Chintan Shah

• shah.ch@northeastern.edu • 857-399-7418 • LinkedIn • Github • chnsh.me

EDUCATION

Khoury College of Computer Sciences, Northeastern University September 2018 - Present, Boston, MA
Master of Science in Computer Science (**GPA: 4.0**) Expected December 2020
Courses: Algorithms, PDP, Machine Learning, NLP, Causal Machine Learning, Advanced ML, Reinforcement Learning

EXPERIENCE

Deep Learning Research Intern, PathAI June 2020 - September 2020, Boston, MA

- Delivered deep convolutional networks for cancerous tissue classification in histopathology images in PyTorch.
- Researched meta-learning approaches to improve out-of-distribution model generation performance by 50%.
- Augmented deep learning pipelines to support higher-order gradients for bi-level optimization techniques.

Machine Learning Research Assistant, Northeastern University November 2019 - Present, Boston, MA

- Won full funding to research the effectiveness of GNNs in locating the source of an epidemic (P0) over a network.
- Led the design of **"model-free"** graph neural network (GNN) architectures to identify P0 and achieved a **100x speed-up in inference time** and **improved accuracy by 20%** in comparison to current state-of-the-art methods.
- Identified the **theoretical** bounds on prediction accuracy and established the importance of early contact-tracing.
- Delivered a data-driven talk on the research outcome at NetSCI 2020! Preprint available here: **paper**

Machine Learning Intern, Apprentice Health (YC 18) May 2019 - December 2019, Boston, MA

- Designed evolutionary algorithms for optimizing doctor schedules to **reduce patient wait-time by over 40%**.
- Developed a deep learning model to learn a **permutation-invariant representation** of the in-clinic state.
- Predicted expected patient wait in real-time and at scale improving **patient satisfaction scores by over 12%**.
- Slashed infrastructure **costs by 30%**, increased hardware **utilization to over 95%**, and **reduced model training time by 70%** by architecting high-throughput, distributed machine learning pipelines using Kubernetes on AWS.

Senior Software Engineer, Machine Learning, Media.net June 2017 - June 2018, Mumbai, India

Mentored a team of 4 software engineers in an entrepreneurial environment to pitch, design, develop and then lead to completion product initiatives in the area of algorithmic revenue optimization, ad-text generation, automated campaign creation, statistical time-series forecasting, anomaly detection, and streamlined high-throughput data pipelines.

- Drove research to develop time-series forecasting systems for optimizing ad bids to **increase daily profit by 22%**
- Strategized an effort to architect horizontally-scalable microservices and set up continuous integration pipelines.

Software Engineer, Media.net June 2015 - June 2018, Mumbai, India

- **Reduce campaign creation and bidding time by 70%** by designing a novel contextual ad-generation system.
- Spearheaded development of new stream-processing architectures to **slash ingestion time by over 90%**

TECHNICAL SKILLS

Programming Languages: Python, Java, R, Kotlin, C++, C
Libraries: PyTorch, Tensorflow, Scikit-Learn, Numpy, Pandas, Matplotlib, Seaborn, Pyro, BNLearn
Other Technologies: Docker, Kubernetes, AWS, Redis, Apache Kafka, Hive, Spark, SQL

ADDITIONAL PROJECTS

Causal Reasoning for Reinforcement Learning Agents, Northeastern University March 2020 - April 2020

- Demonstrated that any non-causal RL agent will lead to biased outcomes in the presence of a confounder.

Deep Semantic Code Search, Northeastern University January 2019 - April 2019
Problem: Can we use deep learning to model the semantics of retrieving code segments given natural language queries?

- Outperformed benchmark scores by learning a joint embedding space for code and natural language queries.