Jennifer Williams

jlw1@alumni.cmu.edu
http://github.com/jennwilliams
jennwilliams.github.io
Los Angeles, California

Highlights

- Research Scientist focused on building solutions using Natural Language Processing (including Large Language Models), multimodal learning, AI agents, and deep learning, to deliver real world impact
- Developed ML models across diverse data types—including text, audio, videos, and timeseries data
- Built and deployed a production-ready Conversational AI system using an agentic framework, transformers, and retrieval-augmented generation (RAG) to improve employee productivity
- Led evaluation of generative AI system, including analyzing AI agent behavior and developing human-in-the-loop protocols; enabled rapid prototyping and adaptive deployment across user-facing systems
- Built, deployed, and optimized large-scale AI models, improving internal workflows for thousands of users and increasing customer satisfaction, in collaboration with cross-functional teams

Technical Skills

- Languages: Python, Bash, Java, R, MATLAB | DL Frameworks: PyTorch, TensorFlow, Hugging Face
- Expertise: generative AI, transformers, foundation models, multimodal learning, synthetic data, deep learning, multi-agent systems, Natural Language Processing (NLP), model and system evaluation
- Data Processing & Deployment: BigQuery, AWS, GCP, Docker, Airflow, high-performance computing

Education

PhD, *School of Computer Science, Carnegie Mellon University (CMU),* 2016 - 2022
Advisor: Dr. Leila Wehbe (Machine Learning Department)
Thesis: Modeling personalized speech and video processing to predict behavior

MS in Natural Science and MSc in Systems Biology, EU-USA Atlantis Dual Degree Program, 2013-2016
Roswell Park Comprehensive Cancer Center, University at Buffalo / University of Luxembourg

Experience

Senior Machine Learning Scientist, *CVS Health*, April 2023 - present

- Developed and deployed a Conversational AI system leveraging hybrid retrieval, boosting employee productivity by 38% and improving customer satisfaction by 2%
- Led design and creation of evaluation framework to assess Gen AI system performance, securing stakeholder approval and enabling deployment while supporting iteration for future releases
- Showcased impact to 1,600+ colleagues and President, through demos of co-developed AI products
- Led an Agile team to build and deploy platform infrastructure for ranking model evaluation and automated deployment, reducing time-to-deployment to 20 minutes, accelerating org-wide iteration
- Drove internal upskilling by co-leading AI education initiatives, founding an "ML Lunch and Learn" lecture series and leading "Generative AI Training" workshops for cross-functional teams

Selected Publications

- **Same cause; different effects in the brain.**
  M. Toneva*, J. Williams* (co-first), A. Bollu, C. Dann, L. Wehbe. *Causal Learning and Reasoning (CLeaR) 2022*
- **Behavior measures are predicted by how information is encoded in an individual's brain.**
  J. Williams, L. Wehbe. *arXiv 2021* (in submission)
- **Discriminative subtyping of lung cancers from histopathology images via contextual deep learning.**
  B.J. Lengerich, M. Al-Shedivat, A. Alavi, J. Williams, S. Labbaki, E.P. Xing. *medRxiv 2020*

Relevant Courses

- Machine Learning
- Statistical Methods in Machine Learning
- Intermediate Statistics
- Probabilistic Graphical Models
- Cognitive Neuroscience

Leadership and Service

- Co-organizer of the ICLR workshop *ICBINB Deep Learning* and the *"brAIn"* seminar series at CMU
- Reviewer for ICLR, NeurIPS (WiML, Learning from Time Series for Health), ICML (New in ML), Nature Scientific Reports, IJCAI, ML4H, MLHC, ECCB
- Mentored CS student (now PhD at Princeton)

Awards

- Top 10 Reviewer Machine Learning for Health (ML4H)
- International Conference on Machine Learning (ICML) Travel Award
- Invited to attend Machine Learning Summer School (MLSS)

Selected Research Projects

Modeling personalized video processing

- Question: Can individual differences in how multimodal video (audiovisual) information is encoded in the brain predict behavior?
- Method Innovation: Built on insights from two sub-fields of neuroscience (brain mapping and behavioral neuroscience), to create the first machine learning framework to identify individual differences in brain encoding and test if these differences predict behavior.
- Scientific Discoveries:
  - Individual differences in how the brain encodes video stimuli can predict behavioral variability.
  - Advised researchers to optimize their choice of neuroimaging task and feature-space for their behavior of interest.
- Paper: arXiv 2112.06048
- Code: github.com/brainML/great-apes

Disambiguating audio processing with causal inference

- Question: Why do neural network derived features of naturalistic audio predict large parts of the brain well? Do these brain zones process the features similarly or differently?
- Method Innovation: Developed a causal inference framework, that includes two new metrics, to provide insights beyond current brain mapping techniques. Specifically, the framework enables researchers to infer if a complex (multivariate and high dimensional) stimulus, such as audio, affects two brain zones similarly.
- Scientific Discovery: Real-world language stimuli (i.e., videos) do not affect all parts of the brain's language network similarly.
- Paper: CLeaR 2022 (arXiv 2202.10376)
- Code: github.com/brainML/stim-effect

Contextual deep learning for multimodal classification

- Question: Can sample-specific models, similarly to subject matter experts, effectively integrate multimodal data for accurate classification?
- Method Innovation: Inspired by contextual deep learning, created sample-specific multimodal models for lung cancer classification, by adapting Contextual Explanation Networks (CENs) (Al-Shedivat et al., JMLR 2020). Integrated both imaging and tabular data into the classification models.
- Scientific Discovery: Sample-specific multimodal models increase classification accuracy and capture the heterogeneity of biological processes underlying lung cancer.
- Paper: medRxiv (DOI: 10.1101/2020.06.25.20140053)

Fellowships

- **Digital Health Fellowship** – Center for Machine Learning and Health (2020 - 2021)
  Full tuition and stipend for 12 months and $3,000 for research-related expenses
- **NIH T32 Training Grant** – National Institute of Biomedical Imaging and Bioengineering (2017 - 2019)
  Full tuition and stipend for 2 years and $6,000 for research-related expenses
- **CanSys MS Scholarship** – Atlantis EU-USA Training Program (2013 - 2015)
  Stipend for 12 months