Yemen Journal of Medicine

https://yemenjmed.com

Article information

DOI: 10.63475/yjm.v5i2.0455

Article history:
Received: 28 June 2026
Accepted: 05 July 2026
Published: 25 August 2026

Correspondence to:
Elmukhtar Habas
Email: Habas1962@gmail.com
ORCID: 0000-0002-7754-8993

How to cite this article
Habas E, Habas A, Rayani A, Habas A. Artificial intelligence and machine learning in kidney disease, dialysis, and transplantation: Opportunities, validation gaps, and equity challenges for Arab and African populations–Narrative review. Yemen J Med. Published online August 25, 2026. doi: 10.63475/yjm.v5i2.0455:1-7.

Review Article

Artificial Intelligence and Machine Learning in Kidney Disease, Dialysis, and Transplantation: Opportunities, Validation Gaps, and Equity Challenges for Arab and African Populations–Narrative Review

Elmukhtar Habas[1], Ala Habas[2], Amnna Rayani[3], Aml Habas[4]

1 Professor/Senior Consultant, Department of Medicine, Hamad General Hospital, Qatar University, Doha, Qatar
2 Resident, Tripoli Central Hospital, University of Tripoli, Tripoli, Libya
3 Professor/Consultant, Department of Hematology, University of Tripoli, Tripoli, Libya
4 Specialist, Tripoli Pediatric Hospital, The Open Libyan University, Tripoli, Libya

ABSTRACT

Artificial intelligence (AI)—including machine learning, deep learning, and large language models—is being rapidly adopted across nephrology, spanning chronic kidney disease and acute kidney injury prediction, dialysis management, organ allocation, graft monitoring, and patient education; reported performance is often strong, yet most models are developed and validated on data from a small number of high-income countries and can degrade substantially in populations that differ from the training distribution. The removal of race from estimated glomerular filtration rate equations offers an instructive precedent for how embedded assumptions in clinical algorithms can entrench disparities. This narrative review synthesizes AI applications across kidney disease, dialysis, and transplantation through the lens of external validity and health equity, with particular attention to Arab and African populations, who remain largely absent from both the relevant datasets and the surrounding discussion. To our knowledge, it is among the first reviews to apply this combined framework to these populations and to translate that analysis into a concrete regional agenda. We argue that the principal barrier to equitable benefit is not algorithmic capability but data representativeness, validation, and governance.

Key words: artificial intelligence, machine learning, deep learning, large language models, chronic kidney disease, dialysis, kidney transplantation, health equity

INTRODUCTION

Chronic kidney disease (CKD), acute kidney injury (AKI), and end-stage kidney disease (ESKD) constitute a growing global health burden, with rising prevalence driven by aging populations and the increasing incidence of diabetes and hypertension. Conventional risk-assessment tools, while clinically useful, are limited in their ability to integrate the high-dimensional, longitudinal, and multimodal data now routinely generated in clinical practice. Against this background, artificial intelligence (AI) has emerged as a means of analyzing complex clinical data to support earlier detection, more precise risk stratification, and individualized management of kidney disease. [1,2]

AI is an umbrella term that encompasses machine learning (ML), in which models learn patterns from data rather than from explicitly programmed rules; deep learning (DL), a subset of ML that uses multilayered neural networks to model complex, non-linear

relationships; and, most recently, generative AI and large language models (LLMs), which generate text and other outputs and are increasingly applied to documentation, triage, and patient education. [1,3] Across each of these categories, nephrology has seen a rapid expansion of published applications spanning the full spectrum of kidney disease, dialysis, and transplantation. [1,2]

A corresponding proliferation of review articles has accompanied this rapid growth. Recent syntheses include a state-of-the-art roadmap for clinical integration of AI across AKI, CKD, dialysis, and transplantation, [1] a scoping review of AI in CKD management, [2] and a focused narrative review of AI for predicting CKD progression to kidney failure. [4] These reviews establish that AI is technically capable across many nephrology tasks. What they address only in passing is a question of growing importance: whether the benefits of these tools will extend equitably to populations that are poorly represented in the data used to build them. The existing equity literature in kidney care has focused predominantly on Black American [5,6] and Hispanic[7] populations. In contrast, Arab and African populations remain largely absent from both the development datasets and the discussion. [8]

In this narrative review, we synthesize AI, ML, and DL applications across kidney disease, dialysis, and transplantation, and we situate them within an explicit framework of external validity and health equity for Arab and African populations. To our knowledge, this is among the first reviews to apply this combined framework specifically to Arab and African populations and to translate it into a concrete research, validation, and governance agenda. We argue that the central barrier to equitable benefit is not algorithmic capability but data representativeness, validation, and governance, and we draw on the removal of race from estimated glomerular filtration rate (eGFR) equations as a cautionary precedent. The aim is not to catalog every reported model, but to identify where current evidence is strong, where it is fragile, and what would be required for AI to serve Arab and African patients with kidney disease rather than to widen existing gaps.

This is a narrative review. We searched PubMed and Google Scholar between January 1, 2019, and June 15, 2026, for English-language literature on AI, ML, DL, and LLMs in CKD, AKI, dialysis, and transplantation, combined with terms for health equity, external validity, and Arab, Middle Eastern, North African, and African populations, supplemented by the reference lists of relevant reviews. We did not undertake a systematic search, predefined eligibility criteria, or formal quality appraisal; studies were selected for relevance to the review's framework, and our novelty assessment reflects this search rather than an exhaustive one. **Figure 1** provides a graphical synthesis of this argument, from the breadth of current applications across the kidney care continuum to the regional research, validation, and governance agenda required for equitable benefit. AI across the kidney care continuum and the equity gap for Arab and African populations are illustrated in **Figure 1**.

AI IN KIDNEY DISEASE

Table 1 summarizes representative AI, ML, and DL applications across kidney disease, dialysis, and transplantation, together with their reported performance and validation status; each domain is considered in turn in the sections that follow.

Detection and Progression of CKD

A substantial body of work has applied ML and DL to the early detection of CKD and to the prediction of its progression. These approaches draw on data from electronic health records (EHRs), laboratory results, and imaging, with emerging interest in additional modalities such as omics data and wearable devices. [2] Across studies summarized in recent reviews, ML and DL models have achieved strong discrimination on internal validation—commonly with AUCs in the approximately 0.81 to 0.93 range for image-based CKD detection—but performance has frequently declined on external validation, falling to roughly 0.73 to 0.87 across independent cohorts. [2] Reported gains have more often concerned CKD detection and risk stratification than direct prediction of progression to ESKD, and the strongest figures are not uniformly expressed as AUCs. Unsupervised techniques, including clustering of eGFR trajectories, have been used to identify subgroups of patients with distinct progression patterns and differing risks of hospitalization and mortality. [2]

A dedicated narrative review of AI for predicting CKD progression to kidney failure concluded that ML-based models can outperform conventional risk equations in selected settings, while emphasizing that most models remain at the stage of retrospective development and internal validation rather than prospective, workflow-integrated deployment. [4] This gap between reported discrimination and demonstrated clinical utility is a recurring theme across the AI-in-nephrology literature and is especially consequential when models are considered for use in populations different from those in which they were trained.

Prediction of AKI

AKI prediction is among the most mature applications of DL in nephrology. In a landmark study, a recurrent neural network was developed on a large, longitudinal dataset of electronic health records comprising more than 700,000 adult patients across diverse inpatient and outpatient sites. It was able to predict a substantial proportion of AKI episodes up to 48 hours before they occurred, with state-of-the-art discrimination. [9] Such continuous, individual-level prediction within an actionable time window represents a genuine advance over static risk scores.

However, the same body of work also illustrates the central caution of this review. The dataset underpinning that landmark model was drawn from a United States veterans population that was approximately 94% male. When the model was subsequently reproduced and evaluated across health systems, it performed worse for women in both the original and an academic-hospital setting. At the same time, the discrepancy could be partly corrected by recalibration using a sex-balanced cohort in the non-veteran population; poorer performance persisted in veterans. [10] Importantly, these discrepancies could not be explained simply by sample size. [10] This experience demonstrates concretely that a model achieving excellent performance in its development population may systematically underperform in subgroups that are underrepresented in the training data, a problem that extends naturally from sex to geography, ancestry, and care setting.

AI IN DIALYSIS

In the dialysis setting, AI has been explored as a means of optimizing the delivery and monitoring of renal replacement

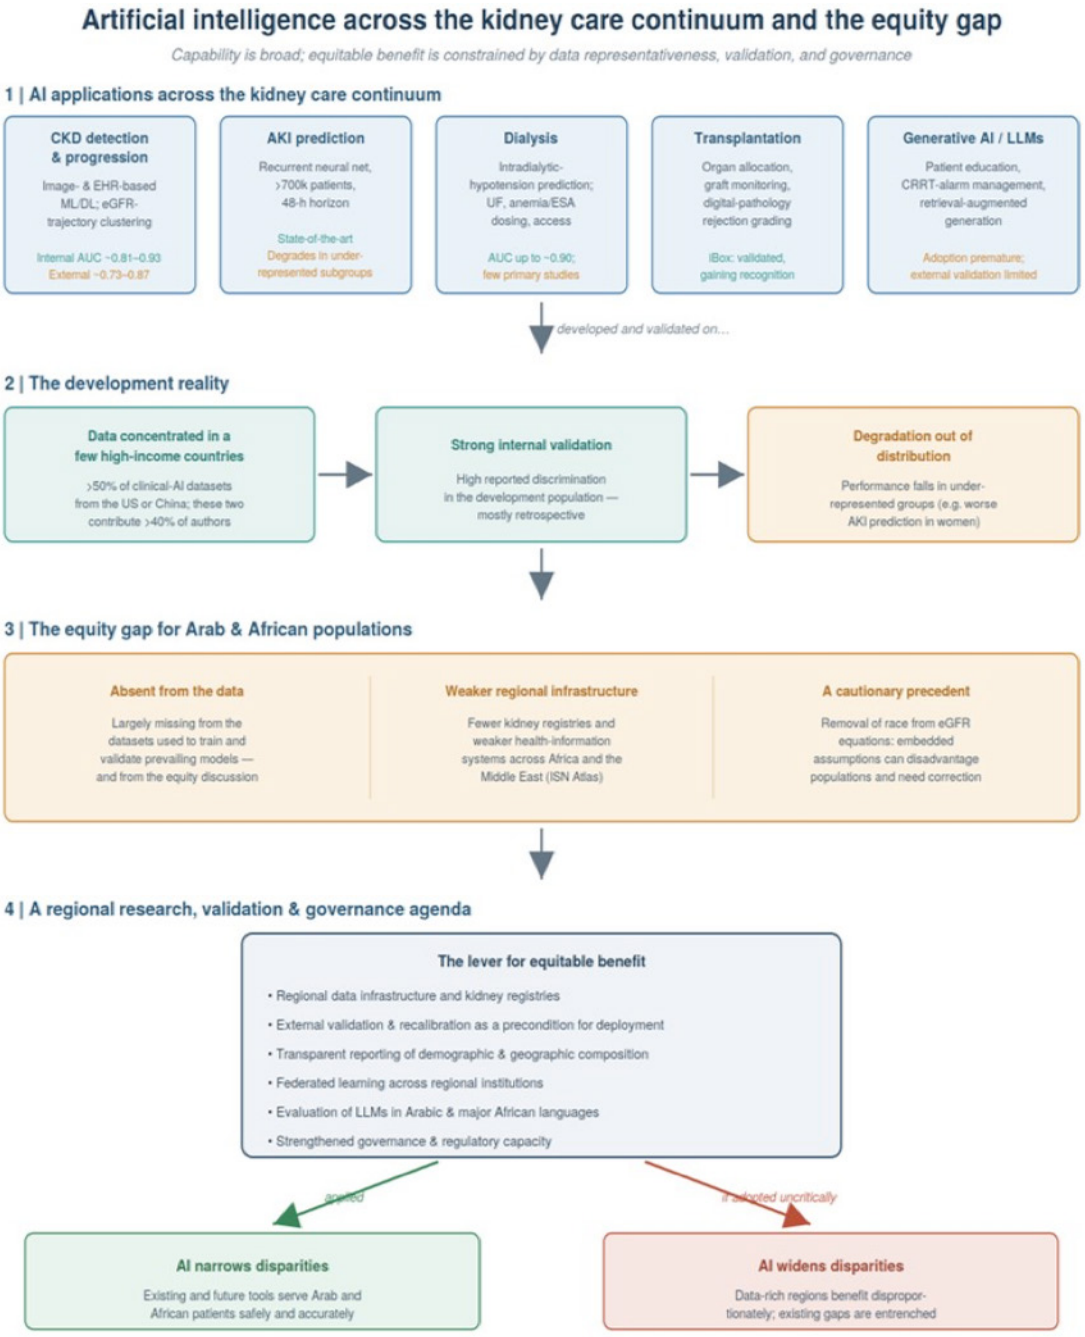

Figure 1: AI applications span CKD detection and progression, AKI prediction, dialysis management, transplantation, and generative AI/LLM-based patient communication (1). These tools are developed and internally validated largely on data from a small number of high-income countries, yet performance can degrade out of distribution and in underrepresented subgroups (2). Arab and African populations remain largely absent from training datasets and the equity discussion, face a weaker regional kidney-care infrastructure, and the race-free eGFR transition offers a cautionary precedent (3). Regional research, validation, and governance agenda are the lever that determines whether AI narrows or widens disparities (4). AKI, acute kidney injury; AUC, area under the curve; CKD, chronic kidney disease; CRRT, continuous renal replacement therapy; DL, deep learning; eGFR, estimated glomerular filtration rate; EHR, electronic health record; ESA, erythropoiesis-stimulating agent; ISN, International Society of Nephrology; LLM, large language model; ML, machine learning; UF, ultrafiltration.

therapy. Applications described in recent reviews include AI-enabled systems for optimizing ultrafiltration, supporting anemia management and erythropoiesis-stimulating agent dosing, and surveillance of vascular access, as well as models intended to predict intradialytic hypotension and other hemodynamic events. [1,2] As a concrete primary example, a DL model developed across two hemodialysis centers predicted intradialytic hypotension in an upcoming 10-minute window using only routinely recorded dialysis-machine measurements, reaching area-under-the-curve values up to approximately 0.90; importantly, it was evaluated in a separate external center rather than only in its development

Table 1: Representative applications of AI across kidney disease, dialysis, and transplantation.

Domain	Representative application(s)	Example model/finding	Reported performance	Validation status	Ref(s)
CKD detection and progression	EHR- and image-based ML/DL; unsupervised eGFR-trajectory clustering	ML/DL prediction of progression to ESKD; clustering into progression subgroups	AUC ~0.81–0.93 (internal); ~0.73–0.87 (external), image-based detection	Predominantly retrospective; external validation limited	[2,4]
AKI prediction	Continuous, individual-level risk prediction	Recurrent neural network; >700,000 patients; up to 48-hours horizon	State-of-the-art discrimination (AUC ~0.92)	Degrades in women on cross-system external evaluation; only partly corrected by recalibration	[9,10]
Dialysis	IDH prediction; ultrafiltration, anemia/ESA dosing, vascular-access surveillance	Two-center DL predicting IDH from hemodialysis-machine data alone	AUROC up to 0.905 (internal); 0.853–0.872 (external)	Externally validated (single study); few primary studies overall	[1,2,11]
Transplantation	Organ allocation, graft-function monitoring, digital-pathology rejection classification	iBox prediction of long-term graft outcomes	Validated; gaining regulatory recognition	Most advanced domain for clinical integration	[1]
Generative AI / LLMs	Patient education, CRRT-alarm management, RAG	LLM plain-language explanations of 100 transplantation terms	High accuracy in selected scenarios	Adoption premature; external validation limited	[12–14]

AI, artificial intelligence; AKI, acute kidney injury; AUC/AUROC, area under the (receiver operating characteristic) curve; CKD, chronic kidney disease; CRRT, continuous renal replacement therapy; DL, deep learning; eGFR, estimated glomerular filtration rate; EHR, electronic health record; ESA, erythropoiesis-stimulating agent; ESKD, end-stage kidney disease; IDH, intradialytic hypotension; LLM, large language model; ML, machine learning; RAG, retrieval-augmented generation.

cohort, and its reliance on machine data alone avoided the use of identifiable patient information. [11] Data-driven dosing models trained on large volumes of routine dialysis data have been reported to support more individualized prescriptions. [2]

The evidence base for dialysis applications is, at present, less consolidated than that for CKD and AKI prediction, and much of it is reported within broader reviews rather than through large, externally validated primary studies. [1,2] This is a particularly important limitation for Arab and African settings, where the dialysis population, vascular access practices, and resource constraints may differ substantially from the high-income settings in which most of these tools were developed. Robust, locally validated evidence will be needed before AI-assisted dialysis management can be recommended for routine use in these contexts.

AI IN KIDNEY TRANSPLANTATION

AI applications in kidney transplantation span the pre-transplant, peri-transplant, and post-transplant phases. Reported uses include support for organ allocation, dynamic monitoring of graft function, and digital pathology-assisted classification of rejection on biopsy specimens. [1] Some tools have begun to attain regulatory recognition; for example, the iBox system for prediction of long-term graft outcomes has been highlighted as a validated tool gaining such recognition. [1] These developments suggest that transplantation, with its rich longitudinal and histopathological data, is a promising domain for clinically integrated AI.

A distinct and rapidly growing strand concerns the use of LLMs for patient-facing communication in transplantation. In a recent evaluation, a contemporary LLM generated plain-language explanations for 100 clinically relevant kidney transplantation terms; the great majority of explanations were rated as highly accurate, and explicit readability-focused prompting substantially improved readability without compromising accuracy. [14] These findings derive from a single evaluation conducted in English, with a limited term set and rater pool, and should be regarded as preliminary. Such tools could help bridge persistent gaps in patient understanding. However, their evaluation to date has been conducted predominantly in English and in high-resource contexts [14]; apart from an evaluation of AI translation for Spanish speakers [7] comparable work in Arabic or in the diverse languages of African transplant populations is lacking.

GENERATIVE AI AND LLMS

The LLM applications introduced in transplantation form part of a broader, fast-moving frontier that spans the whole of nephrology and therefore warrants separate treatment. Generative AI, and LLMs in particular, represent the most recent and most rapidly evolving frontier in nephrology AI. A systematic review of LLM applications in nephrology identified studies addressing patient education, workflow optimization, laboratory data interpretation, and clinical decision support. [3] Reported applications include the generation of readable patient-education materials and the management of continuous renal replacement therapy alarms, where an LLM achieved high accuracy in selected scenarios. [3] The authors nonetheless concluded that widespread adoption remains premature, noting that performance depends heavily on input quality and that external validation is limited. [3]

Complementary reviews have described emerging capabilities such as conversational agents, multimodal integration, and retrieval-augmented generation (RAG), in which an LLM is coupled to a curated knowledge base to improve factual grounding and reduce fabrication. [12,13] RAG and domain-specific fine-tuning are of particular interest because they offer a route to adapting general-purpose models to local guidelines, formularies, and languages. [13] At the same time, these reviews consistently flag concerns regarding patient privacy, data security, model accuracy, algorithmic bias, and ethical accountability, and emphasize that real-world evidence supporting safety and effectiveness remains scarce. [3,12] For Arab and African settings, the dependence of LLM performance on training-data composition, and the present concentration of evaluation in English-language and high-income contexts, are central considerations. [8,12]

ALGORITHMIC BIAS AND HEALTH EQUITY: THE CASE OF ARAB AND AFRICAN POPULATIONS

How Bias Enters Clinical AI

Bias in medical AI arises and compounds throughout the model lifecycle, from data collection and labeling through model development, evaluation, deployment, and publication. [6] Insufficient sample sizes for particular patient groups can produce suboptimal and clinically unhelpful predictions for those groups; non-randomly missing data and labels that reflect prior disparities in care can encode existing inequities into ostensibly objective models. [6] The clinical consequences are not hypothetical. In an influential analysis, a widely used commercial algorithm that guided access to care-management programs was found to disadvantage Black patients systematically, because it used healthcare costs as a proxy for health need and less money had historically been spent on Black patients at equivalent levels of illness. [15] This example illustrates how a model can be technically well-performing on its chosen metric while perpetuating harm.

Race-Free eGFR Equation as a Precedent

Nephrology has already confronted, in a non-AI context, the question of whether embedding population assumptions in a clinical algorithm is appropriate. For more than two decades, eGFR equations incorporated a coefficient for race. In 2021, a National Kidney Foundation–American Society of Nephrology task force recommended adoption of new eGFR equations that do not include race, because race is a social rather than a biological construct and that its inclusion contributed to disparities in CKD care. [5,16] New creatinine- and cystatin C-based equations were developed and validated for this purpose. [16] This episode is directly relevant to the present discussion: it demonstrates both that clinical algorithms can embed assumptions that disadvantage specific populations, and that correcting such assumptions requires deliberate, evidence-based effort rather than the mere availability of better technology.

Data Underrepresentation and the Limits of External Validity

We group Arab and African populations pragmatically, based on a shared structural challenge—marked underrepresentation in the datasets, authorship, and validation of clinical AI—rather than any assumption of biological, genetic, or cultural homogeneity. These populations are highly heterogeneous: the Gulf states, the wider Middle East and North Africa (MENA), and sub-Saharan Africa differ substantially in healthcare infrastructure, disease etiology, and resources, and both "Arab" and "African" encompass enormous ethnic, genetic, and linguistic diversity. Where relevant, we therefore distinguish MENA from sub-Saharan African settings, while treating data underrepresentation as the common thread motivating a shared agenda. Importantly, this grouping is an organizing frame for a shared research, validation, and governance agenda, not an analytic category to be embedded in models: nothing in our argument implies that "Arab" or "African" should serve as a model input, as race once did in eGFR equations. The agenda we set out instead calls for disaggregated, country- and subgroup-specific data and local validation, so that a common starting point of underrepresentation resolves into population-specific evidence rather than hardening into a new coarse category.

The fundamental constraint on equitable AI is that models generalize poorly to populations unlike those on which they were trained. A global review of the clinical-AI landscape found that more than half of the datasets used to develop clinical AI originate from either the United States or China, and that these two countries contribute over 40% of the authors of the corresponding publications. [8] Because data-rich regions stand to benefit disproportionately, the authors warned that AI could entrench, rather than reduce, global health inequity, and they identified development of data infrastructure in data-poor regions, together with external validation and recalibration before deployment, as essential safeguards. [8] The AKI prediction experience described earlier provides a concrete nephrology example of subgroup performance degradation, [10] and a systematic review of clinical AI/ML studies found that demographic composition of training data is frequently not reported at all, making it impossible for clinicians to judge whether a model applies to their patients. [17] The role and use of race and ancestry in health-related AI models remain contested and inconsistently handled. [18]

For Arab and African populations, these concerns are acute. Data from the ISN Global Kidney Health Atlas show that kidney registries, health information systems, and the wider infrastructure for kidney care are markedly less developed across the Africa and Middle East regions than in high-income settings, including fewer kidney registries and weaker health information systems. [19] Against this backdrop, clinical datasets from the MENA and sub-Saharan Africa remain scarce in the corpora used to train and validate prevailing models, [8] and regional differences in disease etiology, comorbidity patterns, and health-system resources mean that performance demonstrated elsewhere cannot be assumed to transfer. A model that has never seen data resembling a Libyan, Qatari, or sub-Saharan African patient population offers no guarantee of accuracy in those settings, however impressive its reported metrics.

Nephrology-Specific Equity Evidence and the Existing Gap

Equity-focused evaluation within nephrology AI is beginning to emerge but remains narrow in scope. A study evaluating how widely used LLMs handle diversity, equity, and inclusion considerations in simulated nephrology decision-making

found that such models can reflect and potentially reinforce disparities, underscoring the need for regulatory oversight of equitable AI deployment. [20] The limited equity work that does exist has concentrated on specific populations and languages; for example, AI translation of kidney-donor information has been evaluated for Spanish speakers. [13] Commentaries on how racial bias becomes built into clinical models reinforce that the problem is structural rather than incidental. [21] What is conspicuously absent is comparable work centered on Arab and African populations, on Arabic and African languages, and on the external validation of existing models in MENA and African cohorts. This absence is precisely the gap this review seeks to make explicit, and it motivates the regionally focused research, validation, and governance agenda set out in the next section.

LIMITATIONS, REGULATION, AND FUTURE DIRECTIONS

Several cross-cutting limitations temper current enthusiasm. Most published models are retrospective and internally validated; prospective, workflow-integrated, and externally validated evidence is comparatively scarce. [1,3] Interpretability, data heterogeneity, privacy, and accountability remain unresolved challenges, particularly for generative models. [3,12] Regulatory frameworks are still maturing: an analysis of AI-based medical devices cleared by the United States Food and Drug Administration found that approvals have been concentrated in image-rich specialties, with comparatively few in nephrology, [22] and professional bodies have begun to articulate principles for responsible use, including a recent statement from the American Society of Nephrology on responsible use of AI to improve kidney care. [23]

Translating these general principles into equitable benefit for Arab and African populations is, in our view, the central contribution this review seeks to advance, and it points to several concrete priorities. First, deliberate investment in regional data infrastructure and registries is needed so that models can be developed and validated on representative data. Second, external validation and recalibration in local cohorts should be a precondition for deployment rather than an afterthought. [8] Third, transparent reporting of the demographic and geographic composition of training and validation data should be expected of all clinical AI studies, enabling clinicians to assess applicability. [17] Fourth, technical approaches such as federated learning may allow institutions across the region to contribute to model development without sharing patient-level data, partially addressing both privacy and representativeness. [24] Fifth, evaluation of LLMs and patient-facing tools should extend to Arabic and to the major languages of African transplant and dialysis populations. [7,14] Sixth, as a narrative review, this synthesis is subject to selection bias, did not apply a systematic search or formal quality assessment, and reflects the authors' framing of a fast-moving field. Finally, governance and regulatory capacity within the region should be strengthened so that equitable deployment is actively overseen. [20,23] These priorities must also confront economic and implementation barriers—the cost of data infrastructure, computing, and clinical integration weighs heavily in lower-resource settings—and the dominance of commercial AI products means partnerships among academic institutions, regional professional societies, governments, and industry are likely needed to ensure equitable benefit rather than market reach drives deployment.

CONCLUSIONS

AI, ML, and DL have demonstrated genuine promise across kidney disease, dialysis, and transplantation, and generative models are extending that promise into documentation, decision support, and patient communication. Nevertheless, the evidence base is dominated by data from a small number of high-income countries, and models that perform well in their development populations can fail in groups they have not adequately seen. The removal of race from eGFR equations shows both the harm that embedded assumptions can cause and the deliberate effort required to correct them. For Arab and African populations, who are at present largely absent from the relevant datasets and from the equity discussion, the central task is not to await better algorithms but to build the data infrastructure, validation practices, reporting standards, and governance that would allow existing and future tools to serve these populations safely. Approached in this way, AI can help narrow disparities in kidney care; approached uncritically, it is likely to widen them.

AUTHORS' CONTRIBUTION

All authors have significantly contributed to the work, whether by conducting literature searches, drafting, revising, or critically reviewing the article. They have given their final approval of the version to be published, have agreed with the journal to which the article has been submitted, and agree to be accountable for all aspects of the work.

DISCLOSURE OF AI ASSISTANCE

The authors disclose the use of AI technologies in the preparation of this manuscript, specifically for linguistic refinement, content restructuring, and figure creation. Despite this assistance, the intellectual framework, analysis, and scholarly conclusions remain the exclusive original work of the authors.

SOURCE OF FUNDING

None.

CONFLICT OF INTEREST

None.

REFERENCES

1. Cheungpasitporn W, Athavale A, Ghazi L, Kashani KB, Colicchio T, Koyner JL. Transforming nephrology through artificial intelligence: A state-of-the-art roadmap for clinical integration. Clin Kidney J. al. 2026;19(2):sfag004. https://doi.org/10.1093/ckj/sfag004
2. Sabanayagam C, Banu R, Lim C, Tham Y, Cheng C, Tan G. Artificial intelligence in chronic kidney disease management: A scoping review. Theranostics. 2025;15(10):4566–78. https://doi.org/10.7150/thno.108552
3. Unger Z, Soffer S, Efros O, Chan L, Klang E, Nadkarni G. Clinical applications and limitations of large language models in nephrology: A systematic review. Clin Kidney J. 2025;18(9):sfaf243. https://doi.org/10.1093/ckj/sfaf243
4. Miller Z, Dwyer K. Artificial intelligence to predict chronic kidney disease progression to kidney failure: A narrative review. Nephrology (Carlton). 2025;30(1):e14424. https://doi.org/10.1111/nep.14424

5. Delgado C, Baweja M, Crews D, Eneanya N, Gadegbeku C, Inker L. A unifying approach for GFR estimation: Recommendations of the NKF-ASN Task Force on Reassessing the Inclusion of Race in Diagnosing Kidney Disease. J Am Soc Nephrol. 2021;32(12):2994–3015. https://doi.org/10.1681/ASN.2021070988
6. Cross J, Choma M, Onofrey J. Bias in medical AI: Implications for clinical decision-making. PLOS Digit Health. 2024;3(11):e0000651. https://doi.org/10.1371/journal.pdig.0000651
7. Garcia Valencia O, Thongprayoon C, Jadlowiec C, Mao S, Leeaphorn N, Budhiraja P. Advancing health equity: Evaluating AI translations of kidney donor information for Spanish speakers. Front Public Health. 2025;13:1484790. https://doi.org/10.3389/fpubh.2025.1484790
8. Celi L, Cellini J, Charpignon M, Dee E, Dernoncourt F, Eber R. Sources of bias in artificial intelligence that perpetuate healthcare disparities-A global review. PLOS Digit Health. 2022;1(3):e0000022. https://doi.org/10.1371/journal.pdig.0000022
9. Tomašev N, Glorot X, Rae J, Zielinski M, Askham H, Saraiva A. A clinically applicable approach to continuous prediction of future acute kidney injury. Nature. 2019;572(7767):116–119. https://doi.org/10.1038/s41586-019-1390-1
10. Cao J, Zhang X, Shahinian V, Yin H, Steffick D, Saran R. Generalizability of an acute kidney injury prediction model across health systems. Nat Mach Intell. 2022;4(12):1121–1129. https://doi.org/10.1038/s42256-022-00563-8
11. Kim H, Heo S, Kim M, Lee J, Park K, Lee G. Deep learning model for predicting intradialytic hypotension without privacy infringement: A retrospective two-center study. Front Med (Lausanne). 2022;9:878858. https://doi.org/10.3389/fmed.2022.878858
12. Hu Y, Liu J, Jiang W. Large language models in nephrology: Applications and challenges in chronic kidney disease management. Ren Fail. 2025;47(1):2555686. https://doi.org/10.1080/0886022X.2025.2555686
13. Miao J, Thongprayoon C, Suppadungsuk S, Garcia Valencia O, Cheungpasitporn W. Integrating retrieval-augmented generation with large language models in nephrology: Advancing practical applications. Med Kaunas. 2024;60(3):445. https://doi.org/10.3390/medicina60030445
14. Garcia Valencia O, Thongprayoon C, Miao J, Craici I, Cheungpasitporn W. AI-generated explanations in kidney transplantation: Accuracy vs. readability and implications for patient education. Front Artif Intell. 2026;9:1806516. https://doi.org/10.3389/frai.2026.1806516
15. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–453. https://doi.org/10.1126/science.aax2342
16. Inker L, Eneanya N, Coresh J, Tighiouart H, Wang D, Sang Y, et al; Chronic Kidney Disease Epidemiology Collaboration. New creatinine- and cystatin C-based equations to estimate GFR without race. N Engl J Med. 2021;385(19):1737–1749. https://doi.org/10.1056/NEJMoa2102953
17. Otokiti A, Shih H, Williams K. Gender and racial bias unveiled: Clinical artificial intelligence (AI) and machine learning (ML) algorithms are fanning the flames of inequity. Oxf Open Digit Health. 2025;3:oqaf027. https://doi.org/10.1093/oodh/oqaf027
18. Were M, Li A, Malin B, Yin Z, Coco J, Collins B. Role and use of race in artificial intelligence and machine learning models related to health. J Med Internet Res. 2025;27:e73996. https://doi.org/10.2196/73996
19. Bello A, Okpechi I, Levin A, Ye F, Damster S, Arruebo S, et al; ISN-GKHA Group. An update on the global disparities in kidney disease burden and care across world countries and regions. Lancet Glob Health. 2024;12(3):e382–e395. https://doi.org/10.1016/S2214-109X(23)00570-3
20. Balakrishnan S, Thongprayoon C, Wathanavasin W, Miao J, Mao M, Craici I. Evaluating artificial intelligence bias in nephrology: The role of diversity, equity, and inclusion in AI-driven decision-making and ethical regulation. Front Artif Intell. 2025;8:1525937. https://doi.org/10.3389/frai.2025.1525937
21. Jindal A. Misguided artificial intelligence: How racial bias is built into clinical models. Brown J Hosp Med. 2023;2(1):38021. https://doi.org/10.56305/001c.38021
22. Benjamens S, Dhunnoo P, Meskó B. The state of artificial intelligence-based FDA-approved medical devices and algorithms: An online database. NPJ Digit Med. 2020;3(1):118. https://doi.org/10.1038/s41746-020-00324-0
23. Tangri N, Cheungpasitporn W, Crittenden S, Fornoni A, Peralta C, Singh K, et al. Responsible use of artificial intelligence to improve kidney care: A statement from the American Society of Nephrology. J Am Soc Nephrol. 2026;37(4):881–90. https://doi.org/10.1681/ASN.0000000929
24. Rieke N, Hancox J, Li W, Milletarì F, Roth H, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3(1):119. https://doi.org/10.1038/s41746-020-00323-1