Check for updates

International Journal of Advanced Research and Interdisciplinary Scientific Endeavours, Vol. 4(11), 2026
DOI: 10.61359/11.2206-2652

Artificial Intelligence Applications in Ayurveda: A Systematic Review

Ankit Raj [*1], Ranjana M [†2], Girish Nath Jha [‡3]

Email Correspondence*: ankitraj0624@gmail.com

[1] Research Scholar, School of Sanskrit and Indic Studies, Jawaharlal Nehru University, New Delhi, India.

[2] Research Scholar, Centre for Studies in Science Policy, School of Social Sciences, Jawaharlal Nehru University, New Delhi, India.

[3] Professor, School of Sanskrit and Indic Studies, Jawaharlal Nehru University, New Delhi, India.

Abstract:

Artificial intelligence is applied to *Ayurveda* in a growing literature, but no systematic review has mapped what that literature contains, which computational methods it uses, or how reliably it reports them, and existing reviews of AI in traditional medicine treat *Ayurveda* alongside Traditional Chinese Medicine and other systems, where it is a small minority of the evidence. This review searched Scopus, PubMed, Web of Science and Google Scholar from database inception to August 2026 for studies applying AI or machine-learning methods to *Ayurveda* or Indian traditional medicine. Records were deduplicated with a four-pass deterministic cascade (DOI, PMID, normalized title, title prefix) followed by manual review of multi-signal fuzzy-matched pairs. Eligibility required an explicit AI or machine-learning method (C1), *Ayurveda* or Indian traditional medicine as the study domain (C2), and at least one of five predefined application domains (C3). Studies whose methodological contribution was recognition of plant, leaf or species from images were excluded, including those that added a therapeutic lookup to the recognized species. Every included study was read in full and coded on ten reporting items, and the 14 studies that predict an individual's disease or constitution were assessed for risk of bias with PROBAST+AI. Screening, extraction and coding were carried out by a single reviewer, the protocol was not registered, and the review is reported according to PRISMA 2020. Of 1,253 records, 635 remained after screening and deduplication, and 42 studies met all criteria (38 from database searching, 4 from hand-searching); 38 of the 42 appeared from 2024 onward. Disease or symptom prediction with therapeutic linkage was the most common domain (22 studies), and drug discovery, phytochemistry and molecular modelling were the least common (8). Scopus returned all 38 database-traced studies, and no other database retrieved the included study that Scopus had missed. Image-based plant recognition accounted for 156 of 597 excluded records, more than any other exclusion reason. Full-text assessment and re-checking removed seventeen studies that had passed title and abstract screening: eight with no learning method, four reviews, three in which *Ayurveda* was not the study domain and two that reported no results, one of the last carrying undisclosed machine-generated text alongside unverifiable references. Eight of the 42 studies showed none of the ten reporting defects and fifteen showed four or more. Molecular-modelling studies carried fewer defects than the rest (mean 1.12 against 2.91, two-sided Mann-Whitney $p = 0.020$), but seven of the eight were journal articles, and the difference could not be separated from publication venue (adjusted rate ratio 0.63, 95% CI 0.29–1.38). All 14 prediction-model studies were at high risk of bias, most often because of their data sources and analysis. AI work on *Ayurveda* is concentrated in symptom-to-remedy recommendation systems, often built on datasets whose

*Research Scholar, School of Sanskrit and Indic Studies, Jawaharlal Nehru University, New Delhi, India.
†Research Scholar, Centre for Studies in Science Policy, School of Social Sciences, Jawaharlal Nehru University, New Delhi, India.
‡Professor, School of Sanskrit and Indic Studies, Jawaharlal Nehru University, New Delhi, India.

Article Timeline: Received: Nov 06, 2026; **Revised:** Nov 10, 2026; **Accepted:** Nov 10, 2026; **Published:** Nov 11, 2026

size is not reported, and is thinly represented in compound-level discovery; much of it is published in venues not indexed by the major bibliographic databases, and fourteen screened records retrieved as AI research contained no learning component.

Keywords: *Ayurveda*; Artificial Intelligence; Machine Learning; Systematic Review; PRISMA 2020; Reporting Quality; PROBAST+AI; Research Integrity; Indian Traditional Medicine.

1. Introduction

Ayurveda documents several thousand plant drugs and compound formulations across more than 150 classical texts, organized by principles that have no direct equivalent in biomedical pharmacology: *rasapañcaka* (the five pharmacological attributes of a drug), action on *doṣa*, *dhātu* and *mala*, and indication by disease entity rather than by molecular target. That knowledge is largely unstructured. It sits in Sanskrit and regional-language texts, in pharmacopeial monographs, and in clinical practice, in forms that resist computational analysis.

Machine learning has been proposed as the route to structuring it. Over the past three years the number of papers applying machine learning, deep learning, knowledge graphs and large language models to *Ayurvedic* material has risen sharply, spread across engineering conference proceedings, pharmacology journals and Indian specialist journals. The work ranges from symptom-to-herb recommendation systems to graph databases linking formulations, plants, phytochemicals and diseases, and to machine-learning-guided virtual screening of phytochemicals against protein targets.

What this literature contains has not been established. Two prior reviews cover adjacent ground. Chu et al. (2022) scoped AI in complementary and alternative medicine across PubMed, Embase and the Cochrane Library, retrieving 1,114 records and including 32 studies, of which *Ayurveda* was a small share. A later narrative review of AI-titled traditional medicine drug discovery retrieved 2,128 records from five sources and retained 12 studies, again dominated by Traditional Chinese Medicine (Chen et al., 2026). Neither was restricted to *Ayurveda*, and neither used Scopus as a primary source, which matters because Indian engineering conference proceedings are indexed there and poorly indexed elsewhere.

Two further problems make the field hard to read. The first is that "AI" is applied loosely: papers titled as artificial intelligence research sometimes describe a database query behind a web form. The second is that a large body of work applies deep learning to photographs of medicinal plant leaves. Species recognition from images is a computer-vision task on a plant specimen, and it tells us nothing about a drug, a formulation or a therapeutic relationship, yet it shares vocabulary and venues with work that does.

This review asks what artificial intelligence has been applied to in *Ayurveda*, which methods have been used, how adequately they are reported, and where the gaps lie. We set out eligibility criteria that separate knowledge-level, therapeutic and compound-level applications from image-based taxonomy, applied them across four databases, and verified the surviving studies against their full texts wherever these could be obtained.

2. Methods

2.1 Protocol and reporting

The review follows PRISMA 2020 (Page et al., 2021) for searching, screening and reporting. The eligibility criteria were fixed before screening and are reproduced in full in Table 2. Where a criterion was refined during screening, the refinement and its effect on the included set are reported in section 3.4 rather than

presented as the original rule. The protocol was not registered in PROSPERO or any other registry, and no protocol was published before the review began. Screening, full-text assessment, data extraction and coding of the reporting items were all carried out by one reviewer (the first author). With no second reviewer there were no disagreements to resolve; where a decision was uncertain, the full text was re-read and the decision and its reason were recorded in the screening workbook (Supplementary File S1).

2.2 Information sources and search strategy

We searched Scopus, PubMed, Web of Science and Google Scholar from database inception to August 2026. The search combined an artificial-intelligence concept block with an *Ayurveda* and Indian traditional medicine block, adapted to each platform's syntax (Table 1).

Table 1 - Search strings and records retrieved by database (searches run to August 2026).

Database	Search string	Records
Scopus	TITLE-ABS-KEY("artificial intelligence" OR "machine learning" OR "deep learning" OR "neural network*" OR "natural language processing" OR "knowledge graph" OR "large language model") AND TITLE-ABS-KEY(ayurved* OR dravyaguna OR rasashastra OR "Indian traditional medicine" OR "Indian system of medicine" OR siddha OR unani)	427
PubMed	("artificial intelligence"[tiab] OR "machine learning"[tiab] OR "deep learning"[tiab] OR "neural network*"[tiab] OR "knowledge graph"[tiab]) AND (ayurved*[tiab] OR "Indian traditional medicine"[tiab] OR dravyaguna[tiab] OR siddha[tiab] OR unani[tiab])	517
Web of Science	TS=("artificial intelligence" OR "machine learning" OR "deep learning" OR "neural network*") AND TS=(ayurved* OR "Indian traditional medicine" OR dravyaguna OR siddha OR unani)	49
Google Scholar	allintitle: ("artificial intelligence" OR "machine learning" OR "deep learning") (ayurveda OR ayurvedic OR "Indian traditional medicine")	260
Total		1,253

Searches were restricted to journal articles and conference papers. No language limit was applied at the search stage; non-English records were excluded at screening. Open-access status was recorded as a study variable rather than applied as a filter. An earlier version of this search did apply an open-access limit, which dropped database-traced studies disproportionately and pushed most of the included set into the hand-searched arm. Because that distorts the PRISMA flow, the limit was removed.

Reference lists of included studies were hand-searched, and citation tracking was used to identify studies in venues the databases do not index.

2.3 Deduplication

Duplicates were identified with a script written in Python 3.13, using pandas (version 3.0.5) to handle the records, RapidFuzz (version 3.14.6) to compute string similarity, the standard-library unicodedata module to normalize text, and openpyxl (version 3.1.5) to export the results. Titles were first normalized by removing publication-type tags, stripping diacritics and reducing each title to lowercase letters and digits.

The 728 records retained after screening were then matched in four deterministic passes: exact DOI (70 duplicates), PMID (none beyond those already matched by DOI), normalized title (21), and a title prefix of at least 40 characters where one title was a truncated form of the other and the two records shared a first author or a publication year within one year (1). Matches were grouped transitively with a union-find structure, so that a record linked to two others in different passes formed a single cluster. Records not yet matched were then compared in a fuzzy pass. In total 93 duplicates were removed, leaving 635 unique records. Deduplication followed screening because each database export was screened as it was retrieved. A record was carried forward if any of its copies passed screening, so screening before deduplication meant that some duplicate copies were assessed more than once, but it could not remove a study that one copy would have admitted. For reference, running the same script on all 1,253 records before screening removes 155 duplicates and leaves 1,098 unique records (Supplementary File S2).

2.4 Eligibility criteria

Table 2 - Eligibility criteria and exclusion codes.

Criterion	Include	Exclude
Relevant publications	Original studies applying AI or machine-learning methods to *Ayurveda* or Indian traditional medicine for a knowledge-level, therapeutic or compound-level purpose	AI applied to *Ayurveda* with no knowledge-level, therapeutic or compound-level outcome
Publication type	Peer-reviewed primary research, conference proceedings, publicly accessible preprints; no year restriction	Reviews, meta-analyses, editorials, commentaries, opinion pieces, letters, perspectives, unless containing original computational data
Study framework	Quantitative and computational designs, including in silico studies with or without experimental validation	Qualitative studies, case reports, case series
C1 AI methodology	Explicit machine-learning, deep-learning, knowledge-graph, natural-language-processing, large-language-model or computational-learning method, or reported model training, validation or performance metrics	No AI or machine-learning component; conventional statistics alone; fuzzy logic alone
C2 Traditional medicine system	*Ayurveda* as the study domain, including its component disciplines *Dravyaguna* and *Rasashastra*, evidenced by classical terminology, Indian pharmacopeial sources or Indian medicinal plant and formulation subject matter	Other AYUSH systems, namely *Unani*, *Siddha*, Homeopathy, Naturopathy and *Yoga*, where *Ayurveda* is not the study domain; Traditional Chinese Medicine, *Kampo*, Korean medicine, *Sowa Ripa*, *Jamu* or other non-Indian systems; modern pharmacology only
C3 Application domain	At least one of five domains: (A1) drug discovery, phytochemistry and molecular modelling; (A2) formulation, prescription and drug-herb interaction modelling; (A3) knowledge bases, ontologies and knowledge graphs; (A4) personalized medicine and	Image-based plant, leaf or species recognition where recognition is the methodological contribution, including studies appending a therapeutic, utility or information-retrieval lookup to the recognized species

	prakriti modelling; (A5) disease or symptom prediction with therapeutic linkage	
Language	English	Non-English
Availability	Full text available; record long enough to assess	Abstract only; record under 250 characters, leaving too little content to judge against the criteria (distinct from X8, where a substantive record passed screening but no full text could be obtained)

C2 was scoped to *Ayurveda* specifically rather than to AYUSH as a whole. *Unani*, *Siddha*, Homeopathy, Naturopathy and *Yoga* are distinct systems with their own diagnostic and therapeutic frameworks, and studies in which one of those systems rather than *Ayurveda* was the study domain were excluded under X3. One record otherwise meeting C1 and C3 was excluded on this basis: a study applying pulse-waveform signal processing within *Unani* medicine, in which *Ayurveda* is not mentioned (Waqar et al., 2025). A check of all included studies confirmed that no other record has a non-*Ayurvedic* system as its domain, though one includes *Unani* and *Siddha* among incidental references to AYUSH.

The exclusion for image-based recognition needs explanation, because it removed more records than any other criterion. Recognizing a species from a photograph and then returning that species documented uses from a lookup table adds no modelled relationship between a drug and a disease. The recognition is a computer-vision result about a specimen, and the therapeutic content is a static mapping that existed before any model was trained. We therefore excluded both pure recognition studies (coded X1a) and recognition studies with a lookup layer appended (X1b) and recorded the two separately so that either can be reinstated by a reader who draws the boundary differently.

Fuzzy logic alone was treated as ineligible under C1. Chu et al. (2022) excluded fuzzy-logic studies on the grounds that the method converts a linguistic variable and is not itself a learning technique. We applied the same rule because a hand-specified fuzzy rule base does not learn from data.

2.5 Screening

Screening ran in two stages. Titles and abstracts were assessed against C1, C2 and C3, with exclusion coded to the first criterion failed, in the order X1 (image recognition), X2 (non-primary literature), X3 (not *Ayurveda*), X4 (no AI component), X5 (no eligible application domain), X6 (insufficient content to assess). Records whose assessable text ran under 250 characters, almost all truncated Google Scholar snippets, were coded X6 rather than judged on partial information.

Full texts were then retrieved for studies that passed and for records whose eligibility could not be settled from the abstract. Full-text assessment applied the same criteria and reversed seventeen earlier inclusions, which are itemized in section 3.4. Two further codes were added during full-text assessment: X7 for records reporting no original results, and X8 for records whose full text could not be obtained from any source.

2.6 Data extraction

For each included study we extracted bibliographic details (authors, year, venue, country, DOI, source database, open-access status), the traditional medicine system and the *Ayurvedic* knowledge source used, the AI method and named algorithms, the application domain, dataset provenance and size, the validation

approach, reported performance, and the study's principal output. Data were extracted by one reviewer into a structured workbook (Supplementary File S1).

2.7 Appraisal and synthesis

The included studies vary too widely in design, data and outcome for meta-analysis, so results are reported in tables and summarized narratively. Reporting was appraised on ten items, coded from the full text of every included study (Table 6). The starting point was the four dimensions used by Chen et al. (2026) in their narrative review of AI-titled traditional medicine discovery: a named algorithm, stated dataset provenance and size, a described validation approach, and performance reported as a held-out or cross-validated estimate rather than an author or practitioner judgement. The remaining six items were added because they recurred during full-text reading. A study is coded for an item only where its full text gives explicit evidence. For the validation item, evaluation on a stated set of test cases, such as a fixed set of queries or case scenarios, was counted as a reported validation approach. The items overlap with those of TRIPOD+AI (Collins et al., 2024) and TRIPOD-LLM (Gallifant et al., 2025), which also ask for dataset size, the validation approach and performance on data not used for training; adherence to either guideline was not scored, because most included studies are not prediction-model studies. These items describe reporting. They are not a risk-of-bias assessment, which was made separately for the prediction-model studies.

Risk of bias and applicability were assessed with PROBAST+AI (Moons et al., 2025) for the 14 studies that develop or evaluate a model predicting an individual's disease or constitution. Both parts of the tool were applied: model development (16 signaling questions, rating quality) and model evaluation (18 questions, rating risk of bias), each across four domains (participants and data sources, predictors, outcome, analysis), with applicability rated for the first three.

The other 28 studies (knowledge bases, molecular modelling, and retrieval or recommendation systems) have no predicted health outcome, and no risk-of-bias tool was applied to them.

Two integrity checks were added after anomalies emerged during full-text reading. We checked each full text for text indicating undisclosed machine generation, and we attempted to locate the cited venues of references in studies whose reference lists appeared generic. Findings are reported in section 3.4.

Defect counts were compared between groups with two-sided Mann-Whitney U tests, and proportions with two-sided Fisher's exact tests. Because studies could carry more than one domain tag, comparisons were made between disjoint groups (a domain against all studies without that tag), and comparisons among the five domains are descriptive. Because publication type could confound the domain comparison, a Poisson regression of defect count on the A1 tag and publication type (journal or conference) was fitted, with rate ratios and 95% confidence intervals reported, and a negative binomial model was fitted as a check on overdispersion. Analyses were run in Python with SciPy and statsmodels.

3. Results

3.1 Study selection

The four databases returned 1,253 records. Title and abstract screening excluded 525, of which 176 were reviews, editorials, letters or retracted items and 349 failed one of the three criteria. Deduplication removed 93 duplicates from the 728 records that remained, leaving 635 unique records for assessment against the full criteria. Of these, 597 were excluded, 595 on the criteria and 2 because no full text could be obtained, and 38 were eligible (Fig. 1).

Hand-searching and citation tracking identified a further 12 records in venues the databases do not index. Four were excluded at record level as not *Ayurveda*-specific, seven were retrieved in full text, of which four met all criteria and three did not, and one could not be obtained and was excluded under X8. Forty-two studies are included.

The largest single exclusion category is image-based recognition of plants, leaves or species, at 156 records, or 26% of all exclusions. Records too short to assess, almost all truncated Google Scholar snippets, account for a further 188.

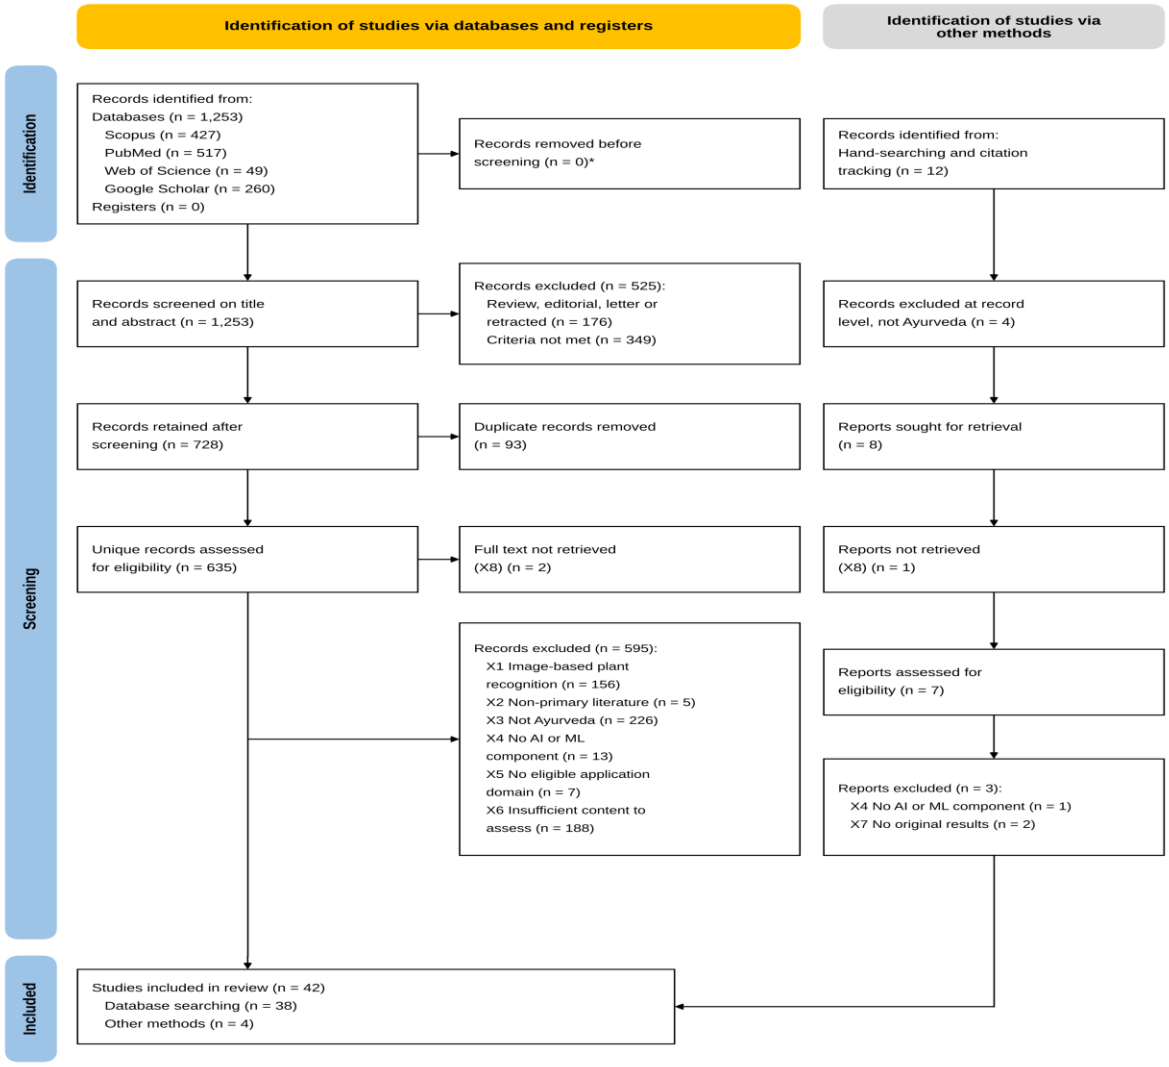


Fig. 1 - PRISMA 2020 flow diagram of study selection. *Duplicates were removed after title and abstract screening (n = 93), so no records were removed before screening. X codes are the exclusion reasons defined in Table 2.

3.2 Characteristics of included studies

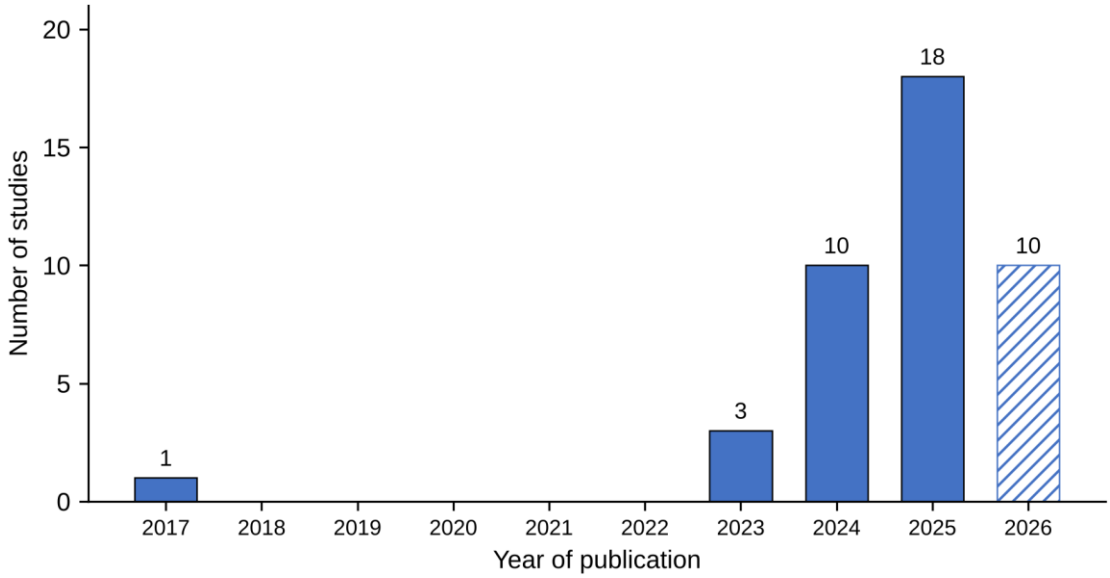


Fig. 2 - Included studies by year of publication (n = 42). The hatched 2026 bar covers January to August only.

One included study predates 2021, and 38 of the 42 appeared from 2024 onward (Fig. 2). The 2026 bar covers January to August only, so the drop from 2025 reflects the search date and says nothing yet about output for the full year.

Conference proceedings outnumber journal articles, 28 to 14 (Table 3). Most of the proceedings are Indian engineering conferences indexed by Scopus and by no other database in this search. Work is overwhelmingly Indian in origin, by author affiliation.

Table 3 - Publication type and database source of the 42 included studies.

Characteristic	Studies (n = 42)
Published 2024 or later	38
Conference proceedings	28
Journal articles	14
Retrieved by Scopus	38
Retrieved by Google Scholar	9
Retrieved by PubMed	6
Retrieved by Web of Science	2
Identified only by hand-searching	4

The retrieval figures are unusual for a four-database search. Every one of the 38 database-traced studies was returned by Scopus, and 24 of them by Scopus alone. No study was found by PubMed, Web of Science or Google Scholar that Scopus had missed. Web of Science returned two included studies and contributed neither uniquely.

3.3 Application domains and methods

Studies could be assigned to more than one domain, giving 77 domain tags across the 42 studies (Table 4).

Table 4 - Included studies by application domain (n = 42). Studies may carry more than one domain, so shares sum to more than 100%.

Application domain	Studies	Share of 42
A5 Disease or symptom prediction with therapeutic linkage	22	52%
A4 Personalized medicine and *prakriti* modelling	17	40%
A2 Formulation, prescription and drug-herb interaction modelling	16	38%
A3 Knowledge bases, ontologies and knowledge graphs	14	33%
A1 Drug discovery, phytochemistry and molecular modelling	8	19%

The field is concentrated at the clinical end. Half the studies predict a disease or classify a symptom pattern and return a herb, formulation, diet or regimen. The typical design takes free-text or checkbox symptoms, classifies them, and looks up a remedy from a curated table. Compound-level work is the smallest category, at 8 studies, about a fifth of the included set.

That ordering is the opposite of what Chen et al. (2026) found, where drug discovery dominated. The difference is explained by scope: that review required a drug-discovery component for inclusion and was dominated by Traditional Chinese Medicine, where the compound-level literature is far larger.

Methods divide into recognizable families, counted here over all 42 included studies. A study may fall into more than one family, so the counts do not sum to 42. Classical supervised learning is the most common, in 17 studies, typically random forests, support vector machines, decision trees, naive Bayes or gradient boosting on tabular symptom data. Deep learning appears in 10, large language models and retrieval-augmented generation in 9, natural-language processing and text mining in 7, and knowledge graphs or ontologies in 7. Three studies use molecular docking, molecular dynamics, QSAR or ADMET prediction, and only one of those combines such methods with a trained learning model. The techniques that need an external molecular referent are the least used, and those that can run on a symptom table without one are the most used.

Several studies are worth naming for what they establish. The largest by sample is a Parkinson's disease prediction model built on the Fox Insight cohort of 80,916 records, using kernel SVM, logistic regression and gradient boosting with 10-fold cross-validation and a best accuracy of 92.6% (Majhi et al., 2023). Against that, the median dataset across the remaining studies is in the hundreds, and many report no sample size at all. One study learns *Ayurvedic* potency, the *śīta* and *uṣṇa* classification of *vīrya*, from molecular descriptors of phytochemicals using a gradient-boosting classifier trained on a 627-herb reference set built from classical texts (Panghalia & Singh, 2026). Two graph resources integrate formulations, plants, phytochemicals and diseases at scale, one connecting 1,039 formulations to 12,743 plants, 129,542 phytochemicals and 13,480 diseases (S. Joshi et al., 2026), the other curating 27,440 phytochemicals from 6,959 plants with ADMET properties and deep-learning compound classification (Kiewhuo et al., 2023).

Table 5 - Characteristics of the 42 included studies, ordered by primary application domain. "Not reported" means the item is absent from the full text.

Study	TM system	AI method	Task	Data source(s)	Validation	Key findings
A1 Drug discovery, phytochemistry and molecular modelling						
Kiewhuo et al. (2023), OSADHI	Indian medicinal plants, traditional knowledge	Random Forest, XGBoost, NPClassifier deep-learning framework, in-house ADMET scripts	Build a pan-India phytochemical database with predicted properties and antiviral potency	6,959 plants across 348 families, 28 states and 8 union territories; 27,440 phytochemicals curated, 22,314 reported	Accuracy of 88% cited from the authors' earlier publication; check against known antivirals	Database integrates traditional knowledge, geographical indication, phytochemistry and chemoinformatics; compounds classified by class, superclass and pathway; antiviral potency predicted
Shi et al. (2025)	*Ayurveda*-inspired	GCN, GraphSAGE, graph attention network; centrality measures	Predict phytochemical-protein interactions for epilepsy	Bipartite graph of phytochemicals and proteins	Link-prediction evaluation	Phytochemical-protein interaction modelled as link prediction on a bipartite graph for epilepsy targets
Kumar et al. (2025)	*Ayurveda* formulation source plant	Artificial neural network	Identify elite genotypes of *Nilgirianthus ciliatus*	GC-MS profiles and genetic markers from Western Ghats accessions	Train and test sets reported	ANN reached MSE 2.43e-02 and R 0.99992; NC 10 yielded 793.0 ng squalene; 140 marker bands, 115 polymorphic (82.14%)
Nikhil et al. (2025)	Indian medicinal plants, *Ayurvedic* literature	RAG with BioBERT and Mistral; Neo4j knowledge graph; FAISS and Pinecone vector stores	Semantic search and drug-candidate prediction	Taxonomic and phytochemical datasets, *Ayurvedic* literature, biomedical research	BERTScore on a single example query; no evaluation set	RAG BERTScore F1 0.789 and LLM BERTScore F1 0.779 for therapeutic-use queries
Niharika et al. (2026)	*Ayurveda* (*Jeevaneey*	ML-QSAR with SVR, NuSVR,	Identify NLRP3	*Jeevaneeya Rasayana*	Internal q2 0.679, r2	SVR predicted pIC50 reliably

	a Rasayana)	LGBMR, HGBR via LazyPredict and scikit-learn; docking; MD; DFT	inflammasome inhibitors for epilepsy	phytochemicals	training 0.797, RMSE 0.454; Golbraikh-Tropsha external criteria; Y-randomisation; applicability domain	and met all external validation criteria; the most completely validated model in the review
S. Joshi et al. (2026), GRAYU	*Ayurveda*	Graph database and knowledge-graph integration	Integrate formulations, plants, phytochemicals and diseases for repurposing	1,039 formulations, 12,743 plants, 129,542 phytochemicals, 13,480 diseases	Not reported	1,370,257 plant-phytochemical, 116,531 plant-disease, 2,389 plant-formulation and 4,087 formulation-disease associations; applied to analogs, plant substitution and disease networks
Panghalia and Singh (2026), SitoshnaPred	*Ayurveda* (*vīrya*)	LightGBM ensemble classifier on molecular descriptors	Predict *śīta* and *uṣṇa* potency of herbs from their phytochemistry	Gold-standard set of 627 herbs built from classical textbooks	Binary classifier evaluation	Herbal potency predicted from phytochemical descriptors, giving a molecular account of a classical *rasapañcaka* attribute
Sil et al. (2026)	*Ayurveda* (single herbal drugs)	Neo4j knowledge graph, pathway enrichment, network pharmacology, docking	Elucidate multi-target mechanisms in diabetes and obesity	*Ayurvedic Pharmacopoeia of India*; IMPPAT; FDA-approved drug targets; disease gene sets	Drug-likeness and bioavailability filtering	11 single herbal drugs and 188 phytochemicals shortlisted; significant pathway overlap with disease sets; PTPN1 among prioritised targets
A2 Formulation, prescription and interaction modelling						
Sukruth Gowda et al. (2025),	*Ayurveda*; *Charaka*, *Suśruta*,	TF-IDF with Multinomial Naive Bayes	Classify free-text symptoms and	Curated *Ayurvedic* symptom set;	Practitioner Likert rating over 50+ case	85-90% relevance as judged by

AYURSOFTAPP	*Aṣṭāṅga*, *Bhāvaprakāśa*	(scikit-learn), Flask	recommend formulations	Formulation-Indications dataset	scenarios; no split	practitioners; 88% of suggestions rated 4 or above. No dataset size, no held-out split
Vyas et al. (2025)	*Ayurveda*; AYUSH, TKDL, IMPPAT, classical repositories	Majority-vote ensemble of Naive Bayes, KNN, Random Forest, SVM, MLP; BioBERT token classifier	Predict disease from symptoms and return *Ayurvedic* treatment	DDXPlus and a Kaggle symptom-disease set: 6,480 samples, 192 features, 54 classes	80:20 stratified split	Ensemble 99.91%; Naive Bayes 98.64%, Random Forest 98.19%, MLP 96.75%, KNN 95.30%, SVM 91.96%. Symptom-disease data are biomedical; the *Ayurvedic* layer is a lookup
Kasera et al. (2025)	*Ayurveda* (*Prakriti*, *Vikriti*)	Decision Tree classifier benchmarked against Random Forest, SVM, Logistic Regression, KNN	Predict *Ayurvedic* drug prescriptions	1,300 patient samples with disease, demographics, severity and treatment	Test-set evaluation	Decision tree best at 95.3% on the test set; model presented as interpretable for clinical use
Nayak et al. (2025)	*Ayurveda* (*Prakriti*, *Vikriti*)	Decision tree	Predict disease then recommend a drug by age, sex, type and severity	Kaggle symptom dataset: 32 symptoms, 1,071 patients	Train and test accuracies reported separately	Prediction 97% accuracy, precision and recall; recommendation model training accuracy 98%, precision 93%, recall 77%
Pujar et al. (2025)	*Ayurveda*	Decision Tree Classifier in scikit-learn	Predict medicine name from symptom indications	Symptom-indication dataset with categorical encoding; size not stated	70:30 train-test split; confusion matrix and classification report on train and test	Classifier learns symptom-to-medicine decision rules; precision, recall and F1 computed per class, no headline figure stated. The published abstract does

						not describe the system
Supal et al. (2025)	*Ayurveda*	RAG with ChromaDB vector retrieval, Neo4j knowledge graph, Hugging Face embeddings, Zephyr-7B	Hybrid chatbot combining retrieval and graph reasoning	*Ayurvedic* texts; 110-query benchmark	Benchmark of 110 user queries	BLEU-4 0.35, ROUGE-1 0.55, ROUGE-2 0.52, ROUGE-L 0.38, METEOR 0.42, F1 0.61
Manikandan et al. (2026)	*Ayurveda* (*Prakriti*, *Vata*, *Pitta*, *Kapha*)	Multi-layered neural network with NLP front end	Suggest remedies with dosage, preparation and contraindications	Curated symptoms, patient profiles and herbal formulations; 10,000 clinical case records	70/15/15 split and 5-fold cross-validation; compared against expert recommendations	Reported as accurate and aligned with expert recommendations; no figures given
A. Joshi et al. (2026), AyurAI	*Ayurveda*; curated classical corpora	Transformer LLMs via Hugging Face and LangChain; vector semantic retrieval	Disease prediction and formulation recommendation with traceability	Curated *Ayurvedic* textual corpora	Not reported	Framed as decision support rather than autonomous diagnosis; recommendations traceable to identifiable textual references
Chaya et al. (2026)	*Ayurveda*	Random Forest, Naive Bayes, SVM	Automate disease prediction and drug recommendation	Symptom dataset of 4,920 records attributed to the AYUSH portal	Separate 41-record test set described; no result reported	Reported as accurate and scalable; no figures given
Vignesh and Muthukumaran (2024)	*Ayurveda*	LeNet for feature extraction with stacked autoencoders	Chatbot for formulation selection	Medical dataset of 4,000 rows, source not specified	3,200/800 train-test split	Accuracy 0.97, precision 0.95; handles synonymy and disease-name variation, and flags contraindications
Mathew et al. (2024)	*Ayurveda* (*Prakriti*)	Clustering with classification, deliberately	Recommend drugs from *Prakriti* and symptoms	Patient symptoms and constitution	Cross-validation named; no	Reported as accurate in identifying *dosha*

		avoiding deep learning			folds or results given	imbalance; no figures given
Dhabe et al. (2024), Medi-Compass	*Ayurveda*	Not named in the record assessed	Symptom-based medication recommendation website	Not reported	Not reported	No method, dataset or metric stated in the record assessed
Khan et al. (2024), MediMatch	*Ayurveda* (*Vikruti*)	Decision trees, cosine similarity, bagging, Random Forest	Detect disease and recommend *Ayurvedic* medicines	Historical patient data linked to *Ayurvedic* interventions	Clinical validation named as a challenge, not performed	Data collection, data quality and clinical validation identified as unresolved problems
Vijay et al. (2024)	*Ayurveda* (*Prakriti*)	Decision tree	Personalised recommendation with clinic location	Symptom and constitution inputs; Google Maps	80:20 split and 10-fold cross-validation; accuracy, precision and recall	Performance reported without figures in the record assessed
Pooja et al. (2025), MediGuide AI	*Ayurveda* alongside conventional drugs	Naive Bayes, TF-IDF, NLTK; regression models for dosage	Dual conventional and *Ayurvedic* recommendation with dosage calculation	Patient demographics and clinical factors	k-fold cross-validation described; no k or result given	Dosage module computes by age, weight and metabolism; *Ayurvedic* alternatives presented alongside conventional drugs
A3 Knowledge bases, ontologies and knowledge graphs						
R. Patil et al. (2024), AyUR-bot	*Ayurveda*	Llama 2 LLM with RAG, LangChain, FAISS vector store, Hugging Face embeddings	Question answering on nutrition, remedies and wellness	Curated *Ayurvedic* document repository	BERTScore	BERTScore used to validate context-aware responses; conversational retrieval chain with recursive character text splitting
Ahmed et al. (2024), AyurChat	*Ayurveda*	NLP with machine learning; no algorithm named as implemented	Conversational *Ayurvedic* guidance	Knowledge base of *Ayurvedic* remedies	Evaluated on 30 sets of input symptoms	Response accuracy 82% against Babylon Health 53%, ADA Health 71% and Your MD 64%. BERT and GPT named only under

						future enhancements
Pavithra et al. (2024)	*Ayurveda*	Knowledge graph with HITS (Hyperlink-Induced Topic Search)	Rank medications by authority within a disease-drug graph	Disease and medication relations with demographic attributes	Not reported	HITS authority and hub scores rank treatments; incompatible drugs filtered by demographic and comorbidity constraints
Gawande et al. (2025), HerbFusionAI	*Ayurveda*	BERT embeddings, graph neural network, generative modelling, cosine similarity	Interpretable herbal recommendation	Disease and symptom descriptions; herb relations	Not reported	Positioned against prior Random Forest, Naive Bayes and collaborative-filtering work as better able to capture disease-symptom-herb structure
Adith and Gupta (2025)	*Ayurveda*; Sanskrit source texts	RAG with OpenAI and MiniLM embeddings, GPT-3.5 Turbo; Deepgram speech-to-text and text-to-speech	Speech-to-speech *Ayurvedic* guidance	Authentic *Ayurvedic* texts	Answer relevancy and context recall	OpenAI embeddings with GPT-3.5 Turbo performed best: answer relevancy 0.94, context recall 0.73
P. Bhatnagar et al. (2026), AyurGenie	*Ayurveda* (*dosha* profiles)	RAG chatbot on pretrained language models with semantic retrieval	Personalised *Ayurvedic* recommendation	Authoritative *Ayurvedic* texts	Benchmarked against ChatGPT-4, Practo Health Chatbot and FAQ systems	Response accuracy 85-90%, query processing 0.8-1.2 s, 75% satisfaction, 52% retention
Dhawan et al. (2026), Ayurveda-Net	*Ayurveda* (*Dosha*, *Prakriti*, *Nidana*, *Samprapti*, *Chikitsa*)	Causal knowledge graph with explainable-AI framing	Represent causal relations among *Ayurvedic* concepts	*Ayurvedic* conceptual framework; *Madhumeha* case study	Not reported	Multi-entity causal graph demonstrated on *Madhumeha* for disease-progression and treatment modelling
Siddireddy et al. (2026)	*Ayurveda*	Custom NER, bespoke WordNet, RAG	Extract *Ayurvedic* entities and	*Ayurvedic* literature	Faithfulness and context recall	Faithfulness 0.8739, context recall 0.8158;

			answer queries			the custom NER outperformed standard NER models on *Ayurvedic* terminology
A4 Personalised medicine and *prakriti* modelling						
P. Patil and Gore (2017)	*Ayurveda* (*Prakriti*)	Decision tree and neural network	Recommend *yoga* and *rāga* by constitution	Expert-defined parameters for disorder, atmosphere, mood, timing and constitution; 250 instances	Not reported	Parameters for recommendation fixed by *Ayurvedic* experts; earliest included study
Parshionikar et al. (2025)	*Ayurveda* (*Prakriti*, *dosha*)	Random Forest classifier with SMOTE class balancing	Classify *Prakriti* from a questionnaire	Questionnaire responses, 1,200 records, balanced across *Prakriti* types	80:20 split after SMOTE on the full dataset	SMOTE used to balance representation across constitution types
Manishramar et al. (2025)	*Ayurveda*	Decision tree classifier	Predict disease from self-reported symptoms and advise treatment	User symptom reports	Model comparison reported	Decision tree accuracy 94.74%, recall 85.71%, F1 87.88%
Suryawanshi et al. (2026), AyurGenixAI	*Ayurveda* (*dosha*, *Prakriti*)	Random Forest, XGBoost, SVM, KNN, CNN, LSTM with an LLM layer	Personalised diagnosis, *dosha* classification, diet and therapy	15,160 records covering 447 diseases and 35 parameters	None reported	XGBoost 85%, Random Forest 55%, overall system 90%; no cross-validation or split reported
A5 Disease or symptom prediction with therapeutic linkage						
Chelluboina and Rao (2023), CATDSNet	Indian *Ayurvedic* medicine (tongue diagnosis)	Hybrid extreme learning machine; FNLM filtering; colour moments and GLCM features	Predict disease from tongue images	Tongue image dataset, 713 images	No split stated; compared against prior systems	Colour and texture features combined to address variation in tongue colour, shape, coating and texture
Shetty et al. (2024)	*Ayurveda* (lifestyle,	Several supervised	Stratify gestational	Clinical risk-factor data	No split or cross-	Classifiers reached 75-

	decoctions, *cūrṇa*)	classifiers, compared	diabetes risk and link to *Ayurvedic* management		validation stated; accuracy, precision, F1	82% accuracy; *Ayurvedic* lifestyle and herbal management discussed for risk reduction
Shinde et al. (2024)	*Ayurveda*	Random Forest, decision tree, CNN; KNN for hospital recommendation	Predict disease, suggest remedies, recommend hospitals	Open dataset: 4,920 records, 41 diseases, 132 symptoms	Train and test accuracy reported	Decision tree 95%, Random Forest 98%, CNN 99% on both train and test
Pradhan et al. (2025)	*Ayurveda*	NLP with machine learning, not further named	Chatbot identifying intent and suggesting remedies	*Ayurvedic* texts; practitioner referral	Not reported	Continuous refinement from user interaction claimed; no figures given
Serrao et al. (2025), AyurMate	*Ayurveda*	Artificial neural network in Keras; NLTK for NLP; Flask deployment	Classify user intent and suggest treatment	Intents and responses dataset	Part of the data held back; size and ratio not stated	Covers fever, indigestion and skin complaints; real-time deployment via Flask
Vayadande et al. (2025)	*Ayurveda* (*dosha*, *Nāḍī Parīkṣā*)	Pulse sensors with classification algorithms	Diagnose heart disease and classify *Prakriti* for diet advice	Pulse-sensor readings; heart-disease table of 1,025 rows	820/205 split stated; headline scores are training-set scores	Couples classical pulse examination with sensor capture for diagnosis and dietary recommendation
Majhi et al. (2023)	*Ayurveda* (*dosha* analysis)	Kernel SVM, SVM, logistic regression, XGBoost, decision tree, Random Forest in scikit-learn	Predict Parkinson's disease through *dosha*-mapped features	Fox Insight dataset, n = 80,916 records	10-fold cross-validation with confusion matrix	Best model 92.6% accuracy with sensitivity 0.856 and false-positive rate 0.045; four algorithms above 92.5%. Largest dataset in the review

In Table 5, no validation approach is reported in 13 of the 42 studies, and dataset size is not stated in 20. Where performance is reported without a stated split or cross-validation, as in several of the A2 recommendation systems, the figure cannot be interpreted as a generalization estimate.

Reporting quality across the included studies

Because every study included was read in full, the same reporting items could be checked across all 42. Ten recur often enough to count. A study is coded as showing an item only where its full text gives explicit evidence, so the figures below are minima rather than estimates, and the per-study coding, with the page on which the evidence appears, is given in the defect matrix of Supplementary File S1.

Dataset size is the most frequent omission, absent from 20 of the 42 studies. Citation defects come next, in 18. These range from references that cannot be identified at all, such as the seven author-initial entries in one related-work section (Shinde et al., 2024), through numbering that skips twenty places mid-paragraph, to a reference list of fourteen entries that is eight unique sources duplicated in pairs (P. Bhatnagar et al., 2026). Thirteen studies involved patients, users or practitioners and reported no ethical approval, consent procedure or institutional review; one of these drew on 10,000 clinical case records obtained from unnamed clinical institutions (Manikandan et al., 2026). Thirteen report no validation protocol of any kind.

Ten studies contain figures that cannot all be true. Five of these are arithmetic. Two papers print an F1 score outside the range their own precision and recall permit (Shetty et al., 2024; Siddireddy et al., 2026), one prints the accuracy equation as its own reciprocal (Pavithra et al., 2024), one states a train, validation and test split summing to 104 per cent (Shi et al., 2025), and one states a split summing to 95. The other five are inconsistencies between parts of the same paper: an abstract against a body, a reported accuracy against a conclusion, a confusion matrix against the accuracy claimed from it, and a precision at five exceeding a precision at three (P. Bhatnagar et al., 2026). In one study the arithmetic reverses the finding (Siddireddy et al., 2026). The printed scores make the authors' purpose-built entity recognizer the best of three; the corrected values make an off-the-shelf baseline the strongest, which is the opposite of what the paper claims in both its abstract and its results.

Nine studies claim superiority over named alternatives without re-implementing them. Two of these state in the text that the comparison was not carried out, while presenting a comparison table.

Eight studies describe their methods or results in future or conditional tense. One has a results section headed "Outcomes expected" (A. Joshi et al., 2026); another writes its entire abstract conditionally and lists data collection as future work (Khan et al., 2024). Seven report no measured quantity at all, offering interface screenshots or prose claims in place of numbers.

Seven studies make an *Ayurvedic* claim their data does not support. A gestational diabetes model is trained on the Pima Indians Diabetes Database, which records type 2 diabetes (Shetty et al., 2024). A knowledge-graph paper states that building the *Ayurvedic* dataset "is the future work of our project" (Pavithra et al., 2024). One precautions dataset was written by the authors from named consumer websites, one of which has an invalid address (Shinde et al., 2024). A heart disease model classifies on age, sex, blood pressure, cholesterol and chest pain type, in a paper whose title and hardware are about *Ayurvedic* pulse examination (Vayadande et al., 2025).

Three studies report a training-set score as the headline result. In one, the authors' own figure caption says so (Vayadande et al., 2025), which also explains why a single decision tree in that paper appears to outperform a random forest.

Eight of the 42 studies show none of these items and fifteen show four or more (Table 6).

Table 6 - Reporting items observed across the 42 included studies. Counts are minima, coded only where the full text gives explicit evidence.

Reporting item	n	% of 42
Dataset size not stated	20	48
Citation defects	18	43
Human subjects or expert evaluation without an ethics statement	13	31
No validation protocol reported	13	31
Internally contradictory numbers	10	24
Baseline comparison asserted but never run	9	21
Methods or results in future or conditional tense	8	19
Ayurvedic claim not supported by the data used	7	17
No quantitative result of any kind	7	17
Leakage or training-set scoring as the headline result	3	7

Reporting quality and the availability of an external referent

The defects are not evenly spread across application domains, and the pattern is suggestive, although it is confounded by publication venue.

Studies in drug discovery, phytochemistry and molecular modelling (A1) carry a mean of 1.12 coded defects, against 2.91 for the other 34 studies (Mann-Whitney U = 64.0, two-sided p = 0.020). The two groups are disjoint, because the comparison is with every study that carries no A1 tag. The gap is widest against personalized medicine and *prakriti* modelling (A4), where the mean is 3.29 and no study is also tagged A1 (U = 24.0, two-sided p = 0.010). None of the eight A1 studies makes an *Ayurvedic* claim its data does not support, against seven of the other 34, but that difference is not significant (Fisher's exact, two-sided p = 0.31) (Fig. 3).

Publication venue is a likely confounder. Journal articles carry fewer coded defects than conference papers (mean 1.29 against 3.21; U = 88.0, two-sided p = 0.004), and seven of the eight A1 studies are journal articles, against seven of the other 34. With both factors in a Poisson regression of defect count, the rate ratio is 0.48 for journal publication (95% CI 0.27–0.87) and 0.63 for the A1 tag (95% CI 0.29–1.38); a negative binomial model gives similar estimates. Among journal articles alone, A1 studies average 0.57 defects against 2.00 for the other seven (U = 11.5, two-sided p = 0.094). On these numbers the A1 difference cannot be separated from venue, and the explanation that follows is offered as a hypothesis.

A1 is also the smallest domain in the corpus at 8 of 42 studies, while A5 and A4 are the largest at 22 and 17. One possible explanation is that A1 work has an external referent. A binding energy, a protein structure and a compound identifier exist independently of the paper reporting them, and a reader can check them against a public source. *Prakriti* classification and symptom-to-disease prediction have no comparable object: the quantity being predicted is established by the same procedure the model is asked to learn, so a reported accuracy has nothing outside the study to be tested against. The alternative, that journal review rather than subject matter accounts for the difference, cannot be ruled out with this corpus.

The proportion of studies with no coded defect is also higher in A1, at 4 of 8 against 4 of 34 (Fisher's exact, two-sided p = 0.03), but five of the eight defect-free studies are journal articles, so this comparison carries the same confound.

If the referent explanation holds, effort in this field is concentrated where claims are hardest to check and is thinnest where they can be checked. More predictive models trained on *prakriti* or symptom tables would not change that, because what those models lack is an external object for the claim to refer to.

As a sensitivity check, we looked at the eight studies with no coded defect: Sil et al. (2026), Niharika et al. (2026), Panghalia and Singh (2026), Kumar et al. (2025), Supal et al. (2025), Vyas et al. (2025), Ahmed et al. (2024) and Majhi et al. (2023). Six of the eight were published in 2025 or 2026. Their domain profile is flatter than that of the full set.

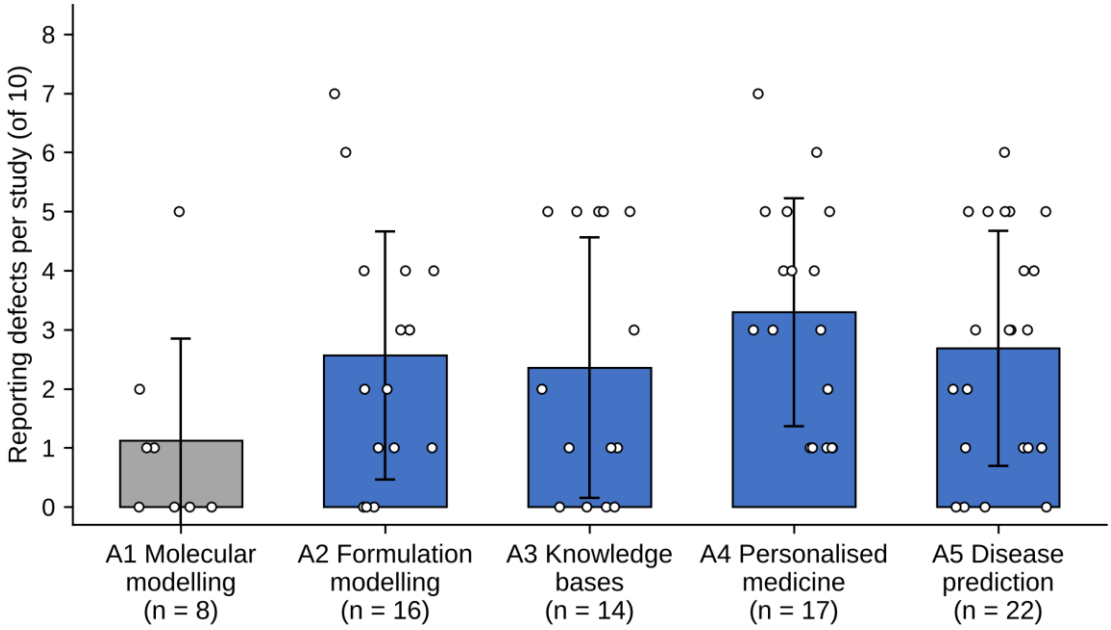


Fig. 3 - Coded reporting defects per study by application domain. Bars show means and error bars one standard deviation; points are individual studies. Studies tagged to more than one domain appear in each. A1 against the other 34 studies: two-sided Mann-Whitney U = 64.0, p = 0.020. Because of multiple tagging the five bars are not independent groups.

A1, A3 and A5 each account for four of the eight and A2 for three, and no A4 study is defect-free, whereas in all 42 studies A5 leads with 22 and A1 trails with 8. The predominance of symptom-to-remedy work reported in section 3.3 therefore comes mostly from the less well-reported studies. This subset is defined by the defect coding, so it cannot be used to re-test the A1 comparison above, and at eight studies it is too small to carry a synthesis of its own. The conclusions of this review rest on all 42 studies.

Risk of bias in prediction-model studies

All 14 prediction-model studies were at high overall risk of bias in the evaluation of their models, and all 14 raised high concern about the quality of model development (Table 7). The participants and data sources domain were rated high in 12. Most models were trained on public symptom files, self-made datasets or internet images whose origin is not described, or on data balanced by construction to about 120 records per disease; one combined curated knowledge entries with synthetic records (Suryawanshi et al., 2026).

The analysis domain was rated high in 10 and unclear in 4. The problems include a headline accuracy that is a training score (Vayadande et al., 2025), synthetic oversampling applied before the train-test split

(Parshionikar et al., 2025), test-set features used to optimize the model (Chelluboina & Rao, 2023), and repeated records of the same participants that could fall on both sides of a random split (Majhi et al., 2023). None of the 14 studies assessed calibration. The outcome domain was unclear in 9 and high in 5, because few studies say how the disease or *Prakriti* label in their data was established. Applicability concern was high in 12 and unclear in 2.

Some studies show real strengths. Two removed duplicate records before splitting (Kasera et al., 2025; Nayak et al., 2025), and one used 80,916 records with 10-fold cross-validation and a held-out test set (Majhi et al., 2023). These did not offset the concerns about where the data came from and how the outcome was defined.

Table 7 - PROBAST+AI ratings for the 14 prediction-model studies (model evaluation). D1 participants and data sources; D2 predictors; D3 outcome; D4 analysis. Development-quality ratings, all signaling-question answers and page references are in the supplementary assessment workbook.

Study	D1	D2	D3	D4	Risk of bias	Applicability
Chaya et al. (2026)	High	Low	Unclear	High	High	High
Chelluboina and Rao (2023)	High	Unclear	High	High	High	High
Kasera et al. (2025)	High	Low	High	Unclear	High	High
Khan et al. (2024)	High	Unclear	Unclear	High	High	High
Majhi et al. (2023)	High	High	Unclear	High	High	High
Manishramar et al. (2025)	High	High	Unclear	High	High	High
Mathew et al. (2024)	Unclear	Unclear	Unclear	High	High	Unclear
Nayak et al. (2025)	High	Low	Unclear	Unclear	High	High
Parshionikar et al. (2025)	Unclear	Low	Unclear	High	High	Unclear
Shetty et al. (2024)	High	Unclear	High	High	High	High
Shinde et al. (2024)	High	Low	Unclear	Unclear	High	High
Suryawanshi et al. (2026)	High	Unclear	High	High	High	High
Vayadande et al. (2025)	High	High	Unclear	High	High	High
Vyas et al. (2025)	High	Low	High	Unclear	High	High

3.4 Methodological reporting and research integrity

Seventeen inclusions that had passed title and abstract screening were reversed at the full-text stage, fourteen from the database arm and three from other sources. They fall into four groups.

Four reviews. Three say so in their abstracts and should have been excluded at title and abstract screening: a narrative review of generative AI in ethnopharmacology (Bhadra et al., 2025), a review of robotic process automation in the pharmaceutical industry (N. Bhatnagar, 2020), and a paper whose abstract states that it "review[s] academic journal and conference papers which adopted Machine Learning (ML) techniques" although its publisher labels it a research article (Keerthi & Sridevi, 2022). Their late removal reflects screening error. The fourth reads as primary research until the full text, which "represents a systematic review and conceptual framework" and presents architectures "outlined in the recent literature" (Charan et al., 2025). This paper and the one by Keerthi and Sridevi (2022) quote performance figures from the studies they survey, which is how they pass an automated screen. The figures are real, but they come from other people's models.

Two reported no results at all, and both were coded X7 (Gani et al., 2026; Kasar et al., 2025). (Kasar et al., 2025) names decision trees and neural networks, specifies a 70:30 split and lists accuracy, precision, recall and F1 as its evaluation plan, then ends without a dataset size or a single performance figure. Both are design proposals published in the form of completed studies.

Eight failed the AI criterion, the largest group. Three are competent pharmacology or chemistry papers whose computational content is a query to a pre-built web server. One combines ADMET profiling, network pharmacology, molecular docking, transcriptomic correlation and in vivo validation in a rat model, with no learning algorithm anywhere in its methods (Talati & Gaikwad, 2026). A second characterizes a mineral drug by ICP-OES, FTIR, XRD and XPS and describes a single ProTox-II query as its AI-assisted component (Bohra et al., 2026). A third describes a "multivariate analysis of chemical properties" that proves on reading to be a projection method in a data-mining package combined with Tanimoto similarity clustering assessed by Kolmogorov-Smirnov tests, with machine learning appearing only in the keyword list (Falah et al., 2024). Querying a server that someone else trained does not meet C1. A fourth is a symptom-to-formulation matcher over a structured database whose reported outcomes are data correctness, search specificity and user contentment, which are software-quality measures (Rafi et al., 2025). A fifth, titled as artificial intelligence research, is a qualitative interview study of ten practitioners analyzed thematically (Srivastava et al., 2024). The remaining three are a docking and molecular-dynamics study whose machine-learning claim rests on a phrase in its abstract (Kumar Verma et al., 2021), a wet-lab antibiotic discovery paper in which no model is trained (Chakraborty et al., 2024), and a mental-wellness application whose language component is a pre-built sentiment lexicon (Chauhan et al., 2025).

Three failed the *Ayurveda* criterion. In two, a respiratory-disease sensor system (Bhuvaneswary et al., 2025) and a healthcare appointment chatbot (Prakasam et al., 2023), *Ayurveda* appears only once or twice in the front matter. The third is the *Unani* pulse-waveform study described in section 2.4 (Waqar et al., 2025).

Two integrity findings came out of this reading.

The first concerns a 2026 paper in a journal not indexed by the four databases searched (Gani et al., 2026), excluded under X7. We record the following observations from its published version, with page numbers from the journal's pagination; the passages are reproduced in Supplementary File S4. At the top of page 14 the sentence "You are being trained on information until October 2023" appears, followed by a figure description of a genetic algorithm with a random forest fitness evaluator "as applied to the credit card fraud detection", in a paper about *Ayurvedic* herbs. The abstract (p. 11) names a support vector classifier, page 15 names decision trees, random forest, naive Bayes and logistic regression as the selected models, and Table 1 on the same page lists a rule-based model, a support vector machine, k-nearest neighbours and neural networks. No results are reported. Page 18 carries an informed-consent statement and a statement of compliance with the Declaration of Helsinki, although no human participants are described, and an author-contributions grid in which no contribution is marked. We could not locate two of the cited venues, given as "Int. J. Ayurveda Comput." (reference 1, p. 18) and "Neurophytotherapy Rev." (reference 8, p. 19). We draw no conclusion about how the paper was prepared beyond these observations.

The second is more general. Seven of the 42 included studies present a machine-learning or AI system but report no measured quantity of any kind (Table 6). Across the whole screened corpus, 14 records were excluded under C1 for having no AI component (13 from the databases and one from other sources), eight of them only after the full text had been read.

Reporting quality among the studies that do pass is uneven. Dataset size is not stated in 20 of the 42 studies, and 13 state no validation approach of any kind. Several studies report accuracy as a practitioner satisfaction rating on a Likert scale rather than as a held-out estimate: one reports 85 to 90% accuracy "as validated by *Ayurvedic* practitioners" over 50 case scenarios, with no dataset size and no split (Sukruth Gowda et al., 2025). Another reports an ensemble accuracy of 99.91% on a 54-class problem, a figure high enough to warrant a check for leakage between training and test partitions, although the paper itself gives no direct evidence of leakage (Vyas et al., 2025).

3.5 Database coverage

The four databases did not contribute equally, and the pattern has consequences for anyone replicating this search.

Scopus returned every one of the 38 database-traced studies and was the sole source for 24 of them. Web of Science returned 49 records in total and two included studies, neither of them uniquely. PubMed returned the largest raw count, 517 records, but only six included studies, because its indexing of Indian engineering conference proceedings is sparse and most of the retrieved records concerned traditional medicine systems other than *Ayurveda*. Google Scholar contributed nine included studies and the great majority of the 188 records excluded as too short to assess, since its exported snippets are truncated at a median of roughly 190 characters.

Judged on PubMed alone, AI research on *Ayurveda* looks almost absent. Once Scopus is added it is present in volume, published in venues PubMed does not index.

Four included studies were reachable by no search string we could construct and were found only by hand-searching and citation tracking. They sit in Indian journals and conference series outside the major indexes. A wider set of candidate records identified the same way came mostly from the same kinds of venue, chiefly Indian multidisciplinary science and engineering journals not indexed by the four databases (listed in Supplementary File S1). This body of work is dominated by clinical advisory software rather than compound-level research, and a review restricted to indexed sources is likely to miss most of it.

One methodological note follows from an earlier version of this search. Applying an open-access filter reduced database-traced included studies by more than half and pushed roughly two thirds of the included set into the hand-searched arm, which distorts the PRISMA flow to the point of misrepresenting where the evidence was found. Open access is recorded here as a study characteristic instead.

4. Discussion

AI work on *Ayurveda* is recent and growing, and most of it builds recommendation systems. Half the included studies take a symptom description and return a herb or formulation. The engineering is often competent, while the datasets are often small or of unstated size (20 of 42 studies do not state it) and are rarely validated against clinical outcome. Eight studies work at the level of compounds and targets, which is where claims about mechanism would have to be settled.

A system that maps symptoms to remedies reproduces an existing mapping, whether that mapping came from a classical text or a curated spreadsheet. It does not test the mapping. The compound-level studies are the ones that could, and they are the minority.

4.1 What counts as AI

The most consistent problem we encountered is the use of "AI" as a label rather than a description. Papers titled as artificial intelligence research turn out to query a database behind a web form, survey other people's models, or run a qualitative interview study. This is why we required a named method or reported metrics under C1, and why we assessed full texts rather than trusting abstracts: seventeen inclusions survived abstract screening and did not survive the full text.

Applying C1 as written has an uncomfortable consequence that we report rather than resolve. Classical molecular docking, molecular dynamics and network pharmacology are not machine learning, and under the criterion as stated they do not satisfy it. One methodologically strong pharmacology paper was excluded on exactly these grounds. A reviewer may reasonably prefer a criterion that admits computational molecular modelling alongside learning methods, which would enlarge the compound-level category and change the balance reported in section 3.3. We have kept the stricter reading, labelled the exclusion, and recorded enough detail for the boundary to be redrawn.

4.2 Image-based plant recognition

The largest coherent activity in this literature is deep learning applied to photographs of leaves, which has no therapeutic content. One hundred and fifty-six records in the screened corpus do this, more than any other excluded category and more than three times the size of the included set.

Species identification matters for raw-drug authentication and for adulteration control, both real problems in *Ayurvedic* supply chains. But the output is a species label, and appending the documented uses of that species from a lookup table adds no modelled relationship between a drug and a disease. Reviews that count these studies as AI applications in *Ayurveda* will overstate how much computational work addresses therapeutics.

4.3 Comparison with earlier reviews

Two prior reviews cover overlapping ground and reached different pictures, for reasons that are mostly structural. The scoping review of AI in complementary and alternative medicine (Chu et al., 2022) searched PubMed, Embase and the Cochrane Library, found 1,114 records and included 32 studies, with *Ayurveda* a small share of a corpus dominated by other modalities. The narrative review of AI-titled traditional medicine discovery (Chen et al., 2026) searched five sources, found 2,128 records and retained 12, almost all Traditional Chinese Medicine, after requiring a drug-discovery component.

Neither used Scopus as a primary source. Our results suggest that this choice largely shapes what a review of this field can see, because Scopus returned every database-traced study we included, and the other three databases contributed nothing uniquely. A review of AI in *Ayurveda* built on PubMed and Embase would probably conclude that the field is small, because those databases index few Indian conference proceedings.

4.4 Implications

For researchers, the gap is at the compound level, and specifically in work that links a phytochemical to a target and a target to a disease with evidence at each step rather than asserting the end-to-end association. The published graph resources provide the data, and validated inference over them is still scarce.

For reviewers and editors, the integrity findings suggest checks that could be made routinely: whether a named method appears in the methods section and not only in the title, whether reported accuracy is a held-out estimate, and whether cited venues exist.

For anyone searching this literature, Scopus was close to necessary in this search, and the other databases added no included study that Scopus had missed, while a meaningful body of work sits outside all of them and is reachable only by hand-searching.

5. Limitations

This review has several limitations. First, screening, full-text assessment, data extraction and coding were carried out by a single reviewer, with no independent second reviewer and no measure of agreement, and the protocol was not registered. The reporting items were coded only where the full text gave explicit evidence, but some judgements, such as whether an *Ayurvedic* claim is supported by the data or whether a set of test cases counts as a validation approach, involve interpretation. A recheck of every study against its full text during revision changed 16 of the 420 study-item codes (Supplementary File S1), which shows the extent of this subjectivity.

Second, although four databases were searched, every database-traced included study was returned by Scopus (section 3.5). In practice the included set therefore depends on the coverage of a single database, and work that Scopus does not index was reachable only by hand-searching, which is unlikely to have been complete.

Third, 188 records were excluded as too short to assess (X6), almost all of them Google Scholar entries whose exported snippets are truncated. Their full texts were not sought, and some eligible studies may be among them.

Finally, the review covers English-language records only. We also did not search Embase or the Cochrane Library.

6. Conclusion

Forty-two studies apply artificial intelligence to *Ayurveda* in a way that produces a knowledge-level, therapeutic or compound-level result, and 38 of them were published from 2024 onward. The work concentrates on symptom-to-remedy recommendation, often on datasets whose size is not reported, and thins out at the compound level where mechanistic claims would have to be tested.

Two findings qualify the rest. Fourteen screened records retrieved as AI research contained no learning method, eight of them identified only at full text, and the single largest activity in the screened literature is image-based species recognition, which answers a question about a specimen rather than about a therapy. Both inflate how much computational work appears to address *Ayurvedic* therapeutics.

The literature is also hard to find. Scopus returned all 38 database-traced studies included here, the three other databases added none, and four studies were reachable only by hand-searching venues that no index covers.

7. Declarations

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflict of interest. The authors declare that they have no conflict of interest.

Author contributions. Ankit Raj: conceptualization, search, screening, data extraction, coding, analysis and writing of the original draft. Ranjana M: Methodology, Formal Analysis, Writing — review and editing. Girish Nath Jha: supervision, and review and editing of the manuscript.

Data availability. The search strings, screening and exclusion decisions, deduplication script and output, extraction table, per-study defect matrix with page-level evidence and PROBAST+AI ratings are provided as Supplementary Files S1 to S3.

Supplementary files. S1: review workbook (eligibility, screening flow, included and excluded records, extraction table, defect matrix with evidence, journal list). S2: deduplication script and results. S3: PROBAST+AI assessment workbook.

8. References

[1] Adith, M., & Gupta, D. (2025). Ayurvedic health and wellness navigator using retrieval augmented generation. In 2025 5th Asian Conference on Innovation in Technology, ASIANCON 2025. https://doi.org/10.1109/ASIANCON66527.2025.11281029

[2] Ahmed, S. A., Gahlawat, G., Imam, F., & Saranya, G. (2024). AyurChat: Empowering health through AI-driven Ayurvedic guidance. In 2024 15th International Conference on Computing Communication and Networking Technologies, ICCCNT 2024. https://doi.org/10.1109/ICCCNT61001.2024.10724133

[3] Bhadra, S., Pundir, C., Mukherjee, M., Kar, A., Banerjee, S., Charoensup, R., Duangyod, T., & Mukherjee, P. K. (2025). Reimagining ethnopharmacology with generative AI: Towards inclusive, ethical, and data-driven traditional medicine. Pharmacological Research. https://doi.org/10.1016/j.phrs.2025.108002

[4] Bhatnagar, N. (2020). Role of robotic process automation in pharmaceutical industries. In Advances in Intelligent Systems and Computing. Springer. https://doi.org/10.1007/978-3-030-14118-9_50

[5] Bhatnagar, P., Sodhi, M., Lokesh, G. H., Koravanghat, G., & Pruthviraja, D. (2026). AyurGenie: An AI genie for Ayurvedic health guidance using RAG and deep learning. In AICCC 2025 - 2025 8th Artificial Intelligence and Cloud Computing Conference. https://doi.org/10.1145/3789982.3790044

[6] Bhuvaneswary, N., Babu, B. N., Pavithra, N., & Reddy, K. I. (2025). Design and development of an IoT and web-integrated respiratory disease prediction system using machine learning. In Proceedings of the 5th International Conference on Evolutionary Computing and Mobile Sustainable Networks, ICECMSN 2025. https://doi.org/10.1109/ICECMSN68058.2025.11383348

[7] Bohra, B., Yadav, B., Kumar, D., & Saxena, B. (2026). From structure to safety: Material characterization and toxicological evaluation of Vanshlochan (Bambusa bambos (L.) Voss [Poaceae]) using AI-assisted and in vivo approaches. Journal of Ethnopharmacology. https://doi.org/10.1016/j.jep.2026.121266

[8] Chakraborty, A. K., Maity, M., & Sahoo, S. (2024). Discovery of new antibiotics using AI-guided spectroscopy and 3D drug-protein computer simulation technologies to combat MDR bacteria-associated mortality. Artificial Intelligence in Health. https://doi.org/10.36922/aih.2284

[9] Charan, K. S., Surya Manikanta, B. D., Sathiya Suntharam, V., & Sai, B. M. (2025). Intelligent symptom conversations: New frontiers in AI-powered preliminary healthcare diagnosis. In 2025 International Conference on Future Technologies, ICFT 2025. https://doi.org/10.1109/ICFT66708.2025.11336630

[10] Chauhan, B., Badnaware, T., Maharana, S., Patil, S., & Kadam, P. (2025). AyurMi: A synergistic approach to mental wellness through AI and Ayurvedic mindfulness. In 2025 9th International Conference on Computing, Communication, Control and Automation, ICCUBEA 2025. https://doi.org/10.1109/ICCUBEA65967.2025.11284029

[11] Chaya, P., Anand, M., Annapurna, S. J., Poornima, M., Sowmya, S., & Yashaswini, S. (2026). An automation for Ayurveda drug recommendation system using ML. Grenze International Journal of Engineering and Technology, 12(1), 290-296.

[12] Chelluboina, S. P., & Rao, K. N. (2023). CATDSNet: Computer aided tongue diagnosis system for disease prediction using hybrid extreme learning machine. International Journal of Intelligent Engineering and Systems, 16(1), 1-12. https://doi.org/10.22266/ijies2023.0228.01

[13] Chen, S., Aslam, M. S., & Liang, H. (2026). AI-powered traditional medicine discovery: A structured narrative review of "AI"-titled studies. Current Topics in Medicinal Chemistry. Advance online publication. https://doi.org/10.2174/0115680266470947260514050449

[14] Chu, H., Moon, S., Park, J., Bak, S., Ko, Y., & Youn, B.-Y. (2022). The use of artificial intelligence in complementary and alternative medicine: A systematic scoping review. Frontiers in Pharmacology, 13, 826044. https://doi.org/10.3389/fphar.2022.826044

[15] Collins, G. S., Moons, K. G. M., Dhiman, P., Riley, R. D., Beam, A. L., Van Calster, B., Ghassemi, M., Liu, X., Reitsma, J. B., van Smeden, M., Boulesteix, A.-L., Camaradou, J. C., Celi, L. A., Denaxas, S., Denniston, A. K., Glocker, B., Golub, R. M., Harvey, H., Heinze, G., . . . Logullo, P. (2024). TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ, 385, e078378. https://doi.org/10.1136/bmj-2023-078378

[16] Dhabe, P., Garje, R., Ghodake, S., Dhoot, H., Ingle, P., & Shaikh, I. (2024). Medi-Compass: Symptom-based Ayurvedic medication recommendation system. In Proceeding of 2024 International Conference on Communication, Computing and Energy Efficient Technologies, I3CEET 2024. https://doi.org/10.1109/I3CEET61722.2024.10994125

[17] Dhawan, S., Gupta, A. D., Lohani, B. P., Kushwaha, P., Khanna, T., & Meenakshi, K. (2026). Ayurveda-Net: A causal knowledge graph for herbal actions, diagnostics, and treatments. In Proceeding - 2026 IEEE 6th International Conference on Computing, Power, and Communication Technologies, IC2PCT 2026. https://doi.org/10.1109/IC2PCT68894.2026.11584320

[18] Falah, K., Zhang, P., Nigam, A. K., Maity, K., Chang, G., Granados, J. C., Momper, J. D., & Nigam, S. K. (2024). In vivo regulation of small molecule natural products, antioxidants, and nutrients by OAT1 and OAT3. Nutrients, 16(14), Article 2242. https://doi.org/10.3390/nu16142242

[19] Gallifant, J., Afshar, M., Ameen, S., Aphinyanaphongs, Y., Chen, S., Cacciamani, G., Demner-Fushman, D., Dligach, D., Daneshjou, R., Fernandes, C., Hansen, L., Landman, A., Lehmann, L., McCoy, L., Miller, T., Moreno, A., Munch, N., Restrepo, D., Savova, G., . . . Bitterman, D. (2025). The TRIPOD-LLM reporting guideline for studies using large language models. Nature Medicine, 31(1), 60-69. https://doi.org/10.1038/s41591-024-03425-5

[20] Gani, A., Pimpalshende, A., Bhoyar, G., Kumbhare, N., & Jadhav, A. (2026). Full-stack symptom driven Ayurvedic herbs and lifestyle recommendation system (Ayurveda herbs recommendations and disease recognition). International Journal of Research in Science & Engineering, 6(1), 11-20. https://doi.org/10.55529/ijrise.61.11.20

[21] Gawande, P., Gawande, K., Sakhare, A., Khobragade, P., & Madankar, M. (2025). HerbFusionAI: Integrating BERT, graph learning & generative modeling for Ayurveda. In 2025 1st International Conference on Advancement in Futuristic Technologies, ICAFT 2025. https://doi.org/10.1109/ICAFT66710.2025.11453196

[22] Joshi, A., Yadav, B., & Kasar, M. (2026). AyurAI: An LLM-based framework for disease prediction and personalized Ayurvedic formulation recommendation. In Proceedings of the 2026 13th International Conference on Computing for Sustainable Global Development, INDIACom 2026. https://doi.org/10.23919/INDIACom70271.2026.11525492

[23] Joshi, S., Pathak, A., Regati, D. R., Menon, R., Ajith, D. S., Sheshadri, A., Viswanathan, N., Ray, P., Koul, V., Panda, P., Bhambore, S. A., Verma, S., Sinha, A., Shafi, K. M., Pavalam, M., & Sowdhamini, R. (2026). GRAYU: Graph-based database integrating Ayurvedic formulations, medicinal plants, phytochemicals and diseases. Frontiers in Pharmacology, 16, Article 1727224. https://doi.org/10.3389/fphar.2025.1727224

[24] Kasar, B. V., Pal, A., Ghotakar, P., Bagade, Y., & Badve, S. (2025). AyurPharma: Ayurvedic drug recommendation system. International Journal for Multidisciplinary Research, 7(1), Article IJFMR250137337.

[25] Kasera, G., Prajapati, H., & Modi, N. (2025). An interpretable decision tree model for Ayurvedic drug recommendation in personalized healthcare. In ICDT 2025 - 3rd International Conference on Disruptive Technologies. https://doi.org/10.1109/ICDT63985.2025.10986329

[26] Keerthi, K. L., & Sridevi, S. (2022). Big data analysis of traditional Indian Ayurveda medicine and treatment process. In AIP Conference Proceedings. AIP Publishing. https://doi.org/10.1063/5.0080305

[27] Khan, S., Saify, A., Gosaliya, S., Jain, D., & Zalte, J. (2024). MediMatch: AI-driven drug recommendation system. In 2nd International Conference on Sustainable Computing and Smart Systems, ICSCSS 2024 - Proceedings. https://doi.org/10.1109/ICSCSS60660.2024.10625049

[28] Kiewhuo, K., Gogoi, D., Mahanta, H. J., Rawal, R. K., Das, D., S, V., Jamir, E., & Sastry, G. N. (2023). OSADHI - An online structural and analytics based database for herbs of India. Computational Biology and Chemistry. https://doi.org/10.1016/j.compbiolchem.2022.107799

[29] Kumar, P. K., Awere, C. O., Kumari, A. R., Çolak, A. B., Bayrak, M., Ogolla, F. O., Govindan, S., & Ramesh, M. (2025). Machine learning-based identification of elite genotypes in the endangered Nilgirianthus ciliatus through qualitative and quantitative trait analysis. Current Research in Biotechnology, 10, Article 100307. https://doi.org/10.1016/j.crbiot.2025.100307

[30] Kumar Verma, A., Kumar, V., Singh, S., Goswami, B. C., Camps, I., Sekar, A., Yoon, S., & Lee, K. W. (2021). Repurposing potential of Ayurvedic medicinal plants derived active principles against SARS-CoV-2 associated target proteins revealed by molecular docking, molecular dynamics and MM-PBSA studies. Biomedicine & Pharmacotherapy, 137, Article 111356. https://doi.org/10.1016/j.biopha.2021.111356

[31] Majhi, V., Choudhury, B., Saha, G., & Paul, S. (2023). Development of a machine learning-based Parkinson's disease prediction system through Ayurvedic dosha analysis. International Journal of Ayurvedic Medicine, 14(1), 180-189.

[32] Manikandan, A., Thendral, R., Naveen, S., Sanjay, S., Raj Kumar, K., & Kumaran, S. (2026). Smart Ayurvedic medicine suggestion system using deep learning. In Proceedings of the 2026 IEEE Nexo-Tech International Conference on Advanced Technologies and Innovations, ICATI 2026. https://doi.org/10.1109/ICATI69399.2026.11604381

[33] Manishramar, N., Saini, S., & Iniyan, S. (2025). Disease prediction and Ayurvedic guidance applying machine learning. In Proceedings of 2025 AI-Driven Smart Healthcare for Society 5.0, AdSoc5.0 2025. https://doi.org/10.1109/IEEECONF64992.2025.10963118

[34] Mathew, J. K., Pandi, J. M., Kumar, V. H., Charan, V. V. S. V., Bindu, S. M., & Bhoomika, H. (2024). Hybrid machine learning for personalized Ayurvedic drug recommendation via prakriti assessment. In 2024 1st International Conference for Women in Computing, InCoWoCo 2024 - Proceedings. https://doi.org/10.1109/InCoWoCo64194.2024.10863221

[35] Moons, K. G. M., Damen, J. A. A., Kaul, T., Hooft, L., Andaur Navarro, C., Dhiman, P., Beam, A. L., Van Calster, B., Celi, L. A., Denaxas, S., Denniston, A. K., Ghassemi, M., Heinze, G., Kengne, A. P., Maier-Hein, L., Liu, X., Logullo, P., McCradden, M. D., Liu, N., . . . van Smeden, M. (2025). PROBAST+AI: An updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. BMJ, 388, e082505. https://doi.org/10.1136/bmj-2024-082505

[36] Nayak, S. K., Garanayak, M., Swain, S. K., Panigrahi, G. R., Sethy, P. K., & Nanthaamornphong, A. (2025). A decision tree-based approach for disease prediction and Ayurvedic drug recommendation. Ingenierie des Systemes d'Information, 30(8), 2001-2010. https://doi.org/10.18280/isi.300806

[37] Niharika, D. G., Salaria, P., Anusha, A., Lakshmi Ganapati, P., & Amarendar Reddy, M. (2026). Integrated machine learning and computational approaches for identifying NLRP3 inhibitors from Jeevaneeya Rasayana in epilepsy. Journal of Molecular Graphics and Modelling, 147, Article 109460. https://doi.org/10.1016/j.jmgm.2026.109460

[38] Nikhil, S., Prasad, P., Monish, T. R., Singh, A., Roopashree, S., & Hema, M. S. (2025). RAG-based multi-agent systems for drug discovery with Indian medicinal plants. In International Conference on Emerging Technologies in Electronics and Green Energy, ICETEG 2025. https://doi.org/10.1109/ICETEG66194.2025.11472810

[39] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., . . . Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

[40] Panghalia, A., & Singh, V. (2026). SitoshnaPred: Learning phytochemical descriptors to elucidate Ayurvedic herbal potency. Journal of Ethnopharmacology, 355, Article 120620. https://doi.org/10.1016/j.jep.2025.120620

[41] Parshionikar, S., Tuscano, S., Mascarenhas, S., Gonsalves, C., & Lopes, A. (2025). An AI-driven approach for predicting prakriti in Ayurvedic healthcare using machine learning. In 2025 3rd World Conference on Communication and Computing, WCONF 2025. https://doi.org/10.1109/WCONF64849.2025.11233348

[42] Patil, P., & Gore, S. (2017). Recommendation system for yoga and raga for personalized health based on constitution. In Proceedings - 2nd International Conference on Computing, Communication, Control and Automation, ICCUBEA 2016. https://doi.org/10.1109/ICCUBEA.2016.7860010

[43] Patil, R., Yeolekar, S., Khade, O., Kadam, Y., Ingole, P., & Patil, S. (2024). AyUR-bot: A revolutionary Ayurvedic chatbot empowered by generative AI. In 2024 8th International Conference on Computing, Communication, Control and Automation, ICCUBEA 2024. https://doi.org/10.1109/ICCUBEA61740.2024.10774721

[44] Pavithra, M. N. S., Harshini, M., Padmadevi, S., Deepalakshmi, R., & Vijayalakshmi, R. (2024). Ayurvedic elixirs for digital generation using machine learning. In 2nd International Conference on Emerging Trends in Information Technology and Engineering, ic-ETITE 2024. https://doi.org/10.1109/ic-ETITE58242.2024.10493477

[45] Pooja, R., Elugubanti, N., Mary Shiba, C., Kalai priya, V., Sasirekha, K., & Sabapathy, S. R. (2025). MediGuide AI using machine learning. In 2nd International Conference on Machine Learning and Autonomous Systems, ICMLAS 2025 - Proceedings. https://doi.org/10.1109/ICMLAS64557.2025.10968033

[46] Pradhan, S., Raizada, I., Kumar, V., Kumar, A., Singh, R., Das, D., & Patra, S. (2025). Automated Ayurvedic healthcare chatbot using NLP and machine learning. In 1st IEEE International Conference on Data Science and Intelligent Network Computing, ICDSINC 2025. https://doi.org/10.1109/ICDSINC66221.2025.11448125

[47] Prakasam, S., Balakrishnan, N., Kirthickram, T. R., Ajith Jerom, B., & Deepak, S. (2023). Design and development of AI-powered healthcare WhatsApp chatbot. In ViTECoN 2023: 2nd IEEE International Conference on Vision Towards Emerging Trends in Communication and Networking Technologies, Proceedings. https://doi.org/10.1109/ViTECoN58111.2023.10157423

[48] Pujar, A. M., Chinchawade, A. R., Pardeshi, K. R., Vibhute, G. A., Babhale, S. R., & Shah, S. M. (2025). Recommendation of Ayurvedic drugs using machine learning. In Proceedings of the 2025 3rd International Conference on Inventive Computing and Informatics, ICICI 2025. https://doi.org/10.1109/ICICI65870.2025.11069926

[49] Rafi, G. M., Khanum, D. A., Rafi, C. M., Siraj, M., & Srinivasulu, T. (2025). VEDA CARE: A digital solution for personalized herbal medicine. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security, ICCAMS 2025. https://doi.org/10.1109/ICCAMS65118.2025.11234051

[50] Serrao, V., Ashwini, K. M., & Varsha. (2025). AyurMate: Blending Ayurveda and artificial intelligence for personalized wellness. In 2025 International Conference on Artificial Intelligence and Data Engineering, AIDE 2025 - Proceedings. https://doi.org/10.1109/AIDE64228.2025.10987375

[51] Shetty, N. P., Shetty, J., Hegde, V., Dharne, S. D., & Kv, M. (2024). A machine learning-based clinical decision support system for effective stratification of gestational diabetes mellitus and management through Ayurveda. Journal of Ayurveda and Integrative Medicine, 15, Article 101051. https://doi.org/10.1016/j.jaim.2024.101051

[52] Shi, X., Sundareswaran, R., Shanmugapriya, M., Jayaguptha, N., & Khan, A. (2025). Graph based link prediction for epilepsy drug discovery. Scientific Reports. https://doi.org/10.1038/s41598-025-14550-7

[53] Shinde, S. K., Buchade, A., Mahalle, P., Raut, P., Dirgule, A., & Mandale, P. (2024). Disease prediction with Ayurvedic remedies suggestion and hospital recommendation using ML and DL. In 2024 International Conference on Intelligent Systems and Advanced Applications, ICISAA 2024. https://doi.org/10.1109/ICISAA62385.2024.10828665

[54] Siddireddy, S. K. R., Sahay, A., & Singh, T. (2026). NLP-based Ayurvedic health recommendation system with chatbot interaction. In Lecture Notes in Electrical Engineering. https://doi.org/10.1007/978-981-96-9971-1_31

[55] Sil, P., Tiwari, R., Garisetti, V., Baskaran, S. P., Dhanaseelan, F. H., Srivastava, S., & Samal, A. (2026). Computational investigation of single herbal drugs for diabetes and obesity using knowledge graph and network pharmacology. Computational Biology and Chemistry. https://doi.org/10.1016/j.compbiolchem.2026.109194

[56] Srivastava, V. D., Kumar, V., & Kausar, K. (2024). Exploration the potential of artificial intelligence in Ayurveda diagnostics through the machine learning approach. Saudi Journal of Economics and Finance, 8(3). https://doi.org/10.36348/sjef.2024.v08i03.001

[57] Sukruth Gowda, M. A., Harika, Y., Shyam, R., Samuel, R., & Praveen Kumar, G. S. (2025). AYURSOFTAPP-software that suggests drugs and formulations for disease/pharmacological property based on the Ayurvedic classical books/repositories. International Journal of Scientific Research in Engineering and Management. https://doi.org/10.55041/IJSREM47138

[58] Supal, S., Anzar, S. M., Sankar, C. P., & Abhilash, O. V. (2025). Retrieval-augmented generation and knowledge graphs for intelligent Ayurvedic chatbots. In 2025 1st International Conference on Smart and Intelligent Systems, SISCON 2025. https://doi.org/10.1109/SISCON66686.2025.11409308

[59] Suryawanshi, R., Panchikattil, S., Jadhav, J. S., Jain, P., Patil, A., & Jain, N. (2026). AyurGenixAI: Integrating generative AI with Ayurvedic knowledge: A data-driven personalized healthcare system. International Journal of Drug Delivery Technology, 16. https://doi.org/10.25258/ijddt.16.36s.26

[60] Talati, Y. K., & Gaikwad, A. B. (2026). Exploring the anti-fibrotic potential of gramine in a UUO-induced rat model of chronic kidney disease: A network pharmacology and molecular docking approach with experimental validation. Journal of Ethnopharmacology. https://doi.org/10.1016/j.jep.2025.120975

[61] Vayadande, K., Sukte, C. D., Bodhe, Y., Jagtap, T., Joshi, A., Joshi, P., Kadam, A., & Kadam, S. (2025). Heart disease diagnosis and diet recommendation system using Ayurvedic dosha analysis. EAI Endorsed Transactions on Internet of Things, 11. https://doi.org/10.4108/eetiot.6016

[62] Vignesh, A., & Muthukumaran, N. (2024). AI-driven Ayurvedic chatbot: Revolutionizing formulation selection and knowledge accessibility in traditional medicine through the integration of LENET and stacked autoencoders. In 10th International Conference on Advanced Computing and Communication Systems, ICACCS 2024. https://doi.org/10.1109/ICACCS60874.2024.10717189

[63] Vijay, M., Mari Selvan, R., Nagavenkat, N., Sree, K. K., Jyothsna, M., & Tejasree, M. (2024). Personalized Ayurvedic medicine recommendations using decision tree algorithms and prakriti analysis. In 2nd International Conference on Emerging Research in Computational Science, ICERCS 2024. https://doi.org/10.1109/ICERCS63125.2024.10895343

[64] Vyas, R., Jain, A., Verma, S., & Sharma, H. (2025). Software for finding Ayurvedic medicines for diseases and properties. In 2025 4th OPJU International Technology Conference (OTCON), IEEE. https://doi.org/10.1109/OTCON65728.2025.11071063

[65] Waqar, S., Qayyum, S., Waheed, T., Shahid, A., Munir, A., Mukhtar, A., & Shahzad, R. (2025). A novel direction for non-invasive diabetes prediction: Integrating Unani wisdom with sensor technology. ECTI Transactions on Electrical Engineering, Electronics, and Communications. https://doi.org/10.37936/ecti-eec.2525233.255146

7.Conflict of Interest

The authors declare that there are no conflicts of interest associated with this article.

8.Funding

No funding was received to support this study.