Position: Opportunities Exist for Machine Learning in Magnetic Fusion Energy

Lucas Spangher [*1] **Allen Wang** [*1] **Andrew Maris** [*1] **Myles Stabelberg** [1] **Viraj Mehta** [2] **Alex Sparstein** [1] **Stephen Lane-Walsh** [1] **Akshata Kishore Moharir** [3] **Alessandro Pau** [4] **Cristina Rea** [1]

Abstract

Magnetic confinement fusion may one day provide reliable, carbon-free energy, but the field currently faces technical hurdles. In this position paper, we highlight six key research challenges in the field of fusion energy that we believe should be research priorities for the Machine Learning (ML) community because they are especially ripe for ML applications: (1) disruption prediction, (2) simulation and dynamics modeling (3) resolving partially observed data, (4) improving controls, (5) guiding experiments with optimal design, and (6) enhancing materials discovery. For each problem, we give background, review past ML work, suggest features of future models, and list challenges and idiosyncrasies facing ML development. We also discuss ongoing efforts to update the fusion data ecosystem and identify opportunities further down the line that will be enabled as fusion and its data infrastructure advance. It is our position that fusion energy offers especially exciting opportunities for ML practitioners to impact decarbonization and the future of energy.

1. Introduction: Magnetic Confinement Fusion

Eras of human history are defined by their energy sources, from the charcoal that forged the iron age to the fossil fuels that shape our modern world. Today, our planet faces the looming consequences of climate change, and energy insecurity continues to catalyze conflicts around the globe. Humanity must urgently tackle the challenge of scaling cheap, clean, and carbon-free energy.

Fusion energy, the nuclear process that powers the stars, has long been sought as a practical energy source (Harms et al., 2000). Fusion uses abundant fuel, emits no carbon, and uses minimal land[1]. For these reasons, fusion is an alluring possibility as a component of future energy grids.

We focus in this paper on the most technically mature magnetic fusion device: the "tokamak." A tokamak is a doughnut-shaped machine that confines a steady-state plasma at fusion-relevant temperatures by applying external magnetic fields and inducing a current in the plasma that circulates around the torus, creating a helical magnetic field that confines the plasma (Wesson & Campbell, 2011; Taylor, 1997).

Fusion entered a new era in the 2020s because of three factors: technical innovation, a growing awareness of the climate crisis, and big bets by venture capital. The development of high temperature superconductors, for example, put previously infeasible fusion concepts within the grasp of commercial-scale labs. At the same time, the cost of climate change and rapid decarbonization is coming into view. These factors brought in billions from established venture capital firms to the nascent fusion start-up landscape (Fusion Industry Association, 2023; Parisi & Ball, 2019).

However, several "showstopping" challenges remain; it is our position that Machine Learning (ML, see common acronyms in Appendix A) may be key to addressing several critical issues for tokamak fusion. Tokamaks can generate massive amounts of data and require fairly sophisticated control schemes. Furthermore, the field has been generating data across an array of modalities for decades. It is our observation that there are low-hanging fruit abound for ML practitioners in fusion, and we strongly believe that collaboration between the two fields will be important and mutually beneficial.

[*]Equal contribution [1]Plasma Science and Fusion Center, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA [2]Department of Electrical Engineering and Computer Science, Carnegie Mellon University, 5000 Forbes Avenue Pittsburgh, PA 15213, USA [3]Robert H. Smith School of Business, University of Maryland, College Park, 7901 Regents Drive College Park, MD 20742-5025, USA [4]Swiss Plasma Center, Swiss Federal Institute of Technology, Chemin du Barrage 16, 1015 Ecublens, Lausanne, Switzerland.. Correspondence to: Lucas Spangher <spangher@psfc.mit.edu>, Allen Wang <awang@psfc.mit.edu>, Andrew Maris <maris@psfc.mit.edu>, Cristina Rea <crea@psfc.mit.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

[1]Nuclear fusion should not be confused with nuclear *fission*, which is a fundamentally different process. No fusion power plants exist today, but fission reactions are responsible for nuclear power plants (Kulcinski et al., 1979).

In this position paper, we lay out the main open problems in fusion that are relevant to ML, and future opportunities that will emerge as fusion data infrastructure is upgrade for the era of ML. Few others have written on the broader role of ML in fusion, and all have address fusion practitioners (Kamal, 2020; Humphreys et al., 2019) or covered narrow areas of ML and fusion (Pavone et al., 2023). Ours is the first to attempt a comprehensive overview of ML in tokamak fusion and to address the ML community directly. **It is the position of this paper that there are several avenues of exciting collaborative research directions, and that the future success of fusion may well hinge on support from the ML community.**

2. Open Problems Relevant to ML

The fusion community has shown an increasing interest in ML-accelerated research. A notable example is the recent creation of a 5-year coordinated research project (CRP) sponsored by the International Atomic Energy Agency (IAEA) to foster a multi-institutional network focused on ML-applied fusion research (AI for Fusion, 2023).

In this section, we will discuss the main open problems in magnetic fusion relevant to ML today; i.e., the tasks on which interested researchers can immediately begin collaboration with domain experts. For each modeling challenge, we will frame the problem, describe existing work, and discuss opportunities for further progress. Since we write for ML practitioners, we put further details on the physics of fusion in Appendix B.

2.1. Disruption Prediction

2.1.1. Background of problem

A "disruption" is a rapid, total loss of plasma confinement in a tokamak. These events can be prompted by a range of factors from human error to internal plasma instabilities, and result in the plasma suddenly colliding with the inside of the confinement chamber. While disruptions in today's experiments are manageable, future power plant-scale tokamaks could face major thermal loads and mechanical stresses during disruptions that jeopardize the machine health and plant's viability (Maris et al., 2023).

Disruptions are extremely challenging to predict in practice. A plasma can become unstable on millisecond timescales, and many contributing factors to disruptions are not well-understood. The lack of first-principle models combined with abundant experimental data across decades of operations lends itself to ML solutions. ML disruption predictors can be integrated into a tokamak's control system to help steer the plasma towards a more stable operating regime, or attempt to preemptively mitigate an upcoming disruption by launching a large amount of cold matter into the tokamak. ML disruption predictors can also supplement or integrate with physics-based models for disruption prediction, such as the DECAF suite of codes (Sabbagh et al., 2023).

2.1.2. Prior and Current Work

Disruption prediction through statistical modeling and machine learning was first explored through in the 2000's (Windsor et al., 2005), though the body of literature has expanded rapidly in recent years (Cannas et al., 2007; Murari et al., 2008; 2009; Rattá et al., 2010; De Vries et al., 2011; Vega et al., 2013; Cannas et al., 2014; Rattá et al., 2014; Aledda et al., 2015; Montes et al., 2019; Murari et al., 2018; Rea & Granetz, 2018). Notable recent applications include more advanced architectures such as temporal (Churchill et al., 2020; Zhu et al., 2021) and spatial (Aymerich et al., 2022) convolutional neural nets (CNNs) and variational autoencoders (VAEs) (Wei et al., 2021). Other example applications span from tree-based models (Zhong et al., 2021; Rea et al., 2018; 2019) to recurrent networks (Kates-Harbeck et al., 2019), manifold learning and generative mapping (Pau et al., 2019), and survival analysis (Tinguely et al., 2019; Keith et al., submitted). Potential future models may incorporate explicit autoregression, attention or long-range convolutions for historic patterns, continuous state space dynamics, and implicit taxonomization of the range of physically possible instability pathways.

2.1.3. Challenges

While many promising advances have been made in disruption prediction, two key challenges stand out.

Low failure requirement: Classical ML shines in "51/49" problems, i.e. those in which one needs only be correct slightly more often than they are incorrect, such as stock sales (Ginsberg, 2012; Ginsberg, 2021). Disruption prediction for a tokamak power plant will be a "99/01" problem, i.e. one in which as high an accuracy as possible (>99%) is important because even a few failures may be catastrophic to the machine. Disruption prediction work should be closely aware of advances in fields of ML safety and robustness. Also, the severity of a disruption may be information incorporated into the training of a model; for example, one may weight the training loss by a measure of observed disruption severity, ensuring the model is more attuned to the especially dangerous cases.

Explainability: "Black-box" ML disruption prediction models often show strong performance, however they have generally been unable to meaningfully improve our understanding of disruption physics. Additionally, deep models that cannot be easily validated for safety have limited applicability for future control systems that must reliably drive the plasma away from disruptive boundaries. Promising solutions need to include a level of explainability, either

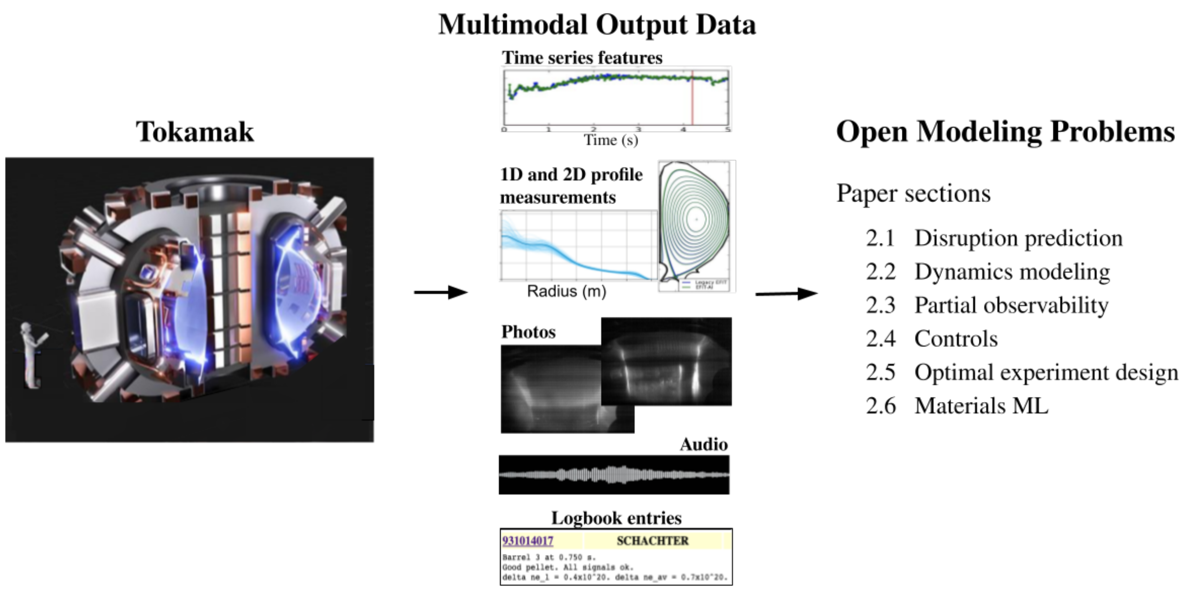

Figure 1. An overview of the types of data tokamaks can generate, and the modeling problems that this contributes to. Profile graphs taken from (Kwak et al., 2021), time series data taken from (Lao et al., 2022), and the tokamak image reproduced with permission from CFS/MIT-PSFC — CAD Rendering by T. Henderson.

through post-hoc models (Rea et al., 2019) or via algorithms that are interpretable by design. Finally, another promising direction follows a shift in paradigm to assess the disruption risk as a function of the plasma survival time (Tinguely et al., 2019; Keith et al., submitted).

2.2. Simulation and Dynamics Modeling

2.2.1. Background of problem

Simulation is crucial for realizing the robust and economical design, operations, and control of a power plant. While the fusion community has made massive strides in developing simulation tools for reactor design, such tools typically require massive computational resources to simulate just one time-step of the plasma (Rodriguez-Fernandez et al., 2022b). Considerable advances in time-varying simulation are needed to advance both the planning of operational trajectories and training of control policies. Regardless of the exact control techniques utilized, which may include anything from classical feedback control to contemporary policy optimization methods, advances in control-oriented simulation and dynamics modelling are needed. Arguably, the first demonstration of reinforcement learning (RL) for plasma control (Degrave et al., 2022) succeeded because the physics relevant to the particular control task, "shape control", are the most well-understood of tokamak physics, with simulations that work well in practice (Carpanese, 2021).

The domain of robotics has blazed a path to follow, with efforts to build simulators explicitly for the purpose of advanced control, such as MuJoCo (Todorov et al., 2012) and Isaac Gym/Sim (Makoviychuk et al., 2021), making a large impact on real-world learning-based control via features including but not limited to: speed, physics fidelity, modularity, an easy-to-use API, automatic differentiation, and GPU parallelization. While plasma simulation faces a considerable set of challenges beyond those of robotics simulation, it is worthwhile to transfer relevant learnings from robotic simulation to plasma simulation.

2.2.2. Prior and Current Work

Simulating tokamak plasmas faces a considerable set of challenges. Physics based approaches suffer from the challenges of incomplete models and extreme computational requirements. Purely data driven sequence-to-sequence approaches have been explored (Abbate et al., 2021; Kolemen et al., 2023; Char et al., 2023), but they face the challenges of data paucity and distributional drift of fusion machines. We elaborate on promising approaches for future work after each relevant section in the Challenges below.

2.2.3. Challenges

Incomplete physics-based models: Important physical phenomena relevant to plasma control are not well-understood from first principles, forcing the use of empirical models in otherwise first-principles based simulations. One notable example is the phenomenon of the H-mode confinement

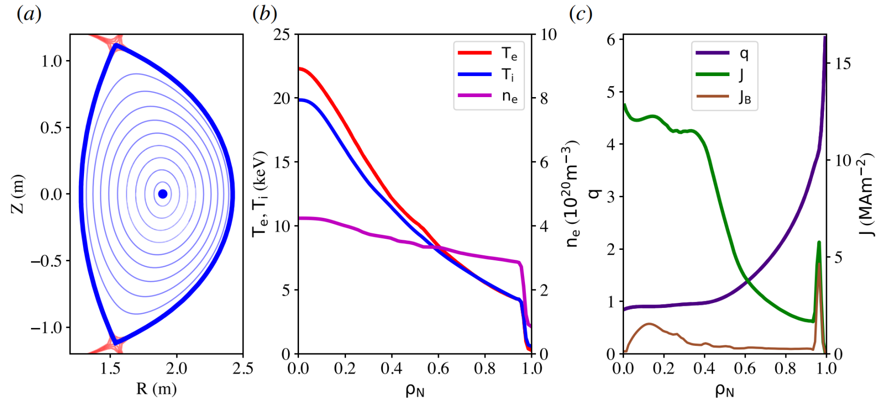

Figure 2. Here we reproduce with permission (Rodriguez-Fernandez et al., 2020) an example of the types of plasma simulation output that ML may enhance. Note that $R\ (m)$ implies linear distance from tokamak center, $Z\ (m)$ the vertical height, ρ_N the normalized minor radial coordinate, T_e and T_i the electron and ion temperatures respectively, q the safety factor (a measure of rotational transform of the magnetic field), and J the toroidal current density. Subfigures: (a) a plasma boundary is used as input to a simulation code (TRANSP, (Hawryluk & Coppi, 1981)) and internal flux surfaces as calculated by the fixed-boundary equilibrium solver. (b) Electron and ion temperature and electron density profiles. (c) Several plasma current-relevant profiles such as the safety factor (q), the total toroidal current density (j), and toroidal current component from bootstrap current (j_B).

regime, which typically doubles a key performance metric (Martin et al., 2008). On this front, considerable opportunities exist both for learning subsets of plasma dynamics from empirical data to use inside physics-based simulators and for developing hybrid physics-ML models that assimilate both physics-based principles and empirical data. Such techniques have already proven crucial in robotics for enabling ML based control; (Hwangbo et al., 2019), for example, learned an actuator dynamics model from robot data to augment a physics-based RL training environment.

Extreme computational requirements: Another considerable challenge is the extreme computational cost of simulating certain subsets of plasma dynamics from first-principles. For example, properly simulating the dynamics in the plasma core from first principles currently takes millions of CPU-hours for a single time slice, as it involves solving a complex set of partial differential equations (PDEs) (Rodriguez-Fernandez et al., 2022b). On this front, ML-based *surrogate models* (also referred to as *emulations*) have already made an impact in enabling simulations at a level of fidelity previously considered intractable (Rodriguez-Fernandez et al., 2022b; Di Siena et al., 2022; van de Plassche et al., 2020). However, even surrogate model accelerated simulations of plasma models, such as nonlinear gyrokinetics, still requires millions of CPU-hours (Rodriguez-Fernandez et al., 2022b). Thus, many orders of magnitude of improvement are still needed. Real-time simulations are another area that will benefit from surrogate modelling (Felici et al., 2018); for example, real-time reconstruction is typically restricted to the 2D case as the 3D case, which exist for offline codes (Hoelzl et al., 2020), are intractable in realtime. Surrogate modelling may also be a point of synergy between plasma simulation and materials design (Sec 2.6).

Stochasticity: There are certain aspects of the plasma dynamics that are important, but seem difficult or impossible to model deterministically. Some examples include the introduction of impurities and the onset of certain instabilities. To address these issues, approaches from uncertainty quantification, generative modelling, and stochastic differential equations should be explored to augment simulations to enable the training of controllers that are robust to these sources of uncertainty.

Data paucity and **distributional drift**: The primary challenges of applying ML to learn plasma dynamics on empirical data are the relative paucity of data and the distributional drift of machines over time. Both are discussed further in Section 3, but here we reiterate the necessity of sample efficiency for upcoming net energy experiments which need to succeed with as little trial and error as possible.

2.3. Partial Observability

2.3.1. Background of problem

Sensing the state of the plasma is a difficult task, constituting its own field of work. Observations of the plasma state are typically incomplete, motivating the application of specialized techniques to infer the underlying plasma state. This section describes the following three problems:

Undetermined feature set: Certain key quantities in fusion plasmas are spatially varying, but are measured by sensors that cannot physical resolve their spatial distribution. For example, measuring magnetic fields inside the plasma would require inserting sensors into the plasma. Thus the typical technique, known as *equilibrium reconstruction*, uses measurements from magnetic field sensors at the edge in conjunction with a physics-based model to infer magnetic fields inside the plasma (Moret et al., 2015; Lao et al., 1985; Li et al., 2013). Such reconstructions can be computationally demanding or inaccurate.

Data assimilation across machines and sensors: Different machines can have varying numbers and types of sensors, and each sensor itself can produce outputs of varying sampling rates and spatial distributions. For example, interferometers measure the electron density at high frequency along a sight line while Thomson Scattering measures plasma density and temperature at individual points in space with low frequency, motivating the application of methods to combine information from both sensor modalities (Pastore et al., 2023). In addition, sensors are often modified or shut off based on the needs of the individual principle investigator responsible for the sensor.

Noisy and missing measurements: Sensors may unintentionally add different types of noise to measurements. Also, quantities of interest may be difficult to measure, resulting in dropped measurements. For example, we cannot easily measure the ion temperature and density, and so the output is generally noisy. To make matters worse, it can be difficult to tell whether those quantities were noisy or missing: in many tokamak monitoring sytems, sensors that fail simply report white noise. Thus, robustness to measurement error is essential for an array of modeling tasks (Shousha et al., 2023).

2.3.2. Prior and Current Work

Much work in feature extraction has focused on physics-informed NN's (PINNs) (Mathews, 2022; Aymerich et al., 2023; Mathews et al., 2020b) primarily through loss function curtailing (Huang & Wang, 2023); more work is needed. Variational Information Bottlenecks have been used for equilibrium reconstruction emulation (Lao et al., 2022; Kruger et al., 2022), as have transformers (Wan et al., 2022) and PINNs (Mathews et al., 2020a). Diffusion models, LSTMs, multi-adaptive auto-regressive splines (MARS) (Rasul et al., 2021) and masked VAEs (Variational Autoencoders) have been used for data reconstruction and denoising (He et al., 2021; Lei et al., 2016). State-space and S4-type (Structured State Space for Sequence) models have been used for data assimilation tasks (Arnold et al., 2023). We expect future models to feature continuous state spaces and increased physics hybridization.

2.3.3. Challenges

In addition to the challenges that partial observability shares with other problems (see Section 3), we expect models addressing to gnenerally require especially low latencies and be robust to an array of boundary conditions. Any model that reconstructs an underdetermined state space should be undertaken with "physicists in the loop": windows into functioning and reconstructed variables are essential to make sure that it is recreating known and expected physics.

2.4. Control

2.4.1. Background of problem

This section provides an abbreviated description of plasma control challenges and opportunities, some of which are highlighted in Figure 3. The reader is referred to (Walker et al., 2020) for a more in-depth introduction to tokamak plasma control. Plasma operation in a fusion device requires active control. During the entire evolution of a plasma discharge, the control system must ensure the desired conditions, defined by physics and machine operators, are met while maintaining plasma stability. A considerable challenge is that many, sometimes conflicting, control objectives and tasks must be met simultaneously to ensure the plasma dynamics evolution is desirable. Net-energy experiments introduce stringent requirements on the robust handling of off-normal events, such as the onset of instabilities, and simultaneous control of multiple plasma quantities, making the development of advanced plasma control systems increasingly urgent.

2.4.2. Prior and Current Work

The most basic operation mode of a tokamak only requires control of the position and current of the plasma, often described as "RZIP" control (i.e. Radial position, Z position, and I_p control). *Magnetic control* typically refers to RZIP control with the addition of control of the plasma shape. While magnetic control is typically done with techniques from classical linear control (De Tommasi, 2019), it has also recently been demonstrated with RL in the TCV (Degrave et al., 2022), KSTAR (Seo et al., 2021), and EAST tokamaks (Wan et al., 2022) (see table of tokamaks, Appendix A). However, a number of outstanding challenges remain for the RL approach to be realized in routine tokamak operations (Tracey et al., 2024). Future works may investigate model features such as planning (Spangher et al., 2020), action-space constraints (Arnold et al., 2021), offline imitation (Jang et al., 2021), and autocurricula generation (Jang et al., 2023).

2.4.3. Challenges in control

While RZIP control enables basic operations, robust high performance control of net energy tokamaks faces many more challenges. All of these control challenges share the common challenge of accurate simulation and dynamics modelling, a topic addressed at length in Sec 2.2. For

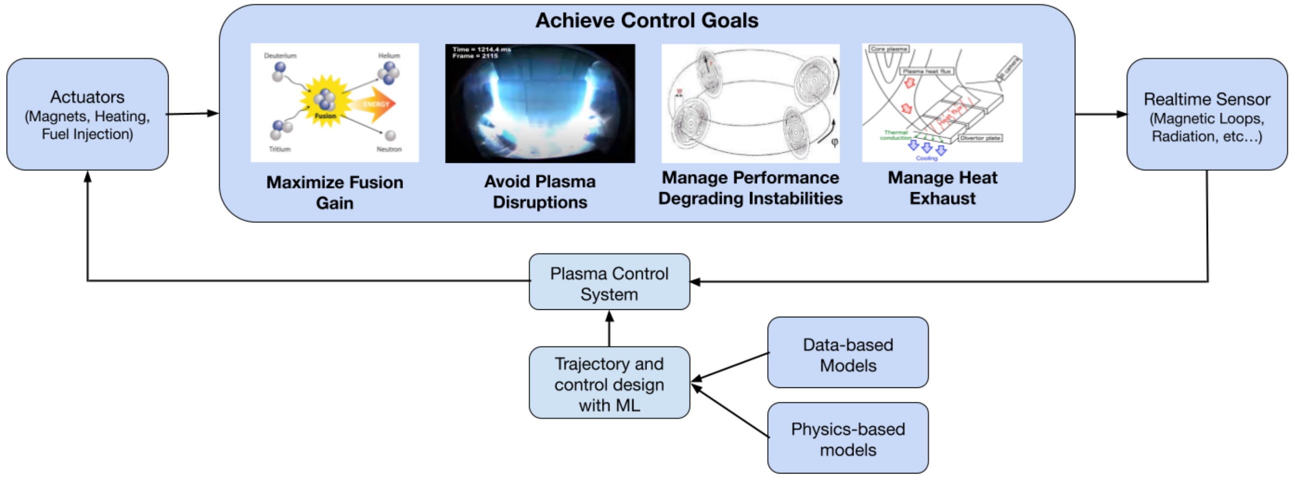

Figure 3. Diagram showing important goals for advanced control research to address along with key relevant components. Inspiration for the diagram was taken from (Humphreys et al., 2019). Credit for embedded figures, from left to right: (Gunn, 2022), (IPP, 2017), (Igochine), (Hayashi et al., 2021).

one, active avoidance of potentially unstable portions of the plasma state space remains an open problem (Boyer et al., 2021). Certain instabilities also degrade plasma performance; a recent work demonstrated the application of RL to avoiding one such instability to increase plasma performance (Kolemen et al., 2023). When off-normal events happen, the control system also needs to make intelligent decisions about how to handle the off-normal event (Vu et al., 2021). Net energy devices such as the SPARC and ITER tokamaks have to contend with excessive heat damaging the wall; techniques such as detachment (Leonard, 2018) and strikepoint control (Kolemen et al., 2010) have been developed to tackle this problem, but it remains far from solved.

2.5. Optimal Design of Experiments

2.5.1. Background of problem

In the context of fusion, there are potentially many applications for optimal experiment design due to the expense of large-scale experiments; for example, at the DIII-D tokamak, experiments are allocated through an annual process awarding blocks of 10-20 experiments. For upcoming experiments attempting net-energy, machine damage is a major concern (Lehnen et al., 2016), making each trial a risk. Since future experiments will be heavily trial-limited, each trial must maximally inform models used for planning, trajectory, and control design to reliably control the plasma.

2.5.2. Prior and Current Work

In fusion research, scientists often conduct 'scans' in order to understand how quantities, such as the level of heat transport in the plasma, change in response to a given operating scenario of interest. In these scans, scientists typically vary a design variable in a linearly spaced manner. This is likely a sub-optimal usage of experimental resources. Though there are many works (Chaloner & Verdinelli, 1995; Gevers & Ljung, 1986) addressing this problem in general and with progressively weakening assumptions (Foster et al., 2019), there has been minimal work applying these ideas to the practical process of running fusion experiments, though it has been discussed in Murari et al. (2021). One notable example is a stochastic perturbation method, which was applied to a magneto-inertial fusion experiment to assist in designing experimental settings (Baltz et al., 2017). Another notable example is the development of a statistical model to correct simulations that was successfully used to significantly increase fusion gain in inertial confinement fusion (Baltz et al., 2017).

Another common experimental goal is to find parameter settings or control trajectories that maximize some objective function, a setting appropriate for *Bayesian Optimization* (Frazier, 2018; Shahriari et al., 2015). Here, typical approaches involve choosing trials which maximize information gain (Hennig & Schuler, 2012; Hernández-Lobato et al., 2014) among expected improvement in the best observed function value (Frazier, 2018), with several proposed alternatives. Recent work (Mehta et al., 2021) generalizes these ideas to design experiments for efficiently identifying an optimal trajectory. One early work in this direction (Mehta et al., 2024) addresses the problem of experimental design of trajectories for the "ramp-down" phase, where the plasma needs to be carefully de-energized. We believe that there could be many applications of similar techniques to improve other control solutions on tokamaks.

2.5.3. Challenges

Modeling and time constraints: There are substantial challenges in optimal experiment design. The ingestion and processing of the results from a particular trial needs to happen in a matter of minutes. As with other problems discussed here, there are substantial difficulties in modeling. With all that said, we believe that this is a fruitful area for research and that it is one which complements existing research processes well in that often, machine operation proceeds as usual.

2.6. Materials Design

2.6.1. Background of the problem

Materials in fusion power plants must survive extreme conditions: high temperatures, large mechanical loads, contact with corrosive media, and/or exposure to high energy neutrons (Cohen-Tanugi et al., 2023). Today, it is experimentally challenging – or impossible – to match the conditions of fusion power plants (Wirth, 2023). As such, modeling techniques based on density functional theory (DFT) and molecular dynamics (MD) play a key role in studying radiation damage (Yip & Short, 2013). Unfortunately, these atomistic models are computationally intractable to apply on macroscopic scales (meters and years).

2.6.2. Prior and Current Work

Machine learning interatomic potentials (MLIPs) have been successful in achieving close to DFT accuracy with 3-4 orders of magnitude speed up, using approaches such as Graph Neural Networks (GNNs) (Chen & Ong, 2022; Deng et al., 2023; Duval et al., 2023b; Bihani et al., 2024), Message Passing Neural Networks (MPNNs) (Musaelian et al., 2022), Transformers (Hedelius et al., 2024; Thölke & De Fabritiis, 2022), Gaussian Approximation Potentials (GAP) (Byggmästar et al., 2019), and Atomic Cluster Expansion (Batatia et al., 2022). Optimal experiment design and active learning add ML to the experimental loop, maximizing the value of each experiment (Damewood et al., 2023; Choudhary et al., 2020). Techniques such Gaussian Processes (GPs) and Bayesian optimization methods approaches show promise in automated lab setups (Ren et al., 2023a;b; Couet, 2022; Vecchio et al., 2021; Szymanski et al., 2023), and we expect future work may build on these advances (Lee et al., 2023). Similarly, GPs, Bayesian Uncertainty Sampling, and Transformers have been used for active learning (Morgan et al., 2022; Ren et al., 2023a).

2.6.3. Challenges

Representing radiation damage in an expressive and efficient manner: GNNs and MPNNs have successfully predicted the energies, forces, and stresses of tens to thousands of atoms (Duval et al., 2023a; Merchant et al., 2023), but struggle (MPNNs less so) to scale to hundreds of thousands of atoms (Zhang et al., 2021; Musaelian et al., 2022). Unfortunately hundreds of thousands to millions of atoms are needed to ensure radiation damage effects are contained in a "representative bulk", i.e. too small of a cell will cause moving atoms to escape the box and travel back in through periodic boundary conditions (for additional detail and references, refer to Appendix C.2). Therefore new MLIPs that can combine the universality of GNNs and the speed of GAPs are needed to enable quantum accurate modeling of systems large enough to learn radiation effects using molecular dynamics.

Fusion-relevant training data for MLIPs: Current datasets (Chanussot et al., 2021) for MLIPs are often insufficient for fusion-relevant materials with non-equilibrium structures due to radiation damage (Hamedani et al., 2021; Jin et al., 2018; Ko & Ong, 2023). Generating all possible defect clusters interacting with all possible elements in a system is computationally infeasible because defect structures and local chemical environments come in many forms (Byggmästar et al., 2019). That being said, more effort can be put into capturing these structures in MLIP training data. Additionally, active learning can guide the direction of DFT calculations to more efficiently compile training datasets (Qi et al., 2023). Additional information about active learning in automated labs is in Appendix C.2.

3. Shared Challenges

Transfer to out of domain distributions: Many ML models are based on categorically different tokamaks. Today's fusion devices operate in certain subsets of the power, temperature, size, magnetic field strength, and plasma discharge duration domains, yet models of disruption prediction will necessarily be implemented by hotter, more powerful machines. Models with weak inductive biases may face challenges extrapolating to the new domains of higher energy devices. Even if ML solutions objectively demonstrate superior performance, "black-box" solutions may face challenges with acceptance. Careful consideration should be paid to inductive bias selection and interpretability. Symbolic regression (Dubčáková, 2011; Udrescu & Tegmark, 2020; Cranmer, 2023), by contrast, could potentially be used to find closed-form expressions for stability limits that can be validated via known physical laws and relationships. More robustness to domain shift could also come from multitask approaches and utilizing simulations.

Distributional drift Another major challenge facing fusion devices is distributional drift across time. This is in part due to hardware upgrades, changes, and recalibrations of sensors and actuators. One notable challenge is the effect of material desorption or ablation off the tokamak wall, which affects plasma dynamics and is sensitive to the maintenance cycle of the wall (Shimada & Pitts, 2011). On this front, studies need to examine the impact of including historical time as an input into learned models. Advances in metadata curation are also needed to better account for machine changes.

Narrow and/or incongruous datasets: Given the expense and labor-intensity of fusion experiments, it is common for datasets in the field to be small and bespoke; certain sensors may not always be running or calibrated correctly, making it challenging to assemble large database with consistent settings data availability. Also, the pressure to publish encourages narrow, poorly documented, and closed datasets. Finally, dataset size and scope can be limited by the need for expert labeling. Many important events are hard to label automatically with thresholds, and so studies often rely on the individual scientist conducting the study to parse through and manually pick out the time and duration of events. Due to these factors, the data landscape in fusion (and in materials science for fusion, see Section 2.6) often looks more like an archipelago than a continent.

Despite these challenges, there are growing efforts in the **unification of disparate datasets**, of which the Open/FAIR data grant (National Science Foundation, 2022; Almada et al., 2020) and the DOE's PuRE grant (D.O.E, 2021) are examples of better practices; for examples in fusion see (Montes et al., 2020). While tearing down the walls between the narrow and incongruous datasets often found in fusion will take significant effort, it will open many doors for ML-based analysis.

The problem of **data scarcity** is particularly pertinent as it is desirable for upcoming net-energy tokamaks such as SPARC (Creely et al., 2020) and ITER (Aymar et al., 2002) to work with as few trial runs as possible. To address the relative paucity of data, a range of data-efficient approaches to dynamics modelling should be explored. Works in other data scarce domains offer some promising starting points. Some examples include GNNs for weather forecasting (Lam et al., 2022), neural differential equations and differentiable simulation in several domains (Djeumou et al., 2023; Kidger, 2022; Wang et al., 2023), and development of a custom architecture like AlphaFold (Jumper et al., 2021).

4. Fusion Data Ecosystems in the Age of ML

Since the 1950s, tokamak experiments have generated increasingly large amounts of data. While Alcator C-mod in 1991 produced around 5 megabytes per day, the ITER tokamak, when it comes online, will produce 2 petabytes of data per day of experiments (ITER Team, 2020).

A significant volume of fusion data is stored in MDSplus (Stillerman et al., 1997; Lane-Walsh et al., 2021), a hierarchical database and data access system. While MDSplus has provided a significant service to the fusion community, it was developed prior to the current era of ML and big data, and is thus currently not optimized for ML-scale workflows. For one, it encounters bottlenecks when retrieving data from multiple experiments (also called "plasma discharges" or "shots") as it is designed for single discharge retrieval. There also lacks a consistent data schema, with each device having its own data structure that evolves over time. On this front, the closed-source ITER Integrated Modeling and Analysis System (IMAS) is working to standardize a data schema (Romanelli et al., 2020).

There are also administrative and legal barriers to fusion data access. Currently, individuals, not institutions, need to request access and sign written agreements with each institution operating the tokamak to access data. While it is therefore not impossible to gain data access to many tokamaks, these restrictions have prevented sharing of multi-machine datasets between researchers, and stand as a major barrier to entry for members of the ML community.

There is ongoing work to provide tools to smoothly aggregate fusion datasets at scale. These tools have several goals in mind (DOE, 2023): (1) organize data retrieval systems around a cache structure for signals relevant to ML-type scale (2) provide a comprehensive mapping between machines for retrieval of similar nodes (3) ease a one-stop-shop between multiple tokamaks for researcher's use, and (4) parsing of stored logbook records with AI (Mehta et al., 2023). Input from ML communities will be invaluable for steering and accelerating this process.

5. Future ML Opportunities on the Horizon

Here, we will speculate on future fusion ML work possible once the presence of other systemic advances are in place.

5.1. ML with ML-Scale Data

As data gathering infrastructure grows, so too will the scale of ML. Models trained across multi-device datasets will become the norm, with care taken to account for data distribution drift that will occur throughout a device's lifetime. We expect models to be sampling rate independent, as sampling rates will be different across machines and sensors.

A future capability enabled by data scale may be foundation-type models for fusion applications. Evidence suggests that multi-task training can benefit individual goals in fusion; prior work has shown that next state prediction pre-training boosts the accuracy of disruption prediction in transformer-like models (Spangher et al., 2023), and other work has shown that multi-headed multi-task transformers have higher disruption prediction capabilities (Wan et al., 2022).

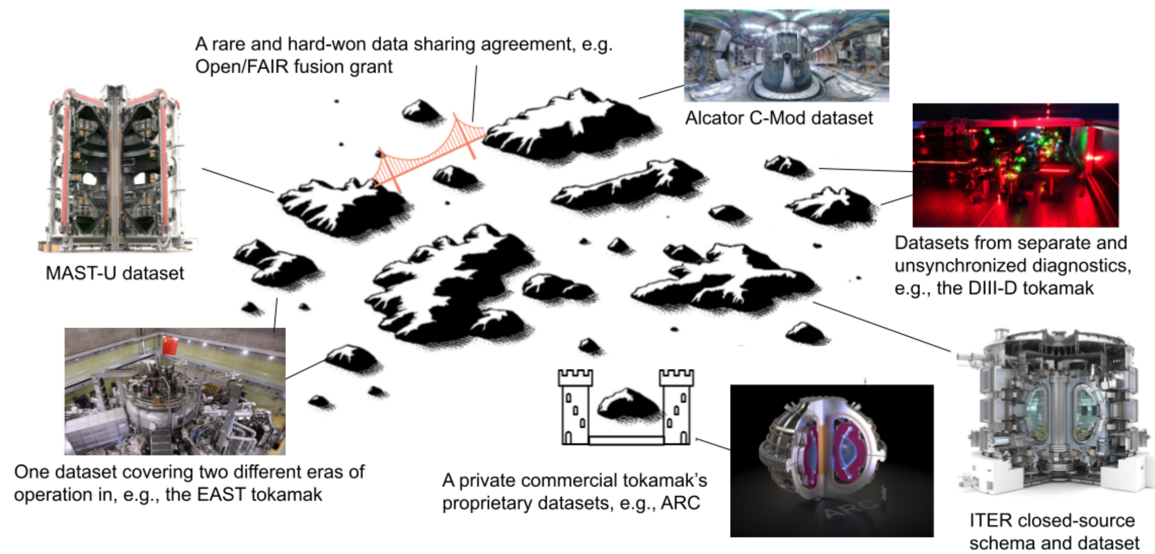

Figure 4. The Archipelago of Magnetic Fusion Data. This illustration highlights key challenges in fusion data: (1) lack of integrated datasets within tokamaks, (2) asynchronous diagnostics and data types, (3) inconsistencies across different operational eras, (4) difficult data sharing agreements, and (5) proprietary nature of upcoming large-scale commercial datasets. Please note that the illustration is symbolic and not intended to critique any specific tokamak.

Other future work may seek to address data privacy and access restriction by implementing types of federated learning to gain information from multi-machines while keeping data on national servers. Such an endeavor would likely require coordination beyond the development of a new data aggregation framework.

5.2. ML with Advanced Multi-Modalities

The ML community is beginning to investigate and incorporate multi-modality in models (Gao et al., 2020; Lahat et al., 2015), also called "data fusion". This advance happens as multi-modal benchmarks are built (Ge et al., 2023; Fu et al., 2023; Liang et al., 2021), multi-modal datasets grow (Gadre et al., 2023), and desire for multi-modal capabilities in large language model (LLM) grow (Yin et al., 2023).

Fusion datasets naturally feature many data modalities. In addition to photos, 2D profiles, time-series of global state variables, static metaparameters, reconstructed data from equilibrium models, fusion datasets also feature detailed text logbook entries about shots, known physics insights, and even audio waves from the tokamak hall. Successfully incorporating multi-modality may make fusion ML far more robust to noise in any one modality, aware of different physical events, and generalizable to new operating regimes of temperature and pressure.

5.3. ML for Grid Integration and Plant Operation

ML work on energy systems writ large has helped address economic dispatch of generators (Chahar et al., 2021), predict grid demand (Chow, 2021), and siting resources (Sun et al., 2023; Petrov & Wessling, 2015).

As fusion plants approach commercial deployment, ML may similarly play a role in hour-by-hour operational decisions (i.e. hourly energy dispatch) as well as actions taken once over a plant's lifetime (i.e. siting, scaling, and messaging to host communities.) AI has been used for automation, optimization, and analysis of various operations of in nuclear power plants under the support of the IAEA (IAEA, 2022), and many applications may transfer well to fusion.

6. Conclusion

We believe the future of nuclear fusion is bright, with the potential to play a significant role in de-carbonization and grid stabilization. Many problems remain on the path to commercial fusion energy, several of which lend themselves naturally to ML solutions, motivating cross-community collaboration. While the barrier to entry into fusion research is currently high, ongoing efforts are working to reduce said barriers to entry; further cross-community collaborations will only help to accelerate this process. We hope this position paper provided a useful starting point to help the ML practitioner understand the landscape of problems and identify opportunities to make an impact on fusion, and thus the future of energy.

Acknowledgements

We gratefully acknowledge those who have advised our work, including Robert Granetz, Alex Tinguely, Sara Ferry, and Michelle Kuchera. We are thankful for the supportive community of Ph.D. students and postdocs at MIT's PSFC, including Rebecca Masline, Zander Keith, Stuart Benjamin, Abby Feyrer, Benjamin Stein-Lubrano, Henry Wietfeldt, Andrew Rothstein, Ian Char, and Amelia Cavallaro. We gratefully acknowledge the plasma physics and nuclear fusion community for their tireless work for a shared vision of clean, compact, waste-free energy.

Impact Statement

The goal of our perspective piece is to inspire more collaborations between ML practitioners and the nuclear fusion community. Such work would enable a faster path to commercialization and deployment of fusion devices, thereby improving our tools to address powerful climate change (Spangher et al., 2019) and energy security.

References

Abbate, J., Conlin, R., and Kolemen, E. Data-driven profile prediction for diii-d. *Nuclear Fusion*, 61(4):046027, 2021.

AI for Fusion. AI for fusion: accelerating fusion r&d with ai, through the creation of a platform and cross-community network for innovation and partnership, September 2023. URL https://nucleus.iaea.org/sites/ai4atoms/ai4fusion/SitePages/AI4F.aspx.

Aledda, R., Cannas, B., Fanni, A., Pau, A., Sias, G., Team, A. U., et al. Improvements in disruption prediction at asdex upgrade. *Fusion Engineering and Design*, 96:698–702, 2015.

Almada, M., Midão, L., Portela, D., Dias, I., Núñez-Benjumea, F., Parra-Calderón, C. L., and Costa, E. [a new paradigm in health research: Fair data (findable, accessible, interoperable, reusable)]. *Acta medica portuguesa*, 33 12:828–834, 2020.

Arnold, W., Srivastava, T., Spangher, L., Agwan, U., and Spanos, C. J. Adapting surprise minimizing reinforcement learning techniques for transactive control. *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*, 2021. doi: 10.48550/arXiv.2111.06025.

Arnold, W., Spangher, L., and Rea, C. Continuous convolutional neural networks for disruption prediction in nuclear fusion plasmas. *ArXiv*, abs/2312.01286, 2023. doi: 10.48550/arXiv.2312.01286.

Aymar, R., Barabaschi, P., and Shimomura, Y. The iter design. *Plasma physics and controlled fusion*, 44(5):519, 2002.

Aymerich, E., Sias, G., Pisano, F., Cannas, B., Carcangiu, S., Sozzi, C., Stuart, C., Carvalho, P., and Fanni, A. Disruption prediction at jet through deep convolutional neural networks using spatiotemporal information from plasma profiles. *Nuclear Fusion*, 62, 2022.

Aymerich, E., Pisano, F., Cannas, B., Sias, G., Fanni, A., Gao, Y., Böckenhoff, D., Jakubowski, M., et al. Physics informed neural networks towards the real-time calculation of heat fluxes at w7-x. *Nuclear Materials and Energy*, 34:101401, 2023.

Baltz, E., Trask, E., Binderbauer, M., Dikovsky, M., Gota, H., Mendoza, R., Platt, J., and Riley, P. Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm. *Scientific reports*, 7(1):6425, 2017.

Batatia, I., Kovacs, D. P., Simm, G. N. C., Ortner, C., and Csanyi, G. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. 2022. URL https://openreview.net/forum?id=YPpSngE-ZU.

Batzner, S., Musaelian, A., Sun, L., Geiger, M., Mailoa, J. P., Kornbluth, M., Molinari, N., Smidt, T. E., and Kozinsky, B. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. 13(1):2453, 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-29939-5. URL https://www.nature.com/articles/s41467-022-29939-5. Number: 1 Publisher: Nature Publishing Group.

Bihani, V., Mannan, S., Pratiush, U., Du, T., Chen, Z., Miret, S., Micoulaut, M., Smedskjaer, M. M., Ranu, S., and Krishnan, N. A. Egraffbench: evaluation of equivariant graph neural network force fields for atomistic simulations. *Digital Discovery*, 3(4):759–768, 2024.

Boyer, M., Rea, C., and Clement, M. Toward active disruption avoidance via real-time estimation of the safe operating region and disruption proximity in tokamaks. *Nuclear Fusion*, 62(2):026005, 2021.

Buttery, R., Covele, B., Ferron, J., Garofalo, A., Holcomb, C., Leonard, T., Park, J., Petrie, T., Petty, C., Staebler, G., et al. Diii-d research to prepare for steady state advanced tokamak power plants. *Journal of Fusion Energy*, 38:72–111, 2019.

Byggmästar, J., Hamedani, A., Nordlund, K., and Djurabekova, F. Machine-learning interatomic potential for radiation damage and defects in tungsten.

100(14):144105, 2019. doi: 10.1103/PhysRevB.100.144105. URL https://link.aps.org/doi/10.1103/PhysRevB.100.144105. Publisher: American Physical Society.

Byggmästar, J., Nordlund, K., and Djurabekova, F. Modeling refractory high-entropy alloys with efficient machine-learned interatomic potentials: Defects and segregation. 104(10):104101, 2021. doi: 10.1103/PhysRevB.104.104101. URL https://link.aps.org/doi/10.1103/PhysRevB.104.104101. Publisher: American Physical Society.

Cannas, B., Delogu, R. S., Fanni, A., Sonato, P., Zedda, M. K., contributors, J.-E., et al. Support vector machines for disruption prediction and novelty detection at jet. *Fusion engineering and design*, 82(5-14):1124–1130, 2007.

Cannas, B., Fanni, A., Murari, A., Pau, A., Sias, G., and Contributors, J. E. Overview of manifold learning techniques for the investigation of disruptions on jet. *Plasma Physics and Controlled Fusion*, 56(11):114005, 2014.

Carpanese, F. Development of free-boundary equilibrium and transport solvers for simulation and real-time interpretation of tokamak experiments. Technical report, EPFL, 2021.

Chahar, R. K., Srivastava, S., Prasad, S., and Singh, A. S. A faster optimal solution for economic dispatch problem in a smart grid environment using tensorflow. *2021 3rd International Conference on Advances in Computing, Communication Control and Networking (ICAC3N)*, pp. 366–369, 2021.

Chaloner, K. and Verdinelli, I. Bayesian experimental design: A review. *Statistical science*, pp. 273–304, 1995.

Chanussot, L., Das, A., Goyal, S., Lavril, T., Shuaibi, M., Riviere, M., Tran, K., Heras-Domingo, J., Ho, C., Hu, W., et al. Open catalyst 2020 (oc20) dataset and community challenges. *Acs Catalysis*, 11(10):6059–6072, 2021.

Char, I., Abbate, J., Bardóczi, L., Boyer, M., Chung, Y., Conlin, R., Erickson, K., Mehta, V., Richner, N., Kolemen, E., et al. Offline model-based reinforcement learning for tokamak control. In *Learning for Dynamics and Control Conference*, pp. 1357–1372. PMLR, 2023.

Chen, C. and Ong, S. P. A universal graph deep learning interatomic potential for the periodic table. 2(11):718–728, 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00349-3. URL https://www.nature.com/articles/s43588-022-00349-3. Number: 11 Publisher: Nature Publishing Group.

Choudhary, K., Garrity, K. F., Reid, A. C., DeCost, B., Biacchi, A. J., Hight Walker, A. R., Trautt, Z., Hattrick-Simpers, J., Kusne, A. G., Centrone, A., et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj computational materials*, 6(1):173, 2020.

Chow, H. F. Short-term electricity grid maximum demand forecasting with the arimax-svr machine learning hybrid model. *HKIE Transactions*, 2021.

Churchill, R., Tobias, B., Zhu, Y., team, D.-D., et al. Deep convolutional neural networks for multi-scale time-series classification and application to tokamak disruption prediction using raw, high temporal resolution diagnostic data. *Physics of Plasmas*, 27(6), 2020.

Coda, S., Ahn, J., Albanese, R., Alberti, S., Alessi, E., Allan, S., Anand, H., Anastassiou, G., Andrèbe, Y., Angioni, C., et al. Overview of the tcv tokamak program: scientific progress and facility upgrades. *Nuclear Fusion*, 57(10):102011, 2017.

Cohen-Tanugi, D., Stapelberg, M. G., Short, M. P., Ferry, S. E., Whyte, D. G., Hartwig, Z. S., and Buonassisi, T. Materials for fusion energy: Key challenges and guiding principles for developing under high design uncertainty. *arXiv preprint arXiv:2311.12187*, 2023.

Couet, A. Integrated high-throughput research in extreme environments targeted toward nuclear structural materials discovery. 559:153425, 2022. ISSN 0022-3115. doi: 10.1016/j.jnucmat.2021.153425. URL https://www.sciencedirect.com/science/article/pii/S0022311521006450.

Cranmer, M. Interpretable machine learning for science with pysr and symbolic regression. jl. *arXiv preprint arXiv:2305.01582*, 2023.

Creely, A., Greenwald, M. J., Ballinger, S. B., Brunner, D., Canik, J., Doody, J., Fülöp, T., Garnier, D., Granetz, R., Gray, T., et al. Overview of the sparc tokamak. *Journal of Plasma Physics*, 86(5):865860502, 2020.

Damewood, J., Karaguesian, J., Lunger, J. R., Tan, A. R., Xie, M., Peng, J., and Gómez-Bombarelli, R. Representations of materials for machine learning. *Annual Review of Materials Research*, 53:399–426, 2023.

De Tommasi, G. Plasma magnetic control in tokamak devices. *Journal of Fusion Energy*, 38(3-4):406–436, 2019.

De Vries, P., Johnson, M., Alper, B., Buratti, P., Hender, T., Koslowski, H., Riccardo, V., Contributors, J.-E., et al. Survey of disruption causes at jet. *Nuclear fusion*, 51(5):053018, 2011.

Degrave, J., Felici, F., Buchli, J., Neunert, M., Tracey, B., Carpanese, F., Ewalds, T., Hafner, R., Abdolmaleki, A., de Las Casas, D., et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.

Deng, B., Zhong, P., Jun, K., Riebesell, J., Han, K., Bartel, C. J., and Ceder, G. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. 5(9):1031–1041, 2023. ISSN 2522-5839. doi: 10.1038/s42256-023-00716-3. URL https://www.nature.com/articles/s42256-023-00716-3. Number: 9 Publisher: Nature Publishing Group.

Di Siena, A., Navarro, A. B., Luda, T., Merlo, G., Bergmann, M., Leppin, L., Görler, T., Parker, J., LoDestro, L., Dannert, T., et al. Global gyrokinetic simulations of asdex upgrade up to the transport timescale with gene–tango. *Nuclear Fusion*, 62(10):106025, 2022.

Djeumou, F., Neary, C., and Topcu, U. How to learn and generalize from three minutes of data: Physics-constrained and uncertainty-aware neural stochastic differential equations. *arXiv preprint arXiv:2306.06335*, 2023.

D.O.E. Public reusable research (pure) data. https://shorturl.at/Dh0WE, 2021.

DOE, 2023. De-foa-0002905 machine learning, artificial intelligence, and data resources for fusion energy sciences, 2023.

Dubčáková, R. Eureqa: software review, 2011.

Duval, A., Mathis, S. V., Joshi, C. K., Schmidt, V., Miret, S., Malliaros, F. D., Cohen, T., Liò, P., Bengio, Y., and Bronstein, M. A hitchhiker's guide to geometric gnns for 3d atomic systems. *arXiv preprint arXiv:2312.07511*, 2023a.

Duval, A. A., Schmidt, V., Hernández-Garcıa, A., Miret, S., Malliaros, F. D., Bengio, Y., and Rolnick, D. Faenet: Frame averaging equivariant gnn for materials modeling. In *International Conference on Machine Learning*, pp. 9013–9033. PMLR, 2023b.

Felici, F., Citrin, J., Teplukhina, A., Redondo, J., Bourdelle, C., Imbeaux, F., Sauter, O., Contributors, J., Team, E. M., et al. Real-time-capable prediction of temperature and density profiles in a tokamak using raptor and a first-principle-based transport model. *Nuclear Fusion*, 58(9): 096006, 2018.

Foster, A., Jankowiak, M., Bingham, E., Horsfall, P., Teh, Y. W., Rainforth, T., and Goodman, N. Variational Bayesian optimal experimental design. *Advances in Neural Information Processing Systems*, 32, 2019.

Frazier, P. I. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Freidberg, J. P. *Plasma physics and fusion energy*. Cambridge university press, 2008.

Fu, C., Chen, P., Shen, Y., Qin, Y., Zhang, M., Lin, X., Qiu, Z., Lin, W., Yang, J., Zheng, X., Li, K., Sun, X., and Ji, R. Mme: A comprehensive evaluation benchmark for multimodal large language models. *ArXiv*, abs/2306.13394, 2023.

Gadre, S. Y., Ilharco, G., Fang, A., Hayase, J., Smyrnis, G., Nguyen, T., Marten, R., Wortsman, M., Ghosh, D., Zhang, J., Orgad, E., Entezari, R., Daras, G., Pratt, S., Ramanujan, V., Bitton, Y., Marathe, K., Mussmann, S., Vencu, R., Cherti, M., Krishna, R., Koh, P. W., Saukh, O., Ratner, A. J., Song, S., Hajishirzi, H., Farhadi, A., Beaumont, R., Oh, S., Dimakis, A. G., Jitsev, J., Carmon, Y., Shankar, V., and Schmidt, L. Datacomp: In search of the next generation of multimodal datasets. *ArXiv*, abs/2304.14108, 2023.

Gao, J., Li, P., Chen, Z., and Zhang, J. A survey on deep learning for multimodal data fusion. *Neural Computation*, 32(5):829–864, 2020.

Ge, W., Chen, S., Chen, G., Chen, J., Chen, Z., Yan, S., Zhu, C., Lin, Z., Xie, W., Wang, X., Gao, A., Zhang, Z., Li, J., Wan, X., and Wang, B. Mllm-bench, evaluating multi-modal llms using gpt-4v. *ArXiv*, abs/2311.13951, 2023.

Gevers, M. and Ljung, L. Optimal experiment designs with respect to the intended model application. *Automatica*, 22(5):543–554, 1986.

Ginsberg, M. *Essentials of artificial intelligence*. Newnes, 2012.

Ginsberg, 2021. Conversation with matt ginsberg. Personal communication, 2021.

Gunn, A. How does nuclear fusion work? *BBC Science Focus*, 2022.

Hamedani, A., Byggmästar, J., Djurabekova, F., Alahyarizadeh, G., Ghaderi, R., Minuchehr, A., and Nordlund, K. Primary radiation damage in silicon from the viewpoint of a machine learning interatomic potential. 5 (11):114603, 2021. doi: 10.1103/PhysRevMaterials.5.114603. URL https://link.aps.org/doi/10.1103/PhysRevMaterials.5.114603.

Harms, A. A., Schoepf, K. F., and Kingdon, D. R. *Principles of fusion energy: an introduction to fusion energy for students of science and engineering*. World Scientific, 2000.

Hawryluk, R. and Coppi, B. Physics of plasmas close to thermonuclear conditions. *Proceedings of the course held in Varenna, Italy*, 1, 1981.

Hayashi, Y., Kobayashi, M., Mukai, K., Masuzaki, S., and Murase, T. Divertor heat load distribution measurements with infrared thermography in the lhd helical divertor. *Fusion Engineering and Design*, 165:112235, 2021. ISSN 0920-3796. doi: https://doi.org/10.1016/j.fusengdes.2021.112235. URL https://www.sciencedirect.com/science/article/pii/S0920379621000119.

He, K., Chen, X., Xie, S., Li, Y., Doll'ar, P., and Girshick, R. B. Masked autoencoders are scalable vision learners. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15979–15988, 2021.

Hedelius, B. E., Tingey, D., and Della Corte, D. TrIP-transformer interatomic potential predicts realistic energy surface using physical bias. 20(1):199–211, 2024. ISSN 1549-9618. doi: 10.1021/acs.jctc.3c00936. URL https://doi.org/10.1021/acs.jctc.3c00936. Publisher: American Chemical Society.

Hennig, P. and Schuler, C. J. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(6), 2012.

Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. *Advances in neural information processing systems*, 27, 2014.

Hoelzl, M., Huijsmans, G., Pamela, S. J. P., Becoulet, M., Nardon, E., Artola, F. J., Nkonga, B., Atanasiu, C. V., Bandaru, V., Bhole, A., Bonfiglio, D., Cathey, A., Czarny, O., Dvornova, A., Fehér, T., Fil, A., Franck, E., Futatani, S., Gruca, M., Guillard, H., Haverkort, J., Holod, I., Hu, D., Kim, S., Korving, S. Q., Kos, L., Krebs, I., Kripner, L., Latu, G., Liu, F., Merkel, P., Meshcheriakov, D., Mitterauer, V., Mochalskyy, S., Morales, J., Nies, R., Nikulsin, N., Orain, F., Penko, D., Pratt, J., Ramasamy, R., Ramet, P., Reux, C., Schwarz, N., Verma, P. S., Smith, S., Sommariva, C., Strumberger, E., vanVugt, D., Verbeek, M. G., Westerhof, E., Wieschollek, F., and Zielinski, J. The jorek non-linear extended mhd code and applications to large-scale instabilities and their control in magnetically confined fusion plasmas. *Nuclear Fusion*, 61, 2020.

Huang, B. and Wang, J. Applications of physics-informed neural networks in power systems - a review. *IEEE Transactions on Power Systems*, 38:572–588, 2023.

Humphreys, D. A., Kupresanin, A., Boyer, D., Canik, J. M., Cyr, E. C., Granetz, R., Hittinger, J. A. F., Kolemen, E., Lawrence, E. C., and Pascucci, V. Advancing fusion with machine learning research needs workshop report. *Journal of Fusion Energy*, 39:123 – 155, 2019.

Hwangbo, J., Lee, J., Dosovitskiy, A., Bellicoso, D., Tsounis, V., Koltun, V., and Hutter, M. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019.

Igochine, V. Investigation of mhd instabilities in conventional and advanced tokamak scenarios on asdex upgrade [elektronische ressource] /.

Ikeda, K. Progress in the iter physics basis. *Nuclear Fusion*, 47(6):E01, 2007.

IPP, 2017. The effect of nitrogen impurity seeding in the compass tokamak. URL https://youtu.be/DAssKNprOGA?si=iqdzEQ9VrU0YrNry.

Jang, D., Spangher, L., Khattar, M., Agwan, U., Nadarajah, S., and Spanos, C. J. Offline-online reinforcement learning for energy pricing in office demand response: lowering energy and data costs. *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 2021. doi: 10.1145/3486611.3486668.

Jang, D., Yan, L., Spangher, L., and Spanos, C. J. Active reinforcement learning for robust building control. *ArXiv*, abs/2312.10289, 2023. doi: 10.48550/arXiv.2312.10289.

Jin, M., Cao, P., and Short, M. P. Thermodynamic mixing energy and heterogeneous diffusion uncover the mechanisms of radiation damage reduction in single-phase ni-fe alloys. 147:16–23, 2018. ISSN 1359-6454. doi: 10.1016/j.actamat.2017.12.064. URL https://www.sciencedirect.com/science/article/pii/S135964541830020X.

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

Kamal, A. AI and machine learning in nuclear fusion. 2020.

Kates-Harbeck, J., Svyatkovskiy, A., and Tang, W. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, 568(7753):526–531, 2019.

Keith, Z., Nagpal, C., Rea, C., and Tinguely, R. Risk-Aware Framework Development for Disruption Prediction: Alcator C-Mod and DIII-D Survival Analysis. submitted.

Kidger, P. On neural differential equations. *arXiv preprint arXiv:2202.02435*, 2022.

Ko, T. W. and Ong, S. P. Recent advances and outstanding challenges for machine learning interatomic potentials. 3 (12):998–1000, 2023. ISSN 2662-8457. doi: 10.1038/s43588-023-00561-9. URL https://www.nature.com/articles/s43588-023-00561-9.

Kolemen, E., Gates, D., Rowley, C. W., Kasdin, N. J., Kallman, J., Gerhardt, S., Soukhanovskii, V., and Mueller, D. Strike point control for the national spherical torus experiment (nstx). *Nuclear fusion*, 50(10):105010, 2010.

Kolemen, E., Seo, J., Conlin, R., Rothstein, A., Kim, S., Abbate, J., Erickson, K., Wai, J., Shousha, R., and Jalalvand, A. Avoiding tokamak tearing instability with artificial intelligence. 2023.

Kruger, S., Lao, L., Akcay, C., Bechtel, T., Liu, Y., McClenaghan, J., Orozco, D., Schissel, D., Smith, S., Sun, X., et al. Overview and progress of the EFIT-AI project. *Bulletin of the American Physical Society*, 2022.

Kulcinski, G. L., Kessler, G., Holdren, J., and Häfele, W. Energy for the long run: Fission or fusion? factors such as hazards, technological costs, and development time are the significant points of comparison for the two most likely long-term energy sources. *American Scientist*, 67 (1):78–89, 1979. ISSN 00030996. URL http://www.jstor.org/stable/27849063.

Kwak, J.-G., Oh, Y., Yang, H., Park, K., Kim, Y., Kim, W., Kim, J., Lee, S., Na, H., Kwon, M., et al. An overview of kstar results. *Nuclear Fusion*, 53(10):104005, 2013.

Kwak, S., Hergenhahn, U., Höfel, U., Krychowiak, M., Pavone, A., Svensson, J., Ford, O. P., König, R., Bozhenkov, S., Fuchert, G., Pasch, E., Brunner, K. J., Knauer, J., Kornejew, P., Mora, H. T., Pedersen, T. S., and Team, W. X. Bayesian inference of spatially resolved zeff profiles from line integrated bremsstrahlung spectra. *The Review of scientific instruments*, 92 4:043505, 2021.

Lahat, D., Adalı, T., and Jutten, C. Multimodal data fusion: An overview of methods, challenges, and prospects. *Proceedings of the IEEE*, 103:1449–1477, 2015.

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.

Lane-Walsh, S., Stillerman, J., Santoro, F., and Fredian, T. W. Introduction to mdsplus using docker. *Fusion Engineering and Design*, 165:112121, 2021.

Lao, L., John, H. S., Stambaugh, R., Kellman, A., and Pfeiffer, W. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear fusion*, 25(11): 1611, 1985.

Lao, L., Kruger, S., Akcay, C., Balaprakash, P., Bechtel, T., Howell, E., Koo, J., Leddy, J., Leinhauser, M., Liu, Y., et al. Application of machine learning and artificial intelligence to extend efit equilibrium reconstruction. *Plasma Physics and Controlled Fusion*, 64(7):074001, 2022.

Lee, G., Kwon, M., Doh, C., Hong, B., Kim, K., Cho, M., Namkung, W., Chang, C.-S., Kim, Y., Kim, J., et al. Design and construction of the kstar tokamak. *Nuclear Fusion*, 41(10):1515, 2001.

Lee, K. L. K., Gonzales, C., Nassar, M., Spellings, M., Galkin, M., and Miret, S. Matsciml: A broad, multi-task benchmark for solid-state materials modeling. *arXiv preprint arXiv:2309.05934*, 2023.

Lehnen, M., Aleynikov, P., and Bazylev, B. Plasma disruption management in iter. *2016 IAEA Fusion Energy Conference*, 2016.

Lei, Y., Jia, F., Lin, J., Xing, S., and Ding, S. X. An intelligent fault diagnosis method using unsupervised feature learning towards mechanical big data. *IEEE Transactions on Industrial Electronics*, 63:3137–3147, 2016.

Leonard, A. Plasma detachment in divertor tokamaks. *Plasma Physics and Controlled Fusion*, 60(4):044001, 2018.

Li, D., Svensson, J., Thomsen, H., Medina, F., Werner, A., and Wolf, R. Bayesian soft x-ray tomography using non-stationary gaussian processes. *Review of Scientific Instruments*, 84(8), 2013.

Liang, P. P., Lyu, Y., Fan, X., Wu, Z., Cheng, Y., Wu, J., Chen, L., Wu, P., Lee, M. A., Zhu, Y., Salakhutdinov, R., and Morency, L.-P. Multibench: Multiscale benchmarks for multimodal representation learning. *ArXiv*, abs/2107.07502, 2021.

Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.

Maris, A. D., Wang, A., Rea, C., Granetz, R., and Marmar, E. The impact of disruptions on the economics of a tokamak power plant. *Fusion Science and Technology*, pp. 1–17, 2023.

Marmar, E. Alcator c-mod, final technical report. Technical report, Massachusetts Inst. of Technology (MIT), Cambridge, MA (United States), 2018.

Martin, Y., Takizuka, T., et al. Power requirement for accessing the h-mode in iter. In *Journal of Physics: Conference Series*, volume 123, pp. 012033. IOP Publishing, 2008.

Mathews, A. Physics-informed machine learning techniques for edge plasma turbulence modelling in computational theory and experiment. *ArXiv*, abs/2205.07838, 2022.

Mathews, A., Francisquez, M., Hughes, J., Hatch, D., Zhu, B., and Rogers, B. Uncovering turbulent plasma dynamics via deep learning from partial observations. *Physical review. E*, 104 2-2:025205, 2020a.

Mathews, A., Francisquez, M., Hughes, J., and Hatch, D. R. Uncovering edge plasma dynamics via deep learning from partial observations. *arXiv: Plasma Physics*, 2020b.

Matteis, T. D., de Fine Licht, J., Beránek, J., and Hoefler, T. Streaming message interface: high-performance distributed memory programming on reconfigurable hardware. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019.

Mehta, V., Paria, B., Schneider, J., Ermon, S., and Neiswanger, W. An experimental design perspective on model-based reinforcement learning. *arXiv preprint arXiv:2112.05244*, 2021.

Mehta, V., Abbate, J., Wang, A., Rothstein, A., Char, I., Schneider, J., Kolemen, E., Rea, C., and Garnier, D. Towards llms as operational copilots for fusion reactors. In *NeurIPS 2023 AI for Science Workshop*, 2023.

Mehta, V., Barr, J. L., Abbate, J., Boyer, M. D., Char, I., Neiswanger, W., Kolemen, E., and Schneider, J. Automated experimental design of safe rampdowns via probabilistic machine learning. *Nuclear Fusion*, 2024. URL http://iopscience.iop.org/article/10.1088/1741-4326/ad22f5.

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023.

Montes, K., Rea, C., Tinguely, R. A., Sweeney, R., Zhu, J., and Granetz, R. A semi-supervised machine learning detector for physics events in tokamak discharges. *Nuclear Fusion*, 61, 2020.

Montes, K. J., Rea, C., Granetz, R., Tinguely, R. A., Eidietis, N., Meneghini, O., Chen, D., Shen, B., Xiao, B., Erickson, K., et al. Machine learning for disruption warnings on alcator c-mod, diii-d, and east. *Nuclear Fusion*, 59(9): 096015, 2019.

Moret, J.-M., Duval, B., Le, H., Coda, S., Felici, F., and Reimerdes, H. Tokamak equilibrium reconstruction code liuqe and its real time implementation. *Fusion Engineering and Design*, 91:1–15, 2015.

Morgan, D., Pilania, G., Couet, A., Uberuaga, B. P., Sun, C., and Li, J. Machine learning in nuclear materials research. 26(2):100975, 2022. ISSN 1359-0286. doi: 10.1016/j.cossms.2021.100975. URL https://www.sciencedirect.com/science/article/pii/S1359028621000784.

Murari, A., Vagliasindi, G., Arena, P., Fortuna, L., Barana, O., Johnson, M., Contributors, J.-E., et al. Prototype of an adaptive disruption predictor for jet based on fuzzy logic and regression trees. *Nuclear Fusion*, 48(3):035010, 2008.

Murari, A., Vega, J., Rattá, G., Vagliasindi, G., Johnson, M., Hong, S., Contributors, J.-E., et al. Unbiased and non-supervised learning methods for disruption prediction at jet. *Nuclear Fusion*, 49(5):055028, 2009.

Murari, A., Lungaroni, M., Peluso, E., Gaudio, P., Vega, J., Dormido-Canto, S., Baruzzo, M., Gelfusa, M., and Contributors, J. Adaptive predictors based on probabilistic svm for real time disruption mitigation on jet. *Nuclear Fusion*, 58(5):056002, 2018.

Murari, A., Peluso, E., Craciunescu, T., Dormido-Canto, S., Lungaroni, M., Rossi, R., Spolladore, L., Vega, J., Gelfusa, M., and Contributors, J. Frontiers in data analysis methods: from causality detection to data driven experimental design. *Plasma Physics and Controlled Fusion*, 64(2):024002, dec 2021. doi: 10.1088/1361-6587/ac3ded. URL https://dx.doi.org/10.1088/1361-6587/ac3ded.

Musaelian, A., Batzner, S., Johansson, A., Sun, L., Owen, C. J., Kornbluth, M., and Kozinsky, B. Learning local equivariant representations for large-scale atomistic dynamics, 2022. URL http://arxiv.org/abs/2204.05249.

Parisi, J. and Ball, J. *The future of fusion energy*. World Scientific, 2019.

Pastore, F., Felici, F., Bosman, T., Galperti, C., Sauter, O., Vincent, B., Vu, N., et al. Model-based electron density estimation using multiple diagnostics on tcv. *Fusion Engineering and Design*, 192:113615, 2023.

Pau, A., Fanni, A., Carcangiu, S., Cannas, B., Sias, G., Murari, A., Rimini, F., et al. A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at jet. *Nuclear Fusion*, 59(10): 106017, 2019.

Pavone, A., Merlo, A., Kwak, S., and Svensson, J. Machine Learning and Bayesian inference in nuclear fusion research: an overview. *Plasma Physics and Controlled Fusion*, 65, 2023.

Petrov, A. N. and Wessling, J. M. Utilization of machine-learning algorithms for wind turbine site suitability modeling in iowa, usa. *Wind Energy*, 18(4):713–727, 2015.

Qi, J., Ko, T. W., Wood, B. C., Pham, T. A., and Ong, S. P. Robust training of machine learning interatomic potentials with dimensionality reduction and stratified sampling, 2023. URL http://arxiv.org/abs/2307.13710.

Rasul, K., Seward, C., Schuster, I., and Vollgraf, R. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*, 2021.

Rattá, G., Vega, J., Murari, A., Vagliasindi, G., Johnson, M., De Vries, P., and Contributors, J. E. An advanced disruption predictor for jet tested in a simulated real-time environment. *Nuclear Fusion*, 50(2):025005, 2010.

Rattá, G., Vega, J., Murari, A., and Contributors, J. E. Simulation and real-time replacement of missing plasma signals for disruption prediction: an implementation with apodis. *Plasma Physics and Controlled Fusion*, 56(11): 114004, 2014.

Rea, C. and Granetz, R. S. Exploratory machine learning studies for disruption prediction using large databases on diii-d. *Fusion Science and Technology*, 74(1-2):89–100, 2018.

Rea, C., Granetz, R., Montes, K., Tinguely, R. A., Eidietis, N., Hanson, J. M., and Sammuli, B. Disruption prediction investigations using machine learning tools on diii-d and alcator c-mod. *Plasma Physics and Controlled Fusion*, 60(8):084004, 2018.

Rea, C., Montes, K., Erickson, K., Granetz, R., and Tinguely, R. A real-time machine learning-based disruption predictor in diii-d. *Nuclear Fusion*, 59(9):096016, 2019.

Ren, Z., Ren, Z., Zhang, Z., Buonassisi, T., and Li, J. Autonomous experiments using active learning and AI. 8 (9):563–564, 2023a. ISSN 2058-8437. doi: 10.1038/s41578-023-00588-4. URL https://www.nature.com/articles/s41578-023-00588-4.

Ren, Z., Zhang, Z., Tian, Y., and Li, J. Accelerated high-entropy alloys discovery for electrocatalysis via robotic-aided active learning. 2023b. URL https://openreview.net/forum?id=YRxd1szajS.

Rodriguez-Fernandez, P., Howard, N., Greenwald, M., Creely, A., Hughes, J., Wright, J., Holland, C., Lin, Y., and Sciortino, F. Predictions of core plasma performance for the sparc tokamak. *Journal of Plasma Physics*, 86, 10 2020. doi: 10.1017/S0022377820001075.

Rodriguez-Fernandez, P., Creely, A., Greenwald, M., Brunner, D., Ballinger, S., Chrobak, C., Garnier, D., Granetz, R., Hartwig, Z., Howard, N., et al. Overview of the sparc physics basis towards the exploration of burning-plasma regimes in high-field, compact tokamaks. *Nuclear Fusion*, 62(4):042003, 2022a.

Rodriguez-Fernandez, P., Howard, N., and Candy, J. Nonlinear gyrokinetic predictions of sparc burning plasma profiles enabled by surrogate modeling. *Nuclear Fusion*, 62(7):076036, 2022b.

Rodríguez Fernández, P. et al. *Perturbative transport experiments and time-dependent modeling in Alcator C-Mod and DIII-D*. PhD thesis, Massachusetts Institute of Technology, 2019.

Romanelli, M., Coelho, R., Coster, D. P., Ferreira, J., Fleury, L., Henderson, S. S., Hollocombe, J., Imbeaux, F., Jonsson, T., Kogan, L. A., Meneghini, O. M., Merle, A., Pinches, S. D., Sauter, O., Tardini, G., Yadykin, D., Smith, S. P., Strand, P., and Team, W. Code integration, data verification, and models validation using the iter integrated modeling and analysis system (imas) in eurofusion. *Fusion Science and Technology*, 76:894 – 900, 2020.

Sabbagh, S., Berkery, J., Park, Y., Butt, J., Riquezes, J., Bak, J., Bell, R., Delgado-Aparicio, L., Gerhardt, S., Ham, C., et al. Disruption event characterization and forecasting in tokamaks. *Physics of Plasmas*, 30(3), 2023.

Seo, J., Na, Y.-S., Kim, B., Lee, C., Park, M., Park, S., and Lee, Y. Feedforward beta control in the kstar tokamak by deep reinforcement learning. *Nuclear Fusion*, 61, 2021.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.

Shimada, M. and Pitts, R. A. Wall conditioning on iter. *Journal of Nuclear Materials*, 415(1):S1013–S1016, 2011.

Shousha, R., Seo, J., Erickson, K., Xing, Z., Kim, S., Abbate, J., and Kolemen, E. Machine learning-based real-time kinetic profile reconstruction in diii-d. *Nuclear Fusion*, 64(2):026006, 2023.

Fusion Industry Association. The global fusion industy in 2023. Technical report, 07 2023.

IAEA. Artificial intelligence for accelerating nuclear applications, science, and technology. Technical report, 2022.

ITER Team. ITER Scientific Data Centre | How to manage 2 petabytes of new data every day. *ITER*, November 2020. URL http://www.iter.org/newsline/-/3534.

National Science Foundation. Nsf 22-553: Findable accessible interoperable reusable open science research coordination networks, 2022.

Spangher, L. personal correspondence with Stephen Lane-Walsh, 2024.

Spangher, L., Vitter, J. S., and Umstattd, R. J. Characterizing fusion market entry via an agent-based power plant fleet model. *Energy Strategy Reviews*, 2019. doi: 10.1016/j.esr.2019.100404.

Spangher, L., Gokul, A., Khattar, M., Palakapilly, J., Agwan, U., Tawade, A., and Spanos, C. J. Augmenting reinforcement learning with a planning model for optimizing energy demand response. *Proceedings of the 1st International Workshop on Reinforcement Learning for Energy Management in Buildings & Cities*, 2020. doi: 10.1145/3427773.3427863.

Spangher, L., Arnold, W., Spangher, A., Maris, A., and Rea, C. Autoregressive transformers for disruption prediction in nuclear fusion plasmas. 2023. doi: 10.48550/arXiv.2401.00051.

Stillerman, J., Fredian, T. W., Klare, K. A., and Manduchi, G. Mdsplus data acquisition system. *Review of Scientific Instruments*, 68:939–942, 1997.

Sun, Y., Zhu, D., Li, Y., Wang, R., and Ma, R. Spatial modelling the location choice of large-scale solar photovoltaic power plants: Application of interpretable machine learning techniques and the national inventory. *Energy Conversion and Management*, 289:117198, 2023.

Szymanski, N. J., Rendy, B., Fei, Y., Kumar, R. E., He, T., Milsted, D., McDermott, M. J., Gallant, M., Cubuk, E. D., Merchant, A., Kim, H., Jain, A., Bartel, C. J., Persson, K., Zeng, Y., and Ceder, G. An autonomous laboratory for the accelerated synthesis of novel materials. 624 (7990):86–91, 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06734-w. URL https://www.nature.com/articles/s41586-023-06734-w. Number: 7990 Publisher: Nature Publishing Group.

Taylor, T. Physics of advanced tokamaks. *Plasma Physics and Controlled Fusion*, 39(12B):B47, 1997.

Thölke, P. and De Fabritiis, G. Torchmd-net: Equivariant transformers for neural network based molecular potentials. *arXiv preprint arXiv:2202.02541*, 2022.

Tinguely, R., Montes, K., Rea, C., Sweeney, R., and Granetz, R. An application of survival analysis to disruption prediction via random forests. *Plasma Physics and Controlled Fusion*, 61(9):095009, 2019.

Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033. IEEE, 2012.

Tracey, B. D., Michi, A., Chervonyi, Y., Davies, I., Paduraru, C., Lazic, N., Felici, F., Ewalds, T., Donner, C., Galperti, C., et al. Towards practical reinforcement learning for tokamak magnetic control. *Fusion Engineering and Design*, 200:114161, 2024.

Udrescu, S.-M. and Tegmark, M. AI feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020.

van de Plassche, K. L., Citrin, J., Bourdelle, C., Camenen, Y., Casson, F. J., Dagnelie, V. I., Felici, F., Ho, A., Van Mulders, S., and Contributors, J. Fast modeling of turbulent transport in fusion plasmas using neural networks. *Physics of Plasmas*, 27(2), 2020.

Vecchio, K. S., Dippo, O. F., Kaufmann, K. R., and Liu, X. High-throughput rapid experimental alloy development (HT-READ). 221:117352, 2021. ISSN 1359-6454. doi: 10.1016/j.actamat.2021.117352. URL https://www.sciencedirect.com/science/article/pii/S135964542100731X.

Vega, J., Dormido-Canto, S., López, J. M., Murari, A., Ramírez, J. M., Moreno, R., Ruiz, M., Alves, D., Felton, R., Contributors, J.-E., et al. Results of the jet real-time disruption predictor in the iter-like wall campaigns. *Fusion Engineering and Design*, 88(6-8):1228–1231, 2013.

Vu, T., Felici, F., Galperti, C., Maraschek, M., Pau, A., Rispoli, N., Sauter, O., and Sieglin, B. Integrated real-time supervisory management for off-normal-event handling and feedback control of tokamak plasmas. *IEEE Transactions on Nuclear Science*, 68(8):1855–1861, 2021.

Walker, M. L., De Vries, P., Felici, F., and Schuster, E. Introduction to tokamak plasma control. In *2020 American Control Conference (ACC)*, pp. 2901–2918. IEEE, 2020.

Wan, C., Yu, Z., Pau, A., Liu, X., and Li, J. A machine-learning-based tool for last closed magnetic flux surface reconstruction on tokamak. *ArXiv*, abs/2207.05695, 2022.

Wang, A. M., Garnier, D. T., and Rea, C. Hybridizing physics and neural odes for predicting plasma inductance dynamics in tokamak fusion reactors. *arXiv preprint arXiv:2310.20079*, 2023.

Wei, Y., Levesque, J., Hansen, C., Mauel, M., and Navratil, G. A dimensionality reduction algorithm for mapping tokamak operational regimes using a variational autoencoder (vae) neural network. *Nuclear Fusion*, 61(12):126063, 2021.

Wesson, J. and Campbell, D. J. *Tokamaks*, volume 149. Oxford university press, 2011.

Windsor, C., Pautasso, G., Tichmann, C., Buttery, R., Hender, T., Contributors, J. E., Team, A. U., et al. A cross-tokamak neural network disruption predictor for the jet and asdex upgrade tokamaks. *Nuclear fusion*, 45(5):337, 2005.

Wirth, B. FINAL REPORT: DE-SC0006661 a multiscale investigation of the mechanisms controlling materials degradation in the fusion energy environment, 2023. URL https://www.osti.gov/biblio/2221775.

Wu, S., team, E., et al. An overview of the east project. *Fusion Engineering and Design*, 82(5-14):463–471, 2007.

Wurzel, S. E. and Hsu, S. C. Progress toward fusion energy breakeven and gain as measured against the lawson criterion. *Physics of Plasmas*, 29(6), 2022.

Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., and Chen, E. A survey on multimodal large language models. *ArXiv*, abs/2306.13549, 2023.

Yip, S. and Short, M. P. Multiscale materials modelling at the mesoscale. 12(9):774–777, 2013. ISSN 1476-4660. doi: 10.1038/nmat3746. URL https://www.nature.com/articles/nmat3746. Number: 9 Publisher: Nature Publishing Group.

Zhang, H., Yu, Z., Dai, G., Huang, G., Ding, Y., Xie, Y., and Wang, Y. Understanding GNN computational graph: A coordinated computation, IO, and memory perspective, 2021. URL http://arxiv.org/abs/2110.09524.

Zhong, Y., Zheng, W., Chen, Z., Xia, F., Yu, L., Wu, Q., Ai, X., Shen, C., Yang, Z., Yan, W., et al. Disruption prediction and model analysis using lightgbm on j-text and hl-2a. *Plasma Physics and Controlled Fusion*, 63(7): 075008, 2021.

Zhu, J., Rea, C., Granetz, R., Marmar, E., Montes, K., Sweeney, R., Tinguely, R., Chen, D., Shen, B., Xiao, B., et al. Scenario adaptive disruption prediction study for next generation burning-plasma tokamaks. *Nuclear Fusion*, 61(11):114005, 2021.

A. Common acronyms and tokamak names

Acronym	Definition
CFS	Commonwealth Fusion Systems
CRP	Coordinated research project
DECAF	Disruption Event Characterization And Forecasting
DFT	Density Functional Theory
EAM	Embedded Atom Model
GAP	Gaussian approximation potentials
GNN	Graph neural network
GP	Gaussian processes
IAEA	International Atomic Energy Agency
IMAS	Integrated Modeling and Analysis System
LLM	Large language model
LSTM	Long short-term memory network
MARS	Multi-adaptive auto-regressive splines
MD	Molecular dynamics
ML	Machine Learning
MLIP	Machine learning interatomic potentials
MPNN	Message passing neural networks
PDE	Partial differential equations
PINN	Physics-informed neural network
RL	Reinforcement learning
RZIP	Control of plasma radial (r) position, vertical (z) position, and current (I_p)
S4	Structured state space for sequence modeling
TRIP	Transformer Interatomic Potential
VAE	Variational Autoencoders

Table 1. Common acronyms used in this paper.

B. Fusion Physics

This appendix section provides a brief introduction to fusion physics. We refer the curious reader to (Freidberg, 2008) for a deeper introduction targeted towards a first year graduate student audience. For an accessible introduction targeted towards a general audience, we recommend (Parisi & Ball, 2019).

Fusion occurs when light nuclei, such as hydrogen isotopes, get close enough that the attractive strong force overcomes the repulsive electric force. The fusion of the nuclei produces an enormous amount of energy per unit mass; for reference, a typical American's lifetime energy consumption could be fueled by the deuterium (hydrogen isotope with one neutron) in a bathtub of water and the lithium content of two lithium-ion laptop batteries.

A number of fuel combinations can be utilized. However, deuterium - tritium fusion, involving two isotopes of Hydrogen, has by far the easiest physical requirements for achieving energy gain. To achieve net energy gain, the

Tokamak	Description
Alcator C-mod	High magnetic-field tokamak decommissioned in 2016, formerly operated by MIT (Marmar, 2018)
DIII-D	Largest tokamak in the United States, operated by General Atomics (Buttery et al., 2019)
EAST	Superconducting tokamak, operated by the Chinese Academy of Sciences Institute of Plasma Physics (Wu et al., 2007)
ITER	Burning plasma tokamak experiment, under construction by international consortium (Ikeda, 2007)
KSTAR	Superconducting tokamak, operated by the Korea Institute of Fusion Energy (Kwak et al., 2013)
SPARC	Compact, high magnetic field tokamak, under construction by CFS (Rodriguez-Fernandez et al., 2022a)
TCV	Small tokamak with flexible shape control, operated by the Swiss Plasma Center (Coda et al., 2017)

Table 2. Tokamaks referenced in this paper.

triple product of density n, temperature T, and confinement time τ, $nT\tau$, must exceed a certain threshold. While more advanced fusion reactions such as deuterium-deuterium, proton-boron-11, and helium-3 fusion have potential advantages over deuterium-tritium, they require substantially higher values of $nT\tau$.

Achieving break-even conditions with deuterium-tritium fusion requires plasmas ten times hotter than the core of the sun, requiring creative solutions to confine and control the plasma. Many approaches to fusion have been conceptualized and attempted, with most falling massively short of break-even conditions. Only laser-based inertial confinement fusion (ICF) and the tokamak approach to magnetic confinement fusion (MCF) have approached break-even conditions. Recent work systematically documented the performance of a wide array of fusion concepts attempted (Wurzel & Hsu, 2022).

While many important phenomena of plasma physics can be understood quite readily from first principles, many other phenomena are emergent from the complex interaction of many particles. For example, the dominant mode of heat loss from the plasma is due to micro-turbulence which emerges from a kinetic description of plasmas, which is extremely computationally intensive (Rodríguez Fernández et al., 2019). The extreme computational requirements of a kinetic treatment, coupled with an extreme separation of scales in both time and space necessitating a hierarchy of models from quantum mechanics to single particle to kinetic, to multi-fluid to single fluid to empirical. As noted in Sec 2.2.3, even today, a number of important plasma phenomena are still not well-understood from first principles. All of this makes the problem of plasma modelling and simulation extremely challenging.

C. Extended Literature reviews

C.1. Equilibrium reconstruction

The fusion community has spent a large effort in reconstructing equilibrium profiles, yet there is still room for improvement. For example, Svennson uses Gaussian Process emulation for tomography (Li et al., 2013), and the EFIT-AI team at the DIII-D tokamak has published notably on an ML-enhanced Bayesian 2D magnetic profile framework, a Model-Order-Reduction version of the two-dimensional Grad-Shafranov equation solver using a neural network, and a three-dimensional perturbed equilibrium reconstruction solver. (Lao et al., 2022; Kruger et al., 2022).

C.2. Machine Learning for Material Design

Graph Neural Networks (GNN), Equivariant Message Passing Networks (MPNN), and more recently Transformers have all been successfully trained from ab initio Density Functional Theory data (Chen & Ong, 2022; Batzner et al., 2022; Hedelius et al., 2024). Graph Neural Networks have been successful at predicting properties of materials, with DFT level accuracy, using tens to hundreds of atoms. Unfortunately, due to the $O(N^2)$ scaling of edges with numbers of atoms, using GNNs on a system with hundreds of thousand or millions of atoms is intractable (Zhang et al., 2021). Only the Equivariant Message Passing Network, Allegro, a speed optimized variant of Nequip has achieved scaling up to 421,824 atoms, but required 64 V100 GPUs to generate 10 ns/day in LAMMPS (Musaelian et al., 2022). For comparison, using a simple Embedded Atom Model Potential for tungsten on a MacBook Pro can also achieve 10 ns/day. On the other spectrum of MLIPs, tabulated Gaussian Approximation Potentials (tabGAP) have shown to scale to over 250,000 atoms, and scale similarly to embedded atom model (EAM) potentials (Byggmästar et al., 2021). Unfortunately the tabGAP potential struggles to handle more than 4 atoms, and long-range effects like defects and dislocations are not properly represented. The Transformer interatomic potential (TRIP) also showed improved accuracy compared to MPNNs, but their scalability was not demonstrated past 312 atoms (Hedelius et al., 2024). The bottom line is while GNNs can most accurately capture complex, long-range interactions for 1000s of atoms, they struggle to scale. MPNNs and TRIPs have the potential to scale, but still are expensive to run in MD simulations, and are about a factor of 5-10 less accurate compared to GNNs

(TRIPs are more accurate than MPNNs like MACE and NeQUIP, but their speed have not been validated). tabGAP potentials can easily scale past hundreds of thousands of atoms, but only remain expressive for simple systems, and struggle to handle long range effects. There is a need for an improved representation of atoms allowing for DFT level accuracy for hundreds of thousands to millions of atoms, in 1000x less time.

Automated labs with active learning to develop nuclear structural materials have also been demonstrated, but the composition control (i.e $\pm 15\%$ errors in the composition, which is about 2 orders of magnitude worse than humans) of the additive-manufacturing based sample synthesis system was 1-2 orders of magnitude off what is needed for a fusion relevant alloy (Couet, 2022; Vecchio et al., 2021). Additionally, other efforts to setup automated labs have struggled with the repeatability and reliability of sample handling and characterization (Szymanski et al., 2023).

D. An extended discussion on realtime fusion device computation

Computer systems for fusion devices are optimized for managing controls and diagnostics, and have changed substantially throughout the years. Throughout Alcator C-Mod's operational lifetime, there was a single computer without an OS, running just a single program. In order to enable speed of reading, there was no network to act as a middle message passer to translate network readings, and thus data transfer was as fast as the CPU possibly could transfer it.

It is very likely that current and future fusion devices, like South Korea's KSTAR, are no different (Lee et al., 2001). Using a network in passing data between device and computer is new and relatively unproven, and there is reason to think that it could be bad: latency is on the order of milliseconds in network message passing, whereas it is on the order of picoseconds for a specialized FPGA (Matteis et al., 2019). However, the benefits of networking the data storage would be increases in networked, failover, flexibility, scalability, reduncancy, etc. For this reason, the SPARC tokamak is being constructed to enable multiple machines to be networked together (Creely et al., 2020).

Realtime model deployment has certain hard limits in the amount of ML development that can be done. Thus, testing of possible online learning, or even of production model performance is limited. Once the realtime system ports data over to offline systems, it's done, and though the data quality itself is mostly the same, there are some differences. For example, NaNs are flagged, created, and filled while ported to the system. Realtime systems do not produce any NaNs, but rather fill their would-be NaN values with white noise. Thus, while we can approximate a realtime test environment fairly closely, we cannot simulate it entirely. Realtime applications being explored, with highly optimized network cards at a certain speed (Spangher, 2024).