Kiran Vodrahalli

https://kiranvodrahalli.github.io
kirannv@google.com

Summary

I am a research scientist at Google DeepMind, where I study long-context sequence modeling (e.g. LLMs) and strategic and interactive machine learning. I am currently particularly interested in resource-efficient approaches to long-context models. I am also very interested in questions at the intersection of multi-agent learning, strategic learning in games, incentives for data collection and collaborative machine learning, and algorithmic game theory.

Education

Columbia University 2017 — 2022
- Ph.D. in Computer Science
- Advisors: Professor Daniel Hsu, Professor Alex Andoni
- Thesis: *Resource-Efficient Methods in Machine Learning*

Princeton University 2012 — 2017
- A.B. Mathematics (cum laude), M.S.E. Computer Science
- Advisors: Professor Sanjeev Arora, Professor Ken Norman
- Master's Thesis: *Temporally Dependent Mappings between fMRI Responses and Natural Language Descriptions of Natural Stimuli*

Industry Research

Research Scientist, **Google DeepMind** 2023 — now
- Core contributor on Gemini: resource-efficient models (Gemini 1.5 Flash and more), leading development of novel long-context evals.
- Research scientist working on long-context sequence modeling, resource-efficient and principled LLMs, and questions at the intersection of multi-agent learning and interactive, strategic learning in games.

Research Scientist, **Google Brain** 2022 — 2023
- Core contributor on Google Bard for both modeling and evaluation.
- Core contributor on the Long Context workstream for Google's PaLM 2 model.

Student Researcher, **Google Brain** 2021 — 2022
- Part-time employment as a student researcher.

Research Intern, **Google Brain** 2021
- Research internship on training deep neural networks with resource constraints.

Publications[1]

Preprints

[18] *Michelangelo: Long Context Evaluations Beyond Haystacks via Latent Structure Queries.*
Kiran Vodrahalli, Santiago Ontanon, Nilesh Tripuraneni, Kelvin Xu, Sanil Jain, Rakesh Shivanna, Jeffrey Hui, Nishanth Dikkala, Mehran Kazemi, Bahare Fatemi, Rohan Anil, Ethan Dyer, Siamak Shakeri, Roopali Vij, Harsh Mehta, Vinay Ramasesh, Quoc Le, Ed Chi, Yifeng Lu, Orhan Firat, Angeliki Lazaridou, Jean-Baptiste Lespiau, Nithya Attaluri, Kate Olszewska.
Corresponding Author.

[1]Note that * indicates equal contribution. In theory publications, the citation order is alphabetical by last name.

[17] *Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context*. Gemini Team, Google. **Core Contributor**.

[16] *Gemini: A Family of Highly Capable Multimodal Models*.
Gemini Team, Google. **Core Contributor**.

[15] *PaLM 2 Technical Report*.
Google. **Core Contributor to Long Context workstream**.

[14] *Nonlinear Initialization Methods for Low-Rank Neural Networks*.
Kiran Vodrahalli, Rakesh Shivanna, Maheswaran Sathiamoorthy, Sagar Jain, Ed H. Chi.

Conference Proceedings

[13] *Online Learning with Bounded Recall*.
Jon Schneider*, **Kiran Vodrahalli***.
Poster. ICML, July 2024.

[12] *Is Learning in Games Good for the Learners?*.
William Brown, Jon Schneider, **Kiran Vodrahalli**.
Spotlight. NeurIPS, December 2023.

[11] *The Platform Design Problem*.
Christos Papadimitriou*, **Kiran Vodrahalli***, Mihalis Yannakakis*.
Oral Presentation. Conference on Web and Internet Economics, December 2021.
Spotlight Oral Presentation (top 10%). StratML Workshop, NeurIPS 2021.
Oral Presentation at NetEcon Workshop, EC 2021. Poster at EC 2021.

[10] *The Logical Options Framework*.
Brandon Araki, Xiao Li, **Kiran Vodrahalli**, Jonathan DeCastro, J. Micah Fry, Daniela Rus.
(Long) Oral Presentation and Poster. ICML, July 2021.

[9] *Deep Bayesian Nonparametric Learning of Rules and Plans from Demonstrations with a Learned Automaton Prior*.
Brandon Araki, **Kiran Vodrahalli**, Thomas Leech, Cristian Ioan Vasile, Mark Donahue, Daniela Rus.
Spotlight Presentation. AAAI Conference on Artificial Intelligence, February 2020.

[8] *Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform*.
Mathias Lécuyer, Riley Spahn, **Kiran Vodrahalli**, Roxana Geambasu, Daniel Hsu.
Oral Presentation. Symposium on Operation Systems Principles, October 2019.

[7] *Learning to Plan with Logical Automata*.
Brandon Araki*, **Kiran Vodrahalli***, Thomas Leech, Cristian Ioan Vasile, Mark Donahue, Daniela Rus.
Spotlight Presentation and Poster. Robotics: Science and Systems, June 2019.
Spotlight Oral Presentation at NeurIPS 2018 Infer2Control Workshop.

[6] *Attribute-Efficient Learning of Monomials over Highly-Correlated Variables.*
Alex Andoni*, Rishabh Dudeja*, Daniel Hsu*, **Kiran Vodrahalli***.
Oral Presentation. Algorithmic Learning Theory, March 2019.

[5] *A Large Self-Annotated Corpus for Sarcasm.*
Mikhail Khodak, Nikunj Saunshi, **Kiran Vodrahalli**.
Poster. Language Resources and Evaluation, May 2018.

[4] *A Compressed Sensing View of Unsupervised Text Embeddings, Bag-of-n-Grams, and LSTMs.*
Sanjeev Arora*, Mikhail Khodak*, Nikunj Saunshi*, **Kiran Vodrahalli***.
Poster. International Conference on Learning Representations, April 2018.
Oral Presentation. ICML 2018 Workshop on Theory of Deep Learning.
Poster at ACL 2018 Workshop on Representation Learning for NLP.

[3] *A Temporal Decay Model for Mapping between fMRI and Natural Language Annotations.*
Kiran Vodrahalli, Cathy Chen, Viola Mocz, Christopher Baldassano, Uri Hasson, Sanjeev Arora, Kenneth A. Norman.
Poster. Cognitive Computational Neuroscience, September 2017.

Journal Publications

[2] *Learning and Planning with Logical Automata.*
Brandon Araki, **Kiran Vodrahalli**, Thomas Leech, Cristian-Ioan Vasile, Mark Donahue, Daniela Rus.
Autonomous Robots, August 2021.

[1] *Mapping between fMRI Responses to Movies and their Natural Language Annotations.*
Kiran Vodrahalli, Po-Hsuan Chen, Yingyu Liang, Christopher Baldassano, Janice Chen, Christopher Honey, Uri Hasson, Peter Ramadge, Kenneth A. Norman, Sanjeev Arora.
Neuroimage, June 2017.
Oral Presentation at NeurIPS 2016 Workshop on Representation Learning in Artificial and Biological Networks.
Oral Presentation at ICML 2016 Workshop on Multi-View Representation Learning.

Invited Talks

Meta Research	May 2022
Google Brain AutoML	May 2022
Berkeley Center for Human-Compatible AI (CHAI) Seminar	May 2022
Google Research NYC	April 2022
Google Brain Neural Modeling Group	February 2022
Simons Flatiron Center for Computational Neuroscience	February 2022
Amazon AWS	February 2022
Simons Theory of Computing, Learning in Games Program, Equilibrium Computation and ML Reading Group	February 2022
Simons Flatiron Center for Computational Mathematics	January 2022
Google Algorithms Seminar	November 2021
Google Learning Theory Group	October 2021
Google Brain	August 2021
NY Academy of Sciences Machine Learning Symposium	March 2020
Yahoo Research	August 2019

NY Academy of Sciences Machine Learning Symposium — March 2019
Princeton Neuroscience Institute — September 2017

Awards

Spotlight Prize at NYAS Annual ML Symposium — 2019, 2020
- Top 10% of posters chosen to give a spotlight presentation.

NSF Graduate Research Fellowship Award — 2016
- Awarded for Computer Science in the subfield Machine Learning.

Teaching

Columbia University
Teaching Assistant, Computation and the Brain (graduate) — Fall 2018

Princeton University
Teaching Assistant, Theoretical Machine Learning (graduate) — Spring 2017
Teaching Assistant, NLP Independent Work Seminar — Fall 2016
Grader, Introductory Algorithms — Spring 2014
Lab Teaching Assistant, Introductory Algorithms and Systems — Fall 2013 — Fall 2014

Mentorship

Advisor, Columbia Undergraduate Theory Seminar — Summer 2021
- Designed the seminar syllabus for the Algorithmic Game Theory Undergraduate Theory Seminar, and ran the seminar.

Advisor, Princeton Junior Independent Work — Spring 2017
- Jointly advised the junior independent work of Cathy Chen (COS'18) with Professor Ken Norman in neuroscience and computer science.

Advisor, Princeton Junior Independent Work — Spring 2017
- Jointly advised the junior independent work of Viola Mocz (NEU'18) with Professor Ken Norman in neuroscience and computer science.

Service

Program Committee
Neural Information Processing Systems (NeurIPS) — 2020 — 2023
International Conference on Machine Learning (ICML) — 2020 — 2023
- Top 33% Reviewer in 2020, Expert Reviewer starting in 2021.

Transactions of Machine Learning Research (TMLR)
Nature Communications — 2023
International Conference on Learning Representations (ICLR) — 2021 — 2023
Symposium on Discrete Algorithms (SODA) — 2023
NeuroImage — 2017

University and Department Service
Pre-Submission Application Review (PAR), Columbia CS Department — 2020
Colloqium Organizer, Columbia CS Department — 2018 — 2019
NLP-ML Reading Group Organizer, Princeton CS Department — 2014 — 2016
- Organized the NLP-ML Reading Group with Dr. Christiane Fellbaum.