A Smart Cyberinfrastructure should use AI/ML to Reduce its Carbon Footprint and Negative Environmental Impacts

Andrew A Chien, University of Chicago
White Paper for the 2020 NSF Smart Cyberinfrastructure Workshop

Computing is one of the largest consumers of electrical power and in most parts of the world, the fastest growing consumer. While the largest fraction of this growth is due to the growth of cloud computing [Nature19], the scientific computing community's carbon footprint has grown rapidly, with a recent estimate for the Top 500 at 3 million metric tons of carbon per year. While the growth is due in part to the end of Moore's Law and the rising appreciation of the value of scientific computing, and there is a growing recognition that expanding use of AI/ML is a major contributor to this rapid increase [Amodei18,Hao19].

We believe that the NSF's Smart Cyberinfrastructure should use AI/ML to reduce its carbon footprint and negative environmental impacts. Specifically,

- AI/ML should be used to schedule workloads both at centers and at the edge to minimize the carbon impact of cyberinfrastructure computing resources, choosing location, time of day, and more to consume the least damaging electric power
- AI/ML should be used to manage and plan cyberinfrastructure resources so as to achieve scientific ends with the minimum quantity of e-waste, maximizing the useful lifetime of equipment and scheduling computations on the resources that minimize the consumption of electric power

In addition, the NSF's Smart Cyberinfrastructure should make aggressive use of renewable energy, exploiting breakthrough insights that show that batch and throughput scientific computing can be delivered economically with zero or near zero-carbon approaches [ZCCloud19,Yang16]. This insight has several implications and opportunities:

- NSF should exploit new types of renewable-energy models (stranded power), and new models of operation to achieve zero-carbon operation for large-scale computing resources, doing so can also reduce the power cost of operating these resources
- NSF Smart Cyberinfrastructure should provide national and international thought leadership and operational leadership for both government, educational, and even commercial community to reduce the carbon impact of computing
- NSF Smart Cyberinfrastructure adoption of these renewable based models represents the **opportunity** to maximize the ability to use compute-intensive AI/ML techniques for scientific and social good as well as to optimize scientific computing resource management

References

[Nature18] Jones, N. How to stop data centres from gobbling up the world's electricity. *Nature* (Sept. 12, 2018).

[Amodei18] Amodei and Hernandex, “Ai and compute,” openai.com/blog/ai-and- compute/, May 2018, 10x growth per year.

[Hao19] K. Hao, “Training a single ai model can emit as much carbon as five cars in their lifetimes,” Technology Review, June 2019.

[ZCCloud19] Chien, Zero-Carbon Cloud Project, http://zccloud.cs.uchicago.edu/

[Yang16] Yang, F. and Chien, A.A. ZCCloud: Exploring Wasted Green Power for High-Performance Computing, IPDPS, May 2016.