
AI ACROSS THE NUCLEAR LIFECYCLE: CURRENT CAPABILITIES AND THE PATH TO DECISION SUPPORT

Majdi I. Radaideh

Department of Nuclear Engineering and Radiological Sciences
Department of Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109, USA

September 25, 2026

ABSTRACT

Artificial intelligence (AI) methods are being investigated across nuclear engineering, yet evidence from individual demonstrations does not necessarily establish readiness for operational use. This work examines nuclear datasets and computational infrastructure, surrogate and physics-informed modeling, digital twins, monitoring and prognostics, lifecycle applications, optimization and control, and trustworthy AI. Reported studies show how learning-based methods can accelerate high-fidelity calculations, extract information from multivariate measurements, and support decisions in reactor operation, maintenance, waste management, and environmental assessment. They also expose recurring limitations: scarce abnormal-condition data, differences between simulated and physical systems, uncertain generalization, and incomplete evaluation of downstream decisions. We argue that progress depends on assessing complete AI-enabled workflows rather than predictive models alone. Representative data, independent testing, calibrated uncertainty, physical and operational constraints, model traceability, and human oversight are especially important when outputs affect safety-relevant decisions. Future research should connect simulation, sensing, digital twins, and uncertainty-aware decision methods while testing their performance under realistic changes in operating conditions. These steps can help establish where AI provides dependable value across the nuclear lifecycle.

Keywords Artificial intelligence · Nuclear reactors · Physics-informed machine learning · Digital twins · Safety-critical systems · Trustworthy AI

1 Introduction

The transition toward sustainable nuclear systems requires balancing safety, reliability, economic performance, resource utilization, waste reduction, and environmental protection across the nuclear lifecycle. This spans reactor and fuel design, operation, component lifetime management, spent fuel, decommissioning, and environmental stewardship. Nuclear circularity therefore extends beyond material recycling to efficient resource use, justified lifetime extension, waste reduction, and better use of information throughout the lifecycle.

At the same time, nuclear science and engineering are undergoing a digital transformation driven by advances in sensing, high-performance computing (HPC), multiphysics simulation, data management, and artificial intelligence (AI). Learning-based methods are being developed and evaluated for reactor analysis, monitoring, diagnostics, and optimization [1], with applications spanning reactor design, operation, maintenance, and safety [2], including nuclear power plant control [3], as well as radiation measurements and radiological science [4]. More broadly, ML is helping connect complex experimental and computational data with physical inference in nuclear physics [5], and international organizations have identified AI as an enabling technology for nuclear science and technology [6].

The opportunity is not to replace physics-based analysis with data-driven models, but to complement established mechanistic models, domain knowledge, and high-fidelity simulations. AI can extract information from complex data,

accelerate expensive calculations, connect simulations with measurements, and support decisions under uncertainty. The integration of simulation, deep learning, and uncertainty quantification (UQ) demonstrates this physics–data connection [7], while ML and surrogate models can further support calibration of uncertain model parameters using experimental or computational information [8].

Figure 1 summarizes the framework adopted in this work: nuclear training data are transformed into models, predictions, quantified uncertainty, decisions, and actions. Experimental, simulation, sensor, and operational data characterize the physical system, while physics-based and data-driven models estimate its state, performance, and health. UQ estimates predictive uncertainty, while validation and applicability assessment help determine how much confidence to place in predictions before they inform optimization, control, planning, or other engineering decisions. The resulting actions generate new information that feeds back into the data and models, enabling continued learning and improvement. NRC-sponsored research has similarly examined AI and ML for operating nuclear plants in the context of computational tools, plant information, safety and risk, and decision-making [9].




Figure 1: High-level definition of AI-enabled nuclear systems and modeling. Data inform physics-based and data-driven models that generate predictions with quantified uncertainty, supporting decisions and actions. Operational outcomes provide feedback for continuous model and data improvement, with the aim of supporting safer and more autonomous nuclear systems.

From this perspective, AI is not an isolated predictive technology but part of a closed digital loop connecting data and physics with uncertainty-aware engineering decisions and their resulting actions. The remainder of this study follows this progression, beginning with the computational and data foundations and advancing through surrogate and physics-informed modeling, digital twins, monitoring, lifecycle applications, optimization and control, trustworthy AI, and deployment for increasingly safe, autonomous, circular, and sustainable nuclear systems.

2 Data, AI, and Computational Foundations

Reliable nuclear AI begins with the data and computational infrastructure from which models learn. Nuclear systems are inherently multiscale and multiphysics, spanning interactions from fuel-rod and subchannel scales to assemblies and full reactor cores, while coupling neutronics, thermal fluids, materials, and other physical phenomena, as illustrated in Figure 2. Their data therefore vary widely in dimensionality, resolution, fidelity, uncertainty, and availability, ultimately determining which AI methods are appropriate and how their predictions should be interpreted.

2.1 Nuclear Datasets and Computational Infrastructure

The multiscale and multiphysics characteristics illustrated in Figure 2 are reflected directly in the diversity of nuclear design and operational datasets. Nuclear reactor or nuclear engineering data originate from experiments, operational instrumentation, inspection and maintenance activities, radiation measurements, and environmental monitoring. They may include time-series sensor signals, spatial fields, images, spectra, maintenance records, and other multimodal information. These measurements are increasingly complemented by simulation data from neutronics, thermal-hydraulics, fuel-performance, structural, and coupled multiphysics simulations.

HPC has substantially expanded this computational source of information. The Multiphysics Object Oriented Simulation Environment (MOOSE) provides a framework for coupled nonlinear multiphysics simulation [10] and has evolved toward massively parallel applications [11]. The Virtual Environment for Reactor Applications (VERA) similarly integrates high-fidelity reactor-analysis capabilities within a coupled computational environment [12] for light water reactor (LWR) analysis. From an AI perspective, such platforms are not only analysis tools but also generators of structured data over large ranges of design and operating conditions.

These datasets can be highly dimensional. Deep neural networks, for example, have been applied to large boiling water reactor (BWR) core neutronics datasets [13], while high-fidelity pressurized water reactor (PWR) flow-field simulations have been developed for downstream data-driven applications [14]. Experimental facilities provide a complementary




Figure 2: Multiscale and multiphysics characteristics of nuclear systems: (a) spatial scales extending from fuel rods and subchannels to assemblies and the reactor core, and (b) coupled neutronics, thermal-fluid, and materials phenomena. These interconnected scales and physics generate heterogeneous, high-dimensional data for modeling and AI applications.

source of dense time-dependent information; recurrent neural networks have, for example, been investigated for sensor prediction and sensor selection using thermal-hydraulic experimental data [15].

Experiments capture real behavior but can be costly or limited in scope. Simulations can explore more conditions, though their results depend on modeling assumptions and numerical accuracy. HPC makes detailed simulations possible, while AI surrogates make their results faster to use repeatedly.

2.2 AI and Machine Learning Methods

AI methods can be organized according to both learning strategy and engineering task. Supervised learning supports regression and classification when labeled examples are available, whereas unsupervised and representation-learning methods support clustering, feature extraction, dimensionality reduction, and anomaly detection. Both approaches have been used for data-driven fault detection and diagnosis in nuclear power plants [16]. Deep neural networks extend these capabilities to strongly nonlinear and high-dimensional problems: feedforward models are commonly used for static mappings, convolutional architectures exploit spatial structure, and recurrent or related sequential models address time-dependent data. Such sequential models have been applied to forecasting nuclear-plant transients [17], while convolutional models have recently been investigated for identifying resonances in experimental neutron transmission data [18].

Generative methods offer an additional capability by learning aspects of the underlying data distribution and producing synthetic samples. This is potentially valuable when observations are scarce, although generated data must remain representative and physically meaningful. The data efficiency of generative adversarial networks has been investigated for energy applications to quantify their performance as training data become limited [19].

The growing number of model architectures and hyperparameters also motivates systematic model-development workflows. Automated ML can assist with model selection, optimization, and evaluation; pyMAISE (Michigan Artificial Intelligence Standard Environment), for example, was developed to accelerate and automate ML development for nuclear power applications [20]. Ultimately, model selection should follow the structure of the problem, available data, physical knowledge, computational requirements, and consequences of prediction error rather than any single preferred algorithm.

2.3 Training Data Challenges

Nuclear AI often remains data-limited in precisely the regimes of greatest interest. Routine operation can generate extensive histories, whereas failures, severe power transients, and unusual operating conditions are intentionally rare. In addition, new reactor concepts may have little or no operational data. Data scarcity, class imbalance, and differences between training and deployment conditions are recognized challenges for nuclear AI [2].

Simulation can generate otherwise unavailable information, but high-fidelity calculations can themselves be too expensive to sample extensively. This produces a tradeoff among data quantity, fidelity, and computational cost. Active learning can reduce this burden by directing new simulations toward informative regions; neural-network ensembles with active learning have, for example, been applied to heat-pipe microreactor modeling under abnormal conditions [21].

Uncertainty further arises from measurement noise, model parameters, nuclear data, boundary conditions, physical approximations, and the ML model itself. Neural networks have been used to model nuclear cross-section uncertainties [22], illustrating that AI can participate in uncertainty analysis rather than merely consume deterministic data. For transient experimental data, Bayesian augmentation of recurrent models provides another route for representing predictive uncertainty alongside time-dependent predictions [23]. Heterogeneous data from different simulations, experiments, facilities, or operating regimes also create distribution shifts that can degrade model performance. Recent neutron-resonance classification studies, for example, illustrate how high apparent accuracy within available data does not necessarily imply reliable generalization to previously unseen physical cases [18].

The goal is to build datasets that reflect the physical system and suit the prediction task, not just to collect more data. Multifidelity learning, physics-informed models, active learning, and UQ can help address data gaps and limitations, as discussed next.

3 Surrogate Modeling, Physics-Informed AI, and Digital Twins

High-fidelity nuclear simulations can be too slow for optimization, UQ, or real-time use, while data-driven models can struggle with limited data or extrapolation conditions. Surrogate models, physics-informed AI, and digital twins may help by speeding up predictions, using physical knowledge, and connecting models to operating systems.

3.1 Surrogate and Multifidelity Modeling

Surrogate models replace repeated execution of expensive simulations with computationally efficient approximations. Polynomial models, Gaussian processes, neural networks, and reduced-order methods can therefore enable sensitivity analysis, UQ, optimization, inverse analysis, and real-time prediction that would otherwise require prohibitive computa-

tional resources. Different learning strategies have been developed for constructing such approximations from nuclear simulation data, including group method of data handling (GMDH)-based models [24] and deep Gaussian-process surrogates [25]. At the reactor scale, ML-assisted surrogates have also been developed for full-core fuel-performance analysis [26], while surrogate models have been integrated directly into design optimization under uncertainty constraints [27].

Reduced-order methods similarly represent high-dimensional behavior in a smaller state space, which is particularly useful for spatial reactor fields and dynamic applications. Reconstruction methods applied to three-dimensional reactor temperature distributions illustrate the potential to recover high-dimensional fields from reduced information [28].

Multifidelity modeling further combines information from models or datasets with different costs and accuracies. Large quantities of low-fidelity information can describe broad system behavior while limited high-fidelity data correct discrepancies between fidelity levels. Nonlinear information fusion provides one framework for learning such relationships [29], while broader multifidelity formulations span statistical and ML approaches [30].

Figure 3 illustrates this concept using a representative multifidelity regression problem [29]. A relatively dense set of inexpensive low-fidelity observations is combined with a much smaller set of high-fidelity observations to learn the high-fidelity response. The example demonstrates the central motivation of multifidelity learning: information from a lower-cost model can reduce the amount of expensive high-fidelity data required when the relationship between fidelity levels can be learned [29].




Figure 3: Illustration of multifidelity information fusion. (a) Low-fidelity data are available at substantially more locations than the high-fidelity observations. (b) The multifidelity model combines the available information to reconstruct the high-fidelity response and its predictive uncertainty. *Reproduced/adapted from Perdikaris et al. [29].*

This concept is particularly well suited to nuclear system analysis, where system models, coarse-mesh or 1D system-level calculations, high-fidelity simulations, and experiments often coexist. Multifidelity surrogates have, for example, been applied to natural-circulation prediction following loss of forced cooling in high-temperature gas reactors, simulating a very long and expensive transient [31].

The computational advantage becomes especially important for high-fidelity CFD calculations. Figure 4 shows a preliminary nuclear application in which multifidelity neural networks (MFNN) and multifidelity Kolmogorov–Arnold networks (MFKAN) use lower-fidelity coarse-mesh CFD information to predict a high-fidelity reactor flow solution. In this example, the high-fidelity CFD calculation requires approximately 20 days compared with approximately 6 days for the lower-fidelity calculation. The assembly-wise prediction errors illustrate the potential for multifidelity learning to recover expensive high-fidelity spatial information with acceptable errors from substantially less expensive simulations.




Figure 4: Preliminary demonstration of multifidelity learning for prediction of a high-fidelity nuclear CFD solution from a lower-fidelity coarse-mesh simulation. Assembly-wise mean absolute percentage errors (MAPE) are shown for (a) a multifidelity neural network (MFNN), with a core-average error of 4.18% and peak error of 9.32%, and (b) a multifidelity Kolmogorov–Arnold network (MFKAN), with a core-average error of 4.74% and peak error of 10.05%. The high- and low-fidelity CFD calculations require approximately 20 and 6 days, respectively. *Results are from preliminary unpublished work by the author.*

3.2 Physics-Informed and Hybrid AI

Computational acceleration alone does not resolve the limitations of purely data-driven models. Nuclear engineering possesses extensive knowledge in the form of conservation laws, governing equations, boundary conditions, symmetries, and engineering constraints. Physics-informed machine learning seeks to incorporate such knowledge into data-driven modeling rather than requiring it to be learned entirely from observations [32].

Physics can enter ML through physically meaningful inputs, constraints, loss functions, model architectures, or direct coupling between mechanistic and learned components. Physics-informed neural networks (PINNs), for example, incorporate governing-equation residuals and initial or boundary conditions into the training objective [33]. Their potential is accompanied by challenges in training stiff, multiscale, or high-dimensional systems [34].

Neural operators provide another scientific-ML paradigm by learning mappings between functions rather than only fixed-dimensional input–output relationships. DeepONet learns nonlinear operators [35], while the Fourier neural operator (FNO) was introduced for learning solution operators of parametric partial differential equations [36]. In broad terms, conventional surrogates approximate particular input–output relationships, PINNs embed governing equations into learning, and neural operators seek to represent families of solutions. Hybrid approaches can combine these ideas according to the application.

For nuclear systems, the principal opportunity is to exploit existing physical knowledge while learning relationships that remain uncertain or computationally expensive. A related study of a thermal energy distribution system combines physical modeling, ML, and active learning, illustrating methods that could also inform integrated nuclear digital-twin development [37].

3.3 Digital Twins

Digital twins extend accelerated and physics-informed modeling toward an evolving representation of a physical system. Unlike a static surrogate, a digital twin maintains a connection with its physical counterpart and can combine simulation, reduced-order or ML models, sensor measurements, state estimation, and decision support. Nuclear digital-twin architectures consequently span different levels of connectivity and capability [38]. A representative closed-loop architecture is illustrated in Figure 5, where simulation and sensor information support model calibration, state prediction, and full-field estimation.




Figure 5: Digital-twin architecture proposed by Burnett et al. [39], connecting a physical component with sensor data and simulation through forward, inverse, and generative loops for prediction, model calibration, and full-field estimation. The cited study implements the inverse loop; the other loops are components of the proposed framework. *Reproduced from [39].*

Computational speed is essential because the twin may need to update faster than a conventional high-fidelity model can execute. Physics-informed reduced-order modeling has been applied to reactor digital twins [40], while ML-based dimensional reduction and latent data assimilation have been used to connect reduced representations with evolving reactor information [41]. Online ML-based calibration provides another mechanism for synchronizing a nuclear power plant model with new measurements [42], and parameter identification and state estimation have been incorporated into reactor-operation digital twins [43].

Uncertainty is equally important because a continuously updated best-estimate prediction may be insufficient for safety-relevant decisions without information about its confidence. In the architecture of Figure 5, the inverse loop uses observations from the physical system to calibrate the digital representation. Variational digital twins provide a probabilistic formulation of this inverse problem, enabling model calibration while quantifying the associated uncertainty [39]. Related Bayesian recurrent formulations have also been investigated for uncertainty-aware prediction of transient thermal-hydraulic experimental data [23]. More broadly, recent nuclear work connects digital twins and AI/ML with condition monitoring and component health [44], providing a direct bridge to the monitoring and diagnostics applications discussed next.

Digital twins bring these tools together. Detailed simulations provide physical insight; surrogates make predictions faster; physics-informed AI incorporates known behavior; and sensor data and state estimation keep the model aligned with the physical system. UQ helps assess uncertainty in its predictions. Figure 5 shows only one example on how these connected loops can support model updates, decisions, and monitoring.

4 Monitoring, Anomaly Detection, and Component Health

Continuous awareness of system and component condition is central to safe and reliable nuclear operation. Nuclear facilities generate multivariate data from instrumentation and control systems, equipment sensors, process measurements, inspections, and maintenance records. AI and data analytics can complement conventional limits and periodic surveillance by identifying relationships across these measurements, recognizing departures from normal behavior, diagnosing faults, and estimating future degradation. These functions form the basis of prognostics and health management (PHM), which connects monitoring, diagnostics, and prognostics with operation and maintenance [45, 46, 47]. The progression from observing system condition to supporting maintenance decisions is summarized in Table 1.

Table 1: Progression of AI-enabled prognostics and health management from system monitoring to maintenance decision support.

PHM stage	Question	Data	Model output	Example
Monitoring and sensor assessment	Is the system and its instrumentation behaving as expected?	Normal-operation histories and correlated sensor measurements	Expected behavior, reconstructed measurements, or sensor-health indicators	[48]
Anomaly detection	Is the observed behavior different from normal?	Predominantly normal data, with abnormal examples when available	Anomaly score, reconstruction error, or abnormal-state indication	[49, 50]
Fault diagnosis	What fault is occurring and where?	Sensor signatures associated with known or inferred fault conditions	Fault class, location, severity, or probable cause	[51]
Prognostics	How is the condition likely to evolve?	Condition histories, degradation trajectories, and operating conditions	Future health state, degradation trajectory, failure probability, or RUL	[52]
Maintenance decision support	What intervention is warranted and when?	Health estimates, prognostic information, maintenance history, and operational constraints	Inspection, maintenance, repair, replacement, or scheduling support	[45, 47]

As Figure 1 shows, measurements help models estimate system health and predict future behavior. Uncertainty helps assess those predictions before they guide maintenance or operations. For safety-related decisions, teams must also consider sensor reliability, false alarms, possible failures, operating procedures, and safety requirements.

4.1 Monitoring and Anomaly Detection

The first layer of component-health assessment is determining whether the system is behaving as expected. Data-driven monitoring extends conventional threshold-based approaches by learning relationships among multiple signals and operating conditions. Models can estimate quantities that are difficult to measure directly, reconstruct expected sensor behavior, and identify multivariate patterns that may be difficult to recognize from individual measurements. Online sensor assessment is particularly relevant because apparent process abnormalities may instead originate from sensor drift, calibration error, or instrumentation problems [48].

Supervised classifiers can recognize known faults when representative labeled examples are available, but nuclear systems present an inherent data imbalance: normal operating data may be abundant while examples of important faults are rare. Unsupervised and semi-supervised methods can instead learn representations of normal behavior and identify deviations from them. Autoencoders, recurrent networks, and convolutional architectures are especially useful for multivariate temporal signals, and anomaly detection has consequently become an important component of ML-based nuclear equipment condition assessment [46, 51].

Applications within the broader nuclear technology domain illustrate this progression. Convolutional and feedforward neural networks have been applied to fault detection in accelerator power systems [53], while recurrent and ConvLSTM autoencoders have been developed for time-series anomaly detection in power-electronics signals [50]. Ensemble learning has further enabled earlier identification of developing faults [54]. Although these examples involve accelerator power electronics rather than reactor components, they demonstrate methods applicable to high-frequency, multivariate equipment signals across complex nuclear facilities.

Monitoring can also be improved through the sensor network itself. Recurrent neural networks have been investigated for sensor optimization and data prediction in thermal-hydraulic experiments [15], illustrating opportunities for virtual sensing, measurement reconstruction, and sensor placement. Practical deployment must nevertheless balance detection sensitivity against false alarms and distinguish equipment degradation from sensor faults, changing operating regimes, and out-of-distribution data.

4.2 Fault Diagnostics and Prognostics

As summarized in Table 1, anomaly detection establishes that observed behavior is unusual, whereas diagnosis seeks to determine the type, location, severity, and possible cause of the underlying fault. Diagnostic models can therefore be

formulated as classification, state-estimation, or inference problems mapping sensor signatures to candidate failure modes. Nuclear fault diagnosis has increasingly incorporated neural networks and other deep-learning methods capable of learning nonlinear relationships directly from multivariate measurements [51].

Generative models offer another option when fault classes are complex or sparsely represented. Conditional variational autoencoders, for example, have been applied to fault prediction in accelerator power systems [55]. Such approaches may help represent under-sampled abnormal conditions, although synthetic or learned fault representations require careful validation before supporting safety-relevant conclusions.

Prognostics extends diagnosis into the future by estimating degradation trajectories, component health, probability of failure, or remaining useful life (RUL). Available information changes over a component’s lifetime, motivating approaches that transition from population-level reliability information toward increasingly component-specific condition measurements [56]. Recent nuclear applications have combined graph and temporal deep-learning models for degradation and RUL prediction in PWR systems [52], while broader reviews identify anomaly detection, health indicators, multisource data fusion, and temporal prediction as interconnected elements of nuclear PHM [57].

Diagnosis and prognosis are therefore complementary: the former identifies what is occurring, while the latter estimates how the condition may evolve. For nuclear applications, these predictions should also quantify uncertainty arising from measurements, operating conditions, model form, and limited failure histories, a topic addressed further in Section 7.

4.3 Predictive Maintenance and Asset Management

The final step in the PHM progression of Table 1 is translating component-health information into maintenance decisions. Condition-based maintenance uses evidence of actual equipment condition, while predictive maintenance additionally considers expected degradation. PHM connects sensing, diagnosis, prognosis, and decision support to enable this transition [45].

Maintenance histories can also update component-health estimates and support long-term operation [58]. Condition assessment may therefore inform inspection intervals, component replacement, outage scheduling, spare-parts planning, and broader asset management. This creates a direct connection to circularity: premature replacement consumes materials and generates waste, whereas unjustified lifetime extension can challenge reliability and safety. AI-enabled condition assessment can support the middle ground of extending asset life when technically justified while maintaining established inspection and safety requirements.

5 AI for Circular and Sustainable Nuclear Systems

AI can support nuclear sustainability across the full lifecycle, from fuel and material design to operation, spent-fuel management, decommissioning, and environmental stewardship. In this context, circularity includes efficient resource use, justified lifetime extension, waste reduction and reuse where appropriate, and informed fuel-cycle and disposition decisions.

AI capabilities can support these goals in different ways. Surrogates accelerate fuel and reactor analyses, optimization explores competing objectives, computer vision assists material and waste characterization, robotics enables work in hazardous environments, and spatial learning supports radiation mapping and remediation.

5.1 Fuel, Materials, and Reactor Operations

Nuclear fuels and structural materials operate under combinations of radiation, temperature, stress, and corrosive environments that produce complex degradation mechanisms. ML has consequently been investigated for material-property prediction, irradiation effects, microstructure–property relationships, and materials discovery [59]. Computer vision provides a complementary capability, with deep semantic segmentation demonstrated for identifying defects such as dislocations, precipitates, and voids in microscopy images of irradiated steels [60]. Data-driven methods have also been used to predict thermodynamic properties of nuclear-fuel compounds and examine the effects of fission products on fuel degradation [61].

At the reactor scale, AI-enabled optimization can explore large design spaces while balancing competing engineering objectives. Earlier physics-based work on BWR fuel demonstrated the importance of simultaneously considering depletion, criticality, and uncertainty throughout fuel analysis [62]. Building on such computational problems, large-scale neuroevolution has been applied to BWR bundle design [63], while physics-informed reinforcement learning has been used for constrained nuclear assembly optimization [64]. Hybrid evolutionary and swarm algorithms have likewise been applied directly to nuclear-fuel optimization [65].

The connection to high-fidelity physics is important because optimized designs must remain grounded in reactor behavior. Tools were developed for example to couple OpenMC particle-transport calculations with optimization algorithms [66], while surrogate-driven optimization can reduce the number of expensive simulations while incorporating uncertainty constraints [27]. Machine learning has also been used to optimize passive molten-salt microreactor configurations [67].

Similar methods apply to reactor operation and fuel management. Deep-learning models combined with multiobjective optimization have been used for core reloading-pattern design [68], while evolutionary and swarm algorithms have been applied to microreactor reactivity-control-system operation [69]. More general frameworks such as NEORL (NeuroEvolution Optimization with Reinforcement Learning) integrate neuroevolution and reinforcement learning for optimization of complex energy systems [70, 71]. Generally, these methods illustrate how AI can accelerate design-space exploration and support improved fuel utilization, operating flexibility, and performance while retaining physical and safety constraints.

5.2 Spent Fuel, Waste Management, and Recycling

The back end of the nuclear lifecycle introduces characterization, inventory, storage, transportation, treatment, and disposition problems that can also benefit from data-driven analysis. For example, criticality and uncertainty analyses of BWR assembly misloading in transportation casks illustrate the importance of accurately characterizing fuel configurations during post-operation handling [72]. AI can complement such physics-based analyses by accelerating characterization and repeated fuel-cycle calculations.

ML applied to gamma-spectrometry measurements has been used to infer spent-fuel properties including burnup, cooling time, initial enrichment, and plutonium content [73]. Neural networks have also been developed to predict used-fuel composition and integrated into the Cyclus fuel-cycle simulator for rapid transmutation calculations [74]. These applications show how data-driven models can reduce computational or measurement burdens in fuel-cycle analysis without replacing the underlying physical basis.

AI can also assist radioactive-waste identification and segregation. Deep object-recognition methods have been investigated for automated waste categorization [75], while ML has been applied to management of disused sealed radioactive sources, including assessment of potential reuse [76]. For high-level waste, data-driven methods have also been investigated as reduced-order tools for complex processes relevant to geological disposal [77].

These capabilities can support broader evaluation of recycling, partitioning, transmutation, and alternative fuel-cycle strategies by accelerating predictions of material inventories, isotopic evolution, waste generation, and long-term consequences. Such tools are best viewed as accelerators of technically grounded scenario analysis; fuel-cycle decisions additionally depend on safety, safeguards, proliferation, economics, and policy considerations beyond the scope of an AI model.

5.3 Decommissioning, Robotics, and Remote Operations

Decommissioning provides a strong motivation for combining AI and robotics because characterization and dismantlement frequently occur in radiological or otherwise hazardous environments. Robotic systems can reduce worker exposure, while AI can support perception, mapping, navigation, manipulation, and planning.

Computer-vision methods have been developed to classify waste materials during nuclear decommissioning [78], supporting improved identification and segregation. Mobile platforms extend these capabilities to facility characterization: the Vega robot, for example, was developed for mapping, radiation surveying, inspection, and sample retrieval in nuclear decommissioning environments [79]. Experience from Three Mile Island, Chernobyl, and Fukushima further demonstrates both the importance and technical difficulty of robotic deployment in high-radiation, uncertain, and cluttered environments [80]. Figure 6 illustrates autonomous robotic navigation through a representative cluttered nuclear laboratory environment at the University of Michigan. The robot dog was able to achieve the inspection mission successfully without human guidance or hitting any objects in the surroundings.

A practical direction is progression from direct teleoperation toward assisted and supervised autonomy. Rather than removing the operator, assisted robotic systems can combine human supervision with perception, digital representations, planning, and control to navigate complex environments, reduce workload, and improve task execution [81]. Better robotic characterization and material identification can simultaneously reduce occupational exposure and improve waste segregation, linking remote operation directly with safety and circularity.




Figure 6: Autonomous robotic navigation through a cluttered nuclear laboratory environment at the University of Michigan, illustrating sequential perception and navigation around equipment and obstacles representative of constrained remote-operation settings. *Preliminary results that will be presented in the 2026 American Nuclear Society Winter Meeting in Phoenix, Arizona.*

5.4 Environmental Monitoring and Remediation

Environmental stewardship extends the digital lifecycle beyond the facility boundary. Radiation detectors, environmental samples, fixed sensors, and mobile surveys provide spatially and temporally distributed measurements from which the location and magnitude of contamination must be inferred. AI can complement physical and statistical models by extracting patterns from these measurements and estimating conditions between observed locations.

At radium-contaminated legacy sites, ML interpretation of gamma-ray measurements has improved identification of radioactive “hot” particles [82], while related methods have been used to estimate the spatial distribution and activity of ^{226}Ra [83]. Artificial neural networks and fuzzy-logic approaches have also been applied to ambient gamma-dose measurements for radiological mapping between monitoring locations [84].

The value of these predictions increases when they are connected to uncertainty and decision support. UQ can identify regions where additional measurements are most valuable, while optimization can support survey planning, monitoring-network design, and prioritization of remediation activities. The framework in Figure 1 therefore extends naturally to environmental stewardship: measurements generate spatial predictions, uncertainty identifies gaps in knowledge, and those predictions inform additional surveys or remediation actions.

The applications in Table 2 show how AI could support circular and sustainable nuclear systems across the lifecycle. Better fuel design can reduce resource demands; condition-informed operation can extend justified asset lifetimes; spent-fuel and waste characterization can improve disposition decisions; robotics can reduce worker exposure; and environmental analytics can guide monitoring and remediation. In each case, predictions must remain grounded in physical understanding, uncertainty, and defensible engineering decisions.

6 Optimization, Control, and Decision Support

The value of AI extends beyond predicting system behavior to supporting decisions about how nuclear systems should be designed, operated, and managed. Optimization and control connect predictive models with engineering actions, while decision support integrates predictions with objectives, constraints, uncertainty, and human judgment. This closes an important part of the digital pathway introduced in Figure 1: predictions become operationally useful when they inform defensible decisions and, where appropriate, actions.

This transition also increases the consequence of model error. A predictive model can be evaluated primarily according to the quality of its outputs, whereas an optimization or control algorithm can influence the state of the system. As AI moves from offline analysis toward decision support and autonomous control, physical constraints, uncertainty, operating limits, and mechanisms for human intervention therefore become increasingly important.

Table 2: Examples of past and ongoing AI efforts across the nuclear lifecycle.

Lifecycle stage	AI capability	Application	Example
Fuel and materials	ML, computer vision	Fuel microstructure analysis and material-property prediction	INL irradiated-fuel analysis [85]
Reactor design and fuel management	Surrogates, optimization	Core loading-pattern and fuel-management design	Westinghouse ACE; INL RAVEN [86, 87]
Spent fuel and waste	ML, spectroscopy, vision	Fuel characterization, waste sorting, and treatment modeling	Gamma-based characterization [73]; image-based sorting [75]; SRNL waste-treatment research [88]
Remote operations and decommissioning	Vision, robotics	Remote surveying and handling in hazardous areas	RadPiper [89, 90]; Vega robot [79]
Environmental stewardship	ML, spatial modeling	Contamination mapping and migration prediction	SRNL AI-ASSET/ALTEMIS [91]; radium mapping [83]

Acronyms: INL, Idaho National Laboratory; ACE, Advanced Core Engine; RAVEN, Risk Analysis Virtual ENvironment; SRNL, Savannah River National Laboratory; DOE-EM, U.S. Department of Energy Office of Environmental Management; AI-ASSET, AI Accelerated Strategies and Solutions in Environmental Technology; ALTEMIS, Advanced Long-Term Environmental Monitoring Systems.

6.1 AI-Enabled Optimization

Many nuclear engineering problems involve computationally expensive models, large design spaces, and competing objectives. Direct optimization using high-fidelity neutronics, thermal-hydraulics, structural, or multiphysics simulations may require thousands of model evaluations. AI-enabled optimization addresses this challenge through two complementary capabilities: (1) intelligent search algorithms explore large design spaces, while (2) surrogate models reduce the cost of evaluating candidate solutions.

Evolutionary and swarm algorithms are useful when objective landscapes are nonlinear, multimodal, black-box, or difficult to differentiate. Reinforcement learning can also be integrated with evolutionary search to exploit information accumulated during optimization. Rule-based reinforcement learning, for example, has been used to inform evolutionary algorithms for constrained engineering optimization [92]. Large-scale population-based search has subsequently been extended through animorphic ensemble optimization [93] and scalable multiobjective formulations [94].

These capabilities are relevant because nuclear design rarely involves a single objective. Reactor and fuel designs may simultaneously consider power distributions, safety margins, fuel utilization, economics, control requirements, and material constraints. Physics-informed RL has been demonstrated for nuclear assembly optimization [64], while neuroevolution has been applied to large-scale BWR bundle design [63]. Hybrid evolutionary and swarm algorithms provide another approach to nuclear-fuel optimization [65], and particle-transport calculations can be coupled directly with optimization through automated frameworks [66].

Surrogate-assisted optimization further reduces computational cost by replacing selected high-fidelity evaluations with learned approximations. This is especially valuable when each candidate requires an expensive Monte Carlo or multiphysics calculation. Surrogate-driven optimization with uncertainty constraints demonstrates how computational acceleration can be combined with model confidence [27]. The surrogate is therefore more useful when treated as part of an adaptive optimization process in which uncertainty or periodic high-fidelity calculations determine when its predictions are sufficiently reliable.

Multiobjective optimization helps weigh how nuclear design choices affect resource use, component life, waste, cost, and safety. AI can quickly explore the trade-offs, but *engineers must choose among the options based on requirements, uncertainty, and operating priorities.*

6.2 Model Predictive Control and Reinforcement Learning

Optimization becomes a dynamic problem when decisions must be made repeatedly as the system evolves. Model predictive control (MPC) and reinforcement learning (RL) provide complementary approaches to this setting. MPC forecasts system behavior over a finite horizon and repeatedly solves an optimization problem to select control actions while explicitly accounting for constraints. Recent work has combined reduced-order reactor physics, state estimation, and MPC for reactor operation during communication transients [95], as illustrated in Figure 7(c).

RL instead learns a policy mapping observed states to actions through interaction with an environment and a reward function. Nuclear applications of deep RL now include reactor control, load following, and transient operation [96, 97]. A trained policy can potentially produce actions rapidly without solving a full optimization problem at each time step, but its behavior depends strongly on the training environment, state and action representations, reward formulation, and conditions encountered during learning.

Microreactors provide a useful test case because future deployments may require flexible operation and reduced operator burden. A reactivity-control system has been developed for microreactor load-following applications [98], while single-agent deep RL has been demonstrated for transient and load-following control [99], as illustrated in Figure 7(a). Related work has extended this formulation to RL load-following policies [100], multi-agent control [101], shown in Figure 7(b), and multistep criticality search and power shaping [102]. These studies illustrate a progression toward increasingly complex sequential decision-making.




Figure 7: Representative learning- and model-based frameworks for autonomous nuclear reactor control: (a) single-agent reinforcement learning, in which a single policy maps reactor states to multiple control actions, reproduced from [99]; (b) multi-agent reinforcement learning, in which multiple agents coordinate control actions through a shared policy, reproduced from [101]; and (c) model predictive control with state estimation, in which a predictive reactor model and optimizer determine control actions using updated system information, reproduced from [95]. Together, the frameworks illustrate complementary approaches to sequential decision-making and autonomous reactor control.

The distinction between learning a high-performing policy and establishing a safe controller remains fundamental. Standard RL can explore poor actions during training, an assumption generally unacceptable on an operating nuclear system. Safe RL therefore seeks to constrain exploration or policy behavior [103]. For nuclear applications, this motivates training in simulation or digital-twin environments, explicit state and action constraints, physics-informed reward functions, independent protection systems, and rigorous testing before connection to plant control.

MPC and RL need not be competing approaches. As Figure 7 illustrates, RL learns control policies from state–action interactions, whereas MPC combines prediction, state estimation, and constrained optimization. Hybrid architectures could therefore use RL for higher-level strategy while MPC or another constrained controller enforces local operating limits. Such hierarchical combinations provide a potential pathway toward nuclear autonomy while separating strategic learning from safety-constrained execution.

6.3 Human-in-the-Loop and Autonomous Decision-Making

AI-enabled decision-making exists along a spectrum rather than as a binary distinction between manual and autonomous operation. AI may first perform offline analysis, then provide predictions, diagnoses, or recommended actions to operators, and eventually perform bounded functions under supervisory control. Reviews of neural-network control and autonomous reactor concepts illustrate this progression and its increasing technical requirements [104, 105].

Human oversight remains particularly important when AI outputs require contextual interpretation. An anomaly detector may identify unusual behavior without determining whether it originates from degradation, instrumentation error, or a new operating condition. Similarly, an optimizer may identify a mathematically favorable solution while omitting practical considerations not represented in its objectives. Decision support can therefore combine computational capability with operator knowledge, procedures, and responsibility for consequential decisions.

The allocation of authority should depend on the application and consequence of failure. Increasing autonomy requires stronger evidence that an AI system remains within its validated domain, recognizes conditions for which it is uncertain, respects physical and operational constraints, and transitions safely when assumptions are violated. Runtime monitoring and fallback controllers can provide additional protection by returning the system to a known safe mode.

Autonomy may also extend beyond individual reactors. Forecasts of minute-level grid conditions using generative models illustrate the broader information environment in which flexible nuclear systems may operate [106]. Future decision systems could integrate reactor state, grid demand, energy storage, and other resources, connecting component-level control with system-level energy management. Such integration makes uncertainty, interpretability, and V&V central requirements for increasingly autonomous operation.

7 Uncertainty Quantification and Trustworthy AI

High predictive accuracy alone is insufficient to establish that an AI model is suitable for safety-relevant nuclear applications. Reliable deployment also requires knowledge of predictive uncertainty, robustness, applicability, failure modes, and the evidence supporting model performance. Trustworthy nuclear AI therefore combines predictive capability with uncertainty quantification, interpretability, verification, validation, and lifecycle management. These requirements become increasingly important as AI progresses from offline analysis toward decision support and control.

7.1 Uncertainty Quantification and Robustness

Uncertainty has long been central to nuclear analysis. Sampling, sensitivity analysis, and uncertainty propagation quantify the effects of uncertain physical parameters and modeling assumptions. Some recent applications include uncertainty analysis of reactor kinetics [107] and natural-circulation behavior during air-ingress accidents [108]. AI introduces additional uncertainty through the training data, domain coverage, model architecture, and training process [109, 110], as illustrated in Figure 8.

The sources in Figure 8 also illustrate the distinction between aleatoric and epistemic uncertainty. Data noise is primarily aleatoric and irreducible, whereas uncertainty from insufficient coverage, extrapolation, and imperfect model architecture is epistemic and can potentially be reduced through additional information or improved modeling [109]. Stochastic training can also make predictions vary across fitted models, while noise in the training observations remains a separate source of aleatoric uncertainty.

Several ML methods estimate these uncertainties. Monte Carlo dropout provides an approximate Bayesian mechanism through repeated stochastic inference [111], while deep ensembles estimate predictive distributions using independently trained models [112]. Variational neural networks have also been developed specifically for UQ in nuclear applications [113], including Bayesian recurrent formulations for uncertainty-aware transient prediction [23].

An uncertainty estimate is useful only if it is meaningful relative to observed error. Calibration therefore examines whether stated confidence corresponds to empirical predictive performance. A systematically overconfident model may be more problematic for decision support than a slightly less accurate model that reliably indicates when it is uncertain. UQ should consequently be evaluated alongside accuracy rather than added only after model development.

A related problem occurs when deployment inputs differ from the training distribution. Neural networks can produce confident predictions for out-of-distribution (OOD) inputs, motivating explicit OOD detection [114]. Distribution-shift tests similarly identify changes between training and deployment data [115]. In nuclear systems, these changes may arise from new operating regimes, component aging, instrumentation changes, plant modifications, or accident conditions poorly represented in historical data.




Figure 8: Sources of uncertainty in data-driven machine-learning models. Panels (a)–(d) use Gaussian-process models to illustrate the effects of data noise, incomplete training-domain coverage, and extrapolation; panel (e) illustrates model-form uncertainty through the bias–variance tradeoff; and panels (f)–(h) illustrate uncertainty associated with stochastic training and deterministic versus probabilistic model parameters. *Adapted and reorganized from Wu et al. [109].*

Robustness therefore extends beyond uncertainty estimation. Models should be tested against measurement noise, missing sensors, corrupted inputs, parameter variations, and conditions near and beyond their training support. For emerging language-based AI, robustness can take different forms: LLM outputs can be sensitive to paraphrasing, sarcasm, and other linguistic variations [116], while fairness and social bias can influence model outputs [117]. Although these issues differ from uncertainty in reactor surrogates, they reinforce the general principle that trustworthy AI requires characterization of task-specific failure modes.

7.2 Explainability and Interpretability

Explainability and interpretability provide mechanisms for examining the relationships learned by AI models and the factors influencing their predictions. In nuclear applications, engineers may need to determine whether a model relies on physically meaningful relationships, identify measurements responsible for a diagnosis, or understand why similar conditions produce different predictions.

Post-hoc explanation methods provide one approach. LIME constructs a locally interpretable approximation around an individual prediction [118], whereas SHAP uses Shapley-value concepts to attribute predictions among input features [119]. Nuclear applications have expanded beyond feature attribution to gradient-based, attention-based, causal, and intrinsically interpretable approaches, as summarized in Table 3.

Nuclear applications also demonstrate the need to evaluate explanations rather than relying only on qualitative plots. Explainable AI (XAI) methods have been quantitatively compared for NPP accident-diagnosis models [120], while explainability has been applied to severe-accident progression [121]. Hybrid approaches have also compared internal attention with post-hoc SHAP explanations, providing a means of assessing consistency between interpretation methods [122].

A complementary strategy is to develop models that are interpretable by construction. Symbolic regression seeks explicit mathematical relationships rather than relying entirely on post-hoc interpretation of black-box predictors. Kolmogorov–Arnold networks combined with symbolic regression have been investigated for advanced energy applications involving nuclear datasets [123], illustrating how prediction can potentially be accompanied by compact analytical representations.

Interpretability nevertheless does not establish physical correctness. Post-hoc explanations can depend on the explanation method, reference distribution, and local approximation. Interpretability therefore requires its own evaluation rather

Table 3: Representative explainability and interpretability approaches in nuclear AI applications.

Approach	Methods	Explanation	Nuclear application	Example
Feature attribution	SHAP, Kernel SHAP, LIME	Input-feature contributions to predictions	Fault/accident diagnosis and severe-accident prediction	[120, 121]
Gradient or activation based	Grad-CAM	Influential regions, signals, or temporal features	Equipment fault and abnormal-state diagnosis	[126, 127]
Attention based	Attention weights, channel attention	Importance of variables, channels, or temporal features	Fault diagnosis and severe-accident prediction	[122, 121]
Causal/graph based	Causal graphs, GNN/GAT	Dependencies, propagation paths, and influential nodes	Fault diagnosis, root-cause localization, fault propagation	[128, 129]
Symbolic/intrinsic	Symbolic regression, KAN	Explicit analytical input–output relationships	Surrogate modeling and thermal-hydraulic prediction	[123, 130]
Hybrid pre-/post-hoc	Attention + SHAP	Comparison of internal and post-hoc importance	Fault diagnosis and explanation validation	[122]

than serving as a substitute for validation [124]. Its strongest role in nuclear applications is as one component of a broader assurance process combining physical consistency, uncertainty, independent testing, and engineering review.

Generative AI expands this challenge because scientifically plausible outputs may not preserve domain knowledge. Comparative studies of text-to-image models in nuclear contexts illustrate the need to evaluate general-purpose generative systems against domain-specific scientific concepts [125]. This concern becomes increasingly important as generative models enter scientific and engineering workflows.

7.3 Verification, Validation, and Safety Assurance

Verification and validation (V&V) should show that an AI system works reliably for its intended use. Because AI learns behavior from data, assurance must cover the data, preprocessing, model design and weights, software, and deployment environment. The evidence needed depends on the role: an exploratory surrogate requires less than a model advising operators, and an autonomous controller requires more [131, 132]. Relevant evidence includes independent testing, uncertainty calibration, sensitivity and stress tests, physical consistency, reproducibility, and performance across the intended operating range.

Datasets also need documented origins and coverage of the conditions for which performance is claimed [133, 134]. Simulation data depend on the credibility of the underlying physics model; plant and experimental data require checks of sensor quality, calibration, missing values, preprocessing, and instrumentation changes.

Deployed models must remain traceable. Recording software versions, data, preprocessing, hyperparameters, and weights allows results to be reproduced and changes to be assessed. Models retrained with new data need controlled requalification so that their safety case remains valid [133, 134].

General frameworks such as the NIST AI Risk Management Framework provide principles for identifying, measuring, and managing AI risks [134]. Nuclear applications must connect these principles with established practices of quality assurance, defense-in-depth, software control, independent review, and regulatory oversight. The objective is therefore not simply to “validate the AI,” but to establish evidence that the complete AI-enabled system performs acceptably within its intended role and that failures are detected, bounded, or mitigated.

8 Deployment Challenges and Future Research Priorities

Despite rapid progress in AI research, a substantial gap remains between demonstrating model performance in controlled studies and deploying AI within operational nuclear systems. Closing this gap requires advances not only in algorithms but also in data, infrastructure, software integration, cybersecurity, assurance, regulation, governance, and workforce

development. These challenges define the transition from individual demonstrations toward integrated digital capabilities across the nuclear lifecycle.

8.1 From Research Demonstration to Nuclear Deployment

A research demonstration typically begins with a well-defined dataset, a selected prediction or control problem, and quantitative performance metrics. Operational deployment introduces a larger system boundary. Data must be acquired reliably from plant instrumentation, transferred through appropriate interfaces, processed consistently, and supplied to models under computational and timing constraints. Predictions must then be communicated to operators or downstream systems in a form that can support action. These differences create two separate points of view of AI in action.

Studies of barriers to AI and ML adoption in nuclear power identify technical, organizational, regulatory, and data-related obstacles between successful demonstrations and operational use [135]. The IAEA similarly identifies data availability and quality, trustworthiness, cybersecurity, organizational readiness, and integration among considerations for AI deployment in the nuclear power industry [133]. Deployment should therefore influence model development from the beginning rather than be considered only after high predictive performance is achieved.

Integration with existing systems is a major practical challenge. Nuclear facilities can contain instrumentation, control, and software developed over multiple generations, whereas modern AI frameworks may depend on computational environments not designed for safety-relevant industrial deployment. Interfaces must preserve data integrity, timing, cybersecurity, and configuration control. For some applications, separating AI from direct safety functions provides a practical initial pathway, with models first operating as advisory or monitoring systems.

Performance monitoring must continue after deployment. Sensor replacement, component aging, changes in operating strategy, software updates, and evolving plant conditions can alter the input distribution. Runtime monitoring should therefore evaluate both system health and AI health, including data quality, distribution shift, prediction uncertainty, and model performance where ground truth becomes available. Digital twins that we described before can support this process by maintaining continuously updated representations of physical systems [38, 44].

Deployment is consequently better understood as a lifecycle than a single transition from laboratory to plant. Models must be qualified, integrated, monitored, maintained, and eventually updated or retired, with evidence and configuration histories remaining traceable throughout that lifecycle.

8.2 Regulation, Governance, and Standards

AI in safety-relevant nuclear applications raises regulatory questions beyond predictive accuracy. Relevant considerations include data provenance, model traceability, uncertainty, cybersecurity, human oversight, configuration management, responsibility for AI-assisted decisions, and treatment of models that evolve after deployment.

The U.S. Nuclear Regulatory Commission has incorporated AI into its strategic planning. NUREG-2261 identifies regulatory readiness, organizational capability for reviewing AI, partnerships, workforce development, and potential agency applications as important objectives [136]. More recent NRC strategy continues this effort through staff development and evolution of the regulatory framework [137]. Regulatory capability must therefore develop in parallel with the technology rather than after AI systems reach operational maturity.

International cooperation is similarly important because many assurance questions are shared across jurisdictions. Joint principles developed by the Canadian Nuclear Safety Commission, UK Office for Nuclear Regulation, and U.S. NRC address considerations for developing AI systems in nuclear applications [138]. Regulatory gap assessments have also examined how existing nuclear requirements apply to AI and where additional clarification may be needed [139]. These efforts suggest that AI regulation need not begin from an entirely separate framework: established principles of nuclear safety, quality assurance, software control, and defense-in-depth remain relevant, although their implementation must account for data-driven behavior.

Governance also extends beyond formal regulation. Organizations must define ownership of models, authority to modify them, evidence required before updates, performance-monitoring responsibilities, and authority to act on AI outputs. For human-facing generative systems, additional issues can arise. LLM-based analysis of nuclear-power discourse illustrates expanding nuclear-domain information applications [140, 141], while studies of model fairness demonstrate the need to examine systematic biases [117].

Governance requirements should therefore scale with the role and consequence of the AI system. A model assisting scientific information processing does not require the same assurance as a diagnostic tool presented to an operator, and neither requires the same evidence as an autonomous control function. Risk-informed governance can preserve

these distinctions while maintaining common expectations for traceability, human responsibility, and controlled model evolution.

8.3 Research and Technology Priorities

Future research should increasingly focus on integrated capabilities rather than isolated improvements in individual algorithms. Surrogate models, physics-informed learning, digital twins, UQ, anomaly detection, optimization, and autonomous control are individually promising. *The larger challenge is connecting them into systems that remain reliable as data, physical conditions, and operational requirements change.*

A first priority is the development of qualified, representative, and shareable datasets. Nuclear AI continues to face limited abnormal and failure data, proprietary information, heterogeneous formats, and differences among facilities and experiments. Generative models can partially address scarcity [19], but synthetic data cannot substitute for validation against representative physical observations. Future datasets should therefore be designed for coverage, provenance, uncertainty, and relevance to deployment conditions rather than volume alone.

A second priority is tighter integration of physics and AI. *Purely data-driven models can perform well under familiar conditions but become unreliable when extrapolating.* Physics-informed models, hybrid architectures, multifidelity methods, and digital twins provide mechanisms for incorporating established knowledge. Physics-informed digital twins with active learning illustrate how these capabilities can form adaptive predictive systems.

A third priority is moving from predictive UQ toward uncertainty-aware decisions. Producing an uncertainty interval is of limited value if a downstream controller, optimizer, or operator ignores it. Future systems should use uncertainty to determine when additional data are required, when a high-fidelity calculation should be invoked, when an automated action should be restricted, or when authority should return to a human operator. *Connecting prediction, uncertainty, and action is therefore central to trustworthy autonomy.*

A fourth priority is the development and evaluation of multimodal and foundation models for scientific applications. Nuclear information spans numerical simulations, time-series measurements, images, maintenance records, reports, diagrams, and natural language. General-purpose generative models provide mechanisms for processing these modalities, but scientific reliability must be established. Foundation models are beginning to find uses across nuclear research and operations. Early work has evaluated how general text-to-image models represent nuclear concepts [125] and developed models tailored to generating nuclear-related images [142]. Language models have also been applied to analyzing public sentiment about nuclear energy [140, 141] and explored as interfaces to digital twins for remote reactor monitoring and thermal-fluid testbeds [143, 144]. These initial efforts show the breadth of possible applications, but they also raise a common question: *how can domain knowledge, data provenance, uncertainty, and verification be built into these systems so that their outputs are useful and dependable in nuclear settings?*

A fifth priority is scalable computational and experimental infrastructure. High-fidelity simulation remains essential for generating training data, testing rare scenarios, and evaluating AI-enabled control systems. High-fidelity Monte Carlo or multiphysics modeling for data-driven applications illustrates the computational scale required to generate physically detailed datasets. *Future infrastructure should increasingly connect HPC, experiments, digital twins, and hardware-in-the-loop environments so that AI systems can be tested under realistic dynamic conditions before deployment.*

These priorities reinforce one another: better data improve models; physics supports generalization; UQ identifies limitations; digital twins provide evolving system representations; and high-fidelity test environments enable realistic evaluation of decision and control systems. The research objective should therefore progress from demonstrating isolated AI capabilities toward establishing integrated and defensible digital systems.

8.4 Workforce and Implementation Pathway

The transition toward AI-enabled nuclear systems requires a workforce capable of working across traditional disciplinary boundaries. Nuclear engineers increasingly require knowledge of data science, ML, HPC, digital instrumentation, optimization, control, and cybersecurity. Conversely, AI specialists working in nuclear applications must understand physical modeling, uncertainty, safety analysis, operational constraints, and the regulatory environment. Regulatory organizations also require sufficient expertise to evaluate emerging applications, a need reflected in NRC strategic planning [136, 137].

Education should therefore emphasize integration rather than treating AI as a standalone computational skill. Students and practitioners should understand how data are generated, how physical models are validated, how uncertainty

propagates into predictions, and how AI outputs affect engineering decisions. Practical training can combine simulation, experimental data, and human–machine interaction so that AI is evaluated within the larger engineering system.

Generative and foundation models further increase the importance of AI literacy. These tools can assist coding, data analysis, information retrieval, visualization, and scientific communication, but can also *produce convincing outputs that are incomplete or incorrect*. Nuclear-specific evaluations of generative models [125, 142] provide examples for teaching both their capabilities and limitations. Workforce preparation should therefore emphasize verification of AI-generated information and continued human responsibility for engineering conclusions.

Implementation is likely to proceed incrementally. AI-assisted analysis can first accelerate simulation, data processing, and engineering workflows without directly influencing operation. Trusted predictive tools can follow as sufficient validation and uncertainty information become available. Integration with digital twins can then provide continuously updated representations of system condition, supporting operator decision-making before progression toward supervised autonomy. Higher levels of autonomy would require correspondingly stronger evidence for robustness, runtime monitoring, fail-safe behavior, cybersecurity, and regulatory acceptance.

This progression does not imply that maximum autonomy is the desired endpoint. The appropriate level depends on the application, the benefits obtained from automation, and the evidence supporting it. For some functions, AI-assisted analysis or decision support may remain preferable; for others, supervised autonomy may improve reliability, flexibility, or worker safety.

9 Conclusions

This study shows that nuclear AI has developed well beyond a single class of prediction problems. Surrogates and reduced-order models accelerate reactor and fuel analyses; physics-informed learning and digital twins connect computation with physical knowledge and measurements; and data-driven methods support anomaly detection, diagnosis, prognosis, optimization, and control. Applications in waste characterization, decommissioning, and environmental monitoring extend these capabilities across the nuclear lifecycle. Collectively, these studies demonstrate substantial technical breadth, although most reported successes remain specific to their datasets, systems, and evaluation settings.

The main unresolved question is how these methods perform when conditions differ from those used to develop them. Scarce fault data, simulation-to-plant differences, changing operating regimes, and uncertain measurements can weaken apparent performance. For decision-making applications, predictive accuracy alone is an incomplete measure: uncertainty must be calibrated, physical and operational limits respected, and failures recognized before model outputs influence consequential actions. Evidence for deployment must also cover data provenance, software and model configuration, cybersecurity, human interaction, and continued performance after updates.

A useful research direction is therefore to evaluate complete AI-enabled workflows rather than isolated models. Shared and well-documented datasets, independent testing across facilities and operating conditions, comparisons with established methods, and realistic simulation and experimental testbeds would make results easier to interpret and reproduce. Future work should connect high-fidelity models, sensing, digital twins, uncertainty quantification, and constrained decision-making while defining when an AI system should defer to a higher-fidelity analysis or a human operator. Progress on these questions, alongside regulatory engagement and interdisciplinary training, will determine which research demonstrations can become dependable tools for safer and more sustainable nuclear systems.

AI Usage Statement

GPT-6 Sol, accessed through the ChatGPT platform, was used to correct grammar, polish incomplete sentences, and remove redundant wording. The manuscript’s scientific content, arguments, and conclusions are the author’s work and responsibility. The AI tool was also used to arrange Figure 8 into panels; the underlying images came from the original published works cited in that figure. AI was not used for any other purpose in preparing this manuscript. The author reviewed the AI-assisted changes and takes full responsibility for the final text, figures, and any remaining errors.

Data Availability

This article presents no new dataset. Data associated with the author’s published work, where publicly released, are available through the AIMS Lab GitHub page: <https://github.com/aims-umich>. Data from preliminary, unpublished work cited in the text may be available from the author upon reasonable request. Availability of data from other cited studies is determined by their respective authors and publications.

References

- [1] Mario Gomez-Fernandez, Kathryn Higley, Akira Tokuhiro, Kent Welter, Weng-Keen Wong, and Haori Yang. Status of research and development of learning-based approaches in nuclear science and engineering: A review. *Nuclear Engineering and Design*, 359:110479, 2020.
- [2] Qingyu Huang, Shinian Peng, Jian Deng, Hui Zeng, Zhuo Zhang, Yu Liu, and Peng Yuan. A review of the application of artificial intelligence to nuclear reactors: Where we are and what’s next. *Heliyon*, 9(3):e13883, 2023.
- [3] Derjew Ayele Ejigu, Yanjie Tuo, and Xiaojing Liu. Application of artificial intelligence technologies and big data computing for nuclear power plants control: A review. *Frontiers in Nuclear Engineering*, 3:1355630, 2024.
- [4] Abbas J. Jinia, Shaun D. Clarke, Jean M. Moran, and Sara A. Pozzi. Intelligent radiation: A review of machine learning applications in nuclear and radiological sciences. *Annals of Nuclear Energy*, 201:110444, 2024.
- [5] Amber Boehnlein, Markus Diefenthaler, Nobuo Sato, Malachi Schram, Veronique Ziegler, Cristiano Fanelli, Morten Hjorth-Jensen, Tanja Horn, Michelle P. Kuchera, Dean Lee, Witold Nazarewicz, Peter Ostroumov, Kostas Orginos, Alan Poon, Xin-Nian Wang, Alexander Scheinker, Michael S. Smith, and Long-Gang Pang. Colloquium: Machine learning in nuclear physics. *Reviews of Modern Physics*, 94(3):031003, 2022.
- [6] International Atomic Energy Agency. Artificial intelligence for accelerating nuclear applications, science and technology. Technical Report IAEA/ART/INT, International Atomic Energy Agency, Vienna, 2022.
- [7] Majdi I Radaideh and Tomasz Kozlowski. Combining simulations and data with deep learning and uncertainty quantification for advanced energy modeling. *International Journal of Energy Research*, 43(14):7866–7890, 2019.
- [8] Majdi I Radaideh, Hoang Tran, Lianshan Lin, Hao Jiang, Drew Winder, Sarma Gorti, Guannan Zhang, Justin Mach, and Sarah Cousineau. Model calibration of the liquid mercury spallation target using evolutionary neural networks and sparse polynomial expansions. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 525:41–54, 2022.
- [9] Zhegang Ma, Han Bao, Sai Zhang, Ming Xian, and Adam Mack. Exploring advanced computational tools and techniques with artificial intelligence and machine learning in operating nuclear plants. Technical Report NUREG/CR-7294, INL/EXT-21-61117, U.S. Nuclear Regulatory Commission, Washington, DC, 2022.
- [10] Derek Gaston, Chris Newman, Glen Hansen, and Damien Lebrun-Grandié. MOOSE: A parallel computational framework for coupled systems of nonlinear equations. *Nuclear Engineering and Design*, 239(10):1768–1778, 2009.
- [11] Cody J. Permann, Derek R. Gaston, David Andrš, Robert W. Carlsen, Fande Kong, Alexander D. Lindsay, Jason M. Miller, John W. Peterson, Andrew E. Slaughter, Roy H. Stogner, and Richard C. Martineau. MOOSE: Enabling massively parallel multiphysics simulation. *SoftwareX*, 11:100430, 2020.
- [12] John A. Turner, Kevin Clarno, Matt Sieger, Roscoe Bartlett, Benjamin Collins, Roger Pawlowski, Rodney Schmidt, and Randall Summers. The virtual environment for reactor applications (VERA): Design and architecture. *Journal of Computational Physics*, 326:544–568, 2016.
- [13] Rabie Abu Saleem, Majdi I Radaideh, and Tomasz Kozlowski. Application of deep neural networks for high-dimensional large bwr core neutronics. *Nuclear Engineering and Technology*, 52(12):2709–2716, 2020.
- [14] Logan A Burnett, Hyungjun Kim, Hsien-Cheng Chou, Arsha Witoelar, Robert A Brewster, Benoit Forget, Emilio Baglietto, and Majdi I Radaideh. High-fidelity modeling of pressurized water reactor flow fields for data-driven applications. *Energy*, 362:142104, 2026.
- [15] Farhana Farha, Meredith Eaheart, Izabela Gutowska, Brendan Kochunas, and Majdi I Radaideh. Sensor optimization and data prediction in the high temperature test facility with recurrent neural networks. In *Proceedings of the 21st International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-21)*, Busan, Korea, 2025.
- [16] Guang Hu, Taotao Zhou, and Qianfeng Liu. Data-driven machine learning for fault detection and diagnosis in nuclear power plants: A review. *Frontiers in Energy Research*, 9:663296, 2021.
- [17] Majdi I Radaideh, Connor Pigg, Tomasz Kozlowski, Yujia Deng, and Annie Qu. Neural-based time series forecasting of loss of coolant accidents in nuclear power plants. *Expert Systems with Applications*, 160:113699, 2020.
- [18] Nataly R Panczyk, Athanasios Stamatopoulos, Josef Svoboda, and Majdi I Radaideh. Learning to resolve neutron resonances with fully convolutional neural networks. *arXiv preprint arXiv:2608.04027*, 2026.

- [19] Umme Mahbuba Nabila, Linyu Lin, Xingang Zhao, William L Gurecky, Pradeep Ramuhalli, and Majdi I Radaideh. Data efficiency assessment of generative adversarial networks in energy applications. *Energy and AI*, 20:100501, 2025.
- [20] Patrick A Myers, Nataly Panczyk, Shashank Chidige, Connor Craig, Jacob Cooper, Veda Joynt, and Majdi I Radaideh. pymaise: A python platform for automatic machine learning and accelerated development for nuclear power applications. *Progress in Nuclear Energy*, 180:105568, 2025.
- [21] Dean Price, Majdi I Radaideh, and Brendan Kochunas. Active learning with deep neural network ensembles for heat pipe microreactor modeling in abnormal operation. In *Proceedings of the 18th International Probabilistic Safety Assessment and Analysis (PSA 2023)*, pages 355–364, Knoxville, TN, 2023.
- [22] Majdi I Radaideh, Dean Price, and Tomasz Kozlowski. Modeling nuclear data uncertainties using deep neural networks. In *EPJ Web of Conferences*, volume 247, page 15016. EDP Sciences, 2021.
- [23] Logan A Burnett, Umme Mahbuba Nabila, Izabela Gutowska, and Majdi I Radaideh. Variational bayesian last layer augmented recurrent neural networks for transient thermal hydraulic experiments. In *Proceedings of the 21st International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-21)*, Busan, Korea, 2025.
- [24] Majdi I Radaideh and Tomasz Kozlowski. Analyzing nuclear reactor simulation data and uncertainty with the group method of data handling. *Nuclear Engineering and Technology*, 52(2):287–295, 2020.
- [25] Majdi I Radaideh and Tomasz Kozlowski. Surrogate modeling of advanced computer simulations using deep gaussian processes. *Reliability Engineering & System Safety*, 195:106731, 2020.
- [26] Yifeng Che, Joseph Yurko, Paul Seurin, and Koroush Shirvan. Machine learning-assisted surrogate construction for full-core fuel performance analysis. *Annals of Nuclear Energy*, 168:108905, 2022.
- [27] Omer F Erdem, David P Broughton, Josef Svoboda, Chengkun Huang, and Majdi I Radaideh. Surrogate-driven design optimization with uncertainty constraints in monte carlo simulations. *Energy and AI*, 22:100655, 2025.
- [28] Dean Price, Majdi I Radaideh, and Brendan Kochunas. Simplified matching pursuits applied to 3d nuclear reactor temperature distribution construction. *Applied Mathematical Modelling*, 131:134–158, 2024.
- [29] Paris Perdikaris, Maziar Raissi, Andreas Damianou, Neil D. Lawrence, and George Em Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751, 2017.
- [30] M. Giselle Fernández-Godino. Review of multi-fidelity models. *Advances in Computational Science and Engineering*, 1(4):351–400, 2023.
- [31] Meredith Eaheart and Majdi I Radaideh. Multifidelity surrogate modeling of onset of natural circulation following a depressurized loss of forced cooling in high-temperature gas reactors. *Applied Thermal Engineering*, 303:132489, 2026.
- [32] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.
- [33] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [34] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, 2022.
- [35] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [36] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- [37] Umme Mahbuba Nabila, Paul Seurin, Linyu Lin, and Majdi I Radaideh. Physics-informed digital twin development of thermal energy distribution systems with active learning. *Applied Energy*, 416:127903, 2026.
- [38] Mengyan Hu, Xueyan Zhang, Cuiting Peng, Yixuan Zhang, and Jun Yang. Current status of digital twin architecture and application in nuclear energy field. *Annals of Nuclear Energy*, 202:110491, 2024.
- [39] Logan A Burnett, Umme Mahbuba Nabila, and Majdi I Radaideh. Variational digital twins. *Energy and AI*, 24:100756, 2026.

- [40] Helin Gong, Sibao Cheng, Zhang Chen, and Qing Li. Data-enabled physics-informed machine learning for reduced-order modeling digital twin: Application to nuclear reactor physics. *Nuclear Science and Engineering*, 196(6):668–693, 2022.
- [41] Helin Gong, Sibao Cheng, Zhang Chen, Qing Li, César Quilodr  n-Casas, Dunhui Xiao, and Rossella Arcucci. An efficient digital twin based on machine learning SVD autoencoder and generalised latent assimilation for nuclear reactor physics. *Annals of Nuclear Energy*, 179:109431, 2022.
- [42] Houde Song, Meiqi Song, and Xiaojing Liu. Online autonomous calibration of digital twins using machine learning with application to nuclear power plants. *Applied Energy*, 326:119995, 2022.
- [43] Helin Gong, Tao Zhu, Zhang Chen, Yaping Wan, and Qing Li. Parameter identification and state estimation for nuclear reactor operation digital twin. *Annals of Nuclear Energy*, 180:109497, 2023.
- [44] Hector Mendoza, Vaibhav Yadav, Syed Bahauddin Alam, Doug Eskins, and Raj Iyengar. Advances in digital twins and AI/ML for condition monitoring in nuclear applications. *Frontiers in Energy Research*, 14:1716514, 2026.
- [45] Jamie B. Coble, Pradeep Ramuhalli, Leonard J. Bond, J. Wesley Hines, and Belle R. Upadhyaya. Prognostics and health management in nuclear power plants: A review of technologies and applications. Technical Report PNNL-21515, Pacific Northwest National Laboratory, Richland, WA, 2012.
- [46] Yong Xu, Yunze Cai, and Lin Song. Condition assessment of nuclear power plant equipment based on machine learning methods: A review. *Nuclear Technology*, 209(7):929–962, 2023.
- [47] Xingang Zhao, Junyong Kim, Kyle Warns, Xinyan Wang, Pradeep Ramuhalli, Sacit Cetiner, Hyun Gook Kang, and Michael Golay. Prognostics and health management in nuclear power plants: An updated method-centric review with special focus on data-driven methods. *Frontiers in Energy Research*, 9:696785, 2021.
- [48] Jamie B. Coble, Pradeep Ramuhalli, Ryan M. Meyer, and Hash Hashemian. Online sensor calibration assessment in nuclear power systems. *IEEE Instrumentation & Measurement Magazine*, 16(3):32–37, 2013.
- [49] Ahmad Y Al Rashdan, Hany S Abdel-Khalik, Kellen M Giraud, Daniel G Cole, Jacob A Farber, William W Clark, Abenezzer Alemu, Marcus C Allen, Ryan M Spangler, and Athi Varuttamaseni. A qualitative strategy for fusion of physics into empirical models for process anomaly detection. *Energies*, 15(15):5640, 2022.
- [50] Majdi I Radaideh, Chris Pappas, Jared Walden, Dan Lu, Lasitha Vidyaratne, Thomas Britton, Kishansingh Rajput, Malachi Schram, and Sarah Cousineau. Time series anomaly detection in power electronics signals with recurrent and convlstm autoencoders. *Digital Signal Processing*, 130:103704, 2022.
- [51] Aicheng Gong, Zhongjian Qiao, Xihui Li, Jiafei Lyu, and Xiu Li. A review on methods and applications of artificial intelligence on fault detection and diagnosis in nuclear power plants. *Progress in Nuclear Energy*, 177:105474, 2024.
- [52] Shadman Ahmad Khattak, Yong-Kuo Liu, Yu-Kun Liu, Jia-Rong Gao, Zhou-Xin Shi, and Ji Liu. A deep learning-based prognostic approach for predicting PWR degradation and remaining useful life using GNN-PTC-LSTM. *Annals of Nuclear Energy*, 230:112172, 2026.
- [53] Majdi Radaideh, Chris Pappas, Pradeep Ramuhalli, and Sarah Cousineau. Application of convolutional and feedforward neural networks for fault detection in particle accelerator power systems. In *Proceedings of the Annual Conference of the PHM Society*, volume 14, pages 1–8. PHM Society, 2022.
- [54] Majdi I Radaideh, Chris Pappas, Mark Wezensky, Pradeep Ramuhalli, and Sarah Cousineau. Early fault detection in particle accelerator power electronics using ensemble learning. *International Journal of Prognostics and Health Management*, 14(1), 2023.
- [55] Yasir Alanazi, Malachi Schram, Kishansingh Rajput, Steven Goldenberg, Lasitha Vidyaratne, Chris Pappas, Majdi I Radaideh, Dan Lu, Pradeep Ramuhalli, and Sarah Cousineau. Multi-module-based cvae to predict hvcm faults in the sns accelerator. *Machine Learning with Applications*, 13:100484, 2023.
- [56] Michael Sharp, Jamie Coble, Alan Nam, J. Wesley Hines, and Belle Upadhyaya. Lifecycle prognostics: Transitioning between information types. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, 229(4):279–290, 2015.
- [57] Shuo Zhang, Wei Cheng, Wentao Wang, Hassaan Ahmad, Le Zhang, Haoyu Liu, Zelin Nie, and Xuefeng Chen. Deep learning-based health monitoring and prognostics for nuclear power plants: A systematic review. *Applied Energy*, 421:128252, 2026.
- [58] Zachary Welz, Jamie Coble, Belle R. Upadhyaya, and Wes Hines. Maintenance-based prognostics of nuclear plant equipment for long-term operation. *Nuclear Engineering and Technology*, 49(5):914–919, 2017.

- [59] Dane Morgan, Ghanshyam Pilania, Adrien Couet, Blas P. Uberuaga, Cheng Sun, and Ju Li. Machine learning in nuclear materials research. *Current Opinion in Solid State and Materials Science*, 26(2):100975, 2022.
- [60] Graham Roberts, Simon Y. Haile, Rajat Sainju, Danny J. Edwards, Brian Hutchinson, and Yuanyuan Zhu. Deep learning for semantic segmentation of defects in advanced STEM images of steels. *Scientific Reports*, 9(1):12744, 2019.
- [61] Denise A. Lopes, Rinkle Juneja, Alicia M. Raftery, J. Matthew Kurley, William F. Cureton, and Andrew T. Nelson. A machine learning approach to quantify degradation of nuclear fuels and the effects of fission products. *Journal of Nuclear Materials*, 616:156053, 2025.
- [62] Majdi I Radaideh, Dean Price, Daniel O’Grady, and Tomasz Kozlowski. Advanced bwr criticality safety part i: Model development, model benchmarking, and depletion with uncertainty analysis. *Progress in Nuclear Energy*, 113:230–246, 2019.
- [63] Majdi I Radaideh, Benoit Forget, and Koroush Shirvan. Large-scale design optimisation of boiling water reactor bundles with neuroevolution. *Annals of Nuclear Energy*, 160:108355, 2021.
- [64] Majdi I Radaideh, Isaac Wolverton, Joshua Joseph, James J Tusar, Uuganbayar Otgonbaatar, Nicholas Roy, Benoit Forget, and Koroush Shirvan. Physics-informed reinforcement learning optimization of nuclear assembly design. *Nuclear Engineering and Design*, 372:110966, 2021.
- [65] Majdi I Radaideh and Koroush Shirvan. Pesa: Prioritized experience replay for parallel hybrid evolutionary and swarm algorithms-application to nuclear fuel. *Nuclear Engineering and Technology*, 54(10):3864–3877, 2022.
- [66] Xubo Gu, Majdi I Radaideh, and Jingang Liang. Openneomc: A framework for design optimization in particle transport simulations based on openmc and neorl. *Annals of Nuclear Energy*, 180:109450, 2023.
- [67] Andrew Larsen, Ross Lee, Braden Clayton, Edward Mercado, Ethan Wright, Brent Edgerton, Brian Gonda, and Matthew Memmott. Optimization of passive modular molten salt microreactor geometric perturbations using machine learning. *Nuclear Engineering and Design*, 424:113307, 2024.
- [68] Muhammad Kamran Butt, Liangzhi Cao, Chenghui Wan, Kaihui Lei, and Izat Khan. Integrating the deep learning and multi-objective genetic algorithm to the reloading pattern optimization of HPR1000 reactor core. *Nuclear Engineering and Design*, 428:113531, 2024.
- [69] Dean Price, Majdi I Radaideh, and Brendan Kochunas. Multiobjective optimization of nuclear microreactor reactivity control system operation with swarm and evolutionary algorithms. *Nuclear Engineering and Design*, 393:111776, 2022.
- [70] Majdi I Radaideh, Katelin Du, Paul Seurin, Devin Seyler, Xubo Gu, Haijia Wang, and Koroush Shirvan. Neorl: Neuroevolution optimization with reinforcement learning. *arXiv preprint arXiv:2112.07057*, 2021.
- [71] Majdi I Radaideh, Katelin Du, Paul Seurin, Devin Seyler, Xubo Gu, Haijia Wang, and Koroush Shirvan. Neorl: Neuroevolution optimization with reinforcement learning—applications to carbon-free energy systems. *Nuclear Engineering and Design*, 412:112423, 2023.
- [72] Majdi I Radaideh, Dean Price, and Tomasz Kozlowski. Criticality and uncertainty assessment of assembly misloading in bwr transportation cask. *Annals of Nuclear Energy*, 113:1–14, 2018.
- [73] Péter Kirchknopf, Bálint Batki, Péter Völgyesi, Zoltán Kató, and Imre Szalóki. Application of machine learning methods for spent fuel characterization based on gamma spectrometry measurements. *Annals of Nuclear Energy*, 205:110601, 2024.
- [74] Jin Whan Bae, Andrei Rykhlevskii, Gwendolyn Chee, and Kathryn D. Huff. Deep learning approach to nuclear fuel transmutation in a fuel cycle simulator. *Annals of Nuclear Energy*, 139:107230, 2020.
- [75] Jeong-Guk Kim, Sung-Chan Jang, Il-Sik Kang, Dong-Ju Lee, Jin-Woo Lee, and Hee-Seoung Park. A study on object recognition using deep learning for optimizing categorization of radioactive waste. *Progress in Nuclear Energy*, 130:103528, 2020.
- [76] Ihsan Aulia Rahman, Zico Pratama Putra, Pendi Rusadi, Kanita Salsabila Dwi Irmanti, Ajrieh Setyawan, Moch Romli, Ayi Muzyawati, Suhartono Suhartono, Hendra Adhi Pratama, Raden Sumarbagiono, Gustri Nurliati, Niken Siwi Pamungkas, and Muhammad Yusuf. Integration of machine learning models for enhancing radioactive waste management of disused sealed radioactive sources. *Nuclear Engineering and Design*, 442:114272, 2025.
- [77] Guang Hu and Wilfried Pfingsten. Data-driven machine learning for disposal of high-level nuclear waste: A review. *Annals of Nuclear Energy*, 180:109452, 2023.
- [78] Affan Shaukat, Yang Gao, Jeffrey A. Kuo, Bob A. Bowen, and Paul E. Mort. Visual classification of waste material for nuclear decommissioning. *Robotics and Autonomous Systems*, 75:365–378, 2016.

- [79] Benjamin Bird, Matthew Nancekievill, Andrew West, Jim Hayman, Chris Ballard, Will Jones, Shaun Ross, Toby Wild, Tom Scott, and Barry Lennox. Vega—a small, low cost, ground robot for nuclear decommissioning. *Journal of Field Robotics*, 39(3):232–245, 2022.
- [80] Yasuyoshi Yokokohji. The use of robots to respond to nuclear accidents: Applying the lessons of the past to the Fukushima Daiichi Nuclear Power Station. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:681–710, 2021.
- [81] Erwin Jose Lopez Pulgarin, Dave Hopper, Jon Montgomerie, James Kell, Joaquin Carrasco, Guido Herrmann, Alexander Lanzon, and Barry Lennox. From traditional robotic deployments towards assisted robotic deployments in nuclear decommissioning. *Frontiers in Robotics and AI*, 12:1432845, 2025.
- [82] Adam Varley, Andrew Tyler, Leslie Smith, Paul Dale, and Mike Davies. Remediating radium contaminated legacy sites: Advances made through machine learning in routine monitoring of “hot” particles. *Science of the Total Environment*, 521–522:270–279, 2015.
- [83] Adam Varley, Andrew Tyler, Leslie Smith, Paul Dale, and Mike Davies. Mapping the spatial distribution and activity of ^{226}Ra at legacy sites through machine learning interpretation of gamma-ray spectrometry data. *Science of the Total Environment*, 545–546:654–661, 2016.
- [84] Cafer Mert Yeşilkanat, Yaşar Kobya, Halim Taşkın, and Uğur Çevik. Spatial interpolation and radiological mapping of ambient gamma dose rate by using artificial neural networks and fuzzy logic methods. *Journal of Environmental Radioactivity*, 175–176:78–93, 2017.
- [85] Idaho National Laboratory. AI for a clean energy future: Researchers use machine learning for advanced fuel development, May 2022.
- [86] Westinghouse Electric Company. Advanced core engine (ACE) for loading pattern optimization. Accessed 2026-09-24.
- [87] Junyung Kim, Mohammad Gamal Abdo, Yong-Joon Choi, Juan Cristhian Luque Gutierrez, Jason Hou, Christopher Gosdin, and Jarrett Valeri. Pressurized-water reactor core design demonstration with genetic algorithm based multi-objective plant fuel reload optimization platform. Technical Report INL/RPT-23-74498, Idaho National Laboratory, 2023.
- [88] U.S. Department of Energy, Office of Environmental Management. SRNL awarded AI-powered cleanup projects through DOE’s genesis mission, August 2026.
- [89] Siri Maley, Heather Jones, Warren Whittaker, David Kohanbash, Josh Spisak, Ralph Boirum, Andrew Zhang, J. David Yesso, Russell Riddle, and William Whittaker. Commissioning of robotic inspection and automated analysis system for assay of gas diffusion piping. In *Proceedings of the Waste Management Symposia (WM2019)*, number 19500, Phoenix, AZ, March 2019.
- [90] Heather Jones, Siri Maley, Mohammadreza Mousaei, David Kohanbash, Warren Whittaker, James Teza, Andrew Zhang, Nikhil Jog, and William Whittaker. A robot for nondestructive assay of holdup deposits in gaseous diffusion piping. *arXiv preprint arXiv:1901.10341*, 2019.
- [91] U.S. Department of Energy, Office of Environmental Management. SRNL harnesses AI to tackle environmental challenges, cut cleanup costs, March 2026.
- [92] Majdi I Radaideh and Koroush Shirvan. Rule-based reinforcement learning methodology to inform evolutionary algorithms for constrained optimization of engineering applications. *Knowledge-Based Systems*, 217:106836, 2021.
- [93] Dean Price and Majdi I Radaideh. Animorphic ensemble optimization: a large-scale island model. *Neural Computing and Applications*, 35(4):3221–3243, 2023.
- [94] Omer F Erdem, Dean Price, Paul Seurin, and Majdi I Radaideh. Maeo: Scalable multiobjective animorphic ensemble optimization for large-scale optimization applications. *Nuclear Engineering and Design*, 458:115156, 2026.
- [95] Zachery Dahm, Konstantinos Vasili, Vasileios Theos, Linyu Lin, Joe Oncken, Robert England, Vivek Agarwal, and Stylianos Chatzidakis. Model predictive control with state estimation and reduced-order physics for nuclear reactor operation during unanticipated communication transients. *Annals of Nuclear Energy*, 240:112733, 2027.
- [96] Yongchao Liu, Bo Wang, Sichao Tan, Tong Li, Wei Lv, Zhenfeng Niu, Jiangkuan Li, Puzhen Gao, and Ruifeng Tian. Applications of deep reinforcement learning in nuclear energy: A review. *Nuclear Engineering and Design*, 429:113655, 2024.
- [97] Aicheng Gong, Yangkun Chen, Junjie Zhang, and Xiu Li. Possibilities of reinforcement learning for nuclear power plants: Evidence on current applications and beyond. *Nuclear Engineering and Technology*, 56(6):1959–1974, 2024.

- [98] Kamal K Abdullaheem, Sooyoung Choi, Qicang Shen, Brendan Kochunas, and Majdi I Radaideh. A load following reactivity control system for nuclear microreactors. *Progress in Nuclear Energy*, 184:105676, 2025.
- [99] Leo Tunkle, Kamal Abdullaheem, Linyu Lin, and Majdi I Radaideh. Nuclear microreactor transient and load-following control with deep reinforcement learning. *Energy Conversion and Management: X*, 27:101090, 2025.
- [100] Leo Tunkle, Kamal Abdullaheem, Maxime Amans, and Majdi I Radaideh. Load-following control in microreactors with deep reinforcement learning. In *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2025)*, pages 845–854, Denver, CO, 2025.
- [101] Leo Tunkle, Kamal Abdullaheem, Linyu Lin, and Majdi I Radaideh. Enhancing autonomous control of microreactors using multi-agent reinforcement learning. In *Nuclear Plant Instrumentation and Control & Human-Machine Interface Technology (NPIC&HMIT 2025)*, pages 40–49, 2025.
- [102] Majdi I Radaideh, Leo Tunkle, Dean Price, Kamal Abdullaheem, Linyu Lin, and Moutaz Elias. Multistep criticality search and power shaping in nuclear microreactors with deep reinforcement learning. *Nuclear Science and Engineering*, 200(sup1):S309–S321, 2026.
- [103] Javier García and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(42):1437–1480, 2015.
- [104] Gang Zhou and Da Tan. Review of nuclear power plant control research: Neural network-based methods. *Annals of Nuclear Energy*, 181:109513, 2023.
- [105] Hui-Yu Hsieh and Pavel Tsvetkov. Advancements and challenges of machine learning and deep learning in autonomous control of nuclear reactors. *Annals of Nuclear Energy*, 223:111643, 2025.
- [106] Umme Mahbuba Nabila, Gavin McDuffee, Xingang Zhao, William L Gurecky, Pradeep Ramuhalli, and Majdi I Radaideh. Predicting minute-level power grid conditions using generative adversarial networks. In *2025 Nuclear Plant Instrumentation and Control and Human-Machine Interface Technology, NPIC and HMIT 2025*, pages 893–902. American Nuclear Society, 2025.
- [107] Majdi I Radaideh, William A Wieselquist, and Tomasz Kozłowski. A new framework for sampling-based uncertainty quantification of the six-group reactor kinetic parameters. *Annals of Nuclear Energy*, 127:1–11, 2019.
- [108] Meredith Eaheart, Jacob Cooper, Molly Ross, Nate See, and Majdi I Radaideh. Sensitivity analysis and uncertainty propagation of the time to onset of natural circulation in air ingress accidents. *Nuclear Engineering and Design*, 445:114510, 2025.
- [109] Xu Wu, Lesego E. Moloko, Pavel M. Bokov, Gregory K. Delipei, Joshua Kaizer, and Kostadin N. Ivanov. Uncertainty quantification for data-driven machine learning models in nuclear engineering applications: Where we are and what do we need? *Nuclear Science and Engineering*, 200(8):1798–1820, 2026.
- [110] Wenchong He, Zhe Jiang, Tingsong Xiao, Zelin Xu, and Yukun Li. A survey on uncertainty quantification methods for deep learning. *ACM Computing Surveys*, 58(7):1–35, 2026.
- [111] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059. PMLR, 2016.
- [112] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, volume 30, pages 6402–6413, 2017.
- [113] Logan A Burnett, Umme Mahbuba Nabila, and Majdi I Radaideh. Development of variational neural networks for uncertainty quantification of nuclear applications. In *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2025)*, pages 784–793, Denver, CO, 2025.
- [114] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.
- [115] Stephan Rabanser, Stephan Günnemann, and Zachary C. Lipton. Failing loudly: An empirical study of methods for detecting dataset shift. In *Advances in Neural Information Processing Systems*, volume 32, pages 1394–1406, 2019.
- [116] Naman Bhargava, Mohammed I Radaideh, O Hwang Kwon, Aditi Verma, and Majdi I Radaideh. On the impact of language nuances on sentiment analysis with large language models: Paraphrasing, sarcasm, and emojis. *IEEE Access*, 14:81885–81900, 2026.

- [117] Mohammed I Radaideh, O Hwang Kwon, and Majdi I Radaideh. Fairness and social bias quantification in large language models for sentiment analysis. *Knowledge-Based Systems*, 319:113569, 2025.
- [118] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. Association for Computing Machinery, 2016.
- [119] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.
- [120] Seung Geun Kim, Seunghyoung Ryu, Kyungho Jin, and Hyeonmin Kim. Quantitative comparison of explainable artificial intelligence methods for nuclear power plant accident diagnosis models. *Progress in Nuclear Energy*, 180:105605, 2025.
- [121] Semin Joo, Yeonha Lee, Seok Ho Song, Kyusang Song, Mi Ro Seo, Sung Joong Kim, and Jeong Ik Lee. Leveraging explainable AI for reliable prediction of nuclear power plant severe accident progression. *Reliability Engineering & System Safety*, 264:111307, 2025.
- [122] Yukun Liu, Yongkuo Liu, Ji Liu, Zhouxin Shi, Zhen Wang, and Guohua Wu. Research on interpretable OCA-MLP fault diagnosis of nuclear power plants based on Kernel SHAP and attention mechanism. *Annals of Nuclear Energy*, 227:111910, 2026.
- [123] Nataly R Panczyk, Omer F Erdem, and Majdi I Radaideh. Opening the ai black-box: Symbolic regression with kolmogorov-arnold networks for advanced energy applications. *Energy and AI*, 22:100595, 2025.
- [124] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- [125] Veda Joynt, Jacob Cooper, Naman Bhargava, Katie Vu, O Hwang Kwon, Todd R Allen, Aditi Verma, and Majdi I Radaideh. A comparative analysis of text-to-image generative ai models in scientific contexts: a case study on nuclear power. *Scientific Reports*, 14(1):30377, 2024.
- [126] Gensheng Qian, Yang Huang, Hang Zhu, Changran Geng, Huimin Yuan, Shaoxian Gu, and Xiaobin Tang. Enhancing transparency of intelligent fault diagnosis models in nuclear power plants via explainability analysis. *Annals of Nuclear Energy*, page 112617, 2026.
- [127] Young Do Koo, Sa Kil Kim, Seung Geun Kim, and Man Gyun Na. A Grad-CAM-based deep learning diagnosis model with operator-centered explanations under abnormal states in NPPs. *Nuclear Engineering and Technology*, 58(6):104217, 2026.
- [128] Jie Liu, Kai Pan, Enrico Zio, and Yuantao Yao. Temporal causal graph-based attention gated recurrent unit for interpretable fault diagnosis in nuclear power plants. *Reliability Engineering & System Safety*, 273:112332, 2026.
- [129] Xin Wang, Hang Wang, and MinJun Peng. Interpretability study of a typical fault diagnosis model for nuclear power plant primary circuit based on a graph neural network. *Reliability Engineering & System Safety*, 261:111151, 2025.
- [130] Wufeng Zheng, Yankuan Li, Weiyi Yao, Kunjie Chen, Zikun Chen, Kai Wang, Ryo Yokoyama, Shuichiro Miwa, and Koji Okamoto. Physics-interpretable symbolic regression for critical heat flux prediction using a wide-range experimental database. *International Journal of Heat and Mass Transfer*, 272:129449, 2027.
- [131] Canadian Nuclear Safety Commission, Office for Nuclear Regulation, and U.S. Nuclear Regulatory Commission. Considerations for developing artificial intelligence systems in nuclear applications. Technical report, Canadian Nuclear Safety Commission, Office for Nuclear Regulation, and U.S. Nuclear Regulatory Commission, September 2024.
- [132] Vivek Agarwal, Cody Walker, Koushik Araseethota Manjunatha, Torrey Mortenson, Nancy Lybeck, Anna Hall, Rachael Hill, and Andrei Gribok. Technical basis for advanced artificial intelligence and machine learning adoption in nuclear power plants. Technical report, Idaho National Laboratory, September 2022.
- [133] International Atomic Energy Agency. Considerations for deploying artificial intelligence applications in the nuclear power industry. Technical Report NR-T-1.26, International Atomic Energy Agency, Vienna, 2025.
- [134] Elham Tabassi. Artificial intelligence risk management framework (AI RMF 1.0). Technical Report NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD, 2023.
- [135] Anna Hall and Vivek Agarwal. Barriers to adopting artificial intelligence and machine learning technologies in nuclear power. *Progress in Nuclear Energy*, 175:105295, 2024.

- [136] Matthew Dennis, Thomas Lalain, Luis Betancourt, Amy Hathaway, and Robert Anzalone. Artificial intelligence strategic plan: Fiscal years 2023–2027. Technical Report NUREG-2261, U.S. Nuclear Regulatory Commission, Washington, DC, 2023.
- [137] U.S. Nuclear Regulatory Commission. Artificial intelligence strategic plan. Technical Report ML25269A196, U.S. Nuclear Regulatory Commission, Washington, DC, 2025.
- [138] Canadian Nuclear Safety Commission, UK Office for Nuclear Regulation, and U.S. Nuclear Regulatory Commission. Considerations for developing artificial intelligence systems in nuclear applications. Technical Report ML24241A252, Canadian Nuclear Safety Commission, UK Office for Nuclear Regulation, and U.S. Nuclear Regulatory Commission, 2024.
- [139] Osvaldo Pensado, Patrick LaPlante, Michael Hartnett, and Kenneth Holladay. Regulatory framework gap assessment for the use of artificial intelligence in nuclear applications. Technical report, Southwest Research Institute, San Antonio, TX, 2024. Prepared for the Office of Nuclear Regulatory Research, U.S. Nuclear Regulatory Commission; ADAMS Accession No. ML24290A059.
- [140] O Hwang Kwon, Katie Vu, Naman Bhargava, Mohammed I Radaideh, Jacob Cooper, Veda Joynt, and Majdi I Radaideh. Sentiment analysis of the united states public support of nuclear power on social media using large language models. *Renewable and Sustainable Energy Reviews*, 200:114570, 2024.
- [141] O Hwang Kwon, Katie Vu, Jacob Cooper, Veda Joynt, and Majdi I Radaideh. Using large language models to classify public sentiment toward nuclear power. In *Proceedings of the Pacific Basin Nuclear Conference 2024 (PBNC)*, pages 640–649, Idaho Falls, ID, 2024.
- [142] Mohammed I Radaideh, Jeremy Moon, Andre Gala-Garza, Emma Son, Yug Shah, and Majdi I Radaideh. Nucleardiffusion: Text-to-image foundation models for learning nuclear energy concepts. *arXiv preprint arXiv:2608.04030*, 2026.
- [143] Xavier Ndim Ndim, Doyeong Lim, John Ford, Simon Adu, Jian Tao, Yassin Hassan, and Yang Liu. Large language model-assisted digital twin for remote monitoring and control of advanced reactors. *Progress in Nuclear Energy*, 192:106172, 2026.
- [144] Doyeong Lim, Xavier Ndim Ndim, Christian Young, Yassin Hassan, and Yang Liu. An ai-driven thermal-fluid testbed for advanced small modular reactors: Integration of digital twin and large language models. *AI Thermal Fluids*, 4:100023, 2025.