

## Invited Commentary

# Is Artificial Intelligence for Retinopathy of Prematurity Ready to Go?

Gil Binenbaum, MD, MSCE

**DeepMind scientists** began with old-school Atari games.<sup>1</sup> Using deep and reinforcement learning, they developed an artificial intelligence (AI) program that learned to play Breakout, Pong, and other games at super-human levels after being pro-



Related article

<https://www.youtube.com/watch?v=EfGD2qveGdQ>.<sup>1,2</sup> They caught the eyes of Google executives, who promptly bought the company and convinced them to tackle the 3000-year-old game of Go. Go was an elusive goal for AI researchers, who had primarily used game-tree-search approaches that could not handle Go's complexity. The number of possible board configurations in this deceptively simple-ruled game is estimated to be  $10^{170}$ , about  $10^{90}$  times greater than the number of atoms in the known universe. Using machine learning similar to the Atari project, DeepMind developed AlphaGo, which learned first from human-expert games and then self-play.<sup>3</sup> In 2015, AlphaGo became the first AI to defeat a human Go master.<sup>1</sup> In 2016, AlphaGo defeated 8-time world champion Lee Sedol in an AI landmark event that inspired a movie.<sup>4</sup>

In this issue of *JAMA Ophthalmology*, Coyner et al<sup>5</sup> report an AI milestone accomplishment of their own: the first validated autonomous AI algorithm for retinopathy of prematurity (ROP) screening. They developed and validated a deep learning model, i-ROP-DL, that assesses plus-like posterior-pole vascular changes in retinal photographs to determine whether an infant requires an ophthalmologist examination. In development, they applied robust AI methodology to 2530 images from 843 infants in the US. In validation, i-ROP-DL flagged in advance 100% of infants requiring treatment among 2244 infants from existing ROP telemedicine programs at 59 hospitals in the US and India. The broader context is important. While there are still details to study about treatment timing and modality, we generally know how to treat severe ROP to acutely reduce risk of retinal detachment. Globally, most blindness due to ROP results from inadequate screening infrastructure and too few ophthalmologists who manage ROP, particularly in low- and middle-income countries. Telemedicine can help address resource limitations, but what if there are not even enough ophthalmologists to read images? The e-ROP study showed trained nonphysicians could do the job,<sup>6</sup> but given steady technological progress, it is likely that AI algorithms eventually will do it better than humans. This report is a fundamental step in that direction.

In the 37th move of its second game against Lee Sodon, AlphaGo made what experts at first thought was a mistake, until they realized it was a completely novel strategy. One Go master commented, "It's not a human move. I've never seen a human play this move. So beautiful. Beautiful.... Beautiful...."<sup>4</sup> As the technology advanced, AlphaGo evolved into AlphaZero, which learned Go better from scratch by play-

ing itself repeatedly, without first studying human games. The AI became simpler, stronger, and more generalizable. After only a few hours of training, AlphaZero was able to defeat all the best Go programs, and it taught itself chess just as quickly.<sup>7</sup>

As with Go, maybe AI will teach us something new about ROP and evolve to interpret images better than humans. At this point, however, how much of the screening process should we entrust to AI? Consider some current limitations. The i-ROP-DL algorithm does not determine stage, just plus-like change, though, the two are closely related. It also does not assess zone or signal mature vasculature. It has not yet been trained to identify a retinal detachment, retinoblastoma, or other unexpected finding that a human would recognize. More generally, telemedicine requires available cameras and photographers, image quality and resolution must be high, and from the infant's point of view, getting photographs is still a repeated intervention. Finally, there is a profound issue related to the way AI represents things, because it turns out that deep learning models do not understand things deeply.

Go continued to be dominated by AI, which only got faster and stronger, until early 2023, when researchers led by Stuart Russell at University of California, Berkeley, surprised the world when their amateur Go player defeated all the leading Go programs using a telling "circular group sandwich" strategy.<sup>8</sup> They reasoned that the algorithms could detect patterns but had limited conceptual representations of what was happening; the AI did not understand, for example, the basic strategic concept of a group of pieces, and this weakness could be exploited.<sup>8</sup> Without a conceptual understanding, deep learning models, which express things using giant, circuit-type designs, are vulnerable to small aberrations, which can result in mistakes or failures, and because we really do not understand what the circuits do, the mistakes are unpredictable.<sup>8</sup> This characteristic has been seen with image recognition, where, for example, an AI algorithm changes an initially correct answer of "panda" to a highly confident but incorrect answer of "gibbon" after the introduction of minor "noise" pixels into the image.<sup>8</sup> What if noise enters a retinal photograph, changing the diagnosis? i-ROP-DL does not understand there are blood vessels, or a retina, or an infant that could go blind. And we fundamentally cannot predict what or when mistakes might happen. Using even immense amounts of training data cannot address this inherent limitation.<sup>8</sup>

So is ROP AI ready to go? Humans are still the ROP masters, and until AI methodology moves beyond data-driven deep learning design, a fundamental vulnerability to make unpredictable mistakes when presented with new circumstances will persist. Therefore, the safest approach is ideally for humans to remain as involved as possible in each step of ROP management. However, where ophthalmologist supply is limited but retinal camera availability is not, iROP-DL could be used to ration examinations to babies who need them most. It will help

to continue updating iROP-DL, add training to identify detachments, retinoblastomas, and other unexpected findings, and incorporate human checks wherever possible. Ultimately, however, as AI keeps advancing, it will eventually

surpass us in diagnostic and prognostic abilities, justifiably be used more broadly, and perhaps teach us something new about ROP. I myself will welcome our new AI overlords with open arms. Whatever is best for the babies.

#### ARTICLE INFORMATION

**Author Affiliation:** Division of Ophthalmology, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania.

**Corresponding Author:** Gil Binenbaum, MD, MSCE, Division of Ophthalmology, Children's Hospital of Philadelphia, 3500 Civic Center Blvd, 11th Floor, Philadelphia, PA 19104 ([binenbaum@chop.edu](mailto:binenbaum@chop.edu)).

**Published Online:** March 7, 2024.  
doi:10.1001/jamaophthalmol.2024.0218

**Conflict of Interest Disclosures:** Dr Binenbaum is a co-investigator on National Institutes of Health grant R21EY034179 that relates to AI and ROP work.

#### REFERENCES

1. Minh V, Kavukcuoglu K, Silver D, et al. Playing Atari with deep reinforcement learning. *arXiv*.
2. Deepmind artificial intelligence @FDOT14. YouTube. Accessed February 3, 2024. <https://www.youtube.com/watch?v=EfGD2qveGdQ>
3. Silver D, Huang A, Maddison CJ, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*. 2016;529(7587):484-489. doi:10.1038/nature16961
4. AlphaGo. IMDb. Accessed January 16, 2024. <https://www.imdb.com/title/tt6700846/>
5. Coyner AS, Murickan T, Oh MA, et al. Multinational external validation of autonomous retinopathy of prematurity screening. *JAMA Ophthalmol*. Published online March 7, 2024. doi:10.1001/jamaophthalmol.2024.0045
6. Quinn GE, Ying GS, Daniel E, et al; e-ROP Cooperative Group. Validity of a telemedicine system for the evaluation of acute-phase retinopathy of prematurity. *JAMA Ophthalmol*. 2014;132(10):1178-1184. doi:10.1001/jamaophthalmol.2014.1604
7. Silver D, Hubert T, Schrittwieser J, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*. 2018;362(6419):1140-1144. doi:10.1126/science.aar6404
8. Marcus G. David beats Go-liath. Substack. 2023. Accessed January 15, 2024. [https://garymarcus.substack.com/p/david-beats-go-liath?utm\\_source=%2Fsearch%2F david%2520beats&utm\\_medium=reader2](https://garymarcus.substack.com/p/david-beats-go-liath?utm_source=%2Fsearch%2F david%2520beats&utm_medium=reader2)