

EDUCATION

New York, NY **New York University** **Sep 2019 – May 2021**

MS in Data Science. GPA: 3.78 / 4.0

Coursework: Probability & Statistics; Linear Algebra; Machine Learning; Deep Learning; Big Data; Computer Vision.

Course Assistant: Introduction to Deep Learning Systems (Fall 2020).

Vellore, India **Vellore Institute of Technology** **Jul 2015 – May 2019**

B. Tech. in Information Technology. GPA: 9.15 / 10.0

EXPERIENCE

Machine Learning Intern **PepsiCo** **Jul 2020 – Aug 2020**

- Trained a Word2Vec-like model on customer product baskets to project ~10,000 unique products to a low-dimensional embedding space (using TensorFlow).
- Evaluated learned representations using the measures of “complementarity” and “substitutability” between different products, inspired by the paper - [Product2Vec](#).
- Used t-SNE to visualize and qualitatively validate the embedding similarity between complementary/substitute products.

Data Science Intern **Noodle.ai** **Jan 2019 – May 2019**

- Developed a modular and reusable python framework for working with sequence-to-sequence models (using Keras).
- Designed a benchmark suite for evaluating different sequence-to-sequence architectures for multivariate time-series forecasting (on internal datasets of the company).
- Standardized and adapted the data preprocessing pipeline across various ongoing projects to handle time-series data.

PROJECTS

Head and Neck CT image segmentation using deep learning [GitHub](#)

- Designed an end-to-end framework for segmenting organs in the head and neck CT images, achieving near-SOTA performance on the MICCAI Auto-segmentation Challenge (72.16 vs. 75.28 dice score).
- Enhanced the baseline U-Net architecture by adding residual connections and a mixup training regime to prevent overfitting on the relatively small dataset (only ~30 CT scans).
- Used Grad-CAM to interpret and understand the model predictions better.

Generating bird’s-eye view from multi-image scene using self-supervised learning [GitHub](#)

- Utilized a pre-trained denoising auto-encoder as a feature extractor for the task of road map construction.
- Combined the feature embeddings of the multi-view scene captured from moving vehicles and used a U-Net to generate a top-down view of the binary road map.
- Achieved a threat score of 0.70 on the hold-out test set (with the best score in the intra-class competition being ~0.81).

Neural style transfer (with improvements in terms of speed and quality) [GitHub](#)

- Reimplemented the seminal work by Gatys et al. titled “A Neural Algorithm of Artistic Style” (using PyTorch).
- Programmed numerous extensions to preserve color¹ and improve the algorithm’s quality² (primarily based on the gram matrix characteristics).
- Implemented the fast style transfer method proposed by Johnson et al. working in real-time (~ seconds) compared to the original optimization-based algorithm (~ minutes).

“Mydia” – A open-source framework for working with videos in python [GitHub](#)

- Developed an efficient and user-friendly python library for reading videos as NumPy tensors with (according to the tracker [PePy](#)) more than 55,000* cumulative downloads.
- Utilized python’s multiprocessing capabilities to read (multiple) videos in parallel, thereby reducing video processing time by 4x compared to other similar tools.

TECHNICAL SKILLS

Proficient	Python, PyTorch, Scikit-learn, NumPy, Pandas, Matplotlib, OpenCV, Weights & Biases, SQL, Git, Linux
Familiar	C++, TensorFlow and Keras, AWS, GCP