Astanga Wellness Education Series

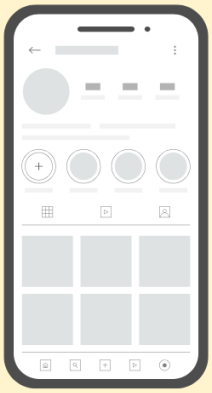

"What would the Ayurvedic equivalent look like?"

THE APP THAT DOESN'T EXIST

Why Ayurveda Has No Clinical Decision Support Tools — and What It Would Take to Build Them

"The system prefers the archive. The patient needs the calculator."

Dr Aakash Kembhavi

2026

The App That Doesn't Exist

Why Ayurveda Has No Clinical Decision Support Tools — and What It Would Take to Build Them

Dr. Aakash Kembhavi

MD(Ayu), PGDMLS, MS(Counseling & Psychotherapy)

Published under the Astanga Wellness Education Series Imprint, May 2026

A One-Page Provocation

It began with an Instagram post.

A doctor — going by the handle @drxshaw — shared a list of apps she actually relies on during duty and study. Simple, practical, worth having. She was careful to note: not sponsored, just the ones that genuinely make her life easier.

The list included Full Code, for practical training in real emergency scenarios. MDCalc, for evidence-based scores and calculators, right when you need them. India Drug Index, for reliable drug information for everyday prescribing. Prognosis: Your Diagnosis, for case-based learning that sharpens clinical thinking. And a 3D anatomy tool for interactive visual learning.

I read that list and felt something I can only describe as productive discomfort.

Not envy, exactly. More like the sensation of holding two realities next to each other and noticing how far apart they are. Because what that list represents is not merely a collection of useful apps. It represents an entire ecosystem — decades of peer-reviewed research, validated clinical decision rules, consensus protocols, and shared data infrastructure — that has been operationalized into tools a junior doctor can use at three in the morning in a district hospital when a patient's life depends on getting the decision right.

Then I asked myself: what would the equivalent list look like for an Ayurvedic practitioner?

The honest answer is that it would not exist. Not because Ayurveda lacks clinical depth. Not because our tradition has nothing to offer in the clinic. But because the upstream conditions that make such tools possible — shared protocols, outcome-linked data, operationalized constructs, reproducible decision rules — have not been built. And in some important ways, our tradition has actively resisted building them.

That realization is what produced this document.

What follows is not a lament. It is an attempt to understand, with as much honesty as I can muster, why Ayurveda finds itself in this position — and what it would actually take to change it. I have written it as someone who has spent 35 years inside this tradition, who loves it, who has staked a career on its reform, and who believes that the most dangerous thing we can do for Ayurveda is protect it from scrutiny.

The tradition does not need our protection. It needs our honesty.

How to Read This Document

This monograph is organized as three connected arguments, not three independent articles. Each builds on the one before it.

Article One asks why Ayurveda has no shared clinical data pool — and traces the answer to a deep structural feature of how our tradition transmits knowledge: the Guru-Shishya model and its contemporary distortions.

Article Two asks the harder question: if the problem is visible, why hasn't it been solved? The answer lies not in ignorance but in interest — in who benefits from the current opacity and what they stand to lose if clinical decision-making becomes transparent and reproducible.

Article Three turns toward what is possible. Not a fantasy roadmap, but an honest account of the upstream infrastructure Ayurveda needs to build before a clinical decision support tool can mean anything — and what a realistic first decade of that effort might look like.

Read them in sequence. The third article will feel naive without the second. The second will feel cynical without the third.

Article One

The Cult of the Guru and the Data Pool That Never Forms

There is a passage in the Charak Samhita that has always struck me as both beautiful and dangerous.

Charaka describes the fully realized physician — the one who has mastered all four *padas* of medical knowledge: *Shastra* (textual learning), *Adhyayana* (study under a teacher), *Loka* (observation of the world), and *Tantra Yukti* (logical synthesis). He then notes, almost in passing, that such a physician is rare. The rarity is not presented as a problem. It is presented as the nature of things — as the inevitable outcome of a tradition that demands total formation of the whole person, not merely the acquisition of technical skills.

I have quoted that passage in classrooms for decades. I have meant it admiringly.

But I have come to believe that we have built an entire professional culture on the wrong reading of it.

We have taken Charaka's description of a rare ideal and turned it into a justification for keeping clinical knowledge personal, non-transferable, and immune from external verification. We have romanticized the exceptional Vaidya to the point where we have no system for the competent average one. And in doing so, we have made it structurally impossible to build the kind of shared clinical data infrastructure that modern healthcare — and modern patients — require.

The Guru-Shishya Model as Epistemological Architecture

The Guru-Shishya model is not merely a pedagogical arrangement. It is an epistemological one. It encodes a specific theory of how knowledge works: that clinical wisdom is fundamentally personal, contextual, and transmitted through proximity rather than protocol. The student learns by watching. By absorbing. By slowly internalizing a way of seeing that cannot be fully articulated, only demonstrated.

This is not wrong. There is genuine clinical wisdom that resists reduction to algorithms. Any experienced clinician — Ayurvedic or otherwise — knows that some of what they do in a consultation cannot be fully explained. It is pattern recognition at a level of complexity that formal language struggles to capture.

The problem is not the model. The problem is what the model has become in its contemporary distortions.

In its classical form, the Guru-Shishya relationship was bounded by rigorous expectations on both sides. The Guru was accountable — to the tradition, to outcomes, to the student's eventual independence. The student was expected to ultimately surpass the teacher, not merely replicate him.

What we have instead, in many Ayurvedic institutions and lineages today, is something closer to a cult of personality. Senior clinicians whose clinical methods are treated as proprietary. Formulations that cannot be questioned. Diagnostic frameworks that are lineage-specific and therefore non-comparable. Students who are trained to revere rather than to interrogate. And — most damaging of all — an implicit understanding that the senior Vaidya's clinical authority rests precisely on its unverifiability.

You cannot challenge what you cannot measure. And in a system where clinical outcomes are never systematically recorded, never compared, never subjected to external review, the Guru's authority is perfectly protected.

What a Data Pool Actually Requires

To understand why this matters, consider what it would take to build something as seemingly simple as an Ayurvedic equivalent of a clinical risk calculator — say, a tool that helps a practitioner assess the severity of a patient's *Agni* dysfunction and suggests an intervention protocol.

Before any such tool can be built, several things must exist.

First, there must be an agreed operational definition of *Agni* dysfunction — one specific enough that two different practitioners examining the same patient would arrive at the same assessment. This sounds elementary. It does not currently exist in any standardized form. Ask ten senior Vaidyas to define *Mandagni* in measurable clinical terms and you will receive ten different answers, each rooted in a different textual reading and a different lineage tradition.

Second, there must be a shared protocol for intervention — a defined set of options, with criteria for choosing between them, that has been applied consistently across a large enough patient population to generate meaningful outcome data.

Third, those outcomes must be recorded. Systematically. Longitudinally. In a format that allows comparison across practitioners and institutions.

Fourth — and this is where the culture problem becomes most acute — those records must be made available for analysis. Not hoarded. Not kept as the private property of the clinician who generated them. Shared, so that the field can learn from them.

None of these conditions currently exist at scale in Ayurvedic practice. Some pioneering institutions are beginning to move in this direction. But the systemic conditions are not yet in place. And they will not be put in place by technology alone, because the obstacle is not technological. It is cultural and philosophical.

The Reproducibility Problem

Modern clinical decision support tools rest on a single foundational assumption: that a well-trained practitioner following a validated protocol should be able to reproduce outcomes across different patient populations and clinical settings. Reproducibility is not a concession to reductionism. It is the minimum condition for knowing whether a treatment actually works — as opposed to working in the hands of one gifted practitioner under specific conditions that no one has documented.

Ayurveda's current culture does not merely fail to meet this standard. It actively celebrates its failure to meet it. We speak of the exceptional Vaidya's clinical outcomes as evidence of the tradition's power. We do not ask whether those outcomes could be reproduced by a different practitioner following the same protocol. We do not ask because the question itself feels like a category error — like asking whether a master painter's brushstroke can be reproduced by a competent technician.

But medicine is not painting. When a patient comes to us in distress, they are not commissioning art. They are asking for help that is reliable, reproducible, and as free from the lottery of individual genius as possible.

Until we internalize that distinction — fully, institutionally, not just rhetorically — the data pool will not form. Not because we lack the tools to collect data. But because we have not yet decided that collecting it is a moral obligation rather than an optional scholarly exercise.

The Question This Leaves Us With

If the problem is cultural and philosophical, then solving it requires cultural and philosophical work — not just research funding and app development. It requires us to ask, honestly, what we are actually protecting when we protect the current system.

Are we protecting the tradition? Or are we protecting the authority structures that have grown up around it?

That question leads directly to the next article.

Article Two

Who Benefits From Opacity?

Let me begin with an uncomfortable observation.

The people who are best positioned to build Ayurveda's clinical data infrastructure — senior clinicians with large patient volumes, institutional access, academic standing, and the professional credibility to convene consensus — are also, in many cases, the people with the least incentive to build it.

This is not a coincidence. It is the predictable outcome of a system in which clinical authority is constructed and maintained through opacity. And understanding it is essential if we want to move beyond good intentions toward actual change.

How Authority Works Without Evidence

In evidence-based medicine, clinical authority is — at least in principle — tethered to demonstrated outcomes. A clinician's reputation is built partly on their track record, which is partially visible: through published cases, audit data, peer review, and the testimony of colleagues who have observed their practice. Authority is still unevenly distributed and still subject to the distortions of hierarchy and prestige. But there is at least a framework in which outcomes can, in principle, be examined.

In Ayurvedic practice as currently constituted, clinical authority works differently. It is built on a combination of lineage, seniority, textual erudition, confident communication, and — most importantly — the impossibility of external verification. A senior Vaidya who claims a high success rate in treating a complex condition cannot be meaningfully challenged, because there is no shared standard against which the claim can be assessed. No common definition of "success." No recorded baseline. No follow-up protocol. No comparison group.

In this environment, the Vaidya who speaks most confidently about their clinical outcomes is not necessarily the one with the best outcomes. They are the one who has most successfully managed the narrative around their practice.

This is a harsh thing to say. But it is, I believe, substantially true. And it has consequences that go far beyond individual reputation.

The Proprietary Formulation Problem

One of the most concrete expressions of the opacity problem is the proprietary formulation — the special churna, the house kwath, the family recipe that has been passed down through generations and is presented as the secret of a lineage's clinical success.

I want to be precise here, because I am not arguing against traditional knowledge or intellectual property. There are legitimate reasons why a practitioner might not wish to share every detail of their practice immediately and unconditionally. The problem arises when proprietary formulations become the primary evidence for a treatment's efficacy — when "my Guru used

this and it worked" is treated as sufficient justification for continued use, and when any request to study the formulation systematically is resisted as a threat to lineage integrity.

What is actually being protected in such cases is not the formulation. It is the authority that derives from exclusive access to the formulation. Standardize the formula, study it rigorously, share the protocol — and the mystique dissolves. The treatment may work just as well. But it works for reasons that can now be understood and replicated by others. And that is precisely what the current system is structured to prevent.

The Institutional Dimension

The opacity problem is not confined to individual practitioners. It is also embedded in institutional structures — including, notably, in AYUSH policy and research infrastructure.

Consider the trajectory of AYUSH-funded research over the past two decades. A substantial portion of it has been directed toward validating what is already believed — toward finding scientific language to describe classical concepts rather than toward genuinely testing them. The research questions are safe. The outcomes are, with depressing regularity, confirmatory. Studies that produce inconvenient findings are rarely followed up. Negative results are almost never published.

This is not a failure of individual researchers. It is a structural outcome. Researchers who challenge established formulations or question received clinical wisdom find their funding, their institutional affiliations, and their relationships with senior colleagues at risk. The incentive to produce comfortable findings is enormous. The incentive to produce disruptive but accurate ones is negligible.

The result is a body of Ayurvedic research literature that has grown substantially in volume while contributing relatively little to the kind of actionable clinical knowledge that would make a decision support tool possible. We have hundreds of studies. We have very few validated protocols.

What the System Prefers

There is a particular irony in the fact that the single most ambitious digitization project in Ayurveda's recent history — the digitization of classical texts, the development of searchable Samhita databases, the archiving of manuscript traditions — has received far more institutional support than any effort to build clinical outcome databases.

Digitizing texts is politically safe. It is culturally celebratory. It does not challenge any existing authority structure. It produces outputs that can be announced in press releases and displayed at conferences. It requires no consensus on clinical protocols, no standardization of diagnostic criteria, no uncomfortable comparison of outcomes across institutions.

Building a clinical outcomes database requires all of those things. It requires institutions to open their records to scrutiny. It requires senior clinicians to accept that their outcomes will be compared to those of their peers. It requires the field to acknowledge, in writing and in public, that some approaches work better than others — which implies, inescapably, that some approaches work worse.

This is why we have digitized the Charaka Samhita seventeen different ways and still cannot tell a patient with any statistical confidence what their probability of improvement is with a given Ayurvedic intervention for a given condition.

The system prefers the archive. The patient needs the calculator.

The Younger Generation and the Same Trap

It would be comforting to believe that the next generation of Ayurvedic practitioners will simply do things differently. And there is genuine cause for hope — younger Vaidyas are more research-literate, more connected to global conversations, more aware of the gap between what Ayurveda claims and what it can currently demonstrate.

But the incentive structures they inherit are largely unchanged. The path to clinical reputation still runs through association with a respected senior. Publications still tend to reward confirmation over disruption. Institutional positions still go to those who navigate hierarchy skillfully rather than those who challenge it productively.

A reformist-minded young Vaidya who decides to collect systematic outcome data, publish protocol-based studies, and openly discuss cases where Ayurvedic intervention did not produce the expected result will find the process professionally costly in ways that are difficult to sustain without institutional support that currently does not exist at scale.

This is not a counsel of despair. It is a description of the terrain that any serious reform effort must navigate. And navigating it requires being honest about where the resistance will come from — not just the inertia, but the active interests.

What This Means for the App

If you have followed the argument so far, you will understand why the question "when will Ayurveda have its MDCalc?" cannot be answered with a technology roadmap.

The app is not the problem. The app is the last mile. The problem is the fifty miles before it — the shared constructs, the outcome data, the validated protocols, the institutional willingness to be scrutinized — none of which will materialize automatically, and none of which will materialize at all without confronting the interests that currently benefit from their absence.

That confrontation is uncomfortable. It implicates people and institutions that many of us respect and are professionally connected to. It requires us to say, clearly and publicly, that the current system is not merely incomplete — it is, in some of its features, actively harmful to the patients it claims to serve.

But that confrontation is also the only honest starting point for the constructive work that comes next.

Article Three

From Evidence-Consuming to Evidence-Generating: What Ayurveda's Clinical Infrastructure Actually Needs

Having named the problem and examined the interests that sustain it, I want to turn toward what is actually possible — not in a spirit of optimism that papers over difficulty, but in a spirit of honest construction.

Because something can be built. Not quickly, not without friction, not without the kinds of institutional and cultural shifts that Articles One and Two have described. But built.

The question is what, exactly, needs to be built — and in what order.

The Last Mile Is Not the First Step

Let me begin by being direct about something. An Ayurvedic equivalent of MDCalc — a validated clinical decision support tool that a practitioner can trust in real time with a real patient — is probably fifteen to twenty years away from being meaningfully possible. Not because the technology is unavailable. The technology exists today. But because the upstream conditions for that tool do not yet exist, and building them takes time even under the best circumstances.

This is not defeatism. Fifteen years is not forever. Fifteen years ago, smartphone-based clinical tools of the sophistication available today would have seemed remarkable. The question is whether the Ayurvedic field will spend the next fifteen years doing the upstream work, or whether it will spend them continuing to produce confirmatory studies and digitizing more manuscripts.

The choice is ours. But it has to be made consciously, institutionally, and soon.

Step One: Operationalizing the Core Constructs

The first and most foundational task is one that requires no technology at all. It requires scholars, clinicians, and researchers to sit together — across lineages, across institutional affiliations, across regional traditions — and agree on operational definitions of Ayurveda's core clinical constructs.

What is *Mandagni*, expressed in terms that two different practitioners can apply to the same patient and reach the same conclusion? What are the observable, recordable clinical signs of *Vata Vriddhi*? What constitutes a measurable improvement in *Sroto-dushti* of the *Raktavaha Srotas*?

These questions sound basic. They are not easy. They will generate genuine disagreement. Lineage loyalties will be invoked. Textual authorities will be cited against each other. The process will be slow and occasionally acrimonious.

It is also completely unavoidable. No clinical data pool can form around constructs that mean different things to different practitioners. No outcome study is interpretable if the diagnosis it is built on has no shared definition. This is the foundation, and it has to be laid before anything else can be built on top of it.

Some work toward this end already exists — the Prakriti assessment tools developed by CSIR-IGIB, some of the diagnostic standardization work done at Gujarat Ayurved University, elements of the AYUSH research guidelines. But it is fragmented, contested, and has not yet produced the kind of broadly accepted operational consensus that would allow multi-site clinical data collection to begin in earnest.

A national consensus process — genuinely inclusive, rigorously facilitated, and committed to prioritizing clinical utility over lineage loyalty — is needed. It does not require government funding to begin. It requires the will of a critical mass of respected clinicians to convene and do the work.

Step Two: Building the Outcome Database

Once core constructs are operationalized, the next task is to begin systematically recording clinical outcomes against them.

This does not require a massive, centralized, government-funded infrastructure to start. It requires a network of willing institutions — perhaps ten to fifteen Ayurvedic hospitals with significant patient volumes — to adopt a common minimum dataset: standardized Ayurvedic diagnosis, intervention protocol, and measurable outcomes at defined follow-up intervals.

The technology for this is entirely available. A well-designed electronic health record system built around Ayurvedic diagnostic categories, with outcome tracking built in, could be developed and deployed within two years at modest cost. Several attempts have been made. None have achieved the critical mass needed to make the resulting data meaningful.

The obstacle, again, is not technological. It is the institutional willingness to share data — to allow a patient's outcome in one institution to be compared, anonymously and in aggregate, with outcomes in another. That willingness requires trust between institutions that are accustomed to competing rather than collaborating. It requires senior clinicians to accept that their outcomes will be part of a dataset in which they are not guaranteed to look good.

This is where leadership matters. A small number of respected senior figures in Ayurvedic medicine choosing to open their institutional data — and publicly calling on others to do the same — could shift the culture faster than any policy mandate.

Step Three: Protocol-Based Research

With operational constructs and an outcomes database in place, the third step becomes possible: protocol-based clinical research that generates the kind of evidence from which validated decision rules can eventually be derived.

This is where the research culture problem identified in Article Two must be directly addressed. Protocol-based research requires pre-registration of studies, commitment to publishing negative results, and peer review processes that reward rigor over confirmation. It requires researchers to design studies that could genuinely disprove their hypotheses — not merely studies that are structured to confirm them.

This is a cultural shift as much as a methodological one. It requires the Ayurvedic research community to collectively decide that the tradition's credibility is better served by honest negative findings than by a literature full of studies that prove what we already believed.

It also requires research mentorship that models this standard. This is where the role of PhD guides, PG supervisors, and journal editors becomes critical. The standards enforced at the point of research training and publication are the standards that will shape the next generation of Ayurvedic researchers. There is no neutral position here. Every guide who accepts a poorly designed confirmatory study and every editor who publishes it is making a choice — and that choice has consequences for the field's epistemic health.

Step Four: The Tool Layer

Only after Steps One through Three have been substantially advanced does the tool layer become meaningful.

But it is worth describing what that tool layer could eventually look like — not as a fantasy, but as a design target that helps orient the upstream work.

A mature Ayurvedic clinical decision support ecosystem would include, at minimum: a validated Prakriti and Vikruti assessment tool with known inter-rater reliability; evidence-linked herb-drug and herb-herb interaction databases; protocol-based management guidelines for high-prevalence conditions — diabetes, hypertension, inflammatory arthritis, irritable bowel, anxiety — expressed in Ayurvedic diagnostic terms and supported by outcome data; a case-based learning platform drawing on real, anonymized clinical cases from the outcome database; and a research integration layer that updates guidance as new evidence accumulates.

Some of these elements are closer to reality than others. The herb-drug interaction work is perhaps most advanced, with serious scholarship already available that could be aggregated and made clinically accessible. The protocol-based management guidelines are furthest away, because they depend most heavily on the upstream outcome data that doesn't yet exist.

A realistic near-term goal — achievable within five years with genuine institutional commitment — would be the first two elements: a validated assessment tool and a reliable interaction database. Those alone would represent a genuine advance on what currently exists and would begin to establish the habit of evidence-based clinical decision support in Ayurvedic practice.

The Role of the Next Generation

Everything described above requires the active participation of younger Ayurvedic practitioners and researchers — not as implementers of a programme designed by their seniors, but as co-architects of a new research culture.

This is particularly important because the cultural shifts required are, in some ways, easier for a generation that has not yet built its authority on the old system. A senior Vaidya who opens their outcomes to scrutiny risks something real. A young researcher who builds their career on rigorous, protocol-based, openly published research risks something too — but has more to gain, because the alternative path to reputation that the old system offered is becoming less viable as global scrutiny of Ayurvedic claims increases.

The young Vaidyas who will build Ayurveda's clinical infrastructure are already in our colleges. Some of them are reading this document. What they need is not inspiration — they are not short of that. What they need is institutional cover, mentorship that models the standards described here, and the knowledge that there is a professional community that will support rather than penalize the kind of honest, rigorous, protocol-based practice they want to pursue.

Building that community is, perhaps, the most urgent task of all.

What We Are Actually Asking For

Not revolution. Not the abandonment of classical knowledge. Not the reduction of Ayurveda to a set of algorithms.

We are asking for the discipline to build the infrastructure that a tradition of this depth and this reach owes to its patients and its practitioners. We are asking for the honesty to acknowledge that clinical authority without verifiable outcomes is not a mark of the tradition's sophistication — it is a gap in its integrity.

The app will come. But only if we do the work that makes it mean something.

What We Are Calling For

A direct statement for practitioners, institutions, regulators, and the next generation.

We are calling for a national consensus process to operationalize Ayurveda's core clinical constructs — across lineages, across institutions, across regional traditions.

We are calling for a network of willing Ayurvedic hospitals to adopt a common minimum clinical dataset and begin building the outcome database that Ayurvedic evidence-based practice requires.

We are calling for research mentors and PhD guides to hold their scholars to the standard of protocol-based, pre-registered, honestly reported research — and to model that standard in their own work.

We are calling for journal editors, including those of us who hold those positions, to publish negative results and to reject studies that are structured to confirm rather than test.

We are calling for senior clinicians who have built large clinical practices to open their outcome data — anonymized, aggregated, and shared — as a contribution to the field rather than as a threat to their reputation.

We are calling for AYUSH policy to redirect a meaningful proportion of research funding from confirmatory studies and manuscript digitization toward clinical outcome infrastructure.

We are calling for younger Vaidyas to build careers on rigor, to publish what is true rather than what is comfortable, and to understand that the tradition's long-term survival depends on their willingness to do so.

And we are calling for all of us — scholars, clinicians, educators, regulators — to accept that the most dangerous thing we can do for Ayurveda is to protect it from the scrutiny that every tradition serving living patients must ultimately face.

The app that doesn't exist is a symptom. The infrastructure that doesn't exist is the problem. And the will to build it — honestly, collaboratively, without regard for whose authority it inconveniences — is the only solution.

Disclaimers

Developmental Intent This document has been written with the sole intention of contributing to the growth, reform, and global credibility of Ayurveda as a living clinical tradition. Every critique offered here is offered from within the tradition, by a practitioner who has spent 35 years committed to its integrity and future. The intention is not to diminish Ayurveda but to hold it to the standard its depth deserves.

Personal Opinion The views expressed in this document are the personal and professional opinions of the author. They do not represent the official positions of any institution, regulatory body, university, or organization with which the author is affiliated. Readers are encouraged to engage critically, to disagree where they disagree, and to contribute their own thinking to the questions raised here.

AI Collaboration Disclosure This document was developed through a collaborative process between the author and Claude, an AI assistant developed by Anthropic. The inquiry, the intellectual framework, the positions taken, and the clinical and institutional judgments expressed are the author's own. The AI served as a thinking partner, a structural co-architect, and a drafting collaborator. This collaboration is disclosed in the spirit of transparency that this document itself advocates.

About the Author

Dr. Aakash Kembhavi is a senior Ayurvedic physician-educator with over 35 years of clinical, academic, and administrative experience. He is Director of Astanga Wellness Pvt. Ltd., Hubli — a publishing and wellness education enterprise. He has previously served as visiting professor at Thames Valley University, London and Europe Ayurveda Academy, France.

His writing on Ayurvedic reform, research culture, and epistemology is available at his **Dr Kembhavi's Ayurveda Unfiltered Blog** at **https://drkembhavi-s.github.io/**