

# WINSTON MOH TANGONGHO

LinkedIn: [/winston-moh-8730b756/](#) • GitHub: [/WinstonMoh](#) • 507-304-5449 • [moh.winston@yahoo.co.uk](mailto:moh.winston@yahoo.co.uk)

## EDUCATION

---

**Northeastern University** *Boston, MA*

**Khoury College of Computer Science**

*Master of Science in Computer Science* / GPA: 3.5/4.0

*May 2021*

Related Courses: Scalable Data Management, Deep Learning, Scalable Distributed Systems, Data Structures and Algorithms, Design Paradigms, Computer Systems, Artificial Intelligence, Machine Learning

**Minnesota State University, Mankato** *Mankato, MN*

*Bachelor of Science in Computer Engineering* / GPA: 3.8/4.0

*May 2018*

Related Courses: Computer Hardware and Organization, Algorithmic Structures, Digital and Electronic Systems Design, Software Engineering, Operating Systems Design, Fundamentals of Software Development

## TECHNICAL KNOWLEDGE

---

- **Operating Systems:** Linux (Ubuntu), Windows OS
- **Languages:** C/C++, Java, Python, Golang, C#, MATLAB, Flutter
- **Software:** Kubernetes, JIRA, Git, Visual Studio Code, Microsoft AZURE, Amazon AWS
- **Databases:** MongoDB, DynamoDB, SQL
- **AI/ML** TensorFlow, Keras, VGG15 Neural Networks

## WORK EXPERIENCE / ACADEMIC PROJECTS

---

**Software Development Engineer Intern, Amazon Web Services**

*June 2020 – August 2020*

- Worked on a feature to integrate Amazon Support with a chat-based system to provide more flexibility for our premium Customers to create and track their support cases. Utilized tools such as Java, CloudFormation and CloudWatch.
- Leveraged the AWS Step-Function together with a state-machine to periodically poll an AWS Service endpoint to get the Status update of a support case. An AWS Lambda function was created to interact with the API endpoint.
- Trained a Deep Reinforcement Learning model for 6 hours by designing a reward function and adjusting its Hyperparameters to compete in an AWS DeepRacer tournament. The model utilized TensorFlow, Amazon SageMaker and the Proximal Policy Optimization Algorithm for training. It was then evaluated on the AWS re:Invent 2018 track during which it completed 3 laps with an average time of 18 seconds.

**Software Engineering Intern, Datadog**

*January 2020 – May 2020*

- Modified different parts of a consumer application (Golang production code) to add a Kubernetes readiness check on application startup. This improved traffic to the pods in case the application crashed during startup.
- Set-up autoscaling on a metric for two services which ran on Kubernetes with extra features to the existing HPA. This Saved the company money, reduced the number of nodes by ~70 and made our systems more reactive to capacity changes.
- Created three new gRPC endpoints for an in-house application in Golang with batching for API calls to improve latency when communicating with the Postgres servers.
- Leveraged a Sarama client which implemented consumer groups by storing consumer offsets in Kafka instead of Zookeeper. This improved the consumer lag and also removed a deprecated feature which stored consumer offsets in a different client.

**Pattern Recognition in Top Trending YouTube videos, Northeastern University**

*September 2019 – December 2019*

- Performed analysis of thumbnail images of the videos by performing Feature extraction, Dimensionality reduction and Unsupervised clustering with a self-organizing map (an Artificial Neural Network) with a dataset size of 5385 entries.
- Applied NLP techniques such as word2vec to obtain a representation of the video descriptions.
- Designed a U-Matrix to display a sample of 1000 images after clustering the thumbnails with the Artificial Neural Network And observed that Talk show videos, basketball videos, CNN News videos and music videos from the 'vevo' channel had the highest Probabilities of being Top Trending videos.

**Solving Lunar Lander with Reinforcement Learning, Northeastern University**

*September 2019 - December 2019*

- Solved the Lunar Lander environment in OpenAI gym by utilizing Reinforcement Learning techniques such as SARSA And Deep Q Learning to train the Agent by using the Keras API.
- Fine-tuned the Network parameters of the Deep Q Network to obtain a top-performing Agent which solved the environment In an average of 200 steps after learning the optimal policy.

## INTERESTS/ACTIVITIES/HONORS

---

• **OvenBot** (*PerkinsHacks 2018 Hackathon: Challenge Award Winner*) *Boston, MA*

Software Program to ease use of smart Home Appliances by visually impaired individuals

• **National Society of Black Engineers**, Treasurer: Organized fund-raising events and participated in meetings with Department heads to gather enough funds for our yearly Conference and on-campus activities