Robik Shrestha

Research Assistant/Ph.D. candidate

(585) 766-5549 | robikshrestha@gmail.com | linkedin.com/in/robikshrestha
github.com/erobic | https://tinyurl.com/ydk7u8m2 | robikshrestha.com

SUMMARY

A Ph.D. candidate with 5+ years of research experience on scrutinizing and improving vision and language systems. The overarching goal of my research is to build deep learning systems that are right for the right reasons, thus behaving reliably when deployed in the wild.

My publications reveal failure modes of "state-of-the-art" deep learning systems and provide elegant solutions to improve the status quo. I am currently seeking research and engineering positions on large-scale deep learning systems.

Research Interests: artificial intelligence, deep learning, computer vision, natural language processing (NLP), vision and language, visual question answering, bias and fairness, continual/lifelong machine learning, causality

WORK AND RESEARCH EXPERIENCE

Ph.D. Candidate

Rochester Institute of Technology | Rochester, NY | Aug 2017 – Present

- Revealed critical issues with state-of-the-art deep learning systems: they can be right for the wrong reasons (ACL '20, NeurIPS '20) and fail on realistic types/levels of distribution shifts (WACV '22)
- Published more elegant alternatives that generalize better (ECCV '22, ACL '20, CVPR '19)
- Provided a fresh perspective on handling distribution shifts by applying Occam's razor to architecture design (OccamNets, ECCV '22)
- Enabled online lifelong learning by developing a model based on hippocampal indexing theory (REMIND, ECCV '20)
- Pioneered VQA models that generalize to natural, synthetic and chart domains (CVPR '19, WACV '20)

Advisor: Dr. Christopher Kanan

Research Assistant

SRI International | Rochester, NY | Apr 2021 – Present

- Broadened the range of cognitive skills tested for VQA models by creating a storybook dataset based on Bloom's Taxonomy
- Assessed if VQA systems can: remember basic concepts, apply commonsense knowledge, reason under new situations, and imagine new storylines given prior context or prompts

Team: Dr. Yunye Gong, Dr. Michael Cogswell, Dr. Christopher Kanan, Dr. Ajay Divakaran

Research Intern

Adobe Inc | **Rochester, NY** | **May 2021 – Oct 2021**

- Trained vision and language embeddings that are sensitive to fine-grained conceptual differences
- Developed contrastive learning methods with algebraic constraints on atomic and compositional concepts
- Built a test set with small conceptual differences to measure fine-grained sensitivity

Mentors: Dr. Kushal Kafle, Dr. Scott Cohen, Dr. Zhe Lin

Research Assistant

ITEL Laboratories | **Rochester, NY** | **Jul 2018 – Mar 2021**

- Helped move from a manual floor material claims processing system to a semi-automated system that processes 85% of images automatically
- Developed the deep learning model and client-side app for floor material classification system
- Improved bias resilience, model calibration, and out-of-distribution detection to ensure reliable semi-automated processing

The project partially funded my Ph.D. research at RIT

Full Stack Developer

Viveka Health | **Lalitpur, Nepal** | **Feb 2014 - Apr 2016**

- First Hire!
- Helped save $1M+ for 25K+ lives by developing a fraud detection engine for U.S. healthcare claims
- Onboarded new clients and implemented the core fraud detection and payment consolidation engine. Designed the frontend too.
- Interviewed and hired backend + frontend personnel to help the company grow

---

Software Engineer

Yomari Incorporated | **Lalitpur, Nepal** | **Nov 2012 - Feb 2014**

- Developed business intelligence solutions, including ETL scripts and analytics reports for Nepal telecom and international retailers

INVITED TALKS AND GUEST LECTURES

- *"OccamNets: Mitigating Dataset Bias by Favoring Simpler Hypotheses"*

  [Oral Presentation], European Conference on Computer Vision (ECCV) 2022, Tel Aviv, Israel

- *"Dataset Bias in Vision and Language Tasks: Problems and Potential Solutions"*

  [Invited Talk], Center for Human Aware AI (CHAI) 2021, Rochester, NY

- *"Dataset Bias and Bias Mitigation Techniques"*

  [Guest Lecture], Deep Learning for Vision 2020/21, RIT

PUBLICATIONS

- **Robik Shrestha**, Kushal Kafle, Christopher Kanan. *"OccamNets: Mitigating Dataset Bias by Favoring Simpler Hypotheses."* European Conference on Computer Vision (2022) (Oral Presentation, Top 2.7%)

- **Robik Shrestha**, Kushal Kafle, Christopher Kanan. *"An investigation of critical issues in bias mitigation techniques."* IEEE/CVF Winter Conference of Applications on Computer Vision (2022)

- Usman Mahmood, **Robik Shrestha** et al. *"Detecting Spurious Correlations With Sanity Tests for Artificial Intelligence Guided Radiology Systems."* Frontiers in Digital Health (2022)

- Damien Teney, Kushal Kafle, **Robik Shrestha** et al. *"On the Value of Out-of-Distribution Testing: An Example of Goodhart's Law."* Neural Information Processing Systems (2020)

- Tyler Hayes*, Kushal Kafle*, **Robik Shrestha***, Manoj Acharya, and Christopher Kanan. *"REMIND Your Neural Network to Prevent Catastrophic Forgetting."* European Conference on Computer Vision (2020). (* = equal contributions)

- **Robik Shrestha**, Kushal Kafle, and Christopher Kanan. *"A negative case analysis of visual grounding methods for VQA."* Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (2020)

- Kushal Kafle, **Robik Shrestha**, and Christopher Kanan *"Answering questions about data visualizations using efficient bimodal fusion."* The IEEE Winter Conference on Applications of Computer Vision (2020)

- **Robik Shrestha**, Kushal Kafle, and Christopher Kanan. *"Answer them all! toward universal visual question answering models."* Proceedings of the IEEE conference on computer vision and pattern recognition (2019)

- Kushal Kafle, **Robik Shrestha**, and Christopher Kanan. *"Challenges and prospects in vision and language research."* Frontiers in Artificial Intelligence (2019).

PATENT APPLICATIONS

- C. Kanan, T.L. Hayes, K. Kafle, and **R. Shrestha**. "*Method for training parametric machine learning systems*", January 28, 2021. US Patent App. 16/938,035

SKILLS

- **Proficient in** PyTorch, PyTorch-Lightning, Python, Scikit-Learn, Git, Shell Scripting
- **Also knows about** Tensorflow, Java, Javascript, MTurk, Matlab, AWS, C/C++, Jira

EDUCATION

Ph.D. in Imaging Science | Rochester Institute of Technology | Aug 2017 - Current

Relevant Courses: Deep Learning Systems for Vision, Image Processing and Computer Vision, Human Visual System, Principles of Statistical Data Mining

B.E. in Computer Engineering | Institute of Engineering, Nepal | Oct 2008 - Dec 2012

Relevant Courses: Image Processing and Pattern Recognition, Artificial Intelligence

Full Merit-Based Scholarship awarded to 200 STEM students across Nepal