

# MUDIT DHAWAN

Pittsburgh, PA | [412-721-8074](tel:412-721-8074)

[mdhawan@cs.cmu.edu](mailto:mdhawan@cs.cmu.edu) | [linkedin.com/mudit-dhawan](https://www.linkedin.com/in/mudit-dhawan) | [mudit-dhawan.github.io/](https://mudit-dhawan.github.io/) | [Google Scholar](#)

## EDUCATION

- Carnegie Mellon University** | Master's in Machine Learning December 2025  
Current Courses: Advanced Introduction to Machine Learning (PhD), Probability and Statistics  
Independent Study with **Prof. Ruslan Salakhutdinov**: Jailbreaking Text-to-Music Diffusion Models
- IIIT Delhi, India** | Advisor: **Prof. Ponnurangam Kumaraguru** [[thesis](#)] May 2022  
BTech. in Electronics and Communication Engineering Dept Rank: 7/81

## EXPERIENCE

- Microsoft Research Lab India** | Research Fellow July 2022 - July 2024  
*Extreme Classification (XC) Group, Advisor: Dr. Manik Varma*
- Deployed a novel reformulation of Query AutoComplete as an XC task on Bing AI Chat and Search platforms. Led to **1% gain in Click Through Rate (CTR)** in production. On Bing AI Chat Platform, it led to an **8% increase in Keystrokes Saved Per Query**, and **6% increase in acceptance rate** for long queries [publication: [ICLR'24](#)].
  - Highlighted in-efficiency of current pair-wise rankers in recommendation systems and devised an algorithm with **30x lower latency** than SOTA theoretical baseline, **10% more accurate** than production system on an offline set, and led to **0.5% increase in Impression Yield Click Yield** in en-markets in online A-B tests.
  - Executed a novel scoring method to use GPT-4 as a large scale oracle to de-noise hard-negatives for retriever training. Led to **30x increase in throughput**, reduced API costs and trained retriever led to **0.6% absolute gain in clicks** during **online A-B tests** in English markets.
- IIIT Hyderabad** | Research Assistant August 2019 - June 2022  
*Precog Research Group, Advisor: Prof. Ponnurangam Kumaraguru*
- Proposed novel algorithms for multimodal fake news detection [publications: [MMAsia'21](#), [SNAM'24](#)].
  - Introduced **Multi-Task Learning Framework** for Bail prediction (**2% more accurate** than SOTA) with an auxiliary extractive summarization task for grounding predictions in Hindi documents [publication: [ACL Findings'22](#)].
  - Collaborated and developed first version of **multimodal search** engine with huggingface and milvus.io framework for Factly serving **thousands** of users worldwide [[LinkedIn Post](#), [Website](#)].
  - Formulated a Weighted sampling strategy for training an XLM model for English-Hindi **code mixed data**. Led to **15% decrease in perplexity** on eval set, and **2% increase in F1 score** on POS Tagging task on LINC Dataset [[code](#)].
- Oracle NetSuite** | Software Development Intern May 2021 - July 2021
- Designed and refactored code to implement **Facade design pattern** for a system to account for variance in multi-currency ERP systems to improve platform stability.
  - Developed a suite of Unit tests utilizing Mockito and PowerMockito for automated testing with **94% code coverage**.

## PUBLICATIONS

- Accurate and Efficient Cross-encoders for Ranking** *NeurIPS ENLSP'24, EMNLP WiNLP'24*  
B. Paliwal\*, D. Saini\* [M. Dhawan](#), et.al.
- Enhancing Tail Performance in Extreme Classifiers by Label Variance Reduction** *ICLR'24*  
A. Buvanesh, R. Chand, J. Prakash, B. Paliwal, [M. Dhawan](#), et.al.
- GAME-ON: Graph Attention NeTwork Based Multimodal Fusion for Fake News Detection** *SNAM'24*  
[M. Dhawan](#)\*, S. Sharma\*, A. Kadam, R. Sharma, P. Kumaraguru
- HLDC: Hindi Legal Documents Corpus** *ACL Findings'22*  
A. Kapoor, [M. Dhawan](#), et.al.
- Inter-modality Discordance for Multimodal Fake News Detection** *ACM MMAsia '21*  
S. Singhal, [M. Dhawan](#), R.R. Shah, P. Kumaraguru

## SKILLS

**Languages:** Python, Java, C/C++, BASH; **Software:** PyTorch, Tensorflow, HuggingFace, Keras, MATLAB  
**ML/DL Skills:** Retrieval Augmented Generation (RAG), Multimodal ML, Large Language Models (LLMs), Natural Language Processing (NLP), Diffusion, Computer Vision, Audio Modelling, Distillation, Retrieval, Search, Recommendation and Ranking, Extreme Classification (XC), Generative Models, Data Structures and Algorithms.

## SELECTED PROJECTS

---

### ***Jailbreaking Text-to-Music Diffusion Models***

Sep'24 - Present

*Advisors: Prof. Ruslan Salakhutdinov — CMU in collaboration with SonyAI*

- Designing a prompt-tuning technique for controllable generation from text-to-music diffusion models (open and closed-source).
- Experimenting with prompting techniques for model jailbreaking, textual inversion, style transfer and conducting research on memorization of training data in diffusion models.

### ***Multi-Intent Session-Based Recommendation Systems using LLMs***

Aug'23 - Present

*Advisors: Dr. Amit Sharma, Dr. Manik Varma— Microsoft Research*

- Implemented and performed task-specific distillation using Low-Rank Adaptation of GPT-4 to increase **inference speedup by ~ 100x** by using smaller LLMs with minimal loss in relevance and diversity of predictions.
- Benchmark open source encoder and decoder based models along the axes of diversity, and combination of intents in session based recommendation using GPT-4 as a pseudo-oracle to understand human preferences.

### **Query Auto-Complete (QAC) using XC and Statistical Language Models**

Aug'22 - Jul'23

*Advisors: Dr. Yashoteja Prabhu, Dr. Manish Gupta, Dr. Manik Varma — Microsoft Research*

- Proposed a novel reformulation of QAC as an XC task, which led to significant improvement in tail queries over NLG based solutions. To overcome the sparsity of training data based on clicks, added popularity weighted augmentation to vanilla XC loss to improve performance on tail and torso prefixes. Released an XC style AutoComplete dataset with less lexical overlap between query and label [ICLR'24].
- Improved over XC based QAC solutions by replacing suffix classifiers with small localized statistical language models to improve performance on head and torso queries.
- Utilized Kneser-Ney smoothing and weighted finite state transducer (FST) to represent the n-gram language model over queries containing a cluster of suffixes to better memorize frequency based patterns in the data. This framework led to a **6% and 5% increase in SR@10 and MRR@10** respectively, on publicly available dataset over XC baseline.

### **Multimodal Fake News Detection [thesis]**

Jul'20 - Jan'22

*Advisors: Prof. Ponnurangam Kumaraguru, Prof. Rajiv Ratn Shah, Prof. Rajesh Sharma— PreCog Research Group*

- Designed a novel graph framework that allows for granular interactions across (inter)- and within (intra)- modalities to fuse them early in the framework, decreasing information loss. The framework utilized scene graph for images and transformer representations for text nodes with graph attention layer, and outperformed sota by ~11% on publicly available dataset, with ~91% fewer parameters than the best comparable baseline. [SNAM'24] [🔗]
- Proposed and implemented an inter-modality discordance based fake news detection framework with sequential encoders based on the hypothesis that fabrication of either modality will lead to dissonance between the modalities. The proposed method increased sota by an average F1-score of 6.3% on multiple datasets. [ACM MMAAsia'21] [🔗]
- Open-sourced well-documented code of Multimodal Fake News detection frameworks to make previous state-of-the-art approaches in the field more accessible [🔗]

## HONORS AND AWARDS

---

Granted a **US Patent** for optimized ranking model: *CROSS-JEM: Cross-encoder Joint Efficient Modeling for ranking in large-scale search and recommendation systems*.

Nominated for **poster presentation** of ACM MMAAsia'21 work at **ARCS 2022, ACM India** [link].

Selected for **Visiting Student Research Programme** at **India Connect@NTU** [link].

## OPEN SOURCE PROJECTS

---

**Digital Image Processing Concepts** [🔗]: Code for Geometric Transformation, Bi-Linear Interpolation, Histogram matching and equalization, Constrained Least Squares Filtering.

**Computer Vision Concepts** [🔗]: Code for 2D Convolution for Edge detection, Spatial Pyramid Pooling, Semantic segmentation, Circular Bounding Boxes, Super-pixel saliency, Interactive saliency, Background subtraction from video.

## REFERENCES

---

**Dr. Manik Varma**, *Distinguished Scientist & Vice President, Microsoft Research, India* [🔗]

[manik@microsoft.com](mailto:manik@microsoft.com)

**Prof. Ponnurangam Kumaraguru**, *Professor, IIT Hyderabad, India* [🔗]

[pk.guru@iiit.ac.in](mailto:pk.guru@iiit.ac.in)

**Dr. Manish Gupta**, *Principal Applied Researcher, Microsoft India R&D Pvt. Ltd.* [🔗]

[gmanish@microsoft.com](mailto:gmanish@microsoft.com)

**Dr. Yashoteja Prabhu**, *Senior Researcher, Microsoft Research, India* [🔗]

[yprabhu@microsoft.com](mailto:yprabhu@microsoft.com)