# DENIZ TURKCAPAR

San Francisco, CA ▪ dturkcapar@gmail.com ▪ https://denizturkcapar.github.io ▪ +1 773 524 8459


## SKILLS

**Programming Languages / Skills:** Fluent in Python, SQL, Java; experienced in C, Ruby, R, JavaScript, C++, statistical modeling, A/B testing, market/user research, data mining, project management
**Technologies and Frameworks:** Apache Kafka, Docker, Kubernetes, AWS (EMR, S3 & Elasticsearch), Git, REST API

## WORK EXPERIENCE

**SALESFORCE** Sep 2021 - Current
**Software Engineer** *San Francisco, CA*

- Owned and led the architecture design and implementation of translation data change capture, enabling reliably captured data with 0 event loss by leveraging Kafka
- Built the Incremental Search Indexing feature, which decreased the search indexing time by up to 90%
- Won company-wide hackathon for two consecutive years with Trending Search and Synonym Search AI-focused projects
- Productionized and implemented Semantic Search feature, which utilizes nearest neighbor estimation, on Salesforce Commerce, increasing click-through-rate by 35%
- Implemented intelligent SKU (stock keeping unit) Search, increasing findability rate by 35% in B2B Commerce Search

**SALESFORCE** Jun 2020 - Aug 2020
**Software Engineer Intern** *San Francisco, CA*

- Developed Sort Rules feature end-to-end for B2B Commerce Search with RESTful APIs for sort rules retrieval, one of the three major features rolled out, within the search indexing schema using Java, SQL, Elasticsearch and Spark, enabling sorting by relevance and alphabetical order for the first time
- Improved the operation time p95 (measured in milliseconds) 24% by optimizing the performance of resource-extensive classes (such as cutting down expensive calls) in the query pipeline
- Wrote automation tests to cover any possible regression with the introduction of the new feature, committing to 100% no regression

**PAYPAL** Jun 2019 - Aug 2019
**Software Engineer Intern** *Chicago, IL*

- Implemented a Cronjob in Kubernetes to send heartbeat messages every 5 seconds and defined a DataDog plug-in, enabling the team to have data on Kafka topic lags in an interactive graph in terms of seconds of lag rather than message count for the first time
- Created and deployed sane memory related defaults and a dynamically adjusting heap size for certain topics in Kafka to eliminate crash loop back off by over 95% in Kubernetes pods
- Maintained data streaming replication over multiple data centers and AWS regions to make all Kafka data (across physical data centers & AWS) readily accessible to all teams and applications

## RESEARCH & PROJECT

- Published research on Fuzzy Matching Problems with Predicate Constraints, such as missing data, using Graph Theory to improve matching accuracy by 60% compared to naive matching
- Trained a logistic regression model using various hypothesized key factors to predict Yelp Elite status with 97.8% prediction accuracy, and a k-nearest-neighbors algorithm to find the most similar users with 95% accuracy in grouping Elite Yelp members together

## EDUCATION

**The University of Chicago**
Bachelor of Science in Data Science Specialization in Economics and Computer Science Jun 2021
**Honors & Accomplishments:** Dean's List, Facebook University Data Analytics Batch Day Finalist 2019, Google Engineering Practicum Finalist 2019, Grace Hopper Scholar 2019-2020, Jeff Metcalf Fellow 2017-2020
**Relevant Courses:** Machine Learning, Data Science for Computer Scientists, Big Data Analysis, Natural Language Processing (NLP), Software Development, Theory of Algorithms, Discrete Math, Statistical Models/Methods, Econometrics, Linear Algebra, Calculus