Pharmacokinetic–Pharmacodynamic Network Modeling of Herb–Drug Interactions and Combination Safety

Fanlong Dong
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
fan.work@ucf.edu

Wayne L. Allen
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
allen522@oregonstate.edu

Abstract

The concurrent use of herbal medicines and conventional pharmaceuticals has expanded globally, but systematic assessment of herb–drug interactions remains constrained by fragmented pharmacological data, heterogeneous product compositions, and limited mechanistic integration across biological scales. This article presents a system-level analysis of pharmacokinetic–pharmacodynamic network modeling as a framework for anticipating herb–drug interactions and evaluating combination safety. Rather than isolating molecular targets, the perspective treats herb–drug interactions as emergent properties of coupled absorption, distribution, metabolism, excretion, and target-response networks. The discussion examines architectural choices in constructing multilayer networks that integrate chemical constituents, membrane transporters, metabolic enzymes, regulatory pathways, and clinical phenotypes. It emphasizes structural trade-offs between mechanistic depth and computational tractability, between data-driven inference and mechanistic identifiability, and between population-level risk estimation and individualized safety prediction. Because herb–drug interaction models depend on heterogeneous data sources, the article analyzes data infrastructure requirements, metadata standards, provenance tracking, and semantic interoperability. Governance and validation challenges are discussed in relation to regulatory guidance, explainability, auditability, and post-market surveillance. Fairness and equity concerns arising from biased training data, population-specific metabolic variation, and uneven representation in herbal medicine knowledge bases are also addressed. Deployment pathways are considered within clinical decision support, pharmacovigilance, and adaptive learning systems. The article argues that robust, sustainable, and equitable network models require coordinated investment in open data, causal inference methods, model stewardship, and regulatory science. The resulting framework can inform safer polypharmacy involving complex herbal products while supporting evidence-based policy.

Keywords

pharmacokinetic-pharmacodynamic networks; herb-drug interactions; combination safety; network pharmacology; systems pharmacology; clinical decision support.

1. Introduction

The growing global use of herbal medicines alongside prescription drugs has created an urgent need for systematic approaches to evaluate interaction risks and combination safety. Herbal products are chemically complex, often containing hundreds of constituents that may simultaneously influence drug metabolizing enzymes, transporters, receptors, and physiological regulatory circuits. Conventional reductionist methods focus on single active compounds and single interaction mechanisms, but they are insufficient for capturing the emergent behavior of multi-component herbal formulations when co-administered with narrow-therapeutic-index pharmaceuticals. Network pharmacology has emerged as a paradigm that shifts attention from individual molecular targets to the topology and dynamics of biological systems [1]. In parallel, methodological advances in traditional medicine network analysis have made it possible to represent multi-herb formulas as interacting modules whose combined effects cannot be inferred from isolated ingredients [2]. The clinical significance of these interactions depends not only on molecular binding but also on pharmacokinetic processes governing absorption, distribution, metabolism, and excretion, as well as pharmacodynamic processes governing target engagement and downstream response [3]. Regulatory guidance on drug interaction evaluation has emphasized the need for systematic assessment of enzyme and transporter mediated pathways, but it has largely been developed for single-agent pharmaceuticals rather than complex botanical mixtures [4].

Network biology provides a formal language for representing the interconnected molecular determinants of drug action and safety. Biological systems are organized as scale-rich networks in which a small number of hub nodes exert disproportionate influence over system behavior [5]. Systems biology further underscores that physiological responses to pharmacological perturbation are not simple linear outputs but emerge from feedback, redundancy, and modularity [6]. When applied to herb–drug interactions, these concepts imply that a herbal constituent may produce negligible direct inhibition but still alter system state through indirect effects on regulatory hubs, metabolic coupling, or transport competition. Consequently, safety evaluation must account for network-level interactions among multiple constituents, multiple drugs, and patient-specific molecular contexts. The present article develops a system-level perspective on pharmacokinetic–pharmacodynamic network modeling as a strategy for anticipating herb–drug interactions and improving combination safety. It examines architecture, data infrastructure, governance, deployment, sustainability, robustness, fairness, and policy. The aim is not to propose a single mathematical model but to articulate the structural and institutional conditions under which network models can become trustworthy components of clinical and regulatory decision making.

2. Background and Conceptual Foundations

Molecular interaction networks have become central to modern drug discovery because they represent the multiple pathways through which a pharmacological agent may produce therapeutic and adverse effects. The structure and dynamics of molecular networks reveal principles such as modularity, redundancy, and crosstalk that are essential for understanding polypharmacology and drug repositioning [7]. In the context of herbal medicines, a network representation can link individual phytochemicals to their protein targets, metabolizing enzymes, transporters, and downstream transcriptional programs. Such networks provide a substrate for evaluating how herbal constituents modulate the disposition of co-administered drugs. However, the construction of these networks depends on high-quality curated knowledge bases that connect chemical structures, molecular targets, metabolic pathways, and drug interactions. Public resources such as DrugBank supply structured information on drug

targets, enzymes, transporters, and drug interactions, thereby enabling computational models of pharmacological networks [8]. Similarly, KEGG organizes pathway maps for metabolism, signaling, and drug action, which facilitate integration of herbal compound data with endogenous metabolic routes [9].

A fundamental challenge in network-based herb–drug interaction modeling is that the underlying evidence is distributed across experimental literature, clinical case reports, in vitro assay repositories, and pharmacovigilance databases. Integration therefore requires not only data aggregation but also harmonization of semantics, identifiers, and confidence levels. The FAIR principles emphasize that scientific data should be findable, accessible, interoperable, and reusable, and they provide a governance framework for the heterogeneous data environments that support network pharmacology [10]. Without such principles, network models are vulnerable to duplication, missing provenance, and incompatible entity resolution. In practice, herb names, chemical identifiers, target names, and interaction labels must be mapped onto shared vocabularies before they can inform safety reasoning. This semantic infrastructure is as important as the computational algorithms used to infer network relationships.

Pharmacokinetic and pharmacodynamic modeling has a mature methodological tradition that is often underutilized in herb–drug interaction research. Population pharmacokinetic analysis provides a statistical framework for characterizing interindividual variability in drug exposure and response using nonlinear mixed-effects models, and software environments have been developed to support reproducible pharmacometric workflows [11,12]. These approaches allow quantitative estimation of how covariates such as age, organ function, genotype, and co-administered substances influence drug clearance and effect. When coupled with network models, pharmacokinetic–pharmacodynamic frameworks can move from static interaction tables to dynamic predictions of concentration-time profiles and response trajectories under combination exposure. Nevertheless, the transition from single-drug pharmacometrics to multi-component herbal network modeling requires careful treatment of uncertainty, because many herbal constituents lack validated quantitative exposure data. Therefore, conceptual grounding in both systems biology and pharmacometrics is necessary before designing network architectures for combination safety.

3. Pharmacokinetic–Pharmacodynamic Network Modeling Architecture

A system-level architecture for herb–drug interaction modeling can be organized as a multilayer network in which pharmacokinetic, pharmacodynamic, and clinical phenotype layers are coupled. The pharmacokinetic layer represents drug and metabolite transport across biological barriers, plasma protein binding, enzyme-mediated biotransformation, and elimination routes. The pharmacodynamic layer represents target binding, signaling pathway modulation, and physiological end points. The clinical layer represents patient covariates, disease states, concomitant medications, and adverse event phenotypes. Edges across layers encode causal or statistical dependencies, including competitive inhibition, allosteric modulation, transporter saturation, and pathway feedback. Such architectures are consistent with recent network community-based strategies that identify synergistic combinations from herbal medicines and complex systems by detecting modules of interacting constituents and targets [13]. This type of modular representation is useful because herbal products frequently operate through coordinated effects across multiple weak interactions rather than through a single dominant mechanism.

The architecture must balance mechanistic detail against computational tractability and data availability. A fully mechanistic model that represents every molecular interaction for every herbal constituent is neither feasible nor necessary. Instead, models can adopt hierarchical abstractions in which high-confidence metabolic pathways are represented mechanistically, while broader regulatory consequences are represented through network propagation or statistical association. For example, a model may include detailed cytochrome P450 inhibition kinetics for major constituents while using a simplified network neighborhood for downstream signaling effects. This hybrid approach reflects a structural trade-off: increasing mechanistic detail improves biological plausibility and extrapolation to unseen combinations, but it also increases parameter uncertainty and demands more experimental data. Conversely, purely data-driven network inference can capture correlations but may confuse confounding with causation and may fail under distribution shift.

Another architectural decision concerns the level at which combination safety is predicted. Some models operate at the level of individual chemical constituents and extrapolate to whole herbal extracts through summing or interaction terms. Others operate at the level of formula modules, treating an herbal formula as a network community with emergent properties that cannot be reduced to individual components. A module-level perspective is attractive for traditional multi-herb formulations, where synergistic and buffering interactions may be encoded in the formula design. However, modular models require reliable mapping from phytochemical composition to network modules, which is complicated by batch variability, seasonal variation, and processing methods. The architectural choice therefore has direct implications for model validation, because a model that predicts formula-level risk may be easier to calibrate against clinical observations, while a constituent-level model may be easier to interpret mechanistically.

Robustness and modularity are critical architectural principles. A robust network model should not collapse when few data are missing or when a particular constituent cannot be identified. It should degrade gracefully under incomplete phytochemical characterization. Modular approaches allow sensitive components, such as transporter inhibition models, to be updated without requiring re-estimation of the entire system. Furthermore, the architecture should support versioning of both model structure and knowledge graphs, because biomedical knowledge changes rapidly. In deployment, the model may need to be re-trained or re-validated as new pharmacovigilance signals emerge. System-level design should therefore anticipate model life cycle management, including model cards, provenance logs, and explicit statements of intended use. These architectural considerations connect computational modeling to governance and regulatory expectations.

4. Structural Trade-offs in Herb–Drug Interaction Inference

Inference of herb–drug interactions from molecular and clinical data involves fundamental trade-offs between coverage, precision, and mechanistic interpretability. Large-scale pharmacogenomic resources have demonstrated the value of systematic screening across many compounds and cell models, revealing that drug action is often mediated by multiple targets and pathway contexts [14]. Yet such data are typically generated in simplified cellular systems that may not reflect the complex exposure profile of a botanical product. Translating these high-dimensional interaction maps into clinical safety predictions requires careful calibration against human pharmacokinetic and phenotypic data. The structural trade-off is that broad coverage can be achieved by high-throughput molecular profiling, but clinical specificity demands targeted pharmacokinetic studies and well-phenotyped cohorts.

The inference problem is further complicated by the non-random nature of herbal product composition. Unlike single-molecule drugs, herbal extracts vary in chemical content depending on geographical origin, harvesting time, extraction solvent, and manufacturing practices. A model trained on a reference extract may not generalize to a product with different phytochemical fingerprints. This distribution shift resembles the challenge of transferring pharmacological knowledge across biological contexts. One approach is to model the sensitivity of safety predictions to compositional uncertainty by propagating variability through the network. However, this increases computational burden and may require probabilistic representations of network edges. A less computationally demanding alternative is to restrict predictions to well-characterized marker compounds and explicitly acknowledge the residual uncertainty from unmeasured constituents. The choice between these approaches reflects a structural trade-off between robustness to composition variability and the ability to deliver precise quantitative statements.

Causal inference is another important structural dimension. Many network edges in public databases are derived from correlational evidence, literature mining, or high-throughput binding assays. Treating these edges as causal can lead to overconfident interaction predictions. For example, a constituent may bind to a cytochrome P450 enzyme in vitro but have negligible in vivo effect because of low tissue concentration or extensive plasma protein binding. Mechanistic absorption, distribution, metabolism, and excretion models help to contextualize such findings, but their parameters are often unknown for herbal constituents. The inference architecture must therefore distinguish among statistical association, binding affinity, functional inhibition, and clinically meaningful interaction. This distinction can be encoded through edge labels, confidence scores, and causal graphs that specify which variables are expected to influence exposure and response under realistic dosing conditions.

Trade-offs also exist between population-level and individual-level inference. Population models estimate average risk and variance across a reference population, which is valuable for regulatory labeling and general clinical guidance. However, individual risk depends on genetic variants in drug metabolizing enzymes and transporters, kidney and liver function, age, diet, and the specific herbal product. A network model may include covariates to personalize predictions, but adding covariates increases model complexity and data requirements. The resulting architecture must decide which covariates are represented explicitly and which are subsumed into random effects or uncertainty terms. This structural decision influences fairness, because omitting an important covariate can exacerbate disparities in safety prediction for underrepresented groups.

5. Data Infrastructure and Integration Challenges

Data infrastructure is a foundational constraint for network-based herb–drug interaction modeling. Gene-expression signatures have been used to connect small molecules and disease states, demonstrating the value of systematic transcriptional profiling for discovering functional links across compounds [15]. Similar resources for herbal constituents are less mature, partly because botanical products contain complex mixtures and because their components are not always represented in standardized chemical databases. Building a data infrastructure that supports herb–drug interaction modeling requires integrating phytochemical databases, target annotations, pharmacokinetic parameters, clinical reports, and adverse event data. A central challenge is entity resolution: the same herbal species may be referred to by multiple names, and a single constituent may be represented by multiple

chemical identifiers. Without rigorous mapping, network edges will be fragmented and inflated across duplicate nodes.

Another data integration challenge arises from the difference between in vitro potency and in vivo relevance. Plasma protein binding is often cited as a determinant of drug interactions, but its clinical relevance depends on the magnitude of free drug changes, extraction ratio, and therapeutic window [16]. When modeling herbal constituents, protein binding data may be incomplete or measured under non-physiological conditions. The infrastructure must therefore support not only data storage but also the annotation of assay conditions, concentrations, and experimental models. Metadata standards are essential for distinguishing a high-confidence clinical interaction from a preliminary in vitro binding event. Provenance tracking allows modelers to trace every network edge to its original evidence, assign confidence levels, and update predictions when new data emerge.

Moreover, the computational infrastructure should enable reproducible workflows that connect raw data to final safety predictions. This includes versioned knowledge graphs, standardized exposure scenarios, and automated quality checks. Without such infrastructure, network models become difficult to audit, and their predictions cannot be compared across institutions. A federated approach may be necessary when raw patient data cannot be shared across jurisdictions, but federated learning for pharmacovigilance raises additional challenges related to heterogeneity, privacy, and model drift. The data infrastructure must be designed to support both centralized curated resources and distributed learning across clinical sites. Such a design requires careful attention to data governance, security, and semantic interoperability.

6. Governance, Validation, and Regulatory Considerations

Pharmacogenomic variation is a critical source of heterogeneity in drug response and safety, and its incorporation into clinical practice requires robust evidence and decision support [17]. In herb–drug interaction modeling, genetic variation in cytochrome P450 enzymes, transporters, and drug targets can modify the consequences of a given herbal exposure. Governance frameworks must specify how genomic covariates are collected, stored, and used in model validation. Regulatory guidance for drug interactions generally expects that clinically relevant interactions are identified before marketing, but herbal products often enter the market with less pre-market safety evaluation. This asymmetry creates a governance gap: conventional drugs are subject to formal interaction studies, while herbal products may be marketed as dietary supplements with limited oversight. Network models can help bridge this gap only if they are supported by transparent validation standards and regulatory acceptance criteria.

Validation of network-based herb–drug interaction models is complicated by the absence of a single reference standard. Model performance can be assessed against clinical pharmacokinetic studies, case reports, pharmacovigilance signals, or prospective interaction trials, but these sources differ in quality and relevance. A model may perform well in predicting metabolic inhibition in healthy volunteers but fail to predict adverse outcomes in older patients with multiple comorbidities. Validation strategies should therefore include multiple levels: internal data fit, temporal validation on held-out interaction reports, external validation in independent populations, and prospective clinical confirmation. The governance challenge is to define what level of evidence is sufficient for a model to inform clinical decisions. High-performance artificial intelligence in medicine has generated enthusiasm for predictive models, but it has also highlighted the need for rigorous external validation and the risks of overfitting to local practice patterns [18].

Auditability and explainability are central governance requirements. A network model that predicts a serious herb–drug interaction should be able to identify the contributing constituents, pathways, and assumptions. Explainability can take different forms, including network path highlighting, sensitivity analysis, and counterfactual reasoning. However, explanation is not a substitute for validation. A model may provide a plausible mechanistic story that is nevertheless incorrect. Regulatory review should therefore examine both the empiricalassociation between model inputs and observed clinical outcomes and the mechanistic plausibility of the inferred interaction paths. A model that predicts a serious interaction solely on the basis of a high-throughput binding screen without considering exposure, competing metabolic routes, and therapeutic index may be biologically plausible but clinically misleading. Conversely, a model that rejects a true interaction because of sparse pathway annotation may fail to protect vulnerable patients. Governance frameworks should therefore require uncertainty quantification and evidence provenance rather than binary pass/fail assessments. Regulatory review should examine whether the model has been validated in populations that resemble the intended use context, whether the model can be updated without disrupting previously validated behavior, and whether the evidence supporting each network edge is traceable to an auditable source. This form of governance is less about certifying a single model and more about certifying a learning system that remains responsive to new pharmacological and epidemiological knowledge.

A useful governance model for network-based herb–drug interaction prediction distinguishes between model development, model validation, and model deployment. During development, standards should require documentation of training data, knowledge graph versions, entity resolution rules, and parameter estimation procedures. During validation, independent data sets should be used to assess calibration, discrimination, and clinical utility. During deployment, post-market surveillance should monitor whether the model continues to perform as expected in real-world settings. This life cycle perspective is necessary because herbal product markets, patient populations, and prescribing practices change over time. Without continuous monitoring, a model that was accurate at the time of development may become systematically biased as new herbal products enter the market or as clinical practice shifts. Governance should therefore include both pre-deployment review and periodic revalidation.

7. Deployment and Clinical Decision Support

The translation of herb–drug interaction network models into clinical decision support requires attention to sociotechnical integration, workflow constraints, and the temporal dynamics of clinical reasoning. Health information technology operates within complex adaptive healthcare systems, where the introduction of a new predictive tool interacts with local culture, workloads, and care pathways [19]. A model that provides accurate interaction predictions may fail in practice if its outputs cannot be interpreted quickly by prescribers, if alert fatigue leads to overrides, or if the required patient data are not available at the point of decision. Deployment therefore demands not only algorithmic performance but also careful design of user interfaces, default presentation of evidence, and integration with electronic health records. The model should be positioned as decision support rather than decision substitution, especially when herbal product use is poorly documented in routine clinical encounters.

A central deployment challenge is the incompleteness of herbal exposure information. Many patients do not disclose herbal medicine use to their clinicians, and standardized documentation of herbal products is rare in electronic health records. Consequently, network

models may be invoked only when a patient mentions a herbal product, leading to selective deployment and missed interactions. One approach is to build passive screening systems that flag potential herb–drug interactions when a patient's medication list includes entities known to be susceptible to common herbal modulation. However, such systems must avoid overwhelming clinicians with low-probability alerts. The deployment architecture therefore needs a tiered strategy: high-specificity alerts for high-risk combinations and lower-urgency informational modules for combinations with uncertain but plausible risk. This tiering should be informed by model confidence, severity of potential adverse effect, and the therapeutic index of the affected drug.

Moreover, deployment should account for the fact that clinical decisions are made under time pressure and often rely on heuristics. A network model that presents a forest of interacting pathways may be less useful than a model that summarizes the highest-leverage intervention points, such as a strong transporter inhibition by a specific herbal constituent. Yet simplification risks losing the contextual nuance that network modeling provides. A balanced deployment design can offer progressive disclosure: an initial concise safety summary followed by on-demand visualization of the underlying pathway evidence. Such an approach respects clinical workflows while preserving the explanatory depth needed for cases that require detailed assessment. The use of patient-specific covariates, such as renal function and genotype, can further personalize the alert, but this requires access to structured data that may not be uniformly available across healthcare systems. Consequently, deployment strategies must be designed for varying levels of data completeness rather than assuming idealized electronic health record environments.

A further deployment consideration is the distinction between synchronous and asynchronous decision support. Synchronous alerts delivered at the moment of prescribing can interrupt workflow and may be ignored if clinicians are under cognitive load. Asynchronous consultation, in which a clinician reviews a structured interaction report before a medication review, may be more appropriate for complex polypharmacy. Network models can support both modes, but the design of their outputs should differ. A synchronous alert should be simple, actionable, and accompanied by a clear recommendation. An asynchronous report can include richer pathway visualizations, uncertainty summaries, and links to evidence sources. The architectural separation between prediction engine and presentation layer is therefore not merely a technical convenience; it is a sociotechnical requirement that enables the same underlying model to serve different clinical contexts safely.

8. Fairness, Equity, and Population Heterogeneity

Network models for herb–drug interactions inherit the biases present in their training data and knowledge bases. Algorithmic bias in healthcare has been shown to arise when prediction targets encode unequal access to care rather than underlying physiological need [20]. In the context of herb–drug interactions, bias can emerge because the published literature may preferentially report interactions involving widely marketed products or populations enrolled in academic medical centers, while interactions involving traditional medicines used by underserved communities remain underreported. The result may be a model that performs well for common herbal products in well-resourced settings but poorly for culturally specific formulations or for populations with different dietary and metabolic backgrounds. Fairness in this domain requires not merely equal predictive accuracy across groups but also equitable attention to the interaction risks that matter most for each population.

Population heterogeneity in drug metabolizing enzymes and transporters is a well-established source of pharmacokinetic variability. Genetic polymorphisms in cytochrome P450 enzymes vary across ancestral groups, and environmental factors such as diet and co-medication further modulate activity. If network models are trained primarily on data from one population, their predictions may be less reliable when applied to another. For example, a herbal constituent that weakly inhibits a metabolic pathway in a population with high baseline enzyme activity may produce a clinically significant interaction in a population with reduced activity. Representing such context dependence requires explicit covariate modeling and the inclusion of diverse data during model development. However, data diversity alone is insufficient if the underlying knowledge graph lacks annotations for population-specific dietary botanical use.

Fairness also concerns the epistemic exclusion of traditional medicine knowledge. Many herbal combinations arise from systems of medicine that are not well represented in Western biomedical databases. The structural architecture of a knowledge graph can inadvertently exclude these systems by relying on chemical identifiers and molecular targets that have not been characterized for those traditions. As a result, the model may assign low risk to combinations simply because of absent data rather than genuine safety. This is a form of epistemic injustice that emerges from the interaction between formal representation systems and cultural knowledge. Addressing it requires collaborations with traditional medicine practitioners, ethnobotanical researchers, and community stakeholders to expand the represented space of herbal products and their use contexts. It also requires model outputs to distinguish between evidence of safety and absence of evidence, rather than conflating the two.

A further equity dimension concerns the distribution of computational and data resources. High-quality network models may be developed in institutions with access to advanced computing infrastructure and comprehensive pharmacovigilance data, but the patients most at risk of herb–drug interactions may be served by health systems with limited information technology capacity. If model deployment requires proprietary data pipelines or expensive integration efforts, the benefits of network-based safety evaluation will accrue unevenly. Open-source knowledge graphs, standardized application programming interfaces, and low-resource deployment modes can mitigate this disparity. Policy incentives may be needed to ensure that safety models for commonly used traditional medicines are maintained as public goods rather than proprietary assets. The governance of herb–drug interaction models therefore intersects with broader debates about digital health equity and global access to medical knowledge.

9. Sustainability, Maintenance, and Policy Implications

Sustaining a network-modeling infrastructure for herb–drug interactions requires continuous curation, versioned updates, and active model stewardship. Biomedical knowledge graphs degrade over time as new enzymes, transporters, and interaction mechanisms are discovered. A model that is not updated will increasingly diverge from current pharmacology and may produce outdated safety recommendations. Stewardship also requires transparent reporting of model assumptions, data limitations, and the boundaries of valid use. The current enthusiasm for explainable artificial intelligence in medicine must be tempered by evidence that explanation quality does not automatically translate into improved clinical decisions [21]. Model owners should therefore document not only what the model predicts but also what evidence supports each network edge and how uncertainty is represented.

Model maintenance also requires pharmacovigilance feedback loops. Post-market surveillance systems can generate signals of adverse events associated with herb–drug combinations, but these signals are often weak and confounded [22]. Network models can be updated using such signals as external validation or as triggers for targeted mechanistic investigation. Conversely, model predictions can be used to prioritize pharmacovigilance monitoring for combinations that are predicted to carry high risk but for which clinical reports are sparse. This bidirectional relationship between model and surveillance system improves both over time. However, the integration of spontaneous reporting data into model updates must account for reporting biases, duplicate reports, and the absence of a reliable denominator for exposure frequency. A governance approach that combines quantitative signal detection with expert review is more likely to maintain model reliability.

Policy implications extend beyond clinical deployment to regulatory science and public health. Regulatory agencies have begun to explore frameworks for adaptive artificial intelligence in medical products, but most existing guidance is not tailored to complex network models of herb–drug interactions. The development of performance standards, audit criteria, and post-market monitoring expectations would support responsible innovation. Precision medicine initiatives have emphasized the need to tailor prevention and treatment to individual variability, and herb–drug interaction models can contribute to this goal by informing personalized medication reviews [23]. However, policy must also address the regulatory status of herbal products, which varies widely across jurisdictions. In some regions, herbal products are regulated as medicines with formal safety requirements; in others, they are dietary supplements with limited pre-market review. Network models may be used to prioritize which products require enhanced post-market surveillance or formal interaction studies.

Finally, the reproducibility crisis in science illustrates that findings may fail to replicate when methods, data processing, and analytical choices are not transparently reported [24]. Network models are similarly vulnerable to irreproducibility if their training data, code, and provenance records are not maintained. Sustainability therefore involves institutional commitment to long-term funding, open data sharing, and governance of model updates. Without such commitment, the field may produce a proliferation of short-lived predictive tools that cannot be compared, replicated, or trusted. A coordinated effort to establish common data standards, benchmark tasks, and evaluation protocols would improve the scientific basis of herb–drug interaction modeling. This effort should include not only computational researchers but also pharmacologists, clinicians, ethnobotanists, and regulatory scientists.

10. Conclusion

Pharmacokinetic–pharmacodynamic network modeling provides a powerful conceptual and computational framework for anticipating herb–drug interactions and improving combination safety at the systems level. The value of this framework lies not in generating a single prediction for a single combination but in making the complex web of molecular, physiological, and clinical dependencies explicit and queryable. The resulting models can support evidence synthesis, prioritize experimental studies, inform clinical decision support, and guide pharmacovigilance. Yet the construction and deployment of such models involve fundamental trade-offs among mechanistic depth, data availability, computational tractability, interpretability, and fairness. These trade-offs cannot be resolved by algorithmic choices

alone; they require coordinated investment in data infrastructure, governance, and regulatory science.

A robust system for herb–drug interaction modeling must treat model architecture as an ongoing design problem rather than a static technical artifact. It must support versioned knowledge graphs, traceable evidence, uncertainty propagation, and modular updates. It must also remain responsive to the social and cultural dimensions of herbal medicine use, ensuring that traditional knowledge systems are represented without being reduced to Western molecular categories. The inclusion of diverse populations and the explicit modeling of population-specific pharmacokinetic variation are essential for equitable safety predictions. Deployment in clinical environments requires attention to workflow integration, alert design, and the sociotechnical realities of healthcare organizations.

Ultimately, the safe combination of herbal medicines and conventional drugs depends on an infrastructure that connects molecular evidence, clinical observation, and regulatory oversight in a continuous learning cycle. Network models are one component of this infrastructure, but they cannot replace clinical judgment or robust pharmacovigilance. Their greatest contribution may be in making the systemic consequences of polypharmacy more legible, thereby enabling clinicians, regulators, and patients to engage in more informed deliberation about risk. The path forward requires interdisciplinary collaboration, open data standards, and a commitment to producing models that are not only predictive but also fair, auditable, and sustainable.

References

1. Hopkins, A. L. (2008). Network pharmacology: the next paradigm in drug discovery. Nature Chemical Biology, 4(11), 682-690.

2. Li, S., & Zhang, B. (2013). Traditional Chinese medicine network pharmacology: theory, methodology and application. Chinese Journal of Natural Medicines, 11(2), 110-120.

3. Rowland, M., & Tozer, T. N. (2011). Clinical pharmacokinetics and pharmacodynamics: concepts and applications (4th ed.). Lippincott Williams & Wilkins.

4. U.S. Food and Drug Administration. (2020). Clinical drug interaction studies — cytochrome P450 enzyme- and transporter-mediated drug interactions guidance for industry. U.S. Department of Health and Human Services.

5. Barabasi, A. L., & Oltvai, Z. N. (2004). Network biology: understanding the cell's functional organization. Nature Reviews Genetics, 5(2), 101-113.

6. Kitano, H. (2002). Systems biology: a brief overview. Science, 295(5560), 1662-1664.

7. Yildirim, M. A., Goh, K. I., Cusick, M. E., Barabási, A. L., & Vidal, M. (2007). Drug-target network. Nature Biotechnology, 25(10), 1119-1126.

8. Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., ... Wilson, M. (2018). DrugBank 5.0: a major update to the DrugBank database for 2018. Nucleic Acids Research, 46(D1), D1074-D1082.

9. Kanehisa, M., Furumichi, M., Tanabe, M., Sato, Y., & Morishima, K. (2017). KEGG: new perspectives on genomes, pathways, diseases and drugs. Nucleic Acids Research, 45(D1), D353-D361.

10. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data, 3, 160018.

11. Sheiner, L. B., & Beal, S. L. (1980). Evaluation of methods for estimating population pharmacokinetic parameters. I. Michaelis-Menten model: routine clinical pharmacokinetic data. Journal of Pharmacokinetics and Biopharmaceutics, 8(6), 553-571.

12. Bauer, R. J. (2019). NONMEM tutorial part I: description of commands and options, with simple examples of population analysis. CPT: Pharmacometrics & Systems Pharmacology, 8(8), 525-537.

13. Wang, Y., Yao, J., Sui, Y., Jiang, H., Ma, B., Lai, S., ... & Tan, N. (2026). HerbSyner_Finder: a network community-based model for identifying synergistic combinations from herbal medicines and complex systems. Targetome, 2(2).

14. Iorio, F., Knijnenburg, T. A., Vis, D. J., Bignell, G. R., Menden, M. P., Schubert, M., ... Garnett, M. J. (2016). A landscape of pharmacogenomic interactions in cancer. Cell, 166(3), 740-754.

15. Lamb, J., Crawford, E. D., Peck, D., Modell, J. W., Blat, I. C., Wrobel, M. J., ... Golub, T. R. (2006). The Connectivity Map: using gene-expression signatures to connect small molecules, genes, and disease. Science, 313(5795), 1929-1935.

16. Benet, L. Z., & Hoener, B. A. (2002). Changes in plasma protein binding have little clinical relevance. Clinical Pharmacology & Therapeutics, 71(3), 115-121.

17. Relling, M. V., & Evans, W. E. (2015). Pharmacogenomics in the clinic. Nature, 526(7573), 343-350.

18. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... Dean, J. (2019). A guide to deep learning in healthcare. Nature Medicine, 25(1), 24-29.

19. Sittig, D. F., & Singh, H. (2010). A new sociotechnical model for studying health information technology in complex adaptive healthcare systems. Quality and Safety in Health Care, 19(Suppl 3), i68-i74.

20. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447-453.

21. Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. The Lancet Digital Health, 3(11), e745-e750.

22. Bate, A., & Evans, S. J. W. (2009). Quantitative signal detection using spontaneous ADR reporting. Pharmacoepidemiology and Drug Safety, 18(6), 427-436.

23. Collins, F. S., & Varmus, H. (2015). A new initiative on precision medicine. New England Journal of Medicine, 372(9), 793-795.

24. Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. Nature, 533(7604), 452-454.