ML Engineering

Hop develops algorithms at scale – reliably, reproducibly and responsibly

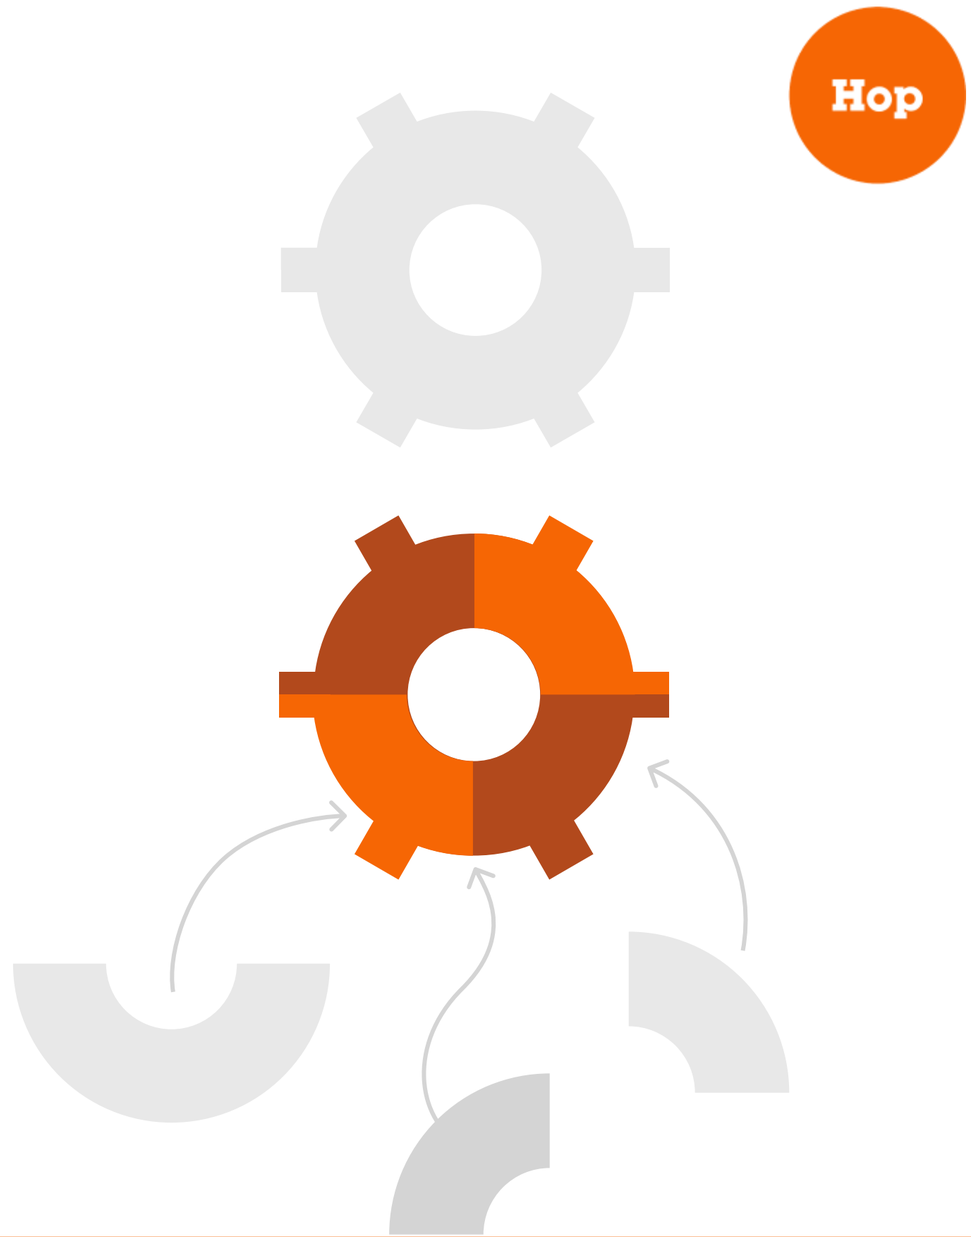

Here at Hop, we build production-scale systems to deploy machine learning at scale – whether that's on premise, in the cloud or on device.

Sometimes, this also requires us to construct novel compute substrates to explore more interesting research questions. Our engineers focus on questions of scale, latency, concurrency and resilience. Though we have preferences (spoiler alert: PyTorch/Python/AWS), we're generally language- and platform-agnostic, and have worked deeply in AWS, GCP, Azure and Heroku, as well as various on-premise installations.

Featured Case Study

TOYOTA RESEARCH INSTITUTE

Accelerating Research in Autonomous Driving

Working closely with Toyota Research Institute's Human Interactive Driving division, we've provided advanced engineering and operations support to scale and accelerate their machine learning research efforts.

SUMMARY

- Toyota Research Institute (TRI) is doing cutting-edge R&D work in the autonomous driving space.
- Hop has worked as an embedded part of TRI's team, building the infrastructure to accelerate their work.
- TRI needed advanced engineering and operations support to scale its machine learning research efforts.
- Our collaboration has resulted in deployment in TRI research vehicles, as well as time and cost savings for TRI.

READ THE FULL CASE STUDY ↗

Does your team need additional ML engineering capacity?

Contact us to learn how Hop can help.