

Riyasat Ohib

CONTACT
INFORMATION

Google DeepMind
New York City, NY, USA
riyasatohib.com
riyasat.ohib@gmail.com

 GitHub
 LinkedIn
 Google Scholar

RESEARCH
INTERESTS

My research examines the computational and representational organization of deep neural networks. I study how internal structure develops through learning and shapes reasoning, adaptation, and model behavior, using efficient methods that extend to large models. A central interest is alignment, especially how stable patterns of behavior and model personas form and persist, and how they can be controlled. I am also deeply interested in the science of human and machine learning.

EDUCATION

PhD, Georgia Institute of Technology Aug 2021 – May 2026
Electrical and Computer Engineering Atlanta, GA
Concentration in Artificial Intelligence · GPA: 4.0/4.0
Dissertation: The Hidden Structure of Deep Models: Discovering and Exploiting Sparse Subnetworks
Advisors: [Vince D. Calhoun](#) and [Sergey Plis](#)

MS, Georgia Institute of Technology Aug 2019 – Aug 2021
Electrical and Computer Engineering Atlanta, GA
Concentration in Artificial Intelligence · GPA: 4.0/4.0
Thesis: Explicit Group Sparse Projection for Machine Learning

EXPERIENCE

Google DeepMind · Research Scientist Jul 2026 – Present
Permanent, full-time New York City, NY
Research on representations, control, alignment and interpretability of large-scale frontier models.

Google DeepMind · Research Intern Sep 2025 – May 2026
Model Alignment, Control, Responsibility and Safety Kirkland, WA
Studied how concepts are represented in the residual stream of diffusion transformers, and how that structure can be used to steer generation (CVPR HOW workshop 2026).

Cohere · Intern of the Technical Staff Sep 2024 – Dec 2024
LLM Efficiency Research Atlanta, GA
Developed inference-time activation sparsity methods for large language models, targeting deployment settings where latency and serving cost dominate.

Dolby Laboratories · PhD Research Intern May 2024 – Aug 2024
Experience Delivery Lab, Advanced Technologies Group Atlanta, GA
Designed a parameter-efficient fine-tuning method for LLMs based on probabilistic layer selection, concentrating updates on the layers most relevant to the downstream task.

Meta AI (FAIR) · Research Scientist Intern May 2022 – Aug 2022
Sparsity and Efficiency Menlo Park, CA
Built *WeiGit*, a git-like library for versioning and compressing neural network weights during training; released as part of facebookresearch/fairscale.
Studied extreme sparsity in deep networks using signal-processing transforms (FFT, DCT) applied during training.

EXPERIENCE
(continued)**TReNDS Center** · Graduate Research Assistant

Aug 2019 – Aug 2026

Georgia Institute of Technology

Atlanta, GA

Introduced sparse adapters as parameter-efficient building blocks for modular architectures, outperforming LoRA and full fine-tuning when merged across 20 NLP tasks; with [Mila](#), Montreal (COLM 2025).

Developed communication-efficient federated learning for the non-IID regime, applied to distributed neuroimaging (TMLR 2026; Frontiers in Neuroinformatics 2024).

Designed single-shot pruning and neural-pathway discovery for offline and multi-task reinforcement learning; with Mila (NeurIPS 2024).

Built Grouped Sparse Projection, a projection-based algorithm that trains sparse deep models to near-baseline accuracy above 90% sparsity (TMLR 2022; ICLR HAET workshop 2021).

PUBLICATIONS

* Equal contribution

- [1] **R. Ohib**, M. Hahn, M. Malek. *Concept Spaces in the Residual Stream of Diffusion Transformers*. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), HOW workshop, 2026. [\[paper\]](#)
- [2] J. Gammell, B. Thapaliya, Y. Jung, **R. Ohib**, B. Fehri, D. Chakrabarti. *Learning to Query History: Non-stationary Classification via Learned Retrieval*. International Conference on Learning Representations (ICLR), TSALM workshop, 2026. [\[paper\]](#)
- [3] **R. Ohib**, B. Thapaliya, G. K. Dziugaite, J. Liu, V. D. Calhoun, S. Plis. *SSFL: Discovering Sparse Unified Subnetworks at Initialization for Efficient Federated Learning*. Transactions on Machine Learning Research (TMLR), 2026. [\[paper\]](#)
- [4] S. Y. Arnob, Z. Su, M. Kim, O. Ostapenko, **R. Ohib**, E. Saleh, D. Precup, L. Caccia, A. Sordoni. *Exploring Sparse Adapters for Scalable Merging of Parameter Efficient Experts*. Conference on Language Modeling (COLM), 2025. [\[paper\]](#)
- [5] B. Thapaliya*, **R. Ohib***, E. Geenjaara, J. Liu, V. Calhoun, S. M. Plis. *Efficient Federated Learning for Distributed Neuroimaging Data*. Frontiers in Neuroinformatics, 2024. [\[paper\]](#)
- [6] S. Y. Arnob, **R. Ohib**, S. M. Plis, A. Zhang, A. Sordoni, D. Precup. *Efficient Reinforcement Learning by Discovering Neural Pathways*. Conference on Neural Information Processing Systems (NeurIPS), 2024. [\[paper\]](#)
- [7] B. Thapaliya, **R. Ohib**, E. Geenjaara, J. Liu, V. Calhoun, S. M. Plis. *Decentralized Sparse Federated Learning for Efficient Training on Distributed Neuroimaging Data*. Conference on Neural Information Processing Systems (NeurIPS), MedNeurIPS workshop, 2023. [\[paper\]](#)
- [8] **R. Ohib**, B. Thapaliya, P. Gaggenapalli, J. Liu, V. Calhoun, S. Plis. *SalientGrads: Sparse Models for Communication Efficient and Data Aware Distributed Federated Training*. International Conference on Learning Representations (ICLR), SNN workshop, 2023. [\[paper\]](#)
- [9] E. Geenjaara, D. Kim, **R. Ohib**, M. Duda, A. Kashyap, S. M. Plis, V. Calhoun. *Uncovering the latent dynamics of whole-brain fMRI tasks with a sequential variational autoencoder*. Conference on Neural Information Processing Systems (NeurIPS), DeepGen workshop, 2023. [\[paper\]](#)
- [10] **R. Ohib**, N. Gillis, N. Dalmaso, S. Shah, V. Potluru, S. Plis. *Explicit Group Sparse Projection with Applications to Deep Learning and NMF*. Transactions on Machine Learning Research (TMLR), 2022. [\[paper\]](#)
- [11] **R. Ohib**, N. Gillis, S. Shah, V. Potluru, S. Plis. *Grouped Sparse Projection for Deep Learning*. International Conference on Learning Representations (ICLR), HAET workshop, 2021. [\[paper\]](#)
- [12] S. Yeasar, **R. Ohib**, S. Plis, D. Precup. *Single-Shot Pruning for Offline Reinforcement Learning*. Conference on Neural Information Processing Systems (NeurIPS), Off-RL workshop, 2021. [\[paper\]](#)
- [13] **R. Ohib**, S. Y. Arnob, M. Muhaisin, R. Arefin, T. Reza, M. R. Amin. *ENF Based Machine Learning Classification for Origin of Media Signals: Novel Features from Fourier Transform Profile*. International Conference on Electrical, Electronics and Computer Science (ICEECS), 2018.
- [14] S. Y. Arnob, **R. Ohib**, M. M. Muhaisin, T. B. Hassan. *Power file extraction process from bangladesh grid and exploring ENF based classification accuracy using machine learning*. 2017 IEEE Region 10

Humanitarian Technology Conference (R10-HTC), 2017. [\[paper\]](#)

TALKS AND PRESENTATIONS	<i>The Hidden Structure of Deep Models: Discovering and Exploiting Sparse Subnetworks</i>		Jul 2026
	TReNDS Center		Atlanta, GA
	<i>What Models Learn and What They Hide: Sparse Structures and Representation Control</i>		Apr 2026
	Google DeepMind		Kirkland, WA
	<i>Navigating a PhD and Doing Scientific Research</i>		Apr 2026
	TReNDS Center		Atlanta, GA
OPEN SOURCE	WeiGit		2022
	Weight versioning for neural networks · Meta AI (FAIR) · facebookresearch/fairscale		
	Version control and compression for model weights during training, exposed through a git-like command line and Python API; merged into the fairscale library.		
HONORS AND AWARDS	Full Funding Award, CIFAR DLRL Summer School		2019
	Selected by research proposal · CIFAR and Amii · Edmonton, Canada		
	Honorable Mention, IEEE Signal Processing Cup		2016
	Team and programming lead		
	Undergraduate Honors List		2013 – 2017
	Named to the Graduating Honors Group for academic excellence		
	OIC Scholarship		2013 – 2017
ACADEMIC SERVICE			
	Program Committee		2026
	ISWC	International Semantic Web Conference	
	Reviewer		2023 – Present
	NeurIPS	Conference on Neural Information Processing Systems	
	ICLR	International Conference on Learning Representations	
	TMLR	Transactions on Machine Learning Research	
	TNNLS	IEEE Transactions on Neural Networks and Learning Systems	
GRADUATE COURSEWORK	Statistical Machine Learning		Advanced Programming Techniques
	Convex Optimization		Fourier Analysis
	Linear Algebra		Information Processing in Neural Systems
	Advanced Digital Signal Processing		Real Analysis
PROGRAMMING	Languages	Python, C++, Matlab, Bash	
	Frameworks	PyTorch, JAX, NumPy, Pandas	
	Infrastructure	Linux, Slurm, distributed and multi-GPU training, cluster computing	