Make it so!

A fresh look on how individual researchers and teams can make large-scale research infrastructures work for them.

Michael Bussmann · Topic Lead DMA

Helmholtz MATTER

Helmholtz MATTER

Large-scale research infrastructures for fundamental research

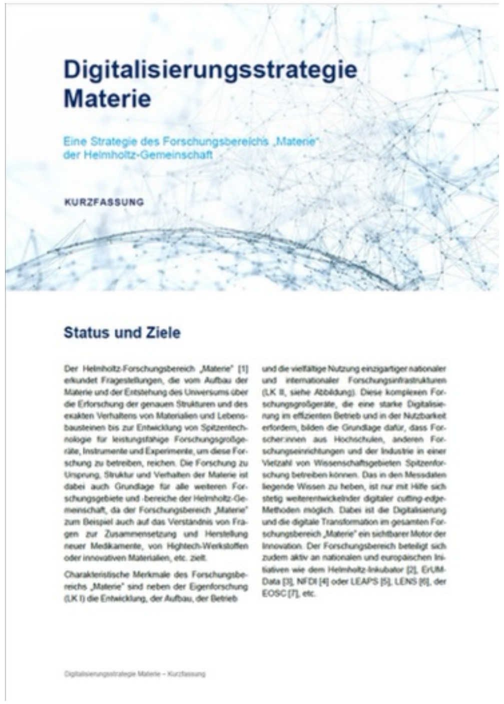
Digitalisierungsstrategie Materie

Eine Strategie des Forschungsbereichs „Materie" der Helmholtz-Gemeinschaft

KURZFASSUNG

Status und Ziele

Der Helmholtz-Forschungsbereich „Materie" [1] erkundet Fragestellungen, die vom Aufbau der Materie und der Entstehung des Universums über die Erforschung der genauen Strukturen und des exakten Verhaltens von Materialien und Lebensbausteinen bis zur Entwicklung von Spitzentechnologie für leistungsfähige Forschungsgroßgeräte, Instrumente und Experimente, um diese Forschung zu betreiben, reichen. Die Forschung zu Ursprung, Struktur und Verhalten der Materie ist dabei auch Grundlage für alle weiteren Forschungsgebiete und -bereiche der Helmholtz-Gemeinschaft, da der Forschungsbereich „Materie" zum Beispiel auch auf das Verständnis von Fragen zur Zusammensetzung und Herstellung neuer Medikamente, von Hightech-Werkstoffen oder innovativen Materialien, etc. zielt.

Charakteristische Merkmale des Forschungsbereichs „Materie" sind neben der Eigenforschung (LK I) die Entwicklung, der Aufbau, der Betrieb und die vielfältige Nutzung einzigartiger nationaler und internationaler Forschungsinfrastrukturen (LK II, siehe Abbildung). Diese komplexen Forschungsgroßgeräte, die eine starke Digitalisierung im effizienten Betrieb und in der Nutzbarkeit erfordern, bilden die Grundlage dafür, dass Forscher:innen aus Hochschulen, anderen Forschungseinrichtungen und der Industrie in einer Vielzahl von Wissenschaftsgebieten Spitzenforschung betreiben können. Das in den Messdaten liegende Wissen zu heben, ist nur mit Hilfe sich stetig weiterentwickelnder digitaler cutting-edge-Methoden möglich. Dabei ist die Digitalisierung und die digitale Transformation im gesamten Forschungsbereich „Materie" ein sichtbarer Motor der Innovation. Der Forschungsbereich beteiligt sich zudem aktiv an nationalen und europäischen Initiativen wie dem Helmholtz-Inkubator [2], ErUM-Data [3], NFDI [4] oder LEAPS [5], LENS [6], der EOSC [7], etc.

Digitalisierungsstrategie Materie – Kurzfassung

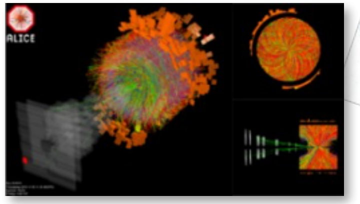

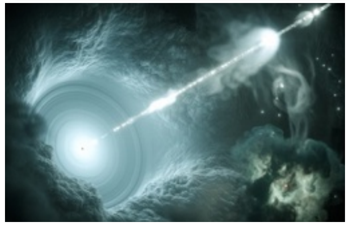

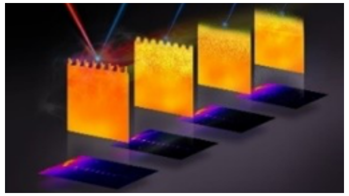

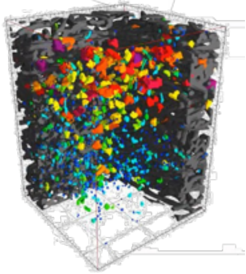

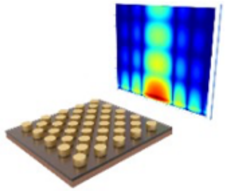

	LK II (User facilities)	International Research Infrastructures	National Research Institutes/ Infrastructures
DESY	FLASH PETRA III IDAF	LHC Belle II CTA (*under construction*) IceCube European XFEL ESRF	CFEL CSSB NanoLab DESY Test Beams DAF HIB@European XFEL PITZ
FZJ	JCNS (in MLZ)	ESS (*under construction*) ILL	(FRM-II)
GSI	UNILAC SIS18 ESR	FAIR (under construction) ALICE@LHC	HI Jena HI Mainz
Hereon	GEMS	ESS (*under construction*)	EMSC
HZB	BESSY II		SupraLab EMIL
HZDR	ELBE HLD IBC	European XFEL EMFL ESRF	HIB@European XFEL DRESDYN
KIT	GridKa	KATRIN Auger IceCube	ATP FLUTE TLK SR Beamlines

Particle Physics, Photon Science, Neutron Science, Astroparticle Physics, Accelerator Physics, Detector Physics, Quantum Physics, Matter under Extreme Conditions, Material science, Life sciences

We unlock the potential of Helmholtz MATTER's data

DMA is the focal point for computational and data science in MATTER

Helmholtz Matter research facilities produce **huge amounts of unique, high-quality data**.
This is both a **challenge** and an **opportunity**.

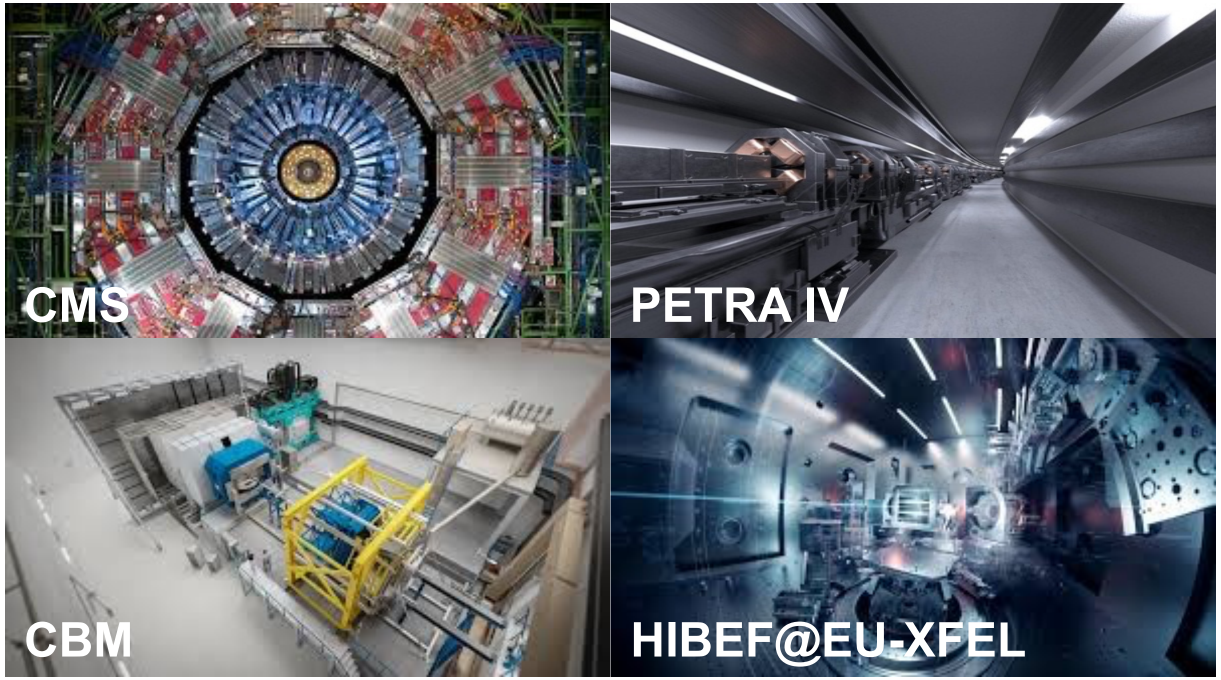

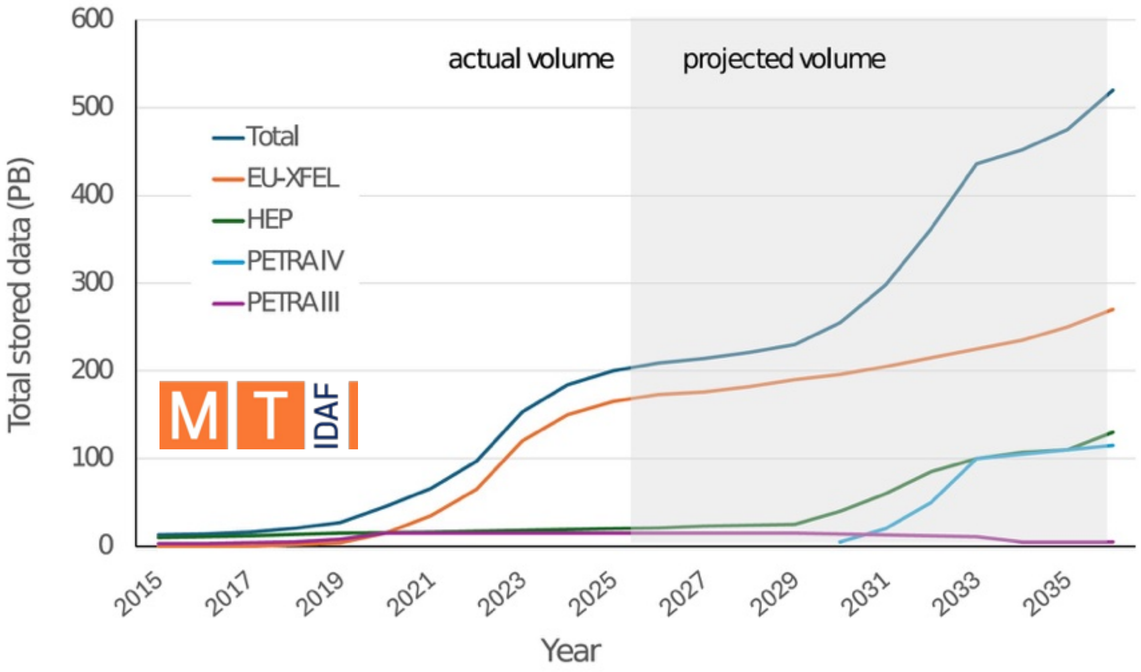

DRESDEN concept

Helmholtz manages & analyses data at high volumes and rates

We are world leading in LHC data management solutions and hold world records in data analysis

dCache **manages** >50% of remote LHC data

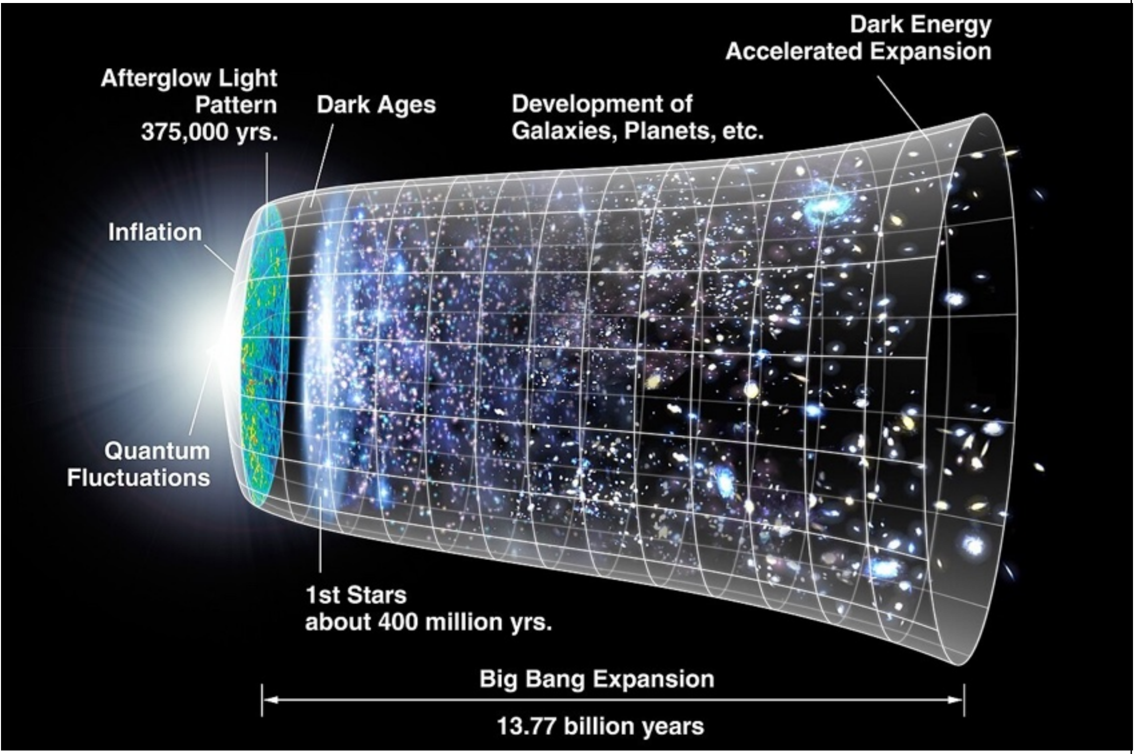

Crystfel **analyses** crystallograms at <500 ms world record speed

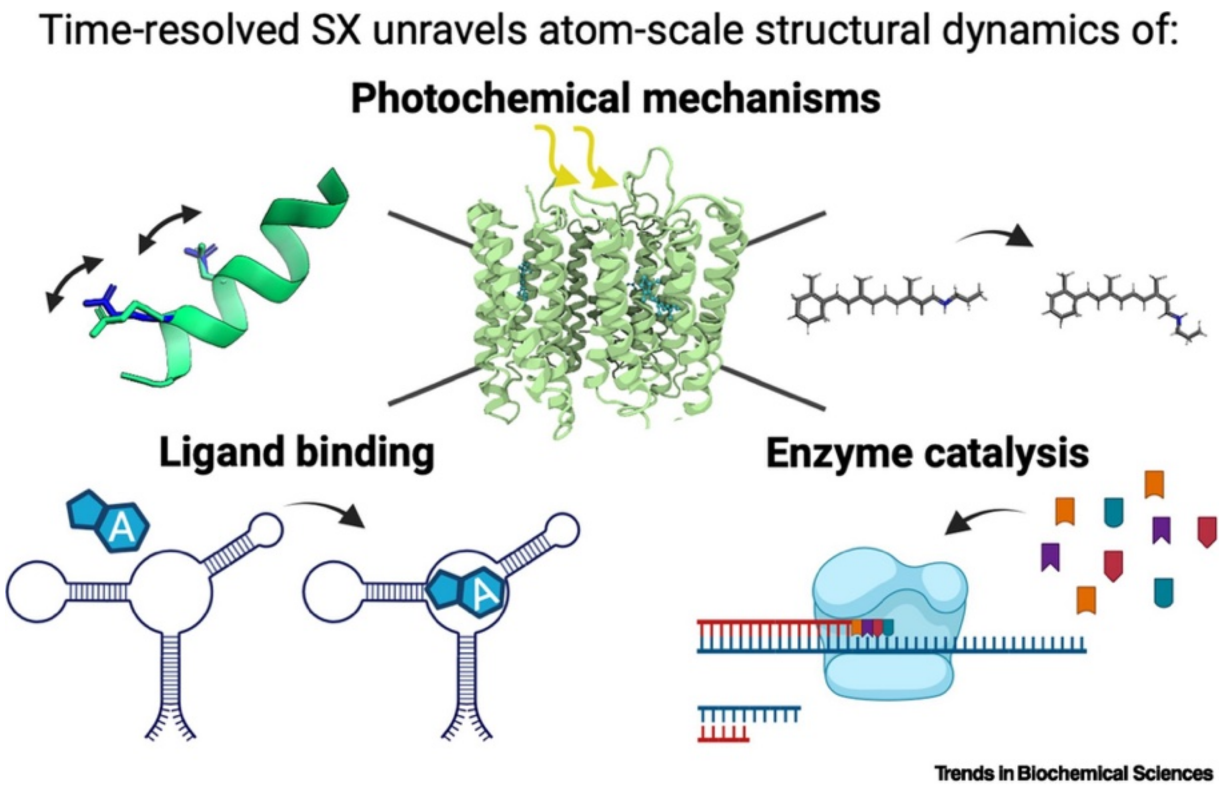

Data Management and Analysis is a Helmholtz cross-center topic

Accelerating technological development transforms science and requires broad adaption

Challenge	Approach
New technologies are needed to understand data	Frontier Methods in Computational and Data Science
Our facilities produce huge volumes of high quality data at high rates	Sustainable Solutions from Data to Understanding
Broad digital innovation at our research facilities is needed	Innovation for Research Infrastructures accelerating Science

Autonomous Research Facilities

When does AI help?

Data + guided search/optimization + feedback (is there really nothing else?)

Lots of high quality data

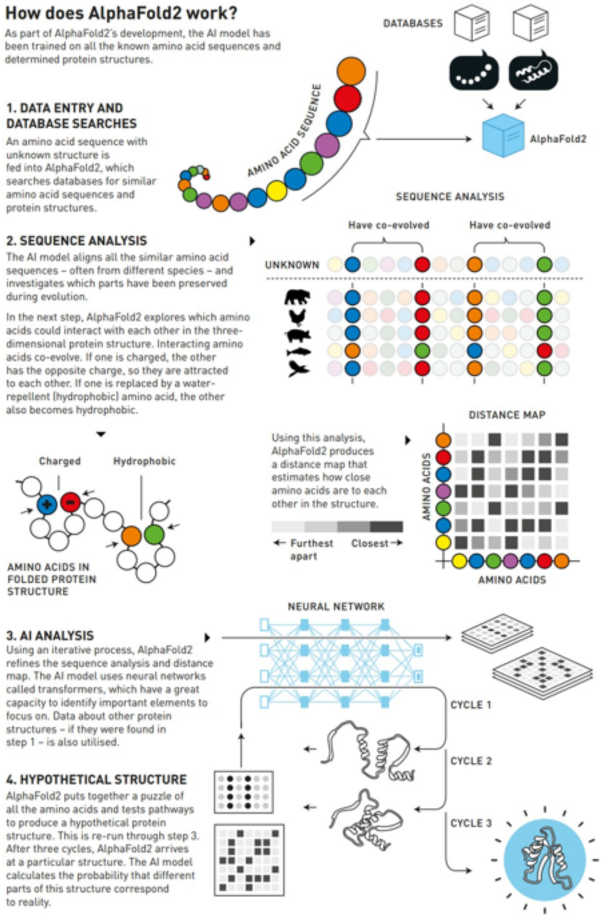

Guided search / optimization

Planar Point Sets with Many Unit Distances

OpenAI

Abstract

For a finite planar set P, let $\nu(P)$ be the number of unordered unit-distance pairs in P, and let $\nu(n)$ be the maximum of $\nu(P)$ over all n-point planar sets. We prove that, for some fixed $\delta > 0$, one has $\nu(n) \geq n^{1+\delta}$ for infinitely many n. This disproves the well-known unit distance conjecture from [Erd46].

The construction passes through an infinite unramified tower of totally real number fields with 3-power Galois groups of growing degree, in which a fixed set of rational primes splits completely. After adjoining i, these fields produce high-dimensional lattices with many elements whose images under every complex embedding have absolute value 1. Golod–Shafarevich theory ensures existence of an infinite such tower, even after a quotient step making the prescribed Frobenius classes trivial. A crucial property of the construction is that all resulting discriminants and class numbers are at most exponential in the extension degree.

1 Main Result

For a finite set $P \subset \mathbb{R}^2$, let $\nu(P) = \#\{\{x,y\} \subset P : |x-y| = 1\}$ and $\nu(n) = \max_{|P|=n} \nu(P)$. The planar unit-distance problem goes back to Erdős [Erd46], who conjectured that there should be an absolute constant C such that, for all sufficiently large n,

$$\nu(n) \leq n^{1+C/\log\log n}. \tag{1}$$

The same 1946 paper also introduced the distinct-distances problem, later resolved up to logarithmic factors by Guth–Katz [GK15]. An elementary upper bound is $\nu(n) = O(n^{3/2})$: the Euclidean unit-distance graph is $K_{2,3}$-free, since two unit circles meet in at most two points, and the Kővári–Sós–Turán theorem applies [KST54]. The best known upper bound $\nu(n) = O(n^{4/3})$ is due to Spencer–Szemerédi–Trotter [SST84]; Székely later gave a short crossing-number proof of this incidence bound [Sze97].

Feedback

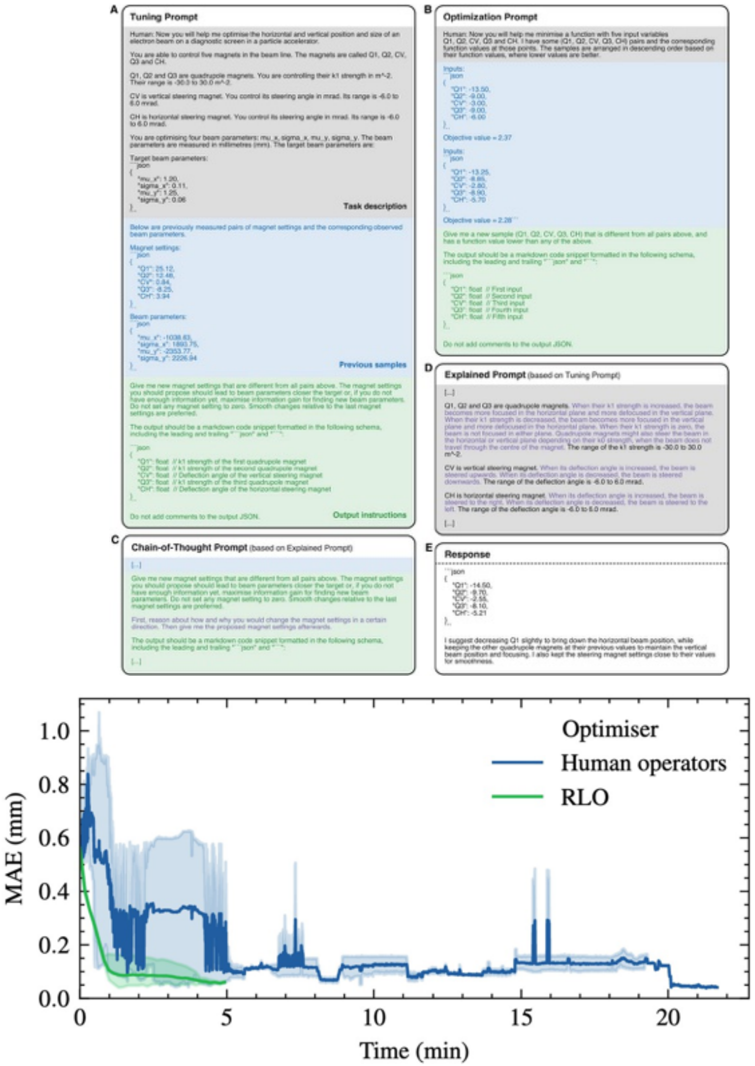

Helmholtz common vision of autonomous facility operation

A shared approach across research centers and disciplines

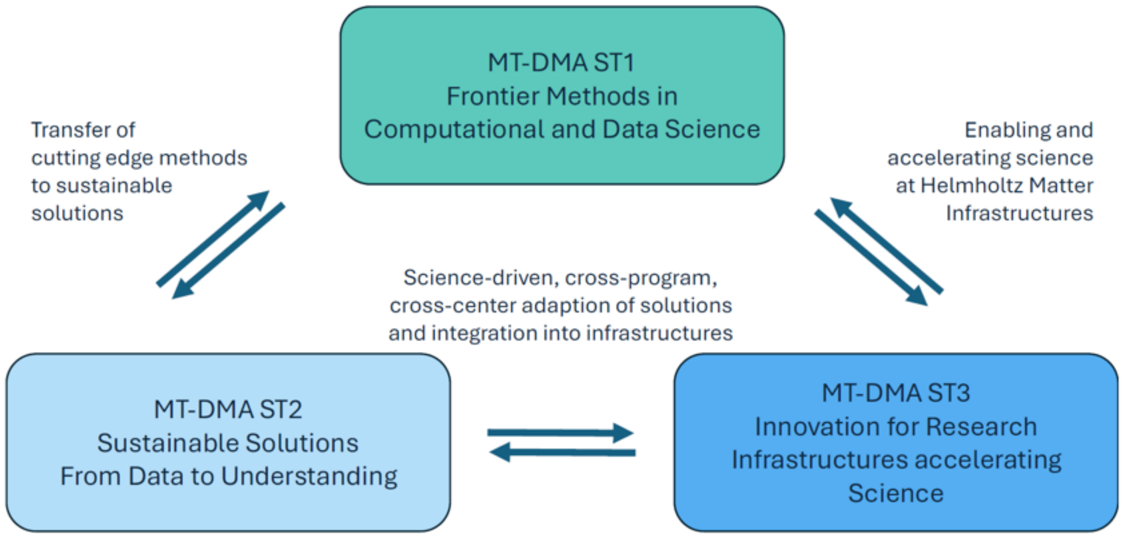

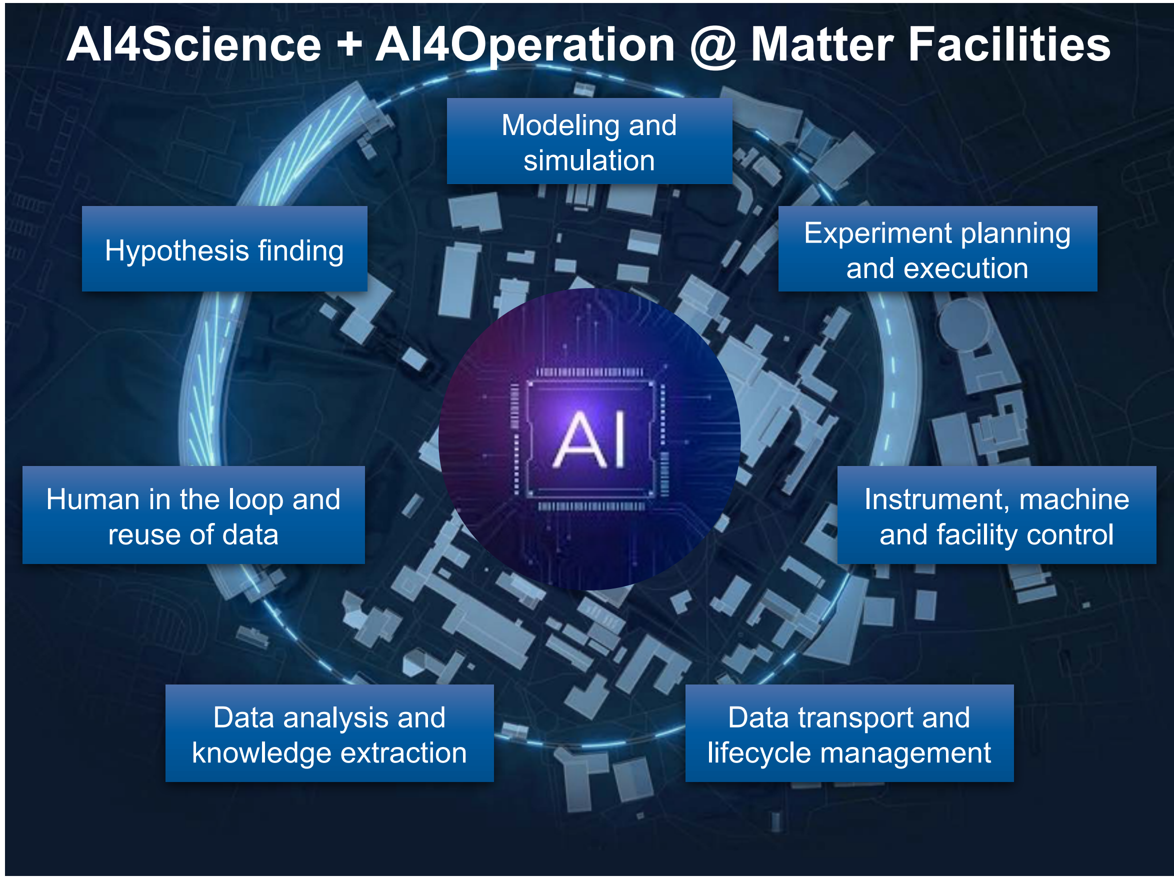

DRESDEN concept HZDR

Autonomous Research Facilities

World models include the machines, detectors, models, scientists, …

- **A well defined understanding of observables**
  - A clear ontology
  - Models with predictive capability
  - How to measure and quantify observables
- **A complete knowledge of your system**
  - Digital twins of systems under investigation
  - Variation over all possible outcomes
- **Digitalization of facility operation and experiments**
  - A quantifiable and steerable model of your facility
  - Include all known sources of errors
  - A complete error analysis
  - Invertible synthetic diagnostics
- **Useable and useful F.A.I.R. workflows**
  - for simulation
  - for data analysis
  - for visualization
  - for operation and experiments

Autonomous Research Facilities

World models include the machines, detectors, models, scientists, …

Lawrence Livermore National Laboratory

NVIDIA

HZDR
HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Dagstuhl Seminar 23132 / Dagstuhl Reports 13
AI-Augmented Facilities:
Bridging Experiment and Simulation with ML

HPC+AI 4 Fusion

Frontier computing for fusion science

Joining forces in Europe, bringing together magnetic and laser-driven fusion

EuroHPC Center of Excellence for Plasma Simulations PLASMA-PEPSC

- Brings **laser fusion and magnetic fusion** community together
- Common modeling efforts for atomic physics, looking at **plasma-wall interaction** (magnetic fusion) and **radiation transport** (laser fusion)

How to run on every TOP 500 System in record tim

Performance-portable parallel programming on heterogeneous systems

Alpaka: The performance portable C++ abstraction library for HPC

- “The European Kokkos”
- SIMD + SIMT
- Write code one time and run it everywhere
- Supports GPU (Nvidia, AMD, Intel) and CPU (x86, ARM, RISC-V)
- Minimum overhead
- Follows the default GPU offload model
- Single source C++20 user friendly API
- Allows to use native vendor APIs like CUDA, HIP and OneAPI

GH 200

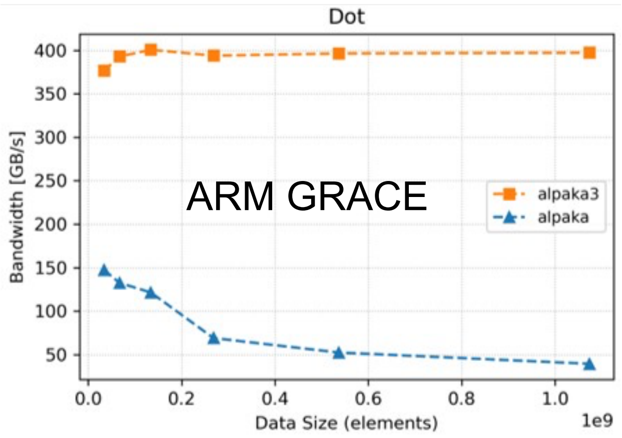

Large-scale simulations for fusion research

Making the best of the worlds fastest supercomputers

JÜLICH Forschungszentrum

Institute

Jülich Supercomputing Centre (JSC)

News Research Systems Services Education About us

JUPITER - Exascale for Europe / **Applications: Make JUPITER shine**

Applications: Make JUPITER shine

Applications are a central aspect in the exascale strategy of the Jülich Supercomputing Centre (JSC). Concrete use cases computed on JUPITER have the potential to advance developments in important areas of society and, in particular, to significantly accelerate change due to JUPITER's enormous total computing power. JUPITER will achieve more than 1 ExaFLOP/s for double-precision computing operations and more than 70 ExaFLOP/s for 8-bit precision, relevant for AI applications in particular.

JUPITER Benchmark Suite

Applications played an important role in the JUPITER procurement process. The JUPITER Benchmark Suite was developed to bring a system into production that not only performs very well in synthetic benchmarks but also in real and relevant use cases. Based on current and projected important applications, a collection of 16 applications was compiled and used for the performance evaluation of possible JUPITER configurations. Andreas Herten, lead author of the corresponding scientific paper, comments: "With the JUPITER Benchmark Suite, we have created a milestone for the a-priori evaluation of supercomputers. The combination of standardization and real applications is unique and allowed a targeted projection of the expected performance."

JUPITER Research and Early Access Program

For applications, JUPITER is not only an opportunity but also a challenge. To use a supercomputer of such size to its full potential, applications must be thoroughly prepared and optimized. This can range from making adjustments to algorithms and methods to optimizing individual settings. Such optimizations are currently being carried out as part of the JUPITER Research and Early Access Program (JUREAP). In JUREAP, experts from JSC are working with more than 100 applications to assess and prepare them for exascale. The applications are highly diverse and have the potential to fundamentally advance many different scientific disciplines. They range from AI applications, such as foundation models and generative video generation, to climate models, particle physics and energy applications, to molecular dynamics simulations relevant for drug development and disease control. Some examples are briefly presented below.

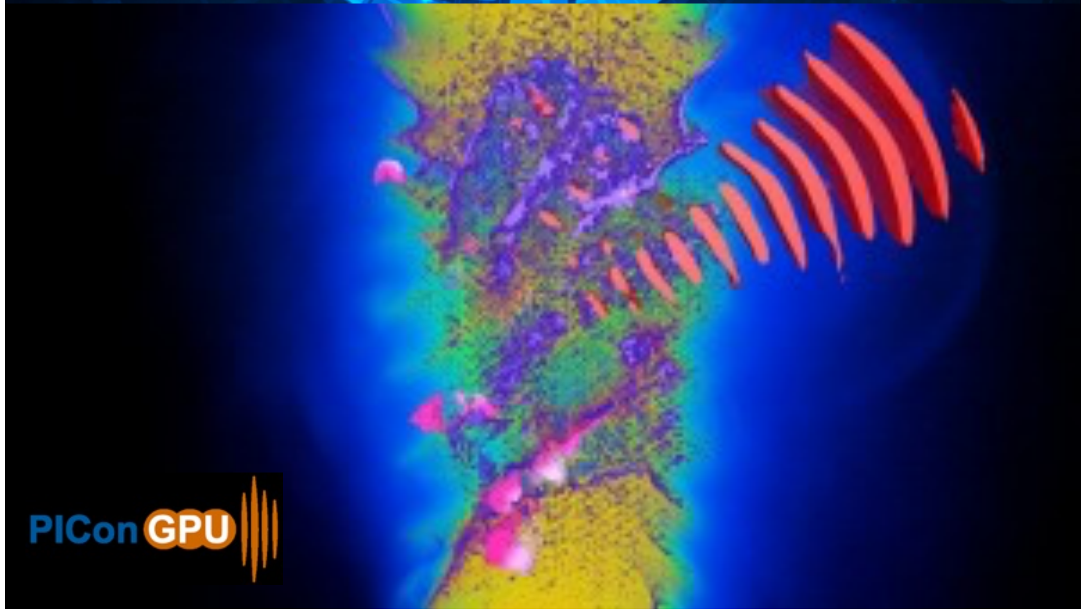

Large-scale simulations for fusion research

Making the best of the worlds fastest supercomputers

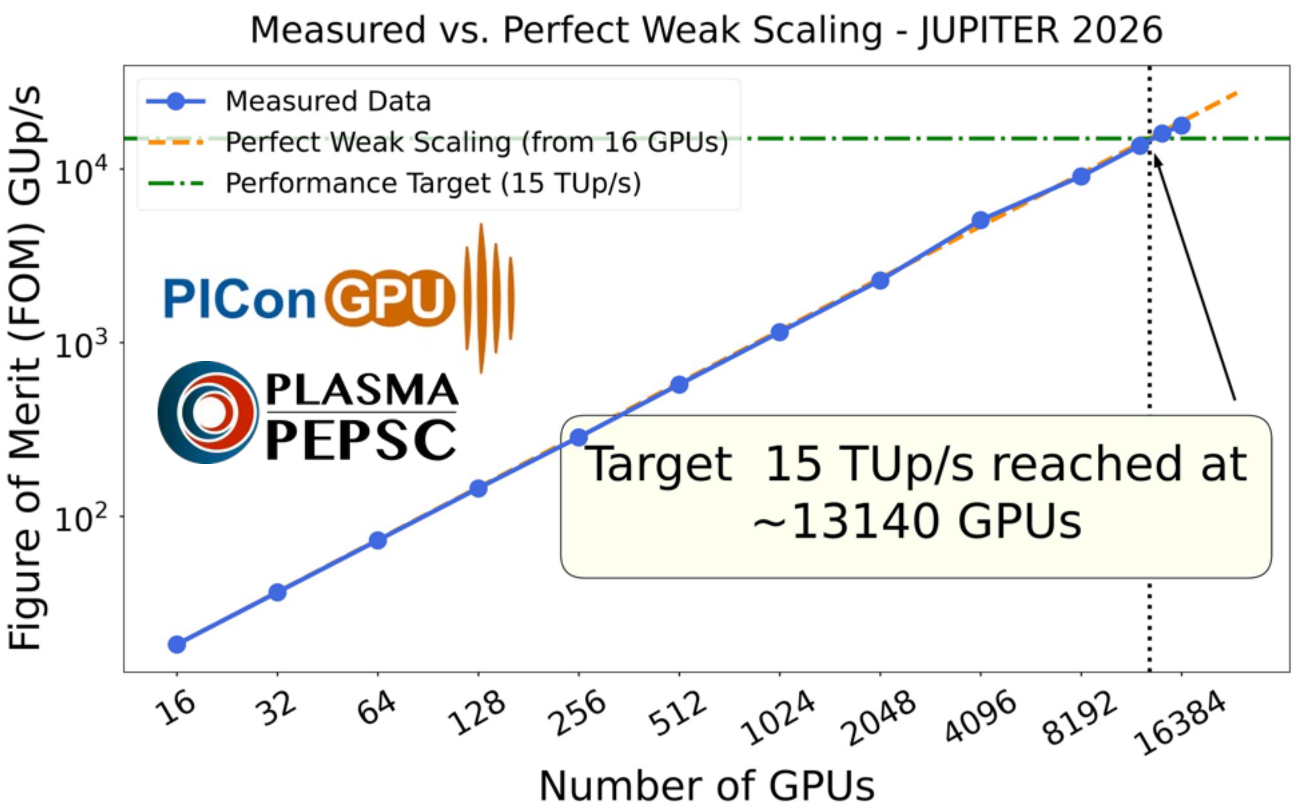

DRESDEN concept HZDR

High energy density physics
@ fs & nm resolution
@ PB/s data rates
@ 20 fps
@ o(10) MW power consumption

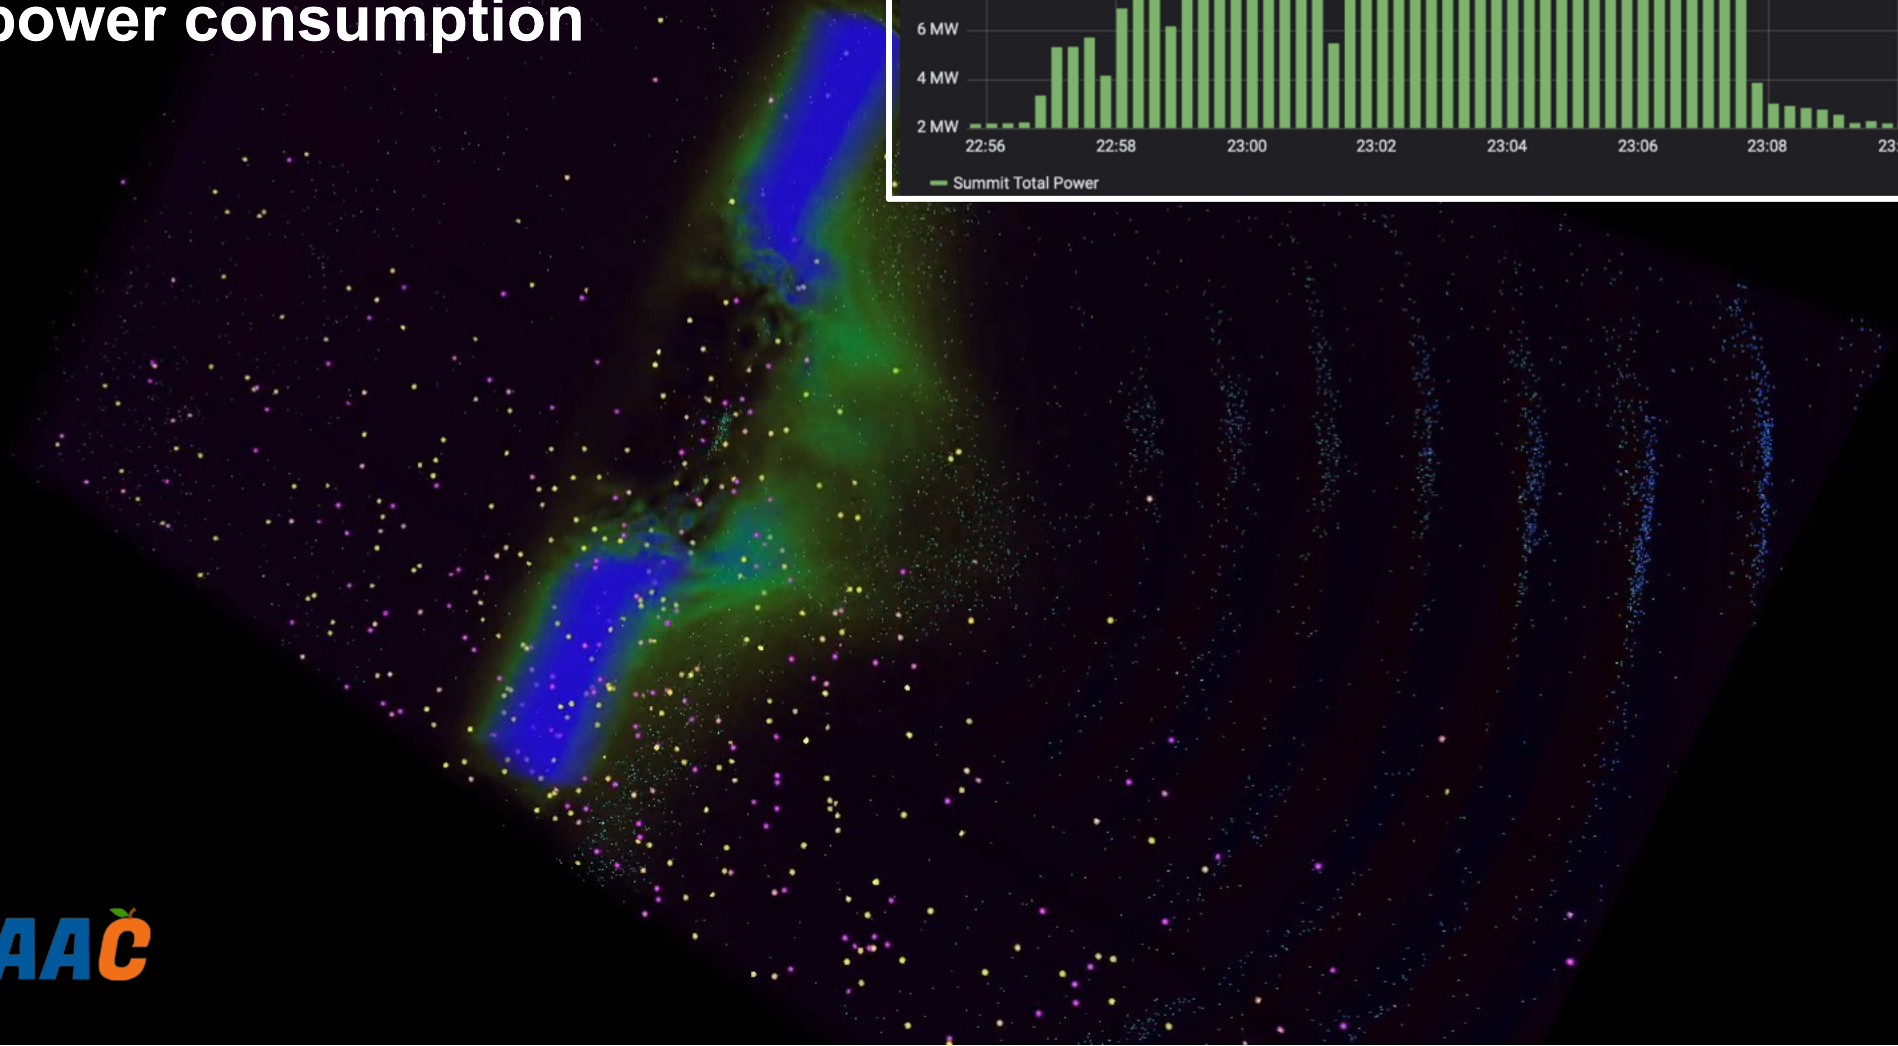

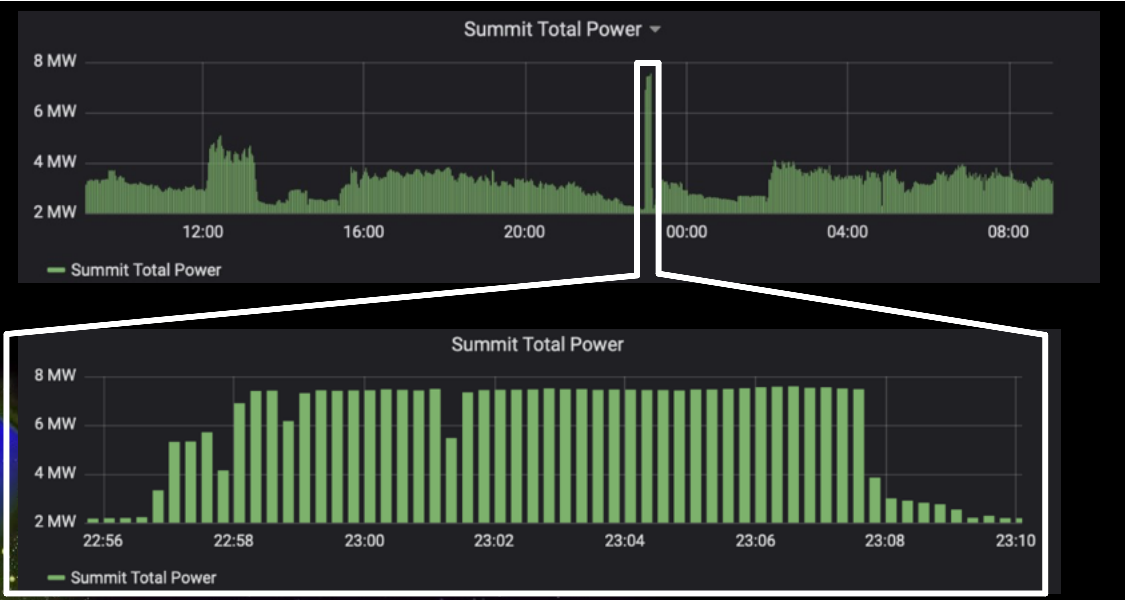

World models combining experiment and simulation

A well defined understanding of observables: Simulate what is measured

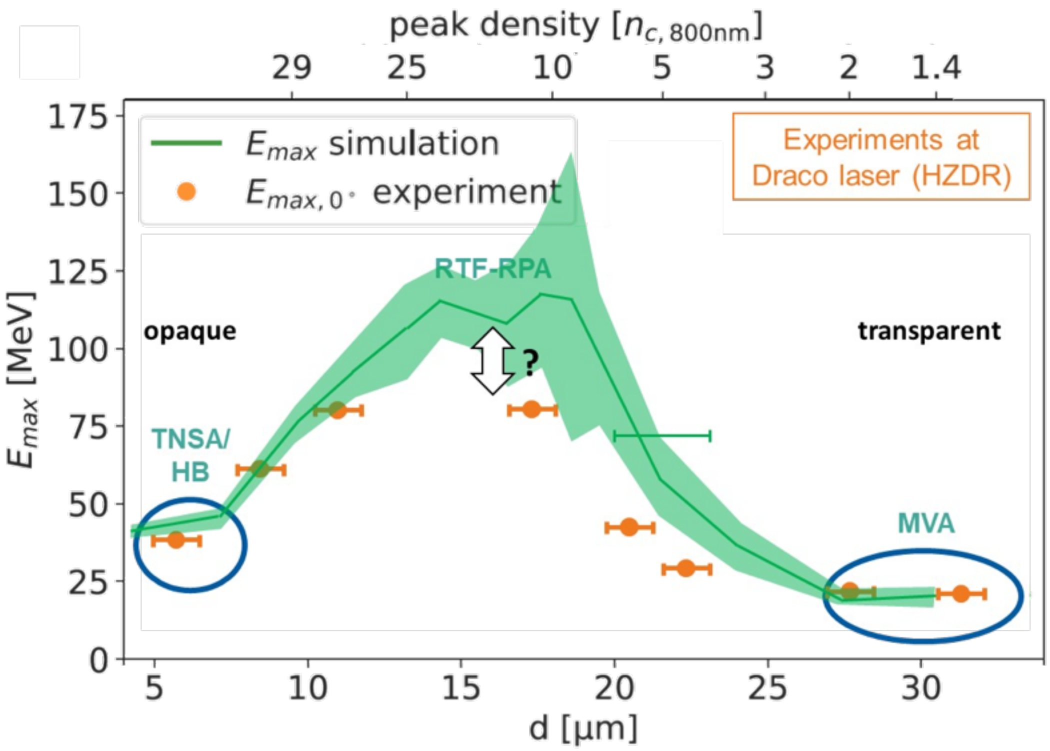

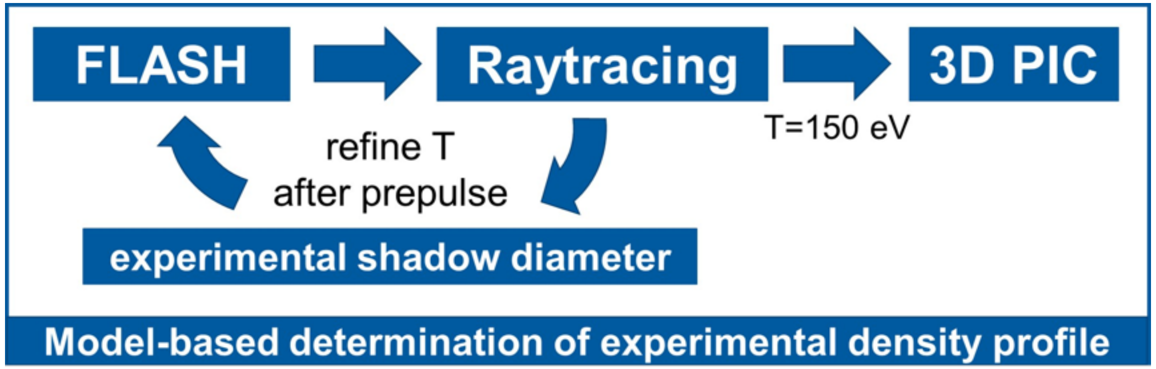

T. Kluge + K. Zeil and many others

HiBEF

HELMHOLTZ INTERNATIONAL
BEAMLINE FOR EXTREME FIELDS

World models combining experiment and simulation

A complete knowledge of your system: Atomic resolution of high energy density matter

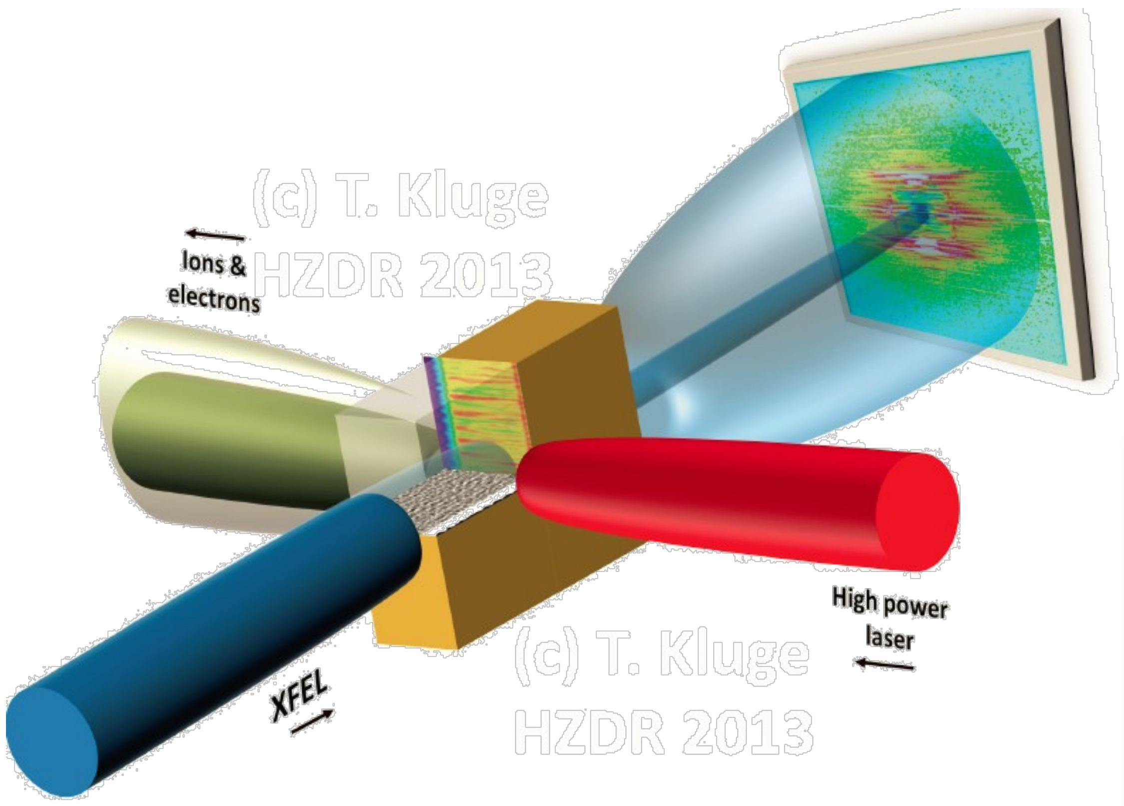

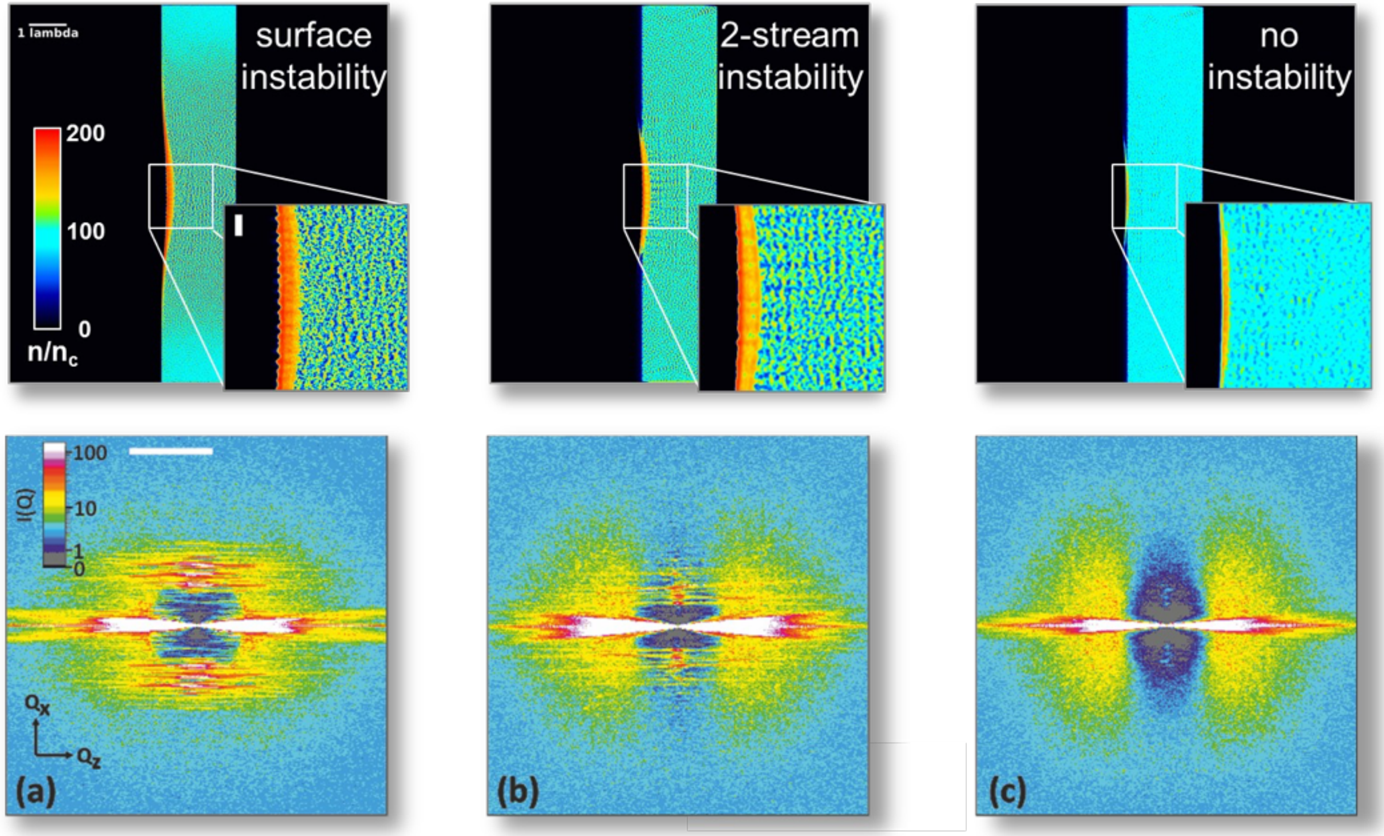

T. Kluge, L. Huang

World models combining experiment and simulation

Digitalization of facility operation and experiments: Invertible models for discovery

Invertible physics-informed neural networks

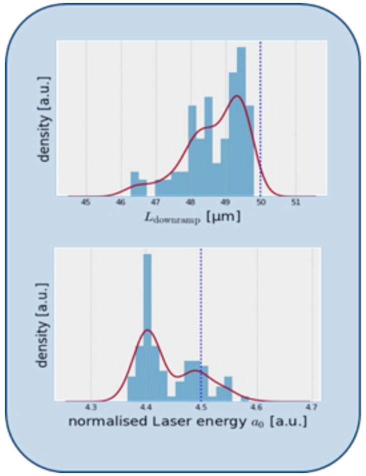

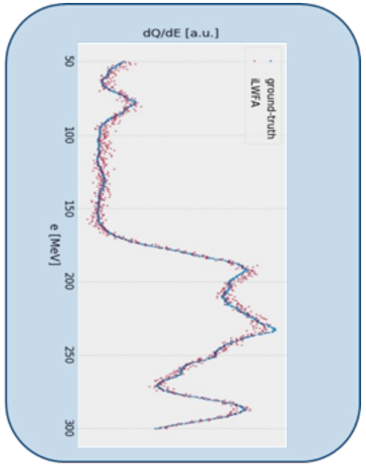

Benefits

- Recover ambiguous mapping
- Uncertainty quantification

Multi-physics, multi-modal surrogate modeling

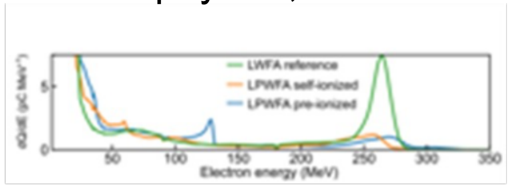

Particle energy spectra

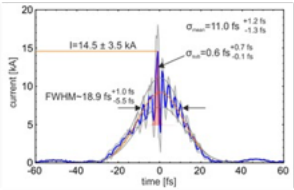

Transition radiation

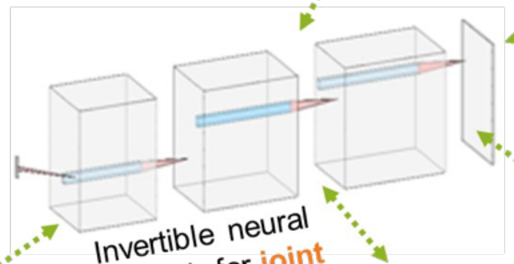

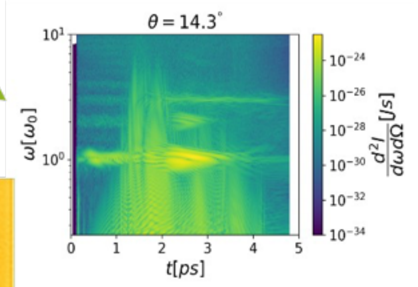

Radiation spectra

ML-based solvers

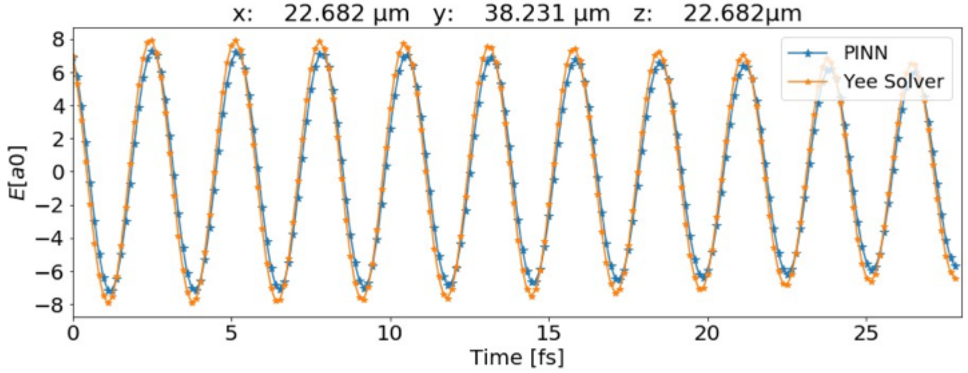

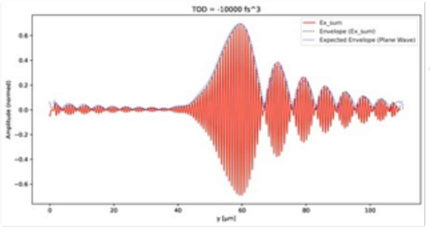

Experimental parameters

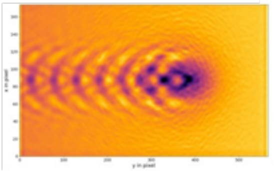

Shadowgraphy

DRESDEN concept HZDR

World models combining experiment and simulation

Digitalization of facility operation and experiments: Campaign-driven optimized operation

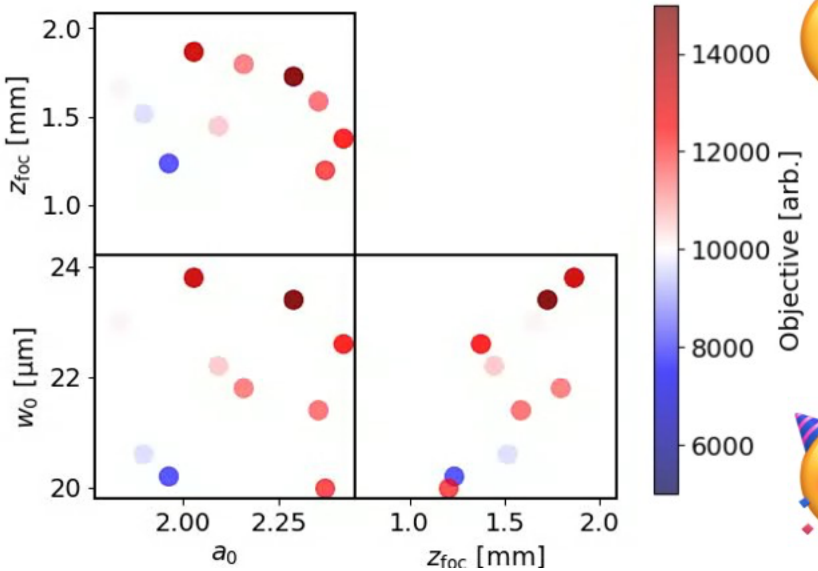

Each dot represents a full-scale HPC simulation

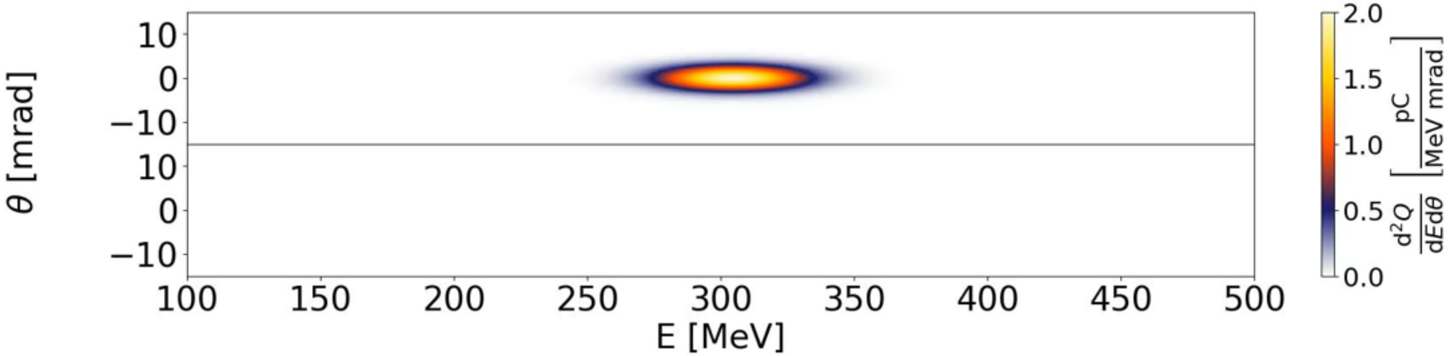

World models combining experiment and simulation

Useful and useable workflows: Community-driven input/output standards for workflows

PICMI Standard → In → PIConGPU → Out → openPMD

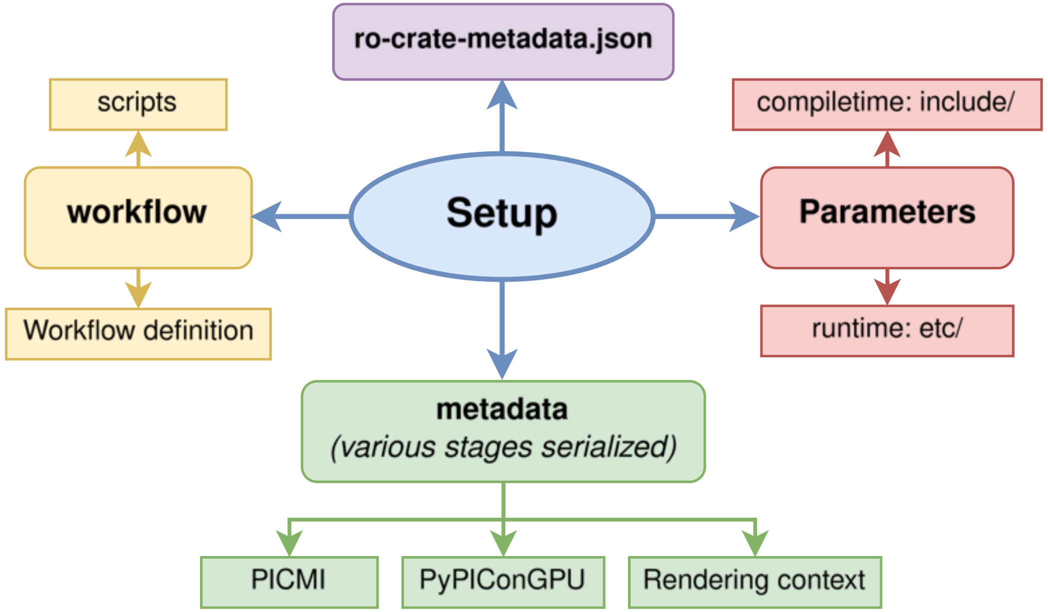


```python
from picongpu import picmi

sim = picmi.Simulation(
    time_step_size=9.65531e-14,
    max_steps=1024,
    solver=picmi.ElectromagneticSolver(
        method="Yee",
        grid=picmi.Cartesian3DGrid(...)
    ),
    ...
)

sim.run()
```

Yet, data storage is a thing of the past!

Data rates become too high, volumes too large for storage

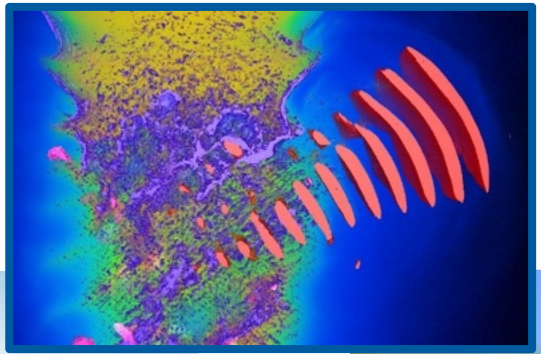

0.6 PetaByte / s

0.1TB/s GPU to GPU 16GB/s internode

file System: 679 PB

full in 100s

but simulation runs longer

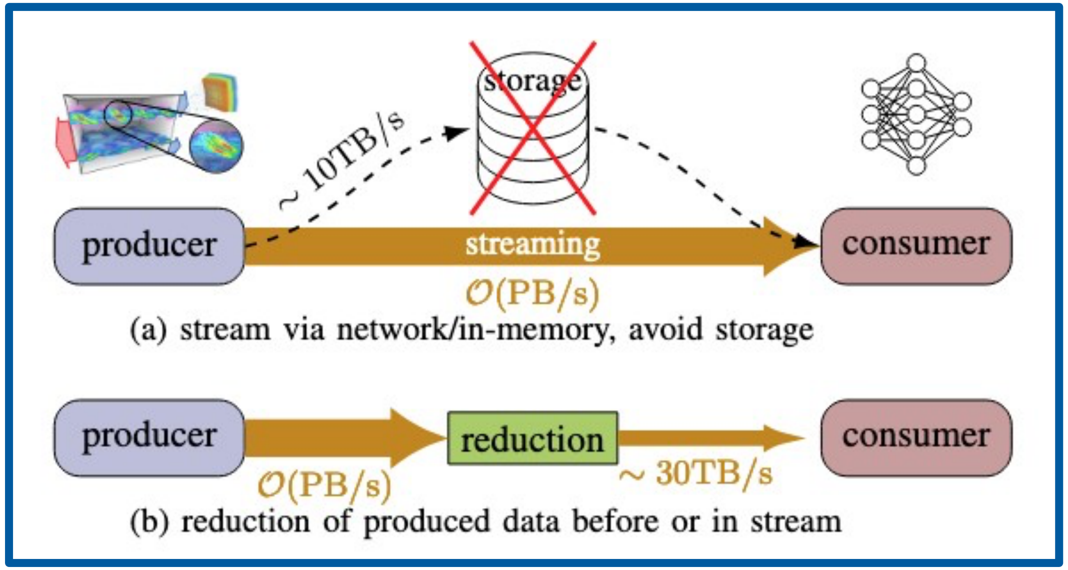

Coupling HPC with AI training at Exascale

Training of a large-scale invertible AI model with streaming simulation data

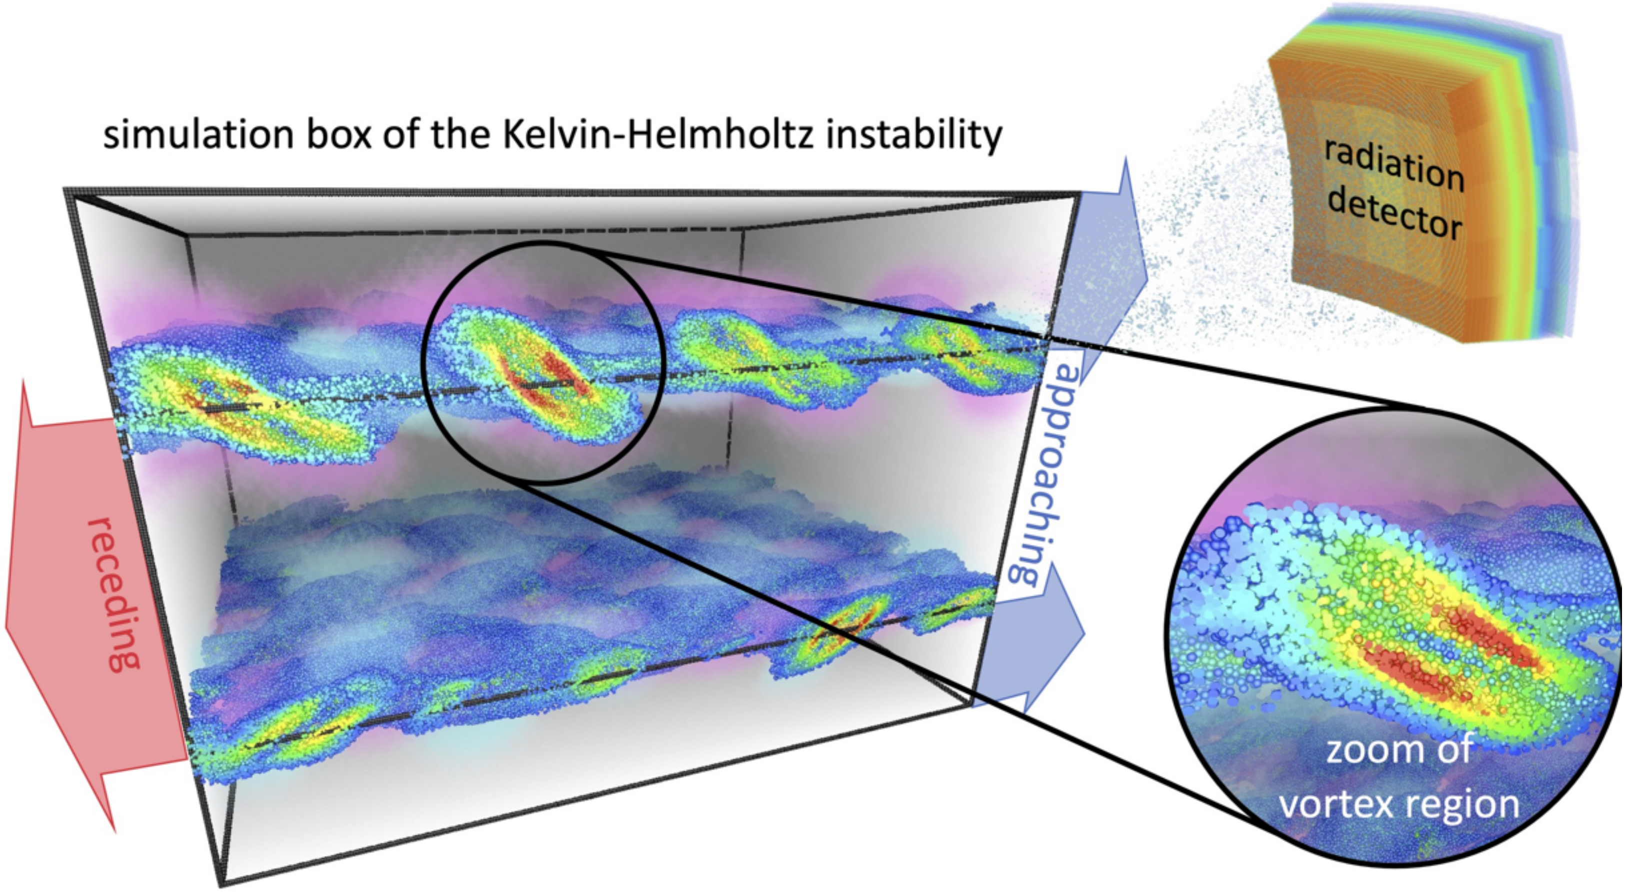

Train neural networks for:

- Forward problem:
  particle to radiation
- Compression:
  particle to particle latent space
- Inverse problem:
  radiation to particle

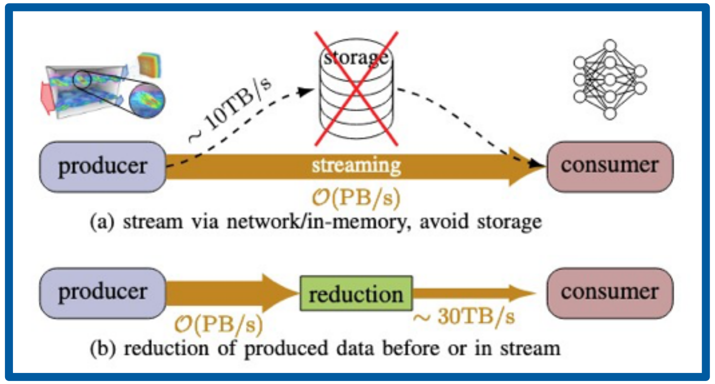

Online training of streaming data @ Exascale

Why use files when you have neural networks?

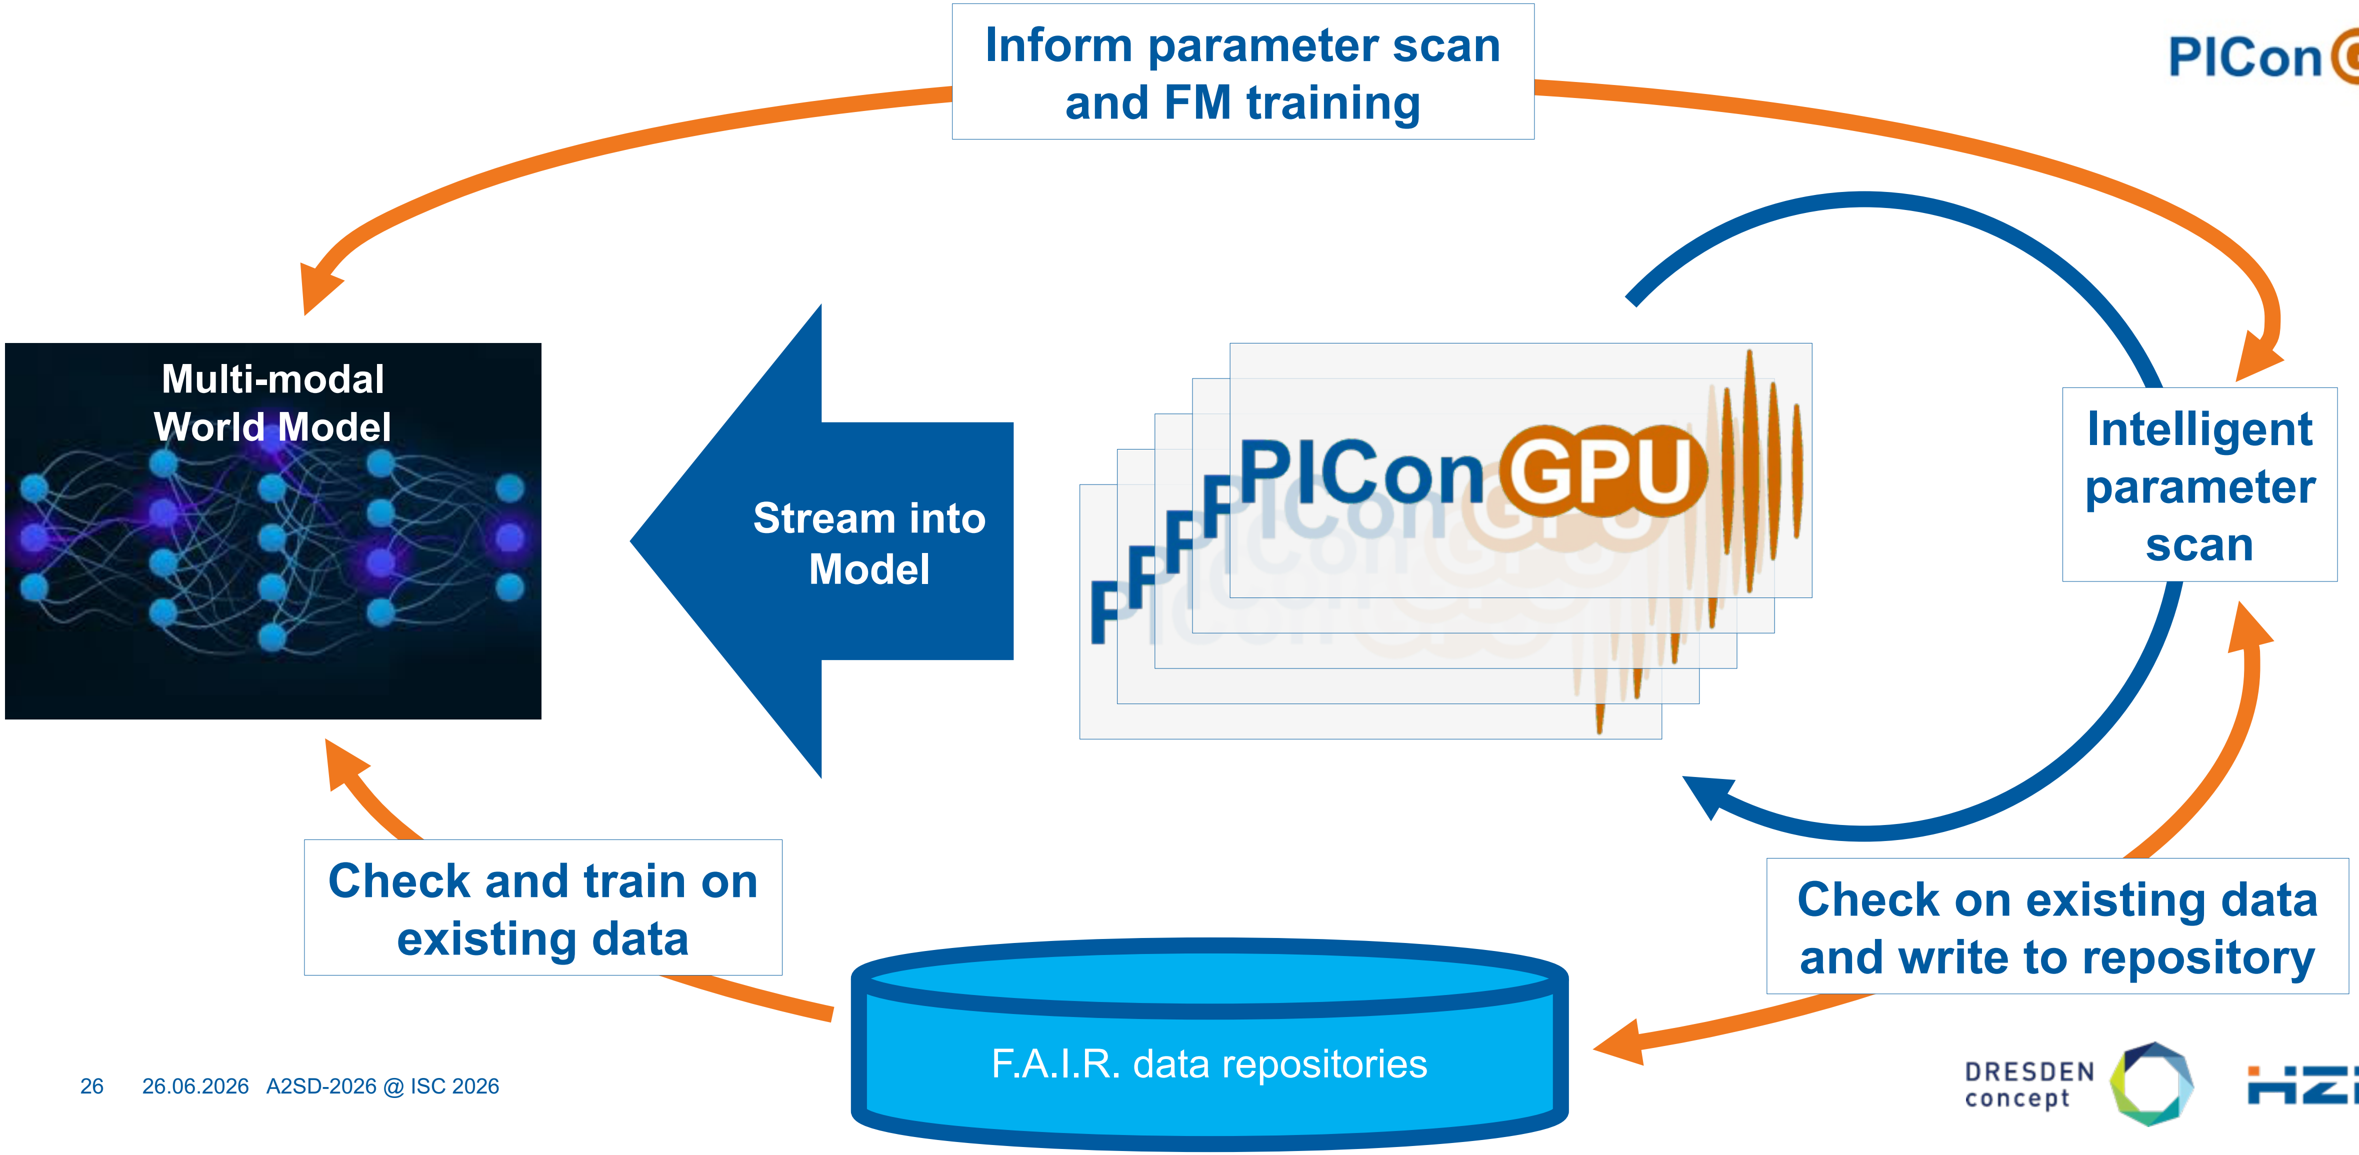

Online training of streaming data @ Exascale

Why use files when you have neural networks?

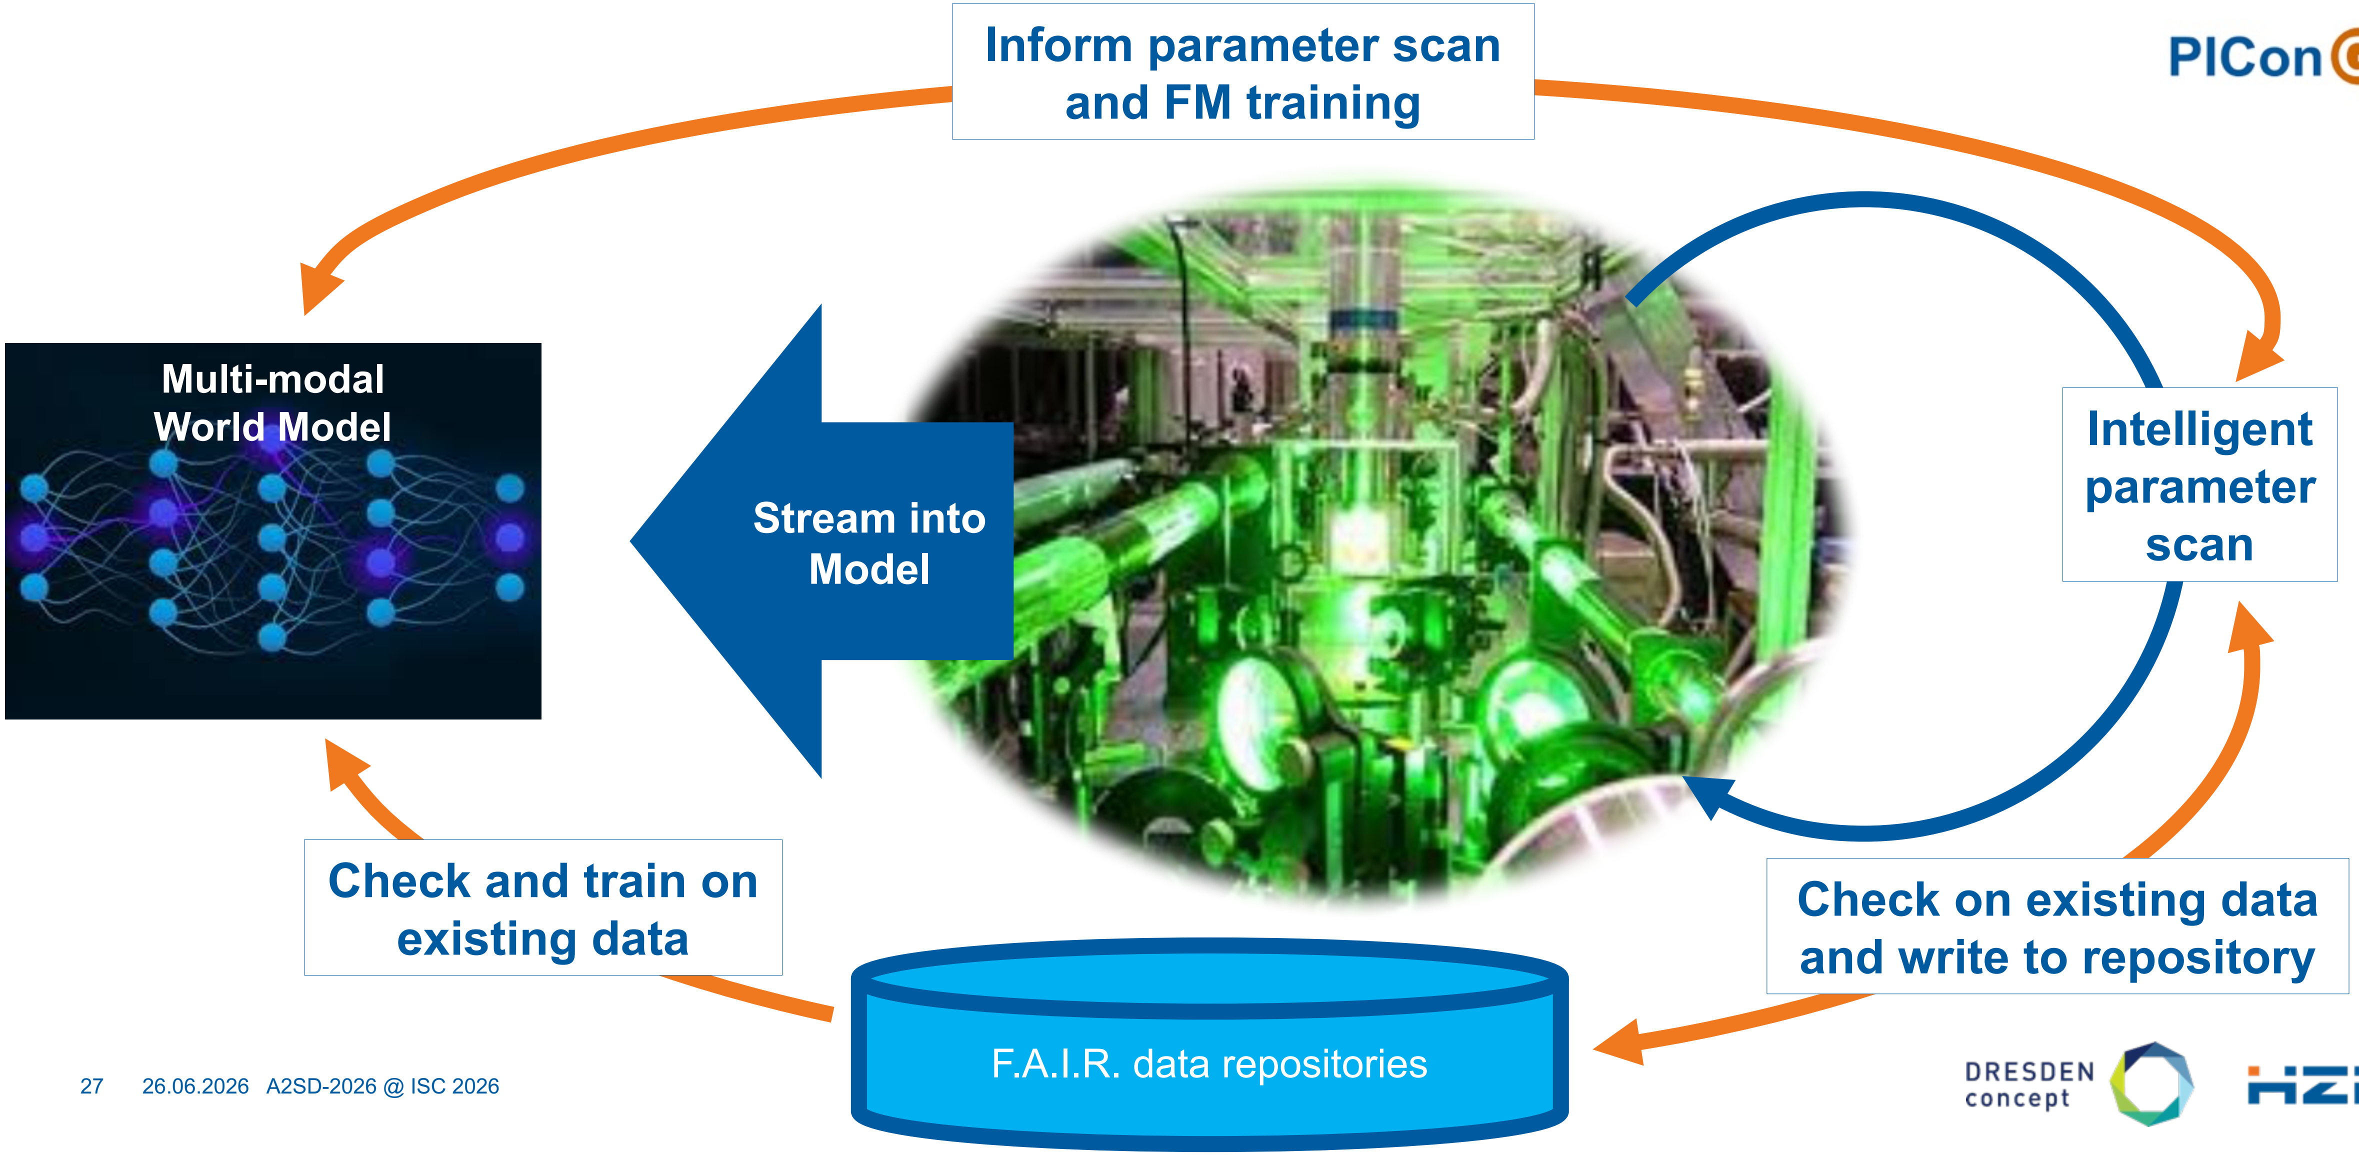

Make it so!

Agentic AI puts the user in the center, not the data

Scientists know stuff and like to tinker

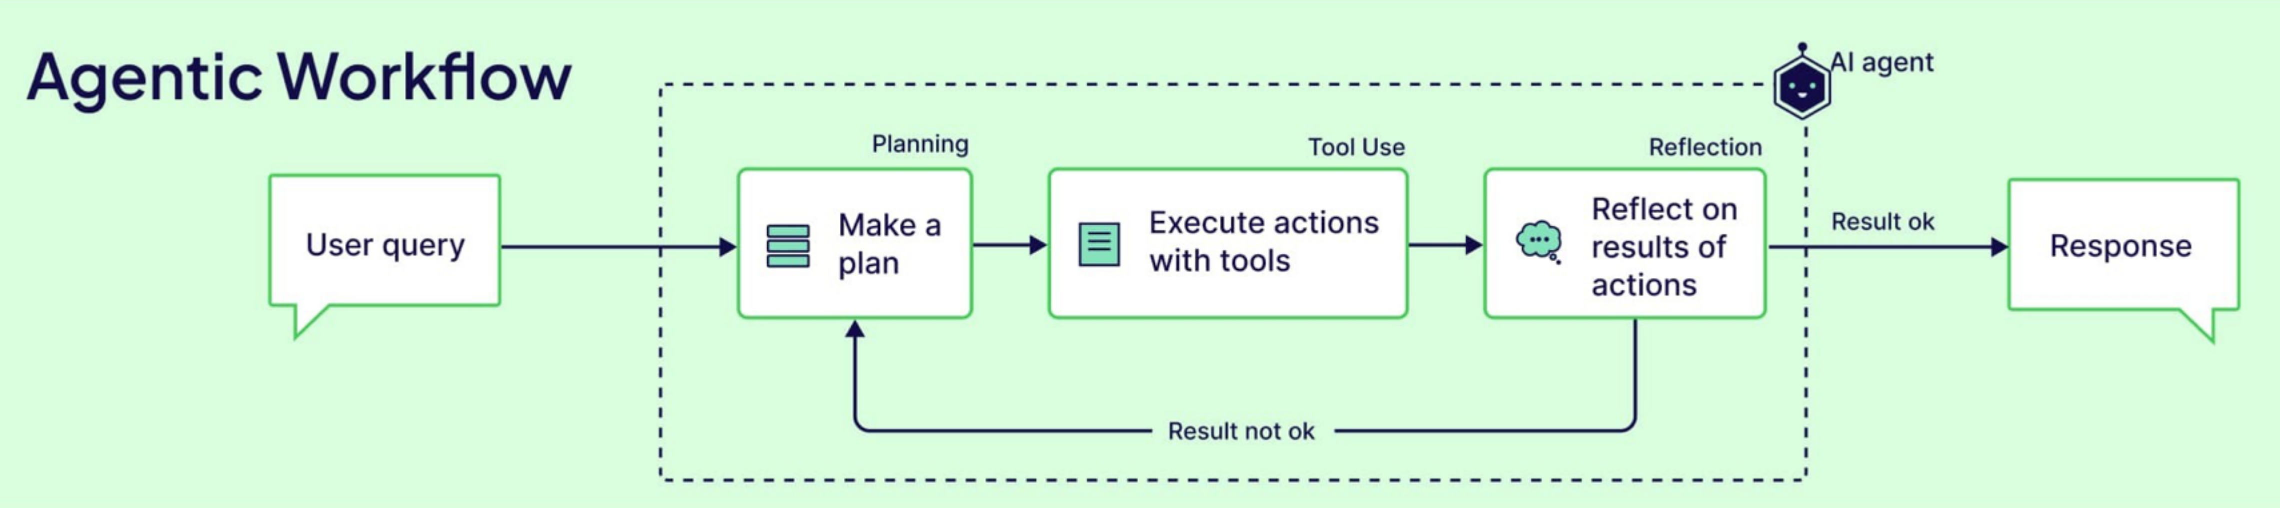

Agentic AI solves:

- Diverse user needs
- Easy usage

Agentic AI helps with:

- Tooling
- Interfaces
- Interoperability

Agentic AI is (currently) limited in:

- Autonomous discovery*
- Optimization**

** Our research facilities are not (yet) flexible enough for agentic planning (all hail our robot overlords!)*
*** In many cases, an optimum operation point stays an optimum operation point*

DRESDEN concept HZDR

Stop selling shells on the seashore… sorry, in HPC!

Once we have clear interfaces, HPC clusters are just tools queried by agents

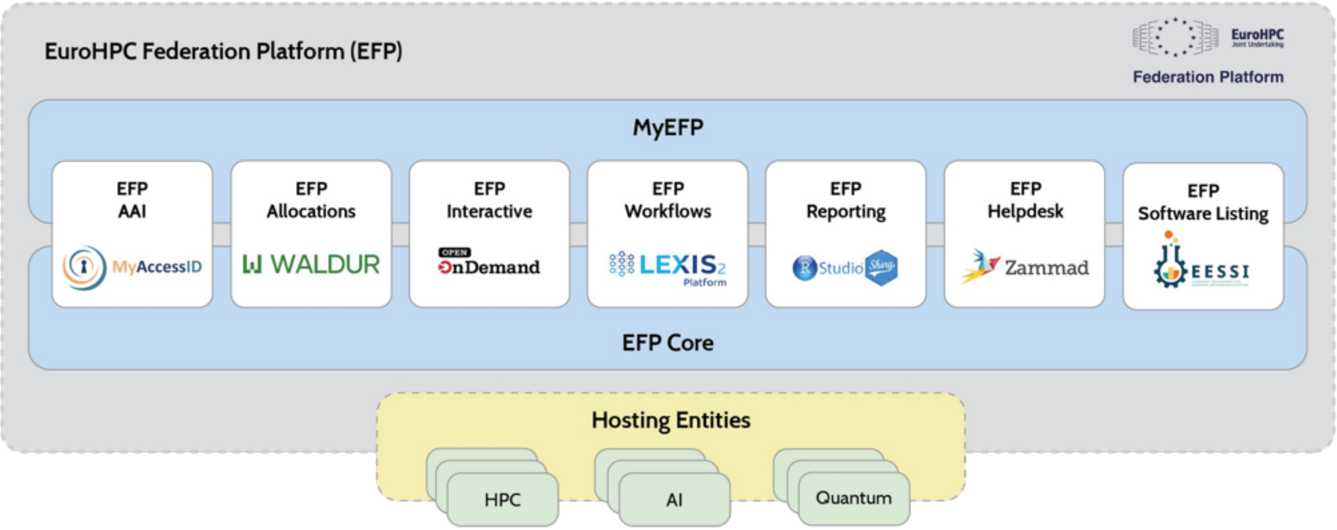

The three facility pillars: Operation, Science, User management

In Europe, there is a TIER-Systems of Compute/Data Centers, EuroHPC is still orthogonal to that

Agentic AI interface and knowledge base

Operation	Science	User management
▪ Quality control ▪ Optimization ▪ Predictive maintenance ▪ Digital twins	▪ Data analysis ▪ Feedback ▪ Mission / Survey ▪ Discovery	▪ Proposal writing ▪ Proposal review ▪ Experiment preparation ▪ Occupancy / usage

Applications and Services

DRESDEN concept

Make it so!

Autonomous user facilities as tools for scientists

Learnings:

- Analysis to Data
- Federation creates trust
- Include user communities
- Start with usage, not data
- Make data intake automatic
- Make it safe & useful

Authentication and authorization

AI-augmented user interface

AI applications: Fusion | Health | Materials | ...

Federation
One portal across facilities: *Applications to data, Interoperability*

Agentic AI Knowledge Base
Sovereignty: *On-premise, self-hosting*

AI Ops: *Reusable, reproducible AI*

AutoML: *Easy AI development*

Data Management

Workflow Management

Secure access and control layer to facility