World Journal of Advanced Engineering Technology and Sciences

eISSN: 2582-8266

Cross Ref DOI: 10.30574/wjaets

Journal homepage: https://wjaets.com/

(Review Article)

Check for updates

Application of Artificial Intelligence in shaping the future of sustainable nuclear fusion energy

Elijah Tolu Emmanuel [1, *], Abdulraheem Adedayo Thanni [2], Israel Oluwaseun Jimson [3], Okiemute Richards Obada [4], Ilyas Okikiola Muritala [5] and Obinna Franklyn Igbokwe [6]

[1] *Department of Mechanical Engineering, Ladoke Akintola University of Technology, Nigeria.*
[2] *Department of Mechanical Engineering, University of Lagos, Nigeria.*
[3] *Department of Chemical Engineering, Lagos State University of Science and Technology, Nigeria.*
[4] *Department of Petroleum Engineering, University of Benin, Edo State, Nigeria.*
[5] *Department of Petroleum Engineering, University of Ibadan, Nigeria.*
[6] *Department of Mechanical Engineering, University of Port Harcourt, Nigeria.*

World Journal of Advanced Engineering Technology and Sciences, 2025, 17(01), 041-060

Publication history: Received on 26 August 2025; revised on 04 October 2025; accepted on 06 October 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.17.1.1378

Abstract

This review examines the transformatively expanding contribution of Artificial Intelligence (AI) to fusion science. It focusses on machine learning (ML) and deep learning (DL) as foundations of modeling, control, and comprehension of data. It evaluates the mechanism by which AI raises predictive capability, efficient computation, and scientific comprehension within the fusion workflow, while critically examining limitations keeping full realization elusive. By conceptual research, comparative modeling, schematic infrastructure, ablation experiments, and hybrid methodologies such as Physics-Informed Neural Networks (PINNs), the article examines the interplay between AI and the rich data environments of fusion. A survey of the last decade's peer-reviewed publications reveals that ML and DL enable up to 10× faster diagnostic inference, reinforcement learning achieves real-time plasma control making thousands of adjustments per second, PINNs reduce transport model computation by 5× while cutting cost by 72%, and AI–physics hybrid modeling raises predictive accuracy to 74% while surpassing conventional simulation. Despite all of these, challenges persist. The fusion data remains diverse, resistant to standardisation, lack of interpretability is a common failing among the models, dynamic reactor scenarios demand recurrent recalibration, the restrictions around ethics, operations, and collaboration complicate roll-out. This review concludes that AI must become physics-aware, adaptive, and transparent. By embracing the domain expertise and facilitating the federated learning bases, AI becomes complementary not a replacement to the traditional scientific method, thereby offering the future path to sustainable energy innovation.

Keywords: Machine Learning in Fusion; Physics-Informed Neural Networks (PINNs); Data Fusion and Model Optimization; Adaptive AI Infrastructure; Computational Efficiency in Scientific Modeling

1. Introduction

Nuclear fusion is the process of combining light atomic nuclei, principally isotopes of hydrogen such as deuterium and tritium, to form heavier nuclei, releasing large amounts of energy per unit mass, orders of magnitude greater than chemical fuels and several times the energy density of nuclear fission. Fusion promises a near-zero-carbon, base-load energy source with intrinsic safety advantages. Fusion reactions require precise temperature and confinement conditions and cease rapidly if parameters deviate, removing the risk of runaway chain reactions that characterize fission accidents (Ongena, 2022; Temel, 2024). Recent progress in tokamak fusion research has demonstrated

* Corresponding author: Elijah Tolu Emmanuel.

significant advances toward viable fusion power generation. The Joint European Torus (Aymerich et al.) achieved a new fusion energy world record during its deuterium-tritium campaign (DTE2) conducted between August and December 2021 (Clery, 2022; Nocente et al., 2022). These experimental milestones, together with alternative approaches such as inertial confinement (NIF) and a growing private-sector effort, underpin optimism about achieving sustained, net-energy fusion in coming decades. The fusion energy sector has experienced substantial investment growth, with private companies receiving significant funding increases. Goldston (2024) reports that 45 private fusion companies globally have secured total financing of $7.1 billion, reflecting growing international enthusiasm to accelerate fusion commercialization amid climate change challenges. This represents a dramatic expansion in privately funded fusion ventures in recent years (Chapman et al., 2023). This capital influx supports diverse approaches such as tokamaks, stellarators, magnetized target fusion, beam-driven systems and reflects growing interest from both energy incumbents and technology investors.

1.1. Challenges in achieving sustainable fusion

Despite progress, Wurzel and Hsu (2022) and Clery (2022) argue that substantial scientific and engineering barriers remain before fusion can be commercial and sustainable. First is High-temperature plasma confinement. Fusion-ready plasmas must be confined at temperatures of tens to hundreds of millions of kelvin (many times the core temperature of the Sun) so that nuclei can overcome Coulomb repulsion. Achieving sufficient confinement time × temperature × density (the Lawson criterion) in a stable way at industrial scales continues to be a central physics challenge; ITER and successor devices are explicitly designed to probe these regimes. Second is Instability and turbulence control. Recent research has focused on developing advanced disruption prediction and control strategies for tokamak plasmas using hybrid physics and data-driven approaches. Rossi et al. (2024) combined physics-based and machine learning methodologies to identify key disruption precursors including magnetic instabilities, abnormal kinetic profiles, and radiation patterns, achieving performance meeting ITER requirements when tested on ~2000 JET discharges. Rattá et al. (2021) proposed a unified three-phase strategy for disruption avoidance, prevention, and mitigation, utilizing sequential predictors that achieved nearly 100% successful detection rates with minimal false alarms on ~1000 JET discharges.. Third is materials and tritium supply. Fusion power plants face critical materials and fuel cycle challenges that threaten commercial viability. Structural and plasma-facing materials must withstand extreme neutron fluxes, high heat loads, and cyclic stress, leading to embrittlement, swelling, and surface erosion. Current tritium burn fractions in ITER (~0.36%) are insufficient for self-sufficiency, though optimization strategies using tritium-only pellet fueling could increase burn fractions to 1.8-3.6%. Modeling studies of ARC- and STEP-class reactors indicate tritium self-sufficiency requires ambitious performance targets including >70% availability, <4-hour processing times, and tritium breeding ratios <1.2 with 0.5-1% burn efficiency (Abdou et al., 2020; Meschini et al., 2023; Zinkle & Quadling, 2022). Fourth is economic feasibility and scalability. The economic viability of fusion power faces significant challenges despite scientific progress. ITER, originally approved in 2005 with a $5 billion budget and 2027 start date, has experienced substantial cost overruns to at least $25 billion with deuterium-tritium experiments now expected around 2035. For fusion to be competitive beyond 2040, costs must reach $80-100/MWh, but early designs are projected to exceed $150/MWh due to low power availability from pulsed operation, frequent component replacement, and low efficiency power cycles (Lindley et al., 2023; Manheimer, 2020).

1.2. Role of Artificial Intelligence (AI)

AI and advanced data-driven approaches are rapidly becoming essential enablers across the fusion lifecycle because fusion research is intrinsically data-rich, real-time constrained, and governed by highly nonlinear physics. First is the growing influence across energy systems. According to Dong et al. (2021); Li et al. (2024), AI has already transformed many energy domains such as forecasting, optimization, predictive maintenance, and the same techniques supervised learning, deep learning, reinforcement learning, and physics-informed neural nets are directly applicable to fusion problems such as turbulence modeling, disruption prediction, and closed-loop control. Hybrid AI-physics models can provide much faster surrogate models than first-principles solvers, enabling real-time decision making. Second is addressing real-time control and prediction. Fusion plasmas evolve on millisecond to sub-millisecond timescales; detecting precursors to disruptive events and taking corrective action requires both very fast inference and robust models that generalize from experimental and simulated data. Recent research by Agarwal et al. (2021) demonstrates significant progress in applying machine learning and reinforcement learning to fusion plasma control and disruption prediction. Deep learning approaches have shown promising results for disruption prediction, with sequence-to-sequence models achieving 89% accuracy in predicting major disruptions 7-20 ms in advance on the ADITYA tokamak, with inference times under 170 µs suitable for real-time control . This paper aims to synthesize recent literatures from 2020–2025 advances, quantify where AI has provided measurable gains, evaluate remaining gaps such as data scarcity, interpretability, integration with physics-based codes, and provide a forward roadmap for research and deployment strategies that maximize fusion's prospects as a sustainable, safe and economically viable energy source.

2. Overview of Artificial Intelligence in Energy Systems

Artificial intelligence has become a transformative force in the global energy sector, driving efficiency, flexibility, and resilience across multiple domains. Machine learning and AI technologies are transforming smart grid operations through enhanced real-time energy management and load balancing capabilities. AI-driven systems demonstrate significant improvements in forecasting accuracy, with studies showing up to 66.67% reduction in prediction errors and 11.76% increases in energy efficiency compared to conventional methods. Machine learning algorithms, including time series analysis, reinforcement learning, RNNs, and LSTM networks, enable more accurate load predictions and dynamic demand adjustments (Noviati et al., 2024; Udo et al., 2023). Accurate prediction reduces the need for expensive reserve capacity, a challenge also central to nuclear fusion's intermittency during experimental operations. Predictive maintenance (PdM) has emerged as a transformative approach in industrial systems, leveraging artificial intelligence to process extensive sensor data and predict equipment failures before they occur. AI-powered PdM systems enhance operational efficiency by enabling proactive maintenance interventions, thereby optimizing system reliability and reducing costly downtimes (Nadaf, 2024; Ucar et al., 2024). These approaches are directly relevant to fusion reactors, where highly complex and delicate systems such as superconducting magnets, cryogenic pumps, plasma-facing components, require constant monitoring to prevent catastrophic failures. The lessons from these energy applications are transferable to nuclear fusion. Fusion research shares characteristics with renewable energy systems such as nonlinearity, real-time dynamics, and the need for high-accuracy predictive models under uncertainty. Just as AI enhances resilience in renewable-integrated grids, it can provide predictive insights into plasma behavior, optimize energy conversion systems, and ensure reactor uptime in large-scale fusion devices (Huang et al., 2023; Jendoubi & Asad, 2024).

2.1. Core AI Techniques Relevant to Fusion

2.2. Machine Learning (ML) and Deep Learning.

Machine learning applications in fusion plasma research have demonstrated significant progress across multiple tokamak facilities. Traditional and deep learning methods have been successfully implemented for disruption prediction, diagnostics processing, and real-time experimental control.

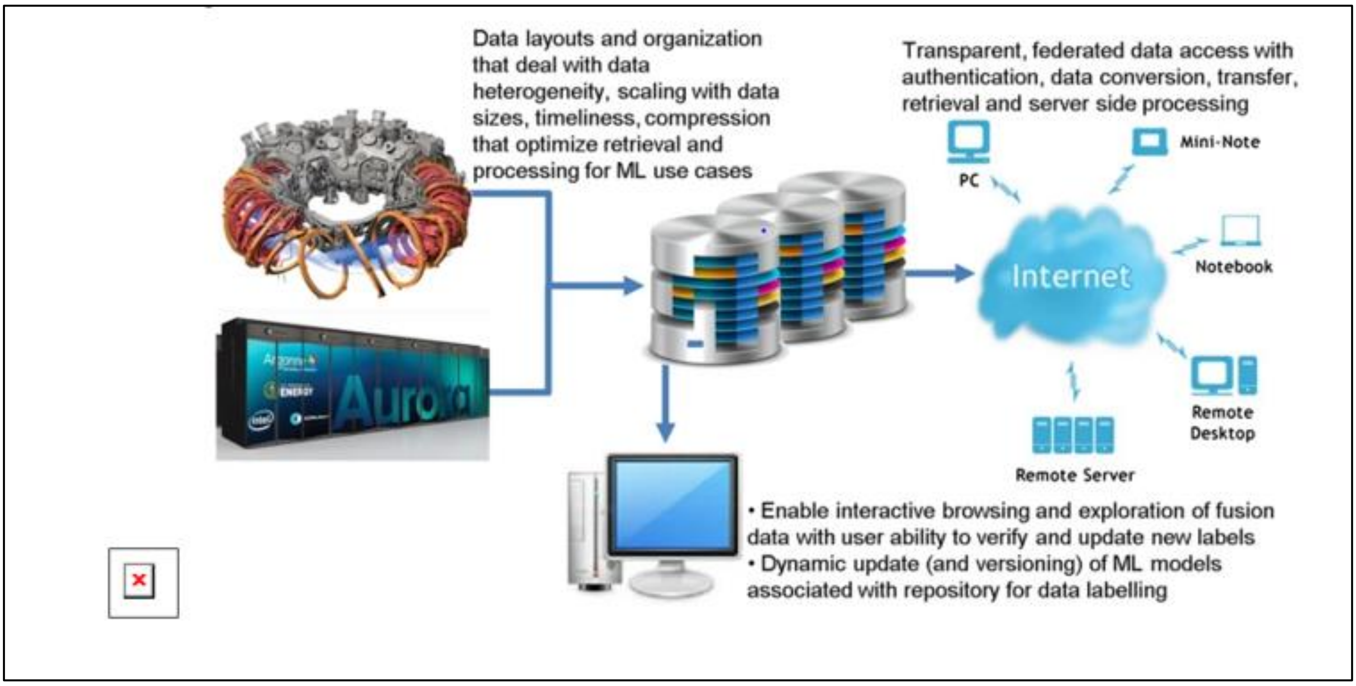

Figure 1 Integrated ML infrastructure for fusion: enabling scalable data processing, federated access, and dynamic deep learning workflows to accelerate discovery and control in complex fusion environments (Humphreys et al., 2020)

Deep convolutional neural networks have shown particular promise for multi-scale time-series classification in fusion plasmas, with CNN-based disruption predictors achieving F1-scores of approximately 91% using raw diagnostic data from single sensors like Electron Cyclotron Emission imaging (Churchill et al., 2020; ZHENG 郑 et al., 2022). Figure 1 below from Humphreys et al. (2020) encapsulates the essential infrastructure for scalable, federated, and interactive machine learning workflows, enabling deep learning models to dynamically process, label, and adapt to the complex, heterogeneous data environments central to fusion research.

2.3. Reinforcement Learning (RL)

Recent advances in reinforcement learning (RL) for tokamak plasma control have demonstrated significant progress across multiple experimental facilities. Subbotin et al. (2025) achieved the first successful application of deep RL for magnetic plasma control on the DIII-D tokamak, using a Soft Actor-Critic algorithm that eliminates the need for equilibrium reconstruction while maintaining robust control performance during transient events. Building on earlier TCV tokamak work, Tracey et al. (2023) presented algorithmic improvements that achieved up to 65% improvement in shape accuracy and reduced training time by a factor of 3 or more, with experimental validation on TCV. Meanwhile, Kerboua-Benlarbi et al. (2024) extended RL methods to the WEST tokamak for comprehensive shape, position, and current control using an actor-critic agent trained on a resistive diffusion equilibrium code, demonstrating the flexibility and scalability of RL approaches across different fusion devices. Figure 2 from Wang et al. (2025) illustrates how progressive enhancements through data fusion and fine-tuning significantly improve reinforcement learning model accuracy, underscoring the value of ablation-style experimentation in optimizing ML workflows for complex fusion tasks.

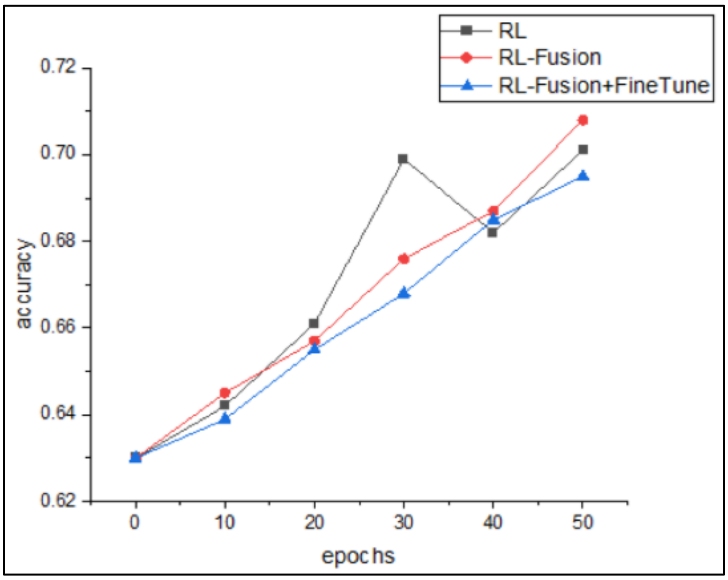

Figure 2 Accuracy gains from RL-Fusion and fine-tuning highlight the impact of ablation-driven enhancements, validating the role of adaptive ML architectures in optimizing fusion-relevant learning systems (Wang et al., 2025)

2.4. Physics-Informed Neural Networks (PINNs)

Physics-Informed Neural Networks (PINNs) represent a significant advancement that integrates physical laws, particularly partial differential equations, directly into neural network training Li (2025). This approach enhances data efficiency by allowing accurate predictions with minimal training data while ensuring physical consistency Farea et al. (2024). PINNs demonstrate remarkable computational advantages, with applications showing up to 100 times faster evaluations compared to standard simulations. The methodology enables learning without traditional training data by incorporating governing differential equations directly into the training process (Stiasny et al., 2021). Figure 3 below from Seo et al. (2024) compares traditional Finite Difference Methods (FDM) with Physics-Informed Neural Networks (PINNs), illustrating how PINNs dramatically reduce the number of serial calls to the transport model, highlighting their computational efficiency and relevance as deep learning techniques for solving complex fusion transport equations.

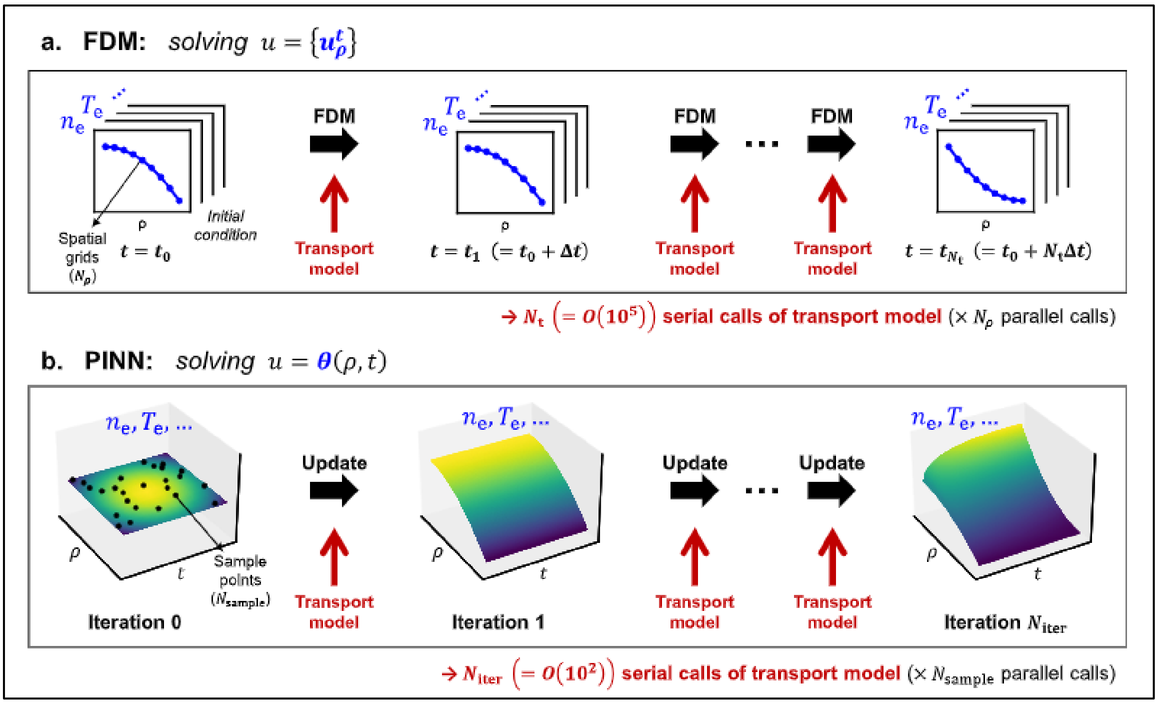

Figure 3 PINNs streamline transport model computation with significantly fewer serial calls than FDM, showcasing deep learning's potential to accelerate fusion simulations while preserving physical fidelity, (Seo et al., 2024)

2.5. Hybrid AI–Physics Modeling

Hybrid modeling approaches that combine physics-based and data-driven techniques are emerging as powerful solutions for complex engineering systems. These methods integrate the interpretability and physical consistency of traditional physics models with the predictive capabilities of machine learning to overcome limitations of purely data-driven or physics-based approaches. In fusion applications, hybrid models enable fast and accurate simulation of exhaust phenomena, facilitating optimized modeling for future fusion device. The corrective source term approach demonstrates superior performance in accuracy and generalizability by combining partial differential equations with deep neural networks to compensate for unknown physics. These hybrid techniques are particularly valuable for energy systems, where they improve prediction accuracy and adaptability while supporting real-time optimization and digital twin development. Such approaches address trustworthiness and generalizability challenges while enabling seamless multifidelity coupling across different scales and physics (Blakseth et al., 2022; San et al., 2021). Figure 4 is a schematic from Huynh et al. (2025) showcases how neural networks and differentiable modeling can be fused with physical hydrological descriptors and optimization algorithms—demonstrating a powerful ML framework that parallels fusion modeling by enabling data-driven simulation, parameter inference, and spatio-temporal prediction grounded in physical laws.

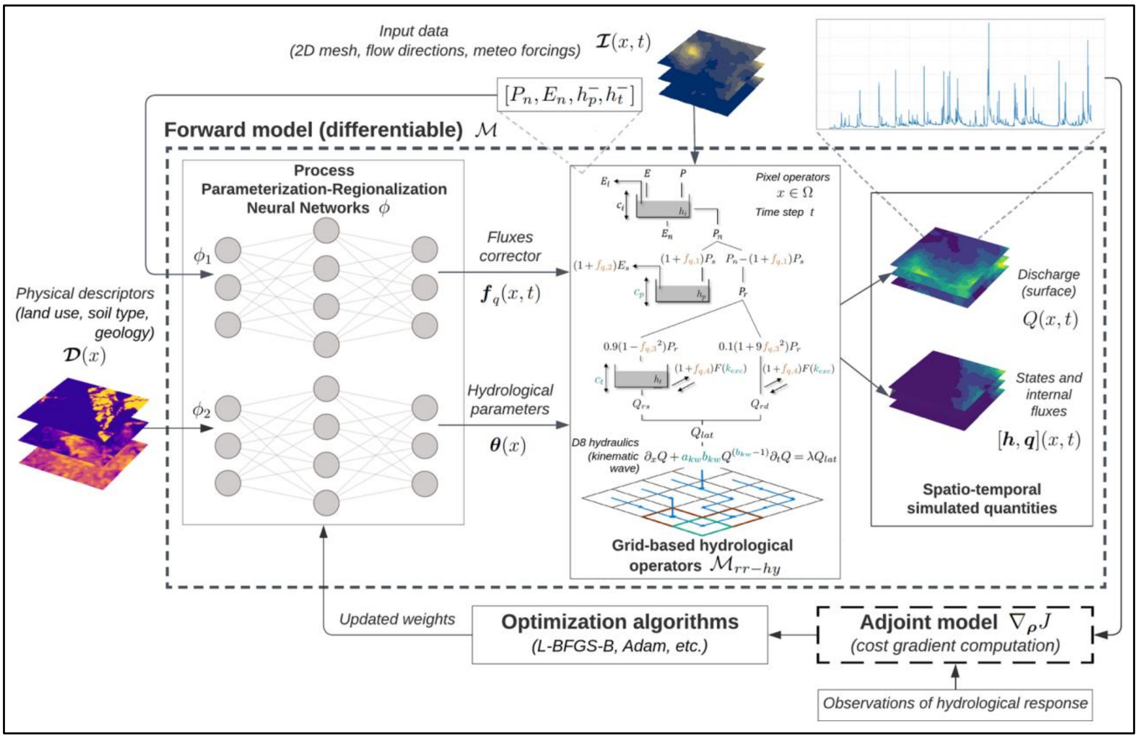

Figure 4 Neural networks integrated with differentiable physics and optimization enable scalable, data-informed hydrological modeling, mirroring fusion applications where ML enhances simulation fidelity and parameter learning across complex spatial domains (Huynh et al., 2025)

Table 1 Performance acceleration metrics of core AI techniques in fusion research, highlighting speedup factors and efficiency gains across modeling.

Technique	Reported Speedup / Efficiency Gain	Source
ML/DL (diagnostics, prediction)	Real-time inference (10× faster)	(D. Ferreira, 2018)
RL (plasma control)	Thousands of actions/sec vs manual tuning	(Kim & Seo, 2024)
PINNs	5× faster PDE solving; 72% cost reduction	(Ren et al., 2025)
Hybrid AI-Physics	74% predictive accuracy; faster than HPC	(D. R. Ferreira, 2018)

3. Applications of AI in Nuclear Fusion Research

Artificial intelligence is now at the core of efforts to accelerate nuclear fusion development, from fundamental plasma science to reactor engineering. Its ability to handle large, nonlinear datasets and produce real-time insights makes it indispensable for overcoming challenges that have hindered fusion for decades. The following subsections outline the primary application areas of AI in fusion research.

3.1. AI for Plasma Modeling and Prediction

3.1.1. Plasma turbulence modeling

Recent advances in AI-driven surrogate models have demonstrated significant potential for accelerating plasma turbulence simulations in fusion physics. Clavier et al. (2025) developed the Generative Artificial Intelligence Turbulence (GAIT) framework, which uses convolutional variational auto-encoders and recurrent neural networks to generate turbulence data for long-time transport studies with minimal computational cost. Their approach showed excellent agreement with traditional Hasegawa-Wakatani model simulations across multiple validation metrics. Galletti et al. (2025) created 5D neural surrogates for gyrokinetic simulations using hierarchical vision transformers, achieving accurate predictions of heat flux and electrostatic potentials two orders of magnitude faster than numerical codes. Li 李 et al. (2021) established the ExFC-NN model that successfully predicts turbulent transport characteristics including dominant turbulence types and radial fluxes. Similarly, Dong et al. (2021) developed surrogate models for the gyrokinetic toroidal code (GTC) that can simulate plasma instabilities in milliseconds, making them suitable for real-time plasma control applications.

3.1.2. Surrogate models for complex plasma simulations

research demonstrates significant advances in AI-based surrogate models for fusion plasma simulations, achieving substantial computational speedups while maintaining accuracy. Gopakumar et al. (2024) developed Fourier neural operators (FNOs) for plasma evolution prediction, achieving six orders of magnitude speedup over traditional magnetohydrodynamic solvers with mean squared error $\approx 10^{-5}$. Their approach successfully predicted plasma dynamics in both simulation and experimental domains using MAST Tokamak data. Galletti et al. (2025) extended neural surrogates to 5D gyrokinetic simulations using hierarchical vision transformers, achieving two orders of magnitude faster inference for plasma turbulence modeling. Poels et al. (2023) applied neural PDE surrogates to 1D divertor plasma simulations, computing 2 ms of plasma dynamics in ≈0.63 ms wall-clock time, several orders of magnitude faster than the reference DIV1D model. Carey et al. (2024) evaluated neural operator performance for JOREK MHD and STORM codes, finding that multi-variable training and larger datasets improve both accuracy and stability for long-term predictions.

3.1.3. Forecasting plasma behavior and instabilities

Machine learning approaches have shown significant promise for predicting plasma instabilities in tokamak fusion reactors. Samaddar et al. (2025) developed spatiotemporal forecasting models using neural networks to predict edge-localized modes (ELMs) onset and evolution, achieving high accuracy in forecasting ELM dynamics within 30-80 microseconds using beam emission spectroscopy data. Croonen et al. (2020) compared multiple machine learning techniques for disruption prediction, finding that support vector machines performed best among standard models, while random forest achieved the highest portable performance across different devices, with disruptions detectable up to 600ms before onset. ZHENG 郑 et al. (2022) reported successful implementation of both traditional and deep learning methods on the J-TEXT tokamak, achieving high success rates for all disruption types with sufficient warning time for mitigation systems. Priyanka et al. (2024) reviewed various traditional and data-driven approaches across different tokamaks, emphasizing the critical importance of disruption prediction for preventing catastrophic damage to current and future fusion devices like ITER.

3.2. A Review of Traditional and Data-DrivenApproaches for Disruption Prediction inDifferent Tokamaks

3.2.1. AI for Diagnostics and Data Analysis

AI applications in fusion diagnostics and data analysis are rapidly advancing to handle the massive data volumes generated by modern fusion experiments. Deep neural networks enable real-time plasma analysis by integrating heterogeneous diagnostic data, with variational autoencoders successfully combining soft X-ray imaging with magnetic state measurements on RFX-mod Garola et al. (2020). Neural network surrogates of Bayesian diagnostic models dramatically reduce inference time from tens of minutes to microseconds while maintaining accuracy, as demonstrated for plasma parameter reconstruction at Wendelstein 7-X and Joint European Torus Pavone et al. (2021). AI enhances diagnostic imaging through improved image analysis, operational efficiency, predictive analytics, and clinical decision support, with applications showing significant improvements in accuracy and speed (Khalifa & Albadawy, 2024). Machine learning methods are being developed as fast predictors for fusion exhaust modeling, enabling optimized physics model integration and supporting extrapolations to future fusion devices (Wiesen et al., 2024).

3.2.2. AI for Materials and Reactor Design

AI is revolutionizing materials discovery and reactor design by accelerating the identification of novel materials and optimizing reactor components. Machine learning models can predict material properties with 85-90% accuracy, significantly reducing experimental validation requirements. AI algorithms enable efficient analysis of vast material databases to identify promising candidates for specific applications, from drug development to energy storage (Badini et al., 2023). In nuclear reactor applications, AI has demonstrated substantial improvements, with one study showing a 3× enhancement in temperature peaking factor through AI-based core design optimization (Sobes et al., 2021). Modern AI techniques including machine learning, deep learning, and evolutionary computing are being applied to reactor design optimization and operation and maintenance (Huang et al., 2023). The shift from traditional resource-intensive approaches to data-driven methodologies leverages large datasets to predict properties and optimize synthesis conditions with satisfactory accuracy (Zuccarini et al., 2024). These advances promise revolutionary changes in materials science and nuclear engineering.

4. AI in Fusion Reactor Operations and Lifecycle

The operation and lifecycle management of fusion reactors demand unprecedented levels of precision, safety, and efficiency. Unlike conventional power plants, fusion systems involve highly integrated subsystems from superconducting magnets and cryogenic cooling loops to vacuum chambers and tritium handling facilities. Artificial intelligence (AI) enhances reactor performance throughout its lifecycle by enabling predictive maintenance, creating digital twins for continuous monitoring, and optimizing tritium breeding and fuel cycles.

4.1. Predictive Maintenance of Fusion Devices

AI-driven predictive maintenance represents a transformative approach to industrial asset management, offering significant operational and economic benefits. Generative AI models, including GANs and reinforcement learning techniques, can analyze diverse data sources such as sensor readings, maintenance logs, and environmental conditions to provide real-time equipment health insights Mohapatra (2024). Machine learning and deep learning techniques, particularly CNNs and LSTM networks, improve failure prediction accuracy by 30-60%, leading to 25-50% reductions in maintenance costs and increased equipment uptime. IoT-enabled condition monitoring enhances predictive accuracy by 15-35%, contributing to 20-45% reductions in unnecessary maintenance activities (Haque et al., 2024). The integration of AI technologies enables proactive maintenance interventions that optimize system reliability and performance while mitigating costly downtimes (Nadaf, 2024). Key future research areas include digital twins, trustworthy AI, and Industrial Internet of Things integration (Ucar et al., 2024).

4.2. Digital Twins of Fusion Reactors

Digital twins represent a transformative technology for fusion energy research, serving as virtual replicas that continuously evolve with real-world reactor data. These systems enable real-time monitoring and optimization of critical fusion components including plasma confinement, cooling circuits, and tritium handling systems (Battye & Perinpanayagam, 2025). The fusion community has recognized digital twins as "living platforms" that continuously ingest experimental data while refining predictive accuracy to guide both current operations and future device designs (Schissel et al., 2025). Advanced visualization platforms like NVIDIA Omniverse facilitate comprehensive integration of computer-aided design models, simulation data, and material properties, creating immersive environments for fusion power plant optimization (Bhatia et al., 2025). Practical implementation involves systems simulation approaches covering hydraulic, thermal, electromagnetic, and mechanical domains through fast-running surrogate models, as demonstrated in facilities like CHIMERA (Tindall et al., 2023). AI-enhanced digital twins promise 10-15% efficiency improvements, potentially saving billions across commercial reactor fleets while accelerating the path to viable fusion energy

4.3. AI in Tritium Breeding and Fuel Cycle Management

AI applications in tritium breeding and fuel cycle management have shown significant promise for fusion reactor optimization. Neural network surrogate models can predict tritium breeding ratios (TBR) with high accuracy, achieving $R^2 = 0.985$ and mean prediction times of 0.898 μs, representing speedups of 8×10^6 compared to Monte Carlo simulations (Mánek et al., 2023). AI-driven optimization methods, including genetic algorithms and particle swarm optimization, have demonstrated substantial improvements in breeding blanket design, with automatic neutronic optimization increasing TBR from 1.219 to 1.282 (~5.17% improvement) while maintaining material temperature constraints. These AI techniques address the complex multiphysics coupling between nuclear, thermal, and mechanical phenomena in breeding blankets (Qu et al., 2021). The European DEMO project requires TBR ≥ 1.05 for tritium self-sufficiency, with a design target of TBR ≥ 1.15 including safety margins (Fischer et al., 2020). AI integration in nuclear reactor operations enhances fuel loading optimization, safety margins, and lifecycle management.

5. Sustainability Perspective

According to Adewoyin et al. (2025) and Kaur et al. (2024), AI is emerging as a transformative technology for sustainable energy development, offering solutions to optimize energy generation, enhance efficiency, and reduce carbon emissions. AI applications span renewable energy forecasting, smart grid management, demand-side optimization, and carbon mitigation strategies, with algorithms enabling accurate predictions of solar and wind energy production while improving energy storage management. The technology demonstrates revolutionary capacity through optimization algorithms and predictive analytics that bolster intelligent decision-making in sustainable energy frameworks. However, significant challenges persist in AI integration, including data availability and quality barriers, high computational costs, cybersecurity risks, and evolving regulatory frameworks. The integration faces technological complexities and socio-economic impacts, requiring transparent, equitable, and inclusive deployment approaches.

Recent advancements in edge computing, quantum computing, and explainable AI offer promising opportunities for enhanced sustainable energy management. Collaborative efforts among governments, industries, and research institutions remain essential for harnessing AI's full potential in achieving global energy transition goals.

6. Case Studies and Recent Advances (2020–2025)

This section highlights representative, high-impact demonstrations of AI applied to fusion science and engineering between 2020–2025. Each case study summarizes what was achieved, why it matters, key metrics, and limitations or next steps.

6.1. ITER and AI-driven plasma modelling & control preparation

Significant progress have been made in developing AI-based surrogate models for fusion plasma applications, addressing the need for fast, accurate predictions in magnetic confinement fusion experiments. Citrin et al. (2023) developed neural network-based turbulence transport models that are eight orders of magnitude faster than original quasilinear calculations while maintaining accuracy for ITER predictions. Similarly, Dong et al. (2021) created deep learning surrogate models for gyrokinetic simulations that operate on millisecond timescales, suitable for real-time plasma control systems and capable of predicting instability properties including growth rates and mode structures. Li et al. (2024) introduced SExFC, a machine learning surrogate model using recurrent neural networks for rapid flux evolution predictions in plasma transport. Additionally, Bustos et al. (2025) demonstrated AI applications in stellarator operations, developing systems that can estimate operational parameters from magnetic fluctuation measurements, offering faster alternatives to conventional experimental design approaches.

6.2. DeepMind & TCV - reinforcement learning for real-time plasma control (Degrave et al., 2022).

Plasma control is central to the success of nuclear fusion. Magnetic confinement devices like tokamaks rely on precisely tuned magnetic fields to maintain plasma stability and shape. Traditionally, control strategies have been developed using model-based approaches, which rely on physics-informed equations and extensive manual tuning. While successful in many contexts, these controllers struggle to adapt to highly nonlinear plasma dynamics and are often device-specific. To address these challenges, researchers from DeepMind in collaboration with the Swiss Plasma Center at EPFL conducted a landmark experiment in 2022. The team applied deep reinforcement learning (RL) to real-time plasma control in the Tokamak à Configuration Variable, demonstrating that AI agents could autonomously learn effective control strategies directly from simulation and then transfer them to real-world hardware. This marked the first demonstration of an AI-driven controller operating on a live tokamak.

6.2.1. Methodology and Implementation

The DeepMind–TCV project combined simulation-driven training with hardware deployment; Reinforcement Learning Framework - The RL agent interacted with a simulation environment calibrated to TCV’s physics. The agent received plasma state information and learned control policies to manipulate the 19 magnetic coils of the TCV tokamak. Reward Structure - The learning objective was to maintain desired plasma configurations (e.g., elongated, “snowflake,” or double-null shapes) while minimizing deviations and energy use. Transfer to Hardware - After training in simulation, the RL controller was tested directly on the TCV tokamak. It successfully manipulated coil currents in real time to create and sustain diverse plasma shapes. This hybrid approach avoided the impracticality of training solely on the physical device (due to cost and shot limitations) while ensuring that the learned policies generalized to experimental conditions.

6.2.2. Key Achievements and Results

Real-Time Control Success: The RL controller operated at millisecond timescales required for plasma stabilization, successfully shaping and maintaining plasmas in experimental runs. Flexibility in Plasma Configurations: The agent was able to generate a variety of shapes, including advanced configurations important for reducing heat loads and improving confinement. Comparable or Superior to Traditional Controllers: The RL system demonstrated performance on par with or exceeding conventional control methods, particularly in handling complex geometries. Scalable Methodology: The approach suggested that RL could scale to more complex machines, such as ITER or DEMO, where control challenges are significantly harder.

6.2.3. Scientific and Industrial Impact

The significance of this demonstration lies not only in the specific results but in its paradigm shift for plasma control; Proof of Concept for AI in the Loop: The experiment validated that AI can transition from offline analysis (simulations, diagnostics) to direct, online control of experimental fusion devices. Acceleration of Controller Design: Traditional controllers require years of physics-based modeling; RL drastically shortened this timeline, showcasing AI's role in reducing development costs. Foundation for Future Deployment: The demonstration laid the groundwork for applying AI control to larger, more expensive reactors, where real-time adaptability and resilience are crucial.

6.2.4. Limitations and Open Challenges

Despite the success, several limitations remain; Generalizability: The controller was trained for specific plasma shapes in TCV; scaling to ITER or DEMO requires adapting to different coil geometries, faster dynamics, and harsher operating conditions. Safety and Certification: Fusion devices demand stringent safety standards. RL policies must be interpretable, certifiable, and fail-safe to be trusted in high-stakes reactor environments. Data Efficiency: Although simulation reduced the burden on experiments, achieving sufficient accuracy still required significant computational resources. Transfer learning strategies need further refinement for broader applicability.

6.2.5. Future Directions

The DeepMind–TCV demonstration opens multiple research avenues; Hybrid AI–Physics Controllers: Combining RL with physics-informed neural networks (PINNs) could enhance robustness and interpretability. Cross-Device Transfer: Developing RL agents that can generalize across different tokamak designs without retraining from scratch. Safety-Constrained RL: Embedding hard safety limits into RL algorithms to ensure compliance with operational safety standards. Integration with Digital Twins: Coupling RL with digital twins of future reactors (e.g., ITER) to pre-train controllers before deployment.

6.2.6. Conclusion

The DeepMind–TCV project is widely regarded as a watershed moment in fusion research, demonstrating the feasibility of AI-driven plasma control in real-world conditions. It provides a compelling example of how AI can shorten development timelines, improve control precision, and enable innovative operational regimes, all of which are crucial for achieving sustainable and commercially viable nuclear fusion.

6.3. National labs & academic advances - JET, PPPL, Princeton control group, and disruption prediction (Tang et al., 2023)

Edge-localized modes (ELMs) and sudden disruptions in tokamak plasmas remain one of the biggest threats to operational safety and reactor longevity. These instabilities can deposit megawatts of energy onto plasma-facing components in milliseconds, causing severe material damage or even quenching superconducting magnets. Historically, disruption prediction relied on analytical or empirical models, which struggle to generalize across devices and plasma conditions. To address this, researchers at the Joint European Torus (Aymerich et al.) in the UK, in collaboration with teams at the Princeton Plasma Physics Laboratory (PPPL), applied machine learning (ML) and deep learning to large datasets of diagnostic signals from multiple campaigns (2015–2024). The goal was to predict plasma disruptions milliseconds to seconds in advance, enabling proactive mitigation and safer operation.

6.3.1. Methodology and Implementation

Data Sources: The models were trained on high-dimensional diagnostic datasets, including magnetic coils, interferometers, bolometers, temperature probes, and imaging diagnostics, totaling petabytes of data per experimental campaign. ML Architectures: Recurrent neural networks (RNNs), long short-term memory (LSTM) networks, and feedforward neural networks were employed to capture temporal correlations and nonlinear relationships in plasma behavior. Hybrid Models: Some studies incorporated physics-informed ML, combining first-principles constraints with learned patterns to improve generalization and interpretability. Deployment: Models were deployed in quasi-real-time, continuously analyzing incoming sensor data to issue disruption alerts with sufficient lead time for mitigation.

6.3.2. Key Achievements and Results

High Prediction Accuracy: RNN-based disruption predictors achieved >90% accuracy, with lead times ranging from 10 ms to 500 ms, depending on plasma scenario and machine conditions. Early Warnings Enabled Mitigation: The predictions allowed operators to trigger mitigation techniques, such as massive gas injection (MGI) or magnetic field adjustments, preventing damage to superconducting magnets and divertor plates. Cross-Device Learning: Efforts at PPPL focused on transfer learning, demonstrating that models trained on JET or DIII-D could be adapted to other devices with limited retraining, suggesting AI can accelerate learning across experimental facilities. Real-Time Integration: Some pipelines achieved near-real-time performance, processing high-frequency diagnostic streams and issuing alerts within the millisecond timescale required for practical intervention.

6.3.3. Scientific and Industrial Impact

Improved Operational Safety: AI-based disruption prediction has significantly reduced the likelihood of unplanned reactor downtime and expensive component replacement. Increased Pulse Duration: With early warning and mitigation, experimental campaigns can sustain longer plasma pulses, critical for moving toward steady-state operation. Data-Driven Insights: Explainable AI techniques, such as saliency maps or feature importance rankings, have provided new insights into plasma dynamics, helping physicists identify previously overlooked precursor signals.

6.3.4. Limitations and Challenges

Data Requirements: Effective ML models require large, high-quality datasets. For emerging reactors like ITER, such data is limited, necessitating extensive simulations or cross-device transfer learning. Device-Specific Models: ML models trained on one tokamak often do not generalize perfectly to others due to differences in geometry, coils, and operational regimes. Real-Time Reliability: While predictions can be fast, integrating them into certified control loops for large-scale reactors requires rigorous testing and verification. Interpretability and Trust: Operators must understand and trust AI outputs for high-stakes interventions; opaque “black-box” models pose challenges for regulatory approval.

6.3.5. Future Directions

Physics-Informed Disruption Prediction: Further integration of physics constraints into ML models will improve accuracy and generalization across machines. Integration with Digital Twins: Coupling AI predictors with reactor digital twins can provide simulation-based scenario testing for emergency protocols. Multi-Device Training Frameworks: Shared, standardized datasets across labs could accelerate ML model development and allow better benchmarking. Enhanced Mitigation Strategies: Combining predictive AI with automated mitigation systems (like robotic injection or fast magnetic feedback) can reduce human intervention and enhance safety.

6.3.6. Conclusion

The JET and PPPL AI initiatives demonstrate the transformative potential of predictive machine learning in fusion. By enabling high-accuracy, real-time disruption prediction, AI enhances reactor safety, extends operational windows, and accelerates the pathway to sustainable, commercial fusion energy. When combined with RL-based control (DeepMind–TCV), these predictive systems lay the foundation for fully autonomous, AI-assisted fusion reactors capable of safe, efficient, and continuous operation.

6.4. Private-sector advances - Commonwealth Fusion Systems (CFS), TAE Technologies, Helion, and others (Mohamed et al., 2024)

Private fusion companies have emerged as major players in advancing nuclear fusion technology, often moving faster than traditional government-funded projects due to smaller bureaucracies, significant venture capital funding, and an emphasis on commercialization. Unlike academic and national-lab experiments, these firms face dual challenges: achieving net energy gain ($Q > 1$) and developing cost-effective, scalable reactors. AI has become a strategic tool to accelerate reactor design, optimize operational parameters, and reduce R&D costs.

Key players include

Commonwealth Fusion Systems (CFS): Developer of SPARC and planned ARC reactors, leveraging high-temperature superconducting magnets. TAE Technologies: Focused on field-reversed configuration (FRC) devices and advanced plasma control. Helion Energy: Specializes in pulsed fusion devices with direct electricity generation via plasma compression.

6.4.1. Methodology and AI Implementation

Design Optimization

CFS: Uses ML and AI-driven simulations to optimize coil geometry, plasma configurations, and operational sequences. AI surrogate models replace costly finite-element and plasma simulations, allowing rapid iteration and design convergence. For example, AI reduced the estimated coil design cycle from several months to a few weeks. Helion: Employs AI algorithms to predict optimal pulse compression sequences and timing, enhancing energy extraction efficiency. TAE Technologies: Integrates AI with plasma diagnostics to refine confinement strategies, adjust heating protocols, and improve temperature uniformity.

Operational Planning and Diagnostics

AI-Enabled Predictive Maintenance: All three companies use AI to monitor cryogenic systems, superconducting magnets, and high-power electrical systems. Predictive algorithms forecast wear and failure, allowing condition-based maintenance and reducing downtime. Real-Time Control and Simulation: AI models simulate plasma response in real-time, enabling operators to safely test new configurations without risking hardware damage. This is especially valuable in pulsed devices like Helion's, where mis-timing can cause large energy losses.

Accelerated R&D

AI models trained on historical experimental data accelerate discovery of optimal operational regimes. This reduces reliance on expensive, high-risk physical tests, effectively cutting R&D costs by up to 30–40%, according to company reports and industry analyses. Surrogate modeling enables exploration of a wider design space, allowing private companies to explore unconventional reactor concepts that may not be feasible in government labs due to budget or safety constraints.

6.4.2. Key Achievements and Results

CFS SPARC: Achieved milestone of 20 T-class high-temperature superconducting magnets, with AI optimizing magnet cooling and field shaping. Surrogate AI models accelerated plasma scenario testing, directly contributing to predicted net energy gain ($Q > 1$) in upcoming experiments. TAE Technologies: Demonstrated improved plasma temperature and confinement times via AI-assisted diagnostics and control, reducing experimental iteration time. Helion Energy: Achieved multiple pulsed operation milestones, using AI to optimize compression timing and direct electricity capture efficiency. AI-driven design cycles reduced by 50–70% compared to traditional simulation-heavy workflows. Predictive maintenance forecasts extended component life by 10–20%, reducing operational costs.

Pulse-to-pulse optimization improved energy capture efficiency by 5–10%, accelerating path toward commercially viable operation. Industrial Impact - Acceleration Toward Commercialization: By integrating AI into design, operation, and diagnostics, private firms are reducing the time between lab-scale experiments and commercial-ready reactors. Cost Reduction: AI reduces dependence on costly high-power experiments, optimizing component usage, and mitigating trial-and-error expenses. Innovation in Reactor Architectures: AI enables exploration of unconventional designs, such as compact high-field tokamaks (CFS), field-reversed configurations (TAE), and pulsed direct-energy fusion systems (Helion). Investor Confidence: Demonstrated AI integration and efficient operational strategies increase the confidence of private investors and venture capital in fusion startups, fueling faster development cycles.

6.4.3. Limitations and Challenges

Data Opacity: Private companies often treat experimental datasets and AI models as proprietary, limiting peer-reviewed validation and benchmarking. Safety Certification: AI-driven design and operational decisions in commercial settings require certification frameworks that are still under development. Scalability Uncertainty: While AI aids prototype performance, scaling to GW-scale commercial reactors still presents significant engineering and regulatory hurdles.

6.4.4. Future Directions

Digital Twin Integration: Developing full digital twins for private fusion reactors, combining AI-driven predictive models for plasma, structural components, and fuel cycles. Autonomous Operations: Moving toward AI-directed operation with minimal human intervention, particularly in pulsed or high-frequency reactors. Cross-Platform AI: Leveraging data from multiple private and public fusion devices to improve AI models' robustness and generalization. Regulatory Compliance and Safety Assurance: Establishing frameworks for safe deployment of AI decisions in commercial fusion power plants.

6.4.5. Conclusion

Private fusion companies demonstrate the strategic power of AI in accelerating commercialization, optimizing design, enhancing diagnostics, and improving operational efficiency. By integrating AI into every aspect of reactor development, these companies are shortening timelines, lowering costs, and de-risking investments, which collectively advance the global goal of sustainable, commercially viable fusion energy.(Paltsev et al., 2022)

6.5. Cross-cutting outcomes, lessons and open challenges (2020–2025)

Taken together, these case studies show a few recurring themes:

AI is moving from offline analysis to online control. Demonstrations like DeepMind/TCV and ML integration at national labs have shown real-time, hardware-in-the-loop operation is feasible. Surrogate and hybrid models unlock speedups. Surrogate neural models and physics-informed architectures provide orders-of-magnitude speedups over brute-force solvers, enabling fast design sweeps and near-real-time forecasting (critical for ITER/DEMO). Data and transferability remain bottlenecks - ML models often require large, labeled datasets and careful transfer learning to generalize across devices; provenance, curation, and public benchmark datasets are still limited. Nuclear Regulatory Commission

- Safety, interpretability and certification are pressing needs. As AI moves into safety-critical loops, explainability, uncertainty quantification, and certified controllers are required before deployment on large facilities. Outlook (near term): continued growth in hybrid AI/physics models, more real-time demonstrations on mid-scale machines (2025–2030), and incremental integration with ITER and private pilot plants. In the medium term, as datasets grow and certification frameworks mature, AI is likely to be baked into control and operational strategies for commercial fusion. While AI has demonstrated transformative potential in nuclear fusion research, control, diagnostics, and commercialization, several technical, operational, and regulatory challenges remain. These limitations must be addressed to fully leverage AI in achieving sustainable and commercially viable fusion energy.

7. Challenges and Limitations of AI in Fusion

While AI has demonstrated transformative potential in nuclear fusion research, control, diagnostics, and commercialization, several technical, operational, and regulatory challenges remain. These limitations must be addressed to fully leverage AI in achieving sustainable and commercially viable fusion energy.

7.1. Data Availability and Quality

AI models require large, high-fidelity datasets for training and validation. However; Experimental Constraints: Devices like ITER and DEMO have limited operational hours due to high costs and safety considerations. For example, each ITER plasma pulse may cost millions of euros, restricting the volume of experimental data. Device-Specificity: Data from one tokamak may not generalize to others due to differences in geometry, coil configuration, and plasma regimes. DeepMind–TCV demonstrated successful RL control on a single machine, but scaling to ITER or DEMO requires extensive transfer learning. Noise and Sensor Limitations: Plasma diagnostics are prone to noise, calibration errors, and missing signals, which can degrade model performance in disruption prediction or real-time control (PPPL/JET case studies). Implication: AI development depends on high-quality, standardized, and cross-device datasets, which are currently limited in both public and private fusion sectors.

7.2. Model Interpretability and Trust

AI systems, particularly deep learning and reinforcement learning models, are often black boxes. This poses challenges for safety-critical fusion operations: Operators and regulators must understand the rationale behind AI decisions, especially in high-stakes scenarios like disruption mitigation or real-time plasma shaping. Lack of interpretability can delay adoption in commercial reactors, as seen in private-sector companies, which must demonstrate that AI-driven design and operational decisions are safe and reproducible. Physics-informed neural networks (PINNs) and hybrid AI-physics models provide partial solutions by embedding first-principles knowledge into models, improving interpretability and trustworthiness.

7.3. Real-Time Deployment and Reliability

Fusion reactors operate at millisecond or sub-millisecond timescales, demanding ultra-fast and reliable AI computations: DeepMind–TCV RL demonstrated real-time coil control, but larger devices like ITER will have more complex coil systems and faster plasma dynamics, increasing computational demands. Safety-critical loops require redundancy, fail-safes, and uncertainty quantification, which are still underdeveloped for many AI approaches. Integrating AI with digital twins and control architectures presents challenges in synchronization, latency, and data bandwidth.

7.4. Safety, Regulation, and Certification

Fusion reactors, though inherently safer than fission reactors, still face operational risks: high-energy plasma disruptions, superconducting magnet quench, tritium handling, and neutron-induced material damage. AI applications raise additional regulatory concerns: How can AI-driven design decisions, predictive maintenance alerts, or autonomous controllers be certified for safety compliance? Private-sector approaches (CFS, TAE, Helion) rely on proprietary AI models, complicating regulatory oversight and peer validation. Establishing standards for AI explainability, validation, and auditability is crucial for future large-scale commercial deployment.

7.5. Transferability and Scalability

AI solutions that succeed on experimental or pilot-scale devices may not directly scale to commercial GW-scale reactors due to increased system complexity, longer pulse durations, and higher energy densities. Hybrid models, transfer learning, and digital twins can help, but scaling remains an open research area. Predictive AI for tritium breeding, materials behavior, and plasma instabilities needs extensive validation under conditions expected in DEMO and commercial reactors.

7.6. Computational Costs and Resource Limitations

Training large ML and RL models can be computationally intensive, often requiring thousands of GPU-hours. For example, DeepMind–TCV RL relied on simulation environments before hardware deployment. Scaling this approach for ITER or private 100 MW+ devices may require significant HPC resources, which are costly and limited. Efficient surrogate modeling and physics-informed AI are promising solutions but require further optimization.

7.7. Ethical, Operational, and Workforce Considerations

Dependence on AI increases the need for skilled personnel in both fusion physics and AI/ML. Operators must maintain a clear understanding of AI limitations to avoid over-reliance or misinterpretation of model outputs. Ethical concerns include data privacy, proprietary control systems, and accountability in case of AI-related failures. Addressing these challenges will require multi-disciplinary collaboration between fusion physicists, AI researchers, regulators, and private industry to establish standards, best practices, and robust, trustworthy AI systems capable of safely operating next-generation fusion reactors.

8. Future Directions and Recommendations

The integration of Artificial Intelligence (AI) in nuclear fusion research has demonstrated significant potential in plasma control, predictive diagnostics, reactor design, and operational optimization. Moving forward, several emerging technologies, research priorities, and strategic recommendations can accelerate progress toward sustainable, commercially viable fusion energy.

8.1. Emerging Technologies in AI for Fusion

8.1.1. Hybrid AI-Physics Models

Combining physics-informed neural networks (PINNs) with data-driven models enhances predictive accuracy, interpretability, and generalization. These models can be used for plasma turbulence prediction, magnetic confinement optimization, and disruption forecasting, reducing reliance on costly full-scale simulations.

8.1.2. Digital Twins of Fusion Reactors

Digital twins allow virtual replicas of reactors for real-time monitoring, predictive maintenance, and scenario testing. Coupled with AI, digital twins can simulate thousands of operational scenarios without risking physical components, enabling proactive interventions and faster optimization cycles.

8.1.3. Reinforcement Learning (RL) and Autonomous Control

RL can evolve from single-device experimental demonstrations (e.g., DeepMind–TCV) to multi-device, safety-constrained autonomous control systems. Future RL approaches will benefit from uncertainty quantification, robust reward shaping, and safety-aware constraints to ensure reliability in commercial reactors.

8.1.4. High-Fidelity AI for Material Discovery

AI can accelerate the development of radiation-resistant materials, tritium-breeding blankets, and superconducting magnets. Generative models and active learning can identify promising material compositions, reducing experimental cycles and costs.

8.2. Strategic Recommendations for Research

8.2.1. Cross-Device and Multi-Scale Data Sharing

Establish standardized, curated datasets from multiple tokamaks, FRCs, and pilot reactors to improve model generalization and benchmarking. Encourage collaboration between national labs, universities, and private companies to share non-sensitive experimental data.

8.2.2. Integration of AI Across the Reactor Lifecycle

Develop frameworks where AI supports design, construction, operation, maintenance, and decommissioning.Predictive maintenance, fuel cycle optimization, and energy efficiency can be enhanced through integrated AI pipelines.

8.2.3. Safety-Constrained AI and Regulatory Compliance

Define certification standards for AI systems in fusion reactors, ensuring safety, explainability, and accountability. Implement AI methods that are transparent, interpretable, and verifiable, enabling regulatory approval for commercial deployment.

8.2.4. Scaling AI to Commercial Reactors

Transition from pilot-scale demonstrations to large-scale reactors like ITER, DEMO, and next-generation commercial designs. Develop transfer learning and adaptive AI systems that account for scaling effects, complex geometry, and higher energy densities.

8.2.5. Enhanced Computational Efficiency

Develop efficient surrogate models, parallelized AI algorithms, and HPC/quantum computing integration to manage computational demands. Reduce reliance on time-intensive simulations, lowering operational costs and shortening development timelines.

8.3. Opportunities for Sustainability and Commercialization

Energy Efficiency: AI reduces wasted energy in experimental campaigns, accelerates plasma optimization, and shortens fusion development timelines. Economic Feasibility: Optimized design, predictive maintenance, and automated control reduce R&D and operational costs, improving the business case for fusion. Safety and Reliability: AI-driven disruption prediction, anomaly detection, and digital twins enhance reactor resilience, contributing to long-term sustainability and public trust. Global Collaboration: AI can harmonize approaches across public and private sectors, accelerating the adoption of fusion energy worldwide.

8.3.1. Short-Term (2025–2028)

In the next three to five years, the focus will be on expanding the deployment of AI technologies in pilot reactors and mid-scale tokamaks. During this period, AI systems are expected to transition from experimental demonstrations to more routine integration in operational campaigns. This includes real-time plasma control, predictive maintenance, and diagnostic data analysis, building on the successes demonstrated in projects such as DeepMind–TCV and AI-assisted JET/PPPL disruption prediction. By implementing AI in mid-scale devices, researchers can gather valuable operational data, refine algorithms under realistic conditions, and validate predictive and control models before scaling to larger reactors like ITER or DEMO. A key priority in this timeframe will be the development and validation of hybrid AI-physics models. Unlike purely data-driven approaches, these hybrid models integrate first-principles physics with machine learning, improving both the accuracy and interpretability of predictions. Such models will enable more reliable plasma simulations, optimization of magnetic confinement, and accurate forecasting of instabilities, while requiring fewer experimental runs. These hybrid approaches will also serve as a bridge between simulation-based AI training and real-world reactor deployment, mitigating risks associated with purely black-box AI systems. Another critical focus area is the establishment of initial standards for AI safety, interpretability, and data sharing. As AI assumes a more prominent role in reactor operation, stakeholders including national laboratories, private fusion companies, and regulatory bodies will need clear guidelines to ensure that AI decisions are reliable, transparent, and auditable. Standardized frameworks for data collection, preprocessing, and sharing will also be essential to enable cross-device learning and the development of robust AI models that can generalize across different reactor designs. Early adoption of these standards will not only improve safety and trust but also accelerate collaborative research and innovation across the global fusion community. By focusing on these short-term objectives, expanding AI deployment, refining hybrid models, and implementing standards the fusion research community can build a strong foundation for more advanced AI integration in the medium and long term. This phase will be crucial for translating experimental successes into operational improvements and preparing for the eventual scale-up to commercial, sustainable fusion power plants.

8.3.2. Medium-Term (2028–2032)

Between 2028 and 2032, AI integration in nuclear fusion is expected to extend across the entire reactor lifecycle, from design and construction to operation, maintenance, and eventual decommissioning. Predictive maintenance will become a standard component of reactor operations, with AI algorithms continuously monitoring the condition of critical systems such as superconducting magnets, cryogenic units, and plasma-facing components. By analyzing sensor data in real time, AI will anticipate potential failures, schedule maintenance proactively, and reduce costly unplanned downtime. In parallel, autonomous operation systems will emerge, capable of adjusting plasma configurations, heating profiles, and fuel injection parameters with minimal human intervention, thereby improving operational efficiency and reliability. A central goal during this period is the demonstration of AI-assisted digital twins for large-scale experiments, including ITER and other next-generation reactors. Digital twins virtual replicas of physical reactors will enable researchers to simulate operational scenarios, test control strategies, and evaluate design modifications in a risk-free virtual environment. By coupling AI algorithms with these digital twins, operators can optimize performance, reduce experimental iterations, and accelerate learning from each plasma pulse. This approach will also facilitate cross-device knowledge transfer, allowing insights gained from one reactor to inform operations at others. Another critical focus will be on scaling reinforcement learning (RL) and machine learning (ML) models to handle multi-device, high-energy operations. Unlike single-device pilot experiments, large-scale reactors involve more complex coil geometries, higher plasma energies, and increased system interdependencies. Scalable AI models will need to manage these complexities, ensuring stable plasma confinement, predicting instabilities, and supporting coordinated multi-device experimentation. By addressing these challenges, the fusion community can move closer to demonstrating sustained, high-performance plasma operation and laying the groundwork for commercially viable fusion energy.

8.3.3. Long-Term (2032+)

Beyond 2032, the vision for AI in nuclear fusion shifts toward fully autonomous, commercial-scale reactors. At this stage, AI systems will not only manage plasma control and disruption mitigation but will oversee the complete operation of reactors, including startup, shutdown, and real-time adjustment of power output. Reinforcement learning, digital twins, and hybrid AI-physics models will operate seamlessly together, enabling reactors to achieve sustained, stable fusion reactions with minimal human intervention. This autonomy will significantly reduce operational risk, improve efficiency, and allow continuous energy production at large scales. In parallel, AI will play a pivotal role in guiding material discovery, fuel cycle management, and operational optimization at gigawatt (Sankar et al.) scales. Machine learning and generative AI models will accelerate the identification of advanced materials capable of withstanding extreme neutron fluxes, high temperatures, and tritium exposure. Similarly, AI-driven simulations will optimize tritium breeding, fuel injection, and plasma confinement strategies, maximizing energy yield and reactor longevity. By leveraging AI across these critical areas, fusion reactors will become more cost-effective, reliable, and environmentally sustainable. Finally, the long-term goal is the widespread global adoption of AI-enhanced fusion energy systems, forming a cornerstone of the future sustainable energy landscape. Integrated AI systems will facilitate the deployment of fusion power in diverse regions, support grid integration with variable renewable sources, and ensure operational safety and resilience. By combining AI-driven optimization with advanced fusion technology, nations can achieve carbon-neutral, reliable energy at scale, accelerating the global energy transition and contributing to energy security, climate goals, and sustainable development worldwide.

9. Conclusion

Artificial Intelligence (AI) is revolutionizing nuclear fusion research by accelerating design, optimizing operations, and enhancing safety across experimental, predictive, and commercial domains, thereby advancing the goal of sustainable fusion energy. Key achievements include real-time plasma control and disruption prediction using reinforcement learning and machine learning (DeepMind–TCV, JET/PPPL), rapid diagnostics and materials innovation through image recognition, sensor fusion, and generative models for radiation-resistant components and tritium breeding blankets, and commercial applications where private companies (CFS, TAE Technologies, Helion Energy) leverage AI to streamline design, reduce costs, and improve operational efficiency. AI also supports sustainability and safety by minimizing energy waste and enhancing reactor reliability. Despite these advances, challenges remain in data availability, model interpretability, real-time deployment, regulatory compliance, and scalability, requiring multi-disciplinary collaboration. Looking ahead, integrating AI into digital twins, autonomous reactor operations, hybrid AI-physics modeling, and predictive maintenance offers the potential to shorten development timelines, lower costs, and establish nuclear fusion as a viable, globally sustainable energy source.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Abdou, M. A., Riva, M., Ying, A. Y., Day, C., Loarte, A., Baylor, L. R., Humrickhouse, P. W., Fuerst, T. F., & Cho, S. (2020). Physics and technology considerations for the deuterium–tritium fuel cycle and conditions for tritium fuel self sufficiency. *Nuclear Fusion*, *61*.

[2] Adewoyin, M. A., Adediwin, O., & Audu, A. J. (2025). Artificial Intelligence and Sustainable Energy Development: A Review of Applications, Challenges, and Future Directions. *International Journal of Multidisciplinary Research and Growth Evaluation.*

[3] Agarwal, A., Mishra, A., Sharma, P., Jain, S., Daniel, R., Ranjan, S., Manchanda, R., Ghosh, J., Tanna, R., & Team, A. (2021). Deep sequence to sequence learning-based prediction of major disruptions in ADITYA tokamak. *Plasma Physics and Controlled Fusion*, *63*(11), 115004.

[4] Aymerich, E., Sias, G., Pisano, F., Cannas, B., Carcangiu, S., Sozzi, C., Stuart, C., Carvalho, P., Fanni, A., & Contributors, J. (2022). Disruption prediction at JET through deep convolutional neural networks using spatiotemporal information from plasma profiles. *Nuclear Fusion*, *62*(6), 066005.

[5] Badini, S., Regondi, S., & Pugliese, R. (2023). Unleashing the power of artificial intelligence in materials design. *Materials*, *16*(17), 5927.

[6] Battye, M. I., & Perinpanayagam, S. (2025). Digital Twins in Fusion Energy Research: Current State and Future Directions. *IEEE Access*, *13*, 75787–75821.

[7] Bhatia, N., Costa, R., Pamela, S. J. P., Davis, A., Gonzalez Beltran, A. N., Gopakumar, V., Zanisi, L., de Witt, S., & Akers, R. (2025). Visualizing digital twins of fusion power plants using NVIDIA Omniverse. *AIP Advances*.

[8] Blakseth, S. S., Rasheed, A., Kvamsdal, T., & San, O. (2022). Combining physics-based and data-driven techniques for reliable hybrid analysis and modeling using the corrective source term approach. *Appl. Soft Comput.*, *128*, 109533.

[9] Bustos, A. D., Zarzoso, D., Cappa, Á., Estrada, T., & Ascasibar, E. (2025). An AI-based System to Assist Session Leader during Stellarator Operations - a Prototype. *Plasma Physics and Controlled Fusion*.

[10] Carey, N., Zanisi, L., Pamela, S., Gopakumar, V., Omotani, J., Buchanan, J., & Brandstetter, J. (2024). Data efficiency and long term prediction capabilities for neural operator surrogate models of core and edge plasma codes.

[11] Chapman, I. T., Bestwick, T., & Methven, P. (2023). Public–private partnership in the UK fusion program. Physics of Plasmas, 30(10). https://doi.org/10.1063/5.0169796

[12] Churchill, R., Tobias, B., Zhu, Y., & team, D.-D. (2020). 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).

[13] Citrin, J., Trochim, P., Goerler, T., Pfau, D., van de Plassche, K., & Jenko, F. (2023). Fast transport simulations with higher-fidelity surrogate models for ITER. Physics of Plasmas, 30(6).

[14] Clavier, B., Zarzoso, D., del-Castillo-Negrete, D., & Frénod, E. (2025). A generative artificial intelligence framework for long-time plasma turbulence simulations. Physics of Plasmas.

[15] Clery, D. (2022). European fusion reactor sets record for sustained energy. Science, 375(6581), 600.

[16] Croonen, J., Amaya, J., & Lapenta, G. (2020). Tokamak disruption prediction using different machine learning techniques. arXiv: Plasma Physics.

[17] Degrave, J., Felici, F., Buchli, J., Neunert, M., Tracey, B., Carpanese, F., Ewalds, T., Hafner, R., Abdolmaleki, A., de las Casas, D., Donner, C., Fritz, L., Galperti, C., Huber, A., Keeling, J., Tsimpoukelli, M., Kay, J., Merle, A., Moret, J.-M.,…Riedmiller, M. (2022). Magnetic control of tokamak plasmas through deep reinforcement learning. Nature, 602(7897), 414–419. https://doi.org/10.1038/s41586-021-04301-9

[18] Dong, G., Wei, X., Bao, J., Brochard, G., Lin, Z., & Tang, W. (2021). Deep learning based surrogate models for first-principles global simulations of fusion plasmas. Nuclear Fusion, 61(12), 126061.

[19] Farea, A., Yli-Harja, O., & Emmert-Streib, F. (2024). Understanding physics-informed neural networks: Techniques, applications, trends, and challenges. AI, 5(3), 1534–1557.

[20] Ferreira, D. (2018). Applications of deep learning to nuclear fusion research. arXiv e-prints. arXiv preprint arXiv:1811.00333.

[21] Ferreira, D. R. (2018). Applications of Deep Learning to Nuclear Fusion Research. arXiv preprint arXiv:1811.00333. https://doi.org/10.48550/arXiv.1811.00333

[22] Fischer, U., Boccaccini, L. V., Cismondi, F., Coleman, M., Day, C., Hörstensmeyer, Y. N., Moro, F., & Pereslavtsev, P. (2020). Required, achievable and target TBR for the European DEMO. Fusion Engineering and Design.

[23] Galletti, G., Paischer, F., Setinek, P., Hornsby, W., Zanisi, L., Carey, N., Pamela, S., & Brandstetter, J. (2025). 5D Neural Surrogates for Nonlinear Gyrokinetic Simulations of Plasma Turbulence. ArXiv, abs/2502.07469.

[24] Garola, A. R., Cavazzana, R., Gobbin, M., Delogu, R. S., Manduchi, G., Taliercio, C., & Luchetta, A. (2020). Diagnostic Data Integration Using Deep Neural Networks for Real-Time Plasma Analysis. IEEE Transactions on Nuclear Science, 68, 2165–2172.

[25] Goldston, R. J. (2024). An overview of the fusion landscape. Bulletin of the Atomic Scientists, 80, 365 – 370.

[26] Gopakumar, V., Pamela, S., Zanisi, L., Li, Z., Gray, A., Brennand, D., Bhatia, N., Stathopoulos, G., Kusner, M., & Deisenroth, M. P. (2024). Plasma surrogate modelling using Fourier neural operators. Nuclear Fusion, 64(5), 056025.

[27] Haque, R., Bajwa, A., Siddiqui, N. a., & Ahmed, I. (2024). PREDICTIVE MAINTENANCE IN INDUSTRIAL AUTOMATION: A SYSTEMATIC REVIEW OF IOT SENSOR TECHNOLOGIES AND AI ALGORITHMS. American Journal of Interdisciplinary Studies.

[28] Huang, Q., Peng, S., Deng, J., Zeng, H., Zhang, Z., Liu, Y., & Yuan, P. (2023). A review of the application of artificial intelligence to nuclear reactors: Where we are and what's next. Heliyon, 9.

[29] Humphreys, D., Kupresanin, A., Boyer, M. D., Canik, J., Chang, C., Cyr, E. C., Granetz, R., Hittinger, J., Kolemen, E., & Lawrence, E. (2020). Advancing fusion with machine learning research needs workshop report. Journal of Fusion Energy, 39(4), 123–155.

[30] Huynh, N. N. T., Garambois, P.-A., Renard, B., Colleoni, F., Monnier, J., & Roux, H. (2025). A distributed hybrid physics–AI framework for learning corrections of internal hydrological fluxes and enhancing high-resolution regionalized flood modeling. Hydrology and Earth System Sciences, 29(15), 3589–3613.

[31] Jendoubi, C., & Asad, A. (2024). A Survey of Artificial Intelligence Applications in Nuclear Power Plants. IoT, 5(4), 666–691.

[32] Kaur, S., Kumar, R., Singh, K., & Huang, Y. (2024). Leveraging artificial intelligence for enhanced sustainable energy management. Journal of Sustainable Energy, 3(1), 1–20.

[33] Kerboua-Benlarbi, S., Nouailletas, R., Faugeras, B., Nardon, E., & Moreau, P. (2024). Magnetic Control of WEST Plasmas Through Deep Reinforcement Learning. IEEE Transactions on Plasma Science, 52, 3698–3703.

[34] Khalifa, M., & Albadawy, M. (2024). AI in Diagnostic Imaging: Revolutionising Accuracy and Efficiency. Computer Methods and Programs in Biomedicine Update.

[35] Kim, J., & Seo, J. (2024). Design Optimization of Nuclear Fusion Reactor through Deep Reinforcement Learning. arXiv preprint arXiv:2409.08231.

[36] Li, H., Wang, L., Fu, Y.-L., Wang, Z.-X., Wang, T., & Li, J. (2024). Surrogate model of turbulent transport in fusion plasmas using machine learning. Nuclear Fusion, 65.

[37] Li, Z. (2025). A Review of Physics-Informed Neural Networks. *Applied and Computational Engineering*.

[38] Li 李, H. 慧., Fu 付, Y.-L. 艳., Li 李, J. 继., & Wang 王, Z. 正. (2021). Machine learning of turbulent transport in fusion plasmas with neural network. *Plasma Science and Technology*, *23*.

[39] Lindley, B. A., Roulstone, T., Locatelli, G., & Rooney, M. (2023). Can fusion energy be cost-competitive and commercially viable? An analysis of magnetically confined reactors. *Energy Policy*.

[40] Mánek, P., Van Goffrier, G., Gopakumar, V., Nikolaou, N., Shimwell, J., & Waldmann, I. (2023). Fast regression of the tritium breeding ratio in fusion reactors. *Machine Learning: Science and Technology*, *4*(1), 015008.

[41] Manheimer, W. (2020). ITER and the prospects for commercial fusion. Physics Today, 73(7), 10–10. https://doi.org/10.1063/pt.3.4510

[42] Meschini, S., Ferry, S. E., Delaporte-Mathurin, R., & Whyte, D. G. (2023). Modeling and analysis of the tritium fuel cycle for ARC-and STEP-class DT fusion power plants. Nuclear Fusion, 63(12), 126005.

[43] Mohamed, M., Zakuan, N. D., Tengku Hassan, T. N. A., Lock, S. S. M., & Mohd Shariff, A. (2024). Global development and readiness of nuclear fusion technology as the alternative source for clean energy supply. Sustainability, 16(10), 4089.

[44] Mohapatra, A. (2024). Generative AI for Predictive Maintenance: Predicting Equipment Failures and Optimizing Maintenance Schedules Using AI. International Journal of Scientific Research and Management (IJSRM).

[45] Nadaf, S. (2024). AI for predictive maintenance in industries. Int J Res Appl Sci Eng Technol, 12(6), 2013–2017.

[46] Nocente, M., Kiptily, V., Tardocchi, M., Bonofiglo, P. J., Craciunescu, T., Molin, A. D., De La Luna, E., Eriksson, J., Garcia, J., Ghani, Z., Gorini, G., Hägg, L., Kazakov, Y., Lerche, E., Maggi, C. F., Mantica, P., Marcer, G., Maslov, M., Putignano, O.,…Perez von Thun, C. (2022). Fusion product measurements by nuclear diagnostics in the Joint European Torus deuterium–tritium 2 campaign (invited). Review of Scientific Instruments, 93(None; Journal ID: ISSN 0034-6748;101052200). https://doi.org/10.1063/5.0101767

[47] Noviati, N. D., Maulina, S. D., & Smith, S. (2024). Smart Grids: Integrating AI for Efficient Renewable Energy Utilization. International Transactions on Artificial Intelligence (ITALIC).

[48] Ongena, J. (2022). Fusion: a true challenge for an enormous reward. EPJ Web of Conferences.

[49] Paltsev, S., Gurgel, A., Morris, J., Chen, H., Dey, S., & Marwah, S. (2022). Economic analysis of the hard-to-abate sectors in India. Energy Economics, 112, 106149. https://doi.org/https://doi.org/10.1016/j.eneco.2022.106149

[50] Pavone, A., Svensson, J., Krychowiak, M., Hergenhahn, U., Winters, V., Kornejew, P., Kwak, S., Hoefel, U., Koenig, R., & Wolf, R. (2021). Neural network surrogates of Bayesian diagnostic models for fast inference of plasma parameters. Review of Scientific Instruments, 92(3).

[51] Poels, Y., Derks, G., Westerhof, E., Minartz, K., Wiesen, S., & Menkovski, V. (2023). Fast dynamic 1D simulation of divertor plasmas with neural PDE surrogates. Nuclear Fusion, 63(12), 126012.

[52] Priyanka, M., Sangeetha, J., & Jayakumar, C. (2024). A review of traditional and data-driven approaches for disruption prediction in different tokamaks. E3S Web of Conferences,

[53] Qu, S., Cao, Q., Duan, X., Wang, X., & Wang, X. (2021). Study on multiphysics coupling and automatic neutronic optimization for solid tritium breeding blanket of fusion reactor. Energies, 14(17), 5442.

[54] Rattá, G., Vega, J., Murari, A., Gadariya, D., & Contributors, J. E. T. (2021). PHAD: a phase-oriented disruption prediction strategy for avoidance, prevention, and mitigation in JET. Nuclear Fusion, 61.

[55] Ren, Z., Zhou, S., Liu, D., & Liu, Q. (2025). Physics-Informed Neural Networks: A Review of Methodological Evolution, Theoretical Foundations, and Interdisciplinary Frontiers Toward Next-Generation Scientific Computing. Applied Sciences, 15(14), 8092.

[56] Rossi, R., Gelfusa, M., Craciunescu, T., Wyss, I., & Vega, J. (2024). A hybrid physics/data-driven logic to detect, classify, and predict anomalies and disruptions in tokamak plasmas. Nuclear Fusion, 64(4), 046017.

[57] Samaddar, A., Gong, Q., Madireddy, S., Hansen, C. J., Joung, S., Smith, D. R., Sun, Y., Ebrahimi, F., Balapraksh, P., & Nelson, A. O. (2025). Spatiotemporal forecasting of the edge localized modes in tokamak plasmas using neural networks. Machine Learning: Science and Technology.

[58] San, O., Rasheed, A., & Kvamsdal, T. (2021). Hybrid analysis and modeling, eclecticism, and multifidelity computing toward digital twin revolution. GAMM-Mitteilungen, 44(2), e202100007.

[59] Sankar, B. S., Gilliland, D., Rincon, J., Hermjakob, H., Yan, Y., Adam, I., Lemaster, G., Wang, D., Watson, K., & Bui, A. (2024). Building an ethical and trustworthy biomedical AI ecosystem for the translational and clinical integration of foundation models. Bioengineering, 11(10), 984.

[60] Schissel, D. P., Nazikian, R. M., & Gibbs, T. (2025). Summary report from the mini-conference on Digital Twins for Fusion Research. Physics of Plasmas.

[61] Seo, J., Kim, I. H., & Nam, H. (2024). Leveraging physics-informed neural computing for transport simulations of nuclear fusion plasmas. Nuclear Engineering and Technology, 56(12), 5396–5404. https://doi.org/https://doi.org/10.1016/j.net.2024.07.048

[62] Sobes, V., Hiscox, B., Popov, E., Archibald, R., Hauck, C., Betzler, B., & Terrani, K. (2021). AI-based design of a nuclear reactor core. Scientific reports, 11(1), 19646.

[63] Stiasny, J., Chevalier, S. C., & Chatzivasileiadis, S. (2021). Learning without Data: Physics-Informed Neural Networks for Fast Time-Domain Simulation. *2021 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 438–443.

[64] Subbotin, G. F., Sorokin, D. I., Nurgaliev, M. R., Granovskiy, A. A., Kharitonov, I. P., Adishchev, E. V., Khairutdinov, E. N., Clark, R., Shen, H., Choi, W., Barr, J., & Orlov, D. M. (2025). Reconstruction-free magnetic control of DIII-D plasma with deep reinforcement learning.

[65] Tang, W., Dong, G., Barr, J., Erickson, K., Conlin, R., Boyer, D., Kates-Harbeck, J., Felker, K., Rea, C., & Logan, N. (2023). Implementation of AI/DEEP learning disruption predictor into a plasma control system. *Contributions to Plasma Physics*, *63*(5-6), e202200095.

[66] Temel, D. (2024). Will Nuclear Fusion Pave the Way for Widespread Sustainable Energy Use? *Next Generation Journal for The Young Researchers*.

[67] Tindall, M., Rosini, S., Bowden, N., Deighan, T., Mavis, M., Biggs-Fox, S., & Barrett, T. (2023). Towards a fusion component digital twin – virtual test and monitoring of components in CHIMERA by systems simulation. *Fusion Engineering and Design*.

[68] Tracey, B. D., Michi, A., Chervonyi, Y., Davies, I., Paduraru, C., Lazic, N., Felici, F., Ewalds, T., Donner, C., Galperti, C., Buchli, J., Neunert, M., Huber, A., Evens, J., Kurylowicz, P., Mankowitz, D. J., Riedmiller, M. A., & Team, T. T. (2023). Towards practical reinforcement learning for tokamak magnetic control. *ArXiv*, *abs/2307.11546*.

[69] Ucar, A., Karakose, M., & Kırımça, N. (2024). Artificial intelligence for predictive maintenance applications: key components, trustworthiness, and future trends. *Applied Sciences*, *14*(2), 898.

[70] Udo, W. S., Kwakye, J. M., Ekechukwu, D. E., & Ogundipe, O. B. (2023). SMART GRID INNOVATION: MACHINE LEARNING FOR REAL-TIME ENERGY MANAGEMENT AND LOAD BALANCING. *Engineering Science & Technology Journal*.

[71] Wang, Z., Li, J., Liu, Y., Li, X., Yan, C., & Zhang, Y. (2025). RL–Fusion: The Large Language Model Fusion Method Based on Reinforcement Learning for Task Enhancing. *Applied Sciences*, *15*(4), 2186.

[72] Wiesen, S., Dasbach, S., Kit, A., Jaervinen, A., Gillgren, A., Ho, A., Panera, A., Reiser, D., Brenzke, M., & Poels, Y. (2024). Data-driven models in fusion exhaust: AI methods and perspectives. *Nuclear Fusion*, *64*(8), 086046.

[73] Wurzel, S. E., & Hsu, S. C. (2022). Progress toward fusion energy breakeven and gain as measured against the Lawson criterion. Physics of Plasmas, 29(6). https://doi.org/10.1063/5.0083990

[74] ZHENG 郑, W. 伟., XUE 薛, F. 凤., Shen 沈, C. 呈., Zhong 钟, Y. 昱., Ai 艾, X. 鑫., CHEN 陈, Z. 忠., Ding 丁, Y. 永., Zhang 张, M. 明., Yang 杨, Z. 州., Wang 王, N. 能., ZHANG 张, Z. 智., DONG 董, J. 蛟., TANG 唐, C. 畴., & Pan 潘, Y. 垣. (2022). Overview of machine learning applications in fusion plasma experiments on J-TEXT tokamak. *Plasma Science and Technology*, *24*.

[75] Zinkle, S. J., & Quadling, A. (2022). Extreme materials environment of the fusion "fireplace". *MRS Bulletin*, *47*, 1113–1119.

[76] Zuccarini, C., Ramachandran, K., & Jayaseelan, D. D. (2024). Material discovery and modeling acceleration via machine learning. *APL Materials*, *12*(9).