Ramprasaath (Ram) RS

Researcher · Computer Vision · Deep Learning

Ph.D Student, Virginia Tech

(+1) 434-616-0082 | ram21@vt.edu | ramprs.github.io

Research Interests

Understanding and Visualizing Deep Neural Networks for Computer Vision and Natural Language.

Education

Virginia Tech — *Blacksburg, VA, USA*

Ph.D in Computer Engineering — *Aug. 2015 - Present*

- Advised by Dr. Devi Parikh and co-advised by Dr. Dhruv Batra.

Birla Institute of Technology & Science (BITS)-Pilani — *Hyderabad, India*

Bachelor of Engineering (Honor) in Electrical and Electronics

Master of Science (Honor) in Physics — *Aug. 2010 - May. 2015*

Publications

1. **Ramprasaath R. Selvaraju**, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. "Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization." arXiv preprint arXiv:1610.02391 (2016).
2. Vijayakumar Ashwin K., Michael Cogswell, **Ramprasaath R. Selvaraju**, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. "Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models." arXiv preprint arXiv:1610.02424 (2016).
3. Chattopadhyay Prithvijit, Ramakrishna Vedantam, **Ramprasaath R. Selvaraju**, Dhruv Batra, and Devi Parikh. "Counting Everyday Objects in Everyday Scenes." arXiv preprint arXiv:1604.03505 (2016).
4. Miksik Ondrej, Vibhav Vineet, Morten Lidegaard, **Ramprasaath R. Selvaraju**, Matthias Nießner, Stuart Golodetz, Stephen L. Hicks, Patrick Pérez, Shahram Izadi, and Philip HS Torr. "The semantic paintbrush: Interactive 3d mapping and recognition in large outdoor spaces." In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 3317-3326. ACM, 2015.
5. Garg Sanyam, **Ramprasaath R. Selvaraju**, Suman Kapur, and Kunda MM Rao. "Automated colorimetric analysis in paper based sensors." In 2014 IEEE International Conference on Image Processing (ICIP), pp. 3607-3611. IEEE, 2014.
6. **Ramprasaath R. Selvaraju**, Spandana P, and Kunda MM Rao. "A novel algorithm for Image fusion and enhancement using Dual Tree Complex Wavelet Transform." In 29th National Convention on Electronics and Telecommunication Engineers at Institute of Engineers, Hyderabad, 2013

Teaching Experience

Teaching Assistant — *Virginia Tech*

Data Structures and Algorighms — *Aug. 2015 - May. 2016*

Research Experience

Virginia Tech — *VA, USA*

Undergrad Thesis working with **Devi Parikh** — *Jan. 2015 - Aug. 2015*

- Worked on building curious systems that ask Natural Language open-ended questions about an image.

University of Oxford — *Oxford, UK*

Undergrad Thesis working with **Philip Torr** and **Stephen Hicks** — *May. 2014 - Dec. 2014*

- Worked on developing an interactive augmented reality system where a carer helps the user understand the scene better through interactive labeling with laser pointer through a shared virtual environment.
- Our work was published as Oral at Computer Human Interaction Conference, CHI 2015.

Brown University — *RI, USA*

Summer Internship working with **Benjamin Kimia** — *May. 2013 - Aug. 2013*

- Worked on designing a vision based navigation system to help the blind/vision impaired people navigate indoor environments, through use of glass mounted stereo/depth haptic belt mounted IMUs.

Course Work

- Computer Vision (Intro and Adv.)
- Adv. Machine Learning
- Deep Learning for Perception
- Optimization in High-dim Spaces
- Bayesian Statistics

Skills

Programming	Python, MATLAB, C/C++
Deep Learning Frameworks	Caffe, Torch, Tensorflow
Operating Systems	Linux (Ubuntu), MacOS, Windows and Android

Projects

Where did you learn that?: Mining training examples that most influence decisions of Deep Models

Virginia Tech

With the guidance of Dhruv Batra and Stefan Lee *2016*

- The goal of this project is to understand the effect of each training example on the decision made by any Deep Neural Network.
- While our previous work, Grad-CAM explored what regions in the input image were most important for a particular prediction, here we explore where the model learned to make that decision.
- This can help us with:
  - Curriculum Learning
  - Understanding if current datasets need to be this large
  - Identifying label noise
  - Machine Teaching

Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization

Virginia Tech

With the guidance of Dhruv Batra and Devi Parikh *2016*

- Developed a Deep Neural Network Visualization technique called, Grad-CAM (Gradient-weighted Class Activation Mapping) that:
  - is class-discriminative and can make any CNN-based model interpretable
  - requires no change in architecture → no need for retraining → no compromize on accuracy
- Grad-CAM provides tools for:
  - understanding networks (eg. debugging) and instill trust in user
- Can visualize models for a variety of applications: Image Classification, Image Captioning and Visual Question Answering
- Code: `https://github.com/ramprs/grad-cam`
- Arxiv Paper: `https://arxiv.org/abs/1610.02391`
- Demo: `https://gradcam.cloudcv.org`

Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models

Virginia Tech

With the guidance of Dhruv Batra *2016*

- Traditional Beam Search explores the search space in a greedy left-right fashion – resulting in sequences that differ only slightly from each other.
- Lack of diversity in the decoded solutions is fundamentally crippling in AI problems with significant ambiguity.
- To overcome this problem, we propose Diverse Beam Search (DBS), an alternative to BS that decodes a list of diverse outputs by optimizing for a diversity-augmented objective.
- In addition to generating diverse predictions, it also helps finding better top-1 solutions.
- Code: `https://github.com/ashwinkalyan/dbs`
- Arxiv Paper: `https://arxiv.org/abs/1610.02424`
- Demo: `dbs.cloudcv.org`

Counting Everyday Objects in Everyday Scenes

Virginia Tech

With the guidance of Devi Parikh and Dhruv Batra *2015*

- The goal of this project is to count the number of occurrences of Common Everyday occurring categories in real-world scenes
- Arxiv paper: `https://arxiv.org/abs/1604.03505`

Reviewing

- Reviewer for Neural Information Processing Systems (NIPS'16)
- Reviewer for Computer Vision and Pattern Recognition (CVPR'16)

Honors

2016 **Accepted**, Deep Learning Summer School (Only 235 out of 775 applicants were accepted) *Montreal, Canada*

Extra-curricular Achievements

2016 **First Place**, Virginia Division Table-Tennis Championship *VA, USA*
2016 **Second Place**, US Mid-Atlantic Region Table-Tennis Championship *NC, USA*
2016 **Represented Virginia Tech**, US-Canada National Table-Tennis Championship *TX, USA*

References

- Dr. Devi Parikh, Assistant Professor, Virginia Tech - parikh@vt.edu
- Dr. Dhruv Batra, Assistant Professor, Virginia Tech - dbatra@vt.edu
- Dr. Philip Torr, Professor, University of Oxford - philip.torr@eng.ox.ac.uk
- Dr. Stephen Hicks, Research Fellow, University of Oxford - stephen.hicks@ndcn.ox.ac.uk
- Dr. Benjamin Kimia, Professor, Brown University - kimia@brown.edu
- Dr. KMM Rao, Deputy Director (rtd), ISRO - kmm@drkmm.com