

LARRY P. BOLSTER II, LCSW

CLINICAL AI SAFETY WORK | Clinician-Led Evaluation, Privacy & Human Authority

Bangor, Maine | 207-944-9055 | BOLSTELP0@GMAIL.COM | Independent project, 2026

I am an LCSW with 20+ years in behavioral health, including emergency psychiatric evaluation, critical-incident review, and clinical supervision. In 2026 I began an independent clinical AI safety project to translate frontline clinical judgment into product constraints, evaluation, privacy boundaries, source discipline, and explicit human authority. The goal is useful clinician support without silently turning AI capability into clinical decision-making authority.

SYSTEM VIEW

CLINICAL JUDGMENT Risk, context, authority	PRIVACY BOUNDARY Local, isolated, gated	SOURCE GROUNDING Approved, traceable sources	BOUNDED AI Draft/support, not decide	SAFETY TESTING Synthetic, adversarial gates	CLINICIAN REVIEW Decide, approve, export
---	--	---	---	--	---

SELECTED WORK

01 CLINICAL RISK & EVALUATION - Red-teamed crisis timing, safety-clearance, PHI, diagnosis, and treatment scenarios; documented where instruction-only safeguards failed. - Built a 9-category risk taxonomy and 34-case gold-standard suite with expected verdicts and safe controls; refusal testing surfaced a guard miss that was preserved and routed to targeted repair.	02 CLINICAL ASSISTANT - Designed a ground-up local clinician workspace around client/session isolation, deny-by-default external egress, minimal retention, and explicit save/export. - Human clinician remains final authority; diagnosis, treatment, crisis, disclosure, communication, and chart-submission decisions are not autonomous.
03 SOURCE & REFERENCE SAFETY - Built source-ingestion controls so clinician support can be traced to approved references: page-preserving extraction, source/chunk provenance, package verification, and rights-aware handling. - Established approved-source discipline using current federal, state, professional, peer-reviewed, and lawfully held licensed references.	04 EVIDENCE & RELEASE GATES - Separate source, mechanical, behavioral, and human evidence; reviewers reconstruct findings from primary evidence rather than inherit prior verdicts. - Synthetic privacy, clinical-fidelity, usefulness, and recovery gates come before any limited-PHI pilot. Real PHI and clinical deployment remain gated.

THE CLINICAL TO AI BRIDGE 20+ years in behavioral health, including three years of emergency psychiatric evaluation, crisis assessment, critical-incident review, chart review, and supervision of interns and clinical staff across three organizations.	Clinical judgment becomes product logic: identify harm, weigh incomplete evidence, define escalation and ownership, and test whether a system is both safe and genuinely useful. The through-line is simple: capability is not authority. Keep limits visible, preserve clinician responsibility, and require evidence before widening what the system can do.
---	--

SYSTEMS FOUNDATION

Local-first design • Client isolation • Approved-source retrieval • Evidence-backed state • Human approval • Fail-closed gates • Failure recovery

TECHNICAL ENVIRONMENT

Local LLMs / Ollama • RAG / ChromaDB • Open WebUI • n8n • Docker • Flask • AI-assisted Python & PowerShell test harnesses • Tailscale

Selected independent clinical AI safety portfolio | Companion to Headway Clinical Product Lead application | Development/test work only