

Frederick Hayes

Data Scientist | Independent Researcher in Artificial Intelligence, Computational Neuroscience & Neuromorphic AI
San Antonio, Texas, USA | frederickhayes3@gmail.com | fhayes3.com | github.com/fehayes | linkedin.com/in/frederickhayes3

RESEARCH PROFILE

My research explores how intelligent systems construct, preserve, and adapt internal representations under uncertainty. Drawing on machine learning, probabilistic inference, and computational neuroscience, I investigate the computational principles that enable robust adaptive behavior in artificial agents. My work combines recurrent and spiking neural networks, sequential decision-making, representation learning, mechanistic interpretability, and controlled experimentation to study latent-state continuity, uncertainty-aware control, and adaptive behavior under information degradation.

Alongside independent research, I have developed applied machine learning systems for industrial diagnostics, probabilistic risk modeling, hyperspectral-imaging evaluation, and retrieval-augmented generation. My long-term goal is to bridge practical AI engineering with foundational questions about intelligence, learning, adaptive cognition, biologically inspired architectures, and post-von Neumann computing. Recent work includes an accepted SBP-BRiMS 2026 poster and speed presentation on internal control regimes for behavior representation, along with public research artifact repositories supporting manuscript drafts, figures, tables, and reproducibility materials.

RESEARCH THEMES

Artificial Intelligence Machine learning; representation learning; mechanistic interpretability; probabilistic AI; foundation models	Computational Neuroscience Adaptive control; homeostatic and allostatic systems; neural dynamics; cognitive modeling
Neuromorphic & Alternative Computing Spiking neural networks; neuromorphic computing; post-von Neumann architectures; energy-efficient AI; event-driven computation	Statistical Learning & Sequential Decision Making Bayesian inference; Markov models; MCMC; Monte Carlo methods; reinforcement learning; recurrent neural networks

EDUCATION

University of California, Berkeley | **Master of Information and Data Science (MIDS) | May 2025**

GPA: 3.90 / 4.00. Selected coursework: Generative AI; Natural Language Processing with Deep Learning; Machine Learning at Scale; Experiment Design and Causal Inference; Applied Machine Learning. Activities: Student Ambassador; Student Mentor; Philosophy in AI Reading Group; GenAI Paper Discussion Group. NeuroBeacon selected for the UC Berkeley School of Information Spring 2025 Capstone Showcase.

University of the Incarnate Word | **B.S., Computer Information Systems | 2022**

RESEARCH PROGRAM

Adaptive Intelligent Systems Under Uncertainty

My current research investigates how intelligent systems preserve, update, and adapt latent internal representations under uncertainty, incomplete information, and changing environments. Current themes include latent-state continuity, internal control regimes under partial observability, hazard-aware adaptive control, representation learning under information degradation, spiking neural networks, and probabilistic inference for adaptive intelligent systems.

INDEPENDENT RESEARCH

Erasure Horizon Research Program | 2025-Present

- Developing a multi-stage experimental program investigating how recurrent agents distinguish temporary information loss from destructive internal-state degradation.
- Designed a reproducible pipeline of progressively validated experiments spanning safe foraging, additive calibration, channel switching, I/O suspension versus Erasure discrimination, hazard-aware control, reward abandonment, and latent-state geometry.
- Implemented hidden-state probes, PCA, centroid displacement, perturbation analyses, ablations, policy sensitivity diagnostics, and multi-seed confirmatory evaluations to test representation dynamics under uncertainty.
- Current work focuses on causal interventions into latent-state geometry, oscillatory gating mechanisms, reward-distractor suppression under Erasure pressure, and allostatic regulation of future control capacity.

Internal Control Regimes for Behavior Representation in Partially Observable Agents | 2025-2026

- Single-author conference paper accepted for poster presentation and speed presentation at SBP-BRiMS 2026; develops a behavior-representation workflow for partially observable simulated agents.
- Combined frozen behavioral evaluation, internal-state probes, compression analyses, feature-family controls, perturbations, and event-level prediction to study recurrent and spiking agents.
- Framed internal control regimes as distributed, predictive, action-relevant states that explain behavior beyond external outcomes alone.

- Released a public research artifact repository containing the camera-ready conference draft, manuscript artifacts, final derived tables, figures, and figure/appendix generation notebooks: github.com/fehayes/predictive-allostatic-organization

SELECTED RESEARCH PROJECTS

Audience-Adaptive Retrieval-Augmented Generation | UC Berkeley Generative AI | 2025

- Solo final project designing a dual-audience RAG system for engineering researchers and marketing professionals using LangGraph orchestration, hybrid retrieval, reranking, and audience-aware prompting.
- Developed the Fred Score, a composite evaluation metric combining fluency, semantic match, retrieval alignment, and responsiveness; the metric was incorporated into graduate Generative AI course curriculum.
- Final configuration achieved Fred-R 0.82 and Fred-M 0.75 on a 75-question test set; error analysis included hallucination diagnostics and the Paper Airplane Test for detecting fluent but unsupported answers.

NeuroBeacon | UC Berkeley MIDS Capstone | 2025

- Team-based capstone project selected for the UC Berkeley School of Information Spring 2025 Capstone Showcase.
- Built an adaptive cognitive-training system using calibrated reinforcement learning, ELO-style difficulty modeling, and game-based interaction design.
- Contributed to UI, game design, reinforcement-learning components, and DQN-related implementation. Project page: ischool.berkeley.edu/projects/2025/neurobeacon

INDUSTRY RESEARCH & PROFESSIONAL EXPERIENCE

Neogen Corporation | Data Scientist | Apr 2023-Present

- Develop applied machine learning and statistical modeling systems for industrial diagnostics, food safety, and regulated operational contexts.
- Built an internal RAG system for document management, knowledge retrieval, and summarization across regulatory, technical, and KPI-focused materials; system reached full internal access across multiple organizational groups.
- Designed statistical modeling evidence chains used to support proposed changes to federal regulatory standards for particle measurement.
- Develop data lakehouse and Azure-based analytics pipelines emphasizing reproducibility, auditability, and cross-functional deployment.
- Conducted hyperspectral-imaging threshold-screening evaluation using grouped sample-level validation, false-negative-first decision analysis, and evidence-boundary reporting.

Industrial Failure Risk Prediction and Operational Monitoring | Internal technical report / client-facility pilot | 2026

- Built a GBM risk model combined with Markov-chain location dynamics and decision-oriented flagging for fixed-budget operational review.
- Combined score improved out-of-time holdout performance over GBM alone (ROC AUC 0.94 vs. 0.93; PR AUC 0.658 vs. 0.64) and reduced expected decision cost under a fixed 2% flag budget.
- Report was circulated internally and presented to R&D; facility identity anonymized in external materials.

University of the Incarnate Word | Web Platform Specialist | Jul 2018-Apr 2023

Supported higher-education web platforms, web governance, accessibility, and stakeholder training across institutional units; provided technical training and support for users with diverse technical backgrounds.

Case Western Reserve University | Applications Developer | Oct 2014-Jul 2018

Developed and maintained higher-education web applications and platform workflows, building foundational experience in production software, stakeholder requirements, and technical communication.

PUBLICATIONS, MANUSCRIPTS & TECHNICAL REPORTS

Conference Presentations and Proceedings

Hayes, F. (2026). Internal Control Regimes for Behavior Representation in Partially Observable Agents. Accepted for poster presentation and speed presentation, 19th International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction and Behavior Representation in Modeling and Simulation (SBP-BRIMS 2026). Online non-archival proceedings.

Manuscripts and Research Artifacts

Hayes, F. (2026). Predictive Allostatic Organization in Recurrent and Spiking Agents Under Partial Observability. Manuscript draft and research artifact repository. github.com/fehayes/predictive-allostatic-organization

Hayes, F. (2026). Erasure Horizon: Reward Subordination and Latent-State Continuity Under Internal-State Threat. Manuscript draft and research artifact repository. github.com/Fhayes3/erasure-horizon

Technical Reports

Hayes, F. (2026). Industrial Failure Risk Prediction and Operational Monitoring: GBM Risk Model + Markov-Chain Location Dynamics + Decision-Oriented Flagging. Internal technical report, Neogen Corporation.

Hayes, F. (2026). Hyperspectral Imaging Threshold-Screening Evaluation for Food-Safety Decision Support. Internal executive summary / technical evaluation, Neogen Corporation.

Hayes, F. (2025). Retrieval-Augmented Generation for Audience-Specific Answering. Final report, UC Berkeley Generative AI.

Selected Technical Writing

Hayes, F. (2026, March 25). Let's Talk About the Work: Building a Mind That Doesn't Just React. Medium.

Hayes, F. (2026, May 1). Let's Talk About the Work: What Neuroscience Taught Me About Synthetic Intelligence. Medium.

Hayes, F. (2025, August 19). Let's Talk About the Work: The Problem with One-Size-Fits-All AI. Medium.

RESEARCH PRESENTATIONS & PROFESSIONAL ACTIVITIES

- Accepted poster presenter and speed presentation speaker, "Internal Control Regimes for Behavior Representation in Partially Observable Agents," SBP-BRiMS 2026, Carnegie Mellon University / Pittsburgh, PA, October 2026.
- Presenter and participant, UC Berkeley Philosophy in AI Reading Group.
- Presenter and participant, UC Berkeley GenAI Paper Discussion Group.
- Student Ambassador and Student Mentor, UC Berkeley MIDS.
- Presented internal technical work at Neogen R&D, including industrial risk monitoring, hyperspectral-imaging threshold-screening evaluation, and statistical modeling evidence chains.

HONORS & RECOGNITION

- UC Berkeley School of Information Spring 2025 Capstone Showcase selection: NeuroBeacon.
- UC Berkeley curriculum adoption: Audience-Adaptive AI / Fred Score evaluation project incorporated into graduate Generative AI coursework.
- Neogen Technical Achievement recognition for deploying an internal RAG system supporting document management and knowledge retrieval.
- Neogen Innovation Achievement recognition for statistical modeling work contributing to evidence used in proposed federal regulatory changes for particle measurement.

RESEARCH METHODS

- Experimental machine learning: controlled simulation; frozen evaluation protocols; multi-seed confirmatory experiments; ablation studies; stress variants; held-out evaluation.
- Mechanistic AI: internal-state probing; representation geometry; PCA; centroid displacement; perturbation analysis; causal intervention.
- Statistical and probabilistic analysis: Markov-chain modeling; calibration; fixed-budget decision analysis; causal inference; Bayesian and Monte Carlo methods.
- Computational neuroscience: recurrent dynamics; spiking neural networks; homeostatic and allostatic regulation; energy-sensitive control.

TECHNICAL EXPERTISE

- Programming and data: Python; R; SQL; NumPy; pandas; SciPy; Jupyter/Colab.
- Machine learning and AI: PyTorch; scikit-learn; gradient boosting / GBM; LangChain; LangGraph; Hugging Face; Qdrant; Stable-Baselines3; Gymnasium; Norse.
- Cloud and engineering: Azure; Azure ML; Azure Synapse; Git; reproducible experiment pipelines; structured logging.
- Research artifacts: GitHub repositories; manuscript artifact packaging; figure-generation notebooks; derived-table workflows; reproducibility documentation.