WILLIAM E. BOEING
DEPARTMENT OF AERONAUTICS & ASTRONAUTICS
UNIVERSITY *of* WASHINGTON

COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Applications and opportunities for fast ML in fusion science and engineering

C. Hansen[1,2], Y. Wei[2], J. Levesque[2], K. Morgan[1],
A. Kaptanoglu[1,3], S. Brunton[1] and others
1- University of Washington
2- Columbia University
3- University of Maryland

Fast Machine Learning for Science Workshop
Southern Methodist University
October 4, 2022

Outline

- Motivation and context
  - Public and private investment in fusion energy has dramatically increased
  - Complexity of fusion plasmas → large gap to necessary latencies
- Offline applications and opportunities in fusion energy
  - Design and engineering: Lightweight models for optimization
- Online applications and opportunities in fusion energy
  - Sensor fusion: Sparse, external, nonlinearly coupled
  - Active control: Tight latency (< ms), advanced control
- Summary

You may have heard about fusion in the news over the last few years

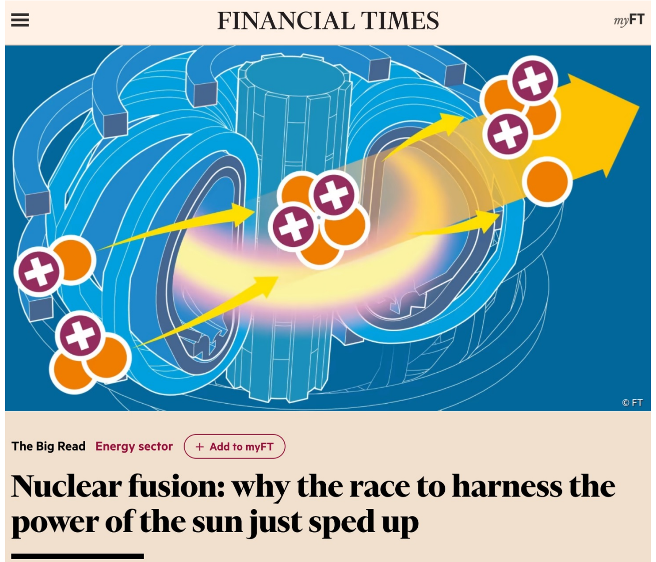
FINANCIAL TIMES

myFT

© FT

The Big Read Energy sector + Add to myFT

Nuclear fusion: why the race to harness the power of the sun just sped up

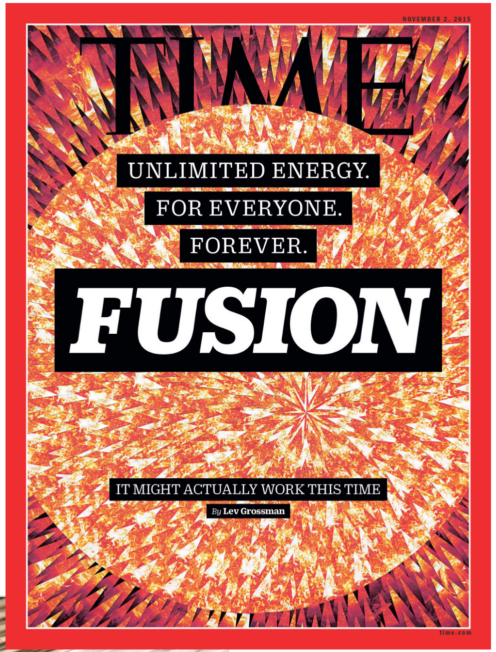
NOVEMBER 2, 2015

TIME

UNLIMITED ENERGY.
FOR EVERYONE.
FOREVER.

FUSION

IT MIGHT ACTUALLY WORK THIS TIME

By Lev Grossman

time.com

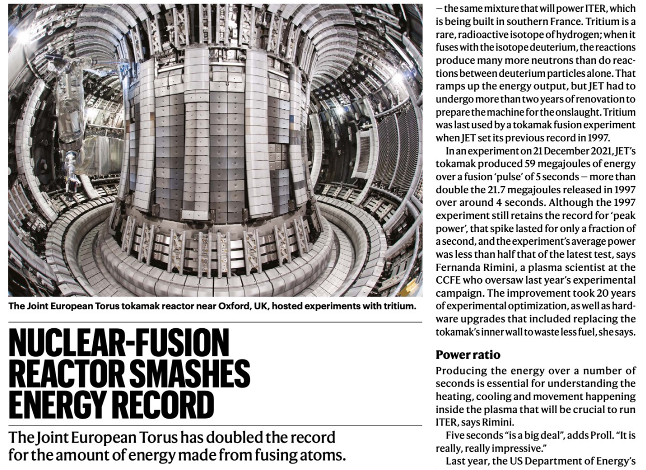
The Joint European Torus tokamak reactor near Oxford, UK, hosted experiments with tritium.

NUCLEAR-FUSION REACTOR SMASHES ENERGY RECORD

The Joint European Torus has doubled the record for the amount of energy made from fusing atoms.

– the same mixture that will power ITER, which is being built in southern France. Tritium is a rare, radioactive isotope of hydrogen; when it fuses with the isotope deuterium, the reactions produce many more neutrons than do reactions between deuterium particles alone. That ramps up the energy output, but JET had to undergo more than two years of renovation to prepare the machine for the onslaught. Tritium was last used by a tokamak fusion experiment when JET set its previous record in 1997.

In an experiment on 21 December 2021, JET's tokamak produced 59 megajoules of energy over a fusion 'pulse' of 5 seconds – more than double the 21.7 megajoules released in 1997 over around 4 seconds. Although the 1997 experiment still retains the record for 'peak power', that spike lasted for only a fraction of a second, and the experiment's average power was less than half that of the latest test, says Fernanda Rimini, a plasma scientist at the CCFE who oversaw last year's experimental campaign. The improvement took 20 years of experimental optimization, as well as hardware upgrades that included replacing the tokamak's inner wall to waste less fuel, she says.

Power ratio

Producing the energy over a number of seconds is essential for understanding the heating, cooling and movement happening inside the plasma that will be crucial to run ITER, says Rimini.

Five seconds "is a big deal", adds Proll. "It is really, really impressive."

Last year, the US Department of Energy's


The international journal of science / 27 January 2022

nature

THE HEAT IS ON

Burning plasma marks key step towards power from nuclear fusion

ffort
omplex research
ss-roots
rative networks

Relocation plans
Genomic data reveal Bronze Age migration to southern Britain

Vol. 601, No. 7894
nature.com

WSJ


◆ WSJ NEWS EXCLUSIVE

Nuclear-Fusion Startup Lands $1.8 Billion as Investors Chase Star Power

No one has been able to generate net energy by combining atoms, yet Commonwealth Fusion Systems has attracted Bill Gates and George Soros

WIRED BACKCHANNEL BUSINESS CULTURE GEAR IDEAS SCIENCE SECURITY

DeepMind Has Trained an AI to Control Nuclear Fusion

The Google-backed firm taught a reinforcement learning algorithm to control the fiery plasma inside a tokamak nuclear fusion reactor.

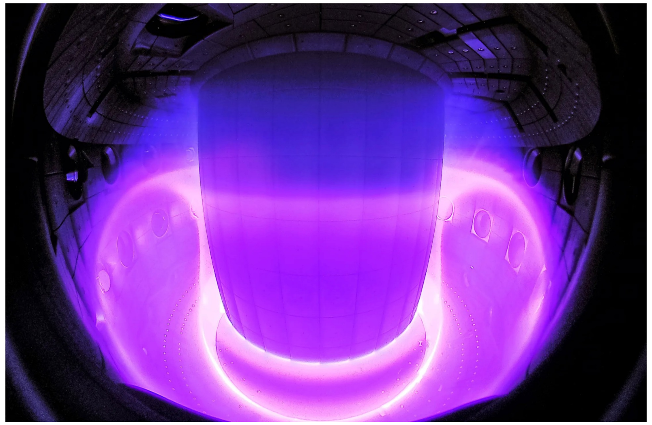

Ok, so what changed?

- Computing power
  - Simulations are crucial to understanding fusion plasmas
  - Many models have only recently become tractable


IT MIGHT ACTUALLY WORK THIS TIME

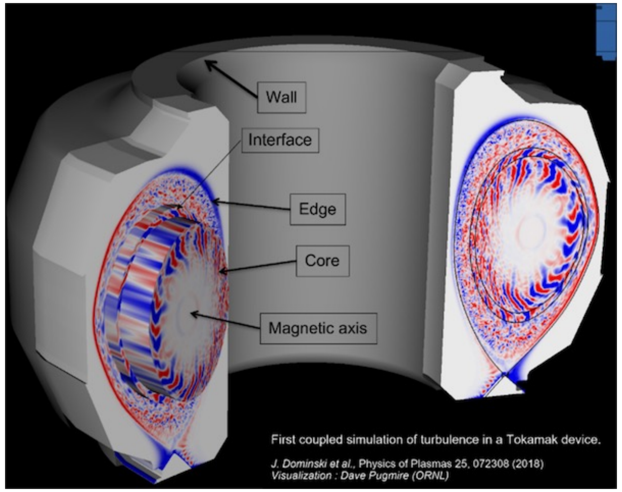

First coupled simulation of turbulence in a Tokamak device.

J. Dominski et al., Physics of Plasmas 25, 072308 (2018)
Visualization : Dave Pugmire (ORNL)

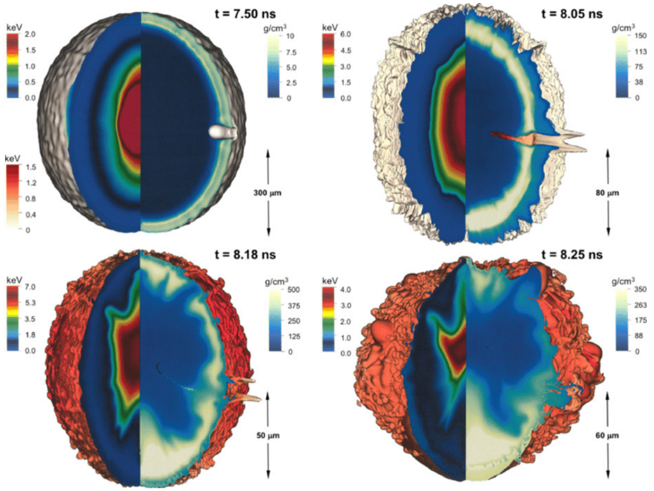

Ok, so what changed?

- Computing power
  - Simulations are crucial to understanding fusion plasmas
  - Many models have only recently become tractable
- New technology
  - HTS magnets, diode lasers, etc.
  - Advanced manufacturing, power electronics, etc.

IT MIGHT ACTUALLY WORK THIS TIME

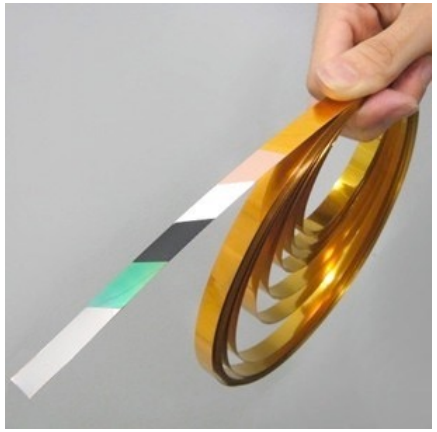

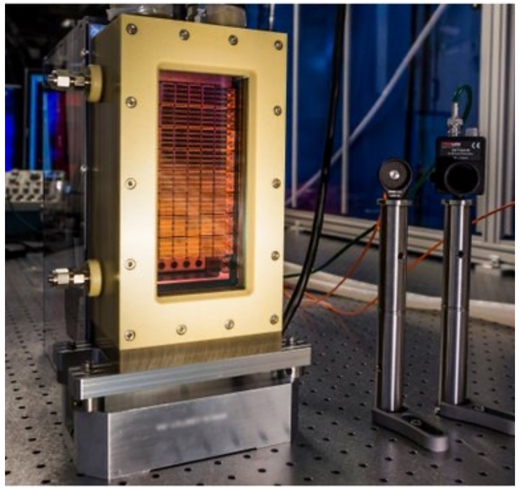

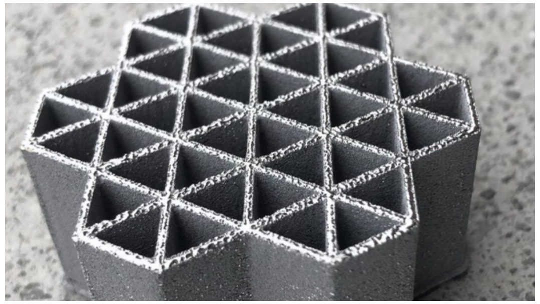

Ok, so what changed?

- Computing power
  - Simulations are crucial to understanding fusion plasmas
  - Many models have only recently become tractable
- New technology
  - HTS magnets, diode lasers, etc.
  - Advanced manufacturing, power electronics, etc.
- Global scientific consensus on readiness
  - Private sector interest (> $4B of investment)
  - Renewed public sector interest

IT MIGHT ACTUALLY WORK THIS TIME

By primary HQ

Source: Fusion Industry Association (2022)

DEPARTMENT OF ENERGY · UNITED STATES OF AMERICA

MILESTONE-BASED FUSION DEVELOPMENT PROGRAM

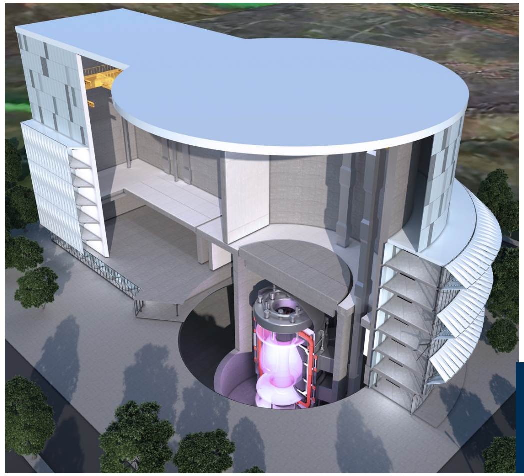

UK Atomic Energy Authority

Fusion is now being actively pursued to accelerate access to a zero-carbon future

- Provide a dispatchable source of carbon-free energy
  - Reduce storage/overbuild requirements to cover seasonality
  - High-quality process heat for non-electrical applications
- Enables a paradigm shift to continued growth in energy usage

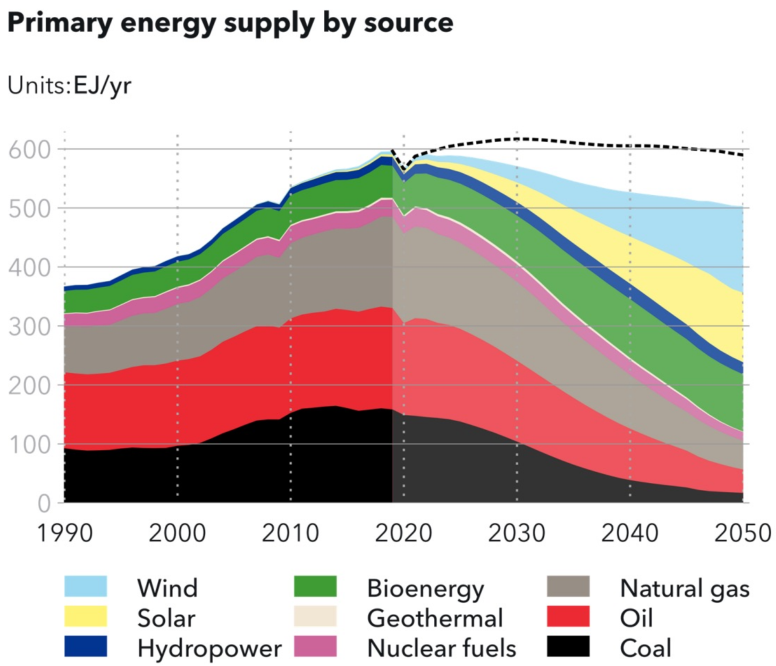

©DNV 2021 Dotted line: ETO forecast. Historical data source: IEA WEB (2020)

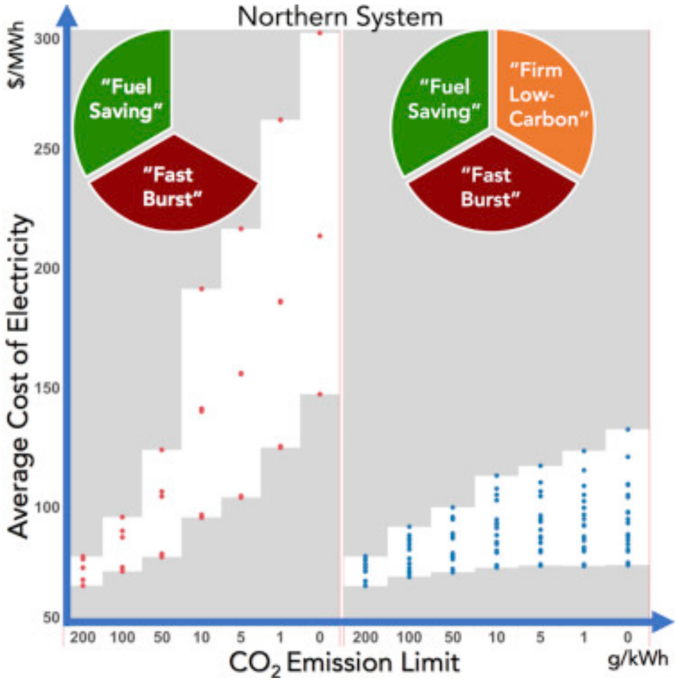

1- N. Sepulveda et al., *Joule* **2**, 2403 (2018)

Many private and government projects are now aiming for fusion demonstrations in the near term

- Both private and public projects are actively building fusion demonstrations
  - Global enterprise with billions of $/year in total investment
- A variety of methods and designs are being pursued
  - Public → lower risk; Private → higher risk
  - Net-energy demonstrations as early 2023

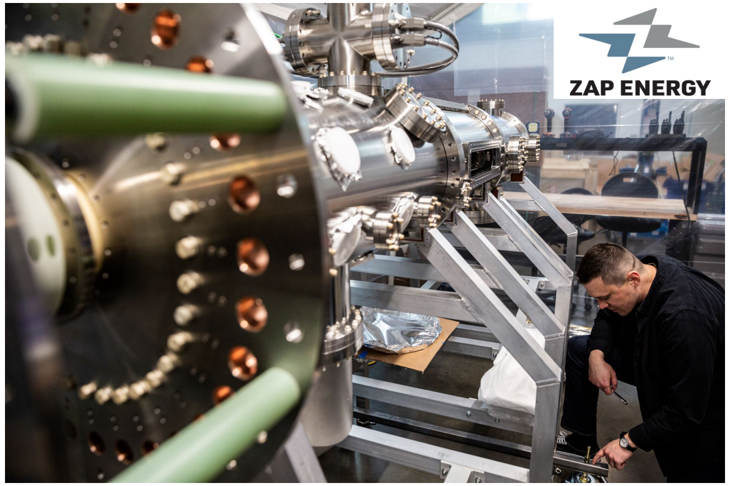

2- "Final report of the committee on a strategic plan of U.S. burning plasma research," *The National Academies Press* (2019)
3- "European Research Roadmap to the Realisation of Fusion Energy," *EUROfusion (2018)*

This talk will focus on considerations for the tokamak magnetic confinement concept

- “Magnetic confinement” is the leading technique
  - Accounts for bulk of private and public investment
- The tokamak is the leading configuration
  - Has demonstrated large amounts of fusion power
- Considerations are common to all fusion concepts with different weights

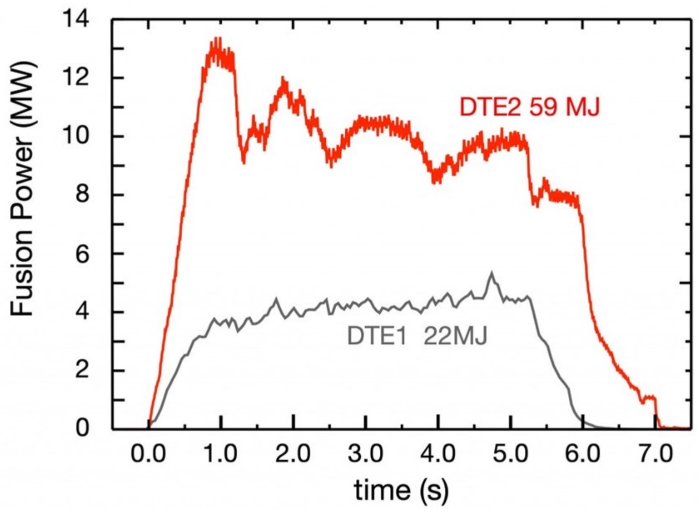

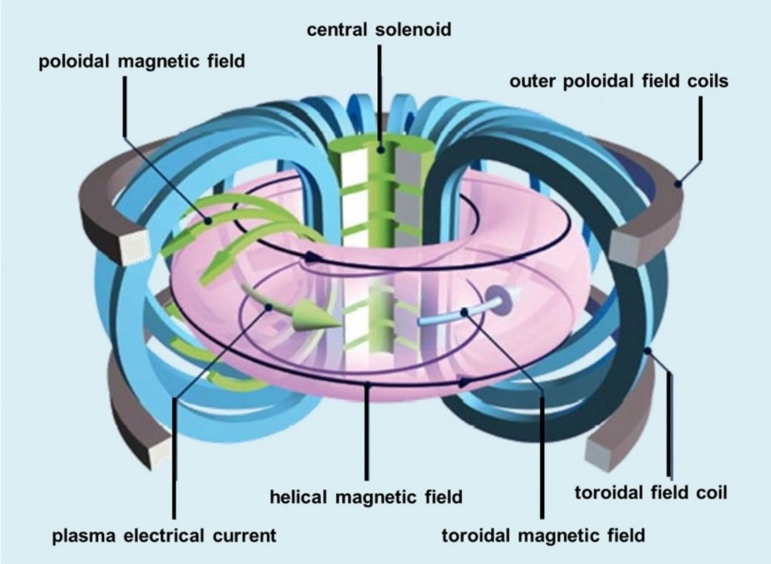

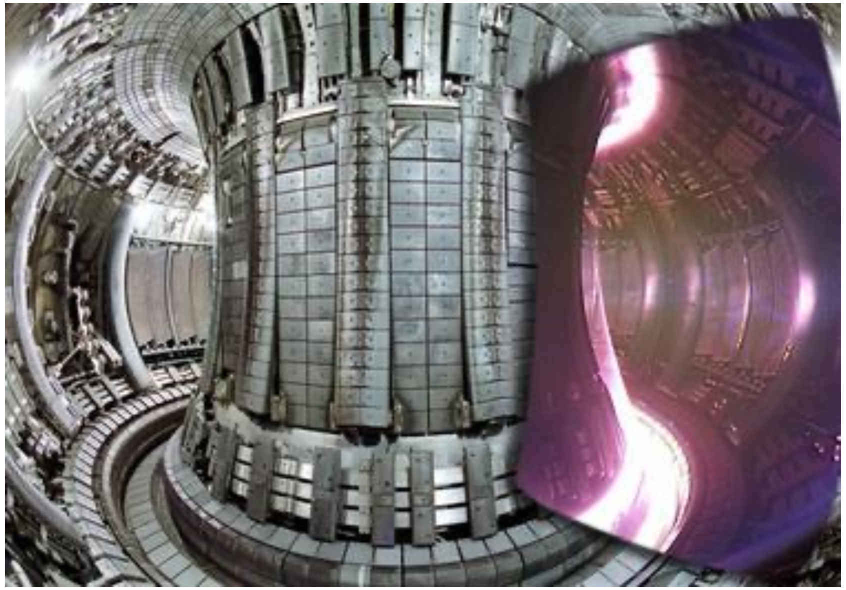

WEDNESDAY, NOVEMBER 9, 1994

Experimental Fusion Reactor At Princeton Sets a Record

By The Associated Press

An experimental fusion reactor at Princeton University has set another world record by generating 10.7 million watts of power, considered a crucial step toward developing new commercial fusion reactors.

The one-second burst of energy from the Tokamak Fusion Test Reactor, enough to power 2,000 to 3,000 homes momentarily, was "a major milestone," Dr. Ronald C. Davidson, director of the university's Plasma Physics Laboratory, said in a telephone interview on Saturday. In May, the reactor produced 9.2 million watts of energy.

"The T.F.T.R. team has demonstrated successful use of practical fusion fuels, exceeded fusion power goals and reached world-record temperatures," Dr. Davidson said.

Officials of the laboratory presented the results yesterday in Minneapolis at the annual meeting of the Division of Plasma Physics of the American Physical Society.

Fusion, the process that powers the Sun, occurs on Earth when special forms of lightweight atoms like hydrogen gas are slammed together at very high temperatures and fused, releasing energy. Fission, which powers existing commercial nuclear reactors, involves breaking apart very heavy atoms like uranium but produces far more dangerous radioactive byproducts than fusion.

The Tokamak reactor at Princeton, designed in 1976, was intended to produce 10 million watts of power. The reactor set its first record in December when it produced the equivalent of 6.4 million watts of power.

Dr. Davidson said the results will help scientists and engineers design future reactors.

"The better information that we can provide as input to the design of these facilities, the better their design will be, and the sooner practical fusion will be developed and the more attractive the power source will be," Dr. Davidson said. "This is all vital information."

The reactor, operated by the Plasma Physics Laboratory, is scheduled to be shut down and dismantled this fall. Some of its auxiliary systems are expected to be reused for the next planned experimental reactor at the laboratory, which is considered the country's leading center for fusion research.

Its replacement, the $750 million Tokamak Physics Experiment, is under design and is scheduled to be in operation in 2001.

A third experimental reactor, the $10 billion International Thermonuclear Experimental Reactor, is a joint project of Russia, Japan, the United States and a European coalition. It is scheduled to be ready for operation by 2005.

Company News:
Tuesday through Saturday,
Business Day

Fusion plasmas are a complex system that challenges understanding, prediction and control

- Multiscale nonlinear dynamics
  - Coupled multispecies system + E-M
  - Turbulence driven by steep gradients
- First-principles models are rarely tractable for online applications and design optimization
  - Even simplified/linearized models too slow
- Diagnostic access is extremely limited
  - Almost exclusively external diagnostics (eg. optical)
- Sensor fusion required for most signals
  - Deconvolve and/or localize measurements
- Strict latency requirements (ms or faster)

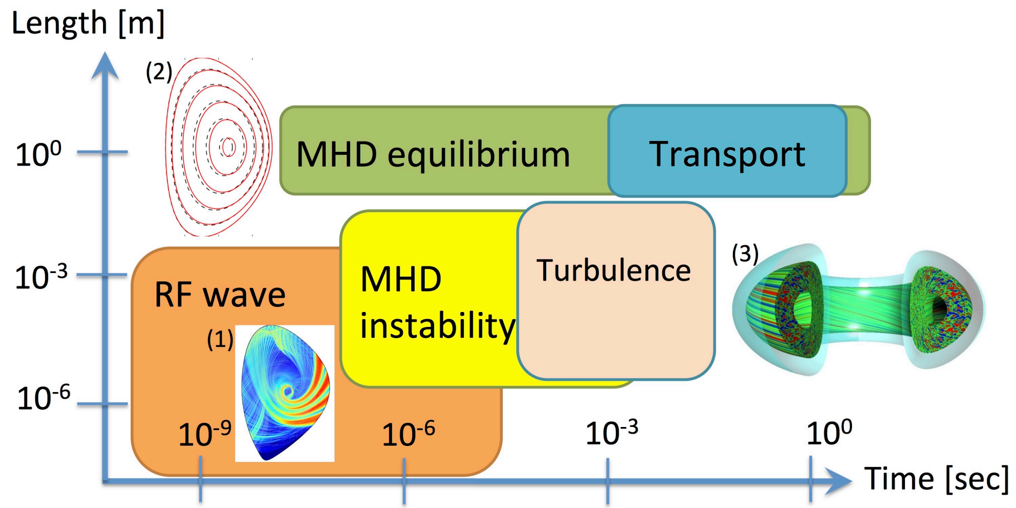

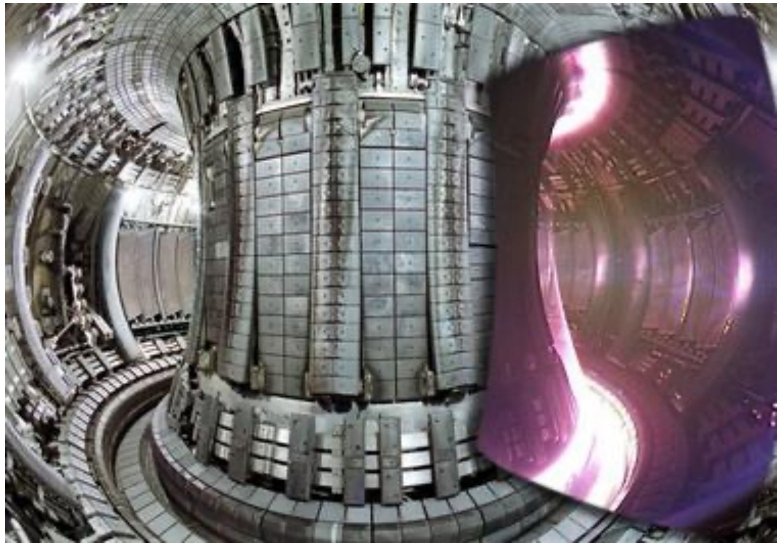

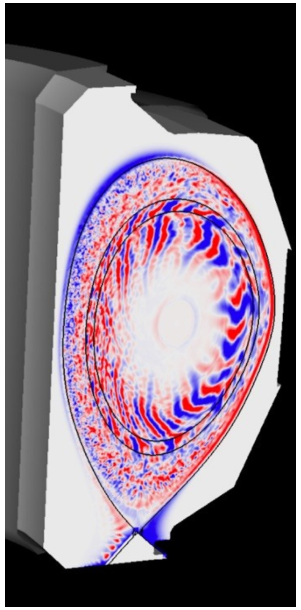

Machine learning holds promise to help accelerate the realization of fusion energy

- Commercial systems have different needs for models used in design and control
  - No more research “compromises”
  - Predict first not model after
- Models must be fast
- Models must be accurate
- Models must be robust
- Fertile ground for this community!

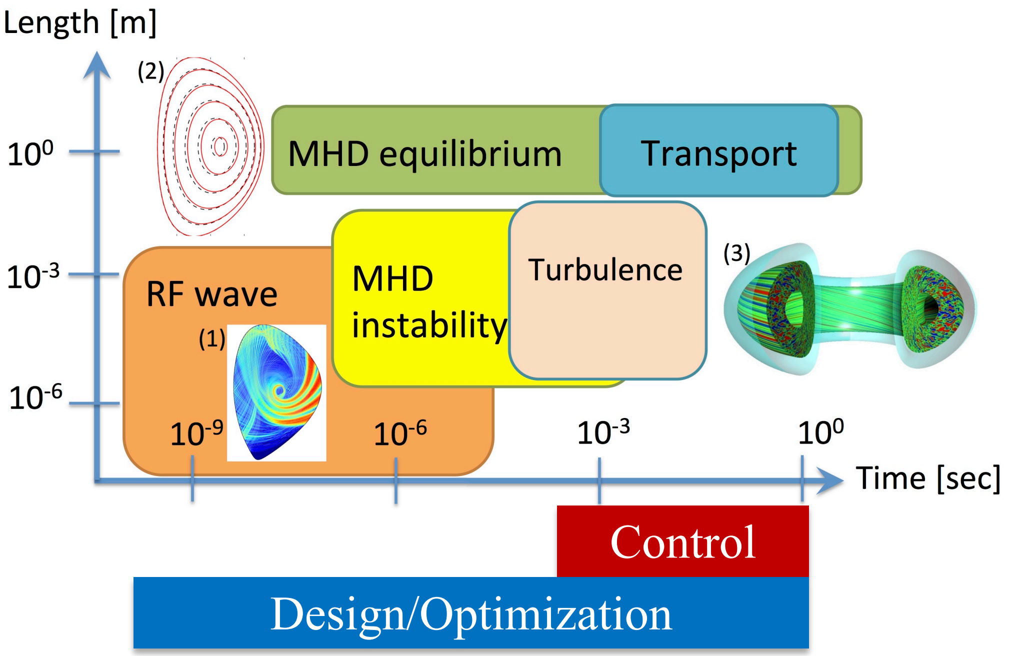

Outline

- Motivation and context
  - Public and private investment in fusion energy has dramatically increased
  - Complexity of fusion plasmas → large gap to necessary latencies
- Offline applications and opportunities in fusion energy
  - Design and engineering: Lightweight models for optimization
- Online applications and opportunities in fusion energy
  - Sensor fusion: Sparse, external, nonlinearly coupled
  - Active control: Tight latency (< ms), advanced control
- Summary

Design point optimization is critical to performance and economics

- Fusion production depends on *T* and *n* profiles
  - Must be in equilibrium (fast and slow)
- Energy deposition and transport is complex
  - Sources: fusion products, external heating, etc.
  - Sinks: turbulent transport, radiation, etc.
  - Many complex bifurcations/phase transitions
- Core and surrounding system are coupled
  - Magnets set plasma shape
  - RF waves refract through plasma
- Large high-dimensional optimization space
  - Blanket required for energy and fuel cycle
  - Maintenance needs to be considered

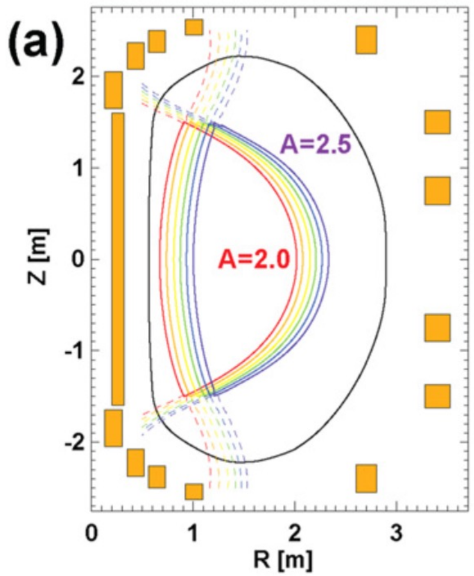

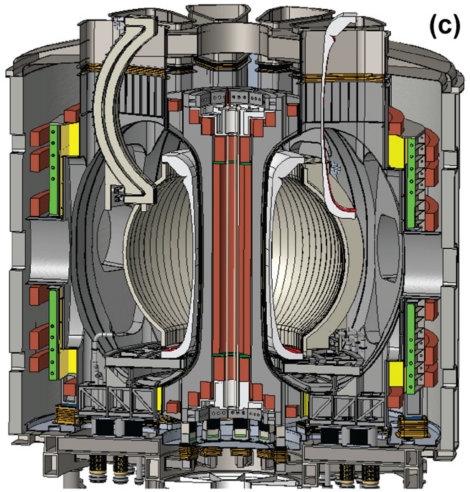

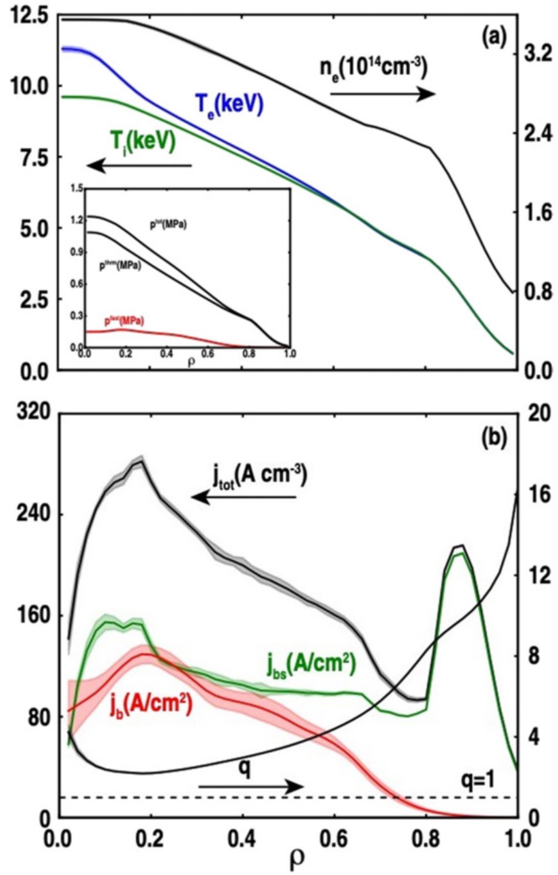

4- J. Menard et al., *Nucl. Fusion* **62**, 036026 (2022)

The stellarator concept focuses heavily on confinement optimization through design

- By introducing 3D fields stellarators can confine fuel without current flowing in the plasma
  - High-dimensionality optimization space
- Transport of heat and mass is very sensitive to magnetic field geometry
  - Neo-classical losses: particle drifts
  - Plasma turbulence: global nonlinear problem
- Machine learning approaches are being investigated to accelerate optimization
  - Transport calculation (or proxy) required

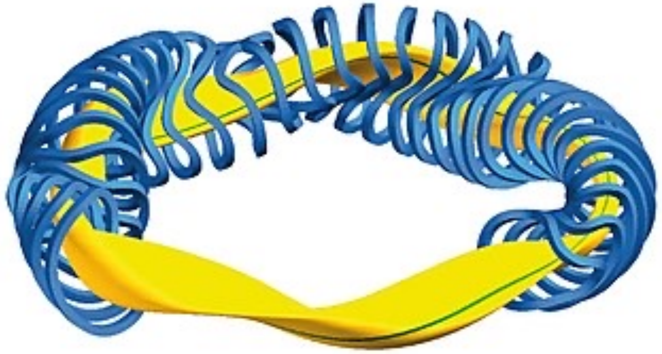

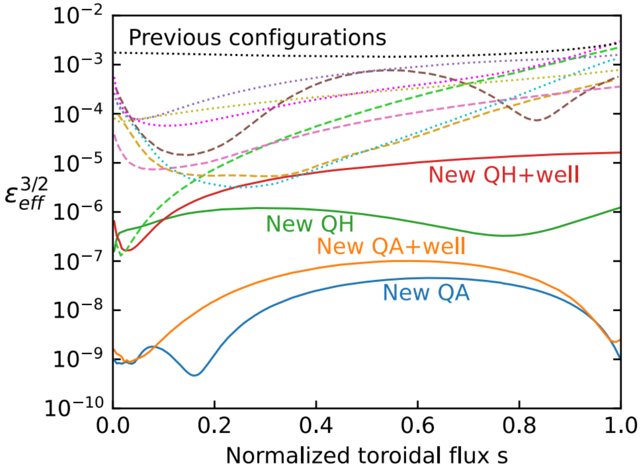

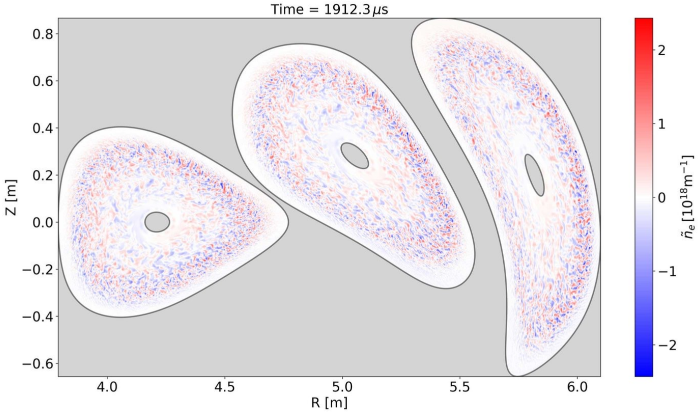

5- E. Paul et al., *PRL* **128**, 035001 (2022)
6- E. Maurer et al., *JCP* **420**, 109694 (2020)

Model-discovery methods can build fast reduced-order models of nonlinear plasmas

- Reduced-order methods have potential for developing fast surrogate models
  - ODE, ML, etc.
- The SINDy method has shown promise on reproducing nonlinear plasma dynamics
  - Build nonlinear ODEs directly from data
  - Naturally supports physics-informed constraints
- Physics-informed constraints can significantly improve model quality
  - Reduce required training data
  - Enforce local or global stability

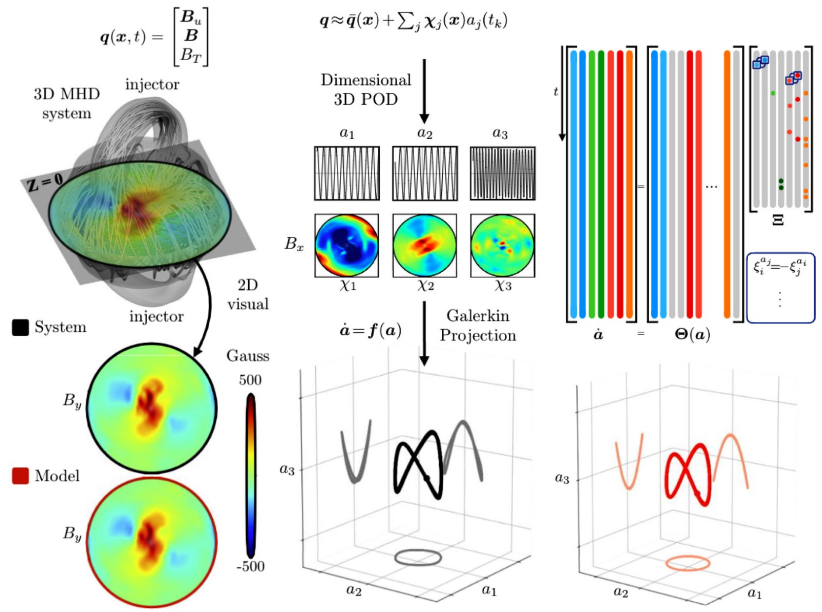

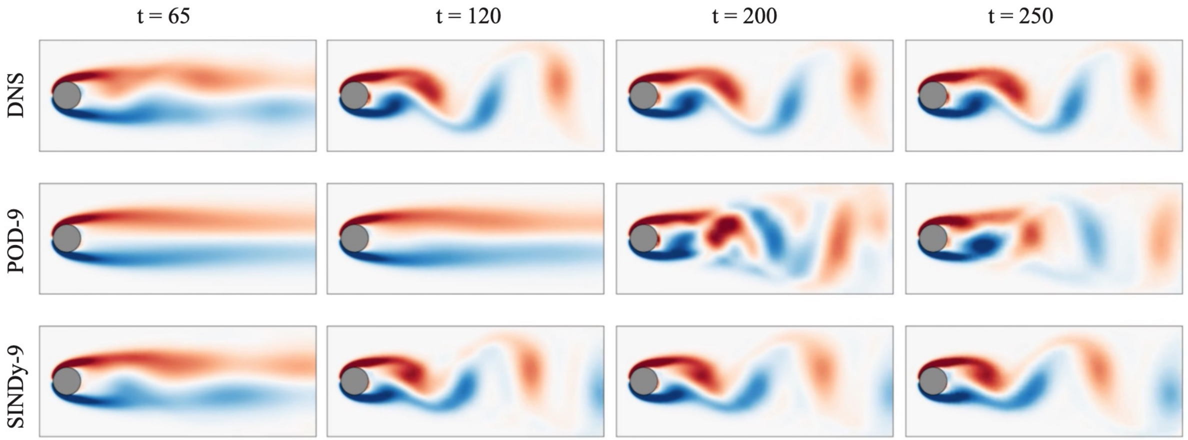

7- A. Kaptanoglu et al., *Phys. Rev. E* **104**, 015206 (2021)

8- A. Kaptanoglu et a., *Phys. Rev. Fluids* **6**, 094401 (2021)

Outline

- Motivation and context
  - Public and private investment in fusion energy has dramatically increased
  - Complexity of fusion plasmas → large gap to necessary latencies
- Offline applications and opportunities in fusion energy
  - Design and engineering: Lightweight models for optimization
- Online applications and opportunities in fusion energy
  - Sensor fusion: Sparse, external, nonlinearly coupled
  - Active control: Tight latency ($<$ ms), advanced control
- Summary

Building a consistent picture of equilibria and dynamics from diagnostics is complicated

- Sparse diagnostic set necessitates integration of multiple signals
  - Plasma state must be reconstructed in time
  - Even fewer signals in future reactors
- Optical and other EM-based diagnostics are complex
  - Nonlocal measurements: equilibrium-dependent sightlines
  - Convolve many fundamental quantities (*T*, *n*)
  - Complex spectrographic landscapes
- Still early days for applying machine learning to this area
  - Restrict ourselves to things we can observe directly (next section)
  - Lump this step in to a larger NN (following section)
- Accurate sensor fusion can make some control easier
  - Ability to work directly with primary variables

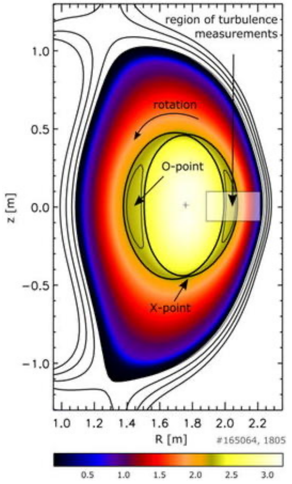

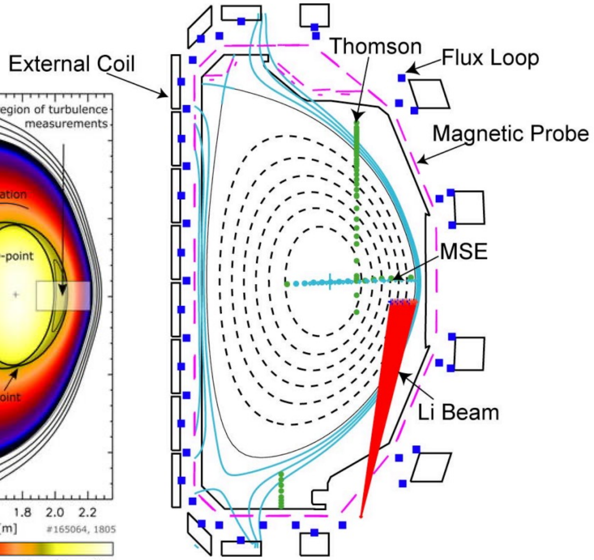

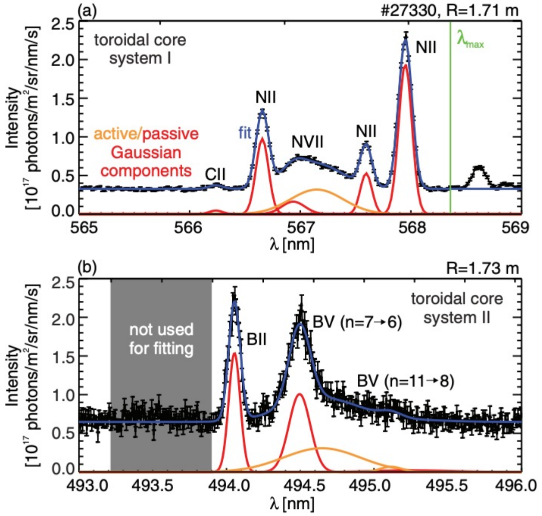

Active control of instabilities in tokamaks requires low-latency pipelines

- Instabilities can degrade performance in fusion plasmas
  - Avoided by reduced performance
  - Suppressed by active control
- Active control requires a low-latency pipeline
  - Mode growth rates are on ms timescales
  - GPUs are frequently used in community
- Plasma response can be complex, motivating optimal control
  - Opportunity for ML system response models

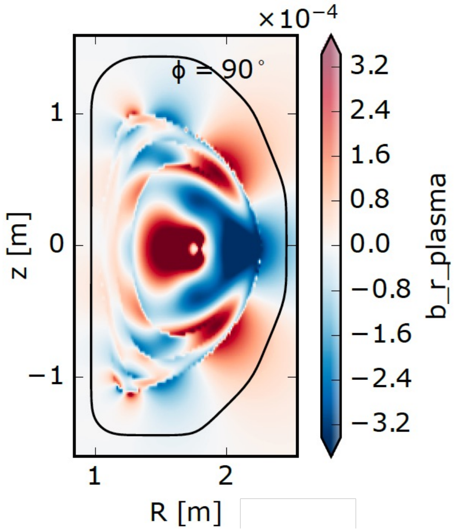

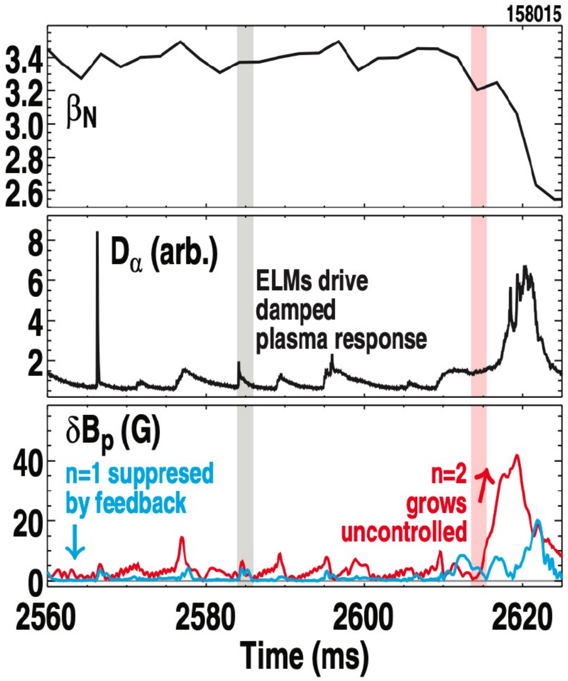

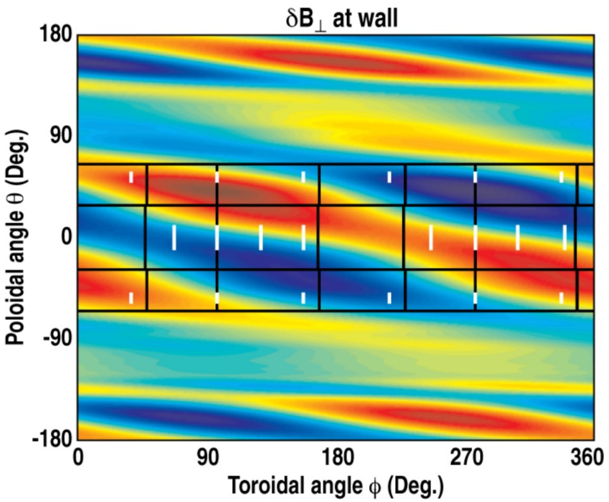

A. Battey, *PhD Thesis: Columbia U* (2022)

We are now working with this community to leverage tools for low-latency ML

- Diagnostic access will become more restrictive in reactor environments
  - Remote observations only, complex emission functions
- ML applies a promising path to fast mapping from signals to desired quantities
- We are working to use optical cameras to observe mode phase/amplitude
  - Convolves mode with edge conditions, reflections, etc.
  - HLS4ML enables implementation on camera FPGA

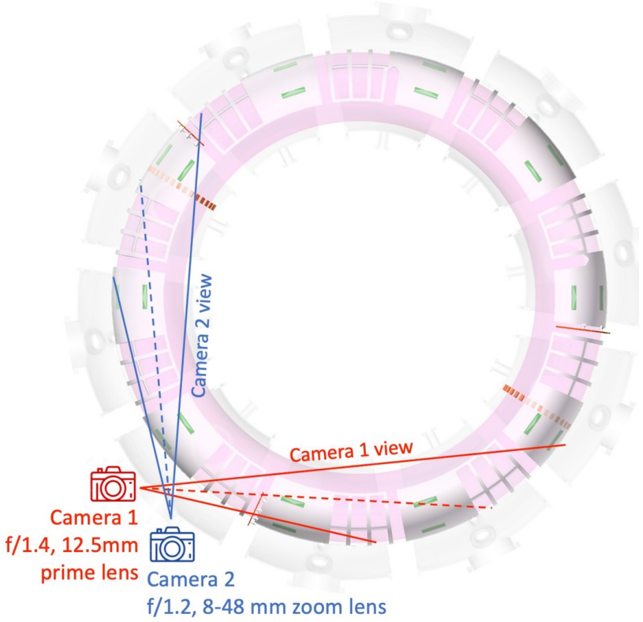

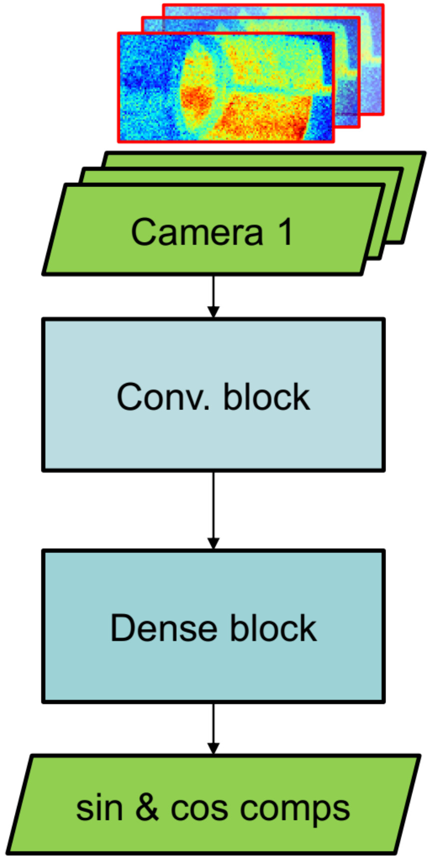

Camera 1

Conv. block

Dense block

sin & cos comps

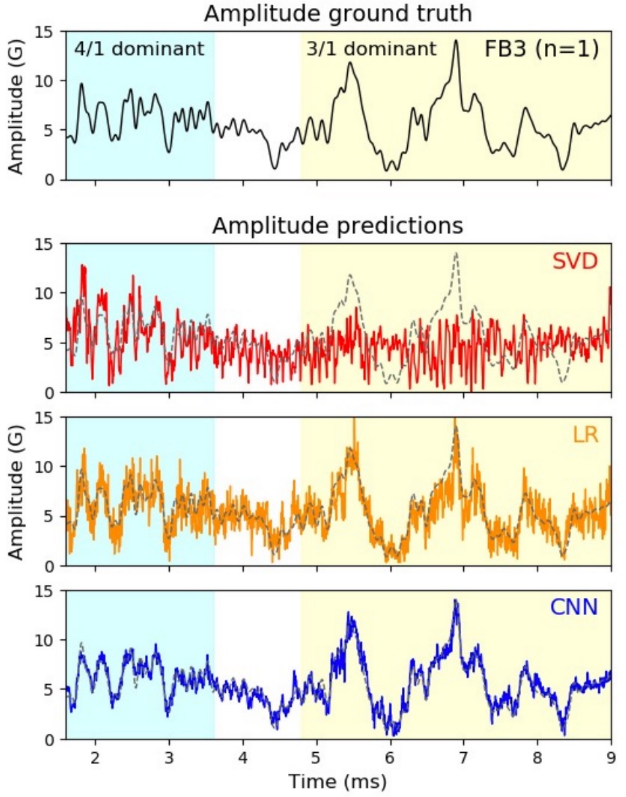

* - Talk by Y. Wei in today’s contributed session

Accurate prediction of plasma disruptions is a significant open question for tokamaks

- A disruption is a rapid termination of plasma current
  - Caused by a range of phenomena: Instabilities, loss of control, etc.
  - Can damage device → more frequent maintenance
- Avoidance and/or mitigation requires long lead time
  - Causality is complex: nonlinearity, multi-event chains
- Machine learning approaches have shown promise
  - Classification: random forest, NN, etc.
  - Online usage requires robust, low-latency implementations
- Very low false negative tolerance (→ 0)
  - Very asymmetric datasets
  - Want as close to day 1 as possible: transfer or virtual learning

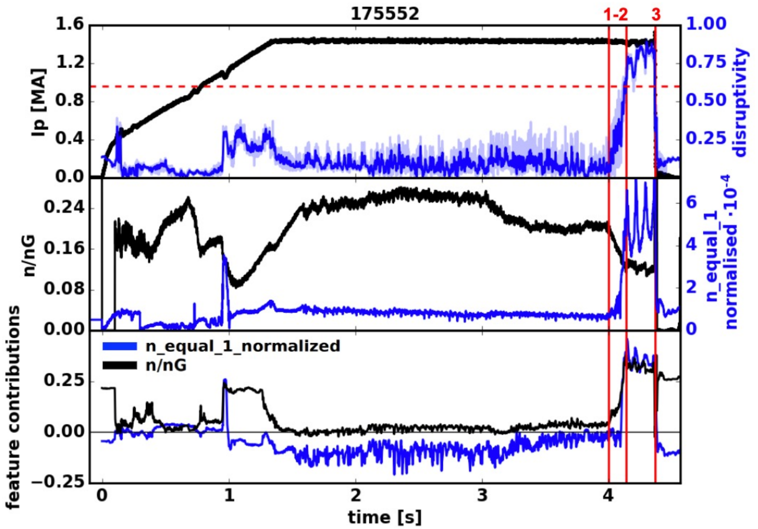

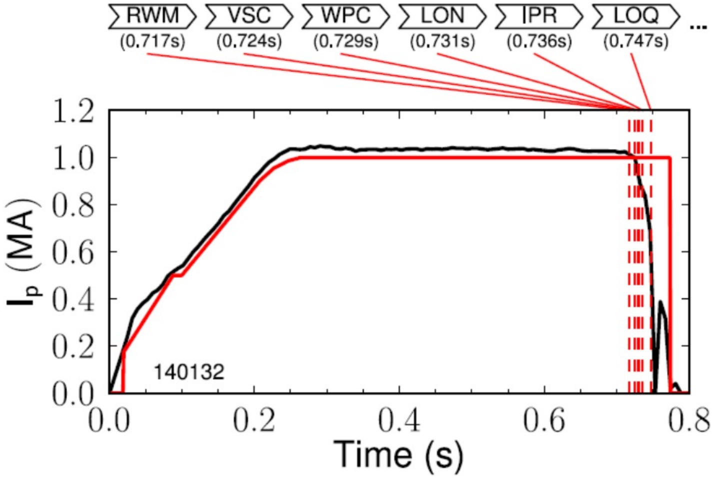

10- C. Rea et al., *Nucl. Fusion* **59**, 096016 (2019)
11- J. Berkery et al., *Phys. Plasmas* **24**, 056103 (2017)

A VAE-based method has shown promise for both detection and avoidance of disruptions

- Diagnostics form a high-dimensional observation space
  - Important dynamics are expected to behave with shared low dimension structure
- A VAE was trained with a 2D latent space representation from 7D input data
- Successfully demonstrated the prediction and avoidance of disruptions
  - Stability boundaries were identified in latent space
  - Local gradients used to identify actuator outputs
- Further study is underway
  - More signals, larger devices

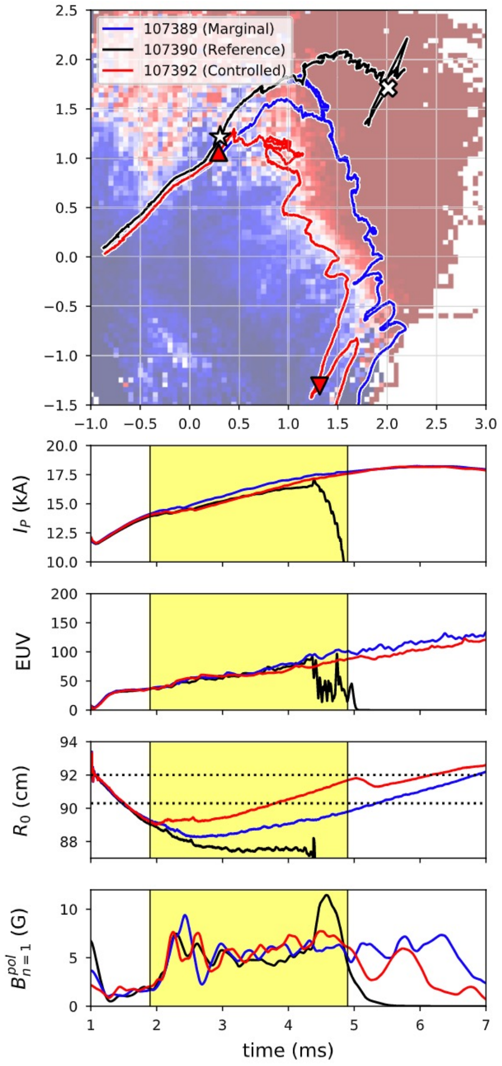

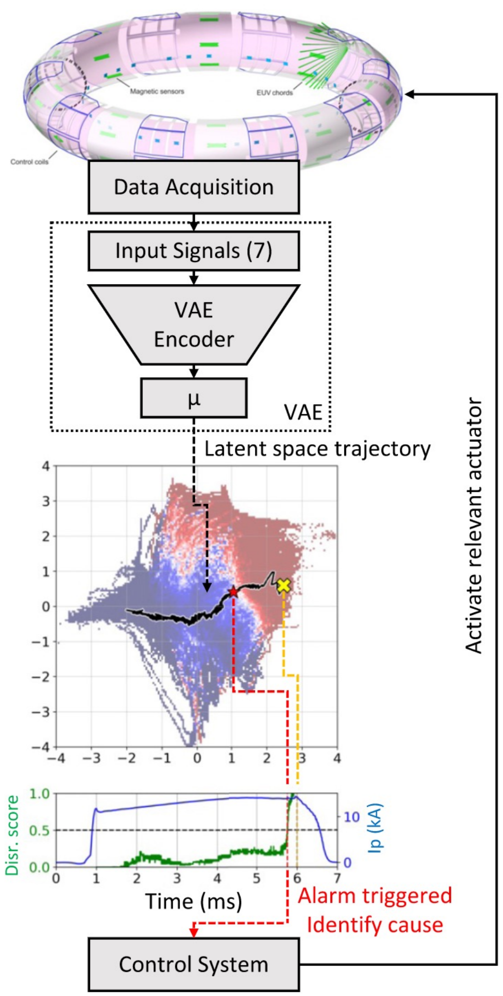

12- Y. Wei et al., *Nucl. Fusion* **61**, 126063 (2021)

Reinforcement learning was recently demonstrated to build an end-to-end controller

- EPFL in collaboration with Google demonstrated a NN-based controller built using reinforcement learning
  - Successfully controlled real plasmas in TCV
  - Performed better than existing hand-tuned controller
- Requires fast, accurate training environment
  - Approaches optimized for data paucity
  - Possible application for reduced-order models
- Online implementation using CPU
  - < 10 ms latency requirements
  - Other applications require lower latency
  - Additional diagnostics require more throughput

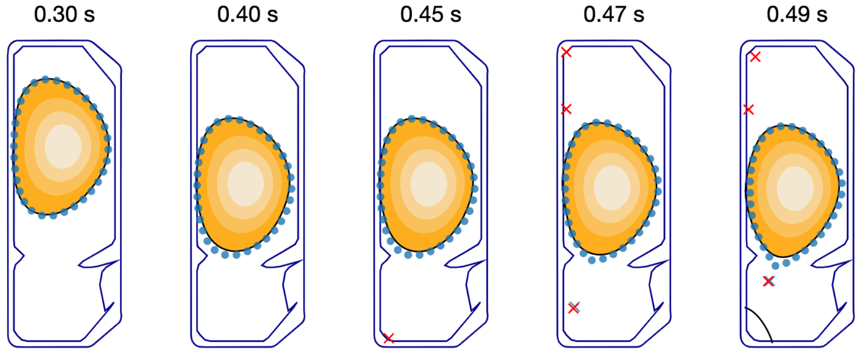

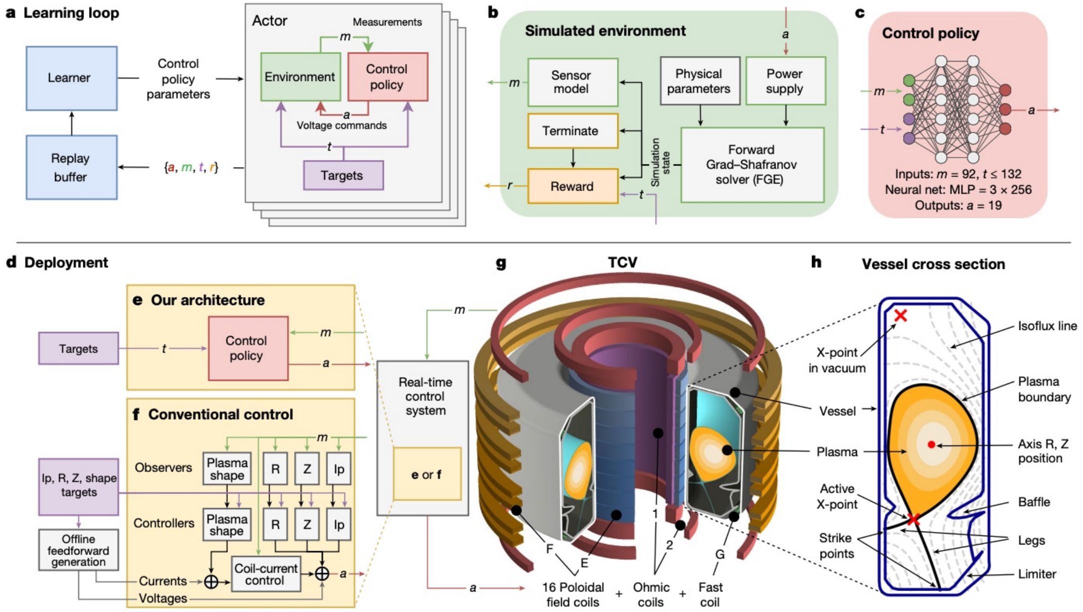

13- J. Degrave et al., *Nature* **602**, 414-419 (2022)

Summary

- The world is now working to realize fusion as an energy source in the near term
  - Public and private sectors are moving together
- Many applications within the fusion space require fast nonlinear models
  - Design optimization: Fast surrogate models and interpolation over high-dimensional datasets
  - Sensor fusion: Integrate multiple signals with nonlinear dependence into unified state
  - Active control: Surrogate models for system response and/or end-to-end controllers
- Machine learning methods can (will?) play an integral part in fusion's realization
  - Models need to be fast, accurate and robust
  - Hardware pipelines will likely be required to satisfy latency requirements
- Interested? Talk to your local fusion scientist or I can point you in the right direction

Funding and collaboration acknowledgements

Funding for this work provided by:

U.S. DEPARTMENT OF ENERGY | Office of Science

NSF

arpa·e
CHANGING WHAT'S POSSIBLE

CTFusion
As Brilliant as the Sun

Work performed in collaboration with:

W

MIT

PPPL

DIII-D
NATIONAL FUSION FACILITY
SAN DIEGO

CTFusion
As Brilliant as the Sun

Thank you for your attention

Questions?