

Tristan Rice

Location

Seattle Area › Kirkland, WA

Software Engineer, ML

rice@fn.lc
github.com/d4l3k

Work Experience

Staff Software Engineer, Machine Learning Infrastructure, [Cruise](#)

January 2023 - Present

- Tech lead for Training Framework and subject matter expert on ML and PyTorch.
- Lead cross-team efforts to make key improvements across Training Framework and partner ML Infrastructure teams, enabling 20x faster development times for business critical models by the Perception and Behaviors orgs.
- Worked within the Vehicle MTL team to overhaul their models and processes, greatly improving development speed and enabling them to ship critical emergency vehicle detection improvements.
- Mentored engineers on the team to help them grow both technical and collaborative skills.

Software Engineer, [Meta](#)

October 2018 - January 2023

PyTorch — Research to Production / Ecosystem

- Tech lead for [TorchX](#), [torch::deploy](#), [torch.package](#) and [TorchElastic](#).
- Grew [TorchX](#) to ~60k production distributed training jobs per day. Collaborated with internal teams and external companies to enable distributed training of PyTorch models.
- Lead the torch::deploy team and shipped the Beta release of torch::deploy with PyTorch 1.13, which is now used as the default inference solution for production TorchRec models.

Distributed Training Platform

- Increased total company revenue by ~0.86% through a novel pipeline parallelism and automatic device placement technique, allowing for 2x model complexity while reducing training time by 30%.
- Designed and led a project to replace the distributed control and data plane for the largest ranking model training jobs accounting for 80% of company revenue. Project running on ~80k machines in production for the past three years.
- Improved per-machine training performance by ~35% by eliminating memory copies, thread communication, thread and batch size tuning, and zero-copy networking.

Performance & Capacity Engineering Intern, [Facebook](#)

Summer 2017

- Created a machine learning model for resource usage prediction of large scale machine learning jobs running on [FBLearner](#). Increased server utilization by up to 400%.

Software Engineering Intern, [Cockroach Labs](#)

Summer 2016

- Worked on stability of CockroachDB's distributed KV-store and the Raft consensus algorithm.

Software Engineering Intern, [Gridspace](#)

Spring 2016

- Added seamless transfer of conference calls between servers to enable scalability.

Software Engineering Intern, [Google](#)

Summer 2015

- Worked on Knowledge Graph tools and data integrity. Increased productivity of thousands of contractors.

Software Engineering Intern, [Socrata](#)

Summer 2014

- Created complex, interactive visualizations for slicing datasets arbitrarily.

Personal Projects

torchdrive github.com/d4l3k/torchdrive

2021 - Present

- Open-source project providing high-level primitives for training self-driving and 3D world models.
- Multi-camera fusion with temporal models using autolabeling and differential rendering techniques. Entire stack created from scratch including data capture, modeling, training, and inference. Rooted a Tesla Model 3 to capture data.
- 3D Transformers, NeRF, Differentiable Rendering, SfM, Object Detection, Semantic Segmentation, CNNs, Voxels, Depth, SimpleBEV, BEVFormer, ThreeJS

Expertise

Languages

- *Proficient:* Go, Python, C++, C, TypeScript, JavaScript, SQL, Bash
- *Working knowledge:* CUDA, Hack/PHP, Java, Scala, R, Ruby

Technologies

- *Proficient:* PyTorch Distributed, Linux, Docker, Kubernetes, PostgreSQL, Redis, Git, AWS, GCP
- *Working knowledge:* TensorFlow, Ray, Ghidra, nginx, Hive/Presto, CockroachDB, Slurm

Education

University of British Columbia

Class of 2018

- BSc. Computer Science. 3.95 GPA