

LARRY P. BOLSTER II, LCSW

CLINICAL AI PRODUCT & SAFETY | Clinician-Centered Design, Evaluation & Human Authority

Bangor, Maine | 207-944-9055 | bolstelp0@gmail.com | larry.bolstercybernetics.com | Independent work, 2026

Licensed Clinical Social Worker with 20+ years in behavioral health, including emergency psychiatric evaluation, critical-incident review, chart review, and clinical supervision. My independent clinical AI work translates frontline clinical judgment into useful product design, explicit boundaries, testable evaluation, and human authority. **The through-line: build tools that help clinicians without quietly becoming the clinician.**

SYSTEM VIEW

CLINICAL JUDGMENT	CLINICIAN NEED	EVIDENCE	BOUNDED PRODUCT	EVALUATION	HUMAN AUTHORITY
Risk, context, authority	Workflow, usefulness, burden	Curated, scoped, traceable	Assist, not decide	Structured, synthetic, recovery	Decide, approve, act

SELECTED WORK

01 THERAPIST SESSION COMPASS - Zero-PHI-by-design first-release concept for in-session clinician guidance across situations, interventions, questions, exercises, psychoeducation, and cautions. 24 underlying evidence sources organized into 17 retrieval-oriented modules with explicit scope rules, supported points, boundaries, and product mappings.	02 CLINICAL SAFETY & EVALUATION - Nine-category clinical risk taxonomy and self-authored 34-case evaluation suite testing both missed risk and unnecessary blocking. - Structured testing across crisis timing, safety-clearance language, privacy/PHI, diagnosis, and treatment boundaries; failed cases remain on record for targeted repair.
03 PRIVACY, SOURCES & CLINICAL ASSISTANT - Local clinician workspace design centered on client/session isolation, minimal retention, explicit save/export, and clinician-final authority. - Traceable source handling with page-preserving extraction, source/chunk provenance, local indexing, and rights-aware handling.	04 AI LAB METHODOLOGY & RELEASE GATES - Local multi-model development/evaluation environment with bounded actions, per-action approval, provenance, failure preservation, stop conditions, and recovery. - Scope widens only after evidence from synthetic privacy/clinical-fidelity testing, usefulness testing, and recovery checks. It is not a clinical product.

THE CLINICAL TO PRODUCT 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 what can go wrong, separate acceptable uncertainty from missing evidence, make escalation and ownership explicit, and test whether a tool is both safe and useful.
---	---

RESEARCH IN PROGRESS

Toward a Breathing AI Lab | Manuscript in preparation for IEEE submission. Examines evidence-governed supervised autonomy in a local multi-model AI environment, including explicit human authority, bounded action, release gates, failure preservation, and recovery. Not yet published or accepted.

Independent clinical AI product & safety work | Synthetic cases and non-client information only | Real-client PHI processing and clinical deployment remain disabled