

Vision 2.0:



Leif Hynnekleiv, MD^{1,4}




Fredrik Andreas Fineide, MD^{1,2,4}



Professor Hugo Lewi Hammer, PhD^{1,2,5}


1. Oslo University Hospital, 2. Oslo Metropolitan University, 3. King's College London, 4. The Norwegian Dry Eye Clinic, 5. SimulaMet, 6. St Thomas' Hospital, London, 7. Department of Ophthalmology, Sørlandet Hospital Trust, Arendal

When Lina, an 85-year-old retiree, registers at the eye clinic for her annual eye check-up, she is met with a surprise. Instead of sitting down in a waiting room, she is asked to stare into a machine for a few seconds. In that time, the machine uses an artificial intelligence (AI)-powered feedback loop to adjust to her eyes' position and movement; acquire multiple, perfectly sharp scans of the anterior and posterior segments; and then preprocess and filter the images and biometrics. Using this and her existing digital health profile, the AI-powered software quickly analyzes all her data. Within minutes, Lina's ophthalmologist is reviewing a comprehensive updated profile of Lina's eye status and function, coupled with a differential diagnosis, risk analysis, and treatment suggestions, including clinical advice and summaries of the most relevant recent research. Lina has a narrowing anterior chamber and a subtle but gradual increase of intraocular pressure, and she will develop sight-threatening posterior cortical cataracts within six months. Subtle signs in the retina and blood vessels suggest a ninety percent likelihood of neovascularization in her left eye before her next visit. After a short talk with her doctor, Lina leaves with new prescriptions and scheduled appointments for robotic cataract surgery, prophylactic anti-VEGF injections, and a follow-up scan.



Lina being introduced to new AI-powered instruments by an ophthalmologist in the not-so-distant future. Illustration by Kristin Skårdal.

The impact of modern artificial intelligence on ophthalmology



Professor Chris Hammond, MD, PhD^{3, 6, 7}



Professor Michael A. Riegler, PhD^{2, 5}

Where AI and ophthalmology collide. Image generated by DALL·E 3 and edited by Kristin Skårdal.

This imagined scenario is not an excerpt from a science fiction novel but rather the imminent status quo of AI-powered modern eye care. The technology described in this scenario already exists in some form, and is slowly moving into the clinic, a situation nearly unimaginable less than a century ago. So how did we get here?

The term *artificial intelligence*, meaning the simulation of human intelligence by machines, was coined and established as an academic discipline in the mid-1950s.¹ A foundational workshop in Dartmouth in 1956 sparked early enthusiasm. This initial drive lost its momentum, in part due to contemporary technology acting as a hindrance. However, throughout AI's history, we have seen continued research efforts into its use in various fields, including medicine and diagnostics.

Among the earliest ventures into AI-assisted diagnostics was CASNET (Causal ASSociational NETworks) for glaucoma, developed in the 1970s at Rutgers University and Mount Sinai Medical School.² The computer program simulated clinical cognition based on expertise from a network

of glaucoma specialists. Using medical history as input, its built-in “understanding” of pathophysiology used a causal network model and associations to present advice and treatment suggestions for a differential diagnosis and provided alternative opinions on contested subjects. Though primarily intended as a research tool, it showed promise when evaluated against human experts.

The field of AI—computer systems capable of performing tasks that typically require human intelligence, such as problem-solving, decision-making, or understanding language—evolved at the intersection of mathematics and computer science. New discoveries and developments were largely driven by academic research.³ With the exponential increase in computational power and improvements in models and algorithms during recent decades, AI has become part of the backbone of our digital daily lives. The integration and dissemination of AI into almost every aspect of modern life has mostly been happening below the public radar, such as the use of predictive

machine learning, image classification, and recommendation systems, integrated into our phones, computers, and TVs.⁴ However, the release of OpenAI's large language model ChatGPT opened the public's eyes overnight. Since then, media coverage has mainly focused on the subcategory of generative AI, such as large language models and image generation models. However, generative AI is just the tip of the AI iceberg, which has been merging for some time with other fields such as medicine and, notably, ophthalmology.

The eye is an ideal and directly accessible organ for obtaining high-quality and safe imaging using a plethora of modalities that ophthalmologists rely on in their daily practice, such as photography, optical coherence tomography (OCT), autofluorescence, and angiography. The eye's structure, appearance, function, biomechanics, pressure, and optics are all readily available for measurement, assessment, and intervention—a unique property compared to internal organs. Retinal imaging even allows for sampling the state of the central nervous system and its

microvasculature,⁵ and the segmentation and measurement of all the retina's layers enable detailed inspection at the microscopic scale. Easy access to data and images from this beautifully complex organ makes ophthalmology a natural choice for AI integration. This article introduces and explores the history, organization, and integration of AI in ophthalmology, giving examples of current applications, challenges, and future directions.

A Brief AI Overview

Machine learning, and particularly deep learning, is the foundation of what we associate with modern AI. A machine-learning algorithm is essentially a piece of mathematics designed to learn from data and produce conclusions. Such algorithms can be relatively simple or extremely complex and can be singular or several models functioning together. Machine-learning models are not explicitly built based on rules but *trained* to learn the patterns in the data. Several modes of training exist, depending on the data available and the end goal of the model. One way of organizing different models into discrete categories is by how the models learn. Check out the AI crash course to learn more (**Box 1**).

This approach to categorizing machine-learning models is an approximation and does not tell the full story. Many other forms of learning exist, such as data-efficient- and interactive-learning, and combined models acting synergistically, including generative adversarial networks. Audio, video, images, text, and numbers can all be used as input to train an AI model. Similarly, an AI can be trained to provide any type of output.

Observation, Hypothesis, Investigation, Evidence

The concept of testing observational hypotheses by clinical investigation to arrive at empirical evidence is typically accredited to the Hungarian physician

Ignaz Semmelweis. Semmelweis refused to believe the popular theory that childbed fever, a common and deadly postpartum infection, was an “inevitability of the harmful influences of atmospheric conditions.”¹¹ In 1847, after rigorous research using cutting-edge empirical methods, Semmelweis famously introduced handwashing with chlorinated lime solutions to doctors attending childbirth. This drastically reduced the incidence of childbed fever and maternal mortality.¹¹

domain, despite a lack of evidence or even in the face of strong counter-evidence. This is exemplified by Andrew Wakefield's infamous retracted 1998 study in *The Lancet* linking the MMR vaccine to autism.¹³

The risk of tunnel vision is high when a hypothesis is developed from personal experience and anecdotal observation—and even higher when the ideas capture the public imagination. Unfortunately, a “theory first, data later” approach, where a theoretical framework is proposed before extensive empirical data can support or refute it, has plagued medicine since its inception.¹⁴

The traditional confirmatory model of research was challenged by the introduction of exploratory data analysis by John Tukey in 1977.¹⁵

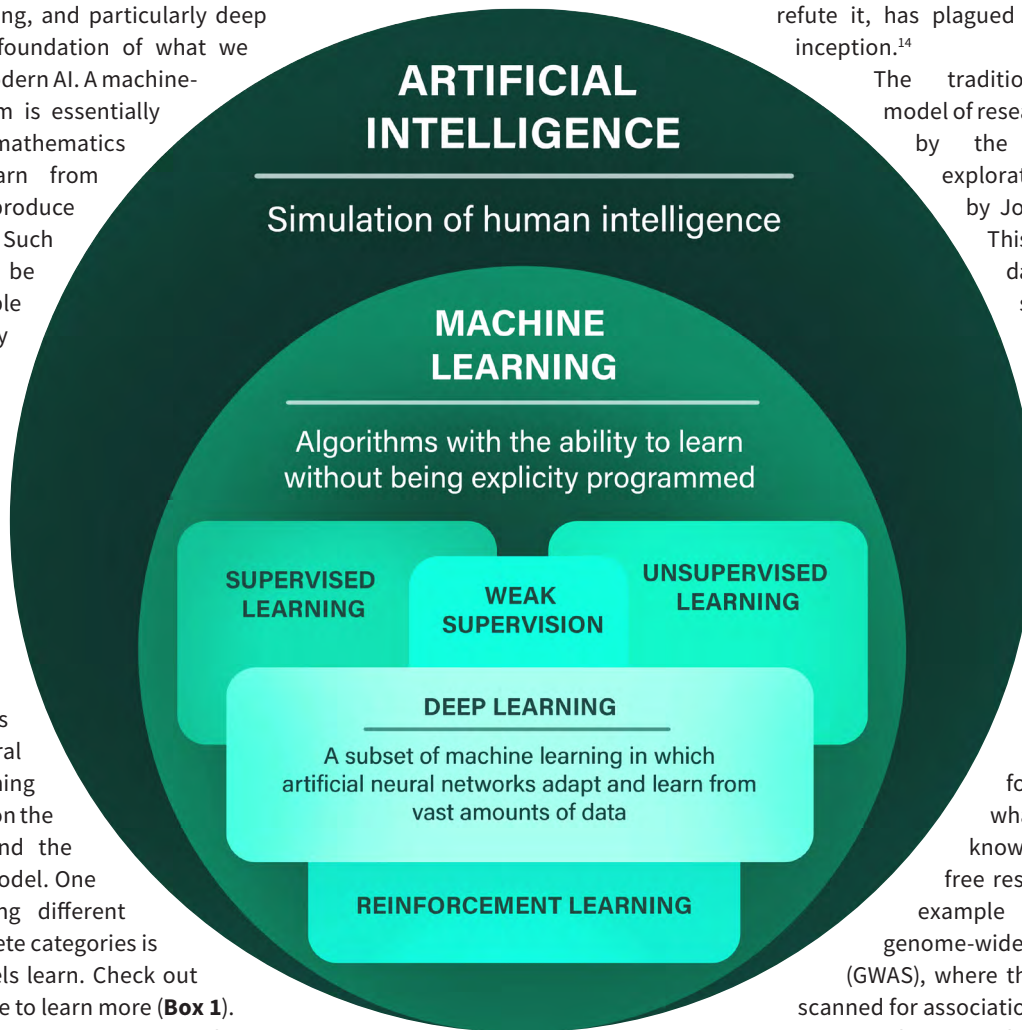
This method allows the data to reveal their structure without imposing strict assumptions.

This created a way of developing research questions based on patterns revealed from data exploration rather than from preconceived

ideas, laying the foundation for what has come to be known as hypothesis-free research.

A prominent example in genetics is genome-wide association study (GWAS), where the entire genome is scanned for associations between genetic variants and traits or diseases. In 2005, one of the first landmark studies to validate the GWAS method identified associations between single-nucleotide polymorphisms and age-related macular degeneration.¹⁶

Machine-learning methods can identify patterns and characteristics in data by analyzing large complex data sets without expectations or assumptions. Hypothesis-free research methods can be improved upon with techniques such as clustering and dimensionality reduction to allow biomarker discovery using high-dimensionality data sets such as spectrometry and microarrays.¹⁷ New hypothesis-free methods exploit the flexibility of machine learning, such as mass analysis of -omics data and deep learning



Organization of the field of artificial intelligence. Deep learning is part of machine learning, which is part of artificial intelligence. All are based on algorithms. Illustration by Kristin Skårdal.

Semmelweis spent most of his life after his discoveries trying to convince his colleagues of his groundbreaking findings. Semmelweis summarized his life's work by saying, “My doctrines exist to rid maternity hospitals of their horror, to preserve the wife for her husband and the mother for her child.”¹¹

The paradigm shift into the systematic use of empirical evidence in clinical practice was finally formalized with the establishment of evidence-based medicine in the early 1990s.¹² However, widely-publicized, unfounded opinions persist in the public

Box 1. A Crash Course in AI Basics

Supervised Learning: Supervised models are trained on a data set in which the training examples have known outcomes. Outcomes are the desired and pre-defined outputs of the model. There are two main types of supervised learning models, classification and regression. When a supervised learning model makes a “guess” at what the real outcome is, it is called a prediction. In this context a prediction is *not* about predicting the future, but about guessing what the real value is.

The outcome assigned to a training example is called a label. A training set of retinal images can be labeled “geographic atrophy” or “healthy,” and the model can be trained to categorize them into these two pre-defined classes. Such “recognition,” or prediction, of a defined outcome is called classification. With enough examples, the supervised learning model learns to classify images with and without geographic atrophy. We provide it with the correct answers so it can compare its predictions against the true values and *learn*.

All classification models assign labels. Some also provide a confidence level, for example, that with 80% confidence the image label should be “geographic atrophy.” Regression models, on the other hand, are trained to predict a continuous value, like the *amount* of geographic atrophy, for example, that 55% of the macula is affected by geographic atrophy.

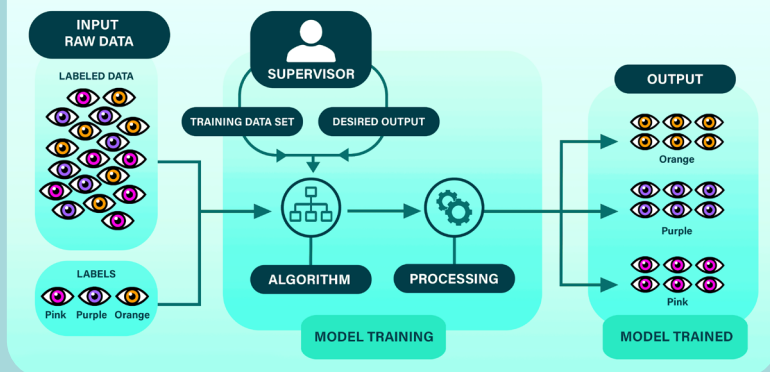
Supervised learning is widely used in applications ranging from email filtering and speech recognition to disease prediction. As a recent example, Dai et al. trained a model using a labeled retinal image database to predict the time to clinical progression of diabetic retinopathy.⁶

Unsupervised Learning: This method involves training models on data without labels. The goal is to uncover relationships or structures within the data. This approach is particularly useful for tasks where the underlying patterns in the data are not known, where AI can help explore or understand the data, such as by organizing it into categories. Applications include user behavior analysis in social media and techniques to simplify bioinformatics data without losing core information, such as principal component analysis. One recent study using unsupervised learning trained a model on 28,000 longitudinal macular OCT images. It identified seven different glaucomatous progression patterns, correlating them with visual field total deviations.⁷

Reinforcement Learning: This type of AI model is characterized by learning from interaction with an environment. Decisions lead to immediate feedback that slowly teaches the model how to behave. In contrast to supervised learning where the model learns from a labeled data set, reinforcement-learning algorithms are looking to maximize their reward for performing correct actions or reaching a defined goal, reminiscent of the biological dopaminergic reward system. This form of AI is historically rooted in behavioral psychology and the concept of operant conditioning, introduced in the 1930s.⁸ Instead of statically evaluating inputs, reinforcement-learning models learn to react to a changing environment by making and evaluating decisions. While training in a simulated dynamic environment, a reinforcement-learning system can make millions of decisions in different situations, evaluate the consequences of those decisions, and reward the best ones, thereby gradually improving the model. The model slowly learns to make better decisions that help achieve its goal while considering the state of the environment. Gaming, robotics, and finance have been among the main frontiers in developing modern reinforcement-learning systems. Within ophthalmology, reinforcement learning has potential in areas such as imaging, intelligent surgery systems, and chronic disease monitoring and treatment.⁹

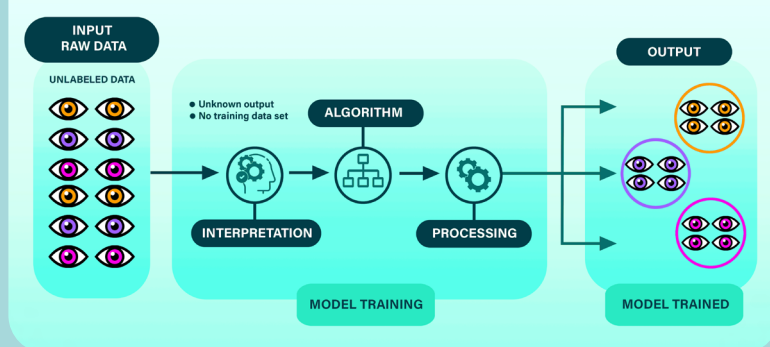
Weak Supervision: Many newer machine-learning models fall between the supervised and unsupervised learning categories. Such models can use a mix of labeled and unlabeled data, use imperfect or noisy data, or learn to assign labels to the data themselves. They are essential for the function of many generative AI models. Weak-supervision models can learn from a wide variety of data, even when most of the data has not been properly organized, and they excel with large and poorly labeled data sets, making them great at tasks such as describing images and translating language. Within ophthalmology, such models have great potential in improving image analysis, such as training a model to identify specific lesions, hemorrhages, or exudates while leveraging unlabeled images to improve generalization and accuracy. A recent study published in *JAMA* described such a model for improved OCT detection of macular telangiectasia type 2.¹⁰

SUPERVISED LEARNING



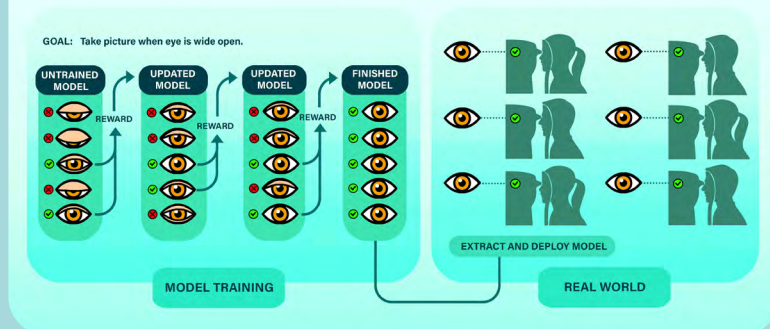
Supervised learning model using labeled images of eyes to learn how to predict eye color. Illustration by Kristin Skårdal.

UNSUPERVISED LEARNING



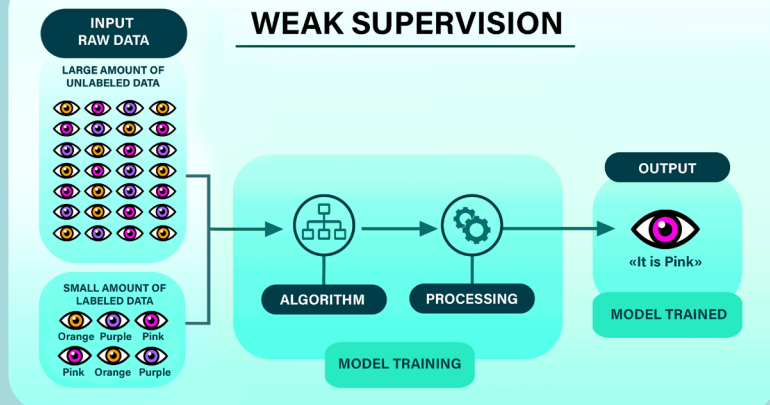
Unsupervised learning algorithm sorting unlabeled images of eyes into newly created categories. It has chosen to separate the images based on eye color. Illustration by Kristin Skårdal.

REINFORCEMENT LEARNING



Reinforcement-learning algorithm training to take pictures at the right moment when the eye is wide open. The best choices are rewarded, and the model is updated. The finished model is extracted and deployed to AI cameras in the real world. Illustration by Kristin Skårdal.

WEAK SUPERVISION



A weak-supervision model learning to identify eye color based on a small amount of labeled data and a large amount of unlabeled data. Illustration by Kristin Skårdal.



Fictional illustration of Semmelweis demonstrating handwashing with chlorinated lime solution. Image generated by DALL-E 3 and edited by Kristin Skårda

for genomic, image, and clinical data sets combined with a subfield of AI known as explainable AI. Machine learning is being used to extract information from multiple -omics data sets simultaneously, such as metabolomics, genomics, and proteomics, to gain more meaningful insights in so-called multi-omics data analysis.¹⁸ The insights gained increase our understanding of the body and disease processes and can be used to help identify biomarkers or as justification for further research and clinical trials. By analyzing large, high-quality data sets without an initial hypothesis, machine learning is increasingly being used to produce results less subject to human biases.

The Really Cool AI Stuff

Artificial neural networks work by sending data through artificial neurons organized in layers, loosely inspired by their biological counterparts. A neural network neuron is a simple decision-making unit that processes information and sends out a signal based on the strength of the inputs. That signal becomes the input to the next layer of neurons. The architecture of a neural network typically consists of an input layer that passes data to the network, one or several hidden layers that process the data, and an output layer presenting the result. Artificial neural networks with more than one hidden layer are referred to as deep neural networks and are used in deep learning.¹⁹ Such models are often termed “black boxes” due to the difficulty in understanding how the model arrived at its answer. Typical deep-learning tasks include text and image recognition. However, what sets deep learning apart from other machine-learning methods is its versatility since it can be configured to solve nearly any problem given enough high-quality training data.²⁰

Explainable AI

When using machine-learning models in research or implementing models in a clinical setting, the researcher or clinician often must understand how the model produced its results. Most machine-learning models do not automatically tell the user which parts of the data were important. For example, with a large, detailed data set from glaucoma patients, you might train a model to predict the risk of rapid glaucoma progression to use for individual patients

in the clinic. But what if you wanted to know exactly why some patients are assigned a higher or lower risk? Is it mainly due to aging, pressure fluctuations, gradual pressure increase, or certain nerve fiber loss patterns? The physician may feel the need to verify how the model prioritizes when making predictions before trusting it or feeling confident about introducing it into clinical practice.²¹ A more transparent or informative model can also provide the physician with valuable feedback and create learning opportunities. This is especially important if the conclusion reached by the clinician differs from that of the model or when investigating the root causes of disease.

Useful information can be extracted by poking and prodding machine-learning algorithms, giving rise to the field of research known as explainable AI. Methods and tools provided by explainable AI help provide model understanding and lead to new insights and ideas. Several approaches to explainable AI are being explored to increase our understanding of AI models.^{22,23} A popular explainable AI method is the local interpretable model-agnostic explanations, which can imitate a black box model such as a deep neural network, with simpler interpretable models. Another method is the Shapley additive explanations, which can tell us which features in the data are the most important.

Current State of AI in Clinical Ophthalmology

Harnessing the power of AI in diagnostics, management, and treatment is about not replacing human capability but improving decision-making, efficiency, and cost-effectiveness. It is a tool to arrive faster at more precise conclusions. In other words, AI is helping us improve patient outcomes and increase turnaround.

The early detection of disease in a clinical setting depends on keen observation and pattern recognition. This comes with the risk of oversight, especially among clinicians with less experience. Even among experts, opinions and methods can vary. Unlike humans, machines are not subject to fatigue, and, with thoughtful design, measurement variability, such as intra-observer variability, can be minimized or eliminated. AI systems have been shown to reduce such variability in clinical practice.^{24,25}

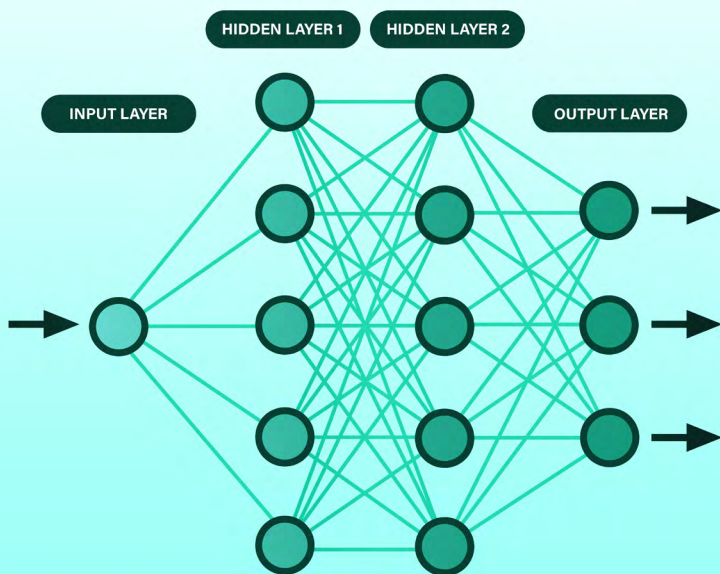
The accuracy of AI models in detecting conditions such as diabetic retinopathy, glaucoma, and age-related macular degeneration based on retinal imaging reached clinical viability years ago. An example is the FDA-approved LumineticsCore that autonomously diagnoses DR and detects macular edema.²⁶ Google’s DeepMind AI for Retinal Diseases has proven, in collaboration with Moorfields Eye Hospital in the UK, that it can identify a wide range of retinal diseases.²⁷ Ting et al. reported that AI detected referable age-related macular degeneration with 93.20% sensitivity.²⁸ Additionally, the i-ROP DL system received FDA breakthrough device status²⁹ for diagnosing retinopathy of prematurity with comparable or better accuracy than examination by expertly trained ophthalmologists.³⁰

Community-based screenings, especially for diabetic retinopathy, have always faced challenges regarding adequate resource availability. AI systems can assist in time-efficient mass screening with minimal resources. Abràmoff et al. developed an AI image system that showed an impressive 96.8% sensitivity and 87.0% specificity in detecting referable diabetic retinopathy in 2016.³¹ Products such as LumineticsCore and EyeArt have earned FDA approval, validating AI’s proficiency in diagnosing more-than-mild and vision-threatening diabetic retinopathy.³²

Glaucoma, the silent blinding disease, is seeing improved early detection with algorithms trained for discerning glaucomatous nerve fiber layer damage, early visual field loss, and more.³³⁻³⁵

Intraocular lens power calculation using machine learning has

DEEP NEURAL NETWORK



Example architecture of a deep neural network with a single input, two hidden layers with five nodes each, and three output layers. The choice of architecture depends on the data and the task, with variations in the number of inputs, outputs, layers, and nodes per layer. Illustration by Kristin Skårdal.

performed better in highly myopic eyes than newer-generation formulas such as Barrett and Olsen.³⁶ Additionally, refractive error prediction based on ocular imaging using AI is showing great promise with high accuracy.³⁷ Moreover, machine learning may be used in the future to predict post-surgical outcomes based on corneal topography and wavefront aberrometry analysis.

Expanding Capabilities


Presently, most AI models are trained for one or a few specific tasks. This implies that for disease screening, one must gather data, annotate, train, test, and predict using a separate model for every disease—a time-consuming and computationally expensive process. To address this, “inverse supervised learning” was recently proposed.³⁸ Rather than training a model on images containing a given disorder, a model was trained on a large data set of healthy brain scans to recognize any aberration from normality *and* highlight the location of interest, for a wide range of pathologies.³⁸ Simply put, this model differentiates normal from abnormal images and shows the location of suspected disease. The same model was fine-tuned and tested on OCT retinal scans, achieving impressive results and showing that this model is generalizable across other image modalities. This approach holds great promise as a clinical supplement, especially for rare disease states, for which gathering large data sets is notoriously difficult, and as a screening tool to quickly filter out patients in need of further diagnostics. Such creative advances also show that we still have much to discover and learn about the potential of AI in ophthalmology.

The annotation process of large image databases for model training requires significant time and expertise and is often a bottleneck in developing reliable image-based clinical AI models. This process is especially labor-intensive when drawing bounds within an image, known as segmentation. Proper segmentation involves not just a label, such as “exudate,” but also a precisely drawn box

around the exudate in the image. Complicated medical images may need segmentation of several different features in the same image. With the advent of segmentation tools such as Meta AI’s Segment Anything, however, this manual labor headache might soon be a thing of the past.³⁹ Pretrained on 11 million images, Segment Anything promises generalization to unfamiliar objects without additional training. The newest medical adaptations, SAM2 and MedSAM 2, have shown strong generalization capabilities. These include promising results for the automatic segmentation of various pathologies in retinal fundus images.⁴⁰

Trust and Ethical Considerations

With an increased reliance on AI, ethical considerations and technological gray areas are among the key hurdles to its future development and implementation in both research and clinical ophthalmology. An ongoing challenge when training and implementing AI models is preventing unwanted prediction variation or bias due to unrepresentative training data. A model needs enough data to produce reliable results, which is challenging to acquire from minority populations, especially for rare diseases or minority ethnic groups.⁴¹ Inclusivity and accessibility are equally important, meaning that future AI tools, whether for professionals or patients, should be built with a variety of end users in mind. This must include the range of physical, mental, and technological abilities found in the real world.⁴²



Explainable AI—opening the “black box” and studying an AI model to learn how it arrives at its conclusions. Illustration by Kristin Skårdal.

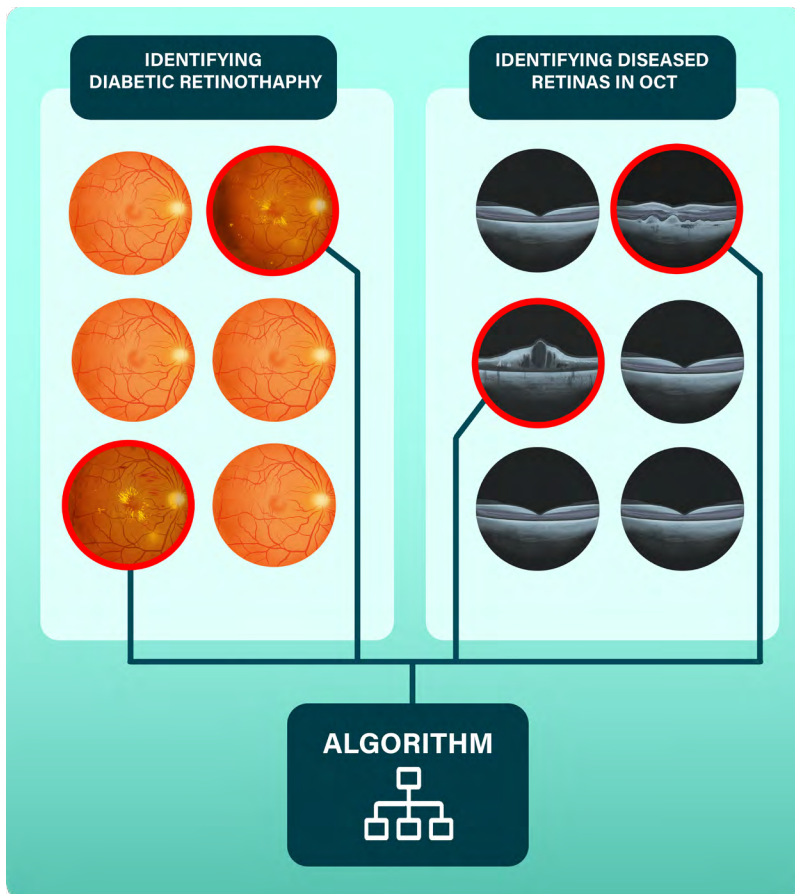


Illustration of a machine-learning algorithm identifying diabetic retinopathy in retinal images and wet maculopathy in OCT images. Illustration by Kristin Skårdal.

Doctors deal with complex, incomplete, and imprecise information daily. Assigning a correct diagnosis and treatment is rarely a clear-cut decision and is often based on objective findings combined with subjective information, such as visual symptoms, pain, and a patient’s memory of events. Building trust in medical AI systems that operate in an environment of imperfect information combined with high stakes, such as the consequences of a wrongful diagnosis, is no simple task. However, clinicians’ and patients’ trust in AI is necessary to allow AI systems into clinical practice. One way to achieve such trust is through interpretability and transparency of the models,⁴³ as well as slow but deliberate implementation while gathering feedback and making improvements.⁴⁴

Dealing with the vast amounts of data needed to train AI models requires staying within legal limits, such as the General Data Protection Regulation regarding data sharing and anonymization of sensitive data, while considering the ethical and human rights ramifications of AI adoption when designing AI models, as highlighted by the World Health Organization.⁴⁵

Regulatory and Profitability Concerns

In the pursuit of profit, some AI systems are developed as closed, proprietary


platforms, limiting transparency and the opportunity for external scrutiny. This raises concerns about the reproducibility of results. Although important steps are being made to improve generalizability, AI models trained on specific data sets may not perform equally well across

different imaging devices or diverse patient populations encountered in the real world.⁴⁶ The solutions to such concerns may include requiring independent validation studies and ensuring interpretable and interoperable AI systems through government regulation.

Existing regulatory systems may not be equipped to handle the novel challenge that AI poses. Regulatory bodies such as the FDA and EMA traditionally approve medical devices based on a fixed state at a specific time.^{47,48} However, the functionality and software of AI-powered devices benefit from real-world use and experience as new data and refined algorithms become available, allowing for incremental improvements after hitting the market. This dynamic nature conflicts with traditional regulatory frameworks that are not equipped to handle continuously learning systems.⁴⁹ In their mission to protect and advance public health, national and international regulatory bodies continue to establish and amend new frameworks regarding AI software and devices. These focus on issues such as transparency, software modifications, good machine-learning practice, and real-world performance monitoring.⁵⁰⁻⁵³ Reasonable regulations must be put in place to protect patient safety without stifling innovation.

Future Predictions for AI in Ophthalmology

The possibilities for AI in ophthalmology are nearly endless. Just a few examples of the multitude of ongoing research efforts



Metaphorical interpretation of human-machine interaction, highlighting the importance of diversity, inclusivity, and building trust. Illustration by Kristin Skårdal.

in this field are early disease detection and classification based on corneal topography and tomography,⁵⁴ predicting or detecting myopia progression and complications,⁵⁵ and the use of AI-assisted augmented reality in vitreoretinal surgery.⁵⁶

We predict an improvement of existing technologies, stronger integration with existing and new technologies, enhanced user experience, and the continued introduction of practical AI applications into nearly every ophthalmological subspecialty. We expect the quality and speed of imaging to improve with AI-enhanced image preprocessing and AI-powered automatic imaging systems, with better diagnostic accuracy, more consistent results, and greater efficiency in the clinic. We expect the gradual implementation of personalized treatment, especially for chronic eye

diseases such as wet age-related macular degeneration, using AI models trained on large data sets of patient information. These models will provide accurate personalized predictions of treatment response and help choose the right medication and treatment protocol, such as the type and frequency of anti-VEGF injections. For telemedicine and the remote monitoring of disease, we expect improved remote imaging solutions integrated with AI image analysis, improving referral systems and increasing the clinical data available during telemedicine consultations. We also expect the automation of routine but time-consuming tasks, such as grading of retinal images for diabetic retinopathy screening programs, to be slowly handed over to AI in the clinical setting. In the slightly more distant future, we expect the clinical implementation of AI-

assisted surgical robots.

Finally, we hope and expect that interacting with differently working systems will become a thing of the past as future clinical AI systems automatically and safely share data and integrate with one another, providing a seamless and effortless experience, as in the imagined story of Lina.

Conclusion

We look optimistically at the reshaping of ophthalmology through research and innovation, as well as the gradual introduction of AI technologies into the clinical setting. Technology optimism must be balanced with responsible and ethical implementation without losing sight of our goal: to improve vision and quality of life for countless individuals like Lina.

References

- McCarthy J, Minsky ML, Rochester N, Shannon CE. A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*. 2006;27(4):12.
- Weiss SM, Kulikowski CA, Amarel S, Safir A. A model-based method for computer-aided medical decision-making. *Artificial Intelligence*. 1978;11(1):145-172.
- Muthukrishnan N, Maleki F, Owens K, Reinhold C, Forghani B, Forghani R. Brief history of artificial intelligence. *Neuroimaging Clinics of North America*. 2020;30(4):393-399.
- Roser M. The brief history of artificial intelligence: the world has changed fast — what might be next? Our World in Data. Published 2022. Accessed October 28, 2024. <https://ourworldindata.org/brief-history-of-ai>
- Snyder PJ, Alber J, Alt C, et al. Retinal imaging in Alzheimer's and neurodegenerative diseases. *Alzheimers Dement*. 2021;17(1):103-111.
- Dai L, Sheng B, Chen T, et al. A deep learning system for predicting time to progression of diabetic retinopathy. *Nature medicine*. 2024;30(2):584-594.
- Kazeminasab Hashemabadi S, Eslami M, et al. An unsupervised deep learning method for identifying glaucoma progression patterns using longitudinal ganglion cell complex (GCC) scans. *IOVS*. 2024;65(7):1626-1626.
- Staddon JE, Cerutti DT. Operant conditioning. *Annu Rev Psychol*. 2003;54:115-144.
- Nath S, Korot E, Fu DJ, et al. Reinforcement learning in ophthalmology: potential applications and challenges to implementation. *The Lancet Digital Health*. 2022;4(9):e692-e697.
- Gholami S, Scheppke L, Kshirsagar M, et al. Self-Supervised Learning for Improved Optical Coherence Tomography Detection of Macular Telangiectasia Type 2. *JAMA ophthalmol*. 2024;142(3):226-233.
- Semmelweis I. *The Etiology, Concept, and Prophylaxis of Childbed Fever*. Madison, WI: The University of Wisconsin Press; 1983.
- Thoma A, Eaves FF, III. A Brief History of Evidence-Based Medicine (EBM) and the Contributions of Dr David Sackett. *Aesthetic Surgery Journal*. 2015;35(8):NP261-NP263.
- Godlee F, Smith J, Marcovitch H. Wakefield's article linking MMR vaccine and autism was fraudulent. *BMJ*. 2011;342.
- Cagnacci A, Venier M. The Controversial History of Hormone Replacement Therapy. *Medicina*. 2019;55(9):602.
- Tukey JW. *Exploratory data analysis*. Reading/Addison-Wesley. 1977.
- Klein RJ, Zeiss C, Chew EY, et al. Complement factor H polymorphism in age-related macular degeneration. *Science*. 2005;308(5720):385-389.
- Ledesma D, Symes S, Richards S. Advancements within Modern Machine Learning Methodology: Impacts and Prospects in Biomarker Discovery. *Curr Med Chem*. 2021;28(32):6512-6531.
- Reel PS, Reel S, Pearson E, Trucco E, Jefferson E. Using machine learning approaches for multi-omics data analysis: A review. *Biotechnology Advances*. 2021;49:107739.
- Ian Goodfellow YB, Aaron Courville. *Deep Learning*. MIT Press; 2016.
- Nielsen M. *Neural Networks and Deep Learning*. Vol 2024: Chapter 4: A Visual Proof That Neural Nets Can Compute Any Function. New York, NY: Determination Press; 2015
- Tan TF, Dai P, Zhang X, et al. Explainable artificial intelligence in ophthalmology. *Current opinion in ophthalmology*. 2023;34(5):422-430.
- Holzinger A. From machine learning to explainable AI. In: *IEEE 2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*. New York, NY: IEEE Computer Society; 2018:55-66. Published August 25, 2018.
- S Band S, Yarhamadi A, Hsu C-C, et al. Application of explainable artificial intelligence in medical health: A systematic review of interpretability methods. *Informatics in Medicine Unlocked*. 2023;40:101286.
- Tizhoosh HR, Diamandis P, Campbell CJV, et al. Searching Images for Consensus: Can AI Remove Observer Variability in Pathology? *The American Journal of Pathology*. 2021;191(10):1702-1708.
- Zhao K, Ma S, Sun Z, et al. Effect of AI-assisted software on inter- and intra-observer variability for the X-ray bone age assessment of preschool children. *BMC Pediatr*. 2022;22(1):644.
- LumineticsCore: Indications for use - US. Digital Diagnostics. April 10, 2024. Accessed October 28, 2024. <https://www.digitaldiagnostics.com/products/eye-disease/indications-for-use-us/>
- Moorfields Eye Hospital. Google DeepMind: A medical research partnership revolutionising the way professionals carry out eye tests, leading to earlier detection of common eye diseases. Moorfields Eye Hospital NHS Foundation Trust. Accessed October 28, 2024. <https://www.moorfields.nhs.uk/research/google-deepmind>
- Ting DSW, Cheung CY, Lim G, et al. Development and Validation of a Deep Learning System for Diabetic Retinopathy and Related Eye Diseases Using Retinal Images From Multiethnic Populations With Diabetes. *JAMA*. 2017;318(22):2211-2223.
- White F. Tech that detects cause of preemie blindness gets federal nod. Oregon Health & Science University. Published January 30, 2020. Accessed August 29, 2024. <https://news.ohsu.edu/2020/01/30/tech-that-detects-cause-of-preemie-blindness-gets-federal-nod>
- Taylor S, Brown JM, Gupta K, et al. Monitoring Disease Progression With a Quantitative Severity Scale for Retinopathy of Prematurity Using Deep Learning. *JAMA ophthalmology*. 2019;137(9):1022-1028.
- Abramoff MD, Lou Y, Erginay A, et al. Improved Automated Detection of Diabetic Retinopathy on a Publicly Available Dataset Through Integration of Deep Learning. *IOVS*. 2016;57(13):5200-5206.
- 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 Network Open*. 2021;4(11):e2134254-e2134254.
- Correia Barao R, Hemelings R, Abegao Pinto L, Pazos M, Stalmans I. Artificial intelligence for glaucoma: state of the art and future perspectives. *Current opinion in ophthalmology*. 2024;35(2):104-110.
- Huang X, Islam MR, Akter S, et al. Artificial intelligence in glaucoma: opportunities, challenges, and future directions. *Biomed Eng Online*. 2023;22(1):126.
- Zhang L, Tang L, Xia M, Cao G. The application of artificial intelligence in glaucoma diagnosis and prediction. *Front Cell Dev Biol*. 2023;11:1173094.
- Zhou Y, Dai M, Sun L, et al. The accuracy of intraocular lens power calculation formulas based on artificial intelligence in highly myopic eyes: a systematic review and network meta-analysis. *Front Public Health*. 2023;11:1279718.
- Yew SME, Chen Y, Goh JHL, et al. Ocular image-based deep learning for predicting refractive error: A systematic review. *Advances in Ophthalmology Practice and Research*. 2024;4(3):164-172.
- He Y, Guo Y, Liu J, et al. Disorder-Free Data Are All You Need—Inverse Supervised Learning for Broad-Spectrum Head Disorder Detection. *NEJM AI*. 2024;1(4):A10a2300137.
- Kirillov A, Mintun E, Ravi N, et al. Segment Anything. In: *Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV)*; 2023; Paris, France. IEEE: 2023:3992-4003.
- Kulyabin M, Zhdanov A, Pershin A, et al. Segment Anything in Optical Coherence Tomography: SAM 2 for Volumetric Segmentation of Retinal Biomarkers. *Bioengineering (Base)*. 2024;11(9).
- Huang J, Galal G, Etemadi M, Vaidyanathan M. Evaluation and Mitigation of Racial Bias in Clinical Machine Learning Models: Scoping Review. *JMIR Med Inform*. 2022;10(5):e36388.
- Stephanidis C. Digital Accessibility and Artificial Intelligence for an Inclusive Future. Paper presented at: 17th Session of the Conference of States Parties to the Convention on the Rights of Persons with Disabilities (CRPD); 2024/06/12; United Nations Headquarters, New York.
- Watson DS, Krutzinna J, Bruce IN, et al. Clinical applications of machine learning algorithms: beyond the black box. *BMJ*. 2019;364:l886.
- Lukyanenko R, Maass W, Storey VC. Trust in artificial intelligence: From a Foundational Trust Framework to emerging research opportunities. *Electronic Markets*. 2022;32(4):1993-2020.
- WHO Health Ethics & Governance (HEG). *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance*. World Health Organization; 2021. ISBN 978-92-4-002920-0
- De Fauw J, Ledsam JR, Romera-Paredes B, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine*. 2018;24(9):1342-1350.
- European Medicines Agency. Human Regulatory Overview: Medical Devices. EMA. Updated October 17, 2024. Accessed October 28, 2024. <https://www.ema.europa.eu/en/human-regulatory-overview/medical-devices>
- US Food and Drug Administration (FDA). Device Approvals and Clearances. FDA. Updated November 6, 2023. Accessed October 28, 2024. <https://www.fda.gov/medical-devices/products-and-medical-procedures/device-approvals-and-clearances>
- FDA. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)—Discussion Paper and Request for Feedback. Published April 2019. <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>
- FDA. Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan. In:2021.
- FDA. Good Machine Learning Practice for Medical Device Development: Guiding Principles. U.S. Food and Drug Administration. Updated October 27, 2021. Accessed October 28, 2024. <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles>
- Council of the European Union, European Parliament. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139, (EU) 2019/2144, and Directives 2014/90/EU, (EU) 2016/797, and (EU) 2020/1828 (Artificial Intelligence Act). Official Journal of the European Union. L 2024/1689. Published July 12, 2024.
- International Medical Device Regulators Forum. Machine Learning-enabled Medical Devices: Key Terms and Definitions. IMDRF/AIMD WG/N67 (Edition 1). Published May 9, 2022. Accessed October 28, 2024. <https://www.imdrf.org/documents/machine-learning-enabled-medical-devices-key-terms-and-definitions>
- Niazi S, Jiménez-García M, Findl O, et al. Keratoconus Diagnosis: From Fundamentals to Artificial Intelligence: A Systematic Narrative Review. *Diagnostics*. 2023;13(16):2715.
- Li Y, Foo L-L, Wong CW, et al. Pathologic myopia: advances in imaging and the potential role of artificial intelligence. *British Journal of Ophthalmology*. 2023;107(5):600-606.
- Poh SSJ, Sia JT, Yip MYT, et al. Artificial Intelligence, Digital Imaging, and Robotics Technologies for Surgical Vitreoretinal Diseases. *Ophthalmology Retina*. 2024;8(7):633-645.