

Taylor Clements, PhD

AI Engineer & Data Scientist

✉ taylor.clements@gmail.com

GitHub: drtclem

Portfolio: taylorclements.com

LinkedIn: /in/drtclements/

Houston, TX

Profile

AI engineer and data scientist with a PhD in Education Sciences and more than a decade leading cross-functional work across analytics, education and product. I build LLM agents, retrieval-augmented generation (RAG) pipelines and multi-agent systems in Python with LangChain and LangGraph, and I prove they work with golden-set evals, held-out test sets and red-team suites that gate releases. Recent work includes an AI call center agent that resolved 81% of eligible calls with zero false containments, and a multi-agent red-teaming platform that found vulnerabilities in a clinical AI assistant and verified the fix. I start with the right question and measure the answer.

Skills

Python | LangChain | LangGraph | LLM Agents | Multi-Agent Systems | RAG | Vector Databases (ChromaDB) | Embeddings (sentence-transformers) | Prompt Engineering | LLM Evals & Guardrails | AI Red Teaming | Anthropic & OpenAI APIs | Local Models (Ollama) | Langfuse | FastAPI | Flask | REST APIs | PostgreSQL | Redis | SQL | Docker | CI/CD (GitLab, GitHub Actions) | Git | TypeScript / React | Machine Learning | PyTorch | TensorFlow | Scikit-Learn | XGBoost | Transfer Learning | NLP | Pandas | NumPy | Statistical Analysis | Data Modeling | Anomaly Detection | Tableau | Databricks | AWS | Microsoft Azure | Agile/Scrum | Stakeholder Alignment | Research Methodology

Education

M.S. Data Science , Eastern University	Expected 2027
AI Engineering Program , Gauntlet AI	2026
<i>Intensive program; designed, built and shipped a new AI system most weeks</i>	
Data Science & Machine Learning Engineering Certification , 4Geeks Academy	2025
Ph.D. Education Sciences , University of Kentucky	2018
<i>Specialization: Educational Leadership</i>	
M.A. Education , Western Kentucky University	2012
B.S. Mathematics, Minor in Music , Murray State University	2011

Professional Experience

Enterprise Analytics Translator , Sysco	08/2025 – Present Houston, TX
- Manage analytics and data science initiatives across merchandising, pricing, supplier funding and operational reporting, focused on high-impact enterprise use cases. - Translate complex business problems into analytical requirements, models and decision frameworks with business leaders, product owners and technical teams. - Design data models, statistical analyses and machine learning approaches to detect anomalies and measure performance, and report results clearly to technical and non-technical audiences.	
AI Engineer (Part-Time) , Cignition, Inc	2026 – Present Remote
Built an automated post-session feedback pipeline (LangChain + DeepSeek V4 Pro via Lightning AI) that turns session transcripts into talk-time and engagement metrics plus structured JSON and Markdown reports, replacing a manual reporting process.	

- Architected a separate RAG-driven tutor coaching pipeline that grounds feedback in internal instructional guidance: chunked ~15 internal documents, embedded them with sentence-transformers (all-MiniLM-L6-v2) and indexed them in ChromaDB for semantic retrieval.
- Engineered prompt templates with structured output, guardrails and JSON-schema validation, and handled absent, partial and malformed transcripts for reliable production output.

Program Manager & Co-Director of Curriculum and Instruction, Cognition, Inc

Spring 2021 – 01/2026 | Remote

- Directed dozens of cross-functional projects across data analysis, research, operations and partnerships using Agile/Scrum, Jira and Trello, and led data collection for school site operations.

Director of Innovation and Design, Christchurch School

Fall 2019 – Spring 2021 | Christchurch, VA

- Directed MakerSpace operations and project-based learning in design thinking, coding and engineering; administered the Canvas LMS.

Instructor, Jefferson County Public Schools, Culver Academies, Bellarmine University, Johns Hopkins CTY, Murray State University

2011 – 2019 | KY / IN / Remote

- Taught Algebra I and II, Geometry and Calculus at Jefferson County Public Schools and Culver Academies, and graduate research methods, building a foundation in experimental design and statistics.

Projects

AI Call Center Agent (Gauntlet AI · agents, RAG, evals)

- Built a deployed agent that answers routine calls from company documents, books and moves appointments in a real scheduling database, and hands everything else to a person with an actionable ticket.
- Classified all 420 calls with two blind labelers from different model families (420/420 agreement) to find the 61% containment ceiling first; kept identity, bookings and handoffs in deterministic code.
- Measured on 360 non-development calls: 96% retrieval accuracy (top 3), 81% of eligible calls resolved, 0 of 139 false containments and 96/96 red-team attacks blocked; chose the stronger model after the cheaper one (0.4¢ vs 2.5¢ per call) produced 2 false containments.

Clinical Co-Pilot for OpenEMR (Gauntlet AI · LangGraph, RAG, healthcare)

- Built an AI assistant inside OpenEMR that gives overnight residents grounded context from FHIR chart data, uploaded lab PDFs and intake forms, and clinical guidelines.
- Designed a LangGraph supervisor routing to chart, document-extraction and guideline-retrieval workers, with per-field source citations, a verification step that strips ungrounded claims and an allergy hard stop in code.
- Gated every merge in CI on a 55-case golden set, 8 graph cases and 288 unit tests, all passing, with 100% on a five-metric rubric; per-user OAuth access and patient identifiers redacted from logs.

Adversarial AI Security Platform (Gauntlet AI · multi-agent, AI security)

- Built a four-agent platform (orchestrator, red team, judge, documentation) that black-box attacks a deployed clinical LLM, with a human-approval gate no model can bypass.
- Ran 31 unique attacks across 6 OWASP LLM risk areas and published 5 vulnerability reports with 0 false findings; validated the judge against 30 human-labeled cases with zero deterministic false negatives.
- Closed the loop: shipped a fix to the target, and the automatic regression replay dropped leaks from 6/6 to 0/6, all at \$0.00 API spend using local models.

Future Coach and Weekly Summary Guardrails (Gauntlet AI · knowledge graphs, LLM evals)

- Future Coach: an LLM proposes workouts but only a deterministic knowledge-graph check can allow them. Weekly Summary: code checks, a judge model and repair gate every AI-written note, with judge models scored against my own labels (github.com/drtclem/weekly-summary-guardrails).

Skin Lesion Classifier (PyTorch, transfer learning)

- Fine-tuned a ResNet50 on 25,000+ dermoscopy images across 8 diagnoses, addressing a 54:1 class imbalance with weighted sampling and loss; macro F1 improved from 0.44 to 0.67 and melanoma recall from 58% to 75%.