

+1 (984) 837-2465
Durham, NC
cs582@duke.edu

Carlos Gustavo Salas Flores

GitHub: cs582
LinkedIn: carlosgustavosalas
Portfolio: cs582.github.io/portfolio/

EDUCATION

Duke University/Duke Kunshan University, B.S. in Data Science & B.S. in Interdisciplinary Studies *Durham, NC + Kunshan, China*

- GPA: 3.7/4.0.
- Dual-degree, Full Scholarship, and Dean's list (2019, 2020).
- Coursework: Data Structures and Algorithms, Data Analysis, Data Visualization, Economics, Econometrics, and Machine Learning.

EXPERIENCE

Amazon **May 2022 — Aug 2022**
Software Development Engineer Intern *Seattle, WA*

- Developed a web application for AWS Lambda customers that helped to increase availability by 5% (Python).
- Handled 10PB+ of data for customer segmentation and created user profiles (SQL, Numpy and Pandas).
- Designed a ML pipeline using time-series to identify low availability customers with a 97% accuracy (Scikit-learn).
- Built and deployed a data analysis package that saved engineers +100 hours/week (EC2, Docker, S3, and Lambda).
- Worked in a Science Research team and wrote research and technical papers for engineers, scientists, and AWS stakeholders.

Sanford School of Public Policy at Duke University **Jan 2022 — May 2022**
Data Analysis Research Assistant *Durham, NC*

- Gathered and cleaned US Census and survey data to design social policies that improved accessibility to non-English speakers.
- Compiled more than 15,000,000 data-points in a database (R).
- Produced info-graphics and dashboards to convey information to the general public.

Data Science Research Center at Duke Kunshan University **May 2021 — May 2022**
Data Science Research Assistant *Shanghai, China*

- Prepared financial data from the S&P 500 for algorithmic trading (NumPy and Pandas) and achieved 22% return of investment.
- Assessed Reinforcement Learning and Supervised Learning algorithms for time-series forecasting (PyTorch and Scikit-learn).
- Evaluated different approaches to statistical arbitrage trading and optimized pairs trading selection time by more than 50%.

Duke Department of Physics **Dec 2020 — May 2021**
Software Engineer *Durham, NC*

- Built an image reconstruction and pattern recognition algorithm to detect Dark Matter particles 1000 times faster.
- Implemented ML algorithms into integrated circuits (FPGAs) to improve the performance of Large Hadron Collider.

PROJECTS

Real-Time Semantic Segmentation for Autonomous Vehicles **Jan 2022 — May 2022**
Duke Kunshan University *Kunshan, China*

- Led research to improve Semantic Segmentation algorithms that optimized the trade-off between accuracy and efficiency.
- Implemented contemporary Machine Learning models in Python (ICNet, BiSeNet, DDRNet, and DFANet).
- Conducted testing on Cityscapes (dataset) and achieved over 65% accuracy and less than 200 ms running time on each model.

Airbnb Customer Segmentation **May 2020 — July 2020**
Duke Kunshan University *Kunshan, China*

- Created 3D charts and graphs that facilitated the interpretation of different groups of suspicious users potentially running hotels.
- Enhanced the data segmentation process by applying dimensionality reduction (SVD) which reduced running time in 10%.

SKILLS

Programming Languages	Python, Java, C/C++, R, and SQL
Technologies	EC2, ECR, Docker, Lambda, S3, Git, and LaTeX
Quantitative Research	ETL, Data Visualization, Clustering, Regression, Statistics, and Time Series
Data Science Libraries	NumPy, SciPy, Pandas, Scikit-learn, PyTorch, Matplotlib, Seaborn, and Ggplot
Analysis Software	QuickSight, Tableau, OpenRefine, and Microsoft Excel
Languages	English (Fluent), Spanish (Native), and Chinese (Intermediate)

ACHIEVEMENTS

- National Finalist at the Alibaba GET Challenge (top 12 out of 250+ teams).
- Professional Certificate on AI Engineering by IBM.