Advancing Science with AI from the Edge to the Datacenter

Middleware and Grid Interagency Coordination (MAGIC) Meeting

August 3 2022

Geoffrey Fox, Biocomplexity Institute and Computer Science department -- University of Virginia

DiscoDiffusion `A beautiful painting of ten small robots under the sun and moon in style of Monet, green color scheme`

Abstract

- We distinguish systems used to support AI and AI supporting data centers (supercomputers, clouds), fog, and edge.
- Today most AI advances are in deep learning which inevitably uses HPC and stretches from edge to data center.
- AI will permeate all parts of the computing continuum software stack, systems and application software
- AI surrogates replace compute intensive libraries by learnt surrogates.
- This support will be at the system level where it will both improve performance and remove the distinction between supercomputers and clouds by carving out application specific compute resources from a heterogeneous networked hyper-hypersystems.
- Meta or Foundation Models could generalize across science domains

- Thanks: This work is partially based on collaborations with Tony Hey, Jack Dongarra, , Shantenu Jha, Pete Beckman, Madhav Marathe, Judy Qiu, Gregor von Laszewski.

Backgrounds from Text to Image Deep Learning (DALL-E, Imagen) Disco Diffusion running in Google Colab

A beautiful painting of ten small robots under the sun and moon in style of Monet, green color scheme

Google Imagen

"A cute corgi lives in a house made out of sushi"

AI for Engineering and Science I

- **Experimental Physics** (accelerators of anything, microscopes, telescopes. Detect gravitational waves): Control and monitor apparatus, remove noise, design optimal apparatus and then analyse resultant data
- **Theoretical Physics**: Speed up simulations -- AI surrogates learns results; new AI-based methods of performing integrals (QCD)
- **Astronomy:** Identify galaxies, stars, black holes, asteroids from observations across multiple wavelengths
- **Material Science:** Predict properties of material from its specification and inversely predict specification from desired properties. Data from simulations and measurements such as Xray and light scattering
- **Drug Discovery:** Similar but properties are not “strength” or “weight” but rather effectiveness against a disease.
    - Also predict protein structure from specification (AlphaFold2)
- **Life Sciences:** AI to analyse omics, patient records, pathology images, epidemiology (Covid), AI+”mechanistic models” to build virtual humans

AI for Engineering and Science II

- **Earth, Agricultural and Environmental Science:** analyze images from satellites and time series from distributed sensors
- **Climate and Weather Science:** AI controls ensembles of simulations and integrates observational and simulation results; AI learns microstructure (clouds)
  - AI for “net zero” will make world green
- **Critical Infrastructure:** Electrical power grids; AI will control the new distributed mix of renewable and traditional power sources
- **Fusion:** Tokomaks use an electron plasma to create “clean nuclear fuel” that needs real time AI control from observation and simulation (surrogates)
- **AI for Computing Infrastructure**
  - Schedule and control computers and networks
  - Learn results of searches, data analyses and simulations (surrogates)
  - Visualization
- **Lots of Social media and e-commerce commercial applications**
- **AI Above Everything:** AI will read all the books and gobble all the data and advise/tell us what to do?

Features of AI Use in Science in Computing Continuum

- **Synergistic with industry** work e.g. **computer vision nets** starting point for remote sensing, pathology; **industry time series** starting point for science sensor analysis
- **Simulation surrogates** used everywhere from material science, climate and QCD to visualization (open questions such as uncertainty quantification)
  - Surrogates often only replace part of simulation
- ~95% of new AI development is **deep learning** with decreasing use of classical ML but substantial **data engineering** before and after the core AI
- **Data** gathered and (partially) analyzed **at the edge**
  - **Realtime** sometimes essential (web shopping, wildfires) and other times data not all streamed (some monitored online) but rather **batched** and analyzed fully at data center.
- **Edge instruments** (from sensors to LHC/SKA/LIGO/Light sources) are designed, monitored and controlled by AI while data analyzed with AI
- **High Performance Computing HPC** essential whether at instrument or at data center and whether latter is a cloud or a supercomputer
  - **Networking** from edge to data center **controlled by AI** as are both **data center and edge/fog computing systems**

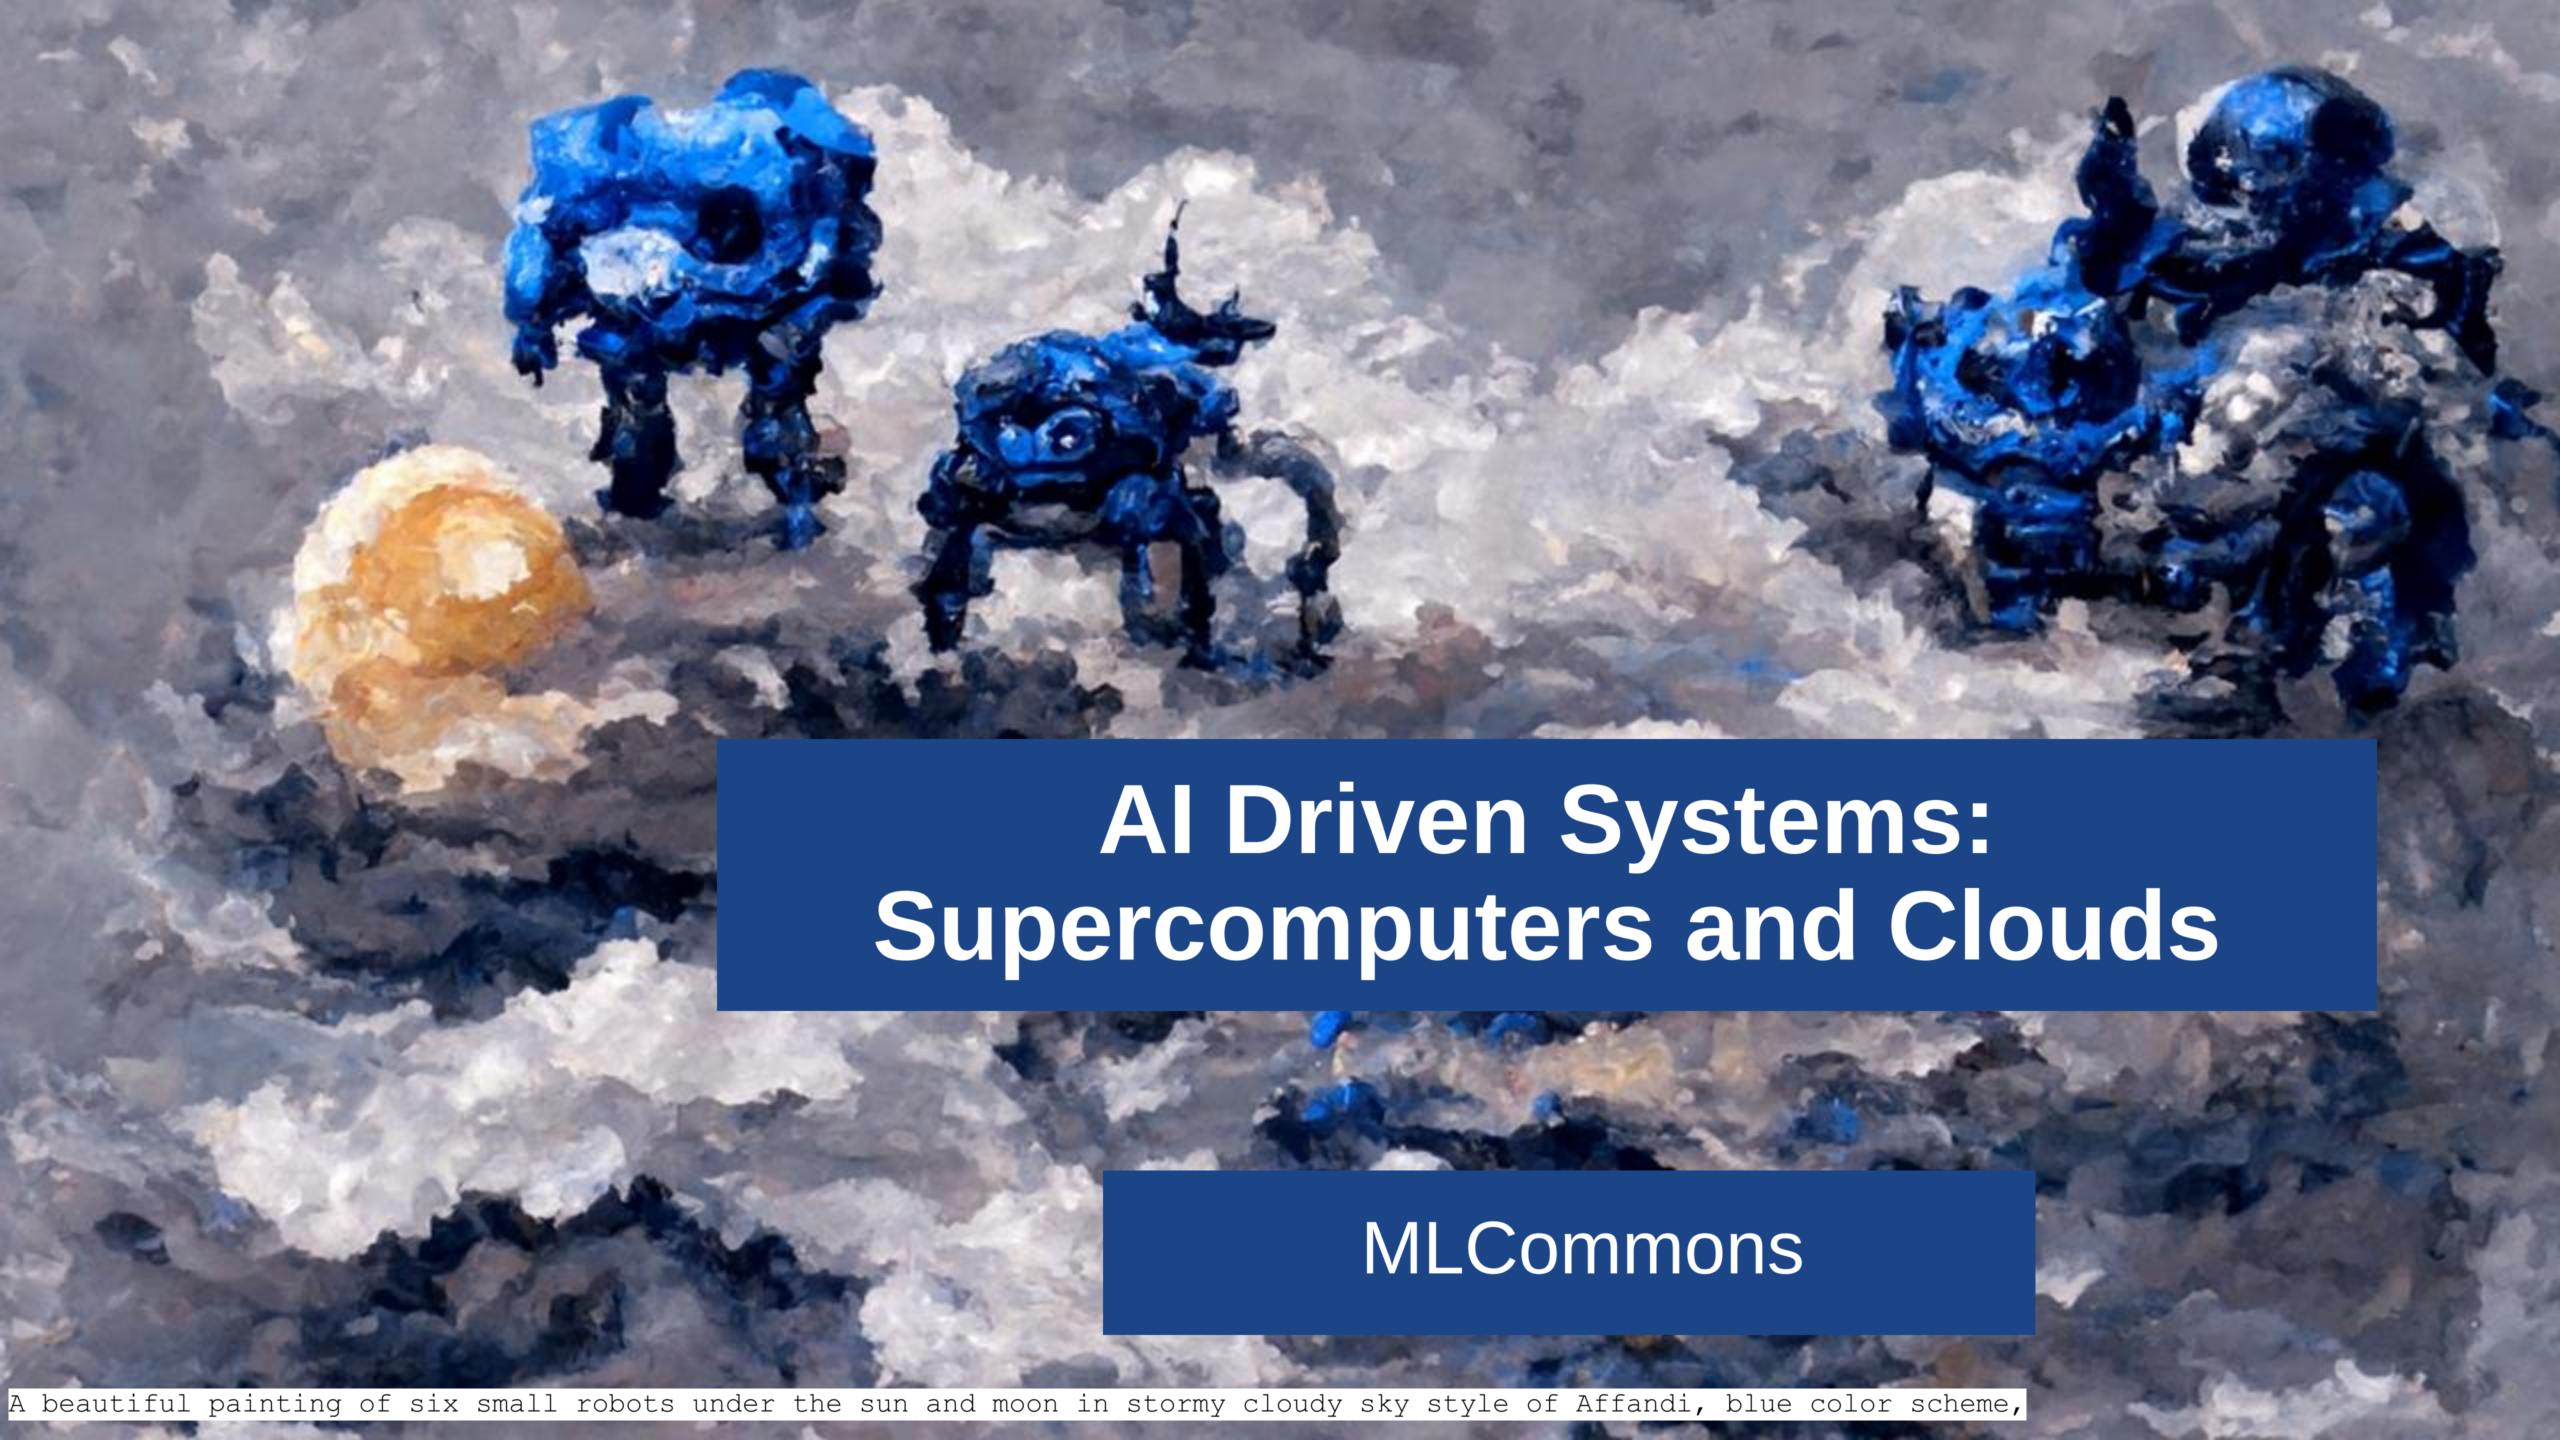

AI Driven Systems: Supercomputers and Clouds

MLCommons

A beautiful painting of six small robots under the sun and moon in stormy cloudy sky style of Affandi, blue color scheme,

MLCommons (MLPerf) Consortium Deep Learning Benchmarks

- MLCommons aims to accelerate machine learning innovation to benefit everyone. Benchmarking, Datasets, Best Practices
- Major effort of 62 companies to produce benchmarks with ongoing challenges
- Training at V2.0 (sixth release)
- Fox set up Science Working Group with co-chair Tony Hey who has a significant benchmarking group SciML.

Some Relevant Working Groups

- Training
- Inference (Batch and Streaming)
- TinyML (embedded)
- Power
- Datasets
- HPC (Supercomputer Implementations)
- Research (Academic-Industry Links)
- Science (AI for Science)
- Storage (new Research Group)
- Best Practice (Software)
- Logging/Infrastructure (metadata)

Total Effort ~50 FTE

MLCommons (MLPerf) Consortium Activity Areas

Benchmarking

Benchmarks provide consistent measurements of accuracy, speed, and efficiency. Consistent measurements enable engineers to design reliable products and services, and enable researchers to compare innovations and choose the best ideas to drive the solutions of tomorrow.

Datasets

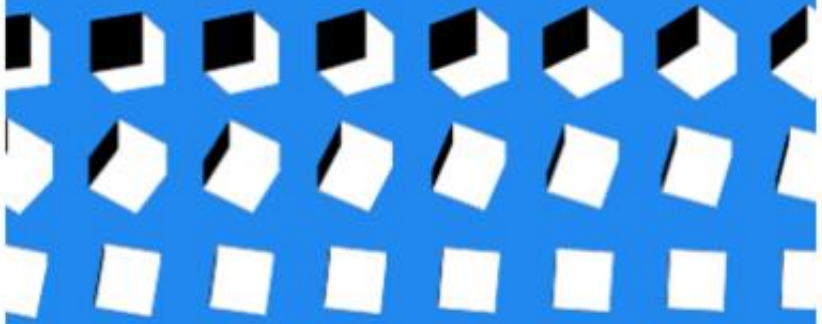

Datasets are the raw materials for all of machine learning. Models are only as good as the data they are trained on. Academics and entrepreneurs in particular depend on public datasets to create new technologies and new companies.

Best Practices

Best Practices empower researchers and engineers to more easily exchange models, reproduce experiments, and build applications that leverages machine learning. Improving best practices accelerates progress in, and grows the market for, machine learning.

MLCommons Research Science Benchmarks

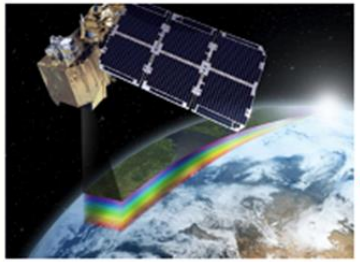

slstr_cloud
Cloud Masking

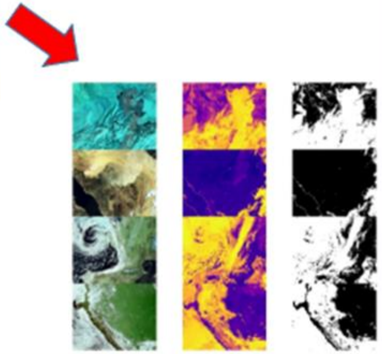

Identifying pixels that are cloud in satellite images

Tevelop
Timeseries Evolution Operator

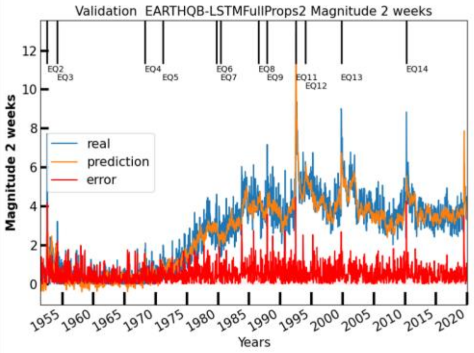

Focusses on extracting how a given timeseries evolves

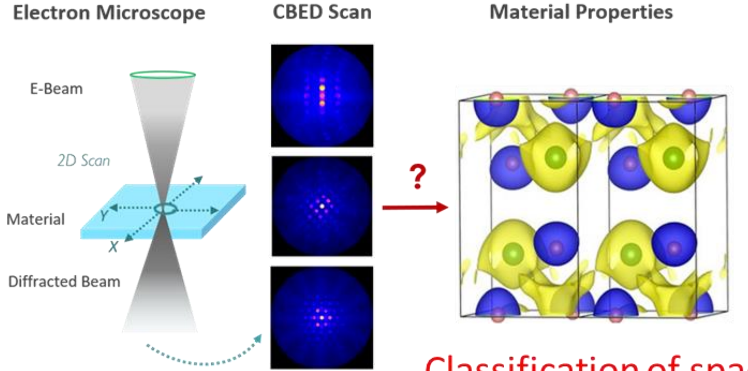

stemdl: *ImageNet for Solid Materials*

Classification of space groups, 40GB dataset with 230 spacegroups

candle-uno
Deep learning-based precision medicine for cancer – drug response measurement

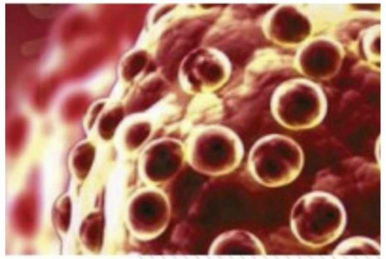

Group info

- https://mlcommons.org/en/groups/research-science/
- Chairs: Geoffrey Fox (Virginia), Tony Hey (RAL, STFC),
  Jeyan Thiyagalingam (RAL, STFC)
- Published and Announced ISC June 2 2022

Science-based metrics

- accuracy, signal-to-noise ratio, etc.

MLPerf HPC v1.0 Benchmarks

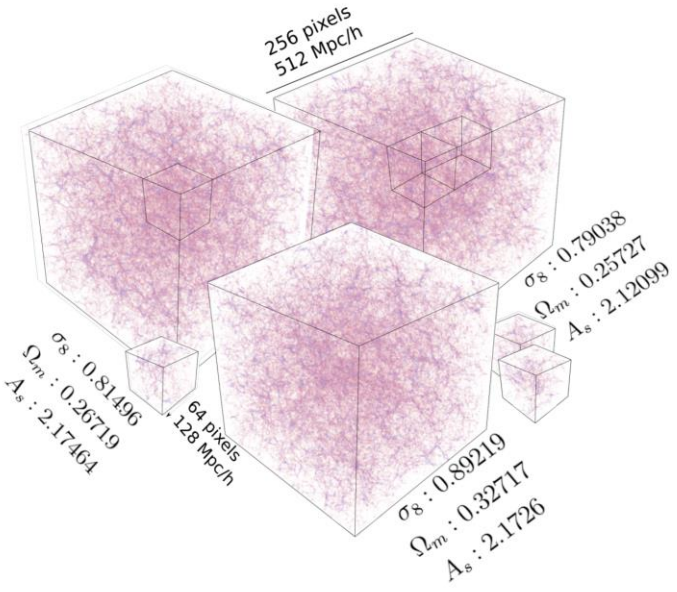

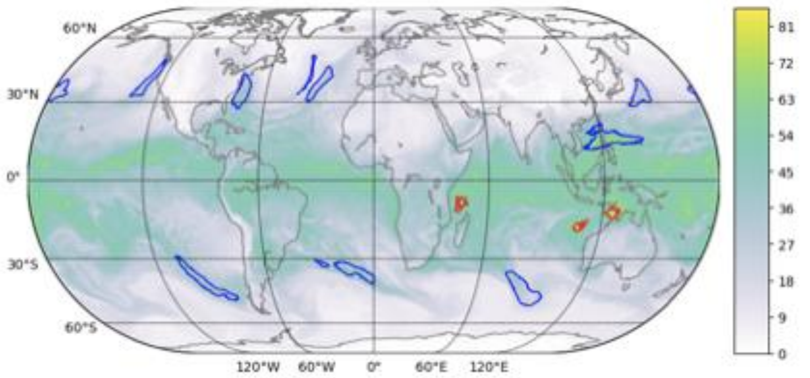

Group info

- https://mlcommons.org/en/groups/training-hpc/
- Chairs: Murali Emani (ANL), Steven Farrell (LBNL)
- Vice chair: Aristeidis Tsaris (ORNL)

CosmoFlow

- 3D CNN regression on cosmology simulations
- Originally published at SC18 total size 10.2 TB

DeepCAM

- 2D CNN segmentation, identifying weather phenomena in climate simulations
- 2018 Gordon Bell prize paper total size 8.8 TB

OpenCatalyst

- Graph NN predicting energy and forces in atomic catalyst systems (material surface + molecule)
- Dataset: Open Catalyst 2020 (OC20), variable system size, 300GB total size
- Reference model: DimeNet++, 1.8M parameters

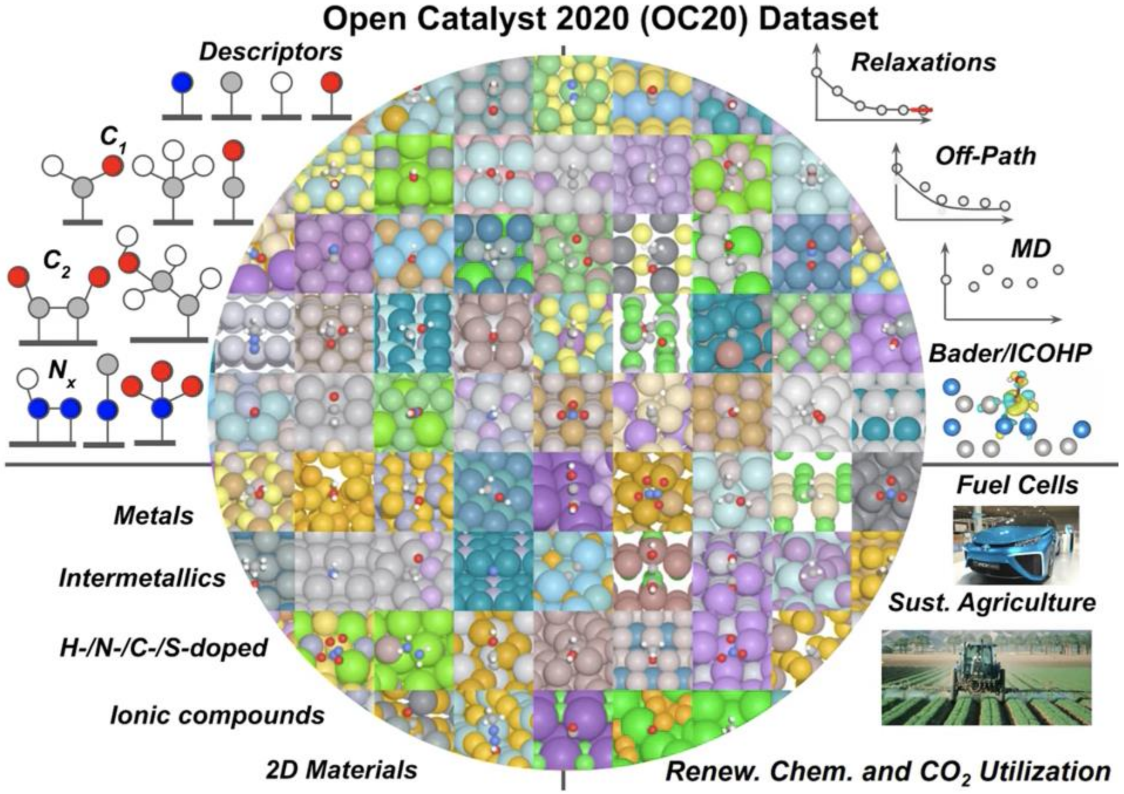

State of the Art AI Models are increasing in Model and Data size

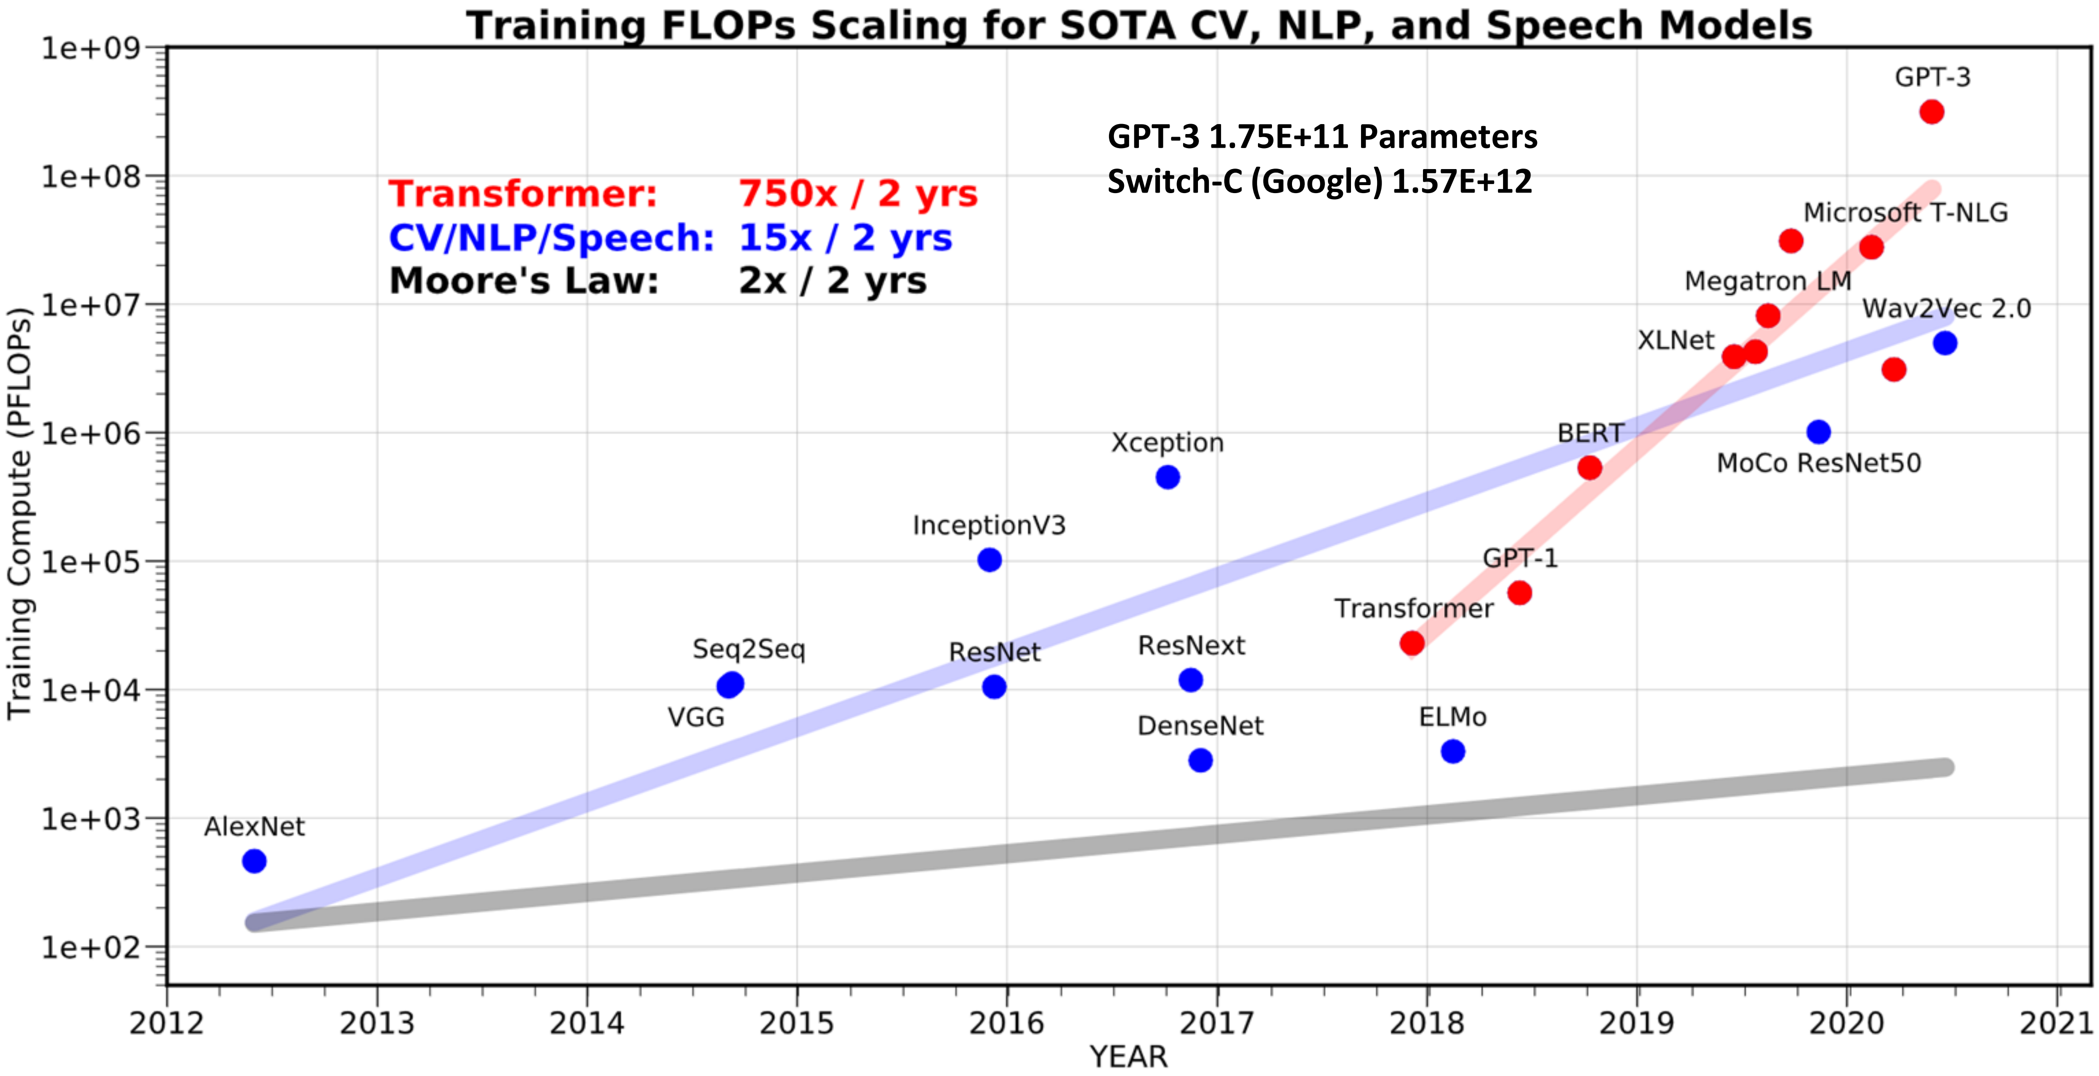

Source: https://medium.com/riselab/ai-and-memory-wall-2cb4265cb0b8

MLCommons Training “Winner” Performance

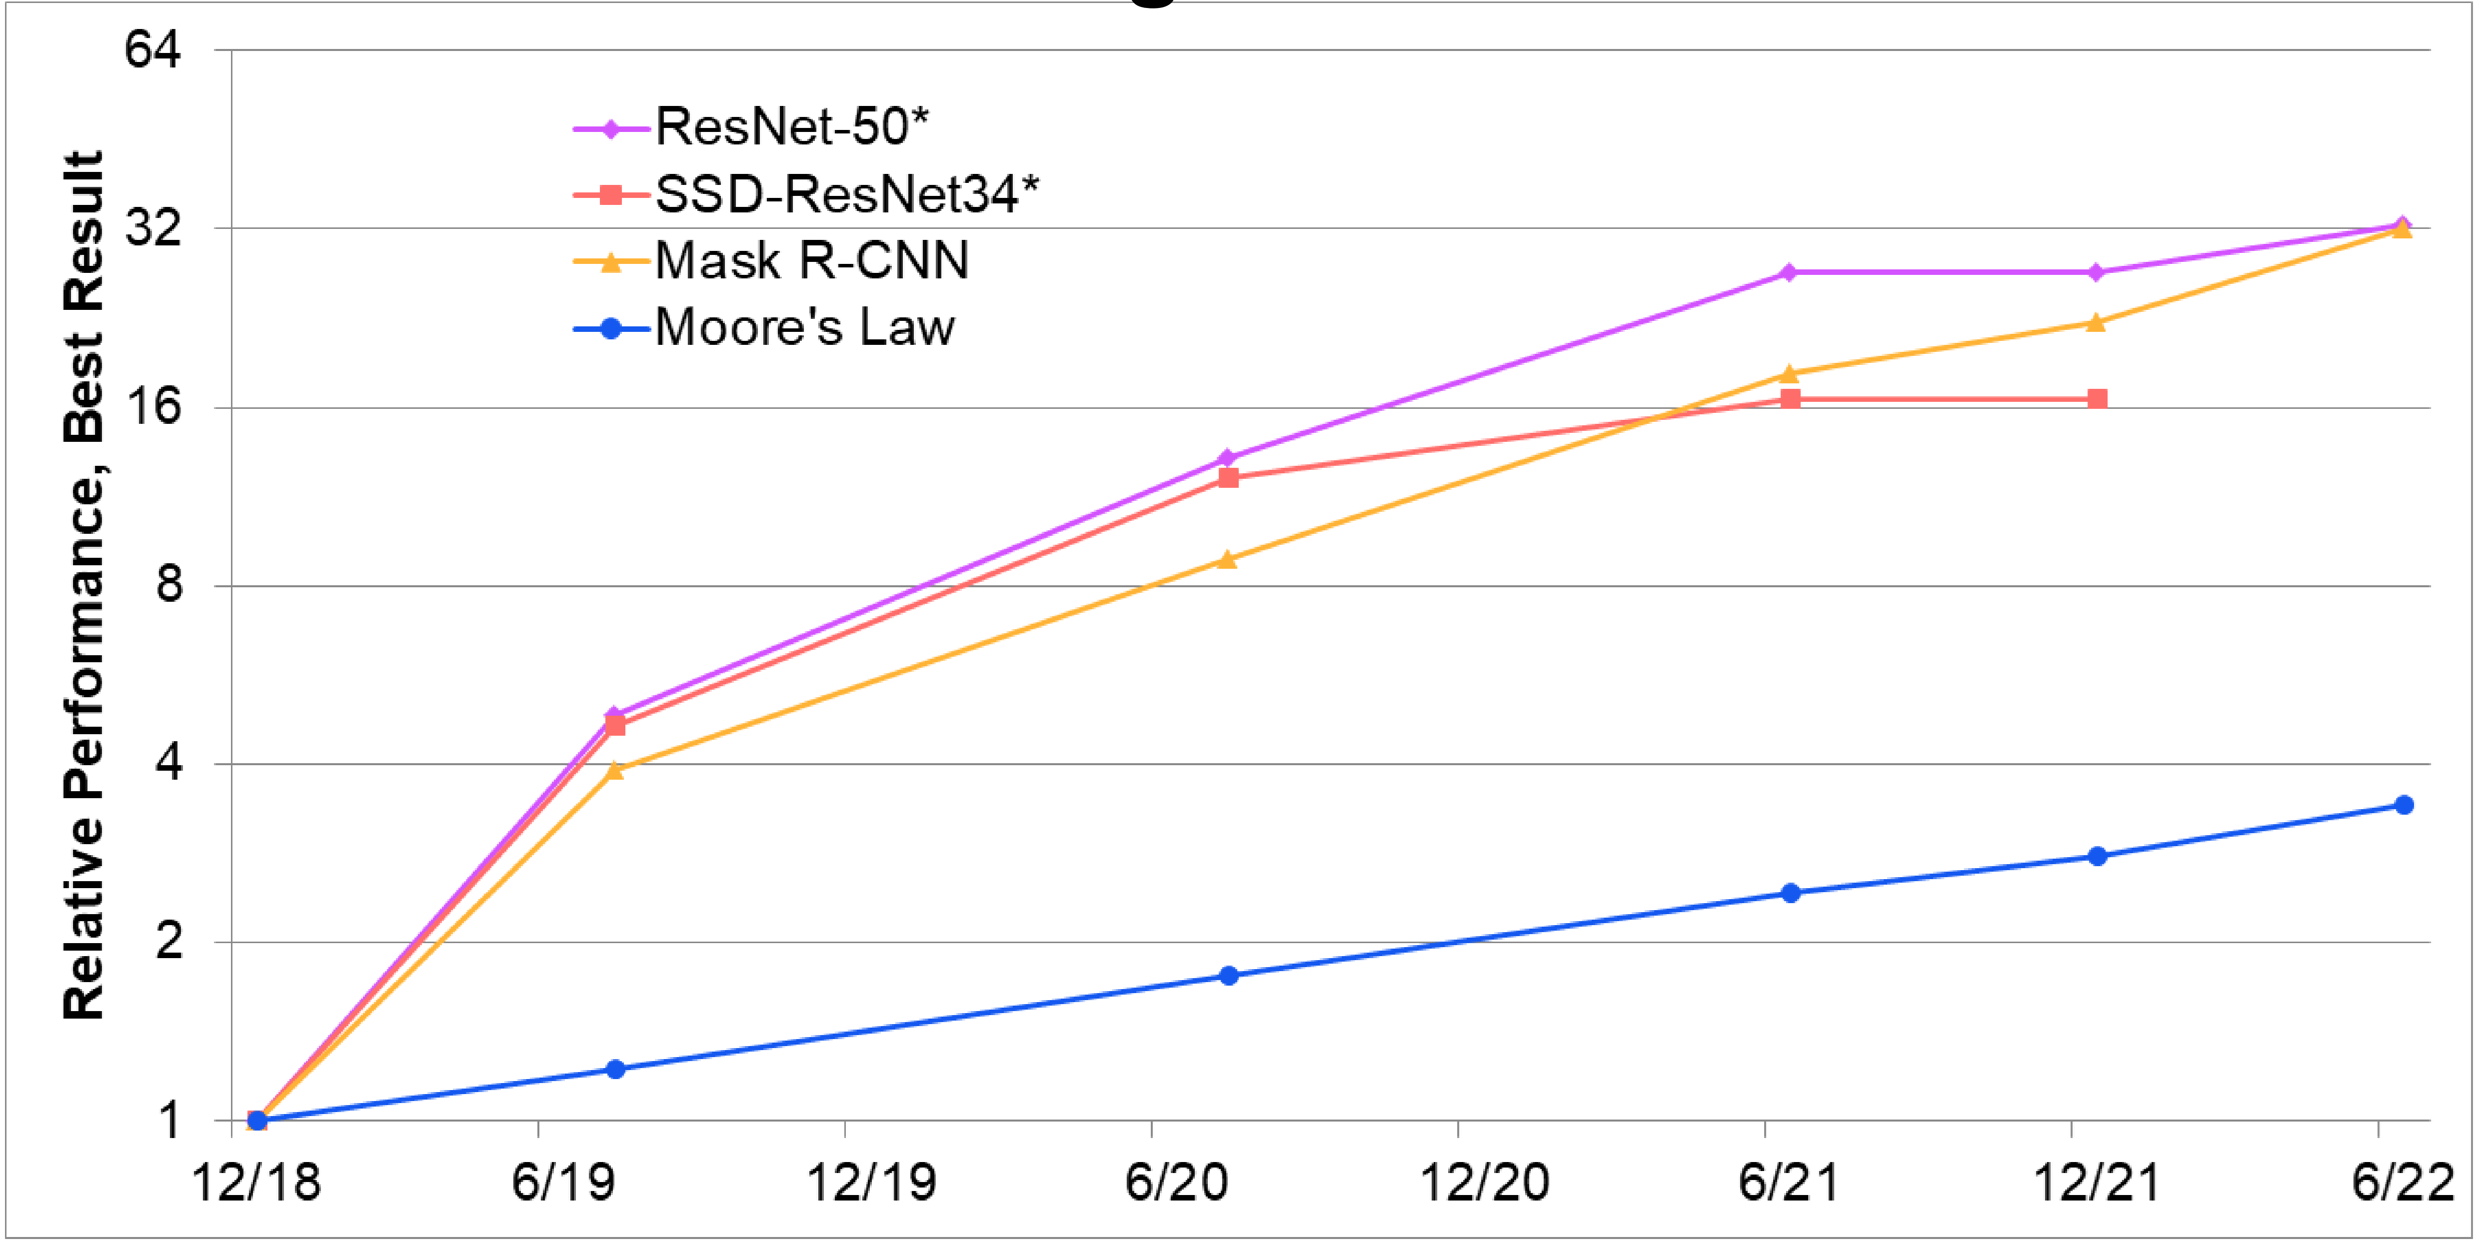

Remarks on MLCommons

- **Highest performance “cloud” systems** are from **Google**(4096 TPU-v4, fastest) and **NVIDIA** (4216 A100’s) but these are surely HPC “clusters”;
  - In fact, larger than supercomputer partitions used in responses to HPC Training MLPerf benchmarks
  - Not very efficient at maximum size!
- MLCommons has as much emphasis on **open datasets** as on models
- Industry participants typically HPC experts with sometimes University or DOE lineage
- Inference has **datacenter, edge, mobile** and **tiny**(embedded) benchmarks
  - MLCommons spans Computing Continuum
- **Power**, **Performance** and **Science Discovery** metrics
- Looking at using open data/benchmarks to produce an **AI for Science Gymnasium** to provide tutorials and examples that can be transferred across domains

Advancing Science with AI from the Edge to the Datacenter

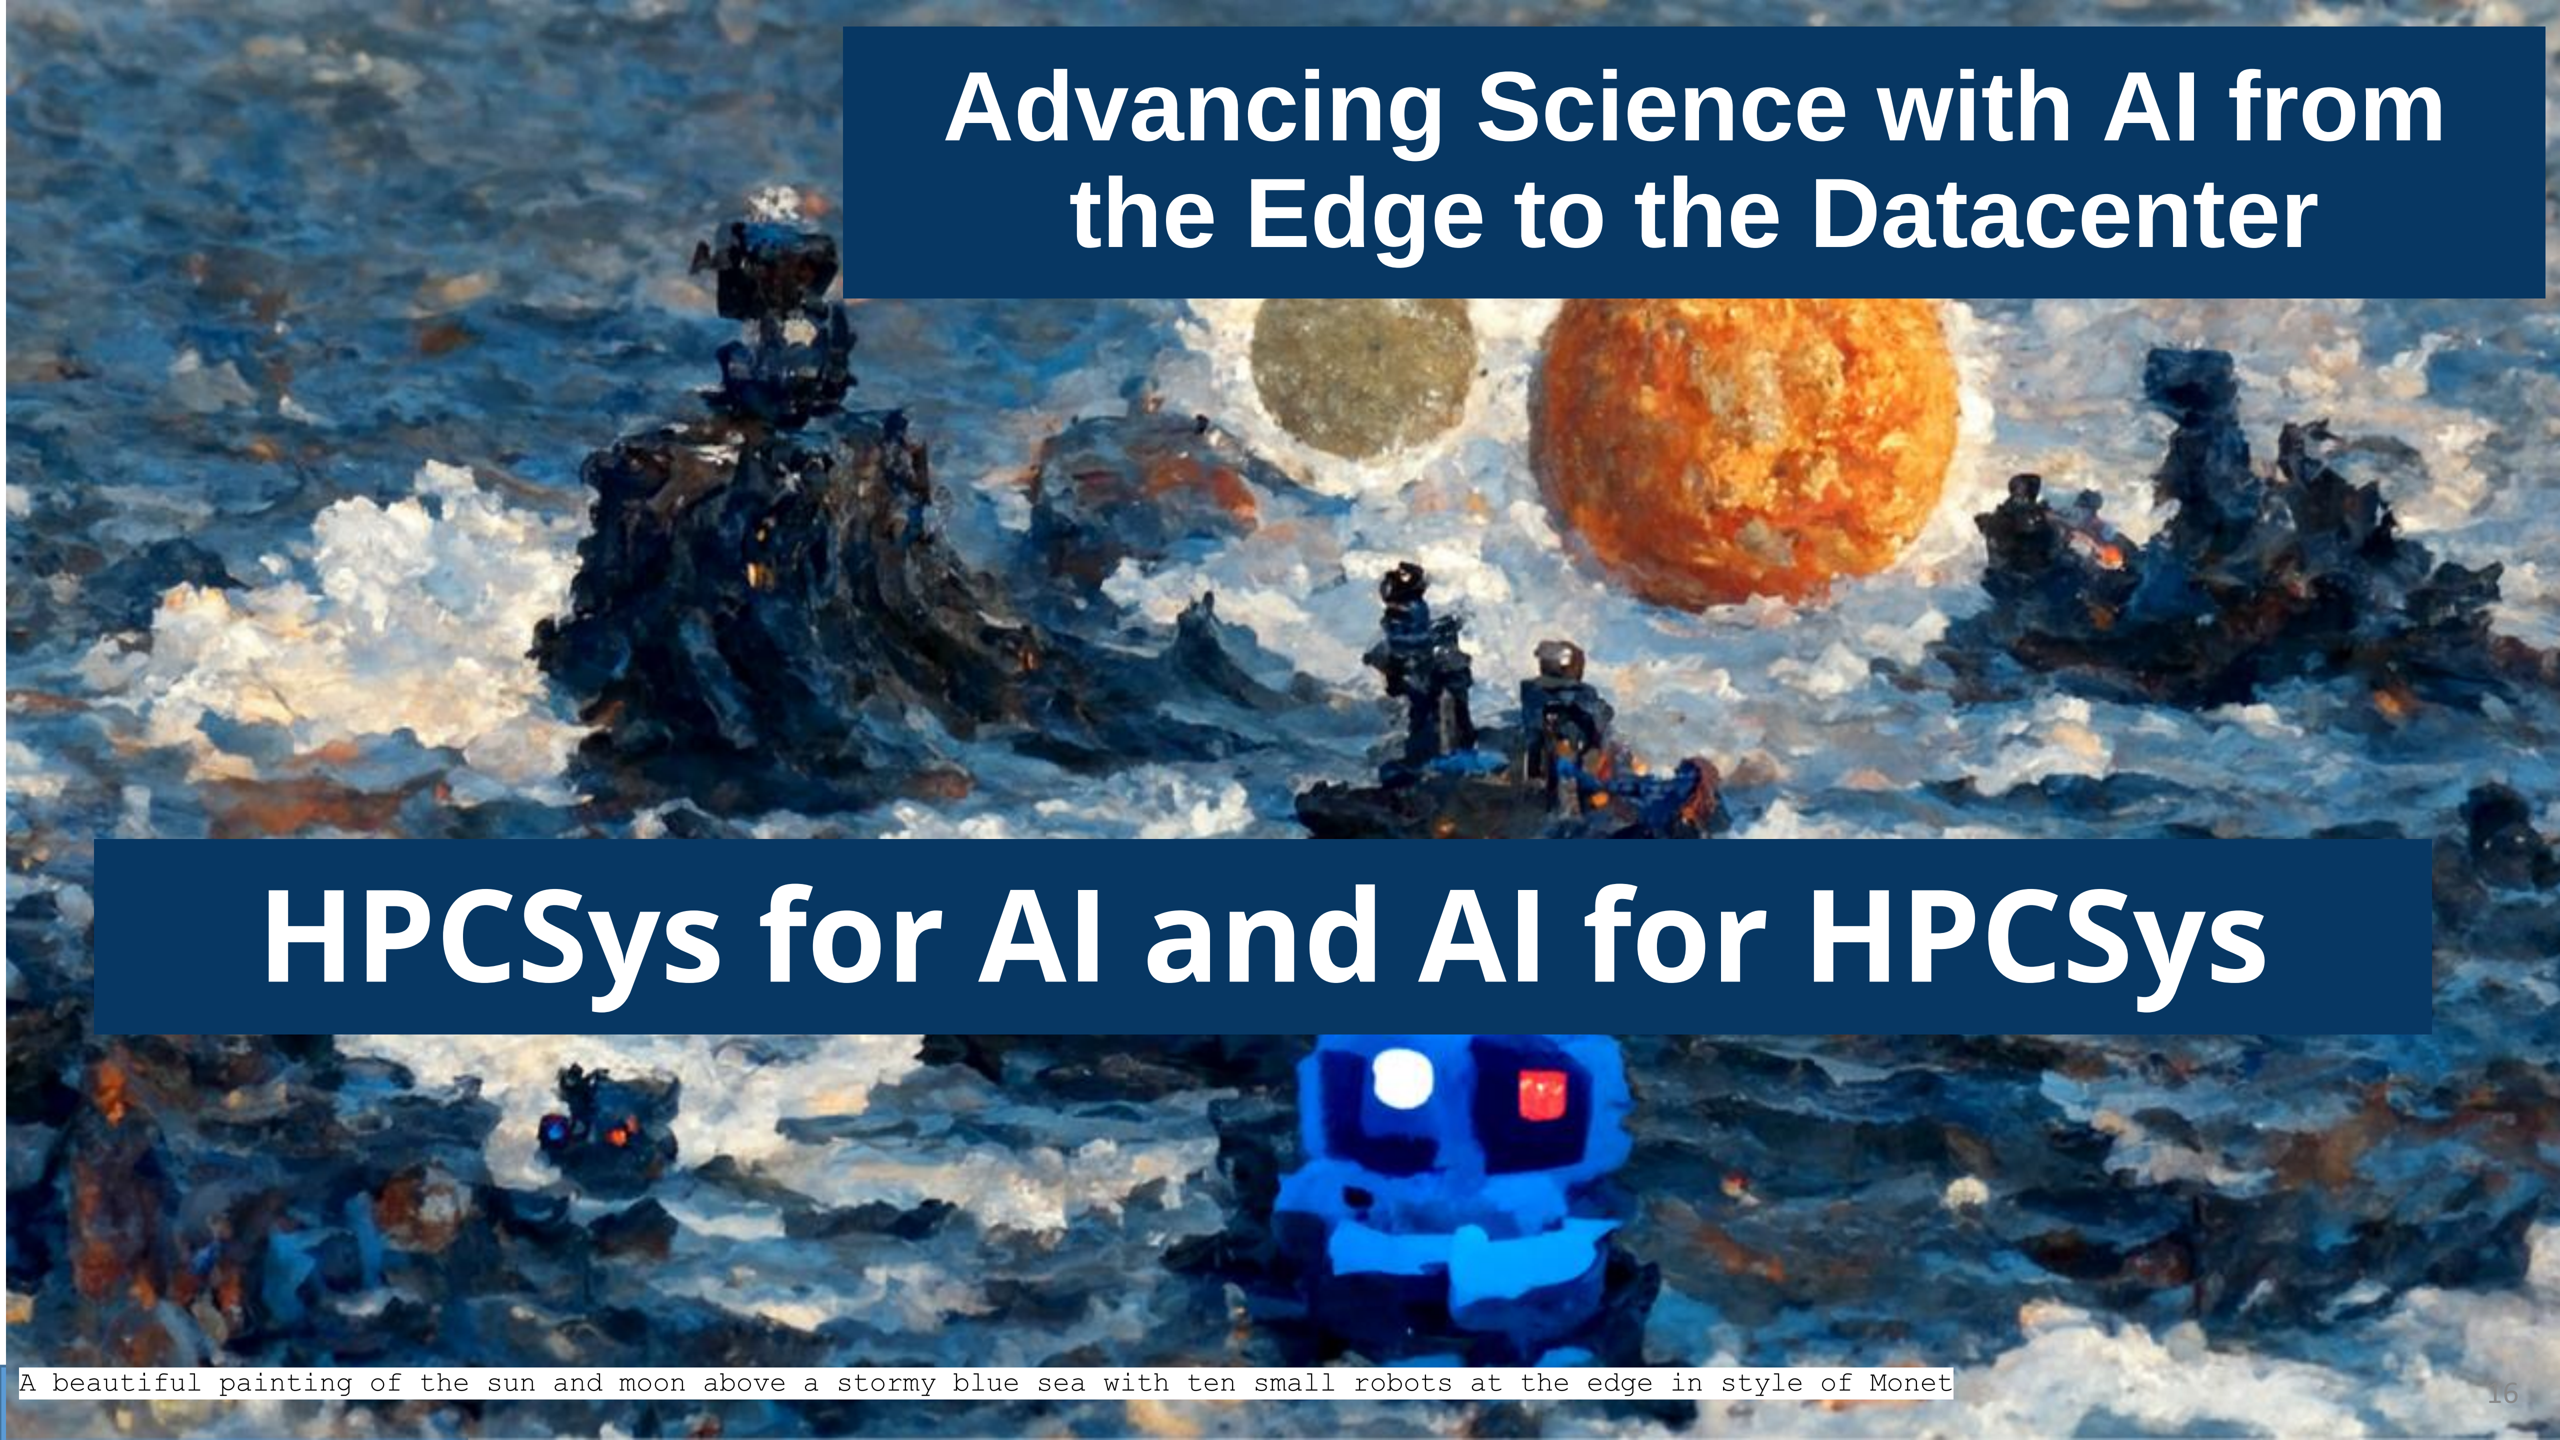

HPCSys for AI and AI for HPCSys

A beautiful painting of the sun and moon above a stormy blue sea with ten small robots at the edge in style of Monet

Microsoft Summer 2018: Global AI Supercomputer

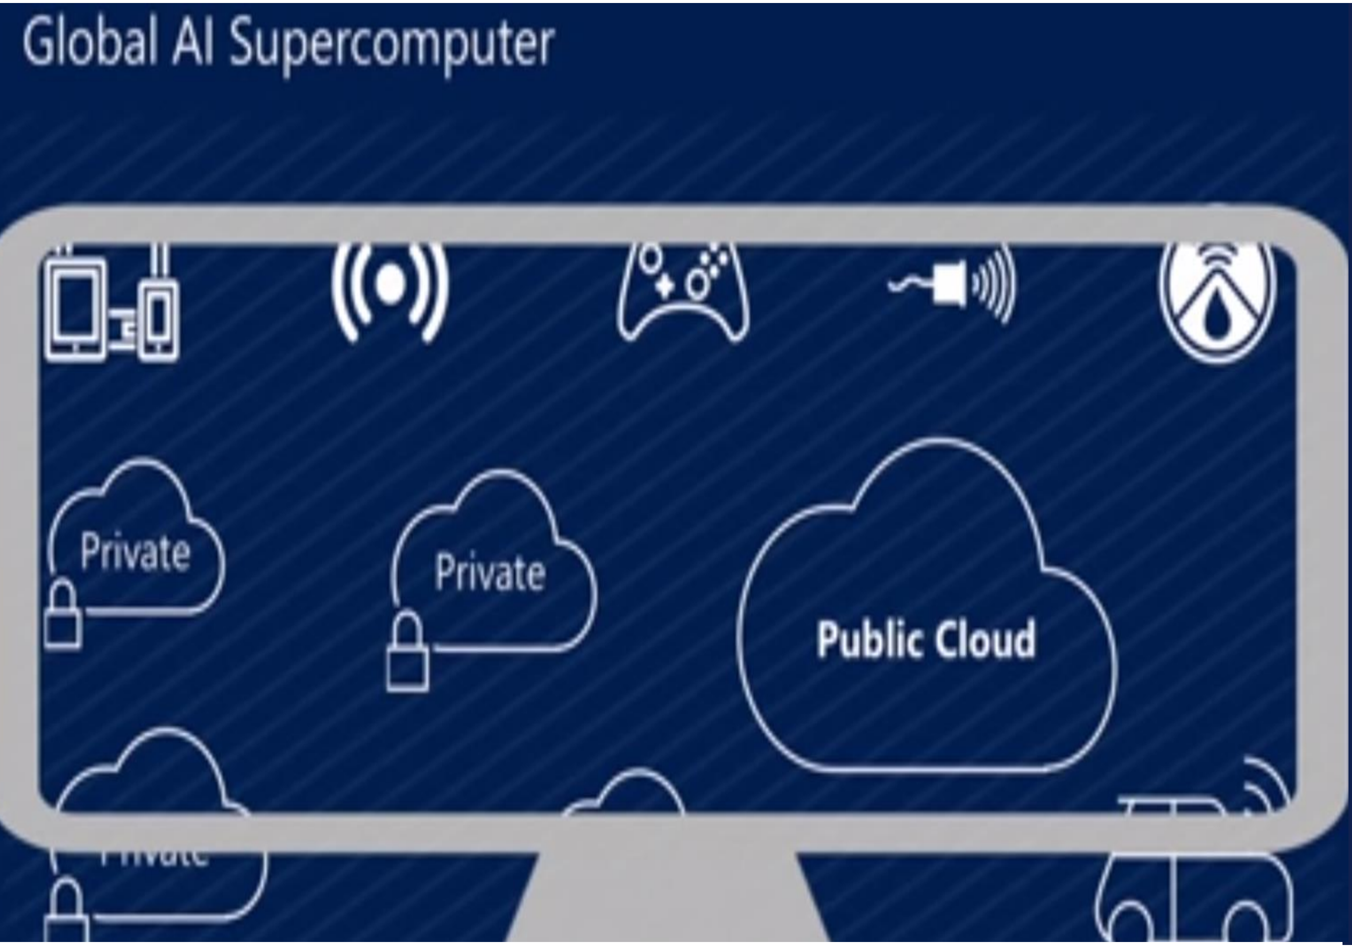

Intelligent Edge & Intelligent Cloud

Intelligent Edge = Collect Data, Make Inference, Take Action

many mobile devices with sensors (cameras, speakers, watches, IoT, ...)
limited processing capabilities
low latency and somewhat trusted

Intelligent Cloud = Aggregate Data, Analyze, and Train Models

virtually infinite computing resources
integrate and correlate data from many sources
mostly untrusted

By Donald Kossmann

Let’s look at AI for HPC or Systems

- Traditionally we did HPC Systems for AI -- run deep learning on a GPU BUT AI for HPC Systems is probably more interesting (stems from Jeff Dean 2017 NeurIPS)
- Currently in science the AI is used to enhance simulations (not data analytics) and dominantly the **AI used is Deep Learning**
- Both SC20 Gordon Bell winners used AI enhanced simulations a) Learn multi-particle potential b) Learn to explore phase space

  - With Shantenu Jha, Introduced a set of major categories of AI/ML for HPC and 8 subcategories but Shantenu usefully simplified to 3.
  - **AI in HPC** is learning and replacing traditional HPC software components; surrogates
  - **AI about HPC** is “in charge” of simulation and optimizing execution e.g. path through phase space
  - **AI out HPC** is AI running synergistically with an HPC job; such as post processing to find particle trajectories

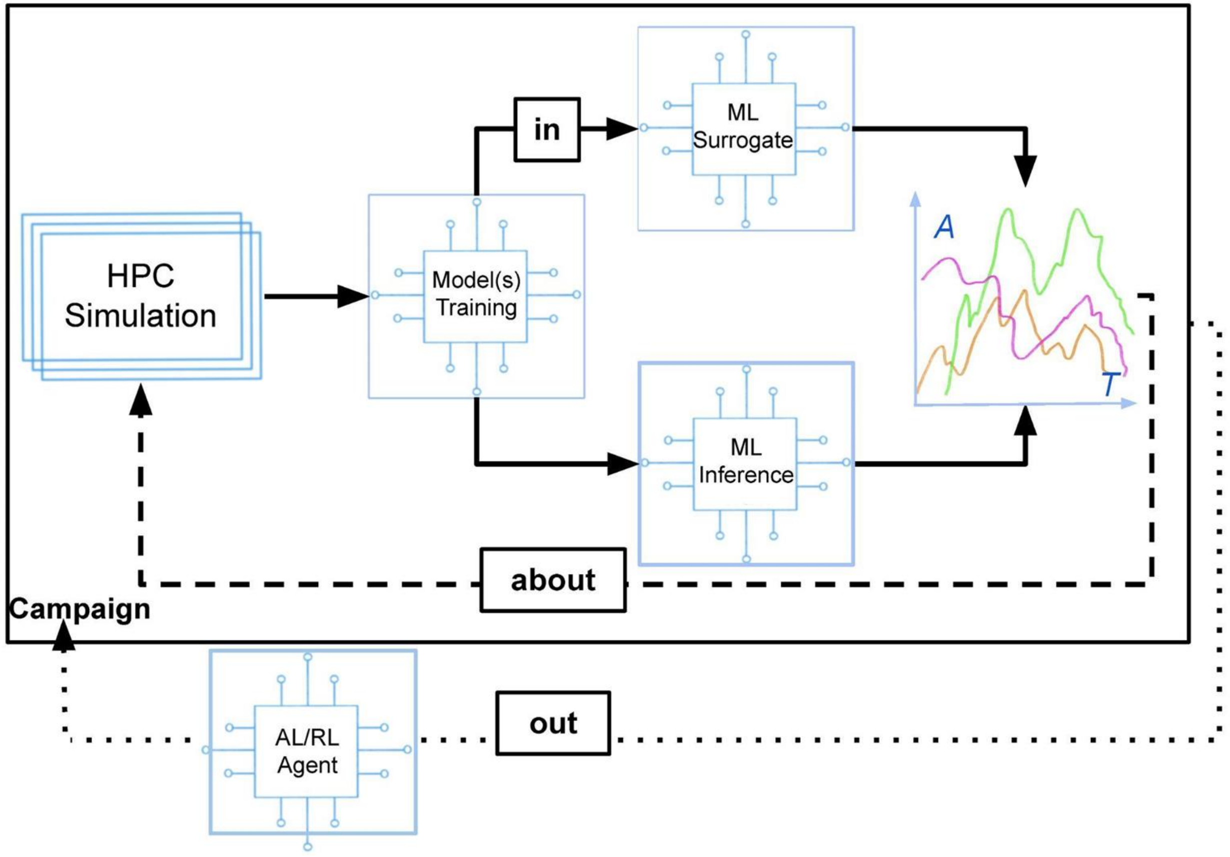

Up to two billion times acceleration of scientific simulations with deep neural architecture search

- January 23 2020 https://arxiv.org/pdf/2001.08055.pdf Updated October 2020
- 10 scientific cases including astrophysics, climate science, biogeochemistry, high energy density physics, fusion energy, and seismology, using the same super-architecture, algorithm, and hyperparameters.
- Approach also **dynamically choses deep network** and provides uncertainty estimation, adding further confidence in their use.

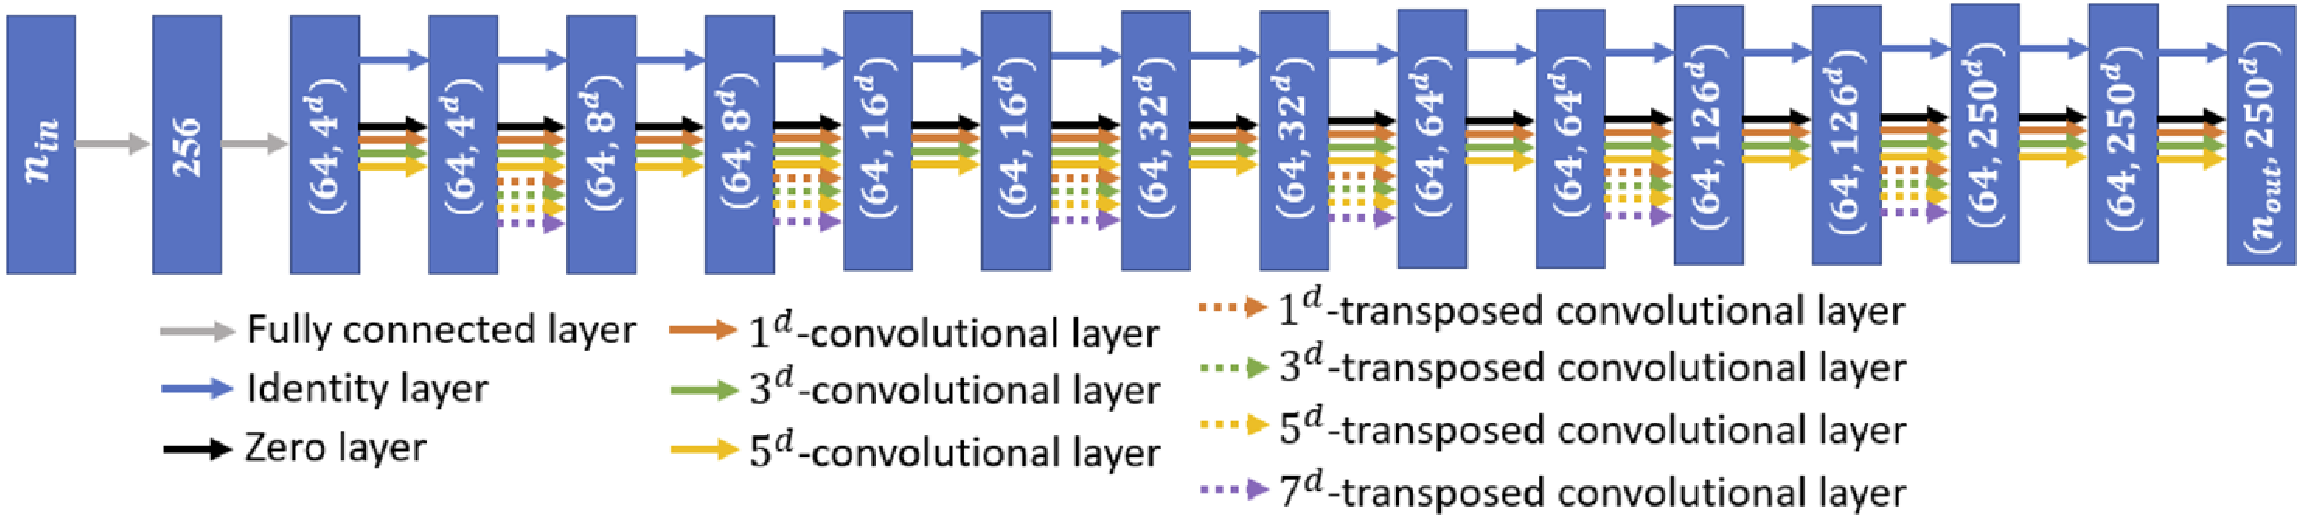

Examples of AI for HPCSys (work with JCS Kadupitiya, Vikram Jadhao)

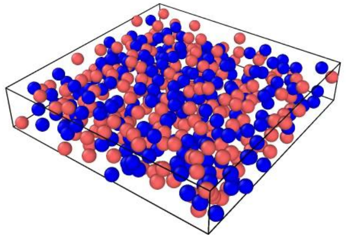

- Extraction of **ionic structure in electrolyte solutions confined by planar and spherical surfaces**.
- Classic HPC code written with C++ and accelerated with hybrid MPI-OpenMP.

The Learning Net

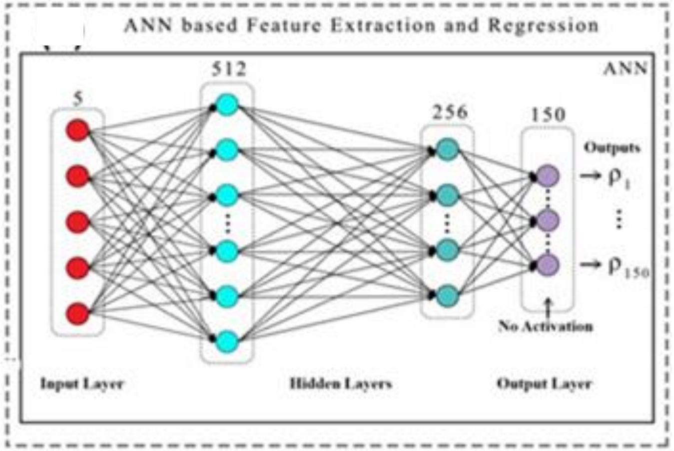

- Uses quite small Multi-layer perceptron AIP to predict 150 observables from 5 input parameters (~5000 in training set)
- AIP outperforms other AI choices
- Deployed on nanoHUB for **education** (an attractive use of surrogates so students get answers fast)
- General Electric uses similar approach to give interactive Engine design options (200 in training set)

Direct simulation compared to Surrogates

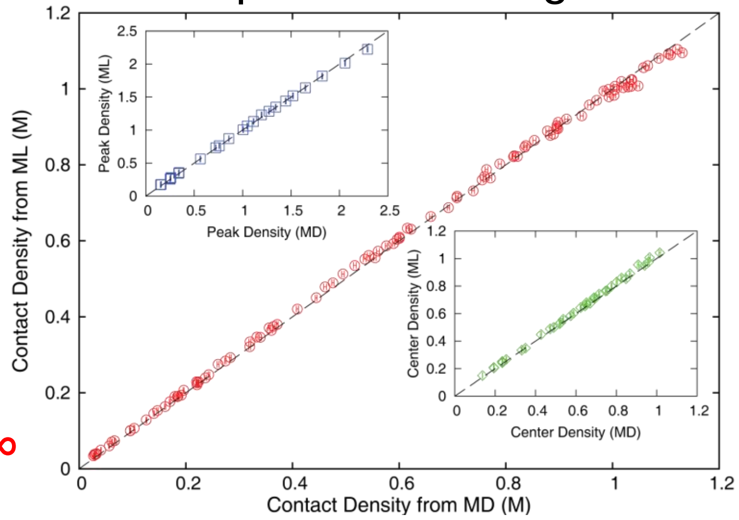

$$\text{Effective Speedup } S = \frac{T_{seq}(N_{lookup} + N_{train})}{T_{lookup}N_{lookup} + (T_{train} + T_{learn})N_{train}} \rightarrow \mathbf{10^6} \textbf{ as } \boldsymbol{N_{lookup}} \rightarrow \infty$$

Learn Newton's laws with Recurrent Neural Networks

JCS Kadupitiya, Vikram Jadhao

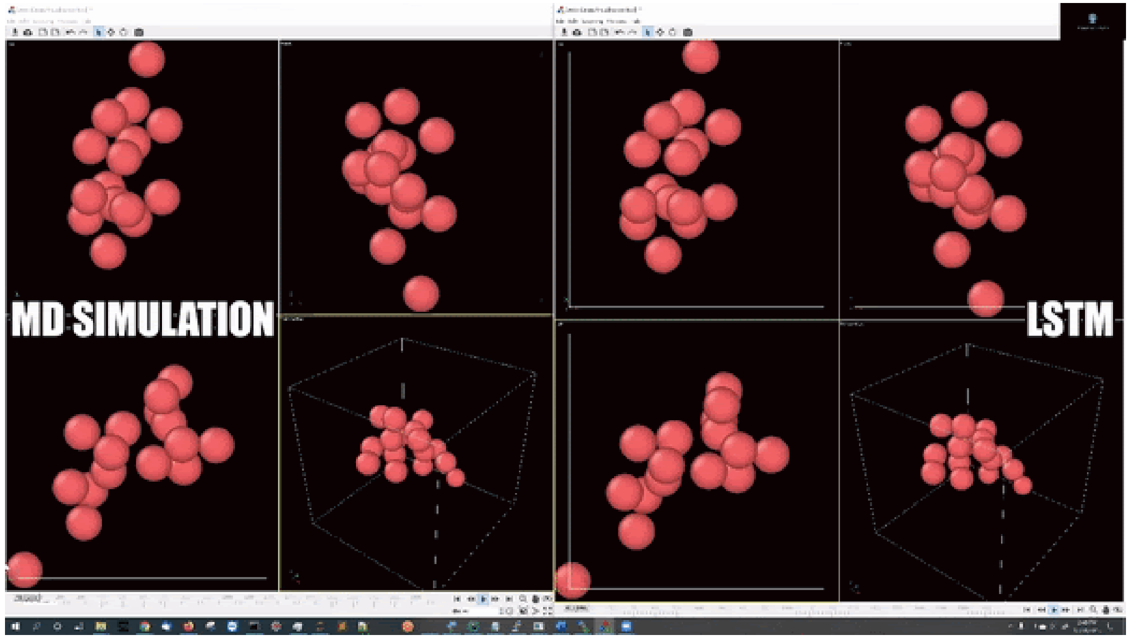

- **Deep Learning is revolutionizing (spatial) Time series Analysis**
- Good example is integrating sets of differential equations
- Train the network on traditional 5 time step series from (Verlet) difference equations
- Verlet needs time step .001 for reliable integration but
- Learnt LSTM network is reliable for time steps which are **4000 times longer** and also learn potential.
- **Speedup is 30000** on 16 particles interacting with Lennard-Jones potentials
- 2 layer-64 units per layer LSTM network: **65,072 trainable parameters**
- 5000 training simulations

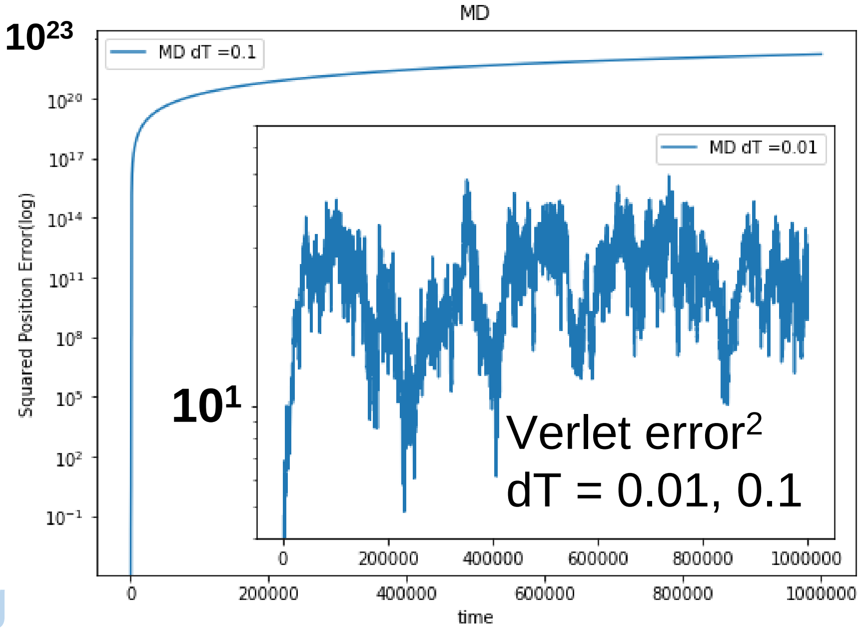

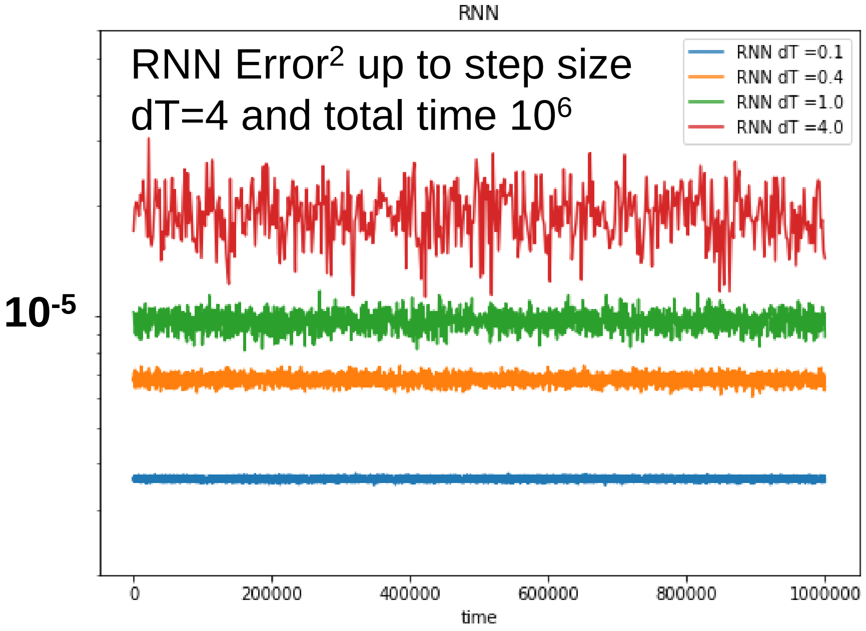

AI Control of Networking Traffic

Work of Mariam Kiran and ESNET
SPA = Shortest Path Algorithm

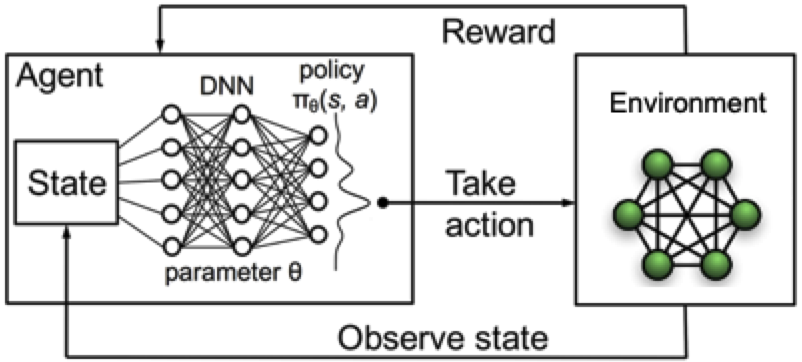

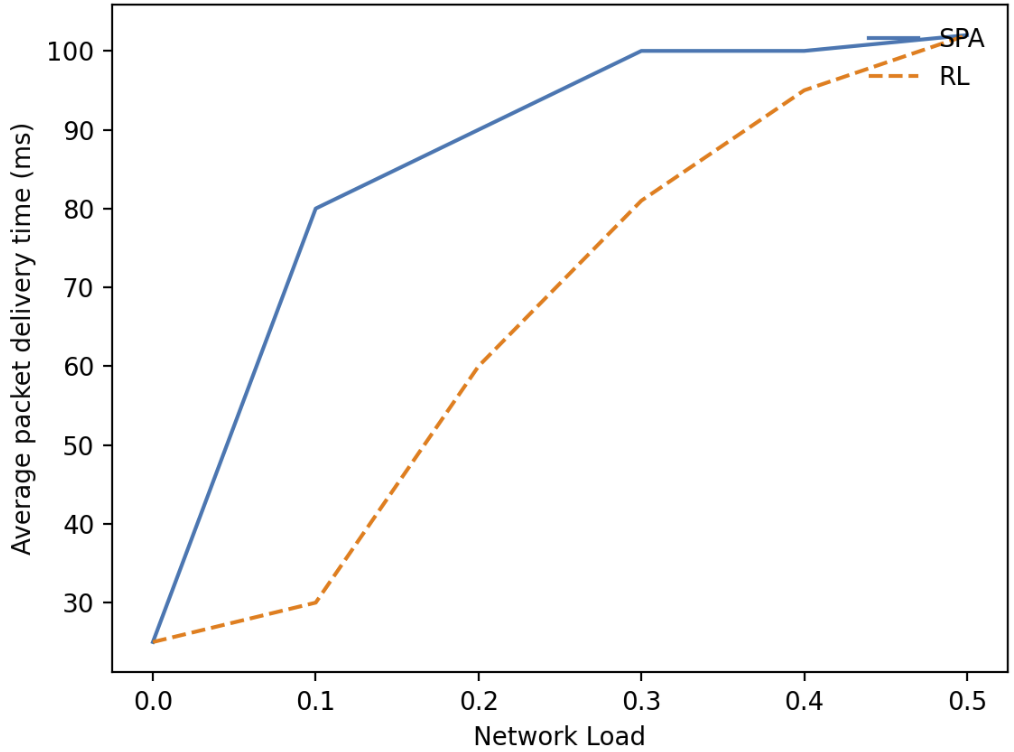

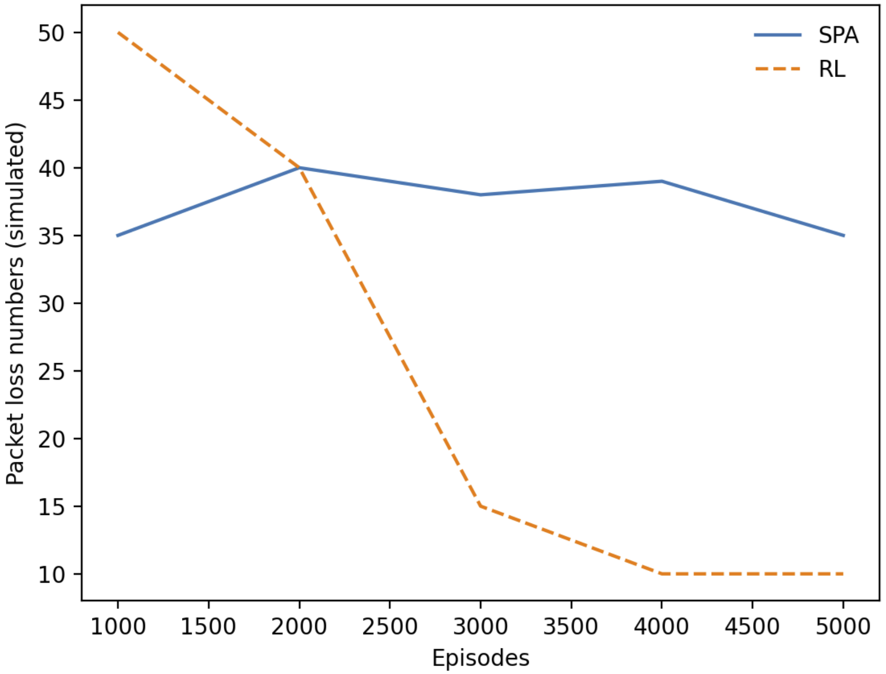

Futures of Surrogates and Computational Science

- Is Surrogate-based Computational Science the future?
  - Traditional simulations used to train very flexible surrogates
  - Some production examples (as in GE’s use in engine design) clear
  - Value of surrogates in education clear
  - Challenge is cases where a **simulation set** (mapping to a single surrogate) is small
  - Surrogates run well on GPU’s; traditional simulations may run well on GPU’s

- Need a neural architecture (**meta-surrogate or Foundation model**) which can both describe individually many different simulation sets and for each set be as broad as possible
  - So a computational scientists can base their work around one (or a few) surrogates

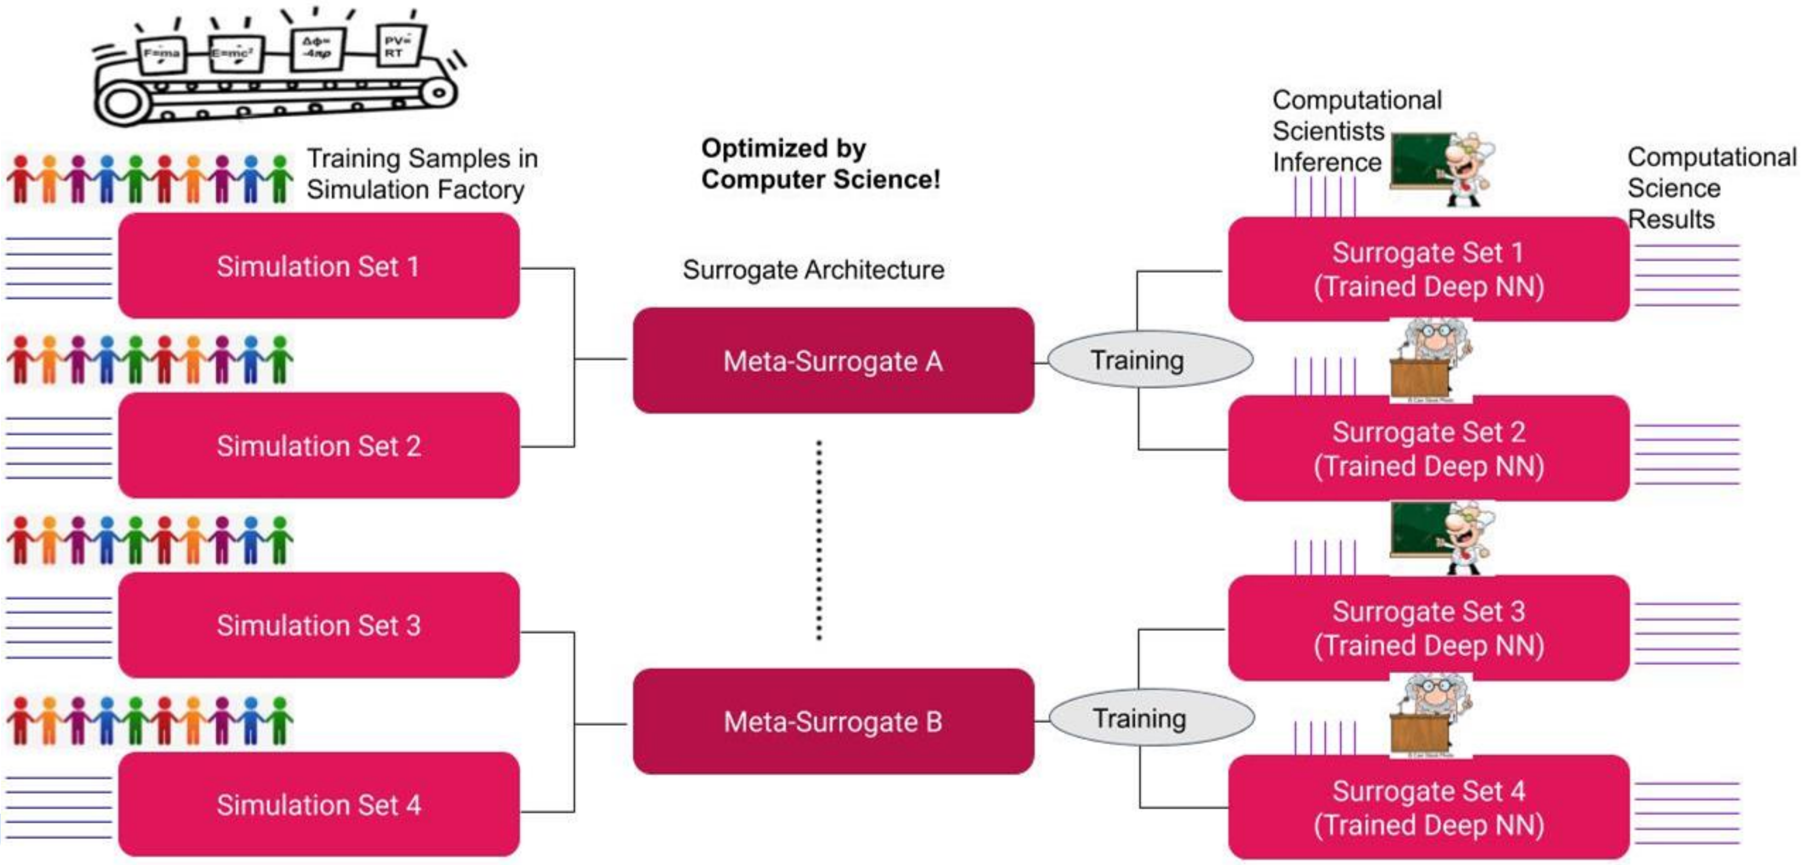

Process for Science Foundation or meta-Models

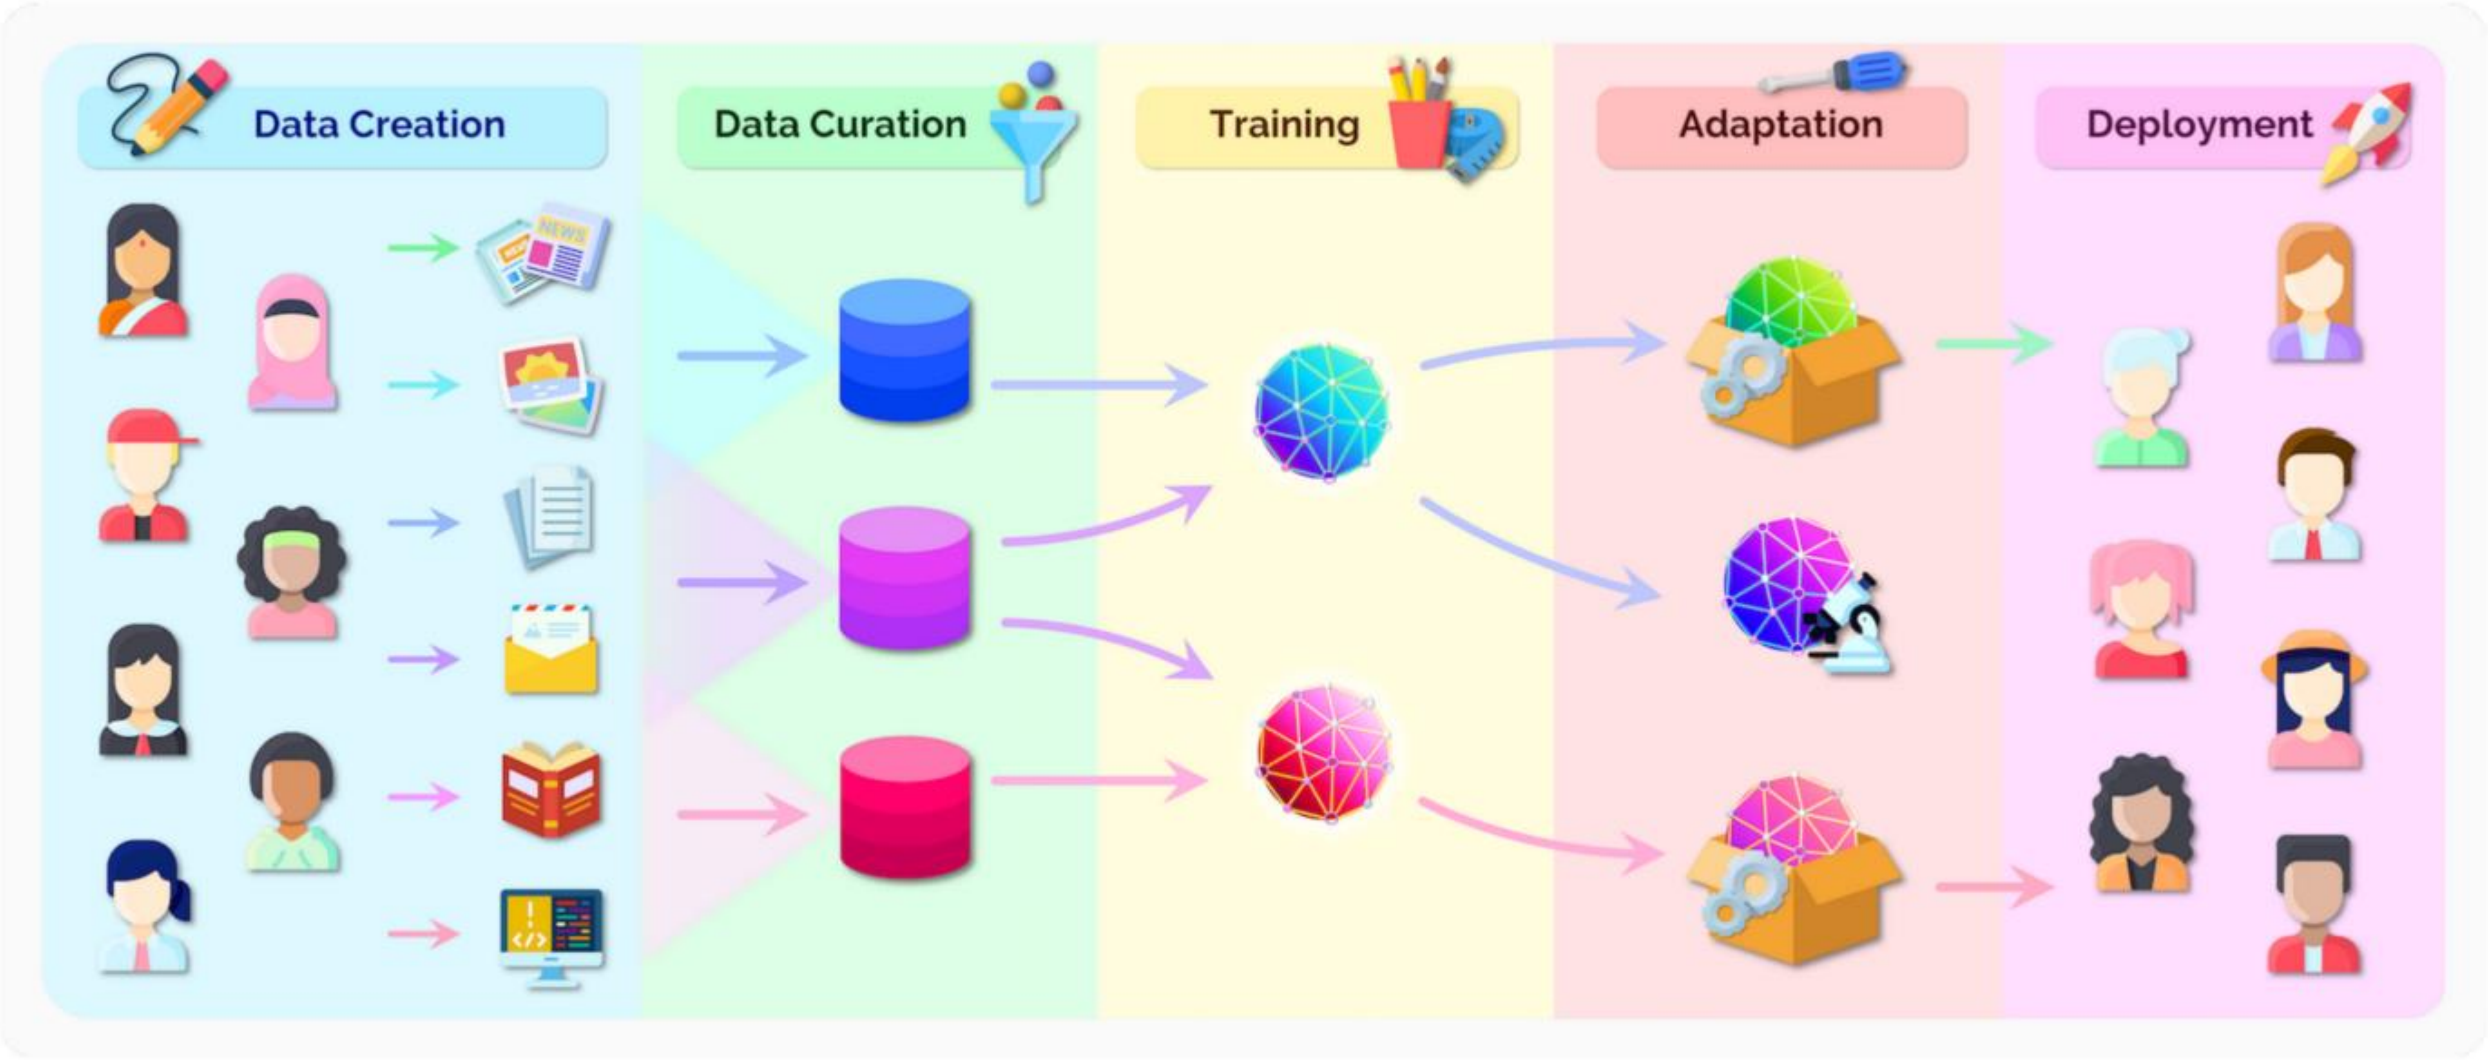

Advancing Science with AI from the Edge to the Datacenter

Way Forward in AI for Science

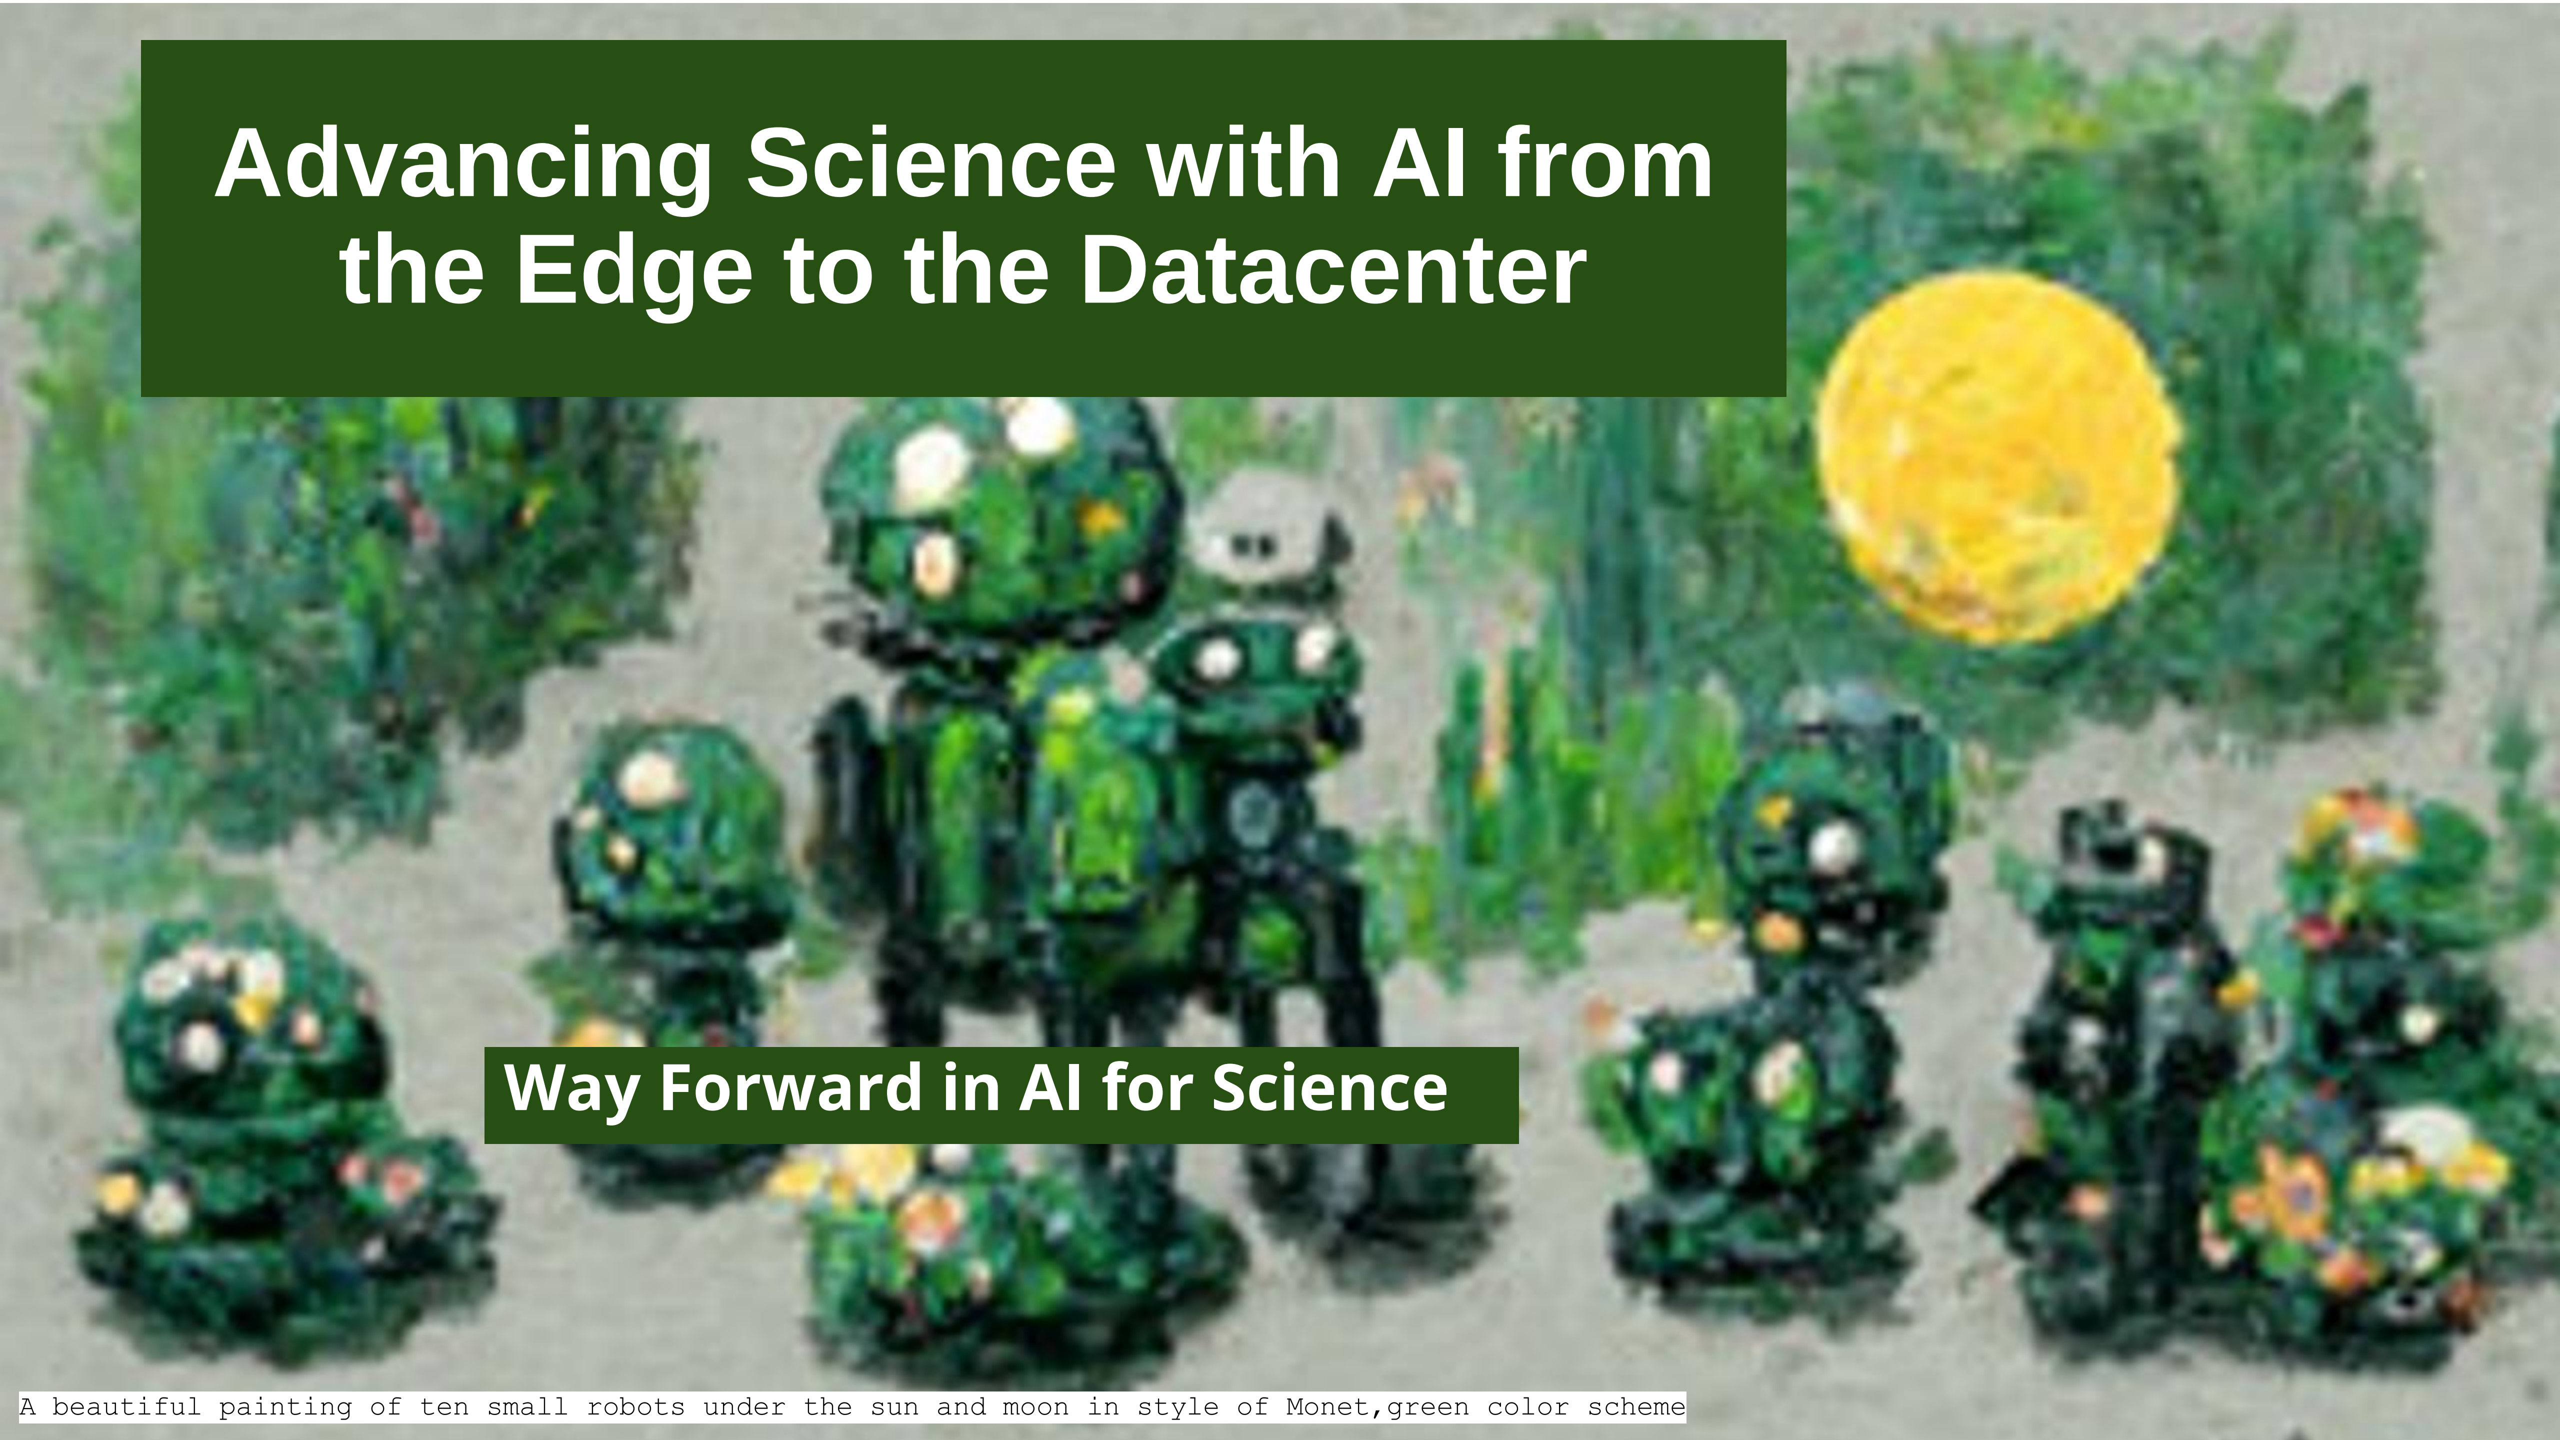

A beautiful painting of ten small robots under the sun and moon in style of Monet,green color scheme

Research Plan: AI for Science Patterns

- Currently, **deep learning** plays a dominant role in science, from data analytics and simulation surrogates to policy decisions.
- We can group the system structure into a “handful” (two hands ≤10) of **patterns** such as
- an **Overall Reasoner** based on **reinforcement learning** and large language models to learn the world’s knowledge and control experiments;
- **Image-based systems** for astronomy, pathology, microscopy, and light scattering;
- **Graph-based systems** such as in social media and traffic studies; represent molecular and other structure;
- **Dense systems to** map structure to properties as in drug discovery;
- **Time series and sequence (Transformer)** models as in language, earth, and environmental science;
- **GAN/Diffusion** models to generate scenarios as in datasets to test experimental system.
- **All Network types** can be mixed together as in text to image systems DALL-E and Imagen.
- We imagine taking good **examples of each pattern** and **supporting them** with high-performance, easy-to-use environments in end-to-end systems, including data engineering.
- This would cover parallelism, storage and data movement, security, and the user interface. Then we explore multiple examples of each pattern, possibly leading to “**Foundation**” models for each of them.
- **Provide AI for Science Gymnasium (in MLCommons) for tutorials, demos**

AI and Systems Interaction

- **The System hosts the AI software**
- AI writes the System Software (Software 2.0)
- AI designs the Systems CPU, Storage, memory etc.
- AI designs the Network Architecture/Hardware

- AI monitors and manages system: **DevOps AIOps** (including **AI for Data center**)

- **AI in real time optimizes system communication on wide area network**
- **AI in real time optimizes job scheduling; Current heuristics**, New DL/RL Based?
- AI in real time optimizes I/O, heterogeneous compute, local area network
- AI learns all possible functions/services/job itself inside a job and supplies very fast results (**surrogates**)

- All of these capabilities are integrated together with new distributed operating system
- Initially runs on single nodes, single clusters, single supercomputer, single cloud; then multiple clouds, clusters, supercomputers

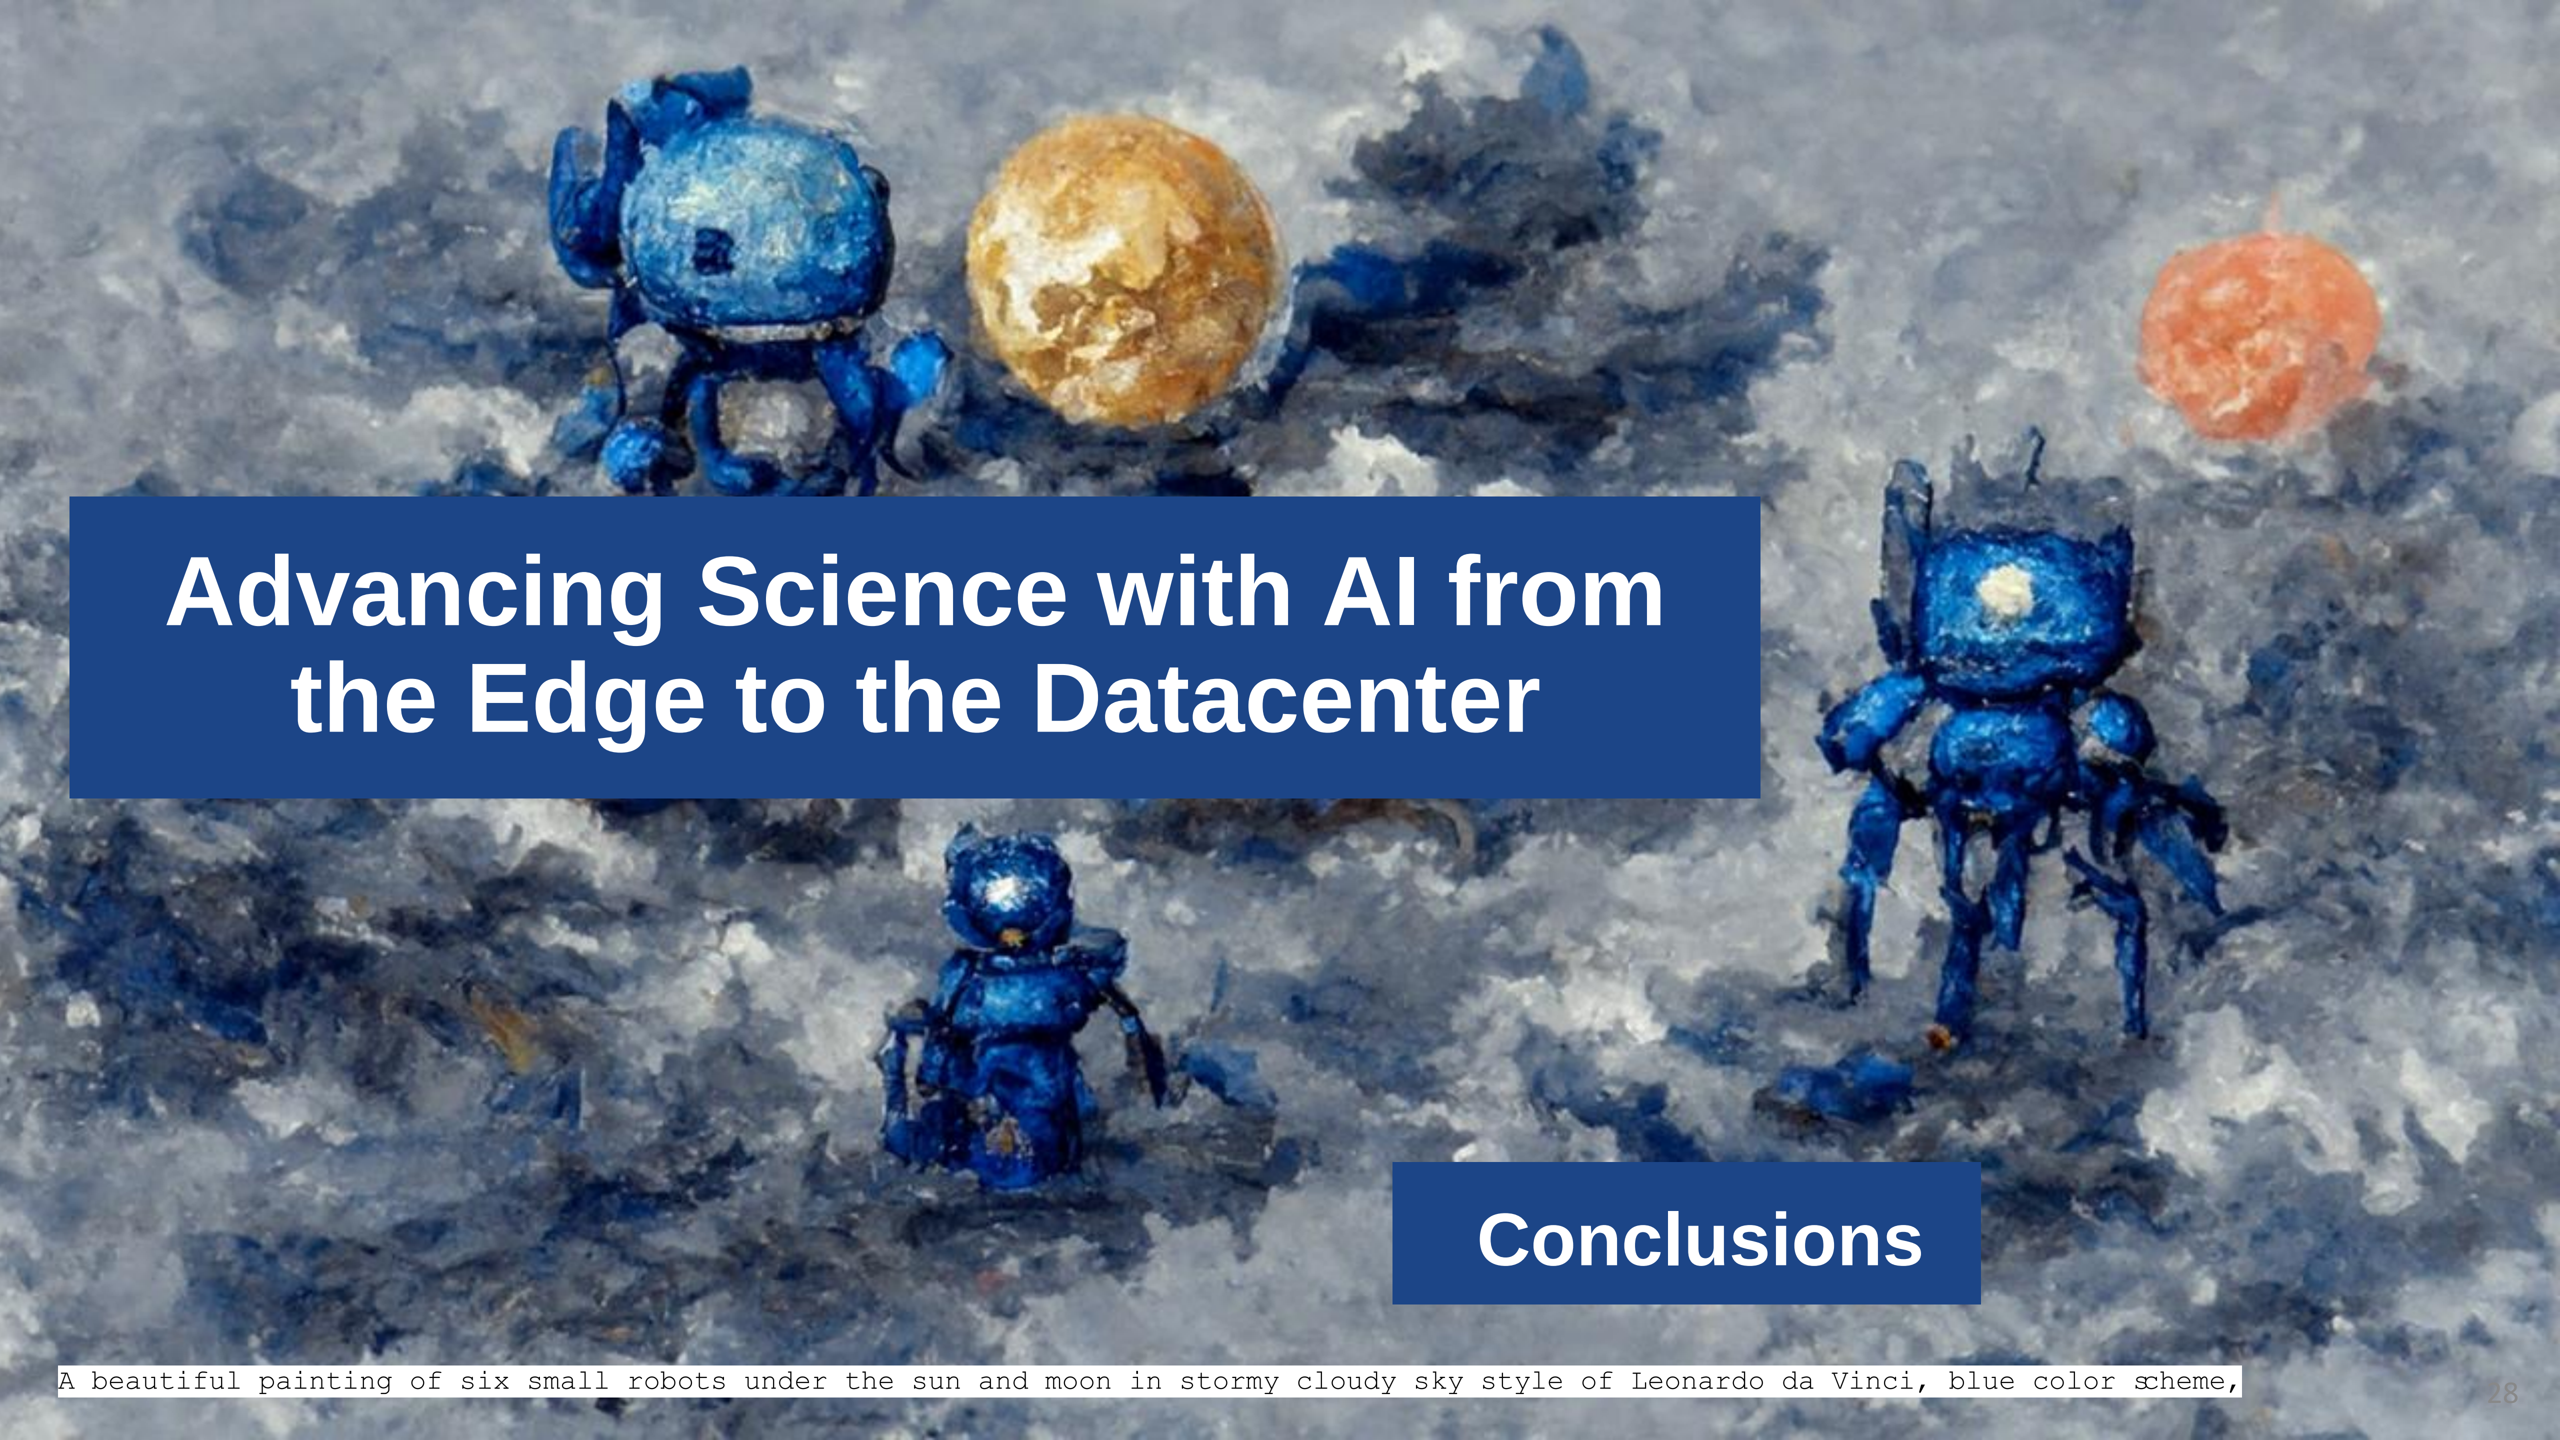

Advancing Science with AI from the Edge to the Datacenter

Conclusions

A beautiful painting of six small robots under the sun and moon in stormy cloudy sky style of Leonardo da Vinci, blue color scheme,

Conclusions

1

AI is promising pervasive approach to continue **increases in system performance**

- The largest factors come from the inner layers but are application specific
- outer layers (scheduling, interoperability of HPCClouds) give useful increases

2

AI for HPC Systems very promising where we could aim at

- First Zettascale **effective** performance in next year or so
- Hardware/Software aimed at **general AI assisted speedup of computation**
- **Your health can be engineered** with AI-assisted **personalized nanodevices designed** based on **the AI-assisted digital twin** of disease in your tissues

3

Global AI and Modeling Supercomputer good framework with HPC Cloud linked to HPC Edge

- Training on cloud; Inference and some training on the edge
- **AI will design, monitor, control experiments**; generate apparatus simulations; analyze its data and propose new theories

4

Need to develop **meta-models** (**Foundation models**) valid across many domains

- Broad use of deep and reinforcement learning

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov