Fake it 'til You Make it: Does Generative Artificial Intelligence Have a Place in Ophthalmology?

First submitted: 16 October 2022
Revision accepted: 23 April 2023
Revision submitted: 19 March 2023
Published: Summer 2023

Peer reviewed

Ivan Potapenko, MD, PhD
Dept. of Ophthalmology, Rigshospitalet, Copenhagen, Denmark

Morten la Cour, Prof., dr. med.,
Dept. of Ophthalmology, Rigshospitalet, Copenhagen, Denmark

Abstract

Artificial intelligence (AI) is currently one of the hottest topics in ophthalmic research. AI has been used with great success for an assortment of imaging-based tasks, such as screening, diagnosis, and staging of ophthalmic diseases. There are, however, other useful ways of employing AI. Instead of simply classifying an image, the so-called generative AI algorithms are able to do the reverse—generate new images based on input. With this approach, it is possible, for example, to predict retinal appearance after treatment, enhance images, and convert between imaging modalities. This article aims to summarize the most promising of such generative AI algorithms in ophthalmology.

Introduction

Artificial intelligence (AI) is broadly defined as the ability of a computer to perform tasks that usually require human intelligence. The most characteristic trait of an AI algorithm (also called a model) is that it starts as a blank slate with no better performance than a random guess. Gradually, it obtains its functionality by being shown many real-world examples of the problem it is designed to solve,[1] in a process called training. Typically, AI models are used to classify images[2] and delineate borders between imaged structures.[3] For these so-called discriminative tasks, AI easily outperforms any manually constructed algorithm, and in some cases even surpasses the accuracy of humans.[4,5] Discriminative AI is already approaching real-world deployment, with the first AI algorithms being approved for screening and clinical decision-making support.[6,7]

Other types of AI include generative models, which are typically designed to output images instead of only evaluating them. They originated in an attempt to understand the internal workings of discriminative algorithms in the late 2000s,[8] and during the following decade were developed further to produce fake (synthetic) images on their own. Some examples of synthetically generated images are shown below: a somewhat

Key concepts:

- *Artificial intelligence (AI)*: ability of a computer program to perform tasks that usually require human intelligence (e.g., recognition of objects, composing text, and interpreting speech)
- *AI model*: an algorithm designed for a specific task (e.g., diagnosis of glaucoma on fundus photos)
- *AI training*: the process where an AI model learns to perform a specific task by being shown multiple correctly performed examples
- *Discriminative AI*: type of AI algorithms that classify data (e.g., diabetic retinopathy grading) or parts of images (e.g., delineating retinal layers and edema on OCT)
- *Generative AI*: type of AI algorithms that generate data, for example, images, resembling those it was shown during training
- *Synthetic image*: an image produced by a generative AI model, that is, arrived at by computation alone and not resulting from an imaging device (camera, scanner etc.)

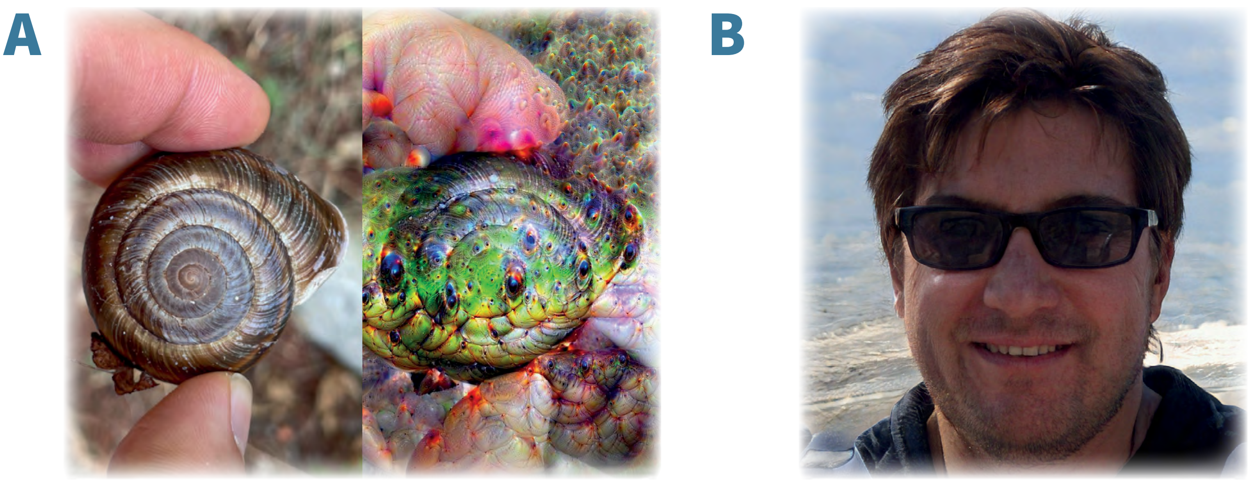

Figure 1. Examples of AI-generated images. (a) A photo processed by Google's DeepDream generator.[9] Original image on the left, and the processed image on the right (private photo). (b) Synthetic photograph of a person that does not exist.[33]

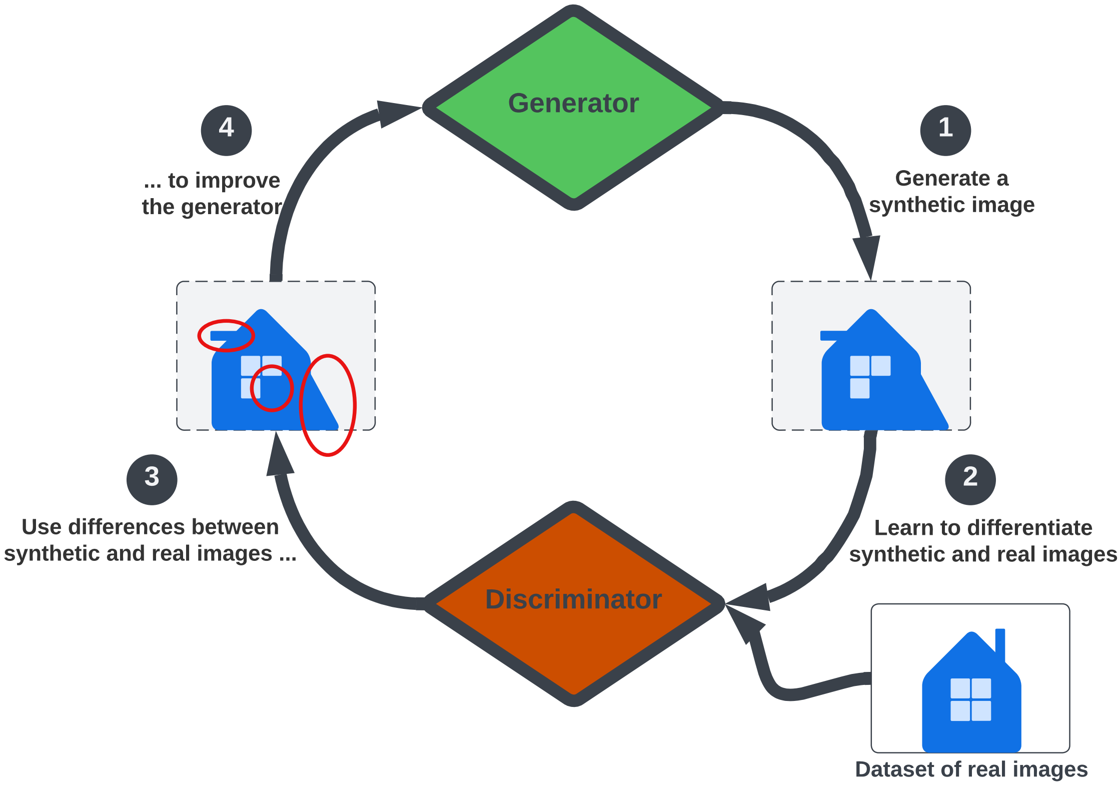

Figure 2. The training process of a generative AI model. The generative capabilities of an AI algorithm are trained by making two models (a generator and a discriminator) compete against each other. While the generator attempts to produce output that is as realistic as possible, the discriminator attempts to find features that distinguish real images from generated ones. In the example above, a generative model that outputs images of houses is trained. The process starts by the generator (green rhomboid shape) outputting an image (1). The discriminator (red rhomboid) then attempts to distinguish it from a dataset of images of real houses (2). Features that were used to single out the generated image (3) are then used as feedback to improve the generator (4). The process is repeated until the discriminator cannot discern between real and generated samples.

nightmarish version of a photo in which unrelated patterns are accentuated by Google's DeepDream[9] (**Figure 1a**), and the extremely photorealistic portrait of a person that does not exist[10,11] (**Figure 1b**). Although these examples might be of limited broader relevance, the usefulness of generative models has been proved for a wide variety of text and imaging tasks.

In contrast, scant attention has been devoted to generative AI in medical research. There is, however, growing recognition that there are useful ways to apply generative approaches in this context. Below, we examine the potential applications of generative models in ophthalmology.

Competition is the key to success

The core concept of training a generative AI model is competition between two models:[12] a so-called generator, which attempts to output data that is as realistic as possible, and a discriminator, which tries to separate the generated samples from real ones. The generator is improved according to the feedback from the discriminator, making its output more realistic. This makes generated images more difficult to distinguish from real ones, forcing further refinement of the discriminator, which is then again used to improve the generator model, and so on. This process repeats until the discriminator can no longer discern real from generated samples, ensuring as realistic an output as possible from the generator. The process is schematically represented in **Figure 2**.

With this basic approach, a number of more complex models with extended functionality can be created, several of which are referenced below. From a technical standpoint, there is nothing limiting generative AI to purely image-based data. This is clearly demonstrated by the recently released ChatGPT model by OpenAI,[13] which has an almost human-like ability to generate cohesive text. Although this model might be used to answer patient questions for several ophthalmic diseases,[14] most generative AI models in ophthalmology are currently image based. The scope of the discussion below is therefore limited to addressing these.

Generating new images

A generative AI algorithm trained on an imaging dataset can produce new images that resemble the original data (**Figure 3**). Generated images are termed synthetic as they do not result from imaging real patients but are rather created using patterns the AI model has learned from the original dataset. Despite this, synthetic images can be convincingly realistic (such as the fake person photo in **Figure 1b**), to a point where experts cannot distinguish between real and artificial data.[15] The process of generating new images has the additional benefit of being highly controllable. Synthetic OCT scans can, for example, be generated with a specific amount of subretinal fluid or a

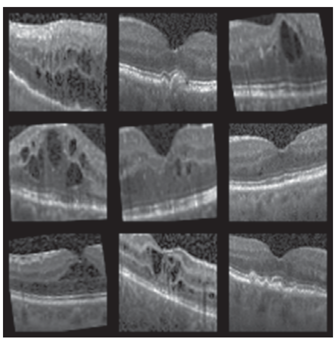

Figure 3. Synthetic OCT images. These realistic-looking OCT scans were not taken from real patients but rather generated by an AI model (own data based on a public OCT dataset, unpublished).

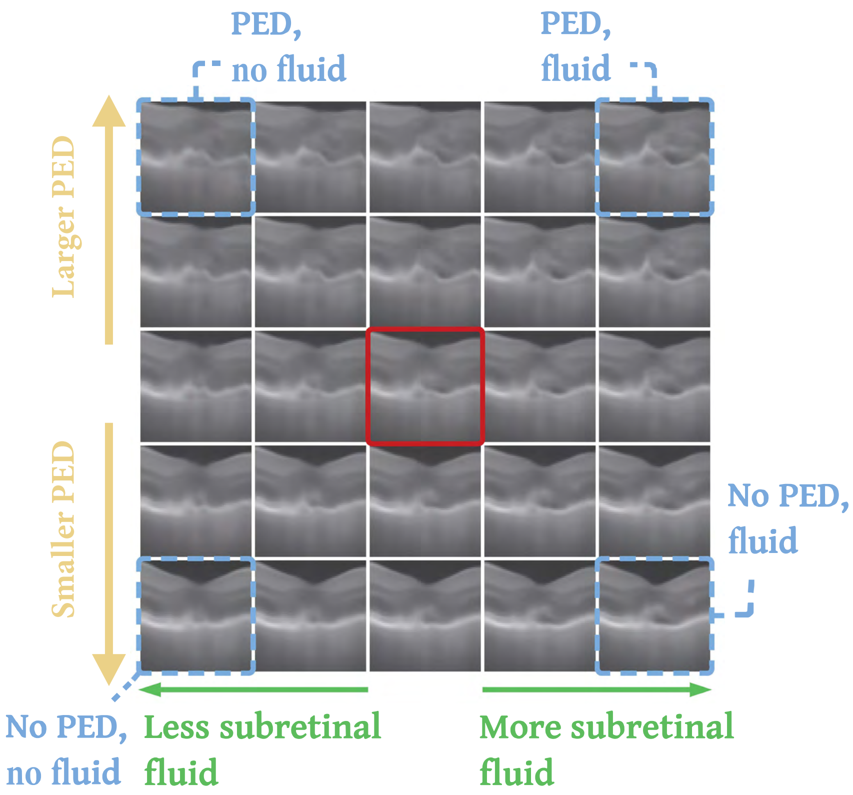

Figure 4. Controlling features of synthetic OCT images. This figure demonstrates that generative AI can produce images with specific features on demand. In the above grid, first, a synthetic image is generated (middle, framed in red). Subsequent images are generated by varying two parameters: the amount of subretinal fluid (x-axis) and size of pigment epithelium detachment (PED, y-axis).

determined size of pigment epithelium detachment (**Figure 4**).

There is no limit to how many new synthetic images a trained generative model can produce. It can thus be a viable option for expansion of existing datasets without recruitment of new patients. This, in turn, can be useful for training discriminative AI models, which require large amounts of data, or providing additional examples for training of clinical personnel—especially when specific disease stages can be generated on demand.

The fact that neither the model itself nor the images it generates contain the training dataset can be exploited for data sharing. Moving patient images across jurisdictions can be a difficult and bureaucratic process. Transferring an AI model, on the other hand, does not reveal any sensitive or identifiable personal information and is thus not subject to regulations.

Predicting the future

Generative models can be trained to predict future appearances based on present imaging results. This has been shown to work well for prediction of visual field progression in glaucoma patients[16,17] or postoperative facial appearance after surgery for thyroid-associated orbitopathy.[18] Another important use is forecasting macular OCT changes after anti-VEGF injections in neovascular AMD.[19] This might be used to create a more precise treatment plan tailored to each patient. Since this is done before any treatment is administered, over-treatment and frequent follow-ups can be avoided.[20,21]

Image enhancement

AI can improve existing images in various ways. For example, a poor-quality OCT image (e.g., due to cataract or vitreous opacities) can be enhanced by noise reduction, producing a clearer depiction of anatomical structures.[22] A synthetic image with higher resolution and better detail can be produced using a smaller and less detailed version.[23] When applied to visual fields, AI denoising can make discovering progression easier.[17,24] Artifacts can also be removed, for example, the crescent shadow on fundus photographs caused by small pupil size,[25] eliminating the need to re-take photos and thus saving time.

Converting between imaging types

One of the most recent developments in generative AI is the ability to convert between imaging modalities. By analyzing one type of imaging, AI models can predict appearance of the same eye when imaged by another modality.

A practical example is widefield imaging with a scanning laser ophthalmoscope (e.g., Optos®). The resulting image is not true color but is rather composed of red and green wavelengths, giving it a greenish tint. This is often cited as being one of the drawbacks of these devices.[26] By using generative AI, a real color image of the central retina can be produced, based solely on an Optos image.[27]

Even more intriguing is the ability to translate between structural and functional imaging—an OCT scan can be converted to an approximation of a visual field,[28] and both OCTs and fundus photos can be used to generate synthetic fluorescein angiograms.[29,30]

A future scenario

All these algorithms are useful on their own, but their combined application may further enhance and simplify the clinical workflow. The following hypothetical case is an example of clinical implementation of the above-mentioned algorithms.

An elderly patient with dry AMD and a suspicion of primary open-angle glaucoma presents with diffuse symptoms of vision deterioration in one eye. Being an elderly woman, she cannot cope with more than a single imaging session, and her pupils dilate poorly. Therefore, only Optos macula OCT and a widefield fundus image are obtained.

With the AI system, artifacts from the eyelashes and small pupil are removed. The widefield fundus image allows for evaluation of the peripheral retina to rule out retinal tears or detachment. A separate image of the optic nerve with greater resolution

Key points:

- Generative artificial intelligence (AI) is a powerful technique that enables generation of realistic images in response to an input.
- Generative AI models can be trained to:
  - improve quality and remove artifacts from images
  - forecast response to treatment
  - infer appearance of different imaging modalities from a single image
  - transfer large datasets safely.
- The potential and pitfalls of generative models are not yet fully explored.

and better detail is produced to evaluate cupping. The visual field is estimated from the OCT, and noise is removed for better comparison with previous examinations. In addition, a visual field prediction 6 months into the future is generated. To evaluate the status of the patient's AMD, the central fundus is reproduced in true color from the Optos image, along with a prediction of how a fluorescein angiogram would appear. As the presence of fluid is noted with the AI system, a prediction of OCT appearance after anti-VEGF treatment is created.

All of this information is available after a 30-second imaging session and before the patient is examined by a physician!

Challenges ahead

As enticing as the scenario above might be, many challenges remain. Accuracy of synthetic data, at least for the time being, cannot be guaranteed as it is not possible to completely rule out erroneous or missing features. This might be a barrier to CE marking and clinical implementation in cases where the output is directly used for diagnosis—for example, for conversion between modalities as described above.

Training generative AI is considered a far greater technical challenge compared with traditional discriminative models,[31] requiring many thousands to millions of samples.[32] Few medical imaging datasets of this size currently exist, potentially limiting the quality of generated images.

Conclusion

Generative models are a fascinating branch of AI research, with several extremely useful applications, such as predicting treatment response, enhancing images, and obtaining more information out of fewer imaging modalities. As with any cutting-edge technology, however, their application in a clinical setting warrants a degree of caution as their adverse effects remain largely unexplored.

However, it is certain that AI comprises an array of powerful tools that are yet to be utilized to their fullest potential. Combined with lowered barriers to entry, nearly limitless computing resources, big data, and increasing healthcare demands, many breakthroughs are surely on the horizon.

References

1. Brown TB, Mann B, Ryder N, et al. Language Models are Few-Shot Learners. *Adv Neural Inf Process Syst*. 2020;2020-December. doi:10.48550/arxiv.2005.14165
2. Potapenko I, Kristensen M, Thiesson B, et al. Detection of oedema on optical coherence tomography images using deep learning model trained on noisy clinical data. *Acta Ophthalmol*. Published online 2021. doi:10.1111/aos.14895
3. Sorkhabi MA, Potapenko IO, Ilginis T, Alberti M, Cabrerizo J. Assessment of Anterior Uveitis Through Anterior-Segment Optical Coherence Tomography and Artificial Intelligence-Based Image Analyses. *Transl Vis Sci Technol*. 2022;11(4). doi:10.1167/TVST.11.4.7
4. Burlina P, Joshi N, Pacheco KD, Freund DE, Kong J, Bressler NM. Utility of Deep Learning Methods for Referability Classification of Age-Related Macular Degeneration. *JAMA Ophthalmol*. 2018;136(11):1305-1307. doi:10.1001/jamaophthalmol.2018.3799
5. Phene S, Dunn RC, Hammel N, et al. Deep Learning and Glaucoma Specialists: The Relative Importance of Optic Disc Features to Predict Glaucoma Referral in Fundus Photographs. *Ophthalmology*. 2019;126(12):1627-1639. doi:10.1016/j.ophtha.2019.07.024
6. Ipp E, Liljenquist D, Bode B, et al. Pivotal Evaluation of an Artificial Intelligence System for Autonomous Detection of Referrable and Vision-Threatening Diabetic Retinopathy. *JAMA Netw open*. 2021;4(11):e2134254. doi:10.1001/JAMANETWORKOPEN.2021.34254
7. Abràmoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *npj Digit Med* 2018 11. 2018;1(1):1-8. doi:10.1038/s41746-018-0040-6
8. Erhan D, Bengio Y, Courville A, Vincent P. Visualizing Higher-Layer Features of a Deep Network Département d'Informatique et Recherche Opérationnelle.; 2009.
9. DeepDream Generator. Accessed October 15, 2022. https://deepdreamgenerator.com/
10. Karras T, Laine S, Aila T. A Style-Based Generator Architecture for Generative Adversarial Networks. *IEEE Trans Pattern Anal Mach Intell*. 2018;43(12):4217-4228. doi:10.48550/arxiv.1812.04948
11. Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T. Analyzing and Improving the Image Quality of StyleGAN. *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit*. Published online December 3, 2019:8107-8116. doi:10.48550/arxiv.1912.04958
12. Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative Adversarial Networks. *Commun ACM*. 2014;63(11):139-144. doi:10.1145/3422622
13. Introducing ChatGPT. Accessed March 19, 2023. https://openai.com/blog/chatgpt
14. Potapenko I, Boberg-Ans LC, Stormly Hansen M, Klefter ON, van Dijk EHC, Subhi Y. Artificial intelligence-based chatbot patient information on common retinal diseases using ChatGPT. *Acta Ophthalmol*. Published online March 13, 2023. doi:10.1111/AOS.15661
15. Burlina PM, Joshi N, Pacheco KD, Liu TYA, Bressler NM. Assessment of Deep Generative Models for High-Resolution Synthetic Retinal Image Generation of Age-Related Macular Degeneration. *JAMA Ophthalmol*. 2019;137(3):258. doi:10.1001/jamaophthalmol.2018.6156
16. Asaoka R, Murata H, Asano S, et al. The usefulness of the Deep Learning method of variational autoencoder to reduce measurement noise in glaucomatous visual fields. *Sci Rep.* 2020;10(1). doi:10.1038/S41598-020-64869-6
17. Berchuck SI, Mukherjee S, Medeiros FA. Estimating Rates of Progression and Predicting Future Visual Fields in Glaucoma Using a Deep Variational Autoencoder. *Sci Rep.* 2019;9(1):18113. doi:10.1038/s41598-019-54653-6
18. Yoo TK, Choi JY, Kim HK. A generative adversarial network approach to predicting postoperative appearance after orbital decompression surgery for thyroid eye disease. *Comput Biol Med.* 2020;118:103628. doi:10.1016/j.compbiomed.2020.103628
19. Liu Y, Yang J, Zhou Y, et al. Prediction of OCT images of short-term response to anti-VEGF treatment for neovascular age-related macular degeneration using generative adversarial network. *Br J Ophthalmol.* 2020;104(12):1735-1740. doi:10.1136/BJOPHTHALMOL-2019-315338
20. Parvin P, Zola M, Dirani A, Ambresin A, Mantel I. Two-year outcome of an observe-and-plan regimen for neovascular age-related macular degeneration treated with Aflibercept. *Graefe's Arch Clin Exp Ophthalmol*. 2017;255(11):2127-2134. doi:10.1007/s00417-017-3762-2
21. Mehta H, Kim LN, Mathis T, et al. Trends in Real-World Neovascular AMD Treatment Outcomes in the UK. *Clin Ophthalmol.* 2020;14:3331-3342. doi:10.2147/OPTH.S275977
22. Cheong H, Devalla SK, Pham TH, et al. DeshadowGAN: A Deep Learning Approach to Remove Shadows from Optical Coherence Tomography Images. *Transl Vis Sci Technol*. 2020;9(2):1-15. doi:10.1167/TVST.9.2.23
23. Ha A, Sun S, Kim YK, et al. Deep-learning-based enhanced optic-disc photography. *PLoS One*. 2020;15(10):e0239913. doi:10.1371/JOURNAL.PONE.0239913
24. Asaoka R, Murata H, Asano S, et al. The usefulness of the Deep Learning method of variational autoencoder to reduce measurement noise in glaucomatous visual fields. *Sci Rep.* 2020;10(1):7893. doi:10.1038/s41598-020-64869-6
25. Yoo TK, Choi JY, Kim HK. CycleGAN-based deep learning technique for artifact reduction in fundus photography. *Graefe's Arch Clin Exp Ophthalmol*. Published online May 2, 2020. doi:10.1007/s00417-020-04709-5
26. Kumar V, Surve A, Kumawat D, et al. Ultra-wide field retinal imaging: A wider clinical perspective. *Indian J Ophthalmol.* 2021;69(4):824. doi:10.4103/IJO.IJO_1403_20
27. Yoo TK, Ryu IH, Kim JK, et al. Deep learning can generate traditional retinal fundus photographs using ultra-widefield images via generative adversarial networks. *Comput Methods Programs Biomed.* 2020;197. doi:10.1016/J.CMPB.2020.105761
28. Hashimoto Y, Kiwaki T, Sugiura H, et al. Predicting 10-2 Visual Field From Optical Coherence Tomography in Glaucoma Using Deep Learning Corrected With 24-2/30-2 Visual Field. *Transl Vis Sci Technol*. 2021;10(13):28-28. doi:10.1167/TVST.10.13.28
29. Lee CS, Tyring AJ, Wu Y, et al. Generating retinal flow maps from structural optical coherence tomography with artificial intelligence. *Sci Rep.* 2019;9(1):5694. doi:10.1038/s41598-019-42042-y
30. Tavakkoli A, Kamran SA, Hossain KF, Zuckerbrod SL. A novel deep learning conditional generative adversarial network for producing angiography images from retinal fundus photographs. *Sci Rep.* 2020;10(1). doi:10.1038/S41598-020-78696-2
31. Yazici Y, Foo CS, Winkler S, Yap KH, Chandrasekhar V. Empirical Analysis of Overfitting and Mode Drop in GAN Training. *Proc - Int Conf Image Process ICIP.* 2020;2020-October:1651-1655. doi:10.48550/arxiv.2006.14265
32. Karras T, Aittala M, Hellsten J, Laine S, Lehtinen J, Aila T. Training Generative Adversarial Networks with Limited Data. *Adv Neural Inf Process Syst*. 2020;2020-December. doi:10.48550/arxiv.2006.06676
33. This Person Does Not Exist. Accessed September 18, 2022. https://this-person-does-not-exist.com/en

Conflict of interest
Nothing to disclose