Christopher Kok

(765) 413-6666 | pherkok@gmail.com | github.com/chriskok | chriskok.site

Education

BACHELOR OF SCIENCE IN COMPUTER SCIENCE | PURDUE UNIVERSITY | GPA: 3.97 **JAN 2017 - MAY 2019**

- Graduated with Highest Distinction (top 3% of graduating class) and a concentration in Machine Intelligence
- Favorite courses: Data Mining & Machine Learning, Data Structures and Algorithms, Web Information Search and Retrieval

Work Experience

LEAD MACHINE LEARNING ENGINEER | RAYTHEON TECHNOLOGIES **JUL 2019 – CURRENT**

- Contributing to the full development life cycle of machine learning solutions for various million-dollar aerospace contracts
- Developing a scalable, containerized, and cloud-based backend infrastructure that serves the FAA's weather camera data collection and analysis system; efficiently processes about 360K images daily
- Building performant end-to-end machine learning systems for several aviation-specific domains (weather forecasting, cabin behavior, pilot health, software requirements, etc.)
- Consulting with clients to prototype, A/B test, and refine programs based on discussed business and technical requirements
- Raised $250K to lead a passionate team of 25 engineers and technical fellows in the AI/ML working group. We are creating platforms and processes for networking, collaboration, information centralization, best practices, and technical road maps.
- Leading the AI/ML Community of Practice (CoP) – the largest active CoP with 1000+ members and 300+ monthly attendees. We are establishing connections and aligning goals of ML/AI leaders across the enterprise

MACHINE LEARNING ENGINEER | REKKU: ANIME RECOMMENDATION SYSTEM **AUG 2020 - CURRENT**

- Developing an interpretable content-based anime recommendation API that is integrated with a popular anime recommendation website (randomanime.org) averaging 75K unique users monthly
- Implementing an extensive anime knowledge graph with state-of-the-art NLP techniques using gensim and networkx
- Establishing a scalable data ingestion pipeline from multiple anime-specific data sources
- Utilizing FastAPI, Docker, AWS EC2, and REACT to create an efficient and portable prototype web application and API

DATA SCIENCE INTERN | ROCKWELL COLLINS **MAY 2018 – AUG 2018**

- Engineered a containerized web application to help student pilots measure their performance individually and against their peers as well as help instructors measure their classes' progress using Flask and Docker
- Developed interactive and dynamic data visualizations of the relevant data with D3.js
- Implemented several machine learning techniques to predict common patterns among students using Keras

Personal Projects

FOUNDER | PROPEL PROJECTS **SEP 2019 - CURRENT**

- Propel is an accessible fellowship program for beginners to learn by building end-to-end applications (propelprojects.org)
- Leading multi-disciplinary teams (of UX, Frontend, and Backend roles) to deploy unique, data-driven applications in 8-week sprints. We're currently on our 4th cohort with over 20 user-validated products in our roster
- Mentoring a diverse group of students and career changers from over 20 countries and 5 continents
- Managing a community of over 700 developers and designers; holding meaningful discussions on research and application

Skills

- **Languages**: Python, SQL, Java, Javascript, HTML, C#, C, C++, R
- **Tools**: Docker, Cassandra, AWS EC2, S3, Lambda, IAM, Spark, Kafka, Linux, SAS, Matlab, Hadoop, Storm, D3.js, REACT
- **Packages**: Scikit-Learn, NumPy, SciPy, Pandas, Tensorflow, Keras, NLTK, Flask, Jinja, Bootstrap, Jupyter Notebook, FastAPI
- **Machine Learning Specialty**: Natural Language Processing, Recommendation Systems, Meta-Learning, Neurosymbolic AI