MAZZE LECZZARE

Durham, North Carolina · mazze@mazzeleczzare.com · mazzeleczzare.com · github.com/mazze93

Applied security research · empirical study design · security infrastructure for civil society

RESEARCH

Counterfeit Banknote Authentication Program — UC Davis, Center for Mind and Brain (Luck Lab), 2018–2020

- Selected by Dr. Steven J. Luck to conduct experimental data collection for a banknote-authentication program supported by the Bank of Canada, the Bank of England, the Reserve Bank of Australia, and the US Federal Reserve.
- Ran the full counterfeit-note paradigm across three modalities — eye-tracking, EEG/ERP, and behavioral — with consented community volunteers (non-student sample).
- Trained visiting researcher Daniel Dodgson on the Luck Lab standard for EEG acquisition and ERP analysis in ERPLAB. Dodgson is first author of *Banknote authenticity is signalled by rapid neural responses* (Dodgson & Raymond, **Scientific Reports** 12:2076, 2022), a program output using the same paradigm. *Acknowledged contributor; not an author.*
- Held custody and nightly physical security of forensically recovered counterfeit currency.

> *Applied security research conducted with cognitive-neuroscience methods: measuring, against real adversarial artifacts, where human detection of a security feature actually succeeds and fails. The direct methodological ancestor of the work proposed here.*

Laboratory Operations Manager (Junior Specialist) — Luck Lab, UC Davis

Participant pipelines, EEG rig maintenance, data-integrity standards, and training incoming researchers to meet them.

Multimodal affect-inference pipeline — UC Davis, Luck Lab

Pipeline combining EEG/ERP, passive eye tracking, representational similarity analysis, and IAPS-grounded affective coding; models continuous three-dimensional affect consistent with Barrett's constructionist framework. Core contribution: a validity argument against behavioral affect research lacking multimodal grounding.

PUBLICATIONS AND WRITING

PAPERS

Intentional Fragility: Bayesian Context Decay and Semantic Rot in Local-First LLM Orchestration (2026). Independent. Formalizes context decay as a Bayesian process with self-correcting anchor protection and a fault-injection calibration protocol.

Epistemic Fault Injection: measuring integrity under adversarial context (v0.1, in preparation). Protocol and five-level scoring rubric measuring whether a model under adversarial pressure admits a limitation or reclassifies it to defend a persona. Multi-model transcripts.

ESSAYS

What We Don't Know Yet: iOS Lockdown Mode Research — security research note that retracts its own prior confidence structure and rebuilds it with assumptions labeled and uncertainties stated as the deliverable.

The Architecture of Forgetting — attribution, memory, and accountable forgetting when the collaborator has no persistence between sessions.

The Lighthouse Earns Its Floor — earned trust in human–AI handoffs; nothing self-promotes to authoritative before verifying the floor it stands on.

The Concept Is Not the State — constructionist critique of emotion-representation claims in interpretability work.

INDUSTRY AND AUDIENCE WRITING

Networking at GDC 2023: Tips for Artists, Game Developers, and Industry Leaders — 80 Level, March 2023.

Three feature articles, LinkedIn Pulse (2023): digital asset management systems; indie publisher relations; industry event coverage.

Clear as Mud — podcast, founder, host and producer. 26+ long-form interview episodes with studio directors, technical designers, and founders. Spotify and Apple Podcasts.

SYSTEMS BUILT

Stratum — epistemic decision ledger `LIVE v0.4.0`

TypeScript core, Python semantics oracle, Cloudflare Workers and Durable Objects, MIT. Append-only event log for human and AI work: evidence (file hashes, test exits, signed approvals) is separated from claim (model prose), and status is a fold over the log rather than a stored field — replay at any epoch is a proof, not a feature. Sixteen-invariant executable contract; single-writer Durable Object gate; fail-closed projection. The system records its own genesis decisions in itself.

STELE — directive-integrity compiler `LIVE`

React 19, TypeScript, Vite. Compiles LLM instruction sets into per-project posture with integrity telemetry, a thirteen-tripwire registry, and a four-state integrity machine. Detects prompt manipulation, semantic drift, and constraint violation; full audit trail and narrative export.

Field Notes — instrument gallery `LIVE`

Self-contained browser instruments, one idea each: a kintsugi circuit breaker that renders its own failure history as visible seams; a Physarum-inspired routing solver that displays every candidate edge it discarded.

`_cert_audit.sh` — macOS certificate and keychain audit system (2026)

Zsh, ~859 lines. Full keychain audit across login, system, and custom tiers. Identifier-level SOGI/GDPR masking. SARIF, JSON, and Markdown output. Launchd deployment with absolute paths written at install time. Runs unattended against a production PKI stack.

802.1x civil-society network infrastructure (2026)

Certificate-based authentication stack: FreeRADIUS, step-ca PKI, Mosyle MDM, UniFi, dynamic VLAN segmentation, full operational documentation. *This is the proposed research testbed.*

MacProbe — macOS security probe (2026)

Seven modules: firmware, hardware, boot, storage, network, persistence, process. Dual JSON/SARIF and Markdown output. Secure Enclave attestation concept; layered backend (Core ML, Ollama, Claude API).

kintsugi — fragility-aware tool orchestrator (2026, Python)

Tool routing for an LLM research harness. Circuit-breaker state machine with an eject-reason taxonomy retained as research payload; ranks tools by value-density (goal advanced per token) rather than reliability. Verified end-to-end.

SECURE PRIDE — FOUNDER AND EXECUTIVE DIRECTOR, 2026–PRESENT

Privacy-first cybersecurity nonprofit building infrastructure, tooling, and training for LGBTQ+ organizations and other at-risk civil-society groups. Zero-trust architecture, WCAG 2.1 AA throughout, no analytics or telemetry brokerage, three-tier data classification with collection minimized at the most sensitive tier. Coordinated-disclosure policy with CVE workflow (72h/14d/90d), MITRE CVE preparation, CycloneDX VEX and CSAF output.
securepride.org · github.com/Secure-Pride · huggingface.co/SecurePride

- **Cyber-Defense Audit Suite** — CLI-first audit framework: network exposure scanning, DNS hardening verification (DNSSEC/SPF/DKIM/DMARC), TLS validation, secrets scanning. Passive by default, screen-reader mode. Python, TypeScript, MIT.
- **Stonewall Standard DPA** — sixteen-section data-processing agreement framework for nonprofits handling community-sensitive data. Parallel legal and plain-English text, AI/LLM vendor onboarding guidance, incident-response templates. CC BY 4.0.
- **AI-Native Threat Toolkit** PRE-RELEASE — prompt-injection detection patterns, LLM output-integrity checking, AI-assisted phishing classification, threat-brief generation from open-source feeds.

OPEN SOURCE

js-yaml — GitHub security advisory credit.

ERPLAB Toolbox — contributor. Open-source MATLAB toolbox for event-related potential analysis, maintained by the Luck Lab and used internationally.

TEACHING AND EDITORIAL

Butte College EOPS (2015–2018) — tutored re-entry students; taught math, science, English, and history to students with learning and developmental disabilities.

Bloom Magazine, Butte College — founding editor-in-chief. Established mission, voice, and editorial standard; built and led a student contributor team; wrote the grants that funded it. Subject: socioeconomic inequity and the communities usually written about rather than by.

Teaching Assistant, Psychology of Emotions (Dr. Eliza Bliss-Moreau) — UC Davis.

Mudstack — Content and Community (2022–2024). Sole content function: 50+ technical posts, 7 webinars produced and hosted, documentation infrastructure, and the *Clear as Mud* podcast.

EDUCATION

University of California, Davis — B.S., Quantitative Psychology and Cognitive Neuroscience, 2020. Regents Scholar. Luck Lab (EEG/ERP) under Dr. Steven J. Luck.

TECHNICAL SKILLS

Python · Zsh/Bash · TypeScript / React 19 · Swift 6 · MATLAB · EEG/ERP acquisition and analysis (EEGLAB, ERPLAB) · eye-tracking · representational similarity analysis · FreeRADIUS · step-ca · Mosyle MDM · UniFi · Wireshark · Cloudflare Workers and Durable Objects · Astro · Docker (ARM-native) · SARIF · CycloneDX · CSAF · OSV · Core ML · Claude API