Shayne Longpre

Email: slongpre@media.mit.edu
Web: shaynelongpre.com

INTERESTS

Data-centric AI, large language models, & their societal impact. Evaluation, transparency, open source, systemic harms, & policies for responsible AI governance.

EDUCATION

Massachusetts Institute of Technology, Cambridge, Massachusetts — Sept 2021 - Present
Ph.D. Candidate, Media Arts & Sciences
Advisor: Prof. Sandy Pentland

Stanford University, Palo Alto, California — 2012 - 2018
M.S. in Computer Science, Artificial Intelligence
B.A. in Economics, minor in History
Advisors: Chris Manning and Danqi Chen

RESEARCH & INDUSTRY EXPERIENCE

Cohere 4 AI, Aya Open Science Initiative — March 2023 - Present, (Remote) US
Co-Leading Multilingual Instruction Tuning
Research Advisors: Sara Hooker

Google Brain, Reasoning Team — May 2022 - Sept 2022, (Remote) Canada
Google Student Researcher
Research Advisors: Jason Wei, Barret Zoph, Denny Zhou, & Adam Roberts

BigScience, BLOOM [10] & ROOTS [9] Teams — 2022, (Remote) US
Volunteer Contributor

Apple, Siri & Information Intelligence Team — Feb 2018 - June 2021, Seattle, Washington
Senior Applied Machine Learning Scientist
Research Advisor: Chris Dubois

Stanford NLP Group — 2016 - 2017, Palo Alto, California
Research Assistant
Research Advisors: Chris Manning & Danqi Chen

Salesforce AI Research, previously Metamind — June 2016 - Sept 2016, Palo Alto, California
Deep Learning Research Intern
Research Advisors: Richard Socher & Caiming Xiong

SELECT AI PUBLICATIONS

[1] The Data Provenance Initiative: A Large Scale Audit of Dataset Licensing & Attribution in AI
Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, Xinyi Wu, Enrico Shippole, Kurt Bollacker, Tongshuang Wu, Luis Villa, Sandy Pentland, Deb Roy, Sara Hooker
Preprint, 2023.

[2] The Foundation Model Transparency Index
Rishi Bommasani, Kevin Klyman, **Shayne Longpre**, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, Percy Liang
Preprint, 2023.

[3] Prometheus: Inducing Fine-grained Evaluation Capability in Language Models
Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, **Shayne Longpre**, Hwaran Lee, Sangdoo Yun,

Seongjin Shin, Sungdong Kim, James Thorne, Minjoon Seo
Preprint, 2023.

[4] OctoPack: Instruction Tuning Code Large Language Models
Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, **Shayne Longpre**
Preprint, 2023.

[5] Mixture-of-Experts Meets Instruction Tuning: A Winning Combination for Large Language Models
Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, **Shayne Longpre**, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuexin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, Denny Zhou
Preprint, 2023.

[6] A Pretrainer's Guide to Training Data: Measuring the Effects of Data Age, Domain Coverage, Quality, & Toxicity
Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, Daphne Ippolito
Preprint, 2023.

[7] The Flan Collection: Designing data and methods for effective instruction tuning
Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, Adam Roberts
ICML, 2023.

[8] Scaling instruction-finetuned language models
{Hyung Won Chung, Le Hou, **Shayne Longpre**}, ... Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, Jason Wei (35 authors)
Under Review, 2022.

[9] The Bigscience Roots Corpus: A 1.6 tb composite multilingual dataset
Hugo Laurençon... **Shayne Longpre**... Margaret Mitchell, Sasha Luccioni, Yacine Jernite (52 authors)
NeurIPS, 2022.

[10] BLOOM: A 176B-Parameter Open-Access Multilingual Language Model
Teven Le Scao, ... **Shayne Longpre**, ... Matteo Manica (128 authors)
ArXiv, 2022.

[11] Combining Compressions for Multiplicative Size Scaling on Natural Language Tasks
Rajiv Movva, Jinhao Lei, **Shayne Longpre**, Ajay Gupta, Chris DuBois
COLING, 2022.

[12] You reap what you sow: On the Challenges of Bias Evaluation Under Multilingual Settings
Zeerak Talat, Aurélie Névéol, Stella Biderman, Miruna Clinciu, Manan Dey, **Shayne Longpre**, Sasha Luccioni, Maraim Masoud, Margaret Mitchell, Dragomir Radev, Shanya Sharma, Arjun Subramonian, Jaesung Tae, Samson Tan, Deepak Tunuguntla, Oskar Van Der Wal
ACL, 2022. *BigScience Workshop*

[13] MIA 2022 Shared Task: Evaluating Cross-lingual Open-Retrieval Question Answering for 16 Diverse Languages
Akari Asai, **Shayne Longpre**, Jungo Kasai, Chia-Hsuan Lee, Rui Zhang, Junjie Hu, Ikuya Yamada, Jonathan H. Clark, Eunsol Choi
NAACL, 2022. *Multilingual Information Access Workshop*

[14] Active Learning Over Multiple Domains in Natural Language Tasks
Shayne Longpre, Julia Reisler, Edward Greg Huang, Yi Lu, Andrew Frank, Nikhil Ramesh, Chris DuBois
NeurIPS, 2022. *Workshop on Distribution Shifts*

[15] Entity-Based Knowledge Conflicts in Question Answering
{**Shayne Longpre**, Kartik Perisetla, Anthony Chen}, Nikhil Ramesh, Chris DuBois, Sameer Singh
EMNLP, 2021.

[16] MKQA: A Linguistically Diverse Benchmark for Multilingual Open Domain Question Answering
Shayne Longpre, Yi Lu, Joachim Daiber
TACL, 2021.

[17] Evaluating Entity Disambiguation and the Role of Popularity in Retrieval-Based NLP
Anthony Chen, Pallavi Gudipati, **Shayne Longpre**, Xiao Ling, Sameer Singh
ACL, 2021.

[18] Evaluating Question Rewriting for Conversational Question Answering
Svitlana Vakulenko, **Shayne Longpre**, Zhucheng Tu, Raviteja Anantha
WSDM, 2021.

[19] Leveraging Query Resolution and Reading Comprehension for Conversational Passage Retrieval
Svitlana Vakulenko, Nikos Voskarides, Zhucheng Tu, **Shayne Longpre**
TREC CAsT, 2021.

[20] A Comparison of Question Rewriting Methods for Conversational Passage Retrieval
Svitlana Vakulenko, Nikos Voskarides, Zhucheng Tu, **Shayne Longpre**
ECIR, 2021.

[21] Open-Domain Question Answering Goes Conversational via Question Rewriting
Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, **Shayne Longpre**
NAACL, 2021.

[22] Pivot Through English: Reliably Answering Multilingual Questions without Document Retrieval
Ivan Montero, **Shayne Longpre**, Ni Lao, Andrew Frank, Christopher DuBois
NAACL, 2021. *Multilingual Information Access Workshop*

[23] On the Transferability of Minimal Prediction Preserving Inputs in Question Answering
Shayne Longpre, Yi Lu, Chris DuBois
NAACL, 2021.

[24] A Wrong Answer or a Wrong Question? An Intricate Relationship between Question Reformulation and Answer Selection in Conversational Question Answering
Svitlana Vakulenko, **Shayne Longpre**, Zhucheng Tu, Raviteja Anantha
EMNLP, 2020. *Search-Oriented Conversational AI Workshop* **Best Paper Award**.

[25] How Effective is Task-Agnostic Data Augmentation for Pretrained Transformers?
Shayne Longpre, Yu Wang, Chris DuBois
EMNLP Findings, 2020.

[26] An Exploration of Data Augmentation and Sampling Techniques for Domain-Agnostic Question Answering
Shayne Longpre, Yi Lu, Zhucheng Tu, Chris DuBois
EMNLP, 2019. *Machine Reading for Question Answering Workshop*

WRITINGS

[27] Lethal autonomous weapons systems & artificial intelligence: Trends, challenges, and policies
Shayne Longpre, Marcus Storm, Rishi Shah
MIT Science Policy Review, Volume III, 2022.

[28] Invigorating Competition in Social Networking: An Interoperability Remedy
Cristian Santesteban, **Shayne Longpre**
Competition Policy International, 2021.

[29] How Big Data Confers Market Power to Big Tech: Leveraging the Perspective of Data Science

Cristian Santesteban, **Shayne Longpre**
The Antitrust Bulletin, 2020.

AWARDS

MIT Generative AI Impact Award, 2023 — 2023
Awarded for the Data Provenance Initiative
Accompanied by a research funding grant of $70,000.
ICLR Highlighted Reviewer (3%) — 2022
EMNLP 2020, *Search-Oriented Conversational AI Workshop* **Best Paper Award** — 2020
EMNLP 2019, *MRQA Workshop Shared Task* **2nd place** — 2019
TREC 2019, *Conversational Assistance Track (CAsT) Shared Task* **1st place ("A Team")** — 2019

TEACHING EXPERIENCE

Instructor, MAS.S68 Generative AI: Evaluation and New Research Methods, MIT — 2023
Instructed research seminar on large language models, and the landscape of socio-political concerns with their adoption.
Instructor, AI4ALL, Stanford University — 2017
Instructed course on Computer Vision fundamentals to young women in STEM.
Teaching Assistant, Natural Language Processing w/ Deep Learning (CS224N), Stanford — 2017
Teaching Assistant, Computer Vision with Deep Learning (CS231N), Stanford — 2017

SERVICE

Leadership & Organization

Instruction Following & Finetuning Workshop (ITIF), NeurIPS 2023 — 2023
Workshop Organizer
Cohere 4 AI Aya Initiative Co-Lead — 2023 - Present
MozFest Facilitator: Bringing Light to Shadow Data — 2023
Workshop Co-Lead Organizer: Defining Transparency Workshop, Brown University — 2022
(*hosted w/ the Algorithmic Transparency Institute and Brown's Information Futures Lab*)
Multilingual Information Access Workshop (MIA), NAACL 2022 — 2022
Workshop Organizer
Shared-task Organizer: MIA 2022 Shared task, NAACL 2022 — 2022

Academic Service

Area Chair (2023): EMNLP, *Large Language Models and the Future of NLP Track*
ACL Professional Conduct Committee — 2021 - 2023
Trained volunteer, responding to complaints from ACL's Anti-Harassment Policy.
Reviewer (2023): ARR, NeurIPS, FAccT, ICML, ICLR, EMNLP
Reviewer (2022): ARR, NeurIPS, ICLR, EMNLP, various workshops

ADVISING

Volunteer at MIT **Students Offering Support Program** — 2022
Advising underrepresented students in MIT graduate applications.

Rajiv Movva, MIT CS → Now Cornell Tech CS PhD student. Published [11]. — 2021 - 2022
Xuhui Zhou, UW MS → Now LTI-CMU CS PhD student. — 2021 - 2022
Erik Jones, Stanford MS → Now UC Berkeley CS PhD student. — 2021
Ivan Montero, UW CS MS student → Now Apple Machine Learning. Published [22]. — 2020 - 2021

INVITED TALKS & PANELS

Mosaic ML Invited Talk. *A Pretrainer's Guide* (Host: Jonathan Frankle) — 2023
Harvard & MIT: Policymaking for AI & Web3 Series Invited Talk & Panel. *A Primer on LLMs for Regulators* (Host: Getting Plurality Research Network) — 2023

University of Washington Invited Talk. *A Pretrainer's Guide* (Hosts: Akari Asai, Sewon Min) 2023
Allen Institute of AI (AI2) Invited Talk. *A Pretrainer's Guide* (Host: Maria Antoniak) 2023
Instituto Superior Técnico Seminar Series *Effective Instruction Tuning* (Host: Nuno Guerreiro) 2023
Amazon Data-centric AI Seminar Series. *Effective Instruction Tuning* (Host: Li Lihong) 2023
Databricks Seminar Series. *Effective Instruction Tuning* (Host: Mike Conover) 2023
Apple Applied ML Reading Group. *Effective Instruction Tuning* (Host: Michael Tu) 2023
MozFest Facilitator. *Bringing Light to Shadow Data* 2023
Kailua Labs AI Seminar Series. *Effective Instruction Tuning* (Host: Pablo Mendes) 2023
Oracle ML Seminar Series. *Effective Instruction Tuning* (Host: Ari Kobren) 2023
Google Research *Effective Instruction Tuning* (Host: Denny Zhou) 2022
Truth & Trust Online 2022 *The Data Access and Transparency (DATA) Index: Evaluating Transparency in Online Social Platforms.* Full Talk 2022
Multilingual Resources panel moderator at the NAACL 2022, *MIA Workshop* 2022
Panelists: Graham Neubig, Holger Schwenk and Alice Oh.
UC Irvine Reading Group. *Knowledge Conflicts in Question Answering* (Host: Sameer Singh) 2021
Question Answering Evaluation panel moderator at EMNLP 2021, *SCAI Workshop* 2021

PRESS & MEDIA

Select Press for Data Provenance Initiative [1] 2023

The Washington Post. *AI researchers uncover ethical, legal risks to using popular data sets (quoted)*

VentureBeat. *MIT, Cohere for AI, others launch platform to track and filter audited AI datasets (quoted)*

TechCircle. *MIT, Cohere for AI, others launch platform to enhance transparency in AI data*

Select Press for Foundation Model Transparency Index [2] 2023

Stanford HAI Blog. *Introducing The Foundation Model Transparency Index (quoted)*
The New York Times. *Maybe we will finally learn more about how AI works*
The Atlantic. *We Don't Actually Know If AI Is Taking Over Everything*
Bloomberg. *Klobuchar Says AI Regulation Still Possible Before End of Year*
The Information. *How Transparent is your model?*
VentureBeat. *How transparent are AI models? Stanford researchers found out.*
The Verge. *The world's biggest AI models aren't very transparent, Stanford study says*
Reuters. *Stanford researchers issue AI transparency report, urge tech companies to reveal more*
Fast Company. *Why everyone seems to disagree on how to define Artificial General Intelligence*

Hacker News Our course on Generative AI trending 2023
Google AI Blog The Flan Collection: Advancing open source methods for instruction tuning 2023
PaLM 2 Technical Report Flan Collection & Methods cited several times as key components. 2023
DAIR.AI Top ML Papers of the Week 2023

LAST UPDATED

October 2023.