Saksham Singhal

Contact

Phone: +1 (412) 636-7043
E-mail: sakshams@andrew.cmu.edu
Linkedin: `https://www.linkedin.com/in/sakshamsinghal`

Education

Carnegie Mellon University, School of Computer Science, Pittsburgh, PA **Dec '17**
Master of Science, Computational Data Science **GPA: 3.96/4.0**
International Institute of Information Technology, Hyderabad, India **May '16**
Bachelor of Technology, Computer Science **GPA: 8.48/10**

Relevant Courses

CMU: Intro to Machine Learning, Machine Learning for Large Datasets, Search Engines, Adv. Multi-modal Machine Learning, Distributed Systems
IIIT H: Artificial Intelligence, Optimization Methods, Information Retrieval and Extraction, Data-warehousing and Data Mining.

Skills

Programming Languages: Python, C, C++, Java, R
Frameworks and Tools: MySQL, AWS, Apache Spark, Apache Hive, Elastic Search, MongoDB

Experience

Machine Learning Intern, Adobe San Jose, United States **Summer '17**
Interned with the **Adobe Search and Sensei** team. Built federated query auto-completion micro-service from scratch which will become a part of Adobe Universal Search Service. Also, built a source affinity model for query classification to understand intent of the query.

Research Intern, IIIT Hyderabad, India **Summer '15**
Interned with the **Center for Data Engineering** lab under Prof. Vikram Pudi. Developed the hypothesis for concept-based communities in citation networks for analyzing similarity in nodes. This work was accepted at **IEEE International Conference on Data Mining 2015 (ICDM)**

Projects

Attention Routing for Fraud Detection **Spring '17**
Built a scalable visualization tool for directing attention towards fraudulent nodes in graphs. Tractability of the tool extends from detecting simple outliers to even micro-clusters of suspicious activity.

Visual Question Answering System **Spring '17**
Improving the state-of-the-art Visual-QA models using deep reinforcement learning over the glimpses in the image. Improved the performance on unseen category of questions using zero shot learning.

Learning to Rank **Fall '16**
Built a complete search engine over ClueWeb09 dataset from scratch. Incorporated additional features like query expansion using pseudo relevance feedback, SVM model for feature based re-ranking retrieval and diversity based retrieval model

Predicting Crime in Pittsburgh **Fall '16**
Modelled crime events in Pittsburgh using Gaussian Processes and ETAS model from both spatial and temporal perspective to predict the amount of crime in an arbitrary neighborhood.

Dynamic Memory Allocator **Summer '16**
Implemented a dynamic memory allocator for C (64 bit systems) using heap's virtual address space. Used segregated free lists to maximize throughput and utilization.

Ranking in Citation Network and Modelling Topic Evolution **Spring '15**
Implemented ranking for research articles, conferences and authors to model topics and their evolution over time on DBLP dataset. Incorporated frequent item-set mining for topic modelling using only the abstract of the article.