AI for Sciences

Into the world of Exascale computing & Beyond

… Digital Biology as Example

Bharat Kaul

Intel Labs

May 2022

Intel Labs - Parallel Computing Lab

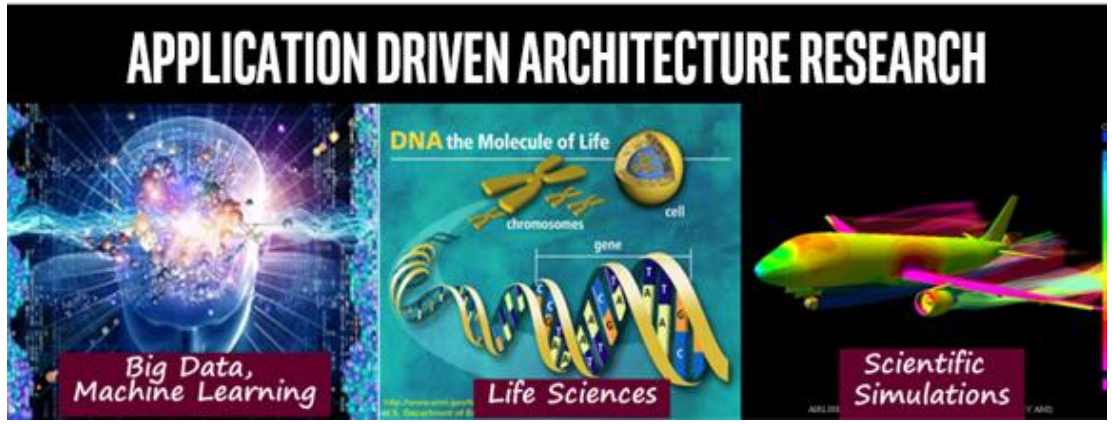

Bring computing to the breakthrough problems at the intersection of domain & computational barrier:

1. Massive computational barrier: novel foundational algorithm, scale
2. Clear goodness metric or benchmark
3. Data, accurate simulator availability

Research to Realization

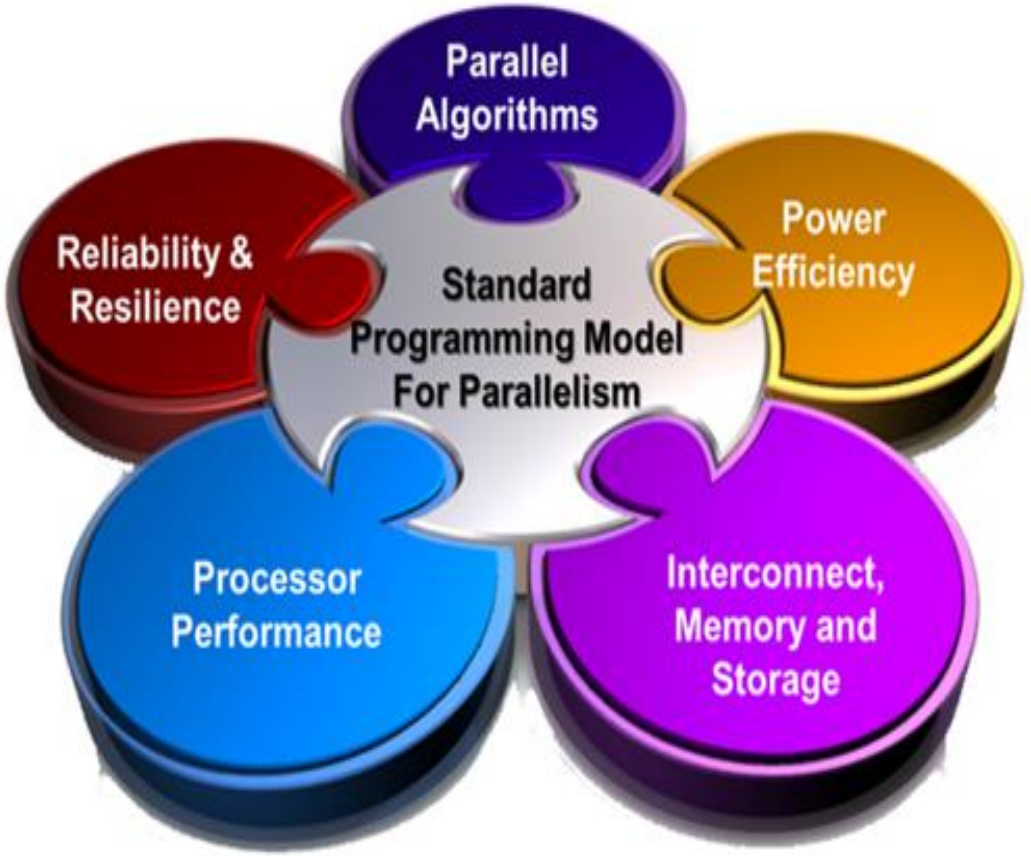

In collaboration with leaders from domain, AI, HPC spanning academia and industry

Balance of decision making is shifting

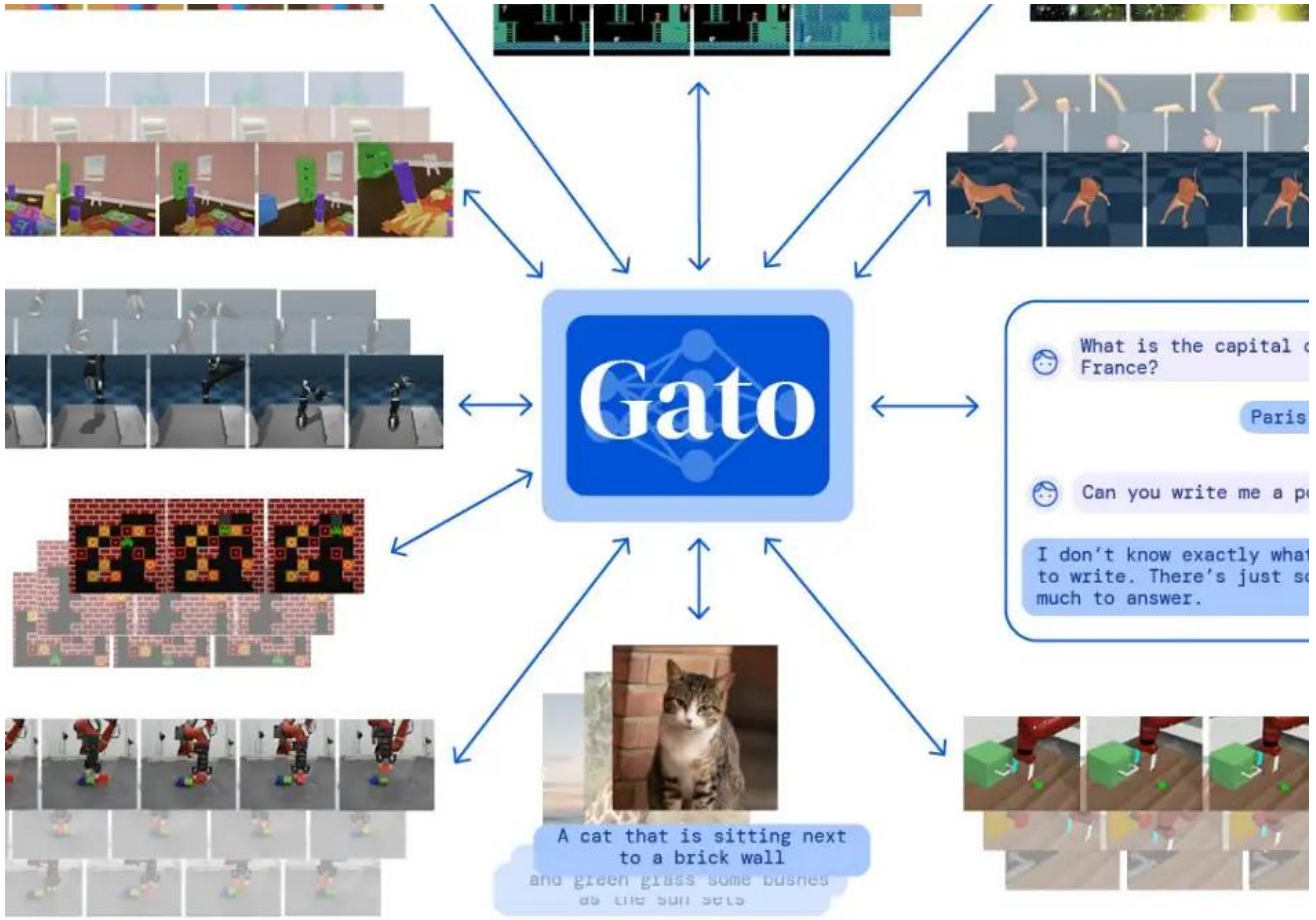

https://www.deepmind.com/publications/a-generalist-agent

When our own breath gonne faint and fade
Up the long stairway of the dead
We shall remember well
The blindness of the light, the whiteness
Of that white land. The footsteps, and
The grasses in that light, and in the shade
The menacing life.

We played, too, once, in the warmest rooms.
To one content, there is one content
For all mankind. The forest glades

https://openai.com/blog/gpt-3-apps/

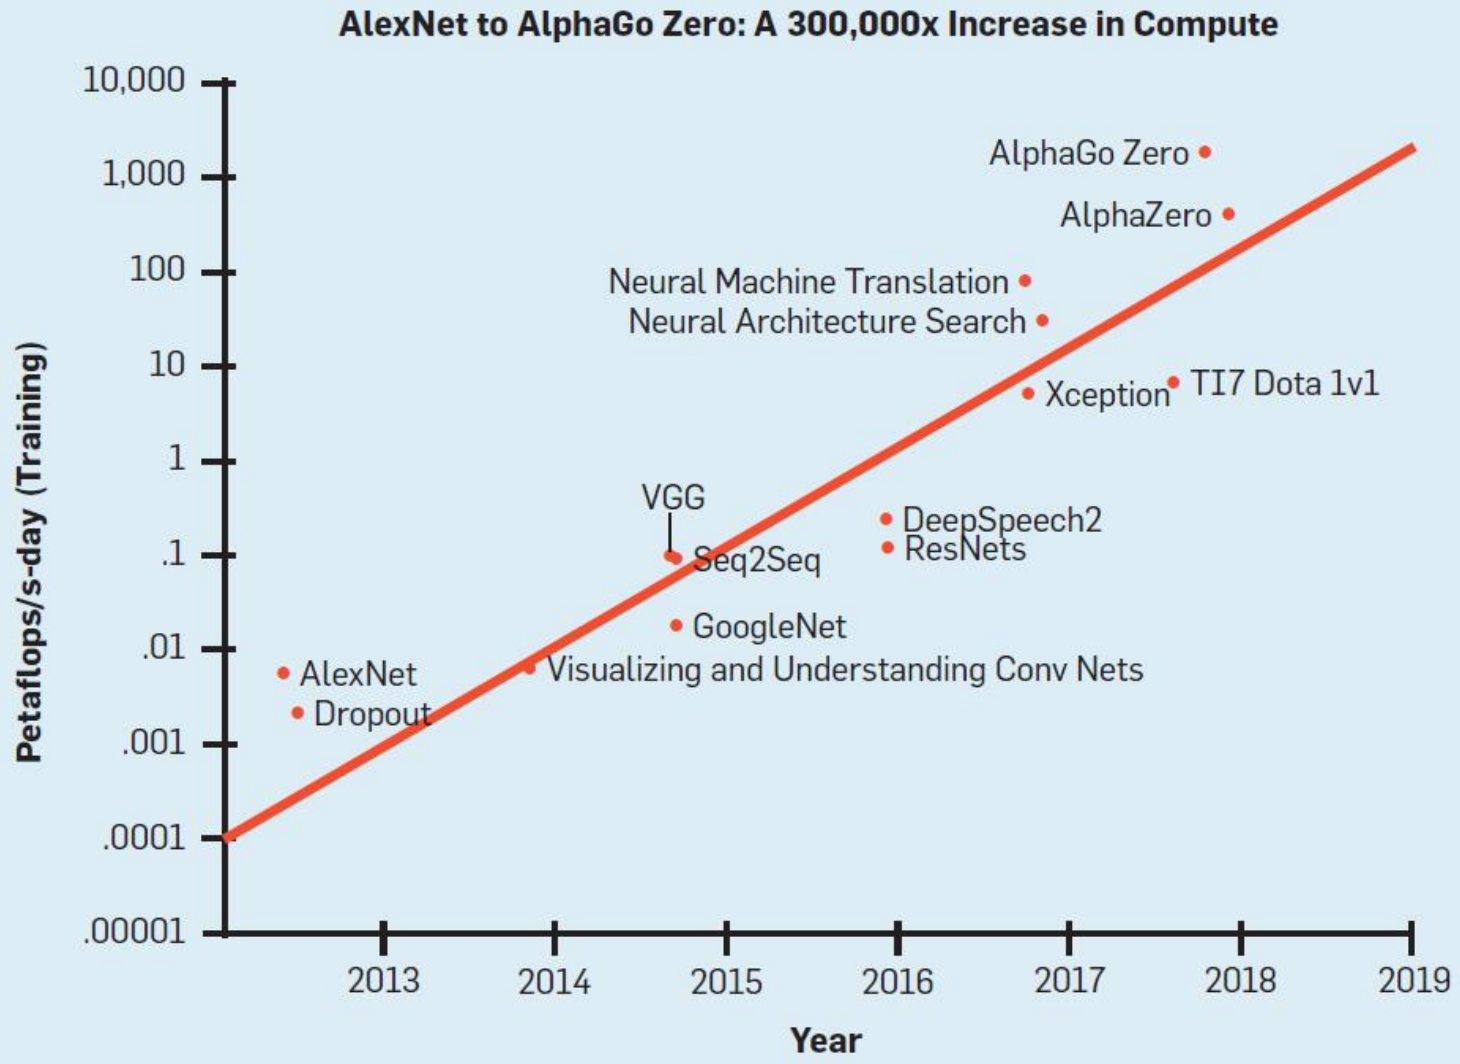

Massive computational appetite … and growing

https://openai.com/blog/ai-and-compute/

Upcoming Shift in AI enables Simulation & modelling complex Systems

AI Capability

Pattern Detection

Trillion Parameter Model

Data Mapped to Euclidean

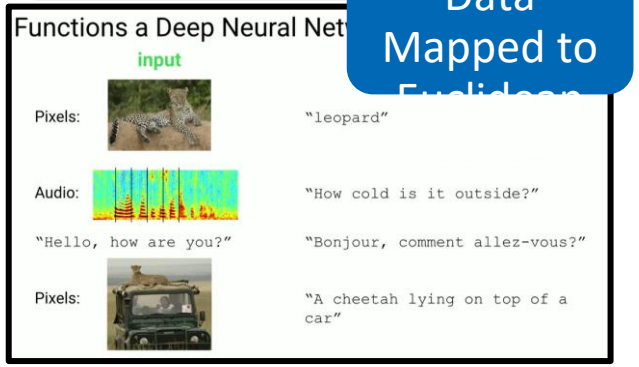

> 100 Trillion Parameter Model in 2022 & growing

Traditional Graph Analytics + AI

AI: Physics Inspired NNs, Meta Learning, Causality, Generative AI, Bayesian + DL, Neural ODE, PDE etc

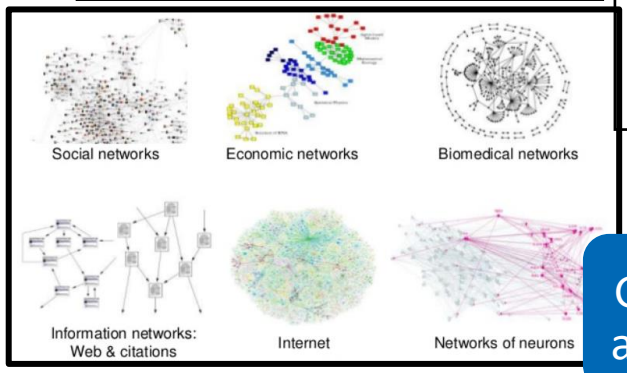

GRAPHS as Input Data

Today

Near Future

> 100 Trillion Parameter Model in 2022 & growing

Traditional Graph Analytics + AI

AI: Physics Inspired NNs, Meta Learning, Causality, Generative AI, Bayesian + DL, Neural ODE, PDE etc

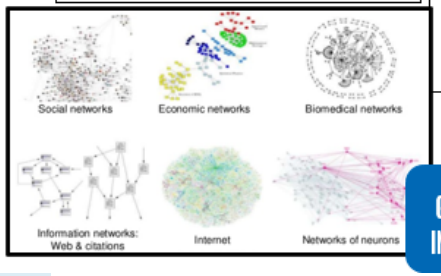

GRAPHS AS INPUT DATA

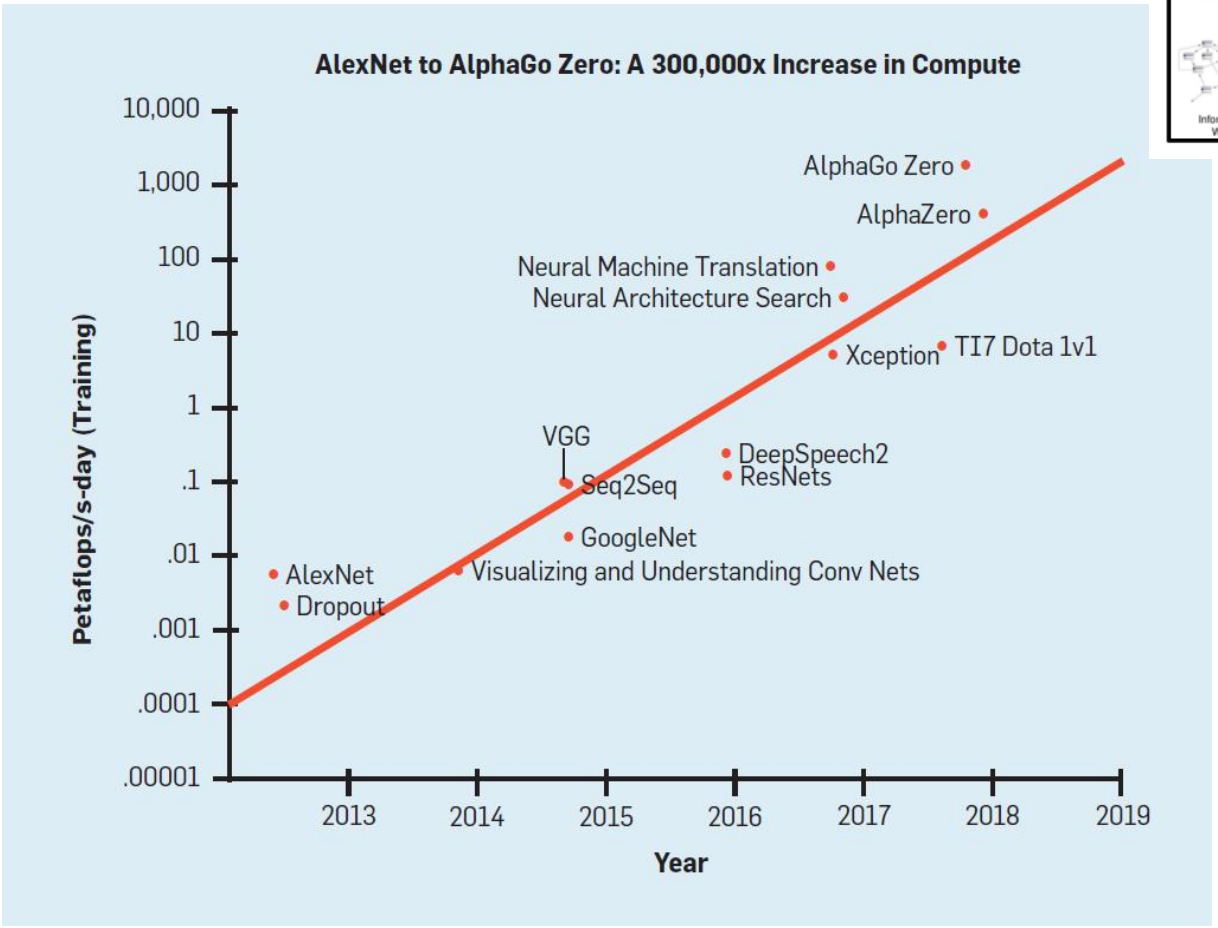

Frontiers in science can build on this foundation

and push this further

https://openai.com/blog/ai-and-compute/

The new frontier

Inside - Out ⇌ *Outside - In*

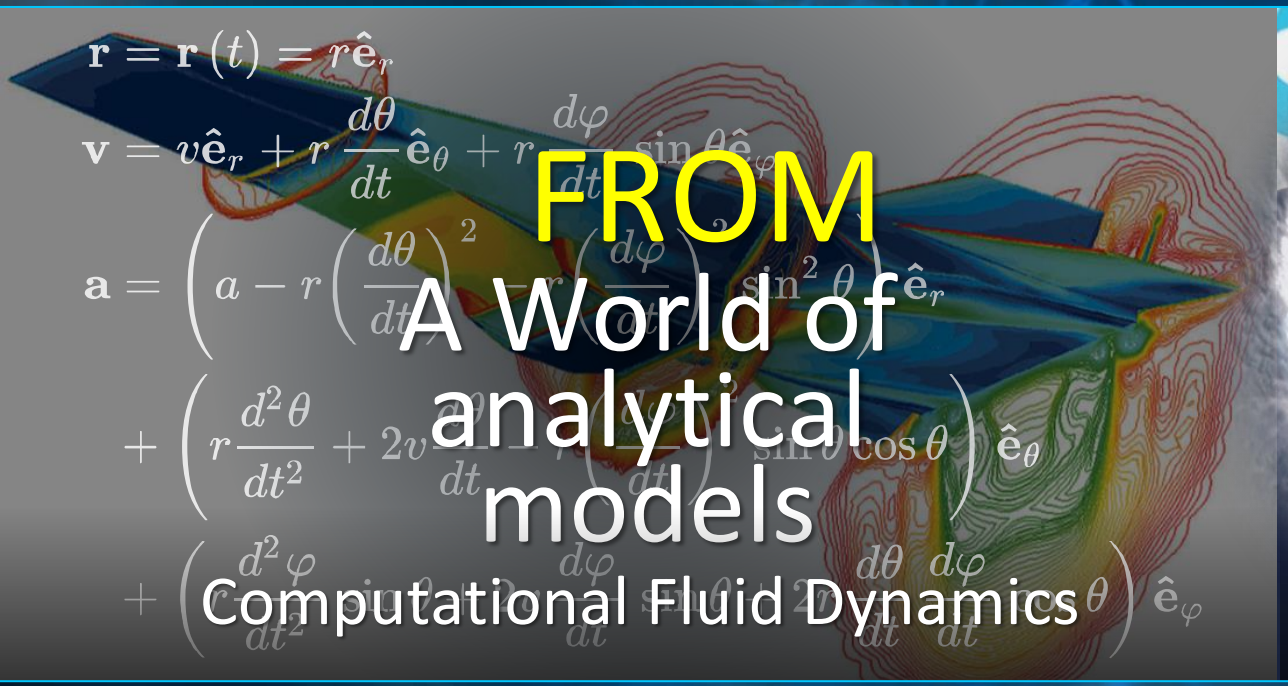

FROM

A World of analytical models

Computational Fluid Dynamics

TO

A World of Data driven Models

Event Detection from Social Media

Start with Mathematical Model

Model → Simulate → Predict

Start with Data

Initial State → Increment → Steer

Simulation To Prediction: AI as a tool to see further

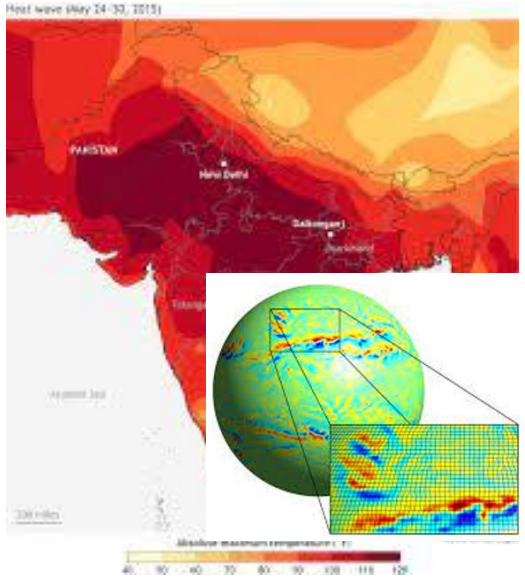

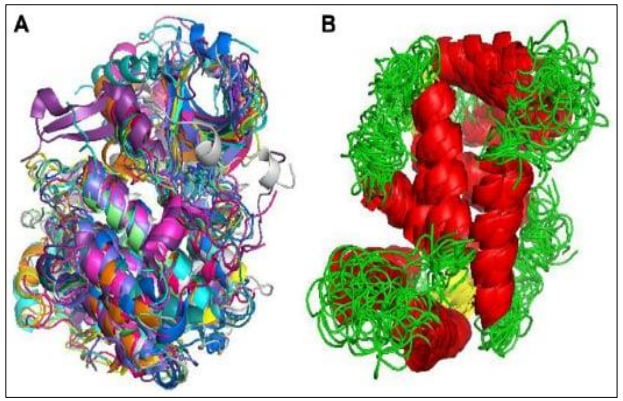

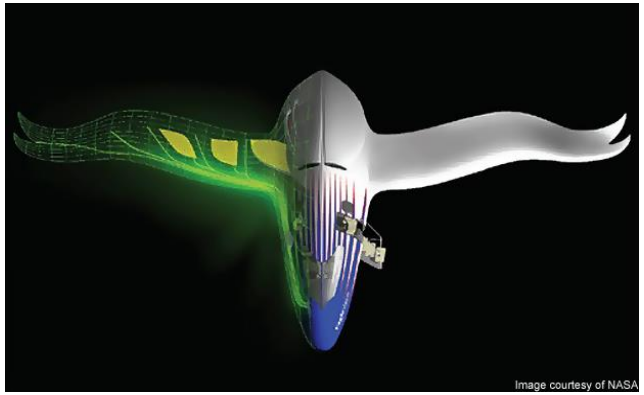

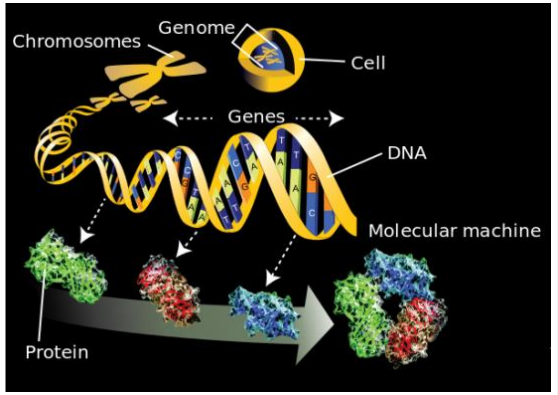

- Complex phenomena: Multi-scale resolution, interaction dynamics
- Brute force compute solving true physics for such problem much farther out
  - 10x more fine grain resolution ➔ 1000x data points ➔ 10^3 computational barrier
- Faster than Moore's law: Novel AI algorithms replacing and/or augmenting brute force computation
  - Domain knowledge critical to algorithm development

Towards better understanding of interiors and

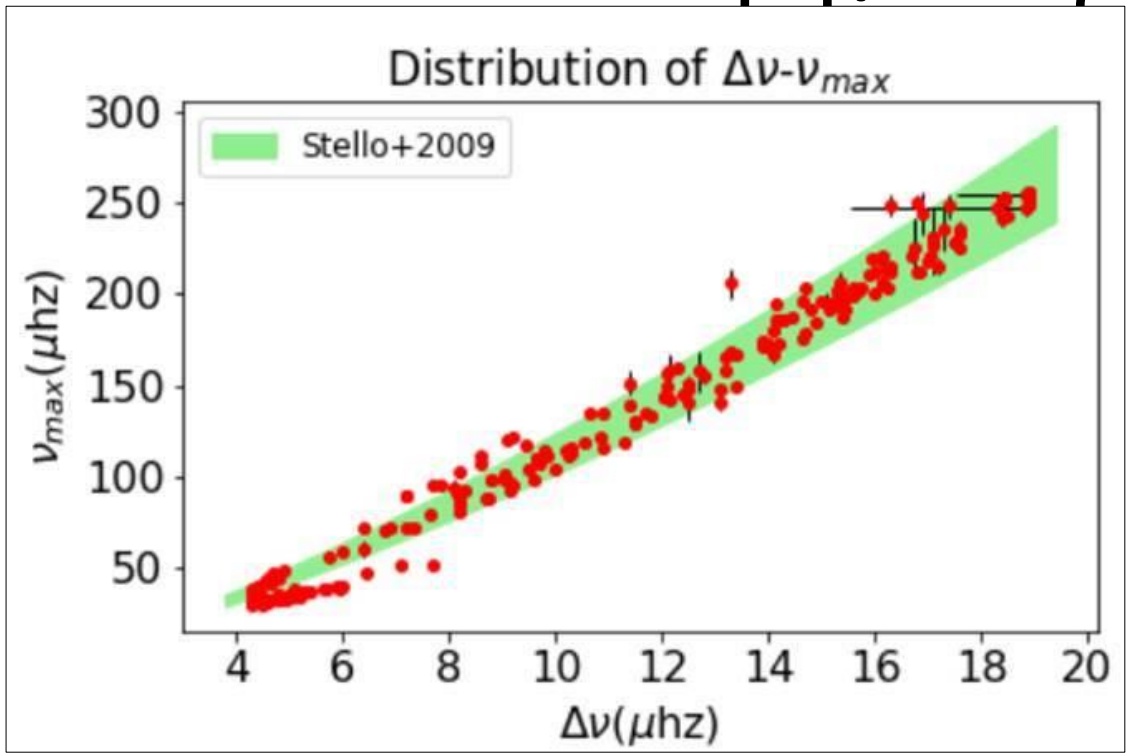

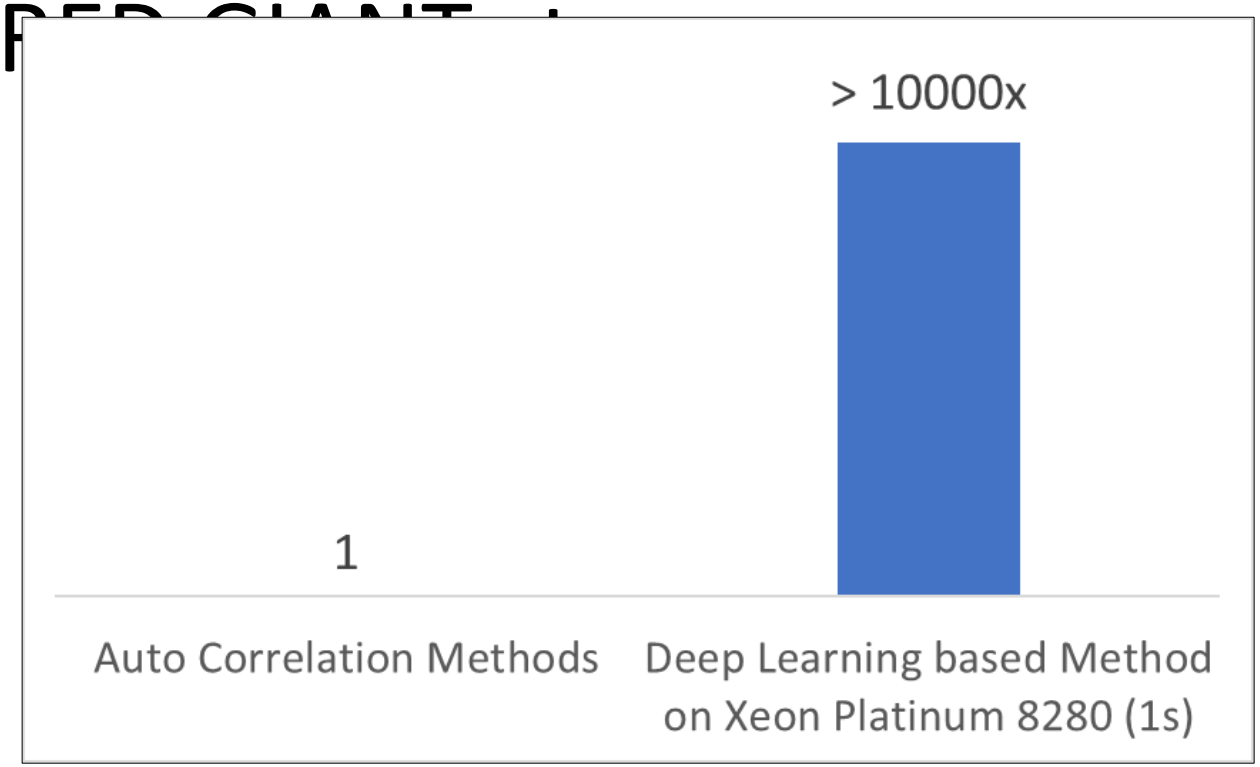

Accurately extract seismic parameters from 1000 spectra in under 10 secs

Measuring Frequency and Period Separations in Red-giant Stars Using Machine Learning
The Astrophysical Journal, Volume 928, Number 2

Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research, Center for Space Science, NYUAD Institute, New York University Abu Dhabi, Division of Solar and Plasma Astrophysics, NAOJ, Mitaka, Tokyo, Japan, **Parallel Computing Lab, Intel Labs, Bangalore, India**

https://arxiv.org/abs/2202.07599v1

Understand mechanisms, Design Interventions

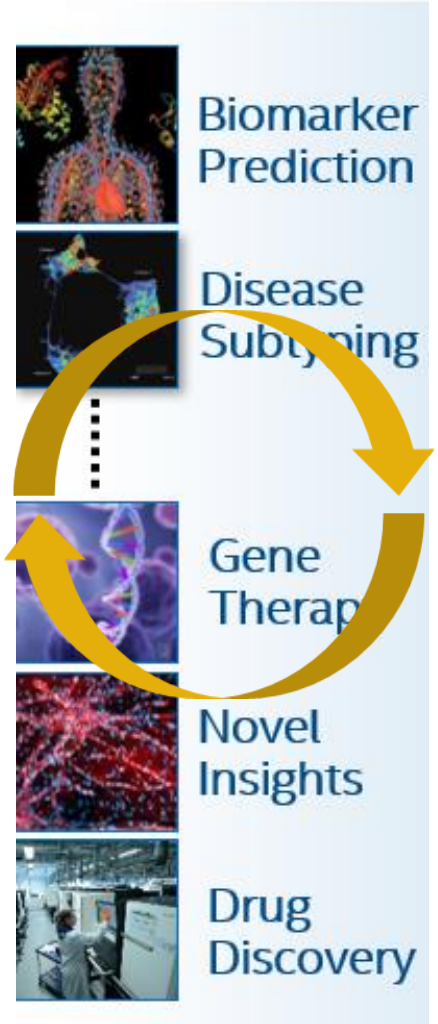

1. Learn Structure and Causality
   - Use of AI in Simulations
   - >>> Trillions of score computations
2. Search a very large hyper dimensional space >> internet O(10^9)
   - Molecule search > 10^60,
   - Protein design > 10^130

* Big Data: Astronomical or Genomical? - PLOS

Novel AI Methods for Breakthrough applications: LambdaZero

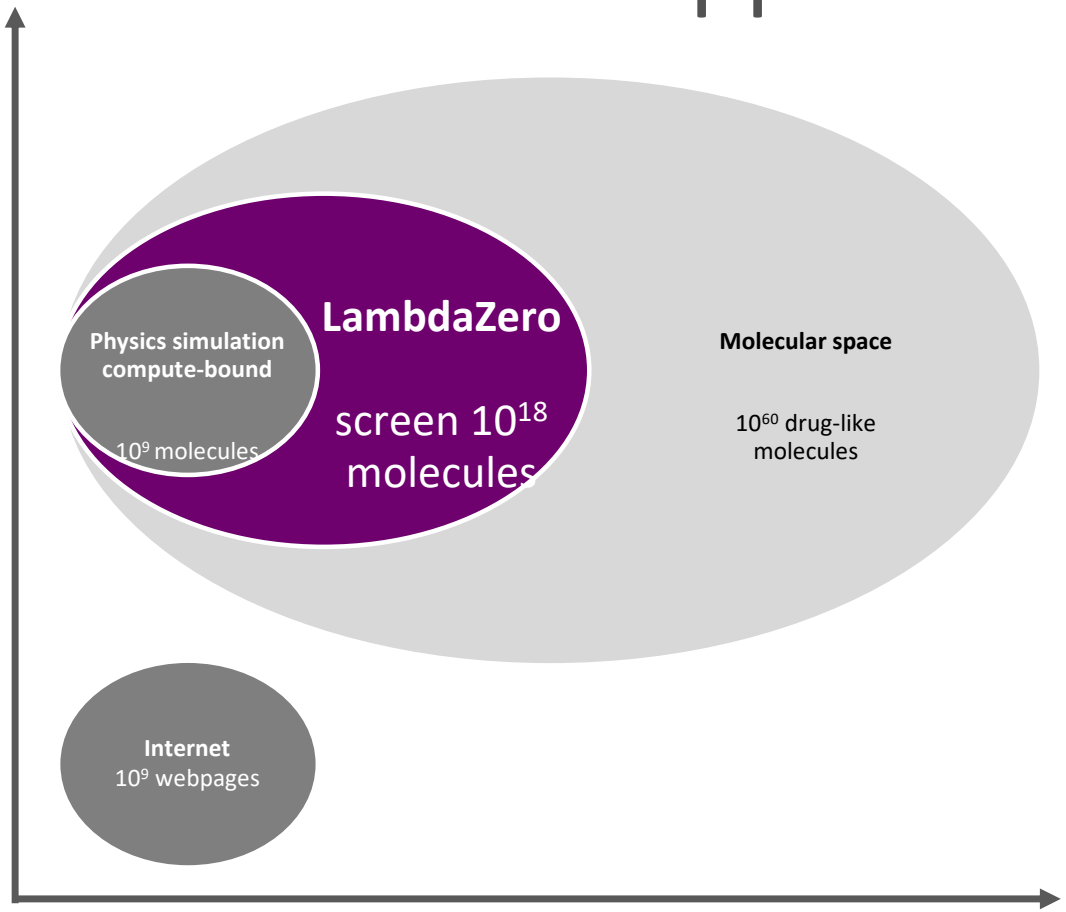

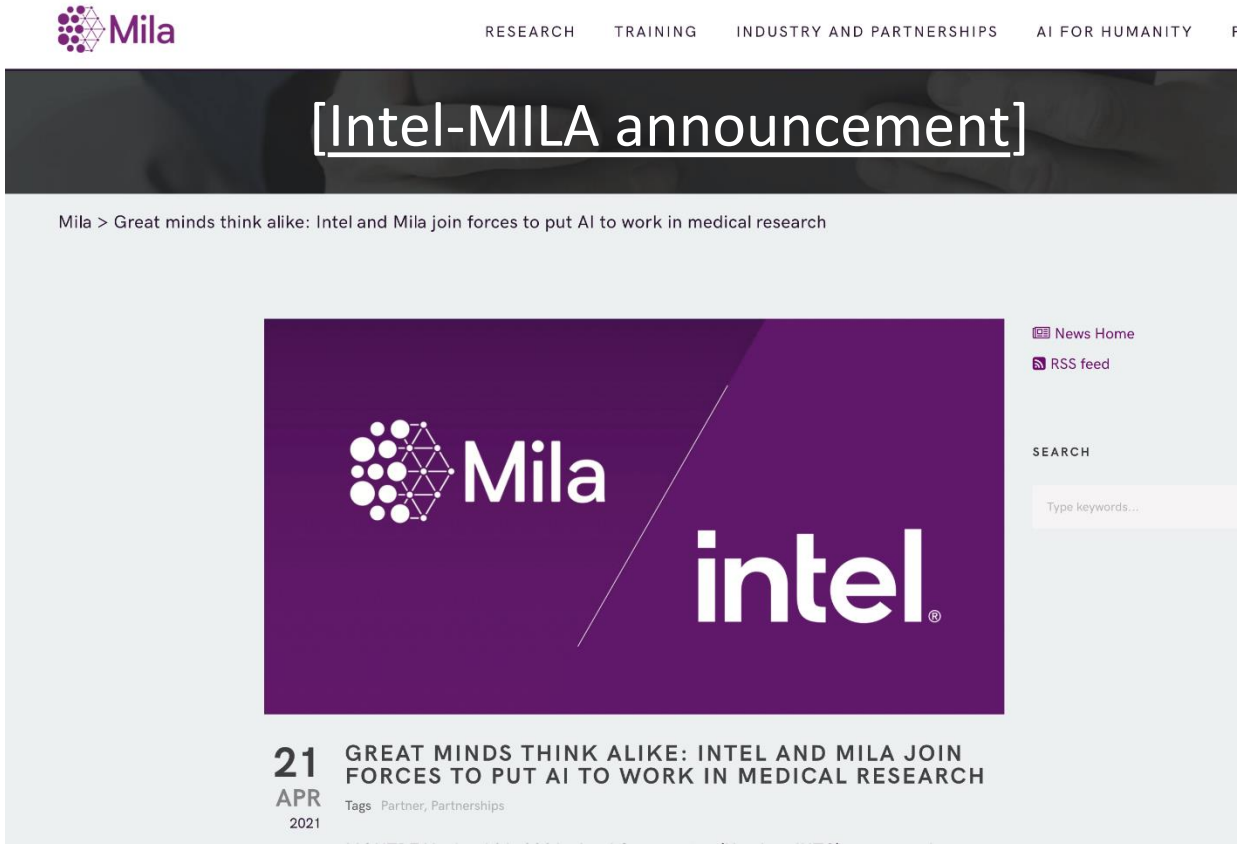

Collaboration Objective: Make it computationally feasible to expand the search space of better molecules (with more negative binding energy) to 10^18 using machine learning and hardware-software codesign *

Graph Neural Network

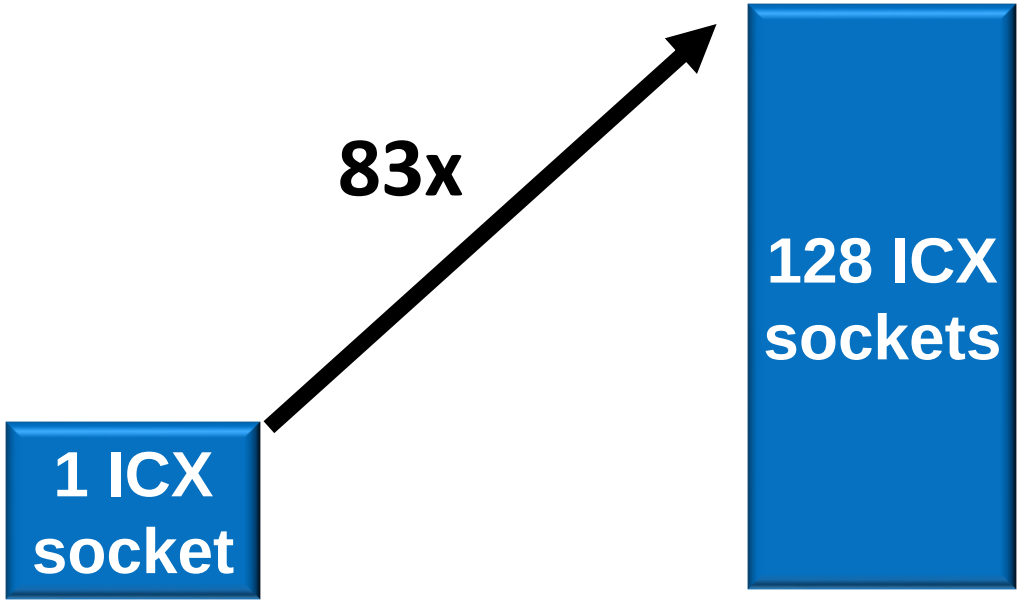

OGB-Papers dataset, 100M Graph Nodes

Full batch training

https://arxiv.org/abs/2104.06700v2

Supercomputing'21 - distGNN: Scalable Distributed Training for Large-Scale Graph Neural Networks

Full batch Training ~2-3.7x faster on 1s-CLX (1s) for GraphSAGE on OGB-Products & Reddit ~83x for distributed training on 128 sockets on OGB-Papers

SCANPY: Analysis of Single Cell RNA-SEQ Data for 1.3 Million mouse brain Cells

Widely used toolkit for analysis of single cell RNA-Seq data with 1300+ citations, ~600K downloads

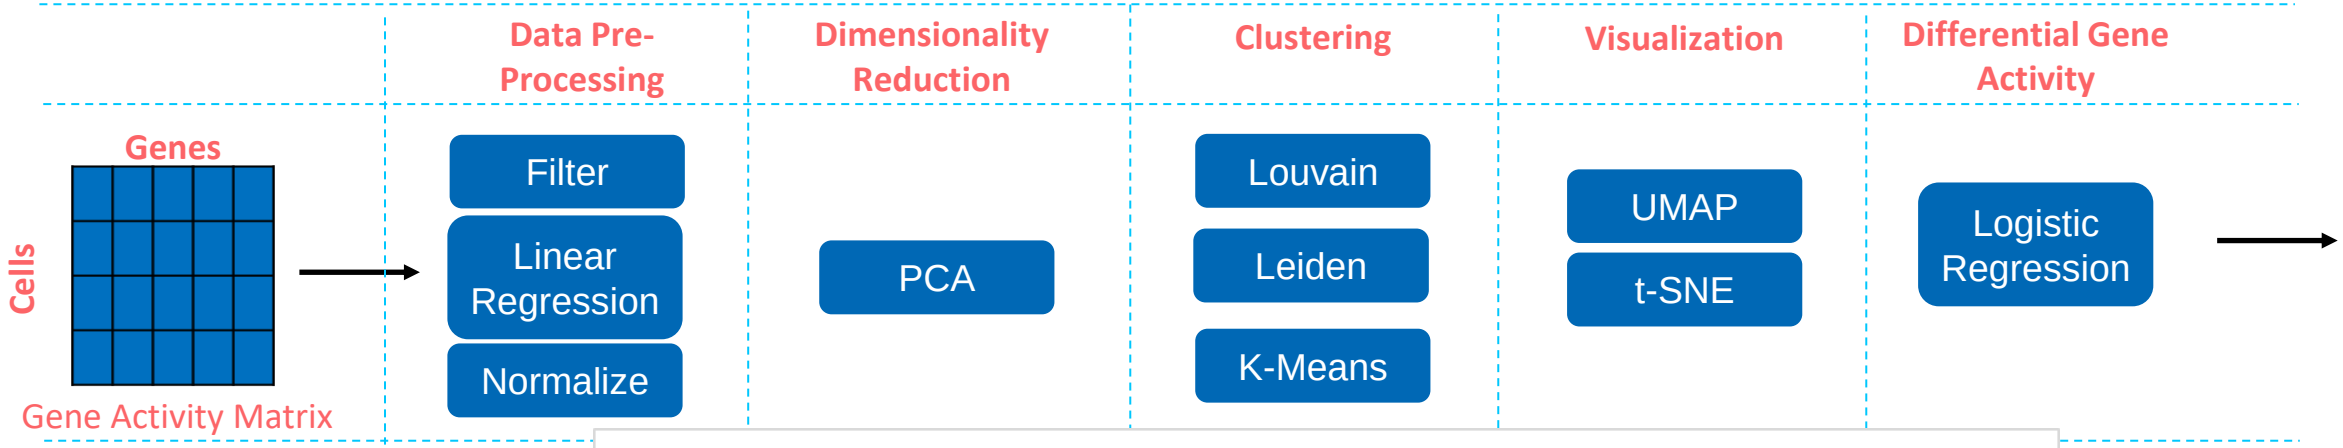

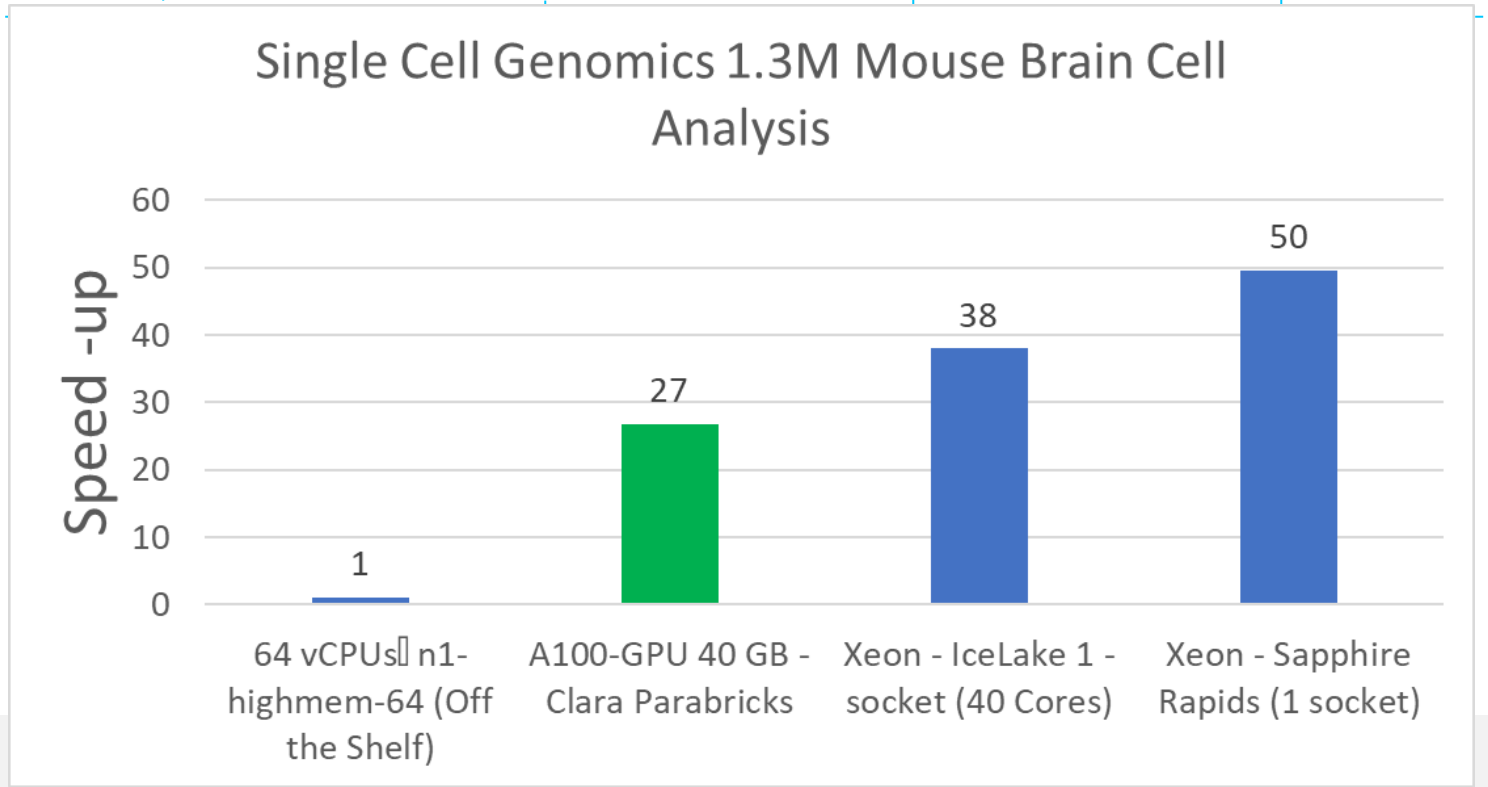

Source: Github repository as of Dec 16, 2020 - Example 2: Single-cell RNA-seq of 1.3 Million Mouse Brain Cells comparing CPU (n1-highmem-64 64 vCPUs) vs GPU (a2-highgpu-1g. https://github.com/clara-parabricks/rapids-single-cell-examples. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
1S Ice Lake: See Backup for workloads and configurations. Results may vary.

Our aim:

Drive Innovation across full stack

Research to Realization

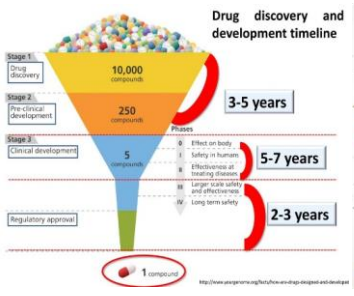

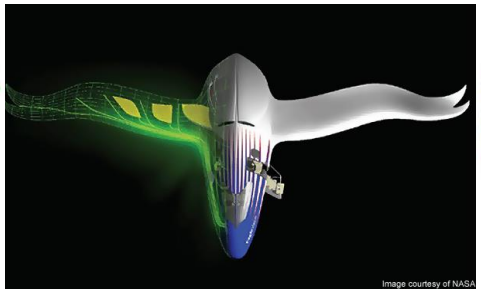

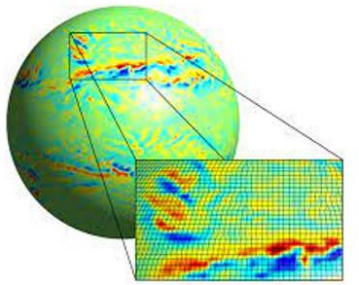

Algorithmic, Computational & Data Management Requirements

e.g., Learning on Large Graphs, structure learning, Search, Causality, Anomaly Detection, Zero shot learning…

Process 100 Terabytes of multi-modal data

Secure, Privacy preserving, Federated

>1000x growth IN COMPUTE TO MATCH DEMAND

100's of TB/s MEMORY BW AT 100'S OF GB CAPACITY

Big Data: Astronomical or Genomical https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002195

Thank you