Carnegie Mellon University
Silicon Valley

ARTIFICIAL INTELLIGENCE CAREER GUIDE

Student Affairs at Carnegie Mellon University Silicon Valley

CONTENTS

ARTIFICIAL INTELLIGENCE CAREER GUIDE

Acknowledgment section

These career guides are adapted each year, and the career services team would like to sincerely thank the 2019 cohort of Peer Career Consultants for creating the first iteration. We appreciate the hard work of Chaitra Dinesh Pai, Nursultan Dyussebayev, Arpit Gupta, Sweta Hari Kumar, and Joseph Konan who researched specific fields of interest and compiled helpful tips and resources for students to enable them to make the most of their CMU-SV experience and navigate the internship and job search process in Silicon Valley. Their valuable contributions made these career guides possible. Additionally, the example resumes are compiled from recent alumni who also wanted to help our future students at CMU-SV. Thank you to everyone who contributed their time and effort.

Introduction

The purpose of this guide is to help you understand what a career in Artificial Intelligence (AI) means, and how to prepare for a job search in various areas of AI. **Example roles within artificial intelligence include but not limited to the roles of Data Scientist, Machine Learning Engineer, Deep Learning Engineer, and Machine Learning Software Engineer.** The materials herein are intended as a starting point for students to use and prompt further research into strategic planning for a career in the field of AI. As the #1 leading school in the field of Artificial Intelligence according to U.S. News and World Report, Carnegie Mellon University is consistently pushing the bounds of what is considered state-of-the-art with world-leading faculty, prestigious research labs, and a diverse network of interdisciplinary researchers, alumni, and industry sponsors.

Image: Carnegie Mellon Artificial Intelligence Beats Chinese Poker Players by $792,327

EXAMPLE RESUME

To give you an idea of what an artificial intelligence focused resume can look like before it's tailored to the job description, please see the example below. *Please note, these roles may not have occurred in 2019-2020. This is a resume example from a real CMU-SV student, however, other data points are edited each year. Please recognize that this is not a real candidate. Please email career-services@sv.cmu.edu if you have any questions about this example resume or if you need to schedule a resume appointment, please do so through Handshake.

Mountain View, CA
linkedin.com/in/cmutartan1

CMU TARTAN

(555) 555-5555
Cmutartan1@andrew.cmu.edu

EXPERIENCE

Graduate Research Assistant — **Emirates Airlines** — **May 2019 – Present**
CMU-Emirates iLab
- Reduced travel time variability of intercontinental flights by 75% by implementing LSTM TensorFlow algorithm.
- Implemented iLab integration with OS X Spotlight Search by creating tool which extracts metadata and uses natural language processing on saved chat transcripts to provide a 148% faster system-wide search database.
- Redesigned chat data pipeline and implemented cloud compatibility for search, scaling to 30,000+ users.

Graduate Research Assistant — **Carnegie Mellon University** — **January 2019 – May 2019**
Accountable Systems Lab
- Identified high influence neurons in quantity and distribution of interest, using influence-directed measures.
- isolated individual features the network uses to make decisions and distinguish related classes using PyTorch.

Graduate Teaching Assistant — **Carnegie Mellon University** — **August 2018– Present**
- Courses: Machine Learning for Product Managers, Intro to Deep Learning, Intro to Machine Learning.
- Mentoring 20+ students on creating state-of-the-art AI. Also, product owner of 4 end-to-end data-driven apps.

EDUCATION

Mountain View, CA — **Carnegie Mellon University** — **Aug 2018 – May 2020**
- M.S. in Electrical and Computer Engineering
- Graduate Coursework: Machine Learning for Signal Processing; Hardware Architectures for Machine Learning; Speech Recognition and Understanding; Machine Learning for Product Managers; Intro to Deep Learning.

Made-in Abyss, Japan — **Orth Prefecture University** — **Aug 2014 – May 2018**
- B.S. in Computer Science and Engineering with Minor in Mathematics
- Undergraduate Coursework: Artificial Intelligence; Mathematical Statistics; Advanced Calculus; Advanced Linear Algebra; Operating Systems; Databases; Algorithms; Programming Languages; Computer Architecture.

TECHNICAL PROJECTS

- **Applications of Deep Learning in Activity Recognition** (2018). Python, PyTorch, TensorFlow, ONNX. Center Loss & Wav2Letter, classifying 9 activity vibrations. 80% raw signal accuracy. 90% feature engineering accuracy.
- **Self-Attention Networks for Connectionist Temporal Classification** (2002). Python, PyTorch. Speech recognition for transcript generation using Bi-LSTM recurrent neural net with string edit difference of 16.8.
- **Deep Residual Learning for Image Classification and Verification** (2003 – 2004). Python, PyTorch, OpenCV. Leveraged CNN residual networks, augmentation, alignment. 69% classification accuracy. 0.93 verification AUC.
- **Sparse Multilayer Perceptron for Phoneme Recognition** (2004). Python, PyTorch. Created high-efficiency low-parameter network to identify phoneme state labels using raw mel spectrogram frames. 65% accuracy.
- **Deep Learning Automatic Differentiation Libraries** (2002). Python. Developed functions and methods from scratch for MLP, CNN, and RNN forward/backward-propagation, gradient descent, and attention mechanism.

WORKING PAPERS AND PUBLICATIONS

- **NIPS 2018.** *Generalized Inverse Optimization through Online Learning*. Nanachi Hollow, Riko Delver, Reg Abyss.
- **IEOM 2017.** *Neural Networks Applications in Manufacturing Processes*. Mitty Narehate , Nanachi Hollow.

Languages and Technologies

- Python, R, SQL, C/C++ | SciKit-Learn, TensorFlow, PyTorch, Keras, ONNX, Matplotlib, ggplot2, Seaborn, Plotly
- Jupyter, RStudio, PyCharm, Ubuntu, Linux | AWS, EC2, SageMaker, Google Cloud, CUDA

CAREER GUIDE STARTER KIT

This guide is broken into five parts, which aim to outline: career expectations, academic preparation, career strategies, presentation, and a case-study of how to make the most of your time at CMU.

What Does a Career in AI Look Like?

The field of artificial intelligence uses machine learning to directly help companies realize higher profits or to provide additional features to a product. Practitioners are able to predict things in reality by having a machine learn parameters that give the best results for a provided mathematical model. Example applications include predicting ad revenue in any company, face or voice recognition on iPhones, language translation on YouTube, medical diagnostics, and more. Most practitioners specialize in one or more of the following: computer vision, natural language processing, signal processing, financial modeling, recommendation system, robotics, and/or theory.

How Do I Prepare for a Career in AI?

The most common roles in artificial intelligence involve data science, machine learning, and more recently the subfield of deep learning. Of the three, deep learning is one of the most important and often required skills that companies leveraging artificial intelligence are looking for. If you don't know deep learning, getting an internship or full-time role working on artificial intelligence is significantly more difficult. Employers like to see relevant, high impact projects that have been deployed, published, or have won Kaggle competitions.

Students who are serious about getting involved in this field are expected to take Introduction to Deep Learning 11-785. Successful students gain access to the network of other students who have completed the course -- connections directly relevant to the field and industry. The course is well rounded in terms of concepts. It helps us understand the fundamentals of Deep Learning. The course starts off gradually with MLPs and progresses towards more complicated concepts such as attention and sequence-to-sequence models. We get a complete hands on experience with PyTorch which is very important to implement Deep Learning models. As a student, you will learn the tools required for building Deep Learning models. The homework assignments usually have 2 components which is Autolab and Kaggle. The Kaggle components allow us to explore multiple architectures and understand how to fine-tune and continuously improve models. The tasks for all the assignments were similar and it was interesting to learn how the same task can be solved using multiple Deep Learning approaches. Overall, at the end of this course you will be confident enough to build and tune Deep Learning models.

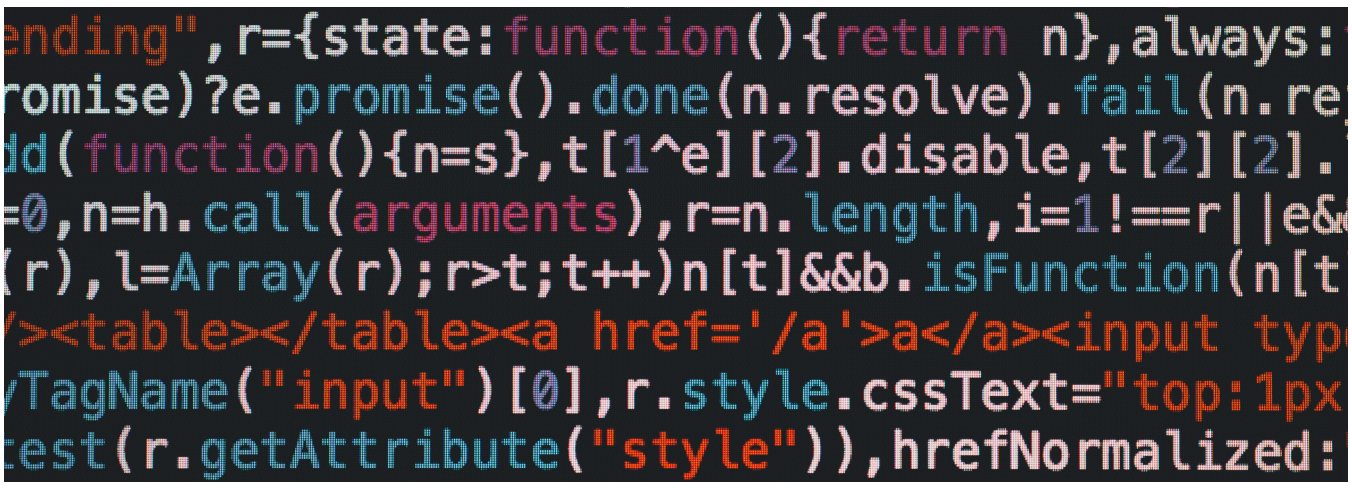

COURSEWORK AND PROJECTS

Having a strong project portfolio is vital to having a strong application. At Carnegie Mellon, you have access to world class faculty who will help you recreate state-of-the-art results, design scalable infrastructure, or implement complex algorithms. Relevant courses include:

Group 1: Foundations (1 Required)
18661 Intro to Machine Learning for Engineers
11785 Introduction to Deep Learning (Recommended)

Group 2: Working with data (1 Required)
18656 Data-Intensive Workflow Development
14736 Distributed Systems
14848 Cloud Infrastructure
15619 Cloud Computing

Group 3: Specialization 1 (1 Required)
18781 Speech Recognition and Understanding
18793 Image and Video Processing
18797 Machine Learning for Signal Processing
11411 Natural Language Processing

Group 4: Specialization 2
18663 Hardware Architectures in Machine Learning
18739 Special Topics in Security: Security and Fairness of Deep Learning
49795 Artificial Intelligence for Product Managers

Group 5: Research Experience (2 Required)
18980 M.S. Graduate Project (repeatable)
NA Course Project with Advisor
NA AI Internship

Group 6: Other
18645 How to Write Fast Code
18751 Applied Stochastic Processes

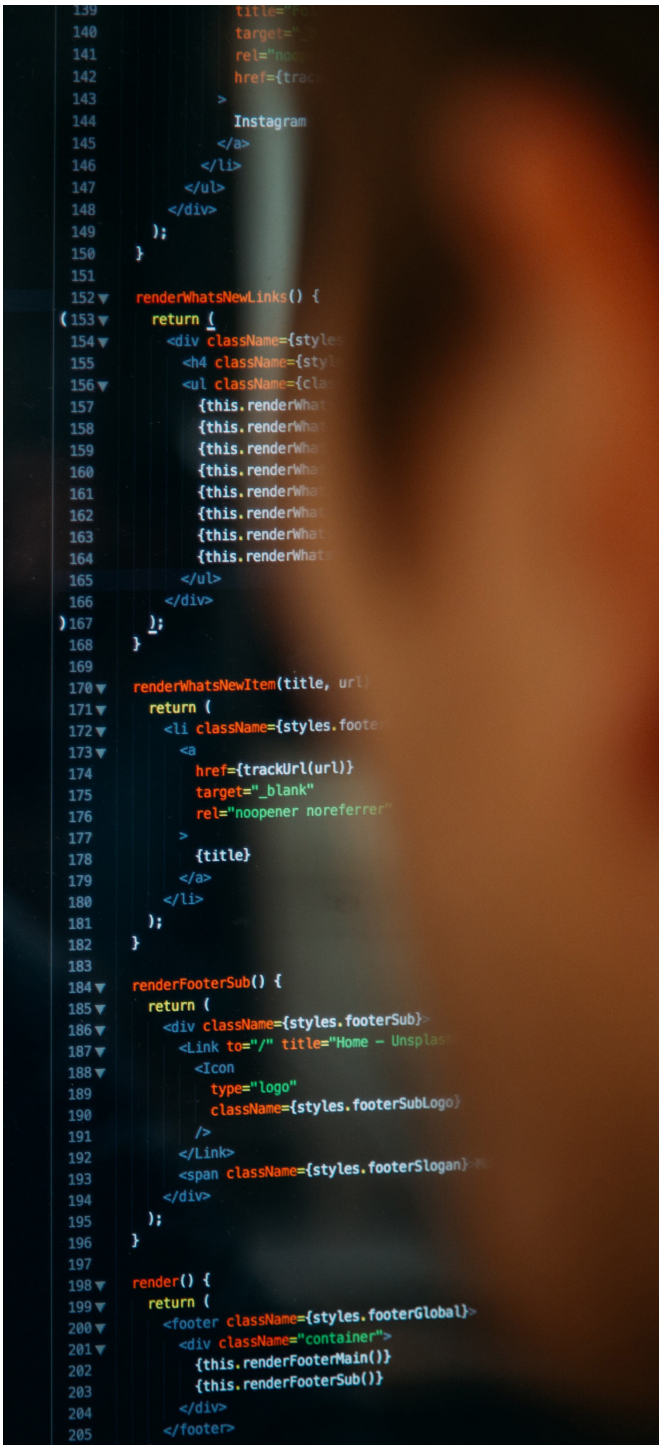

COURSEWORK AND PROJECTS

Below are concrete examples of course projects as they would appear on your resume.

Self-Attention Networks for Connectionist Temporal Classification
• Developed end-to-end speech recognition encoder-decoder RNN. Generated transcripts, surpassing human-level performance w/ Levenshtein distance of XX.X . PyTorch, Python.

Deep Residual Learning for Image Classification and Verification
• Developed deep residual learning CNN architecture for face identification with XX.X% accuracy and 0.XX AUC, leveraging center loss and triplet loss. PyTorch, Python.

Sparse Multilayer Perceptron for Phoneme Recognition
• Created high-efficiency low-parameter MLP network to identify phoneme state labels using raw mel spectrogram frames. XX.X% accuracy. PyTorch Python

Deep Learning Automatic Differentiation Libraries
• Implement low-level logic of an end-to-end deep learning framework, together with automatic differentiation and optimization or MLPs, CNNs, and RNNs. PyTorch, Python.

Eigenfaces for Face Recognition and Bounding
• Implemented eigenfaces as an unsupervised dimensionality reduction algorithm, leveraging results by reshaping eigenvectors for principal component analysis. Matlab.

Source Separation w/ Single Channel Independent Component Analysis
• Implemented my own ICA based on freeing fouth moments method, extracting independent components to simultaneously speaking sources from a recording. Matlab.

Signal Separation using Non Negative Matrix Factorization
• Separated mixed signal into component signals, extracting speech from recordings with background sound and music by minimizing KL divergence between spectrograms. Matlab.

Sparse Recovery of Streaming Signals Using Iterative Hard Thresholding
• Implemented iterative hard thresholding with projected gradient descent algorithm, recovering sparse image transform from 3 compressed measurements. Python, Matlab.

Influence-Directed Explanations for Deep Convolutional Networks
• Developed tool explaining neural networks by identifying influential concepts generalized across instances and isolated individual features uses for decisions. TensorFlow, Python.

Multi-Targeted Adversarial Evasion Attack of Deep Neural Networks
• Implemented multi-targeted adversarial example that attacks multiple models within each target class with single modified image with 100% attack success. TensorFlow, Python.

Below are concrete examples of course projects as they would appear on your resume.

Model Memorization for Calibrated White-Box Membership Inference
• Implemented Membership Inference (MI) attacks, exploiting a learned model's lack of generality to infer whether given sample was in a model's training set. TensorFlow, Python.

Speech Differentiation via Naive Bayes for Text Classification
• Developed a classifier that distinguishes speeches given by Republicans (Red) and Democrats (Blue) running for president. XX.X% accuracy. Python.

Translation via Part-of-speech Tagging and Hidden Markov Models
• Generated and evaluated first order bigram HMMs by leveraging Viterbi algorithm to calculate transmission and emission probabilities, deriving parts-of-speech. Python.

Deep Learning for Sentiment Analysis
• Developed deep neural network to understanding feedback in product reviews and user's opinions in tweets, identifying positive, negative, or neutral text. PyTorch, Python.

Big Data Twitter Analytics Web Service Design and Implementation
• Developed data warehouse complete with front-end query response system, backend database, ETL tool, to process terabytes of tweets on AWS. Java, Tomcat, Python, Hbase.

Blockchain Ledger Design and Implementation
• Developed blockchain mining application with ledger that will award coins if a block is successfully solved given a difficult work and add a new block into the network. Go, Java.

Distributed File System Design and Implementation
• Implement a simple distributed file system using my own RMI library, allowing files to be hosted remotely on one or more storage servers complete with API. Go, Java.

Distributed Key-value Store Design and Implementation Using Raft
• Implement Raft algorithm as a replicated state machine protocol, achieving fault tolerance by storing copies of data on multiple replica servers. Go, Java.

Bigram ETL Tool Design and Implementation Using Apache Spark
• Counts number and frequency of bigrams, as well as the number of bigrams required to add up to 10% of all bigrams within corpus. Apache Spark, Java.

BigTable, a Scalable NoSQL Database Design and Implementation
• Implemented a distributed storage system for managing structured data, designed to scale to a very large size: petabytes of data across thousands of commodity servers. Java.

Multi-path TCP Analytic Tool of Instant Virtual Networks in Mininet
• Analyzed Three-Tier and Fat Tree network topologies as well as east-west congestion at the transport level using MPTCP approach by simulating network traffic. Mininet, Python.

Remote Method Invocation API Tool for Distributed Microsystems
• Implemented remote method invocation (RMI) library for calls over network connections, permitting objects located on one Java virtual machine to call methods on another. Java.

RESEARCH EXPERIENCE

Students who are interested in careers in AI are expected to have done research with a faculty member that specializes in the domain. In particular, research assistantships are easily attainable and lead to high impact projects given you are working with faculty members at Carnegie Mellon University. For perspective, a Research Scientist in industry will have hundreds of citations. In contrast, ECE and SCS professors will have thousands of citations. High-tier internships have "a publication in top journals" as a sought after or required qualification, and you will be more successful if you publish in your interest area.

Many courses, including 11-785, 14-736, 14-848, 18-781, 18-793, and 18-797, require you to create a project proposal and realize that idea by getting results. These courses, in addition to others not listed, allow you to have a faculty advisor on your project. This means, you are able to choose a research idea that the faculty member is interested in working on and get research experience while doing your coursework. In addition, 18-980 is an option for students who have a research project not linked to a traditional course, allowing you to get core credit for doing research with faculty. There is no excuse for not having at least one research assistantship working with faculty that specialize in your area of interest. Further, if you are starting at CMU with no relevant experience in the field, it is readily attainable.

If you want to work at a specific company doing AI research and development, then cut to the chase. People who want to work at Amazon should have research projects improving the products or experience of Amazon's engineers or customers. A simple google search yields multiple big-data resources that are used by Amazon and researchers:

• SNAP Web Data: Amazon Reviews (https://snap.stanford.edu/data/web-Amazon.html)

• Topical Chat Dataset: Amazon Alexa (https://github.com/alexa/Topical-Chat)

• Sales Rank Data: Amazon Print and Kindle Books (https://www.kaggle.com/ucffool/amazon-sales-rank-data-for-print-and-kindle-books)

• Movie Review Web Data: Amazon Prime Video (https://www.kaggle.com/dm4006/amazon-movie-reviews)

• Stock Price Data: AMZN (Amazon) (https://finance.yahoo.com/quote/AMZN/history/)

Applying to Amazon and having multiple projects on Amazon will most likely get your application noticed by Amazon. Moreover, a code you write for these specific datasets relevant to Amazon can be ported to get similar insights and results for other companies. To emphasize this point, here are some equivalent Google datasets for those listed above:

• Google Play Store App Data (https://www.kaggle.com/lava18/google-play-store-apps)

• TensorFlow Speech Recognition Challenge: Google Home (https://www.kaggle.com/c/tensorflow-speech-recognition-challenge/data)

• Google Analytics Customer Revenue Prediction (https://www.kaggle.com/c/ga-customer-revenue-prediction/data)

• Video Comments and Metadata: YouTube (https://developers.google.com/youtube/v3/docs/comments)

• Stock Price Data: GOOG (Google) (https://finance.yahoo.com/quote/GOOG/history/)

And more datasets for Amazon (https://www.kaggle.com/datasets?search=Amazon) or datasets for Google (https://www.kaggle.com/datasets?search=Google) are out there, and you can simply search for your top company's data so you can make the best resume for your application. Joining Kaggle competition is also a great way to build your resume: (https://www.kaggle.com/competitions).

Another application of ML is Computer Vision, which is widely used in the self driving and robotics industries.
Popular datasets: Waymo One (https://waymo.com/open/data/perception/). They also hold self-driving specific competitions (https://waymo.com/open/challenges/)

PATHWAYS

Depending on the amount of AI related coursework you take, and the extent to which you engage with faculty, it may make sense for you to consider a variety of options.

Data Science Bootcamps. These, in terms of competitiveness, are among the least competitive given you are an engineering student from Carnegie Mellon University. Bootcamps often cost money, so, like school, you are paying to build your human capital.

Data Science Fellowships. These are like Data Science Bootcamps but are more competitive and are often sponsored by employers. These opportunities are free, and, for this reason, there is a more competitive process getting in. People who did their Ph.D. in Physics or some non-related STEM field really like these options, because they are often overqualified in terms of foundational understanding. Being a student who has more relevant coursework will give you more intuition than the people who are very smart but want a career change.

AI Residencies. These are ideal if you are shooting for a competitive AI role, but have either published in a low- to mid-tier journal or have not published at all. Even if you have published in a high-tier journal, these are still worth considering because you are able to do similarly impactful research but in an industry setting. AI Residencies are full-time roles that pay like full time roles. This is becoming more favorable among recent Ph.D. grads who do not want to do a Post-Doc that pays subsistently while still getting more specialized research experience. For this reason, AI Residencies, though open to B.S. to Ph.D. students, are often more difficult to get into than Data Science Fellowships.

Consider this list if you are interested in AI Residencies. (https://github.com/dangkhoasdc/awesome-ai-residency)

DS and ML Roles at Mid-tier Companies. These roles are not hard to find because Data science and AI is so popular. There are many people who want to break into data science, and they tend not to first consider bootcamps or fellowships (though they should) so the applicant pool is very large. To stand out at mid-tier companies, it is not necessary that you have published, but you should have relevant project experience on the same problem they are working on. Even better, if you got results using their data or improved their product, then you will be taken more seriously by the company you are applying to. The prestige of this is between Fellowships and AI Residencies, and pays similar to the latter.

Consider these websites if you want to find mid-tier Data Science Jobs. (https://www.springboard.com/blog/data-science/data-science-jobs/)

AI Roles at High-tier Companies. These are competitive. You should have: an AI residency on a related subject, an internship (think short residency) with the company and published, won a major competition by the company and published, a Ph.D. dissertation on improving their product, or equivalent. Talk to your advisors and mentors who make these possible.