the fine vessel structures appeared in a wide variety, making the general appearance similar to the original data distribution. However, upon close examination, the generated vasculature was not always intact (e.g. last example in Fig. 1b). We visualized the latent space structure using a t-SNE projection of the intermediate latent space (Fig. 2). We found that fundus images clustered in latent space by the patients' ethnicity (Fig. 2 a), with ethnicity affecting the retinal pigmentation. We also found the latent space to be structured with regards to the patients' age (Fig. 2 b). From young to old patients, the fundus images showed less bright regions and spots, but more reddish features. We quantitatively assessed the degree of structure in the latent space by applying a k-NN classifier to classify ethnicity based on the t-SNE representation and comparing it to its baseline performance on randomly permuted labels. We obtained an accuracy of 0.66 (baseline 0.59 ± 0.001) for $k = 19$. The evaluation for age with a k-NN regression led to a mean absolute error of 7.5 (baseline 9.06 ± 0.02) years for $k = 13$.

4. Discussion and Conclusion

With this work, we showed how a normative, generative model of retinal fundus images can be used to study the influence of varying patient attributes. Specifically, we demonstrated

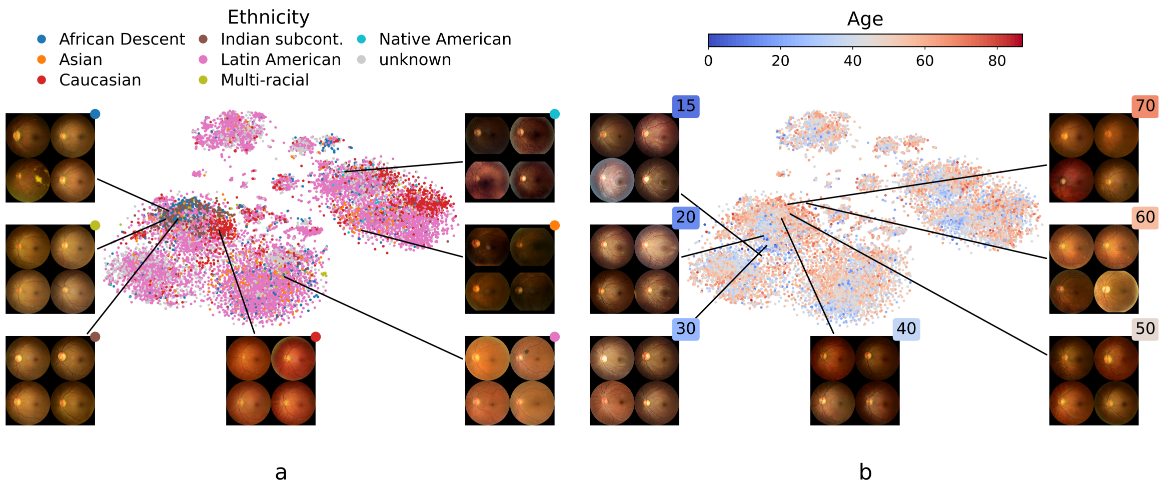

Figure 2: T-SNE visualization of latent space with two patient attributes highlighted. Each arrow is targeting to the mean positions of the images in latent space.

that patient ethnicity and age lead to appearance changes in fundus images. In future work, we plan to address disentangling other attributes in the latent space, e.g., through contrastive approaches or a conditional GAN.

Acknowledgments

We thank Hanna Faber, Valentyn Boreiko and Julius Gervelmeyer for discussion. We thank the German Science Foundation (BE5601/4-2 and EXC 2064, project number 390727645) and the German Ministry of Science and Education (BMBF, 01IS18039A) for funding.

References

Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proc. NeurIPS*, 2020.

Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, Yun Liu, Michael V. McConnell, Greg S. Corrado, Lily Peng, and Dale R. Webster. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering*, 2(3): 158–164, 2018.

Adalberto Claudio Quiros, Roderick Murray-Smith, and Ke Yuan. Pathologygan: Learning deep representations of cancer tissue. *Machine Learning for Biomedical Imaging (MELBA)*, 2021. MIDL 2020 Special Issue.

Sourya Sengupta, Amitojdeep Singh, Henry A. Leopold, Tanmay Gulati, and Vasudevan Lakshminarayanan. Ophthalmic diagnosis using deep learning with fundus images – a critical review. *Artificial Intelligence in Medicine*, 102:101758, 2020.