

# KAAN ERTAS

Stanford, CA ■ kkaanertas@gmail.com ■ github.com/kaanertas ■ +1 650 495-4294

## EDUCATION

Stanford University, M.S. in Computer Science – AI Track

Expected, June 2021

B.S. in Mathematical and Computational Science

Expected, June 2021

Cumulative GPA: 4.03/4.00

**Relevant Coursework:** Machine Learning, Artificial Intelligence, Mining Massive Data Sets, Algorithms, Data Structures, Statistical Inference, Time Series Analysis, Linear Regression and ANOVA, Experimental Design and Causal Inference, Stochastic Processes, Probability, Linear Algebra, Optimization, Real Analysis

## TECHNICAL SKILLS

**Programming Languages:** Python, R, Java, C++, C, SQL, logic programming

**Machine Learning Libraries:** Keras, Pandas, NumPy, SciPy, Seaborn, sklearn, Matplotlib

**Technologies and Frameworks:** Apache Spark (pyspark), AWS, Git, Linux, Hadoop, Jupyter Lab, REST API

## WORK EXPERIENCE

Stanford Network Analysis Project

September 2019 – March 2020

**Data Scientist**

Stanford, CA

- Developed pipeline for generation of metabolic graphs for human gut microorganisms in Python
- Built predictive model for cooccurrence and compatibility of bacteria pairs using graph embeddings, achieving 73% accuracy in the models overall
- Achieved 98% out-of-sample accuracy for predicting vitamin biosynthesis pathways using LASSO regression
- Investigated research literature on biocomputational methods for microbial community analysis

Citadel Securities

June-August 2020

**Quantitative Trading Intern**

Chicago, IL

- Created custom metrics to be used as input to linear models that drive trading, employing concepts such as smoothing, autocorrelation, and price elasticity of demand; and communicated the results of analyses to leadership in order to influence trading strategies
- Automated custom querying of a REST API used for market data, saving upfront work of up to 5 hours per project
- Built a parser for XML query responses to query market data critical for real-time trading and pricing of ~200 assets
- Improved library calculating daily and minutely volatility estimators using calculations in academic literature such as Garman-Klass estimation, using kdb+ queries to make the querying and estimation 10x faster

Tibra Capital

June-August 2019

**Quantitative Trading Intern**

London, UK

- Adapted existing market data pipelines to new stock exchanges, and added filters for instrument types
- Performed data cleaning, linear regression, correlation analysis and signal processing on market data using Python
- Implemented a covariance estimator for high-frequency, noisy and asynchronous data (project won 1<sup>st</sup> place at intern hackathon)

Stanford University Computational Logic Group

June-December 2018

**AI Research Engineer**

Stanford, CA

- Identified higher level structures in logical descriptions of games using Java and Answer Set Programming
- Improved the space complexity of a Monte Carlo Tree Search algorithm used for AI game playing agents
- Worked with convoluted Java codebase and tackled technical debt by increasing the code documentation level to 90% and modularizing the architecture

## PROJECTS

Deep Q-Learner for Modified 1010! (code and report on GitHub)

March-June 2019

- Implemented a Deep Q-Learner for the Tetris-like game of 1010!, using deep learning with Keras and employing decaying epsilon-greedy policy, target networks and experience replay
- Experimented with different neural network architectures (convolutional, fully connected) and hyperparameters, achieving slightly better than average human performance

Yelp Restaurant Photo Classification (code and report on GitHub)

January-March 2019

- Implemented a multiclass, multilabel classifier for the Yelp Restaurant Photo Classification task on Kaggle, achieving Mean F1 Score within 2% of Kaggle competition winner
- Used convolutional neural networks with Keras by implementing methods to tackle data imbalance and weakly labeled data