# Study of Thermodynamics with Artificial Intelligence: A Structured Literature Review

## Summary see:

## https://circularastronomy.com/2025/10/05/the-ai-thermodynamics-revolution-accelerating-scientific-discovery-and-engineering-a-sustainable-future/

## I. Introduction: The Synergistic Intersection of AI and Thermodynamics

### 1.1. The Role of Thermodynamics in Modern Science and Engineering

Thermodynamics, the science governing energy, entropy, and the criteria for phase equilibrium, provides the fundamental physical constraints underpinning material behavior and system operations across engineering and natural sciences. Its principles are indispensable for optimizing industrial operations, advancing novel materials discovery, and achieving energy efficiency and sustainability goals.[1] However, traditional methodologies for thermodynamic analysis face persistent limitations when applied to systems of increasing complexity, size, or time duration. High-fidelity techniques, such such as Density Functional Theory (DFT) or

*ab initio* Molecular Dynamics (AIMD), while accurate, are computationally prohibitive. These computational barriers prevent the comprehensive exploration of high-dimensional parameter spaces or the simulation of phenomena over the long timescales necessary to capture meaningful thermodynamic behavior.[3]

The need to overcome these limitations—specifically, the high cost of achieving high accuracy—has created a compelling impetus for integrating advanced computational tools.

This critical juncture has driven the adoption of Artificial Intelligence (AI) and Machine Learning (ML), which offer the prospect of retaining high accuracy while drastically accelerating the speed of simulation and prediction. The consequence of this integration is the rapid development of ML models that function as high-speed surrogate potentials, enabling the exploration of regimes previously deemed intractable due to time and resource constraints.

## 1.2. Defining the AI/ML Landscape in Scientific Computing

The modern integration of AI into scientific computing is rooted in the computational revolution of the 21st century, characterized by the substantial proliferation of data and enhanced hardware capabilities.[5] Conceptually, the groundwork for AI exploration was laid much earlier, notably with Alan Turing's early work and the development of the Stochastic Neural Analog Reinforcement Calculator (SNARC), one of the first artificial neural networks.[5]

The application of AI in thermodynamics represents a significant methodological paradigm shift—from traditional, experience-driven modeling techniques to a data-driven approach that leverages high-throughput data to analyze complex system behaviors.[6] AI primarily serves three crucial, interlinked functions in this domain: enhancing predictive modeling, accelerating complex simulations, and optimizing intricate physical processes. This transformative movement has profoundly improved the efficiency and accuracy of fundamental materials innovation.[6]

## 1.3. Dual Modalities of AI-Thermodynamics Integration

The literature concerning AI and thermodynamics can be broadly divided into two complementary modalities, reflecting both the utilitarian application of AI as a tool and the theoretical investigation of AI itself through a thermodynamic lens:

1. **AI *for* Thermodynamics:** This dominant modality focuses on applying ML algorithms to solve classical thermophysical problems, such as predicting material properties, modeling complex phase equilibria, or optimizing energy management systems.
2. **Thermodynamics *of* AI:** This emerging theoretical field applies statistical mechanics and thermodynamic principles, such as entropy and stochastic processes, to quantify, understand, or improve the structure and training dynamics of AI algorithms.[7]

The evolution of the field, starting with applied problems (AI *for* prediction) and progressing

toward a fundamental investigation (Thermodynamics *of* AI), indicates a maturation where physical laws are not merely inputs but are increasingly utilized to inform the architecture and resilience of the ML algorithms themselves. This convergence suggests a movement toward a unified theoretical framework.

# II. Evolution of Data-Driven Thermophysical Modeling: From Classification to Prediction

The field's chronological development illustrates a clear progression, moving from utilizing neural networks for simple classification tasks to employing them for generating high-fidelity physical interaction potentials.

## 2.1. Early Statistical and Neural Network Applications (Pre-2007)

Initial applications of machine learning in physics demonstrated its power in tackling complex identification problems. Seminal studies focused on classifying phases of matter and accurately identifying phase transitions, particularly in systems where the Hamiltonian was ill-defined or the order parameters were unknown.[9] For instance, researchers successfully employed deep learning and Convolutional Neural Networks (CNNs), inspired by their success in image recognition, to map out complete two-dimensional topological phase diagrams of quantum systems.[10] By training networks on momentum-space density images of ultracold quantum gases, these methods achieved accurate characterization of transitions, such as the superfluid-to-Mott-insulator transition, demonstrating results that were not feasible using conventional analytical methods.[10]

## 2.2. The Breakthrough: Machine Learning Force Fields (MLFFs)

The central challenge in computational thermodynamics is bridging the gap between the speed of classical MD, which relies on less accurate empirical potentials, and the fidelity of computationally expensive *ab initio* methods.[3] The solution materialized in the form of Machine Learning Force Fields (MLFFs), which were developed to emulate the potential

energy surface derived from high-level electronic structure calculations (e.g., DFT).

### 2.2.1. Seminal Study: High-Dimensional Neural Network Potentials (HDNNPs)

The breakthrough in MLFFs is widely attributed to the work by Behler and Parrinello in 2007, who introduced the concept of High-Dimensional Neural Network Potentials (HDNNPs).[11] The core finding of their work was the formulation of the total energy (

) of a chemical system as a sum of individual atomic energy contributions (), where each  is determined solely by the local chemical environment of atom .[13] By utilizing atom-centered symmetry functions as input, the neural network could treat all atoms of the same type identically, ensuring physical invariance.[15]

This invention addressed the fundamental trade-off, delivering the energetic and force accuracy of electronic structure calculations but computed orders of magnitude faster. This acceleration is critical: MLFF simulations, even for large systems (e.g., 10,000 atoms), can achieve high sampling rates, enabling the simulation of sufficient periods necessary to obtain converged bulk thermodynamic properties.[16]

### 2.2.2. Subsequent Advancements

Following the HDNNP architecture, subsequent advancements have focused on expanding the complexity and transferability of these models. This includes the development of increasingly sophisticated architectures to handle multicomponent systems and incorporate polarizable effects.[18] Furthermore, deep learning frameworks, such as CGnets, have been developed to tackle the challenge of coarse-graining large macromolecular systems. CGnets reformulate coarse-graining as a supervised ML problem, enabling the learning of coarse-grained free energy functions via force-matching schemes. Crucially, these deep learning approaches successfully capture multibody interaction terms that emerge from dimensionality reduction, which classical coarse-graining methods often fail to represent, thereby allowing accurate free energy surface predictions without explicit solvent models.

Table 1 summarizes key studies that established critical methodologies in the fusion of AI and thermodynamics.

Table 1: Seminal Studies in AI-Enhanced Thermodynamics

| Study (Author, Year) | Core Thermodynamic Problem | AI Methodology | Core Finding (IEEE Citation Placeholder) |
|---|---|---|---|
| Behler & Parrinello (2007) [11] | Potential Energy Surface Modeling for MD | High-Dimensional Neural Network Potential (HDNNP) | Established the first 'second-generation' ML potential by summing atomic energy contributions based on local environment, enabling *ab initio* accuracy at MD speeds. |
| Sohl-Dickstein et al. (2015) [19] | Generative Modeling and Probability Distribution | Diffusion Models (Non-Equilibrium Physics Inspiration) | Developed highly flexible, tractable generative models by simulating iterative forward diffusion and a learned reverse diffusion process. |
| Rosenberger et al. (2022) [21] | Equation of State (EOS) Consistency | Free Energy Neural Network (FE-NN) & Automatic Differentiation (AD) | Ensured exact preservation of Maxwell relations and fundamental thermodynamic consistency by training the model on the derivatives of a single learned free energy function . |
| Shams & Tiwary (2024) [8] | AI Interpretability | Thermodynamics-Inspired Explainable | Introduced interpretation |

| | | | |
|---|---|---|---|
| | (XAI) | Representations (TIERA) | entropy, drawing inspiration from classical thermodynamics, to quantify and generate optimally human-interpretable explanations for complex black-box AI models. |

# III. Core Research Themes: AI *for* Predictive Modeling and System Optimization

The practical applications of AI in thermodynamics fall largely under two umbrella themes: accelerating the accurate prediction of properties and enhancing the optimization and control of macroscopic systems.

## 3.1. Theme 1: Predictive Modeling and Simulation Acceleration

The core utility of ML in predictive modeling lies in its capacity to process high-throughput data—often generated by materials genome projects and quantum chemistry simulations—to deeply mine complex structure–performance correlation laws.[6] These hybrid models, which employ diverse techniques such as linear, tree, and ensemble ML algorithms, are trained on input features derived from computational data sets (e.g., DFT data) to accelerate the study of materials and predict previously inaccessible thermal properties.[23]

One key area of acceleration is force field parameterization. For critical applications like binding free energy calculations in drug discovery, the accurate generation of parameters such as partial charges is essential but traditionally time-consuming. ML models trained on DFT-based atomic charges can rapidly and accurately predict these charges, reducing the required time from hours to less than a minute. This capability allows researchers to synthesize the predicted parameters with other neural network outputs (e.g., atom types and phase angles) to produce complete, high-accuracy topologies for small drug-like molecules.[26]

The successful deployment of AI in predicting thermodynamic behavior is fundamentally an economic and temporal efficiency multiplier. By simulating various thermodynamic conditions and predicting material performance, ML algorithms allow researchers to systematically filter and focus development resources on only the most promising candidates. This iterative process of virtual screening and validation significantly shortens the timeframe required to bring new materials to market, a critical advantage in fast-paced industries such as electronics and renewable energy.

### 3.2. Theme 2: Optimization and Control of Industrial Thermodynamic Systems

AI has transformed the study of thermodynamics from a passive analytical pursuit into an active mechanism for system control and improvement. AI-driven optimization techniques, notably genetic algorithms and reinforcement learning, are instrumental in managing complex industrial processes to achieve high energy efficiency and reliability. This methodology extends to crucial areas like demand-side energy management, where AI is used to optimize energy consumption scheduling, leading to enhanced energy efficiency, cost reduction, and operational sustainability.

Furthermore, AI integration significantly enhances simulation methodologies through surrogate modeling. Surrogate models provide rapid evaluations of complex thermodynamic processes, which is particularly valuable in high-stakes environments, such as nuclear power plants or chemical processing facilities, where traditional experimentation is impractical due to cost or safety concerns. This capability allows engineers to conduct extensive “what-if” analyses quickly, leading to improved system dynamics understanding and enhanced operational safety. The ultimate implication of this trend is the fusion of AI and thermodynamics acting as a pivotal force in fostering a more sustainable and efficient global future through optimized energy systems and materials discovery.

## IV. Addressing Consistency: Physics-Informed and Theory-Constrained AI

A primary philosophical and practical challenge in applying data-driven models to physical science is ensuring that predictive accuracy does not come at the expense of fundamental physical consistency. Black-box models often learn data patterns without inherent regard for

physical laws.

## 4.1. The Challenge of Thermodynamic Inconsistency in Black-Box Models

Standard machine learning techniques, such as Multi-Task Neural Networks (MT-NNs) or Kernel Ridge Regression, are designed to minimize the error between predicted and target properties. When multiple thermodynamic properties (e.g., pressure, chemical potential, and internal energy) are treated as independent targets, these models inherently risk violating the rigorous differential relationships dictated by the calculus of thermodynamics.[22] This leads to models that are physically inconsistent and violate core constraints like the Maxwell relations, rendering them unreliable, particularly when extrapolating beyond the training domain.[22]

## 4.2. Physics-Informed Neural Networks (PINNs) in Thermal Modeling

Physics-Informed Neural Networks (PINNs) emerged as a significant paradigm shift, offering a hybrid approach that enforces consistency by embedding known physical laws directly into the model's structure. Unlike conventional deep learning, which relies solely on labeled data, PINNs incorporate governing Partial Differential Equations (PDEs)—such as those describing energy conservation, fluid dynamics, or heat transfer—into the neural network's loss function.[27]

This mechanism compels the neural network to find solutions that are not only accurate but also physically consistent, even in regimes of limited data.[27] This makes PINNs highly valuable for complex thermal modeling, such as in battery systems and electronics.[29] Furthermore, research has focused on adapting PINNs to solve challenging problems involving discontinuities, such as heat transfer across imperfect contact interfaces governed by Kapitza thermal resistance, by employing augmented-variable formulations that maintain network differentiability while capturing complex jumps in the solution.[30]

## 4.3. Preserving Maxwell Relations via Automatic Differentiation

A more fundamental approach to enforcing consistency in equilibrium thermodynamics was

demonstrated by Rosenberger et al. (2022) with the Free Energy Neural Network (FE-NN).[21] The underlying principle of thermodynamics dictates that all state variables are derivatives of a single fundamental thermodynamic function, such as the Helmholtz free energy (

).

The FE-NN models this fundamental function  directly using an Artificial Neural Network.[22] The key enabling technology is

**Automatic Differentiation (AD)**, which is used in a novel way to define the predicted properties, such as pressure () and chemical potential (), as the *analytical derivatives* of the learned free energy function () with respect to the state variables (volume, particle number, or temperature).[22]

The network's parameters () are optimized by minimizing a loss function that compares the actual properties to the derivatives of . By forcing the relationship between properties to originate from a single underlying function, the FE-NN ensures the exact preservation of the Maxwell relations—the inherent consistency requirements of thermodynamics—which traditional MT-NNs fail to maintain. This sophisticated use of AD ensures high accuracy and physical coherence simultaneously, marking a necessary evolution in the development of scientifically valid AI models.[22]

# V. Conflicting Viewpoints and Ongoing Theoretical Debates

The integration of AI into complex physical systems has introduced critical trade-offs and sparked theoretical debates regarding accountability, model veracity, and fundamental scientific interpretation.

## 5.1. Debate I: The Accuracy-Interpretability Trade-Off

A persistent conflict in machine learning literature is the trade-off between predictive accuracy and model interpretability. Deep learning models often achieve superior accuracy by learning intricate, high-dimensional patterns, but their decision-making processes remain opaque—the "black-box" dilemma.[32] Conversely, simpler, more interpretable models often lack

the capacity to capture complex, non-linear physical relationships.[32]

In the context of thermodynamics, particularly in critical industrial or scientific applications, interpretability is essential for fostering user trust, meeting compliance standards, and ensuring accountability.[33] If a complex AI model predicts a catastrophic material failure or a suboptimal reactor condition, engineers require transparent reasoning to diagnose the cause and implement corrective measures.[34] This tension means that in regulated fields, researchers may choose simpler, rule-based or inherently interpretable models—even if they yield marginally lower predictive accuracy—to guarantee traceability and clear criteria.[32] Consequently, significant effort is now directed toward developing inherently interpretable models, such as Explainable Boosting Machines (EBMs), or post-hoc Explainable AI (XAI) frameworks to mitigate this trade-off.[34]

## 5.2. Debate II: Theoretical Unification and the Thermodynamics *of* AI

Beyond using AI as a tool, a growing body of work explores the theoretical connection between machine learning and thermodynamics itself. This research draws heavily on statistical mechanics, viewing algorithms and data transformations through the lens of physical processes.

### 5.2.1. Non-Equilibrium Inspiration for Generative Models

The concept of non-equilibrium statistical physics has fundamentally inspired advanced machine learning architectures, particularly generative models. The development of diffusion models, for example, is based on the essential idea of modeling complex data distribution by simulating an iterative *forward diffusion process* that systematically destroys structure. The subsequent step involves learning the *reverse diffusion process* that restores the structure, yielding a flexible and tractable generative model of the data.[19] This physical analogy provides a rigorous, principled framework for developing powerful algorithms capable of rapid sampling and evaluation of complex probability distributions.

### 5.2.2. Quantifying Interpretability with Entropy

An evolving area of XAI utilizes classical thermodynamic principles to solve the interpretability problem. Shams and Tiwary (2024) introduced the concepts of **Interpretation Entropy** and the **Thermodynamics-inspired Explainable Representations of AI (TIERA)**.[8] This approach draws inspiration from classical thermodynamics to provide a universal metric for evaluating and generating optimally human-interpretable explanations for predictions made by black-box AI models across diverse domains, including molecular simulations and classification tasks.[8]

### 5.2.3. AI Optimization as a Non-Equilibrium Process

The training dynamics of neural networks, often governed by algorithms like stochastic gradient descent (SGD), exhibit close parallels to natural non-equilibrium processes, such as protein folding.[36] Analysis using a Fokker-Planck approach, adapted from statistical physics, demonstrates that conventional SGD often settles into a non-equilibrium stationary state characterized by persistent currents in the parameter space. Critically, this non-equilibrium state exhibits a measurable entropy production rate for any given training trajectory, and the distribution of these rates adheres to the integral and detailed fluctuation theorems—non-equilibrium generalizations of the Second Law of Thermodynamics.[36] This theoretical analysis is crucial for understanding the optimization process and engineering equilibrium stationary states for specific applications, such as Bayesian machine learning.[36]

## 5.3. Debate III: Generalizability vs. Interpolation

A persistent practical limitation across many Deep Learning (DL) applications is the critical difference between interpolation (predicting within the range of training data) and reliable extrapolation (predicting outside the known domain).[37] While ML models demonstrate immense success in accurately interpolating properties, their reliability significantly degrades when they are required to predict the behavior of novel materials or thermodynamic states far beyond what they have been trained on. Since the inherent goal of scientific AI is often the autonomous discovery of unprecedented materials or phases, this limitation severely curtails the generalizability and real-world applicability of currently trained models.[37] Overcoming this bottleneck requires innovative modeling methods that inherently possess better extrapolation capabilities, rather than relying solely on perpetually expanding the training dataset.[37]

# VI. Identified Gaps and Deficiencies in the Current Literature

The deficiencies in the current body of work highlight opportunities for fundamental and applied advancements, forming the natural foundation for future research efforts.

## 6.1. Gap 1: Robust Integration of Non-Equilibrium Statistical Mechanics (NESM)

Despite the theoretical inspiration drawn from non-equilibrium thermodynamics for specific generative models (e.g., Diffusion Models) [19], a substantial gap exists in developing a generalized, systematic mathematical and computational framework for integrating generative modeling with the full scope of Non-Equilibrium Statistical Mechanics (NESM).[38] The field currently lacks comprehensive models capable of simulating, predicting, and understanding complex dynamical systems that operate far from thermodynamic equilibrium.

Furthermore, a significant theoretical gap concerns the management and understanding of stochasticity. Real-world systems and computational models inevitably carry inherent uncertainties. At present, minimal research has addressed the specific thermodynamic consequences arising from this unavoidable uncertainty in dynamic systems, which is a key requirement for bridging the gap between theoretical stochastic thermodynamics and practical AI applications.[39] The future requires dedicated effort to build coherent systematic methodologies, analogous to the FE-NN for equilibrium systems, to handle time-dependent, dissipative processes with physical rigor.

## 6.2. Gap 2: Overcoming Extrapolation Limitations via Hybrid Active Learning

The inability of existing DL models to reliably extrapolate remains the primary technological bottleneck limiting the progress toward truly autonomous discovery in thermophysical science.[37] Relying solely on vast static datasets for high-dimensional design space exploration is inefficient.[40]

The critical solution lies in the robust deployment of Active Learning (AL), a technique that intelligently guides data collection or high-fidelity simulation towards the most informative, unexplored regions of the parameter space.[40] However, the systematic implementation of theory-informed machine learning within AL workflows is underdeveloped. This includes the use of physics-infused kernels, discrepancy models that account for simulation-experiment differences, or Bayesian conavigation frameworks.[40] The necessary step forward involves perfecting the autonomous implementation of closed-loop systems—often termed "self-driving labs"—that seamlessly integrate targeted experimentation or simulation with theory-informed Bayesian optimization and ML model refinement. This fusion is essential to move beyond passive prediction and into an era of autonomous material discovery.

### 6.3. Gap 3: Data Scarcity, Standardization, and Model Transferability

While high-throughput screening initiatives have generated significant data, data scarcity remains a critical issue for highly specialized chemical systems or complex, difficult-to-measure kinetic phenomena. This scarcity is compounded by the poor transferability of models; a DL model trained on one chemical system frequently fails when applied to a chemically dissimilar system.[37]

To address these issues and maximize the utility of existing high-throughput databases (e.g., the Computational Materials Repository, CMR) [41], there is a pressing need for the development of systematic and standardized benchmarking frameworks. These frameworks are required to consistently evaluate the efficiency, precision, and generalizability of various ML algorithms across distinct domains of physics and materials science.[5] Enhanced standardization and a focus on transfer learning methods are necessary to accelerate the predictive capabilities of the field.

## VII. Future Research Trajectories and Recommendations

Based on the synthesis of the literature, future research in AI and thermodynamics must prioritize three core trajectories: methodological development focused on physical consistency and interpretability, expansion into industrial-scale sustainable systems, and theoretical breakthroughs in non-equilibrium processes.

## 7.1. Advancing Hybrid Theory-Informed Modeling and Explainable AI (XAI)

Continued research must focus on generalizing and scaling theory-constrained models such as PINNs and FE-NNs to manage increasingly complex systems. The principle of enforcing theoretical consistency—by embedding Maxwell relations, conservation laws, and other fundamental constraints into the loss function—must be rigorously applied to improve model generalizability and physical plausibility.[31] Furthermore, the dual challenge of accuracy and interpretability necessitates continued innovation in XAI. Future work should refine methods like TIERA and interpretation entropy to provide quantifiable, model-agnostic explanations, thereby building essential trust and accountability, particularly in high-stakes thermodynamic applications.[8]

## 7.2. Large-Scale Simulation and Sustainable Systems Optimization

The application of AI must be significantly expanded to address macro-level sustainability challenges. This includes leveraging reinforcement learning and complex systems theory for the dynamic optimization of large-scale energy systems, smart grids, and renewable energy infrastructure. Surrogate modeling will remain crucial for streamlining the design and analysis phases of complex chemical and thermal processes, accelerating the deployment of next-generation energy-efficient solutions and materials.

## 7.3. Pioneering Research in Non-Equilibrium Thermodynamics

The most challenging and potentially transformative research direction involves filling the theoretical gap in Non-Equilibrium Statistical Mechanics (NESM). Future research must develop systematic computational methodologies for dynamic thermodynamic processes, including those involving turbulence, chemical kinetics, and complex dissipative phenomena.[38] This work should focus on understanding the role of stochasticity, quantifying entropy production rates in realistic physical systems, and developing AI frameworks that are inherently grounded in the fluctuation theorems of non-equilibrium physics.[36] This theoretical unification holds the key not only to modeling realistic kinetic phenomena but also to

advancing the fundamental theoretical understanding of AI itself.

**Works cited**


1. (PDF) Optimization of Thermodynamic Systems Using Machine Learning Algorithms, accessed on October 4, 2025, https://www.researchgate.net/publication/387802836_Optimization_of_Thermodynamic_Systems_Using_Machine_Learning_Algorithms
2. Integrating artificial intelligence into thermodynamics: A new paradigm for sustainable future, accessed on October 4, 2025, https://pubs.aip.org/aip/adv/article/15/6/060701/3348321/Integrating-artificial-intelligence-into
3. Efficient Machine Learning Force Field for Large-Scale Molecular Simulations of Organic Systems | CCS Chemistry - Chinese Chemical Society, accessed on October 4, 2025, https://www.chinesechemsoc.org/doi/10.31635/ccschem.024.202404785
4. Machine Learning of Coarse-Grained Molecular Dynamics Force Fields - PMC, accessed on October 4, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC6535777/
5. Revolutionizing physics: a comprehensive survey of machine learning applications, accessed on October 4, 2025, https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2024.1322162/full
6. The AI-Driven Transformation in New Materials Manufacturing and the Development of Intelligent Sports - MDPI, accessed on October 4, 2025, https://www.mdpi.com/2076-3417/15/10/5667
7. [2302.06584] Thermodynamic AI and the fluctuation frontier - ar5iv - arXiv, accessed on October 4, 2025, https://ar5iv.labs.arxiv.org/html/2302.06584
8. Thermodynamics-inspired explanations of artificial intelligence, accessed on October 4, 2025, https://pubmed.ncbi.nlm.nih.gov/39251574/
9. Replacing Neural Networks by Optimal Analytical Predictors for the Detection of Phase Transitions | Phys. Rev. X, accessed on October 4, 2025, https://link.aps.org/doi/10.1103/PhysRevX.12.031044
10. Identifying Quantum Phase Transitions using Artificial Neural Networks on Experimental Data - arXiv, accessed on October 4, 2025, https://arxiv.org/pdf/1809.05519
11. Iterative charge equilibration for fourth-generation high-dimensional neural network potentials - arXiv, accessed on October 4, 2025, https://arxiv.org/html/2502.07907v1
12. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer - PMC, accessed on October 4, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC7811002/
13. Machine Learning Force Fields - PMC - PubMed Central, accessed on October 4, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC8391964/
14. High-dimensional neural-network potentials for multicomponent systems: Applications to zinc oxide | Phys. Rev. B - Physical Review Link Manager, accessed

on October 4, 2025, https://link.aps.org/doi/10.1103/PhysRevB.83.153101
15. Atom-centered symmetry functions for constructing high-dimensional neural network potentials - PubMed, accessed on October 4, 2025, https://pubmed.ncbi.nlm.nih.gov/21341827/
16. Machine learning force fields for improved materials modeling - Schrödinger, accessed on October 4, 2025, https://www.schrodinger.com/materials-science/learn/white-papers/machine-learning-force-fields-for-improved-materials-modeling/
17. [2409.11037] A High-Dimensional Neural Network Potential for Co$_3$O$_4$ - arXiv, accessed on October 4, 2025, https://arxiv.org/abs/2409.11037
18. Force fields and molecular dynamics simulations, accessed on October 4, 2025, https://www.neutron-sciences.org/articles/sfn/pdf/2011/01/sfn201112009.pdf
19. [1503.03585] Deep Unsupervised Learning using Nonequilibrium Thermodynamics - arXiv, accessed on October 4, 2025, https://arxiv.org/abs/1503.03585
20. Deep Unsupervised Learning using Nonequilibrium Thermodynamics, accessed on October 4, 2025, https://proceedings.mlr.press/v37/sohl-dickstein15.html
21. Machine learning of consistent thermodynamic models using automatic differentiation, accessed on October 4, 2025, https://laro.lanl.gov/esploro/outputs/journalArticle/Machine-learning-of-consistent-thermodynamic-models/9916396738703761
22. Machine learning of consistent thermodynamic models using automatic differentiation | Phys. Rev. E - Physical Review Link Manager, accessed on October 4, 2025, https://link.aps.org/doi/10.1103/PhysRevE.105.045301
23. Machine Learning Aided Design and Optimization of Thermal Metamaterials - PMC, accessed on October 4, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC11009967/
24. Machine learning prediction of thermodynamic stability and electronic properties of 2D layered conductive metal–organic frameworks | APL Materials | AIP Publishing, accessed on October 4, 2025, https://pubs.aip.org/aip/apm/article/13/9/091103/3361508/Machine-learning-prediction-of-thermodynamic
25. Machine Learning for the Prediction of Thermodynamic Properties in Amorphous Silicon, accessed on October 4, 2025, https://www.mdpi.com/2076-3417/15/10/5574
26. Quantum-based machine learning and AI models to generate force field parameters for drug-like small molecules - Frontiers, accessed on October 4, 2025, https://www.frontiersin.org/journals/molecular-biosciences/articles/10.3389/fmolb.2022.1002535/full
27. Physics-Informed Neural Networks for Advanced Thermal Management in Electronics and Battery Systems: A Review of Recent Developments and Future Prospects - MDPI, accessed on October 4, 2025, https://www.mdpi.com/2313-0105/11/6/204
28. Physics-Informed Fully Convolutional Networks for Forward Prediction of

Temperature Field and Inverse Estimation of Thermal Diffusivity | J. Comput. Inf. Sci. Eng. | ASME Digital Collection, accessed on October 4, 2025, https://asmedigitalcollection.asme.org/computingengineering/article/24/11/111004/1194685/Physics-Informed-Fully-Convolutional-Networks-for
29. Physics-Informed Neural Networks (PINNs) for Heat Transfer Problems, accessed on October 4, 2025, https://pics.upenn.edu/physics-informed-neural-networks-pinns-for-heat-transfer-problems/
30. Physics-informed neural network for the heat equation under imperfect contact conditions and its error analysis - AIMS Press, accessed on October 4, 2025, https://www.aimspress.com/article/doi/10.3934/math.2025364?viewType=HTML
31. Bridging Machine Learning and Thermodynamics for Accurate pKa Prediction | JACS Au, accessed on October 4, 2025, https://pubs.acs.org/doi/10.1021/jacsau.4c00271
32. What are the trade-offs between explainability and accuracy in AI models? - Milvus, accessed on October 4, 2025, https://milvus.io/ai-quick-reference/what-are-the-tradeoffs-between-explainability-and-accuracy-in-ai-models
33. Demystifying the Accuracy-Interpretability Trade-Off: A Case Study of Inferring Ratings from Reviews - arXiv, accessed on October 4, 2025, https://arxiv.org/html/2503.07914v1
34. Interpretable or Accurate? Why Not Both? - Towards Data Science, accessed on October 4, 2025, https://towardsdatascience.com/interpretable-or-accurate-why-not-both-4d9c73512192/
35. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics - By Alex Nguyen, accessed on October 4, 2025, https://dev.to/alex-nguyen-duy-anh/deep-unsupervised-learning-using-nonequilibrium-thermodynamics-by-alex-nguyen-3hcb
36. [2306.03521] Machine learning in and out of equilibrium - arXiv, accessed on October 4, 2025, https://arxiv.org/abs/2306.03521
37. Challenges and Opportunities for Machine Learning in Multiscale Computational Modeling, accessed on October 4, 2025, https://asmedigitalcollection.asme.org/computingengineering/article/23/6/060808/1163269/Challenges-and-Opportunities-for-Machine-Learning
38. Non-Equilibrium Foundations Should Underpin Generative Processes in Complex Dynamical Systems - arXiv, accessed on October 4, 2025, https://arxiv.org/html/2505.18621v1
39. Nonequilibrium thermodynamics of uncertain stochastic processes | Phys. Rev. Research, accessed on October 4, 2025, https://link.aps.org/doi/10.1103/PhysRevResearch.6.013021
40. Limitations of theory-informed machine learning algorithms for the prediction and exploration of molecular properties: Solvation free energy as a case study - AIP Publishing, accessed on October 4, 2025, https://pubs.aip.org/aip/aml/article/3/3/036107/3356268/Limitations-of-theory-info

rmed-machine-learning
41. Artificial Intelligence-Powered Materials Science - PMC, accessed on October 4, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC11803041/