



Tristan Thrush

Ph.D. Student in Computer Science, admitted Autumn 2023

Bio

BIO

I'm a Computer Science PhD student at Stanford in the NLP group and AI lab, supervised by Tatsunori Hashimoto and Christopher Potts. Previously, I was a founding member of the technical staff at Contextual AI (a startup working on retrieval augmented generation). Before that, I was a research engineer at Hugging Face. Before that, I was a research associate at Facebook AI Research, supervised by Douwe Kiela and then Adina Williams. And before that, I was a research associate at MIT Brain and Cognitive Sciences, supervised by Roger Levy. I Received my MEng in computer science with a concentration in artificial intelligence under Patrick Winston at the MIT Computer Science and Artificial Intelligence Lab. I received my BS also at MIT in computer science, with a minor in linguistics and a minor in math. While I was an undergrad, I did research with the Perception Systems Group at NASA's Jet Propulsion Lab.

I'm interested in AI. Specifically: natural language processing, computer vision, high-dimensional statistics, and data-centric AI methods. I have done several large-scale projects with a focus on the data side, which is so intertwined with the model side that it is sometimes hard to tell where one ends and the other begins.

Here are three of my favorite papers:

Perplexity Correlations: <https://arxiv.org/abs/2409.05816>

(This one has some fun math and is useful for pretraining data selection)

Multimodal Evaluation: <https://arxiv.org/abs/2204.03162>

(This one poses a still open challenge for word-order understanding in vision-language models)

Rover Relocalization for Mars Sample Return: <https://ieeexplore.ieee.org/abstract/document/9381709>

(There is nothing cooler than robots in space)