WILLIAM C. STANFORD

+1(254) 224-7469 ◇ Houston, TX
william.stanford.c@gmail.com ◇ linkedin.com/in/william-c-stanford/ ◇ github.com/william-c-stanford

EXPERIENCE

Director of AI 09/2024 - present
WasteLinq *Houston, TX*

- **Shipped an AI waste-classification agent for Veolia ($3B revenue customer)** end-to-end, from customer discovery through deployment, fine-tuning Gemini on Vertex AI and deploying a FastAPI service on AWS (ECS, RDS, Cognito, ALB) now generating 20K+ structured waste profiles per month in their portal.
- **Architected and built a multi-service AI platform from scratch across 7 microservices:** LLM prediction engine (Lambda, SQS, DynamoDB, Fargate, API Gateway), multimodal PDF extraction (Gemini vision OCR, 23 parallel extraction pipelines), manifest processing (Step Functions, Terragrunt), and disposal-portal automation, all with CI/CD, Terraform IaC, and production observability.
- **Designed a human-in-the-loop manifest workflow that cut document processing time 85%:** automated PDF extraction, AI-driven validation, EPA e-Manifest auto-submission, and a profile-matching dashboard, enabling operators to process hazardous waste manifests at scale instead of re-keying data by hand.
- **Built a portal automation product from zero** (Chrome Extension + Lambda backends) that maps internal waste profiles to vendor-specific form schemas and auto-fills across 4 disposal facility portals (Tradebe/SAP, Republic Services), eliminating hours of manual data entry per profile.
- **Fine-tuned Gemini models on Vertex AI** for waste classification, improving precision and recall 50%+ over foundation model baselines; designed a multi-provider LLM fallback system with two-pass prediction graphs and structured-output validation via Pydantic.
- **Built a hybrid semantic + keyword search service** (FastAPI, sentence-transformers, ECS) over 500+ internal waste product codes and federal regulatory classifications, enabling dynamic routing of waste profiles to appropriate vendor disposal facilities by matching waste characteristics to facility acceptance criteria.
- **Full-stack contributor across the entire platform:** extended the Django REST API (PostgreSQL, row-level security multi-tenancy, Celery), built React/TypeScript features integrating 5+ AI services, and redesigned UX workflows for time-intensive compliance processes in the web app.

Fractional CTO 11/2025 - present
Elaiya *Houston, TX*

- **Founded and built an enterprise AI evaluation platform from zero as solo CTO:** managers practice coaching with voice-based AI "employees" and are scored on 7 leadership competencies, mapped to a Dreyfus skill model with longitudinal pre/post-training competency tracking.
- **Designed an LLM-as-judge evaluation engine** using the OpenAI Realtime API for live voice agents and structured-output rubric grading, with deterministic scoring pipelines and assessment-state integrity triggers across the in-progress → session → feedback → completed lifecycle.
- **Architected a multi-tenant SaaS on Supabase** (Postgres + row-level security, 5-role RBAC, org-scoped data isolation) with 18+ Deno edge functions spanning assessment lifecycle, real-time coaching, report generation, Stripe billing, and ROI computation; hardened against RLS recursion and cross-tenant data leakage.
- **Shipped a quantitative ROI engine** computing cohort-prorated L&D spend, avoided-cost savings, and competency-improvement deltas, giving enterprise buyers measurable evidence tying training spend to behavioral change.
- **Production engineering rigor solo:** 83-test suite (unit, RPC integration, Playwright E2E), CI/CD, security-definer RLS helpers, and full-stack ownership across a React/TypeScript + Vite frontend and TypeScript serverless backend.

CTO 01/2023 – 09/2024
VisualLabs AI *Chicago, IL*

- **Selected for Techstars '23 Accelerator Cohort,** securing pre-seed funding and mentorship from industry leaders; and established a strong technical and product foundation.
- **Architected and deployed a multi-modal video generation platform on AWS** leveraging LLMs, diffusion, and temporal consistency models, reducing video rendering time by 65% and supporting 100k+ monthly render requests.
- **Developed proprietary fine-tuning pipelines** for LLMs and diffusion models, including custom LoRAs and control nets, achieving 30% improvement in content fidelity and personalization quality across image and video outputs.
- **Led applied research and integration of emerging AI trends** (e.g., multimodal reasoning, agentic workflows, and video-conditioned diffusion), directly driving a 40% increase in feature adoption and positioning VisualLabs as a first mover in generative video.
- **Established MLOps and DevOps infrastructure** with CI/CD pipelines, containerized training environments, and distributed inference clusters, cutting deployment cycles from weeks to days.

Graduate Research Assistant 08/2019 – 07/2023
University of North Carolina at Chapel Hill *Chapel Hill, NC*

- **Published four peer-reviewed papers** in computational neuroscience and medical AI, shipping models from raw data to published results across brain network analysis, disease forecasting, and clinical NLP.
- **Trained deep learning and NLP models on multi-modal medical data** (MRI, PET, clinical text, speech) for Alzheimer's disease progression forecasting, subtype classification, and early detection via linguistic biomarkers, achieving state-of-the-art accuracy on predicting disease prgression.
- **Designed custom graph-theoretic algorithms from scratch** to quantify brain network robustness, producing new metrics adopted by collaborating research groups.
- **Built automated data pipelines for 5TB+ neuroimaging datasets,** handling normalization, segmentation, and feature extraction across thousands of subjects.

Undergraduate Research Assistant 04/2016 – 05/2019
University of Texas at Dallas *Richardson, TX*

- **First-author publication with 90+ citations** on a novel AAV-based viral vector for genome editing, presented at four conferences including the Society for Neuroscience Annual Meeting.
- **Independently designed and led a CRISPR gene delivery project from concept to results,** engineering a novel Cas9-p65 fusion construct to solve a targeted integration problem, demonstrating zero-to-one research ownership as an undergraduate.

SKILLS

Languages	Python, TypeScript, JavaScript, SQL, Terraform, BAML
Frameworks	FastAPI, Django, React, LangChain, Harbor, Pydantic AI, PyTorch, Pulumi, Terraform
AI/ML	LLM Fine-Tuning, RAG, Agents (ADK, MCP, A2A), Knowledge Graphs, Structured Output
Infrastructure	AWS (Lambda, ECS, SQS, DynamoDB, RDS, Step Functions), GCP (Vertex AI, Gemini), Supabase,
Methodologies	AI-native Development, MLOps, LLMOps, DevOps, CI/CD, Test-Driven Development
Soft Skills	Technical & Product Leadership, Product Psychology, UI/UX, Self-Learner, Scientific Writing

EDUCATION

Doctor of Philosophy in Neuroscience, The University of North Carolina - Chapel Hill 08/2019 - 07/2023
Advisor(s): Eran Dayan, Co-Advisor: Peter J. Mucha
Bachelor of Applied Mathematics, The University of Texas at Dallas 08/2017 - 05/2019
Bachelor of Neuroscience, The University of Texas at Dallas 08/2013 - 05/2017

CONTINUING EDUCATION

Techstars Startup Accelerator — Mentorship-driven startup accelerator (Chicago)	Fall 2023
AI Product Leadership — by Miqdad Jaffer (Product Leader at OpenAI)	Dec 2025
UX and Product Psychology — Applying psychology to build better products by Growth.Design	Sept 2025
ProductLed MBA — How to drive Product Led Growth in your company by ProductLed Inc.	Oct 2025
Senior Engineer to Lead — Course on engineering leadership by Gregor Ojstersek	June 2025

PUBLICATIONS

Stanford, W., Mucha, P. J. & Dayan, E. Redundancy protects processing speed in healthy individuals with accelerated brain aging. *bioRxiv,* 2024–07 (2024) [paper].

Stanford, W., Mucha, P. J. & Dayan, E. Age-related differences in network controllability are mitigated by redundancy in large-scale brain networks. *Communications Biology* **7,** 701 (2024) [paper].

Kozak, S., Dezachyo, O., Stanford, W., Bar-Haim, Y., Censor, N. & Dayan, E. Elevated integration within the reward network underlies vulnerability to distress. *Cerebral Cortex* **33,** 5797–5807 (2023) [paper].

Iyer, A., Stanford, W., He, T., Dayan, E., Xu, J. & Xavier, R. Salience Network Segregation in Youth and Young Adults With Prodromal Psychosis Symptoms. In *NEUROPSYCHOPHARMACOLOGY* **48,** 352–352 (2023) [paper].

Stanford, W. C., Mucha, P. J. & Dayan, E. A robust core architecture of functional brain networks supports topological resilience and cognitive performance in middle-and old-aged adults. *Proceedings of the National Academy of Sciences* **119,** e2203682119 (2022) [paper].

Kwak, K., Stanford, W., Dayan, E. & Initiative, A. D. N. Identifying the regional substrates predictive of Alzheimer's disease progression through a convolutional neural network model and occlusion. *Human Brain Mapping* **43,** 5509–5519 (2022) [paper].

Kumar, N., *Stanford, W., De Solis, C., Aradhana, Abraham, N. D., Dao, T.-M. J., Thaseen, S., Sairavi, A., Gonzalez, C. U. & Ploski, J. E. The development of an AAV-based CRISPR SaCas9 genome editing system that can be delivered to neurons in vivo and regulated via doxycycline and Cre-recombinase. *Frontiers in Molecular Neuroscience* **11,** 413 (2018) [paper].