

BHAVYA KOHLI

 [bhavyakohli.github.io](https://github.com/bhavyakohli) |  bhavyakohli@u.nus.edu |  BhavyaKohli |  Google Scholar

RESEARCH INTERESTS

My research interest lies in solving interesting problems in machine learning. My recent works range from optimizing masked diffusion models using multivariate noise schedules, deriving a closed-form solution to a constrained linear assignment problem which could be optimized using machine learning to substitute key steps, deriving an optimal perturbation scheme to spotlight a key vulnerability in MPNNs. In my work, I like to maintain an explainable framework, with theory covering all core bases.

EDUCATION

National University of Singapore

PhD in Electrical and Computer Engineering

Singapore
Aug'25 – Present

Indian Institute of Technology (IIT) Bombay

B.Tech in Electrical Engineering and M.Tech in Artificial Intelligence (**Cumulative GPA: 9.13/10.0**)

Mumbai, India
Nov'20 – Aug'25

• **Key Coursework:** Machine Learning, Optimization, Graph Models, Graph Learning, Probability, Linear Algebra

PUBLICATIONS

1. **Bhavya Kohli**, Biplab Sikdar “**PEANUT: Perturbations by Eigenvector Alignment for Attacking Graph Neural Networks Under Topology-Driven Message Passing**”, currently under review. Preprint available at [arXiv:2603.26136](https://arxiv.org/abs/2603.26136).
2. **Bhavya Kohli**, S. Sarangi, A. Shameem, Abir De “**Learning Right Monotone Permutation Matrices for Neural Subsequence Search**”, *accepted* at The Twenty-Ninth Annual Conference on Artificial Intelligence and Statistics (AISTATS 2026).
3. Prateek Garg*, **Bhavya Kohli***, Sunita Sarawagi “**Masked Diffusion Models are secretly Learned-Order Autoregressive Models**”, *published* at the EurIPS 2025 Workshop on Principles of Generative Modeling (PriGM). Preprint available at [arXiv:2511.19152](https://arxiv.org/abs/2511.19152).
4. **Bhavya Kohli**, Varun Kohli, M. N. Aman, Biplab Sikdar “**Swarm-Net: Firmware Attestation in IoT Swarms using Graph Neural Networks and Volatile Memory**”, *published* at the IEEE IoT Journal. Preprint available at [arXiv:2408.05680](https://arxiv.org/abs/2408.05680).
5. Julia Werner, **Bhavya Kohli**, Paul Palomero Bernardo, Christoph Gerum, Oliver Bringmann “**Energy-Efficient Seizure Detection Suitable for Low-Power Applications**”, *published* at the 2024 International Joint Conference on Neural Networks (IJCNN), IEEE, 2024. Preprint available at [arXiv:2406.16948](https://arxiv.org/abs/2406.16948).

RESEARCH EXPERIENCE

Exploiting Structural Vulnerabilities of Graph Neural Networks

Guide: Prof. Biplab Sikdar

Aug'25 – Present
ECE, National University of Singapore

- ♦ **Introduction:** This work aims to exploit GNN architectures which explicitly consume the graph topography matrix as an input for their message passing. We derive a simple closed-form solution for maximizing output norm differences which works in a black-box setting. The manuscript for this work is **under review**.
- Studied injection-based GNN attacks and discovered a potentially novel attack surface over multiple types of GNNs
- Performed extensive experiments with the derived optimization solution on various GNN architectures

Optimizing Masked Diffusion using Multivariate Noise Schedules

Guide: Prof. Sunita Sarawagi

Mar'25 – Oct'25
CSE, IIT Bombay

- ♦ **Introduction:** Masked diffusion models (MDMs) are generally trained in a way that assigns equal weight to all generation orders, despite certain orders performing better during inference. We prove that MDMs are not only capable of learning these orders during training, but that these ordering structures can be extracted and utilized, netting near-SOTA performance with over 110X fewer parameters
- Proved that noise schedules define a preferred unmasking order and designed training strategies to learn these schedules
- Implemented and evaluated simple MDMs with augmented losses on tabular data against state-of-the-art baselines

Graph Summarization for Tackling Graph-Level Tasks

Guide: Prof. Abir De

Aug'24 – Aug'25
CSE, IIT Bombay

- ♦ **Introduction:** Many algorithms for difficult graph-level tasks scale poorly with the size of the input graph(s). The aim of the project is to devise algorithms to compress graphs to aid in both efficient execution, and training new models on these tasks.
- Organized existing literature and implemented a set of baseline algorithms to specific tasks to gain deeper insights into the problem

Sequence Retrieval using Regularized Soft Permutation Matrices

Guide: Prof. Abir De

Aug'23 – Aug'25
CSE, IIT Bombay

- ♦ **Introduction:** Information retrieval often involves complex and intensive computations for getting a ranked list of corpus elements with decreasing relevance to a given query. The proposed plug-and-play method uses permutation matrices as a basis for query-corpus alignment which is followed by simple metric computations. This work has been *accepted* at AISTATS 2026.
- Modeled an optimization problem to learn unknown permutations using the linear sum assignment problem as a basis
- Derived a concise representation of the objective which allowed for appropriate neural replacements of key optimization steps
- Implemented workflows for testing the end-to-end system using multiple sequence datasets with different embedding strategies

Firmware Attestation in IoT Swarms using Volatile Memory

Guide: [Prof. Biplab Sikdar](#)

May'24 – Aug'24
ECE, National University of Singapore

- ◇ **Introduction:** Malicious activity on one node in a swarm can propagate to larger network sections, possibly impacting other clusters of devices. The proposed method works by leveraging the IoT devices' Static random-access memory (SRAM) making it lightweight, efficient, and unaffected by the usual drawbacks of attestation methods. This work has been published at the IEEE IoT Journal.
- Conducted an extensive review on modeling IoT swarms to best utilize graph models in the restricted-information setup
- Designed and implemented a graph transformer network among other popular baseline graph convolution operators
- Evaluated different adversarial scenarios inspired from real-world applications and commonly encountered edge cases

Online Hardware-aware Seizure Detection

Guide: [Prof. Oliver Bringmann](#)

May'23 – Aug'23
CSE (Embedded Systems), University of Tübingen, Germany

- ◇ **Introduction:** Hardware-aware machine learning allows for seamless deployment of algorithms on low-power edge devices. The proposed seizure detection approach involving a low-footprint TC-ResNet combined with time-series analysis is extremely energy efficient and is suitable for real-time seizure detection using neural implants and wearable devices. This work has been published at IJCNN 2024.
- Performed a thorough literature survey on classical time-series processing methods and classification approaches
- Contributed to an open-source hardware acceleration library (HANNAH) for hardware-aware machine learning
- Implemented a time-series-specific data processing pipeline, leading to a > 92% sensitivity on unseen seizures

Security of Graph Neural Networks in IoT

Guide: [Prof. Biplab Sikdar](#)

May'22 – Jul'22
ECE, National University of Singapore

- ◇ **Introduction:** Machine Learning models trained without explicit defense mechanisms are susceptible to adversarial attacks. We explore graphical models, their benefits over others in the IoT domain, which often deals with networks of interconnected sensors, and their robustness against black-box adversarial attacks.
- Studied and evaluated black-box adversarial attacks and defenses on convolutional neural networks and graph neural networks
- Surveyed and implemented query-efficient black-box gradient estimation based adversarial attack on Spatio-Temporal GCNs
- Performed an in-depth attack efficacy comparison with auxiliary model attacks and a other baseline black-box methods

SCHOLASTIC ACHIEVEMENTS

- One of twenty students accepted into the IDDD Programme at the **Center for Machine Intelligence and Data Science** (2023)
- Secured rank **675** in the highly coveted JEE Advanced (Among top 0.33% in the country) (2020)
- Part of the **99.85th percentile** in the JEE Mains 2020 (Among top 0.3% in the country) (2020)
- Awarded the **Kishor Vaigyanic Protsahan Yojana (KVPY)** fellowship by the Department of Science and Technology, India (2020)
- Scored **383/450** in the Birla Institute of Technology and Science Admission Test (BITSAT) (2020)

KEY TECHNICAL PROJECTS

Local Augmentation for Graph Neural Networks

Course Project, Guide: [Prof. Abir De](#) (CS768: Learning with Graphs)

Aug'23 – Nov'23
IIT Bombay

- Implemented the local augmentation method described in the assigned paper, and extended the method to other graph tasks—edge prediction and graph-level classification. Observed an improvement over non-augmented baselines
- Implemented a normalizing flow model demonstrating superior generative capability at the cost of compute-efficiency
- Studied the over-smoothing effect in GNNs and experimented with probabilistic message-passing to combat it

Representation Learning using Spiking Autoencoders

Course Project, Guide: [Prof. Udayan Ganguly](#) (EE746: Neuromorphic Engineering)

Sep'23 – Nov'23
IIT Bombay

- Studied spike and rate and spike coding techniques used in neuromorphic neural models for tasks involving structured data
- Studied the possible benefits of neuromorphic models over existing paradigms in energy efficiency and robustness
- Implemented a spiking neural network (SNN) based autoencoder as a generative model for a few baseline datasets

Variational Thompson Sampling

Course Project, Guide: [Prof. Jayakrishnan Nair](#) (EE6106: Online Learning and Optimisation)

Jan'23 – May'23
IIT Bombay

- Explored variational inference for approximating posteriors in the **Multi-Armed Contextual Bandit** problem
- Studied a Gaussian mixture model (GMM) to model the distribution of parameters over arms
- Studied trade-offs between the variational approach and using Gibbs sampling for this model

Unsupervised Representation Learning using Tensorized Autoencoders

Adapted using work shared by [Prof. Debarghya Ghoshdastidar](#) (CIT, TUM)

Nov'22 – Dec'22
Remote

- Studied methods for unsupervised representation learning using autoencoders and other generative models (VAEs)
- Implemented a new "tensorized autoencoder" architecture regularized using clustering algorithms and benchmarked it against standard AEs on clustering, data denoising, and vision tasks such as image inpainting and denoising

A study on the Learning with Errors problem

Course Project, Guide: [Prof. Bernard Menezes](#) (CS741: Advanced Network Security and Cryptography)

Oct'22 – Dec'22
IIT Bombay

- Studied the Learning with Errors problem and its versatility of use in designing cryptographic schemes
- Implemented a few attacks on binary LWE and a public key based cryptographic scheme using text as the test medium

Predicting Remaining Useful Life (RUL) of EV Batteries

Course Project, Guide: [Prof. Amit Sethi](#) (DS203: Programming for Data Science)

Nov'21 – Jan'22
IIT Bombay

- Designed and extracted temporal features from the NASA prognostic battery dataset after thorough data analysis
- Compared the performance of an ensemble of models for predicting RUL using the engineered features

TECHNICAL SKILLS

Programming Languages Python, Julia, C/C++
Machine Learning Libraries PyTorch, Pytorch-geometric, Tensorflow, scikit-learn, Keras, snnTorch
Misc. Software and Tools $\LaTeX$, Git, Secure Shell (SSH), MATLAB, Docker, Bash scripting

TEACHING EXPERIENCE AND MENTORSHIP

- Graduate Teaching Assistant**
*TA for the course **CG2023: Signals and Systems** for [Prof. Lin Feng, ECE, NUS](#)* *Spring'26*
- Mentored and aided a class of **200+** students with their lab experiments and projects involving hardware implementations of common signal processing workflows such as signal modulation and encoding
- Graduate Teaching Assistant**
*TA for the course **CS419: Introduction to Machine Learning** for [Prof. Abir De, CSE, IIT Bombay](#)* *Spring'24*
- Mentored a class of **100+** cross-department students as part of a 7-member team
 - Assisted the instructor by setting assignments, mentoring students, and conducting tutorials and help sessions
- Student Mentor**
Institute Student Mentorship Program: Selected through rigorous interviews and peer reviews *2023-2024*
- Responsible for guiding 10 freshmen through their first year at IITB as an **Institute Student Mentor**
 - Mentored a group of 5 undergraduates in their sophomore year with their academics as a **Department Academic Mentor**

EXTRACURRICULAR ACTIVITIES

Sports	• Tennis and table tennis enthusiast • Completed a trek spanning 6+ hrs to Anjaneri Hill Fort, Nashik, Maharashtra 4263 ft above the sea level
Volunteering	• Volunteered to teach python to over 500 undergraduates as part of a team over the summer in 2023 • Devoted 80+ hours of volunteering work and awarded a special mention for Exemplary Volunteering Work with under the National Service Scheme, IIT Bombay for the duration Dec'20 to Jul'21
Achievements	• Certified for the Management Bootcamp and $\LaTeX$ courses conducted under the Non-Technical Summer School (NTSS) initiative by UGAC, IIT Bombay in 2021 • Won 2nd place as one of eight contingents in the Inter IIT Scrabble League (IISL) conducted in July 2021