

# Behnam Neyshabur

Member of Technical Staff  
Anthropic  
San Francisco, CA 94105

<http://www.neyshabur.net>  
[bneyshabur@gmail.com](mailto:bneyshabur@gmail.com)  
[bneyshabur](#)

---

## Professional Career

- 2024–present **Member of Technical Staff (L8) & Team Lead**, Anthropic.
- 2019–2024 **Senior Staff Research Scientist & Gemini Team Lead**, Google DeepMind.
- 2018–2019 **Postdoctoral Associate**, Computer Science Department, New York University.  
Advisor: Yann LeCun
- 2017–2018 **Research Scholar**, School of Mathematics, Institute for Advanced Study (IAS), Princeton.  
Advisor: Sanjeev Arora
- 2011–2017 **Ph.D. in Computer Science**, Toyota Technological Institute at Chicago (TTIC).  
Advisor: Nathan Srebro - Thesis: Implicit Regularization in Deep Learning.
- 2009–2011 **M.S. in Computer Engineering**, Sharif University of Technology.
- 2005–2009 **B.S. in Computer Engineering**, Sharif University of Technology.

---

## Work towards building an AI Scientist/Engineer

Since 2021, I've pursued the north-star goal of developing a generic AI Scientist/Engineer—capable of the long-term reasoning and fundamental skills necessary to advance scientific and technological frontiers. At various stages, this translated into strategic focus on critical blocking capabilities. A brief timeline:

- 2024–present I am **co-leading the Discovery team** with the north star goal of building an AI scientist
  - 2025 Our work led to threefold improvement in computer-use capability of **Claude Sonnet 4.5**.
- 2022–2024 I **co-led the Blueshift team**, focusing on improving reasoning of foundation models.
  - 2024 Blueshift developed **Gemini 1.5 math-specialized model** with superior math performance.
  - 2024 Blueshift's work led to improved reasoning in **Gemini 1.5 (Pro & Flash) and Gemma 2**.
  - 2023 Blueshift was responsible for the reasoning capabilities of the **first version of Gemini**.
  - 2022 Blueshift introduced **Minerva**, the 1st language model with strong mathematical skills.
  - 2022 Presented four LLM-related papers and gave two invited workshop talks at NeurIPS 2022.
- Pre-2022 Academic research aimed at understanding deep learning and large language models.

---

## Publications (citations: 30K+, h-index: 50+)

- Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities**  
By Gemini Team, et al.. In: *arXiv:2507.06261* (2025).
- Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context**  
By Gemini Team, et al.. In: *Technical Report* (2024).
- Gemma 2: Improving open language models at a practical size**  
By Gemini Team, et al.. In: *arXiv e-prints* (2024).

- 4 **Gemini: a family of highly capable multimodal models**  
By Gemini Team, et al.. In: *arXiv preprint arXiv:2312.11805* (2023).
- 5 **REPAIR: REnormalizing Permuted Activations for Interpolation Repair**  
By K. Jordan, Sedghi, O. Hanie Saukh, R. Entezari, and B. Neyshabur. In: *ICLR*, 2023.
- 6 **Long Range Language Modeling via Gated State Spaces**  
By H. Mehta, A. Gupta, A. Cutkosky, and B. Neyshabur. In: *ICLR*, 2023.
- 7 **Beyond Human Data: Scaling Self-Training for Problem-Solving with Language Models**  
By A. Singh, J. D. Co-Reyes, R. Agarwal, A. Anand, P. Patil, P. J. Liu, J. Harrison, J. Lee, K. Xu, A. Parisi, et al.. In: *arXiv preprint arXiv:2312.06585* (2023).
- 8 **Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models**  
By A. Srivastava, A. Rastogi, A. Rao, A. A. M. Shueb, A. Abid, A. Fisch, A. R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, et al.. In: *TMLR* (2023).
- 9 **Exploring the Limits of Large Scale Pre-training**  
By S. Abnar, M. Dehghani, B. Neyshabur, and H. Sedghi. In: *ICLR*, Nov. 2022 (spotlight).
- 10 **Revisiting Neural Scaling Laws in Language and Vision**  
By I. Alabdulmohsin, B. Neyshabur, and X. Zhai. In: *NeurIPS*, Nov. 2022.
- 11 **The evolution of out-of-distribution robustness throughout fine-tuning**  
By A. Andreassen, Y. Bahri, B. Neyshabur, and R. Roelofs. In: *TMLR* (May 2022).
- 12 **Exploring Length Generalization in Large Language Models**  
By C. Anil, Y. Wu, A. Andreassen, A. Lewkowycz, V. Misra, V. Ramasesh, A. Slone, G. Gur-Ari, E. Dyer, and B. Neyshabur. In: *NeurIPS*, 2022 (oral).
- 13 **Data Scaling Laws in NMT: The Effect of Noise and Architecture**  
By Y. Bansal, B. Ghorbani, A. Garg, B. Zhang, M. Krikun, C. Cherry, B. Neyshabur, and O. Firat. In: *ICML*, 2022.
- 14 **The Role of Permutation Invariance in Linear Mode Connectivity of Neural Networks**  
By R. Entezari, H. Sedghi, O. Saukh, and B. Neyshabur. In: *ICLR*, 2022.
- 15 **Convexifying Transformers: Improving optimization and understanding of transformer networks**  
By T. Ergen, B. Neyshabur, and H. Mehta. In: *arXiv preprint arXiv:2211.11052* (2022).
- 16 **Leveraging Unlabeled Data to Predict Out-of-Distribution Performance**  
By S. Garg, S. Balakrishnan, Z. C. Lipton, B. Neyshabur, and H. Sedghi. In: *ICLR*, 2022.
- 17 **A Loss Curvature Perspective on Training Instability in Deep Learning**  
By J. Gilmer, B. Ghorbani, A. Garg, S. Kudugunta, B. Neyshabur, D. Cardoze, G. Dahl, Z. Nado, and O. Firat. In: *ICLR*, 2022.
- 18 **Block-Recurrent Transformers**  
By D. Hutchins, I. Schlag, Y. Wu, E. Dyer, and B. Neyshabur. In: *NeurIPS*, 2022.
- 19 **Layer-Stack Temperature Scaling**  
By A. Khalifa, M. C. Mozer, H. Sedghi, B. Neyshabur, and I. Alabdulmohsin. In: *arXiv preprint arXiv:2211.10193* (2022).
- 20 **Solving Quantitative Reasoning Problems with Language Models**  
By A. Lewkowycz, A. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, et al.. In: *NeurIPS*, 2022.
- 21 **Teaching Algorithmic Reasoning via In-context Learning**  
By H. Zhou, A. Nova, H. Larochelle, A. Courville, N. Behnam, and H. Sedghi. In: *arXiv preprint arXiv:2211.09066* (2022).

- 22 Deep Learning Through the Lens of Example Difficulty**  
By R. J. Baldock, H. Maennel, and B. Neyshabur. In: *NeurIPS*, 2021.
- 23 Sharpness-Aware Minimization for Efficiently Improving Generalization**  
By P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur. In: *ICLR*, 2021 (spotlight).
- 24 Are wider nets better given the same number of parameters?**  
By A. Golubeva, B. Neyshabur, and G. Gur-Ari. In: *ICLR*, 2021.
- 25 Methods and Analysis of The First Competition in Predicting Generalization of Deep Learning**  
By Y. Jiang, P. Natekar, M. Sharma, S. K. Aithal, D. Kashyap, N. Subramanyam, C. Lassance, D. M. Roy, G. K. Dziugaite, S. Gunasekar, et al.. In: *NeurIPS 2020 Competition and Demonstration Track*, PMLR, 2021.
- 26 Extreme Memorization via Scale of Initialization**  
By H. Mehta, A. Cutkosky, and B. Neyshabur. In: *ICLR*, 2021.
- 27 Understanding the failure modes of out-of-distribution generalization**  
By V. Nagarajan, A. Andreassen, and B. Neyshabur. In: *ICLR*, 2021.
- 28 The Deep Bootstrap: Good Online Learners are Good Offline Generalizers**  
By P. Nakkiran, B. Neyshabur, and H. Sedghi. In: *ICLR*, 2021.
- 29 When Do Curricula Work?**  
By X. Wu, E. Dyer, and B. Neyshabur. In: *ICLR*, 2021 (oral).
- 30 The intriguing role of module criticality in the generalization of deep networks**  
By N. S. Chatterji, B. Neyshabur, and H. Sedghi. In: *ICLR*, 2020 (spotlight).
- 31 Neurips 2020 competition: Predicting generalization in deep learning**  
By Y. Jiang, P. Foret, S. Yak, D. M. Roy, H. Mobahi, G. K. Dziugaite, S. Bengio, S. Gunasekar, I. Guyon, and B. Neyshabur. In: *arXiv preprint arXiv:2012.07976* (2020).
- 32 Fantastic Generalization Measures and Where to Find Them**  
By Y. Jiang, B. Neyshabur, H. Mobahi, D. Krishnan, and S. Bengio. In: *ICLR*, 2020.
- 33 Towards learning convolutions from scratch**  
By B. Neyshabur. In: *NeurIPS*, 2020.
- 34 What is being transferred in transfer learning?**  
By B. Neyshabur, H. Sedghi, and C. Zhang. In: *NeurIPS*, 2020.
- 35 Observational Overfitting in Reinforcement Learning**  
By X. Song, Y. Jiang, Y. Du, and B. Neyshabur. In: *ICLR*, 2020.
- 36 Towards Understanding the Role of Over-Parametrization in Generalization of Neural Networks**  
By B. Neyshabur, Z. Li, S. Bhojanapalli, Y. LeCun, and N. Srebro. In: *ICLR*, 2019.
- 37 Stronger Generalization Bounds for Deep Nets via a Compression Approach**  
By S. Arora, R. Ge, B. Neyshabur, and Y. Zhang. In: *ICML*, 2018.
- 38 Predicting Protein–Protein Interactions through Sequence-Based Deep Learning**  
By S. Hashemifar, B. Neyshabur, A. A. Khan, and J. Xu. In: *Bioinformatics* 34.17 (2018).
- 39 A PAC-Bayesian Approach to Spectrally-Normalized Margin Bounds for Neural Networks**  
By B. Neyshabur, S. Bhojanapalli, and N. Srebro. In: *ICLR*, 2018.
- 40 Corralling a Band of Bandit Algorithms**  
By A. Agarwal, H. Luo, B. Neyshabur, and R. E. Schapire. In: *COLT*, 2017.
- 41 Implicit Regularization in Matrix Factorization**  
By S. Gunasekar, B. E. Woodworth, S. Bhojanapalli, B. Neyshabur, and N. Srebro. In: *NeurIPS*, 2017 (spotlight).

- 42 **Implicit Regularization in Deep Learning**  
By B. Neyshabur, PhD thesis, TTIC, 2017.
- 43 **Stabilizing GAN Training with Multiple Random Projections**  
By B. Neyshabur, S. Bhojanapalli, and A. Chakrabarti. In: *arXiv preprint* (2017).
- 44 **Exploring Generalization in Deep Learning**  
By B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro. In: *NeurIPS*, 2017.
- 45 **Geometry of Optimization and Implicit Regularization in Deep Learning**  
By B. Neyshabur, R. Tomioka, R. Salakhutdinov, and N. Srebro. In: *arXiv preprint* (2017).
- 46 **Global Optimality of Local Search for Low Rank Matrix Recovery**  
By S. Bhojanapalli, B. Neyshabur, and N. Srebro. In: *NeurIPS*, 2016.
- 47 **Data-Dependent Path Normalization in Neural Networks**  
By B. Neyshabur, R. Tomioka, R. Salakhutdinov, and N. Srebro. In: *ICLR*, 2016.
- 48 **Path-normalized Optimization of Recurrent Neural Networks with ReLU Activations**  
By B. Neyshabur, Y. Wu, R. Salakhutdinov, and N. Srebro. In: *NeurIPS*, 2016.
- 49 **Joint Inference of Tissue-Specific Networks with a Scale Free Topology**  
By S. Hashemifar, B. Neyshabur, and J. Xu. In: *BIBM*, 2015.
- 50 **Path-SGD: Path-Normalized Optimization in Deep Neural Networks**  
By B. Neyshabur, R. R. Salakhutdinov, and N. Srebro. In: *NeurIPS*, 2015.
- 51 **On Symmetric and Asymmetric LSHs for Inner Product Search**  
By B. Neyshabur and N. Srebro. In: *ICML*, 2015.
- 52 **In Search of the Real Inductive Bias: On the Role of Implicit Regularization in Deep Learning**  
By B. Neyshabur, R. Tomioka, and N. Srebro. In: *ICLR workshop*, 2015.
- 53 **Norm-Based Capacity Control in Neural Networks**  
By B. Neyshabur, R. Tomioka, and N. Srebro. In: *COLT*, 2015.
- 54 **Clustering, Hamming Embedding, Generalized LSH and the Max Norm**  
By B. Neyshabur, Y. Makarychev, and N. Srebro. In: *ALT*, 2014.
- 55 **Sparse Matrix Factorization**  
By B. Neyshabur and R. Panigrahy. In: *arXiv preprint* (2014).
- 56 **NETAL: A New Graph-Based Method for Global Alignment of Protein-Protein Interaction Networks**  
By B. Neyshabur, A. Khadem, S. Hashemifar, and S. S. Arab. In: *Bioinformatics* 29.13 (2013).
- 57 **The Power of Asymmetry in Binary Hashing**  
By B. Neyshabur, P. Yadollahpour, N. Srebro, R. Salakhutdinov, and Y. Makarychev. In: *NeurIPS*, 2013.

---

## Internships

- 2016 **Research Intern, Microsoft Research**, New York, NY.
- 2013 **Research Intern, Microsoft Research**, Mountain View, CA.

---

## Past Academic Service

**Editorial Board/Area Chair:** *NeurIPS*, *ICLR*, *JMLR*, *TMLR*.

**Reviewer:** *NeurIPS*, *ICLR*, *JMLR*, *ICML*, *COLT*.