Sociodemographic Biases in Medical Decision-Making

Arieh Riskin MD PhD[1,2] and Ayala Gover MD[1,2]

[1]Department of Neonatology, Bnai Zion Medical Center, Haifa, Israel
[2]Rappaport Faculty of Medicine, Technion–Israel Institute of Technology, Haifa, Israel

ABSTRACT Sociodemographic biases, which are related to gender, race and ethnicity, socioeconomic status, age, weight, and sexual orientation, permeate medical decision-making and contribute to health disparities that affect diagnosis, treatment allocation, and patient outcomes. In this narrative review, we address sociodemographic bias in general, treating implicit bias as one of several mechanisms. Drawing on peer-reviewed studies, systematic reviews, and meta-analyses published mainly between 2010 and 2025, we examined evidence for these biases, their mechanisms, bias in artificial intelligence (AI) and large language model (LLM) clinical tools, and interventions for bias reduction. We integrated literature usually reviewed separately and distinguished evidence of outcome disparities from evidence that directly measures clinician bias. Robust evidence documents gender disparities in cardiovascular care and pain management, racial and ethnic disparities in procedural access and analgesia, and socioeconomic gradients in treatment intensity. Experimental studies that hold the clinical presentation constant confirm that these patient characteristics can influence clinicians' decisions. LLMs replicate and sometimes amplify these biases. Up to 91.7% of reviewed studies identified biases in medical AI, with LGBTQ+ patients receiving mental health referrals six to seven times more frequently than clinically warranted. Mechanisms operate at the clinician level (prototype-based reasoning, implicit and explicit attitudes, cognitive load), the knowledge level (non-representative research and training data), and the system level (structural inequity). Multilevel interventions, such as individual bias awareness, structural protocols, and algorithmic fairness techniques, show promise but require sustained institutional commitment. Equity in healthcare is not only an ethical imperative but a clinical and scientific one.

IMAJ 2026; 28: 709–716

KEY WORDS: gender bias, health disparities, implicit bias, racial bias, socioeconomic status

Medical decision-making does not occur in a social vacuum. Every clinical encounter is shaped not only by the patient's physiology and the clinician's knowledge, but also by complex social dynamics, including the patient's gender, race, socioeconomic status (SES), age, and sexual orientation. The systematic, often unconscious influence of these characteristics on clinical judgment constitutes sociodemographic bias: a deviation from evidence-based care that correlates with patient identity rather than clinical need [1,2].

The consequences are well documented. In the United States, women with chest pain wait longer for cardiac workup and receive less aggressive treatment than men [3]. When primary care physicians in the United States assessed actor-portrayed patients with identical scripted presentations, Black patients and women were less likely to be referred for cardiac catheterization, and Black women least of all [4]. Lower SES is associated with poorer access to care, delayed diagnosis, and less intensive treatment [5]. These are not isolated anecdotes. They represent patterns replicated across specialties and care settings, although most of the evidence comes from the United States.

Such differences in care occur against a background of broader, systemic health inequities. In Israel, despite universal health insurance, life expectancy is lower among Arab than among Jewish citizens, age-adjusted mortality from heart disease, stroke, and diabetes is higher among Arabs. Inequalities are also evident by socioeconomic position and in subpopulations such as Bedouin Arabs, ultra-Orthodox Jews, and immigrants from the former Soviet Union [6]. In the United States, pregnancy-related mortality among Black women was approximately three times that of White women in 2007–2016. Black women with at least some college education had higher pregnancy-related mortality than women of other racial groups without a high school diploma [7]. Historical undertreatment has also left its mark. False beliefs about biological differences between Black and White people, long embedded in medicine, are still endorsed by some medical trainees and are associated

with lower ratings of Black patients' pain and less accurate treatment recommendations [8]. These outcome gaps have many causes, including social determinants of health, access to care, and health behaviors, and cannot be attributed to clinician bias alone. They form the background against which the clinical decisions reviewed here are made.

In this review, we examine the evidence base for sociodemographic biases in clinical decision-making, the cognitive and systemic mechanisms that sustain them, their reproduction in artificial intelligence (AI) clinical tools, and evidence-based strategies for reduction. Its scope is sociodemographic bias in general; that is, any systematic, clinically unwarranted influence of a patient's group membership on clinical decisions. Implicit bias is treated as one important mechanism among several rather than as the sole explanation. We aimed to add to previous reviews, which have typically addressed a single patient characteristic, a single mechanism such as implicit bias, or bias in AI alone, in three ways: by bringing this literature together, by distinguishing throughout between evidence of outcome disparities and evidence that directly measures clinician bias, and by linking each level of mechanism to a corresponding level of intervention. Understanding these biases is a prerequisite for equitable practice.

SOCIODEMOGRAPHIC BIASES IN CLINICAL DECISION-MAKING, WHICH ARE RELATED TO GENDER, RACE, SOCIOECONOMIC STATUS, AGE, WEIGHT, AND SEXUAL ORIENTATION, ARE DOCUMENTED ACROSS MULTIPLE SPECIALTIES AND CONTRIBUTE TO MEASURABLE HEALTH INEQUITIES.

METHODS

We conducted a narrative review. We searched PubMed/MEDLINE and Google Scholar for English-language articles, with an emphasis on empirical studies published between 2010 and 2025, using combinations of terms for bias (bias, implicit bias, stereotyping, disparities, discrimination), patient characteristics (sex, gender, race, ethnicity, SES, age, weight, obesity, sexual orientation, LGBTQ+), clinical decisions (diagnosis, treatment, referral, pain management), and artificial intelligence (machine learning, large language models [LLMs]). The reference lists of key reviews were also screened. Priority was given to systematic reviews and meta-analyses, experimental studies, and large observational studies. Earlier landmark studies were included where they remained the primary evidence. Because reviews in the *Israel Medical Association Journal (IMAJ)* are limited to 30 references, systematic reviews were cited in preference to individual primary studies where possible. The country in which each primary study was conducted is indicated in the text.

Two types of evidence are distinguished throughout this review. Disparity studies document differences in care or outcomes between groups in real-world data. Such differences may reflect clinician bias, but also differences in clinical presentation, symptom reporting, patient preferences, access, or co-morbidity, which statistical adjustment can reduce but not exclude. Bias-measurement studies isolate the clinician's contribution, either experimentally by presenting clinically identical cases that differ only in the patient's sociodemographic characteristics using vignettes or standardized patients or by measuring clinicians' implicit or explicit attitudes and relating them to their decisions. Only the second type can attribute a difference in care to clinician bias with confidence.

CONCEPTUAL FRAMEWORK: CLINICIAN-LEVEL MECHANISMS

Sociodemographic bias can arise at three nested levels: the individual clinician, whose perception and judgment are shaped by the patient's characteristics; the knowledge base, which is the evidence, clinical prototypes, and algorithms on which clinicians rely; and the healthcare system, which determines access to care and the conditions under which decisions are made. These are interacting layers rather than alternative explanations. This section addresses the clinician level, at which implicit bias operates. The knowledge and system levels are addressed under the section on mechanisms sustaining sociodemographic bias later.

The dual-process account of clinical cognition includes System 1, which includes fast and automatic processing and System 2, which is slow, deliberate analysis. These systems are explained in detail by Flugelman and Riskin [9] in this issue of *IMAJ*. What matters for the present argument is a single feature of that architecture: System 1 pattern recognition is the point at which demographic schemas are activated. When a clinician forms an initial impression of a patient, demographic characteristics activate pre-existing mental prototypes of who typically presents with which condition, often shaped by a medical literature historically derived from non-representative samples [10]. Prototype matching then leads to diagnostic anchoring. Symptoms that do not fit the expected demographic template are underweighted or reframed. The anchoring bias itself is defined and illustrated in the review of cognitive biases and heuristics in this issue of *IMAJ* [11].

Bias can be explicit (conscious, self-reported attitudes and beliefs) or implicit (automatic associations that may operate outside awareness). The Implicit Association Test (IAT) has been widely used to measure implicit bias in healthcare providers. Systematic reviews of studies conducted predominantly in the United States found low-to-moderate pro-White implicit preferences among most providers, at levels similar to those of the general population [12,13]. Implicit biases related to gender, age, and weight have also been reported [2]. Evidence that these implicit attitudes translate into clinical decisions is, however, mixed. In a systematic review of 37 studies, six of 14 studies using vignettes or simulated patients found that higher implicit bias was associated with disparities in treatment recommendations, pain management, or empathy, whereas eight found no significant association. Associations were more consistent for patient–provider communication [13].

The IAT itself has been criticized on methodological grounds. A meta-analysis of criterion studies found that IAT scores were poor predictors of discriminatory behavior and performed no better than simple explicit measures [14]. A review of its construct validity found little evidence that it measures a distinct implicit construct, estimating that only about one fifth of the variance in the race IAT reflects racial preferences [15]. In this review, therefore, IAT findings are taken as evidence that biased associations are common among clinicians, not as proof that they explain any particular disparity. The strongest evidence of bias in clinical decisions comes from experimental studies in which only the patient's characteristics are varied.

LARGE LANGUAGE MODELS AND CLINICAL AI TOOLS REPLICATE AND SOMETIMES AMPLIFY HUMAN SOCIODEMOGRAPHIC BIASES, PARTICULARLY AGAINST LGBTQ+ PATIENTS AND RACIAL MINORITIES, UNDERSCORING THE NEED FOR ALGORITHMIC FAIRNESS STANDARDS IN HEALTHCARE AI.

Cognitive load may amplify these effects. Under time pressure, fatigue, or high patient volume (e.g., conditions characteristic of emergency departments and understaffed wards) clinicians rely more heavily on System 1 heuristics, which favor prototype-based reasoning that penalizes atypical presentations. In a U.S. study of pediatric emergency medicine residents, implicit racial bias scores increased after shifts with greater emergency department overcrowding or a load of more than 10 patients [16], although whether such increases alter clinical decisions has not been tested directly. This finding suggests that systemic factors such as staffing and workflow design may themselves amplify bias.

GENDER BIAS IN CLINICAL DECISION-MAKING

Gender bias is among the most extensively documented forms of sociodemographic bias in medicine. Disparity studies across multiple specialties show that women receive less intensive workup and treatment than men with similar presentations, and vignette studies, reviewed systematically by FitzGerald and Hurst [2], indicated that patient gender can itself influence clinicians' diagnostic and treatment decisions.

CARDIOVASCULAR DISEASE

Cardiovascular disease remains one of the leading causes of death in women worldwide, yet women are systematically underdiagnosed and undertreated. A U.S. national study of emergency department visits found that young women aged 18–55 years presenting with chest pain waited 29% longer for physician evaluation than men (48.1 vs. 37.2 minutes) and were less likely to be admitted, independent of clinical features [3,17]. Women with acute myocardial infarction were more likely to be attributed a non-cardiac etiology for their symptoms, delaying reperfusion.

A systematic review of 19 studies by Al Hamid and colleagues [18] documented that women are offered fewer diagnostic tests, including coronary angiography; are referred less often to cardiologists; and are prescribed fewer cardiovascular medicines than men. A meta-analysis of 19 observational studies (241,613 patients) found that women were less likely than men to be referred to cardiac rehabilitation (39.6% vs. 49.4%; odds ratio 0.68) [19]. These are disparity data. Although the gaps persist after adjustment for clinical severity, they may also reflect milder or atypical symptoms and women's own lower perception of their cardiovascular risk [18]. Experimental evidence isolating the clinician's decision comes from a U.S. study by Schulman and colleagues [4], in which physicians were less likely to recommend cardiac catheterization for women than for men with identical scripted presentations.

Part of the mechanism is prototype mismatch. The prototypical myocardial infarction patient in medical education has historically been male, leading to under-recognition of atypical presentations, which are disproportionately seen in women, including jaw pain, fatigue, nausea, and dyspnea without classic chest pressure [3,18].

PAIN MANAGEMENT

Pain management is another domain in which gender disparities have been rigorously documented. In emergency

department records of 21,851 patients from Israel and the United States, women were less likely than men to be prescribed analgesics (38% vs. 47%), a gap that persisted after adjustment for recorded pain scores and numerous patient, physician, and department variables. This disparity was observed among both male and female physicians. Women's pain scores were also less often recorded, and women spent longer in the emergency department [20]. Because such observational data cannot exclude the possibility that women and men describe or report pain differently, the same investigators added a controlled experiment, in which nurses who read identical clinical vignettes judged the pain of female patients to be less intense than that of male patients [20]. The combination of a consistent real-world disparity with an experiment that holds the patient's report constant makes clinician bias a likely contributor.

MULTILEVEL INTERVENTIONS COMBINING INDIVIDUAL BIAS LITERACY TRAINING, STANDARDIZED CLINICAL PROTOCOLS, AND EQUITABLE AI DEVELOPMENT SHOW THE GREATEST PROMISE FOR REDUCING DISPARITIES, BUT REQUIRE INSTITUTIONAL COMMITMENT AND ONGOING OUTCOME MONITORING.

The proposed mechanisms are numerous. Clinicians may perceive women's pain reports as emotionally inflated or rely on stereotyped assumptions about pain tolerance by gender, while differences in how men and women express and report pain may also be a contributing factor, which is why designs that hold the report constant are needed to separate these explanations.

RACIAL, ETHNIC, AND SOCIOECONOMIC BIASES

In the United States, racial and ethnic disparities in medical care were documented comprehensively in the landmark Institute of Medicine report *Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care* (2003). Studies from the last decade, conducted mostly in the United States, consistently found implicit pro-White bias among healthcare providers across disciplines and specialties, but its association with treatment decisions is inconsistent [12,13]. The evidence on racial bias in clinical decisions therefore rests mainly on disparity studies and on experimental and attitude studies, which are considered in turn below.

Disparity studies document racial gaps in treatment. In a U.S. national sample of children with appendicitis, Black children were less likely than White children to receive opioid analgesia for severe pain, after adjustment for pain score [21]. Experimental and attitude studies point to the clinician's contribution. In the U.S. study by Schulman and colleagues [4], physicians were less likely to refer Black than White patients with identical scripted presentations for cardiac catheterization [4]. In a U.S. study of White medical students and residents, half endorsed false beliefs about biological differences between Black and White people (for example, that Black people's skin is thicker), and those who endorsed these beliefs rated a Black patient's pain as lower and made less accurate treatment recommendations [8]. This finding is an example of explicit, measurable beliefs, rather than implicit associations, shaping clinical judgment.

SES operates through overlapping but partially distinct mechanisms. In a U.S.-focused review, McMaughan and colleagues [5] documented that lower SES predicted delayed diagnosis, less intensive treatment, and lower adherence. This lower adherence was often a consequence of cost barriers and healthcare access rather than patient preference. Experimental evidence comes from a U.S. national survey of 4603 senior medical students at 84 medical schools, who chose between two clinically equivalent management options for cardiac patient vignettes. Patients described as having the highest SES were more likely than those with the lowest SES to receive procedural recommendations (50.3% vs. 43.2%), showing that SES-based bias is present even before independent practice begins [22]. Notably, the same study found no overall racial or gender bias, although race–gender interactions and regional variation were observed [22], a reminder that bias is not uniform across settings.

AGE, WEIGHT, SEXUAL ORIENTATION, AND INTERSECTING IDENTITIES

Ageism in medicine manifests as under-treatment of older patients by assuming that aggressive workup is futile or unwanted and as dismissal of symptoms attributed to normal aging. Older patients are also excluded from clinical trials at rates disproportionate to their disease burden. In an analysis of trials registered at ClinicalTrials.gov from 2010 to 2021, an upper age limit was set in 90% of phase 1 type 2 diabetes trials and 51% of phase 1 cardiovascular trials [23]. As a result, the evidence base is thinnest for the patients who bear much of the disease burden. Evidence on age bias in individual clinical decisions, as distinct from such structural exclusion, is comparatively sparse.

Weight bias and discrimination within healthcare represent a pervasive challenge with serious consequences for patient outcomes. A systematic review and meta-analysis of 41 studies found that physicians, nurses, dietitians,

and other health professionals hold implicit and explicit weight-biased attitudes, although the overall quality of the evidence was rated as very low [24]. This stigma can compromise the quality of care patients receive and erode therapeutic trust. These are, however, attitude data. Studies testing whether such attitudes change clinical decisions are few, and this remains an important gap in the evidence. Addressing weight bias among clinical personnel is nonetheless essential for equitable care.

LGBTQ+ is an umbrella term describing people of diverse sexual orientations and gender identities. LGBTQ+ patients experience distinctive barriers. Providers may lack knowledge of relevant health risks (e.g., higher rates of depression, specific sexually transmitted infections, elevated cardiovascular risk in transgender individuals), make erroneous assumptions about sexual behavior, or apply heteronormative diagnostic frameworks. In a U.S. survey, non-binary patients reported feeling less respected by providers and being misgendered by hospital staff [25]. A narrative review of 107 studies documented stigmatization in healthcare settings, barriers to accessing care, and higher rates of mental ill health and unmet mental health needs in LGBTQ+ communities [26]. These issues are mainly disparity and patient-experience data, in which clinician bias is one contributing factor among several.

Intersectionality, the compounding of multiple marginalized identities, amplifies disparities. In an experimental study by Schulman and colleagues [4], Black women were the least likely of all groups to be referred for cardiac catheterization. An older Black woman from a low SES background experiences the compounded, not merely additive, effect of gender, racial, age, and socioeconomic biases. Intersectional analysis is therefore essential for understanding the full scope of health inequities.

SOCIODEMOGRAPHIC BIASES IN CLINICAL ARTIFICIAL INTELLIGENCE AND LLMS

The rapid integration of AI into clinical practice has introduced a new vector for sociodemographic bias. Machine learning models trained on electronic health records and clinical datasets inherit the biases embedded in those datasets, which reflect decades of inequitable care. The general mechanisms by which this occurs, including representation bias in training data, biased outcome labels and proxy targets, and the amplification of inequity through feedback loops, are set out in the review of the clinical risks of artificial intelligence in this issue of *IMAJ* [27].

A landmark study in *Nature Medicine* [28] evaluated nine LLMs on 1000 emergency department cases (500 real and 500 synthetic), each presented in 32 versions that differed only in the patient's sociodemographic characteristics. The models demonstrated consistent sociodemographic bias: certain LGBTQ+ subgroups were recommended mental health assessments approximately six to seven times more often than clinically indicated, cases labeled as Black or unhoused were more often directed toward urgent care or invasive interventions, and cases labeled as high-income received more recommendations for advanced imaging. The authors noted that these recommendations were not random noise but systematic, reproducible patterns, suggesting that bias is encoded in model weights rather than resulting from stochastic variation. Notably, some LLMs exhibited biases exceeding typical human bias levels, raising the concern that AI-augmented decision support could amplify rather than mitigate disparities. Because the clinical content is held constant, such studies provide unusually direct bias-measurement evidence.

A systematic review by Omar and colleagues [29] found that 91.7% of 24 studies examining LLMs in medical contexts identified biases, with gender bias present in 93.7% of those studies. Bias against LGBTQ+ patients was particularly severe and consistent across models and studies, and racial or ethnic biases were observed in 90.9% of studies [29].

MECHANISMS SUSTAINING SOCIODEMOGRAPHIC BIAS

Beyond clinician-level mechanisms such as prototype activation, implicit and explicit attitudes and beliefs, and cognitive load, sociodemographic bias is sustained at two other levels. These levels are not alternatives to clinician-level bias but nested contributors to it: knowledge- and training-level factors determine which prototypes and attitudes clinicians acquire, and system-level factors determine access to care and the conditions under which those prototypes are applied.

Knowledge level: non-representative prototypes

Diagnostic criteria and clinical prototypes were historically derived from studies predominantly enrolling White male patients. Presentations by women, racial minorities, and older adults frequently deviate from these prototypes, leading to under-recognition [3,10].

Knowledge level: underrepresentation in research

Women represent approximately one-third of cardiovascular trial participants despite bearing an equivalent dis-

ease burden [18]. Minority groups are similarly underrepresented. This deficit generates systematic knowledge gaps that propagate into clinical guidelines [18,23].

Training level: the hidden curriculum

Informal socialization during medical training can reinforce stereotypes. Observational learning from senior clinicians who exhibit biased behavior normalizes it for trainees, creating intergenerational transmission of bias. In a prospective study of 3547 students at 49 U.S. medical schools, having heard negative comments about Black patients from attending physicians or residents was associated with an increase in implicit racial bias over the course of medical school [30].

System level: structural inequity

Language barriers, lack of culturally competent interpreters, and geographic concentration of specialty care in affluent areas create access inequities that interact with clinical bias to compound disparities [1]. The population-level inequities in chronic disease and maternal mortality described in the Introduction [6,7] illustrate how such structural factors shape outcomes independently of any individual clinical decision.

INTERVENTIONS FOR BIAS REDUCTION

Consistent with the division of debiasing content across this issue of *IMAJ*, the strategies described later address the mechanism specific to sociodemographic bias: prototype override and counterfactual substitution.

INDIVIDUAL LEVEL: BIAS LITERACY AND AWARENESS

Implicit bias training, helping clinicians recognize their own biases through IAT feedback, case-based reflection, and structured discussion, increases awareness but has modest effects on clinical behavior change alone [1]. The most effective programs combine awareness with skill building: teaching perspective taking, counterfactual reasoning ("What would I recommend if this patient were male, or White, or higher income?"), and structured communication.

SYSTEM LEVEL: PROTOCOLS AND STANDARDIZATION

Algorithmic decision aids, standardized triage protocols, and checklists that require explicit documentation of reasoning reduce reliance on intuitive judgment and create accountability structures [11]. When diagnostic criteria are applied uniformly regardless of demographic characteristics, as in structured ST-elevation myocardial infarction protocols, gender and racial gaps in reperfusion times narrow substantially. Because cognitive load may amplify bias, staffing and workflow designs that reduce time pressure are equally part of the system-level response. Diverse hiring practices and representative leadership structures also reduce organizational bias [1].

ALGORITHMIC FAIRNESS IN AI

Technical approaches to reducing AI bias include counterfactual data augmentation (generating training examples with reversed demographic attributes to test for differential model behavior), fairness-constrained model training, disaggregated evaluation by subgroup, and routine bias audits before clinical deployment [27,29]. Regulatory frameworks requiring bias testing as a condition of clinical AI certification are under development in multiple jurisdictions.

HEALTH EQUITY AS A SYSTEMIC GOAL

Ultimately, bias reduction requires embedding health equity as an institutional value through leadership commitment, equity metrics in performance evaluation, dedicated equity offices, and community partnerships with underserved populations [1]. Incremental individual interventions without structural change have limited sustained impact.

DISCUSSION

The cumulative evidence we reviewed demonstrates that sociodemographic biases in medical decision making are pervasive, measurable, and consequential. They operate through both conscious and unconscious mechanisms, are amplified by systemic pressures, and are now being replicated in AI clinical tools. The cardiovascular case illustrates the stakes: delays in reperfusion due to gender or racial bias translate directly into myocardial loss, functional impairment, and excess mortality.

The strength of this evidence differs by type. Most of the literature consists of disparity studies, which establish that groups receive different care but cannot by themselves establish why. Evidence that clinicians' decisions are themselves influenced by patient characteristics comes from a smaller body of experimental studies, in which clinically identical cases differ only in sociodemographic attributes [4,20,22], and from studies relating clinicians' measured beliefs or attitudes to their decisions [8,13]. This experimental evidence is reasonably robust for gender, race, and socioeconomic status, whereas for age, weight, and sexual orientation the evidence consists

mainly of attitude and disparity data. The specific contribution of implicit bias is less certain than is often assumed, given the inconsistent association between IAT scores and treatment decisions [13] and the methodological critiques of the IAT [14,15]. LLM studies, by contrast, provide unusually clean bias-measurement evidence, because clinical content can be held perfectly constant [28].

By considering this literature together, across patient characteristics, across human and algorithmic decision makers, and across levels of mechanism, we identified several gaps. Few studies linked measured clinician bias to real-world decisions and outcomes. Evidence on age, weight, and sexual-orientation bias in clinical decisions is sparse. Intervention studies usually measured changes in attitudes rather than in clinical decisions and most data came from the United States, with relatively few studies from Israel and other countries [6,20]. Future research should combine experimental designs with real-world outcome data and test whether interventions change clinical decisions and patient outcomes, not only attitudes.

Several caveats warrant acknowledgment. Most bias research relies on vignette or simulation methodologies, which may not fully capture real-world decision complexity. Most studies were also conducted in the United States, whose healthcare system, population composition, and history differ substantially from those of Israel and other countries, so that their findings may not generalize directly. Effect sizes are heterogeneous across studies, settings, and outcomes, and causal inference is challenging when observational data confound clinical need, patient preference, and provider bias.

Nevertheless, the convergence of evidence across methodologies, countries, and clinical contexts, including the replication of biases in AI systems, argues for urgent action. The emergence of LLMs as clinical decision support tools already being deployed in high-stakes settings makes the problem of algorithmic bias not a future concern but a present one.

Studies demonstrated that structural interventions, such as standardized protocols, diverse training data, algorithmic fairness constraints, and can reduce measurable gaps. The question is whether healthcare systems will commit the organizational will and resources to implement them at scale.

CONCLUSIONS

Sociodemographic biases in medical decision-making affect patients based on gender, race and ethnicity, SES, age, weight, and sexual orientation. These biases are documented across specialties, are demonstrated most convincingly by experimental studies that hold the clinical presentation constant; arise from clinician-level, knowledge-level, and system-level mechanisms of which implicit bias is only one; and are now reproduced in clinical AI tools. Multilevel interventions combining individual bias literacy, structural standardization, and algorithmic fairness are required. Commitment to health equity should be an integral part of clinical excellence.

Correspondence

Prof. A. Riskin
Dept. of Neonatology, Bnai Zion Medical Center, Haifa 31048, Israel
Phone: (972-4) 835-9063
Fax: (972-4) 835-9206
Email: arik.riskin@gmail.com

References

1. Vela MB, Erondu AI, Smith NA, et al. Eliminating explicit and implicit biases in health care: evidence and research needs. *Annu Rev Public Health* 2022; 43: 477-501.
2. FitzGerald C, Hurst S. Implicit bias in healthcare professionals: a systematic review. *BMC Med Ethics* 2017; 18: 19.
3. Mehta LS, Beckie TM, DeVon HA, et al. Acute myocardial infarction in women: a scientific statement from the American Heart Association. *Circulation* 2016; 133 (9): 916-47.
4. Schulman KA, Berlin JA, Harless W, et al. The effect of race and sex on physicians' recommendations for cardiac catheterization. *N Engl J Med* 1999; 340 (8): 618-26.
5. McMaughan DJ, Oloruntoba O, Smith ML. Socioeconomic status and access to healthcare: interrelated drivers for healthy aging. *Front Public Health* 2020; 8: 231.
6. Muhsen K, Green MS, Soskolne V, Neumark Y. Inequalities in non-communicable diseases between the major population groups in Israel: achievements and challenges. *Lancet* 2017; 389 (10088): 2531-41.
7. Petersen EE, Davis NL, Goodman D, et al. Racial/ethnic disparities in pregnancy-related deaths—United States, 2007–2016. *MMWR Morb Mortal Wkly Rep* 2019; 68 (35): 762-5.
8. Hoffman KM, Trawalter S, Axt JR, Oliver MN. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proc Natl Acad Sci U S A* 2016; 113 (16): 4296-301.
9. Flugelman MY, Riskin A. Clinical reasoning and medical decision-making: science, art, and challenges in the era of artificial intelligence. *IMAJ* 2026; 28 (11): 724-9.
10. van Ryn M, Burke J. The effect of patient race and socio-economic status on physicians' perceptions of patients. *Soc Sci Med* 2000; 50 (6): 813-28.
11. Mintz A, Riskin A. Cognitive biases and heuristics in medical decision-making: mechanisms, clinical manifestations, and debiasing strategies. *IMAJ* 2026; 28 (11): 690-700.
12. Hall WJ, Chapman MV, Lee KM, et al. Implicit racial/ethnic bias among health care professionals and its influence on health care outcomes: a systematic review. *Am J Public Health* 2015; 105 (12): e60-76.
13. Maina IW, Belton TD, Ginzberg S, Singh A, Johnson TJ. A decade of studying implicit racial/ethnic bias in healthcare providers using the implicit association test. *Soc Sci Med* 2018; 199: 219-29.

14. Oswald FL, Mitchell G, Blanton H, Jaccard J, Tetlock PE. Predicting ethnic and racial discrimination: a meta-analysis of IAT criterion studies. *J Pers Soc Psychol* 2013; 105 (2): 171-92.
15. Schimmack U. The Implicit Association Test: a method in search of a construct. *Perspect Psychol Sci* 2021; 16 (2): 396-414.
16. Johnson TJ, Hickey RW, Switzer GE, et al. The impact of cognitive stressors in the emergency department on physician implicit racial bias. *Acad Emerg Med* 2016; 23 (3): 297-305.
17. Banco D, Chang J, Talmor N, et al. Sex and race differences in the evaluation and treatment of young adults presenting to the emergency department with chest pain. *J Am Heart Assoc* 2022; 11 (10): e024199.
18. Al Hamid A, Beckett R, Wilson M, et al. Gender bias in diagnosis, prevention, and treatment of cardiovascular diseases: a systematic review. *Cureus* 2024; 16 (2): e54264.
19. Colella TJ, Gravely S, Marzolini S, et al. Sex bias in referral of women to outpatient cardiac rehabilitation? A meta-analysis. *Eur J Prev Cardiol* 2015; 22 (4): 423-41.
20. Guzikevits M, Gordon-Hecker T, Rekhtman D, et al. Sex bias in pain management decisions. *Proc Natl Acad Sci U S A* 2024; 121 (33): e2401331121.
21. Goyal MK, Kuppermann N, Cleary SD, Teach SJ, Chamberlain JM. Racial disparities in pain management of children with appendicitis in emergency departments. *JAMA Pediatr* 2015; 169 (11): 996-1002.
22. Williams RL, Romney C, Kano M, et al. Racial, gender, and socioeconomic status bias in senior medical student clinical decision-making: a national survey. *J Gen Intern Med* 2015; 30 (6): 758-67.
23. Nguyen D, Mika G, Ninh A. Age-based exclusions in clinical trials: a review and new perspectives. *Contemp Clin Trials* 2022; 114: 106683.
24. Lawrence BJ, Kerr D, Pollard CM, et al. Weight bias among health care professionals: a systematic review and meta-analysis. *Obesity (Silver Spring)* 2021; 29 (11): 1802-12.
25. Nowaskie DZ, Menez O. Healthcare experiences of LGBTQ+ people: non-binary people remain unaffirmed. *Front Sociol* 2024; 9: 1448821.
26. Silveri G, Schimmenti S, Prina E, et al. Barriers in care pathways and unmet mental health needs in LGBTIQ+ communities. *Int Rev Psychiatry* 2022; 34 (3-4): 215-29.
27. Glebov M, Mintz A. Clinical risks of artificial intelligence in medicine. *IMAJ* 2026; 28 (11) 717-23.
28. Omar M, Soffer S, Agbareia R, et al. Sociodemographic biases in medical decision making by large language models. *Nat Med* 2025; 31 (6): 1873-81.
29. Omar M, Sorin V, Agbareia R, et al. Evaluating and addressing demographic disparities in medical large language models: a systematic review. *Int J Equity Health* 2025; 24 (1): 57.
30. van Ryn M, Hardeman R, Phelan SM, et al. Medical school experiences associated with change in implicit racial bias among 3547 students: a Medical Student CHANGES study report. *J Gen Intern Med* 2015; 30 (12): 1748-56.

Man is least himself when he talks in his own person. Give him a mask, and he will tell you the truth.

Oscar Wilde (1854–1900), Irish author, playwright, and poet

Forget not that the earth delights to feel your bare feet and the winds long to play with your hair.

Kahlil Gibran (1883–1931), Lebanese-American poet and artist

Capsule

Epitope-diverse autoantibodies induce neutrophil–epithelial niches driving kidney injury in ANCA-associated glomerulonephritis

Nishide and co-authors investigated how antineutrophil cytoplasmic autoantibodies (ANCA) initiate kidney inflammation and tissue injury in autoimmune glomerulonephritis. The study focused on antibodies directed against myeloperoxidase (MPO), a major antigen associated with ANCA-mediated vasculitis. Researchers examined kidney biopsies from eight newly diagnosed, untreated patients with MPO-ANCA-associated glomerulonephritis and four healthy kidney donors. Single-cell spatial transcriptomics identified distinct inflammatory microenvironments containing neutrophils, injured tubular epithelial cells, lymphocytes, and TREM2-positive macrophages. Injured epithelial cells expressed inflammatory chemokines, including CXCL6 and CXCL16, associated with local immune-cell recruitment. Neutrophils within these inflammatory niches displayed increased type II interferon signaling and Fcγ receptor-associated activation pathways. Experiments demonstrated that combinations of anti-MPO antibodies recognizing different antigenic epitopes enhanced neutrophil activation. In mouse models, these antibody combinations induced crescentic glomerulonephritis, reproducing important features of human autoimmune kidney injury. Engineered one-armed anti-MPO antibodies, designed to reduce Fcγ receptor engagement, attenuated neutrophil activation and kidney injury in experimental models. Scientific significance: This study provides a mechanistic connection between autoantibody recognition, localized neutrophil activation, and renal tissue damage. It suggests that selectively blocking pathogenic antibody-mediated neutrophil activation could offer a more targeted treatment approach, although clinical effectiveness in humans remains to be established.

J Exp Med 2026; 223 (11): e20260506

Eitan Israeli