Science and Technology Facilities Council

Scientific Computing

ML Benchmarks for Scientific Applications

Tony Hey

With Contributions from:

*Jeyan Thiyagalingam, Kuangdai Leng, Sam Jackson,
Juri Papay, and Keith Butler*

SciML Group,
Rutherford Appleton Laboratory,
Science and Technology Funding Council (STFC)

Rutherford Appleton Laboratory (RAL)

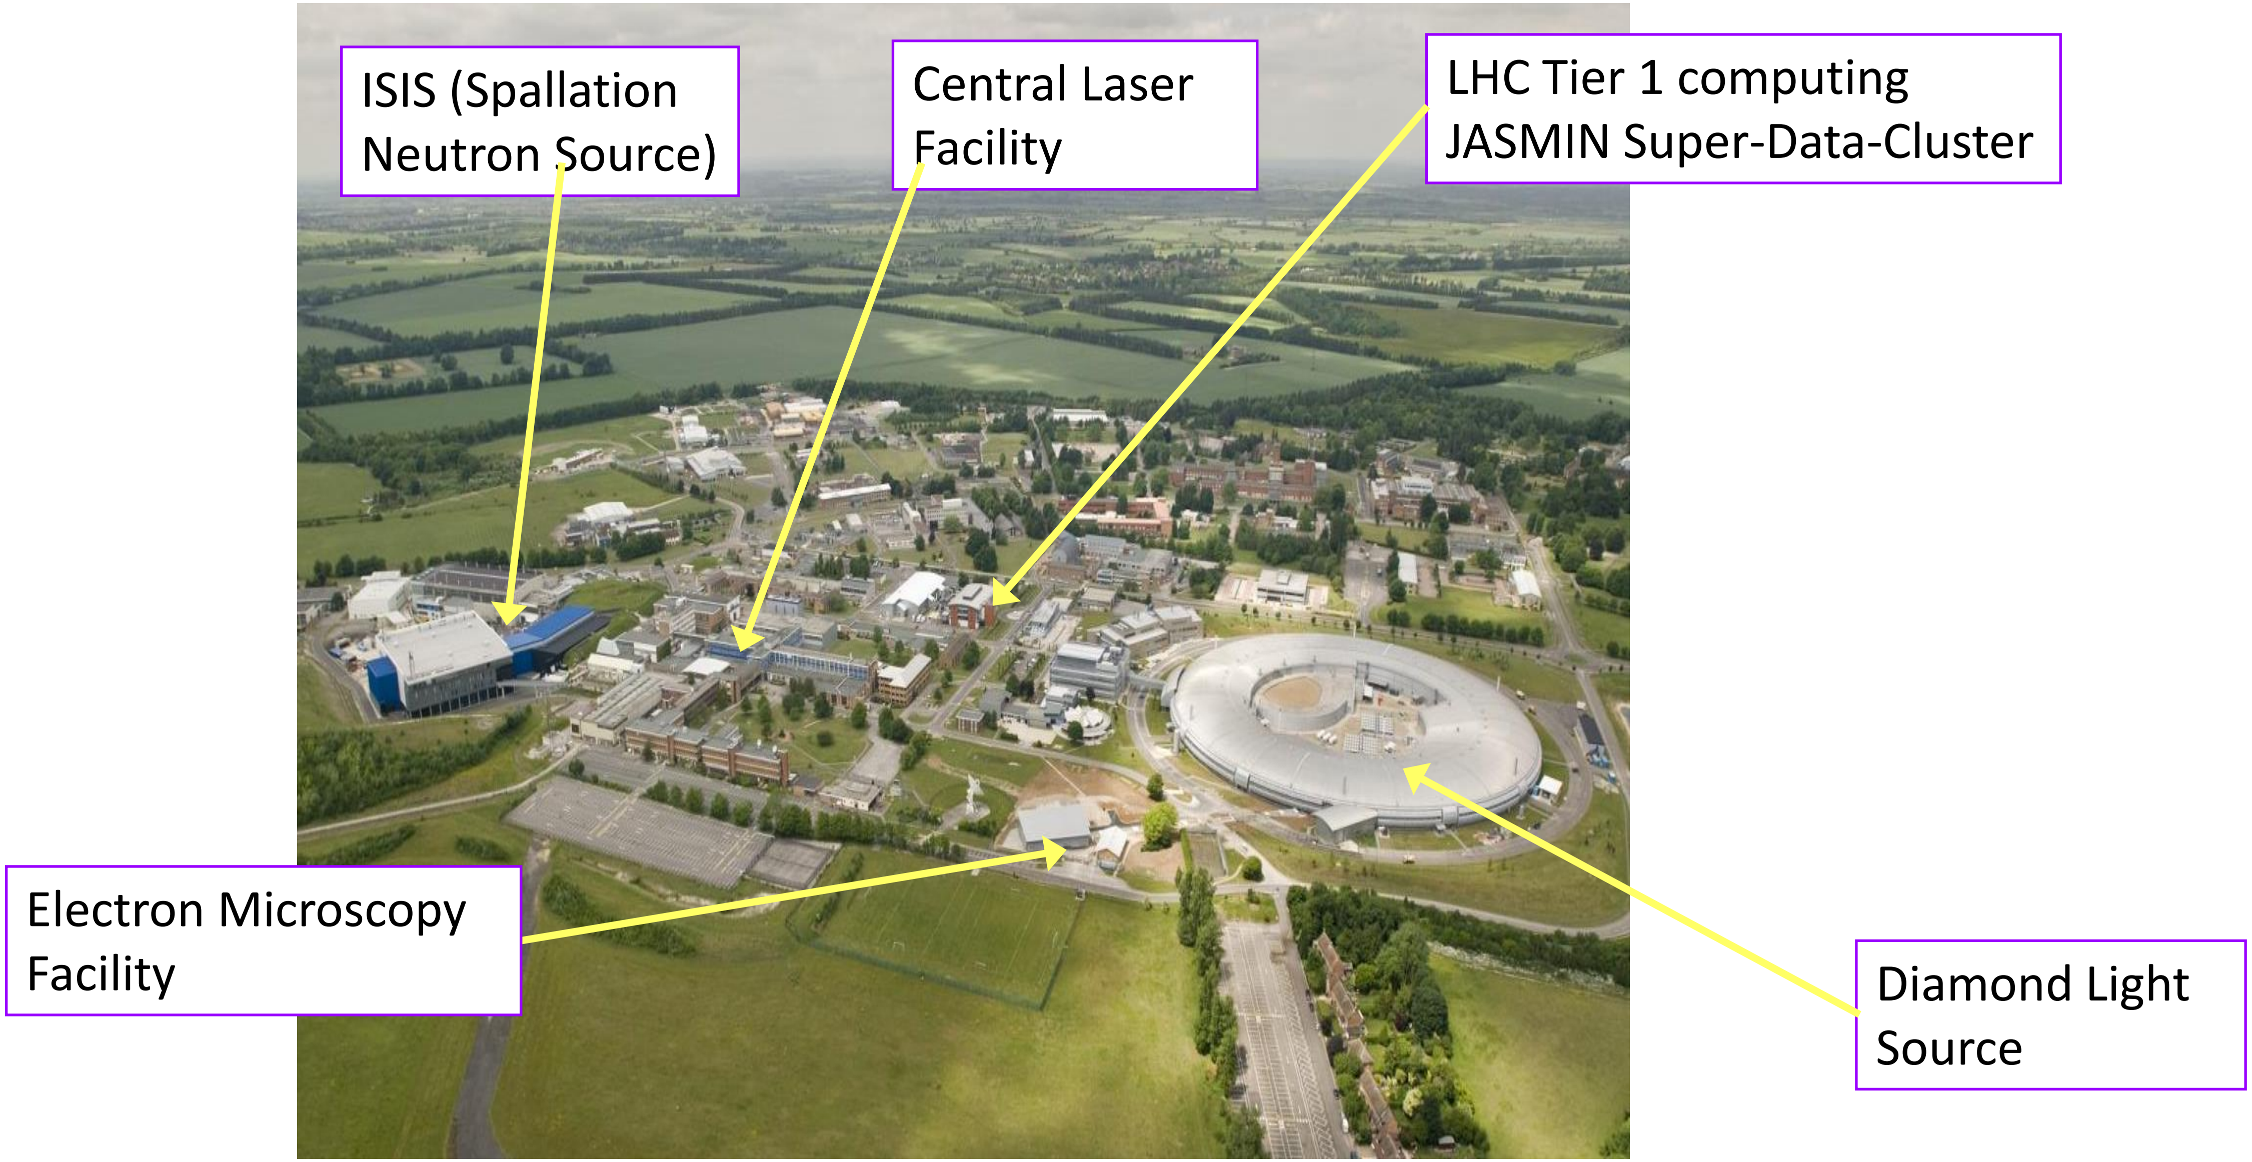

SciML Focus Areas

Research to Advance AI

- Use Science applications to improve AI

AI for Science

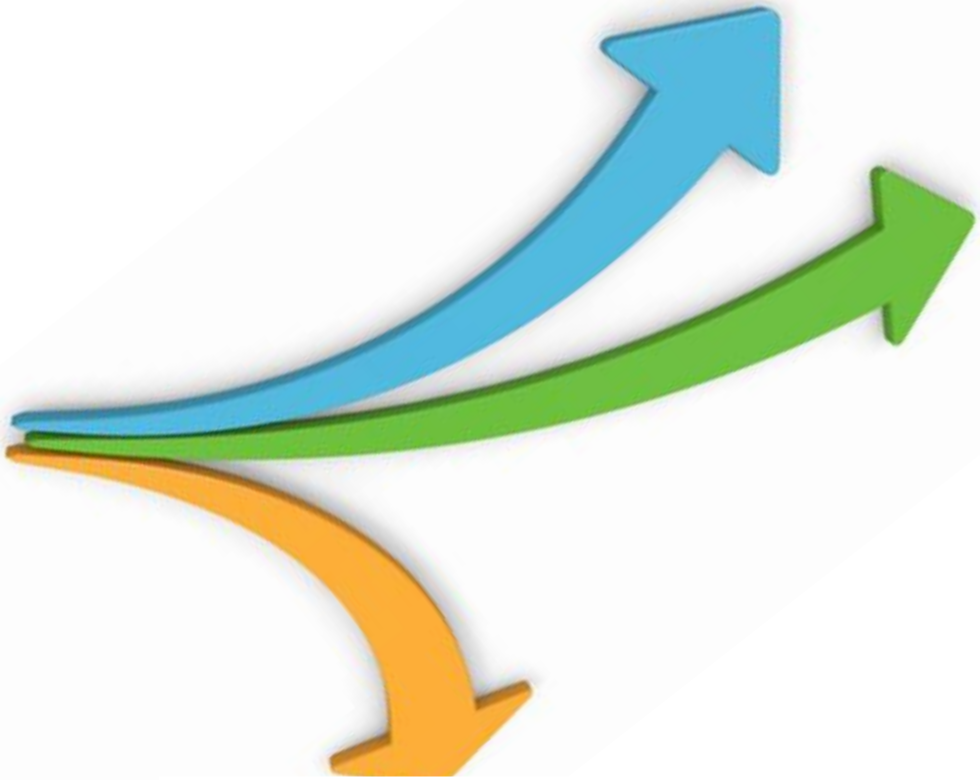

Use of AI in Science

- Use AI to understand experimental results (from facilities)

Smart Facilities

- Smart facilities can improve science – high quality data, faster results etc.
- About embedding AI at the heart of facilities operations

AI in Science: Current Landscape at RAL

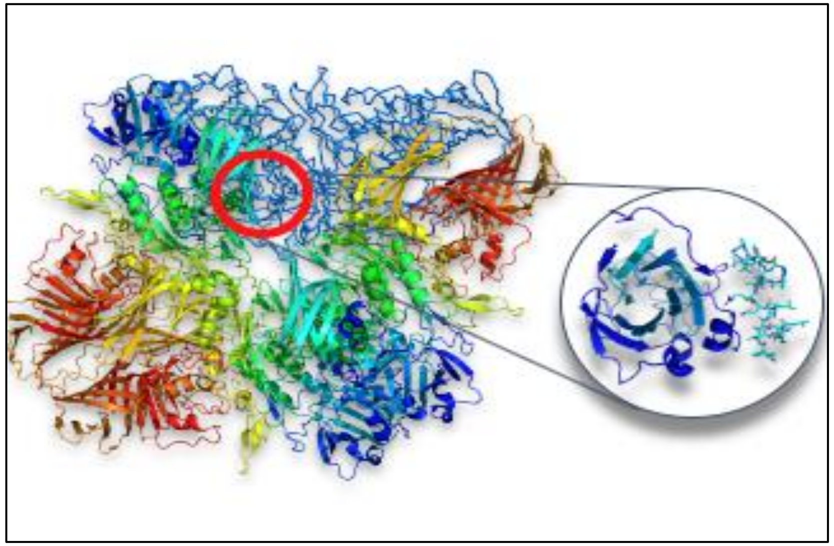

Cryo-EM

Cell Migration

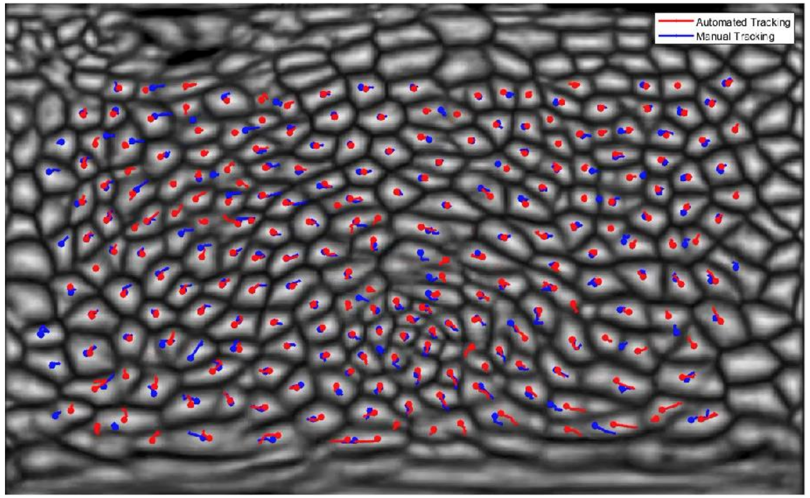

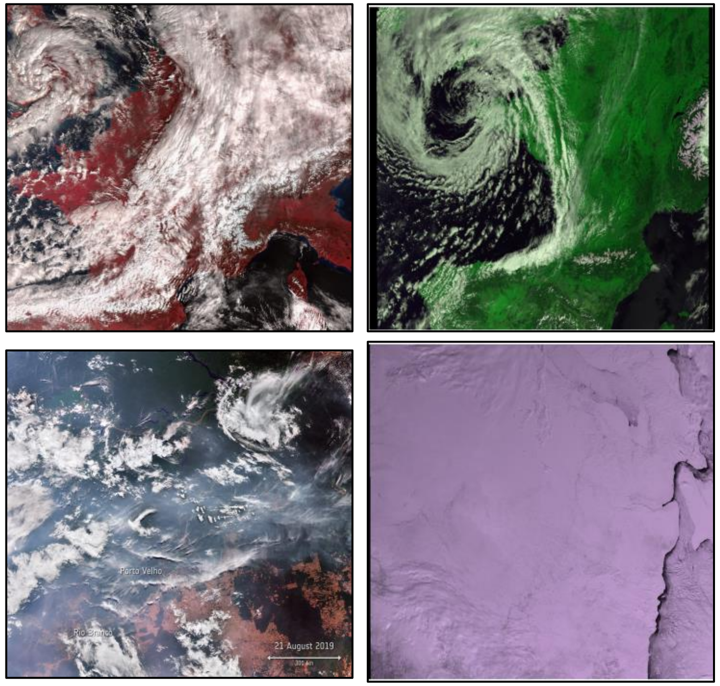

Cloud Masking

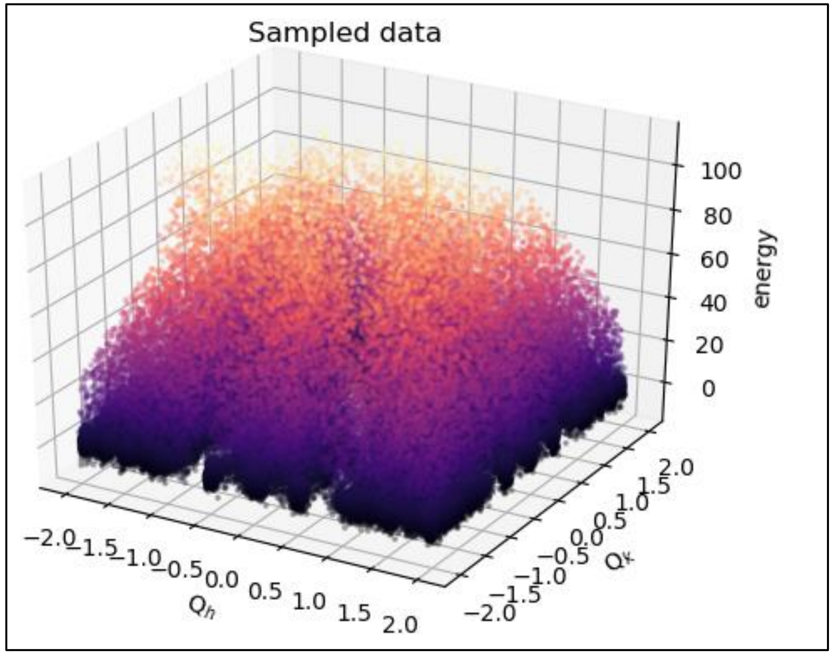

Inelastic Neutron Scattering

Denoising EM Data

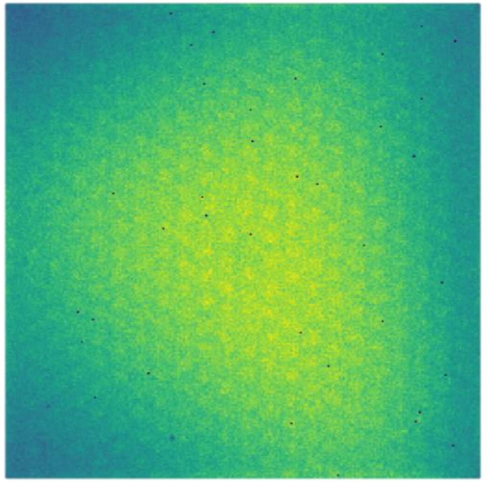

and …. many more

Smart Facilities: Current Landscape at RAL

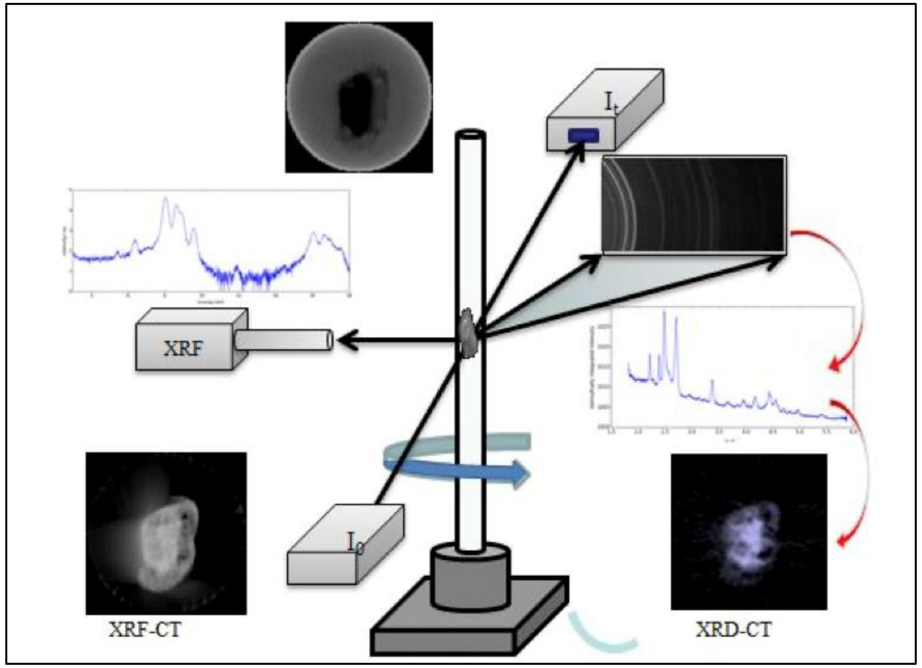

On-the-fly data-driven control of X-ray tomography

On-the-fly pulse shape discrimination

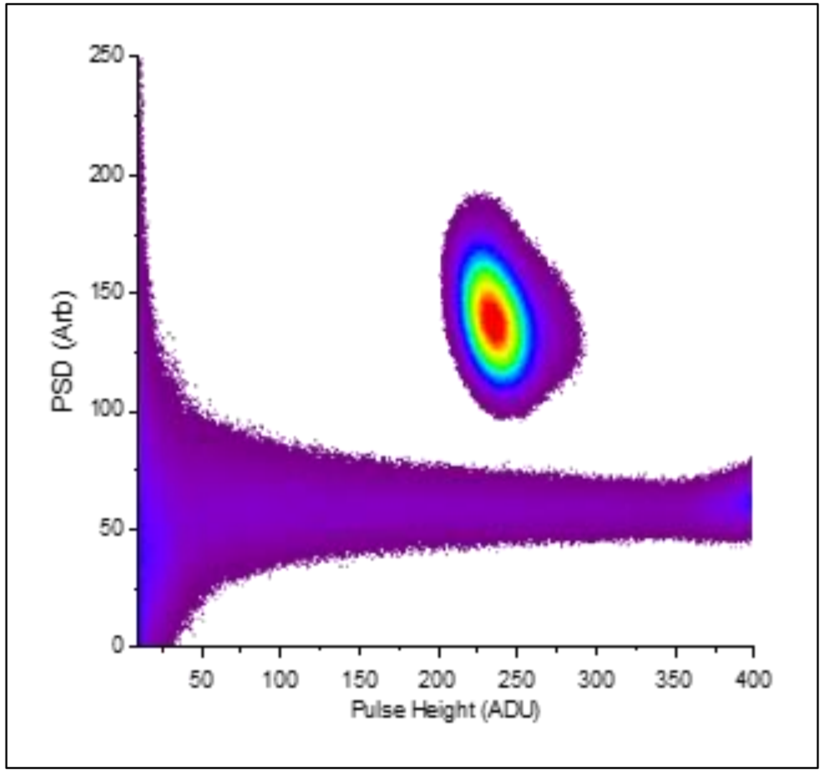

Assessing damages to optics in CLF

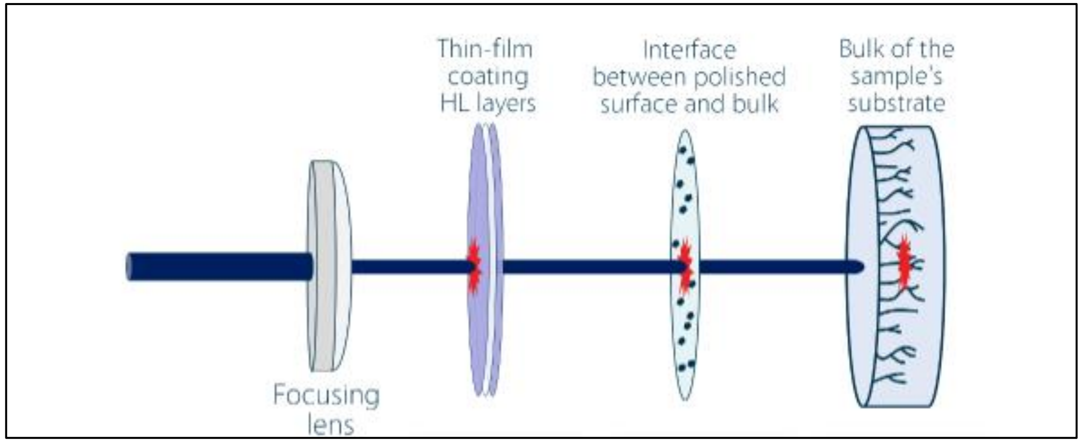

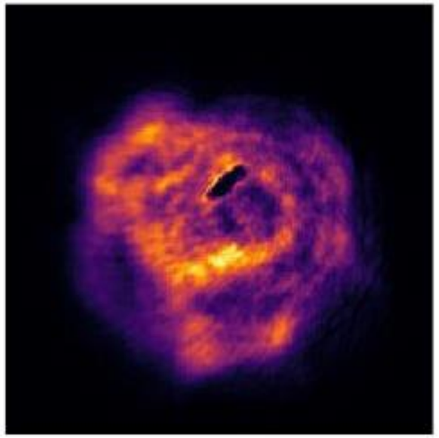

Credits: Sam Jackson – SciML

Mark Basham – RFI

Sion Richards – TECH

Rajeev Pattahil - CLF

Research to advance AI: Scientific ML Benchmarks

Many ML methods

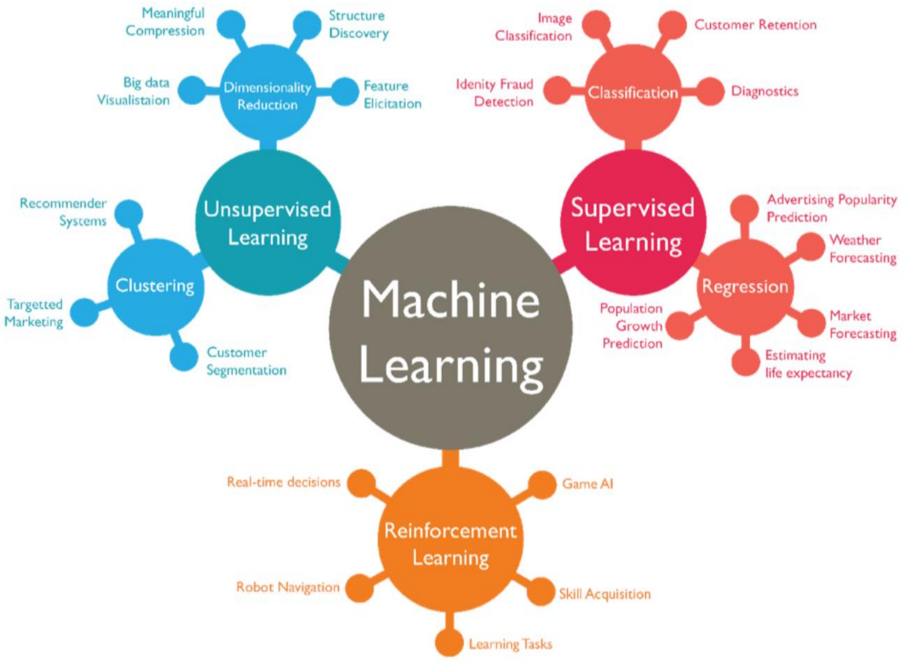

Many different hardware platforms

AI Labs.tw	Alibaba	AMD	Andes Technology	Aon Devices	Arm	Baidu
BAAI	Cadence	Calypso AI	Centaur Technology	Cerebras	Ceva	Cirrus
Cisco	Code Reef	Cray	CTuning Foundation	Dell	Dividiti	DDN Storage
Edgify	Enflame Tech	Esperanto	Facebook	FuriosaAI	Google	Groq

Many scientific domains

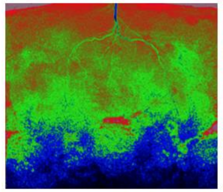
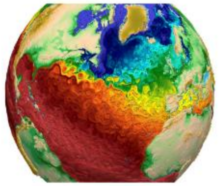
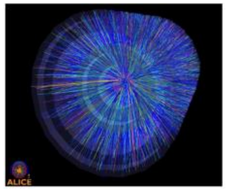
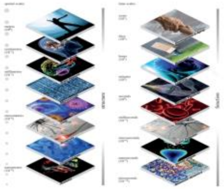
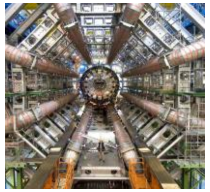


Our Domain Coverage

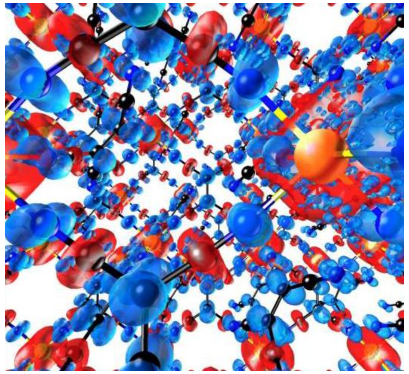

Material Sciences

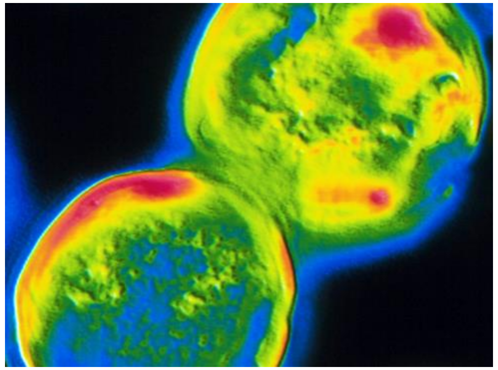

Environmental Sciences

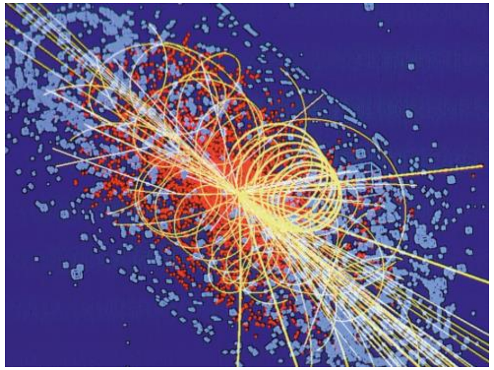

Particle Physics

Astronomy

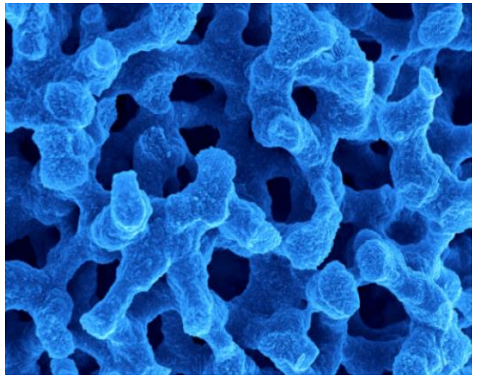

Life Sciences

Benefits to Science

- Baseline performance on a given problem sets a target
- Encourages the community to ‘develop’ better methods
- Availability of curated datasets increase the openness
- Fosters more efficient data-intensive scientific discovery

SciML Benchmarks

SciML-Bench: Table of Contents

GitHub - stfc-sciml/sciml-bench: SciML Benchmarking Suite for AI for Science

SciML Benchmark Suite

Benchmark Suite = Framework + Benchmarks + Datasets

Version 1.0 released with three initial benchmarks:

- **em_denoise** (Material Sciences)
  – 5GB dataset
- **dms_scatter** (Material Sciences)
  – 9GB dataset
- **slstr_cloud** (Environmental Sciences)
  – 187 GB and 2.6 TB datasets

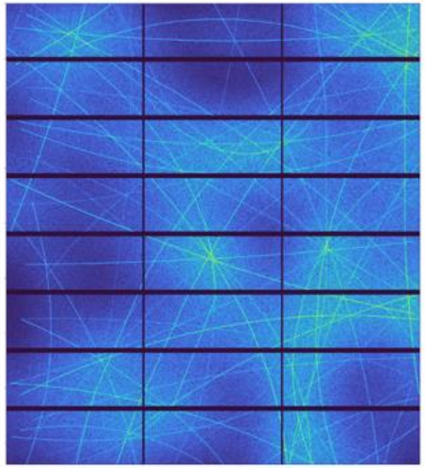

dms_scatter

Diffuse Multiple Scattering

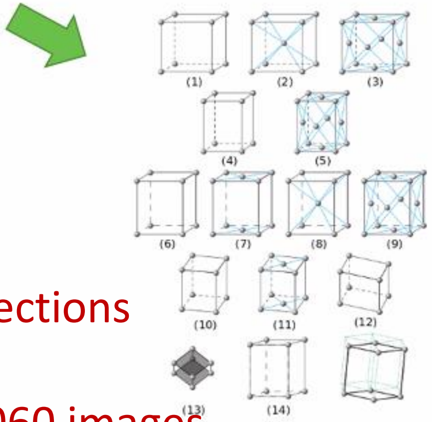

Identifying triple intersections in X-ray images

Dataset of 9GB with 8,060 images

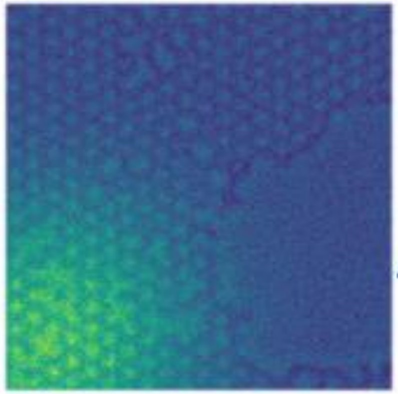

em_denoise

Removing Noise from Electron Microscopy Data

Removing noise from EM images

5GB dataset with 10,000 images of 256x256x1

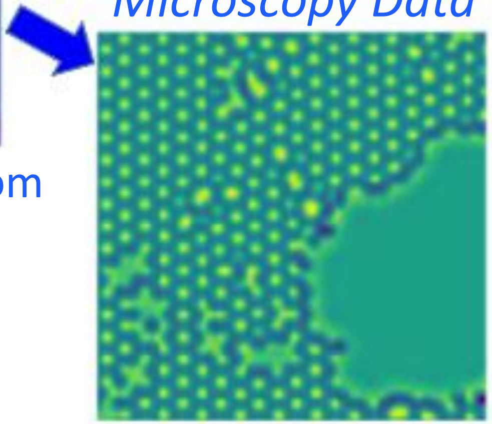

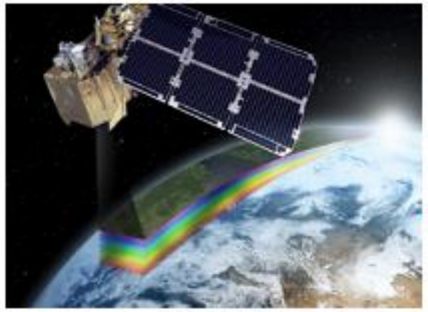

slstr_cloud

Cloud Masking

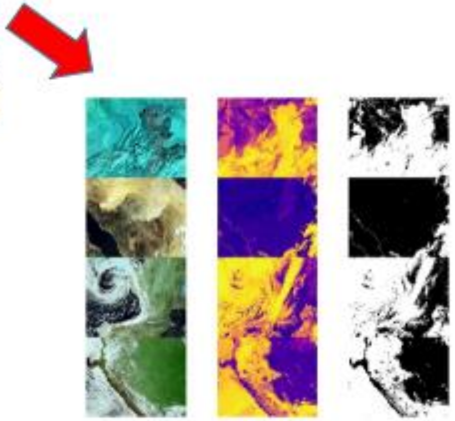

Identifying pixels that are cloud in satellite images

Benchmarking requirements

- Measure application and computer parameters
- Representative of real applications
- Reflect the interaction between the application and architecture
- Allow parallelisation and scalability studies
- Be easy to deploy and run

"Whatever we run should be simple to explain and implement"

Jack Dongarra

Science and Technology Facilities Council

Scientific Computing

SciML Benchmarks

The world of accelerators …

CPU

+ Large memory capacity
+ High clock frequency
+ Large caches (to mask latency)
+ Cores < 100
+ Optimised for serial computation
- Relatively low memory bandwidth
- Cache miss very costly
- Low performance/watt

GPU

+ High memory bandwidth
+ Relatively low clock frequency
+ Cores > 5k
+ Optimised for parallel computation
+ High performance/watt
- Low memory capacity
- Low per-thread performance

Science and Technology Facilities Council
Scientific Computing

SciML Benchmarks

… and two computing worlds to manage

- Copy data from CPU to GPU
- Copy code (kernel) from CPU to GPU
- Launch kernel on GPU
- Copy results from GPU to CPU

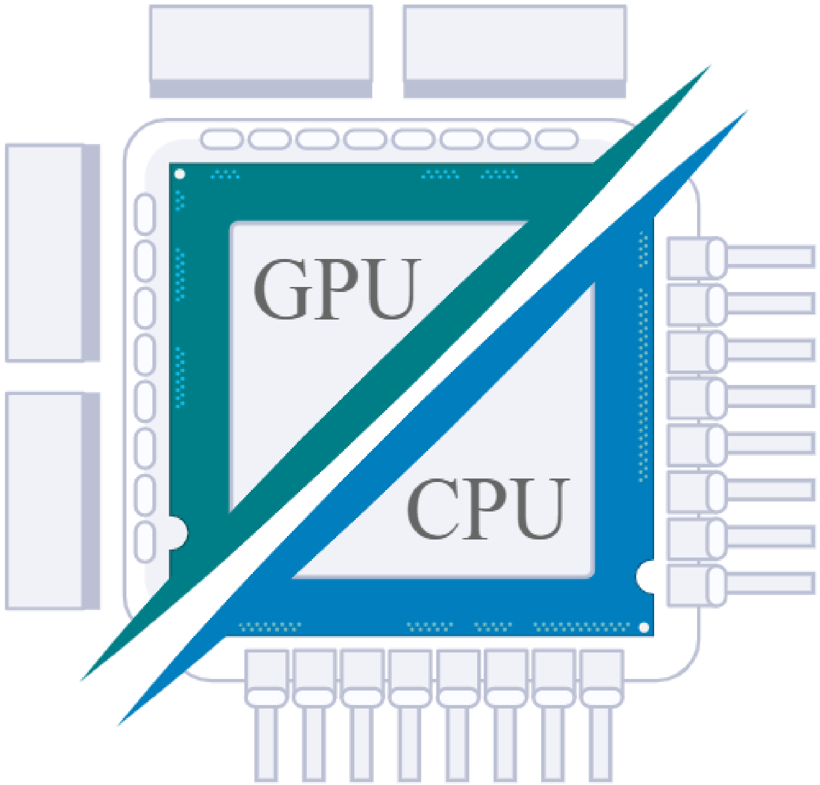

How to accelerate:

- Libraries
- Compiler annotations
- Programming language

GPU vs CPU: Diffuse Multiple Scattering benchmark (dms_scatter)

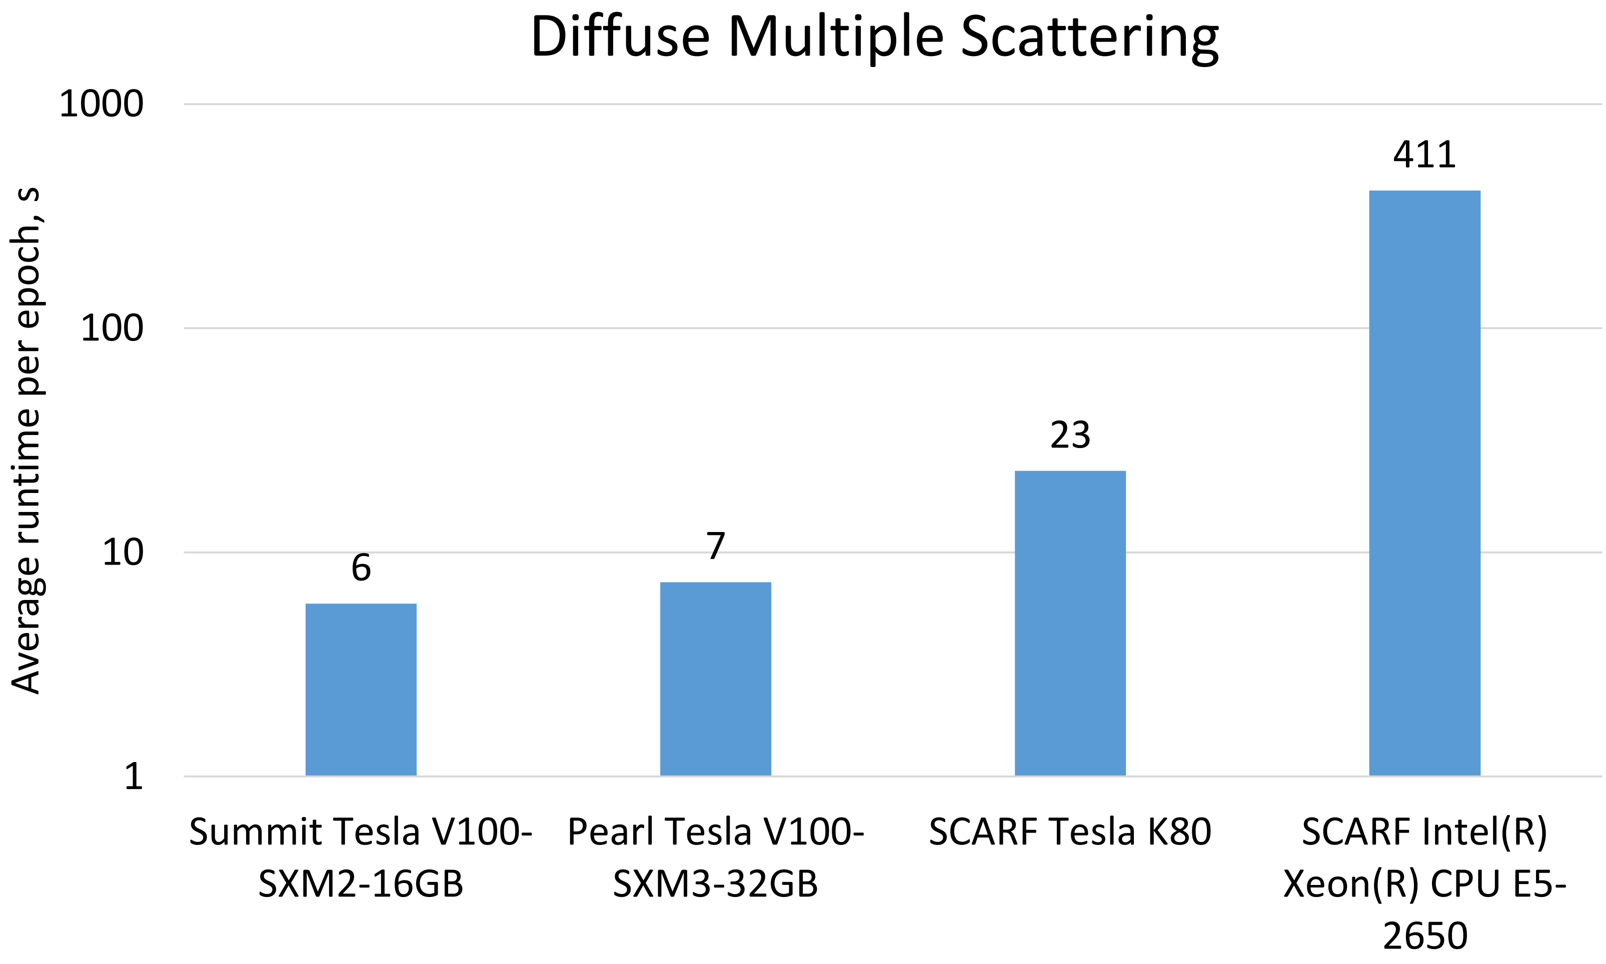

NVIDIA K80 vs V100

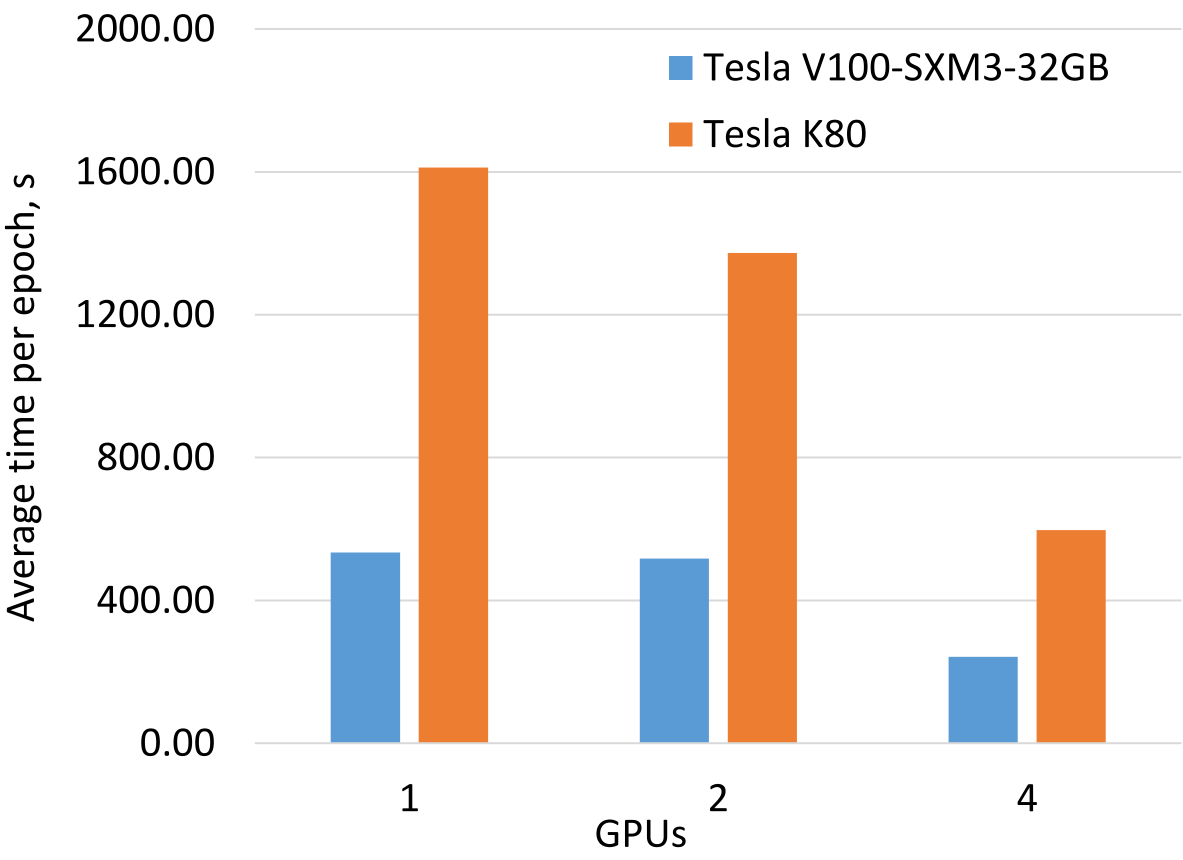

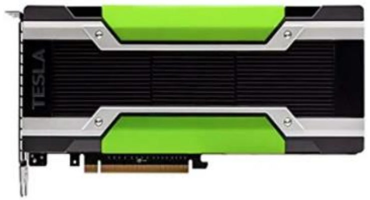

	V100 SXM3 32 GB	Tesla K80
Technology	12 nm	28 nm
Chip area	815 mm²	561 mm²
GPU type	GV100	2 x GK210
Peak Single Precision	14 TFLOPS	5.60 TFLOPS
Peak Double Precision	7 TFLOPS	1.87 TFLOPS
Transistors	21,100 million	7,100 million
CUDA cores	5120	2x2496
Tensor cores	640	n/a
SMs	80	2x13
Bus width	4096 bit	2x384 bit
Memory BDW	1134 GBytes/sec	480 Gbytes/sec
GPU frequency	1290-1530 MHz	562-824 MHz
Max power draw	300 W	300 W
Price:	~$10k	~$0.5k

CloudMask benchmark: slstr_cloud

- Estimation of sea surface temperature
- Sentinel-3 satellite: Sea and Land Surface
- Temperature Radiometer (SLSTR) instrument
- Determine whether the individual pixels of satellite images contain cloud or a clear sky
- Traditional solution: thresholding or Bayesian methods
- U-Net deep neural network
- Two datasets of DS1-Cloud (180GB) and DS2-Cloud (2.3+2.6TB)
- Reflectance (6 channels, 2400 x 3000 pixels)
- Brightness temperature (3 channels, 1200 x 1500 pixels)

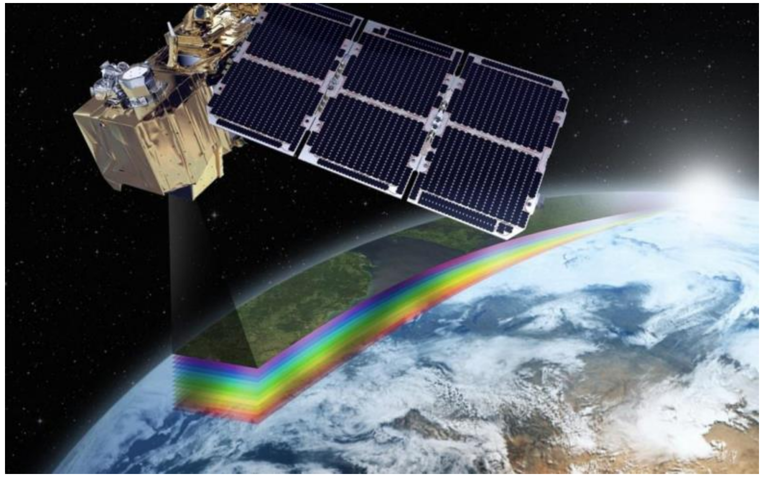

SLSTR = Sea and Land Surface Temperature Radiometer

Science and Technology Facilities Council

Scientific Computing

SciML Benchmarks

Scalability and accuracy on PEARL (DGX-2)

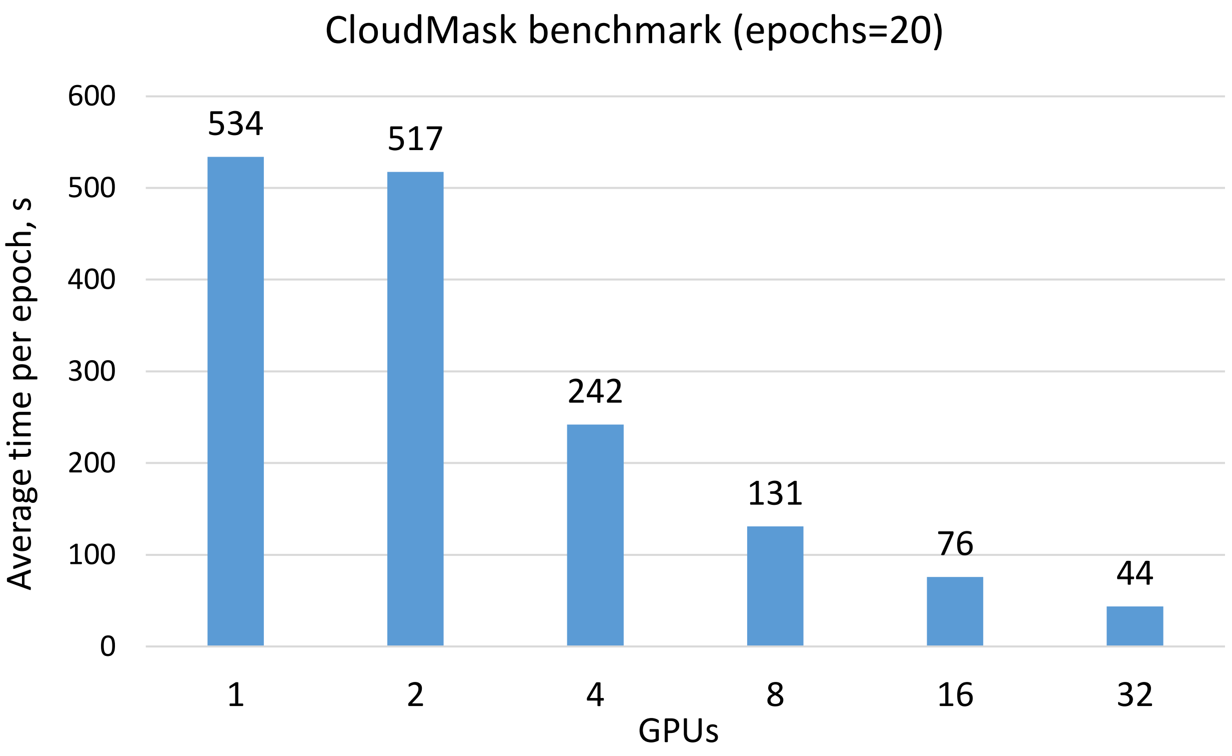

180 GB Dataset

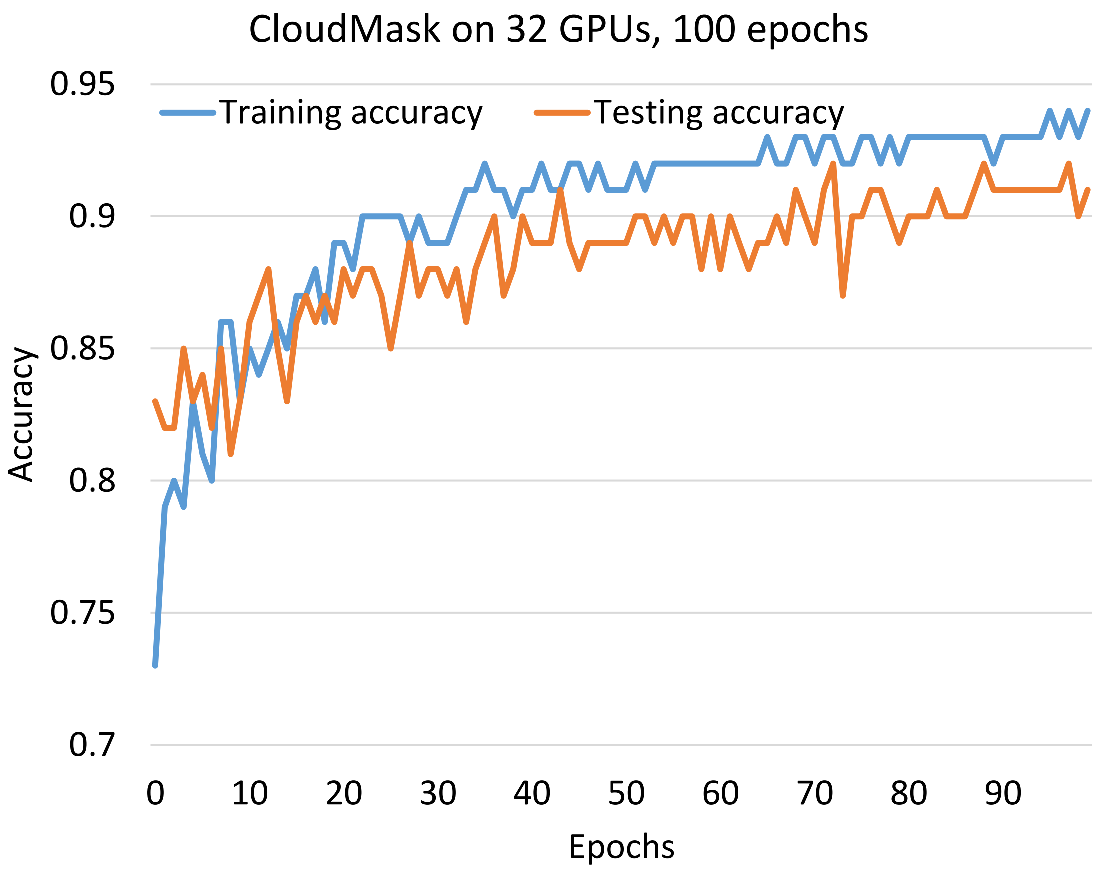

Batch size vs runtime

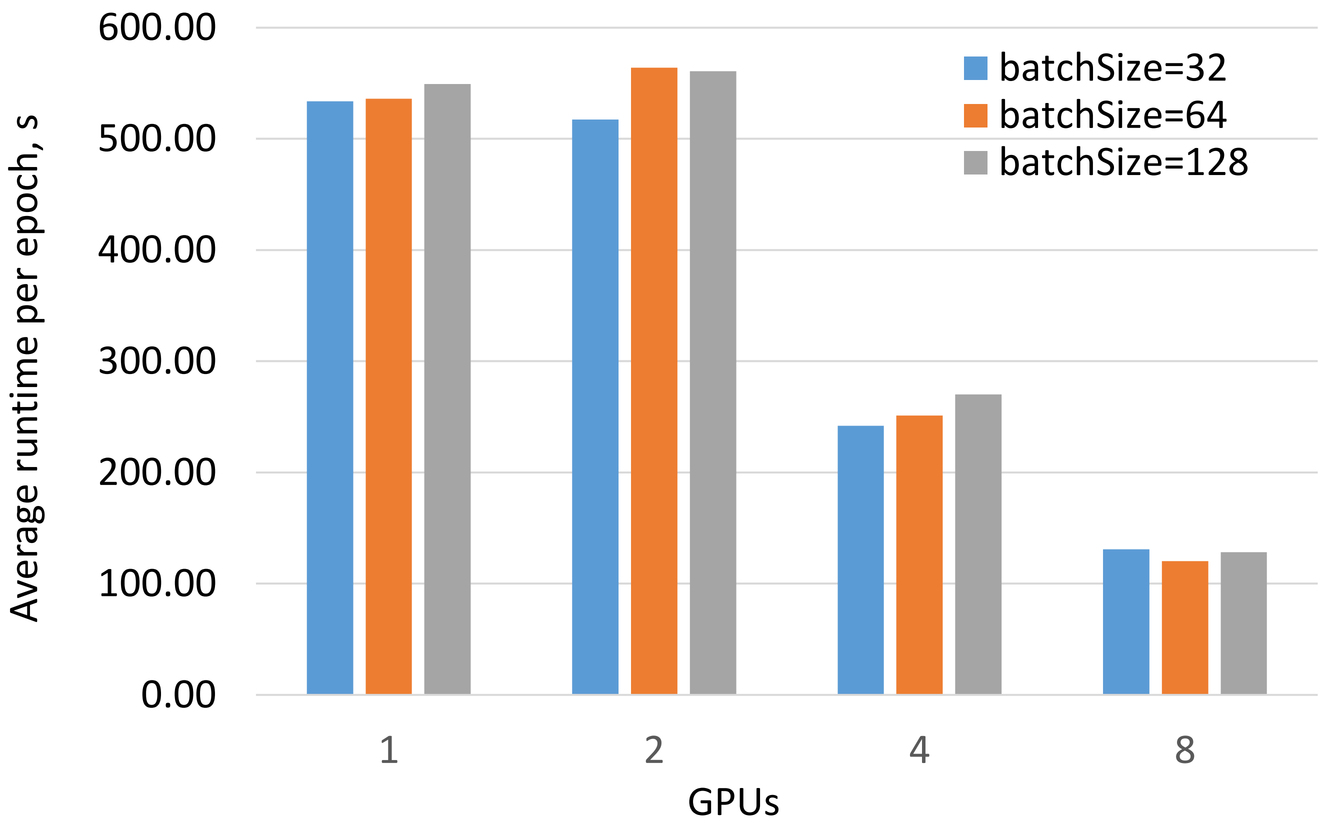

Science and Technology Facilities Council
Scientific Computing

SciML Benchmarks

Big data: GB's to TB's

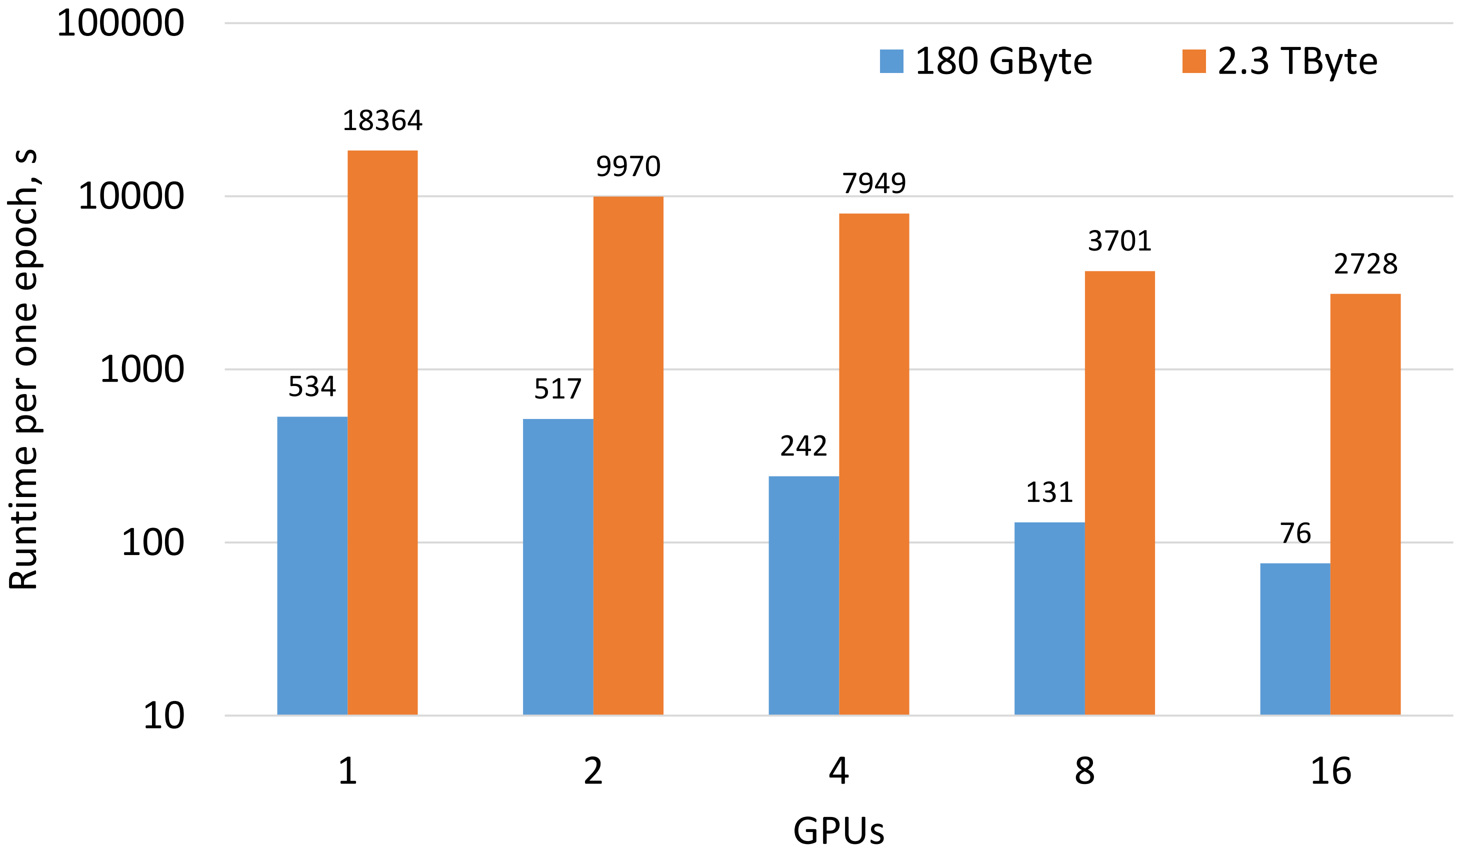

A month of images for:
days: 13542 files, 2.3TB
nights: 15506 files, 2.6TB
Image size: 172MByte

Science and Technology Facilities Council
Scientific Computing
SciML Benchmarks

Summit @ ORNL

© ORNL

Feature	Summit
Peak FLOPS	200 PF
Max possible Power	13 MW
Number of Nodes	4,608
Node performance	42 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2
NV memory per Node	1.6 TB
Total System Memory	2.8 PB + 7.4 PB NVM
System Interconnect	Dual Port EDR-IB (25 GB/s)
Interconnect Topology	Non-blocking Fat Tree
Bi-Section Bandwidth	115.2 TB/s
Processors on node	2 IBM POWER9™ 6 NVIDIA Volta™
File System	250 PB, 2.5 TB/s, GPFS™

Some practical issues – Juri Papay at RAL

- RSA fob – identity check
- No root access
- Singularity container from PEARL does not run (x86 vs ppc64le)
- No IBM Power 9 machines at RAL
- Needed to build container for ppc64le architecture
- Limited disk quota ~52GB, (CloudMask ~200GB – 2.6TB)
- Data transfer nodes ~ 10-20MB/sec
- Max time quota 2 hours

PEARL vs Summit

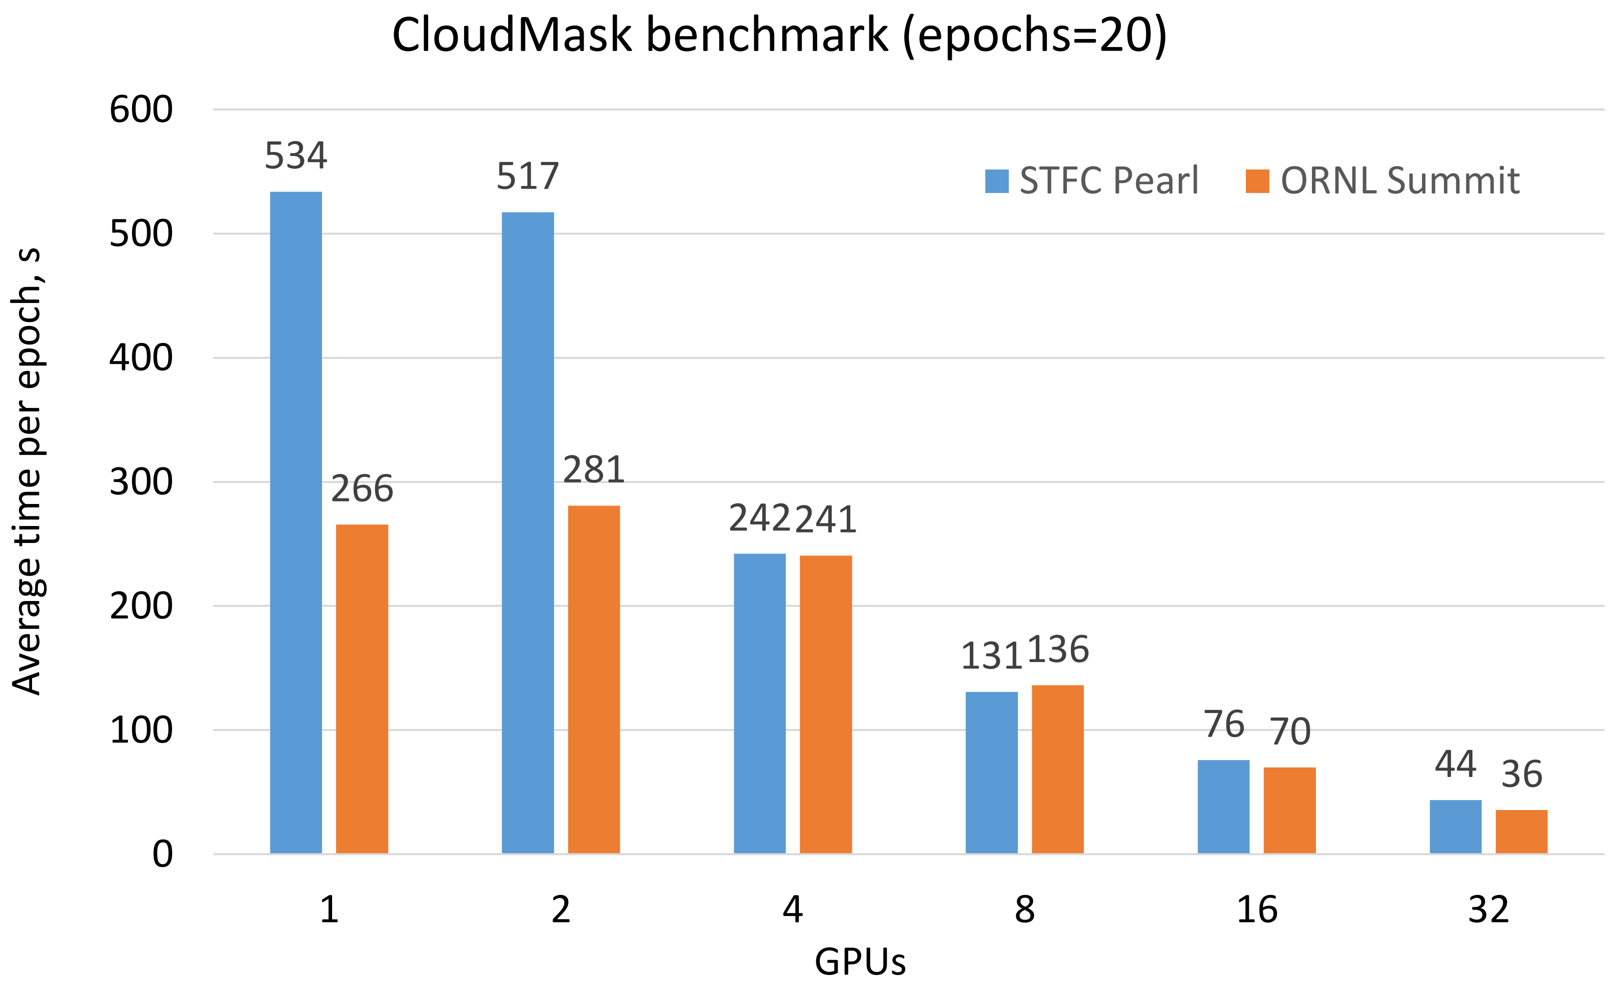

Future work

- Additional benchmarks in the pipeline:
  - Deep-Halo, LIGO, Optical damage, PhotoZ, …
- MLCommons Science Working Group
  - UNO (CANDLE benchmark from ANL)
  - STEM DL (benchmark from ORNL)
- Additional support for distributed training libraries
- More platforms: ORNL, ANL, NVIDIA A100, Cerebras, Groq, Graphcore, …
- Code profiling and Inter-GPU communications
- Develop containers for different architectures

"AI won't replace the scientist, but scientists who use AI will replace those who don't."

Adapted from a Microsoft report, "The Future Computed"

With thanks to David Womble (ORNL)