Aashish Adhikari

being-aerys.github.io | linkedin.com/in/aashish-adhikari | adhikara@oregonstate.edu

Education

Oregon State University, Corvallis, OR — Sept 2018 - June 2021
Master of Science in Computer Science — Cumulative GPA: 3.63

Tribhuvan University, Nepal — Nov 2012 - Dec 2016
Bachelor's Degree in Computer Engineering

Skills and Tools

Languages	Python, R, SQL, Java, Haskell
Data Tools	NumPy, Pandas, sklearn, Jupyter Notebook, PyTorch, skorch, NLTK, OpenCV
Data Skills	Data Cleaning, Preprocessing, Feature Selection, Data Analysis, A/B Testing
Visualization	matplotlib, seaborn, tensorboard, wandb, PowerBI
Machine Learning	Linear Regression, Logistic Regression, GDA, Naïve Bayes, KNN, Decision Tree, Bagging, Boosting, SVM, k-means, kdTree, GMM, EM, PCA, LDA, t-SNE
Deep Learning	MLP, CNN, RNN, GRU, LSTM, Transformer, Transfer Learning, AutoEncoder
Computer Vision	Image Classification, Object Detection, Image Segmentation, Human Pose Estimation
Reinforcement Learning	Value/Policy Iteration, Monte Carlo Learning, Q Learning, SARSA, DQN, REINFORCE, A2C, A3C, DDPG, TD3, Imitation Learning, Random Shooting, CEM
Tools/ Skills	Docker, AWS, SageMaker, Lambda, S3, flask, Heroku, REST, Ray, Git, Linux, MySQL, Unit Testing, AGILE, OOP, Algorithms, Data Structures, Problem Solving

Experience

Graduate Research Assistant, STAR Lab, OSU — Sept 2018 - Jun 2021

- Reduced energy cost by up to **33.9%** at Google data center environment using **deep RL**.
- Explored **~500GB** of raw data center usage data to gain insights and extract artifacts.
- Designed a data center **MDP** with a novel reward function for cost minimization.

Project Supervisor/ TA, Tribhuvan University — Nov 2017 – Sept 2018

- Supervised a capstone ML project - **Human Activity Recognition** using sensor data.

Publication

- A. Adhikari et al., "Improving Data Center Peak Shaving with Reinforcement Learning," in ICAI '21

Graduate Projects

OOP Pipelines of Machine Learning Algorithms [code] — Jul 2020 - Present

- Maintaining a public repository of supervised and unsupervised learning algorithms.
- Implemented OOP pipelines for data preprocessing and model training pipelines using only python.

Distributed Learning with Ray on Intel DevCloud [code] — Apr 2019 – Present

- Trained tabular and deep RL agents in OpenAI environments on DevCloud using Ray.
- Leveraged distributed implementation to reduce convergence time by up to **10 folds**.

ML Deployment on AWS SageMaker [code] — Dec 2020 - Present

- Deployed sentiment analysis, population segmentation, and time-series forecasting models.
- Trained XGBoost, k-means, and DeepAR utilizing REST API, lambda functions, and gateways.

Dynamics Modelling of Self-Driving Cars [code] — Apr 2020 – Jun 2020

- Learned a large neural network state space model that predicts the future latent state of an F1/10 car.
- Used deep Cross Entropy Method to obtain feasible motion plans to follow the centerline of a track.

Sequence Models [code] — Sept 2018 – Jul 2019

- Trained LSTM networks for character-level text generation and bitcoin price prediction.
- Trained an image captioning network using a CNN encoder and an LSTM decoder for Nepali.

Real-Time Pose Guidance on Low-End CPUs [code] — Apr 2019 – Jun 2019

- Designed a ConvNet for human pose detection and guidance on low-end CPUs.
- Accomplished a **four-fold** inference time reduction compared to Lightweight OpenPose.

Latent State Learning for Multiagent Coordination [code] — Jan 2019 – Mar 2019

- Achieved a **four-fold** reduction in state representation of a rover environment using an autoencoder.
- Trained DDPG on the latent representation to achieve an optimal multiagent policy to observe POIs.