

# RISHAB BALASUBRAMANIAN

Email : rishab.edu@gmail.com

Mobile : +1-541-250-7330

Github: <https://github.com/rishabbala>

Website: <https://rishabbala.github.io/>

Linkedin: <https://www.linkedin.com/in/rishab-bala-b01110142/>

## ABOUT ME

---

Second year PhD student in Computer Science at Virginia Tech (advised by Dr. Tu Vu), specializing in efficient model development and alignment. My research focuses on data selection for pre-training, model interpretability, distillation and merging, and architectures for scalable and reliable AI.

## EDUCATION

---

- **Virginia Tech** Blacksburg, VA  
*Ph.D in Computer Science (Advisor: [Tu Vu](#))* Aug 2024 - June 2029 (Expected)
- **Oregon State University** Corvallis, OR  
*M.S in Computer Science & Artificial Intelligence (Advisor: [Huazheng Wang](#), Thesis)* Sep 2021 - May 2024
- **National Institute of Technology** India  
*B.Tech in Instrumentation And Control Engineering (Thesis)* Aug 2016 - May 2020

## PUBLICATIONS AND PREPRINTS

---

- **EMNLP 2025:** Pin-Jie Lin, **Rishab Balasubramanian**, Fengyuan Liu, Nikhil Kandpal, and Tu Vu. Efficient model development through fine-tuning transfer, 2025
- **ICML 2024:** **Rishab Balasubramanian**, Jiawei Li, Prasad Tadepalli, Huazheng Wang, Qingyun Wu, and Haoyu Zhao. Adversarial attacks on combinatorial multi-armed bandits. In *Forty-first International Conference on Machine Learning*, 2024
- **TMLR:** Zichen Wang, **Rishab Balasubramanian**, Hui Yuan, Chenyu Song, Mengdi Wang, and Huazheng Wang. Adversarial attacks on online learning to rank with stochastic click models. *Transactions on Machine Learning Research*, 2024
- **NeurIPS 2024:** Zeyu Zhang, Yi Su, Hui Yuan, Yiran Wu, **Rishab Balasubramanian**, Qingyun Wu, Huazheng Wang, and Mengdi Wang. Unified off-policy learning to rank: a reinforcement learning perspective. *Advances in Neural Information Processing Systems*, 36, 2024
- **arXiv:** **Rishab Balasubramanian** and Kunal Rathore. Contrastive learning for object detection. *arXiv preprint arXiv:2208.06412*, 2022
- **IROS 2021:** **Rishab Balasubramanian**, Lifeng Zhou, Pratap Tokekar, and PB Sujit. Risk-aware submodular optimization for stochastic travelling salesperson problem. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4720–4725. IEEE, 2021
- **AIAA 2021:** **Rishab Balasubramanian** and Sujit PB. A cooperative framework for autonomous landings of quadrotors using vision on a moving ugv. In *AIAA Scitech 2021 Forum*, page 1880, 2021

## RESEARCH EXPERIENCE

---

- **Efficient Model Development Through Model Diffing** Virginia  
*Research with Dr. Tu Vu (VT)* 2024 - Present
  - Working on efficiently transferring post-training updates between model versions via model diffing.
  - Utilizing model weights and steering vectors for interpretability across models - accepted to (**EMNLP 2025**)

- Adversarial Attacks on Multi-Armed Bandits** Oregon  
*Research with Dr. Huazheng Wang (OSU)* 2022 - 2024
  - o Led the work to design a novel adversarial attack algorithm on combinatorial multi-armed bandits (CMAB)
  - o Showed the theoretical difficulty in attacking CMAB environments in white-box and black-box settings, and provided the *first ever* algorithm to attack CMAB instances with sublinear cost.
  - o Reformulated online learning-to-rank (OLTR) as an offline reinforcement learning problem.
  - o Three papers accepted to ([ICML 2024](#)) ([Neurips 2023](#)) and ([TMLR](#)).
- Block-wise Trainable Neural Networks ([code](#))** Remote  
*Research with Dr. Beidi Chen (Meta)* 2022 - 2023
  - o Developed a method to train neural network architectures sequentially to reduce training time, memory consumption, and FLOPs.
  - o Enforced gradient blocking and data pruning methods to further improve efficiency, reducing training time by  $\sim 30\%$  while limiting drop in test accuracy to  $< 2\%$ .
  - o Tested on a variety of networks including Transformers, ResNets and VGG models
- 3D Reconstruction from Endoscopy Images** Remote  
*Research Scientist at [EndovisionAI](#)* 2021
  - o Trained an encoder-decoder model using transformation consistency losses for unsupervised depth estimation from RGB endoscopy images.
  - o Created an API for visualizing 3D pointclouds from the output depthmap.
- Risk Averse Travelling Salesman Problem** Remote  
*Research with Dr. Sujit and Dr. Tokekar* 2020
  - o Developed a greedy algorithm to maximize the Conditional Value at Risk (CVaR) using a submodular function.
  - o We further prove that the algorithm has a polynomial runtime lower than current SOTA and an approximation factor which is proportional to the optimal solution.
  - o Finally, we also show that the suboptimality gap of the solutions is bound by a constant factor.
  - o Paper accepted to ([IROS 2021](#))

## TEACHING EXPERIENCE

---

- TA for CS 5624: Natural Language Processing at Virginia Tech (Spring 2025)
- TA for CS 5814: Machine Learning at Virginia Tech with Dr. Chandan Reddy (Fall 2024)
- TA for AI 534: Machine Learning at Oregon State University with Dr. Xiaoli Fern (Fall 2023)
- Created material for CS 514: Algorithms at Oregon State University with Prof. Samina Ehsan (Spring 2022)
- TA for CS 325: Analysis of Algorithms at Oregon State University with Prof. Umma Reddy (Winter 2022)
- TA for ENGR 201: Electrical Fundamentals I at Oregon State University with Dr. Pallavi Dhagat (Fall 2021)
- Mentor for freshmen in the last two years of undergraduate