IT Infrastructure of the Future

Envisioning the Future of Computing

Dr. Oliver Oberst
IBM HPC/AI Architect
IBM Q Ambassador

 #IBMQ

IBM Research: Diversity of Disciplines

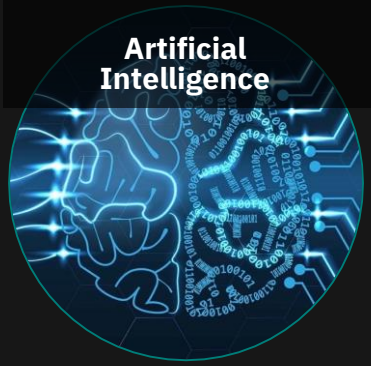

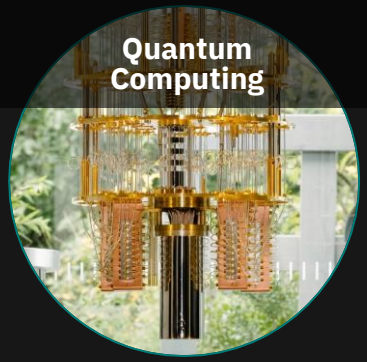

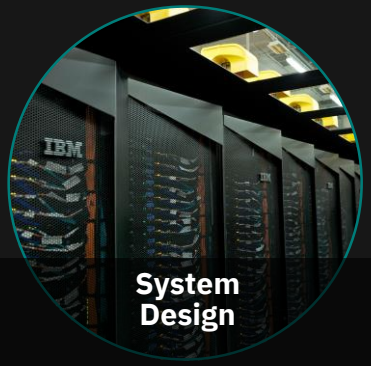

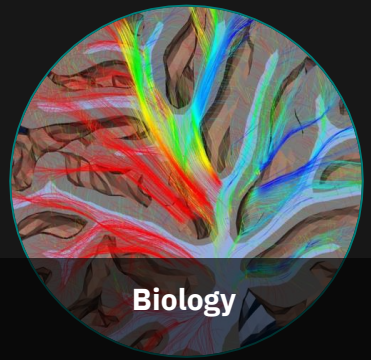

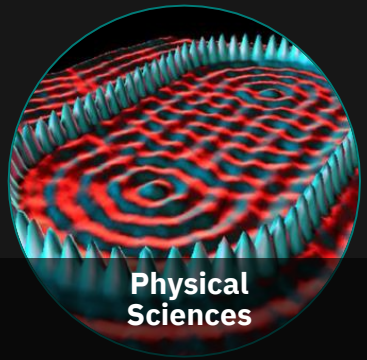

Strategic focal area: Reimagining computing

A quest for new materials, devices and architectures to radically change what it means to compute

The future of computing

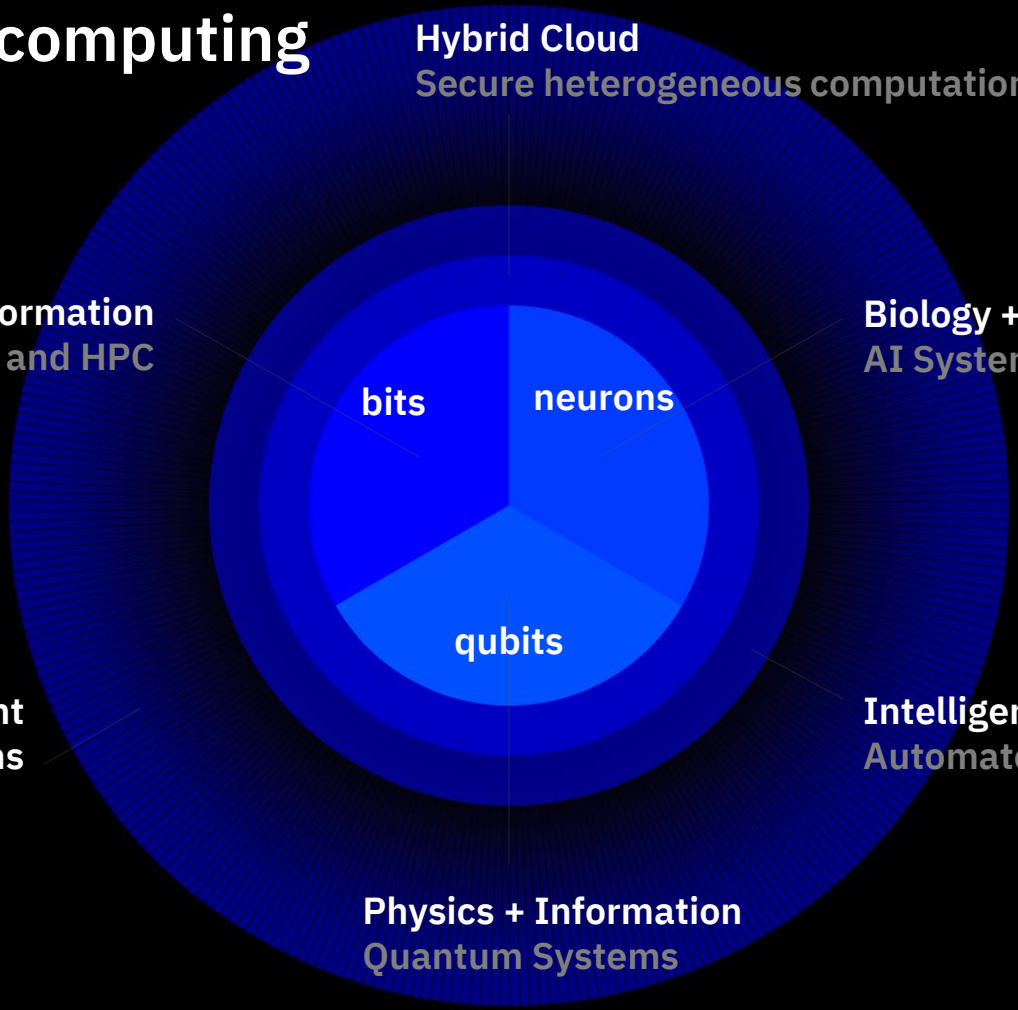

The future of computing

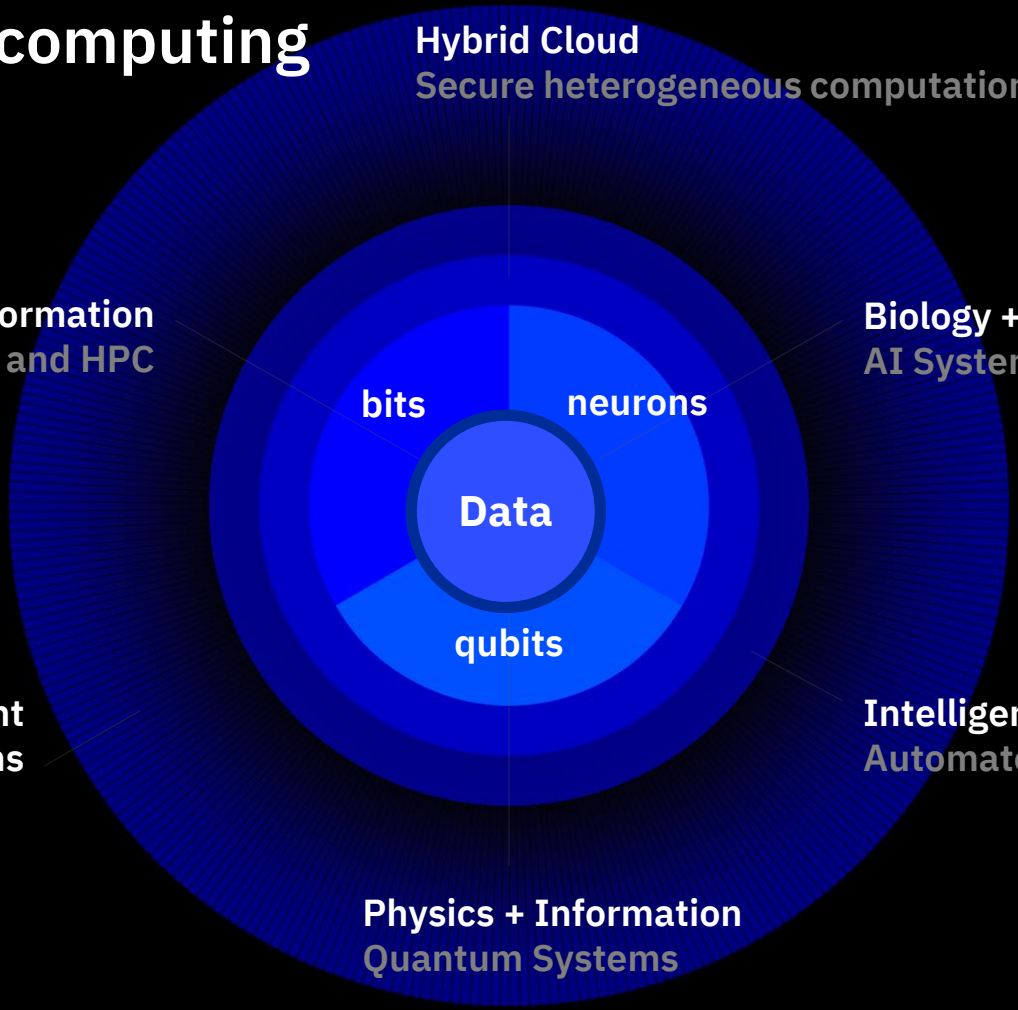

The world’s most powerful “bits + neurons” system

Oak Ridge National Laboratory
US Department of Energy

Summit supercomputer specs

200,000
trillion calculations
per second

9216
IBM Power 9 processors

27,648
NVIDIA GPUs

250 PB
File System

 #IBMQ

Convergence of HPC, AI and HPDA

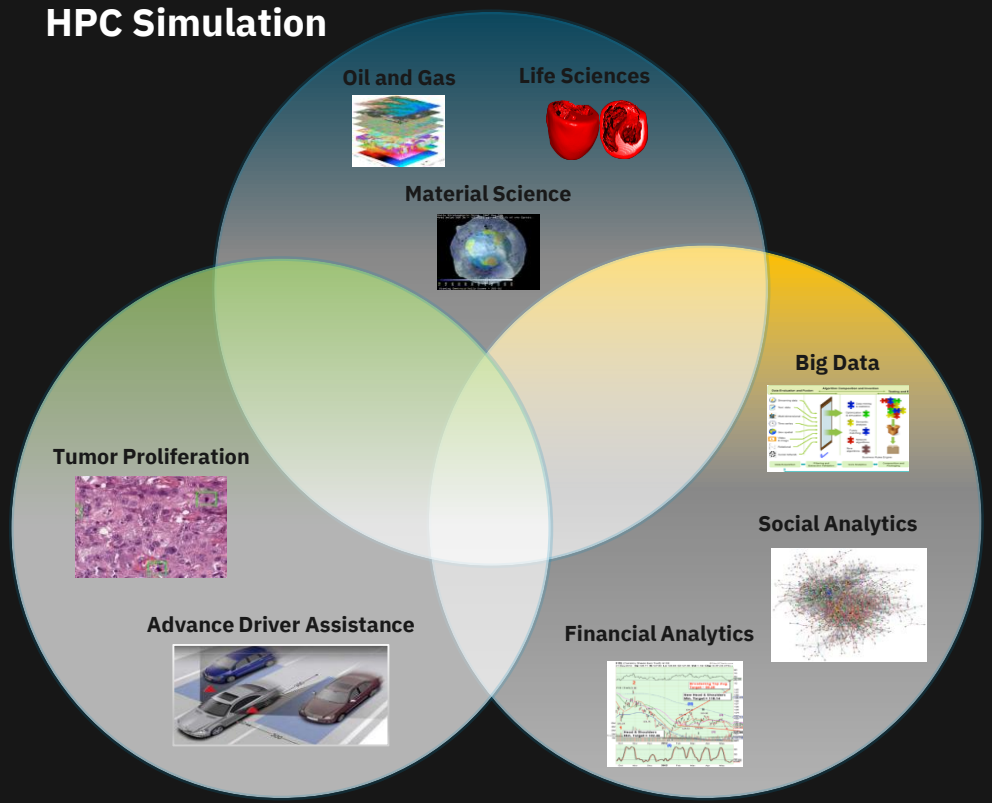

 #IBMQ

Science/R&D workflow

Knowledge

Researchers need to easily access quickly growing and widely diverse information sources.

- Highly unstructured/dark
- Current human based approach not scalable

Inference & Simulation

Domain related inference is largely missing. Setting up and deploying the right simulations is very hard.

- Human capital intensive, non scalable

Evidence & Experiments

Internal evidence and experiments are driven primarily empirically, often brute force, and their results are isolated from wider knowledge space.

Computing Reimagined: Knowledge Discovery Pipeline

Ingest structured and unstructured data to create massive knowledge spaces

- Complex documents
- Structured DBs
- Simulation/Previous runs
- Public and proprietary

Enable contextual based search, based on meaning, not keywords

- Tables, Images, Formulas, Diagrams
- *"What are all the properties of a chemical" -- "What is the cheapest health contract for non smokers younger than 25"*

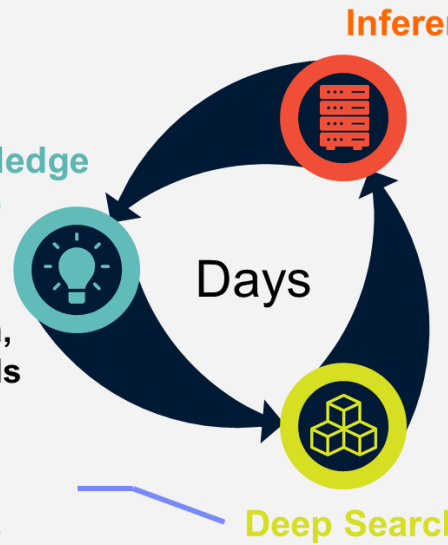

Get deep insights by ML/DL on the Knowledge Space

- Identify trends
- Discover gaps in the knowledge
- Explore "what if" scenaria
- **Run the "right" simulations**
- Enrich knowledge space

- *"How has the use of copper in alloys in the auto industry evolved in the last 20 years?"*
- *"What is the most likely use of a certain chemical by-product of an organic synthesis route next year? Which company is the most likely to want to buy it?"*
- *"How do I come up with the most relevant questions to ask a client for an insurance contract based on their answers so far?"*

Three Pillars of data-intensive Transformation using AI

Intelligent Simulation and Cognitive Discovery

Intelligent Simulation with IBM BOA

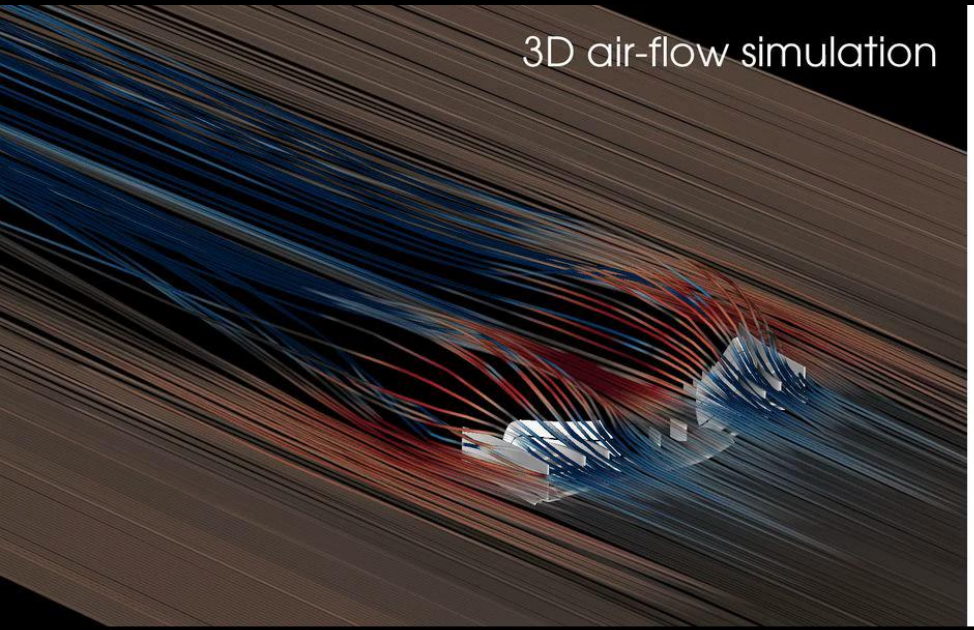

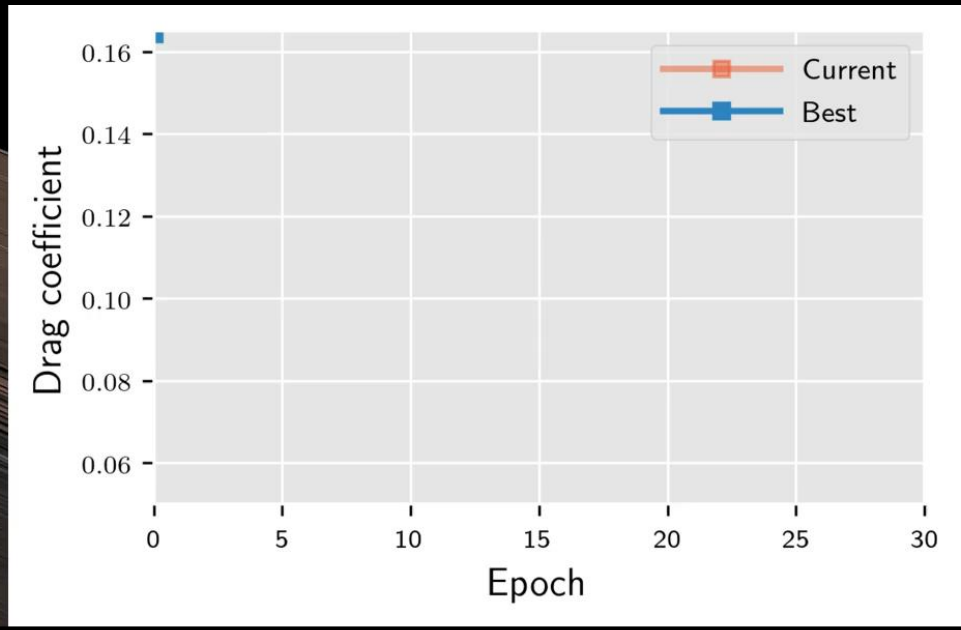

Technical R&D today: Disruption opportunity

Today

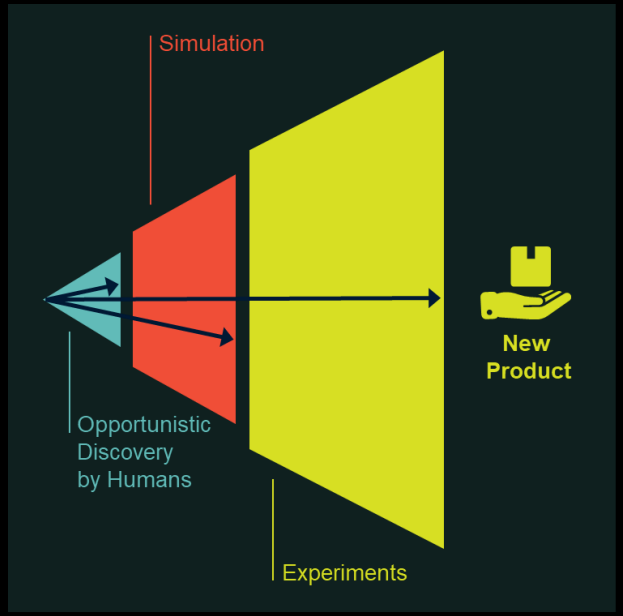

Cognitive Discovery

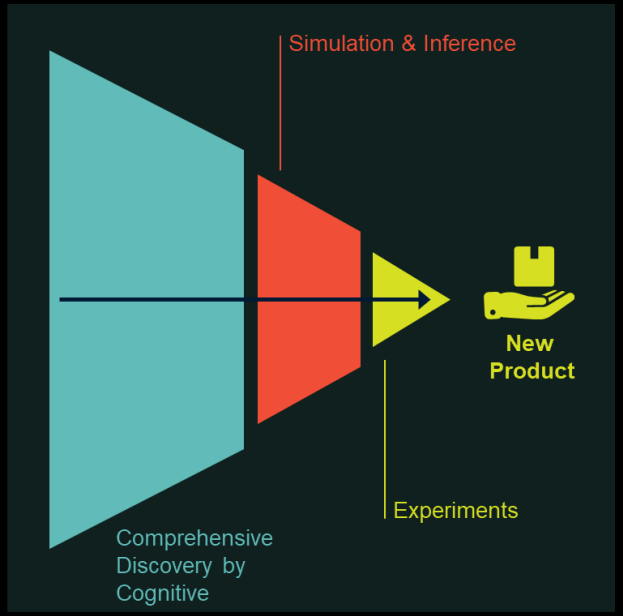

Cognitive Discovery Example App: IBM RXN

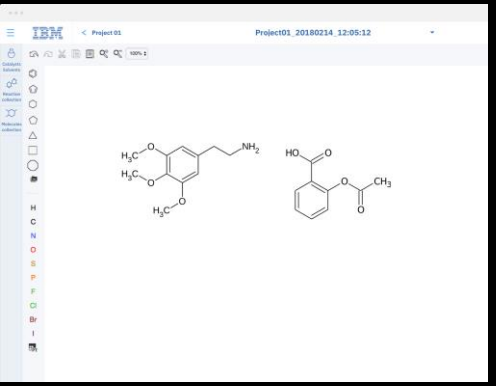

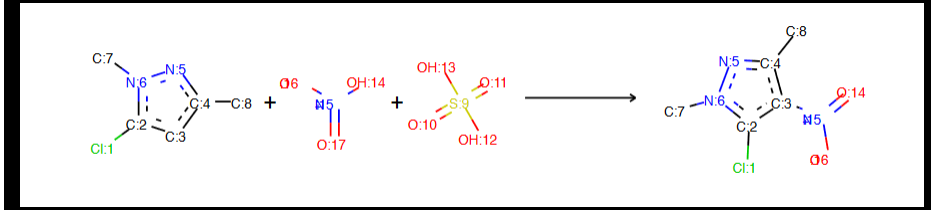

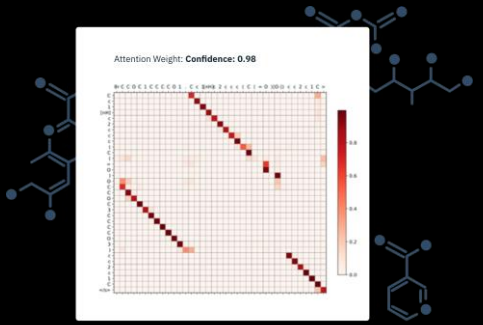

https://rxn.res.ibm.com/

Combining the new Set of Tools for Data-Intensive Science

Data-Intensive Science - High Level Workflow

In 2015 IBM visited eight European scientific & research customers:

JSC – Juelich, KIT – Karlsruhe, MPCDF – Garching, CERN – Geneve, ASTRON – Dwingeloo, DESY – Hamburg, SURFsara – Amsterdam, INAF – Rome

Main Challenges:

Real-time data acquisition

Effective metadata management

Federated IT

Find (keep) the IT Experts

Budget for Computing/Storage

Acquisition

- Data Generation:
  - Measurement
  - Simulation
- Data Curation/Filtering

Distribution

- Moving Data from generation to Analysis Infrastructure
- Inline Filtering

Analysis

- Using HPC/AI clusters and services
- Optimizing workflows and applications for efficiency (Accelerators, I/O etc)

Archival

- Active Archive
  - Actively using huge amount of tape stored data
- Backup

Data Intensive Science: Infrastructure

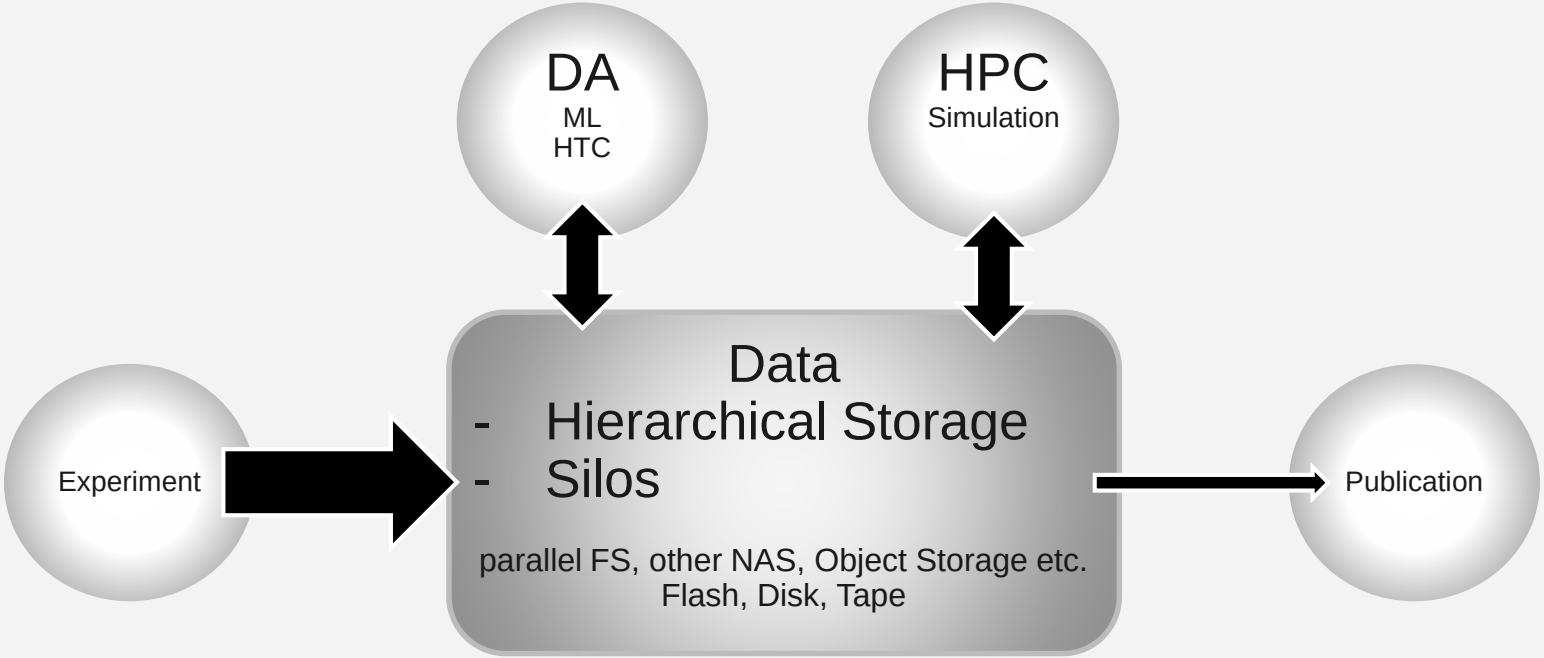

Data Intensive Science: Typical Science Workflow

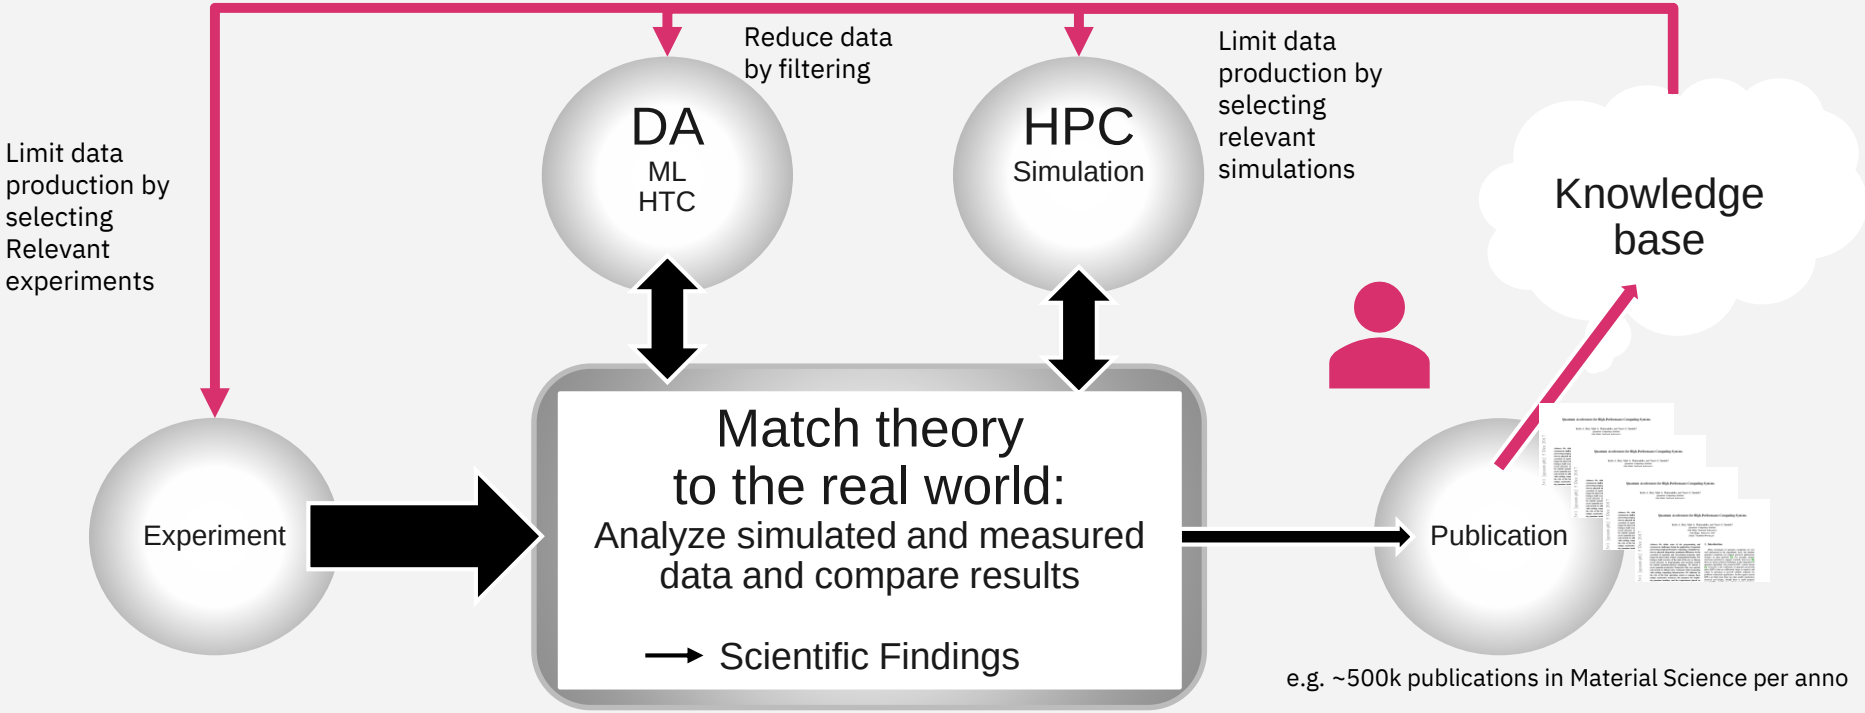

e.g. ~500k publications in Material Science per anno

Data Intensive Science, Future: Augmented Science Workflows

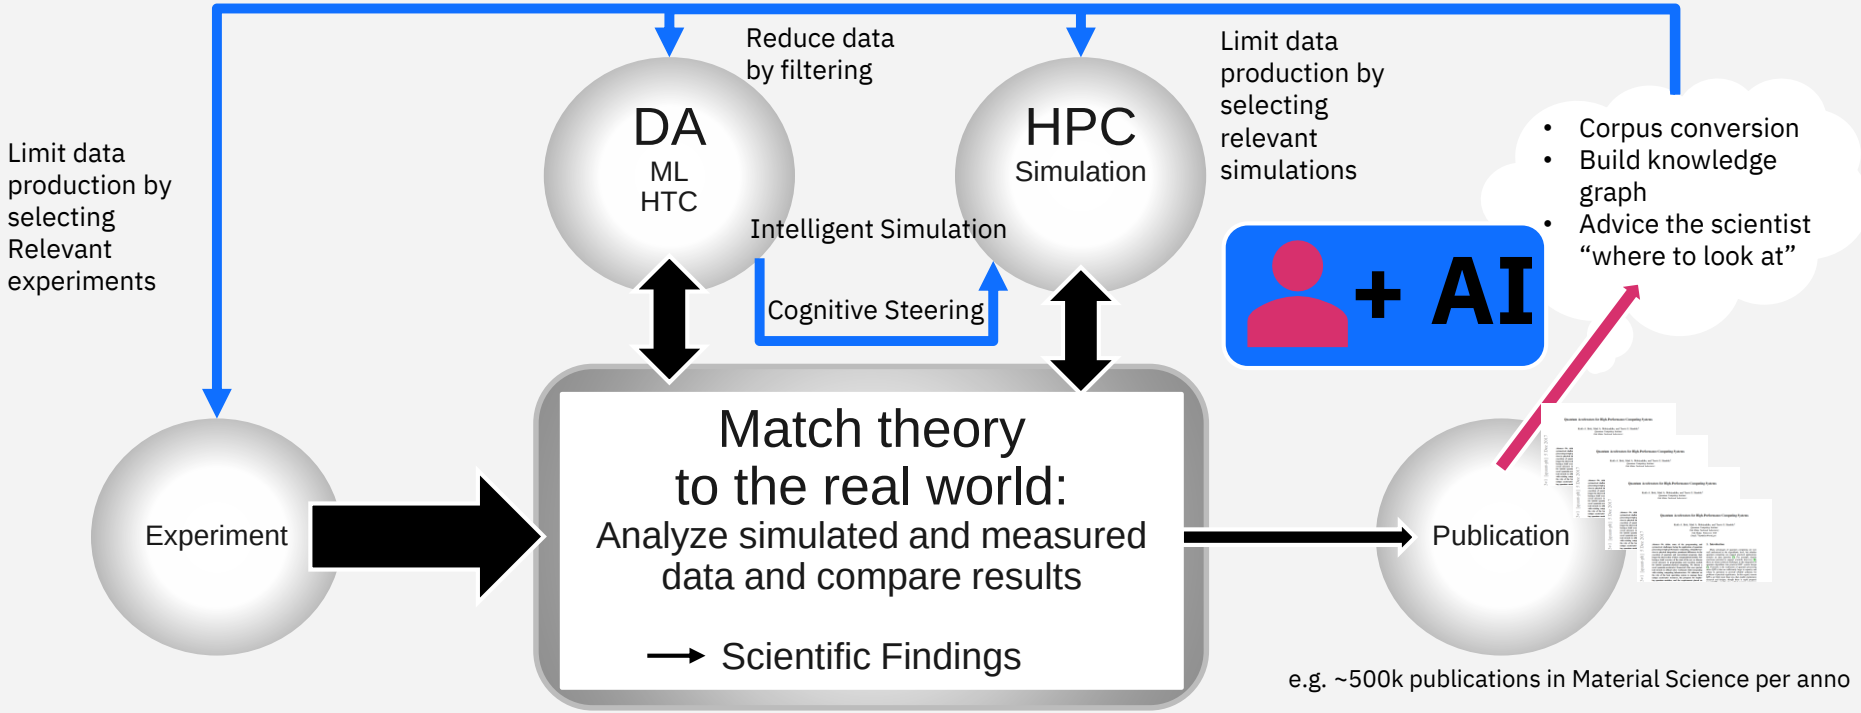

Data Intensive Science: Why IBM?

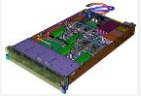

IBM AC922 Power9 + GPU Worker nodes
(NVLINK also between GPU+CPU)

PowerAI

DA
ML
HTC

IBM Spectrum LSF

OpenPOWER

HPC
Simulation

Corpus Conversion Service: A machine learning platform to ingest documents at scale
Peter W J Staar et al. IBM Research

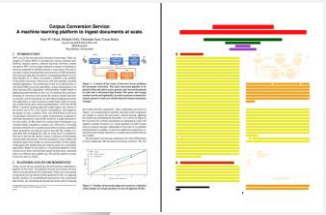

Experiment

aspera
an IBM company

Data

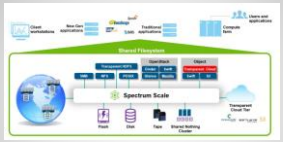

IBM
Spectrum
Scale

HPSS

IBM Elastic Storage Server (ESS) Family,
IBM Flash Storage

IBM TS4500 Tape Library

Are there still intractable problems?

Your next steps to getting Quantum Ready

All information form the Quantum Computing slides of the session can be found at:

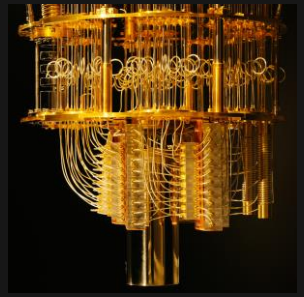

Discover more about IBM's quantum computing initiative

http://www.research.ibm.com/ibm-q/

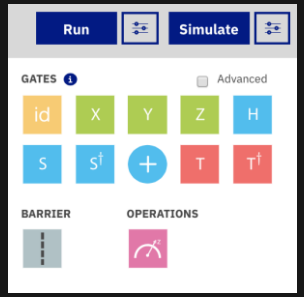

Explore the IBM Quantum Experience and start using real machines today

https://quantum-computing.ibm.com/login

Learn about and start using the Qiskit software development kit

http://www.qiskit.org/

Collaborate, research, and start applying quantum computing through the IBM Q Network

http://ibm.biz/IBMQNetwork

Contact oliver.oberst@de.ibm.com to discuss next steps

Summary

- Main challenge of todays science is the huge amount of generated data and information in structured and unstructured form
- Continue to create tightly coupled Heterogenous systems (Combination of Von Neuman Model and NON Von Neuman Models) -> **New (even more) programming models**
- Using a combination of HPC, HPDA and AI can help to drastically improve the time to solution (or better time to scientific insight)
- Solving Data management by applying data centric strategies is essential
- Quantum Computing (as an accelerator) has the potential to drive scientific insight within the coming decade

Thank you

Dr. Oliver Oberst
HPC/AI Architect
IBM Q Ambassador
—
oliver.oberst@de.ibm.com
ibm.com

 #IBMQ