


REVIEW

Unmasking biases and navigating pitfalls in the ophthalmic artificial intelligence lifecycle: A narrative review

Luis Filipe Nakayama^{1,2*}, João Matos^{2,3,4}, Justin Quion², Frederico Novaes¹, William Greig Mitchell⁵, Rogers Mwavu⁶, Claudia Ju-Yi Ji Hung^{7,8}, Alvina Pauline Dy Santiago^{9,10,11,12}, Warachaya Phanphruk¹³, Jaime S. Cardoso^{3,4}, Leo Anthony Celi^{2,14,15}

1 Department of Ophthalmology, Sao Paulo Federal University, Sao Paulo, Sao Paulo, Brazil, **2** Laboratory for Computational Physiology, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, **3** Faculty of Engineering (FEUP), University of Porto, Porto, Portugal, **4** Institute for Systems and Computer Engineering (INESC TEC), Technology and Science, Porto, Portugal, **5** Department of Ophthalmology, Royal Victorian Eye and Ear Hospital, Melbourne, Australia, **6** Department of Information Technology, Mbarara University of Science and Technology, Mbarara, Uganda, **7** Department of Ophthalmology, Byers Eye Institute at Stanford, California, United States of America, **8** Department of Computer Science and Information Engineering, National Taiwan University, Taiwan, **9** University of the Philippines Manila College of Medicine, Manila, Philippines, **10** Division of Pediatric Ophthalmology, Department of Ophthalmology & Visual Sciences, Philippine General Hospital, Manila, Philippines, **11** Section of Pediatric Ophthalmology, Eye and Vision Institute, The Medical City, Pasig, Philippines, **12** Section of Pediatric Ophthalmology, International Eye and Institute, St. Luke's Medical Center, Quezon City, Philippines, **13** Department of Ophthalmology, Faculty of Medicine, Khon Kaen University, Khon Kaen, Thailand, **14** Department of Biostatistics, Harvard TH Chan School of Public Health, Boston, Massachusetts, United States of America, **15** Department of Medicine, Beth Israel Deaconess Medical Center, Boston, Massachusetts, United States of America

* luisnaka@mit.edu



OPEN ACCESS

Citation: Nakayama LF, Matos J, Quion J, Novaes F, Mitchell WG, Mwavu R, et al. (2024) Unmasking biases and navigating pitfalls in the ophthalmic artificial intelligence lifecycle: A narrative review. *PLOS Digit Health* 3(10): e0000618. <https://doi.org/10.1371/journal.pdig.0000618>

Editor: Alexander Wong, University of Waterloo, CANADA

Published: October 8, 2024

Copyright: © 2024 Nakayama et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: The authors received no specific funding for this work.

Competing interests: LAC is the Editor-In-Chief of PLOS Digital Health.

Abstract

Over the past 2 decades, exponential growth in data availability, computational power, and newly available modeling techniques has led to an expansion in interest, investment, and research in Artificial Intelligence (AI) applications. Ophthalmology is one of many fields that seek to benefit from AI given the advent of telemedicine screening programs and the use of ancillary imaging. However, before AI can be widely deployed, further work must be done to avoid the pitfalls within the AI lifecycle. This review article breaks down the AI lifecycle into seven steps—data collection; defining the model task; data preprocessing and labeling; model development; model evaluation and validation; deployment; and finally, post-deployment evaluation, monitoring, and system recalibration—and delves into the risks for harm at each step and strategies for mitigating them.

Author summary

In recent years, the surge in data availability, computational power, and AI techniques has sparked interest in using AI in fields like ophthalmology. However, before widespread AI deployment can happen, we must carefully navigate its lifecycle, comprising 7 key steps:

data collection, task definition, data preparation, model development, evaluation, deployment, and post-deployment monitoring. This review article stands out by identifying potential pitfalls at each stage and offering actionable strategies to address them. Our article serves as a guide for harnessing AI effectively and safely in ophthalmology and related fields.

Introduction

The use of computers simulating humans in performing cognitive tasks was first described in 1943 [1], but the term “Artificial Intelligence” (AI) would not be coined until 1956. At that time, the fundamental premise was that “every aspect of learning or any other feature of intelligence can, in principle, be so precisely described that a machine can be made to simulate it” [2]. Over the last 2 decades, the exponential growth of data availability and computational power has led to an expansion in interest, investment, and research in AI applications [3].

Machine Learning (ML) is a subfield of AI that focuses on “teaching” computers how to make predictions, classify or optimize some function from data sets. While traditional ML requires a fair amount of work with feature engineering [4], in Deep Learning (DL) computers utilize multiple processing layers and multiple levels of abstraction to produce predictions, classifications, or optimal policies [5,6]. DL models have been used in several healthcare processes, including electronic medical records analysis, predicting patient outcomes, and image analysis [7].

In ophthalmology, the advent of telemedicine screening programs and the widespread use of ancillary imaging examinations have created large volumes of ophthalmic data, enabling the development of AI algorithms [8]. The ophthalmic community has since witnessed the application of AI to a wide range of diagnostic examinations with specialist-level performance [9–12]. Previous studies report the ability of AI-based algorithms to detect diabetic retinopathy [13–15], diabetic macular edema [16], retinopathy of prematurity [17,18], age-related macular degeneration [19–22], glaucoma [23–26], and uveitis [27], among others.

However, biased AI models can perpetuate existing health disparities. For example, algorithms designed to predict COVID-19 outcomes have been found to harbor implicit biases stemming from their reliance on non-representative data sets [28]. Similarly, facial recognition systems employed in criminal detection have demonstrated poorer accuracy rates when identifying females and individuals of black descent [29,30]. Additionally, language models have been shown to inadvertently perpetuate racism and stereotypes through their word choices [31,32]. Furthermore, commercial diabetic retinopathy screening algorithms have exhibited notable discrepancies in performance when applied across different populations within the same country, which can lead to delayed diagnosis and treatment in groups not represented in the training data set [33].

One huge concern is the unintentional encryption and propagation of biases at every stage of the AI lifecycle from data collection to post-deployment and recalibration [34]. In addition to data and algorithmic bias, other issues in the life cycle include reliance on irrelevant evaluation metrics, and a lack of post-deployment studies [28,35,36].

Biased predictions in artificial intelligence are a growing concern; however, the challenge is multifaceted. Bias can arise from every step of the AI lifecycle. This manuscript is the first to explore the biases that can arise in ophthalmological AI systems from data collection to deployment. Identifying pitfalls at every step of the AI pipeline is fundamental to breaking the cycle of health inequities.

The AI lifecycle in ophthalmology

Previous articles have explored the biases and pitfalls inherent in the AI development lifecycle, particularly in domains such as medical imaging [37–39]. However, this is the first to explore the ophthalmology specialty.

Ophthalmology relies on multiple imaging modalities to furnish supplementary data for clinical decision-making. Within the field, AI holds promise across various domains. For example, in diabetic retinopathy screening, AI-enabled analysis of retinal fundus photos can effectively discern patients who require specialist evaluation referrals, with deployed systems that provide autonomous recommendations based on retinal fundus analysis [40]. Moreover, in diagnosing and monitoring age-related macular degeneration, the utilization of retinal fundus photos, fundus autofluorescence, and optical coherence tomography are fundamental. Similarly, in the context of glaucoma, integrating fundus photos, visual field exams, and optical coherence tomography facilitates efficient screening and patient follow-up.

In this narrative review, we divided the ophthalmic AI lifecycle into 7 steps consisting of data collection; defining the model task; data preprocessing and labeling; model development; model evaluation and validation; deployment; and finally, post-deployment evaluation, monitoring, and system recalibration. In every step, we identify pitfalls and challenges (Fig 1).

Data collection

Clinical and imaging data needs to be collected to train AI algorithms. During this process, it is important to understand and document the population characteristics, data modalities, and included features and variables. In ophthalmology, images are collected during clinical practice and in screening programs (Fig 2).

Assembling new data sets is a laborious and expensive process. As a result, the ML community has mostly relied on publicly available data sets [41,42]. When employing these data sets, it is necessary to consider the data generation process, including but not limited to the demographic diversity of the patients, the examination criteria, and the disease distribution within the cohort. It is vital that the target population is adequately represented in the data during the development of ophthalmological models that detect ocular abnormalities in real-world settings [9].

Pitfalls

Unbalanced data sources. Representation is an important attribute to consider when assessing data to train algorithms. The use of overly restrictive inclusion criteria can result in the exclusion of several subgroups, while applying too broad criteria can lead to an increased number of “false positive” patients within a cohort [43]. Moreover, underrepresented populations traditionally are not included in data sets, and even in perfectly representative data, historical biases can lead to misrepresentation [44].

To date, ophthalmic data sets used in AI research are significantly imbalanced with no retinal imaging data sets from middle-low or low-income countries [42]. This imbalance has significant implications for the development of generalizable models.

In ophthalmology, including only images from screening programs and patients from a prespecified population, demographic, or health system will introduce bias and lead to a disparity in ML performance across subpopulations [45]. While it may be impossible to fully represent all subpopulations, efforts to understand the social patterning of the data generation process and engage with marginalized populations to develop strategies for inclusive data collection and analysis are required. At the minimum, awareness of the limitations of data sets should be reported.

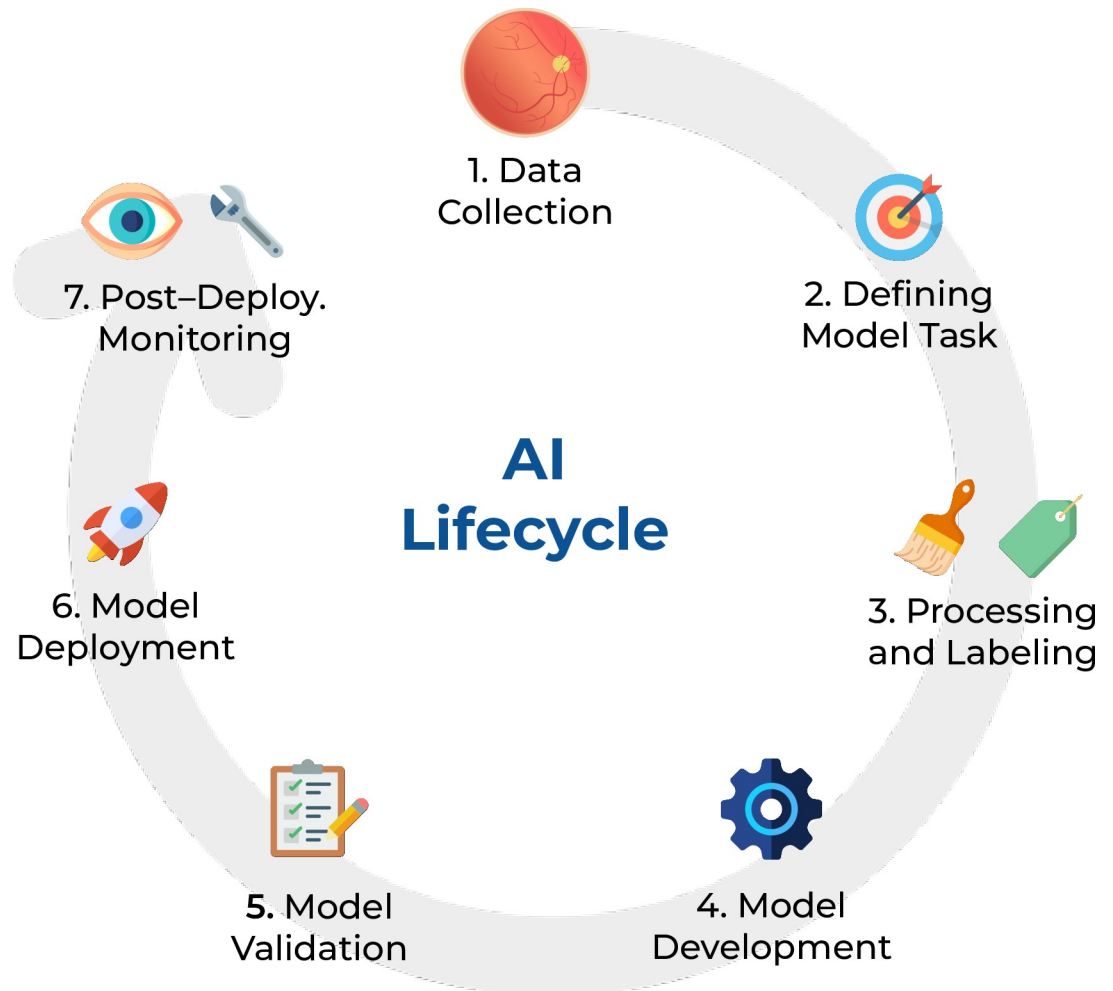


Fig 1. The 7 steps of the Ophthalmologic AI lifecycle. Diagram designed using icons from *Flaticon.com*.

<https://doi.org/10.1371/journal.pdig.0000618.g001>

When faced with imbalanced data sets that cannot be improved with more data, machine learning techniques offer strategies to mitigate biases. One approach involves adjusting class weights in the loss function, giving more weight to underrepresented classes and vice versa for the majority class. Oversampling and undersampling techniques can also be employed to balance the data set, although they come with their own set of challenges. Specifically, oversampling, which creates (equal or similar) copies of the minority class data points, can make the model more prone to overfit and computationally expensive to train; on the other hand, undersampling, whereby samples from the majority class are discarded, can result in a model less robust that does not leverage the data set to its fullest. When applying these techniques, it is fundamental to be cognizant of their limitations [46].

Image acquisition, protocols, and image quality. The image acquisition process involves the implementation of an imaging protocol, which can vary depending on the specific disease being investigated. In the case of diabetic retinopathy screening, it initially entailed capturing seven 30-degree fundus photos per eye, but evolved to either two 45-degree photos or a single image, depending on the screening program [47].

Image quality standards for data sets were established by experts from a handful of high-income countries, with little, if any, contribution from low- and middle-income countries.


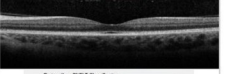
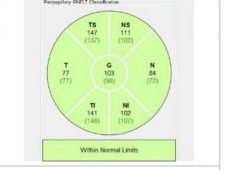

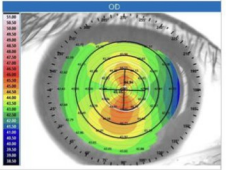
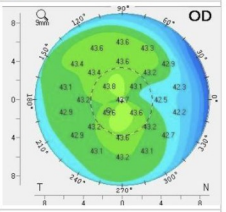
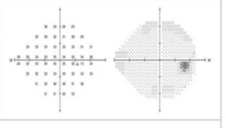
| Data modality | Exam Characteristics | Aim | Image | | |
|-------------------------|---|--|---|--|---|
| Imaging | Retinal fundus photo | Single image to evaluate posterior segment. |  | | |
| | | | | <i>Diabetic retinopathy</i> <i>Age-Related Macular Degeneration</i> <i>Retinopathy of Prematurity</i> | |
| | | | | <i>Glaucoma</i> | |
| | Optical Coherence Tomography | Cross-sectional high-resolution images that evaluate posterior or anterior ocular segment. | |   | |
| | | | | | <i>Diabetic Macular Edema</i> <i>Age-Related Macular Degeneration</i> <i>Macular diseases</i> |
| | | | | | <i>Glaucoma</i> |
| Anterior Segment photos | Single image to evaluate ocular anterior segment. | |  | | |
| | | | | <i>Corneal keratitis</i> <i>Cataract</i> | |
| | | | | <i>Neoplasia</i> | |
| Imaging | Corneal topography | Single image that provides a map of the corneal surface. |  | | |
| | Corneal tomography | Cross-sectional high-resolution images that evaluate cornea thickness and shape. |  | | |
| | | | | | <i>Corneal ectasias and irregularities</i> <i>Glaucoma</i> |
| | Visual Field | Assessment of visual sensitivity at different locations within visual field. | |  | |
| Tabular | Electronic Medical Records | | <i>Myopia</i> <i>Visual acuity</i> | | |
| | Clinical notes | | <i>Ocular lesions characteristics</i> | | |

Fig 2. Ophthalmological data modalities, subtypes, and diagnostic aims.

<https://doi.org/10.1371/journal.pdig.0000618.g002>

Such data standards may discourage creators of small or “low-quality” data sets from contributing to data repositories. Additionally, the same rich countries have the infrastructure and resources to create, curate, and maintain data warehouses and AI pipelines. Setting high standards for image quality in ophthalmological examinations may lead to better model performance but trades off real-world applicability [48]. It is important to ensure that initiatives to establish dataset standards involve more inclusive teams that consider the realities in most countries.

Spectrum bias. Spectrum bias occurs when the diseases studied within the data set are not fully representative of the spectrum of disease severity in the target population, leading to poor model performance [49]. Ophthalmological data sets skewed toward mild cases of diabetic retinopathy may lead to algorithms that are better at detecting mild cases while missing severe cases and vice versa [42,49].

Therefore, it is crucial to consider the model's performance across the full spectrum of the disease. The AI model should be trained and validated on diverse patient data, including different ages, ethnicities, disease stages, and comorbidities, to ensure applicability and reliability in ophthalmology.

Data standardization

A standardized data format is required across different systems to enable AI and should occur during the initial development phase [50,51]. Electronic healthcare record platforms are often fragmented and have limited interoperability [51]. In ophthalmology, data standardization and interoperability are challenging and require integrating multimodal data, including imaging, tabular, and clinical notes.

Defining the model task

In ophthalmology, understanding the target populations' clinical context, examination modalities, and disease prevalence and distribution is essential. Moreover, a comprehensive understanding of the socioeconomic context, the healthcare delivery system, and its resources are important for identifying ocular conditions to target and suitable tasks for training algorithms [52].

Pitfalls

Defining the appropriate disease target. The prevalence of ocular diseases varies significantly across geography and is shaped by several demographic and socioeconomic factors. Defining the disease target must be tailored according to the specific population and the healthcare context. While AI has seen applications in various ophthalmological diseases, diabetic retinopathy screening has been the most extensively explored to date [53]. While diabetic retinopathy is a global epidemic [2], there are many other causes of preventable blindness such as unaddressed refractive errors, affecting 914 million people worldwide; cataracts, affecting 94 million individuals; age-related macular degeneration, affecting 8 million; and glaucoma, affecting 7.7 million people [54]. Several other critical causes of blindness, such as keratoconus, infectious keratitis, and deficiency neuropathies, remain uncharted territories for AI technologies. These conditions are especially prevalent in economically disadvantaged regions where access to diagnostic laboratory services is limited [6]. Resources must be allocated based on the specific burden of ocular diseases in different regions.

Data sparsity bias. Data sparsity bias can manifest when a data set lacks sufficient case numbers for specific combinations of exposure and outcomes [55]. This limitation often occurs despite the use of large databases, resulting in a restricted sample of cases, risk factors, variables, and outcomes.

Traditionally, 2 approaches are employed to mitigate this: adjustment and penalization. While these methods may be effective, a deep understanding of the clinical context and underlying social determinants is necessary [55]. Failing to incorporate this information in the models will exacerbate existing biases. Researchers should actively seek to enhance data collection efforts, particularly in underrepresented or vulnerable populations.

Data preprocessing and labeling

Data preprocessing is an important step in building AI algorithms. Computer vision models, especially with traditional image analysis techniques, require preprocessing, consisting of image enhancement, filtering, and normalization. Additionally, data augmentation can enhance the size and quality of training data sets. This improves the robustness of models, particularly when dealing with limited data.

In ophthalmology, the preprocessing stage is even more critical, given the number of distinct capturing processes and retinal fundus photography devices. Ensuring accurate and reliable results necessitates accounting for the inherent technical differences among images.

Pitfalls

Bias in the handling of missing data. Missing data are due to many reasons, including human and machine limitations, errors during data collection, limited accessibility to clinical screening and examinations, and respondents' refusal to participate in research [56]. In ophthalmology data sets, missing information on comorbidities, demographics, and/or ophthalmological examination data is common [42]. Several methods exist for dealing with missingness, such as deleting instances with missing values or imputing the missing data using estimated values. However, the choice of imputation method should be made carefully, as different methods may introduce additional layers of bias.

Handling missing data with data imputation requires careful consideration of the context in which it is applied. For example, when dealing with a data set that disproportionately represents a certain group, such as the male population, employing the average weight or height from the entire data set to fill in missing values for the female population can lead to skewed results.

A similar issue arises in other domains, like ophthalmological imagery, where the use of generative AI techniques to augment data sets can inadvertently perpetuate biases if the training data is not representative of all groups. To mitigate such biases, employing stratified data imputation methods—where missing values are imputed based on the subgroup they belong to (e.g., using male data to impute missing values for males)—can be a more accurate and equitable approach. This strategy helps ensure that the imputed values are more representative of the specific subgroups, thereby reducing the introduction of bias during the imputation process.

To enhance the quality of data sets during development and reduce potential biases, it is essential to engage a diverse team of ophthalmologists, statisticians, data scientists, and even patient representatives as multi-stakeholder discussions will lead to more informed decisions [56,57].

Labeling. In supervised machine learning, accurate and reliable labels are necessary for training algorithms to predict diagnoses and outcomes. The labeling process in ophthalmology is complex and challenging, involving varying application of grading criteria and standards. Errors and inconsistencies in the labeling process can lead to biases in the models, ultimately impacting the performance and reliability of algorithms [58].

To address this issue, alternatives such as graders' consensus and expert adjudication can be employed to enhance the reliability of ophthalmological labeling [58]. The use of weakly supervised learning techniques offers a promising approach to improve the annotation process, making labeling more efficient and effective [59,60].

Model development

The model development phase consists of data curation and representation, and knowledge creation and validation. Avoiding biases during modeling is an arduous task [61,62] and

requires an in-depth understanding of data disparities, clinical confounders, and the socio-economic context [62].

Pitfalls

Flawed feature engineering or selection. The relevance of features when predicting or optimizing outcomes may vary across patient subgroups. When dealing with tabular data or traditional imaging techniques that require feature engineering and selection, it is essential to acknowledge the risk of both under- and overreliance on certain features. Not all features that can be extracted from EHRs or imaging exams are relevant and their inclusion in modeling may worsen task performance. The input of domain experts in this complex process is paramount.

Diagnostic suspicion bias. The knowledge of a patient's prior exposure or preexisting conditions can introduce diagnostic suspicion bias in data sets [63,64]. For example, when a patient has multiple known comorbidities, it may lead to a more extensive work-up, affecting the data collected for model development. Careful consideration is necessary when incorporating variables related to the outcome, such as the frequency of optical coherence tomography (OCT) and retinal fundus photos, as it could inadvertently bias the model towards specific diseases and risk factors.

Data leakage. Data leakage occurs when information from the training data is also captured in the validation cohort, or information about the outcome is reflected in the features, resulting in a falsely accurate model performance [65]. Notable examples from previously published models include using an antibiotic prescription to predict a sepsis diagnosis, or the blessing of a hospital chaplain to predict mortality. However, data leakage may be subtle and more difficult to detect [66].

In the context of ophthalmology, data leakage can occur when 3D volumetric OCT images are split on a per-image basis [65]. In this scenario, valuable spatial information from the volumetric data may inadvertently influence the model's predictions in ways that compromise the model's true ability to generalize to new data.

Shortcuts. Model shortcuts pertain to features that are learned during model development that are not clinically related to a prediction or classification, such as the hospital where a patient is seen or the medical equipment used for imaging [67]. Studies have shown that algorithms can perform seemingly impossible tasks, such as determining a patient's sex from a retinal fundus photography [68] or identifying race from a chest X-ray [69]. Such findings raise concerns about the use of such "invisible" features for diagnosis and treatment recommendation, rather than those pertaining to the clinical features of the disease.

In ophthalmology, small data sets, image details, and class imbalance can contribute to shortcut learning and measurement bias [34,70]. A thorough data analysis, careful interpretation of results, investigation of model explainability, and generalizability of a test model are needed to assess the impact of shortcut features [67].

Model evaluation and validation

AI algorithms are usually tested on external data sets to evaluate their real-world performance. Typical metrics include accuracy, sensitivity, specificity, the receiver operating characteristic (ROC) curve, and F1-score. It is crucial to choose the appropriate metrics and to perform a thorough analysis of downstream consequences of errors.

Pitfalls

Wrong evaluation metric. Using inappropriate evaluation metrics can hide the underperformance of models in certain patient subgroups and widen outcome disparities [34]. To

ensure algorithmic fairness, it is imperative to assess the performance across marginalized cohorts [71]. It is also essential to acknowledge that all evaluation metrics are only estimations of a construct and, therefore, may not fully capture the true extent of algorithmic fairness [67]. Metrics relying solely on accuracy, or other characteristics based on historical data, run the risk of perpetuating health inequities present in the data [71].

Wrong evaluation method. Evaluation bias arises when the methods used to assess model performance themselves are biased. Inadequate external and benchmark data sets that fail to represent the population accurately can result in a biased evaluation of the model's performance. To improve the robustness and generalizability of algorithms, and avoid shortcut learning, it is important to use external validation data sets that are representative of the target population [34,67].

Model deployment

The deployment of AI systems in real-world settings is the culmination of understanding the data characteristics and the social patterning of the data generation process, model task definition, data preprocessing, and model evaluation. Deployment and post-deployment monitoring and recalibration may introduce additional bias and require as much risk mitigation.

Pitfalls

Defining the model threshold. Defining the decision thresholds depends on the model's purpose (e.g., screening, diagnosis, triage) and the health system it is to be deployed in.

For example, in diabetic retinopathy screening, setting lower thresholds may increase the number of false positive results and unnecessary referral cases and workup. While this could lead to a higher sensitivity for detecting potential cases, it will also burden the healthcare system. Conversely, setting higher thresholds might reduce false positives but could lead to missed cases, compromising the sensitivity of the screening program and potentially delaying necessary interventions.

Comprehending the setting in which the model will be deployed is pivotal to its impact [49]. Deploying AI models, typically developed on data from well-resourced health systems, in limited-resource settings necessitates careful consideration of the available resources within the healthcare system [15]. To strike a balance between sensitivity and specificity, it is essential to consider factors such as the prevalence of the condition, the healthcare resources available, and the anticipated impact on patient outcomes. Multi-stakeholder engagement is crucial in this process.

Covariate and data set shift. Covariate shifts can arise when there are changes in the distribution of the input variables between the training and deployment phase. In the field of ophthalmology, the generalization capabilities of AI models can be compromised due to shifts that occur when transitioning from the training data set to the actual population they are being deployed to. These shifts can be attributed to various factors, including changes in patient demographics, changes in disease prevalence, differences in ophthalmological assessment timing, or variations in equipment use. Furthermore, population drift may occur as a result of a practice pattern change or enactment of new health policies [72,73].

Post-deployment evaluation, monitoring, and system recalibration

It is required to conduct post-deployment monitoring of algorithms and analyze their impact on clinical outcomes with special attention to historically disenfranchised groups within the health system. Continuous measurement of the impact disaggregated across patient subgroups and recalibration of the model as necessary will ensure equitable benefit of AI systems [74,75].

Moreover, it is important to preempt unintended consequences, including heightening social stigma when outcomes among certain groups are highlighted, before integrating the algorithm into clinical practice.

Pitfalls

Model updates. The decision of when and how to update AI systems can be a potential source of bias. The procedures for updating AI systems may differ across different settings, which can lead to differential performance and impact on patient outcomes. Consequently, it becomes essential to thoughtfully assess the potential biases that may arise during the updating process and implement appropriate measures to mitigate them. Timely updates to AI systems are essential for preserving their relevance as new clinical insights and practice patterns emerge.

Conclusion

The implementation of AI systems in healthcare holds great promise for enhancing clinical decision-making. Ophthalmology, in particular, employs imaging ancillary examinations for clinical practice, and telemedicine screening programs enable the development of automated systems. Future directions include incorporating autonomous systems into ophthalmological practice, thereby improving accessibility to ocular care. These initiatives include autonomous screening programs for diabetic retinopathy, age-related macular degeneration, and retinopathy of prematurity among others. Additionally, it can improve the decision-making on ocular treatments such as intravitreal injections and during ophthalmological surgeries.

However, it is important to be aware that bias can be introduced at every step of the AI lifecycle. To address this concern, we propose a seven-step process that encompasses data collection; defining the model task; data preprocessing and labeling; model development; model evaluation and validation; deployment; and finally, post-deployment evaluation, monitoring, and system recalibration. Understanding the impact of data disparities, modeling decisions, and evaluation metrics is vital to developing models that are accurate, equitable, and clinically relevant.

Collaboration between ophthalmologists, data scientists, social scientists, community representatives, and other stakeholders is crucial to this endeavor. In resource-limited settings, the tailoring of AI systems to address practical challenges becomes essential for successful real-world deployment. Moreover, the continuous monitoring of AI systems post-deployment is paramount to identify any biases that may emerge over time and ensure ongoing fairness and transparency. Vigilance in identifying and addressing biases, combined with preemptive measures to prevent unintended consequences, enables AI systems to be leveraged responsibly and ethically in clinical practice.

Author Contributions

Conceptualization: Luis Filipe Nakayama, Frederico Novaes, William Greig Mitchell, Alvina Pauline Dy Santiago, Leo Anthony Celi.

Data curation: Luis Filipe Nakayama, João Matos.

Formal analysis: Luis Filipe Nakayama, João Matos.

Investigation: Frederico Novaes.

Methodology: Luis Filipe Nakayama, Justin Quion, Frederico Novaes.

Project administration: Luis Filipe Nakayama.

Supervision: Leo Anthony Celi.

Visualization: João Matos.

Writing – original draft: Luis Filipe Nakayama, João Matos, Justin Quion, Frederico Novaes, William Greig Mitchell, Rogers Mwavu, Claudia Ju-Yi Ji Hung, Alvina Pauline Dy Santiago, Warachaya Phanphruk, Jaime S. Cardoso, Leo Anthony Celi.

Writing – review & editing: Luis Filipe Nakayama, João Matos, Justin Quion, Frederico Novaes, William Greig Mitchell, Rogers Mwavu, Claudia Ju-Yi Ji Hung, Alvina Pauline Dy Santiago, Warachaya Phanphruk, Jaime S. Cardoso, Leo Anthony Celi.

References

1. McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys.* 1943;115–133. <https://doi.org/10.1007/bf02478259>
2. Schmidt-Erfurth U, Sadeghipour A, Gerendas BS, Waldstein SM, Bogunović H. Artificial intelligence in retina. *Prog Retin Eye Res.* 2018; 67:1–29. <https://doi.org/10.1016/j.preteyeres.2018.07.004> PMID: 30076935
3. Zhang D, Maslej N, Brynjolfsson E, Etchemendy J, Lyons T, Manyika J, et al. The AI Index 2022 Annual Report. In: *aiindex.stanford.edu* [Internet]. [cited 2022 Dec 16]. Available from: https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.
4. Tong Y, Lu W, Yu Y, Shen Y. Application of machine learning in ophthalmic imaging modalities. *Eye Vis (Lond).* 2020; 7:22. <https://doi.org/10.1186/s40662-020-00183-6> PMID: 32322599
5. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* 2015; 521:436–444. <https://doi.org/10.1038/nature14539> PMID: 26017442
6. Kuo M-T, Hsu BW-Y, Yin Y-K, Fang P-C, Lai H-Y, Chen A, et al. A deep learning approach in diagnosing fungal keratitis based on corneal photographs. *Sci Rep.* 2020; 10:14424. <https://doi.org/10.1038/s41598-020-71425-9> PMID: 32879364
7. Esteva A, Robicquet A, Ramsundar B, Kuleshov V, DePristo M, Chou K, et al. A guide to deep learning in healthcare. *Nat Med.* 2019; 25:24–29. <https://doi.org/10.1038/s41591-018-0316-z> PMID: 30617335
8. Kapoor R, Walters SP, Al-Aswad LA. The current state of artificial intelligence in ophthalmology. *Surv Ophthalmol.* 2019; 64:233–240. <https://doi.org/10.1016/j.survophthal.2018.09.002> PMID: 30248307
9. Lu W, Tong Y, Yu Y, Xing Y, Chen C, Shen Y. Applications of Artificial Intelligence in Ophthalmology: General Overview. *J Ophthalmol.* 2018; 2018:5278196. <https://doi.org/10.1155/2018/5278196> PMID: 30581604
10. Du X-L, Li W-B, Hu B-J. Application of artificial intelligence in ophthalmology. *Int J Ophthalmol.* 2018; 11:1555–1561. <https://doi.org/10.18240/ijo.2018.09.21> PMID: 30225234
11. Lee J. Is Artificial Intelligence Better Than Human Clinicians in Predicting Patient Outcomes? *J Med Internet Res.* 2020; 22:e19918. <https://doi.org/10.2196/19918> PMID: 32845249
12. Korot E, Gonçalves MB, Khan SM, Struyven R, Wagner SK, Keane PA. Clinician-driven artificial intelligence in ophthalmology: resources enabling democratization. *Curr Opin Ophthalmol.* 2021; 32:445–451. <https://doi.org/10.1097/ICU.0000000000000785> PMID: 34265784
13. Abràmoff MD, Lou Y, Erginay A, Clarida W, Amelon R, Folk JC, et al. Improved Automated Detection of Diabetic Retinopathy on a Publicly Available Dataset Through Integration of Deep Learning. *Invest Ophthalmol Vis Sci.* 2016; 57:5200–5206. <https://doi.org/10.1167/iov.16-19964> PMID: 27701631
14. Gargeya R, Leng T. Automated Identification of Diabetic Retinopathy Using Deep Learning. *Ophthalmology.* 2017; 124:962–969. <https://doi.org/10.1016/j.ophtha.2017.02.008> PMID: 28359545
15. Ruamviboonsuk P, Tiwari R, Sayres R, Nganthavee V, Hemarat K, Kongprayoon A, et al. Real-time diabetic retinopathy screening by deep learning in a multisite national screening programme: a prospective interventional cohort study. *Lancet Digit Health.* 2022; 4:e235–e244. [https://doi.org/10.1016/S2589-7500\(22\)00017-6](https://doi.org/10.1016/S2589-7500(22)00017-6) PMID: 35272972
16. Li H-Y, Wang D-X, Dong L, Wei W-B. Deep learning algorithms for detection of diabetic macular edema in OCT images: A systematic review and meta-analysis. *Eur J Ophthalmol.* 2023; 33:278–290. <https://doi.org/10.1177/11206721221094786> PMID: 35473414

17. Bai A, Carty C, Dai S. Performance of deep-learning artificial intelligence algorithms in detecting retinopathy of prematurity: A systematic review. *Saudi J Ophthalmol*. 2022; 36:296–307. https://doi.org/10.4103/sjopt.sjopt_219_21 PMID: 36276252
18. Chen JS, Coyner AS, Ostmo S, Sonmez K, Bajimaya S, Pradhan E, et al. Deep Learning for the Diagnosis of Stage in Retinopathy of Prematurity: Accuracy and Generalizability across Populations and Cameras. *Ophthalmol Retina*. 2021; 5:1027–1035. <https://doi.org/10.1016/j.oret.2020.12.013> PMID: 33561545
19. Grassmann F, Mengelkamp J, Brandl C, Harsch S, Zimmermann ME, Linkohr B, et al. A Deep Learning Algorithm for Prediction of Age-Related Eye Disease Study Severity Scale for Age-Related Macular Degeneration from Color Fundus Photography. *Ophthalmology*. 2018; 125:1410–1420. <https://doi.org/10.1016/j.ophtha.2018.02.037> PMID: 29653860
20. 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; 136:1305–1307. <https://doi.org/10.1001/jamaophthalmol.2018.3799> PMID: 30193354
21. Burlina PM, Joshi N, Pacheco KD, Freund DE, Kong J, Bressler NM. Use of deep learning for detailed severity characterization and estimation of 5-year risk among patients with age-related macular degeneration. *JAMA Ophthalmol*. 2018; 136:1359–1366. <https://doi.org/10.1001/jamaophthalmol.2018.4118> PMID: 30242349
22. Pramil V, de Sisternes L, Omlor L, Lewis W, Sheikh H, Chu Z, et al. A Deep Learning Model for Automated Segmentation of Geographic Atrophy Imaged Using Swept-Source OCT. *Ophthalmol Retina*. 2023; 7:127–141. <https://doi.org/10.1016/j.oret.2022.08.007> PMID: 35970318
23. Noury E, Mannil SS, Chang RT, Ran AR, Cheung CY, Thapa SS, et al. Deep Learning for Glaucoma Detection and Identification of Novel Diagnostic Areas in Diverse Real-World Datasets. *Transl Vis Sci Technol*. 2022; 11:11. <https://doi.org/10.1167/tvst.11.5.11> PMID: 35551345
24. Thompson AC, Jammal AA, Medeiros FA. A Review of Deep Learning for Screening, Diagnosis, and Detection of Glaucoma Progression. *Transl Vis Sci Technol*. 2020; 9:42. <https://doi.org/10.1167/tvst.9.2.42> PMID: 32855846
25. Thompson AC, Jammal AA, Berchuck SI, Mariottoni EB, Medeiros FA. Assessment of a Segmentation-Free Deep Learning Algorithm for Diagnosing Glaucoma From Optical Coherence Tomography Scans. *JAMA Ophthalmol*. 2020; 138:333–339. <https://doi.org/10.1001/jamaophthalmol.2019.5983> PMID: 32053142
26. Liu H, Li L, Wormstone IM, Qiao C, Zhang C, Liu P, et al. Development and Validation of a Deep Learning System to Detect Glaucomatous Optic Neuropathy Using Fundus Photographs. *JAMA Ophthalmol*. 2019; 137:1353–1360. <https://doi.org/10.1001/jamaophthalmol.2019.3501> PMID: 31513266
27. Nakayama LF, Ribeiro LZ, Dychiao RG, Zamora YF, Regatieri CVS, Celi LA, et al. Artificial intelligence in uveitis: A comprehensive review. *Surv Ophthalmol*. 2023. <https://doi.org/10.1016/j.survophthal.2023.02.007> PMID: 36878360
28. Roberts M, Driggs D, Thorpe M, Gilbey J, Yeung M, Ursprung S, et al. Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. *Nat Mach Intell*. 2021; 3:199–217. <https://doi.org/10.1038/s42256-021-00307-0>
29. Klare BF, Burge MJ, Klontz JC, Vorder Bruegge RW, Jain AK. Face Recognition Performance: Role of Demographic Information. *IEEE Trans Inf Forensics Secur*. 2012; 7:1789–1801. <https://doi.org/10.1109/TIFS.2012.2214212>
30. Najibi A. Racial discrimination in face recognition technology. *Harvard Online: Science Policy and Social Justice*.
31. Caliskan A, Bryson JJ, Narayanan A. Semantics derived automatically from language corpora contain human-like biases. *Science*. 2017; 356:183–186. <https://doi.org/10.1126/science.aal4230> PMID: 28408601
32. Zhang H, Lu AX, Abdalla M, McDermott M, Ghassemi M. Hurtful Words: Quantifying Biases in Clinical Contextual Word Embeddings. *arXiv [csCL]*. 2020. Available from: <http://arxiv.org/abs/2003.11515>
33. Lee AY, Yanagihara RT, Lee CS, Blazes M, Jung HC, Chee YE, et al. Multicenter, Head-to-Head, Real-World Validation Study of Seven Automated Artificial Intelligence Diabetic Retinopathy Screening Systems. *Diabetes Care*. 2021; 44:1168–1175. <https://doi.org/10.2337/dc20-1877> PMID: 33402366
34. Suresh H, Gutttag J. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. *Equity and Access in Algorithms, Mechanisms, and Optimization*. New York, NY, USA: Association for Computing Machinery; 2021. p. 1–9. <https://doi.org/10.1145/3465416.3483305>
35. Habib AR, Lin AL, Grant RW. The Epic Sepsis Model Falls Short—The Importance of External Validation. *JAMA Intern Med*. 2021; 181:1040–1041. <https://doi.org/10.1001/jamainternmed.2021.3333> PMID: 34152360

36. Wong A, Otlies E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Intern Med.* 2021. <https://doi.org/10.1001/jamainternmed.2021.2626> PMID: 34152373
37. Gichoya JW, Thomas K, Celi LA, Safdar N, Banerjee I, Banja JD, et al. AI pitfalls and what not to do: mitigating bias in AI. *Br J Radiol.* 2023; 96:20230023. <https://doi.org/10.1259/bjr.20230023> PMID: 37698583
38. Maier-Hein L, Reinke A, Godau P, Tizabi MD, Buettner F, Christodoulou E, et al. Metrics reloaded: recommendations for image analysis validation. *Nat Methods.* 2024; 21:195–212. <https://doi.org/10.1038/s41592-023-02151-z> PMID: 38347141
39. Reinke A, Tizabi MD, Baumgartner M, Eisenmann M, Heckmann-Nötzel D, Kavur AE, et al. Understanding metric-related pitfalls in image analysis validation. *Nat Methods.* 2024; 21:182–194. <https://doi.org/10.1038/s41592-023-02150-0> PMID: 38347140
40. Nakayama LF, Zago Ribeiro L, Novaes F, Miyawaki IA, Miyawaki AE, de Oliveira JAE, et al. Artificial intelligence for telemedicine diabetic retinopathy screening: a review. *Ann Med.* 2023; 55:2258149. <https://doi.org/10.1080/07853890.2023.2258149> PMID: 37734417
41. Decenci ere E, Zhang X, Cazuguel G, Lay B, Cochener B, Trone C, et al. Feedback on a publicly distributed image database: The Messidor database. *Image Anal Stereol.* 2014; 33:231. <https://doi.org/10.5566/ias.1155>
42. Khan SM, Liu X, Nath S, Korot E, Faes L, Wagner SK, et al. A global review of publicly available datasets for ophthalmological imaging: barriers to access, usability, and generalisability. *Lancet Digit Health.* 2021; 3:e51–e66. [https://doi.org/10.1016/S2589-7500\(20\)30240-5](https://doi.org/10.1016/S2589-7500(20)30240-5) PMID: 33735069
43. Sauer CM, Chen L-C, Hyland SL, Girbes A, Elbers P, Celi LA. Leveraging electronic health records for data science: common pitfalls and how to avoid them. *Lancet Digit Health.* 2022; 4:e893–e898. [https://doi.org/10.1016/S2589-7500\(22\)00154-6](https://doi.org/10.1016/S2589-7500(22)00154-6) PMID: 36154811
44. Garg N, Schiebinger L, Jurafsky D, Zou J. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proc Natl Acad Sci U S A.* 2018; 115:E3635–E3644. <https://doi.org/10.1073/pnas.1720347115> PMID: 29615513
45. Burlina P, Joshi N, Paul W, Pacheco KD, Bressler NM. Addressing Artificial Intelligence Bias in Retinal Diagnostics. *Transl Vis Sci Technol.* 2021; 10:13. <https://doi.org/10.1167/tvst.10.2.13> PMID: 34003898
46. He H, Garcia EA. Learning from Imbalanced Data. *IEEE Trans Knowl Data Eng.* 2009; 21:1263–1284. <https://doi.org/10.1109/TKDE.2008.239>
47. Huemer J, Wagner SK, Sim DA. The Evolution of Diabetic Retinopathy Screening Programmes: A Chronology of Retinal Photography from 35 mm Slides to Artificial Intelligence. *Clin Ophthalmol.* 2020; 14:2021–2035. <https://doi.org/10.2147/OPTH.S261629> PMID: 32764868
48. Heaven WD. Google’s medical AI was super accurate in a lab. Real life was a different story. *MIT Technology Review.* 27 Apr 2020. Available from: <https://www.technologyreview.com/2020/04/27/1000658/google-medical-ai-accurate-lab-real-life-clinic-covid-diabetes-retina-disease/>. Accessed 2023 Feb 6.
49. Faes L, Liu X, Wagner SK, Fu DJ, Balaskas K, Sim DA, et al. A Clinician’s Guide to Artificial Intelligence: How to Critically Appraise Machine Learning Studies. *Transl Vis Sci Technol.* 2020;7. <https://doi.org/10.1167/tvst.9.2.7> PMID: 32704413
50. He J, Baxter SL, Xu J, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nat Med.* 2019; 25:30–36. <https://doi.org/10.1038/s41591-018-0307-0> PMID: 30617336
51. Kruse CS, Goswamy R, Raval Y, Marawi S. Challenges and Opportunities of Big Data in Health Care: A Systematic Review. *JMIR Med Inform.* 2016; 4:e38. <https://doi.org/10.2196/medinform.5359> PMID: 27872036
52. Burton MJ, Ramke J, Marques AP, Bourne RRA, Congdon N, Jones I, et al. The Lancet Global Health Commission on Global Eye Health: vision beyond 2020. *Lancet Glob Health.* 2021; 9:e489–e551. [https://doi.org/10.1016/S2214-109X\(20\)30488-5](https://doi.org/10.1016/S2214-109X(20)30488-5) PMID: 33607016
53. Ting DSW, Pasquale LR, Peng L, Campbell JP, Lee AY, Raman R, et al. Artificial intelligence and deep learning in ophthalmology. *Br J Ophthalmol.* 2019; 103:167–175. <https://doi.org/10.1136/bjophthalmol-2018-313173> PMID: 30361278
54. Blindness and vision impairment. [cited 2023 Jan 26]. Available from: <https://www.who.int/news-room/fact-sheets/detail/blindness-and-visual-impairment>.
55. Greenland S, Mansournia MA, Altman DG. Sparse data bias: a problem hiding in plain sight. *BMJ.* 2016; 352:i1981. <https://doi.org/10.1136/bmj.i1981> PMID: 27121591
56. Emmanuel T, Maupong T, Mpoeleng D, Semong T, Mphago B, Tabona O. A survey on missing data in machine learning. *J Big Data.* 2021; 8:140. <https://doi.org/10.1186/s40537-021-00516-9> PMID: 34722113

57. Langkamp DL, Lehman A, Lemeshow S. Techniques for handling missing data in secondary analyses of large surveys. *Acad Pediatr*. 2010; 10:205–210. <https://doi.org/10.1016/j.acap.2010.01.005> PMID: 20338836
58. Krause J, Gulshan V, Rahimy E, Karth P, Widner K, Corrado GS, et al. Grader Variability and the Importance of Reference Standards for Evaluating Machine Learning Models for Diabetic Retinopathy. *Ophthalmology*. 2018; 125:1264–1272. <https://doi.org/10.1016/j.ophtha.2018.01.034> PMID: 29548646
59. Wang J, Li W, Chen Y, Fang W, Kong W, He Y, et al. Weakly supervised anomaly segmentation in retinal OCT images using an adversarial learning approach. *Biomed Opt Express*. 2021; 12:4713–4729. <https://doi.org/10.1364/BOE.426803> PMID: 34513220
60. Playout C, Duval R, Cheriet F. A Novel Weakly Supervised Multitask Architecture for Retinal Lesions Segmentation on Fundus Images. *IEEE Trans Med Imaging*. 2019; 38:2434–2444. <https://doi.org/10.1109/TMI.2019.2906319> PMID: 30908197
61. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019:447–453. <https://doi.org/10.1126/science.aax2342> PMID: 31649194
62. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat Med*. 2021; 27:2176–2182. <https://doi.org/10.1038/s41591-021-01595-0> PMID: 34893776
63. Delgado-Rodríguez M, Llorca J. Bias. *J Epidemiol Community Health*. 2004; 58:635–641. <https://doi.org/10.1136/jech.2003.008466> PMID: 15252064
64. Hegedus EJ, Moody J. Clinimetrics corner: the many faces of selection bias. *J Man Manip Ther*. 2010; 18:69–73. <https://doi.org/10.1179/106698110X12640740712699> PMID: 21655388
65. Tampu IE, Eklund A, Haj-Hosseini N. Inflation of test accuracy due to data leakage in deep learning-based classification of OCT images. *Sci Data*. 2022; 9:580. <https://doi.org/10.1038/s41597-022-01618-6> PMID: 36138025
66. Kaufman S, Rosset S, Perlich C, Stitelman O. Leakage in data mining: Formulation, detection, and avoidance. *ACM Trans Knowl Discov Data*. 2012; 6:1–21. <https://doi.org/10.1145/2382577.2382579>
67. Geirhos R, Jacobsen J-H, Michaelis C, Zemel R, Brendel W, Bethge M, et al. Shortcut learning in deep neural networks. *Nat Mach Intell*. 2020; 2:665–673. <https://doi.org/10.1038/s42256-020-00257-z>
68. Korot E, Pontikos N, Liu X, Wagner SK, Faes L, Huemer J, et al. Predicting sex from retinal fundus photographs using automated deep learning. *Sci Rep*. 2021; 11:10286. <https://doi.org/10.1038/s41598-021-89743-x> PMID: 33986429
69. Gichoya JW, Banerjee I, Bhimireddy AR, Burns JL, Celi LA, Chen L-C, et al. AI recognition of patient race in medical imaging: a modelling study. *Lancet Digit Health*. 2022; 4:e406–e414. [https://doi.org/10.1016/S2589-7500\(22\)00063-2](https://doi.org/10.1016/S2589-7500(22)00063-2) PMID: 35568690
70. Robinson C, Trivedi A, Blazes M, Ortiz A, Desbiens J, Gupta S, et al. Deep learning models for COVID-19 chest x-ray classification: Preventing shortcut learning using feature disentanglement. <https://doi.org/10.1101/2021.02.11.20196766> PMID: 33594382
71. Mbakwe AB, Lourentzou I, Celi LA, Wu JT. Fairness metrics for health AI: we have a long way to go. *EBioMedicine*. 2023; 90:104525. <https://doi.org/10.1016/j.ebiom.2023.104525> PMID: 36924621
72. Subbaswamy A, Saria S. From development to deployment: dataset shift, causality, and shift-stable models in health AI. *Biostatistics*. 2020; 21:345–352. <https://doi.org/10.1093/biostatistics/kxz041> PMID: 31742354
73. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The Clinician and Dataset Shift in Artificial Intelligence. *N Engl J Med*. 2021; 385:283–286. <https://doi.org/10.1056/NEJMc2104626> PMID: 34260843
74. Jacoba CMP, Celi LA, Lorch AC, Fickweiler W, Sobrin L, Gichoya JW, et al. Bias and non-diversity of big data in artificial intelligence: Focus on retinal diseases. *Semin Ophthalmol*. 2023;1–9. <https://doi.org/10.1080/08820538.2023.2168486> PMID: 36651834
75. Iqbal U, Celi LA, Hsu Y-HE, Li Y-CJ. Healthcare artificial intelligence: the road to hell is paved with good intentions. *BMJ Health Care Inform*. 2022;29. <https://doi.org/10.1136/bmjhci-2022-100650> PMID: 35940638