

# Dustin Wright

TT Assistant Professor at Aalborg University Copenhagen  
From San Diego, CA, USA

[✉ dbw@cs.aau.dk](mailto:dbw@cs.aau.dk)  
[🔗 dustinwright.com](https://dustinwright.com)  
ORCID: 0000-0001-6514-8733

I am a tenure track assistant professor at Aalborg University Copenhagen. My research centers around reliable and sustainable AI. Core to my work is understanding when and how can NLP systems make the world's information reliable and accessible, while mitigating the environmental and social risks these systems can bring.

## Scientific interests

- Natural Language Processing
- Machine Learning
- Fact Checking
- Reliable NLP

## Education

Mar 2023. **PhD Computer Science**

University of Copenhagen

June 2019. **MSc in Computer Science**

UC San Diego

Dec 2014. **BS in Computer**

Engineering

San Diego State University

## Topics

Misinformation	Scholarly Text
Sustainability	Robustness
Few-Shot Learning	Bayesian Methods

## Languages

English	Native
Danish	B2

## Current

Feb. 2026 – **Tenure Track Assistant Professor**

Aalborg University Copenhagen  
TT Assistant Professor in the Computer Science Department focused on Natural Language Processing.

## Experience

Feb. 2024 – **DDSA Postdoc**

University of Copenhagen & University of Michigan  
2-year Danish Data Science Academy (DDSA) postdoctoral grant focused on automatic summarization and faithful science communication.

May 2024 – Oct. 2024 **Visiting Postdoc**

UMSI, University of Michigan  
Worked with David Jurgens and Lu Wang on query focused summarization of long documents. One paper in submission.

Feb. 2023 – Feb. 2024 **Postdoc**

DIKU, University of Copenhagen  
Performed research at the intersection of machine learning sustainability and resource efficiency. Spotlight paper at NeurIPS 2024 and one paper accepted to Communications of the ACM.

Oct. 2019 – Dec 2022 **PhD Fellow**

DIKU, University of Copenhagen (Adv: Isabelle Augenstein)  
NLP researcher for scientific fact checking and domain adaptation. Published 7 first author papers at ACL/EMNLP.

Jun. 2021 – Aug. 2021. **Research Intern**

Allen Institute for Artificial Intelligence; Remote  
Published a paper at ACL on generating scientific claims for automatic scientific fact checking.

Jun. 2019 – Aug. 2019. **Research Intern**

IBM; San Jose, CA  
Developed an active learning method to perform relation extraction and knowledge graph construction for biomedical text.

## Grants and Honors

Carlsberg Internationalization Postdoc Grant\* (1.7M DKK), 2025

\*Awarded but declined in order to take tenure track position instead

Danish Data Science Academy Postdoc Fellowship (1.2M DKK), 2023

## Grants and Honors (continued)

NeurIPS Spotlight Paper (top 2.5% out of 16,000 submissions), NeurIPS 2024

Honorable mention (top 5 submission), International Conference on Computational Social Science (IC2S2) 2023

Marie Skłodowska-Curie TALENT PhD fellowship, 2019

Best application paper, AKBC 2019

SDSU College of Engineering Valedictorian, 2015

## Invited Talks

*The Many Faces of Science Misinformation* (ICWSM Misinformation Workshop Keynote, June 23, 2025)

*Revealing Political Opinions in Large Language Models* (Last Friday Talks, Pioneer Center for AI, April 25, 2025, Denmark)

*Socially Sustainable NLP* (Research Colloquium, University of Minnesota, April 7, 2025, USA)

*LLM Tropes: Revealing Fine-Grained Values and Opinions in Large Language Models* (DIKU Bits, University of Copenhagen, November 19, 2024, Denmark)

*Revealing Fine-Grained Values and Opinions in Large Language Models* (University of Michigan CSS Working Group, August 9, 2024, USA)

*Fighting Misinformation in Science Communication with NLP* (University of Cambridge, June 2, 2023, UK)

*Automatically Ensuring Information Quality in Scientific Writing* (Elsevier, March 15, 2022, The Netherlands)

*Semi-Supervised Exaggeration Detection in Science* (University of Amsterdam, January 28, 2022, The Netherlands)

*Detecting and Generating Claims for Fact Checking* (IT University Copenhagen, March 23, 2021, Denmark)

*Cite-Worthiness Detection for Improved Scientific Document Understanding* (ETH Zürich, March 1, 2021, Switzerland)

*NormCo: Deep Disease Normalization for Biomedical Knowledge Base Construction* (IBM Research, July 10, 2019, USA)

## Press

*Efficiency is Not Enough: A Critical Perspective of Environmentally Sustainable AI* Montreal AI Ethics Institute, December 2023

*Interview on Scientific Misinformation Detection, Nvidia AI Podcast* Nvidia, February 16, 2022

“*Exaggeration Detector Could Lead to More Accurate Health Science Journalism*” Nvidia Blog, October 1, 2021

“*An NLP Approach to Exaggeration Detection in Science Journalism*”, unite.ai, September 15, 2021

## Supervision

*Clinical Note Error Detection with Large Language Models* Master’s project advisor, UCPH, until Mar 2025  
Wenkun Zhang

*Factual Text Generation*. PhD Co-advisor, UCPH, Oct 2024 onward  
Zain Muhammad Mujahid

*Compression Methods for Efficient Vision Transformers.* Bachelor's project advisor, UCPH, Feb 2023

Olga Henrietta Ptacek

*Uncertainty and Exaggerations of Scientific Findings in Social Media.* Master's project advisor, UCPH, June 2022

Jimmie Jin, Asger Thorleif Knudsen, Sylvester Leonhard Gorm Errebo Lee

*Stance Detection of Attitudes Toward Climate Change on Social Media* Bachelor's project advisor, UCPH, June 2020

Jimmie Jin, Asger Thorleif Knudsen, Sylvester Leonhard Gorm Errebo Lee

## Teaching

*Co-Instructor, Empirical Research Methods* University of Copenhagen, Fall 2025

*Tutorial Organizer, Addressing AI Driven Misinformation* IC2S2 2025

*Tutorial Organizer, Addressing AI Driven Misinformation* ICWSM 2025

*Teaching Assistant, Introduction to Natural Language Processing* University of Copenhagen, Fall 2021

*Introduction to University Pedagogy* University of Copenhagen (Course), Spring 2021

*Teaching Assistant, Introduction to Natural Language Processing* University of Copenhagen, Fall 2020

*Teaching Assistant, Web Science* University of Copenhagen, Spring 2020

*Teaching Assistant, Principles of AI: Probabilistic Reasoning and Decision Making* UC San Diego, Fall 2017

## Service

*Area Chair* Association for Computational Linguistics (ACL) 2026

*Reviewer* European Association for Computational Linguistics (EACL) 2026

*Reviewer* Journal of AI Research (JAIR) 2025

*Area Chair* Empirical Methods in Natural Language Processing (EMNLP) 2025

*Area Chair* Association for Computational Linguistics (ACL) 2025

*Reviewer* WNUT 2024

*Reviewer* Empirical Methods in Natural Language Processing (EMNLP) 2024

*Reviewer* European Association for Computational Linguistics (EACL) 2024

*Reviewer* North American Association for Computational Linguistics (NAACL) 2024

*Area Chair* Empirical Methods in Natural Language Processing (EMNLP) 2023

*Reviewer* BioNLP Workshop (@ACL) 2023

*Reviewer* Association for Computational Linguistics (ACL) 2022

*Publication Chair* Conference for Truth and Trust Online 2021

*Reviewer* Empirical Methods in Natural Language Processing (EMNLP) 2021

# Dustin Wright

Tenure-Track Assistant Professor at Aalborg University Copenhagen

## Publications

ORCiD ID: <https://orcid.org/0000-0001-6514-8733>

Google Scholar Bibliometrics (<https://scholar.google.com/citations?user=0Gk5UnYAAAAJ&hl=en>)

Publications: 29; First Authorships: 15; Corresponding Authorships: 15; Citations: 681; h-index: 15; i10-index: 18

### Papers in Conference Proceedings

**Wright, D.**, Mujahid, Z. M., Wang, L., Augenstein, I., & Jurgens, D. (2025). Unstructured Evidence Attribution for Long Context Query Focused Summarization. In *EMNLP 2025*. Association for Computational Linguistics.

Roitero, K., **Wright, D.**, Soprano, M., Augenstein, I., & Mizzaro, S. (2025). Collecting Cost-Effective, High-Quality Truthfulness Assessments with LLM Summarized Evidence. In *ACM Conference on Research and Development in Information Retrieval (SIGIR) 2025*. Association for Computing Machinery.

**Wright, D.**, Igel, C., & Selvan, R. (2024). BMRS: Bayesian Model Reduction for Structured Pruning. In *Neural Information Processing Systems (NeurIPS) 2024 (Spotlight)*. Neural Information Processing Systems Foundation

**Wright, D.**\*, Arora, A.\*., Borenstein, N., Yadav, S., Belongie, S., & Augenstein, I. (2024). LLM Tropes: Revealing Fine-Grained Values and Opinions in Large Language Models. In *EMNLP 2024 (Findings)*. Association for Computational Linguistics.

\* denotes equal contribution

Wuhrl, A., **Wright, D.**, Klinger, R., & Augenstein, I. (2024). Understanding Fine-grained Distortions in Reports of Scientific Findings. In *ACL 2024 (Findings)*. Association for Computational Linguistics.

**Wright, D.**\*, Pei, J.\*., Jurgens, D., & Augenstein, I. (2022). Modeling Information Change in Science Communication with Semantically Matched Paraphrases. In *EMNLP 2022*. Association for Computational Linguistics.

\* denotes equal contribution

**Wright, D.**, Wadden, D., Lo, K., Kuehl, B., Cohan, A., Augenstein, I., & Wang, L. L. (2022). Generating Scientific Claims for Zero-Shot Scientific Fact Checking. In *ACL 2022*. Association for Computational Linguistics.

**Wright, D.**, & Augenstein, I. (2021). Semi-Supervised Exaggeration Detection of Health Science Press Releases. In *EMNLP 2021*. Association for Computational Linguistics.

**Wright, D.**, & Augenstein, I. (2021). CiteWorth: Cite-Worthiness Detection for Improved Scientific Document Understanding. In *ACL 2021 (Findings)*. Association for Computational Linguistics.

Lima, L. C., **Wright, D.**, Augenstein, I., & Maistro, M. (2020). University of Copenhagen Participation in TREC Health Misinformation Track 2020. In *TREC*.

**Wright, D.**, & Augenstein, I. (2020). Transformer based multi-source domain adaptation. In *EMNLP 2020*. Association for Computational Linguistics.

Atanasova, P.\*., **Wright, D.**\*, & Augenstein, I. (2020). Generating label cohesive and well-formed adversarial claims. In *EMNLP 2020*. Association for Computational Linguistics.

\* denotes equal contribution

**Wright, D.**, & Augenstein, I. (2020). Claim check-worthiness detection as positive unlabelled learning. In *EMNLP 2020 (Findings)*. Association for Computational Linguistics.

**Wright, D.**, Katsis, Y., Mehta, R., & Hsu, C. (2019). Normco: Deep disease normalization for biomedical knowledge base construction. In *AKBC 2019*.

### Best Application Paper

**Wright, D.**, Yan, X., Srinivas, P., Kashani, A., & Ozturk, Y. (2015). A cloud to mobile application for consumer behavior modification. *Procedia Computer Science*, 62, 343-351.

Yan, X., **Wright, D.**, Kumar, S., Lee, G., & Ozturk, Y. (2015). Real-time residential time-of-use pricing: a closed-loop consumers feedback approach. In *Green Technologies Conference (GreenTech), 2015 Seventh Annual IEEE* (pp. 132-138). IEEE.

Massai, S., Routhu, S., **Wright, D.**, Moon, K. S., Ozturk, Y., & Lee, S. Q. (2015). A Wireless Visual Attention Brain Signal Monitoring System. In *MATEC Web of Conferences* (Vol. 32, p. 04005). EDP Sciences.

### Papers in Workshop Proceedings

Holm, A., Plank, B., **Wright, D.**, & Augenstein, I. (2022). Longitudinal citation prediction using temporal graph neural networks. In Proceedings of the AAAI 2022 Workshop on Scientific Document Understanding (SDU 2022).

Koh, E.S., Dubnov, S., & **Wright, D.** (2018). Rethinking recurrent latent variable model for music composition. *IEEE 20th International Workshop on Multimedia Signal Processing (MMSP)*, IEEE.

Yan, X., **Wright, D.**, Kumar, S., Lee, G., & Ozturk, Y. (2015). Enabling consumer behavior modification through real time energy pricing. In *Pervasive Computing and Communication Workshops (PerCom Workshops), 2015 IEEE International Conference on* (pp. 311-316). IEEE.

### Papers in Journals

**Wright, D.**, & Augenstein, I. (2025). Aggregating Soft Labels from Crowd Annotations Improves Uncertainty Estimation Under Distribution Shift. *PLOS ONE*. 2025.

**Wright, D.**, Igel, C., Samuel, G., & Selvan, R. (2025). Efficiency is Not Enough: A Critical Perspective of Environmentally Sustainable AI. *Communications of the ACM*. Association for Computing Machinery. 2025.

Holm, A. N., **Wright, D.**, & Augenstein, I. (2023). Revisiting Softmax for Uncertainty Approximation in Text Classification. *Information*, 14(7), 420.

Badal, V. D., **Wright, D.**, Katsis, Y., Kim, H. C., Swafford, A. D., Knight, R., & Hsu, C. N. (2019). Challenges in the construction of knowledge bases for human microbiome-disease associations. *Microbiome*, 7(1), 1-15.

Bhide, A., **Wright, D.**, & Ozturk, Y. (2016). Per-packet rate adaptation for wireless video. *Signal, Image and Video Processing*, 10(7), 1273-1278.

### Preprints

Mujahid, Z.M., **Wright, D.**, & Augenstein, I. (2025). Stress Testing Factual Consistency Metrics for Long-Document Summarization. arXiv preprint arXiv:2511.16622.

**Wright, D.**, Masud, S., Moore, J., Yadav, S., Antoniak, M., Park, C.Y. & Augenstein, I. (2025). Epistemic Diversity and Knowledge Collapse in Large Language Models. arXiv preprint arXiv:2510.04226.

Pei, J., **Wright, D.**, Augenstein, I., & Jurgens, D. (2025). Modeling Public Perceptions of Science in Media. arXiv preprint arXiv:2506.16622.

Ivey, J., Kumar, S., Liu, J., Shen, H., Rakshit, S., Raju, R., Zhang, H., Ananthasubramaniam, A., Kim, J., Yi, B.,  
**Wright, D.**, Israeli, A., Moeller, A., Zhang, L., & Jurgens, D. (2024). Real or Robotic? Assessing Whether LLMs  
Accurately Simulate Qualities of Human Responses in Dialogue. arXiv preprint arXiv:2409.08330.