Application of Artificial Intelligence in Cold Nuclear Fusion

Todor Malchev

Abstract— **Cold Nuclear Fusion, also known as Low Energy Nuclear Reactions (LENR), has attracted renewed attention as a potential source of clean and sustainable energy. Despite decades of research, reproducibility and theoretical understanding remain major challenges. Artificial Intelligence (AI) and Machine Learning (ML) have recently emerged as transformative tools across many scientific fields, capable of extracting patterns from complex data, guiding experiments, and accelerating discovery. This paper provides an overview of how AI techniques ranging from classical machine learning models to deep learning and generative methods are being applied to Cold Nuclear Fusion research. It discusses progress in data-driven modeling, simulation, parameter optimization, anomaly detection, and material discovery, drawing parallels to applications in conventional nuclear fusion and plasma physics. The paper concludes with future directions for integrating AI into Cold Nuclear Fusion experimental workflows, theoretical modeling, and autonomous laboratory systems.**

Keywords—**Cold Fusion, LENR, Machine Learning, Deep Learning, Synthetic Data, Parameter Optimization**

I. INTRODUCTION

Low Energy Nuclear Reactions (LENR) refer to nuclear-scale energy processes occurring under conditions far milder than those required for conventional fusion. Since the 1989 announcement by Fleischmann and Pons, LENR has remained controversial due to inconsistent reproducibility and lack of a widely accepted theoretical model. Nevertheless, experimental efforts have continued, with researchers reporting anomalous heat production, isotope shifts, and other nuclear signatures [1].

Cold Nuclear Fusion research explores the possibility of nuclear reactions occurring at or near room temperature, typically involving hydrogen or deuterium interacting with metal lattices such as palladium, nickel, or constantan. While experimental evidence has occasionally suggested excess heat or nuclear byproducts, reproducibility issues and insufficient theoretical explanations have hindered mainstream acceptance.

AI excels at uncovering nonlinear correlations in noisy, high-dimensional data – a hallmark of LENR experimentation. Recent advances in Artificial Intelligence (AI) provide new opportunities to revisit Cold Nuclear Fusion phenomena using data-driven approaches that can extract insights from complex, nonlinear, and often noisy datasets. The increasing availability of sensor data (temperature, pressure, gas flow, radiation, etc.) creates opportunities for AI-assisted data analysis, modeling, and experiment control.

In nuclear fusion research, AI has already demonstrated its power in plasma control, predictive maintenance, and diagnostic interpretation. The same principles can be adapted to Cold Nuclear Fusion systems.

AI methods such as neural networks, reinforcement learning, and generative models have proven highly effective in other domains of nuclear research, including plasma control and material discovery for fusion reactors [2, 3, 4]. These advances motivate their adaptation to LENR, where data scarcity, nonlinearity, and hidden dependencies between parameters make conventional modeling challenging.

II. ARTIFICIAL INTELLIGENCE IN NUCLEAR FUSION RESEARCH

AI encompasses computational methods that enable machines to learn from data and make predictions or decisions without explicit programming. ML is a subset of AI focused on pattern recognition and predictive modeling, while Deep Learning (DL) extends ML using neural networks with multiple layers that capture hierarchical representations of data [5]. In nuclear fusion research, AI has achieved notable successes in controlling plasma behavior, predicting disruptions, and optimizing magnetic confinement conditions.

AI is successfully applied in high-temperature (hot) fusion systems:

Plasma control and stability: deep reinforcement learning has been used to maintain plasma equilibrium in tokamaks.

Predictive maintenance and diagnostics: AI models predict equipment failures and plasma disruptions.

Data analysis and simulation: Machine learning accelerates the interpretation of experimental diagnostics and simulates plasma dynamics.

Material discovery: ML-driven materials science identifies alloys and composites resilient to extreme heat and radiation.

A primary driver for AI in fusion is the need to manage the massive data streams from diagnostics and the extreme complexity of plasma physics. For instance, predicting plasma disruptions-sudden terminations of the plasma that can severely damage the reactor-is a critical challenge.

AI models, particularly deep learning approaches like Recurrent Neural Networks (RNNs), are trained on data from past experiments (e.g., from the JET or DIII-D tokamaks) to recognize subtle patterns that precede a disruption, enabling real-time mitigation systems.

Todor Malchev is with the Technical University of Sofia, Sofia, 1000, 8 Kl. Ohridski Blvd, Bulgaria (t.o.malchev@gmail.com)

Beyond safety, AI accelerates progress by creating "surrogate models". These models emulate computationally expensive, high-fidelity physics simulations, allowing researchers to explore vast parameter spaces for optimal reactor designs and control strategies much faster than traditional methods [6].

These achievements form the conceptual and methodological foundation for AI applications in LENR. For example, DeepMind and the École Polytechnique fédérale de Lausanne (EPFL) demonstrated in 2022 that deep reinforcement learning could control plasma configurations in a tokamak reactor in real-time [7]. Similarly, institutions such as ITER and EUROfusion employ ML techniques for diagnostics, predictive maintenance, and reactor simulation [8, 9]. These achievements form the conceptual and methodological foundation for extending AI to LENR research.

The application of AI in LENR is still emerging, with most research focusing on experimental data analysis, simulation, and parameter optimization.

III. APPLICATIONS OF AI IN LOW ENERGY NUCLEAR REACTIONS (LENR)

The experiments conducted at Lakehead University investigate Cold Nuclear Fusion in solid-state systems by examining the interaction of deuterium gas with metal lattices such as constantan and palladium (Fig. 1). One of the key contributions of the study is the establishment of reproducibility, thereby providing an opportunity to advance our understanding of the conditions under which Cold Nuclear Fusion may be reliably observed. The results are very promising and are subject of further research. The recorded dataset includes temperature, pressure, gas flow volume, gas injection period, relaxation time and others [10].

Neural networks are used to model the multivariate relationships among these factors to identify conditions that maximize anomalous heat signals and other nuclear signatures, optimize operational parameters for repeatable outcomes, and achieving higher temperature yield.
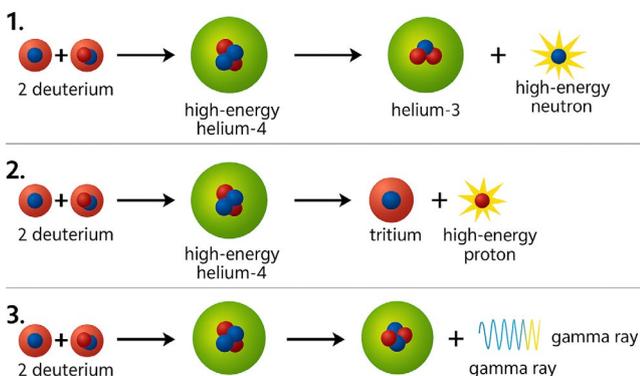
Three deuterium fusion reactions:

Fig 1. Deuterium fusion reactions

This research presents a three-stage framework combining synthetic time-series generation, parameter optimization, and model-based temperature forecasting (Fig. 2). A TimeGAN-based module learns temporally coherent representations and generates multivariate sequences that augment limited observations. The generator is trained in phased fashion (embedding/recovery, supervision, and joint objectives), and synthetic outputs are normalized-compatible with the real pipeline, enabling analysis and export alongside real data.

A parameter prediction and optimization stage employs an LSTM surrogate that consumes historical sequences together with impulse-like inputs. It is trained on a blend of real and synthetic data with configurable weighting to leverage augmentation while controlling its influence. The stage includes a batch-wise optimization routine that directly searches for parameter vectors consistent with observed dynamics, with optional bounds for constrained optimization and utilities to track optimization history.

The final temperature forecasting stage is a sequence-to-one LSTM that predicts the terminal temperature of each window from historical multivariate signals. When available, the optimized parameters are provided as conditioning context to improve accuracy and stability. The predictor yields point estimates (not probabilistic outputs), supports standard validation splits, and provides utilities for learning-curve visualization, performance evaluation via mean-squared error, and denormalization back to the physical scale. Collectively, the pipeline enables data augmentation, parameter-space search guided by a learned surrogate, and temperature forecasting with optional parameter conditioning. It supports validation on held-out data and facilitates comparison of training with versus without synthetic augmentation.
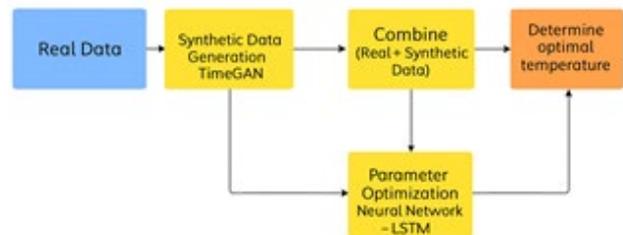
Fig 2. Framework flow

A. Data-Driven Modeling of LENR Experiments

Artificial intelligence provides an empirical modeling framework for experimental systems, such as those exhibiting cold nuclear fusion effects, where a complete, first-principles physics model is unavailable.

Neural networks serve as robust, nonlinear function approximators, adept at mapping the complex relationships among critical operational parameters—like temperature, input power, pressure, and others. When trained on experimental data, these models identify subtle, non-obvious correlations between environmental and material factors and the generation of excess heat.

This capability transforms the model into a powerful predictive tool for experimental optimization. By applying different algorithms to the provided data, researchers can efficiently explore the vast parameter space, prioritizing novel combinations of conditions (e.g., specific thermal or pressure protocols) predicted to maximize the desired heat output (maximize heat yield). This greatly accelerates the research cycle.

B. Synthetic Data Generation

Data scarcity is a persistent obstacle in LENR due to limited experimental replication, limited funding and others.

Having a large, diverse, and high-quality dataset is crucial for the successful application of machine learning, as it improves models' generalization ability, reduces the risk of overfitting, and enables more reliable performance evaluation. Synthetic data can improve model robustness and enable supervised training where experimental data are limited.

To address the fundamental challenge of dataset sparsity, advanced generative architectures such as Generative Adversarial Networks (GANs), the temporally focused TimeGAN, and Variational Autoencoders (VAEs) have been proposed.

These models are leveraged to augment limited experimental archives by producing high-fidelity synthetic data that is statistically indistinguishable from the original measurements. This process is engineered to not only replicate static data points but also to successfully emulate the complex temporal dynamics and sequential dependencies of the source.

TimeGAN (Time-series Generative Adversarial Network) is a specialized deep learning architecture designed to overcome the limitations of standard Generative Adversarial Networks (GANs) when handling sequential data. It generates highly realistic synthetic time-series data by not only capturing the overall statistical distribution of the dataset but, critically, by preserving the inherent temporal dynamics (the sequential dependencies and trends over time). TimeGAN achieves this through a novel framework that combines a four-component architecture—an Embedder and Recovery network (forming an Autoencoder) with a standard Generator and Discriminator—and introduces a key innovation: a supervised loss that explicitly forces the Generator to mimic the step-by-step conditional transitions observed in the real data.

A modified TimeGAN architecture with an added pre-generator appears most appropriate for this study (Fig. 3). The pre-generator boosts the realism of synthetic time series and improves the effectiveness of subsequent LSTM-based parameter optimization.

Fig. 3. Structure of the TimeGAN model

C. Parameter Optimization and Control

To optimize complex experimental systems, researchers are increasingly turning to advanced AI architectures that combine Long Short-Term Memory (LSTM) neural networks with dense layers (Fig. 4).

The network could be trained on both real and synthetic data. Applying different weights to them, that gives more control and is based on the quality of the generated data. This hybrid model is designed to identify the most suitable input parameters by processing time-series data.

The outputs from both branches are merged and passed through additional dense layers to produce refined parameter predictions. Through iterative loss minimization, the system continuously improves its accuracy, enabling efficient exploration of vast parameter spaces that would be impractical to navigate manually.

This approach not only enhances the identification of optimal conditions correlated with increased heat output and system stability but also lays the foundation for adaptive control systems.

These systems can respond in real time to changing experimental conditions by adjusting parameters dynamically, making them ideal for environments where external factors can unpredictably affect outcomes.

Optimization algorithms can explore large parameter spaces far more efficiently than manual tuning, reducing development time and uncovering solutions humans might miss.

Fig. 4. LSTM diagram

D. Temperature and Output Prediction

Deep learning models trained on both real and synthetic data can predict temperature dynamics and output power in LENR systems. Incorporating optimized parameters enhances accuracy and may reveal causal relationships between operating conditions and reaction outcomes.

A possible solution is an LSTM-based predictor that models multivariate time-series windows to forecast a single scalar temperature at the end of each sequence. The model ingests fixed-length, normalized sensor and process features and learns temporal dependencies that drive the terminal temperature. Optimized parameters produced by an upstream module can be fused with the sequence representation to provide additional physical or control-context information, improving training stability and accuracy. Training follows standard protocols with configurable epochs, batch size, validation split, and early stopping, and supports MSE and optional quantile losses for probabilistic outputs. The system includes utilities for rescaling predictions to physical units, evaluating MSE on

hold-out data, visualizing learning curves, and saving model weights for reproducibility and deployment. This solution is suitable for manufacturing and laboratory settings where temperature results from interacting dynamic signals and can operate as a purely data-driven forecaster or a hybrid model enhanced by parameter optimization.

E. Data Fusion and Multi-Modal Analysis

Combining complementary measurements from multiple instruments yields a more comprehensive and reliable characterization of low energy nuclear reaction (LENR) behavior than any single sensor alone. Integrating data from diverse instruments (e.g., RGA for gas analysis, EDX for surface composition, calorimetry for heat measurement) enables a comprehensive understanding of LENR behavior.

Appropriate fusion strategies depend on the goal and data volume: feature-level fusion assembles interpretable descriptors (e.g., peak integrals, isotopic ratios, elemental fractions, integrated heat) into unified vectors for statistical analysis, decision-level fusion combines modality-specific detectors for robust event calls, and hybrid schemes use exploratory feature analysis followed by ensemble monitoring.

A suite of AI techniques can be applied explicitly: unsupervised methods (k-means, DBSCAN, PCA, UMAP) to discover recurring multi-modal signatures; supervised classifiers (random forest, XGBoost, SVM, feedforward or convolutional neural nets) for labeled-event detection; time-series models (AR, HMM, LSTM/GRU, Temporal Convolutional Networks) for lead–lag dynamics; deep fusion and representation learning (modality encoders, cross-modal attention, multimodal transformers) to learn joint embeddings; anomaly detectors (isolation forest, autoencoders, variational autoencoders) for rare events; causal and information-theoretic tools (Granger causality, transfer entropy, causal graphs) to probe directional links; and uncertainty quantification (Bayesian neural nets, Gaussian processes, bootstrapping) plus explainability methods (SHAP, LIME, feature importance) to relate model outputs to physical observables.

Validation is mandatory: perform cross-run cross-validation, report probabilistic uncertainty and confidence intervals, and run synthetic injection tests (controlled gas additions, calibrated heat pulses) to assess sensitivity and false-alarm rates. Candidate LENR events should be confirmed by at least two independent modalities, documented with archived raw traces and preprocessing provenance, and reported with explicit statements separating exploratory correlations from causal claims.

IV. CHALLENGES AND LIMITATIONS

Integrating artificial intelligence into Low Energy Nuclear Reactions and cold fusion research presents significant technical hurdles and meaningful scientific promise. LENR experiments are often small-scale, unevenly instrumented, and recorded with limited metadata, producing sparse, noisy, and inconsistent datasets that make off-the-shelf machine learning unreliable and increase the risk of spurious findings. Nonuniform measurement standards, missing provenance, and variable experiment fidelity further undermines reproducibility and complicate model validation.

Despite these limits, AI can add value by surfacing subtle patterns in noisy records, detecting anomalies and equipment drift, and prioritizing high-impact parameter regimes for follow-up. Realizing this requires rigorous data curation, rich metadata and calibration logs, uncertainty-aware models, and physics-informed algorithms that respect known constraints. Coupling these models to automated or semi-automated experimental platforms and establishing shared benchmarks and open repositories will help move AI from retrospective analysis to an active partner in hypothesis generation and reproducible discovery.

A. Data Availability and Quality

One of the foremost challenges in integrating AI into LENR research is the scarcity of reliable and high-quality data. Unlike conventional nuclear fusion research, which benefits from large, well-instrumented facilities such as ITER and JET, LENR experiments are typically conducted at smaller scales with limited instrumentation. Consequently, only a few datasets are publicly available, and these often lack the volume, variety, and veracity necessary for robust machine learning model training.

Moreover, reported LENR data frequently suffer from inconsistencies in measurement accuracy, incomplete metadata, and noise introduced by environmental fluctuations. Inadequate sensor calibration, unstandardized data acquisition rates, and incomplete logging of control parameters further complicate model development. Without access to comprehensive, high-fidelity datasets, AI models struggle to generalize beyond narrow experimental conditions, leading to unreliable predictions and limited scientific reproducibility.

B. Lack of Standardized Experimental Protocols

Another major barrier is the absence of standardized experimental methodologies across research groups. LENR studies often vary in sample preparation, material purity, gas composition, calorimetry techniques, and environmental control, making direct comparison across datasets difficult. These variations create heterogeneous data structures that hinder the training of generalized AI models capable of capturing universal patterns.

The lack of shared data formats, metadata standards, and benchmarking protocols prevents the establishment of large-scale databases suitable for supervised learning. Developing an internationally accepted LENR data standard—similar to those used in genomics or materials science—would significantly enhance the interoperability of datasets and facilitate collaborative AI model development.

C. Overfitting and Model Interpretability

Due to the limited volume of LENR data, AI models and particularly deep neural networks are at risk of overfitting, where the system memorizes training data rather than learning underlying physical relationships. Overfitted models can show high accuracy on known data but fail to predict new or unseen experimental results accurately. This issue is compounded by the inherently noisy and complex nature of LENR phenomena, which often lack clearly defined input–output mappings.

Interpretability also remains a critical issue. Most AI models operate as "black boxes," offering little insight into how input parameters contribute to predicted outcomes. In a field like LENR, where the underlying physics remains contested, understanding the rationale behind AI predictions is as important as achieving numerical accuracy. Methods such as Explainable AI (XAI), SHAP, LIME and Physics-Informed Neural Networks (PINNs) are promising solutions, as they combine data-driven modeling with known physical constraints to improve both transparency and scientific trust. Combining AI with physical modeling may help reconcile data-driven predictions with theoretical frameworks.

D. Skepticism and Funding Constraints

The scientific skepticism surrounding LENR research presents sociotechnical challenges that extend beyond data and modeling. Since the controversy of the Fleischmann–Pons announcement in 1989, LENR has struggled to secure sustained institutional support, leading to fragmented research efforts and limited funding availability. This lack of mainstream backing has slowed progress in acquiring high-quality data, building modern experimental facilities, and developing interdisciplinary collaborations that would facilitate AI integration.

Furthermore, many funding agencies and academic publishers remain cautious about LENR-related proposals and publications, which limits visibility and slows the dissemination of legitimate findings. As AI techniques mature and demonstrate concrete contributions to LENR analysis—such as pattern detection, anomaly recognition, and parameter optimization—greater recognition and support may follow. Nevertheless, overcoming this skepticism requires consistent transparency, rigorous experimental validation, and open scientific collaboration to establish LENR as a credible research frontier for AI applications.

V. Emerging Trends and Future Directions

Emerging trends point to several promising directions for AI-driven LENR research, including physics-informed machine learning to guide theory, active-learning–led experimental design to maximize scarce data value, automated instrumentation for reproducible trials, and generative methods for discovering novel materials and catalysts. These approaches promise faster hypothesis generation, clearer signal extraction from noisy datasets, and more efficient pathways from lab results to robust, independently validated findings.

A. Autonomous Experimentation and Closed-Loop Systems

Autonomous experimentation, leveraging reinforcement learning and robotic control, could enable closed-loop systems that iteratively design, perform, and analyze LENR experiments. Combining AI with robotics and laboratory automation could lead to "self-driving" LENR experiments. Reinforcement learning agents could iteratively design, execute, and evaluate experimental configurations.

Integrating multi-modal data sources such as calorimetry, spectroscopy, and residual gas analysis could improve anomaly detection and parameter correlation.

B. Explainable AI (XAI) in Scientific Discovery

Interpretability is crucial in controversial scientific areas like LENR. XAI tools (e.g., SHAP, LIME) can help link AI predictions with physical parameters, enhancing scientific credibility and hypothesis formation.

C. Large Language Models (LLMs) in LENR Research

The application of Large Language Models (LLMs) trained on scientific literature can accelerate hypothesis generation and experimental design by summarizing prior research, identifying gaps, and suggesting new parameter ranges. These hybrid systems combining symbolic reasoning from LLMs and numerical modeling from neural networks could help bridge the gap between empirical data and theoretical understanding of LENR phenomena.

D. Integration with Theoretical Models

AI can serve as a bridge between phenomenological observations and emerging LENR theories (e.g., lattice confinement, quantum tunneling effects). Hybrid models combining physics-based equations with neural network components offer a balanced approach between interpretability and predictive power [11].

VI. Conclusion

The intersection of AI and LENR represents a novel and promising research frontier. AI offers tools for managing complexity, generating data, and optimizing experimental outcomes in systems where traditional models fall short.

While challenges such as data scarcity and model interpretability persist, growing computational resources, improved experimental instrumentation, and interdisciplinary collaboration are poised to accelerate progress. As AI-driven methodologies mature, they may help reveal the underlying mechanisms of LENR and unlock new avenues toward sustainable energy generation.

Establishing shared data standards, clear benchmarks, and reproducible workflows will be critical to turn AI-driven signals into credible LENR findings. Prioritize physics-informed, interpretable models and routine cross-lab replication to reduce false positives and guide targeted experiments.

References

[1] E. Storms, "Science of low energy nuclear reaction, the: a comprehensive compilation of evidence and explanations about cold fusion," World Scientific, 2007.

[2] C. Rea, "Machine learning applications enabling fusion energy: recent developments," Journal of Fusion Energy, vol. 44, no. 39, 2025.

[3] M. Becerra, J. Pettersson, S. Franke, S. Gortschakow, "Temperature and pressure profiles of an ablation-controlled arc plasma in air," Journal of Physics D: Applied Physics, vol. 52, no. 43, 2019.

[4] D. Humphreys, et al., "Advancing Fusion with Machine Learning Research Needs Workshop Report," Journal of Fusion Energy, vol. 39, pp. 123-155, 2020.

[5] I. Goodfellow, et al., "Deep learning," vol. 1, no. 2, Cambridge: MIT Press, 2016.

[6] A. Bari, S. Garg, M. Z. E. H. Masud, "Exploring artificial intelligence techniques to research low energy nuclear reactions," 2024.

[7] J. Degrave, et al., "Magnetic control of tokamak plasmas through deep reinforcement learning," Nature, vol. 602, pp. 414-419, 2022.

[8] X. Litaudon, et al., "EUROfusion contributions to ITER nuclear operation," Nuclear Fusion, vol. 64, no. 11, 2024.

[9] Q. Huang, et al., "A review of the application of artificial intelligence to nuclear reactors: Where we are and what's next," Heliyon, vol. 9, no. 3, 2023.

[10] D. Alexandrov, "Low energy nuclear fusion reactions in solids: experiments," International Journal of Energy Research, vol. 45, no. 8, 2021, pp. 12234-12246

[11] S. S. J. and H. S, "Physics-informed neural networks for solving the time-dependent Schrödinger equation," Journal of Physics A: Mathematical and Theoretical, vol. 55, no. 21, p. 215302, 2022.