

Curt Tigges

ct@curttigges.com | [in/curttigges](https://in.linkedin.com/in/curttigges) | github.com/curt-tigges

PAPERS

Language Models Linearly Represent Sentiment (under review for ICML) | 2023 [\[arxiv\]](#)

Investigation of how LLMs build representations and world models of sentiment.

When Bigger Isn't Faster: Learning Speed Plateaus in LLMs (under review for ACL) | 2024

Short paper investigating task-dependent ceiling on learning speeds beyond which further increases to model scale do not result in faster capability acquisition

RELEVANT EXPERIENCE

Research Scientist | EleutherAI Institute | Jan 2023 - Present

- Conduct and assist in LLM mechanistic interpretability & other empirical research projects, including:
 - Circuits Over Time: forthcoming research paper on circuit formation and evolution in different model settings
 - Polygraph: generated dataset for investigating instrumental deception in fine-tuned and RLHF models
- Train LLMs on GPU cluster for various projects as needed, and maintain and improve the GPT-NeoX library

Data Scientist (part time) | NCSU Ops Research & Education Lab | Raleigh, NC | Jun 2022 - Nov 2022

- Built models to predict student population changes and to predict optimal locations for new schools

Senior Data Analyst | Taroko.io | Taipei, TW & Raleigh, NC | Jan 2020 - Dec 2022

- Built & deployed ML & statistical solutions for churn & revenue prediction, key conversion path analysis, etc.
- Planned & built out data warehouse in BigQuery, integrating PostgreSQL database, Heap Analytics & more

Business & Search Engine Marketing Analyst | Taroko.io | Taipei, TW | Jun 2016 - Dec 2019

- Analyzed customer data & produced SQL queries and reports to produce various metrics & ad-hoc analyses
- Developed bidding algorithm that rescued failing products, decreasing CPA by 21% and increasing volume

Software Engineer | KPIT Extended PLM | Raleigh, NC | Mar 2014 - Aug 2015

- Deployed enterprise software (PTC Windchill) customizations to product lifecycle management systems
- Planned & conducted detailed QA tests for software products

PROJECTS

Various Foundational Model Implementations (ARENA Program)

From-scratch implementations of various transformers (GPT-2, BERT); optimizers (SGD, Adam, LAMB); and RL algorithms (DQN, PPO).

DeepMind Perceiver Implementation

Implemented the DeepMind Perceiver architecture from scratch, trained on a variety of image datasets and authored a 3,000-word tutorial on how to build one. [\[Blog post\]](#), [\[Github\]](#)

Vision Transformer Implementation

Implemented vanilla vision transformer from scratch and wrote a detailed tutorial. [\[Blog post\]](#), [\[Github\]](#)

EDUCATION

SERI MATS (Neel Nanda Interpretability Stream) | 2023

Alignment Research Engineer Accelerator (ARENA) | 2022

Master of Computer Science (Data Science Track) | University of Illinois Urbana-Champaign | 2021

Bachelor of Science in Science, Technology and Society | NC State University | 2012

SKILLS

Programming Languages - Proficient In: Python | SQL - **Familiar With:** R | C++ | PHP

Packages: PyTorch | PyTorch Lightning | Transformers | Scikit-Learn | Matplotlib | Pandas | Numpy

Higher-Level Skills: Mechanistic Interpretability | Deep Learning | Fine Tuning | Prompt Engineering | Dataset Generation | Software Engineering