Alec **Hodgkinson**

SENIOR SOFTWARE ENGINEER

☐ (+1) 650-255-2445 | ✉ hodgka@gmail.com | 🏠 alechodgkinson.com | ⬤ hodgka | in hodgka

A backend software engineer with extensive experience in machine learning, adept at designing and deploying scalable solutions across diverse domains.

Experience

Perimetrics.ai *Jan 2023 - Present*

SOFTWARE ENGINEER

- Implemented Kubernetes company wide and set up helm charts for a variety of services.
- Architected and implemented comprehensive machine learning infrastructure, encompassing data ingestion and inference pipelines.
- Created a variety of transformer models for both image classification and natural language translation tasks.
- Collaborated closely with subject matter experts to annotate data and establish a feedback loop, resulting in rapid model improvements.

Thumbtack *Feb 2022 - Dec 2022*

SOFTWARE ENGINEER

- Led a cross-functional team in migrating multiple services to utilize machine learning models in production.
- Took user facing features through the whole scoping to production life cycle.
- Built thrift APIs in Go to expose machine learning models for use by other teams.

Dawnlight *Aug 2021 - Dec 2021*

SENIOR MACHINE LEARNING ENGINEER, PROJECT LEAD

- Designed and implemented a prototype automated pipeline for data annotation and data pre-processing capable of handling 50 Tb per week.
- Led and coached junior engineers in writing maintainable code that minimized technical debt.
- Designed and implemented scalable ML Ops infrastructure to support a team of data scientists.
- Orchestrated, vetted, and managed multiple 3rd party vendors for data annotation.

MACHINE LEARNING ENGINEER *Aug 2020 - Jul 2021*

- Developed a flask and C++ based data collection tool for capturing synchronized data from multiple RealSense cameras.
- Designed and built an in-house machine learning platform on top of Kubeflow and Kubernetes, with plugins for DVC, Tensorboard, and Ray.
- Created prototype visualization tools to for researchers to use in their model development.
- Conducted 30+ interviews for a variety of roles including Research Scientist, ML Engineer, Frontend Developer, etc.

Panasonic Beta *Mar 2018 - Jul 2020*

HOME ACTION GENOME, PROJECT LEAD

- Managed activities in a team of seven, including engineering, process development, and data recording, and proposal writing.
- Prototyped a sensor suite and wrote custom firmware to communicate between the Pi and sensors over I2C.
- Wrote software to cleanly record and package video and sensor data into a video container for easy distribution.
- Developed a Flask-based front-end for recording synchronized data across multiple sensor suites, speeding up recording time by 50x.
- Created a baseline activity recognition model using PyTorch.

DOORGYM: A SCALABLE DOOR OPENING ENVIRONMENT AND BASELINE AGENT, CO-LEAD

- Developed initial proof of concept for door opening simulation.
- Developed a novel CNN-based vision pipeline in PyTorch for regressing 3D door knob location from multiple views. Achieved an accuracy of ±1.7 cm in simulation, accurate enough for an agent to open a door.
- Transferred vision neural network trained in simulation to real doorknobs with an accuracy of ±4.95cm.
- Performed ablation tests to verify the necessity of each component in the vision network pipeline.
- Created infrastructure for controlling Baxter Robots with neural networks.

Panasonic Silicon Valley Labs *Mar 2017 - Mar 2018*

DEEP LEARNING ENGINEER

- Designed a post-processing filter in PyTorch as a replacement for the SAO deblocking filter in H265, resulting in a 20% bitrate reduction.
- Implemented end-to-end deep learning models for image compression using Tensorflow and PyTorch.
- Filed three image compression related patents that are undergoing filing.
- Evaluated and adapted state of the art deep learning models for alleviating catastrophic forgetting in neural networks.
- Developed face tracking/identification software for a robotic platform running in real time at 30 FPS.
- Developed proof of concept demonstration using deep learning based monocular image height estimation running on a Raspberry Pi for presentation to CEO.
- Created a raindrop removal and object detection demo to be run embedded devices, such as a Jetson Nano.
- Implemented production ready multi-label classification models using ONNX and TensorRT, and various model compression schemes such as techer student distillation and weight pruning.

Skills

Languages	Python, C++, Go, LaTeX, Odin, JavaScript
Machine Learning	PyTorch, TensorRT, ONNX, NumPy, OpenCV, Scikit-Learn, Ray, Huggingface Transformers
Other Technologies	Flask, Git, Docker, Bazel, Azure, AWS, GCP, ZeroMQ, Jenkins, Kubernetes