Apoorv Agnihotri

Last Updated on 4th December 2022

apoorvagnihotri.github.io @ apoorvagni@gmail.com

Education

IIT, GANDHINAGAR

B.Tech. in Computer Science

Jul 2020

- Overall grade: 8.25 / 10

Certifications

Coursera

- Neural Networks and Deep Learning
- Convolutional Neural Networks
- Intro. to ML in Production
- PGMs: Representation ← Stanford

Selected Courses

UNCERTIFIED

- CS231n: Deep Learning for Computer Vision @ Stanford
- CS 188: Intro. to AI @ UCB
- COMPM050: RL Course @ UCL
- Full-Stack Deep Learning

Achievements

- **Dean's List Awardee Sem 1 & Sem 3**: An award given to IITGN students for achieving certain level of academic excellence in a semester.
- **Kishore Vaigyanik Protsahan Yojana Scholar 2016**: A fellowship for students that poses aptitude for research, granted by the Indian Institute of Sciences (IISc), India.
- **Pratibhashali Vidyarthi Protsahan Yojana Awardee 2016**: A cash award given to high-school students for achieving certain academic excellence in board examinations.

Teaching

- Gave a guest lecture on Diffusion Model in Probabilistic ML course (2022) @ IITGN.
- Gave research seminars at Wadhwani AI and Rephrase AI
- TA for TEQIP, Machine Learning (2019) @ IITGN.
- Gave a talk (2018) on rapids.ai at PyData, Gandhinagar.
- Started an informal job interview group, with over 40 interviews conducted personally.

Experience

DEEP LEARNING RESEARCHER

Rephrase AI - Series A startup creating realistic avatars for personalized videos.

Apr 2022 – Present — Bangalore

- improving data efficiency and quality of voice cloning using deep probabilistic graphical models (PGMs).
- using off-the-self speech content encoders (hubert-soft) to flag incorrect voice cloning. Later, the team improved our lip-sync model with hubert-soft features owning to my earlier exploration.
- improved team's experimental throughput by integrating ML experimenting frameworks (hydra/pytorch-lightning) in legacy code.

ASSOCIATE ML SCIENTIST I

Wadhwani AI - Gates Foundation supported non-profit building AI solutions for the under-served.

Jun 2020 – Mar 2022 — Bangalore

- led the development of a research prototype for scanning and interpreting medical docs using template matching and classical computer vision.
- surveyed and implemented SOTA techniques for object detection for edge deployment on mobile applications.
- experimented with sample rejection of adversarial inputs in the wild. **(NeurIPS 2022 workshop)**
- contributed to publishing a white paper detailing potential venues for AI disruption in Indian agriculture.
- represented Wadhwani AI in HUL, Google and My Gov's *AI for Agriculture* Hackathon and won a cash reward of 1 million rupees.

SUMMER RESEARCH INTERN

IIT, Gandhinagar

May 2019 – Jul 2019 — Gandhingar

- wrote an expository article on Bayesian Optimization. **(Distill 2020)**
- created an open-source repository (called polire) for spatial interpolation and sensor placement. **(Poster at SenSys 2020)**
- experimented using active learning for sensor placement recommendation to reduce costs. **(CoDS COMAD 2020, ICML Workshop 2020)**

ACCELERATED HPC & MACHINE LEARNING INTERN

NVIDIA

May 2018 – Aug 2018 — Bangalore

- contributed GPU-optimized Kalman filters to rapids.ai, an open-sourced data science framework for ML algorithms.
- created a library to generate multivariate Gaussian random numbers on GPU, subsequently reused by various projects.

Competitions

KDD CUP 2019

Ranked 7th among the 259 teams at a reinforcement learning competition hosted by SIGKDD.

ICPC REGIONALS 2019

Represented IITGN at ICPC Regionals in IIT Kharagpur and IIIT Pune, India

TECH MEET 2018

Represented IITGN in a timeseries classification competition at the inter-IIT technical competition. Fortuitously, our method earned the appellation for being the most novel approach.

Selected Projects

Big-Little Nets | link $\star^{8}$

- implemented a novel CNN architecture, using Pytorch as a part of ICLR Reproducibility Challenge 2019.
- primary idea of the paper: to learn features at different image scales, improving accuracy by 1% at one-third the computations.

Polire: Spatial Interpolation in Python | link $\star^{10}$

- co-developed an open-sourced a python library for spatial interpolation.

Machine Learning Library @ ML Course | link

- designed and implemented a ML library from scratch; implementations include common ML algorithms such as Random Forests, Decision Trees, and Support Vector Machines.
- the library is a distillation of programming assignments from the machine learning course at IITGN.

Temporal Epipolar Regions @ Computer Vision (CV) Course | link

- designed and implemented a ML library from scratch, including ML algorithms such as Random Forests, Decision Trees, and Support Vector Machines.
- distilled programming assignments from the machine learning course at IITGN into a library.

Reinforcement Learning (RL) in Games @ ML Course | link

- implemented value based algorithms such as classical Q Learning, Deep Q Learning.
- analyzed their efficiency in environments provided by Open AI's RL toolkit called gym.

Publications

- Agnihotri, A., & Batra, N. (2020). Exploring bayesian optimization. **Distill**, 5(5), e26.
- Narayanan, S. D., Patel, Z. B., Agnihotri, A., & Batra, N. (2020, November). A toolkit for spatial interpolation and sensor placement. In Proceedings of the 18th Conference on Embedded Networked Sensor Systems. **Poster at SenSys** (pp. 653-654).
- Narayanan, S. D., Agnihotri, A., & Batra, N. (2020). Active learning for air quality station location recommendation. In Proceedings of the 7th **ACM IKDD CoDS and 25th COMAD** (pp. 326-327).
- Narayanan, S. D., Agnihotri, A., & Batra, N. (2020). Active learning for air quality station deployment. Real World Experiment Design and Active Learning **Workshop at ICML**
- White, J., Madaan, P., Shenoy, N., Agnihotri, A., Sharma, M., & Doshi, J. (2022). A Case for Rejection in Low Resource ML Deployment. arXiv preprint arXiv:2208.06359. Challenges in deploying and monitoring Machine Learning Systems **Workshop at NuerIPS**