

Akrit Mohapatra

2714 Wellesley Court, Blacksburg, VA 24060
+1 (864)-517-0130
akrit@vt.edu
<https://filebox.ece.vt.edu/~akrit/>

EDUCATION

- AUG 2016 - May 2018 M.S. in COMPUTER ENGINEERING
Virginia Tech, Blacksburg
Specialization: Machine Learning, Computer Vision, Natural Language Processing
Advisor: [Prof. Dhruv Batra](#)
- AUG 2012 - May 2016 B.S. in COMPUTER ENGINEERING
Virginia Tech, Blacksburg
Dean's List - Spring 2013, Fall 2014, Fall 2015, Spring 2016

PUBLICATIONS

The Promise of Premise: Harnessing Question Premises in Visual Question Answering

Aroma Mahendru*, Viraj Prabhu*, **Akrit Mohapatra***, Dhruv Batra, Stefan Lee
(* Equal contribution)

Conference on Empirical Methods in Natural Language Processing (EMNLP), 2017

Towards Transparent AI Systems: Interpreting Visual Question Answering Models

Yash Goyal, **Akrit Mohapatra**, Devi Parikh, Dhruv Batra





Best Student Paper at International Conference on Machine Learning (ICML) Workshop on Visualization for Deep Learning, 2016

CloudCV: Large-Scale Distributed Computer Vision as a Cloud Service

Harsh Agrawal, Clint Solomon Mathialagan, Yash Goyal, Neelima Chavali, Prakriti Banik, **Akrit Mohapatra**, Ahmed Osman, Dhruv Batra

Book Chapter, Mobile Cloud Visual Media Computing
Editors: Gang Hua, Xian-Sheng Hua. Springer, 2015

RESEARCH EXPERIENCE

- | | |
|-------------|--|
| SUMMER 2017 | Research Intern at CREATIVE TECHNOLOGIES LAB,  Adobe Research
Developed a framework that enables users to be able to perform fine-grained image edits in the form of open-ended natural language commands. This introduces a layer of processing the language into sequential executable commands, adding interpretability to the framework which allows users to supervise the edits possibly through an interface that are then be executed by state of the art deep models to achieve accurate intended results. |
| SUMMER 2016 | Research Intern at MACHINE LEARNING & PERCEPTION LAB, 
Worked on interpretability of Visual Question Answering (VQA) models. Mentored students on the CloudCV project for Google Summer of Code 2016. |
| SUMMER 2015 | Research Intern at MACHINE LEARNING & PERCEPTION LAB, 
Worked on set-up and launch of Visual Question Answering (VQA) website. Created interactive D3 visualizations for the VQA dataset. Also mentored students on the CloudCV project for Google Summer of Code 2015. |
| SUMMER 2014 | Research Intern at MACHINE LEARNING & PERCEPTION LAB, 
Worked on creating and tuning models for object detection using the R-CNN algorithm. Also participated in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) . |

SKILLS

Torch, TensorFlow, PyTorch, Caffe, Python, C++, MongoDB, HTML, D3.js, Linux, Apache Spark, Docker, Hadoop, SQL, Scikit, NLTK

SELECTED PROJECTS

FALL 2016	Exploring performance on CIFAR-10 Explored performance of common ML algorithms such as Nearest-Neighbor, Naive Bayes and SVMs using different feature sets such as HOG and CNN features.
FALL 2015, SPRING 2016	Qualcomm Real-time Mosaicking with Snapdragon Developed an application that stitches incoming video streams from an aerial survey drone into a mosaicked image in real-time. The application required configuring various hardware and software integrations and all processing was done on a Snapdragon DragonBoard 410c.
SPRING 2016	Marco-Polo - Embedded Rover Project Embedded systems project consisting of two rovers made by integrating PIC32s and Zumo motors and operated using a real time operating system. One rover is completely autonomous and is tasked with detection of the second rover using IR sensors mounted on a servo. The second rover is user controlled using a PyQT interface and is used to pursue the first rover. The goal of the second rover is to "catch" the first rover while it avoids being caught thereby creating the famous Marco-Polo gaming experience. (https://www.youtube.com/watch?v=d65eue3KGOw)
FALL 2015	Exploring Nearest-Neighbor Approach on VQA Analyzing performance of the nearest-neighbor algorithm along with a consensus approach on the Visual Question Answering (VQA) dataset. https://filebox.ece.vt.edu/~akrit/Exploring_NN
FALL 2015	Content-Aware Image Resizing Implemented a version of the content-aware image resizing technique described in Shai Avidan and Ariel Shamir's paper, "Seam Carving for Content-Aware Image Resizing". This method involves computing the energy of an image using gradients which generates seams to achieve content-aware resizing. (Matlab)
FALL 2015	Image Stitching Implemented an image stitcher algorithm that computes a homography matrix and uses it to warp multiple overlapping similar images to create a single panoramic image.

SELECTED COURSEWORK

GRADUATE	Advanced Machine Learning , Advanced Computer Vision , Deep Learning, Bayesian Statistics, Convex Optimization, Data Analytics
UNDERGRADUATE	Computer Vision , Artificial Intelligence, Embedded Systems, Network Applications, Digital Design, Applied Software Design, Data Structures & Algorithms

TEACHING EXPERIENCE

FALL 2016	Intro to Computer Vision Instructor - <i>Jia-Bin Huang</i>
SPRING 2017	Microcontroller Programming and Interfacing Instructor - <i>Jason Thweatt</i>
FALL 2017	Fundamentals of Information Security Instructor - <i>Jung-Min (Jerry) Park</i>

LEADERSHIP EXPERIENCE

2016	<i>President, Indian Students' Association</i>
2014-2016	<i>Corporate Relations, IEEE, Virginia Tech Chapter</i>
2016	<i>Vice-President, Council of International Students Organizations</i>
2016	<i>Team Coordinator, Virginia Tech Cricket Club</i>