

An Explainable AI Framework Integrating Traditional Chinese Medicine Knowledge Graphs and Digital Biomarkers for Personalized Edible Herbal Formula Recommendations in Chronic Disease Management

Isaac D. Sanders

School of Computing, Clemson University, Clemson, SC, USA.
isaacdsanders155@clemson.edu

Manav Gendhe

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
contactmanav@unh.edu

Cesar A. Fox

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
cesar.work@buffalo.edu

Abstract

Chronic diseases remain the leading cause of mortality and disability worldwide, demanding a shift from standardized treatment toward precision health strategies that account for inter-individual variability. Traditional Chinese Medicine (TCM) offers a holistic paradigm of syndrome differentiation and personalized herbal formulations that has been refined over millennia, yet its integration into modern clinical workflows is hindered by a lack of quantitative evidence, methodological standardization, and computational tools that can reconcile its conceptual depth with contemporary biomedical data. This paper presents a system-level framework for an explainable artificial intelligence (XAI) engine that combines structured TCM knowledge graphs with longitudinally acquired digital biomarkers to generate personalized edible herbal formula recommendations. We critically examine the architectural design choices, data governance models, interoperability requirements, fairness considerations, and sustainability challenges associated with deploying such a system at scale. The proposed framework adopts a layered architecture in which a dynamically curated TCM knowledge graph, populated with herb-compound-target relationships, syndrome associations, and safety constraints, interacts with a multi-modal biomarker ingestion pipeline that processes data from wearable sensors, continuous metabolic monitors, and self-reported outcomes. Explainability is achieved not through post-hoc approximation alone but through an inherently interpretable inference backbone that traces evidential paths across the knowledge graph and communicates therapeutic rationales aligned with TCM diagnostic principles. The discussion extends to regulatory alignment, the mitigation of cultural and demographic biases, and the long-term institutional and environmental costs of maintaining a living knowledge ecosystem. By articulating these cross-cutting concerns, the paper offers a blueprint for future translational research that bridges ancient medical wisdom with data-driven personalization in chronic disease management.

Keywords

explainable AI, traditional Chinese medicine, knowledge graphs, digital biomarkers, chronic disease management, personalized herbal formulas.

1. Introduction

The global burden of chronic non-communicable diseases, including type 2 diabetes, cardiovascular disorders, and metabolic syndrome, has intensified the search for interventions that transcend the conventional one-size-fits-all paradigm. Although pharmacotherapy and lifestyle modifications remain foundational, their effectiveness is frequently attenuated by heterogeneous patient physiologies, psychosocial contexts, and environmental exposures. Within this landscape, artificial intelligence has been positioned as a transformative force capable of uncovering latent patterns in high-dimensional clinical and lifestyle data to tailor therapies [1]. Concurrently, the maturation of wearable sensing technologies has given rise to the concept of digital biomarkers, objective, quantifiable physiological and behavioral measures collected outside clinical settings, which can offer a granular, longitudinal window into an individual's health trajectory [2]. When properly harnessed, these data streams can reveal dynamic homeostatic perturbations that precede overt pathology, enabling a shift from reactive to proactive management.

To translate such rich data into actionable therapeutic recommendations, robust knowledge representation is indispensable. Knowledge graphs have emerged as a compelling paradigm for integrating heterogeneous biomedical information, linking entities such as diseases, genes, metabolites, drugs, and phenotypes in a machine-traversable semantic network [3]. Their capacity to encode relational knowledge makes them especially suitable for representing the multi-factorial and interconnected nature of chronic disease pathophysiology. However, the clinical deployment of AI-driven recommendation systems is frequently obstructed by their opacity. Clinicians and patients alike demand justification for therapeutic suggestions, particularly when they involve natural products with complex pharmacodynamic profiles. Explainable AI has therefore become a central pillar of medical AI research, aiming to render model behavior transparent without sacrificing predictive accuracy [4].

Traditional Chinese Medicine offers a distinctive lens through which personalization can be reimagined. Rooted in over two thousand years of empirical practice, TCM classifies individual constitution and disease states through pattern differentiation, and prescribes herbal formulas that are tailored to the individual's dynamic balance of vital substances. Recent advances in computational TCM have begun to explore how large language models and knowledge-driven systems can assist in generating edible herbal formula recommendations, signaling a convergence of classical wisdom and modern machine intelligence [5]. Concurrently, the broader field of AI-driven drug discovery and delivery has illuminated pathways for computational repositioning of bioactive compounds and the optimization of therapeutic regimens, demonstrating the feasibility of algorithmic augmentation in natural product research [6].

Yet the leap from isolated algorithmic experiments to a fully operational, trustworthy, and sustainable clinical decision support infrastructure remains vast. The design of such a system must navigate tensions between model complexity and interpretability, between the fluid subjectivity of TCM diagnostics and the quantitative rigidity of digital biomarkers, and between the imperative to personalize and the need to ensure safety across diverse populations. It also raises profound ethical, regulatory, and structural questions that are rarely

addressed within a single architectural discourse. This paper develops a high-level system framework that weaves together TCM knowledge graphs, digital biomarkers, and explainable AI into a coherent recommendation engine for edible herbal formulas. Our contribution is not a singular algorithm but an integrative architectural analysis that foregrounds the governance, fairness, deployment, and sustainability dimensions that will ultimately determine whether such systems achieve clinical credibility and population-level impact.

2. Conceptual Framework and System Architecture

The proposed framework is organized as a layered service-oriented architecture that cleanly separates data acquisition, knowledge curation, inference, explanation, and user interaction. Adopting a layered design is a deliberate architectural choice that promotes modular verifiability, enables independent evolution of each subsystem, and accommodates the heterogeneous update cycles associated with rapidly evolving digital biomarker technologies and comparatively stable TCM knowledge structures. The lowest layer, the sensing and ingestion tier, interfaces with multiple wearable and nearable devices, smartphone-based ecological momentary assessments, and electronic health record application programming interfaces. It performs signal quality assessment, artifact rejection, and time-series alignment before forwarding streams to the digital biomarker computation engine. By abstracting device-specific protocols behind a unified data model, this tier ensures that the system remains vendor-agnostic and can incorporate future sensing modalities without disrupting downstream components.

Above the ingestion layer, the digital biomarker engine transforms raw physiological and behavioral streams into clinically meaningful features. Parameters such as heart rate variability circadian amplitude, nocturnal skin temperature regulation, continuous glucose variability indices, and accelerometry-derived activity fragmentation are computed using validated algorithms. Importantly, the engine also integrates contextual metadata concerning medication timing, dietary intake, and sleep schedules, as these confounders must be disentangled to attribute biomarker deviations to underlying TCM-relevant homeostatic shifts. The output is a set of time-stamped, normalized digital biomarker vectors that represent the patient's dynamic phenotype, ready to be mapped to TCM syndrome patterns.

The central reasoning core of the system is a hybrid recommender that couples a structured TCM knowledge graph with a multi-modal inference module. The knowledge graph stores entities such as herbs, formulas, chemical constituents, protein targets, syndrome types, organ systems, and flavors and properties, along with their weighted relationships derived from classical pharmacopoeias, modern biochemical assays, and observational studies. The inference module accepts both the digital biomarker profile and patient-reported symptoms and performs a context-aware traversal of the knowledge graph, identifying formulas that are semantically proximal to the patient's current constitutional state. Explainability is woven into the inference process itself rather than appended afterwards: the traversal generates explanatory chains that can be rendered as coherent therapeutic narratives, citing which pattern indicators triggered the selection of a cooling or warming herb, or why a specific combination was favored to address coexisting deficiency and excess patterns.

Microservice decomposition underpins the deployment architecture, allowing the knowledge graph, inference engine, and explanation generator to scale independently. Container orchestration platforms facilitate elastic scaling in response to query load, while message brokers decouple the streaming data pipeline from the recommendation service. Edge computing emerges as a pragmatic complement to cloud-centric processing; sensitive

biomarker profiles can be preprocessed and anonymized on the patient's smartphone or a local gateway before transmission, reducing latency and strengthening privacy. The trade-off between on-device model inference and centralized computation must be calibrated according to model size and explainability completeness, since generating rich narrative explanations often requires access to the full knowledge graph, which may be too large to reside locally.

Interoperability with hospital information systems is achieved through standardized health informatics interfaces such as HL7 FHIR, which enable the system to consume structured problem lists and laboratory results and to write back recommendation logs into the clinical record. This bidirectional integration is essential not only for clinical workflow acceptance but also for longitudinal safety surveillance, as adverse events potentially attributable to herbal formulations can be captured and fed back into the knowledge graph as new constraints. The architectural decision to expose recommendation rationales via a dedicated explanation RESTful endpoint allows third-party clinical decision support platforms to embed TCM-informed advice within broader care orchestration dashboards, thus lowering the barrier for adoption in conventional medical settings that may lack deep TCM expertise.

3. Traditional Chinese Medicine Knowledge Graph Construction and Integration

Constructing a machine-traversable representation of TCM knowledge is a formidable knowledge engineering undertaking, primarily because TCM concepts are deeply contextual, often metaphorical, and vary across schools of thought. The knowledge graph must accommodate a taxonomy in which syndromes such as Qi deficiency or damp-heat accumulation are not reducible to single biomarkers but represent emergent gestalts inferred from constellations of signs. To capture this complexity, the ontology design embraces a multi-granular formalization. At the foundational layer, a canonical syndrome taxonomy aligned with World Health Organization International Classification of Diseases extensions for TCM is mapped to lower-level physiological and behavioral proxies. Syndromes are modeled not as static labels but as process-oriented nodes that can be activated transiently and possess temporal dynamics, an essential feature for integrating with continuous digital biomarker streams.

Herbal entities are annotated with an extensive set of attributes covering botanical identification, part used, processing method, organ tropism, flavor, thermal nature, and known chemical constituents. Formula nodes connect constituent herbs through weighted composition edges, capturing the jun-chen-zuo-shi hierarchy that defines the principal, assistant, adjuvant, and guiding roles within a formula. The relational fabric is enriched by importing molecular interaction data from established biomedical knowledge bases, linking herb-derived compounds to protein targets, metabolic pathways, and gut microbial transformations. A paradigmatic case is the encoding of specific polysaccharides that have demonstrated regulation of glycolipid metabolism and modulation of the gut microbiome, thereby creating a traceable chain from herb, through molecular mechanism, to metabolic digital biomarkers [10]. This granularity enables the inference engine to ground recommendations not merely in empirical TCM rules but also in measurable biochemical evidence, satisfying the growing demand for convergent validation across paradigms.

Maintaining the knowledge graph demands a governance framework that orchestrates continuous ingestion from diverse sources. Automated natural language processing pipelines mine classical texts, modern clinical trial registries, and pharmacovigilance databases to extract herb-syndrome-outcome triples. Ontology-based curation workflows, as previously formalized in TCM informatics research, provide a structured mechanism for manual expert

review and conflict resolution [11]. The system must be capable of aging out outdated assertions, flagging contradictory evidence, and tracking provenance to support iterative refinement. Scalability is addressed through graph database engines that support labeled property graphs, enabling efficient traversal of multi-hop paths even as the entity count grows into the tens of millions. Indexing strategies that precompute embedding-based similarity measures for herbs and syndromes allow rapid approximate retrieval while preserving the ability to fall back on exact symbolic reasoning for explanation generation.

An additional layer of integration concerns the reconciliation of TCM knowledge with existing biomedical ontologies such as the Gene Ontology, Human Phenotype Ontology, and the Experimental Factor Ontology. Aligning TCM syndrome nodes with these open-domain ontologies is non-trivial because direct equivalence is rare. Instead, the integration relies on bridging rules that associate TCM patterns with clusters of biomedical phenotypes, for instance linking liver Qi stagnation to a composite of stress-related neuroendocrine indicators. These bridging rules are themselves treated as first-class knowledge elements that carry confidence scores and literature references, enabling the system to modulate the weight given to TCM-derived insights depending on the strength of evidential convergence. The curated graph thus becomes a living, interdisciplinary asset that can both guide therapeutic reasoning and serve as a substrate for hypothesis generation in integrative medicine research.

4. Digital Biomarkers for Personalizing Herbal Formulas

Digital biomarkers provide the quantitative substrate that allows the abstract patterns of TCM syndrome differentiation to be anchored in objective, continuous measurements. A core design challenge is the mapping between the high-level, qualitative language of TCM diagnostics and the numerical time series generated by modern sensors. For example, the TCM concept of Qi deficiency, which encompasses symptoms of fatigue, shortness of breath, and spontaneous sweating, can be probabilistically linked to reduced heart rate variability, diminished accelerometer-measured physical activity bout duration, and altered electrodermal activity patterns. Rather than attempting a reductionist equivalence, the framework treats these mappings as probabilistic likelihoods, learned from cohort data and refined through clinician feedback, and represents them as fuzzy membership functions within the inference engine.

The acquisition pipeline is designed to accommodate both consumer-grade wearables, which offer broad accessibility but lower accuracy, and medical-grade biosensors required for certain metabolic indicators. Continuous glucose monitors, for instance, yield glucose variability and time-in-range metrics that are particularly relevant for the TCM conceptualization of middle burner dampness and heat, conditions often associated with glycemic dysregulation. Similarly, peripheral skin temperature rhythms captured by ring or patch sensors can serve as proxies for the Yang-Qi warming function, with flattened amplitudes suggesting a deficiency pattern. Gut microbiome profiling, while not continuous in the strict sense, provides a periodically sampled digital biomarker that links directly to the TCM notion of spleen-stomach function and to metabolic phenotypes modulated by polysaccharide-rich herbs [10]. The framework incorporates a scheduling module that determines the optimal cadence and combination of measurements based on the patient's risk profile, balancing data richness against measurement burden.

Multi-modal fusion is technically demanding due to heterogeneous sampling rates, varying degrees of missingness, and differential noise characteristics. The adopted strategy is a hierarchical latent variable model that learns a shared low-dimensional representation of an

individual's constitutional state over time. This representation is deliberately interpretable by design, with dimensions constrained to approximate TCM-relevant axes such as cold-heat, deficiency-excess, and interior-exterior. By projecting the fused sensor data into this space, the system creates a dynamic phenotype that can be compared against the syndrome activation patterns encoded in the knowledge graph. Importantly, the fusion model is not a monolithic black box; its factor loadings are examined through post-hoc explainability techniques to ensure that the inferred constitutional state is grounded in recognizable physiological variables, a transparency requirement that directly supports clinical trust and regulatory acceptance.

Privacy is a paramount concern in digital biomarker systems, given the intimate and potentially re-identifiable nature of time-series physiological data. The architecture adopts privacy-by-design principles, performing initial signal processing and feature extraction on-device wherever computationally feasible, and employing federated learning protocols when cross-site model refinement is necessary [12]. Differential privacy guarantees can be applied to the contributed model updates, preventing the leakage of individual-level trajectories while still allowing the central fusion model to benefit from population-wide learning. Consent management is implemented as a granular, revocable scheme that distinguishes between data used for real-time recommendation, data retained for model improvement, and data shared with research repositories, thereby respecting the autonomy of patients whose cultural and personal attitudes toward data sharing may vary widely.

5. Explainable AI for Herb Formula Recommendations

Explainability in the context of TCM herbal formula recommendation carries unique requirements that extend beyond the standard mandates of algorithmic transparency in medical AI. The explanations must not only make the model's reasoning intelligible to clinicians but also resonate with the TCM diagnostic framework, employing terminology and causal narratives that are congruent with pattern differentiation. A purely feature-attribution explanation, for instance highlighting that reduced skin temperature variance contributed 0.34 to a recommendation score, would be semantically impoverished and clinically unpersuasive. Instead, the framework generates what we term therapeutic rationales: coherent paragraphs that trace the evidential path from specific digital biomarker perturbations, through activated syndrome patterns, to the roles of individual herbs within the chosen formula.

To accomplish this, the inference engine is constructed as an inherently interpretable neuro-symbolic system rather than a post-hoc-explained black box. The symbolic backbone leverages the knowledge graph's structure to perform rule-governed retrieval and composition of formula candidates, while learned neural components estimate the strength of association between biomarker profiles and syndrome nodes. The hybrid design ensures that every decision can be audited by replaying the sequence of graph traversals and logical unifications that led to the output. When additional nuance is required, game-theoretic explanation methods such as SHapley Additive exPlanations can be employed to attribute the contribution of specific input features to the syndrome activation scores, but these attributions are then translated into human-readable statements about organ system involvement or pathogenic factor predominance [15].

The generation of narrative explanations depends on a template-based natural language generation module that draws on a curated lexicon of TCM descriptors. For a recommended formula that strengthens spleen Qi and resolves dampness, the explanation might state that the formula was selected because the patient's continuous glucose variability and reported

postprandial heaviness collectively indicate spleen Qi deficiency with dampness encumbrance; the inclusion of *Atractylodes macrocephala* as the sovereign herb serves to tonify spleen Qi, while the addition of *Poria cocos* as an adjuvant leaches out dampness. These statements are not generated arbitrarily but are constrained by the formula's jun-chen-zuo-shi structure encoded in the knowledge graph. Counterfactual explanations further enhance transparency by illustrating what would change if a particular digital biomarker returned to a normal range, showing how the syndrome profile would shift and consequently how the formula would be simplified or altered, thus empowering patients to understand the behavioral and physiological levers that influence their therapy.

Trade-offs between model fidelity and interpretability are acknowledged and managed through a multi-tiered reasoning strategy. For straightforward presentations, the symbolic inference engine suffices and produces complete, transparent explanations. For ambiguous cases where multiple syndromes interact in complex ways, a latent-variable neural scorer ranks formula candidates, and the symbolic layer reconstructs the best possible post-hoc justification that is both faithful to the scorer's ranking and internally consistent with TCM theory. Clinician-facing interfaces expose a confidence indicator alongside each recommendation, reflecting both the degree of evidential support from digital biomarkers and the coherence of the TCM rationale. This design choice directly confronts the tension that complete model opacity would undermine professional autonomy, while excessive rigidity in adhering to predefined rules would prevent the system from capturing the subtleties that make TCM personalization powerful.

6. System Governance, Fairness, and Ethical Considerations

The deployment of an AI system that recommends ingestible herbal products for chronic disease management enters a regulatory and ethical arena of considerable complexity. Herbal formulas, though often classified as food supplements or traditional remedies in many jurisdictions, can exert potent pharmacological effects and interact with prescribed medications. Consequently, the system must be designed from the outset as a safety-critical medical device under frameworks such as the International Medical Device Regulators Forum guidelines, with rigorous pre-market validation and post-market surveillance. The governance structure must delineate clear lines of accountability, determining whether the system serves as a clinical decision support tool that leaves final authority with a qualified practitioner, or whether it aspires to a higher degree of autonomy that would shift liability models. Even in an assistive role, there exists a risk of automation bias, wherein clinicians or patients may over-rely on the recommendation without critical appraisal. Mitigating this requires embedding mandatory confirmation dialogs, requiring explicit acknowledgment of herb-drug interaction warnings, and maintaining an audit trail that logs every overridden recommendation.

Fairness is a multi-dimensional concern in this setting. The knowledge graph may inadvertently codify biases present in the historical TCM literature, which has predominantly reflected the pathophysiological expressions and dietary practices of East Asian populations. A recommendation engine trained or structured around such a knowledge base might systematically underperform for individuals of other ancestries whose biomarker baselines and syndrome presentations differ. Addressing this requires deliberate diversification of the knowledge graph through targeted curation of evidence from multi-ethnic cohorts and the implementation of fairness constraints during inference. For instance, the system can monitor and adjust for performance disparities across demographic subgroups, ensuring that the probability of receiving a high-confidence recommendation does not systematically vary with

race or ethnicity. Moreover, the digital biomarker algorithms themselves can perpetuate inequities if sensor accuracy differs across skin tones or if socioeconomic factors limit access to the required devices. These infrastructural fairness challenges cannot be solved algorithmically alone but demand a governance framework that incorporates equity impact assessments and inclusive deployment strategies.

Informed consent and data sovereignty are especially salient given the cultural significance of TCM and the deeply personal nature of biomarker data. Many potential users may hold traditional beliefs about the energetic properties of herbs that do not neatly align with biochemical explanations, and the system must respect these perspectives by offering multifaceted rationale displays that include classical TCM references alongside molecular evidence. Data governance policies must also navigate cross-jurisdictional legal frameworks, particularly if knowledge graph curation relies on federated learning across institutions in different countries with divergent privacy regulations. The establishment of a multi-stakeholder ethics advisory board, comprising TCM scholars, biomedical ethicists, patient advocates, and regulatory scientists, is advocated as a structural mechanism to provide ongoing oversight, resolve value conflicts, and guide the evolution of the system's consent and fairness protocols.

7. Implementation, Deployment, and Sustainability Challenges

Translating the proposed framework from a conceptual design into a widely deployed clinical service requires confronting a host of practical implementation challenges that are often underappreciated in academic prototyping. The first of these is the integration into real-world clinical workflows. A TCM herbal recommendation system will only achieve impact if it operates within the tight time constraints of a consultation, whether conducted face-to-face or via telemedicine. This demands user experience design that distills complex biomarker dashboards and graph-based rationales into actionable summary views that a practitioner can interpret in under a minute. Extensive participatory design sessions with both TCM-trained physicians and conventional primary care providers are necessary to refine interface metaphors that bridge the two paradigms without forcing either into an alien conceptual frame.

Long-term sustainability requires a viable economic model and a technical architecture that can be continuously maintained without accumulating prohibitive technical debt. The knowledge graph, in particular, constitutes a living artifact that must be regularly updated as new herbal safety signals emerge, as novel phytochemical mechanisms are elucidated, and as TCM diagnostic standards evolve. The cost of expert curation is substantial and must be justified through evidence that the system improves patient outcomes, reduces adverse events, or enhances clinical efficiency. Cloud-based inference services that charge on a per-consultation basis could provide a revenue stream, but open-source community curation and public-academic partnerships represent alternative models that align with the global public health orientation of chronic disease management. The environmental footprint of AI systems is an additional sustainability dimension; large-scale knowledge graph embedding training and continuous sensor data processing entail significant energy consumption, and the adoption of efficiency-conscious design patterns such as sparse graph attention and quantized neural inference can help align the system with green AI principles [16].

Robustness in the face of real-world data distribution shifts is a critical deployability requirement. A digital biomarker pattern that signified spleen Qi deficiency in a temperate, urban population may not generalize to a tropical, rural population with different dietary staples and circadian rhythms. The system must therefore incorporate out-of-distribution

detection mechanisms that flag when a patient's sensor profile falls outside the training envelope and triggers a fallback to conservative, rule-based recommendations grounded in canonical TCM principles. Model updating must also be carefully orchestrated to avoid catastrophic forgetting or abrupt changes in recommendation behavior that would erode user trust. Versioned, shadow deployment of updated inference models, coupled with continuous monitoring of recommendation concordance and clinical outcomes, provides a governance mechanism for safe iteration.

Beyond technical and economic dimensions, deployment success hinges on cultural acceptance and health system readiness. In regions where TCM is an integral part of the national health system, the digital augmentation of traditional diagnostic processes can be positioned as a modernization initiative that enhances, rather than replaces, existing expertise. In settings where TCM is considered complementary, the system may first need to establish epistemic credibility by demonstrating outcome improvements in randomized pragmatic trials. The framework's emphasis on transparency and explanation serves a dual purpose here: it provides clinicians with the reassurance needed to adopt the tool, and it generates the detailed audit artifacts necessary for regulatory approval and reimbursement decisions. Ultimately, realizing the vision of computationally personalized TCM requires an interdisciplinary coalition that spans informatics, clinical medicine, anthropology, and health policy, capable of sustaining the sociotechnical infrastructure over the decades that chronic disease management demands.

8. Conclusion

This paper has articulated a comprehensive system-level framework that integrates traditional Chinese medicine knowledge graphs with digital biomarkers and explainable artificial intelligence to enable personalized edible herbal formula recommendations for chronic disease management. By grounding the discussion in architectural trade-offs, knowledge engineering pragmatics, and governance imperatives, we have highlighted that the primary challenge is not algorithmic novelty alone but the construction of a resilient, fair, and interpretable infrastructure that can earn the trust of clinicians, patients, and regulators. The proposed layered architecture, with its privacy-preserving biometric pipeline, dynamically curated multi-scale knowledge graph, and inherently interpretable neuro-symbolic inference engine, offers a blueprint that balances the depth of TCM personalization with the evidentiary standards of modern medicine. Critical unresolved issues include the operationalization of fairness across culturally heterogeneous populations, the development of sustainable curation ecosystems capable of absorbing rapid advances in both phytochemistry and sensor technology, and the design of reimbursement models that recognize value created through prevention rather than acute intervention. As digital health and traditional medical systems converge, interdisciplinary collaboration will be the most decisive factor in determining whether such frameworks transition from research prototypes to trusted components of planetary health infrastructure.

References

1. Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44-56.
2. Coravos, A., Khozin, S., & Mandl, K. D. (2019). Developing and adopting safe and effective digital biomarkers to improve patient outcomes. *npj Digital Medicine*, 2, 40.

3. Nicholson, D. N., & Greene, C. S. (2020). Constructing knowledge graphs and their biomedical applications. *Computational and Structural Biotechnology Journal*, 18, 1414-1428.
4. Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE Transactions on Neural Networks and Learning Systems*, 32(11), 4793-4813.
5. Li, X., He, H., Lu, G., Yue, P., Chen, J., Yang, Z., & Hon, C. (2025). TCM-DS: a large language model for intelligent traditional Chinese medicine edible herbal formulas recommendations. *Chinese Medicine*, 20(1), 191.
6. Liu, S. (2025). Revolutionizing healthcare: The role of artificial intelligence in drug discovery and delivery. *Journal of Angiotherapy*, 9(1), 1-16.
7. 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.
8. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347-1358.
9. Onnela, J. P., & Rauch, S. L. (2016). Harnessing smartphone-based digital phenotyping to enhance behavioral and mental health. *Neuropsychopharmacology*, 41(7), 1691-1696.
10. Zhao K, et al. *Phyllostachys nigra* (Lodd ex Lindl) derived polysaccharide with enhanced glycolipid metabolism regulation and mice gut microbiome. *Int J Biol Macromol*. 2024;257:128588.
11. Zhou, X., Chen, S., Liu, B., Zhang, R., Wang, Y., Li, P., ... & Jia, W. (2010). Development and evaluation of an ontology-based knowledge base for traditional Chinese medicine. *AMIA Annual Symposium Proceedings*, 2010, 930-934.
12. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, 119.
13. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25(1), 30-36.
14. Li, S., Zhang, B., Jiang, D., Wei, Y., & Zhang, N. (2012). Network-based elucidation of TCM syndrome patterns. *Evidence-Based Complementary and Alternative Medicine*, 2012, 396830.
15. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
16. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54-63.
17. Murray, E., Burns, J., See Tai, S., Lai, R., & Nazareth, I. (2012). Interactive health communication applications for people with chronic disease. *Cochrane Database of Systematic Reviews*, (1), CD004274.

18. Xue, R., Fang, Z., Zhang, M., Yi, Z., Wen, C., & Shi, T. (2013). TCMID: Traditional Chinese Medicine integrative database for herb molecular mechanism analysis. *Nucleic Acids Research*, 41(D1), D1089-D1095.
19. Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning. fairmlbook.org.
20. Kawamoto, K., Houlihan, C. A., Balas, E. A., & Lobach, D. F. (2005). Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *BMJ*, 330(7494), 765.