

Aniruddh Chandratre

Palo Alto, CA · +1 (480) 584-2618 · achandratre22@gmail.com · github.com/C-Aniruddh · [Google Scholar](#)
0-1A (current) · EB-1A approved

ML Infrastructure engineer. I own the internal data processing platform behind Tesla AI: every data operation for **FSD, Optimus, and Digital Optimus** — running **10B+ jobs** and **3,000+ years of compute** a month across a fleet of **10 large GPU clusters**, roughly **250K H100-equivalent GPUs** commanding **~\$200M/month** of compute at list price. Re-architecture, scheduling, and cluster efficiency work took throughput up **17×** and effective utilization to **98%**, unlocking hundreds of millions of dollars of AI compute without buying hardware.

EXPERIENCE

Senior Engineer, ML Infrastructure — Tesla AI Feb 2024 – Present · Palo Alto, CA
Data platform & GPU cluster efficiency · ~250K H100-equivalent GPUs (A100 · H100 · H200 · GB300 · MI325)

- Own end to end the internal data processing platform running **10B+ jobs** and **3,000+ years of compute** per month across **10** large GPU clusters — **all** data operations (acquisition, curation, mining, auto-labeling, training, evaluation) for the FSD, Optimus, and Digital Optimus tracks.
- Proposed and executed a ground-up re-architecture — lower-level core, distributed worker coordination in place of centralized services, rebuilt memory model. Throughput **940M → 15B+** data items/month (**17×**); maximum dataset per run **1M → 1B+ clips**, three orders of magnitude, making dataset sizes that previously could not complete routine.
- Built the cluster scheduling and live defragmentation layer that compacts allocations back into contiguous topological blocks while clusters stay hot, across a heterogeneous fleet normalized to H100 equivalents. Usable compute capacity rose **24% → 88%** and effective utilization now averages **98%** fleet-wide.
- Compute commanded by the platform grew **~\$5M → ~\$200M per month** at public cloud list price — **hundreds of millions of dollars of AI compute unlocked through software rather than hardware** — plus an estimated **\$15M/year** in recovered engineering productivity. Tesla's VP of AI Software credited the redesign with “*pushing Tesla's AI infrastructure to a new level*”.
- Built lineage tracking across artifacts, datasets, training jobs, and models, giving multi-stage pipelines end-to-end provenance for the first time.

Software Engineer, Automation Software — Tesla Jan 2023 – Jan 2024 · Fremont, CA
Edge-to-cloud data infrastructure for the gigafactories

- Designed and shipped an any-to-any edge-to-cloud data broker moving **10B+ manufacturing data points/day** — MQTT into InfluxDB, Kafka into ClickHouse — in production across the Austin, Berlin, and Fremont gigafactories.
- Wrote its low-level primitives from scratch when available options could not sustain the required throughput on this storage profile: an mmap-backed durable queue sustaining **5M writes/s** on NFS-backed Kubernetes volumes roughly **100×** slower than an SSD, and a distributed cache at **10M entries/s** that deduplicated across **50** deployments and cut storage **~80%**.
- Built and operated **20** high-availability services for factory insights across global sites, owning the full deployment lifecycle.

Graduate Research Assistant — ASU, Cyber-Physical Systems Lab Jan 2021 – Dec 2022 · Tempe, AZ

- DARPA **ARCOS** (Automated Rapid Certification of Software) with Lockheed Martin. Co-authored PSY-TaLiRo, a falsification toolbox for cyber-physical systems still in use in Lockheed's CertGATE pipeline. Ran the bare-metal Kubernetes cluster behind high-volume formal-verification simulations.

Founding ML Engineer — AfterShoot Aug 2020 – Jan 2021 · Remote

- Built the image-ranking model still at the product's core six years later. AfterShoot now serves **~250K** photographers, has processed **9B+** images, and is valued at **\$55M**.

EDUCATION

M.S. Robotics & Autonomous Systems (AI) — Arizona State University · 4.0/4.0, Distinction Medal 2022

B.Tech Electronics & Telecommunications — NMIMS MPSTME, Mumbai · Dean's List 2020

RESEARCH, AWARDS & TOOLING

PUBLICATIONS 8 peer-reviewed papers · 174 citations · h-index 5 — HSCC, EMSOFT, FMICS, ARCH. First author, HSCC 2023: *Stealthy Attacks Formalized as STL Formulas for Falsification of CPS Security*.

AWARDS 5× national hackathon winner in India — Smart India Hackathon 2019, Mastek Deep Blue (×2), IET Hack N Code, RAENG.

LANGUAGES Go, Python, C, C++, JavaScript/TypeScript, Java

SYSTEMS Slurm, Kubernetes, Helm, Docker, AWS, GCP, Prometheus, Grafana, Kafka, ClickHouse, InfluxDB, MQTT, SQL