Kevin Stowe, Ph.D.

kevincstowe@gmail.com
+1 517 648 8953
Chicago, IL, USA
LinkedIn
https://kevincstowe.github.io/

Summary

Research scientist with 10+ years experience in natural language processing, machine learning, LLMs, and more. I've undertaken exceptional research projects combining linguistic and machine learning approaches to many tasks and domains: building fraud detection tools, automating the generation of educational content, analyzing social media in crises, generating and detecting figurative language, and many more. I'm passionate about responsible and effective machine learning, and I'm seeking a senior research role where I can drive positive societal impact through cutting-edge NLP/AI research.

Skills

Science	Rapid prototyping, scientific deep dives, paper writing and presentation. Track record of excellent publications and products.
Professional	Client communications, team/project leadership, interdisciplinary collaboration. I can invent, lead, and/or contribute as needed.
Software	Python, Java; Deep Learning with Torch/HuggingFace; NLP w/ SpaCy, NLTK; Git; AWS (EC2, Sagemaker, S3, etc). Comfortable picking up any languages, tools, etc. as needed.
Languages	Native English, fluent German, beginner Russian

Employment History

2025 – · · · ·

Research Scientist, Pindrop Security, Atlanta, GA

- Researcher in fraud prevention across audio and text, implementing transformer- and LLM-based methodology to improve identify verification while minimizing fraud risk.
- Lead research into identifying bias in deepfake detection, fairness metrics, and evaluation strategies for deepfake detection systems, resulting in multiple publications.
- Delivered improved transformer-based classification of authentication events and robotic dialing attacks for customer calls for call authentication.
- Contributed to deepfake text detection shared task competitions.

2024

Research Scientist, Educational Testing Service (ETS), Princeton, NJ

2022 – 2024

Associate Research Scientist

- Fairness, bias, and accessibility project leader: explored socio-technical solutions for the responsible generation of educational content.
- Developed research and infrastructure for the generative AI team, incorporating responsible AI, natural language generation, and model evaluation.
- Automated Content Generation (ACG) team member: built natural language generation systems (LLM and others) for dozens of item types for testing/educational applications.
- Member of the Responsible AI working group, driving company policy in the responsible use of AI tools.

Employment History (continued)

2019 – 2022 **Postdoctoral Researcher**, Ubiquitous Knowledge Processing (UKP), Technical University of Darmstadt, Darmstadt, Germany
• Instructor for a graduate seminar on responsible AI focusing on the ethical use of AI, positive societal impacts, and public policy.
• Lead research in metaphor generation, dataset creation for figurative language inferences, dataset creation during crises, yielding top-tier conference publications.

2013 – 2019 **Research Assistant**, University of Colorado, Boulder, CO
• NLP expert for the Empowering the Public with Information in Crisis (EPIC) group, focusing on positive applications of NLP during crises.
• Research in annotation, lexical resources, figurative language, named entity recognition, and computational social science.

Service

• Session chair: Ethics, bias, and fairness, EMNLP 2025
• Area Chair, Association for Computational Linguistics (ACL), Feb 2025-present
• Reviewer, ACL Rolling Review, ACL, EMNLP, StarSEM, CoNLL, 2017-present.
• Computer science tutor for Linguistics students, focusing on Python and C++ essentials.

Education

2013 – 2019 **Ph.D., University of Colorado**
Joint Degree in Linguistics and Computer Science
Thesis: *Syntactic and semantic improvements to computational metaphor processing*

2009 – 2011 **Master of Arts, Indiana University**
Computational Linguistics Track

2004 – 2009 **Bachelor of Arts, Michigan State University**
Linguistics

Select Publications

1. Stowe, K., Afanaseva, S., Raimundo, R., Sun, Y. & Patil, K. Identifying Bias in Machine-generated Text Detection. *accepted to The 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026).* https://arxiv.org/abs/2512.09292 (2026).

2. Stowe, K. & Patil, K. Spotlights and Blindspots: Evaluating Machine-generated Text Detection. *accepted to The 15th edition of the Language Resources and Evaluation Conference (LREC 2026).* https://arxiv.org/abs/2604.16607 (2026).

3. Stowe, K. *et al.* Identifying Fairness Issues in Automatically Generated Testing Content. *19th Workshop on Innovative Use of NLP for Building Educational Applications.* https://arxiv.org/abs/2404.15104 (2024).

4. Stowe, K., Ghosh, D. & Zhao, M. Controlled Language Generation for Language Learning Items. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP): Industry Track* (eds Li, Y. & Lazaridou, A.) 294–305. https://aclanthology.org/2022.emnlp-industry.30 (2022).

5. Stowe, K., Utama, P. & Gurevych, I. IMPLI: Investigating NLI Models' Performance on Figurative Language. *Proceedings of the Association for Computational Linguistics (ACL),* 5375–5388. https://aclanthology.org/2022.acl-long.369 (2022).