Towards Autonomous and Intelligent Research in Laser-Plasma Physics

State of the art and a look forward

Michael Bussmann · HZDR

Autonomous, intelligent research facility operation

A strategy across centers and programs in Helmholtz

We integrate digital solutions for innovating facility operation and accelerating science

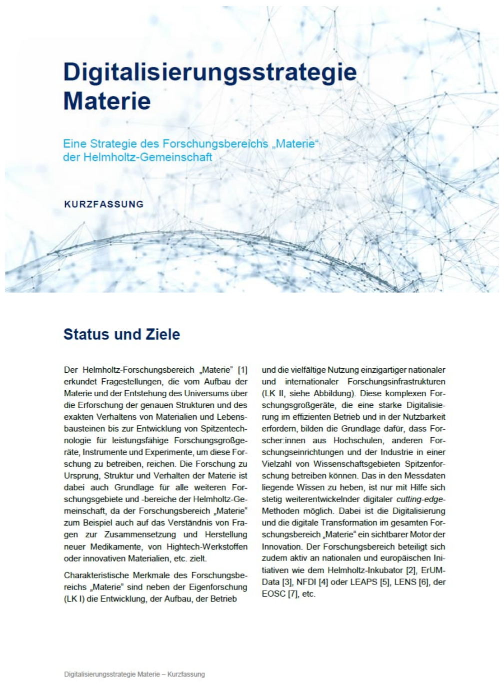

Digitalisierungsstrategie Materie

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

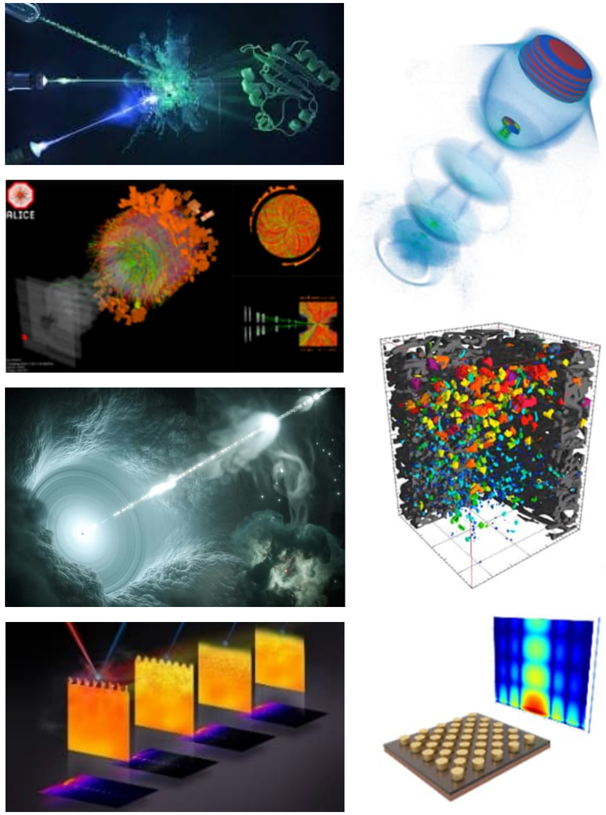

	LK II (User facilities)	International Research Infrastructures	National Research Institutes/ Infrastructures
DESY	FLASH PETRA III IDAF	LHC Belle II CTA (*under construction*) IceCube European XFEL ESRF	CFEL CSSB NanoLab DESY Test Beams DAF HIB@European XFEL PITZ
FZJ	JCNS (in MLZ)	ESS (*under construction*) 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

DRESDEN concept HZDR

Integration of digital solutions in the research life cycle

Driving scientific discovery and creating a competitive advantage for our research facilities

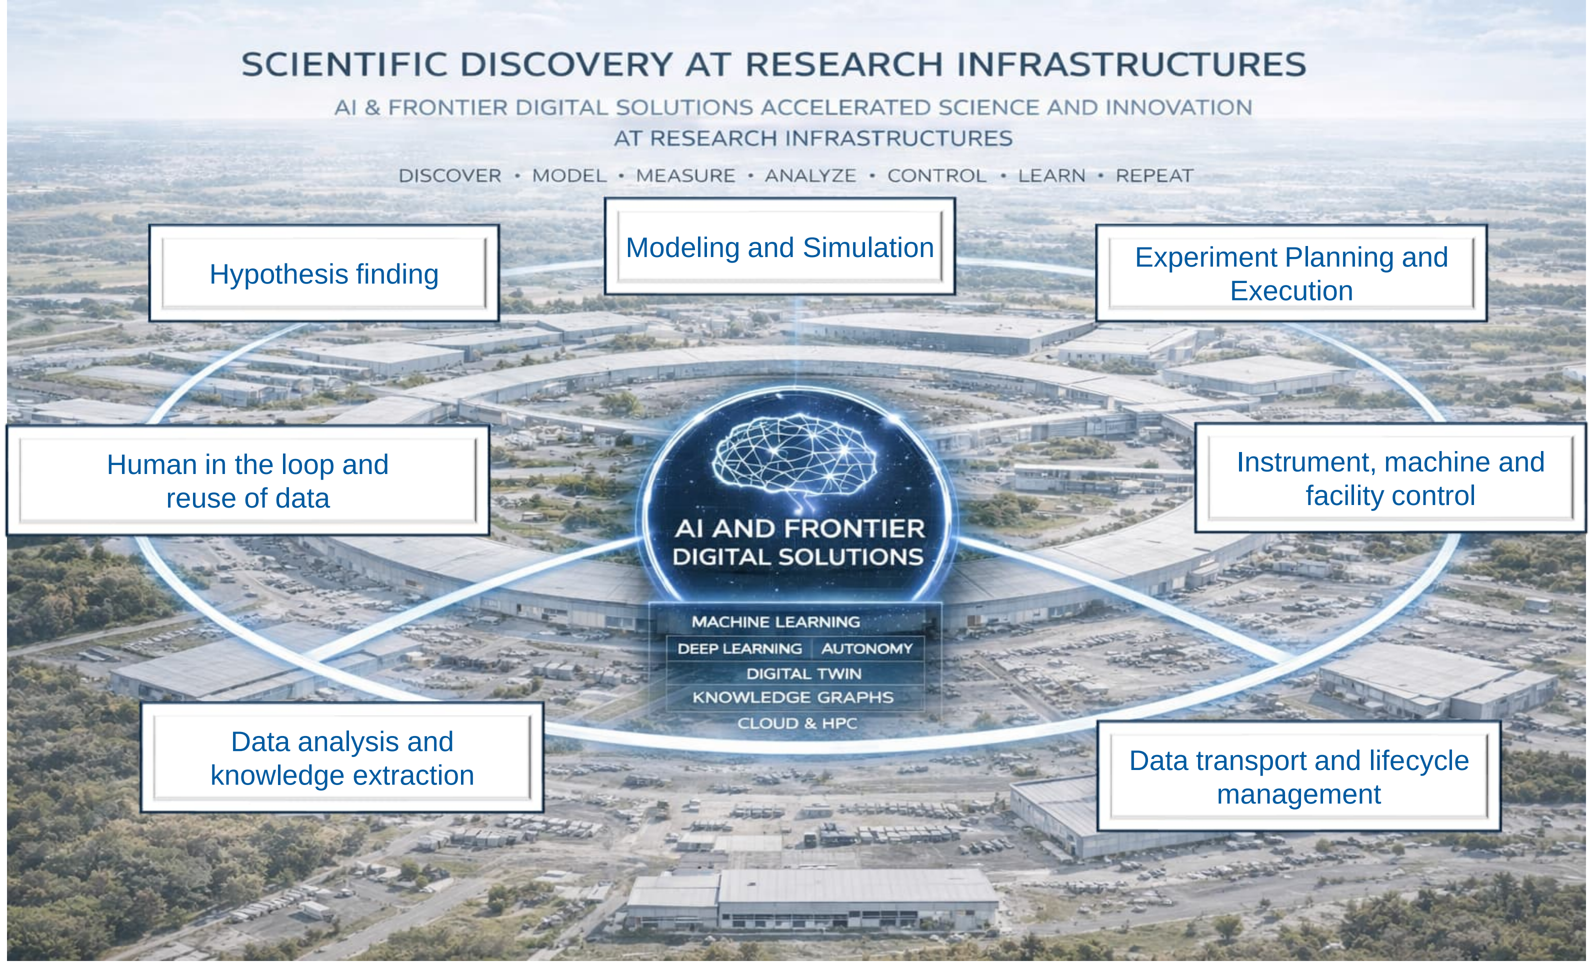

Joining forces on the national and international level

Helmholtz, NFDI, ErUM, EOSC and beyond

HELMHOLTZ-INKUBATOR INFORMATION & DATA SCIENCE

HELMHOLTZAI

HIFIS

<HMC>

HELMHOLTZ IMAGING

HiDA

nfdi Nationale Forschungsdaten Infrastruktur

DAPHNE 4NFDI

PUNCH 4NFDI

ErUM-Data IDT

EUROPEAN OPEN SCIENCE CLOUD

ExPaNDS European Open Science Cloud Photon and Neutron Data Services

panosc photon and neutron open science cloud

ESCAPE

WLCG Worldwide LHC Computing Grid

LEAPS

LENS

EuroHPC Joint Undertaking

RDA RESEARCH DATA ALLIANCE

HSF

Laser fusion success increases interest

A joint vision for research facilities

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

März 2023

Following up at ISC 2026 (Where are the Germans?)

AI and facility advances set a high pace

The A2SD-2026 Workshop brings together researchers and practitioners shaping the future of AI-driven, interconnected, and automated scientific ecosystems. As instruments produce complex data streams at rates beyond human response, A2SD-2026 focuses on connected, intelligent, and verifiable autonomous systems that can operate across scientific facilities.

Building on the A2SD-2025 community roadmap and follow-up activities, the workshop highlights five key areas:

1. Architectures and standards for interconnected autonomous laboratories at scale;
2. Agent-based AI for scientific orchestration (including LLM agents with grounding and verification);
3. Distributed data fabrics and AI-ready infrastructure (FAIR-aligned, real-time metadata, interoperability);
4. Digital twins and closed-loop experiment-to-compute integration leveraging HPC; and
5. Standards, trust, safety, and workforce development for responsible autonomous science.

Interconnected autonomous labs | Agent-based AI orchestration | AI-ready data fabrics (FAIR) | Digital twins + closed loop HPC | Trust, safety + workforce

ISC High Performance — The HPC Event.

Hamburg, Germany

Date: **June 26, 2026**
Time: **2:00pm-6:00pm CEST**
Room: **Hall 10 - 1st Floor**

All workshops

Organizers

Rafael Ferreira da Silva
ORNL

Tom Gibbs
NVIDIA

Michela Taufer
UTK

Frontier computing for frontier 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)

Plasma Accelerators and Free Electron Lasers

Laser-Wakefield Acceleration

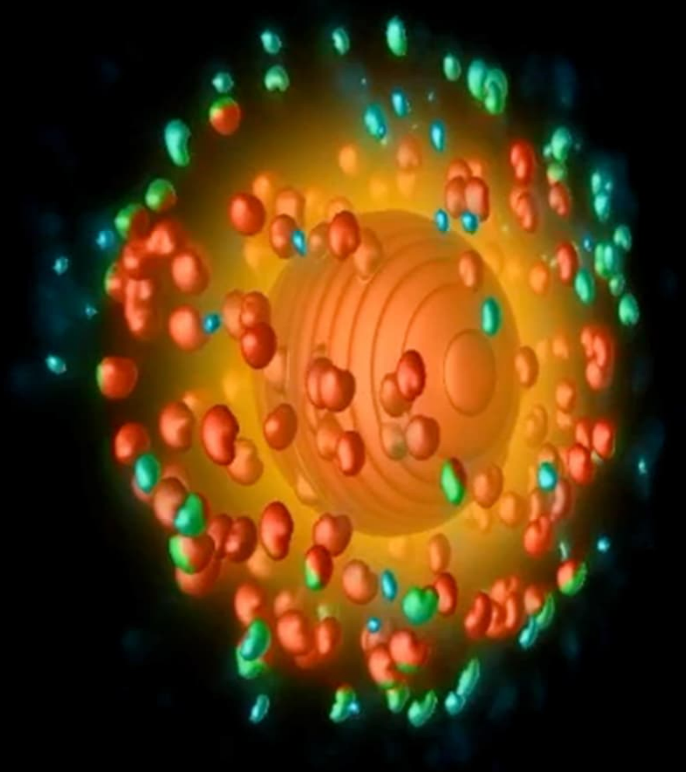

ISAAC PIConGPU

Combining laser and electron generated wakefields

High peak current, high quality electron beams as drivers for compact light sources

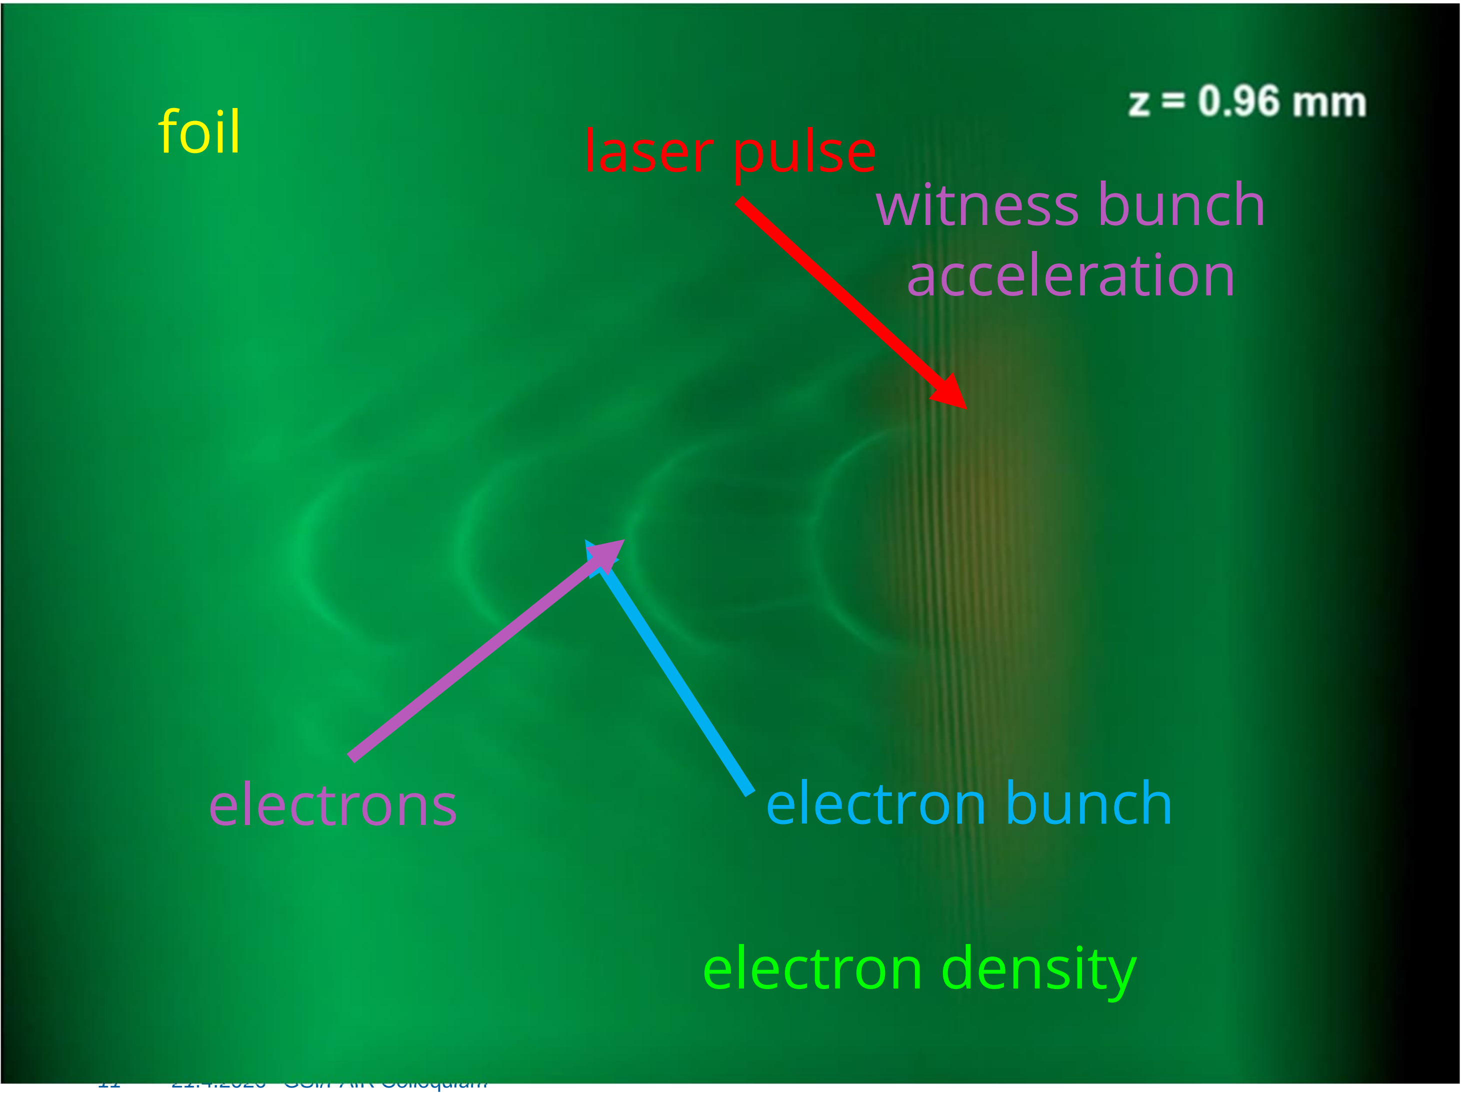

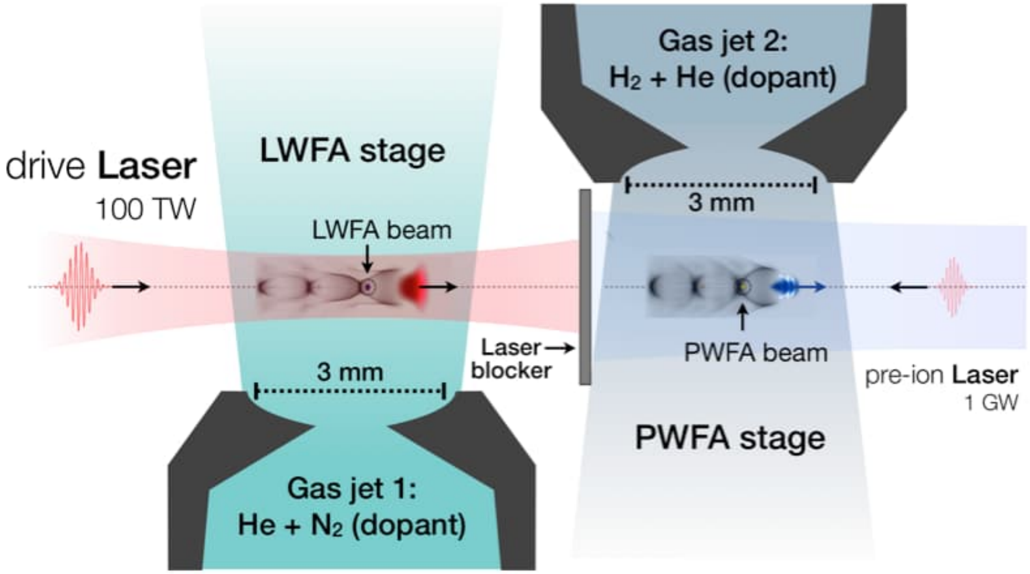

R. Pausch, A. Debus

Labat et al., Seeded free-electron laser driven by a compact laser plasma accelerator

Shrinking particle accelerators from km to cm

Plasma-accelerated of electrons drive compact X-ray free electron lasers (XFELs)

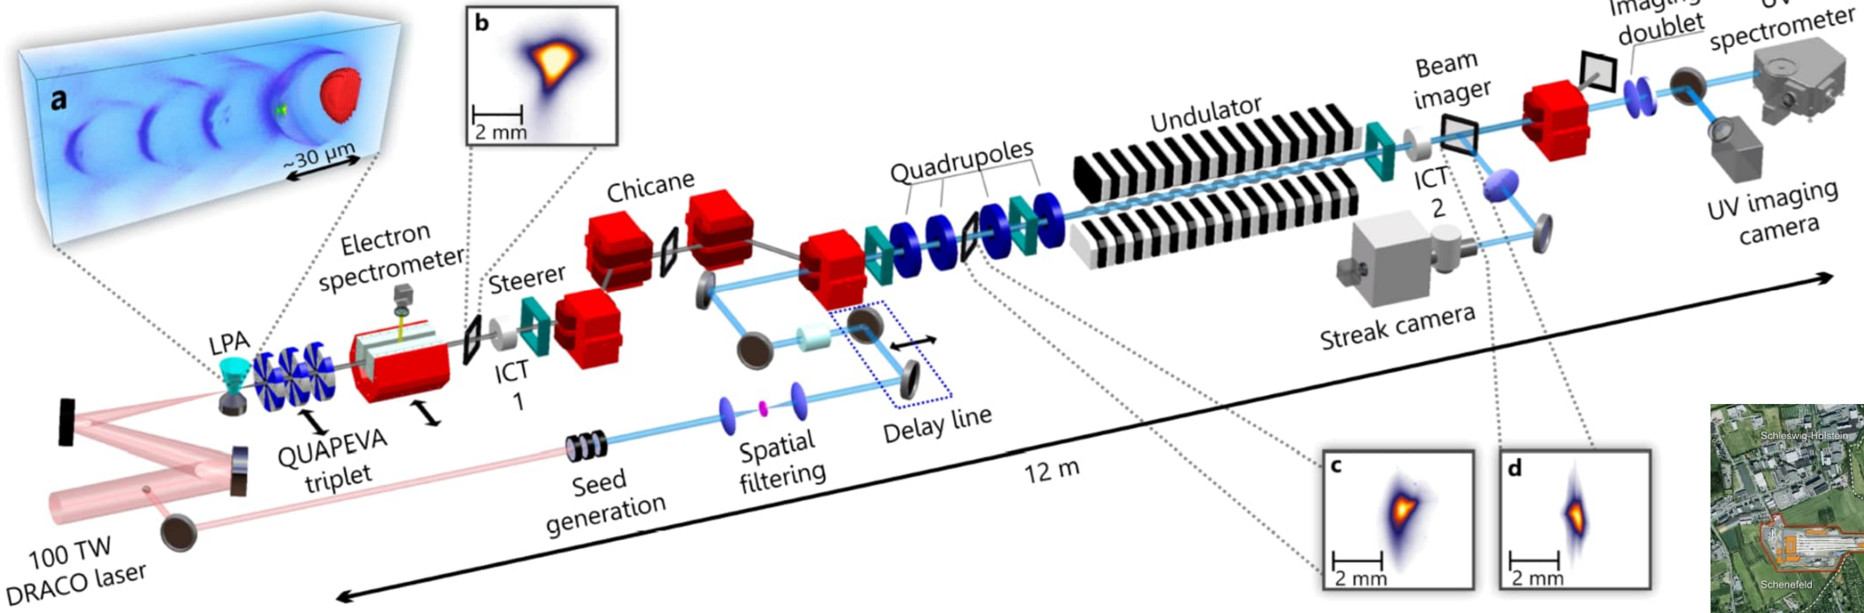

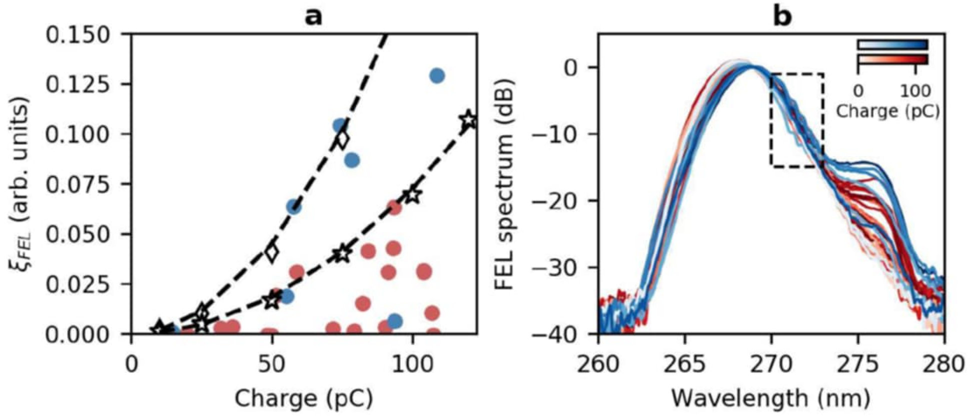

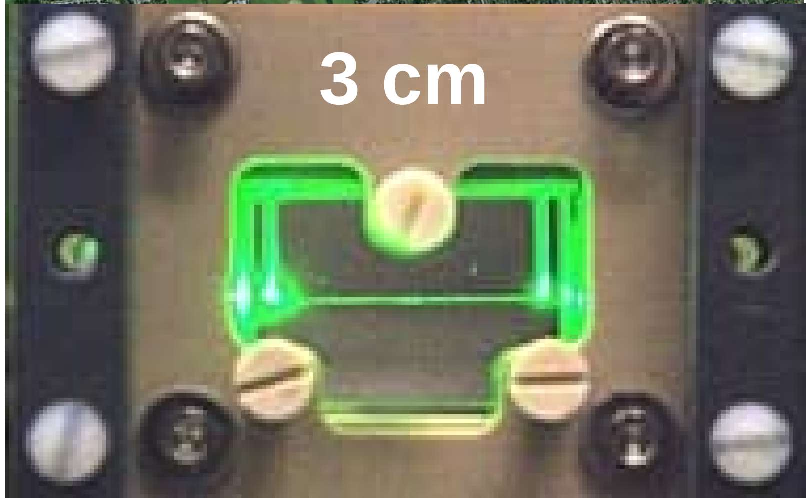

Labat, M., Nat. Photon. 17, 150–156 (2023)

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation of plasma accelerators (via SBI)

We must include all relevant information such as laser fields

The interplay between diagnostics and predictive simulation capabilities is complex

- Including the laser near field in simulations gives better comparison to measurements
- Laser quality control is key to application readiness
- Plasmas act as nonlinear, resonant amplifiers of laser field variations. This is detremental for applications, but also interesting for diagnostics

PIConGPU openPMD

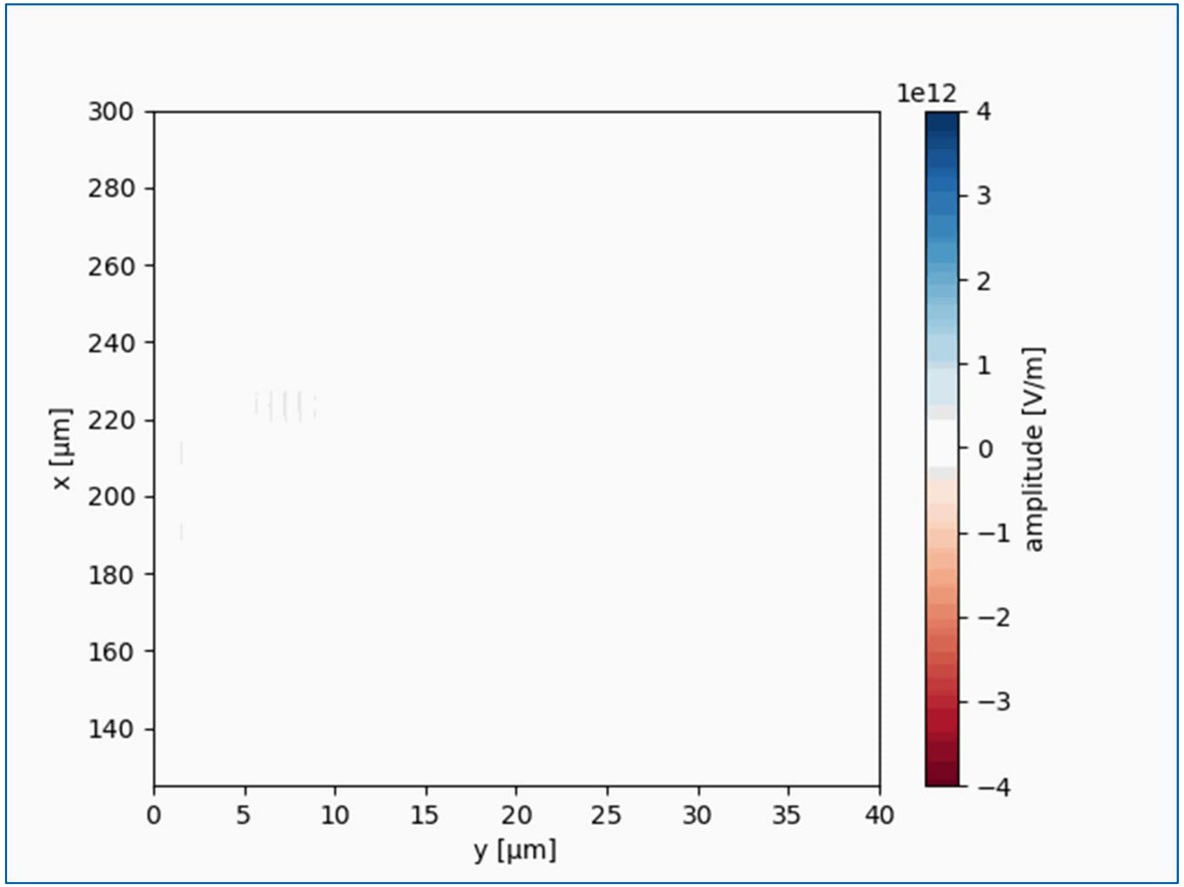

K. Steiniger

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation of plasma accelerators (via SBI)
2. High-quality diagnostics and control are needed to know and set the operational parameters

Finding the optimum parameters for operation with machine learning

ML-guided simulation-based inference is a necessity to ensure optimum XFEL operation

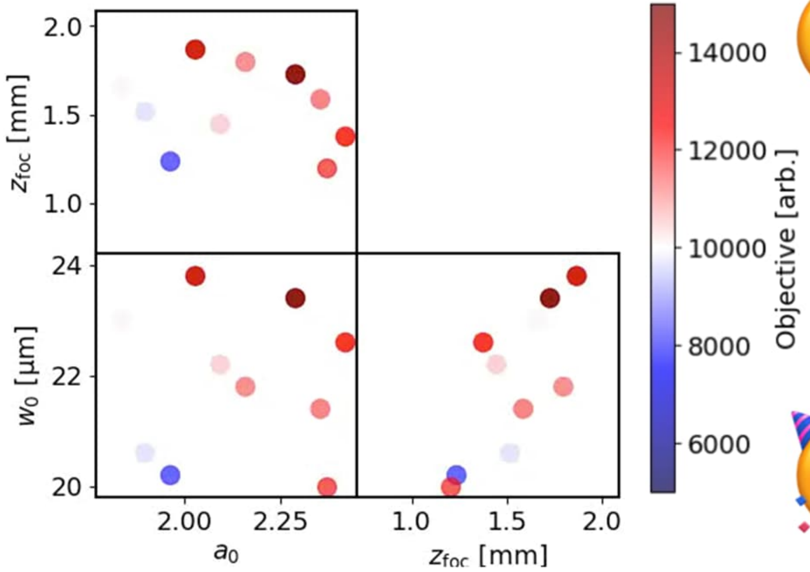

Each dot represents a full-scale HPC simulation

J. Tiebel, R. Pausch

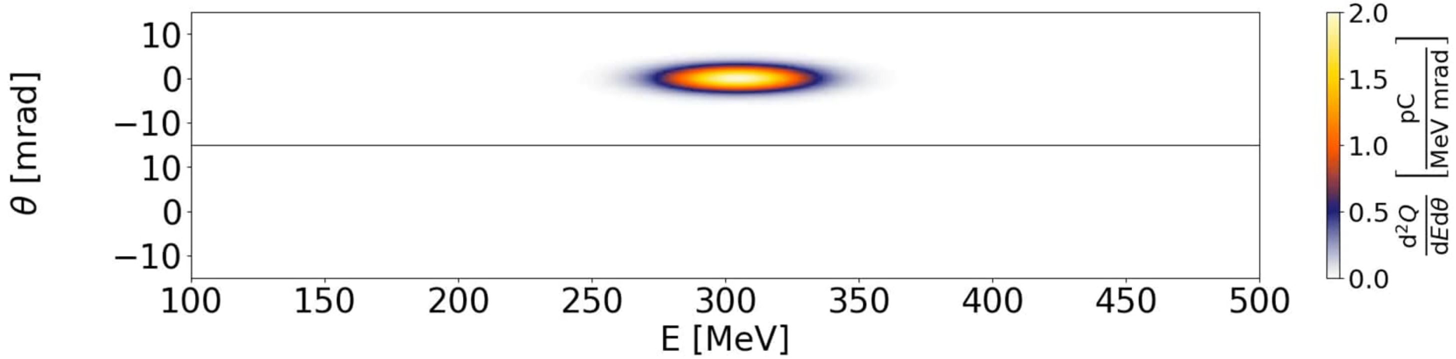

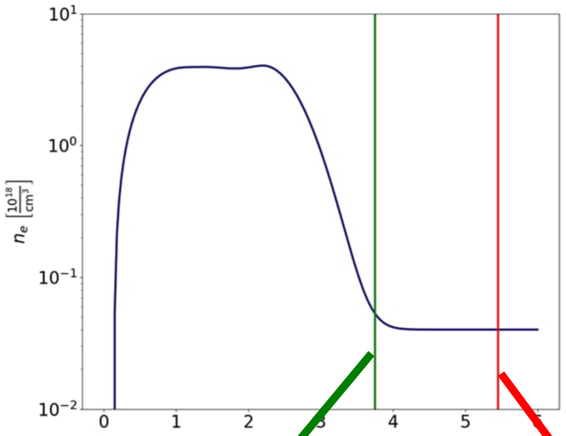

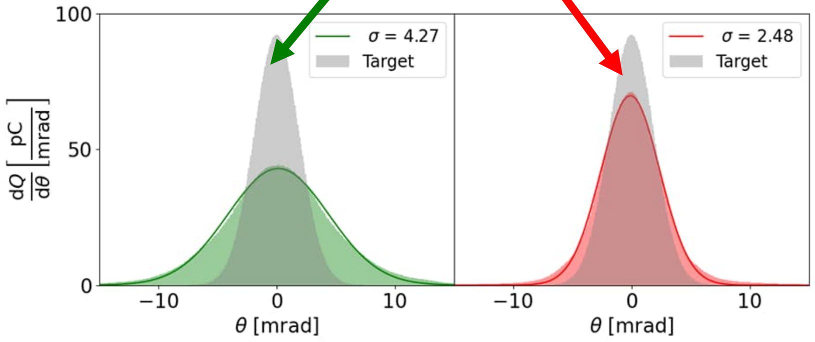

A vision: Staging of plasma accelerators

This has A LOT of challenges, e.g. dephasing, depletion, inter-stage coupling, etc.

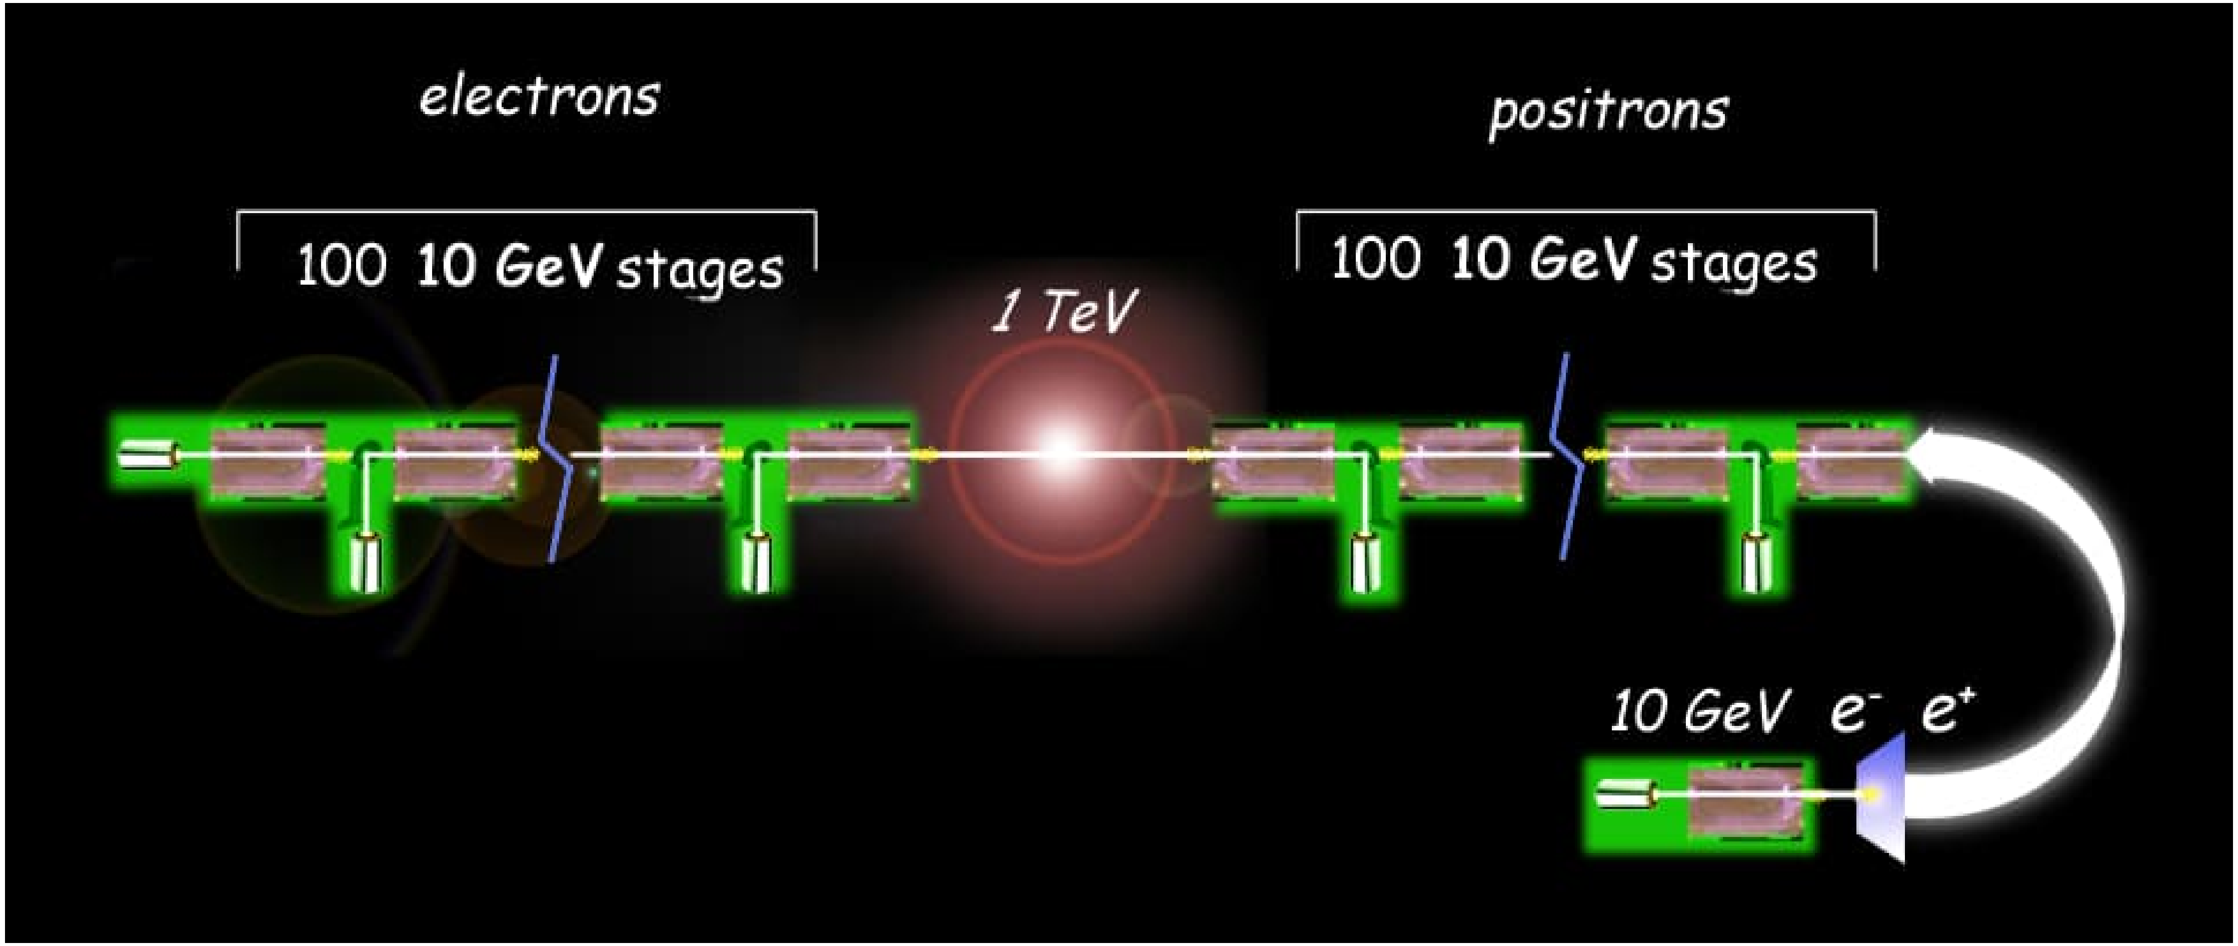

Circumventing dephasing and depletion limits

Travelling-wave electron acceleration (TWEAC)

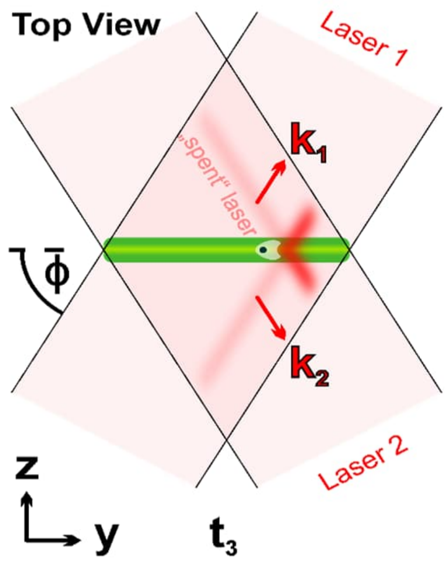

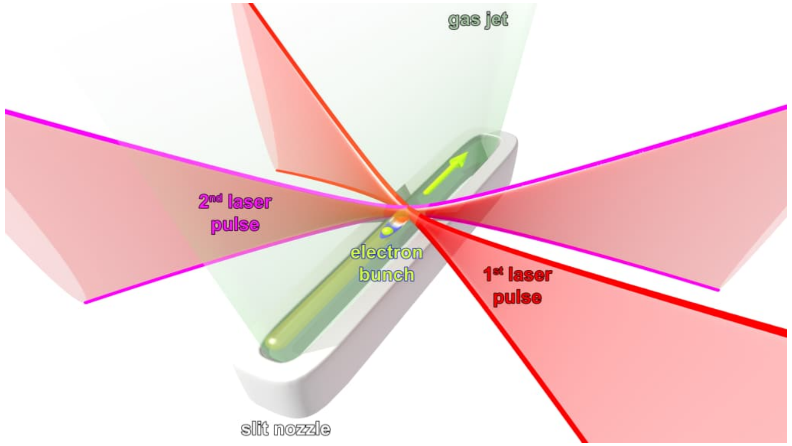

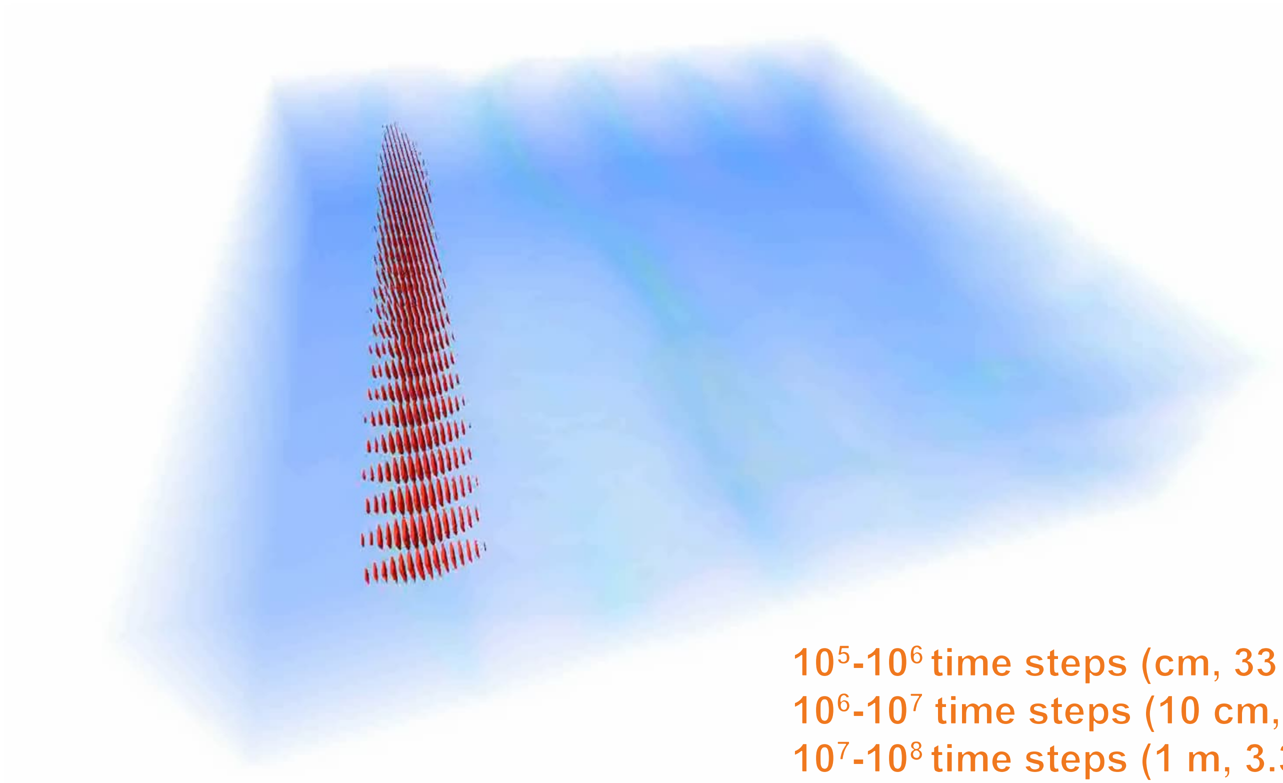

10^5-10^6 time steps (cm, 33 ps)
10^6-10^7 time steps (10 cm, 330 ps)
10^7-10^8 time steps (1 m, 3.3 ns)

- No radial symmetry
- Large transverse extent
- Beyond dephasing/depletion length

A. Debus

Long-scale simulations require higher quality schemes

Obvious lessons with less obvious consequences

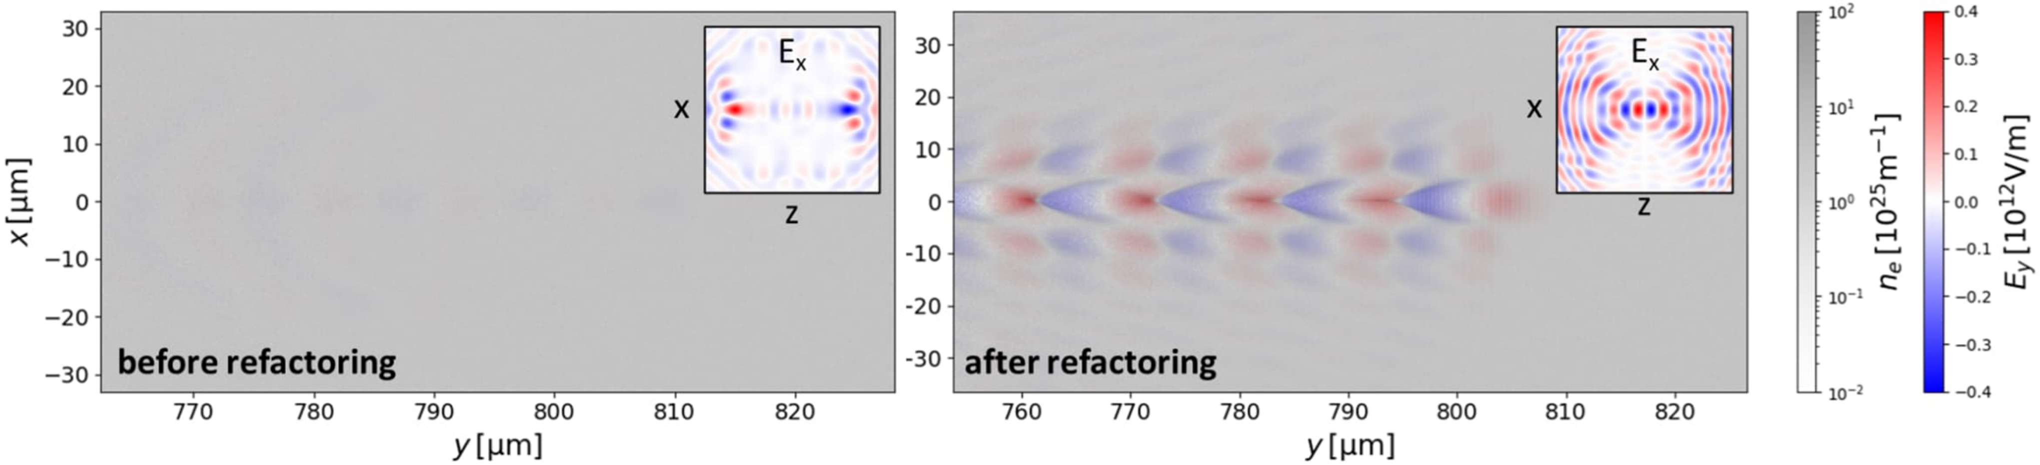

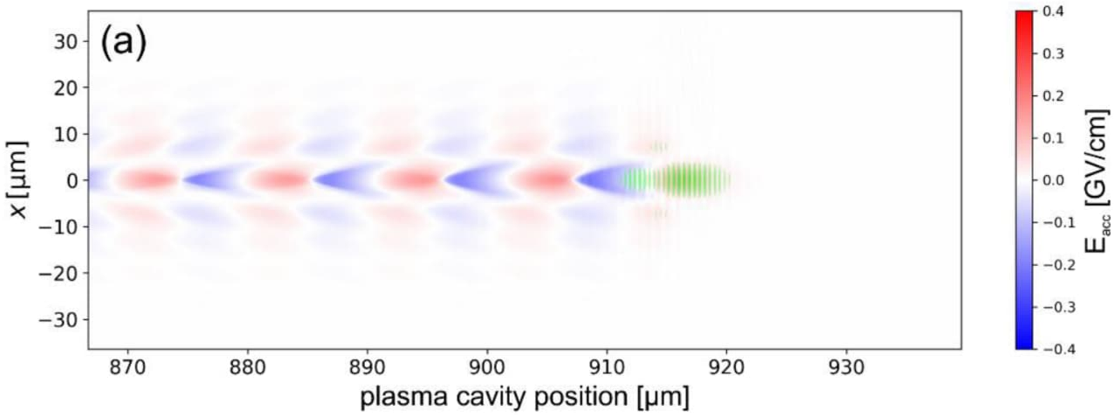

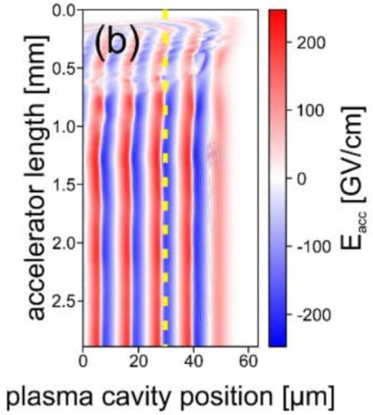

A. Debus

DRESDEN concept HZDR

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation of plasma accelerators (via SBI)
2. High-quality diagnostics and control are needed to know and set the operational parameters
3. At the same time, we must have high-quality, predictive simulations validated in experiments

Matter under extreme Conditions
From Accelerators to Fusion

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

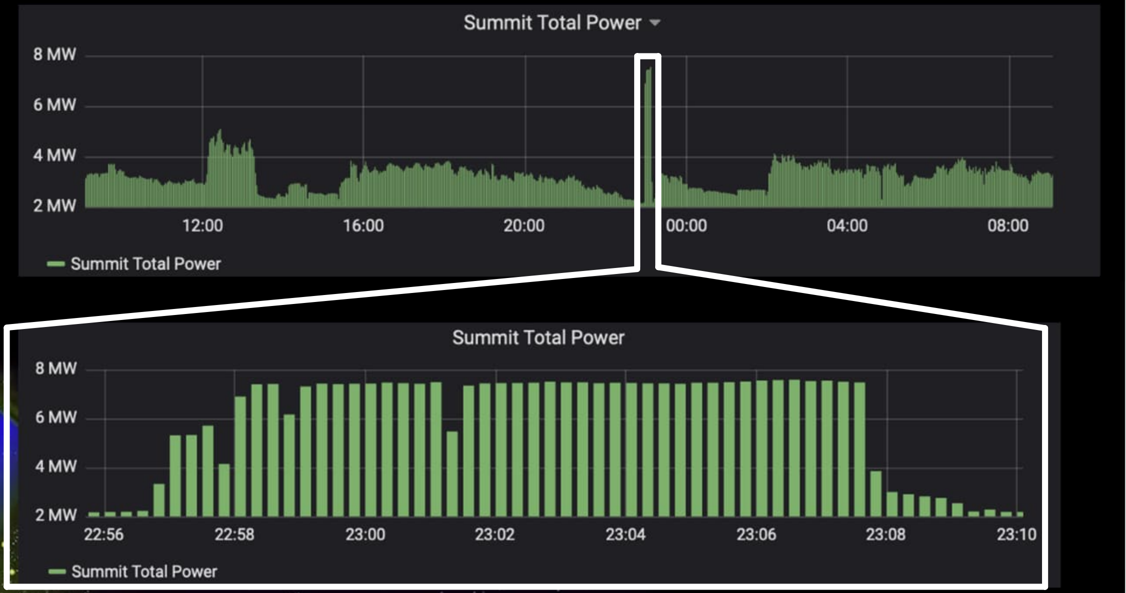

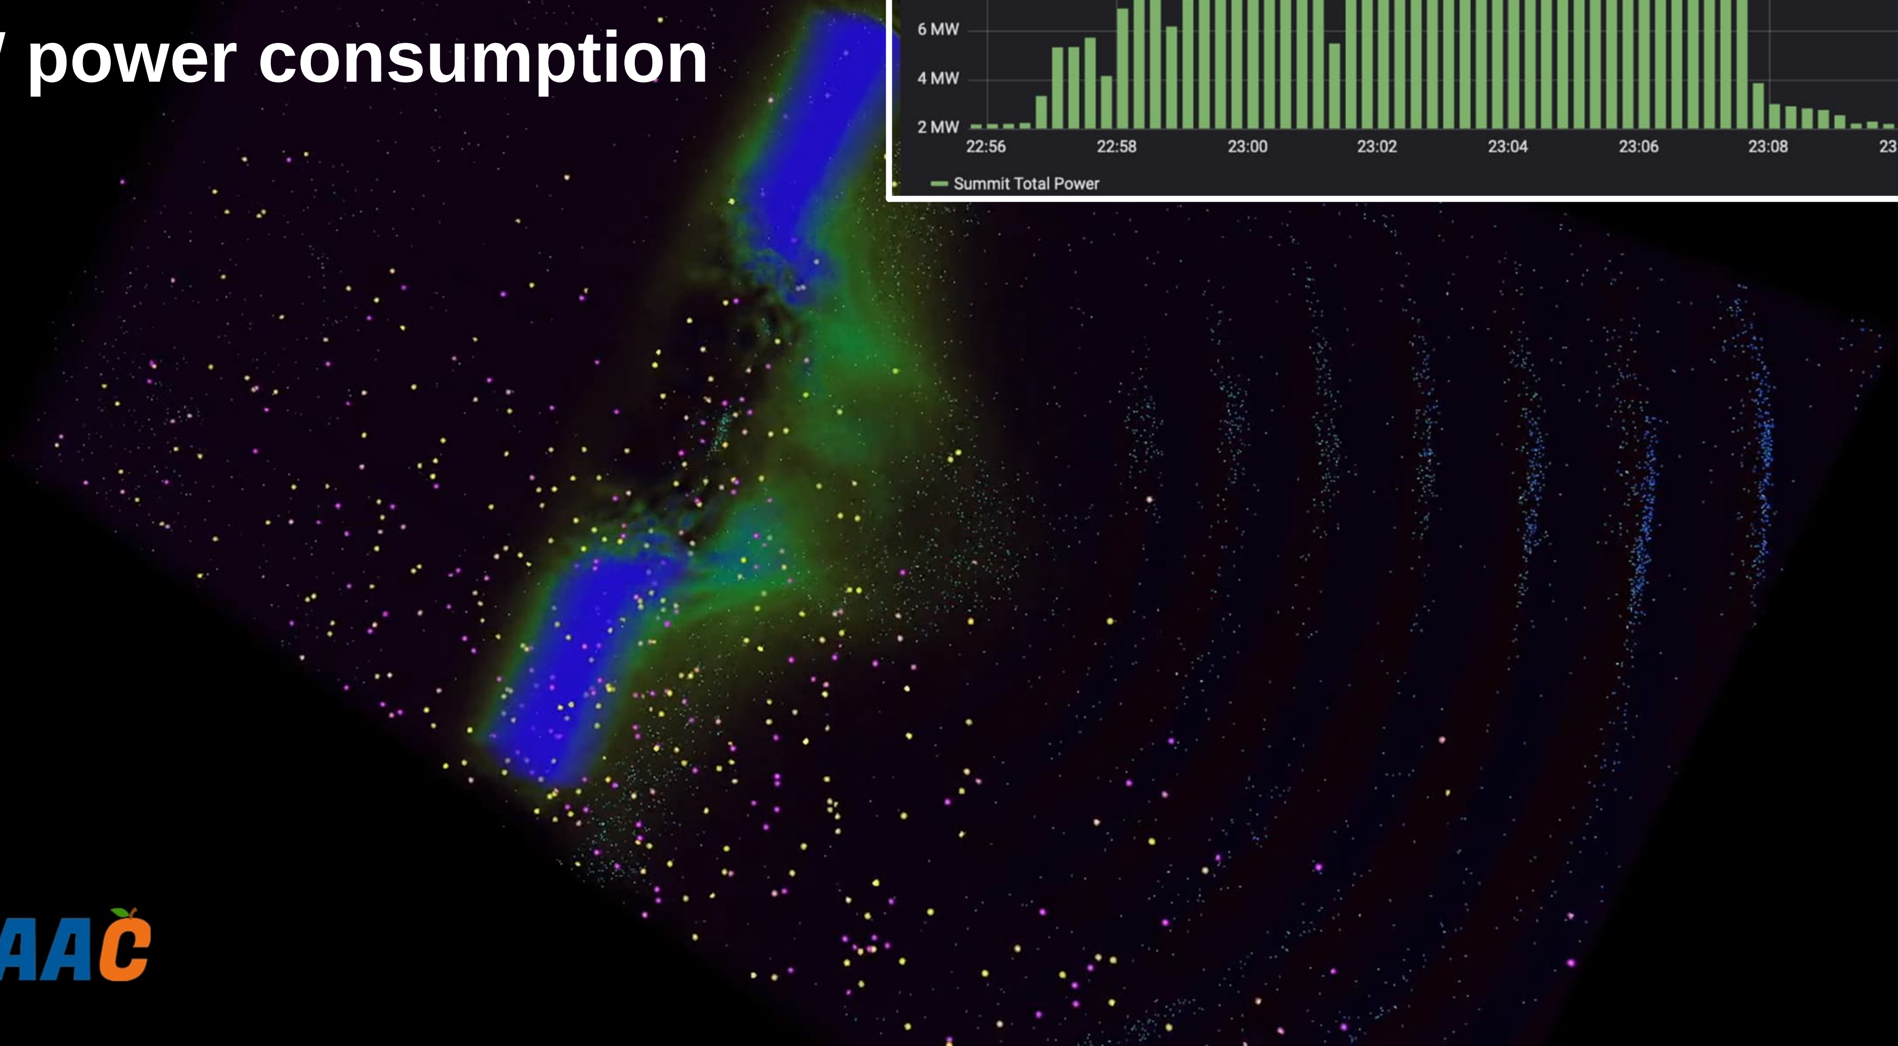

PIConGPU ISAAC

Laser-solid interaction as a multi-scale, multi-physics challenge

Predictive capabilities and code validation by experiment are necessary for excellency

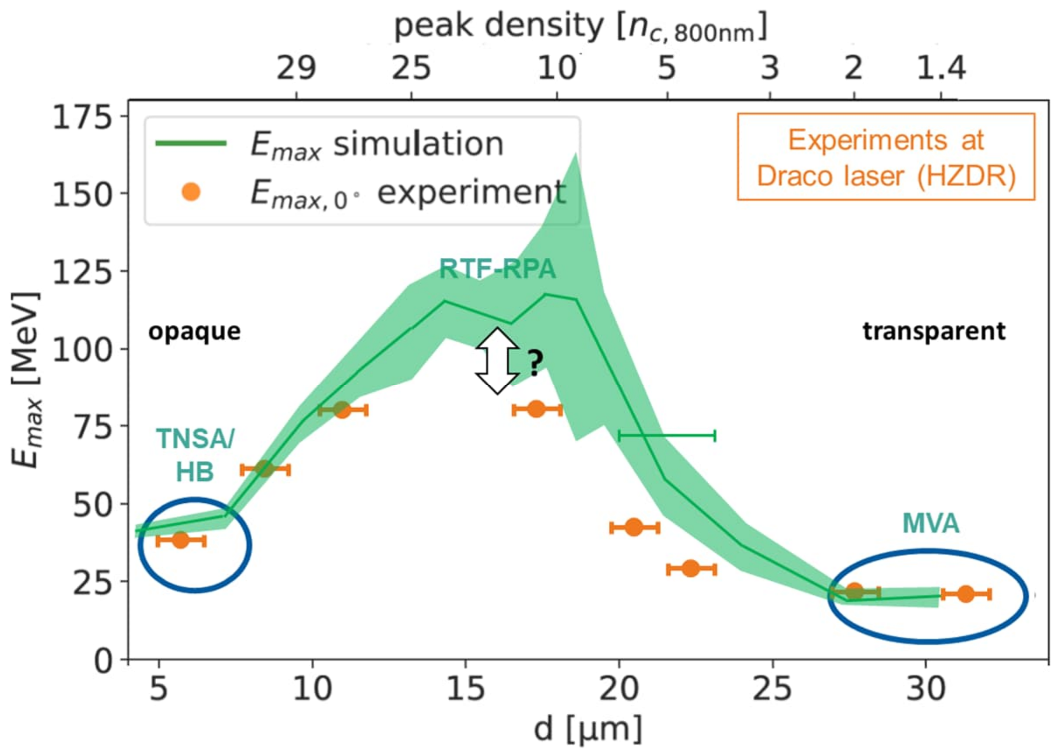

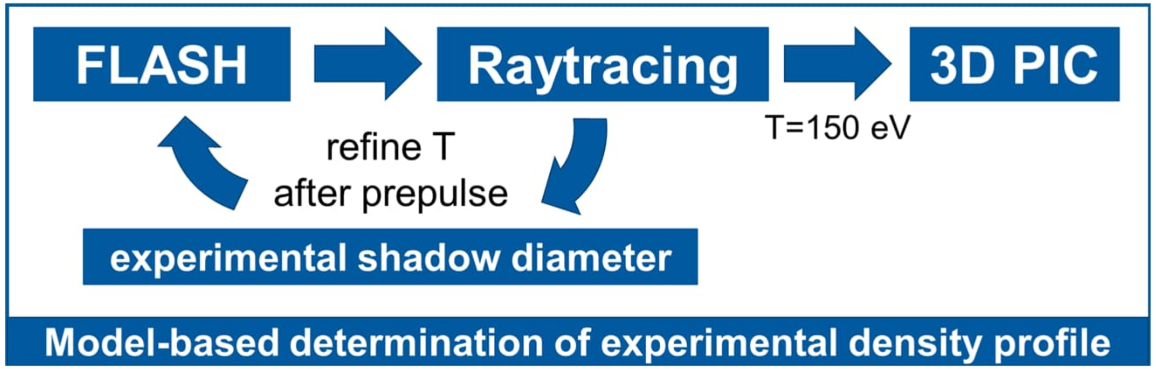

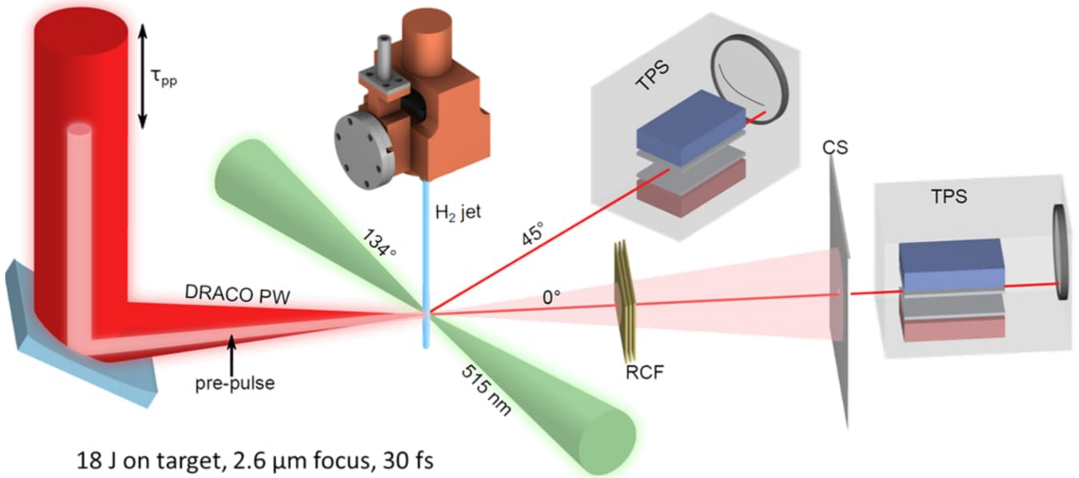

T. Kluge + K. Zeil and many others

DRESDEN concept HZDR

The nonlinear nature of plasma dynamics is a curse and a blessing

We need full spatial resolution + long simulation times + many simulations

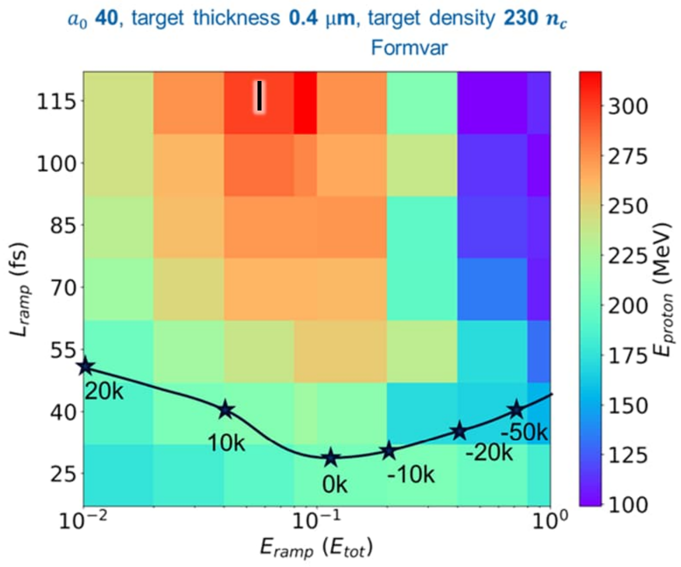

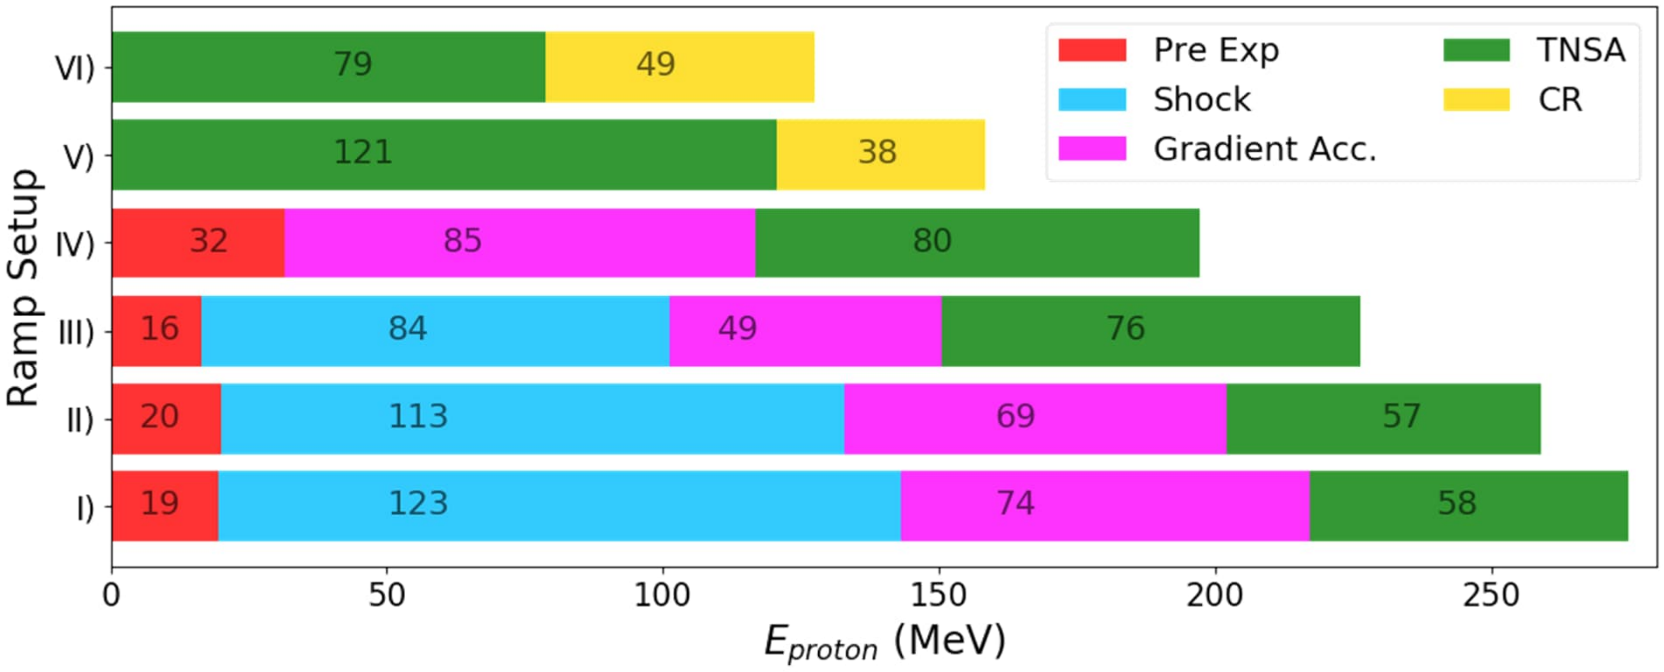

T. Kluge + team

DRESDEN concept HZDR

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation of plasma accelerators (via SBI)
2. High-quality diagnostics and control are needed to know and set the operational parameters
3. At the same time, we must have high-quality, predictive simulations validated in experiments
4. We lose our high fidelity knowledge as simulation data is becoming too large

Comparing simulation to experiment at atomic resolution

HiBEF HELMHOLTZ INTERNATIONAL BEAMLINE FOR EXTREME FIELDS

Plasma probing reaches fs/nm resolution

New experimental capabilities enable unprecedented comparison to models

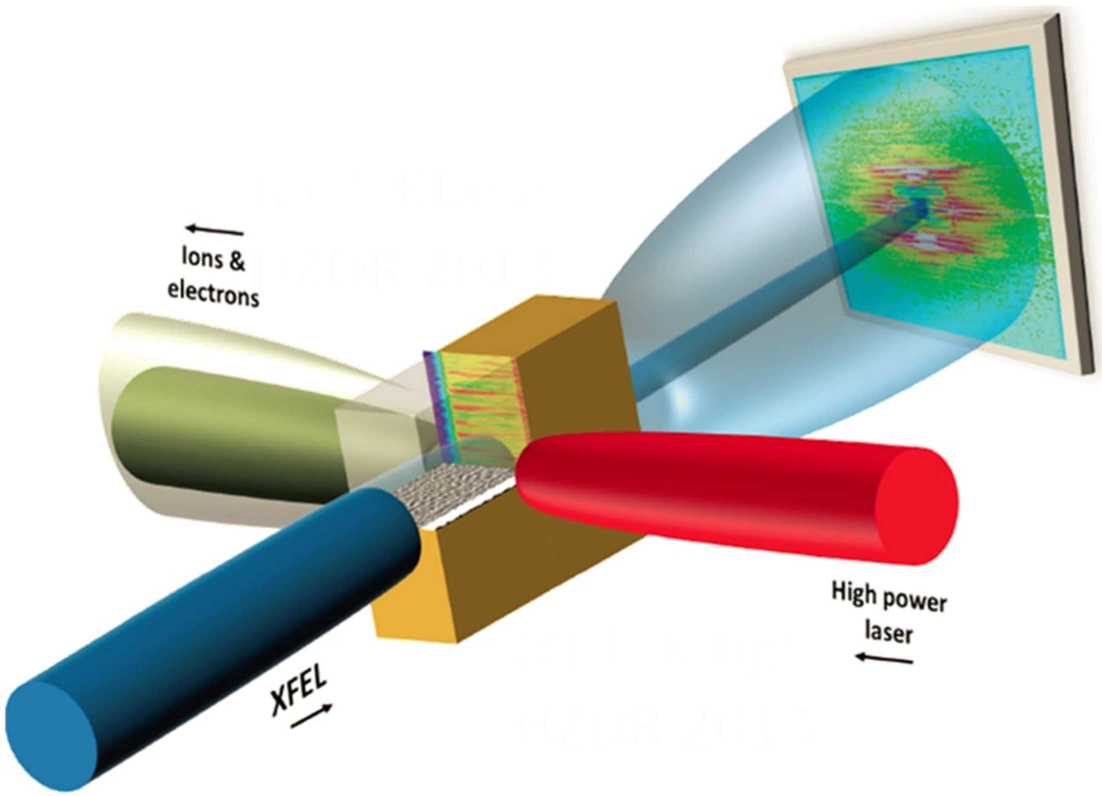

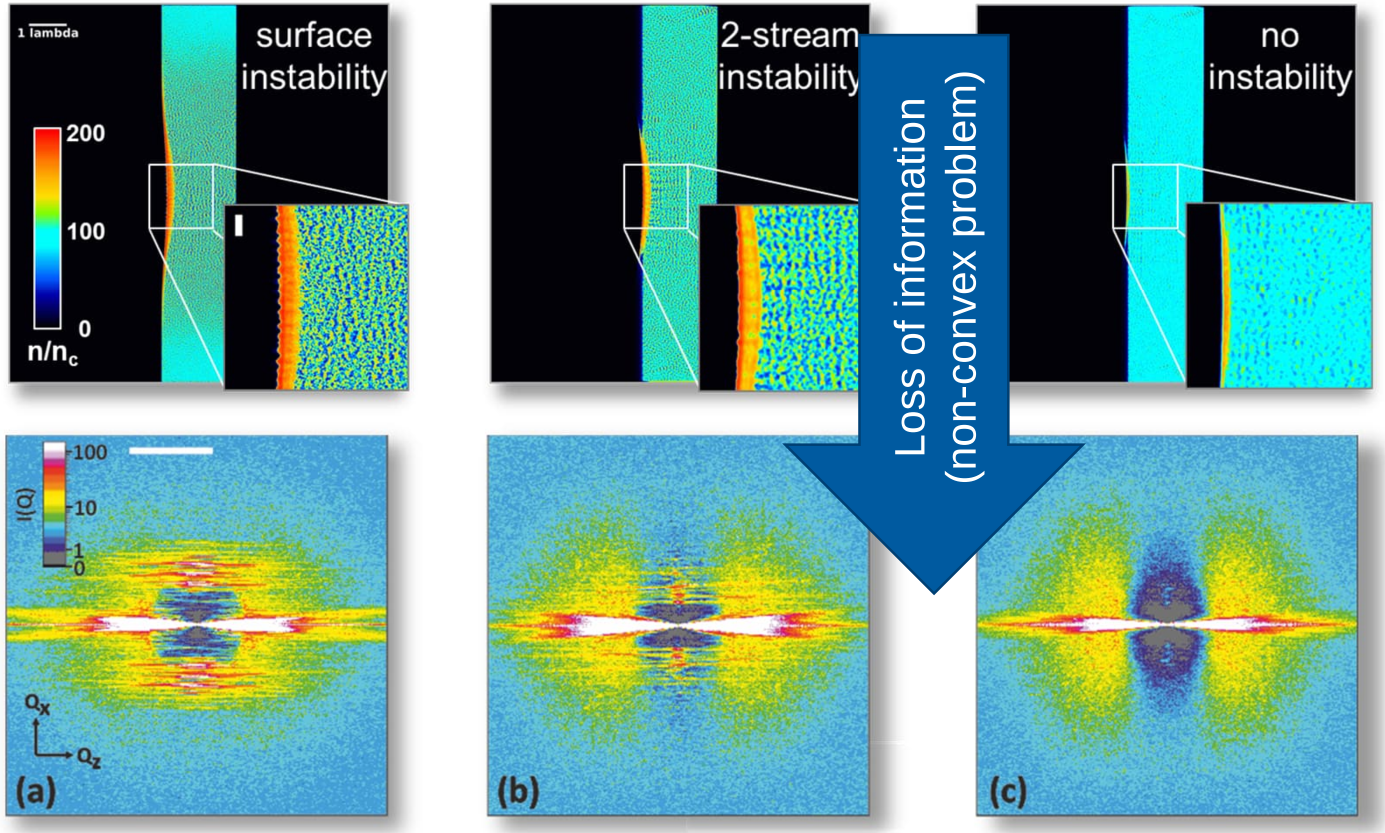

T. Kluge, L. Huang

Plasma probing reaches fs/nm resolution

Coherent photon by photon radiation transport, per-atom atomic state calculation

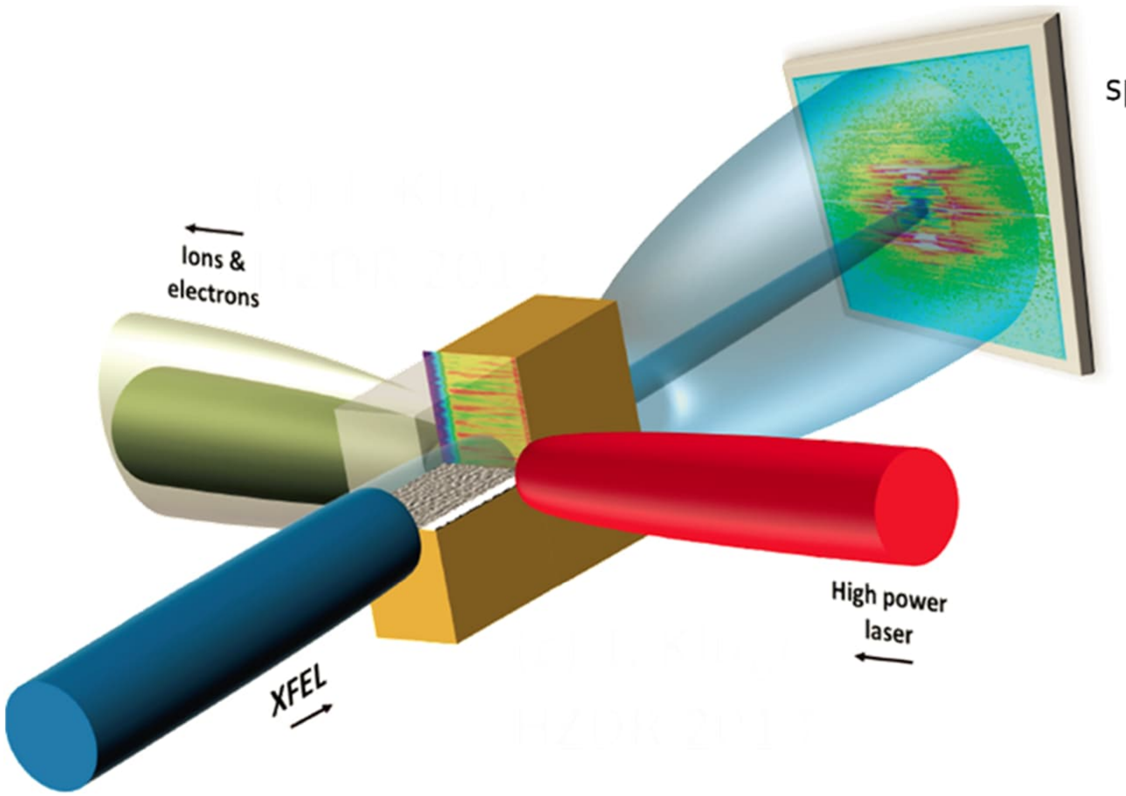

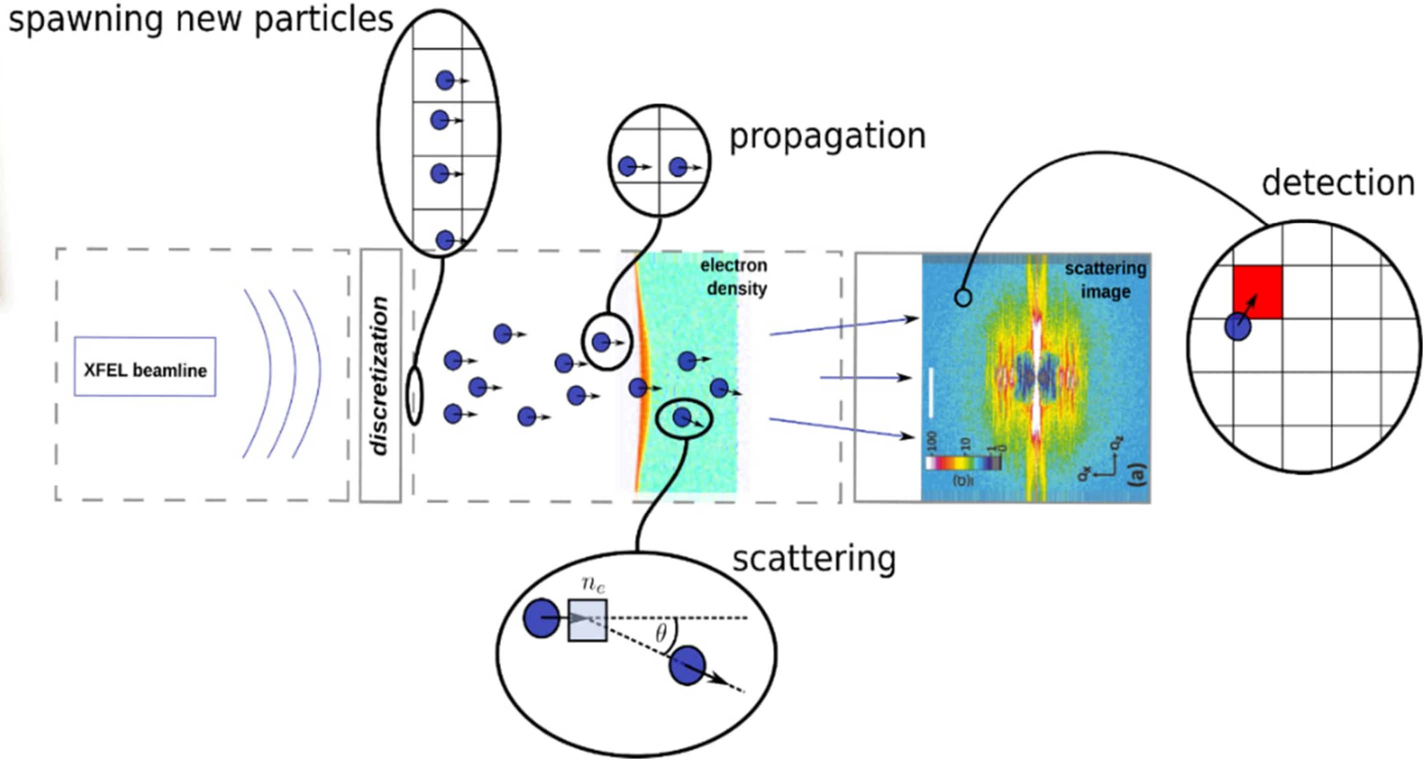

DRESDEN concept HZDR

Plasma probing reaches fs/nm resolution

Coherent photon by photon radiation transport, per-atom atomic state calculation

Electron density
double slit

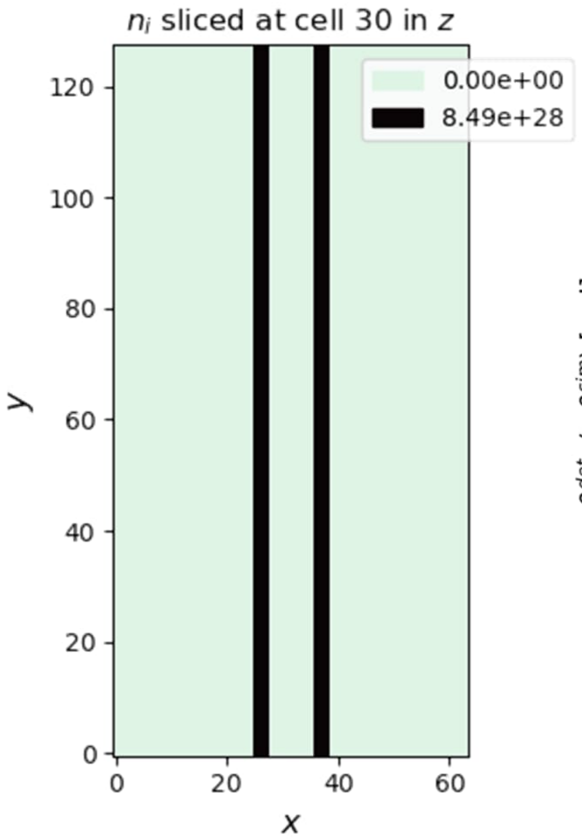

In-situ scattering image on detector

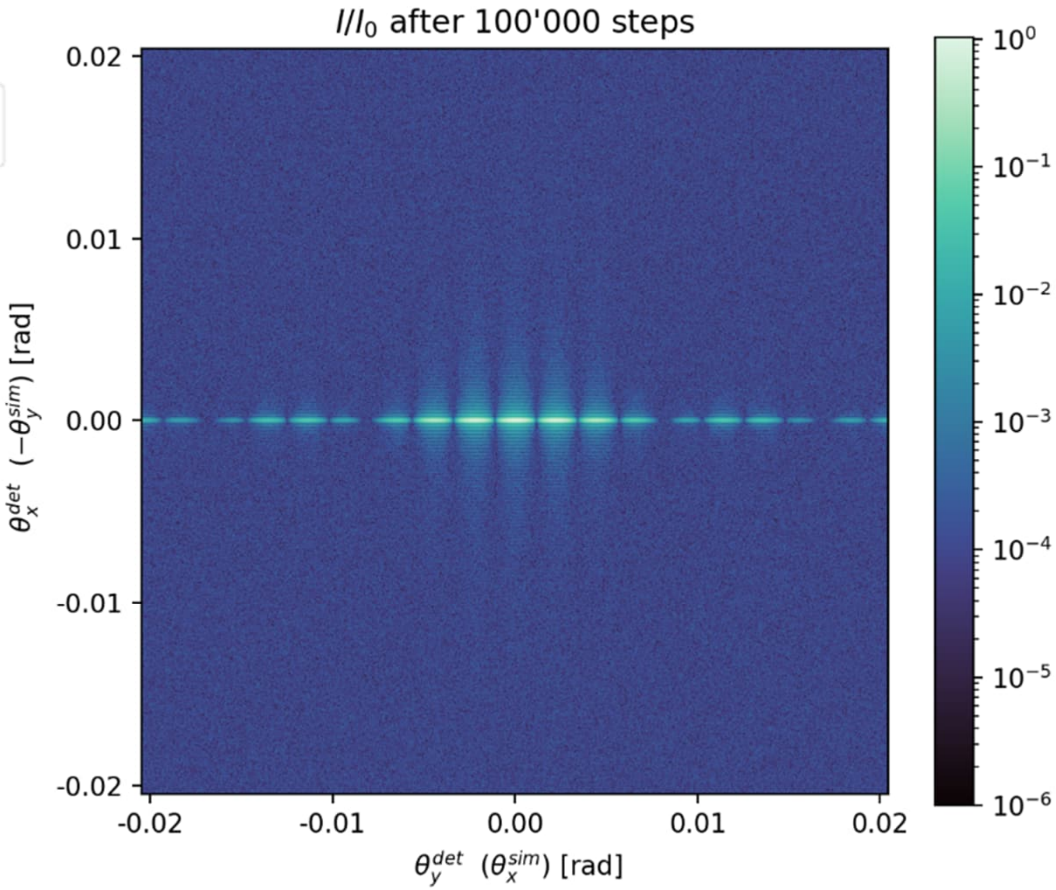

Plasma probing reaches fs/nm resolution

Well-designed targets let us investigate dynamics of laser surface interaction

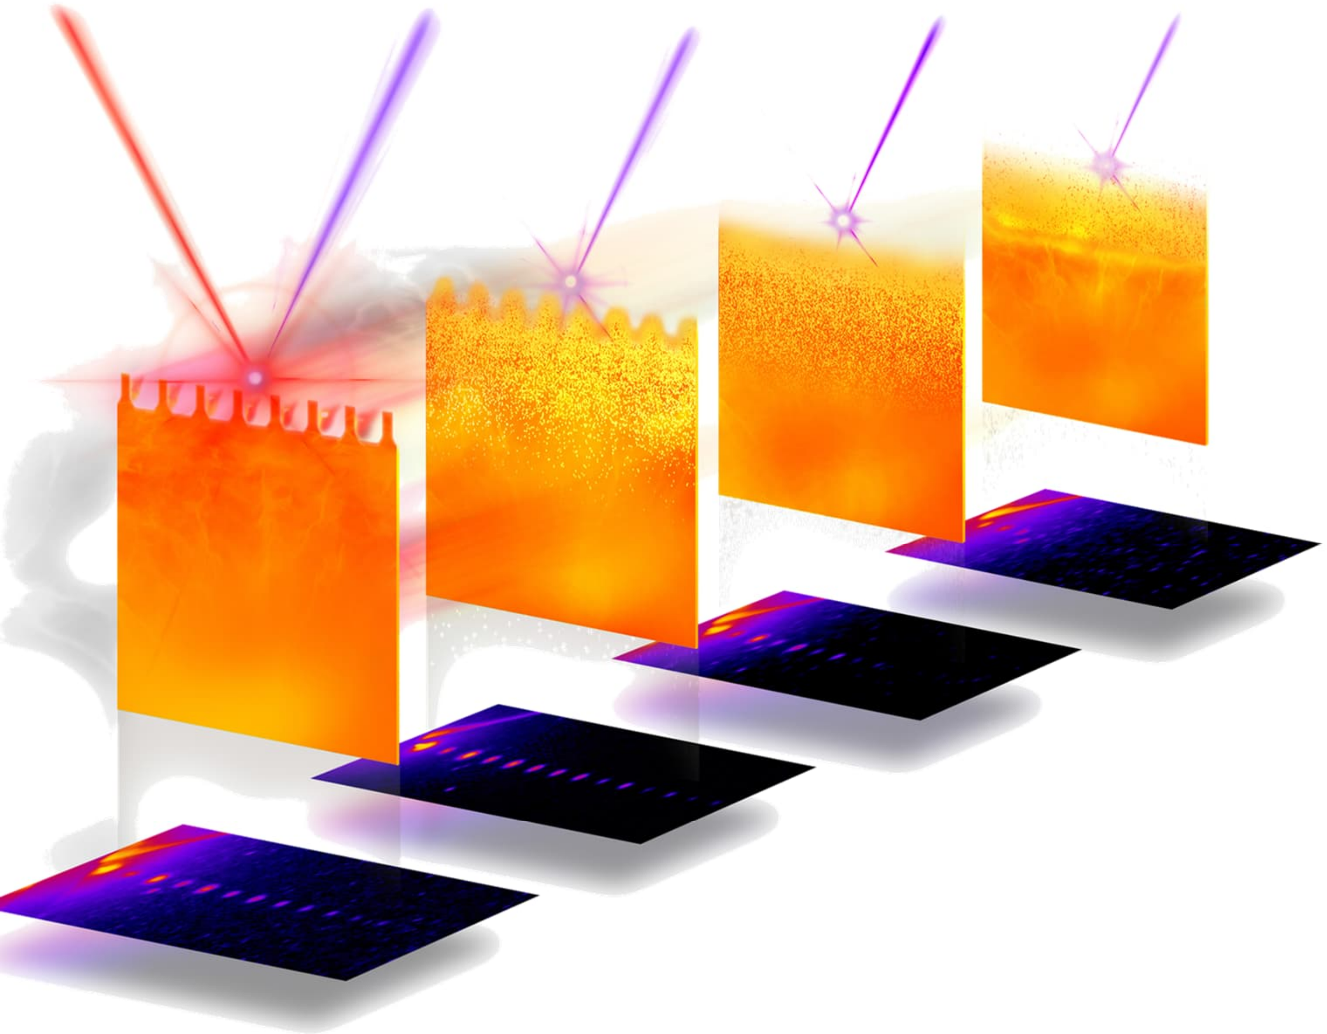

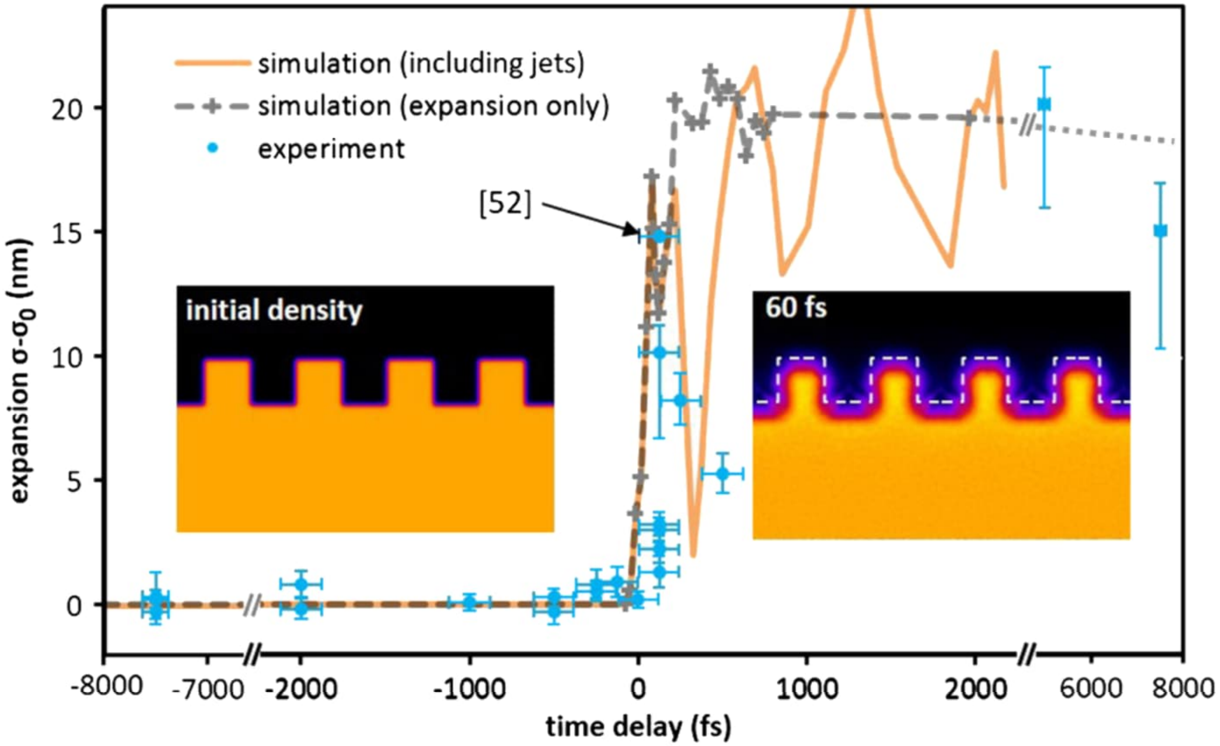

T. Kluge

Plasma probing reaches fs/nm resolution

Including per atom reduced atomic states for simulating non-equilibrium radiation transport

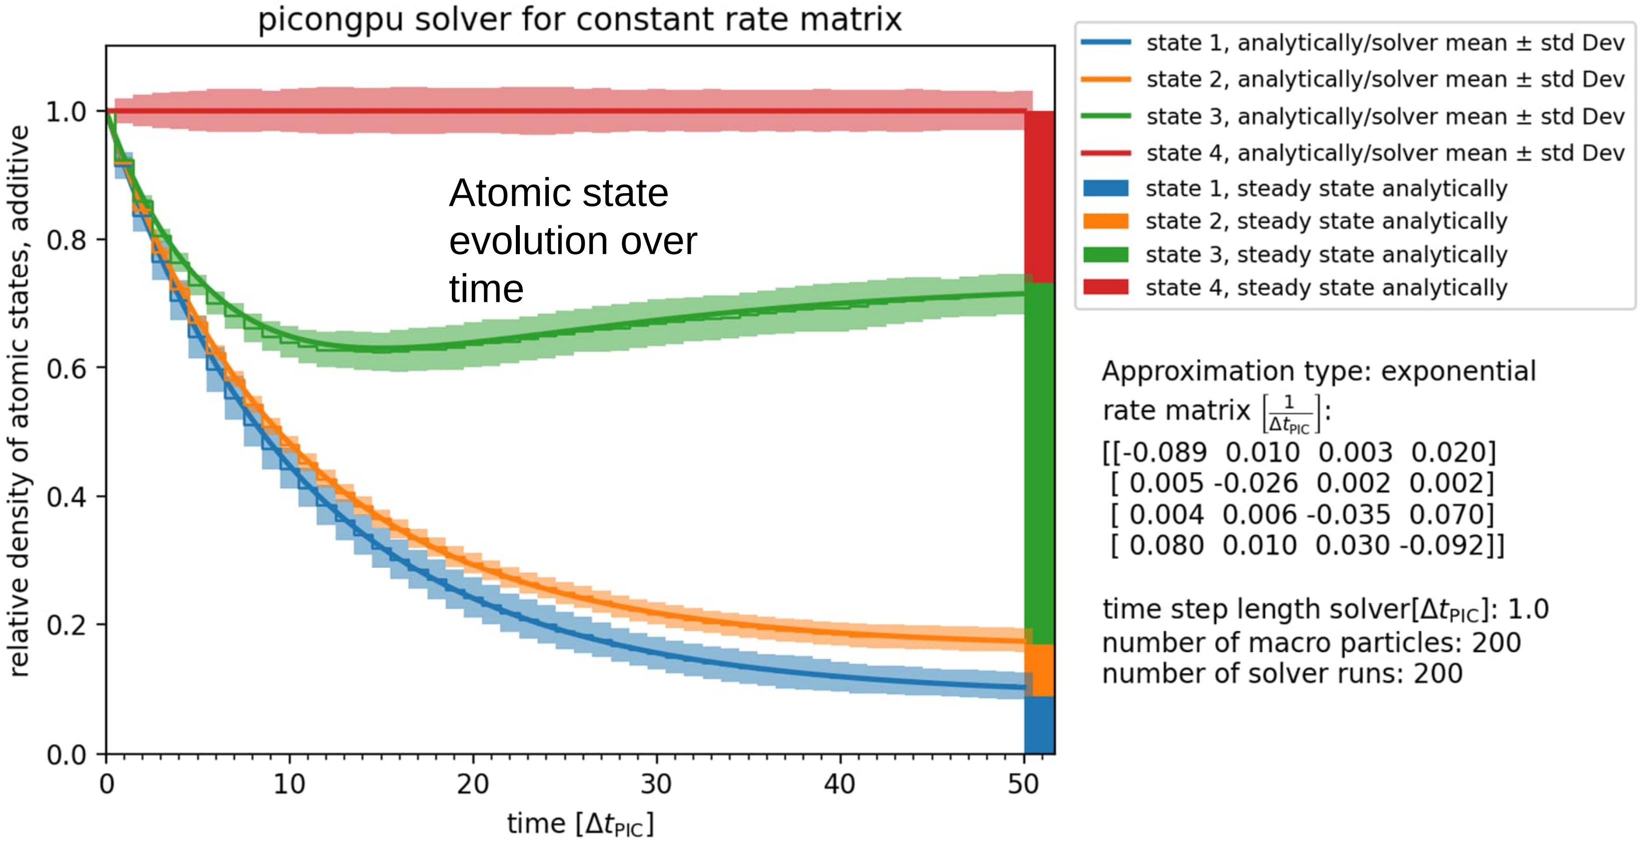

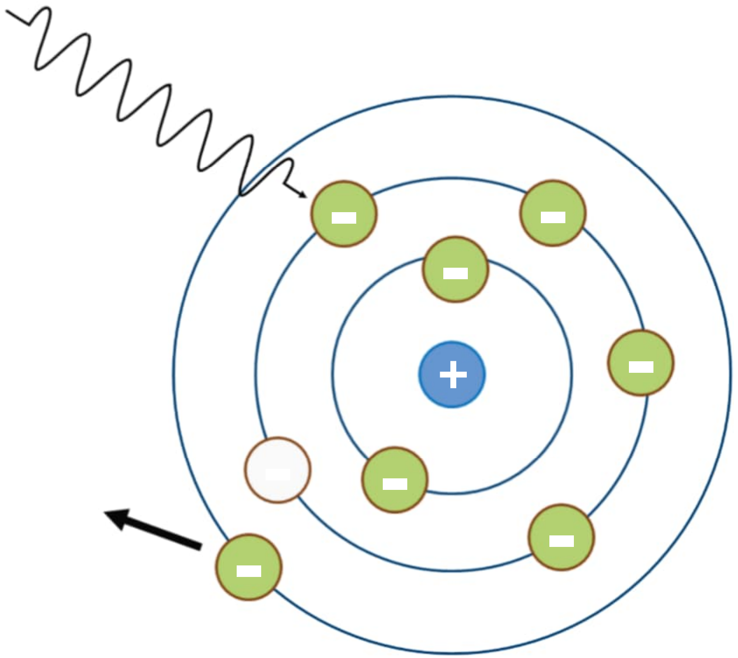

B. Marré

Plasma probing reaches fs/nm resolution

We now have spatially and temporally resolved ionization dynamics in experiment

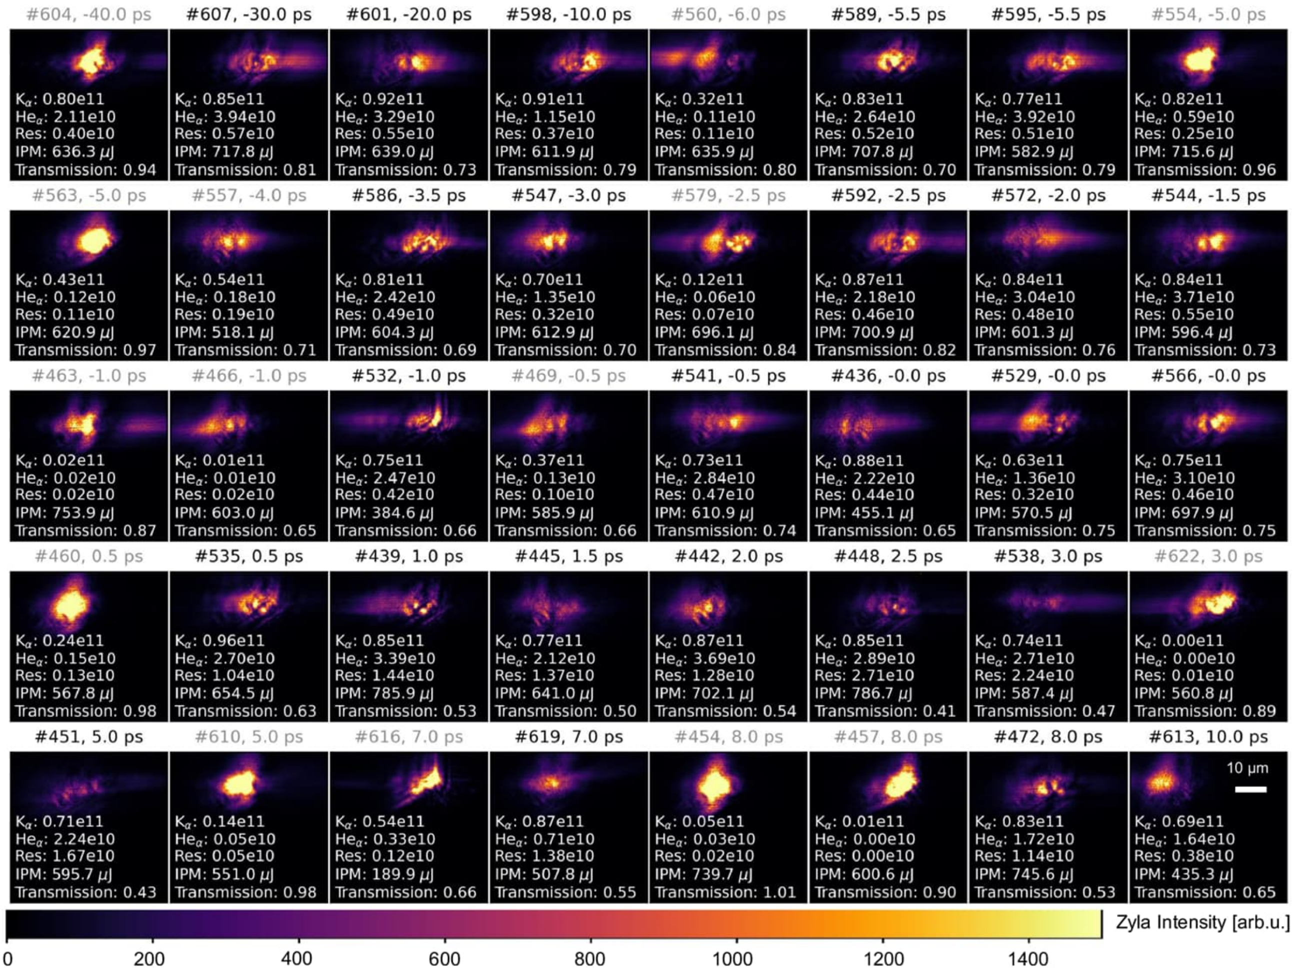

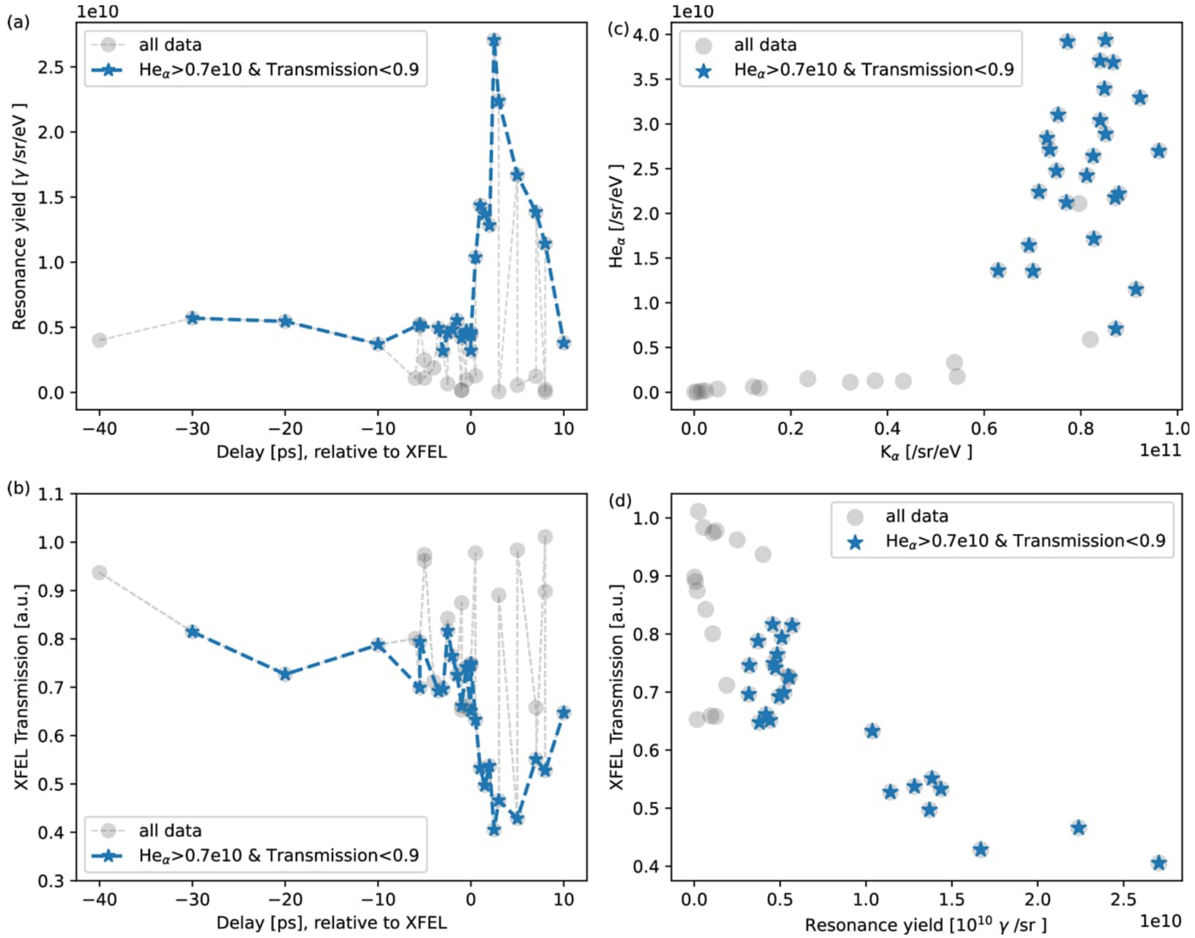

Huang, L., et al. Nat Commun 17, 3219 (2026)

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation for applications
2. High-quality diagnostics and control are needed to know and set the operational parameters
3. At the same time, we must have high-quality, predictive simulations validated in experiments
4. We lose our high fidelity knowledge as simulation data is becoming too large
5. Our experiments now reach similar fidelity as our simulations, with similar data problems ahead

Coupling experiment and simulation

ML-acceleration of scattering image computation (forward)

Physics-informed ML models allow for fast synthetic diagonstics

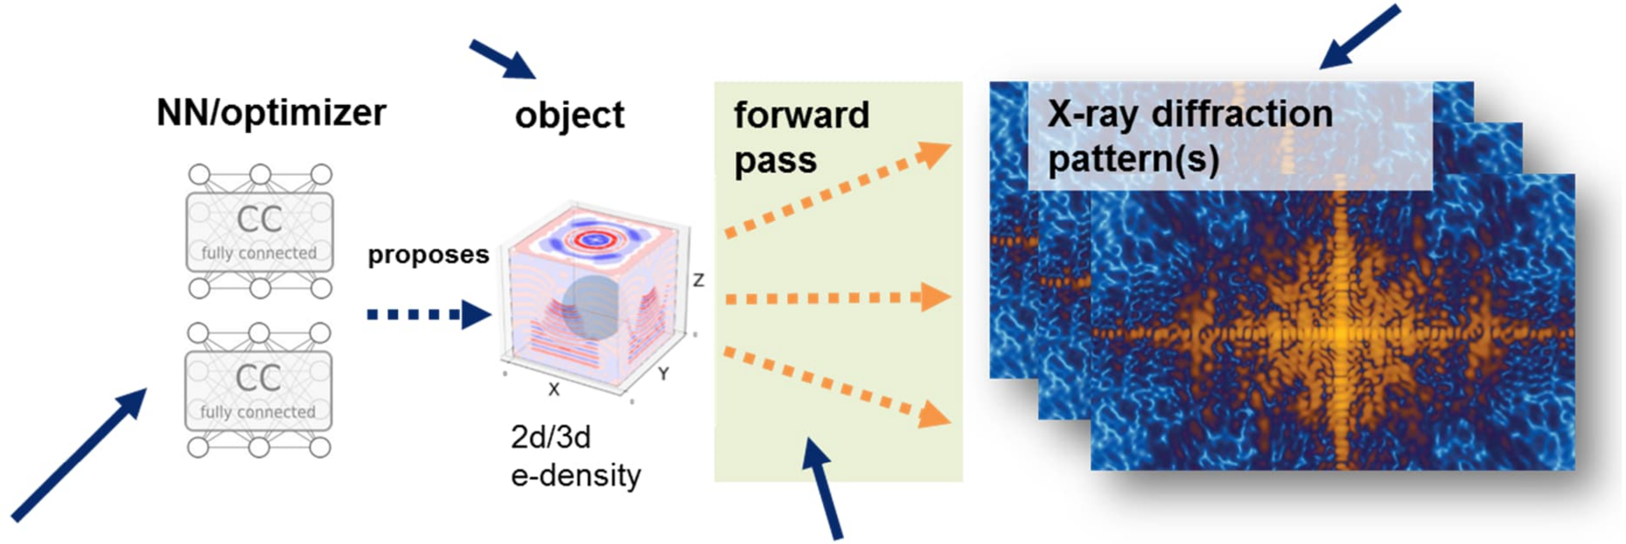

Solver of IIP
generative model of
solution space

(Deep) Physics Prior
analytical expression or surrogate model*

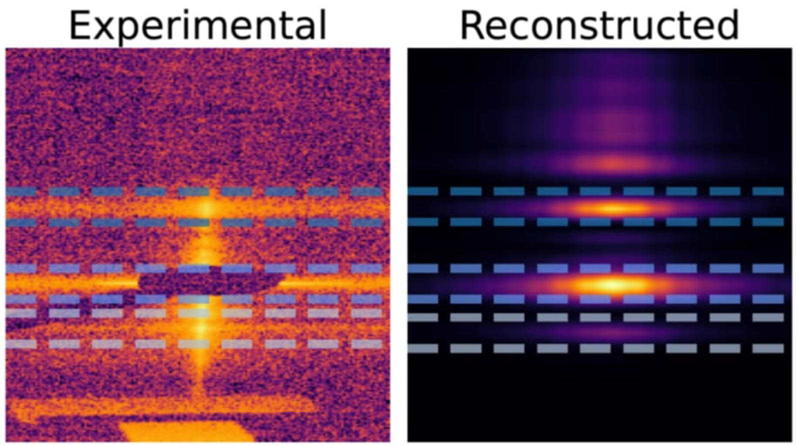

GISAXS-Data from
Laser Plasma Experiment

J.Kelling, N. Hoffmann

ML-acceleration of scattering image reconstruction (inverse)

Physics-informed ML models allow for better data reconstruction

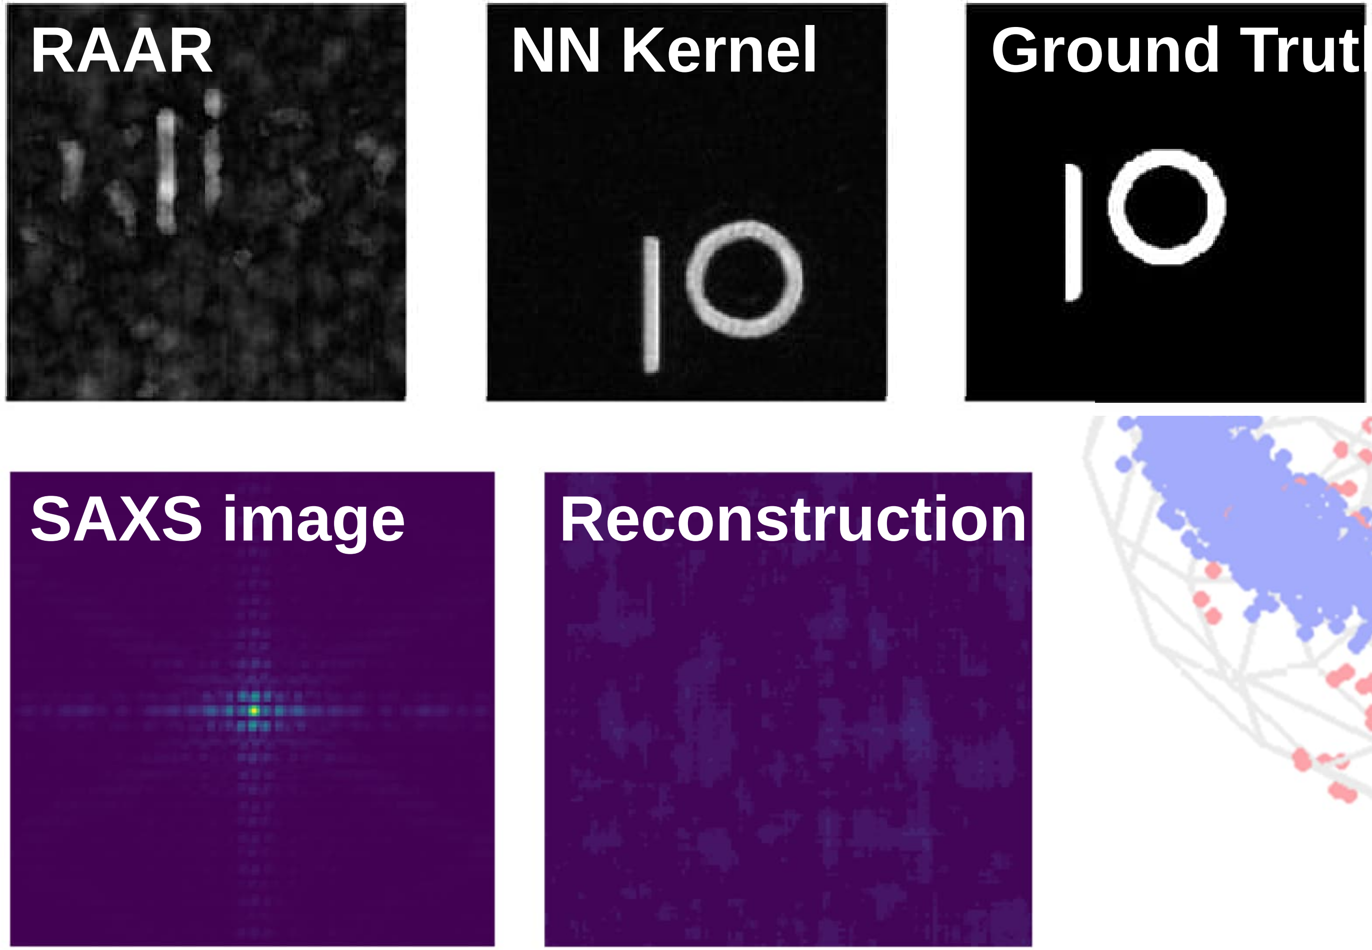

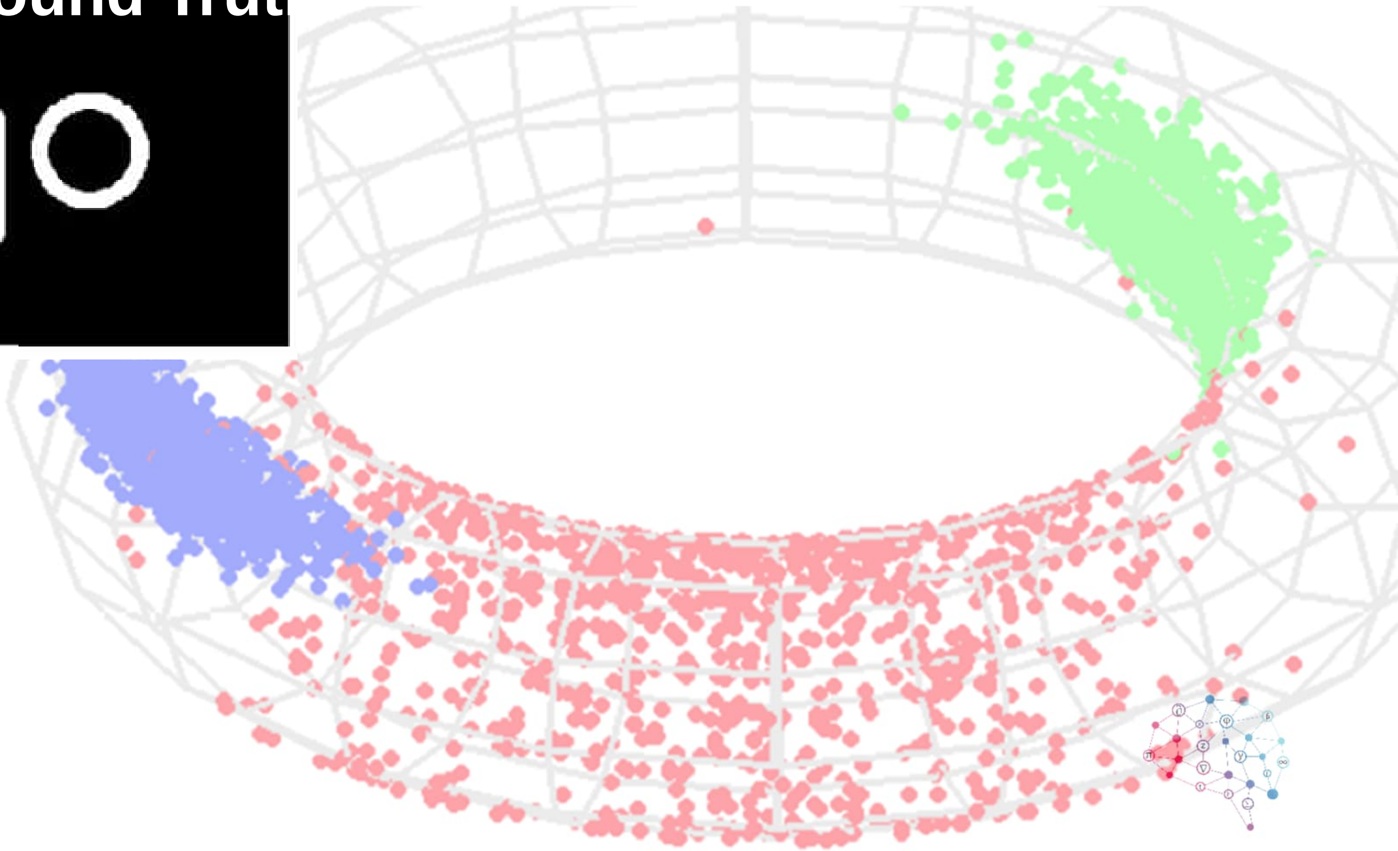

NEURAL SOLVERS

ML surrogate models must use a single system description

Multi-modal AI is a necessity, as every diagnostic tells us something else about the system

J.Kelling, N. Hoffmann

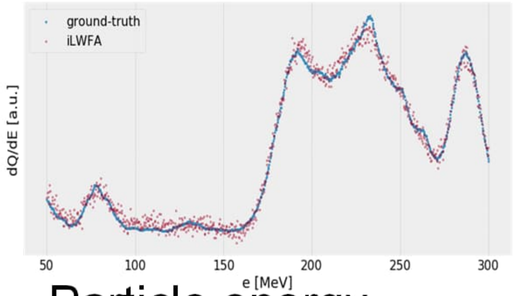

Particle energy spectra

Shadowgraphy

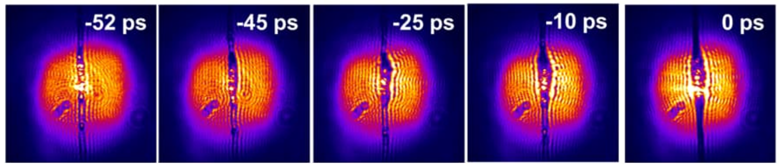

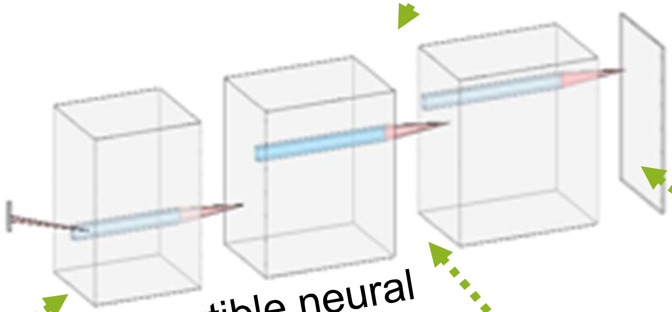

Invertible neural network for **joint surrogate modelling**

Radiation spectra

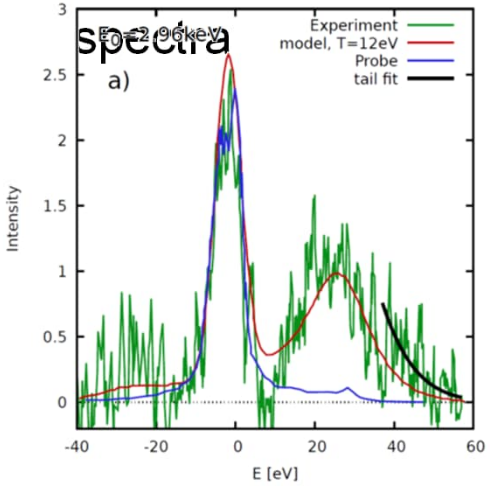

Laser properties

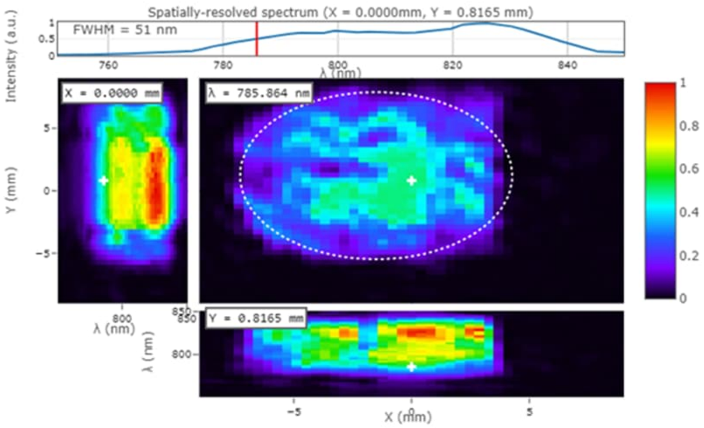

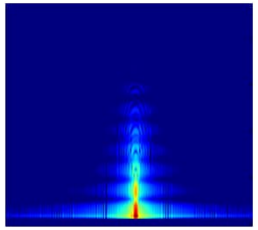

(GI-)SAXS, WAXS

We can couple the forward and the inverse by AI

Invertible neural networks help to couple high fidelity plasma dynamics with observation

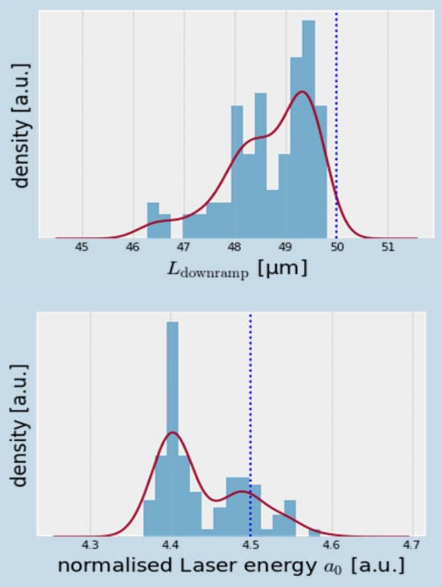

Invertible
Neural Network

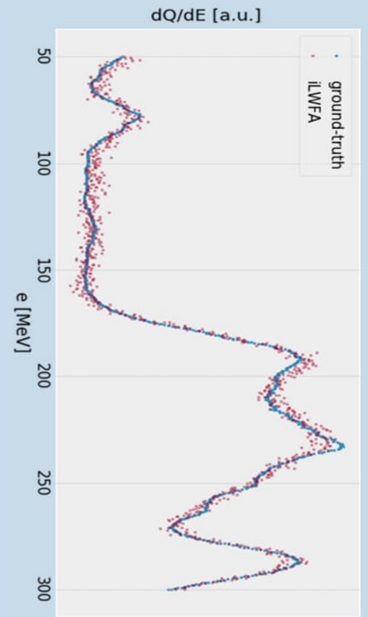

Benefits

- Recover ambiguous mapping
- Uncertainty quantification

Inputs:
Laser Energy
& Plasma Density Profile

J.Kelling, N. Hoffmann

Output:
Particle Energy Spectrum

NEURAL SOLVERS

DRESDEN concept

HZDR

We produce PB/s data with our simulations at high fidelity

We cannot store this and thus need to stream, reduce and keep the relevant information

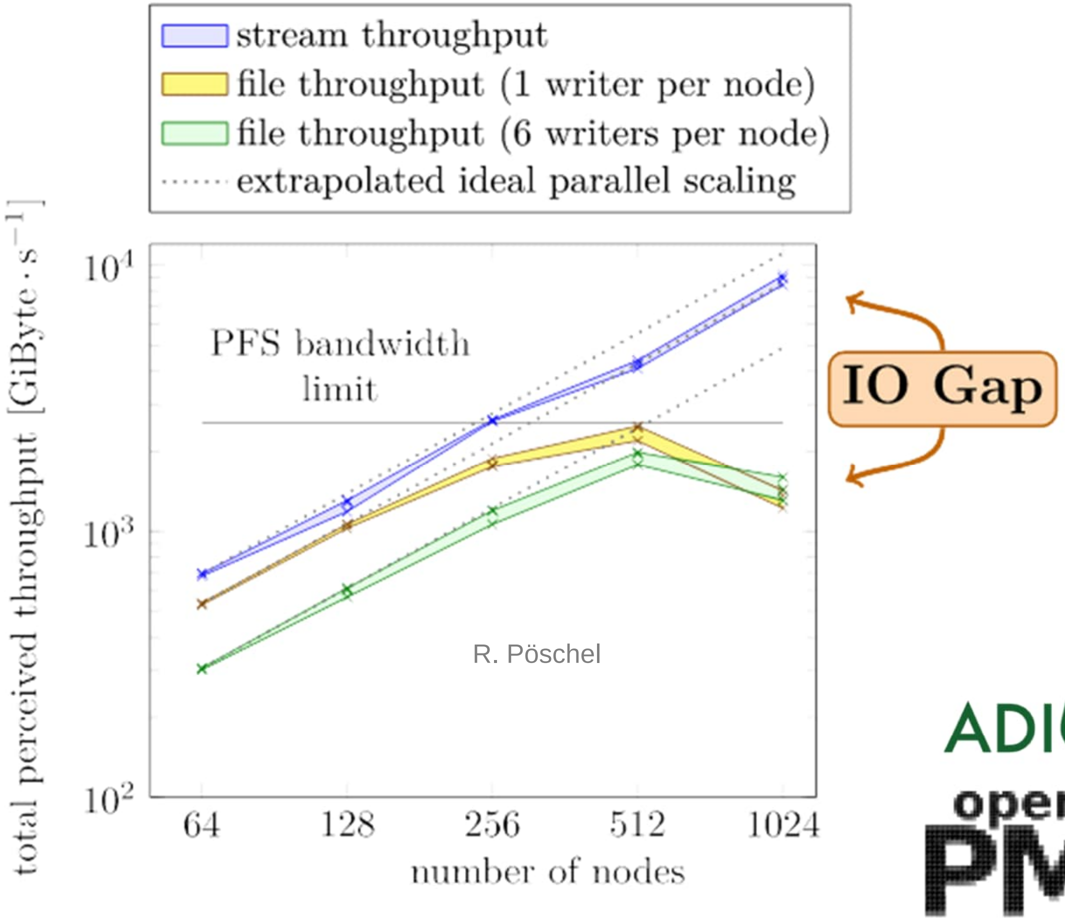

Interoperable:
Data exchange spans applications, platforms and teams

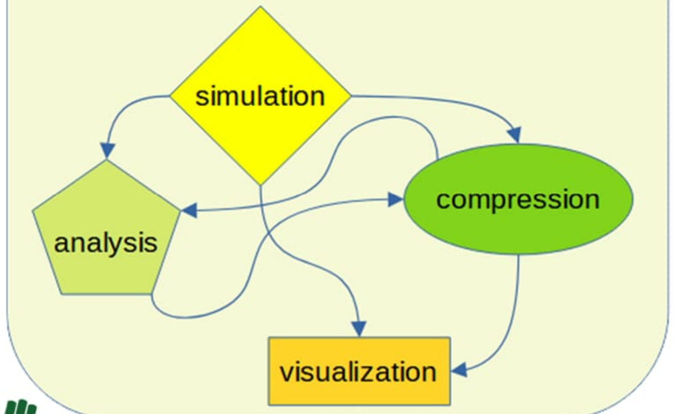

Reusable:
Rich and standardized description for physical quantities

Name	Value
axisLabels	[b'z' b'y' b'x']
dataOrder	b'C'
fieldSmoothing	b'none'
geometry	b'cartesian'
gridGlobalOffset	[0. 0. 0.]
gridSpacing	[4.252342 1.0630856 4.252342]
gridUnitSI	4.1671151662e-08
position	[0. 0. 0.]
timeOffset	0.0
unitDimension	[-3. 0. 1. 1. 0. 0. 0.]
unitSI	15399437.98944343

F. Poeschel

ML surrogate models are now an HPC task by themselves

ML scalability becomes important as high fidelity is needed

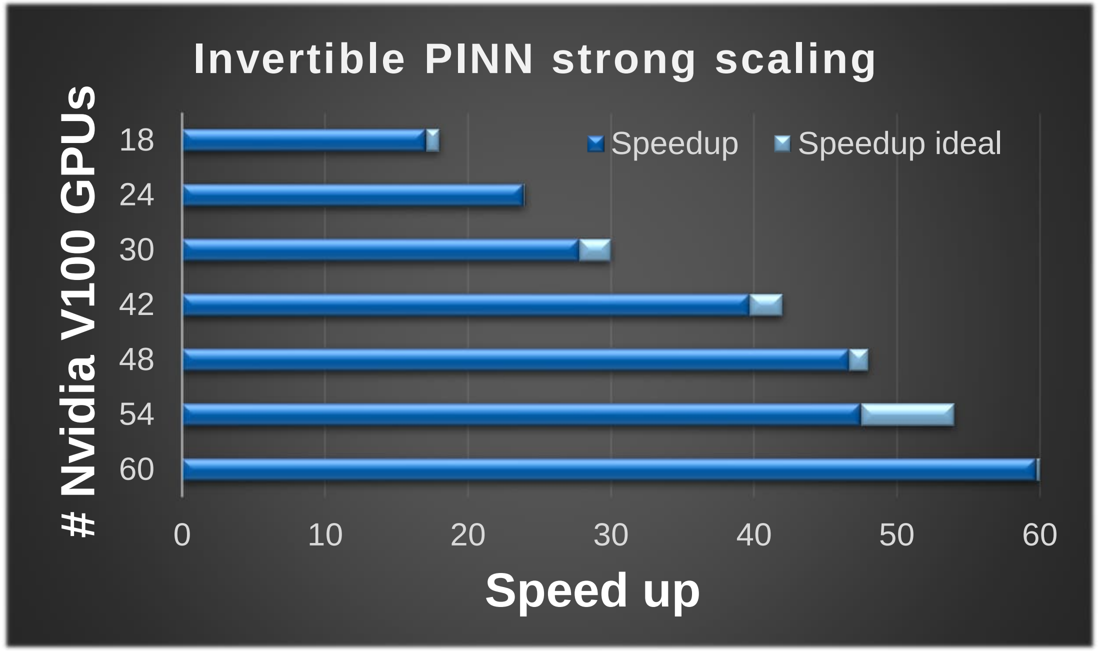

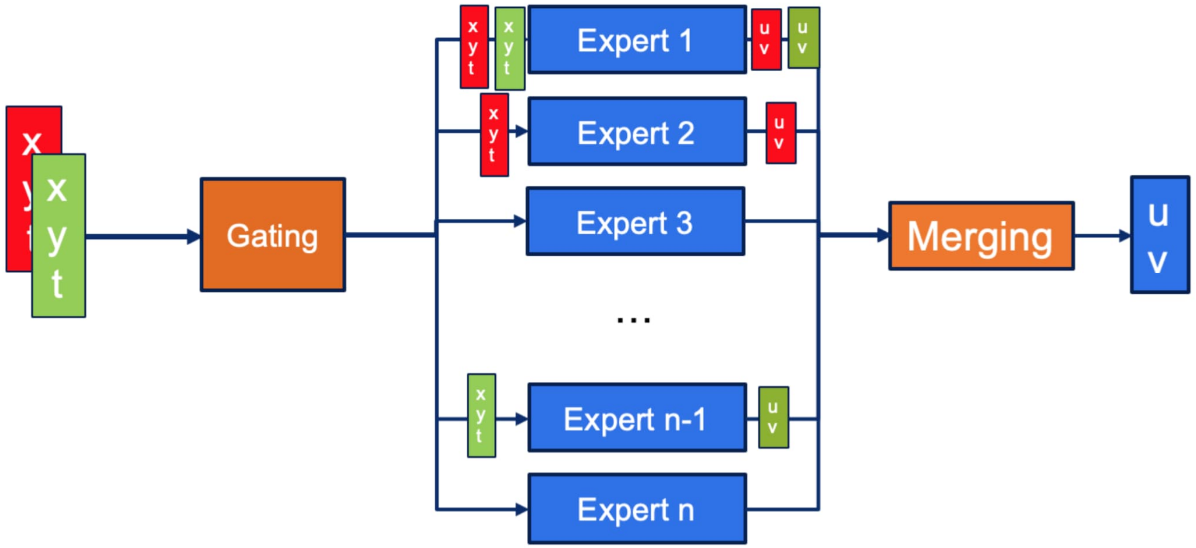

N. Hoffmann

NEURAL SOLVERS
DRESDEN concept
HZDR

Reconstructing electron dynamics from radiation spectra @ PB/s

Radiation from plasmas tells us their composition, structure and dynamic evolution

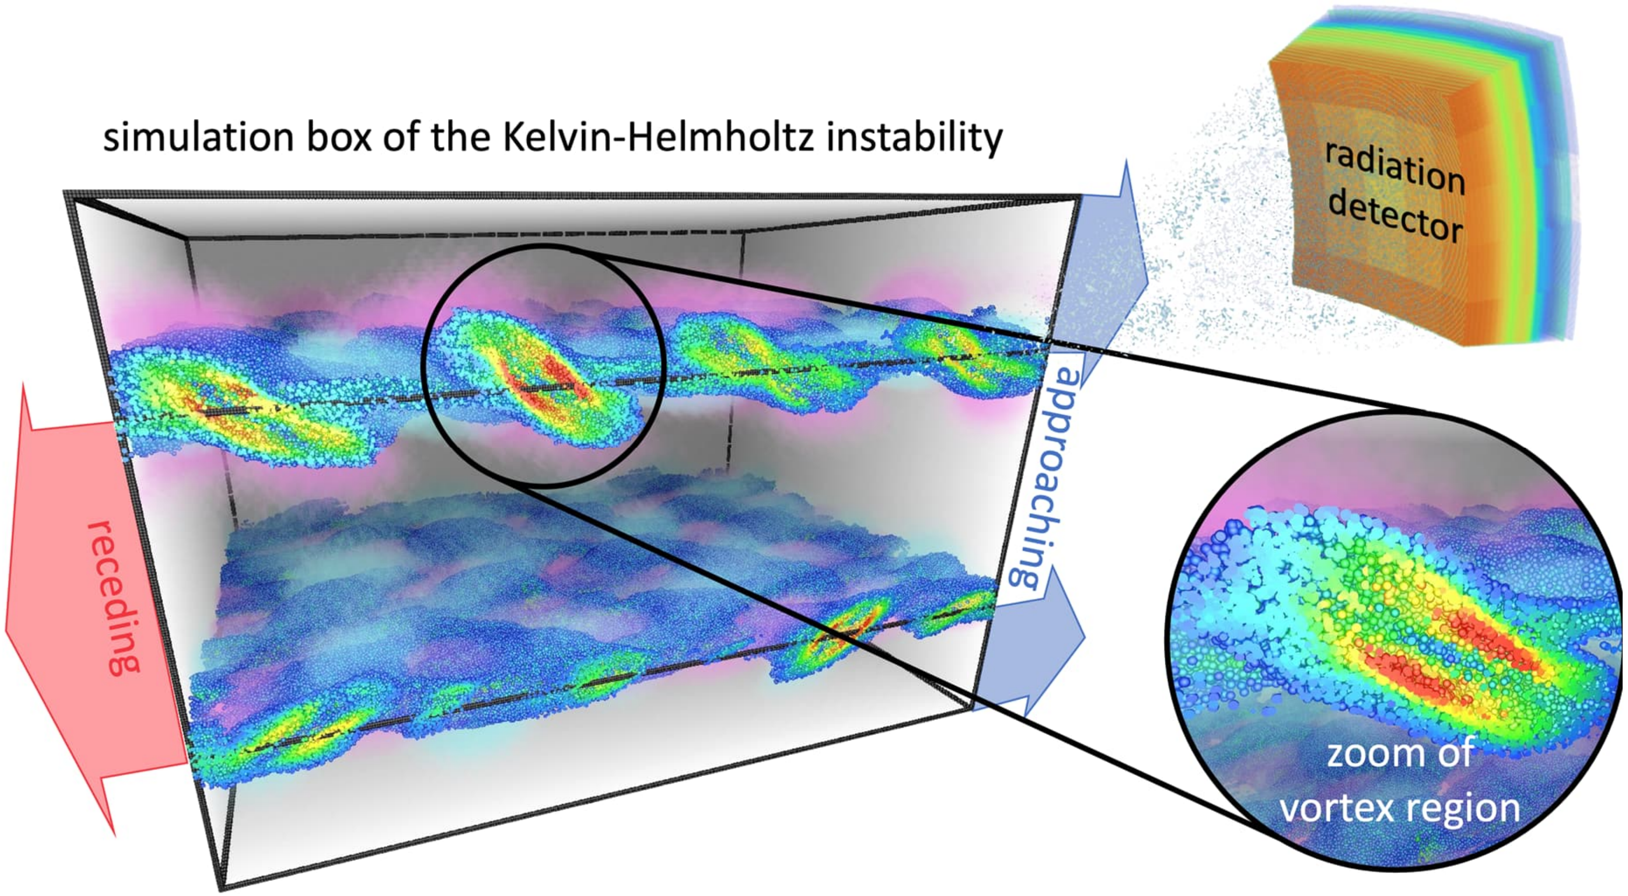

J. Kelling, R. Pausch

Train neural networks for:

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

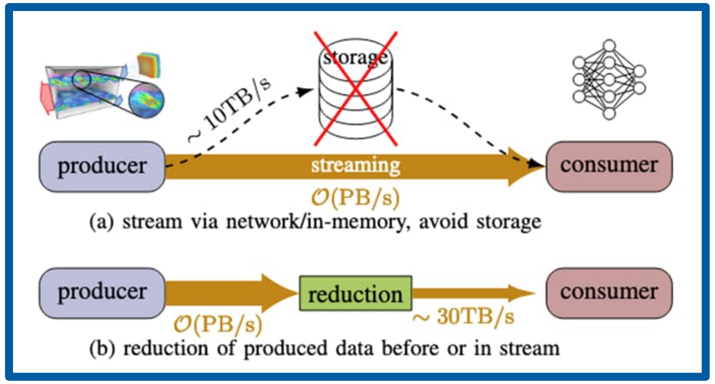

Reconstructing electron dynamics from radiation spectra @ PB/s

Radiation from plasmas tells us their composition, structure and dynamic evolution

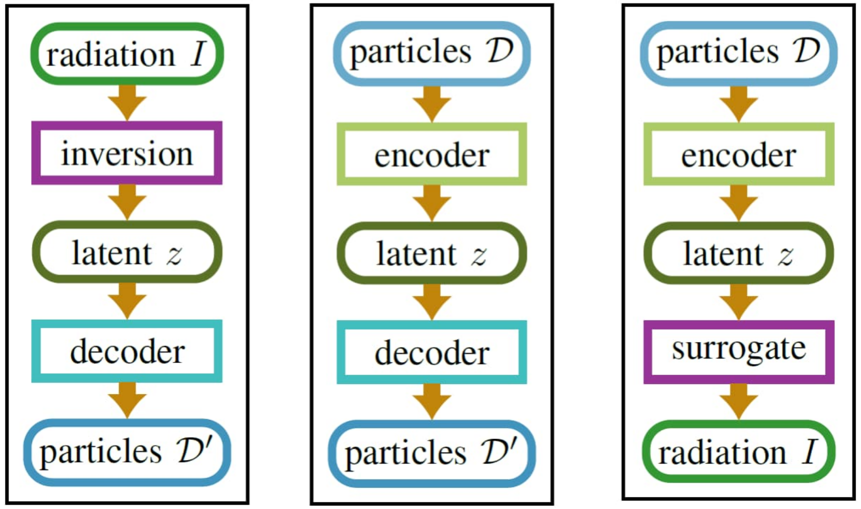

(a) inversion of physical reduction

(b) compression – decompression

(c) physical forward prediction

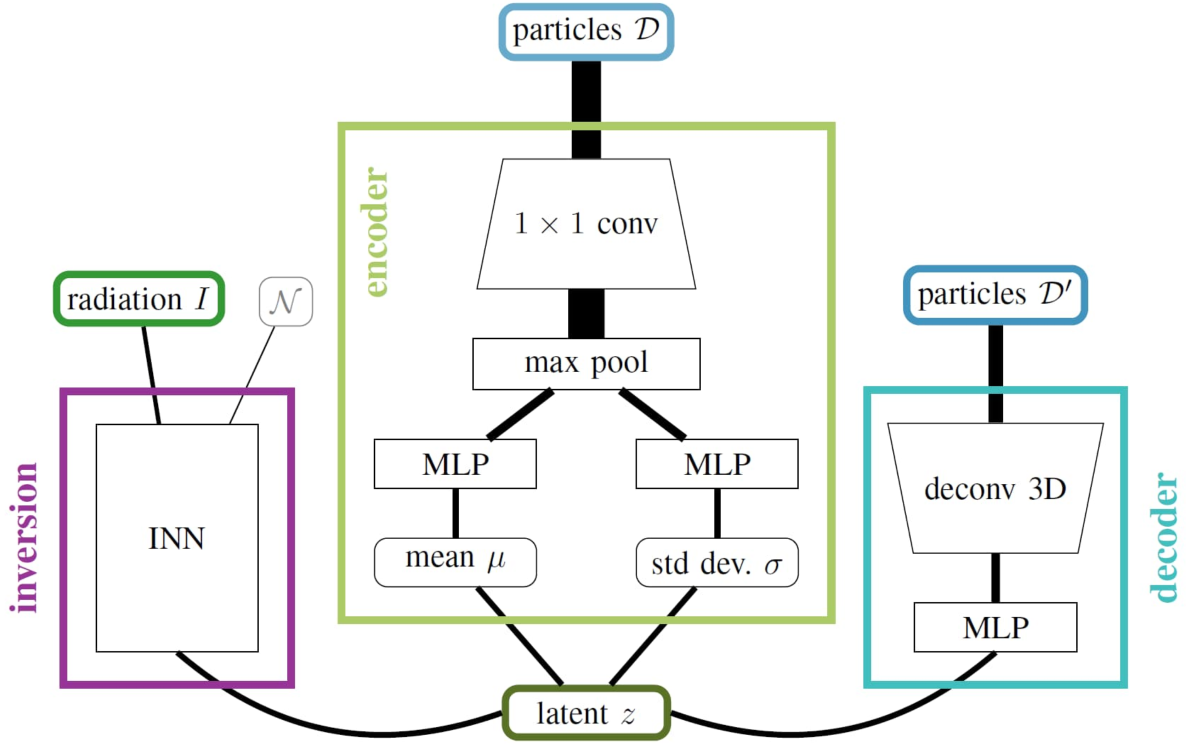

J. Kelling, R. Pausch

Standardization of data and metadata for LPI simulations

Input / output data workflows

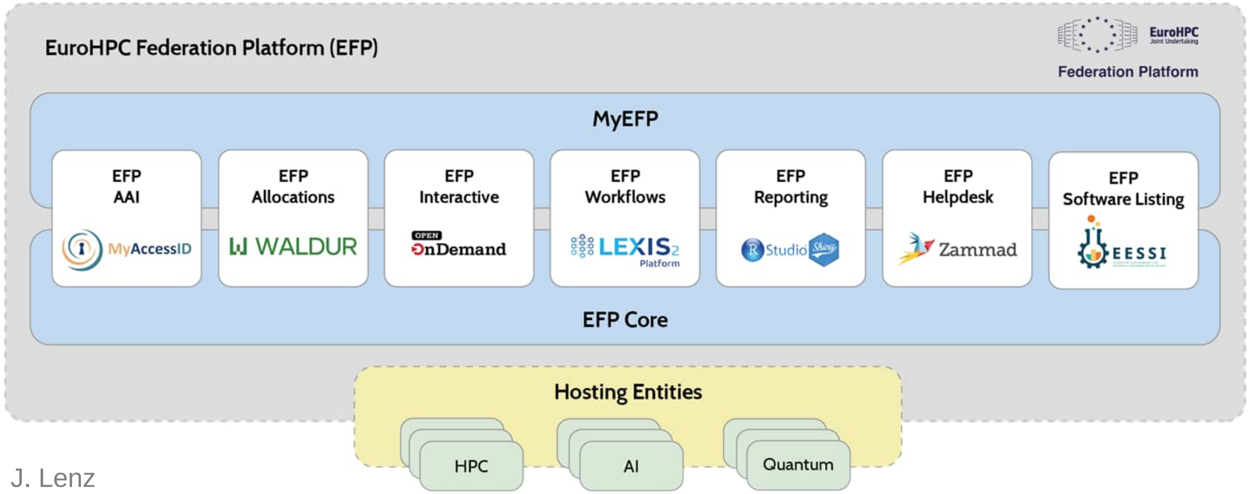

J. Lenz

- Pydantic input interface
- Research object crate support
- Common Workflow Language support
- F.A.I.R. input/output data ecosystem
- Ready for agentic AI workflows

DRESDEN concept HZDR

Complex AI/HPC workflows

Input / output data workflows

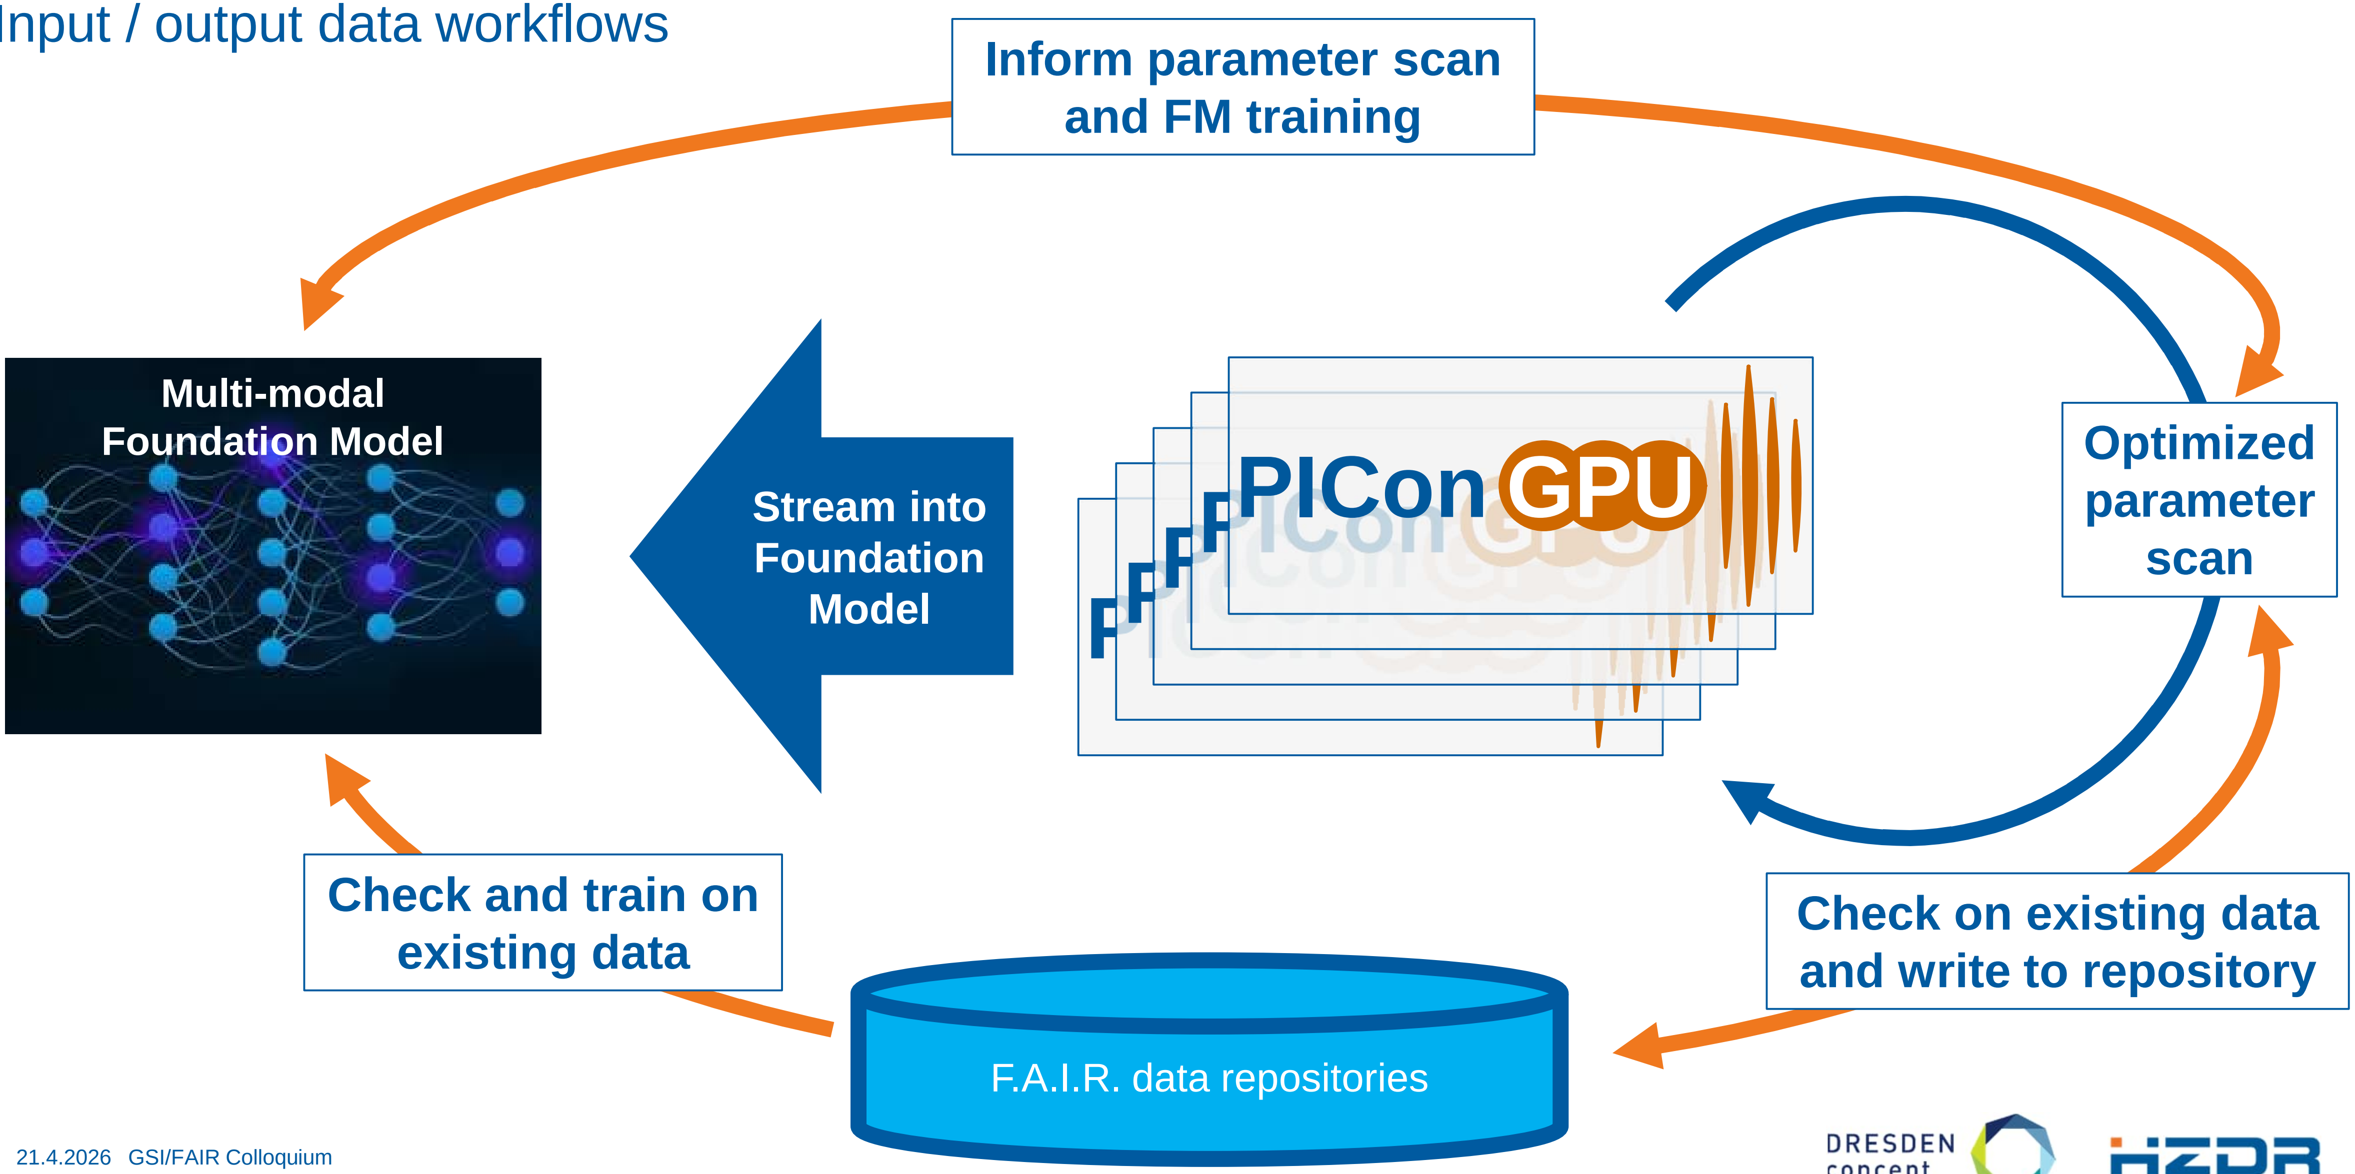

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation for applications
2. High-quality diagnostics and control are needed to know and set the operational parameters
3. At the same time, we must have high-quality, predictive simulations validated in experiments
4. We lose our high fidelity knowledge as simulation data is becoming too large
5. Our experiments now reach similar fidelity as our simulations, with similar data problems ahead
6. AI helps us connect the microscopic system state to observation and vice versa

Lessons learned

Obvious lessons with less obvious consequences

1. We must find the optimum parameters for the intended operation for applications
2. High-quality diagnostics and control are needed to know and set the operational parameters
3. At the same time, we must have high-quality, predictive simulations validated in experiments
4. We lose our high fidelity knowledge as simulation data is becoming too large
5. Our experiments now reach similar fidelity as our simulations, with similar data problems ahead
6. AI helps us connect the microscopic system state to observation and vice versa

Lessons learned

Obvious lessons with less obvious consequences

1. Reproducible, systematic, validated parameter scans in experiment and simulation
2. Automation and digitalization of processes
3. Coupling of experiment and simulation via synthetic diagnostics
4. Create optimized, data-driven, physics-informed surrogate models
5. Create a data and workflow infrastructure that is able to handle this automatically
6. Go for an AI-first strategy, not an AI-last/maybe strategy

HELMHOLTZ Matter MT DMA

HZDR
HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Building blocks for autonomous and intelligent research in Laser-Plasma Physics

DRESDEN concept

The pillars of modern AI-first data platforms

These building blocks are universal, yet often with specifics for each facility

Digitalization

Data Taking
- Automatic data taking
- Meta data enrichment
- Data harmonization (common data models)
- Electronic lab notebooks

Control
- lab equipment
- process paramters
- Ingredients

Interface to lab

Data Lifecycle Management

Pipeline
- Taking
- Processing
- Analysis
- Provenance
- Preservation

F.A.I.R
- Metadata
- Repositories
- Knowledge graphs
- Persistent identifiers

Persistent data object

Digital Twins

Workflows
- Digital process representation
- Directed acyclic graph
- Input/output paradigme
- State machine

Twinning
- Data-based
- Physics-informed
- Simulations
- Multi-scale
- Multi-physics
- What if?

Workflows

Access and Authorization

Data
- Federation
- Data Access
- Data Sharing
- Data Lifetime

Security
- Authentication
- User personae
- Team access
- Tokenized access

Regulatory
- Rule adherence

Restriction / visibility

AI

Optimization
- ROI
- Time to solution
- Throughput
- Quality
- Resources
- Multi-parametric

Simulation
- What-if scenarios

Knowledge Representation

Discovery

ML-Ops

The computer science view on reality

Finite state machines as atoms of digital processes

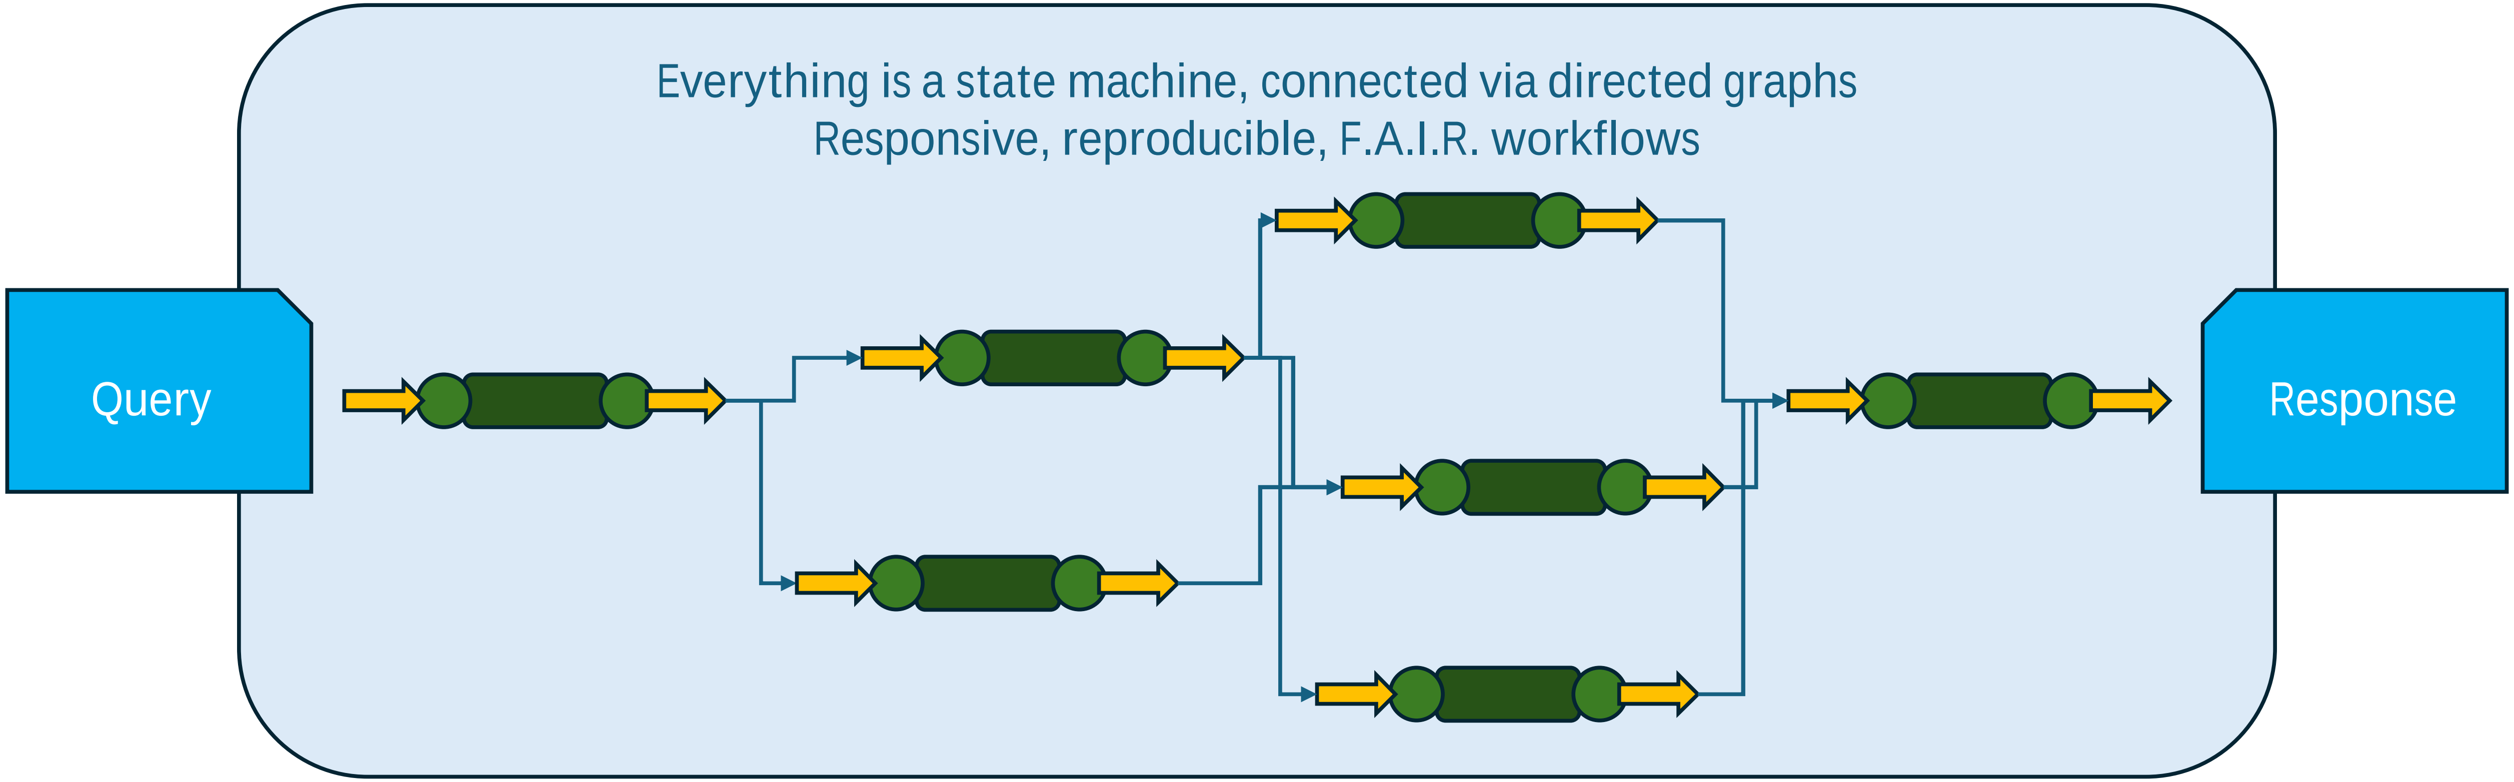

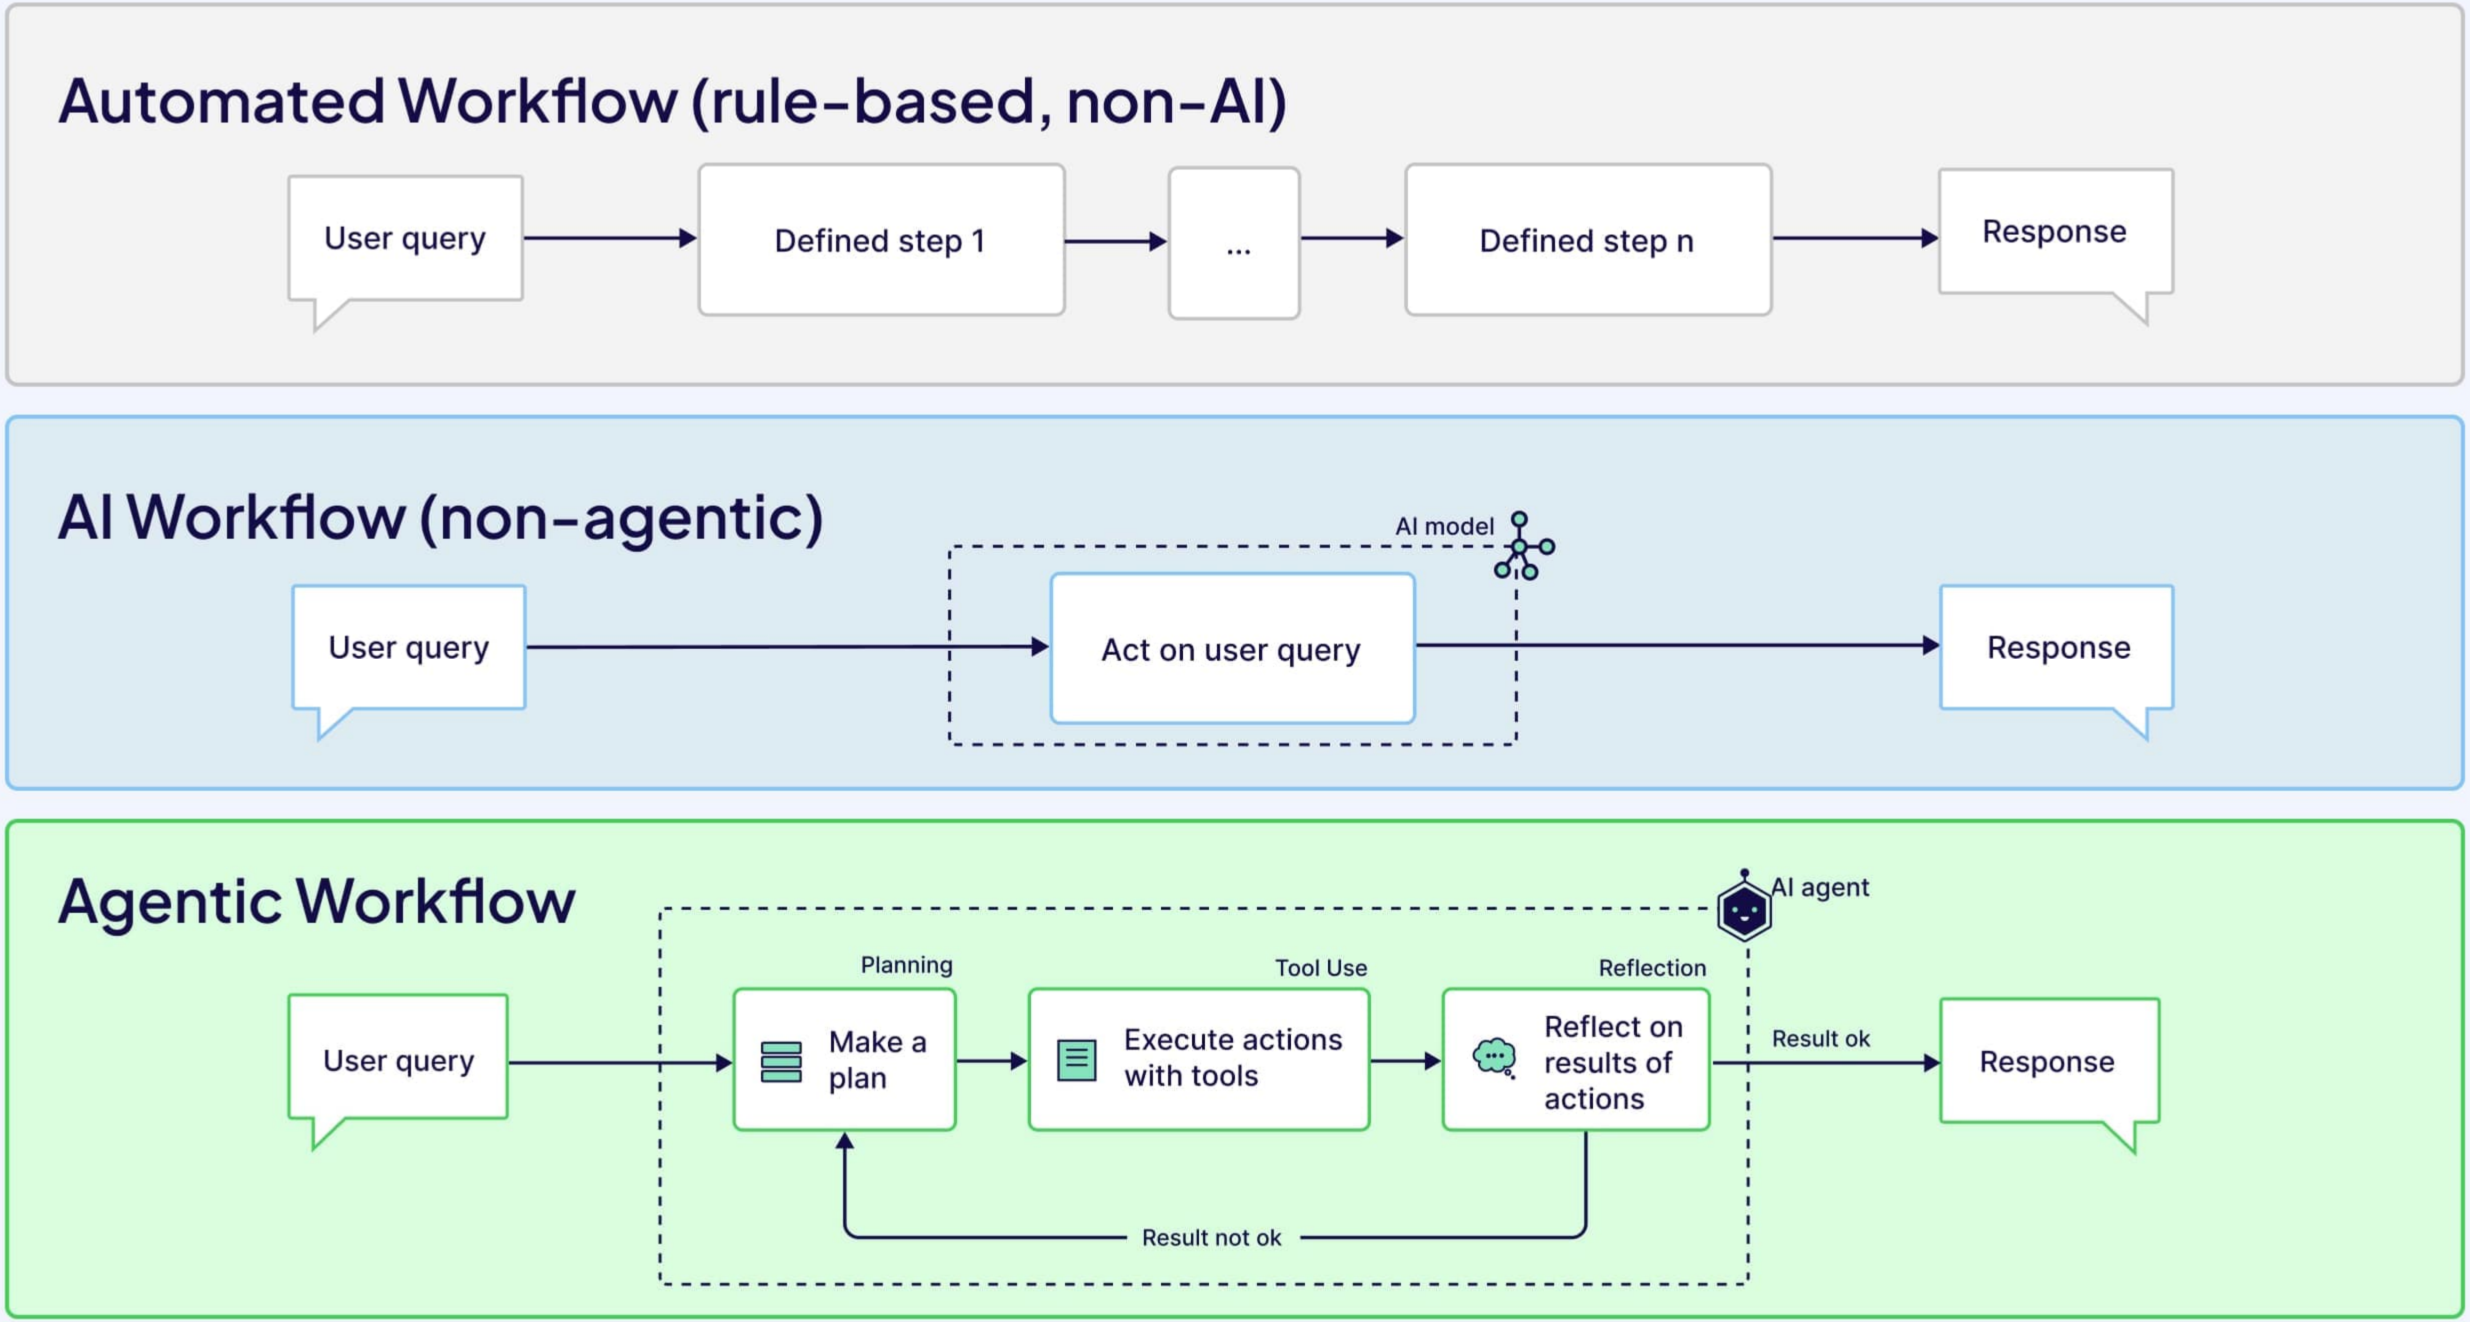

Automated Workflow (rule-based, non-AI)
User query
Defined step 1
...
Defined step n
Response
AI Workflow (non-agentic)
AI model
User query
Act on user query
Response
Agentic Workflow
AI agent
Planning
Tool Use
Reflection
User query
Make a plan
Execute actions with tools
Reflect on results of actions
Result ok
Response
Result not ok

Components of AI Agents

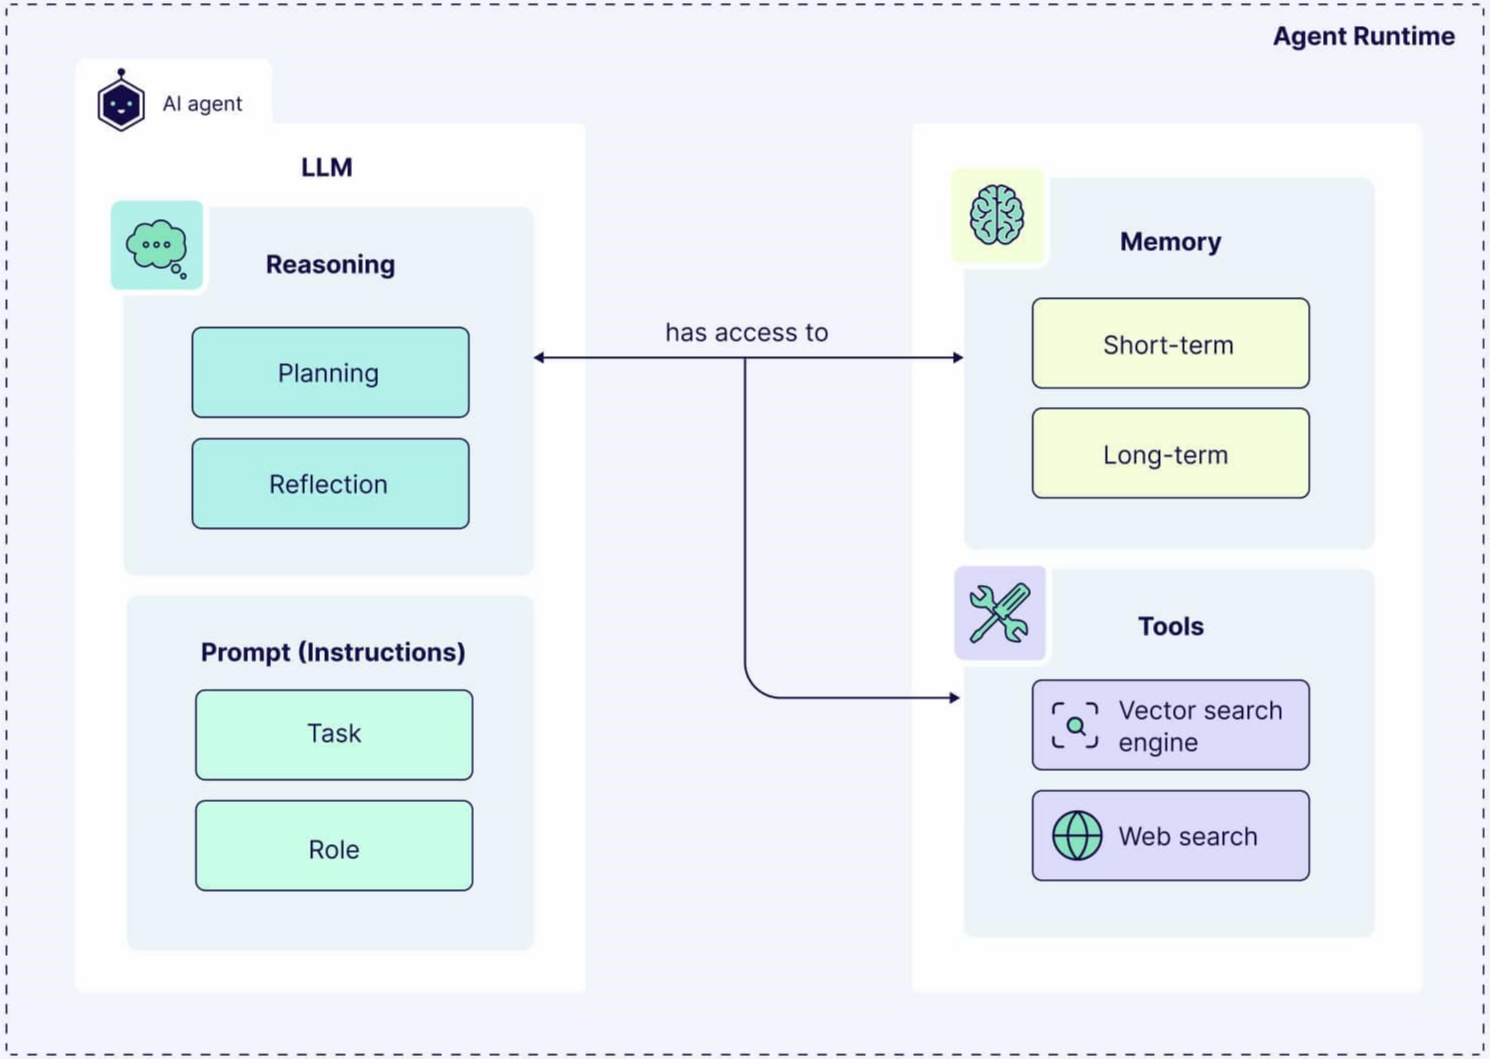

Planning Pattern

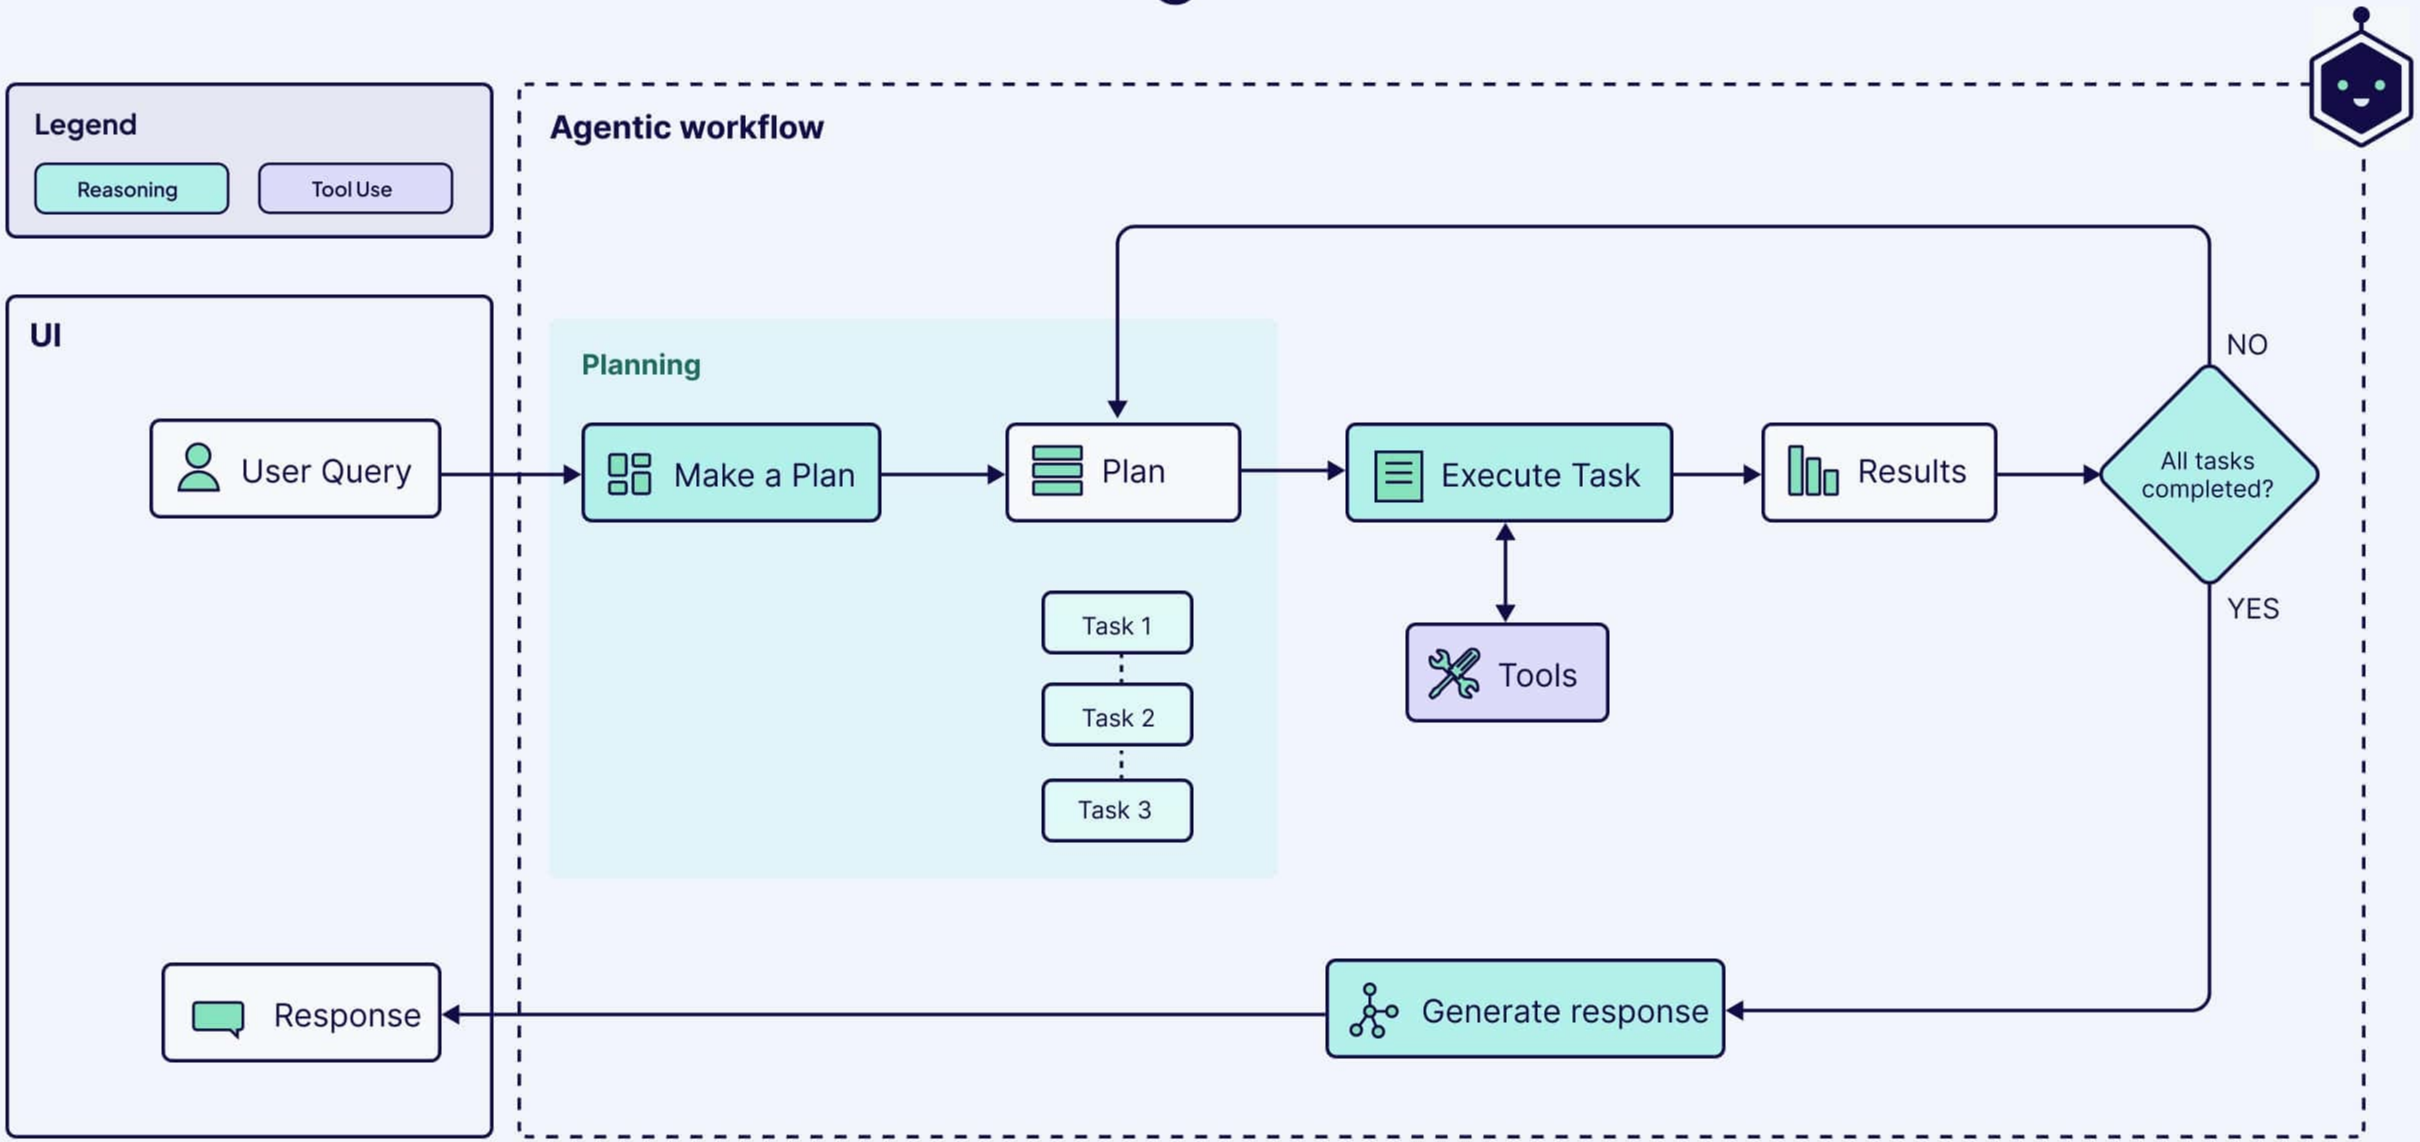

Tool Use Pattern

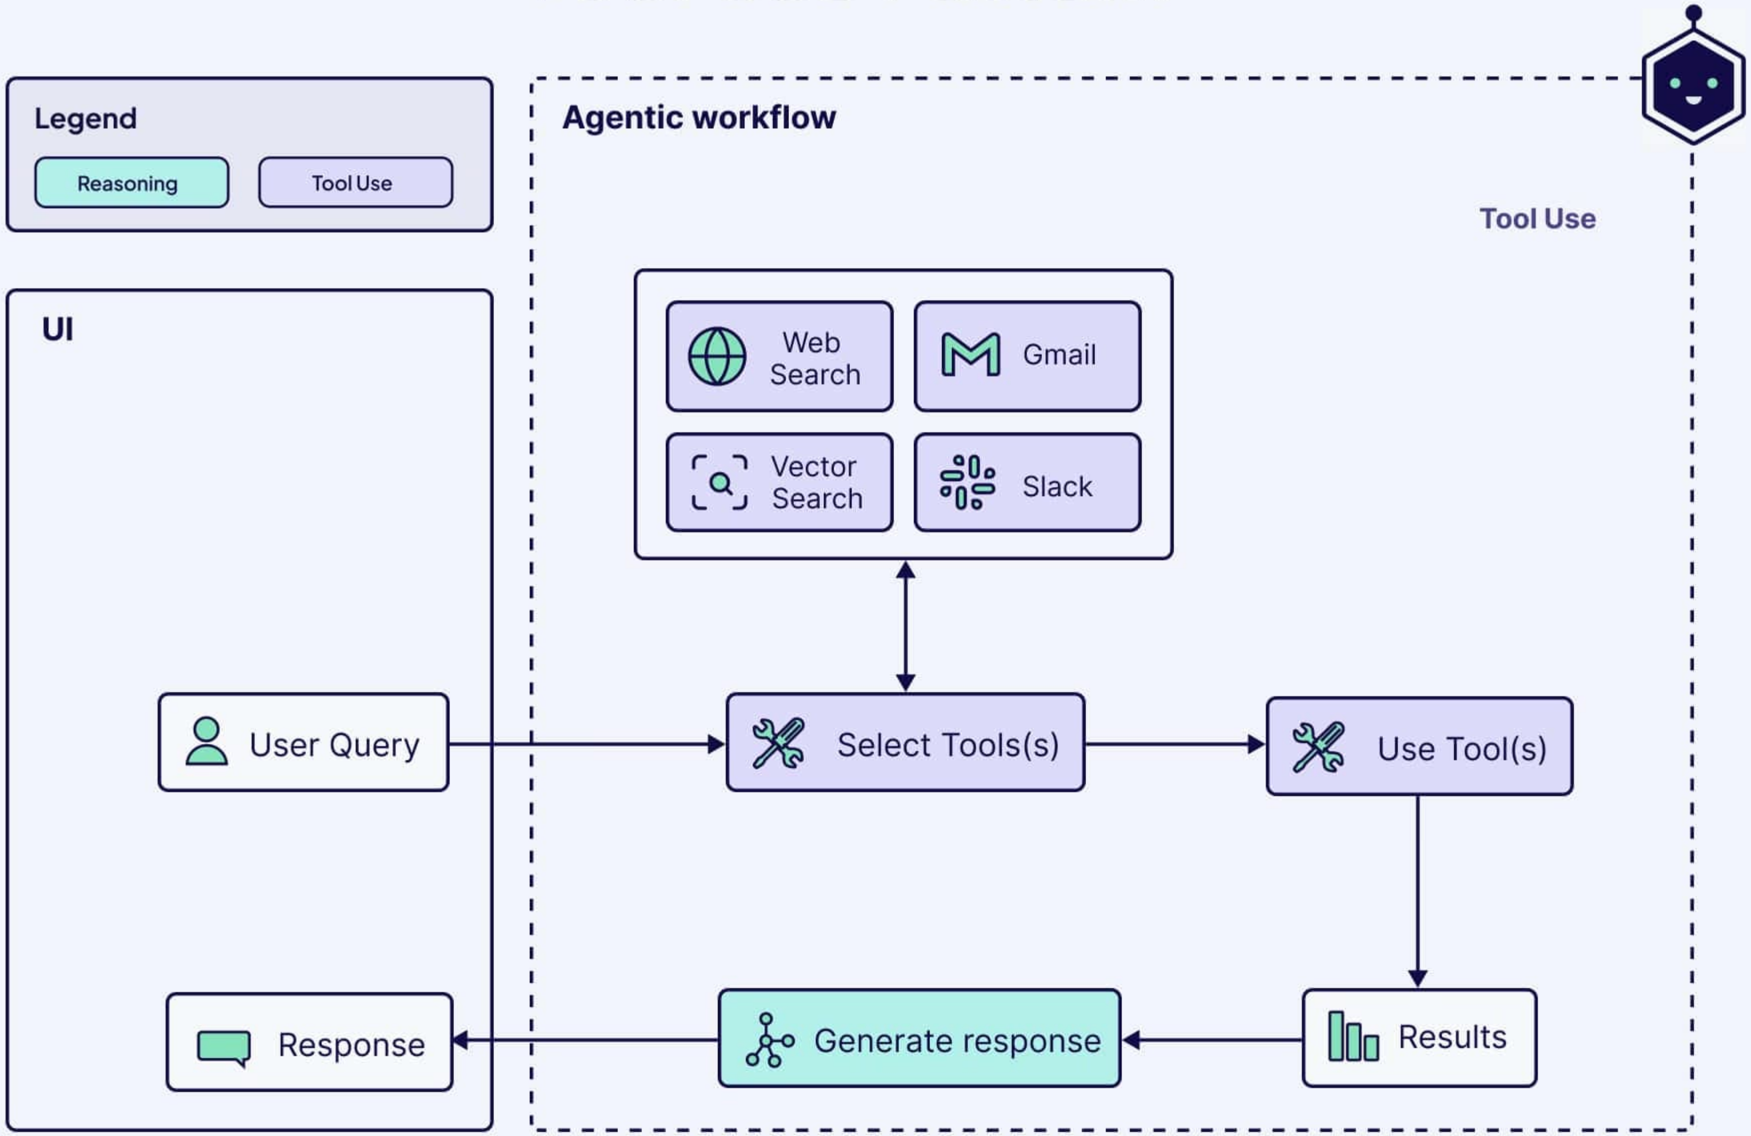

Reflection Pattern

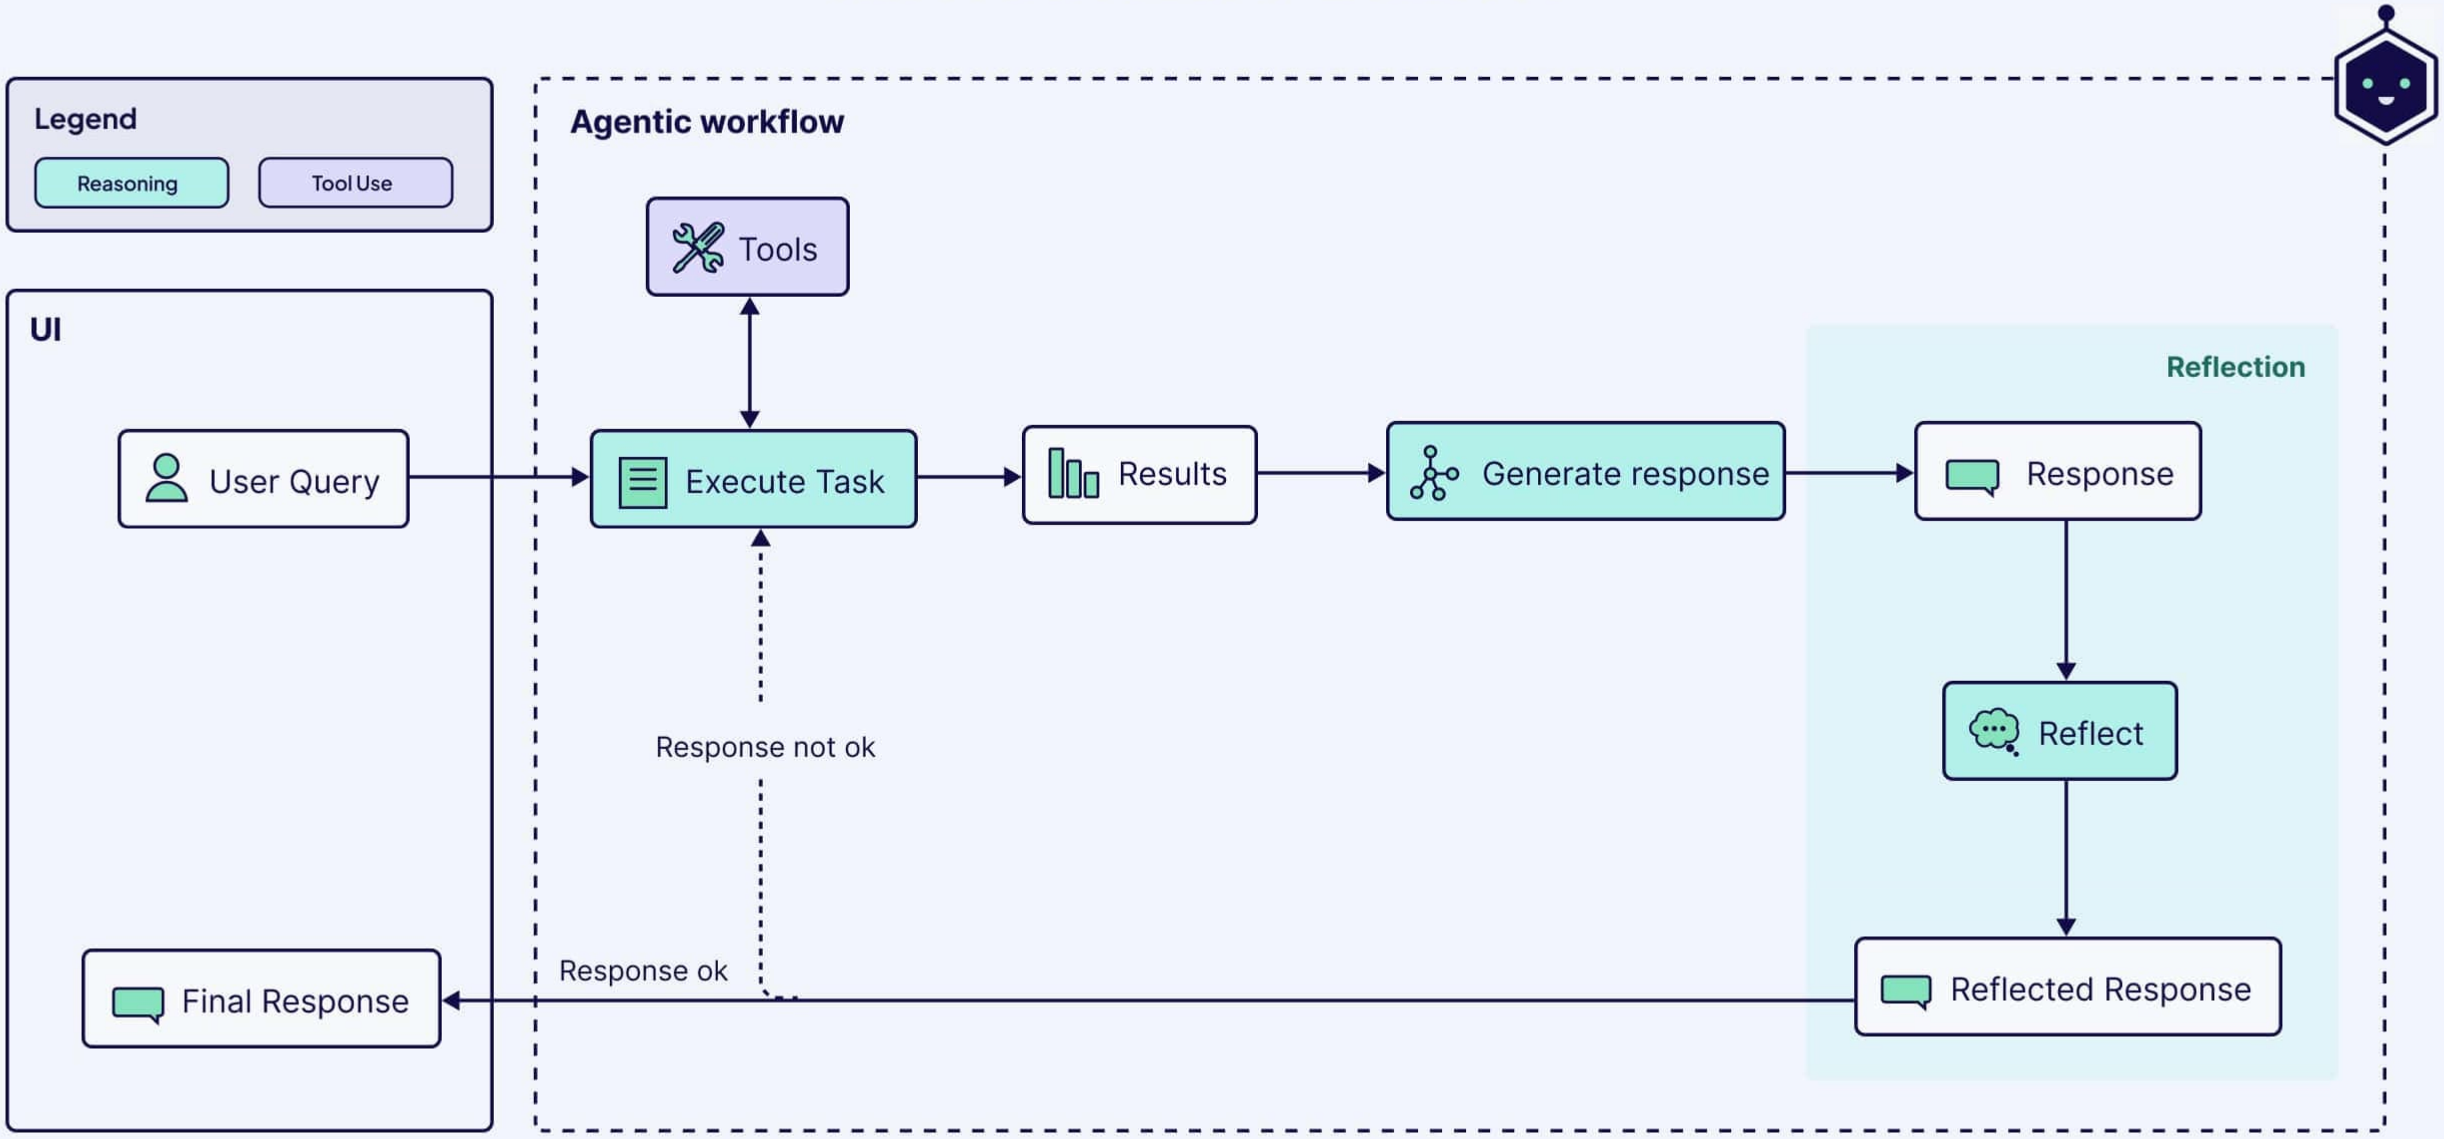

The many layers of AI

AI understands and acts on inputs, producing actionable outputs

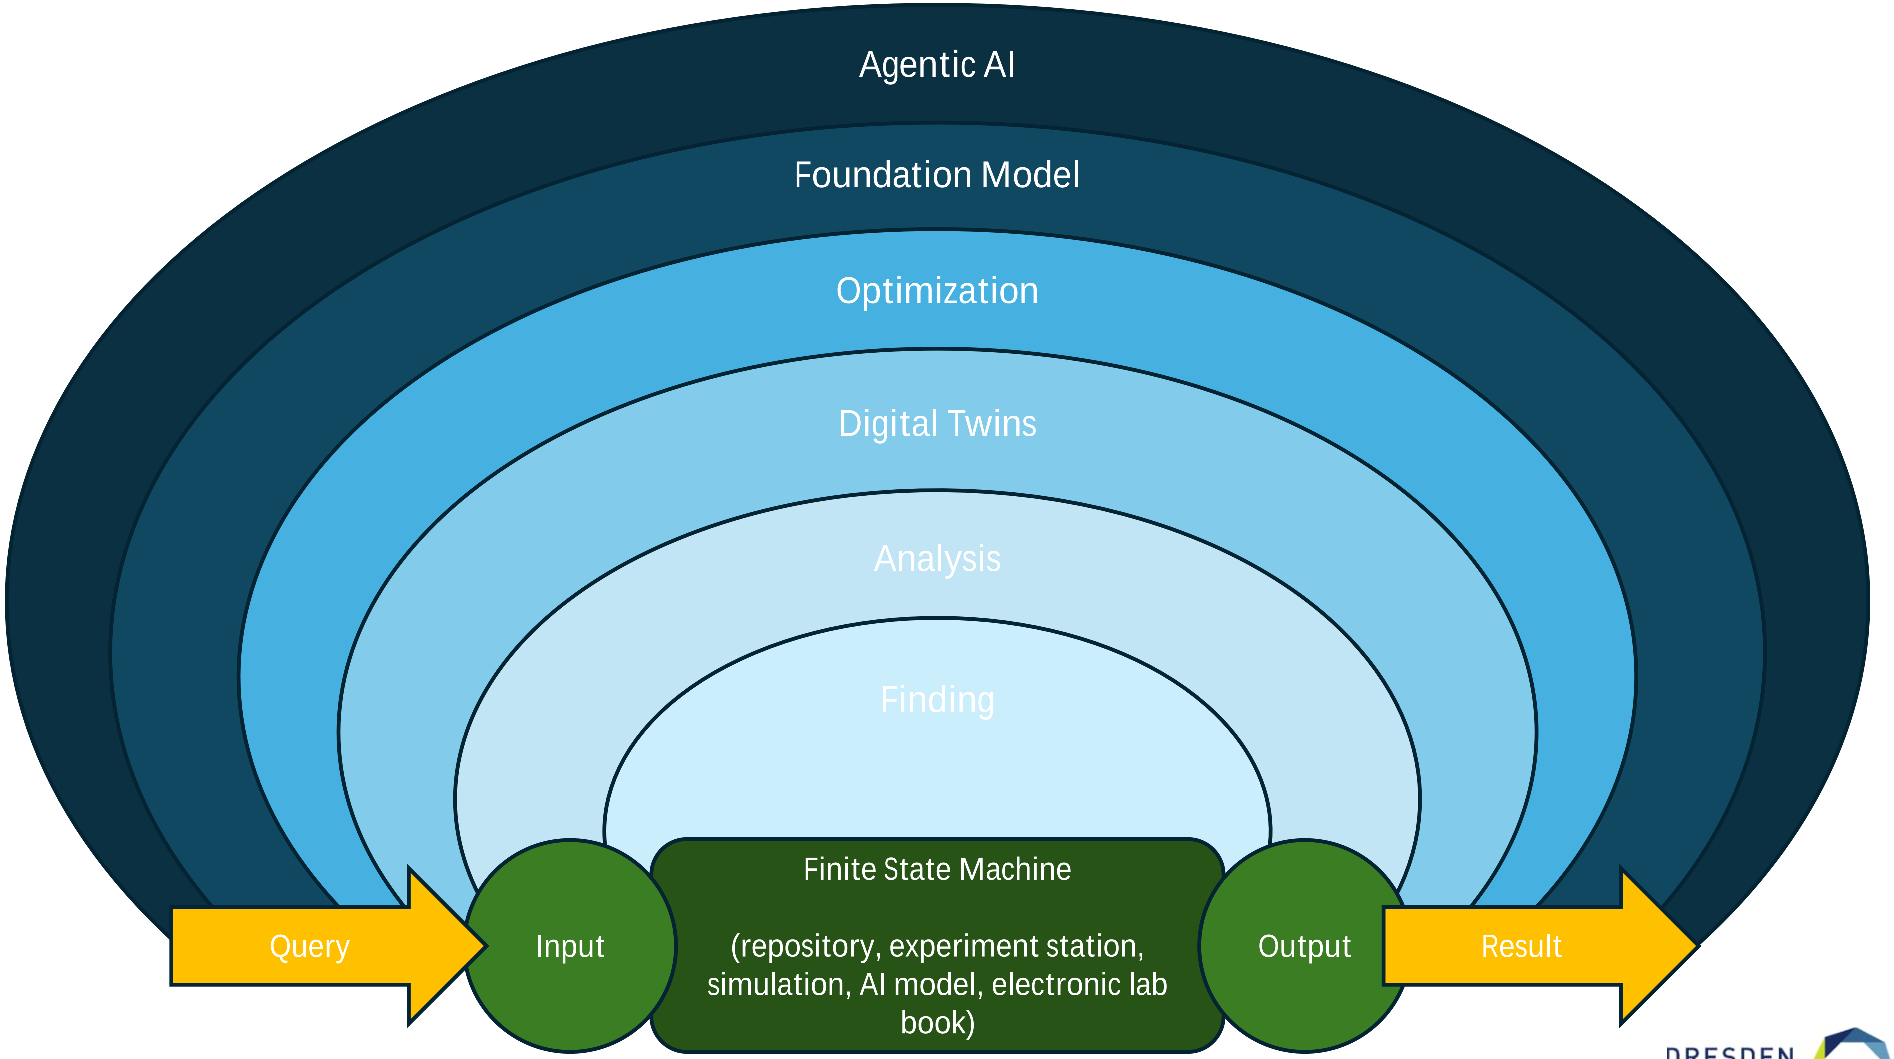

ML-Ops and Agent-Ops as the next step in evolution of workflows

AI models and agents combine processes and data

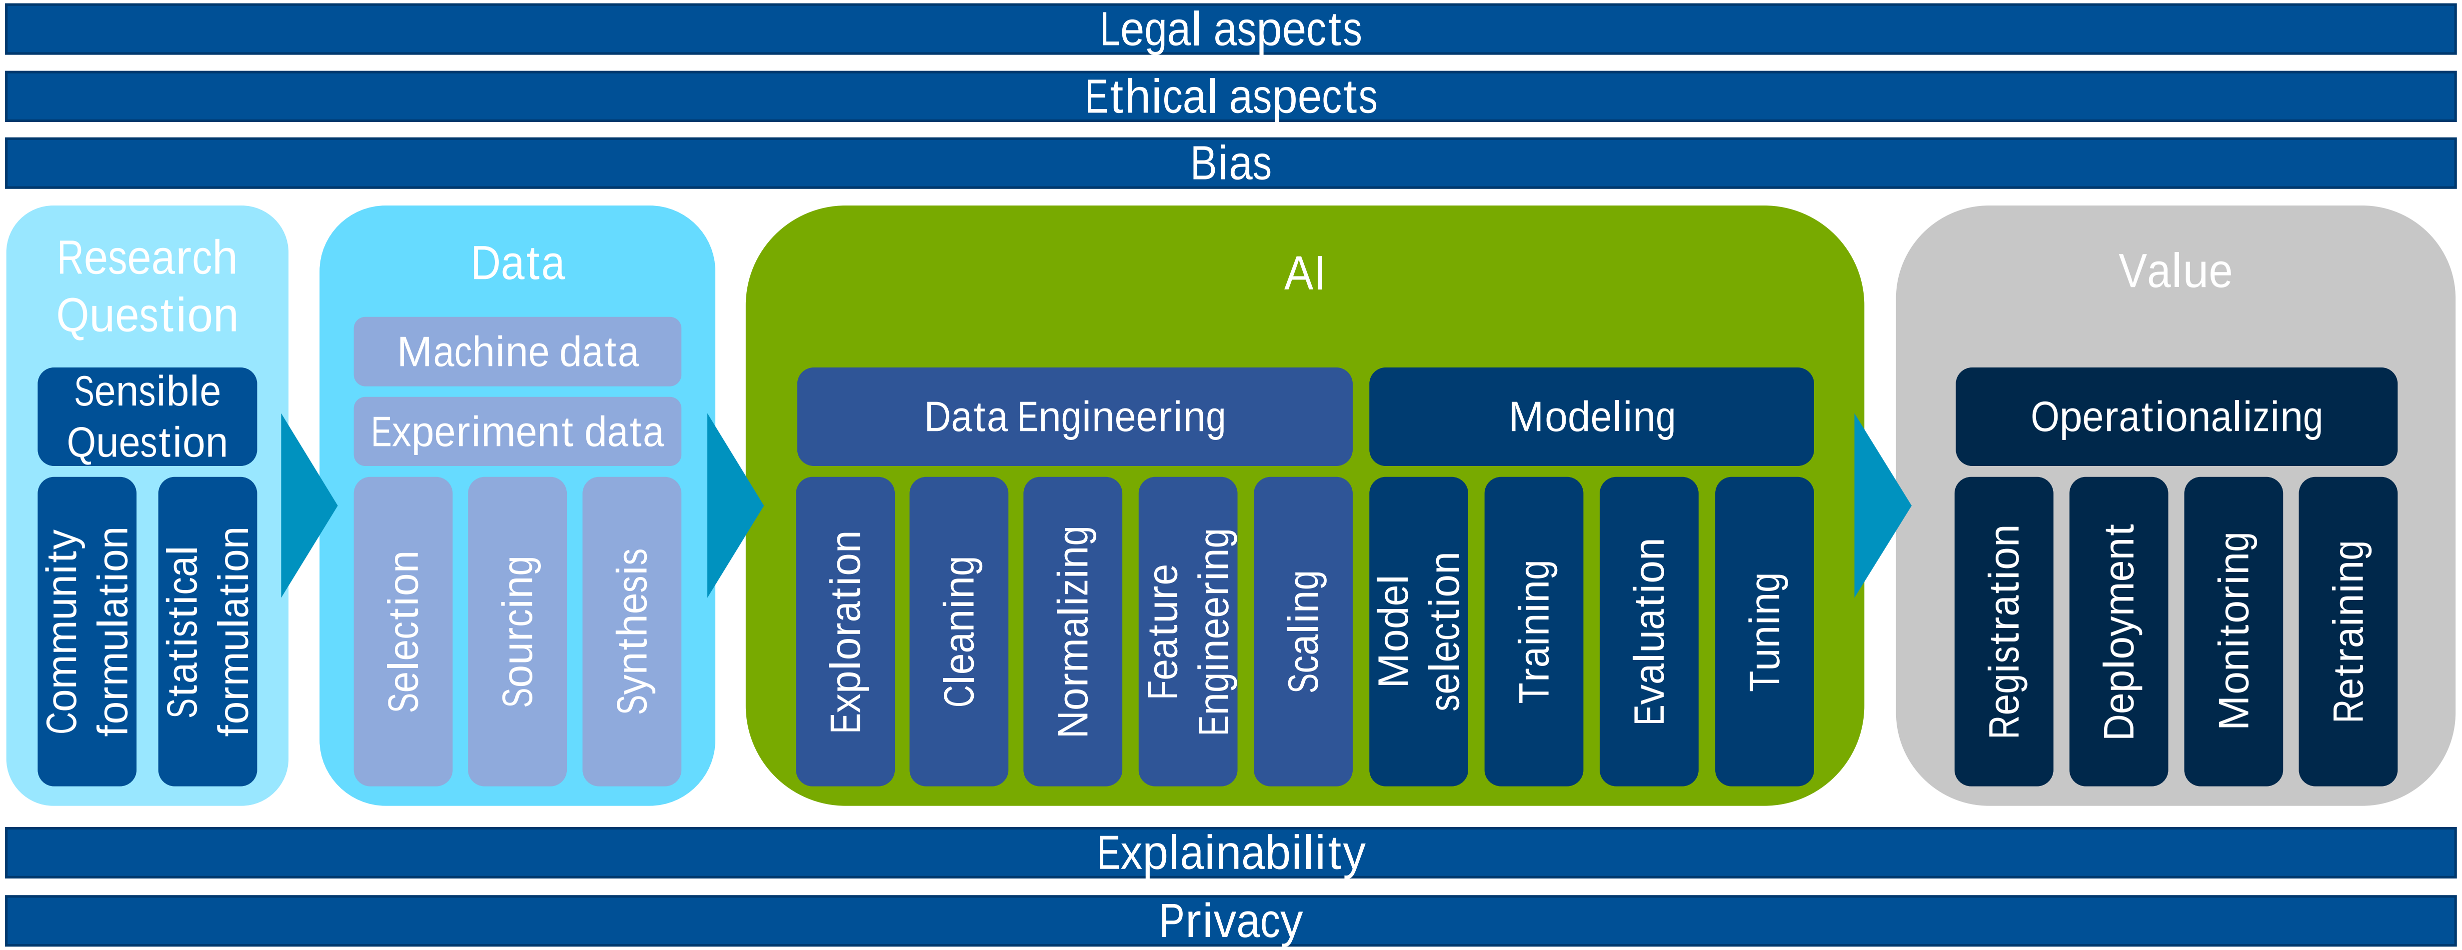

How should we build this?

A birds-eye view on the digital infrastructure we will need

(Agentic) AI layer

Federated AI layer

ML-Ops / Agent-Ops infrastructure

Access-controlled ML-model repositories

Workflows as digital twins of experimental processes

Access-controlled workflow repositories

F.A.I.R., semantic data and metadata

Access-controlled data repositories

Secure data access and experimental control layer / facility middleware as a sandbox

Can we build this?

We align well with national and international research strategies

- AI-sovereignity builds on data from infrastructures
- Fusion strategy needs a path from research to power plant
- Laser-plasma research is key ingredient
- We have the building blocks, but we must put them together