

Senseye pushes machines to the limit by training models with high volume HD video footage. With Saturn Cloud, their pipeline runtime was **reduced from 60 days to 11 hours**.

Taking runtime down from 60 days to 11 hours is such an incredible improvement. We are able to fit in many more iterations on our models. This has a significant positive impact on the effectiveness of our product, which takes many iterations to perform at the standard necessary for our customers."



SETH WEISBERG, Ph.D.
Principal ML Scientist, Senseye

DATA VOLUME

900 ultra-high-definition videos

Every 10 minutes
in duration

Recorded at 120Hz,
nearly 7.5 TB on disk


BEFORE SATURN CLOUD

- ML pipeline using an on-premise GPU machine (Titan RTX)
- **~90 minutes per video per GPU, or a total of 60 days**

WITH SATURN CLOUD

- ML pipeline using Saturn Cloud: A distributed cluster of 160 T4 GPUs on AWS
- Only 10 new lines of Python code
- **40 seconds per video, or a total of ~11 hours**

SUMMARY ML PIPELINE STATS

System	Runtime Per Video	Total Runtime
On-Prem RTX Titan GPU	90 Minutes	60 Days
 Saturn Cluster (160 T4 GPUs)	40 Seconds	11 Hours