

# Mikhail Khodak

✉ khodak@cmu.edu • 🌐 cs.cmu.edu/~mkhodak

## Positions

---

- **Carnegie Mellon University** **Pittsburgh, PA**  
*Doctoral Student, Computer Science Department* 2018–Present
- **Google Research - New York** **New York, NY**  
*Research Intern, Modeling and Data Science* 2022
- **Microsoft Research Lab - New England** **Cambridge, MA**  
*Research Intern, Machine Learning* 2020
- **Princeton University** **Princeton, NJ**  
*Research Assistant, Department of Computer Science* 2016–2018
- **Lawrence Livermore National Laboratory** **Livermore, CA**  
*Summer Intern, Inertial Confinement Fusion* 2015, 2016
- **Princeton Plasma Physics Laboratory** **Princeton, NJ**  
*Research Assistant, Plasma Science & Technology* 2011–2014

## Education

---

- **Princeton University** **Princeton, NJ**  
*M.S.E. in Computer Science* 2016–2018
- **Princeton University** **Princeton, NJ**  
*A.B. in Mathematics with Honors* 2012–2016

## Awards

---

- Selected to attend the 9th Heidelberg Laureate Forum 2022
- Facebook Fellowship Award 2021
- Two Sigma PhD Fellowship First Runner-Up Prize 2020
- IBM Poster Award Presentation 2nd Place Prize, NYAS Machine Learning Symposium 2019

## Publications

---

### Refereed Conference Proceedings

- [1] M.-F. Balcan, M. Khodak, D. Sharma, A. Talwalkar. [Provably Tuning the ElasticNet Across Instances](#). To appear in *Advances in Neural Information Processing Systems (NeurIPS)*. 2022.
- [2] J. Shen\*, M. Khodak\*, A. Talwalkar. [Efficient Architecture Search for Diverse Tasks](#). To appear in *Advances in Neural Information Processing Systems (NeurIPS)*. 2022.

- [3] M. Khodak, M.-F. Balcan, A. Talwalkar, S. Vassilvitskii. [Learning Predictions for Algorithms with Predictions](#). To appear in *Advances in Neural Information Processing Systems* (NeurIPS). 2022.
- [4] R. Tu\*, N. Roberts\*, M. Khodak, J. Shen, F. Sala, A. Talwalkar. [NAS-Bench-360: Benchmarking Diverse Tasks for Neural Architecture Search](#). To appear in *Advances in Neural Information Processing Systems* (NeurIPS): Datasets and Benchmarks Track. 2022.
- [5] C. White, M. Khodak, R. Tu, S. Shah, S. Bubeck, D. Dey. [A Deeper Look at Zero-Cost Proxies for Lightweight NAS](#). In *Proceedings of the 9th International Conference on Learning Representations* (ICLR): Blog Track. 2022.
- [6] N. Roberts\*, M. Khodak\*, T. Dao, L. Li, C. Ré, A. Talwalkar. [Rethinking Neural Operations for Diverse Tasks](#). In *Advances in Neural Information Processing Systems* (NeurIPS). 2021.
- [7] M.-F. Balcan, M. Khodak, D. Sharma, A. Talwalkar. [Learning-to-Learn Non-Convex Piecewise-Lipschitz Functions](#). In *Advances in Neural Information Processing Systems* (NeurIPS). 2021.
- [8] M. Khodak, R. Tu, T. Li, L. Li, M.-F. Balcan, V. Smith, A. Talwalkar. [Federated Hyperparameter Tuning: Challenges, Baselines, and Connections to Weight-Sharing](#). In *Advances in Neural Information Processing Systems* (NeurIPS). 2021.
- [9] L. Li\*, M. Khodak\*, M.-F. Balcan, A. Talwalkar. [Geometry-Aware Gradient Algorithms for Neural Architecture Search](#). In *Proceedings of the 9th International Conference on Learning Representations* (ICLR). 2021. **Spotlight**.
- [10] M. Khodak, N. A. Tenenholz, L. Mackey, N. Fusi. [Initialization and Regularization of Factorized Neural Layers](#). In *Proceedings of the 9th International Conference on Learning Representations* (ICLR). 2021.
- [11] N. Saunshi, Y. Zhang, M. Khodak, S. Arora. [A Sample Complexity Separation between Non-Convex and Convex Meta-Learning](#). In *Proceedings of the 37th International Conference on Machine Learning* (ICML). 2020.
- [12] J. Li, M. Khodak, S. Caldas, A. Talwalkar. [Differentially Private Meta-Learning](#). In *Proceedings of the 8th International Conference on Learning Representations* (ICLR). 2020.
- [13] M. Khodak, M.-F. Balcan, A. Talwalkar. [Adaptive Gradient-Based Meta-Learning Methods](#). In *Advances in Neural Information Processing Systems* (NeurIPS). 2019.
- [14] S. Arora, H. Khandeparkar, M. Khodak, O. Plevrakis, N. Saunshi. [A Theoretical Analysis of Contrastive Unsupervised Representation Learning](#). In *Proceedings of the 36th International Conference on Machine Learning* (ICML). 2019. **Oral**.
- [15] M. Khodak, M.-F. Balcan, A. Talwalkar. [Provable Guarantees for Gradient-Based Meta-Learning](#). In *Proceedings of the 36th International Conference on Machine Learning* (ICML). 2019.
- [16] M. Khodak\*, N. Saunshi\*, Y. Liang, T. Ma, B. Stewart, S. Arora. [A La Carte Embedding: Cheap but Effective Induction of Semantic Feature Vectors](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (ACL). 2018.
- [17] M. Khodak, N. Saunshi, K. Vodrahalli. [A Large Self-Annotated Corpus for Sarcasm](#). In *Proceedings of the 11th Language Resources and Evaluation Conference* (LREC). 2018.
- [18] S. Arora, M. Khodak, N. Saunshi, K. Vodrahalli. [A Compressed Sensing View of Unsupervised Text Embeddings, Bag-of-n-Grams, and LSTMs](#). In *Proceedings of the 6th International Conference on Learning Representations* (ICLR). 2018.

- [19] M. Khodak, L. Zheng, A. S. Lan, C. Joe-Wong, M. Chiang. [Learning Cloud Dynamics to Optimize Spot Instance Bidding Strategies](#). In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 2018.

### Refereed Journal Articles

- [1] P. Kairouz\*, H. B. McMahan\*, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, R. G. L. D'Oliveira, S. El Rouayheb, D. Evans, J. Gardner, Z. Garrett, A. Gascón, B. Ghazi, P. B. Gibbons, M. Gruteser, Z. Harchaoui, C. He, L. He, Z. Huo, B. Hutchinson, J. Hsu, M. Jaggi, T. Javidi, G. Joshi, M. Khodak, J. Konečný, A. Korolova, F. Koushanfar, S. Koyejo, T. Lepoint, Y. Liu, P. Mittal, M. Mohri, R. Nock, A. Özgür, R. Pagh, M. Raykova, H. Qi, D. Ramage, R. Raskar, D. Song, W. Song, S. U. Stich, Z. Sun, A. T. Suresh, F. Tramèr, P. Vepakomma, J. Wang, L. Xiong, Z. Xu, Q. Yang, F. X. Yu, H. Yu, S. Zhao. [Advances and Open Problems in Federated Learning](#). In *Foundations and Trends in Machine Learning (FTML)* **14** 1-2, pp. 1-210. 2021.
- [2] S. Cohen, C. Swanson, N. McGreivy, A. Raja, E. Evans, P. Jandovitz, M. Khodak, G. Pajer, T. D. Rognlien, S. Thomas, M. Paluszek. [Direct Fusion Drive for Interstellar Exploration](#). In *Journal of the British Interplanetary Society (JBIS)* **73**, pp. 37–50. 2019.

### Abstracts and Workshops

- [1] M. Khodak, L. Li, N. Roberts, M.-F. Balcan, A. Talwalkar. [A Simple Setting for Understanding Neural Architecture Search with Weight-Sharing](#). In *Proceedings of the 7th ICML Workshop on Automated Machine Learning (AutoML)*. 2020.
- [2] O. Adewale, A. Beatson, D. Buniatyan, J. Ge, M. Khodak, H. Lee, N. Prasad, N. Saunshi, A. Seff, K. Singh, D. Suo, C. Zhang, S. Arora. [Pixie: A Social Chatbot](#). In *1st Proceedings of Alexa Prize (Alexa Prize)*. 2017.
- [3] M. Khodak, A. Risteski, C. Fellbaum, S. Arora. [Automated WordNet Construction Using Word Embeddings](#). In *Proceedings of the 1st Workshop on Sense, Concept, and Entity Representations and their Applications (SENSE)*. 2017.
- [4] M. Khodak, R. L. Berger, T. Chapman, J. A. F. Hittinger. [Development and Application of a Multi-Fluid Simulation Code for Modeling Interpenetrating Plasmas](#). In *57th Annual Meeting of the APS Division of Plasma Physics (DPP)*. 2015.
- [5] S. A. Cohen, M. Chu-Cheong, R. Feder, K. Griffin, M. Khodak, J. Klabacha, E. Meier, S. Newbury, M. Paluszek, T. Rognlien, S. Thomas, M. Walsh. [Reducing Neutron Emission from Small Fusion Rocket Engines](#). In *Proceedings of the 66th International Astronautical Congress (IAC)*. 2015.

### Manuscripts and Preprints

- [1] M.-F. Balcan, K. Harris, M. Khodak, Z. S. Wu. [Meta-Learning Adversarial Bandits](#). 2022.
- [2] L. Dery, P. Michel, M. Khodak, G. Neubig, A. Talwalkar. [AANG: Automating Auxiliary Learning](#). 2022.
- [3] M. Khodak, N. Saunshi. [A Case Study Critique and Principled Alternatives for Equilibrium-Seeking Congestion Control Protocols](#). 2018.
- [4] M. Khodak, A. Risteski, C. Fellbaum, S. Arora. [Extending and Improving WordNet via Unsupervised Word Embeddings](#). 2017.

- [5] M. Chang, M. Khodak. [Computation and Design of Locally-Connected Internet Exchange Networks](#). 2016.

## Talks

---

*Federated Hyperparameter Tuning: Challenges, Baselines, and Connections to Weight-Sharing*  
MedAI Group Exchange Session, Stanford. 17 March 2022.

*Towards Automatic Architecture Design for Emerging Machine Learning Tasks*  
DDPS Webinar, LLNL. 4 November 2021.

*Factorized Layers Revisited: Compressing Deep Nets without Playing the Lottery*  
DS-BOND Meeting, Microsoft ILDC. 8 June 2021.  
Artificial Intelligence Seminar, CMU. 20 April 2021.

*Geometry-Aware Gradient Algorithms for Neural Architecture Search*  
Workshop on Learning Network Architecture During Training, AAAI. 8 February 2021.

*ARUBA: Efficient and Adaptive Meta-Learning with Provable Guarantees*  
Federated Learning One World Seminar. 9 September 2020.  
Artificial Intelligence Seminar, CMU. 3 December 2019.  
Microsoft Research - New England. 22 November 2019.  
Adaptive and Multi-Task Learning Workshop, ICML. 15 June 2019.

*A Compressed Sensing View of Unsupervised Text Embeddings, Bag-of-n-Grams, and LSTMs*  
Theory of Deep Learning Workshop, ICML. 14 July 2018.  
Seminar on Theoretical Machine Learning, IAS. 5 April 2018.

*A Multi-Fluid Model for Interpenetrating Plasma Flows*  
Hohlraum Science Campaign Meeting, LLNL. 17 August 2016.

## Service

---

### Committee Positions

Co-Organizer of the AutoML Decathlon Competition at NeurIPS 2022.

CMU Computer Science Department 2022 PhD Admissions Committee.

Workflow Assistant for the Thirty-Fourth Conference on Neural Information Processing Systems (NeurIPS 2020).

Co-Organizer of the Workshop on Wordnets and Word Embeddings at the 9th Global Wordnet Conference (GWC 2018).

### Journal Reviewing

TMLR, IEEE TPAMI, SIAM Journal on Mathematics of Data Science, Language Resources and Evaluation, Machine Learning

### Conference Reviewing

ICML, NeurIPS, ICLR, AAAI, EMNLP, ACL, NAACL, KDD

### Outreach

Mentorship Lead, CSD PhD Student Council

2022-Present

Editor, CMU Machine Learning Blog 2020-Present  
Research Team Leader, CMU OurCS 2019  
Site Coordinator, North American Computational Linguistics Olympiad (NACLO) 2016-2018

## Mentoring

---

Samuel Guo (with Ameet Talwalkar and Junhong Shen) 2021-Present  
Renbo Tu (with Ameet Talwalkar) 2020-2022  
    Current position: PhD student in computer science at the University of Toronto  
Nicholas Roberts (with Ameet Talwalkar) 2020-2021  
    Current position: PhD student in computer science at the University of Wisconsin-Madison  
Jeffrey Li (with Ameet Talwalkar) 2019-2020  
    Current position: PhD student in computer science at the University of Washington

## Teaching

---

### **Carnegie Mellon University**

*Teaching Assistant, Machine Learning Department*

10-315: Introduction to Machine Learning Spring 2019, Spring 2021

### **Princeton University**

*Assistant in Instruction, Department of Computer Science*

COS 495: Special Topics in CS – Natural Language Processing Spring 2018

COS 397: Junior Seminar – Natural Language Processing Fall 2017

COS 340: Reasoning About Computation Fall 2016, Spring 2017

*Course Assistant, Department of Mathematics*

MAT 385: Theory of Games Spring 2015, Spring 2016