

# Alex Somerville

## Software Engineer — Applied AI & Full-Stack

Toronto, ON · Remote · j.alex.somerville@gmail.com · github.com/mrlovelies · alexsomerville.dev

### SUMMARY

---

Software engineer with 15+ years building and shipping web, e-commerce, and game products — front-end architecture, Magento/Shopify, Vue, Firebase, Unity — now focused hands-on on applied AI. Recently built eval-first LLM tooling, retrieval/RAG systems on Cohere's Embed + Rerank, and a multi-machine local-LLM infrastructure. After a creative period (voice acting + self-directed AI engineering), looking to contribute as a builder on an AI-native product team.

### APPLIED AI — RECENT WORK

---

**Coherence Keeper** — retrieval + grounded contradiction detection (open source). Hybrid retrieval + Cohere Embed/Rerank cross-encoder + an LLM contradiction judge, built eval-first: a planted-contradiction harness (precision/recall, MRR/NDCG, calibration) gates the build in CI. A real-model run lifted contradiction F1 from 0.50 (baseline) to 1.00 at a 0% false-positive rate. *Python*. [github.com/mrlovelies/coherence-keeper](https://github.com/mrlovelies/coherence-keeper)

**Anton Scout** — LLM technique-discovery scout with a decoy-injection eval (open source). Extracts the transferable idea from candidate AI techniques and scores it against a system's real open problems; ships a decoy-injection eval that fails CI if hype leaks into the digest (1.0 decoy-catch / 0.82 recall over 27 candidates). Also packaged as an installable agent skill. *Python*. [github.com/mrlovelies/anton-scout](https://github.com/mrlovelies/anton-scout)

**Son of Anton** — local-LLM infrastructure (running in production). A three-machine Tailscale GPU mesh with a tiered local-LLM router (smallest model that can answer, escalating only when needed) running the inference and overnight automation behind the projects above. *Ollama, grammar-constrained structured output, Python*.

### EXPERIENCE

---

#### Developer — Relish Studios (contract) · Jul 2022 – May 2023

Built a web-based product-sales dashboard for the *Storylines* mobile game in Vue 3 + Firebase SDK.

#### Developer — Tangent Interactive Media · Jul 2021 – Jul 2022 · Toronto

Owned Firebase liveops for the *Storylines* mobile game; integrated with a Unity client.

#### Demac Media → Bounteous (acquired) · 2014 – 2021 · Toronto

Two tracks across the agency's e-commerce practice:

- **Front-End Architect & Practice Lead** — led a ~7-person front-end team; set architecture, standards, and mentoring across the practice; shipped large-scale e-commerce on Magento 1/2 & Shopify for Government of Alberta, Shake Shack, and a dozen-plus other brands.
- **Senior Solutions Engineer** — a 3-person "special ops" team that owned the hardest cross-project problems: built scalable shared tooling and troubleshooted critical issues across the client base.

#### Creative Technologist · Senior Front-End Developer — Cossette · 2012 – 2014 · Toronto

Delivered campaign technology at one of Canada's largest agencies — matched the right tools to each project and guided the front-end team.

### SKILLS

---

**Languages:** Python, JavaScript / TypeScript, PHP

**AI / LLM:** RAG & retrieval, reranking & embeddings (Cohere Embed/Rerank), LLM-as-judge evals, agentic tooling, prompt & context engineering, structured / grammar-constrained output, MCP, Ollama, local-LLM serving

**Web:** Vue, React / Next.js, front-end architecture, REST APIs

**E-commerce:** Magento 1/2, Shopify

**Games / realtime:** Unity, Firebase (liveops)

**Infra & tooling:** Tailscale mesh, SQLite, CI/CD (GitHub Actions), Git, multi-machine GPU orchestration

## **BACKGROUND**

---

Self-taught — 15+ years of on-the-job engineering, growing from front-end developer to architect and practice lead. Working voice actor (ACTRA) since 2019 — the creative period during which the AI work above was built.