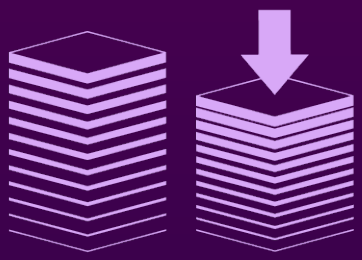

Lightscline

Data Reduction AI

SWaP-C/edge AI by processing on 10x reduced datasets

March 2024

Lightscline, State College, PA, USA

About Lightscline

Lightscline's lightweight AI solves the foundational problem of efficiently analyzing Terabytes of data being generated from dual use applications like satellites, drones, industrial assets, and underwater vehicles. Using Lightscline AI's 4 lines of code that can be setup within 10 minutes, customers can reduce 90% of their sensor data infra & human time and costs by selectively focusing on just 10% of the raw data. By exploiting the redundancy of real-world data, Lightscline AI makes real-time predictions using just 10% important data.

Executive Summary

In this whitepaper, we analyze Lightscline AI's real-time edge computing and SWaP-C capabilities for real-world AI applications generating large amounts of sensor data. We benchmark the performance metrics on a small single board computer: Nvidia Jetson Nano, and a more powerful Intel i7 processor. We demonstrate that selectively analyzing a small fraction of raw data leads to 10-74x benefits in edge compute efficiency and speed, bandwidth requirements, and enables applications which are not even possible today. We also show how we enable transfer learning on the edge and speed up the workflow by 40x per dataset.

Applications Overview

Lightscline's lightweight AI unlocks several capabilities across space, aerial, terrestrial, and underwater applications which are currently unobtainable. Some applications include:

1. Quickly analyze 200+ hours of acoustic data for anomaly detection and ATR using just 10% important data
2. Order-of-magnitude reduction in the amount of training data without degradation to identification performance (Pid)
3. Prioritize training data by selectively focusing on the 10% important data
4. Enable transfer learning on the edge to quickly train for complementary tasks
5. Real-time inference on 15+ classes of data for on-board Human Activity Recognition (HAR)
6. On-board satellite-based AIS signal validation using 10% of the raw RF data

Lightscline
Data Reduction AI

Satellites

On-board computing enables thermal anomaly detection, AIS signal validation, etc.

UAVs

On-board computing enables real-time ATR

Drones

On-board computing enables real-time ATR

Aerial fleet

Quickly identify anomalies from 300+ channels of data / flight

Human Activity Recognition

SWaP-C analytics capabilities

Naval fleet

Efficiently analyze TBs of data from large fleets

Key highlights:

1. We benchmark Lightscline AI for edge computing use-cases on Nvidia Jetson Nano and Intel i7.
2. Model training and transfer learning on the edge (Jetson Nano) are only possible when we use Lightscline's data reduction technique (10-50% of raw data). Both capabilities are unobtainable if we use >50% of the raw data, due to memory and compute restrictions.
3. Training times increase non-linearly with % data used and Nano runs out of resources when training for >50% data used.
4. Edge inference times are <0.20 seconds on Jetson Nano for all models that can be trained on it.
5. 10x reduction in bandwidth required for data transmission
6. 40x faster workflow / dataset leads to massive productivity increase (data scientists / machine learning / embedded engineers - more deployment)

Edge devices used for benchmarking:

We use an Nvidia Jetson Nano and an Intel i7 processor (11th Gen Intel(R) Core(TM) i7-1165G7 @ 2.80GHz 2.80 GHz) for benchmarking Lightscline AI for different edge computing use-cases. Their specifications are mentioned below:

	Jetson Nano	Intel i7
		
Power rating	5-10 W	125-190 W
Memory	4 GB 64-bit LPDDR4 25.6 GB/s	16 GB

Compute & latency:

We successfully train models that give >90% test accuracy on a 4-class classification problem on the edge (Jetson Nano) using a fraction of the raw data. In the figure below, we can see that model training is unobtainable on the edge when % data used >50%, due to the exhaustion of compute and memory. Edge compute time refers to the model training time.

Lightscline's selective AI (LS AI) unlocks model training capabilities on SWaP-C edge devices.

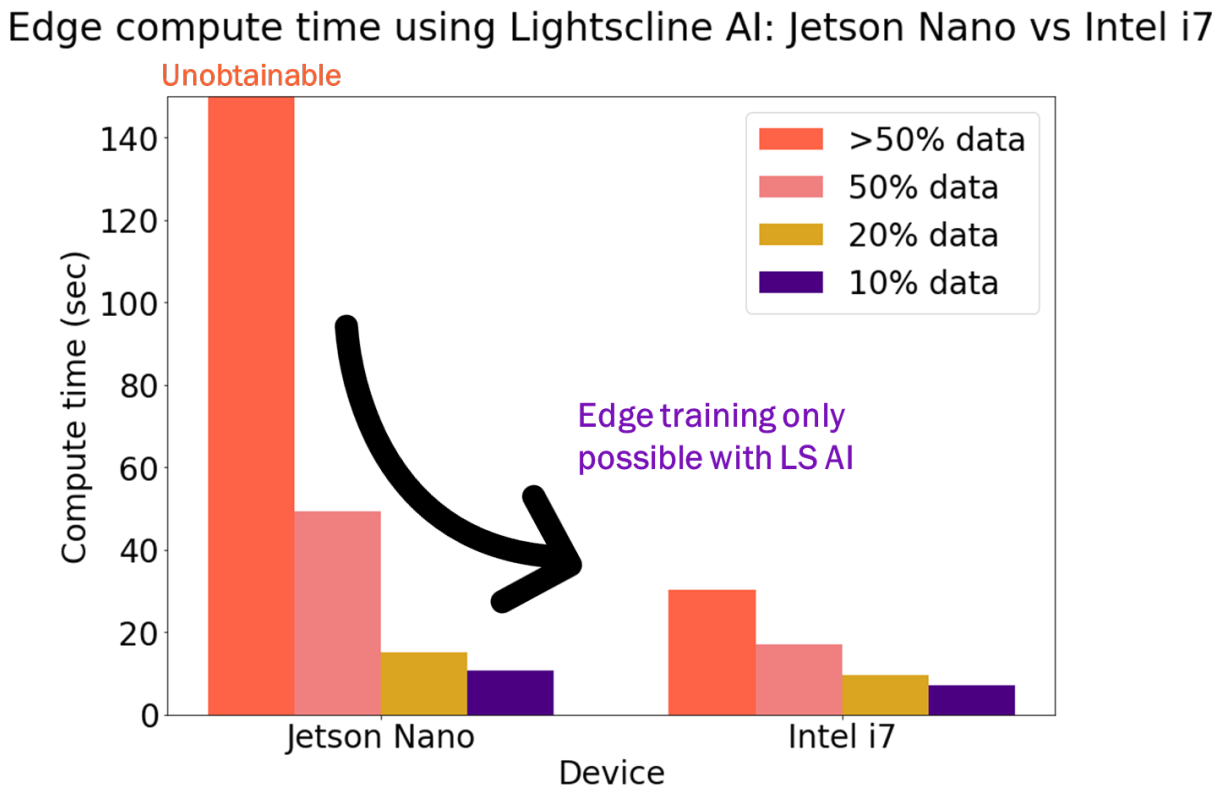

We successfully train transfer learning giving >90% test accuracy on a binary classification problem on the edge (Jetson Nano) using a fraction of the raw data. In the figure below, we can see that model training is unobtainable on the edge when % data used >50%, due to the exhaustion of compute and memory. Edge compute time refers to the model training time.

Lightscline's selective AI (LS AI) unlocks transfer learning capabilities on SWaP-C edge devices.

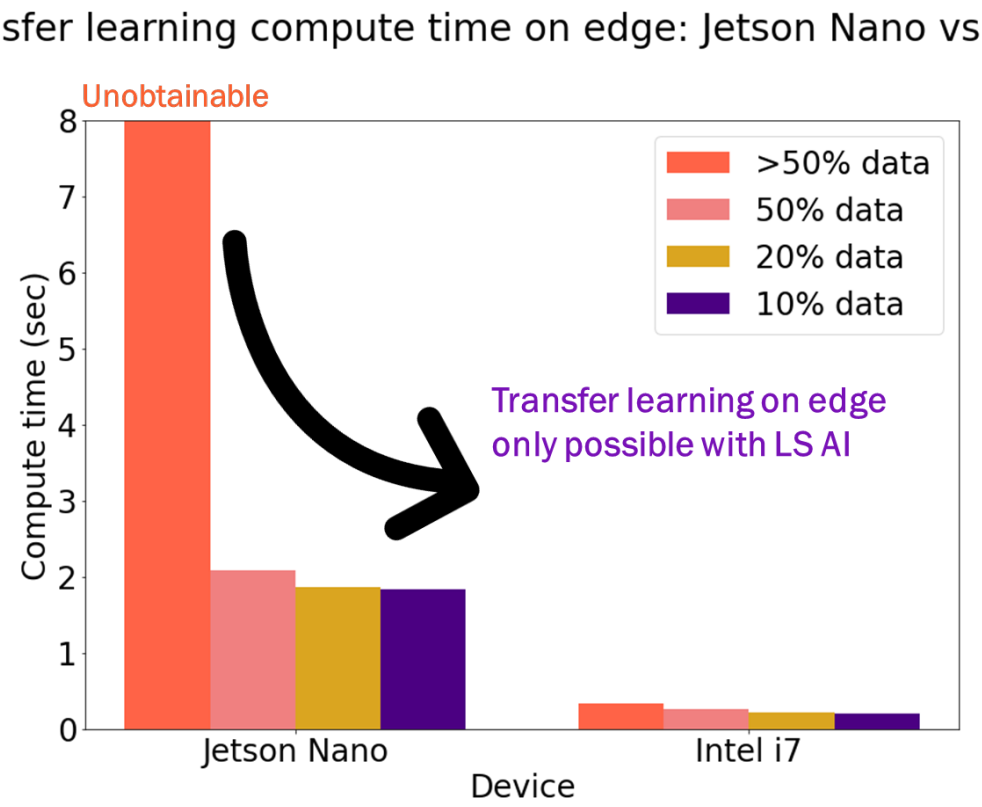

LS AI is up to 74x more efficient than a multi-layer perceptron while using up to 100x less data. These tests were conducted on the vibration data from the industry standard Case Western Reserve University bearing fault dataset.

Bandwidth:

We obtain a 10x reduction in the bandwidth requirements as Lightscline AI learns to selectively focus on just the 10% important data. This is especially useful in off-shore, underwater, lunar, and in-orbit transmission use-cases. Reduction in bandwidth significantly reduces the power consumed by transmission modules present on edge devices.

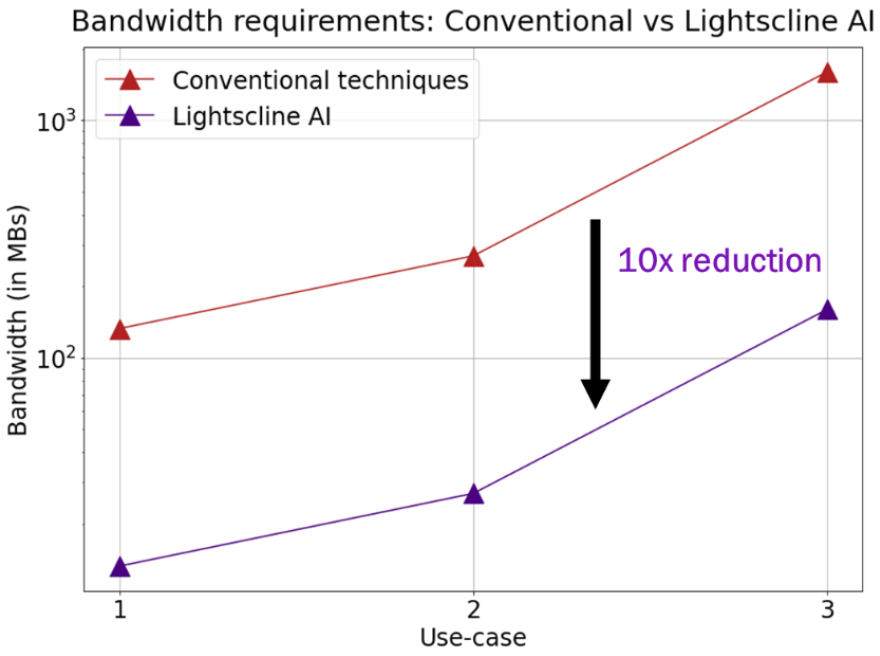

People time:

Conventional machine learning techniques require manual feature extraction and model training leading to 40+ people hours / dataset. Lightscline's selective AI makes end-to-end predictions using just 10% of the raw data resulting in a 40x faster workflow / dataset. This enables end-users to focus on deploying and scaling the number of use-cases rather than data pre-processing and manual feature extraction.

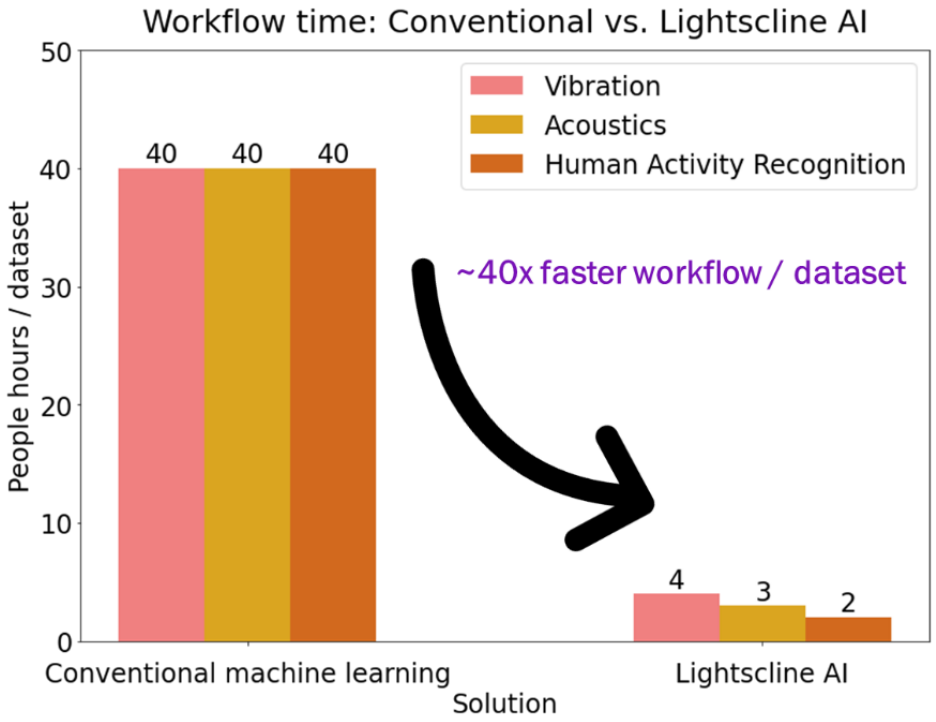

Product:

Users can get started with Lightscline AI using just 4 lines of code that can be setup within 10 minutes, without any need for external data sharing. The product can be run on the customer's cloud environment or on-prem.

Lightscline AI

```
from lightscline.lightscline import LightsclineEdge
## Load data into Lightscline
ls = Lightscline(data=data,fs = SAMPLING_FREQUENCY)
## Reduce the amount of data by 70% of the original
ls.reduce_and_preprocess_data(per_reduction=70)
## Train the model
ls.train_model(verbose=True,n_iters = 1000)
## checking the results
ls.test_model()
```

- 4 lines of code to get started
- Setup within 10 mins
- No data sharing required

Conclusion

Using just 4 lines of code, Lightscline AI performs real-time data collection through prediction with >10x speed and energy efficiency over conventional approaches governed by the Shannon-Nyquist sampling theorem. This leads to orders of magnitude savings in (i) data infrastructure costs and time, and (ii) human resource efficiency. Additionally, this enables several new applications not possible today due to extreme SWaP-C requirements of real-world AI applications. Try Lightscline's AI now. For full capabilities, click here.

Reach out to info@lightscline.com for any queries.