



Introduction

The call for sustainability!

- DL **energy consumption & carbon footprint** skyrocketing

Not all Joules are born equal.

- Grams of CO₂ emitted per Watt of electricity generated
- Varies depending on **time** and **location**


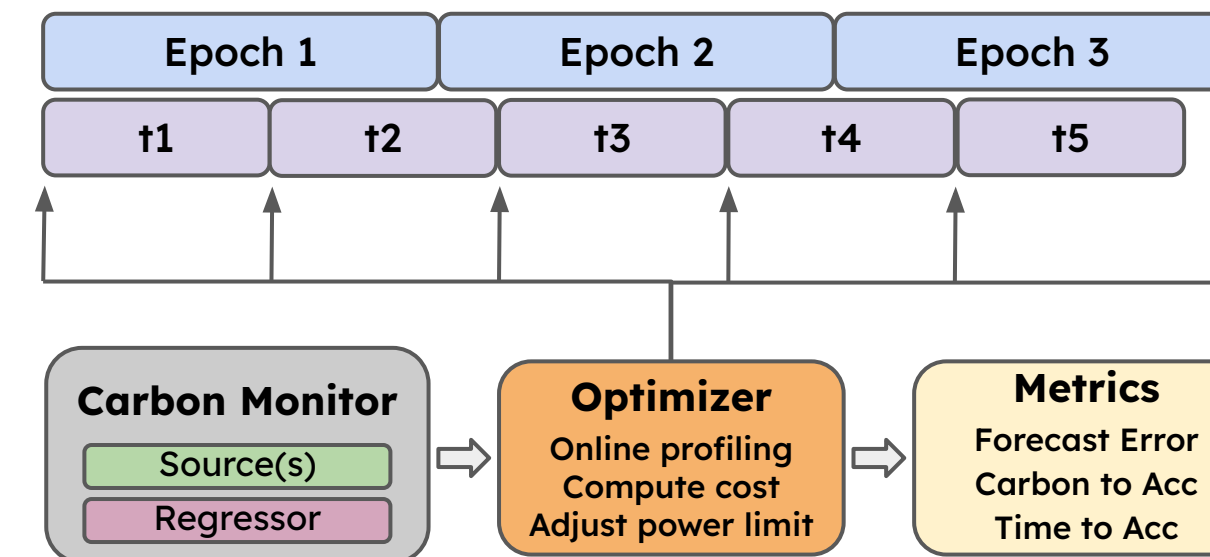
How do we chase clean electricity?

- **Make more progress** when carbon intensity is **low**
- Adjust **GPU power limit** as we train

Methodology

Carbon Intensity Forecasting

1. Retrieve **historical carbon data**.
2. Train a **regression model**.
3. Forecast carbon intensity for the next time period.



Carbon-Aware DNN Training

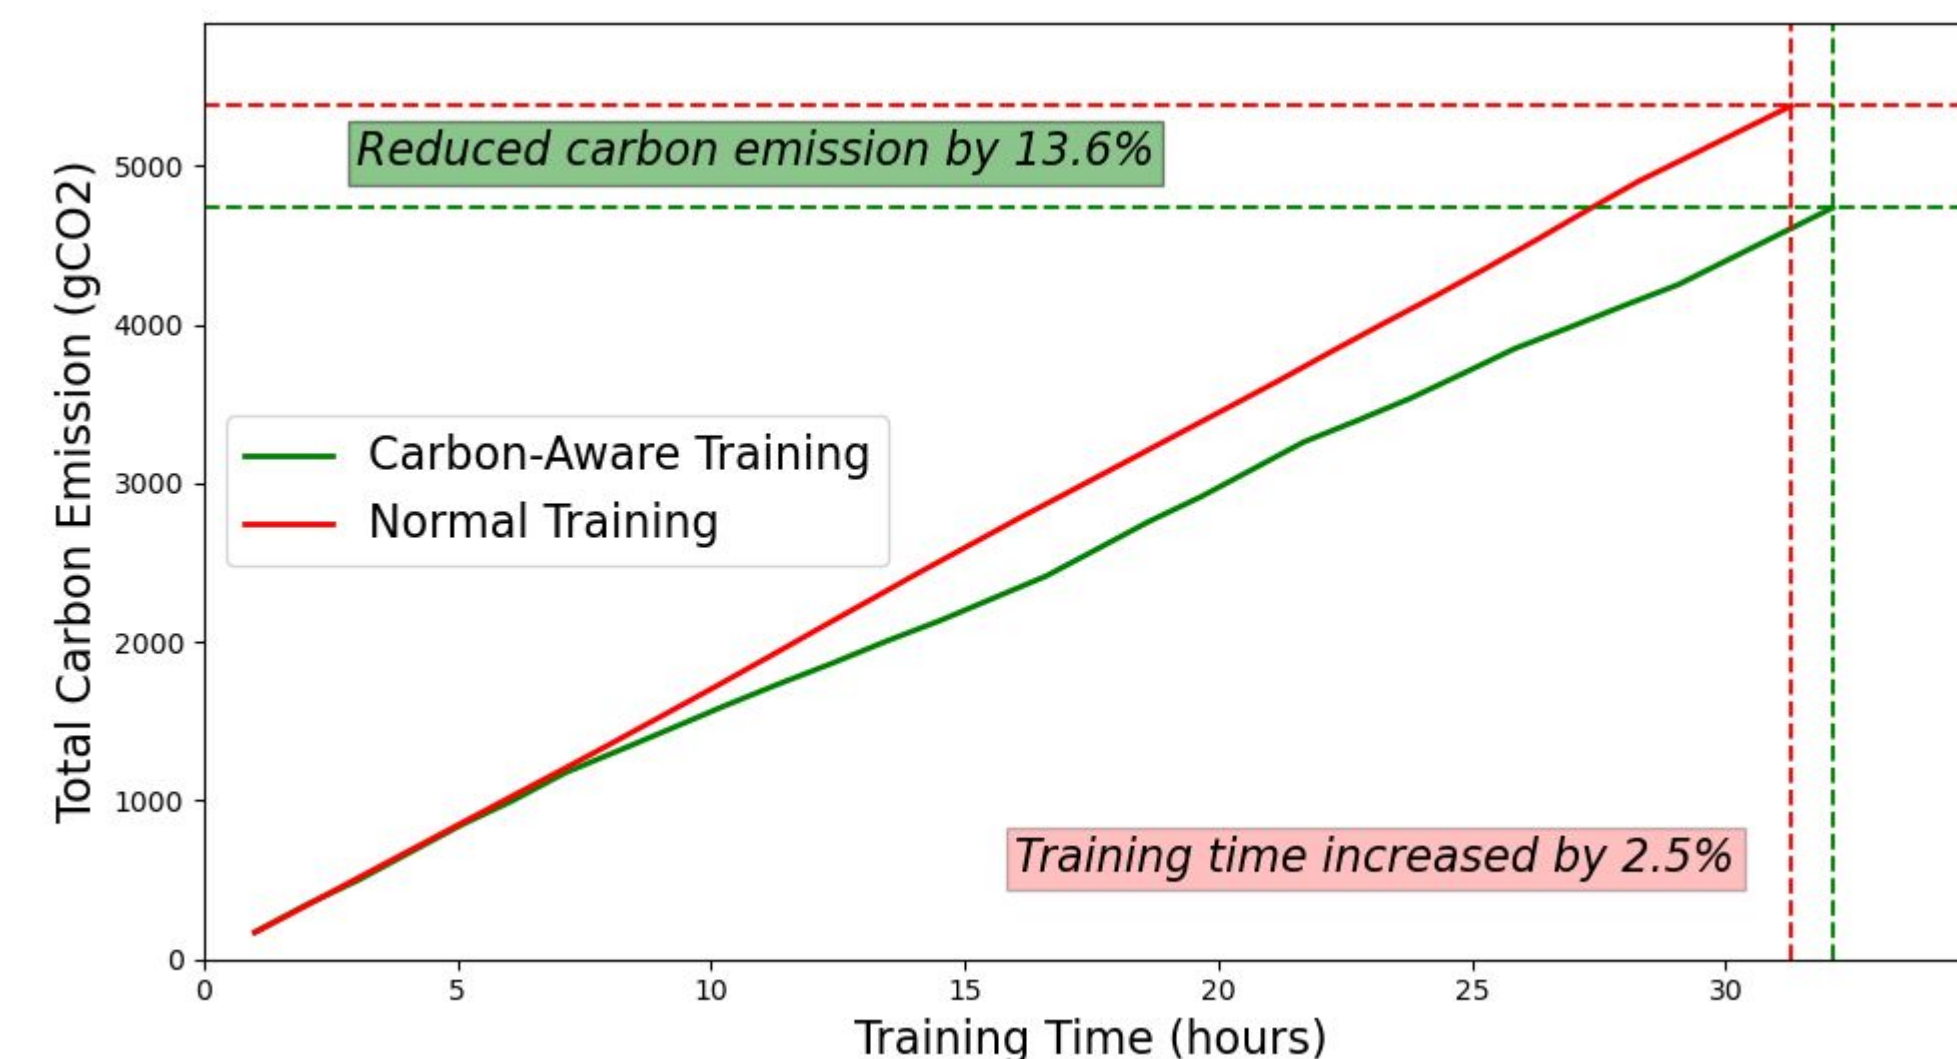
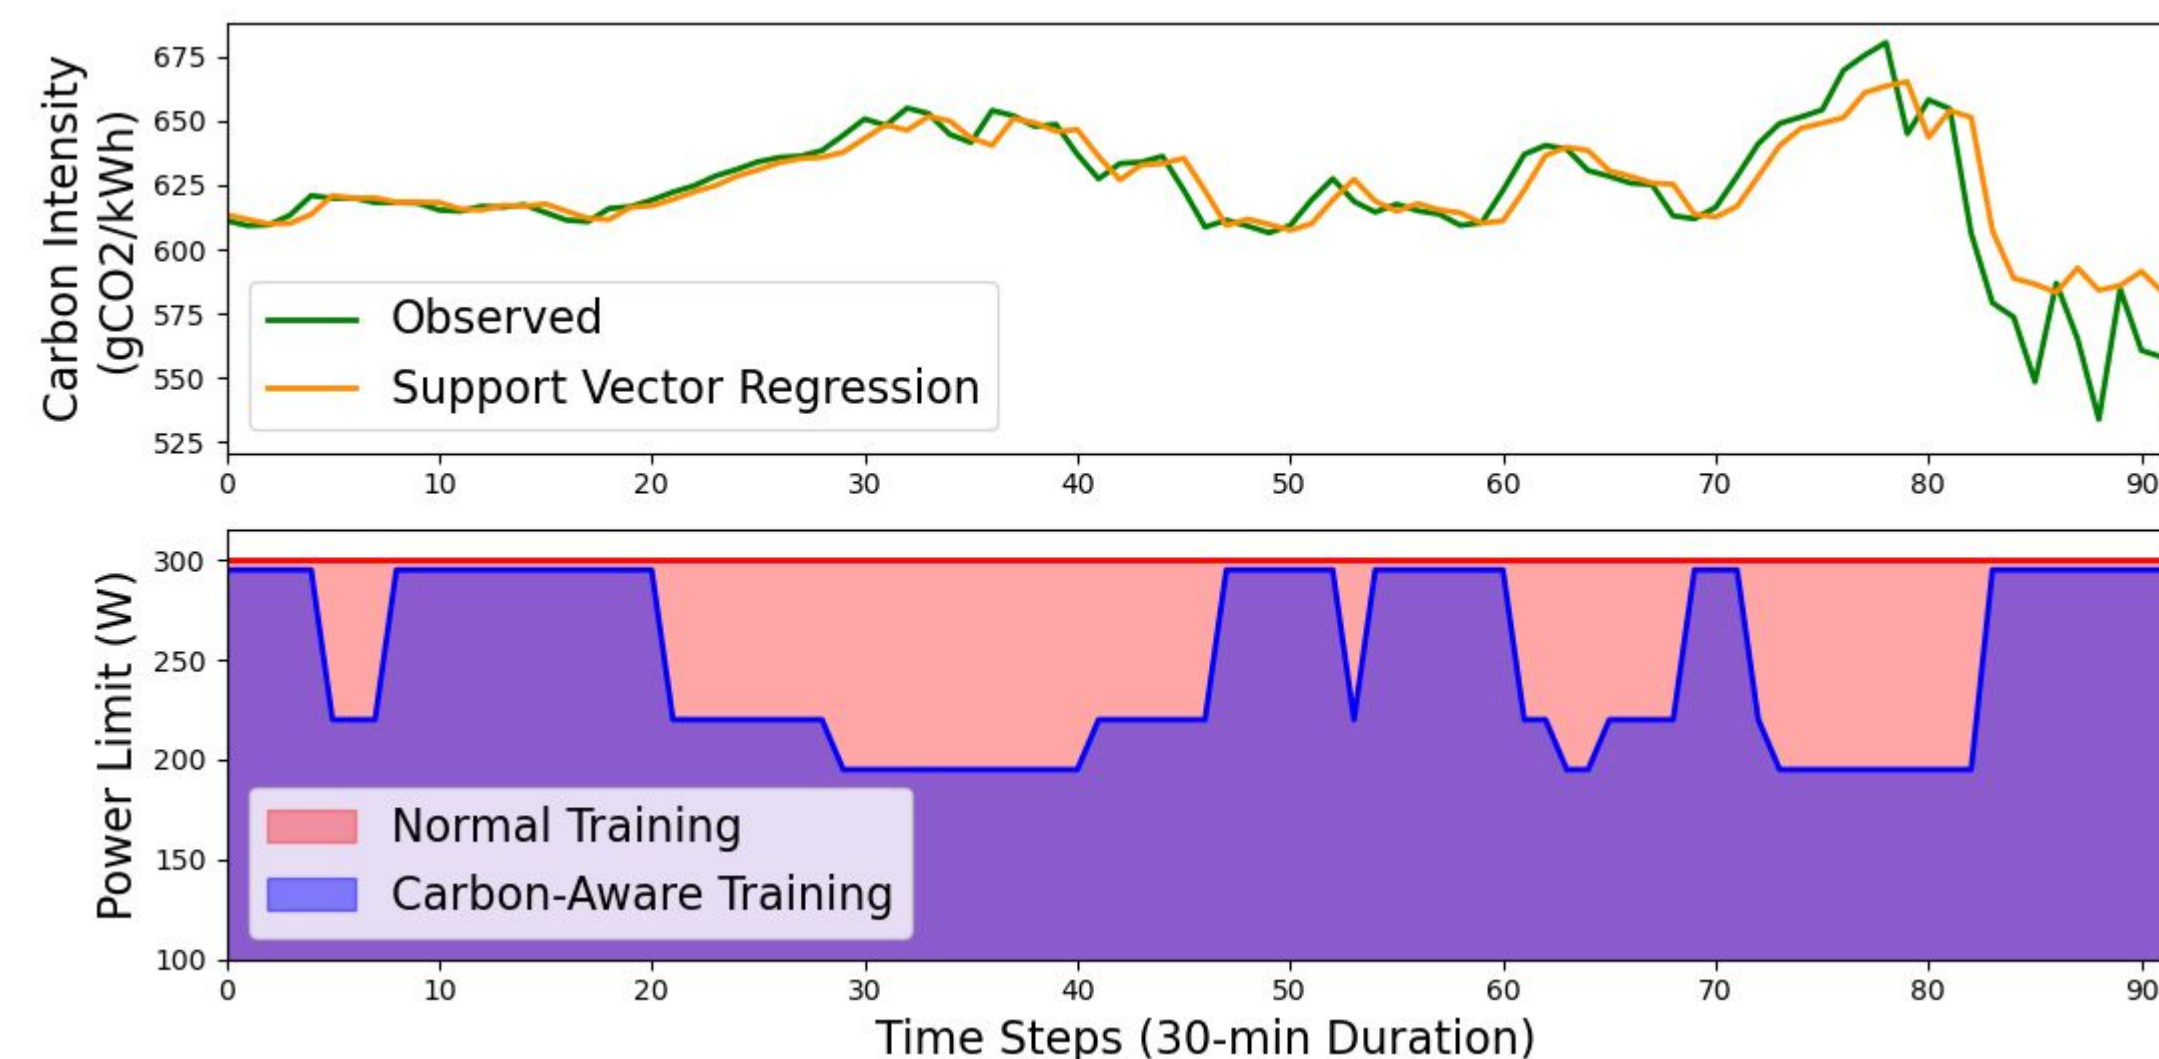
1. Profile **GPU throughput and power** for all power limits.
2. Compute **carbon cost** and determine the **optimal power limit** for current period.
3. Adjust the GPU power limit!

Carbon Forecasting Performance

Model	MAPE %
Support Vector Regression	0.94
Linear Regression	1.57
GradientBoosting	2.23
AdaBoost	2.51
Random Forest	1.76

252 hours of carbon intensity data points for testing

DNN Training Results (ResNet-50 on ImageNet)



Conclusion

- Chase dynamically adjusts the GPU power limit in real time;
No job migration or deferral
- Chase reduces carbon emissions of DNN training with
negligible training time degradation
- Chase works with **any DNN training jobs** with **minimum code modifications**

Contact Information

- Email: znyang@umich.edu