IJAYUSH
International Journal of AYUSH
AYURVEDA • YOGA • UNANI • SIDDHA • HOMEOPATHY
https://internationaljournal.org.in/journal/index.php/ijayush/
ISSN: 2349-7025
OPEN ACCESS
★ PEER REVIEWED ★
MONTHLY

Review Article	Volume 15 Issue 09	September 2026

DOI: https://doi.org/10.5281/zenodo.22284152

THE FUTURE OF PRECISION HOMOEOPATHY: A MULTIMODAL AGENTIC INTELLIGENCE ARCHITECTURE FOR GENERATIVE CASE RECONSTRUCTION, COUNTERFACTUAL REMEDY REASONING AND ADAPTIVE LONGITUDINAL INDIVIDUALISATION

***Dr. Piyushkumar Dholariya[a], Dr. Prakruti Mandaliya[b]**

[a] AI Expert – Technical Architect, Hexaware Technologies, India

[b] Post Graduate Resident, Department of Materia Medica, Baroda Homoeopathic Medical College, Vadodara, Gujarat, India

Corresponding Author's Email ID: drpiyushdholariya@gmail.com

Abbreviations: AI = artificial intelligence; LLM = large language model; NLP = natural language processing; RAG = retrieval-augmented generation. Each abbreviation is defined at first use in the text and used consistently thereafter.

ABSTRACT

Background: Homoeopathic practice is built on individualisation, totality-of-symptoms case analysis, and cross-referencing of rubrics against repertory and materia medica sources — a demanding process that is, as a documented weakness of the field's evidence base, inconsistently recorded. Generative artificial intelligence (AI) — systems, typically built on large language models (LLMs), that produce novel text or hypotheses from a prompt — and agentic AI — systems that autonomously plan and execute multi-step tasks by invoking external tools or databases in sequence — are increasingly used as documentation and research aids across biomedicine, and are beginning to be explored, in a still-nascent literature, within AYUSH (Ayurveda, Yoga and Naturopathy, Unani, Siddha, Homoeopathy) systems.

Objective: To synthesise the current and plausible near-future intersection of generative and agentic

AI with homoeopathic research, and propose a conceptual multi-agent architecture spanning case reconstruction, repertorisation support, remedy differentiation and longitudinal follow-up.

Methods: A narrative, non-systematic search of PubMed, Google Scholar, IEEE Xplore, Scopus, the Cochrane Library and the AYUSH Research Portal was conducted for literature on AI/machine learning in AYUSH systems, LLM agents, and computational tools in homoeopathy.

Findings: Computerised repertorisation software is a direct historical precursor to current tools. Peer-reviewed applications of generative or agentic AI to homoeopathy specifically remain few, concentrated in expert systems, information-extraction pilots and exploratory AI/machine-learning remedy-matching frameworks. Substantial opportunities exist for AI-assisted literature mining, review automation, and case-record standardisation, counterbalanced by tension between homoeopathy's individualising ethic and AI's generalising tendencies, by data scarcity and hallucination risk, by unresolved explainability, regulatory and liability questions, and by the risk that AI outputs reinforce unverified claims.

Conclusion: Generative and agentic AI hold genuine, presently realisable value for homoeopathic research and documentation, but clinical deployment ahead of prospective validation and regulatory clarity is not supported by current evidence. A human-in-the-loop, decision-support framing is essential throughout.

Keywords: Homeopathy; Generative artificial intelligence; Agentic artificial intelligence; Repertorization; Materia medica; Integrative medicine

INTRODUCTION

Homoeopathy remains among the most widely practised complementary systems of medicine worldwide and, in India, is one of the six recognised AYUSH systems with its own statutory education, registration and research infrastructure. Whatever position a reader takes on the underlying efficacy debate, discussed later in this review, tooling that improves the consistency, traceability and documentation quality of homoeopathic research is of practical value to the field. Over the same period that AYUSH systems have pursued digital modernisation, generative and agentic artificial intelligence have matured rapidly across biomedicine more broadly: large language models now draft clinical summaries, propose differential considerations from structured input, and, increasingly, operate as autonomous or semi-autonomous agents that plan multi-step tasks, query external databases, and return a reasoned output rather than a single answer.[1] This convergence of a data-hungry technology with a data-scarce, individualisation-first clinical system is the subject of this review.

Reviews of AI/machine-learning applications across Indian traditional-medicine systems have appeared, but they typically treat homoeopathy as one minor sub-case among Ayurveda-and yoga-dominated literatures, and they precede the shift, general across the AI field since 2023, from single-shot generative tools toward multi-step agentic architectures.[15] General clinical-AI reviews, meanwhile, rarely engage with homoeopathy's specific methodological character at all. This review's objective is therefore to (i) explain, for a readership that does not need convincing of homoeopathy's clinical framework but may be less familiar with current AI terminology, the distinction between generative and agentic AI; (ii) map, narratively, where these technologies currently stand in relation to homoeopathic research and practice; (iii) propose a conceptual multi-agent architecture for case reconstruction, repertorisation support and remedy differentiation; and (iv) set out, without editorialising on the efficacy debate itself, the methodological opportunities and the specific limitations this combination raises. The wider AI-in-medicine literature has already argued that the central task is convergence rather than replacement — using computation to extend, not substitute for, clinical judgement[2] — and that framing is adopted throughout this review: agentic AI is treated strictly as a decision-support layer, never as an autonomous prescriber.

This paper does not argue for or against homoeopathic efficacy, a question addressed only insofar as it bears on how AI outputs built on homoeopathic data should be interpreted and curated. It is organised as a narrative review: a search-strategy statement is followed by thematic sections on foundational concepts, the current literature, methodological opportunities, critical limitations and future directions, before a discussion that weighs benefits against risks and a short conclusion.

METHODOLOGY / SEARCH STRATEGY

This is a narrative, not a systematic, review; it was not registered prospectively, no formal risk-of-bias or quality appraisal was performed on included sources, and no meta-analysis was undertaken. Its purpose is to map an emerging, fast-moving interdisciplinary field rather than to provide an exhaustive or quality-weighted census of it, and its conclusions should be read with that scope in mind.

Databases searched: PubMed/MEDLINE, Google Scholar, IEEE Xplore, Scopus, the Cochrane Library, and the AYUSH Research Portal maintained by India's Ministry of AYUSH.

Search terms, combined with Boolean operators, included: “homeopathy” OR “homoeopathy”; “repertorization” OR “repertorisation”; “materia medica”; combined with “artificial intelligence”; “machine learning”; “large language model”; “generative AI”; “agentic AI” OR “AI agent”; “natural language processing”; “retrieval-augmented generation”; “knowledge graph”; “clinical decision support"; "systematic review automation"; and "pharmacovigilance."

Date range: no lower bound was applied to classical homoeopathic source texts or to the mainstream evidence-base literature; literature on AI, LLM and agentic-AI methods was drawn predominantly from January 2015 to July 2026, with a deliberate emphasis on the last five years given the pace of change in generative and agentic AI specifically.

Inclusion: peer-reviewed journal articles, indexed conference proceedings, and primary regulatory or health-technology-assessment reports, in English, addressing AI/machine-learning/LLM methods applied either to homoeopathy or AYUSH systems, or to general clinical documentation, evidence synthesis, or agentic workflows with plausible relevance by analogy. **Exclusion:** promotional or vendor material lacking methodological description, and any source whose bibliographic details could not be independently verified by the authors.

As a narrative review, source selection is subject to the authors' judgement and to the limits of the databases searched, and the fast-moving AI literature means some recent preprints will inevitably have been missed; this is a limitation returned to in the Discussion. Reference verification was prioritised over exhaustiveness throughout: every reference in this manuscript was independently located and checked against a primary source, and any claim for which the authors could not locate a verifiable citation is flagged in the text as such rather than left unsupported or attached to an invented source.

1. FOUNDATIONS

1.1 Principles of Homoeopathic Practice

Homoeopathic prescribing rests on individualisation: treatment is directed at the patient's unique, complete symptom picture rather than at a diagnostic category as such.[3] Organon-based case-taking aims to elicit the totality of symptoms — the hierarchically weighted set of mental, general and particular symptoms, together with modalities (factors that make a complaint better or worse) — that together individuate the case.[3] Repertorisation then cross-references the elicited rubrics against a repertory, classically Kent's, to generate a shortlist of candidate remedies,[4] which are differentiated against materia medica sources describing each remedy's characteristic symptom picture in depth.[5] This is, in its own terms, a structured but effort-intensive and expertise-dependent information task — precisely the category of workflow for which generative and agentic AI tools are frequently proposed in other fields. This section describes that internal logic descriptively, without adjudicating homoeopathy's clinical efficacy, which is addressed explicitly and separately in Section 4.5.

1.2 Generative AI and Agentic AI: Conceptual Distinctions

Generative AI denotes models, typically large language models (LLMs), that produce novel text or content conditioned on a prompt — for example, drafting a case summary, suggesting differential remedies from a symptom list, or paraphrasing a materia medica entry. Contemporary LLMs are trained on very large text corpora and have been shown, in general clinical contexts, to encode substantial biomedical knowledge retrievable through natural-language prompting.[6] **Agentic AI** denotes a further step: systems in which a

controlling model plans a multi-step task, invokes external tools, databases or software, observes the results, and iterates before returning an output. A widely used general-purpose pattern for this, ReAct, interleaves explicit reasoning traces with tool-invoking actions so that each step in an agent's plan is both executed and explained.[7] The essential distinction for this review is that generative AI produces an output, whereas agentic AI executes a workflow of which generative content is only one component, alongside retrieval, computation and tool use.

A specific technique of direct relevance here is retrieval-augmented generation (RAG), in which model outputs are grounded in passages retrieved from a defined document or knowledge base rather than drawn solely from the model's parametric memory, reducing (though not eliminating) the risk of fabricated content.[8] Applied to homoeopathy, RAG implies grounding any generative output in a curated repertory or materia medica corpus rather than in unverified web text, a point developed in Sections 4.2 and 5.1. A further relevant technique is counterfactual explanation: rather than presenting a single output, a system identifies what would need to change in the input for a different output to result — for instance, which modality, if altered, would favour a different remedy — which has been proposed in the wider explainable-AI literature as a way of making model reasoning auditable to a human expert.[9] Two points should be stated plainly and carried through the rest of this review. First, real homoeopathic case records are multimodal in practice — typed notes, scanned or handwritten records, and increasingly audio-recorded interviews — so any applied architecture must handle text, image (via optical character recognition) and speech input, not text alone. Second, nothing in this review implies that agentic AI systems of the kind described currently operate autonomously in clinical homoeopathic practice; throughout, agentic AI is discussed strictly as a decision-support layer requiring qualified human oversight at every step, never as an autonomous prescriber.

Table 1. Conceptual distinction between generative and agentic AI, as applied throughout this review.

Attribute	Generative AI	Agentic AI
Core function	Produces text or content from a prompt (e.g., a draft case summary or differential list)	Plans and executes a multi-step workflow, invoking tools or databases and iterating on results
Typical output	A single response — text, summary or suggestion	A completed task with an audit trail of intermediate steps
Example in this review	LLM-assisted materia medica synthesis (Section 2.2)	Multi-agent case-to-follow-up pipeline (Section 2.3; Figure 1)
Grounding	May rely on parametric memory alone unless separately paired with retrieval	Typically incorporates retrieval or tool use (e.g., RAG) as part of its plan
Current maturity in homoeopathy	Early pilots (Section 2.2)	Conceptual; no validated deployed system located (Section 2.3)
Required safeguard	Human review of generated content before use	Human-in-the-loop sign-off at every stage, not only the last

2. CURRENT LANDSCAPE: A NARRATIVE REVIEW

Table 2. Maturity of AI and agentic-AI applications across the homoeopathic research-to-practice pipeline, as identified in this narrative review.

Pipeline stage	Representative technique	Current maturity[a]	Key source(s)
Case intake & transcription	Multimodal input handling (text, OCR, speech-to-text)	Conceptual	—
Symptom / rubric extraction	NLP-based information extraction	Pilot	[12]
Repertorisation	Database retrieval (classical); RAG-grounded agent (proposed)	Established (classical) / Conceptual (agentic)	[10]
Remedy differentiation	Multi-agent comparison; counterfactual reasoning	Conceptual	[14]
Literature / evidence synthesis	AI-assisted screening & risk-of-bias tagging	Established (general biomedicine) / Emerging (AYUSH)	[16] [15]
Pharmacovigilance surveillance	Agentic literature / case-report monitoring	Conceptual	—
Follow-up / outcome tracking	Longitudinal-adaptation agent	Conceptual	—

[a] *Maturity is assessed by the authors from the literature located in this narrative search (Section "Methodology / Search Strategy") and should be read as illustrative, not as a formal technology-readiness audit.*

2.1 Historical Precursors: Repertorisation Software and Expert Systems

Computer-assisted repertorisation predates the generative-AI era by decades. Established database-driven repertory tools — including RadarOpus, MacRepertory, ReferenceWorks and Synergy — digitised Kent-style repertories to support rapid rubric cross-referencing.[10] Theseare rule-based retrieval systems: they rank pre-existing rubric–remedy associations rather than generating new text or autonomously executing a multi-step plan, and so are neither generative nor agentic AI in the sense defined in Section 1.2, though they are the direct technological ancestor of both. A further bridge technology is the rule-based expert system: one recent example built a forward-chaining expert system, following a conventional expert-system development lifecycle, to suggest homoeopathic remedies from patient-reported symptoms and complaints.[11] Tools of this kind matter for this review not because they are generative or agentic themselves, but because they have already normalised computer-assisted repertorisation among practitioners, which plausibly eases the path toward the natural-language and multi-step tools discussed below.

2.2 Emerging Generative-AI Applications

A small but directly relevant empirical literature has begun to apply natural-language processing to homoeopathic text. One early study extracted disease and symptom information directly from patient-authored queries in order to propose candidate homoeopathic remedies, describing itself as the first published attempt to do so in this domain.[12] More recently, a proposed framework has explored turning materia medica into a structured, queryable symptom–remedy dataset and applying natural-language processing to extract usable information from free-text symptom descriptions, with the stated aim of narrowing remedy selection in fewer attempts.[10] Beyond homoeopathy specifically, machine-learning methods have been applied to other AYUSH domains — for example, predicting Ayurveda-based constitutional (dosha) balance from patient data using standard classifier methods[13] — illustrating that the broader AYUSH research community is already engaging with AI/ML tooling even where homoeopathy-specific applications remain sparse. General-purpose LLM summarisation capability, well documented in mainstream biomedicine,[6] is a natural extension for literature surveillance relevant to homoeopathic pharmacovigilance, though the authors are not aware of a validated, homoeopathy-specific published implementation of this kind at the time of writing. Across this subsection, the applications described are best characterised as small-scale demonstrations rather than validated clinical tools, and no claim is made that any deployed, clinically validated generative-AI system is in routine homoeopathic use.

2.3 Emerging Agentic-AI Applications

No published, validated agentic system spanning the full homoeopathic case-to-follow-up pipeline was located in this search; the architecture proposed in Section 3 below and illustrated in Figure 1 is accordingly a conceptual synthesis, built from general agentic-AI design patterns[7] applied to the homoeopathic workflow described in Section 1.1, not a description of an existing deployed system. Components of such a pipeline do, however, have precedent elsewhere in clinical AI. Multi-agent conversational frameworks, in which separate model instances argue independently for competing conclusions before a synthesis step, have been piloted in general clinical decision-making specifically to counteract single-model cognitive bias,[14] a pattern directly transferable to remedy differentiation, where independent agents could argue for competing remedy candidates ahead of a human-reviewed synthesis. Agentic automation of research-data curation and systematic-review screening is well established generically (see Section 3.2) and is beginning to be proposed, though not yet widely validated, for AYUSH evidence synthesis specifically.[15] Network-pharmacology-style analysis — mapping constituents to biological targets and pathways — is well established for Ayurvedic and other herbal formulations but has essentially no published precedent for homoeopathic proving data, since the ultra-dilute preparation central to homoeopathic pharmacy removes the direct molecular-target mapping on which network pharmacology depends; this is best treated as an open methodological question rather than an existing application. As throughout this review, every step of the pipeline concept summarised here terminates in clinician review before any patient-facing use: the architecture is proposed strictly as a decision-support and documentation layer.

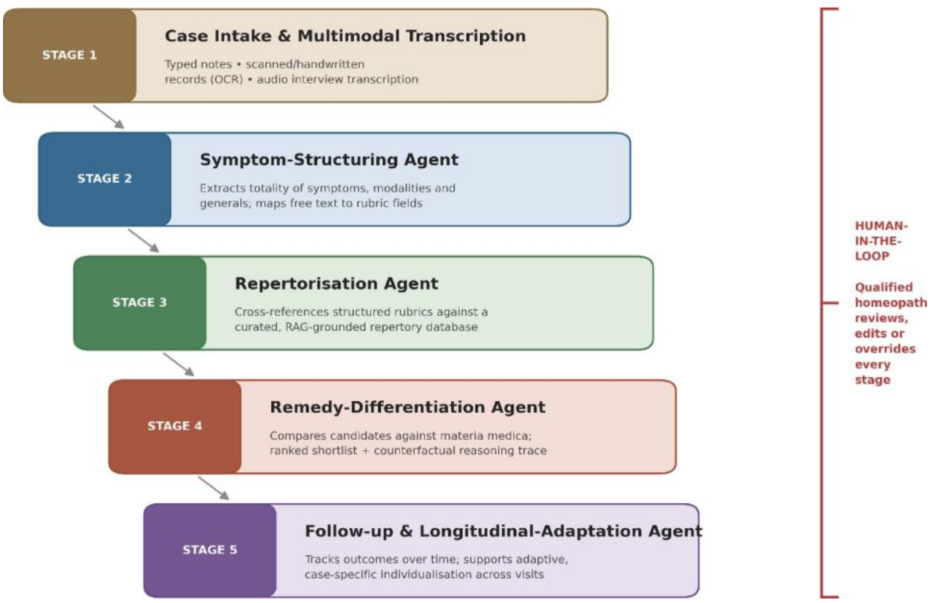

Figure 1. *Proposed five-stage multi-agent architecture for homoeopathic case reconstruction, repertorisation support and longitudinal follow-up. The architecture is conceptual (Section 2.3); every stage is subject to review, edit or override by a qualified homeopath before it informs the case (Section 4.6).*

3. METHODOLOGICAL OPPORTUNITIES FOR HOMOEOPATHIC RESEARCH

3.1 Mining Provings and Materia Medica for Pattern Discovery

Historical proving records and materia medica texts constitute a large, richly symptomatic, but under-digitised corpus. Natural-language methods for symptom and modality extraction, of the kind piloted on patient-authored homoeopathic text,[12] could plausibly be adapted to support rubric-consistency checking and pattern discovery across historical sources, treated as a method transfer from general clinical NLP rather than an established homoeopathy-specific pipeline at present. It is important to keep the research use distinct from the clinical use: pattern discovery for generating hypotheses is a materially lower-stakes application than remedy recommendation at the point of care, and the two should not be conflated in study design or in any future regulatory treatment.

3.2 Automating Systematic Review and Meta-Analysis Components

AI-assisted screening tools are validated, with appropriate human oversight, in general biomedical evidence synthesis. Robot Reviewer, for instance, uses machine learning to flag risk-of-bias domains in randomised-trial reports for reviewer verification rather than as a replacement for reviewer judgement.[16] A prospective evaluation of AI-assisted screening in systematic reviews concluded the approach was promising specifically "when appropriately used" — that is, with human verification of AI-flagged decisions.[17] A dedicated scoping review catalogues dozens of such tools across the review pipeline, from search through screening to data extraction.[18] Homoeopathy trial systematic reviews and meta-analyses, an area with a substantial existing literature (Section 4.5), are a direct candidate for adopting these tools for initial screening and risk-of-bias tagging, always with human adjudication of contested inclusion decisions and of any AI-proposed rating, and reported according to standard guidelines such as PRISMA.[20] Ethical governance of any such AI use should follow established international guidance on AI in health research and practice.[19]

3.3 AI-Assisted Hypothesis Generation for Proving Studies

LLMs can plausibly propose candidate rubric refinements, novel symptom clusters, or under-studied remedy comparisons for human experts to evaluate and, where warranted, take forward into prospectively designed observational or proving studies. This role should be understood strictly as hypothesis-generation support: an LLM-suggested pattern has no evidentiary status in itself until independently tested by qualified homeopaths and researchers. This use also directly intersects the risk discussed in Section 4.1: hypothesis generation of this kind must be designed to preserve, rather than dissolve, case-level individuality, for instance by surfacing candidate patterns for expert scrutiny rather than by silently averaging across cases.

3.4 AI-Supported Standardisation of Case Recording

Inconsistent, non-standardised case documentation is a recognised weakness of the homoeopathic evidence base, complicating both retrospective analysis and trial replication.[31] Structured-input tools, template-driven transcription support, and NLP-assisted coding of free-text case notes into standard rubric and modality fields could plausibly improve documentation consistency without constraining the clinician's qualitative case-taking process itself. This is best framed as a dual-benefit, evidence-neutral improvement: better documentation serves research quality regardless of a reader's prior view on efficacy, because it is an improvement to research infrastructure rather than a claim about outcomes.

4. CRITICAL CHALLENGES AND LIMITATIONS

4.1 Individualisation versus Generalisation: the "Average-Casing" Risk

Homoeopathy's stated method is explicitly anti-average: the same named condition may warrant different remedies in different individuals, based on the total symptom picture rather than the diagnostic label.[3] Statistical and machine-learning pattern-matching, LLMs included, is fundamentally built on regularities learned across many cases; naively applied, it risks collapsing individualised case pictures toward a population-level modal answer — in effect, "average-casing" a therapy whose internal logic explicitly rejects the average case. This is not a peripheral technical inconvenience but a direct conceptual tension that any architecture claiming to support "precision" or "individualised" homoeopathy must address explicitly — for instance, through case-specific retrieval and counterfactual differentiation rather than classification against population norms — rather than assume away.

4.2 Data Scarcity, Quality, and Hallucination Risk

Large, structured, standardised homoeopathic datasets suitable for training or rigorously grounding AI systems are scarce. A systematic review of AI/machine-learning applications across Indian traditional-medicine systems located only a small handful of homoeopathy-specific studies against a much larger number in yoga and Ayurveda — direct, quantified evidence of this data gap.[15] LLMs are well documented to hallucinate — to produce fluent but false content — and this risk rises precisely where a query falls outside dense training-data coverage,[21] a description that fits niche domains such as homoeopathic materia medica closely; reviews aimed at clinicians have accordingly flagged hallucination as a primary risk requiring active mitigation before clinical use of LLMs of any kind.[22,23] Retrieval-augmented grounding in a verified corpus mitigates, but cannot

eliminate, this risk,[8] and is only ever as reliable as the underlying knowledge base against which it retrieves
— itself, per the paragraph above, a scarce resource in this domain.

4.3 Explainability and Clinician Trust

Black-box remedy suggestions are unlikely to be adopted by, or appropriate for, practitioners trained to reason explicitly from symptom to rubric to remedy; the profession's own methodology is inherently explanation-oriented. This is consistent with the wider human-in-the-loop clinical-AI literature, which links structured human oversight and transparent system design to improved outcomes and higher clinician trust relative to either fully automated or fully manual workflows.[24] Any agentic system in this space should therefore be required to expose a reasoning trace — which rubrics and symptoms drove which remedy candidates, and why one was preferred over another, including counterfactual differentiation of the kind introduced in Section 1.2 — rather than a bare ranked list, so that a clinician can audit, rather than simply accept or reject, its output.

4.4 Regulatory, Ethical and Consent Issues; Liability

Deployment of AI decision support in clinical homoeopathic practice raises informed-consent questions (patients should know when AI has contributed to a recommendation), data-protection questions (case records are sensitive health data), and liability questions: responsibility for an AI-influenced remedy choice remains, under current professional norms, with the licensed practitioner, but how liability should be allocated across developer, institution and clinician is not yet settled in most jurisdictions[32]. Regulatory pathways for AI-based decision support are still maturing even in conventional biomedicine, and general international guidance on the ethics and governance of AI in health is available and directly applicable to research use in this domain,[19] but AYUSH-specific regulatory guidance for AI decision-support tools is not yet established to the authors' knowledge — a further reason for caution before any clinical, as opposed to research, deployment.

4.5 Risk of Reinforcing Unverified or Non-Evidence-Based Claims

AI systems trained or grounded on uncurated web text risk amplifying folklore, marketing claims or outdated clinical assertions circulating about specific remedies, and presenting them with the same fluent confidence as well-supported statements. This risk is sharpened by the fact that homoeopathy's own evidence base for clinical efficacy remains a live, unresolved debate in the mainstream biomedical literature, which this review states plainly rather than adjudicates. Government and academic health-technology assessments — including a 2015 Australian National Health and

Medical Research Council review of systematic reviews across sixty-one conditions[25] and a 2010 UK House of Commons Science and Technology Committee evidence check[26] — concluded there was no reliable evidence of homoeopathic effect beyond placebo, a conclusion also reached by an influential 2005 comparative meta-analysis of homoeopathy and conventional-medicine trials[27] and by subsequent Cochrane reviews of specific conditions.[28,29] These conclusions have themselves been disputed on methodological grounds by homeopathy-research bodies and some independent commentators, and separately, basic-science investigation continues into the physicochemical properties of high-dilution preparations, including reports of nanoparticulate retention of starting material at dilutions beyond the Avogadro limit[30] — findings that remain a matter of active scientific discussion and that do not, on their own, establish clinical efficacy. The debate is accordingly presented here as unresolved in the literature rather than settled in either direction [33]. The implication for this review is direct: any AI tool built on homoeopathic case or remedy data inherits this underlying epistemic uncertainty regardless of its computational sophistication. Better symptom processing does not, by itself, resolve open questions about the clinical efficacy of the remedies being recommended, and AI-generated content in this domain must therefore be curated by qualified homeopaths rather than presented as independently validating the therapy.

4.6 Necessity of Human-in-the-Loop Design

Synthesising Sections 4.1–4.5: every risk identified above is mitigated, though not eliminated, by keeping a qualified clinician as the final decision-maker and by architecting agentic systems explicitly as decision support rather than as autonomous-prescribing tools. This is consistent with the wider human-in-the-loop clinical-AI literature, which links structured human oversight to improved diagnostic accuracy and higher clinician trust relative to either fully automated or fully manual workflows, while also documenting that clinician-in-the-loop evaluation is still inconsistently reported in the published literature.[24] As a design principle for the remainder of this review and for any system it informs: agentic AI, wherever discussed, denotes a decision-support layer that proposes, explains and defer- never one that prescribes.

5. FUTURE DIRECTIONS

5.1 Validated, Ethically-Sourced Homoeopathic Knowledge Graphs for Retrieval-Augmented Generation

A practical roadmap has three steps. First, qualified homeopaths and information scientists should jointly curate a structured knowledge graph linking rubrics, remedies, materia medica sources and provenance or evidence-grade metadata. Second, this curated graph, not open web text, should serve as the grounding corpus for any retrieval-augmented tool,[8] directly

addressing the hallucination risk described in Section 4.2. Third, the graph itself should be versioned, audited and published as a citable research artefact independent of any downstream AI tool built on it, so that its provenance and quality can be assessed on its own terms. Ethical sourcing should be treated as a first-order design requirement: proper attribution of classical texts, respect for institutional and practitioner data ownership, and avoidance of uncurated scraped forum content as a knowledge source, given the risks set out in Section 4.5.

5.2 Multi-Agent Research Assistants for Pharmacovigilance and Post-Marketing Surveillance

The screening-automation opportunity described in Section 3.2 extends naturally to continuous, agentic monitoring of literature and case reports for adverse-event signals associated with homoeopathic products, again framed as decision support for a researcher or pharmacovigilance officer, with mandatory human sign-off before any signal is escalated. The natural institutional home for such a system is existing AYUSH pharmacovigilance and research infrastructure rather than a standalone commercial tool, so that institutional accountability for any escalated signal is preserved.

5.3 Collaborative Frameworks Between AYUSH Institutions and AI Developers

Structured collaboration — data-sharing agreements, joint methodology committees, and shared evaluation benchmarks — between AYUSH research institutions and regulators on one side and AI developers on the other would allow domain expertise to shape system design from the outset rather than being retrofitted afterward. Such collaboration sits naturally within the wider international policy direction toward evidence-based integration of traditional and complementary medicine into health systems,[1] but its specifically AI-related components require homoeopaths, not only technologists, to define what “success” and “safety” mean for any given system, given the domain-specific risks identified in Section 4.

5.4 Call for Prospective Validation Studies

Before any clinical, as opposed to research-support, adoption, the field needs prospectively designed, pre-registered comparative studies of AI-assisted versus conventional repertorisation and case management — for example, comparing time-to-shortlist, inter-practitioner agreement, and downstream outcome measures, with independent oversight. Until such evidence exists, the appropriate role for these tools remains research acceleration and documentation support, delivered through a clinician-supervised decision-support role, not autonomous or lightly supervised clinical decision-making.

DISCUSSION

The throughline across Sections 2 to 5 is that generative and agentic AI are, at present, infrastructure technologies for homoeopathic research and documentation rather than validated clinical tools: the literature located in Section 2 is consistently pre-clinical and small in scale, or transferred by analogy from general biomedical AI rather than being homoeopathy-specific and independently validated. Weighed against one another explicitly, the benefits identified in this review cluster around research acceleration (Sections 3.1–3.3), documentation standardisation (Section 3.4), and pharmacovigilance capacity (Section 5.2)
— improvements to the evidentiary infrastructure of the field that are valuable regardless of a reader's prior stance on efficacy. The risks cluster around the individualisation–generalisation tension (Section 4.1), which is specific to, and especially acute for, homoeopathy among medical systems, alongside hallucination, explainability, regulatory and misinformation-amplification risks that are shared with AI-in-medicine generally but land with particular force in a low-data, individualisation-first domain.

There is an important asymmetry here. The benefits described are largely available now, at comparatively low clinical risk, if AI use is confined to research and documentation roles — screening, curation, structuring. The risks escalate sharply, and the evidence for benefit becomes essentially absent, once AI outputs approach the point of care and remedy selection without a qualified human as the final decision-maker. The "average-casing" tension (Section 4.1) is worth addressing constructively rather than only flagging: architectures that retrieve and differentiate among case-specific precedents — case-based reasoning, counterfactual explanation — are conceptually better aligned with homoeopathic individualisation than architectures that classify a new case against a population-trained model. This is a first-order design requirement, not an afterthought, for any AYUSH-facing AI system.

On the evidence-base question addressed in Section 4.5, this review takes no position beyond reporting the state of the debate accurately: whatever a reader concludes about homoeopathic efficacy, better AI-assisted documentation and more efficiently screened, more rigorous evidence synthesis serve the shared goal of that debate eventually being resolved on stronger evidence — itself a reason for cautious, well-governed adoption of these tools in the research domain even amid ongoing disagreement about the underlying therapy. Relative to prior AYUSH-AI surveys, which have concentrated mainly on image- and sensor-based machine learning in Ayurveda and yoga,[15] this review's emphasis on agentic, multi-step, tool-using architectures specifically reflects how quickly the general AI field has moved from generative to agentic paradigms since most existing AYUSH-AI literature was written, and points to a genuine gap for future homoeopathy-specific work to fill.

This review has the limitations inherent to its narrative design: it is not systematic, it is restricted to English-language sources, and it addresses a literature that is moving quickly enough to date parts of this synthesis within a short period. A formal systematic or scoping review, ideally conducted jointly by AI researchers and qualified homeopaths and pre-registered in advance, would be a valuable and logical next step.

CONCLUSION

Generative and agentic AI offer genuine, presently realisable value to homoeopathic research: faster literature screening, more consistent case documentation, and structured hypothesis generation for future proving and observational studies. That promise is real but bounded — it applies most clearly to research and documentation infrastructure, not to autonomous or lightly supervised clinical remedy selection, for which neither the evidence base nor the regulatory framework yet exists. The design requirement running through this review is architectural rather than merely procedural: systems in this space should be built from the outset as human-in-the-loop decision support, with transparent and auditable reasoning traces, grounded in curated and ethically sourced knowledge rather than uncurated text — precisely because homoeopathy's individualising ethic and its contested evidence base both counsel particular caution. The authors regard this as a moment for careful, collaborative, evidence-generating exploration between AYUSH institutions and AI developers, not for premature clinical deployment.

ACKNOWLEDGEMENTS

The authors thank the respective institutions of both authors for infrastructural and academic support during the preparation of this manuscript.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

REFERENCES

1. World Health Organization. *Global traditional medicine strategy 2025–2034*. Geneva: World Health Organization; 2025. ISBN: 978-92-4-011317-6.
2. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56.
3. Hahnemann S. *Organon of medicine*. 6th ed. Boericke W, translator. New Delhi: B. Jain Publishers; 2002.
4. Kent JT. *Repertory of the homoeopathic materia medica*. New Delhi: B. Jain Publishers; 2005.
5. Boericke W. *Pocket manual of homoeopathic materia medica and repertory*. 9th ed. New Delhi: B. Jain Publishers; 1996.
6. Singhal K, Azizi S, Tu T, et al. Large language models encode clinical knowledge. Nature. 2023;620(7972):172-180.
7. Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, et al. ReAct: synergizing reasoning and acting in language models. In: Proceedings of the 11th International Conference on Learning Representations (ICLR); 2023 May 1-5; Kigali, Rwanda.
8. Lewis P, Perez E, Piktus A, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. Adv Neural Inf Process Syst. 2020;33:9459-9474.
9. Ribeiro MT, Singh S, Guestrin C. “Why should I trust you?”: explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. p. 1135-1144.
10. Ghadiyaram R, Morusu V, Krishnamoorthy D, Eripilla J. Revolutionizing homeopathy: integrating data analytics and AI/ML for precision remedy. In: Bhattacharya S, editor. *ICT for Global Innovations and Solutions (ICGIS 2025)*. Advances in Computer Science Applications and Research. Vol. 1. Cham: Springer; 2026. p. 572-588.
11. Mulyani A, Kurniadi D, Ahmad M, Fatimah DDS. Expert system development for homeopathy medicine. IOP Conf Ser Mater Sci Eng. 2021;1098:032058.
12. Priyadarshi A, Saha SK. Web information extraction for finding remedy based on a patient-authored text: a study on homeopathy. Netw Model Anal Health Inform Bioinforma. 2020;9:9.
13. Madaan V, Goyal A. Predicting Ayurveda-based constituent balancing in human body using machine learning methods. IEEE Access. 2020;8:65060-65070.

14. Ke Y, Yang R, Lie SA, et al. Mitigating cognitive biases in clinical decision-making through multiagent conversations using large language models: simulation study. J Med Internet Res. 2024;26:e59439.
15. Muthuperumal P, Karmegam D, Mappillairaju B. Status of artificial intelligence and machine learning in Indian traditional medicine systems: a systematic review. J Med Pharm Chem Res. 2024;6(8):1173-1187.
16. Marshall IJ, Kuiper J, Wallace BC. RobotReviewer: evaluation of a system for automatically assessing bias in clinical trials. J Am Med Inform Assoc. 2016;23(1):193-201.
17. van Dijk SHB, Brusse-Keizer MGJ, Bucsán CC, et al. Artificial intelligence in systematic reviews: promising when appropriately used. BMJ Open. 2023;13:e072254.
18. Khalil H, Ameen D, Zarnegar A. Tools to support the automation of systematic reviews: a scoping review. J Clin Epidemiol. 2022;144:22-42.
19. World Health Organization. *Ethics and governance of artificial intelligence for health: WHO guidance*. Geneva: World Health Organization; 2021.
20. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ. 2021;372:n71.
21. Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of hallucination in natural language generation. ACM Comput Surv. 2023;55(12):1-38.
22. Azamfirei R, Kudchadkar SR, Fackler J. Large language models and the perils of their hallucinations. Crit Care. 2023;27:120.
23. Roustan D, Bastardot F. The clinicians' guide to large language models: a general perspective with a focus on hallucinations. Interact J Med Res. 2025;14:e59823.
24. Olawade DB, Plabon SB, Ojo A, Ogunbona MA, Makanjuola BD, Olasilola OR. Human in the loop artificial intelligence in healthcare: applications, outcomes, and implementation challenges. Int J Med Inform. 2026;213:106362.
25. National Health and Medical Research Council. *NHMRC information paper: evidence on the effectiveness of homeopathy for treating health conditions*. Canberra: National Health and Medical Research Council; 2015.
26. House of Commons Science and Technology Committee. *Evidence check 2: homeopathy*. Fourth report of session 2009-10. HC 45. London: The Stationery Office; 2010.
27. Shang A, Huwiler-Müntener K, Nartey L, Jüni P, Dörig S, Sterne JAC, et al. Are the clinical effects of homoeopathy placebo effects? Comparative study of placebo-controlled trials of homoeopathy and allopathy. Lancet. 2005;366(9487):726-732.

28. Mathie RT, Frye J, Fisher P. Homeopathic Oscillococcinum for preventing and treating influenza and influenza-like illness. Cochrane Database Syst Rev. 2015;(1):CD001957.
29. Mathie RT, Ramparsad N, Legg LA, et al. Randomised, double-blind, placebo-controlled trials of non-individualised homeopathic treatment: systematic review and meta-analysis. Syst Rev. 2017;6(1):63.
30. Chikramane PS, Suresh AK, Bellare JR, Kane SG. Extreme homeopathic dilutions retain starting materials: a nanoparticulate perspective. Homeopathy. 2010;99(4):231-242.
31. Teut M, van Haselen RA, Rutten L, Lamba CD, Bleul G, Ulbrich-Zürni S. Case reporting in homeopathy: an overview of guidelines and scientific tools. Homeopathy. 2022;111(1):2-9.
32. Bhargava M, Bhardwaj P, Dasgupta R. Artificial intelligence in biomedical research and publications: it is not about good or evil but about its ethical use. Indian J Community Med. 2024;49:777-779. doi:10.4103/ijcm.ijcm_560_24.
33. Wilhelm M, Hermann C, Rief W, Schedlowski M, Bingel U, Winkler A. Working with patients' treatment expectations – what we can learn from homeopathy. Front Psychol. 2024;15. doi:10.3389/fpsyg.2024.1398865.

How to cite this article:

Dholariya P, Mandaliya P. The future of precision homoeopathy: a multimodal agentic intelligence architecture for generative case reconstruction, counterfactual remedy reasoning and adaptive longitudinal individualisation. International Journal of AYUSH. 2026;15(9):1-18**.**