Suriya Gunasekar

suriya@ttic.edu or suriyag@microsoft.com
http://sgunasekar.github.io
July 8, 2025

Employment

2019–Present **Principal Research Manager.**
Microsoft Research, Redmond, USA.
Joined Microsoft as Senior Researcher.
Since March 2023, I have primarily focused on the pretraining of the Phi-family of language models at Microsoft. As a core contributor and a founding member of the team, I have led key initiatives, including:
- Developing and evaluating strategies for pretraining data mixtures, data validation, and data curricula, optimizing across diverse data sources with varying levels of quality, diversity, and coverage.
- Executing the pretraining runs for Phi-1 and Phi-1.5.
- Leading the pretraining efforts for Phi-3 (mini & medium) and Phi-4 models.
- Technical lead for data and SFT for Phi-4-reasoning posttraining.

2016–2019 **Research Assistant Professor.**
Toyota Technological Institute at Chicago, USA.

Education

2010–2016 **MS—PhD in Electrical and Computer Engineering.**
The University of Texas at Austin, USA. Advisor: Prof. Joydeep Ghosh.
PhD Thesis: Mining Structured Matrices in High Dimensions.

2006–2010 **B. Tech in Electronics and Communication Engineering.**
National Institute of Technology, Warangal, India.

Mentorship and Teachings

2024-Present **Manager of PhiScale Team at Microsoft.**

2019-2023 **(Co)-mentorship of interns.** Meena Jagadeesan, Omar Montassar, Ruoqi Shen, Ananya Kumar, Mathieu Even, Sadhika Malladi.

2018 **Senior Research Personnel** National Science Foundation award AF: RI: Medium: Collaborative Research: *Understanding and Improving Optimization for Deep and Recurrent Networks.*

2018 **Co-instructor.** *Introduction to Machine Learning Summer School* jointly organized by TTI Chicago and the University of Chicago as part of an NSF Research and Training Grant.
- Two week long ML course. Teaching material at http://www.ttic.edu/intromlss2018.

2017–2018 **Guest Lectures at TTI Japan** on *Introduction to Machine Learning.*
- Topics: Classification, AdaBoost, generative models, mixture models, Expectation-Maximization (EM).

2016–2017 **Guest lectures at TTI Chicago** on *Introduction to Statistical Machine Learning.*
- Topics: Regularization, gradient descent, logistic regression, support vector machines, kernels.

Publications

Technical Reports & Preprints

⋆ = alphabetical author ordering

- Phi-4-reasoning: Technical report. Marah Abdin, Sahaj Agarwal, Ahmed Awadallah, Vidhisha Balachandran, Harkirat Behl, Lingjiao Chen, Gustavo de Rosa, **Suriya Gunasekar**⋆, Mojan Javaheripi, Neel Joshi, Piero Kauffmann, Yahs Lara, Caio C. T. Mendes, Arindam Mitra, Besmira Nushi, Dimitris Papailiopoulos, Olli Saarikivi, Shital Shah, Vaishnavi Shrivastava, Vibhav Vineet, Yue Wu, Safoora Yousefi, Guoqing Zheng (2025).
- Phi-4: Technical report. Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, **Suriya Gunasekar**⋆, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, Yi Zhang (2024).
- Phi-3 technical report: A highly capable language model locally on your phone. GenAI team at Microsoft (2024).
- Phi-2: The surprising power of small language models. Blog Post. Phi team at Microsoft (2023).
- Textbooks Are All You Need II: phi-1.5 technical report. Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, **Suriya Gunasekar**⋆, Yin Tat Lee (2023).
- Textbooks Are All You Need. **Suriya Gunasekar**, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendez, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam T. Kalai, Yin Tat Lee, Yuanzhi Li (2023).
- Unveiling Transformers with LEGO: a synthetic reasoning task. Yi Zhang, Arturs Backurs, Sébastien Bubeck, Ronen Eldan, **Suriya Gunasekar**, Tal Wagner (2022).
- Generalization to translation shifts: a study in architectures and augmentations. **Suriya Gunasekar** (2022).

Peer-reviewed full length conference publications

1. KITAB: Evaluating LLMs on Constraint Satisfaction for Information Retrieval. Marah I Abdin, **Suriya Gunasekar**, Varun Chandrasekaran, Jerry Li, Mert Yuksekgonul, Rahee Ghosh Peshawaria, Ranjita Naik, Besmira Nushi. *In* 12*th International Conference on Learning Representations (ICLR) 2024.*
2. Attention Satisfies: A Constraint-Satisfaction Lens on Factual Errors of Language Models. Mert Yuksekgonul, Varun Chandrasekaran, Erik Jones, **Suriya Gunasekar**, Ranjita Naik, Hamid Palangi, Ece Kamar, Besmira Nushi. *In* 12*th International Conference on Learning Representations (ICLR) 2024.*
3. How to Fine-Tune Vision Models with SGD. Ananya Kumar, Ruoqi Shen, Sebastien Bubeck, **Suriya Gunasekar**. *In* 12*th International Conference on Learning Representations (ICLR) 2024.*
4. (S) GD over Diagonal Linear Networks: Implicit Regularisation, Large Stepsizes and Edge of Stability. Mathew Even, Scott Pesme, **Suriya Gunasekar**, Nicholas Flammarion. *In* 37*th Conference on Neural Information Processing Systems (NeurIPS) 2023.*
5. Neural-Sim: Learning to Generate Training Data with NeRF. Yunhao Ge, Harkirat Behl, Jiashu Xu, **Suriya Gunasekar**, Neel Joshi, Yale Song, Xin Wang, Laurent Itti, Vibhav Vineet. *In* 17*th European Conference on Computer Vision (ECCV) 2022.*
6. Data Augmentation as Feature Manipulation. Ruoqi Shen, Sébastien Bubeck, **Suriya Gunasekar**. *In* 39*th International Conference on Machine Learning (ICML) 2022.*

7. Inductive bias of multi-channel linear convolutional networks with bounded weight norm. Meena Jagadeesan, Ilya Razenshteyn, **Suriya Gunasekar**. *In 35th Conference on Learning Theory (COLT) 2022.*
8. Methods and Analysis of The First Competition in Predicting Generalization of Deep Learning. Y. Jiang, P. Natekar, M. Sharma, S. Aithal, D. Kashyap, N. Subramanyam, C. Lassance, D. Roy, G. K. Dziugaite, **S. Gunasekar**, I. Guyon, P. Foret, S. Yak, H. Mobahi, B. Neyshabur, S. Bengio. *In NeurIPS 2020 Competition and Demonstration Track 2021.*
9. Mirrorless mirror descent: A natural derivation of mirror descent. **Suriya Gunasekar**, Blake Woodworth, Nathan Srebro. *In 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021.*
10. Implicit bias in deep linear classification: Initialization scale vs training accuracy. Edward Moroshko, Blake Woodworth, **Suriya Gunasekar**, Jason D. Lee, Nathan Srebro, Daniel Soudry. *In 34th Conference on Neural Information Processing Systems (NeurIPS) 2020.*
11. Implicit regularization and convergence for weight normalization. Xiaoxia Wu, Edgar Dobriban, Tongzheng Ren, Shanshan Wu, Zhiyuan Li, **Suriya Gunasekar**, Rachel Ward, Qiang Liu. *In 34th Conference on Neural Information Processing Systems (NeurIPS) 2020.*
12. Kernel and Deep Regimes in Overparametrized Models. Blake Woodworth, **Suriya Gunasekar**, Pedro Savarese, Edward Moroshko, Itay Golan, Jason D. Lee, Daniel Soudry, Nathan Srebro. *In 33rd Conference on Learning Theory (COLT) 2020.*
13. Lexicographic and Depth-Sensitive Margins in Homogeneous and Non-Homogeneous Deep Models. Mor Shpigel Nacson, **Suriya Gunasekar**, Jason D. Lee, Nathan Srebro, Daniel Soudry. *In 36th International Conference on Machine Learning (ICML) 2019.*
14. Convergence of gradient descent on separable data. Mor Shpigel Nacson, Jason D. Lee, **Suriya Gunasekar**, Pedro Savarese, Nathan Srebro, Daniel Soudry. *In 22nd International Conference on Artifical Intelligence and Statistics (AISTATS) 2019.*
15. Implicit bias of gradient descent on linear convolutional networks. **Suriya Gunasekar**, Jason D. Lee, Daniel Soudry, Nathan Srebro. *In 32nd Conference on Neural Information Processing Systems (NeurIPS) 2018.*
16. On preserving non–discrimination when combining expert advice. Avrim Blum, **Suriya Gunasekar**⋆, Thodoris Lykouris, Nathan Srebro. *In 32nd Conference on Neural Information Processing Systems (NeurIPS) 2018.*
17. Characterizing implicit bias in terms of optimization geometry. **Suriya Gunasekar**, Jason D. Lee, Daniel Soudry, Nathan Srebro. *In 35th International Conference on Machine Learning (ICML) 2018.*
18. Implicit regularization in matrix factorization. **Suriya Gunasekar**, Blake Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, Nathan Srebro. *In 31st Conference on Neural Information Processing Systems (NIPS) 2017.* ***(spotlight presentation)****.*
19. Learning non–discriminatory predictors. Blake Woodworth, **Suriya Gunasekar**, Mesrob Ohannessian, Nathan Srebro. *In 30th Conference on Learning Theory (COLT) 2017.*
20. Preference completion from partial rankings. **Suriya Gunasekar**, Oluwasanmi Koyejo, Joydeep Ghosh. *In 30th Conference on Neural Information Processing Systems (NIPS) 2016.*
21. Identifiable phenotyping using constrained non–negative matrix factorization. Shalmali Joshi, **Suriya Gunasekar**, David Sontag, Joydeep Ghosh. *In 1st Machine Learning for Healthcare Conference (MLHC) 2016.*
22. Unified view of matrix completion under general structural constraints. **Suriya Gunasekar**, Arindam Banerjee, Joydeep Ghosh. *In 29th Conference on Neural Information Processing Systems (NIPS) 2015.*

23. Consistent collective matrix completion under joint low rank structure. **Suriya Gunasekar**, Makoto Yamada, Dawel Yin, Yi Chang. *In 18th International Conference on Artificial Intelligence and Statistics 2015.*
24. Face detection on distorted images using perceptual quality–aware features. **Suriya Gunasekar**, Joydeep Ghosh, Alan Bovik. *In 14th IS&T/SPIE Human Vision and Electronic Imaging Conference 2014.*
25. Exponential family matrix completion under structural constraints. **Suriya Gunasekar**, Pradeep Ravikumar, Joydeep Ghosh. *In 31st International Conference on Machine Learning (ICML) 2014.*
26. Noisy matrix completion using alternating minimization. **Suriya Gunasekar**, Ayan Acharya, Neeraj Gaur, Joydeep Ghosh. *In Machine Learning and Knowledge Discovery in Databases (ECML/PKDD) 2013.*
27. Review quality aware collaborative filtering. Sindhu Raghavan, **Suriya Gunasekar**, Joydeep Ghosh. *In 6th ACM Conference on Recommender Systems (RecSys) 2012.*

Journal publications

28. The implicit bias of gradient descent on separable data. Daniel Soudry, Elad Hoffer, Mor Scpigel Nacson, **Suriya Gunasekar**, Nathan Srebro. *In Journal of Machine Learning Research 2018.*
29. Face detection on distorted images augmented by perceptual quality–aware features. **Suriya Gunasekar**, Joydeep Ghosh, Alan Bovik. *In IEEE Transactions on Information Forensics and Security 2014.*

Professional activities

Organization/Committees

2022 **Organizer** for *Deep Learning Theory Workshop and Summer School at Simons Institute at Berkeley.*

2021 **Social Chair** for *Let-All Mentorship Workshop* at ALT-2021

2019-2021 **Publications Chair** for *Conference on Learning Theory (COLT) 2019, 2020, 2021.*

2019 **Organizer** of *Workshop on Theory of Deep Learning: Where next? at Institute of Advanced Study (IAS).* Website: https://www.math.ias.edu/wtdl.

2018 **Organizer and co-instructor** for the *Introduction to Machine Learning Summer School* jointly organized by TTIC and UChicago. Website: http://ttic.edu/intromlss2018

2018 **Organizing committee member** for the *2nd Midwest Machine Learning Symposium (MMLS).* Website: http://midwest-ml.org/2018/

2018 **Committee member** on drafting of the *Policy on Discrimination, Harassment, and Abusive Behavior at TTIC.*

2016-2019 **Organizer** for *Machine Learning Seminar Series at TTIC.* Website: http://www.ttic.edu/mls/.

Journal reviewing

Journal of the ACM, Journal of Machine Learning Research, Journal of Neurocomputing, Mathematics of Operations Research, Journal of the Royal Statistical Society, IEEE Signal Processing Letters, IEEE Transactions on Image Processing, IEEE Transactions on Knowledge and Data Engineering, IEEE Transactions on Signal Processing.

Conference reviewing

Senior Area Chair for NeurIPS 2023, 2024, 2025. Area Chair for NeurIPS, ICLR (some past years). Reviewer (past years) for International Conference on Machine Learning (ICML), Neural Information Processing Systems (NeurIPS, previously NIPS), International Conference on Learning Representations (ICLR), Uncertainty in Artificial Intelligence (UAI), Conference on Learning Theory (COLT), Algorithmic Learning Theory (ALT), Association for the Advancement of Artificial Intelligence Conference (AAAI), Conference on Data Sciences (CODS).

References

1. **Sébastien Bubeck** sebastien.bubeck@gmail.com
   OpenAI (previously at Microsoft).
2. **Yin Tat Lee** yintat@uw.edu
   OpenAI & University of Washington (previously at Microsoft).
3. **Nathan Srebro** nati@ttic.edu
   Toyota Technological Institute at Chicago.
4. **Avrim Blum** avrim@ttic.edu
   Toyota Technological Institute at Chicago.