

Frans Moore

Agent systems • Verification and evaluation • Applied AI • Global health

Atlanta, GA • hi@fransmoore.com • fransmoore.com • linkedin.com/in/frans-moore • github.com/FransDevelopment

I build AI systems that institutions can check. In early 2025 I turned back to AI itself. Working with AI coding agents, since April 2026 inside a verification-gated pipeline I designed, I have built a control plane that runs a company with AI agents and treats every agent's account of its own work as a claim, a live Ebola forecasting system that scores its misses in public, and open agent-trust infrastructure. Both ends of the problem matter to me: the data that goes into AI systems, and the evidence that what comes out of them is true.

I came to that work from the input side. After an MSc in International Public Health, I spent the years that followed on data collection, field research and health programs in eastern DR Congo, Tanzania and South Sudan, working on the lack of high-quality, organized data in niche, underserved contexts. A case count or a clinic register is data, and when it is wrong, the people it undercounts are usually the ones with the least room for error. Models are only as good as the data they learn from and the evaluations they are checked against, and data like this, from the places that need it most, is among the scarcest there is. That work gave me a multi-sectoral and humanitarian understanding that is rare in AI, and it is why I build AI for society's direct benefit.

The question behind all of it started at the University of Georgia. From 2015 to 2019 I took courses through UGA's Institute for Artificial Intelligence, including AI, discrete mathematics, cognitive psychology, sensation and perception, and mathematical modeling. I came away asking what the most consequential challenges are that AI, and thinking machines more broadly, could help us solve, and that question is part of why I went into public health.

SELECTED WORK

Earth: control plane for an AI-agent organization Python, Rust, TypeScript

Apr 2026 – present

Built on one premise: an agent's report of its own work is a claim, not evidence

- Designed and operate the system through which a fleet of AI agents runs research, engineering and operations for my company every day, coordinating through shared memory and structured handoffs rather than through each other.
- **Agents' claims of "done" are checked, not trusted.** In May 2026 a rule-based check on my automated dispatcher refused 98 of 143 AI-agent claims of success (69%), most after my own instructions dropped the required documents, a gap no agent flagged.
- **Safeguards are measured before they are enforced.** Each new gate runs in record-only mode first, so its false-block cost is known before it can stop work. One promising gate stayed off because the measurement showed it would have refused most completions.
- **Safety constraints live in the tool layer, not the prompt**, so no chain of agent reasoning can argue past them. The production runtime is fail-closed: if anything in it has been altered, it will not start.
- Audits itself: my own Critical audit found the system's "completed" status was hiding failures, and I fixed it at the root. An internal red-team agent found a credential exposure; it was fixed and closed.

LOVS: outbreak evidence and forecasting system Bundibugyo Ebola, DRC and Uganda

May 2026 – present

Public data and methods: github.com/ArcedeDev/bdbv-2026-lovs

- Built the first release solo in about a week with AI coding agents, then kept it current through **127 official situation reports** from multiple publishers, reconciling conflicting sources into one cited picture.
- Every published figure traces to its source and the date the data describes. An evidence service for AI clients refuses to serve any number without a citation, so a model answering questions about the outbreak cannot state an unsourced figure.
- Pre-registered **87 public, dated, thresholded forecasts** before outcomes were known (one block archived on Zenodo); 54 have resolved and 31 resolve on 1 Oct 2026. Scores are published against a base-rate baseline, misses beside hits, and I retired a model input when its pre-registered ranking came out inverted.
- Engaged UNICEF's Eastern and Southern Africa regional leadership from May to August 2026; the Deputy Regional Director assigned regional staff to review the work in detail. Featured as an **Apolitical Government AI Navigator** case study and covered by Scrolla.Africa.

Verified engineering with AI coding agents

2025 – present

- Earth, LOVS and the agent-trust tools were built with AI coding agents working inside a verification-gated pipeline I designed: requirements are interrogated, assumptions checked against ground truth, and every change independently reviewed and red-teamed before it ships. The review stage has blocked merges and caught Critical bugs, including a data sync that silently dropped every row.

Agent trust infrastructure open source, MIT

Mar 2026 – present

- Author of **Composable Trust Infrastructure for the Agent Internet** (Zenodo, CC BY 4.0), on verifiable agent identity, capability discovery and authorization. Core claim: every consequential action must prove who authorized it, what constraints applied, and what happened.
- **Open Agent Trust Registry**: 13 forks, 9 outside contributors, 11 registered issuers. **AIR SDK** (@arcede/air-sdk), an agent web-perception SDK with its own MCP server: 30 releases, 5,200+ npm downloads.

- 83 pre-registered tests and 87 logged corrections in a locked register, with **power analysis before any null**: underpowered results are reported with their minimum detectable effect, never as absence. Central finding: the failures traced to the observation model, not the biology.

EXPERIENCE

Founder Arcede, Atlanta, GA

Mar 2025 – present

AI, evidence and coordination systems for institutions making consequential decisions. A one-person company run with the agent fleet above; engagements with global health and biodiversity philanthropy, including an AI policy options map for a U.S. biodiversity foundation.

Project, Finance and Research Manager Lake Tanganyika Floating Health Clinic / WAVE, DR Congo

Jan 2022 – present

- Coordinated a three-year mixed-methods human health and biodiversity survey across South Kivu and Tanganyika provinces in a conflict-affected setting: household surveys, focus groups, entomological and fish sampling. The 2022 South Kivu wave alone covered 160 households on a 116-item instrument, 111 focus-group participants and 561 fish.
- Raised about \$500,000 in grants and operating funds with the founder; managed organization and research finances; digitized provincial health-system administrative registers; published the program's fisheries data to GBIF as originator.

Program Advisor (volunteer) Doctors Actively Serving in Africa, Juba, South Sudan

Sep 2022 – present

Program evaluation and board governance; contributed to the 2022–2026 strategic plan and the MOU with South Sudan's Ministry of Health.

Software Engineer Freelance Development

Jan 2021 – Feb 2024

Freelance software engineering, mostly video game development, with web development alongside. Received an Epic MegaGrants award.

External Researcher NIMR Amani Centre / IVCC / WHO, Muheza, Tanzania

Aug 2017 – May 2019

Designed and ran an independent mixed-methods study of staff through the centre's transition to Good Laboratory Practice certification (30 of 36 staff; Swahili instrument, translated and back-translated). Assisted malaria bed-net quality surveys. Two Best Poster awards. Also research assistant on NASA-funded Project FUSION (multiteam systems, LINC Lab) and a workaholism study (Healthy Work Lab), UGA.

EDUCATION

MSc International Public Health: Planning and Management Liverpool School of Tropical Medicine

2019 – 2021

Dissertation: a rapid evidence assessment of how immunization programs in low- and middle-income countries adapt during infectious disease crises (5,806 records screened, 19 studies included). Coursework in statistical methods for epidemiological and clinical research.

BSc Psychology, minor in Cognitive Science University of Georgia

2015 – 2019

Honors thesis on job satisfaction at an East African medical research centre. Coursework included artificial intelligence, taught as CSCI 2670, Introduction to Theory of Computing; cognitive science ran through UGA's Institute for Artificial Intelligence.

PUBLICATIONS AND WRITING

- Moore, F. (2026). Composable Trust Infrastructure for the Agent Internet. Zenodo. doi:10.5281/zenodo.19263547
- Mushagalusa, C. D., Lehman, A. G., Moore, F., et al. (2024). Critical biodiversity, fisheries status and need for inshore fish communities conservation in Lake Tanganyika. Journal of Great Lakes Research, 50(3), 102296.
- Moore, F. (2026). Hugging Face "Hack": Six Questions Nobody Could Answer. Apolitical. When agents can alter the logs they touch, a system's own log is a claim, not evidence.
- Moore, F. (2026). AI Oversight and Verification Saturation. Apolitical. On why oversight is the first institutional capacity AI saturates, and why it saturates silently.
- Arcede (2026). BDBV 2026 pre-registered forecast block (41 forecasts). Zenodo. doi:10.5281/zenodo.21233092
- Talks: Global Bioeconomy Summit panel, Nairobi (2024); Species in Ancient Lakes, Kigoma (2022). Dataset: Fisheries of Lake Tanganyika, GBIF (2022).

SKILLS

AI engineering Agentic coding, MCP servers, multi-agent orchestration, tool-layer guardrails, LLM output validation, evaluation and calibration harnesses.

Software Python, TypeScript, React and Next.js, Rust (Tauri), SQLite, GCP Cloud Run, Vercel.

Methods Forecasting and calibration (Brier, PIT, base rates), pre-registration, power analysis, systematic review, survey design, mixed methods, cross-language instrument translation.

Languages English (native), French (working).