Nicholas Bode

Toronto, Ontario, Canada · nickbode71@gmail.com · GitHub · LinkedIn

PROFILE

AI engineer focused on secure AI integrations, authorization, authenticated agent tools, evaluation workflows, and applied ML. Background in data-quality systems and original time-series research.

EXPERIENCE

Environics Analytics · AI Engineer | April 2026–present

- Implemented authorization enforcement for AI-connected search and tools, including hierarchy-aware access, audited administration, and access revocation.
- Designed and tested a remote agent-tool integration using OAuth2/PKCE, scoped access, and protected token storage; this work was a tested prototype, not a production launch.
- Contributed authenticated remote access and reporting SDK functionality to a team-owned agent-tool platform.
- Built human-review and evaluation workflows for model-assisted data matching and analytical changes, combining deterministic criteria, sample audits, and escalation.

Environics Analytics · Research Associate | February 2024–April 2026

- Built recurring analytics quality checks, anomaly investigation tools, and release reporting workflows using Python and Snowflake.
- Compared data-product releases against external benchmarks and quality criteria, documented tradeoffs, and supported release recommendations.
- Investigated anomalous mobility trends by tracing upstream methodology and validation changes.

EDUCATION & RESEARCH

MSc, Computer Science — Ontario Tech University. Thesis: Discovery of Trend Dependencies Over Time-Series (2023). Studied approximate and conditional relationships in noisy time-series data using symbolic regression, segmentation, and policy-gradient methods.

BSc, Computer Science, Mathematics minor — Ontario Tech University (2021).

Publication: Bode, N. & Sukhai, M. (2024), Automated Image Captioning for the Visually Impaired, Encyclopedia of Computer Graphics and Games, Springer. DOI.

SELECTED INDEPENDENT PROJECTS

AgentShield — LLM Agent Security Evaluation Framework — github.com/simplysoup/agent-shield

Built a paired prompt-injection evaluation framework for tool-using agents, with fixed attack replay, runtime tool-authorization checks, and structured audit traces.

Evaluated RAG + Agentic System — github.com/simplysoup/evaluated-rag-agentic

Synthetic-document retrieval benchmark comparing deterministic extraction with a tool-using agent; includes programmatic grounding checks, correction/abstention logic, and documented evaluation limitations.

TECHNICAL SKILLS

Languages & data: Python, SQL, Snowflake, dbt, Pandas, scikit-learn, FastAPI. **AI systems:** MCP, LangGraph, OAuth2/PKCE, agent-tool infrastructure, LLM evaluation. **Engineering:** Azure, Docker, CI/CD, APIs, orchestration.