

Kaiser Stefan Pister

A founding engineer and consultant with proficiency in building projects and teams from the ground up. I've started 3 companies (failed, sidelined, succeeding). My passion is language: both natural and robot.

EXPERIENCE

Machine Learning Consultant *July '21 - Present; Python*

Multimodal Layout Understanding: Reimplemented Microsoft's LayoutLMV3 and LayoutReader papers to solve classification and reading order problems on a novel dataset. Achieved **87% μ F1** on classification, significantly outperforming legacy systems.

Closed Domain Restaurant Service Bot: Built an annotation system to log and process large unstructured data captured during drive-thru conversations to build a service AI which can accurately take a customer's fast food order.

Language Clustering: Built a tagging system to annotate data in a multi-class, multi-label system. Used research techniques such as deep networks, masking to prevent label bleed, active learning, and hard negative mining to improve model performance.

MLOps Migration: Joined a team to build out their AWS EMR + PySpark infrastructure. Migrated their services from Azure to AWS in less than a week, organizing all infrastructure into Terraform.

Semantic Search: Worked with a small startup to help guide their data collection plans. Designed large scale experiments to prove and disprove potential ideas in their business model. Built data pipelines for complicated free form text data. Quickly implemented proof of concept ideas and proposed best next steps.

Lecturer — UW Madison & UC San Diego *June '19, Sept '22 - Present*

98% Recommendation rating from students (n=50)

Programming Languages (Principles and Paradigms): I teach a fourth year course on the theory and design of programming languages covering fundamentals of functional and procedural programming, the meaning of object oriented programming and other PL topics like memory management and program synthesis.

Founding Engineer — Audiodub *March '20 - March '22*

Amberola: Designed & built systems that served **2M daily users** and generated 10k audio articles daily. Increased engagement rate to **50% at 50%**. Built world class Hebrew text to speech engine, improving phoneme pronunciation from 60% to **93% accuracy**.

Video Translation: Built speech-to-speech translation engine. Worked with YouTubers to translate their videos into foreign languages. Channel reach **10k total subscribers and 1 million views**.

CONTACT

<https://github.com/kpister>
kaiser@pister.dev

EDUCATION

UC San Diego

M.S. NLP & System Security
2018 - 2019

B.S. Computer Science
Highest Distinction & cum laude
2014 - 2018

LANGUAGES

[Python, Javascript, Haskell]
+ [Spanish, Greek, Latin]

2019 and before

Research Assistant in a NLP lab. Focused on normalizing protein databases using siamese LSTM networks ([published](#)). Also worked in a security lab, building distributed secure systems.

Project Lead for a course project. Managed meetings, organized teams, and achieved 1st place in cohort, 3rd place in class, for our app in school competition

Built a simple C compiler from scratch.

Survived a Summer using only Arch linux

Magitech: Developed a CNN to analyze and predict attributes of Magic the Gathering playing cards.

Frontiers of Innovation Scholars Fellow, 2017 (digital archaeology)