HPC, Big Data, and Machine Learning Convergence

Washington DC

Geoffrey Fox | July 16, 2019

gcf@indiana.edu, http://www.dsc.soic.indiana.edu/, http://spidal.org/

Charge

Supported by National Science Foundation through Awards

- 1443054 CIF21 DIBBs: Middleware and High Performance Analytics Libraries for Scalable Data Science
- 1720625 Network for Computational Nanotechnology - Engineered nanoBIO Node

- Primary questions:
  - What are the challenges and opportunities presented by the convergence of HPC, big data, and machine learning?
  - What is driving this convergence and what capabilities might it provide over the current scope/timescale of traditional HPC?
- Secondary questions:
  - What are the main drivers or killer apps for convergence? Does that differ between the US and abroad?
  - How will the integration of machine learning in to modeling and simulation workflows change or improve simulation performance/fidelity/speed capability?
  - How is convergence impacting the development and acquisition process for leadership HPCs?
  - Are there separate technical paths/discussions happening in the cloud/grid computing space vs. the stand-alone supercomputer space?

Evolution of Interests, Technologies and Communities

- AI/ML
- Systems
- HPC
- Cloud
- Big Data
- Edge
- Data Science v. AI First

Papers Submitted: Comparing 4 Conference Types

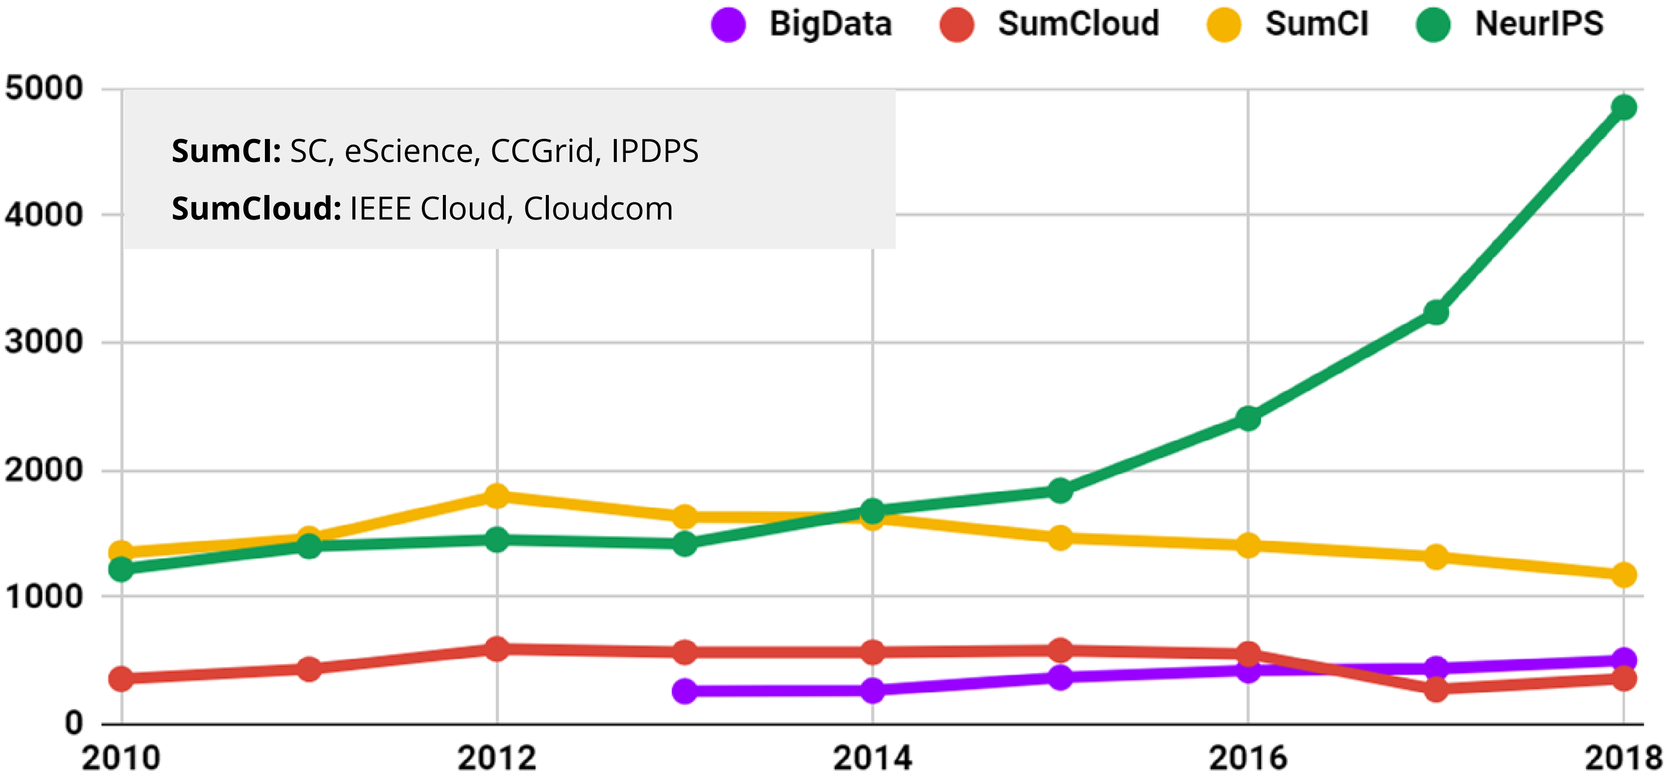

Trends in AI, Big Data, Clouds, Edge, HPC over last 5 years

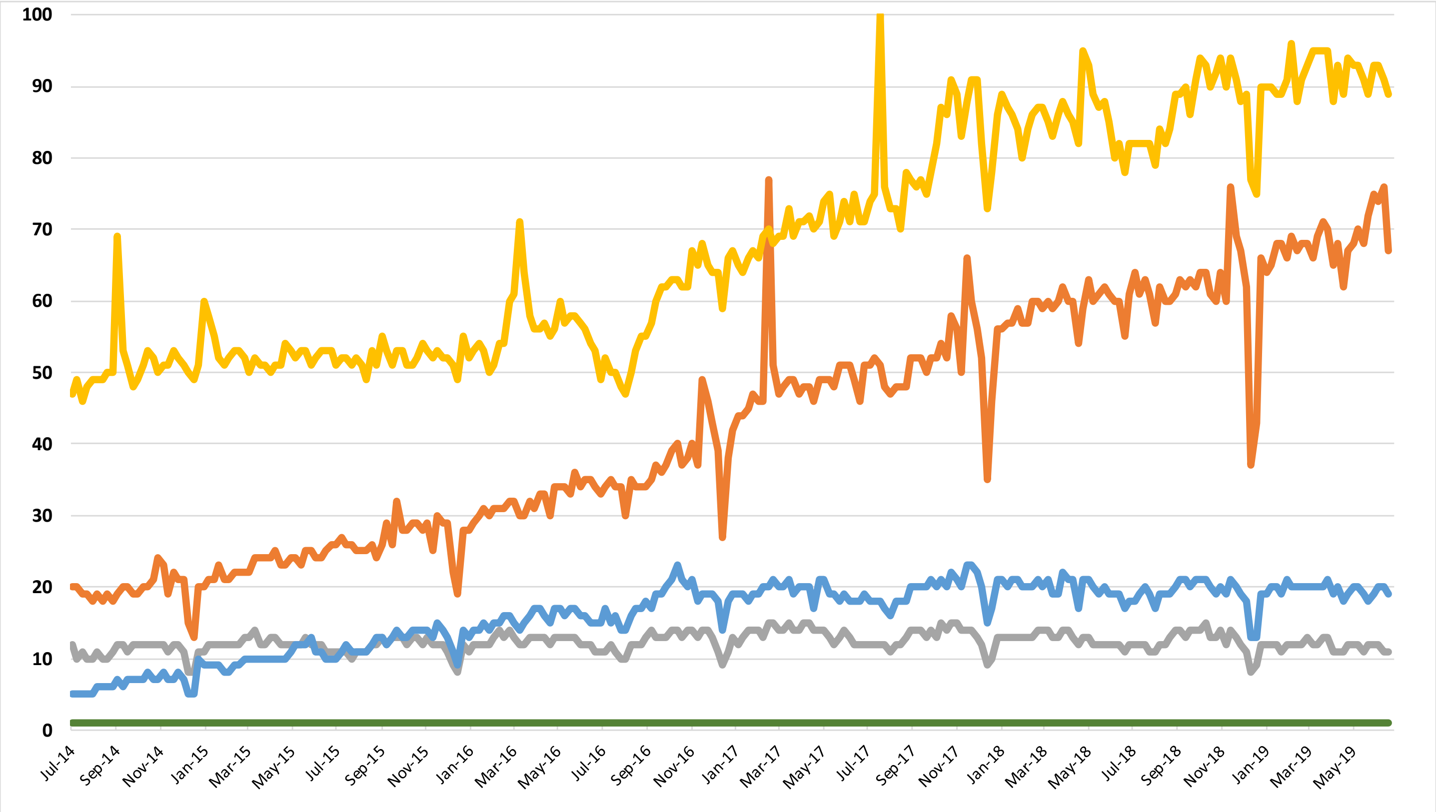

- → Artificial Intelligence
- → Amazon Web Services (Proxy for cloud computing)
- → Internet of Things
- → Big Data
- → High Performance Computing (1%)

Importance of HPC, Cloud, Edge and Big Data Community

- **HPC Community** not growing in terms of obvious metrics such as new faculty advertisements, student interest, papers published
- **Cloud** Community quite strong in Industry; relatively small academically as Industry has some advantages (infrastructure and data)
- **Big data and Edge** communities strong in Academia and Industry
  - Big Data definition unclear but it is growing although still quite small in terms of dedicated activities
- At IEEE services federation in Milan just completed; **Cloud Edge IoT** and **Big Data** conferences had significant overlap – not surprising as most IoT/Edge systems connect to Cloud and essentially all Big Data computing uses cloud.
- All these academic fields need to align with mainstream (**Industry**) systems
- **Microsoft described the AI Supercomputer linking Intelligent Cloud to Intelligent Edge**

Importance of AI and Data Science

- **AI (and several forms of ML) will dominate the next 10 years** and it has distinctive impact on applications whereas HPC, Clouds and Big Data are important and essential enablers

- **AI First** popular with Industry with 2017 Headlines
    - *The Race For AI: Google, Twitter, Intel, Apple In A Rush To Grab Artificial Intelligence Startups*
    - *Google, Facebook, And Microsoft Are Remaking Themselves Around AI*
    - *Google: The Full Stack AI Company*
    - *Bezos Says Artificial Intelligence to Fuel Amazon's Success*
    - *Microsoft CEO says artificial intelligence is the 'ultimate breakthrough'*
    - *Tesla’s New AI Guru Could Help Its Cars Teach Themselves*
    - *Netflix Is Using AI to Conquer the World... and Bandwidth Issues*
    - *How Google Is Remaking Itself As A “Machine Learning First” Company*
    - *If You Love Machine Learning, You Should Check Out General Electric*

- Could **refine emphasis on data science** as **AI First X**
    - where X runs over areas where AI can help
    - e.g. **AI First Engineering**; AI First Cyberinfrastructure; AI First Social Science etc.

ML/AI needs Systems and HPC

- **HPC** is part of **Systems** Community and includes parallel computing
- Recently most technical progress from **ML/AI** and **Big Data Systems**
- At IU, **Data Science** students emphasize ML over systems
- Applications are **Cloud**, **Fog**, **Edge** systems
- Any real Big Data or Edge application needs **High Performance Big Data computing** with systems and ML/AI expertise
  - **Distributed** big data management (not AI) maybe doesn't need HPC
  - **HPC** and **Cyberinfrastructure** are critical technologies for analytics/AI but mature so innovation and h-index not so high

Aligning with Industry

- Clouds dominate
- HPC offered by Public Clouds
- MLPerf
- Global AI Supercomputer
- Intelligent Cloud and Intelligent edge

Dominance of Cloud Computing

- **94 percent** of workloads and compute instances will be processed by **cloud data centers** (22% CAGR) by 2021-- only six percent will be processed by traditional data centers (-5% CAGR).
- **Hyperscale data centers** will grow from 338 in number at the end of 2016 to 628 by 2021. They will represent 53 percent of all installed data center servers by 2021. They form a distributed Compute (on data) grid with some 50 million servers
- Analysis from CISCO https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html updated November 2018

Number of instances per server

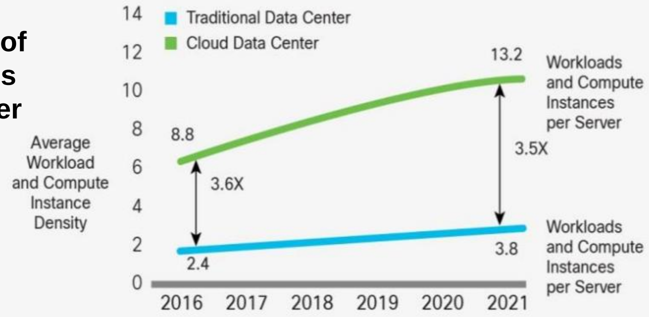

Source: Cisco Global Cloud Index, 2016-2021.

Number of Cloud Data Centers

Source: Cisco Global Cloud Index, 2016-2021.

Number of Public or Private Cloud Data Center Instances

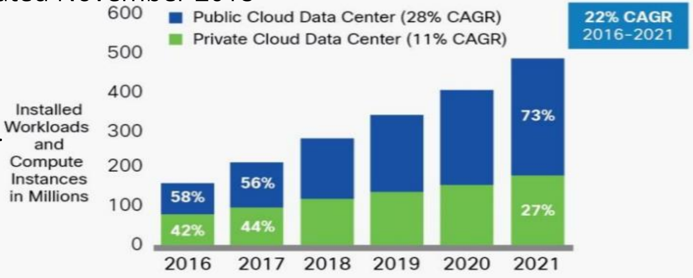

Source: Cisco Global Cloud Index, 2016-2021.

HPC is available on Public Clouds

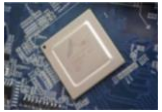

AWS runs everyday HPC for logistics, ML, Data Center, Consumer Product design, Robotics, Semiconductor design, Retail and Financial analytics

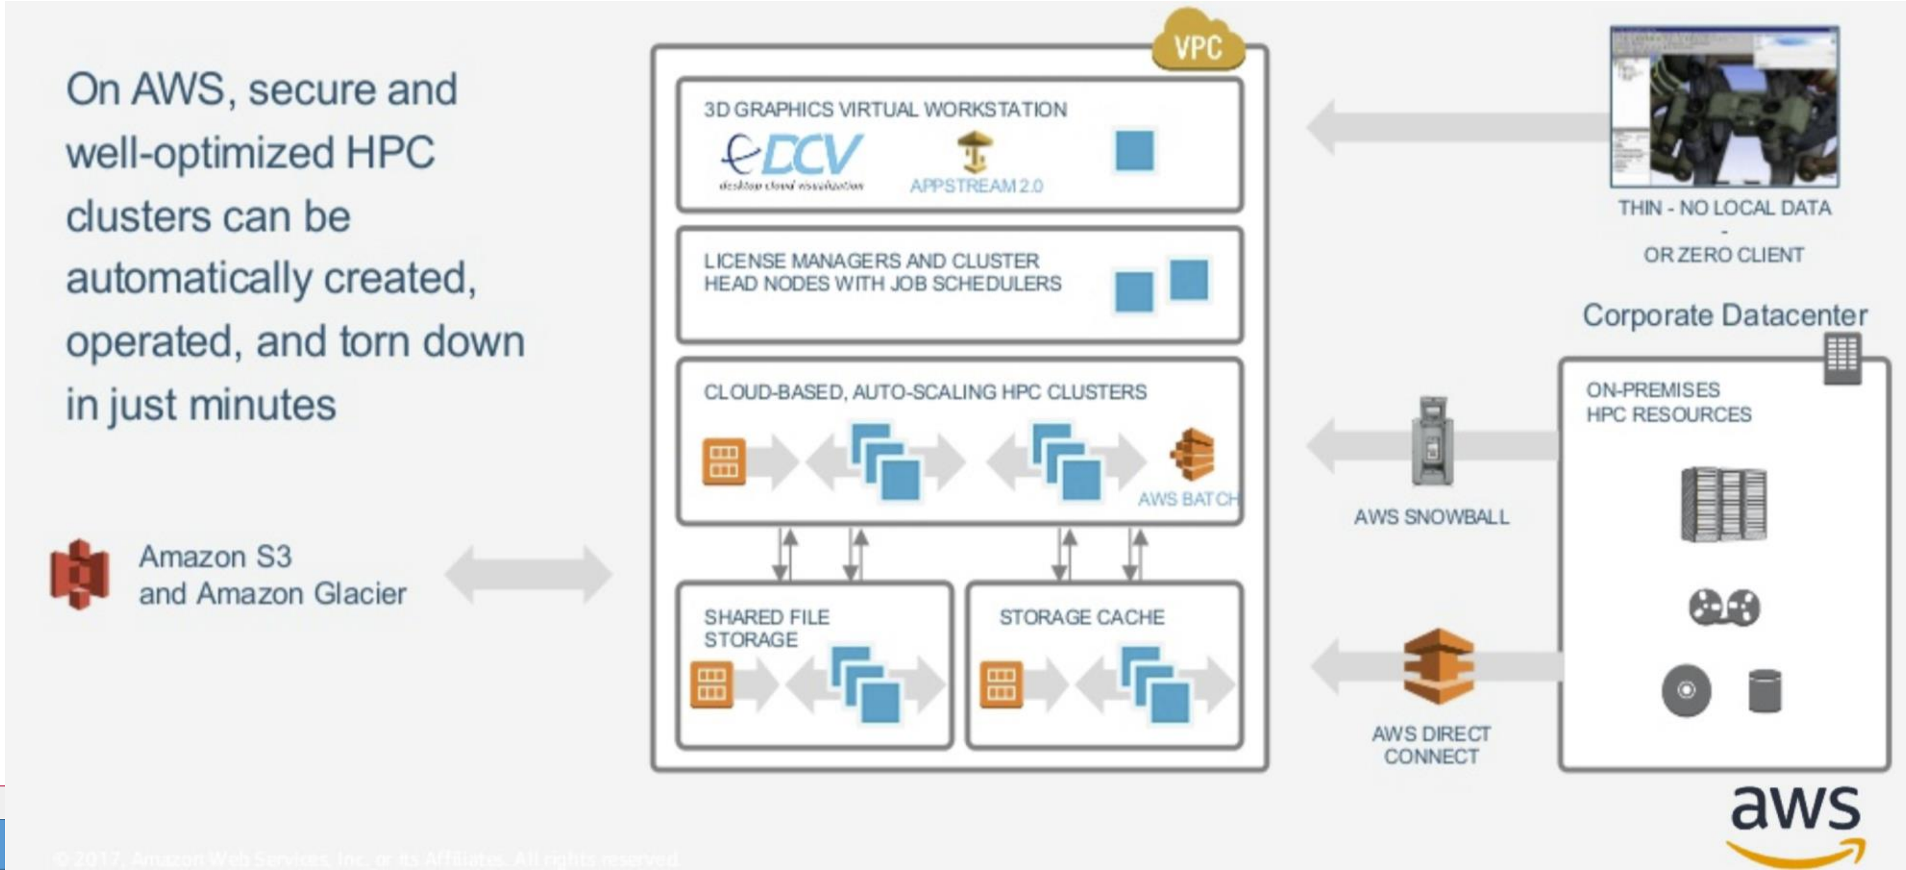

MLPerf

A broad ML benchmark suite for measuring performance of ML software frameworks, ML hardware accelerators, and ML cloud platforms.

MLPerf's mission is to build fair and useful benchmarks for measuring training and inference performance of ML hardware, software, and services. MLPerf was founded in February, 2018 as a collaboration of companies and researchers from educational institutions. MLPerf is presently led by volunteer working group chairs. MLPerf could not exist without open source code and publically available datasets others have generously contributed to the community.

What's New

- 6/24/19: MLPerf Inference v0.5 launched. Results due 9/6.
- 2/14/19: MLPerf Training v0.6 launched. Results due 5/24.
- 12/12/18: MLPerf Training v0.5 results are available.
- 5/2/18: MLPerf Training v0.5 launched. Results due 11/9.

Get Involved

- Join the forum
- Join working groups
- Attend community meetings
- Make your organization an official supporter of MLPerf
- Ask questions, or raise issues

Supporting companies

Alibaba, AMD, Arm, Baidu, Cadence, Cerebras, Cisco

Cray, Dividiti, Enflame Tech, Esperanto, Facebook, Google, Groq

Habana, Huawei, Iluvatar, Intel, In-Q-Tel, MediaTek, Mentor Graphics

Microsoft, Myrtle, Mythic, NetApp, NVIDIA, One Convergence, Qualcomm

Rpa2ai, Sambanova, Samsung S.LSI, Sigopt, Synopsys, Tensyr, Teradyne

VMware, Wave Computing, WekaIO, Xilinx

Note Industry Dominance

Contributions by researchers from

Performance of Time Series Machine Learning Algorithms (MLPerf)

Areas	Applications	Model	Data sets	Papers
Transportation 	Cars, Taxis, Freeway Detectors	TT-RNN , BNN, LSTM	Caltrans highway traffic [1], Taxi/Uber trips [2-5]	[6-8]
Medical 	Wearables, Medical instruments: EEG, ECG, ERP, Patient Data	LSTM, RNN	OPPORTUNITY [9-10], EEG [11-14], MIMIC [15]	[16-20]
Cybersecurity 	Intrusion, classify traffic, anomaly detection	LSTM	GPL loop dataset [21], SherLock [22]	[21, 23-25]
General Social Statistics 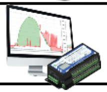	Household electric use, Economic, Finance, Demographics, Industry	CNN, RNN	Household electric [26], M4Competition [27],	[28-29]
Finance 	Stock Prices versus time	CNN, RNN	Available academically from Wharton [30]	[31]
Science 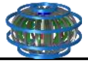 	Climate, Tokamak	Markov, RNN	USHCN climate [32]	[33-35]
Software Systems 	Events	LSTM	Enterprise SW system [36]	[36-37]
Language and Translation 	Pre-trained Data	Transformer [38]	[39-40]	[41-42]Mesh Tensorflow
Google Speech	All-Neural On-Device Speech Recognizer	RNN-T		[43]
IndyCar Racing 	Real-time car and track detectors	HTM		[44]
Social media 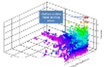	Twitter	Online Clustering	Available from Twitter	[45-46]

Microsoft Summer 2018: Global AI Supercomputer

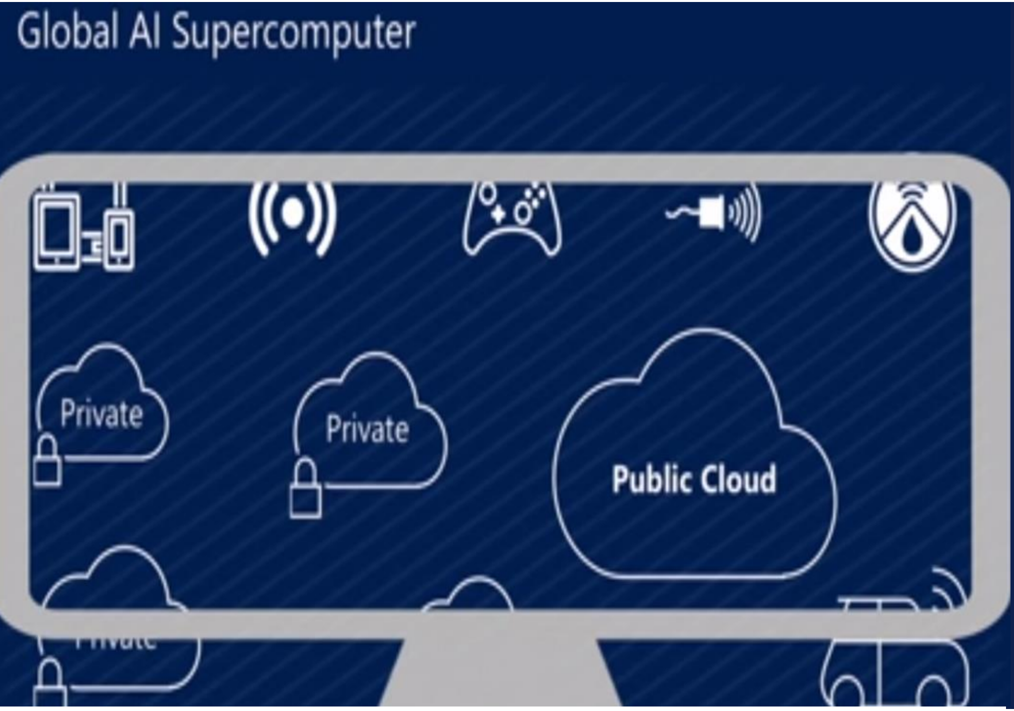

Intelligent Edge & Intelligent Cloud

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

Overall Global AI and Modeling Supercomputer GAIMSC Architecture

- **Global** says we are all involved - it is an HPDC system
- I added “Modeling” to get the Global AI and **Modeling** Supercomputer GAIMSC
- There is only a cloud at the logical center but it’s physically distributed and domated by a few major players
- Modeling was meant to include classic simulation oriented supercomputers
- Even in Big Data, one needs to build a model for the machine learning to use
- GAIMSC will use classic HPC for data analytics which has similarities to big simulations (**HPCforML**)
- GAIMSC must also support I/O centric data management with Hadoop etc.
- Nature of I/O subsystem controversial for such **HPC clouds**
    - Lustre v. HDFS; importance of SSD and NVMe;
- HPC Clouds would suggest that MPI runs well on Mesos and Kubernetes and with Java and Python

HPCforML and MLforHPC

- **Technical aspects of converging HPC and Machine Learning**
- **HPCforML**
  - Parallel high performance ML algorithms
  - High Performance Spark, Hadoop, Storm
- **9 scenarios for MLforHPC**
  - Illustrate 3 scenarios
  - Transform Computational Biology
  - Research Issues

Dean at NeurIPS

DECEMBER 2017

- ML for optimizing parallel computing (load balancing)
- Learned Index Structure
- ML for Data-center Efficiency
- ML to replace heuristics and user choices (Autotuning)

Machine Learning for Systems
and
Systems for Machine Learning

Jeff Dean
Google Brain team
g.co/brain

Presenting the work of **many** people at Google

Implications of Machine Learning for Systems and Systems for Machine Learning

- We could replace “Systems” by “Cyberinfrastructure” or by “HPC” and/or “HPDC”
- I use HPC as we are aiming at systems that support big data or big simulations and almost by (my) definition could naturally involve HPC.
- So we get **ML for HPC** and **HPC for ML**
- **HPC for ML** is very important but has been quite well studied and understood
  - It makes data analytics run much faster
- **ML for HPC** is transformative both as a technology and for application progress enabled
  - If it is ML for HPC running ML, then we have the creepy situation of the AI supercomputer improving itself
  - Microsoft 2018 faculty summit discussed ML to improve Big Data systems e.g. configure database system.

Big Data and Simulation Comparison of Difficulty in Parallelism

Parallel Big Data Algorithms many issues in common with Parallel Simulations

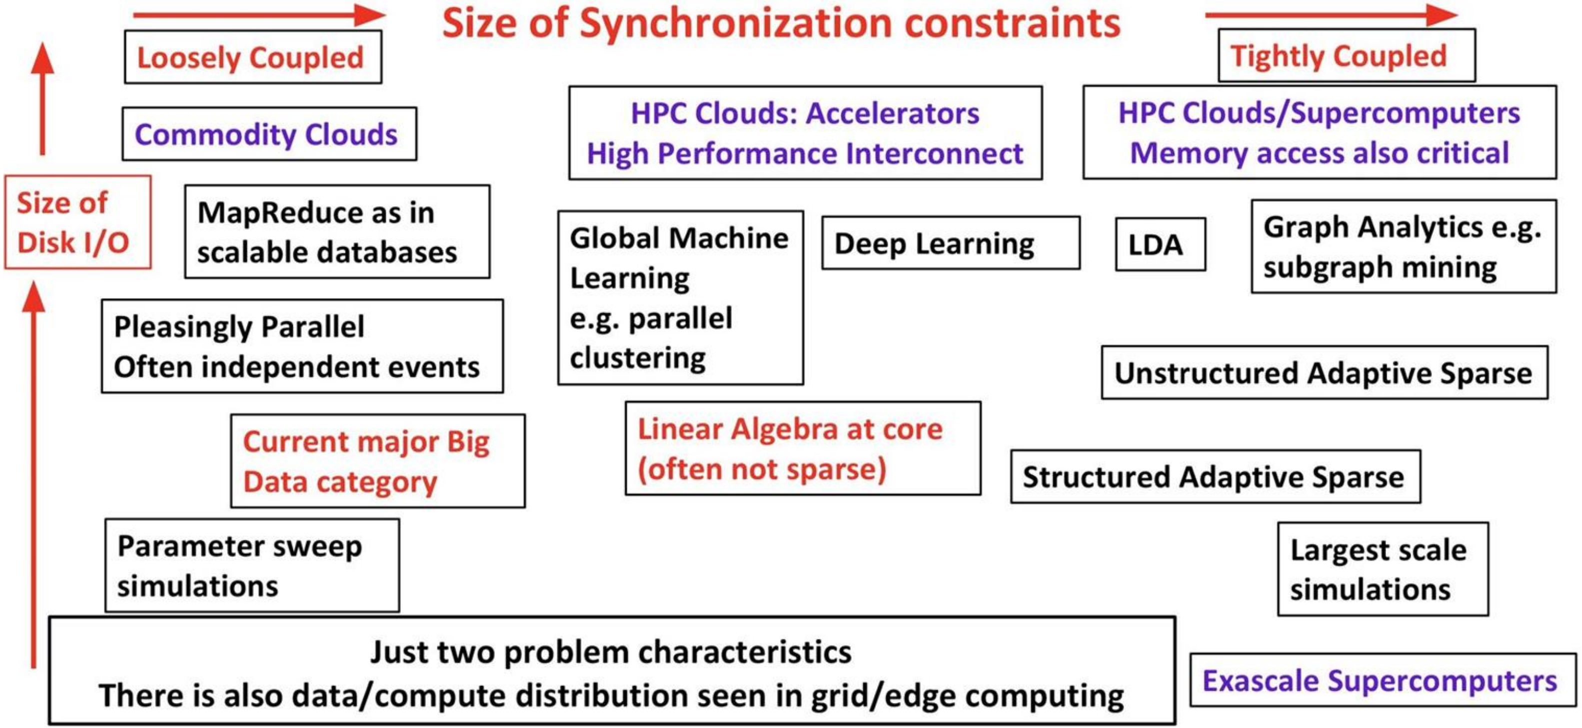

MLaroundHPC
MLAutotuning

MLforHPC (ML for Systems) in detail

- **MLforHPC** can be further subdivided into several categories:
  - **MLafterHPC**: ML analyzing results of HPC as in trajectory analysis and structure identification in biomolecular simulations. Well established and successful
  - **MLControl**: Using simulations (with HPC) and ML in control of experiments and in objective driven computational campaigns. Here simulation surrogates are very valuable to allow real-time predictions.
  - **MLAutotuning**: Using ML to configure (autotune) ML or HPC simulations.
  - **MLaroundHPC**: Using ML to learn from simulations and produce learned surrogates for the simulations or parts of simulations. The same ML wrapper can also learn configurations as well as results. **Most Important.**

3 MLAutotuning 6 MLaroundHPC Scenarios

INPUT

OUTPUT

1. MLAutotuningHPC – Learn configurations

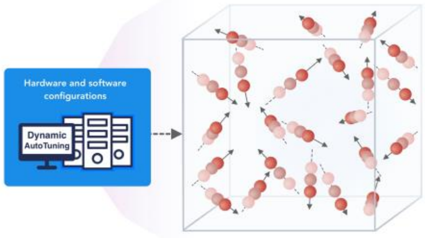

2. MLAutotuningHPC – Smart ensembles

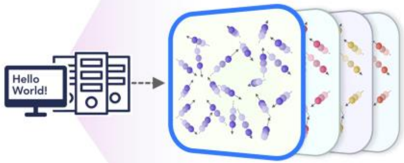

3. MLAutotuningHPC – Learn models from data

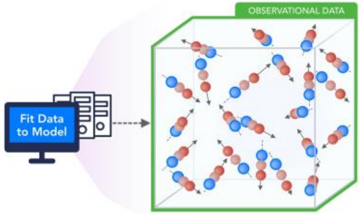

4. MLaroundHPC: Learning Outputs from Inputs (parameters)

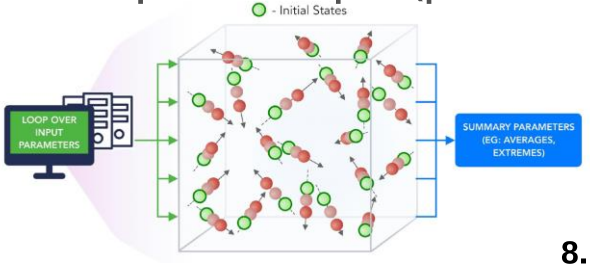

5. MLaroundHPC: Learning Outputs from Inputs (fields)

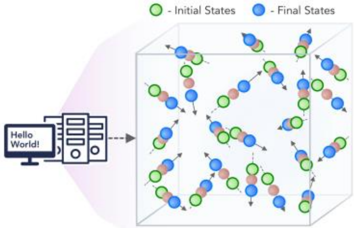

6. MLaroundHPC: Learning Model Details (agents)

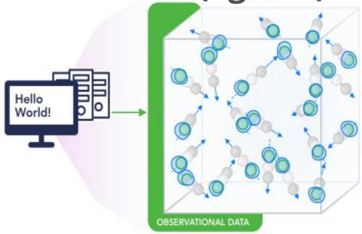

7. MLaroundHPC: Learning Model Details (effective potentials)

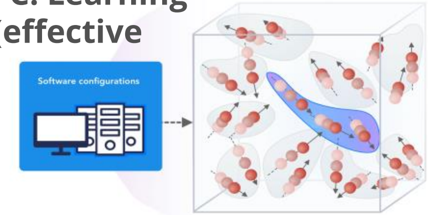

8. MLaroundHPC: Learning Model Details (ML based data assimilation)

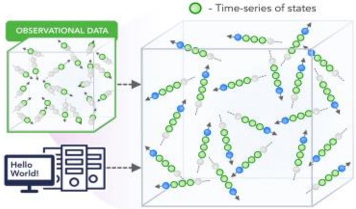

9. Take scenario 8 and use to infer improved model structure

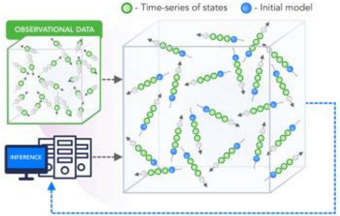

MLAutoTuningHPC: Learning Configurations

- This is classic Autotuning and one optimizes some mix of performance and quality of results with the learning network inputting the configuration parameters of the computation.
- This includes initial values and also dynamic choices such as block sizes for cache use, variable step sizes in space and time.
- It can also include discrete choices as to the type of solver to be used.

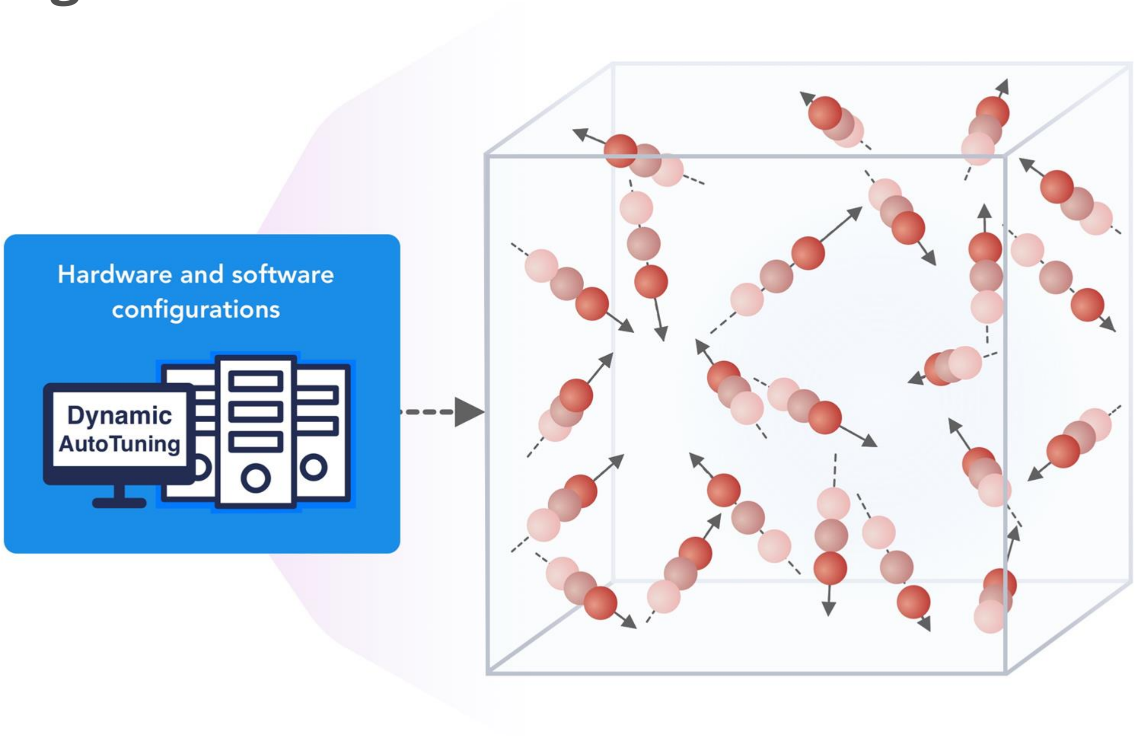

MLAutoTuningHPC: Learning Configurations

MLAutoTuningHPC: Smart Ensembles

- Here we choose the best set of parameters to achieve some computation goal
- Such as providing the most efficient training set with defining parameters spread well over the relevant phase space.

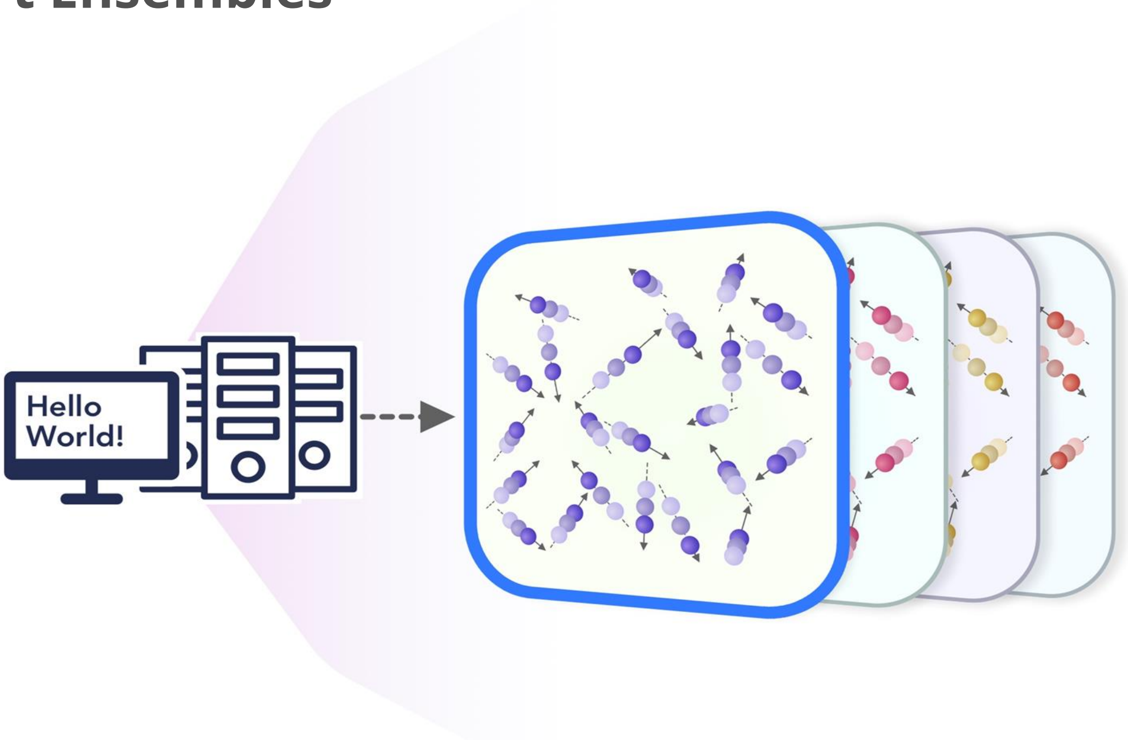

Smart Ensembles

MLforHPC Simulation Surrogates

MLaroundHPC: Learning Outputs from Inputs:

a) Computation Results from Computation defining Parameters

- Here one just feeds in a modest number of meta-parameters that the define the problem and learn a modest number of calculated answers.
- This presumably requires fewer training samples than "fields from fields" and is main use so far

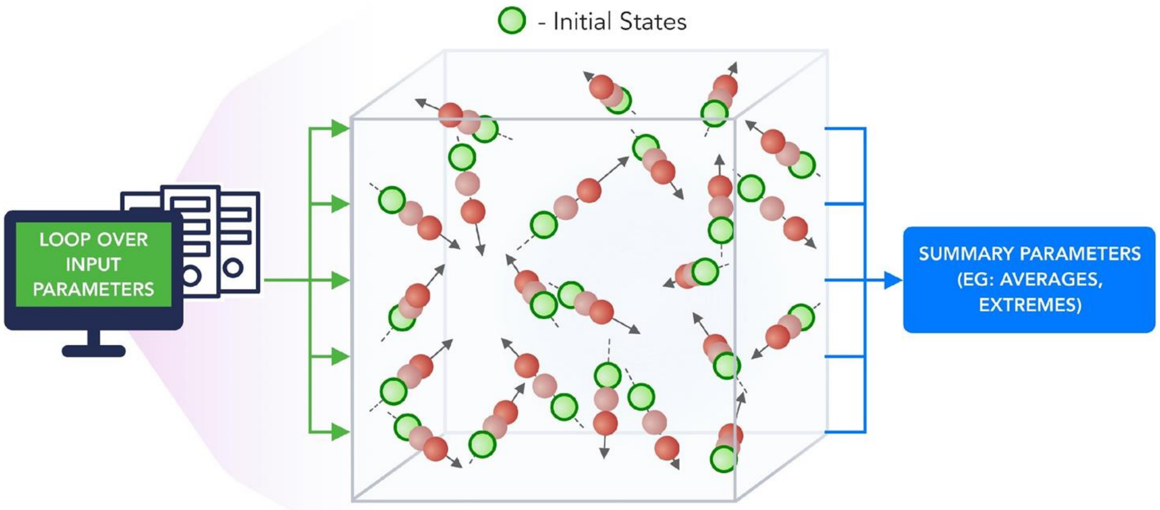

Learning Outputs from Inputs: Computation Results from Computation defining Parameters

Operationally same as **SimulationTrainedML** but with a different goal: In **SimulationTrainedML** the simulations are performed to directly train an AI system rather than the AI system being added to learn a simulation.

MLaroundCI: Machine learning for performance enhancement with Surrogates of molecular dynamics simulations

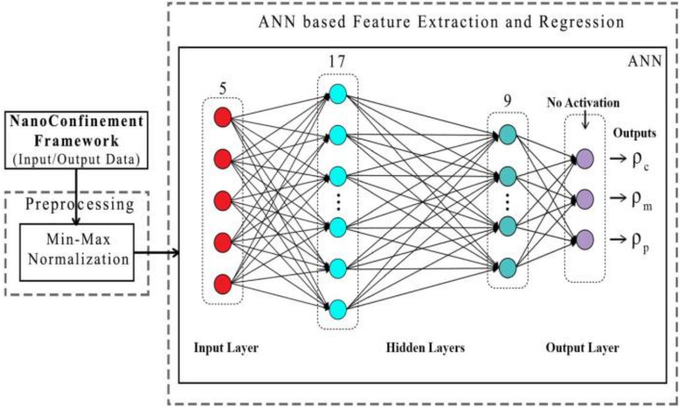

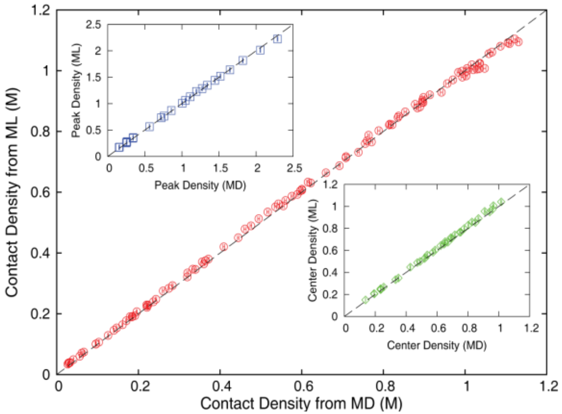

- Employed to extract the ionic structure in electrolyte solutions confined by planar and spherical surfaces.
- Written with C++ and accelerated with hybrid MPI-OpenMP.
- MLaroundHPC successfully learns desired features associated with the output ionic density that are in excellent agreement with the results from explicit molecular dynamics simulations.
- Speed up 10^5

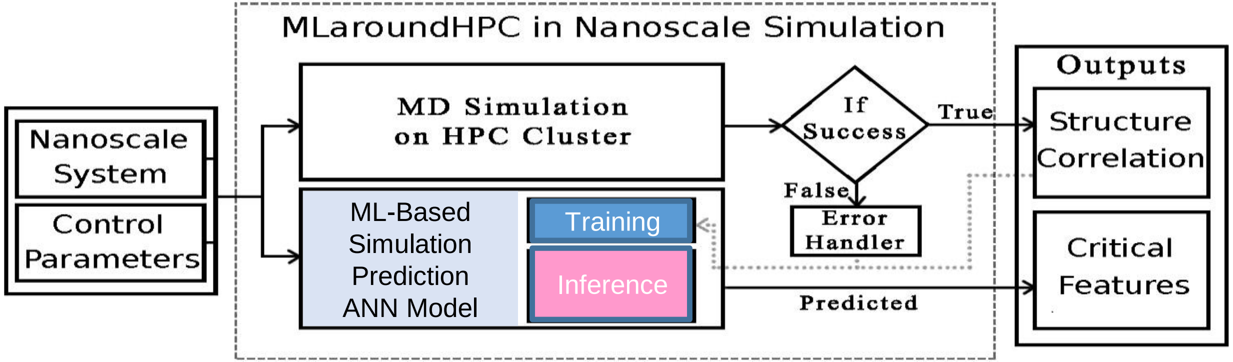

Speedup of MLaroundHPC

- T_{seq} is sequential time
- T_{train} time for a (parallel) simulation used in training ML
- T_{learn} is time per point to run machine learning
- T_{lookup} is time to run inference per instance
- N_{train} number of training samples
- N_{lookup} number of results looked up

N_{train} is 7K to 16K in our work

$$\text{Effective Speedup } S = \frac{T_{seq}(N_{lookup} + N_{train})}{T_{lookup}N_{lookup} + (T_{train} + T_{learn})N_{train}}$$

- Becomes T_{seq}/T_{train} if ML not used
- Becomes T_{seq}/T_{lookup} (10^5 faster in our case) if inference dominates (will overcome end of Moore's law and win the race to **zettascale)**
- Another factor as inferences uses one core; parallel simulation 128 cores
- **Strong scaling as no need to parallelize on more than effective number of nodes**

MLaroundHPC: Learning Model Details

a) Learning Agent Behavior One has a model such as a set of cells as agents modeling a virtual tissue. One can use ML to learn dynamics of cells replacing detailed computations by ML surrogates.

- As can be millions to billions of such agents the performance gain can be huge as each agent uses same learned model..
- This is MLaroundHPC for cells but MLAutotuning for multi-cell (tissue) phase

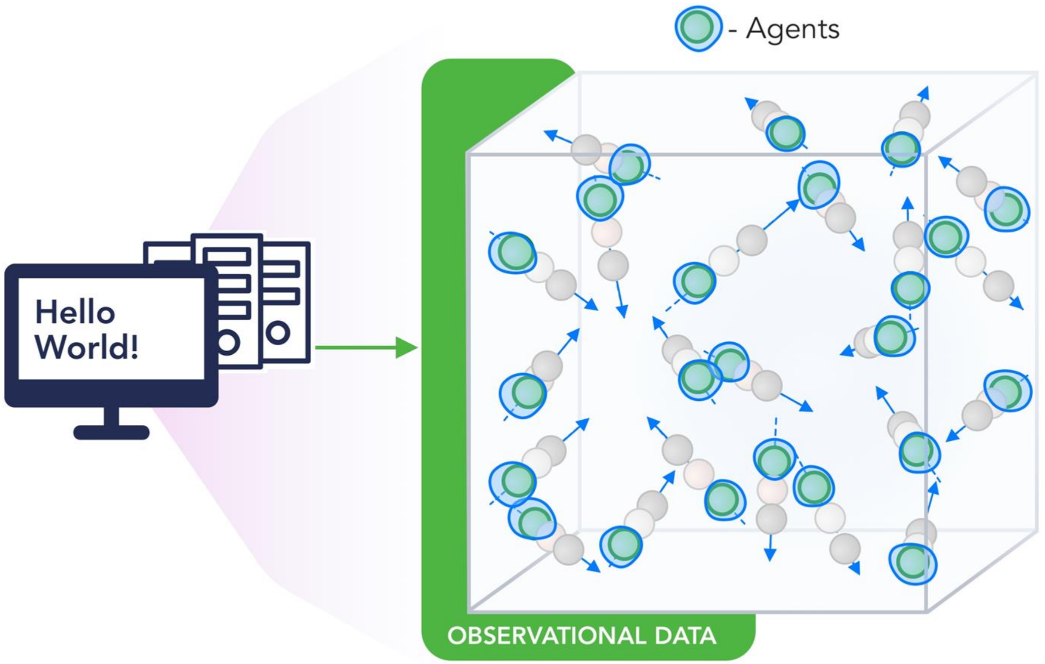

Learning Model Details: Learning Agent Behaviour

MLaroundHPC: Learning Model Details ML for Data Assimilation

Take the case where we have "videos" recording observational data i.e. data is a high dimensional (spatial extent) time series

(c) Learning Agent Behavior – a Predictor-Corrector approach Here one time steps models and at each step optimize the parameters to minimize divergence between simulation and ground truth data.

- **Example:** produce a generic model organism such as an embryo. Take this generic model as a template and learn the different adjustments for particular individual organisms.
- **Build on Ride hailing work**
- Current state of the art expresses spatial structure as a convolutional neural net and time dependence as recurrent neural net (LSTM)

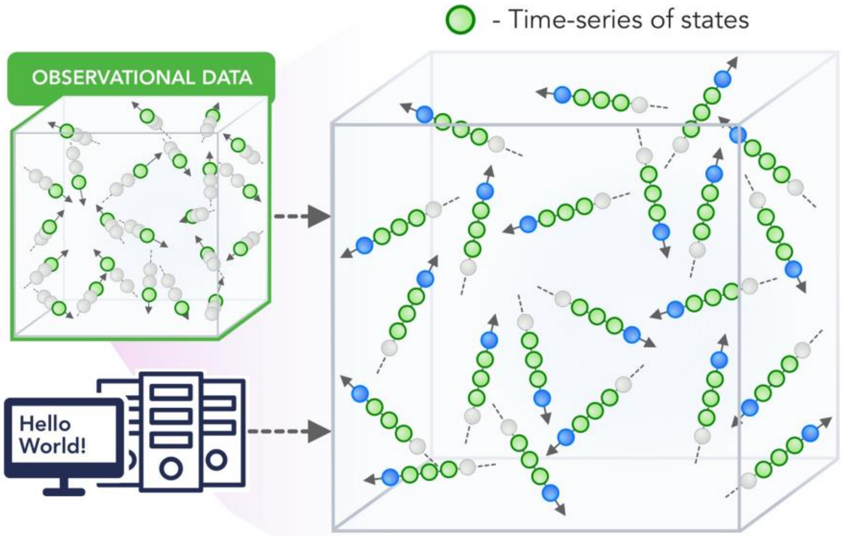

Learning Agent Behaviour: A Predictor-Corrector Approach

Simulating Biological Organisms (with James Glazier @IU)

Learning Agent Behavior

Replace components by learned surrogates
(Reaction Kinetics Coupled ODE's)

All steps use MLAutotuning

Smart Ensembles

Predictor-Corrector

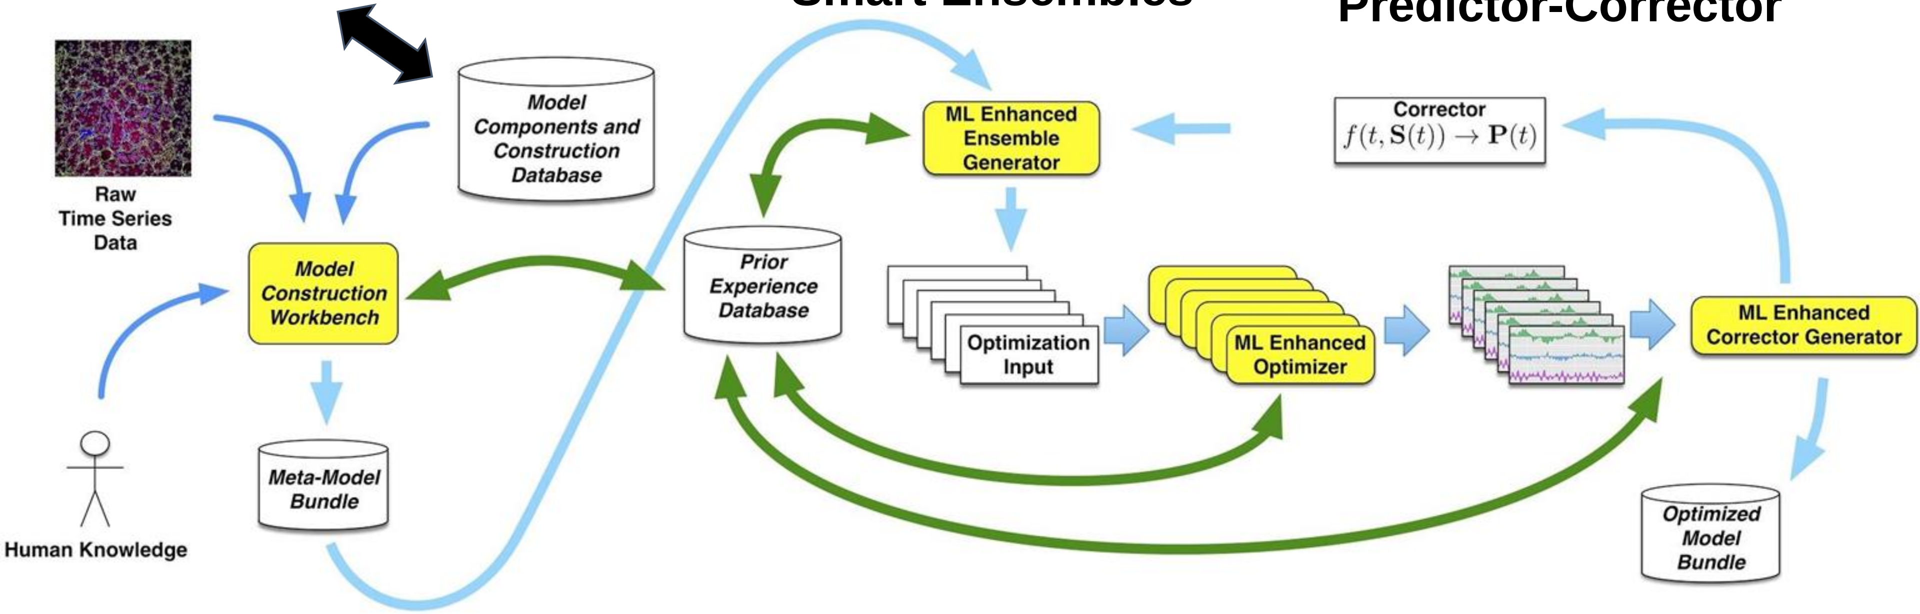

Challenges and Opportunities

Computer Science Issues

Computer Science Issues I

- Hundreds of Ph. D. theses!
- What computations can be assisted by what ML in which of 9 scenarios
  - What is performance of different DL (ML) choices in compute time and capability
  - What can we do with zettascale computing?
- Redesign all algorithms so they can be ML-assisted
- Dynamic interplay (of data and control) between simulations and ML seems likely but not clear at present
- There ought to be important analogies between time series and this area (as simulations are 4D time series?)
  - Exploit MLPerf examples?

Computer Science Issues II

- Little study of best ANN structure especially for hardest cases such as “predict fields from fields” and ML assisted data assimilation
  - Not known how large training set needs to be.
  - Most published data has quite small training sets
- I am surprised that there is not a rush to get effective performances on simulations of exascale and zettascale on current machines
- Interesting load balancing issues if in parallel case some points learnt using surrogates and some points calculated from scratch
- Little or no study of either predicting errors or in fact of getting floating point numbers from deep learning.

Conclusions

Conclusions

1

HPC is essential (**HPCforML**) although innovation in classic areas limited

- Everybody needs systems
- Need to align communities to ensure HPC importance recognized

2

Good to work closely with **industry**

- Student Internships, Collaborations such as Contributions to MLPerf

3

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

- Training on cloud; Inference and some training on the edge

4

MLforHPC very promising where we could aim at

- First Zettascale **effective** performance in next 2 years
- Hardware/Software aimed at **general ML assisted speedup of computation**
- **Your health can be engineered** with ML-assisted **personalized nanodevices designed** based on **the ML-assisted digital twin** of disease in your tissues