Tanner O'Rourke

Seattle, WA | (425) 786-6688 | tannero@live.com | tannerorourke.dev

ML Research Engineer

ML rarely intersects with engineers who've shipped production systems at-scale. Fewer still bring a background in UI design and cognitive psychology. Research Engineer with 5+ years of experience, formal ML training through UT Austin's MS in AI, and the cross-disciplinary range to work fluently across systems, human behavior, and model internals, in service of AI that genuinely works for people.

TECHNICAL SKILLS

Interpretability: Mechanistic Interpretability, Sparse Autoencoders (TopK), Linear/Logistic Probing, Activation Extraction, Manifold Geometry Analysis, Diagnostic Probing & Behavioral Evaluation

ML/AI: PyTorch, Hugging Face Transformers/Datasets, Transformer Fine-Tuning (LoRA/QLoRA), spaCy, A100 GPU Training, Agentic Workflows (Claude Code, hooks, structured memory)

Languages: Python, SQL, TypeScript/JavaScript, C++

Data & Infrastructure: Pandas, NumPy, xarray, Git, BigQuery, AWS S3

EXPERIENCE

Lead Engineer | DirecTV

October 2023 – August 2025 | *Promoted from Software Engineer in 3 months*

2024 Q3 Delivery Excellence award – top individual contributor across all customer-facing web teams.

- Functioned as de-facto technical lead for directv.com, drove feasibility and planning with PM's, design, and VP-level stakeholders, mentored junior engineers, and delivered high-profile releases.
- Designed CI/CD automations for the core brand component library, enabling independent dependency management across department teams.
- Diagnosed rendering bottlenecks on DirecTV's 2nd highest-traffic page: 149% faster load time, 37% smaller bundle, 93% reduction in interaction latency.
- Independently built a component preview GUI for pre-staging review - adopted org-wide.

Interactive Developer | Warner Bros. (via Cognizant)

October 2020 – October 2023

- Sole owner of requirements, design, and delivery for 7 enterprise apps - casting, script management, musical clearance, and production assets - serving 1,000 users across studio operations.
- Directed an 8-person offshore development team as sole onshore technical voice, owning UI decisions, implementation guidance, and PR review across all releases.
- Reduced data entry time by ~2× by digitizing manual workflow across studio operations.

EDUCATION

Master of Science: Artificial Intelligence | **UT Austin** | 2025-2026 | *Completed in 3 semesters*
Focus: Interpretability, NLP, Applied ML systems

Bachelor of Science: Computer Science | **CU Boulder** | 2016-2020 | Chancellor's Scholar
Minor: Psychology | Focus: Human-Computer Interaction

ML PROJECTS

Interpretability of Clinical Sequence JEPA models: First study to apply LLM-based interpretability methods and manifold analysis to a JEPA model's components. Applied TopK sparse autoencoders, linear/logistic probes, and novel geometric techniques (SAE boolean-tree decomposition, label-subspace reconstruction from SAE decoder directions).
Tech: PyTorch, MIMIC-IV, BigQuery, A100 [Code] [*arXiv - submission in progress*]

Syntactic Negation Probing: Built a dependency-parse diagnostic (NPAS) that localizes negation cues to seven syntactic regions, exposing where a fine-tuned ELECTRA-small leans on negation as a shortcut in SNLI. Defined a Negation Reliance Index separating dataset bias from model bias, applied contrast-set augmentation and slice-aware re-weighting – lifting the hardest negation slice from 50% to 85%.
Tech: PyTorch, Hugging Face, spaCy [Code] [Paper]

FireFusion: ConvFormer-based spatiotemporal model that adds attention over the feature axis to learn which feature combinations drive ignitions. Integrated 24 features across meteorological, fuel, topographic, and human-risk inputs. Implemented physics-aware masking and dual-head decoder for ignition probability and cause. Achieved 0.62 AUPRC vs. 0.08 baseline (7.75×) on WA test set.
Tech: PyTorch, xarray, MODIS, gridMET. [Code] [Paper]