Clinician AI Assist

A clinical decision-support tool that shows its work — built for licensed clinicians, designed around FDA non-device CDS expectations.

The problem

A clinician at the bedside — especially in home health, hospice, or a busy clinic — sees a set of findings and has minutes to decide what matters. Most AI tools answer with a confident-sounding conclusion and no way to check it. That is exactly what a clinician cannot safely use, and exactly what regulators do not allow.

Clinician AI Assist takes the opposite approach. It never tells the clinician what to do. It lays out what the findings could mean, why, what argues against it, what is missing, and where every claim came from — so the clinician can agree or disagree in seconds.

What it does for a clinician

Organizes the picture Turns a free-text patient story into a clean list of findings, vital signs, and derived values — with the arithmetic shown, not hidden.	**Builds a differential** Presents multiple candidate explanations, ranked, each tied to the specific patient findings behind it.
Shows the counter-case Runs a separate scan for evidence that argues against each pathway — not just the evidence that supports it.	**Names what's missing** Flags the exact piece of information that would settle the question, instead of quietly guessing.
Cites real sources Checks live medical literature during every analysis and names the exact publication, organization, and section used.	**Explores what could follow** Conditional scenarios — "if this pathway is operative" — covering complications, monitoring needs, and possible interventions.

Every pathway carries a standard **Basis for this consideration** panel: patient findings used, what supports it, what argues against it, what's missing, the exact source, how the patient's data matched that source's criteria, and the limitations. One consistent review unit, every time.

Designed for the regulatory line, not around it

The FDA draws a clear boundary between software that *supports* a clinician and software that *makes the call*. This system was deliberately built on the supporting side of that line.

No urgency verdicts The software never declares a condition time-critical. It surfaces abnormal findings for clinician review; criticality stays a human determination.	**No orders or dosing** It presents options and rationale. It does not instruct, order, triage, or dispose.
No black-box scores Any number shown comes with the inputs that produced it. Unexplained confidence figures were removed on purpose.	**Patient-facing text is gated** Plain-language summaries are drafts labeled for clinician review and cannot be released without explicit approval.

How the quality control actually works

Four independent layers run on every single case, in order:

1. Pattern engine

Recognized clinical patterns are matched and scored in plain code — fixed rules, same input, same result. This gives the reasoning a factual scaffold instead of a blank page.

2. Live cross-check

Current medical literature is searched for this specific case. Results are screened for actual relevance; anything off-topic is discarded rather than padded into the report.

3. Pre-synthesis validator

Before anything is written, an adversarial review re-derives the conclusions from scratch, recomputes the arithmetic, re-runs the elimination logic, and reports the strongest surviving objection.

4. Release gate

A final automated check reads the finished report for prohibited language — urgency verdicts, unexplained scores, predictive phrasing, ungated patient text — and blocks release if any appear.

It gets sharper with use

Patterns, chains, and relationships learned from previously cross-checked analyses are stored and reused as starting hypotheses on future cases — along with the sources they originally rested on. Crucially, a reused pattern never gets a free pass: it is re-validated against fresh literature every time it appears, and dropped if current evidence contradicts it. Recall is a hint; evidence is still required.

Where it fits

Home health & hospice
Complex, multi-morbid patients seen alone in the field, where a second set of eyes is not down the hall.

Clinics & hospital floors
Fast structuring of a messy presentation before a handoff, a call, or an escalation conversation.

Teaching & review
The reasoning is fully visible, which makes it usable for case review and clinician training.

Documentation support
Every analysis exports to a clean, printable report with sources and limitations attached.

What this project demonstrates

For hiring teams, in plain terms:

- **Domain fluency.** Clinical reasoning modeled correctly enough that practicing clinicians can follow and challenge it.
- **Regulatory literacy.** Product decisions driven by a real reading of FDA clinical-decision-support criteria, not a disclaimer bolted on at the end.
- **Judgment about AI.** Deterministic code where determinism belongs — arithmetic, rules, release checks — and generative reasoning only where the problem genuinely requires it.
- **Safety engineering.** Independent verification layers that can actually block the product's own output.
- **End-to-end delivery.** Interface, reasoning pipeline, live literature retrieval, learning store, and export — designed, built, and shipped as one working product.

Intended use. Clinician AI Assist is clinical decision-support software for licensed healthcare professionals. Its output is reference information to support independent clinical review — it is not a diagnosis, not a treatment order, and not a substitute for clinician judgment. It does not determine urgency, direct care, or predict what an individual patient will do. This software has not been cleared or approved by the FDA.