Capabilities with HPC, GPU, and QIS

Workshop for Applied Nuclear Data Activities

February 10, 2025

Kostas Kravvaris | Nuclear and Chemical Sciences Division

Massive thanks to: K. Wendt, N. Schunck, S. Quaglioni

Prepared by LLNL under Contract DE-AC52-07NA27344.

Maintaining a strong deterrent looks vastly different through the years

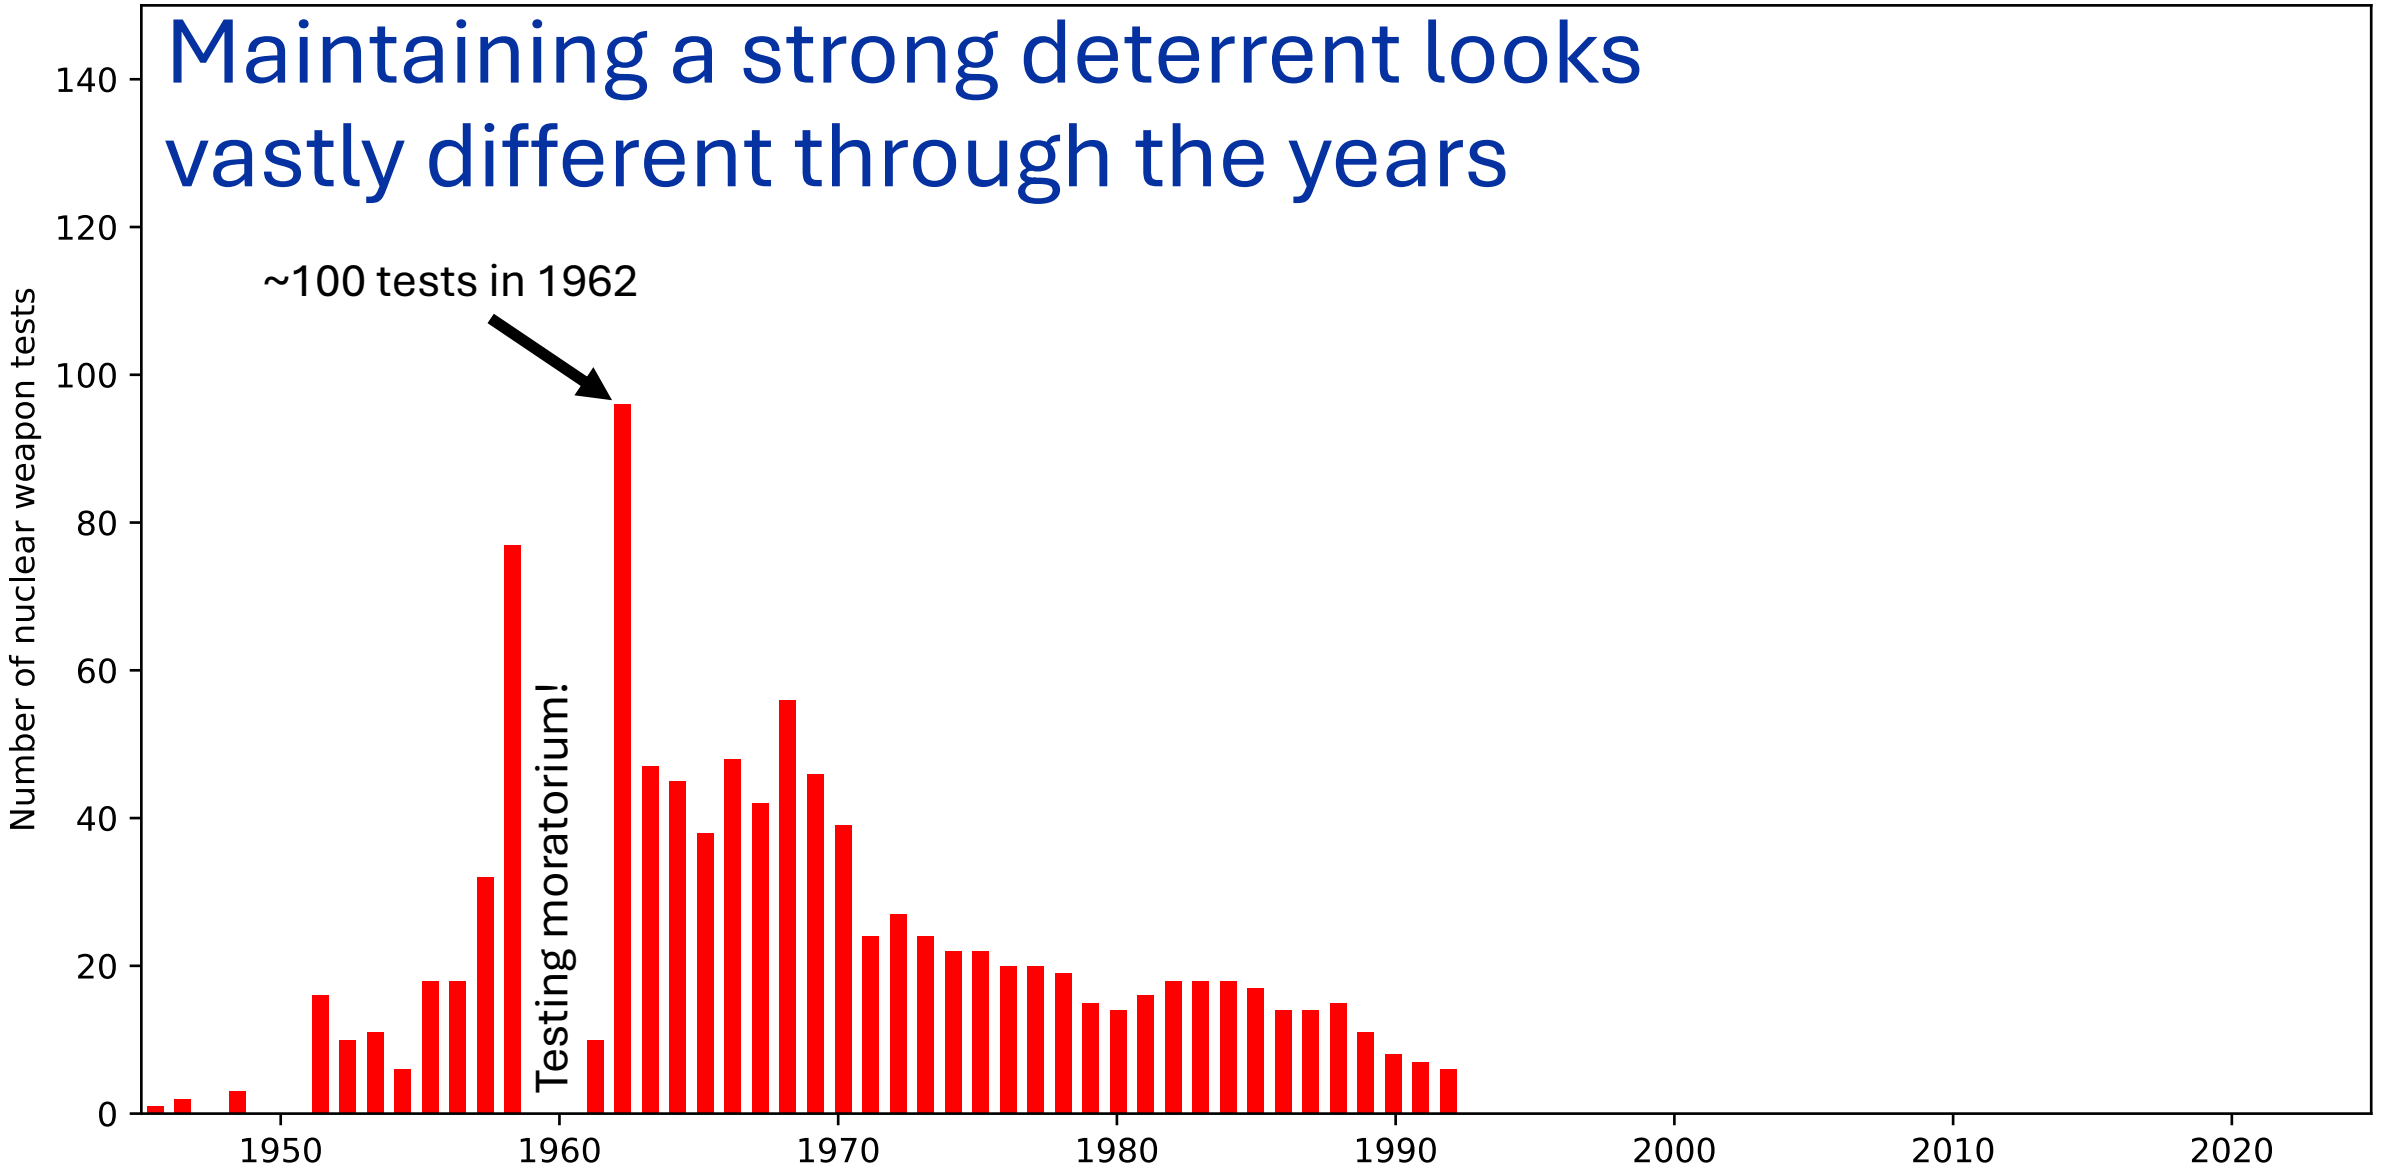

Maintaining a strong deterrent looks vastly different through the years

Maintaining a strong deterrent looks vastly different through the years

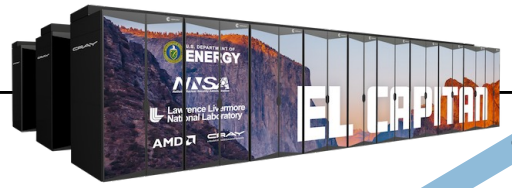

~100 tests in 1962

Testing moratorium!

Sequoia

Roadrunner

Blue-Gene/L

ASCI White

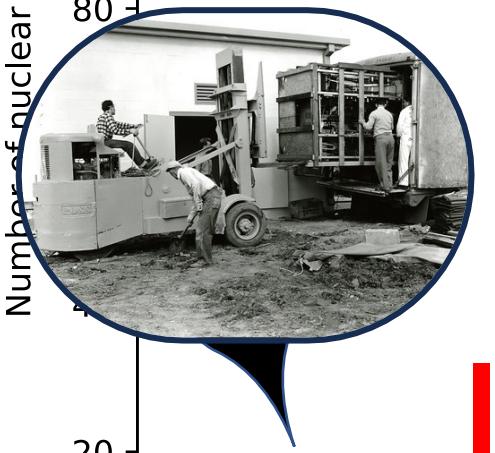

Number of nuclear weapon tests

Computational Power @ LLNL [FLOPS]

https://computing.llnl.gov/about/machine-history

Maintaining a strong deterrent looks vastly different through the years

https://computing.llnl.gov/about/machine-history

A cylindrical conduit that supports—and ensures delivery of accurate nuclear data to—end users.

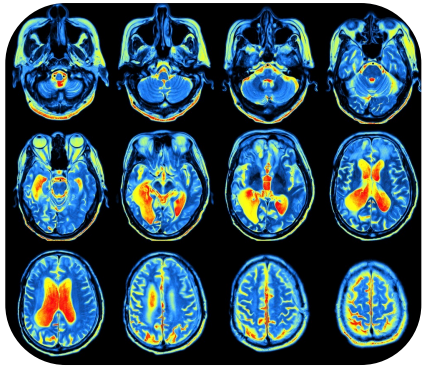

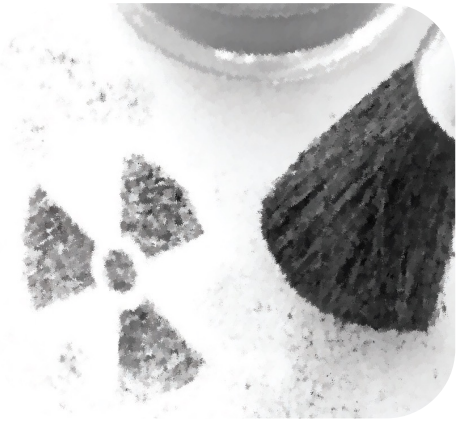

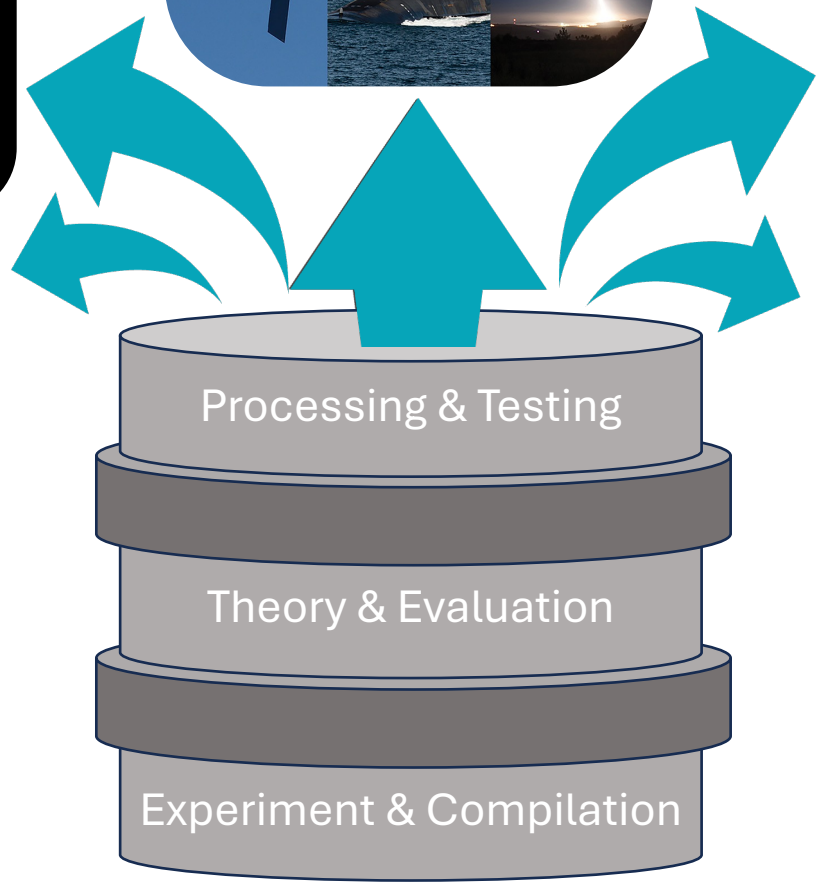

HPC allows for predictive modeling of nuclear data and realistic estimation of modelling uncertainties

The nuclear theory loop of frustration

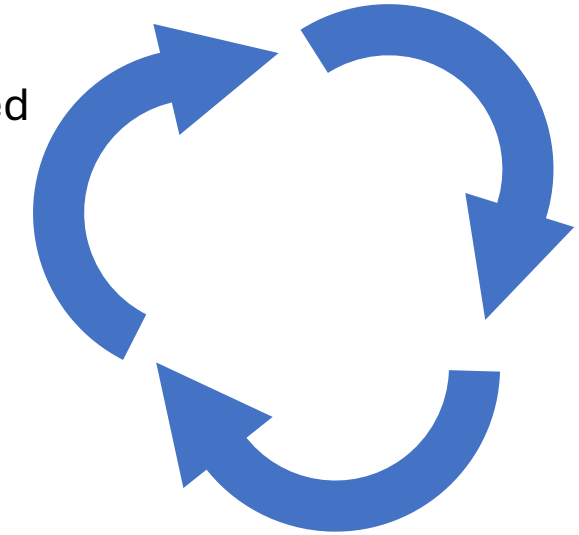

My model does not match my data (to within an acceptable degree)

I need to include more physics in my model

More compute power is needed to include said physics

HPC allows for predictive modeling of nuclear data and realistic estimation of modelling uncertainties

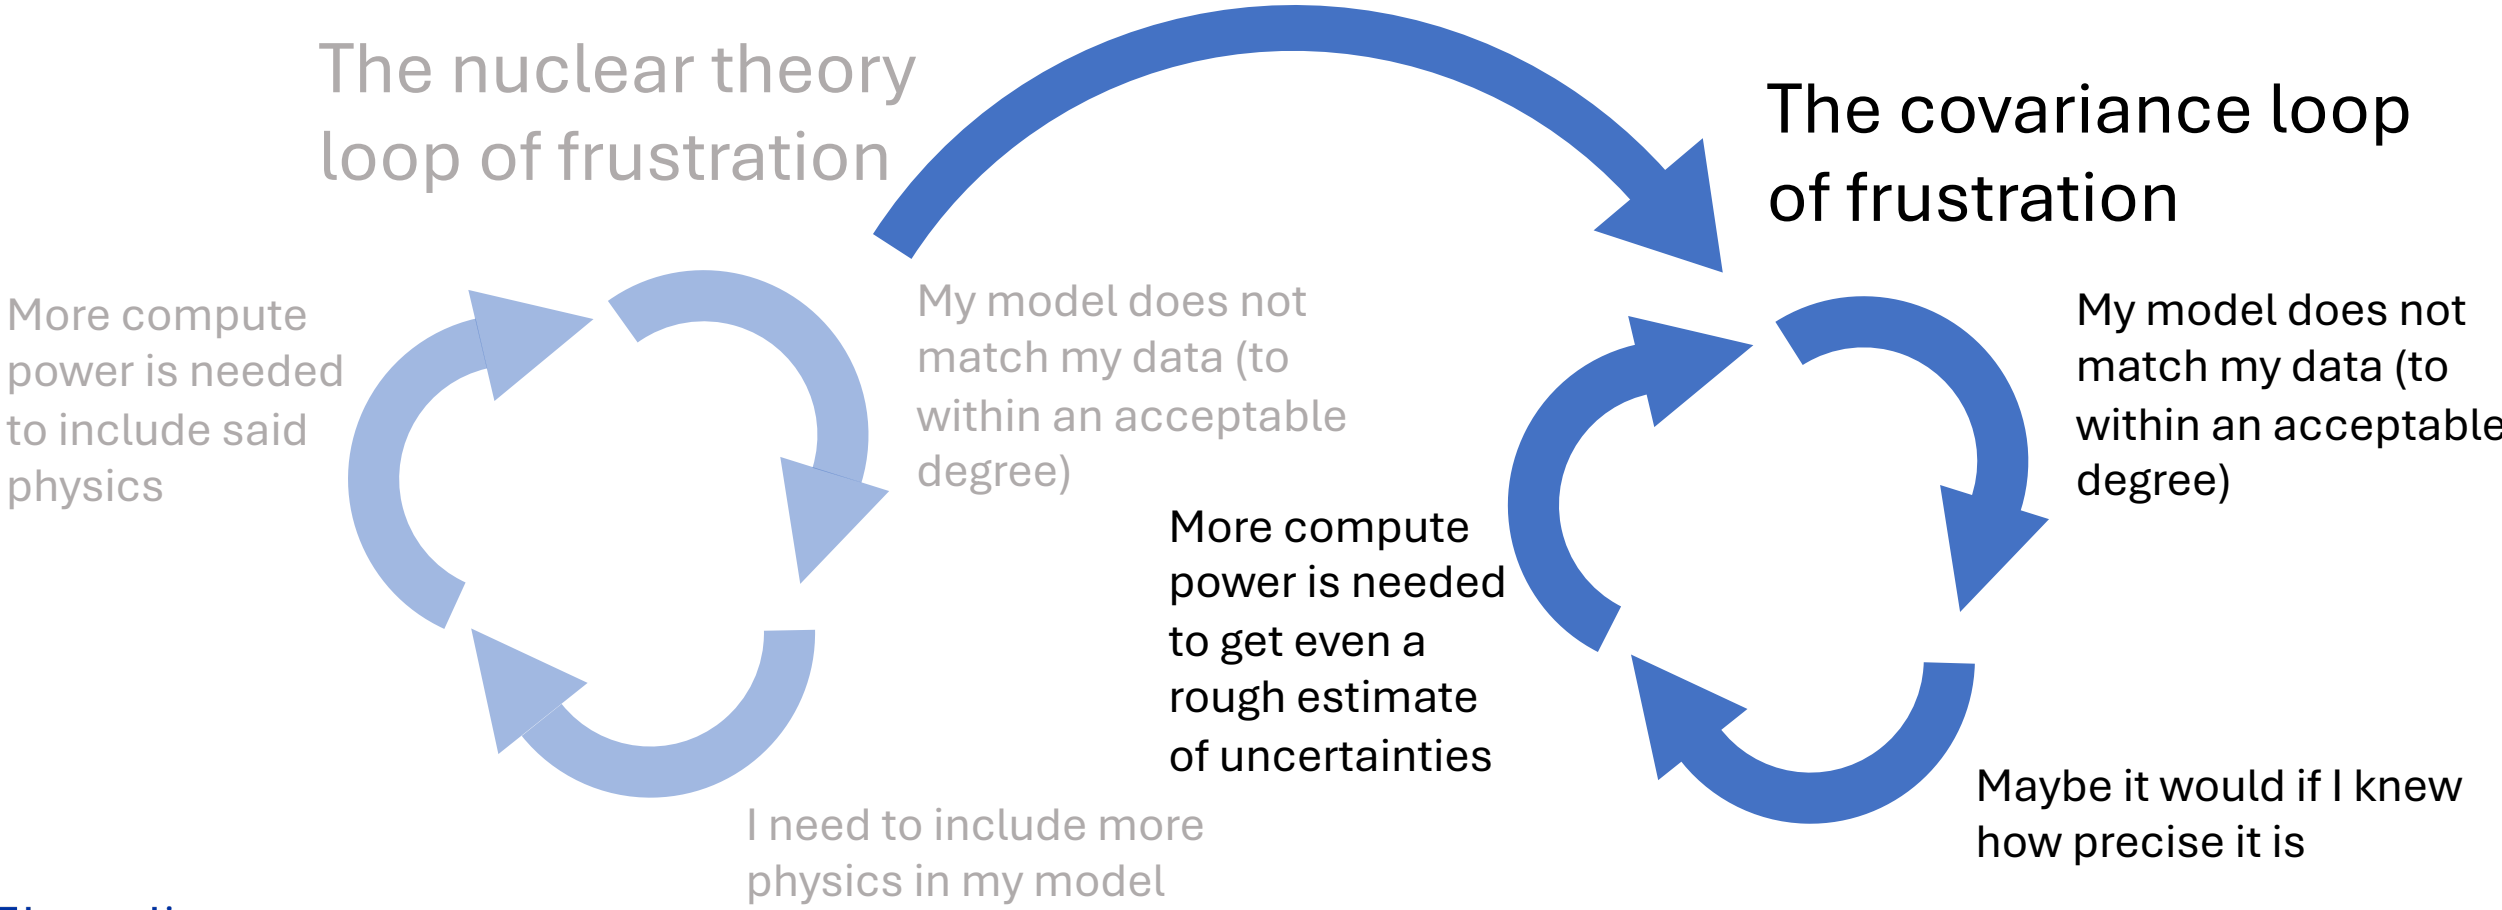

Removing all assumptions in nuclear data models results in accurate evaluations at a high computational cost

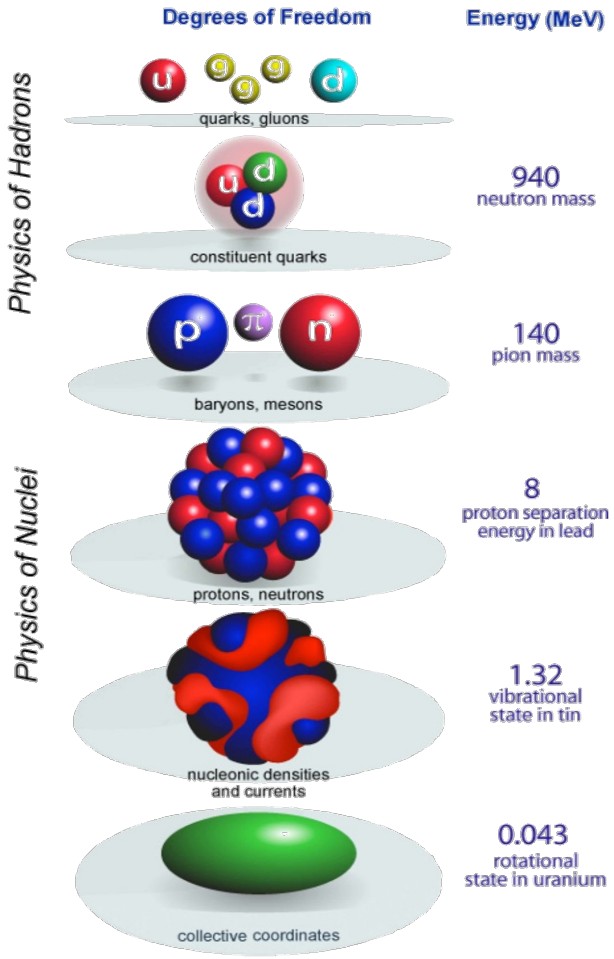

Removing all assumptions in nuclear data models results in accurate evaluations at a high computational cost

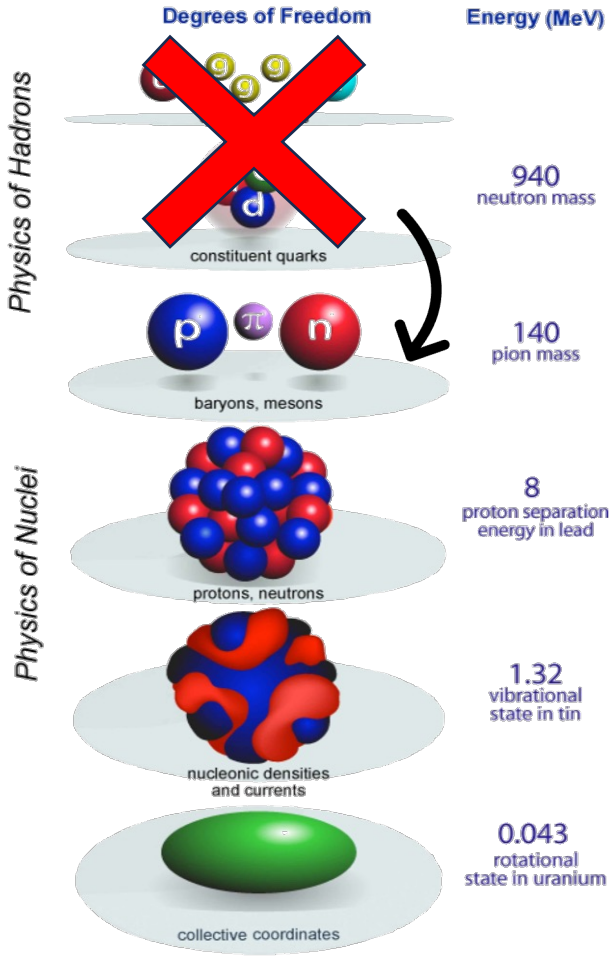

Removing all assumptions in nuclear data models results in accurate evaluations at a high computational cost

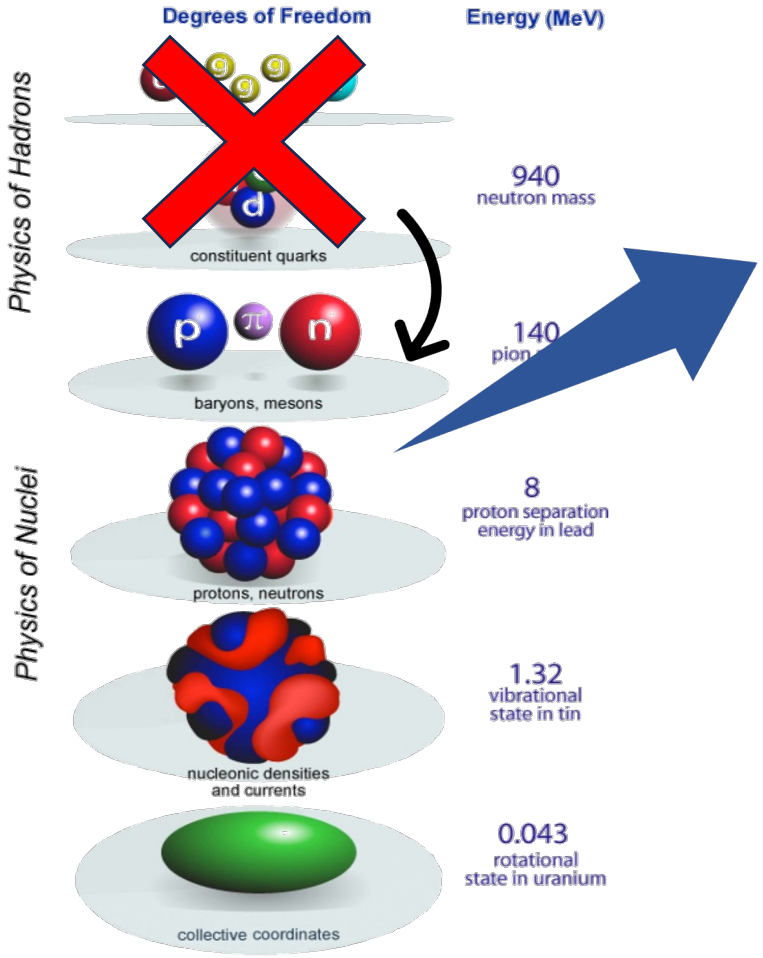

Removing all assumptions in nuclear data models results in accurate evaluations at a high computational cost

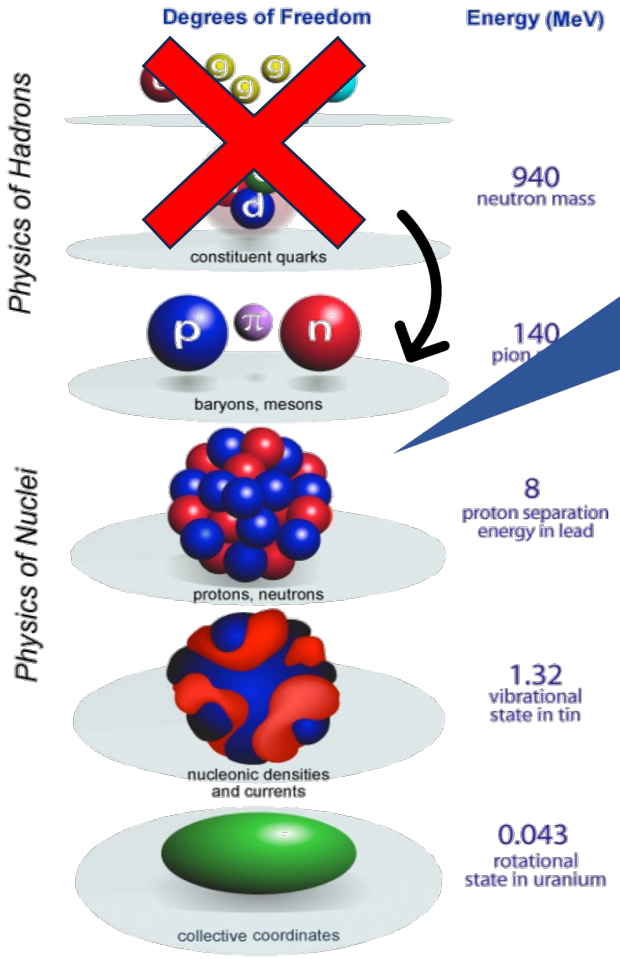

Predicting polarized DT fusion.
From scratch!

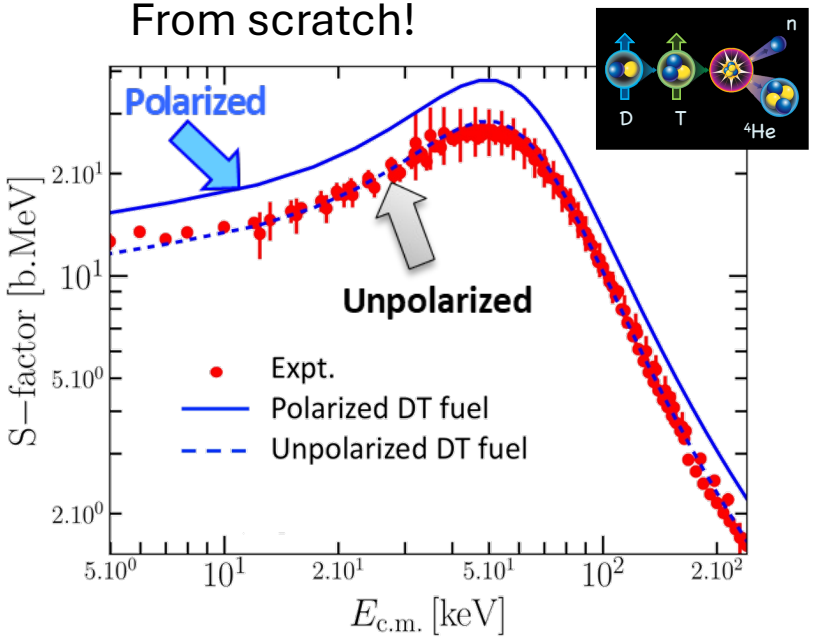

Removing all assumptions in nuclear data models results in accurate evaluations at a high computational cost

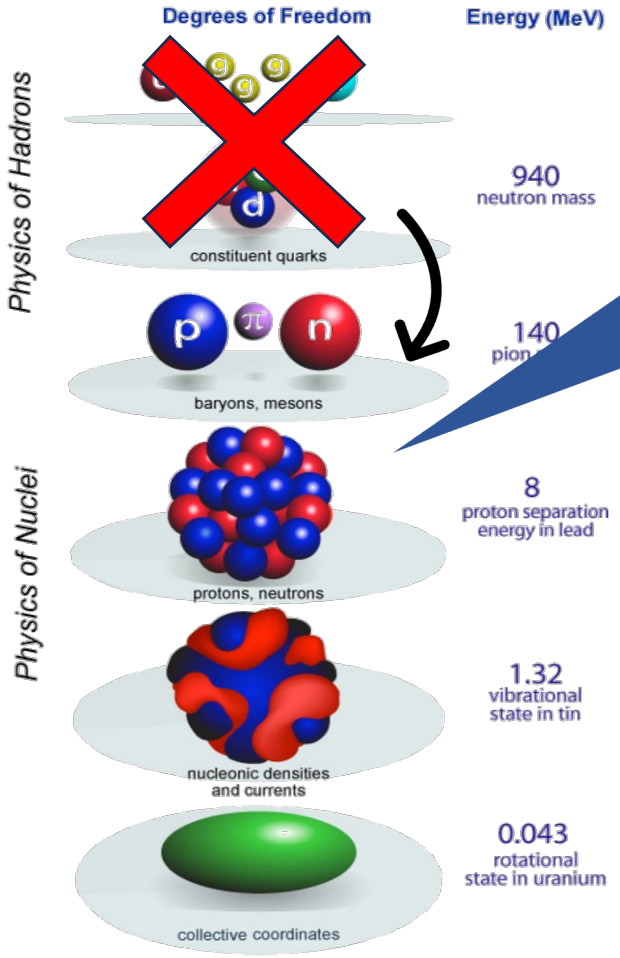

Predicting polarized DT fusion. From scratch!

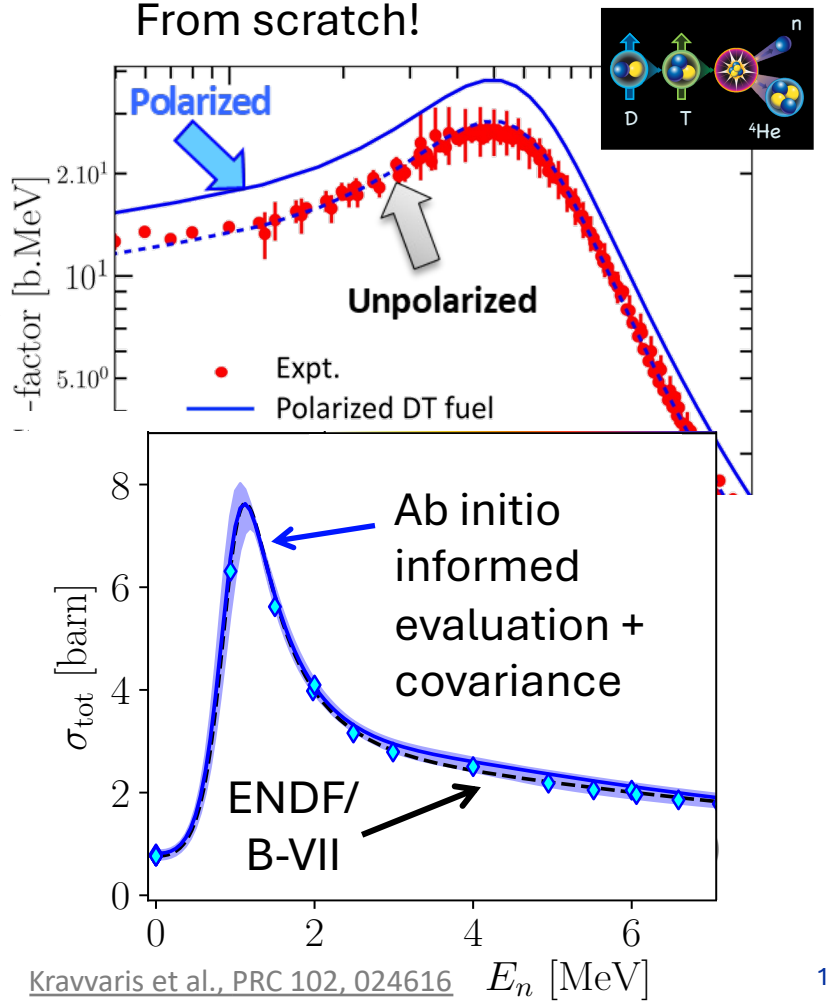

Kravvaris et al., PRC 102, 024616

Calculation adjusted to reproduce reaction cross section yields improved predictions for angular distributions

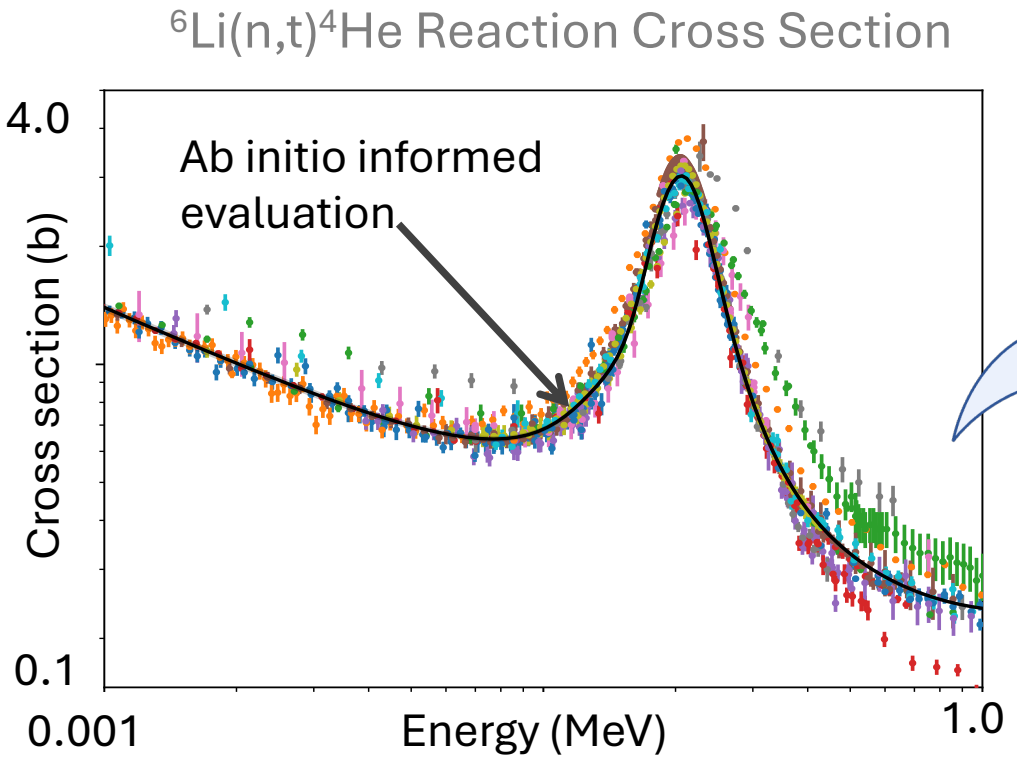

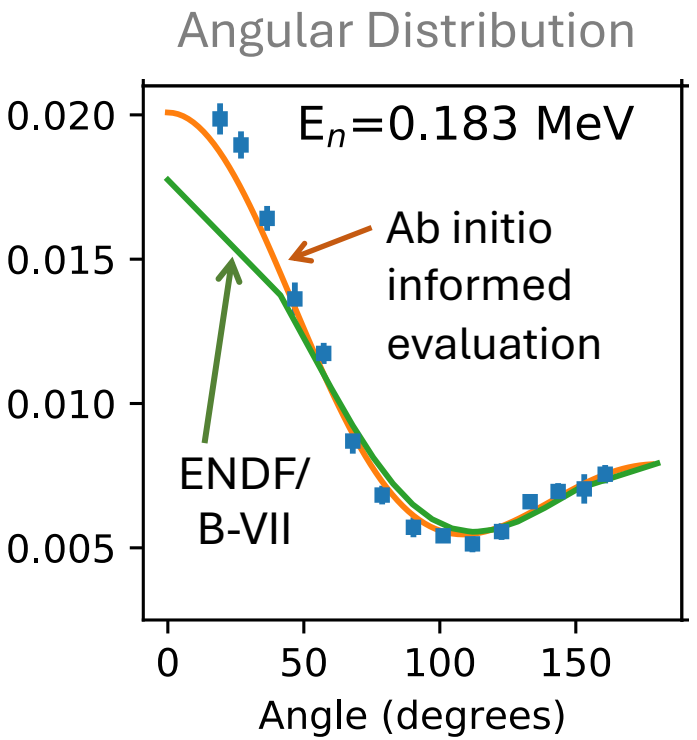

Cost: ~2years of development, demonstrating & implementing new algorithms, some light ML, a significant amount of compute time.

Calculation adjusted to reproduce reaction cross section yields improved predictions for angular distributions

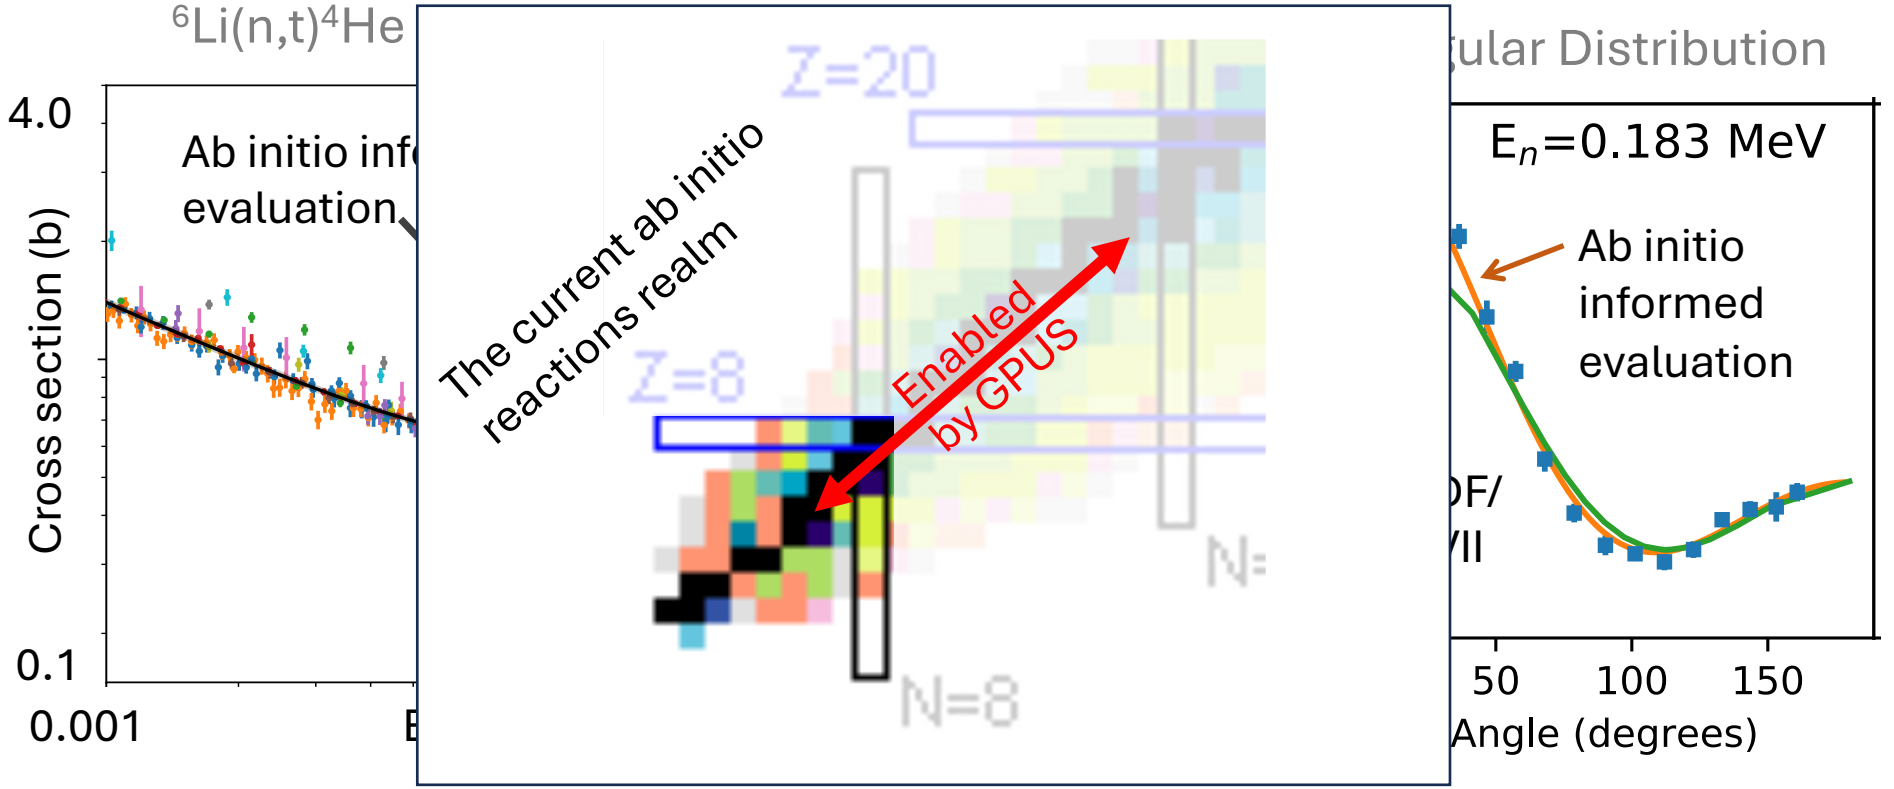

Cost: ~2years of development, demonstrating & implementing new algorithms, some light ML, a significant amount of compute time.

Many applications require properties of thousands of different nuclei

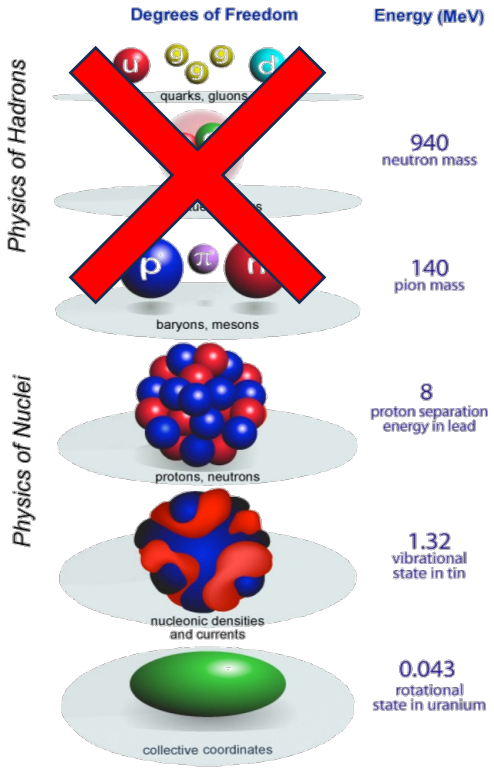

- Density functional theory is a quantum-mechanical theory of atomic nuclei that scale across the entire chart of isotopes
- DFT provides a **consistent** framework for calculations of Q values, decays, fission based on a single set of parameters

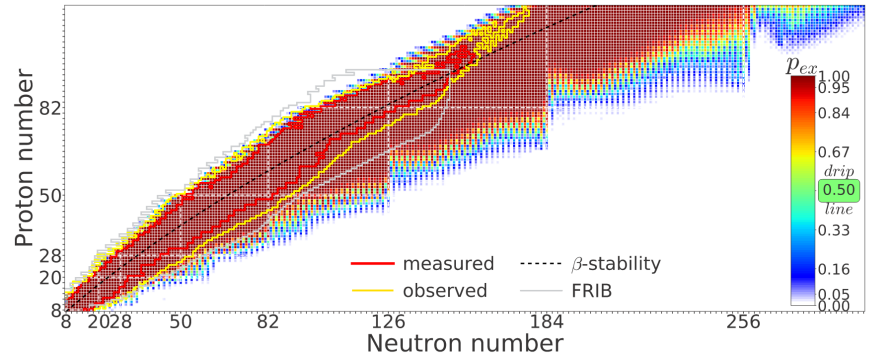

2011	2012	2020	2020	??
DFT mass table	Limits of existence	β-decay rates	Probability of existence	??
	Nature **486**, 509	PRC **102**, 034326	PRC **101**, 044307	

HPC has introduced a paradigm shift in less than 15 years

With HPC, DFT has a shot at cracking fission (really…)

Spin distributions and average spins of FFs (photon spectrum)

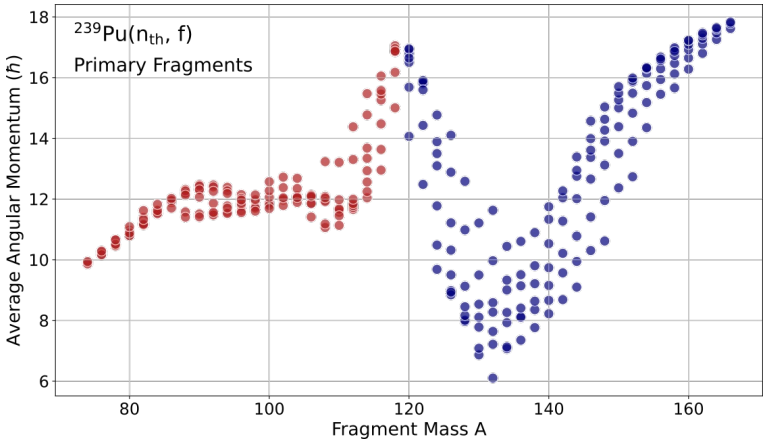

Angular momentum projection and spin distributions

PRC **106**, 014624; PRL **128**, 022501
PRL **126** 142502; PRC **104**, L021601

Two-dimensional yields Y(Z,A)

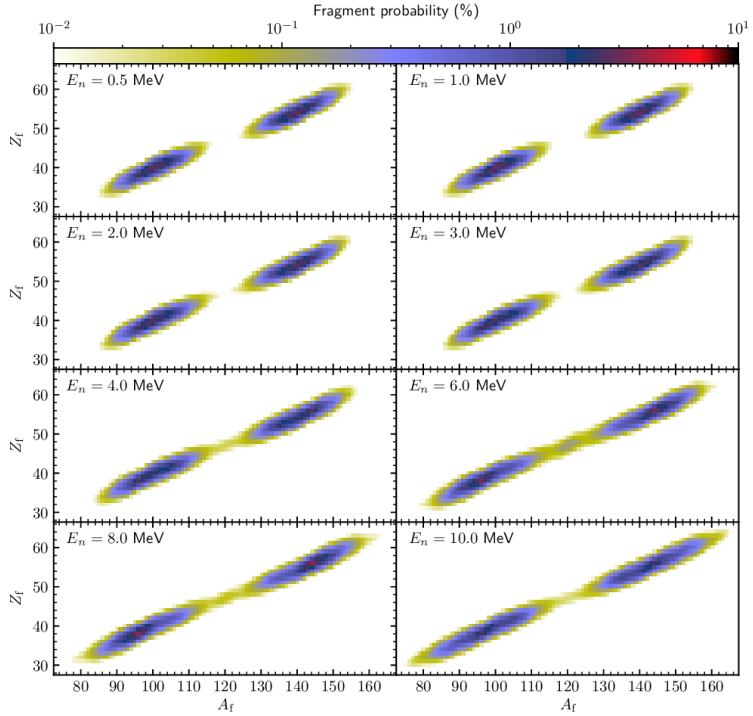

FPYs from time-dependent generator coordinate method

PRC **107**, 044312; PRC **103**, 054602; PRC **99**, 024611;
CPC **225**, 180; PRC **93**, 054611

Total kinetic energy and excitation energy of FFs (prompt spectrum)

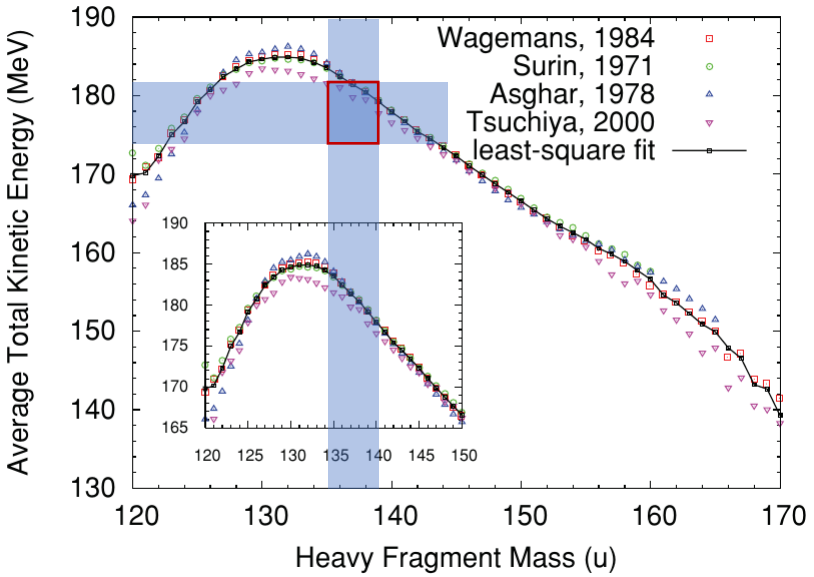

Real-time fission dynamics from time-dependent DFT

CPC **269**, 108130; PRC **100**, 034615; PRL **116**, 122504

Peta- and exascale computers made such simulations possible

Uncertainty quantification and propagation cannot happen without HPC systems

Port computationally-expensive kernels to GPU

Generate real data from models running on GPU-enabled HPC systems

Train emulators with AI/ML models

Propagate theoretical uncertainties to prediction of observables

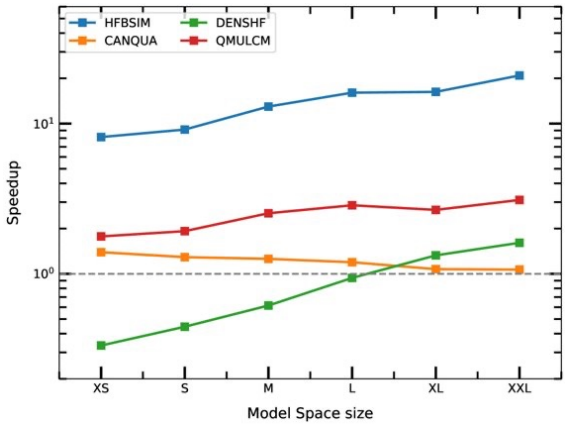

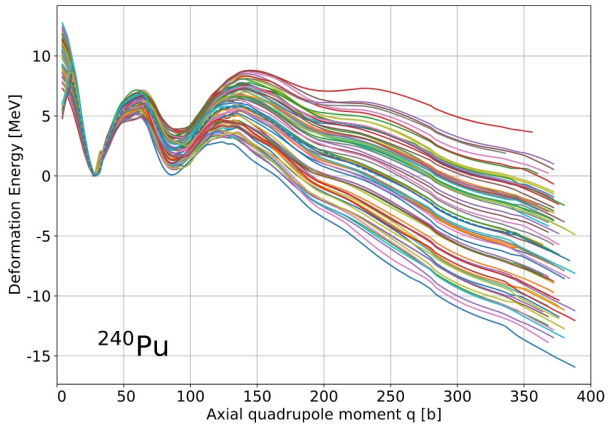

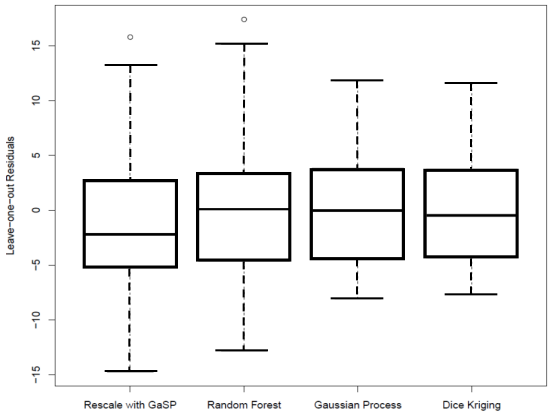

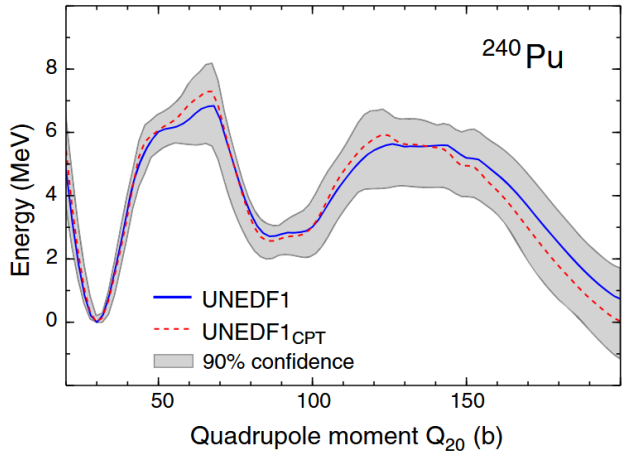

Generating quality data to build reliable and efficient emulators is where HPC is needed

With great power comes a great utilities bill. Machines become more efficient by shifting architectures.

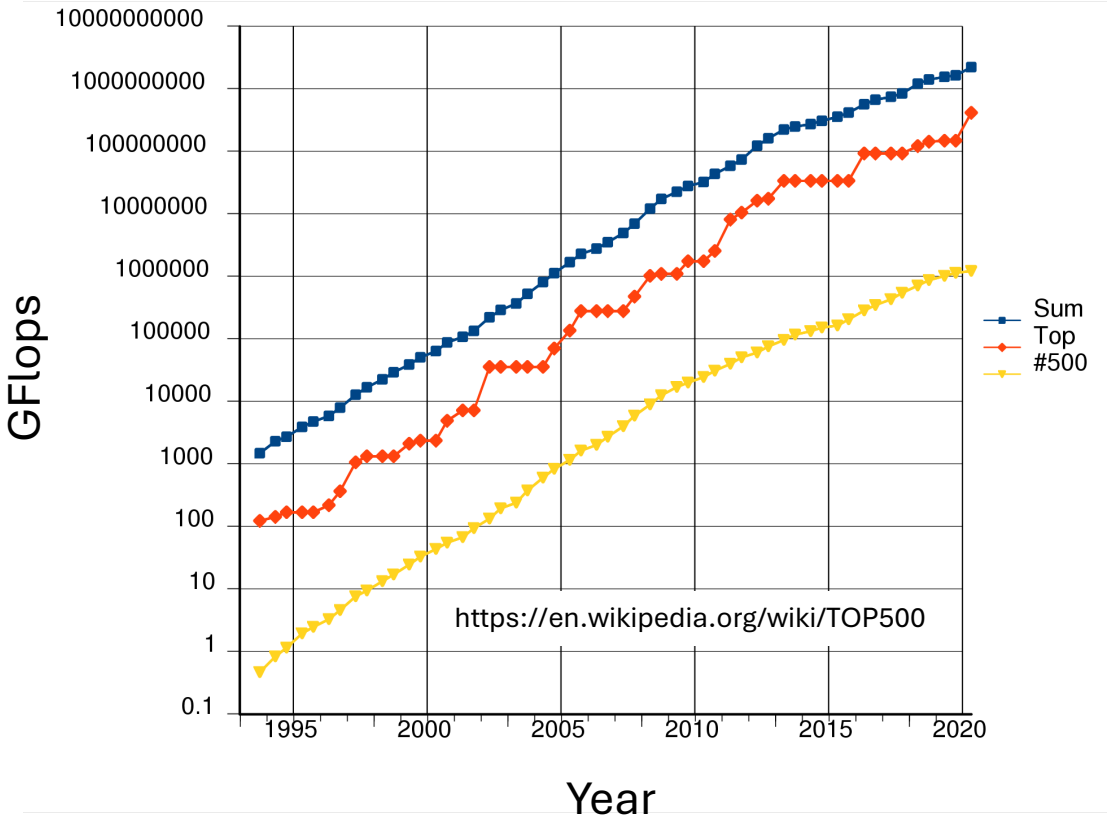

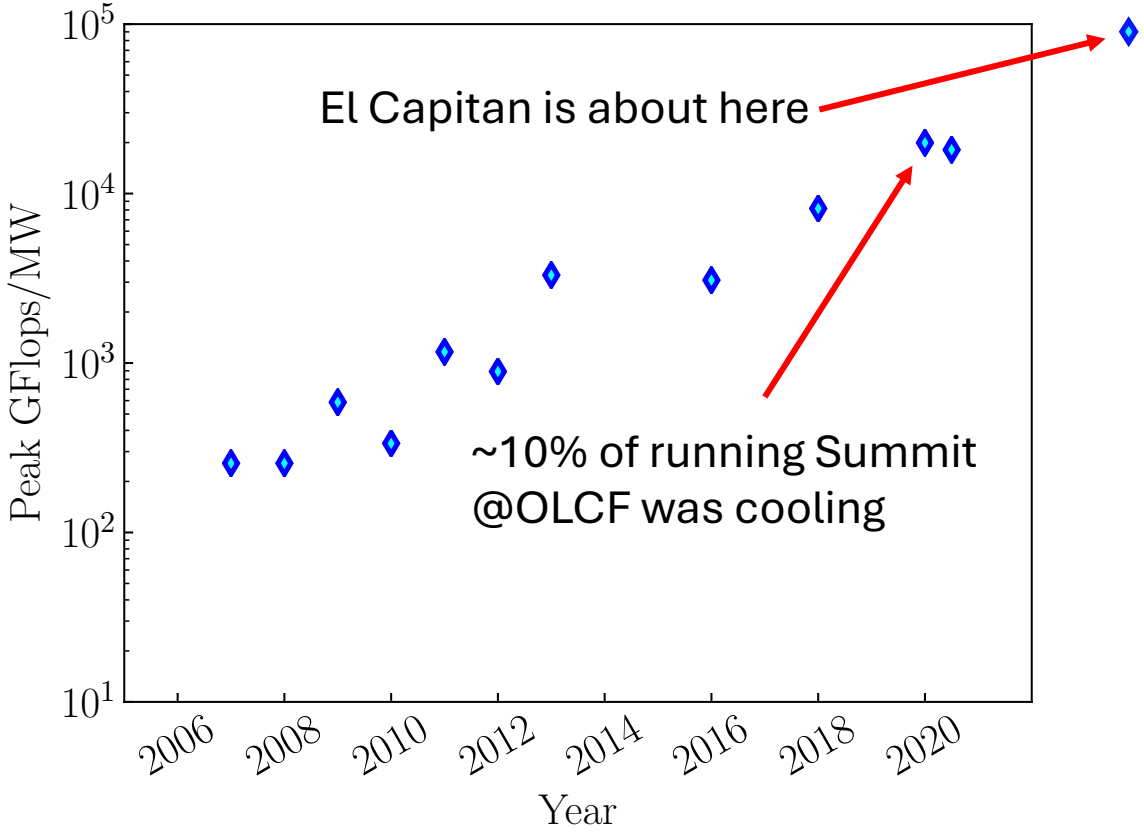

Going beyond general-purpose GPUs: hardware specialized for machine learning.

World's Largest computer chip

SambaNova Systems DataScale™

Great if your problem is training a neural network. Not so great if your problem cannot be mapped to it.

What is the next step for HPC when it comes to modeling quantum systems?

The next (on-board?) accelerators in HPC are going to be **Q**PUs

bit

Either 0 or 1

qubit

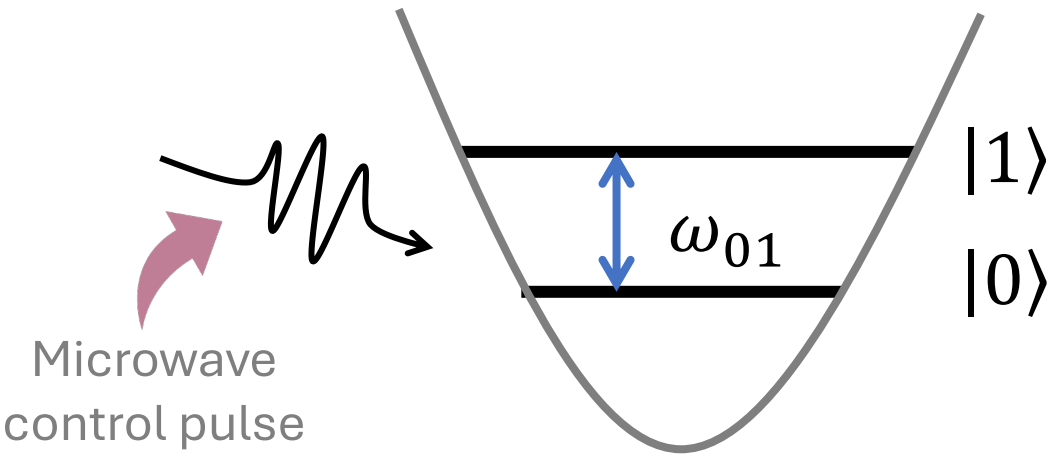

$$|\psi_{qubit}\rangle = \alpha|0\rangle + \beta|1\rangle$$

Pulse-level control of hardware allows constructing operations that will be specific to the problem solution.

Similar to current compute architectures

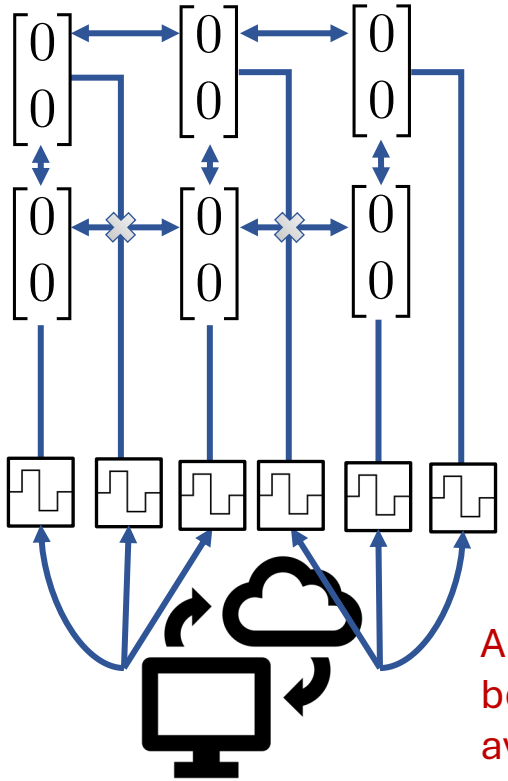

Uses only two modes (**one qubit**) on **multiple** connected **quantum devices**

Discrete set of qubit operations

Algorithms can only be built using available operations

Uses multiple modes (**multi-qubit**) within a **single quantum device** controlled by software

Continuous software-**customizable** operation

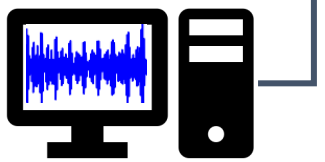

Ability to rethink algorithms at the machine-instruction level

With custom gates, designed noise-resilient algorithm for quantum simulation of multi-nucleon spin dynamics

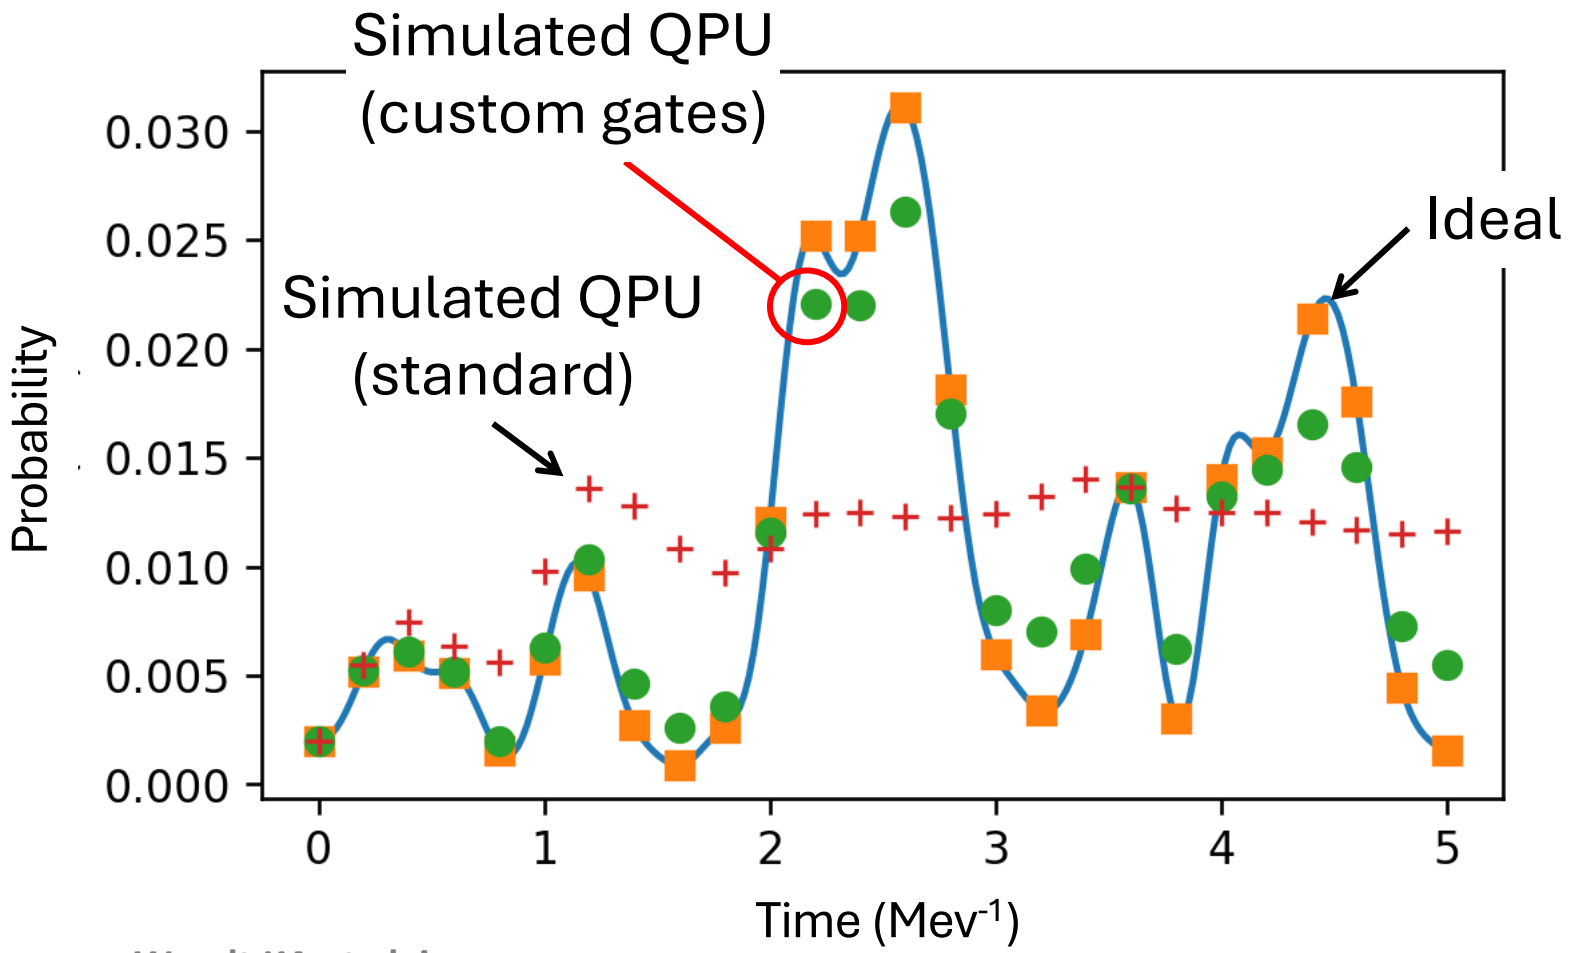

Wendt KA et al, in prep

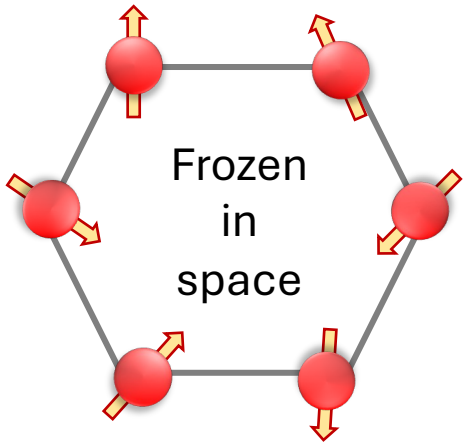

≡ 6 qubits

Polynomial scaling
with number of particles

Gate-based QC still useful despite decoherence, significantly lowers barrier to entry

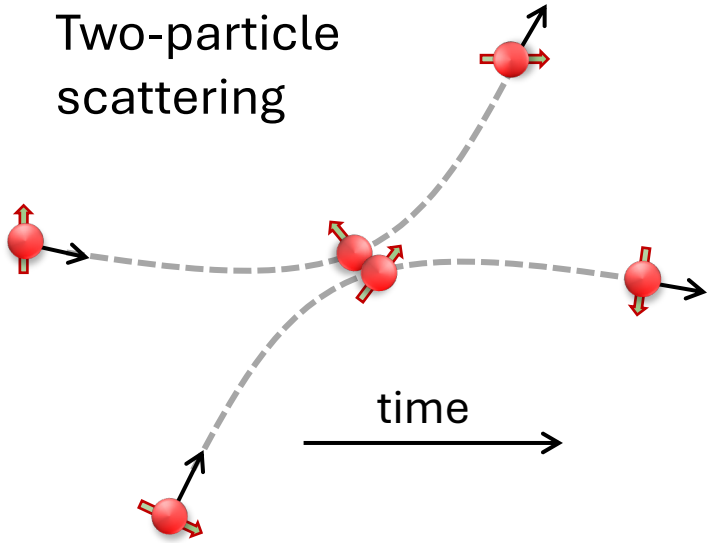

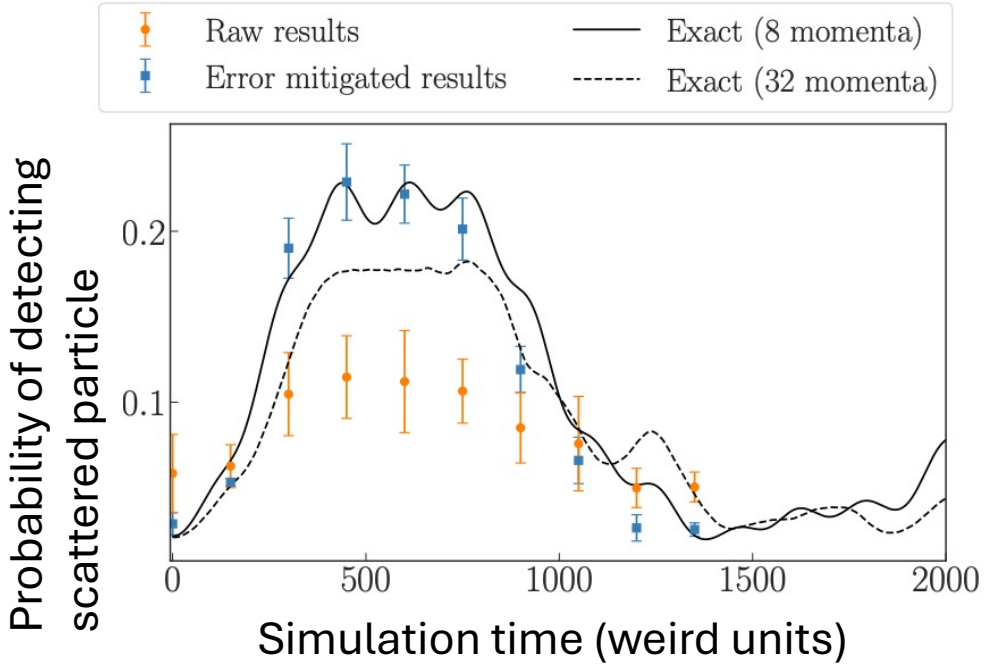

What's the lesson, what is the take-away?

Near-future advanced computing architectures:

1. Theory model outputs for ML training.
2. Increased precision and predictive capabilities for nuclear models/Microscopic inputs for phenomenological ones
3. Improved uncertainty estimates for theory & models.
4. Uncertainty propagation from evaluation to users
5. Leveraging HPC in nuclear experiments for quick turnaround & better/more data

What is stopping us?

1. High barrier to entry/Current resources not nuclear-specific-–Training
2. May need to fundamentally rethink how problems are solved—Training
3. Coding language support varies for different platforms—Software
4. Codes require fine-tuning for specific platforms—Software

Maybe-not-so-near-future advanced computing architectures:

1. Hardware tuned to the solution of a given type of problem
2. Previously-impossible calculations & modeling approaches in sight