

Neelabh Madan

+1-9177495848
neelabh.madan@gmail.com
in neelabh-madan
neelabh17
UDJZ6bYAAAAJ

Education

- 2024–Present **New York University, Courant Department**, NY, USA
Ph.D. in Computer Science with Prof. Lakshminarayanan Subramaniam - GPA 4.0
- 2018–2022 **Indian Institute of Technology Delhi, Hauz Khas**, Delhi, India
B.Tech (Bachelors) in Mechanical Engineering – CGPA: 9.3/10 (Rank: 8/75)
with Minor Degree in Computer Science – CGPA 9.75/10

Publications

*C = Conference, P = Preprint/Submitted — * implies equal/core contribution*

- P.1 **In-Context Alignment at Scale: When More is Less**
Neelabh Madan*, Lakshminarayanan Subramaniam
Accepted [ICML Workshop MoFA 2025](#) [Poster Accepted]
- P.2 **Teaching LLMs to Search and Backtrack: Reinforcement Learning with Trace Bootstrapping on Graph Tasks (2025)**
Neelabh Madan*, Jatin Prakash, Lerrel Pinto
[\[Code\]](#)
- P.3 **Domain Faithfulness Through Counterfactually Robust Learning (2025)**
Ananth Balashankar*, Ankit Bhardwaj*, Neelabh Madan*, Lakshminarayanan Subramaniam
Under review at [CLear](#)
- C.1 **Enhancing Tail Performance in Extreme Classifiers by Label Variance Reduction**
Anirudh Buvanesh*, Rahul Chand*, Jatin Prakash*, Bhawna Paliwal, Mudit Dhawan, Neelabh Madan, Deepesh Hada, Vidit Jain, Manish Gupta, Ramchandra Ramjee, Manik Varma
[ICLR 2024](#) [Paper]
- C.2 **A Stitch in Time Saves Nine: A Train-Time Regularizing Loss for Improved Neural Network Calibration**
Neelabh Madan*, Ramya Hebbalaguppe*, Jatin Prakash*, Chetan Arora
[CVPR 2022 \(Oral presentation\)](#) [Paper]

Awards and Scholastic Achievements

- Received **GSAS TELL: Audience Favorite Award, 2025 (NYU)** and a cash prize of 1000\$
- Received **MacCracken PhD Fellowship (2024-2029)**
- Received **Engineering Excellence Award (2023-2024)** from **Microsoft** for outstanding product impact
- INOI 2017:** Cleared ZIO and qualified for INOI 2017 (Indian National Olympiad in Informatics)
- Received \$20 000 from Micron Technology for winning the International Micron UV Design Challenge

Research Experience / Industry Experience

- Jun 25-Sept 25 **Amazon Science**, Seattle, USA
Applied Science Intern (most details under NDA)
Advisor: [Elad Liebman](#)
- Developed a foundation model from scratch (pretraining, alignment) for learning unified associate representation.
 - Led end-to-end data preparation, model training (RLFT), and inference serving of the model for 2M users.
- July 22-July 24 **Microsoft Research, eXtreme Classification (XC) group**
Pre-Doctoral Research Fellow
Advisors: [Manik Varma](#), [Amit Sharma](#)
- Worked on Personalized Recommendation by introducing eXtreme classifiers (XC)
 - Enhance tail performance of XC models [[ICLR 2024](#)]
 - Distilling GPT4 to highly-efficient XC models
 - Deploying large-scale XC models on Microsoft Audience Network (MSAN) platform

- May 22-Jul 22 **Tata Consultancy Services (TCS) Research**, Noida, India
Research Intern
 Advisor: [Ramya Hebbalaguppe](#)
 ○ Knowledge Distillation and Calibration
- May 21-Aug 21 **Adobe Research (Media and Data Science Research Lab - MDSR)**, Noida, India
Research Intern
 Advisor: [K Balaji](#)
 ○ Worked on Document Visual Question Answering (DocVQA); Augmented LayoutLMV2 architecture with visual, textual, and layout modalities for a high-level understanding of documents.
- Jul 21-Jul 22 **Indian Institute of Technology**, Delhi, India
Undergraduate Student
 Advisors: [Chetan Arora](#), [Arnob Ghosh](#)
 ○ Worked with Prof. Chetan on "A Stitch in Time Saves Nine" as a project work [Accepted **CVPR 2022 (ORAL)**]
 ○ Investigated, with Prof Arnob, the consequences of partial and full information on Contextual Multi-Armed Bandit as part of [B.Tech Thesis](#); Studied the effect of various non linear loss functions on NeuralUCB algorithm

Real World Deployments

- Jul 22-Jul 24 **Improved Personalized Ad Recommendation using XC models on MSAN**, (*Microsoft Research*)
 ○ Modeled Personalized Ad recommendation as an XC task; trained large eXtreme Classifier models to improve performance over previously deployed Siamese style models on Microsoft Audience Network (MSAN) platform.
 ○ Achieve **200-400% gains** in offline recall metrics, which resulted in **20+ mainstreamings**, generating approx. 200\$ revenue for Microsoft and higher click-through-rates (CTR)
 ○ Deployed models were able to predict ads, extending beyond just lexical matches; **Extended Ads coverage to approx 1 Billion Ads** by training classifiers on meta-clusters of Ads
 ○ Received the **Engineering Excellence Award** for the Financial Year of 2023-2024

Selected Research Projects

- Feb 25-Jul 25 **In-Context Alignment at Scale: When More is Less**, (*New York University*)
 Advisor: [Prof. Lakshminarayanan Subramanian](#)
Accepted @ ICML MoFA Workshop 2025
 ○ Investigated how LLMs handle dynamic rule-following using a synthetic dataset with controllable constraints.
 ○ Found that model performance drops sharply when transitioning from 1 to 2 active rules, and observed "cheating" behaviors via memorization.
 ○ Proposed a structured retrieval mechanism for personalized, scalable rule-following in LLMs.
 ○ [\[Poster Accepted\]](#)
- Feb 25-Present **Teaching LLMs to Search and Backtrack via RL Fine-Tuning**, (*New York University*)
 With: [Jatin Prakash](#), [Prof. Lerrel Pinto](#)
 ○ Created a synthetic graph search benchmark to study exploration and backtracking in LLMs.
 ○ Showed that vanilla RLFT fails on small models (1.5B), but using SFT on 14B traces followed by RLFT boosts correct generations from 2 to 96.
 ○ Achieved better generalization on large graphs compared to even the 14B teacher model.
 ○ [\[Code\]](#)
- Aug 24-Present **Domain Faithfulness through Counterfactually Robust Learning**, (*New York University*)
 Advisors: [Prof. Lakshminarayanan Subramanian](#) , [Ananth Balashankar \(Google\)](#)
Under Review CLear
 ○ Aimed at incorporating domain expertise and rules into deep neural networks (DNNs) by leveraging domain-specific declarative knowledge by implementing modular components for regularization and data augmentation
 ○ Achieved significant improvements on synthetic datasets and real-world tasks, including MIMIC-III medication recommendation and reinforcement learning.
- Sept 23-Jul 24 **Data Distillation for Sequential Recommendation using LLMs**, (*Microsoft Research*)
 Advisors: [Amit Sharma](#), [Manik Varma](#)
 ○ Used LLM (GPT-4) to distill down a smaller LM to summarize long sequential user histories (made up of events such as bing searches, websites visited, etc) into shorter text representation for personalized recommendation.
 ○ Shorter yet diverse textual representation of long user histories enables retention of multiple intent in the user embedding generated by BERT like encoders. This is an improvement over previously used methods that embed-then-aggregate individual user events.

- Dec 22-May 23 **Enhancing XC models on tail labels**, (*Microsoft Research*)
 Advisors: [Manik Varma](#) — [ICLR 2024]
- XC models (consisting of one-vs-all linear classifiers) are notorious for sub-par performance on labels that do not have enough training points (tail labels). Siamese models, however, have superior performance on tail labels
 - Exploiting the above, we proposed a framework that improves tail performance of XC classifiers by distilling tail-robust encoders in the form of soft-label targets, resulting in SOTA (upto 5% absolute increase in metrics)
- Jul 21-May 23 **Uniting Knowledge Distillation (KD) with Calibration**, (*Indian Institute of Technology Delhi*)
 Advisor: [Prof. Chetan Arora](#)
- Proposed a simple framework to calibrate models by distilling calibrated teachers in order to get calibrated students. These calibrated students produced have SOTA calibration as compared to previous methods.
 - Explored KD with the lens of calibration and verified that only certain teachers can distill calibrated students. Additionally, we showed that even distilling from smaller teachers can lead to calibration in larger students.
- Jul 21-May 23 **Calibration of deep neural networks**, (*Indian Institute of Technology Delhi*)
 Advisor: [Prof. Chetan Arora](#) — [CVPR 2022 ORAL]
- Proposed a novel trainable calibration method that calibrates all predicted classes (unlike other methods that focuses only on the maximum one). Introduced an auxiliary loss term that can be used in a modular manner.
 - MDCA (proposed auxiliary loss) performs better or on par with the then SOTA methods (Focal Loss, Label Smoothing, etc.) by reducing both top-class (ECE) and multi-class (SCE) miscalibration metrics.

Skills

Technical Areas	LLMs, Agents, Deep Research Systems, Long-Horizon Reasoning, Verification, Structured Retrieval, Rule-Following, Chain-of-Thought, Graph Search, Backtracking, Knowledge Distillation
Training & Alignment	Pretraining, SFT, RLHF, RLFT (Reinforcement Learning Fine-Tuning), DPO, PPO, GRPO, Reward Modeling, Trace Bootstrapping, Parameter Efficient Finetuning (PEFT), Python, PyTorch, verl
Systems & Infrastructure	Architectural Innovation, Large-Scale Distributed Data Processing, Distributed Training, HPC/Slurm, vLLM, CUDA, Ray, AWS, Azure ML, Data Pipelines, Experiment Tracking, WandB, Streamlit

Teaching

Jan-Jul'2020 **Introduction to Computer Science**, *Academic Mentor*

Extra Curricular

Sports	Secured 3rd Position (2019) and 1st Position (2020) at Inter Hostel Football Championship Won the overall General Championship, IITD (2020)
Music	Played Sitar and performed in School Orchestra for 4 years
Coding	Team Member at Algorithms and Coding Club, IITD and a regular competitive programmer
Robotics	Participated in DD ABU ROBOCON (2019) , DD ABU ROBOCON (2020) , DD ABU ROBOCON (2021) Won the International Micron UV Design Challenge against MIT and Georgia Tech (2020)
Leadership	Overall Coordinator (Robotics Club) : Supervised 150 members and revised Club structure Lead the club through collaborations with Innovation Hub for Cobotics (IHFC), Lumos Labs and Tech Analogy and organised robotics workshops and talks. Organised 6 freshmen summer projects on various fields of robotics. View Control Simulator Project