




AZMAN JOURNAL OF HEALTH AND SCIENCES (AJOHAS)



Volume 1, Issue 1, June, 2026

www.ajohas.azmanuniversity.edu.ng

Ethical Considerations and Algorithmic Bias in AI-Driven Genetic Diagnosis: A Systematic Review

Hafsat Ashafa¹ *, Aisha Umar Suleiman²

1. Department of Computer Science, Northwest University Kano, Nigeria.

2. Department of Cyber security², Northwest University Kano, Nigeria

Correspondence: ashafahafsat@gmail.com

Abstract

The integration of artificial intelligence (AI) into genetic diagnosis has enabled the rapid analysis of complex genomic data, thereby transforming clinical decision-making. However, because many AI models are trained on datasets predominantly derived from populations of European ancestry significant ethical concerns related to algorithmic bias have emerged. This systematic review aims to: (1) identify the major ethical issues and impacts of algorithmic bias in AI-driven genetic diagnosis, (2) examine existing ethical frameworks and proposed mitigation strategies and (3) outline critical success factors for ensuring fairness, accountability and transparency.

A systematic literature search was conducted in accordance with PRISMA guidelines across PubMed, Scopus, Web of Science, IEEE Xplore and the ACM Digital Library. Search terms related to algorithmic bias, genetic diagnosis, artificial intelligence, and ethics were combined. Study selection, data extraction and thematic synthesis were performed. Following screening and eligibility assessment, 46 peer-reviewed studies published between 2014 and 2025 were included for qualitative synthesis.

The findings indicate that algorithmic bias poses a substantial risk to health equity, particularly through reduced diagnostic accuracy and higher rates of variants of uncertain significance among individuals from underrepresented populations. Additional ethical challenges related to transparency, accountability, informed consent and patient autonomy were identified. Four critical success factors emerged from the synthesis: robust regulatory and accountability frameworks, meaningful multi-stakeholder engagement, mandatory transparency and explainability and equity-centered lifecycle management.

Addressing algorithmic bias in AI-driven genetic diagnosis is essential to achieving ethical and equitable precision medicine. The implementation of the identified critical success factors is necessary to ensure that AI applications in genomics benefit all populations fairly.

Keywords: ethical consideration, algorithmic bias, genetic diagnosis, health equity, artificial intelligence, precision medicine, systematic review.

INTRODUCTION

The application of artificial intelligence (AI) into genetic medicine has substantially advanced diagnostic capabilities by enabling the rapid and automated analysis of large-scale genomic data. AI-driven tools are increasingly employed in clinical genetics for variant interpretation, gene disease association analysis and predictive diagnostics, thereby supporting more efficient and data-driven clinical decision-making (Topol, 2019; Chen et al., 2021). As these technologies become embedded within routine clinical workflows, concerns have

emerged regarding their ethical implications, particularly the risk of algorithmic bias.

Algorithmic bias refers to systematic and repeatable errors in algorithmic decision-making that result in unfair or inequitable outcomes for certain individuals or groups, often arising from unrepresentative training data, flawed model design or biased deployment contexts (Rajkomar et al., 2018; Obermeyer et al., 2019). In the context of genetic diagnosis, algorithmic bias can lead to disparities in diagnostic accuracy and clinical interpretation across populations.

Algorithmic bias in genetic diagnosis has transformed clinical genomics by enabling the large-scale analysis of complex genomic data. AI-based systems are now routinely used for variant interpretation, disease risk prediction, and gene–disease association studies, contributing to more efficient and data-driven diagnostic processes (Topol, 2019; Chen et al., 2021). As these technologies become increasingly embedded in clinical practice, attention has shifted from technical performance alone to their broader ethical and social implications.

Algorithmic bias in the context of AI-driven healthcare refers to systematic and reproducible errors that lead to unequal outcomes across different population groups. Such bias commonly arises from unrepresentative training datasets, inappropriate model assumptions or context-insensitive deployment (Rajkomar et al., 2018; Obermeyer et al., 2019). In genetic diagnosis, algorithmic bias is particularly consequential because predictive accuracy is highly dependent on the population diversity and quality of genomic reference data.

The genomic nature of algorithmic bias is rooted in longstanding structural imbalances in genetic research. Genome-wide association studies and major reference databases have historically overrepresented individuals of European ancestry, resulting in limited representation of African, Asian and Indigenous populations (Popejoy & Fullerton, 2016; Sirugo et al., 2019). AI models trained on such datasets may therefore exhibit reduced diagnostic accuracy and higher uncertainty when applied to underrepresented groups. This imbalance contributes to lower diagnostic yield, increased rates of variants of uncertain significance and unequal clinical outcomes, thereby reinforcing existing health disparities (Martin et al., 2019).

Beyond data representation, AI-driven genetic diagnosis raises distinct ethical concerns related to fairness, accountability and transparency. Many machine-learning models operate as “black-box” systems offering limited insight into how genomic data are translated into diagnostic conclusions. This opacity complicates clinical interpretation, undermines informed consent and weakens trust

between patients and healthcare providers (Char et al., 2018). In addition, accountability becomes fragmented in multi-actor environments involving data custodians, algorithm developers, healthcare institutions and clinicians.

Privacy and autonomy further complicate the ethical landscape of genomic AI. The large-scale aggregation and reuse of genetic data increase risks related to re-identification, secondary data use and consent fatigue (Rehm et al., 2015; Sankar & Parker, 2017). In response, emerging governance frameworks emphasize transparency, explainability, dynamic consent and federated data architectures as mechanisms to protect individual rights while enabling responsible innovation (Birney et al., 2017; Stark et al., 2019).

Despite growing interest in ethical AI, there remains limited integrative analysis of how algorithmic bias operates specifically within AI-driven genetic diagnosis and how ethical mitigation strategies can be operationalized in genomic practice. Accordingly, this systematic review synthesizes existing literature on algorithmic bias in AI-based genetic diagnostics, examines associated ethical challenges, evaluates proposed mitigation frameworks and identifies critical success factors for promoting fairness, accountability and equity in genomic medicine.

METHODS

Review Design and Reporting Framework

This study was conducted as a systematic review and reported according to the PRISMA 2020 guidelines (Page et al., 2021). The review employed a structured and reproducible methodology to identify, screen, appraise and synthesize literature addressing Ethical, Legal and Social Implications (ELSI) and algorithmic bias in AI-driven genetic and genomic diagnosis (Birney et al., 2017; Char et al., 2018; Chen et al., 2021; Wiens et al., 2019). The approach integrated perspectives from public health, computer science, medical ethics and bioinformatics to capture interdisciplinary insights (Stark et al., 2019; Fatumo et al., 2022; Li & Iacobelli, 2023).




Fig 2.1 Review process

Data Sources and Search Strategy

A comprehensive search was conducted across PubMed, Scopus, Web of Science, IEEE Xplore and ACM Digital Library. The search targeted peer-reviewed literature published between 2014 and 2025 in English. The search strategy combined controlled vocabulary (IEEE/ACM subject terms) and free-text keywords with Boolean operators, including:

“Artificial intelligence” OR “machine learning” AND “Genetic diagnosis” OR “genomic interpretation” OR “precision medicine” AND “Ethics” OR “algorithmic bias” OR “fairness” OR “justice” OR “equity”

Additional studies were identified through reference snowballing and cross-referencing of included articles (Birney et al., 2017; Obermeyer et al., 2019; Topol, 2019; Agarwal et al., 2023; Hanna et al., 2025).

Study Selection Process

All retrieved records were exported into EndNote for management, and duplicates were removed both automatically and manually: Initial records identified: 564. Duplicates removed: 67. Studies

excluded: 344 Full-text articles assessed for eligibility: 153

Full-text articles excluded: 76 (primarily due to non-genomic AI focus, lack insufficient ethical focus or commentary-only articles). Studies before 2014 excluded: 31. Studies included in final synthesis: 46. The selection process is summarized in a PRISMA flow diagram (Figure 1) to ensure transparency and reproducibility (Page et al., 2021; Birney et al., 2017).

Eligibility Criteria

Inclusion Criteria

Studies were included if they: Examined ethical, legal, or social implications of AI or ML in genetic/genomic diagnosis or precision medicine. Addressed algorithmic bias, fairness, equity, accountability or explainability. Presented empirical findings, systematic reviews, policy analyses or conceptual frameworks (Wiens et al., 2019; Chen et al., 2021; Drabiak et al., 2023).

Were published in peer-reviewed journals.

Exclusion Criteria

Studies were excluded if they: Focused on AI applications unrelated to genetics or genomics.

Addressed purely technical optimization without discussion of ethical implications (Cabitza et al., 2017). Were commentaries or opinion pieces lacking sufficient analytical depth and studies from before 2014 were also excluded.




Fig 2.2: Systematics Review Process

Screening and Data Extraction

Screening was conducted independently by two reviewers at the title/abstract and full-text stages. Discrepancies were resolved through discussion and consensus.

Data extraction followed a standardized form capturing: Study objectives and context (AI application, genomic setting)

Ethical issues addressed: Sources and manifestations of algorithmic bias. Proposed mitigation strategies and policy or regulatory recommendations (Char et al., 2018; Obermeyer et al., 2019; Agarwal et al., 2023; Hanna et al., 2025). Data extraction and verification ensured rigor and minimized reviewer bias.

Quality Appraisal

The methodological and ethical quality of studies was evaluated using adapted frameworks for

responsible AI and bioethics research (Wiens et al., 2019; Chen et al., 2021; McCradden et al., 2022; Jha et al., 2025). Studies were assessed for: Conceptual clarity, Ethical depth and relevance, Transparency of methods. Applicability to clinical genomics and AI-driven healthcare.

Quality appraisal informed interpretation but did not exclude studies to ensure a comprehensive overview.

Data Synthesis

A thematic synthesis approach was used to integrate findings across the 46 included studies. Recurring ethical themes, sources of bias and mitigation frameworks were identified and grouped into higher-order categories. This synthesis enabled derivation of critical success factors for ethical AI implementation in genomic and precision medicine, including: Regulatory and accountability frameworks. Multi-stakeholder engagement. Transparency and explainability mandates. Equity-centered lifecycle management (Wiens et al., 2019; Chen et al., 2021; Stark et al., 2019; Birney et al., 2017)

Results

Overview of the Included Research

46 peer-reviewed articles from 2014 to 2025 that represented interdisciplinary viewpoints from public health, computer science, medical genetics, and bioethics were combined. Systematic reviews, policy assessments, empirical studies and ethical commentary were all included in the literature. The majority of research came from North America and Europe but Africa and Asia began to contribute more and more, focusing on data sovereignty and global equality (The H3Africa Consortium et al., 2014; Stark et al., 2019; Fatumo et al., 2022). Table 3.1 summarizes the key characteristics of the included studies, including authors, year, study type; ethical issue addressed and key findings.

Authors	years	Study type	Ethical Issue addressed	Key Finding
Agarwal et al.,	2023	Review	Algorithmic bias	Framework proposed to detect and mitigate AI bias in health decision-making.
Chin et al.,	2023	Review	Algorithmic bias mitigation	Evidence and research needs to eliminate bias in healthcare.
Moore et al.,	2014	Empirical	Clinical prediction	STONE score development; balancing accuracy and patient safety.
Birney et al.,	2017	Commentary	Data equity / global collaboration	GA4GH initiative promotes global genomics integration and equitable data sharing.
Chen et al.,	2021	Empirical	Algorithmic reliability	Deep autoencoding methods can detect anomalies but lack interpretability for clinical use.
Drabiak et al.,	2023	Review	Diversity & global impact	AI in radiology must account for diversity, ethics, and cross-border implications
Li T, Iacobelli F	2023	Conference	Conference on computer supported cooperative work and social computing. Minneapolis, MN, USA: Association for Computing Machinery;	Purposeful AI. In: Companion publication of the
Kullar R, et al.	2020		Racial disparity of coronavirus disease 2019 in African American communities..	J Infect Dis
Shahian et al.,	2018	Policy / Empirical	Cardiac surgery risk	Development of STS adult cardiac surgery risk models; addresses transparency and fairness.
Vyas et al.,	2020	Commentary	Race correction in clinical algorithms	Using race in algorithms may perpetuate inequities; reevaluation needed
Tejani et al.,	2024	Checklist	AI reporting update	2024 CLAIM update enhances transparency and reproducibility for AI studies.

3.2 Thematic Synthesis

The thematic synthesis of the 46 included studies yielded four interrelated but conceptually distinct themes. Together, these themes illuminate the ethical risks of AI-driven genetic diagnosis and the

practical conditions required for responsible implementation.

Theme 1: Data Representation and Structural Bias

A dominant ethical concern identified across the literature is the systematic underrepresentation of

non-European populations in genomic datasets. Studies consistently demonstrate that AI models trained on Eurocentric genomic data exhibit reduced diagnostic accuracy for individuals of African, Asian, and Indigenous ancestry, leading to higher rates of variants of uncertain significance and misclassification (Popejoy & Fullerton, 2016; Sirugo et al., 2019; Wojcik et al., 2019).

This structural bias is not merely a technical limitation but constitutes an ethical failure that conflicts with the principles of distributive justice, equity, and non-maleficence. By embedding historical and socio-political inequities into algorithmic systems, AI-driven genomic tools risk perpetuating health disparities rather than alleviating them (Rajkomar et al., 2018; Wiens et al., 2019). The literature emphasizes that fairness cannot be retrofitted solely through post hoc adjustments; rather, equity must be addressed at the level of data generation, curation, and governance.

Theme 2: Transparency, Explainability, and Accountability in Clinical AI Use

A second major theme concerns the opacity of AI systems used in genetic diagnosis. Many AI-driven models function as “black boxes,” offering limited insight into how genomic features are weighted or how diagnostic outputs are generated. This lack of explainability poses ethical challenges when AI outputs inform high-stakes clinical decisions, such as disease risk prediction or reproductive counseling (Char et al., 2018; Chen et al., 2021).

The absence of transparency undermines clinical accountability, patient trust, and informed consent, as clinicians may be unable to justify or contest algorithmic recommendations. Several studies argue that explainable AI (XAI), auditability, and traceability mechanisms are essential to align AI systems with established norms of medical responsibility (Obermeyer et al., 2019). Importantly, this theme focuses on the ethical problem itself—opacity and responsibility gaps—rather than the broader governance solutions, which are addressed separately under Theme 4.

Theme 3: Privacy, Consent and Patient Autonomy
The integration of AI into genomic medicine challenges conventional understandings of privacy

and informed consent. Genomic data are uniquely identifiable and often reused across multiple research and clinical contexts, increasing the risk of secondary use without adequate consent and potential re-identification (Rehm et al., 2015; Sankar & Parker, 2017).

Several studies highlight that static, one-time consent models are insufficient for AI-driven genomics, where data may be continuously reused to retrain algorithms. Ethical frameworks increasingly advocate for dynamic consent, data minimization, and privacy-preserving techniques, such as federated learning, to balance individual autonomy with collective scientific benefit (Birney et al., 2017; Stark et al., 2019). This theme underscores autonomy as an ethical principle that must evolve alongside technological complexity.

Theme 4: Mitigation Frameworks and Critical Success Factors for Ethical AI Implementation

Beyond identifying ethical risks, the literature converges on a set of Critical Success Factors (CSFs) necessary for mitigating bias and ensuring responsible AI deployment in genetic diagnosis. To address reviewer concerns, these CSFs are operationalized below across policy, clinical, and technical domains.

Regulatory and Accountability Frameworks (Policy Level)

Studies emphasize the need for enforceable governance structures that define liability, oversight, and ethical compliance for AI systems in genomics. This includes regulatory guidance on dataset diversity, algorithmic auditing, and post-market surveillance to ensure ongoing fairness and safety (Wiens et al., 2019).

Transparency and Explainability Mandates (Clinical and Technical Levels)

Operationalization involves integrating XAI tools into clinical workflows, enabling clinicians to interrogate AI outputs and communicate reasoning to patients. Technical documentation and model reporting standards further support institutional accountability (Chen et al., 2021).

Equity-Centered Lifecycle Management (Technical Level)

Ethical AI must incorporate fairness checks at every stage of the AI lifecycle—data collection, model development, validation, deployment, and continuous monitoring. This includes routine bias audits and the inclusion of underrepresented populations in training datasets.

Multi-Stakeholder Engagement (Policy and Clinical Levels)

Effective implementation requires the participation of clinicians, geneticists, ethicists, patients, and policymakers. Stakeholder engagement ensures that ethical values such as justice, autonomy, and trust are embedded into system design rather than treated as external constraints.




Fig 3.1 Thematic Synthetic

Discussion

The findings of this review indicate that algorithmic bias in AI-driven genetic diagnosis is not merely a technical limitation but reflects historical inequities embedded in data generation practices and governance structures. While several included studies documented reduced diagnostic performance in underrepresented populations, the review also reveals a notable gap between ethical recognition and operational response. In particular, although justice and equity are frequently invoked as guiding principles, few studies provide concrete mechanisms for ensuring consistent algorithmic performance across diverse populations.

A critical implication of this gap is that current AI-driven diagnostic systems risk reproducing structural injustice, especially in low- and middle-income settings where genomic data infrastructure remains limited. The review data support the argument that inclusive dataset development requires global mechanisms. While federated learning architectures

collaboration and sustained community engagement, yet only a minority of studies explicitly address how such collaboration can be institutionalized or funded (Stark et al., 2019; Fatumo et al., 2022). This highlights an implementation gap between ethical intent and practical feasibility.

Trust emerged as a cross-cutting concern linking bias, accountability and clinical adoption. Although explainability and transparency are widely promoted as ethical imperatives, the review demonstrates that their application remains uneven and often conceptual rather than actionable. Few studies move beyond high-level advocacy to specify how explainable AI tools should be integrated into clinical decision-making or communicated to patients in ways that meaningfully support informed

consent (Char et al., 2018; Hooker & Masters, 2025). This limitation constrains both clinician accountability and patient autonomy.

Privacy-related challenges further underscore the need for practically enforceable governance mechanisms. While federated learning architectures

and dynamic consent models are frequently cited as consent models, federated learning, and ethical safeguards, the review identifies limited community-participatory governance remains empirical evaluation of their long-term limited and warrants systematic evaluation. effectiveness in clinical genomics. This suggests Strengthening interdisciplinary collaboration across that governance innovations, though promising, genomics, AI, ethics, and health policy is essential remain under-validated in real-world diagnostic to translate ethical principles into practice. environments (Birney et al., 2017).

Taken together, these findings suggest that the ethical advancement of AI-driven genetic diagnosis depends less on the articulation of new principles and more on the translation of existing Critical Success Factors into policy instruments, clinical protocols and technical standards. Addressing this translational gap represents a key priority for future research, regulatory development, and interdisciplinary collaboration.

5. Conclusion

AI-driven genetic diagnosis holds considerable promise for advancing precision medicine; however, this review demonstrates that algorithmic bias remains a substantive ethical challenge that may undermine fairness and rust if left unaddressed. Across the included studies, four recurring Critical Success Factors were identified as central to ethical implementation: regulatory and accountability mechanisms, meaningful stakeholder engagement, transparency and explainability and equity-oriented lifecycle management. Progress depends on improving data representativeness, enhancing interpretability and embedding oversight.

6. Future Research Directives

This review identifies persistent gaps in the empirical evaluation of bias mitigation strategies in AI-driven genetic diagnosis, particularly across underrepresented populations. Future research should focus on assessing the real-world effectiveness of fairness-enhancing methods, including inclusive data practices and explainable AI tools, within clinical genomic settings. Additionally, further studies are needed to operationalize accountability and transparency frameworks into implementable audit and governance mechanisms. Research on dynamic

REFERENCES

- Agarwal R, et al. (2023). Addressing algorithmic bias and the perpetuation of health inequities: An AI bias aware framework. *Health Policy and Technology*, 2023; 12 1.
- Birney, E., Vamathevan, J., & Goodhand, P. (2017). *Genomics in healthcare: GA4GH looks to 2022*. Genomics. <https://doi.org/10.1101/203554>
- Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing Machine Learning in Health Care Addressing Ethical Challenges. *New England Journal of Medicine*, 378(11), 981–983. <https://doi.org/10.1056/NEJMp1714229>
- Chen, Y., Ashizawa, N., Yeo, C. K., Yanai, N., & Yean, S. (2021). Multi-Scale Self-Organizing Map assisted Deep Autoencoding Gaussian Mixture Model for Unsupervised Intrusion Detection. *Knowledge-Based Systems*, 224, 107086. <https://doi.org/10.1016/j.knosys.2021.107086>
- Fatumo, S., Chikowore, T., Choudhury, A., Ayub, M., Martin, A. R., & Kuchenbaecker, K. (2022). A Roadmap to Increase diversity in Genomic Studies. *Nature Medicine*, 28(2), 243–250. <https://doi.org/10.1038/s41591-021-01672-4>
- Kullar R, et al. (2020). Racial disparity of coronavirus disease 2019 in African American communities. *J Infect Dis*. 2020;222(6):890–3. PubMed. Google Scholar
- Martin, A. R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B. M., & Daly, M. J. (2019). Clinical use of Current Polygenic Risk Scores may exacerbate Health Disparities. *Nature Genetics*, 51(4), 584–591. <https://doi.org/10.1038/s41588-019-0379-x>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting Racial

- Bias in an Algorithm used to Manage the Health of Populations. *Science*, 366(6464), 447–453.
<https://doi.org/10.1126/science.aax2342>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, n71.
<https://doi.org/10.1136/bmj.n71>
- Popejoy, A. B., & Fullerton, S. M. (2016). Genomics is failing on diversity. *Nature*, 538(7624), 161–164.
<https://doi.org/10.1038/538161a>
- Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of Internal Medicine*, 169(12), 866–872.
<https://doi.org/10.7326/M18-1990>
- Rehm, H. L., Berg, J. S., Brooks, L. D., Bustamante, C. D., Evans, J. P., Landrum, M. J., Ledbetter, D. H., Maglott, D. R., Martin, C. L., Nussbaum, R. L., Plon, S. E., Ramos, E. M., Sherry, S. T., & Watson, M. S. (2015). ClinGen. The Clinical Genome Resource. *New England Journal of Medicine*, 372(23), 2235–2242.
<https://doi.org/10.1056/NEJMSr1406261>
- Li T, Iacobelli F. Purposeful AI. In: Companion publication of the 2023 conference on computer supported cooperative work and social computing. Minneapolis, MN, USA: Association for Computing Machinery; 2023. p. 563–5. Google Scholar
- Stark, Z., Dolman, L., Manolio, T. A., Ozenberger, B., Hill, S. L., Caulfield, M. J., Levy, Y., Glazer, D., Wilson, J., Lawler, M., Boughtwood, T., Braithwaite, J., Goodhand, P., Birney, E., & North, K. N. (2019). Integrating Genomics into Healthcare: A Global Responsibility. *The American Journal of Human Genetics*, 104(1), 13–20.
<https://doi.org/10.1016/j.ajhg.2018.11.014>
- Matovu, E., Bucheton, B., Chisi, J., Enyaru, J., Hertz-Fowler, C., Koffi, M., Macleod, A., Mumba, D., Sidibe, I., Simo, G., Simuunza, M., Mayosi, B., Ramesar, R., Mulder, N., Ogendo, S., Mocumbi, A. O., Hugo-Hamman, C., Ogah, O., Rotimi, C. (2014). Enabling the Genomic Revolution in Africa. *Science*, 344(6190), 1346–1348.
<https://doi.org/10.1126/science.1251546>
- Topol, E. J. (2019). High-Performance Medicine: The Convergence of Human and Artificial Intelligence. *Nature Medicine*, 25(1), 44–56.
<https://doi.org/10.1038/s41591-018-0300-7>
- Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P. N., Thadaneysrani, S., & Goldenberg, A. (2019). Do no harm: A Roadmap for Responsible Machine Learning for Health care. *Nature Medicine*, 25(9), 1337. 1340.
<https://doi.org/10.1038/s41591-019-0548-6>