Canada’s Future DRI Ecosystem: AI Research Needs

Vector Institute Response

Authors: Ron Bodkin VP AI Engineering and CIO Vector Institute, Garth Gibson CEO Vector Institute

Overview

A huge driver of competitiveness in the increasingly dominant knowledge economy is computational excellence. In the past decade the role of computational artificial intelligence (AI) in economic development and market competitiveness has gone from small to dominating. The economic importance is hard to overstate — driven by innovations in neural network model training originating in Canada, massive changes have disrupted market sector leadership in many fields, including information search, speech recognition, natural language understanding, navigation assistants, self-driving vehicles, litigation preparation, manufacturing qualification, and others. It is essential to the Canadian economy that AI play a big role in the virtuous cycle of academic computational innovation to industrial competitiveness that advanced research computing (ARC) provides.

The impact on resourcing for ARC when AI’s outsized role is fully recognized is well outside the scope of past plans. In particular, spurred by the economic sea change AI has caused, the rate of change in AI technology is faster than in most science. Even as hardware improvements slow down from Moore’s Law (60% annual improvements in cost effectiveness) the increase in resources used for record breaking AI demonstrations has accelerated — OpenAI reports that the computational requirements of new record breaking demonstrations in important AI technologies like natural language processing are doubling in under 4 months, that is more than 700% per year — an increase much too fast to be achieved at constant cost. Countries like China and America are fuelling this spending with the national economic benefits from AI leadership in mind.

Canada has the advantage of its deep brain trust in AI, especially machine learning, deep learning and reinforcement learning. Deep Learning was essentially created in Canada. The Canadian government has sought to strengthen Canada’s leadership with its 2017 Pan-Canadian AI Strategy (PCAIS) and the three AI Institutes it instigated. The computational strategies of the AI Institutes is more nuanced than simple outspending as is done in other countries. The AI Institutes develop new foundational technologies from innovative approaches, demonstrating advantage on scales too small to attract the popular press, but large enough to change technology directions at research organizations around the world, and increasingly, through Canadian corporate partnerships with the AI Institutes, in the private sector.

This white paper provides an overview of the digital research infrastructure that is provided by the Vector Institute for its AI scientists, future needs and recommendations for how the Canadian DRI can enable the important mission of Canadian leadership in AI.

Existing State of DRI at Vector Institute

The Vector Institute has two computation clusters that are available for over 500 AI researchers at the institute. All Vector researchers are supervised by Vector faculty and start with equal priority so they can be productive immediately without requiring them to define a project and wait for an application to be granted. We are in the process of expanding access to this environment to nearly 100 faculty affiliates and their 500 students and we have extended access to industry collaborators as part of technology transfer projects.

In aggregate, Vector provides 1163 GPUs with 12.5 Petaflops of computational power for its researchers, with associated storage, CPU processing, and high bandwidth networking. Figure 1 illustrates the primary shared computational cluster that Vector provides for its faculty and researchers.

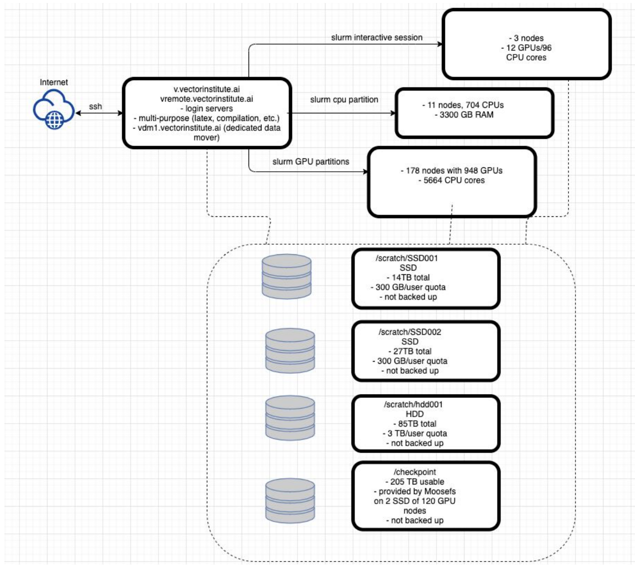

Figure 1. Vector’s Primary Computation Cluster as of Q4 2020 (960 of 1163 GPUs)

To support researcher effectiveness, Vector’s operations staff have developed and integrated numerous enhancements to the cluster environment and SLURM scheduler:

- A distributed file system to allow efficient checkpointing and means to integrate checkpointing into diverse AI tools and workloads
- Calibration of scheduling to support interactive and high priority needs while allowing researchers to run large-scale ongoing experiments and to manage peak loads prior to major conference deadlines
- GPU profiling and physical state analysis to understand job characteristics and to help users understand optimization needs
- Troubleshooting and resolution tools, e.g., to deal with running out of memory, stale processes, exhausting shared storage, excess storage or networking bandwidth usage, GPU error states, system errors, and excess use of memory or CPU on head nodes
- Integration of new GPU drivers and other software libraries for use after testing to avoid conflicts
- Providing frequent updates to shared infrastructure, e.g., adding a data mover utility to preserve network bandwidth on head nodes and integrated common datasets for shared usage
- Customized fair share scheduling (which was provided as a contribution to the upstream SLURM distribution) to allow long-running deep learning jobs to run while supporting preemption

Current Issues

The Vector Institute has spent $6M of seed funding to create this DRI. A key source of this investment has been funding from the Province of Ontario. However, this initial funding source runs out in 2022, but Vector anticipates continuing to need to expand computing resources to support additional researchers and to sustain increases in computation to keep making impactful contributions.

The computational needs of Canada’s AI innovators have grown massively in the past decade, and the PCAIS emphasis on recruiting more world-class AI researchers is compounding this growing computational need. The PCAIS research community is over 1000 researchers, most of whom have not been significant users of Canadian ARC facilities or were recently attracted to Canada, in part because of the computational aspirations of the AI Institutes. It is difficult to accurately estimate this impact, but Vector’s existing environment has about 45% of the 2576 GPUs available in Canadian ARC facilities.

The challenge for NDRIO can be summarized as the impact on computing allocations for the rest of its science workload when Vector’s startup computing becomes obsolete, as its community continues to grow, and Mila and Amii’s researcher needs are also factored in. The workload load from all three AI Institutes is conservatively estimated as more than twice the workload at Vector, suggesting that the AI demand for GPU computing is comparable to the entire GPU capability of NDRIO’s facilities today.

The process of AI research frequently involves changing network architecture and resource requirements, which may result in changes to the overall system requirements. For example, consider the rapid shift in AI research to the use of *transformers* and increasing model scale. The initial Attention is All You Need paper was published in 2017 and quickly changed the field - with 15000 citations and most AI language models now using the technique. It is now widely used in other subdomains of AI including computer vision and time series models, and in 2020 there have been multiple innovations to allow much larger amounts of context (e.g., Longformer). It has also been scaled to a series of increasingly large models, most recently GPT-3 which has 175 billion parameters. GPT-3 is estimated to require hundreds of years of GPU capacity to train just once (i.e., about ⅓ of all GPUs owned by Compute Canada and about ½ of all GPUs owned by the Vector Institute). Of note is also Deep Speed, which is an optimization framework designed to efficiently train trillion parameter models including a large transformer language model.

This example of rapid changes in the AI field also highlights the importance of research institutes having enough resources to test and improve innovations that deliver better results and to be able to work on relevant systems problems like distributed training that are representative of industry applications. However, it also shows the rapid increase of resources that hyper-scalars like Google, OpenAI and Microsoft are applying to achieve these results. Fortunately, it is possible to produce valuable research without entering into this arms race. The rate of increase of computation for state of the art model performance can be seen in this diagram from Measuring the Algorithmic Efficiency of Neural Networks. Rapid improvements in hardware and algorithms have resulted in a 200x improvement over six years whereas increased investments in this arms race have resulted in a 37,500x increase.

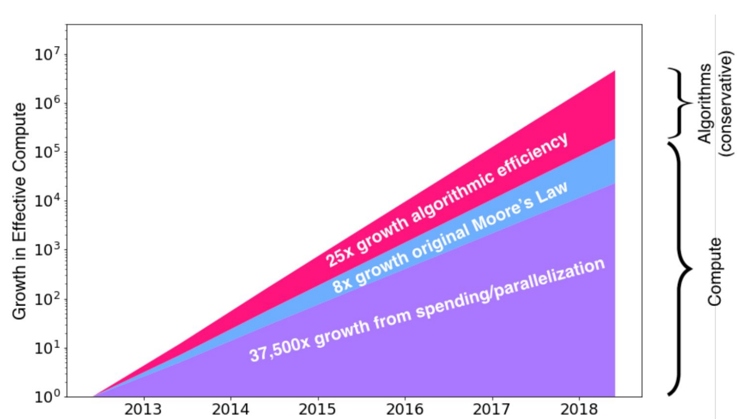

Upgrading a traditional data center is typically done in cycles that are planned and budgeted for long in advance. An AI data center must be more flexible in upgrading to the latest technology. For example, NVIDIA GPUs are the most commonly used in AI research. Between 2016 and 2020 NVIDIA has released three new architectures (Pascal, Volta, and Ampere), resulting in a 60x increase in peak operations per second on a single GPU (P100 performance vs A100 performance).

This makes it important to make frequent purchases of hardware that offers the best aggregate capacity for a budget, rather than competing with hyper-scalars to set supercomputing results (like MLCommons benchmarks). These same investments also support mainstream applications of AI in industries outside technology, where there are also significant constraints on costs and achieving AI results within a budget is important.

It's critical to be able to add systems software and configure systems for the rapid changes required to support fast changing research needs, e.g., updating cluster scheduling policies to allow using the right resources for each purpose, adding GPU profiling to allow optimization, adding specialized libraries, and taking advantage of virtual GPUs for efficiency. These make a big difference to research productivity but require specialized knowledge and support that isn't efficient to provide and may be disruptive for general purpose HPC environments.

Similarly, it's important that AI research environments allow quick experimentation without requiring a lot of up front engineering while still ensuring that users don't degrade the overall environment. A recent analysis highlights how this flow varies between AI systems and traditional HPC systems, especially in data access patterns.

Because of the rapid expansion of the field of AI, researchers are frequently adapting and learning new techniques and responding to recent innovations. This results in a frequent process of exploring new ideas that requires large-scale computation as well as agile project formation. This process is highly dynamic and it requires responsive support and updates to computation environments. The explosion in AI research is illustrated below:

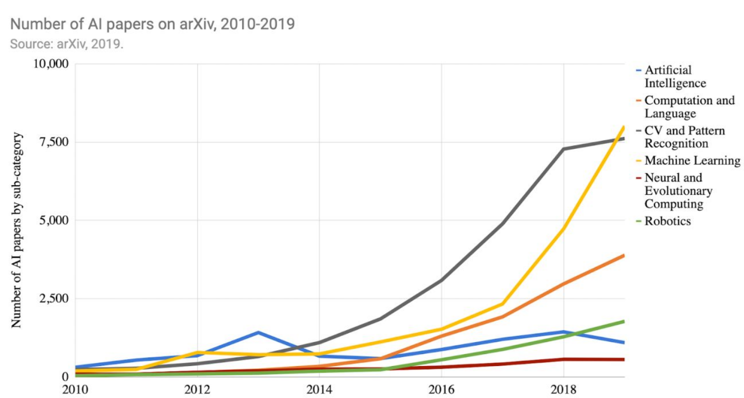

Source: AI Index 2019

Vector's current DRI environment does not meet the needs of some specialized but important research topics. E.g., it's not optimized for systems research into distributed training and neural architecture search. Similarly, Vector doesn’t have an environment that’s certified to hold sensitive healthcare or biomedical data and its researchers rely on third party partners for DRI in these areas.

Future DRI State

We strongly recommend that NDRIO increases ARC infrastructure spending to accommodate the expected workload from Canadian AI researchers without drastically reducing the actual compute available to both AI researchers and the rest of the NDRIO researcher community. We believe this infrastructure should continue and scale the direction that Vector has established with its pilot environment:

- Adding specialized capacity for AI research
- Making frequent incremental purchases of the latest hardware
- Having close collaboration with researchers and specialized operations and software development teams to update the research environment for evolving needs in an agile manner
- Ensuring there is an effective resource economy and appropriate allocation of resources
- Ensuring there is an effective way for researchers to explore new ideas in an agile manner
- Adopting cybersecurity practices that allow research on sensitive data through entitlements and encryption rather than restricted or isolated environments

We also believe that there would be value in NDRIO investing in and modernizing DRI to better enable emerging healthcare research needs for AI. We support the submissions of "A Standards-Based Digital Infrastructure for Secure Sharing of Human Biomedical Research Data" (Stein et al) and "Digital research infrastructure to support federated computing on large scale biomedical datasets" (Brudno et al).

How to Bridge the Gap

We believe that these gaps can best be bridged by providing resources to fund additional DRI that’s prioritized for AI research while allowing use for other purposes following these principles:

- Extend, don’t replace, existing NDRIO funding plans to support HPC in Canada
- Allocate funds into regional computing centers to allow optimization in consulting with their local PCAIS institute
- Establish a joint governance model for operations and systems capabilities between the regional computing centers and the AI institutes
- Budget for regular incremental refreshes to continue to improve capabilities twice per year
- Allow for shared RFP’s for procurement but individualized priority criteria for regions

Conclusion

Canada has realized tremendous benefits by investing in a Pan-Canadian AI Strategy. With the commitment shown by chartering NDRIO, there is a great opportunity to invest in specialized resources for AI researchers. This will align Digital Research Infrastructure with a proven approach to building HPC hardware and software that enables AI researchers to excel and continue to advance Canadian leadership in the field. This promises to result in significant economic, healthcare, and societal benefits.