

SAVYA KHOSLA

Urbana, IL • savyak2@illinois.edu • [LinkedIn](#) • [GitHub](#) • [Google Scholar](#) • [Personal Webpage](#)

EDUCATION

University of Illinois Urbana-Champaign

MS in Computer Science (CGPA: 4.0 / 4.0)

Urbana, IL

Aug 2022 - May 2024

- Thesis advisor: Prof. Derek Hoiem
- Activities: Graduate Teaching Assistant and Graduate Research Assistant

Delhi Technological University

B.Tech. in Computer Engineering (CGPA: 9.40 / 10.0)

New Delhi, DL

Aug 2017 - July 2021

- Awards: Received Commendable Research Award and INR 50,000 for noteworthy contributions to machine learning research

RESEARCH EXPERIENCE

University of Illinois Urbana-Champaign

Research Assistant (Guide: Prof. Derek Hoiem)

Urbana, IL

Jan 2023 - Present

- Working on learning efficient representations for videos of arbitrary lengths

Allen Institute for AI

Research Intern

Seattle, WA

Oct 2022 - Dec 2023

- Worked on a memory-augmented multimodal encoder for understanding videos ranging from a few seconds to tens of minutes
- Contributed to Unified-IO 2, an instruction-following model that can parse and generate multimodal data and perform 120+ tasks

National University of Singapore

Research Assistant (Guide: Prof. Kenji Kawaguchi)

Remote

Apr 2022 - Aug 2022

- Developed robust active learning algorithm for handling heteroskedastic noise, resulting in 10% accuracy boost over baselines
- Demonstrated 15% accuracy improvement in other state-of-the-art algorithms by incorporating a simple self-supervised approach

Mila - Quebec AI Institute

Research Intern (Guide: Prof. Yoshua Bengio)

Montreal, QC

Apr 2021 - Nov 2021

- Demonstrated catastrophic failure of uncertainty-based active learning algorithms by proposing 3 heteroskedastic data distributions
- Proposed adversarial training method that gives 48% reduction in error rate on clean data while preserving adversarial robustness

Delhi Technological University

Undergraduate Researcher (Guide: Prof. Anil Singh Parihar)

New Delhi, DL

Apr 2021 - Nov 2021

- Worked on improving object recognition systems in the presence of adversaries like occlusion and blurriness
- Used image-based representation of malware binaries and leveraged ensembling to develop SOTA model for malware classification

INDUSTRY EXPERIENCE

Google

Software Engineer

Bangalore, KA

Aug 2021 - Mar 2022

- Improved Google Search's web ranking infrastructure using deep learning for better multimodal document understanding
- Enhanced precision and recall in salient entity extraction from webpages by transitioning from traditional ML methods to LLMs

Google

Software Engineering Intern

Bangalore, KA

May 2020 - Jul 2020

- Initiated the development of MuRIL, a BERT-based multilingual language model for 17 Indian dialects
- Achieved 10.42% F1 improvement in sentiment analysis and 9.87% in named entity recognition for Indian languages

Cadence Design Systems

Python Developer Intern

Noida, UP

Dec 2018 - Jan 2019

- Streamlined complex multi-step process of fetching file revisions from 2 version control systems to a single bash command

TEACHING EXPERIENCE

CS 445: Computational Photography

Teaching Assistant (Instructor: Prof. Derek Hoiem)

Urbana, IL

Aug 2023 - Dec 2023

CS 225: Data Structures and Algorithms with C++

Teaching Assistant (Instructor: Prof. Carl Evans and Prof. Brad Solomon)

Urbana, IL

Aug 2022 - May 2023

PUBLICATIONS & PREPRINTS

(† denotes alphabetical order, * denotes equal contribution)

1. **Unified-IO 2: Scaling Autoregressive Multimodal Models with Vision, Language, Audio, and Action** [Link](#)
Jiasen Lu*, Christopher Clark*, Sangho Lee*, Zichen Zhang*, [Savya Khosla](#), Ryan Marten, Derek Hoiem, Aniruddha Kembhavi
arXiv:2312.17172 [cs.CV], 2023
2. **Survey on Memory-Augmented Neural Networks: Cognitive Insights to AI Applications** [Link](#)
[Savya Khosla](#)*, Zhen Zhu*, Yife He*
arXiv:2312.06141 [cs.AI], 2023
3. **Understanding and Improving Neural Active Learning on Heteroskedastic Distributions** [Link](#)
[Savya Khosla](#), Chew Kin Whye, Jordan T. Ash, Cyril Zhang, Kenji Kawaguchi, Alex Lamb
European Conference on Artificial Intelligence (ECAI), 372:1248-1255, 2023
4. **Interpolated Adversarial Training: Achieving Robust Neural Networks Without Sacrificing Too Much Accuracy** [Link](#)
Alex Lamb, Vikas Verma, Kenji Kawaguchi, Alexander Matyasko, [Savya Khosla](#), Juho Kannala, Yoshua Bengio
Neural Networks, 154:218–233, 2022
5. **S-DCNN: Stacked Deep Convolutional Neural Networks for Malware Classification** [Link](#)
Anil Singh Parihar*, Shashank Kumar*, [Savya Khosla](#)*
Multimedia Tools and Applications, 81:30997–31015, 2022
6. **Catastrophic Failures of Neural Active Learning on Heteroskedastic Distributions** [Link](#)
[Savya Khosla](#), Alex Lamb, Jordan Ash, Cyril Zhang, Kenji Kawaguchi
NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications, 2021
7. **AE-DCNN: Autoencoder Enhanced Deep Convolutional Neural Network For Malware Classification** [Link](#)
Shashank Kumar*, [Savya Khosla](#)*, Shivangi Meena, Anil Singh Parihar
International Conference on Intelligent Technologies (CONIT), 2021
8. **MuRIL: Multilingual Representations for Indian Languages** [Link](#)
Simran Khanuja, Diksha Bansal†, Sarvesh Mehtani†, [Savya Khosla](#)†, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, Partha Talukdar
arXiv:2103.10730 [cs.CL], 2021
Media Coverage: [Economic Times](#), [Indian Express](#), [Google AI Blog](#)

PROJECTS

Occluded Facial Expression Recognition

- An occluded facial expression recognition framework that leverages non-occluded images as privileged information
- The technique rendered an average gain of 3.90% over the baseline for 3 standard benchmarking datasets
- Technical stack used: Learning Using Privileged Information, Convolutional Neural Networks, TensorFlow

Image Captioning

- A CNN and RNN-based model for generating a textual description of an image based on the objects and actions in it
- Technical stack used: Convolutional Neural Networks, Recurrent Neural Networks, Beam Search Algorithm, Keras, Python

Text to Image

- A conditional GAN for synthesizing 256x256 dimensional photo-realistic images given textual descriptions
- Technical stack used: Conditional Generative Adversarial Networks (used the StackGAN architecture), TensorFlow, Python

AgroAI

- A group project to build an unbiased platform for farmers to predict the quality and price of the crops
- Presented this project in Google's Explore ML Bootcamp
- Technical stack used: React, NodeJS, Mongo, Flask (Python)

SKILLS

Languages: Python, C++, C, JavaScript, Bash

Frameworks: PyTorch, TensorFlow, JAX, Flax, OpenCV, GradIO

Tools: Git, Visual Studio, Google Cloud Platform

Others: Data Structures, Algorithms, Machine Learning, Computer Vision, NLP, Multimodal Learning, Data Handling

COURSES & CERTIFICATIONS

Graduate Courses

- CS 598: Vision by Prof. Svetlana Lazebnik
- CS 588: Autonomous Vehicle System Engineering by Prof. David Alexander Forsyth
- CS 543: Computer Vision by Prof. Svetlana Lazebnik
- CS 445: Computational Photography by Prof. Derek Hoiem
- CS 444: Deep Learning for Computer Vision by Prof. Svetlana Lazebnik
- CS 410: Text Information System by Prof. ChengXiang Zhai

Relevant Undergraduate Courses

- CO 201: Data Structures
- CO 202: Database Management System
- CO 203: Object-Oriented Programming
- CO 206: Algorithm Design and Analysis
- CO 304: Artificial Intelligence
- CO 404: Data Warehousing and Data Mining
- CO 407: Distributed Systems
- CO 414: Big Data Analytics
- CO 423: Swarm and Evolutionary Computing
- IT 420: Computer Vision

Online Courses & Certifications

- Deep Learning Specialization by Andrew Ng
 - Neural Networks and Deep Learning
 - Improving Deep Neural Networks: Hyperparameter tuning, Regularization and Optimization
 - Structuring Machine Learning Projects
 - Convolutional Neural Networks
 - Sequence Models
- Machine Learning by Stanford University (CS229 Lectures by Andrew Ng)
- Introduction to TensorFlow for Artificial Intelligence, Machine Learning, and Deep Learning by deeplearning.ai
- C++ Bootcamp by Coding Blocks
- Competitive Programming Bootcamp by Coding Blocks
- Machine Learning Master Course by Coding Blocks