

## Work Experience

---

**Data Scientist Intern** **Adobe Systems** **Summer 2018**

Machine Learning Core Services

- Developed models for anomaly detection and cause analysis for Creative Cloud marketing data.

**Machine Learning Engineer** **Snapdeal** **Dec 2015 – Mar 2017**

Multimedia Research Group

- Developed a text-based classification system for products in the catalog to predict categories and subcategories.
- Used deep learning models to create an image quality control system that classifies and filters images based on blurriness, image quality and presence of objectionable content.
- Built an Apache Spark pipeline for feature extraction, clustering and nearest neighbor computation of catalog images, reducing the processing time by over 50%.
- Implementations were done using Python, Caffe, Scikit-learn, Keras and OpenCV.

**MTS, Software Development** **Adobe Systems** **Aug 2013 – Nov 2015**

- Core member of the team that created Adobe Captivate Prime, a cloud based learning management system.
- Designed and developed a content storage server, which imports, converts and stores content and metadata.
- Used Java, Spring Framework, JavaScript for implementations.

**R&D Engineer** **Tejas Networks** **Aug 2012 – Aug 2013**

- Implemented MPLS-TP protocol on network switches.
- Used C, C++ for implementations.

## Education

---

**Amherst, MA** **University of Massachusetts** **Sept 2017 – May 2019**

- M.S. in Computer Science.
- Coursework: Machine Learning; NLP; Reinforcement Learning; Information Retrieval; Algorithms.

**Calicut, India** **National Institute of Technology** **July 2008 – May 2012**

- B.Tech. in Computer Science & Engineering.
- Coursework: Artificial Intelligence; Image Processing; Data Structures and Algorithms.

## Research & Projects

---

- **Information Extraction & Synthesis Lab** (Sept 2017 - Present). Working with **Prof. Andrew McCallum** on Structured Prediction Energy Networks (SPENs), a novel architecture for structured prediction, where the prediction is done by energy minimization using neural networks and gradient based methods. Currently working on using SPENs for unsupervised tasks in NLP such as dependency parsing.
- **Question Answering on SQuAD** (Fall 2017). Developed a machine learning model for question answering using a bi-LSTM with self-attention mechanism and answer pointers. Achieved an F1 score of 50.7 on the Stanford Question Answering Dataset (SQuAD).
- **Quasi-random methods for hyper-parameter optimization** (Fall 2017). Investigated quasi-random approaches such as Poisson Disc Sampling for hyper-parameter optimization as an alternative to grid search. These methods provide a better sampling of the important hyper-parameters than grid search, while giving a better coverage of the hyper-parameter space than pure random search.

## Languages and Technologies

---

- Python; Java; C++; JavaScript
- Tensorflow; PyTorch; Caffe; Keras; Scikit-Learn; OpenCV; Apache Spark