

# Sara Kangaslahti

✉ sarakangaslahti@g.harvard.edu

🌐 skangasl.github.io

📄 Sara Kangaslahti

in linkedin.com/in/sara-kangaslahti

🐙 github.com/skangasl

## Education

---

### Harvard University

Doctor of Philosophy Candidate in Computer Science

Cambridge, MA

Expected Graduation: May 2028

### Harvard University

Master of Science in Computer Science

Cambridge, MA

May 2026

### California Institute of Technology

Bachelor of Science in Computer Science

Minor in Information and Data Sciences

GPA: 4.2/4.3

Pasadena, CA

June 2023

## Research Experience

---

### Graduate Research Fellow

School of Engineering and Applied Sciences, Harvard University

Advised by Professor David Alvarez-Melis

Cambridge, MA

August 2023 - Present

### Undergraduate Researcher

Department of Computing & Mathematical Sciences, California Institute of Technology

Advised by Professor Anima Anandkumar and Professor R. Michael Alvarez

Pasadena, CA

October 2019 - June 2023

**Research Interests:** science of large language models, model representation geometry, model interpolation, knowledge distillation, training dynamics, interpretability, model pruning, model merging

## Publications

---

### Inverse Depth Scaling From Most Layers Being Similar

International Conference on Machine Learning, 2026 [\[link\]](#)

Yizhou Liu, [Sara Kangaslahti](#), Ziming Liu, Jeff Gore.

### Boomerang Distillation Enables Zero-Shot Model Size Interpolation

International Conference on Learning Representations, 2026 [\[link\]](#)

Also presented at the NeurIPS UniReps Workshop, 2025.

[Sara Kangaslahti](#), Nihal V. Nayak\*, Jonathan Geuter\*, Marco Fumero, Francesco Locatello, David Alvarez-Melis.

### Hidden Breakthroughs in Language Model Training

International Conference on Learning Representations, 2026 [\[link\]](#)

Also presented as **Spotlight** at the ICML Workshop on Mechanistic Interpretability, 2024.

[Sara Kangaslahti](#), Elan Rosenfeld, Naomi Saphra.

### Activating Identity and Political Action in the #MeToo Era

Computational Communication Research, 2026 [\[link\]](#)

Melina Much, Daniel Ebanks, [Sara Kangaslahti](#), Jean Kossaifi, R. Michael Alvarez, Anqi Liu, Anima Anandkumar.

### GenAI for Systems: Recurring Challenges and Design Principles from Software to Silicon

Preprint, 2026 [\[link\]](#)

Arya Tschand, Chenyu Wang, Zishen Wan, Andrew Cheng, Ioana Cristescu, Kevin He, Howard Huang, Alexander Ingare, Akseli Kangaslahti, [Sara Kangaslahti](#), Theo Lebryk, Hongjin Lin, Jeffrey Jian Ma, Alexandru Meterez, Clara Mohri, Depen Morwani, Sunny Qin, Roy Rinberg, Paula Rodriguez-Diaz, Alyssa Mia Taliotis, Pernille Undrum Fathi, Rosie Zhao, Todd Zhou, Vijay Janapa Reddi.

## Continuous Language Model Interpolation for Dynamic and Controllable Text Generation

*Transactions on Machine Learning Research, 2025* [\[link\]](#)

Also presented at the NeurIPS Workshop on Adaptive Foundation Models, 2024.

[Sara Kangaslahti](#), David Alvarez-Melis.

## Topic models estimators are statistically inconsistent and substantively misleading: evidence and solutions

*Preprint, 2025* [\[link\]](#)

Danny Ebanks, [Sara Kangaslahti](#).

## Analyzing Political Text at Scale with Online Tensor LDA

*Political Analysis, 2025* [\[link\]](#)

[Sara Kangaslahti](#), Danny Ebanks, Jean Kossaifi, R. Michael Alvarez, Anima Anandkumar.

## Fast Forwarding Low-Rank Training

*Empirical Methods in Natural Language Processing, 2024* [\[link\]](#)

Adir Rahamim, Naomi Saphra, [Sara Kangaslahti](#), Yonatan Belinkov.

## Can you label less by using out-of-domain data? Active & transfer learning with few-shot instructions

*NeurIPS Transfer Learning for Natural Language Processing Workshop, 2023* [\[link\]](#)

Rafal Kocielnik\*, [Sara Kangaslahti](#)\*, Shrimai Prabhumoye, Meena Hari, Michael Alvarez, Anima Anandkumar.

## Adaptive Sampling: Algorithmic vs. Human Waypoint Selection

*Preprint, 2021* [\[link\]](#)

Stephanie Kemna, [Sara Kangaslahti](#), Oliver Kroemer, Gaurav S. Sukhatme.

## Calibrating User Bias In Sentiment Estimation From Social Media

*IEEE International Conference on Big Data (Big Data) Workshop on Big Social Media Data Management and Analysis (BSMDMA), 2018* [\[link\]](#)

[Sara Kangaslahti](#), Anoop Kumar.

\* Equal contribution.

## Fellowships and Awards

---

### NSF Graduate Research Fellow

August 2023 - Present

National Science Foundation

### Doris S. Perpall SURF Speaking Competition Finalist

2022

California Institute of Technology

Top 3% undergraduate research presentation

### Carolyn Ash Summer Undergraduate Research Fellow (SURF)

June 2022 - September 2022

California Institute of Technology

### Jack E. Froehlich Memorial Award Nominee

2022

California Institute of Technology

Top 5% GPA in Caltech junior class

### Beckman Leadership Award Nominee

2022, 2023

California Institute of Technology

### Summer Undergraduate Research Fellow (SURF)

June 2020 - August 2020

California Institute of Technology

## Industry Experience

---

### Research Intern

March 2022 - June 2023

NVIDIA Research

Mentored by Professor Anima Anandkumar

**Software Engineer Intern**  
AI Integrity Team, Facebook  
Mentored by Dr. Liang Tan

June 2021 - August 2021

## Teaching Experience

---

**Teaching Fellow** September 2024 - December 2024  
School of Engineering and Applied Sciences, Harvard University  
*Course:* Computational Science and Engineering Capstone Project, taught by Dr. Weiwei Pan

**Teaching Assistant** September 2022 - December 2022  
Department of Computing & Mathematical Sciences, California Institute of Technology  
*Course:* Distributed Systems, taught by Professor Kianthra M. Chandy

## Talks

---

<b>Boomerang Distillation Enables Zero-Shot Model Size Interpolation</b> <i>Invited Talk at ELLIS x UniReps Seminar Series</i> <a href="#">[link]</a>	<b>Remote</b> February 2026
<b>Boomerang Distillation Enables Zero-Shot Model Size Interpolation</b> <i>Invited Talk at MIT PLS AI Journal Club</i>	<b>Cambridge, MA</b> November 2025
<b>Hidden Breakthroughs in Language Model Training</b> <i>Invited Talk at MIT PLS AI Journal Club</i>	<b>Cambridge, MA</b> October 2025
<b>Analyzing Social Media Conversations at Scale with Tensor LDA</b> <i>Invited Talk at University of California Los Angeles (UCLA) Biostats Seminar</i>	<b>Los Angeles, CA</b> April 2022

## Reviewing

---

Reviewer for ICLR (2025, 2026), NeurIPS (2026), JMLR (2022)

## Outreach

---

- **Volunteer Math and Science Tutor**, Caltech Rise Program, September 2022 - June 2023
- **Volunteer Math and Science Tutor**, Youth of Promise Program, June 2022 - July 2023
- **Volunteer Math Tutor**, NCNW Young Legends Program, January 2022 - May 2022

## Language Skills

---

English (first language), Finnish (fluent), Spanish (excellent), Japanese (beginning)