

Alexander Miserlis Hoyle

Andreasstrasse 5, OAT X11
Zürich, Switzerland 8092

alexander.hoyle@ai.ethz.ch
alexanderhoyle.com

ACADEMIC APPOINTMENTS

AI Center Postdoctoral Fellow, ETH Zürich

Fall 2024 – Present

EDUCATION

PhD in Computer Science

University of Maryland

Dissertation: Developing and Measuring Latent Constructs in Text

Advisor: Philip Resnik

Committee: Jordan Boyd-Graber, Michelle Mazurek, Kris Miler, Aaron Schein

MSc in Computational Statistics and Machine Learning

University College London

Thesis: Citation Detected: Automated Claim Detection through Natural Language Processing

Advisors: Sebastian Riedel, Jeff Mitchell

Received Distinction

BA in Mathematics

Wesleyan University

REFEREED PUBLICATIONS

- H. Licht*, R. Sarkar*, P. Yu, P. Goel, N. Stoehr, E. Ash, **A. Hoyle***
Measuring scalar constructs in social science with LLMs. *EMNLP*. 2025.
- Y. Fan, Y. Tian, S. Ravfogel, M. Sachan, E. Ash*, **A. Hoyle***
The Medium Is Not the Message: Deconfounding Document Embeddings via Linear Concept Erasure. *EMNLP*. 2025.
- T. Nguyen, K. Du, **A. Hoyle**, R. Cotterell
How Persuasive Is My Context? *EMNLP*. 2025.
- A. Hoyle***, L. Calvo-Bartolomé*, J. Boyd-Graber, and P. Resnik
PROXANN: Use-Oriented Evaluations of Topic Models and Document Clustering. *ACL*. 2025.
- Z. Li, L. Calvo-Bartolomé, **A. Hoyle**, D. K. Stephens, P. Xu, J. F. Fung, A. Dima, and J. Boyd-Graber
LLMs Struggle to Describe the Haystack without Human Help: A Social Science-Inspired Evaluation of Topic Models. *ACL*. 2025.
- S. Wyckoff-Gaynor, K. Miler., P. Goel, **A. Hoyle**, and P. Resnik.
Express Yourself (Ideologically): Legislators’ Ideal Points Across Audiences. *The Journal of Politics*. 2025.
- R. Sarkar, P. Wu, K. Miler, **A. Hoyle***, P. Resnik*.
PairScale: Analyzing Attitude Change with Pairwise Comparisons. *Findings of NAACL*. 2025.
- E. Cho, **A. Hoyle**, Y. Hermstrüwer
Modeling Motivated Reasoning in Law: Evaluating Strategic Role Conditioning in LLM Summarization. *Natural Legal Language Processing Workshop*. 2025.
- C. Xiong, J. Ni, Y. Fan, V. Zouhar, D. Rooein, L. Calvo-Bartolom, **A. Hoyle**, Z. Jin, M. Sachan, M. Leippold, D. Hovy, M. El-Assady, E. Ash
Co-DETECT: Collaborative Discovery of Edge Cases in Text Classification. *EMNLP: System Demonstrations*. 2025.
- N. Balepur, M. Shu, **A. Hoyle**, A. Robey, S. Feng, S. Goldfarb-Tarrant, J. Boyd-Graber
A SMART Mnemonic Sounds like “Glue Tonic”: Mixing LLMs with Student Feedback to Make Mnemonic Learning Stick. *EMNLP*. 2024.
- A. Leto, S. Roy, **A. Hoyle**, D. Acuna, M. L. Pacheco
A First Step towards Measuring Interdisciplinary Engagement in Scientific Publications: A Case Study on NLP+ CSS Research. *NLP+CSS Workshop*. 2024.
- C. M. Pham, **A. Hoyle**, S. Sun, and M. Iyyer.
TopicGPT: A Prompt-based Topic Modeling Framework. *NAACL*. 2024

- A. Hoyle*, R. Sarkar*, P. Goel, and P. Resnik.
Natural Language Decompositions of Implicit Content Enable Better Text Representations. *EMNLP*. 2023.
- D. Stammbach, V. Zouhar, A. Hoyle, M. Sachan, and E. Ash.
Re-visiting Automated Topic Model Evaluation with Large Language Models. *EMNLP*. 2023.
- A. Hoyle, P. Goel, R. Sarkar, and P. Resnik.
Are Neural Topic Models Broken? *Findings of EMNLP*. 2022.
- A. Hoyle*, P. Goel*, D. Peskov*, A. Hian-Cheong*, J. Boyd-Graber, and P. Resnik.
Is Automated Topic Model Evaluation Broken?: The Incoherence of Coherence. *NeurIPS* (Spotlight Presentation). 2021.
- A. Hoyle, A. Marasović, and N. A. Smith.
Promoting Graph Awareness in Linearized Graph-to-Text Generation. *Findings of ACL*. 2021.
- P. Rodriguez, J. Barrow, A. Hoyle, J. P. Lalor, R. Jia, and J. Boyd-Graber.
Evaluation Examples are not Equally Informative: How should that change NLP Leaderboards?. *ACL*. 2021.
- A. Hoyle*, P. Goel*, and P. Resnik.
Improving Neural Topic Models using Knowledge Distillation. *EMNLP*. 2020.
- A. Hoyle, L. Wolf-Sonkin, H. Wallach, I. Augenstein, and R. Cotterell.
Unsupervised Discovery of Gendered Language through Latent-Variable Modeling. *ACL*. 2019.
- A. Hoyle, L. Wolf-Sonkin, H. Wallach, R. Cotterell, and I. Augenstein.
Combining Sentiment Lexica with a Multi-View Variational Autoencoder. *NAACL*. 2019.

WORKING PAPERS AND PRESENTATIONS

- A. Hoyle, P. Resnik.
Language Models as a Partner in Construct Development. Presented at the Annual Meeting of the *American Political Science Association*. 2024.
- H. Licht, A. Hoyle, R. Sarkar, P. Goel, N. Stoehr, P. Wu.
Scoring political texts by automating pairwise comparison with transformers. Presented at the Annual Meeting of the *American Political Science Association*. 2024.
- S. Schulhoff, M. Illie, N. Balepur, K. Kahadze, [24 other authors omitted], A. Hoyle, P. Resnik.
The Prompt Report: A Systematic Survey of Prompting Techniques. *Working paper*. 2025.
- A. Hoyle, N. Farra, H. Wallach, M. Ribeiro, and A. Olteanu.
Diagnostic Tests for Long-Tail Failures in Language Generation Models. *In preparation*.

* Denotes equal contribution / advising.

INVITED TALKS

Computational Social Science as a Problem Space for NLP

University of Massachusetts-Amherst, November 2023
Stanford University, November 2023
Bocconi University, November 2023
ETH Zürich AI Center, January 2024
University of Chicago CHAI Lab, February 2024
Microsoft Research Seminar, March 2024
ETH Zürich NLP Group, March 2024
Royal Holloway, University of London, March 2025
University of Minnesota, March 2025
TU Wien, October 2025
Télécom Paris, October 2025

Testing Natural Language Generation Models for Undesirable Behavior in the Long Tail

Microsoft Experiences + Devices Team, October 2022

Structuring Posterior Inference

University of Maryland CLIP Colloquium, February 2020

Some things worth knowing about Large Language Models

University of Zürich Digital Society Initiative Seminar, October 2025

RESEARCH EXPERIENCE

PhD Intern, Microsoft Research

Fairness, Accountability, Transparency, and Ethics group, mentored by Alexandra Olteanu, Marco Ribeiro, and Hanna Wallach (publication in preparation)

PhD Intern, Allen Institute for Artificial Intelligence

AllenNLP Group, mentored by Noah A. Smith and Ana Marasović

Master's Student and Intern, University College London

UCL Machine Reading Group, advised by Sebastian Riedel and Jeff Mitchell

Quantitative Analysis Center Research Apprenticeship, Wesleyan University

Developmental Psychology Lab, advised by Anna Shusterman

TEACHING AND SERVICE

Teaching Experience

- **Projects in Machine Learning Research**, *Joint Lecturer*, ETH Zürich. Master's course, approx. 100 students. Spring 2025.
- **Data Science**, *Graduate Teaching Assistant*, University of Maryland. Advanced undergraduate, approx. 200 students. Spring 2020.
- **Machine Learning**, *Graduate Teaching Assistant*, University of Maryland. Advanced undergraduate, approx. 200 students. Fall 2019.

Guest Lectures

- **Language Models for Law and Social Science**, ETH Zürich. Title: *Word Embeddings*. Spring 2025.
- **Practical Approaches to Data Science with Text**, Emory University. Title: *Some ways to figure out what people are talking about online*. Fall 2024
- **Topics in the Computational Cognitive Neuroscience of Language**, University of Maryland. Title: *Bayes, the Free Energy Principle, and Predictive Coding*. Spring 2022
- **Machine Learning**, University of Maryland. Title: *Gaussian Mixture Models and Expectation Maximization* (two sections). Fall 2019

Mentoring

- Yu Fan, PhD at the ETH Zürich Center for Law and Economics (2024–present). Supervising multiple projects in legal NLP.
- Lehan Zhang, PhD at the ETH Center for Law and Economics. Supervising project on causes and dynamics of online misogyny
- Lorena Calvo-Bartolomé, PhD in NLP at Universidad Carlos III de Madrid (2024–present). Supervised publication on topic model evaluation, hosted as visitor at ETH.
- Eunjung Cho, PhD at the ETH Center for Law and Economics (2024–present). Co-supervising automated legal summarization project.
- Jakob Johannes Bauer, Master's in Computer Science at ETH. (2025–present). Co-supervising Master's thesis on verifiable reasoning in LLMs.
- Tu Nguyen, Master's in Data Science at ETH (2025–present). Co-supervising Master's project on measuring effect of codebooks on model reasoning.
- Konstantina Timoleon, Master's in Artificial Intelligence at the University of Zürich (2025–present). Advising master's thesis and project on gender bias in generated text.
- Eleftherios Kanavakis, Master's in Data Science at ETH (2024–present). Supervising Master's thesis on legal annotation for asylum cases.
- Jonathan Quinn, Bachelor's in Computer Science at ETH (2024–2025). Supervised Bachelor's thesis on word embeddings for transcribed speech.
- Sander Schulhoff, Undergraduate in CS at UMD (2023–2024). Developed survey paper on prompt engineering techniques.
- Pranav Dupelet, Undergraduate in CS at UMD (2023). Supervised re-implementations of social science publications reliant on sentence embeddings
- Joseph Martinez, through the Graduate Student Mentorship Initiative (GSMI, see below) (2024).
- Aditya Mandke, through GSMI (2021). Accepted into a master's in Computer Science at USC.
- Keren Fuentes, through GSMI (2020). Accepted as predoctoral intern at Facebook AI Research.

Reviewing

- Area Chair for ACL Rolling Review (2024–present)
- Reviewer for ACL Rolling Review (EMNLP, NAACL, ...) (2022–present)
- Reviewer for TACL (2023)
- Reviewer for FAccT (2024)

Service

- Mentor to computer science graduate school applicants from nontraditional backgrounds as part of the *Científico Latino* Graduate Student Mentorship Initiative (2020, 2021, 2024)
- Graduate Assistant Advisory Committee, Director of Data & Research. Spring 2020–2024.
Developed comparative analysis of graduate stipends and surveyed graduate assistants about working conditions. Delivered presentation of results to Provost and Dean of Grad. School; to University Faculty Senate; as testimony before state legislature.
- Graduate Student Government Representative for Computer Science. Fall 2020–Spring 2021.
- Panelist for Undergraduate Datafest Competition at Wesleyan University. Spring 2016.

GRANTS AND AWARDS

Co-authored Grants

- Effective Few-Shot Learning for Constructs in Psychological and Social Science. *UMD Social Data Science Center Seed Grant*
- Modeling Co-Decisions: A Computational Framework Using Language and Metadata. *NSF (Award No. 2008761)*
- Advanced Topic Modeling Methods to Analyze Text Responses in COVID-19 Survey Data. *NSF (Award No. 2031736)*

Awards

Best Reviewer Award in the Information Retrieval and Text Mining Track, EMNLP 2023
Accepted to German Academic Exchange Service (DAAD) as an AInet fellow in Human-centered AI
Goldhaber Travel Award, University of Maryland
International Conference Student Support Award, University of Maryland
Computer Science Department Travel Award, University of Maryland
Dean's Fellowship for Outstanding Academic Achievement, University of Maryland
Dean's List, University College London
Quantitative Analysis Center Research Apprenticeship, Wesleyan University

PROFESSIONAL EXPERIENCE

Research Analyst, The Brattle Group

- Performed analysis, created reports, and delivered weekly presentations to U.S. Department of Justice to assess conditions of NYC public housing. Results of this work eventually led to restructuring of the housing authority's management, increased oversight, and a \$2.2 billion-dollar settlement.
- Initiated and later led development of large-scale automated document retrieval tool; created search backend, user interface, classification methods. Tool was used across several white-collar crime projects with government clients.
- Replicated and extended published econometric models for client analyses.

SELECTED SKILLS

- *Machine Learning and Data Analysis:* TensorFlow, PyTorch, Pyro, R, Pandas, Scikit-Learn, Shiny, SQL, Apache Spark, Excel, SPSS, Stata
- *Development tools:* Python, JavaScript, Django, Apache Solr, Amazon Web Services, Google Cloud Platform
- Fluent in French