

# Talia Chopra

San Diego, California • (831) 332-2819 • talia.chopra@gmail.com • <https://techopra1000.github.io/index.html>

---

## Professional Experience

### **Machine Learning Engineer, January 2022 - Present**

Gemini, LLC

- Owned end to end design, development and deployment of both batch and low-latency (less than 800 ms) model inference pipelines and APIs. Architected pipelines using AWS MWAA, EKS, Lambda, SQS, EFS, S3, and API Gateway. Created alarms and API monitors using AWS Cloudwatch and DataDog to make API and underlying pipelines observable and robust.
- End to end model development, deployment, and monitoring for an asset ranking problem. The model was built using asset price, market factor, and sentiment tabular data. The data this model generates is used by business groups in the company for asset characterization.
- Developed a pipeline to scale an image search model to over 10 million images, with a current goal of reaching over 100 million images using Docker, AWS EC2, MWAA Airflow, EKS, and S3.
- Exploratory data analysis and feature engineering using SQL and PySpark on large tabular data (i.e. 300 million+ rows). Performed this type of analysis on various types of tabular data, including time series, finance, fraud, and customer service data.
- Worked directly with stakeholders across the company to shape business requirements of projects (i.e. those described above), identified accuracy metrics and acceptable thresholds, and reviewed model results on an ongoing basis. Built Looker dashboards to visualize and communicate model accuracy metrics.

### **AI/ML Programmer Writer, October 2019 - May 2021**

Amazon Web Service (AWS) at Amazon, Inc.

- Trained machine learning models in the cloud using Amazon SageMaker to create tutorials and demos. See examples [here](#).
- Used and wrote about AWS ML services, including Amazon SageMaker distributed (a deep learning distributed training library), MXNet, Amazon S3, Amazon IAM, Amazon S3 and Amazon Augmented AI to solve business use-cases.

### **Business Development Lead, June 2016 - October 2019**

Schneider Estates, Inc.

- Research and integrate new technologies into business workflows to improve lead acquisition and customer service, resulting in 2x more leads in the company pipeline, monthly.

## Education

### **Master of Science in Computer Science, Concentration in Machine Learning, Jan 2019 – December 2021**

Georgia Institute of Technology (4.0 GPA)

**Key Coursework:** Deep Learning, Machine Learning, Machine Learning for Trading, Computer Vision, Knowledge Based AI, Operating Systems, Graduate Algorithms, AI and Ethics, AI for Robotics, Software Architecture and Design

#### **Project Highlights:**

- Tuned and compared the performance of multiple machine learning models, including neural networks, decision trees, SVMs, and decision tree ensemble learners on two binary classification datasets. (Python)
- Configured and trained a multi-model CNN model using transfer learning to 1) create embeddings for images of food and 2) categorize food images. Used embeddings and categorizations to perform recipe retrieval for new images. (PyTorch, Python, AWS)
- Created a distributed file system using protocol buffers and gRPC. Incorporate a weakly consistent synchronization system to manage cache consistency between multiple clients and a single server. (C++)

### **Computer Science Coursework 2017 – 2018**

Foothill College (4.0 GPA)

**Key Coursework:** Object Oriented Programming in C++, Advanced Data Structures and Algorithms in C++, Linux Shell Programming, Software Design in C++.

### **Bachelor of Arts in Economics, Minor in Mathematics 2012 – 2016**

University of California, Santa Cruz • Summa Cum Laude, Highest Honors (3.97 GPA)

## Languages and Technologies

**Languages:** Python, C++, C, HTML, CSS, SQL | **Machine Learning Frameworks:** Pytorch, Tensorflow, Scikit-learn

**Tools and OS:** AWS services (i.e. EKS, Lambda, S3, SageMaker, EventBright, MWAA), Kubernetes, Docker, PySpark, pandas, matplotlib, Jupyter, DataBricks, Snowflake, Looker, Git, Linux