

Satya Krishna Gorti

647-702-0349 | satyag@cs.toronto.edu | [Github](#) | [Webpage](#) | [Linkedin](#)

EDUCATION

University of Toronto

Master of Science in Applied Computing, GPA: 4/4

Toronto, ON

Aug. 2017 – Dec 2018

MSRIT

Bachelor of Engineering in Computer Science, GPA: 9.57/10

Bangalore, India

Aug. 2012 – June 2016

EXPERIENCE

Layer6 AI - Senior Machine Learning Scientist

Jan 2019 – Present

- Working on joint language-video understanding in large datasets, 1 publication.
- Worked on weakly-supervised and semi-supervised methods for action localization in videos, 2 publications.
- Led a project for vehicle image analysis of auto-insurance claims. Awarded most innovative usecase across TD bank teams, patent in progress.
- Worked on large-scale image retrieval research, 2 publications.
- Worked on predicting credit risk on mortgages and credit cards, a high-risk use-case with projected annual savings upwards of 20 million dollars, patented.
- Led the efforts for setting up distributed training functionality in home grown ML library.
- **Technologies used:** PyTorch, Python, Java, Scala, MXNet

Uber (Advanced Technologies Group) - Machine Learning Research Intern

May 2018 – Dec 2018

- Deep learning models for object tracking with LIDAR and RADAR sensors for self-driving vehicles.
- Bayesian filtering for obtaining smooth velocity estimates for tracked objects.
- **Technologies used:** PyTorch, Python, C++

VoterCircle - Software Engineer

Nov 2016 – July 2017

- Worked on a product that enables friend to friend voter outreach for US based campaigns.
- Led the efforts on building a real time engine on websockets that drives live campaign visualizations.
- Implementing fast name matching in large datasets.
- **Technologies used:** NodeJS, MongoDB, Elasticsearch, RabbitMQ, Angular

Morgan Stanley - Technology Analyst

Aug 2016 – Oct 2016

- Worked on building an API server used by many clients within the organization.
- **Technologies used:** Java, TypeScript, DB2

Morgan Stanley - Technology Analyst Intern

Jan 2016 – July 2016

- Developed an entitlement tool that enables traders to request and grant access to specific applications used within the organization.
- **Technologies used:** Java, Spring, TypeScript, Angular

TCS Innovation Labs - Research Intern

June 2015 – Aug 2015

- Worked on static program analysis. Implemented algorithms for lazily decomposing aggregates, such as arrays and records, into simpler components based on access patterns for aggregate structure identification.
- **Technologies used:** Java, Standard ML, GraphViz

RESEARCH

- **XPool: Cross-Modal Language-Video Attention for Text-Video Retrieval***: We propose a cross-modal attention model called XPool that reasons between a text and the frames of a video for text-video retrieval. We evaluate our method on three benchmark datasets of MSRVT, MSVD and LSMDC, achieving new state-of-the-art results by up to 8% in relative improvement in Recall@1.
CVPR 2022

- **Weakly Supervised Action Selection Learning in Video***: We propose Action Selection Learning (ASL), an approach to temporally localize actions in untrimmed videos using video level class labels as weak supervision. Empirically, we show that ASL outperforms leading baselines on two popular benchmarks THUMOS-14 and ActivityNet-1.2, with 12.3% and 5.7% relative improvement respectively.
CVPR 2021
- **Cross-Class Relevance Learning for Temporal Concept Localization***: We present a framework for temporal concept localization and hold state-of-the-art results on Youtube-8M dataset.
ICCV 2019 - The 3rd Workshop on YouTube-8M Large-Scale Video Understanding
- **Guided Similarity Separation for Image Retrieval**: We propose a graph convolutional network to directly encode neighbour information into image descriptors for image retrieval. We further leverage ideas from clustering and manifold learning, and introduce an unsupervised loss based on pairwise separation of image similarities.
NeurIPS 2019
- **Semi-Supervised Exploration in Image Retrieval***: A novel semi-supervised graph traversal approach for image retrieval.
CVPR 2019 - Landmark Recognition workshop
- **Text to image to text synthesis using cycle consistent adversarial networks***: Improving text to image synthesis using cycle consistent adversarial networks.
- **Gradient descent revisited via an adaptive online learning rate***: Online algorithm for learning the learning rate in stochastic gradient descent using first order and second order approximation methods and studying its effects on convex and non-convex machine learning problems.
- **ReGAN: Study on Sequence generation using REINFORCE, REBAR and RELAX gradient estimators using generative adversarial networks***: Comparative study of effectiveness of gradient estimators for text generation using generative adversarial networks.

* indicates primary authorship

TECHNICAL SKILLS

Languages: Python, Java, JavaScript, C++, Rust
Frameworks: PyTorch, Tensorflow, MXNet, Express, Angular
Databases: MongoDB, MySQL, Redis

ACHIEVEMENTS

- 1st place (Kaggle gold medal), The 3rd YouTube-8M Video Understanding Challenge. [Leaderboard](#)
- 3rd place (Kaggle gold medal), Google Landmark Retrieval 2019 challenge. [Leaderboard](#)
- Winner of Hack59, an 18-hour hackathon. Developed a web application that gives automatic feedback of products bought by customers to the respective companies based on the sentiments of their tweets on those products, May 2014
- Winner CodeMSRIT 24-hour hackathon. Developed a web application that showcases how trending a particular topic on Twitter is, based on the sentiments of tweets, March 2014.
- TCS Codevita 2015 finalist (top 15), a global algorithmic competition where over 46,000 teams were registered.
- Advanced to ACMs International Collegiate Programming Contest (ICPC) regionals - 2014 and 2015.