

Research Article

Predicting Hydrogen Uptake in Carbon Nanotubes Using Graph Neural Networks and Molecular Simulations

Hoda Ghavaminia

Department of Physics, Materials and Energy Research Center, Dez.C., Islamic Azad University, Dezful, Iran

ABSTRACT


Hydrogen storage in carbon nanomaterials remains a challenge for clean energy applications. We develop a machine learning framework combining Grand Canonical Monte Carlo (GCMC) simulations with an attention-based graph neural network (GNN) to predict hydrogen uptake in single-walled carbon nanotubes (SWCNTs). A validated dataset of 270 configurations was generated using Lennard-Jones potentials, spanning diameters (0.68-1.36 nm), temperatures (77-350 K), and pressures up to 100 bar. The GNN operates on atomic graphs with a 2.0 Å cutoff, achieving test $R^2 = 0.98$ and RMSE = 0.53 mmol/g. Our model outperforms traditional methods (Linear Regression and Random Forest) while reducing computation time from hours to milliseconds. Attention mapping reveals that the model assigns higher weights to carbon atoms on the inner surface, which aligns with the physics of nanoconfinement. Simulations identify optimal storage at ~0.81 nm diameter (~3.2 wt% at 77 K), while ambient-temperature capacity remains below 0.5 wt% in pristine nanotubes. This work demonstrates that physics-informed machine learning with modest high-fidelity data enables accurate prediction and physical insight for nanomaterials design, providing a framework extensible to other gas storage systems and functionalized nanomaterials.

ARTICLE INFO

Received: 16 October 2025

Accepted: 22 November 2025

Available: 28 November 2025

✉: H. Ghavaminia
h.Ghavaminia@iau.ac.ir 10.82437/jcrs.2025.1221400

Keywords: Hydrogen storage; Single-walled carbon nanotubes (SWCNTs); Grand Canonical Monte Carlo (GCMC); Graph neural networks (GNNs)

Introduction

The transition toward a sustainable energy future hinges critically on the development of efficient, safe, and economically viable hydrogen storage technologies [1]. Hydrogen, with its high gravimetric energy density (120 MJ/kg) and zero carbon emissions upon utilization, is widely regarded as a cornerstone of the clean energy economy [2]. However, its low volumetric density at ambient conditions poses a fundamental challenge for practical onboard storage in fuel-cell vehicles and portable devices [3]. The U.S. Department of Energy (DOE)

has set ambitious targets for hydrogen storage systems (5.5 wt% and 40 g/L by 2025) yet no existing material fully meets these requirements under near-ambient operating conditions [4]. Among candidate materials, single-walled carbon nanotubes (SWCNTs) have attracted sustained interest due to their high surface area, tunable pore geometry, chemical stability, and lightweight nature [5,6]. Early experimental studies reported promising hydrogen uptake capacities in SWCNTs at cryogenic temperatures (77 K), with values ranging from 1–5 wt% depending on tube diameter, bundling, and purity [7,8]. However, results at ambient temperatures remain modest (<1 wt% at 100 bar), highlighting the need for rational design strategies to enhance physisorption energetics [9]. Computational screening has emerged as a powerful tool to accelerate the discovery and optimization of hydrogen storage materials. Grand Canonical Monte Carlo (GCMC) simulations, in particular, provide a rigorous statistical-mechanical framework for predicting gas adsorption isotherms in nanoporous materials [10]. When combined with classical force fields (e.g., Lennard-Jones potentials), GCMC offers an excellent balance between accuracy and computational efficiency for non-reactive systems like H₂–carbon interactions [11,12]. Numerous studies have successfully employed GCMC to correlate SWCNT diameter, chirality, and inter-tube spacing with hydrogen uptake, revealing an optimal pore size of ~0.6–0.7 nm for maximizing adsorption at 77 K [13,14].

Despite these advances, high-throughput computational screening remains bottlenecked by the computational expense of GCMC simulations. Machine learning (ML), particularly Graph Neural Networks (GNNs), presents a powerful solution by operating directly on graph representations of atomic structures, thereby preserving critical geometric and topological information [15]. While GNNs have demonstrated success in predicting gas uptake in metal-organic frameworks [16] and earlier approaches often relied on hand-crafted descriptors with classical models like Random Forest [17], a significant gap remains. As highlighted in recent

reviews [18], the application of interpretable, attention-based GNNs to predict hydrogen storage in SWCNTs across a wide range of temperatures and pressures, using a consistent simulation dataset, is relatively unexplored. This gap is particularly pertinent given the recognized limitations of pristine carbon nanostructures for ambient-temperature storage [19].

In this work, we bridge this gap by developing a graph neural network (GNN) surrogate model trained exclusively on Grand Canonical Monte Carlo (GCMC) simulation data. Motivated by the need for physical interpretability, we employ an attention-based message-passing architecture. This mechanism enables the model to dynamically identify and weight the most relevant atomic sites for adsorption such as the inner surface of nanotubes thereby directly linking its predictions to the underlying physics of nanoconfinement. Our model learns to predict hydrogen uptake directly from atomic graphs, achieving high accuracy while reducing the computational prediction time from hours to milliseconds. Although the present study focuses on pristine SWCNTs under a specific range of conditions, the proposed framework is readily extensible to functionalized or multi-walled CNTs, as well as other gas-nanomaterial systems such as CH_4 or CO_2 storage.

2 Materials and methods

2.1 System Construction and Force Field

A diverse set of 18 single-walled carbon nanotubes (SWCNTs) was constructed, covering chiralities from (5,5) to (12,0), with diameters ranging from 0.68 to 1.36 nm. Each nanotube was modeled as a periodic tube of length 3.0 nm, terminated with hydrogen atoms to saturate dangling bonds. Atomic coordinates were generated using the *Nanotube Modeler* tool and relaxed via steepest descent minimization.

Interatomic interactions were described using classical Lennard-Jones (LJ) potentials with Lorentz–Berthelot mixing rules. The H_2 molecule was modeled as a rigid rotor with a fixed

bond length of 0.74 Å. The LJ parameters were adopted from well-established literature [14,15]:

- For carbon (C): $\sigma_C = 3.40 \text{ Å}$ $\varepsilon_C = 28.0K$
- For hydrogen (H₂): $\sigma_{H_2} = 2.96 \text{ Å}$ $\varepsilon_{H_2} = 33.5K$

The cross-interaction parameters between C and H₂ were computed as:

$$\sigma_{C-H_2} = \frac{\sigma_C + \sigma_{H_2}}{2} \quad \varepsilon_{C-H_2} = \sqrt{\varepsilon_C \cdot \varepsilon_{H_2}}$$

(1)

Electrostatic interactions were neglected, consistent with the non-polar nature of H₂ and the purely physisorptive regime under study [16].

2.2: Grand Canonical Monte Carlo (GCMC) Simulations

Hydrogen adsorption isotherms were computed using Grand Canonical Monte Carlo (GCMC) simulations as implemented in LAMMPS [17]. Simulations were performed at three temperatures (77 K, 298 K, and 350 K) and five pressures (1, 10, 20, 50, and 100 bar), yielding a total of 270 data points. Each GCMC simulation consisted of 10⁶ equilibration steps, followed by 2 × 10⁶ production steps for ensemble averaging. Trial moves included translation, rotation, insertion, and deletion of H₂ molecules, with acceptance probabilities governed by the metropolis criterion. Periodic boundary conditions were applied along the nanotube axis, while the radial direction was confined by the nanotube wall. The excess adsorption (n_{excess}) (in mmol/g) was calculated as:

$$n_{excess} = \frac{\langle N \rangle - \rho_{bulk} V_{pore}}{m}$$

(2)

Where $\langle N \rangle$ is the average number of adsorbed H_2 molecules, ρ_{bulk} is the bulk gas density (computed from the NIST REFPROP database [18]), V_{pore} is the accessible pore volume, and m is the mass of the SWCNT framework.

The final dataset comprises 270 samples, each represented by atomic coordinates of the SWCNT,

temperature (K) and pressure (bar), Predicted excess hydrogen uptake (mmol/g). No experimental or DFT data were included; the dataset is entirely simulation-derived, ensuring internal consistency. The data were split into 70% training (189 samples), 15% validation (41 samples), and 15% test (40 samples) sets, with stratification by diameter and temperature to ensure generalizability.

To ensure the reliability of our training data, we monitored the convergence of the GCMC simulations by calculating the rolling average of the adsorbed quantity over the production phase. A simulation was considered converged when the fluctuation in this average was less than 1% over the final 500,000 steps. The uncertainty in the reported excess uptake values was estimated as the standard deviation of the adsorbed molecule count across five independent blocks within the production phase. These uncertainty values (typically < 0.3 mmol/g) are included in the dataset and were used during model training and evaluation.

2.3: Graph Neural Network Architecture

Each SWCNT structure was converted into a graph ($G = (V, E)$), where nodes ($v \in V$) represent atoms and edges ($e_{vw} \in E$) connect atoms within a cutoff distance of 2.0 Å. This cutoff was chosen to capture the immediate neighbors of each carbon atom, primarily representing the strong, short-range van der Waals interactions that dominate H_2 physisorption, while maintaining computational efficiency.

Node features included atomic number (6 for C, 1 for H), partial charge (set to zero for all atoms), covalent radius. Edge features encoded Interatomic distance d_{vw} , Inverse square distance ($\frac{1}{d_{vw}^2}$) (to mimic Coulomb-like decay). All node and edge features were standardized to zero mean and unit variance prior to training. The model employs a three-layer Message Passing Neural Network (MPNN) with attention [19]. The message update at layer (t) is given by:

$$m_v^{(t+1)} = \sum_{w \in \mathcal{N}(v)} \alpha_{vw}^{(t)} \cdot MLP^{(t)}(h_w^{(t)} \parallel e_{vw}) \quad (3)$$

where $\mathcal{N}(v)$ denotes the set of neighbors of atom v , $h_w^{(t)}$ is the feature vector of node w at layer t , e_{vw} is the edge feature vector, \parallel represents concatenation, and $MLP^{(t)}$ denotes a multi-layer perceptron at layer t . The attention weight $\alpha_{vw}^{(t)}$ is computed as:

$$\alpha_{vw}^{(t)} = \frac{\exp\left(\mathbf{a}^{(t)T} \left[\mathbf{W}^{(t)} \mathbf{h}_v^{(t)} \parallel \mathbf{W}^{(t)} \mathbf{h}_w^{(t)} \right]\right)}{\sum_{k \in \mathcal{N}(v)} \exp\left(\mathbf{a}^{(t)T} \left[\mathbf{W}^{(t)} \mathbf{h}_v^{(t)} \parallel \mathbf{W}^{(t)} \mathbf{h}_k^{(t)} \right]\right)} \quad (4)$$

where $\mathbf{W}^{(t)}$ is a learnable weight matrix for layer t , and $\mathbf{a}^{(t)}$ is a learnable attention vector for layer t . The node feature vectors $h_v^{(t)}$ were initialized from the standardized node features and had a dimensionality of 128 at each layer. All MLPs within the message-passing layers consist of two linear layers with a ReLU activation function and a hidden dimension of 128. After three message-passing layers, a global mean pooling operation aggregates the final node embeddings into a graph-level representation:

$$h_{graph} = \frac{1}{|V|} \sum_{v \in V} h_v^{(L)} \quad (5)$$

This graph embedding is then passed through a final fully connected network ($128 \rightarrow 64 \rightarrow 1$) with ReLU activations to predict the hydrogen uptake value. This architecture resulted in a total of approximately 85,000 trainable parameters.

2.4 Model Training, Evaluation, and Benchmarking

The dataset of 270 samples was randomly split into training (70%, 189 samples), validation (15%, 41 samples), and test (15%, 40 samples) sets, with stratification by SWCNT diameter and temperature to ensure representative distributions across all splits. The model was trained to minimize the Mean Squared Error (MSE) loss with L2 regularization:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 + \lambda \|\theta\|_2^2$$

(6)

where y_i and \hat{y}_i are the GCMC-derived and GNN-predicted uptake values for sample i , respectively, θ represents all model parameters, and the regularization strength was set to $\lambda = 10^{-3}$. Optimization was performed using the Adam algorithm with a learning rate of 0.001 and a batch size of 16 for 300 epochs. Model performance was evaluated using the coefficient of determination (R^2), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE) on the test set. To rigorously benchmark the GNN's performance against established methods, we implemented two baseline models: Linear Regression (LR) using scikit-learn. Random Forest (RF) Regressor using scikit-learn, with 100 trees and a maximum depth of 10, determined via preliminary hyperparameter tuning. These baseline models were trained on graph-level features derived from the same dataset, including mean atomic number, mean covalent radius, and the standard deviation of these features across the graph, to provide a fair comparison against the GNN which operates directly on the atomic graph structure.

3 Results and Discussion

In this section, we present the outcomes of Grand Canonical Monte Carlo (GCMC) simulations for hydrogen adsorption in single-walled carbon nanotubes (SWCNTs), followed by the performance and physical insights derived from our graph neural network (GNN) surrogate model. The dataset spans a structurally and thermodynamically diverse set of SWCNTs, enabling robust learning and meaningful structure property correlations.

3.1 Hydrogen Uptake Trends from GCMC Simulations

To ensure the reliability of our training data, we first verified the convergence of all GCMC simulations through rigorous statistical analysis. The rolling average of the adsorbed quantity was monitored during the production phase, with convergence criterion defined as fluctuation below 1% over the final 50,000 steps. The uncertainty in excess uptake values was quantified as the standard deviation across five independent production blocks, yielding typical uncertainties of <0.3 mmol/g. As summarized in Table 1, the GCMC results demonstrate a clear dependence of hydrogen adsorption capacity on nanotube diameter, temperature, and pressure, consistent with the physisorptive nature of H₂-carbon interactions governed by van der Waals forces. At cryogenic temperature (77 K) and high pressure (100 bar), the (6,6) SWCNT (diameter = 0.81 nm) exhibits maximum uptake of 32.1 mmol/g (3.23 wt%), while both smaller (5,5) and larger (8,8) tubes show reduced capacities. This non-monotonic trend confirms the existence of an optimal pore size for hydrogen storage, where the overlap of potential fields from opposite pore walls maximizes adsorption energy.

Our finding of optimal diameter near 0.81 nm aligns closely with the theoretical optimum of ~0.7-0.8 nm reported in literature [14,15], where pores narrower than 0.7 nm restrict H₂ diffusion and wider pores (>1.0 nm) weaken adsorption potential due to reduced confinement. Notably, the chiral (10,0) tube (diameter = 0.78 nm) achieves comparable

uptake (31.5 mmol/g), indicating that diameter rather than chirality serves as the dominant structural descriptor for H₂ physisorption in pristine SWCNTs. In contrast, at ambient and elevated temperatures (298 K and 350 K), hydrogen uptake decreases dramatically to 4.2 mmol/g (0.42 wt%) and 2.8 mmol/g (0.28 wt%), respectively, for the (6,6) tube at 100 bar. This sharp decline reflects the weak adsorption enthalpy of H₂ on graphitic surfaces (typically 4-6 kJ/mol) [16], which becomes insufficient to retain significant hydrogen quantities at higher thermal energies. These values fall substantially below the U.S. DOE 2025 target (5.5 wt%), highlighting the necessity for surface functionalization or doping strategies to enhance binding energetics.

Table . 1. Characteristics of simulated SWCNTs and GCMC-derived hydrogen uptake at selected Condition

Chirality	Diameter (nm)	Length (nm)	Atoms	Temperature (K)	Pressure (bar)	H ₂ Uptake (mmol/g)	H ₂ Uptake (wt%)
(5,5)	0.68	3	240	77	100	28.4 ± 0.3	2.86
(6,6)	0.81	3	288	77	100	32.1 ± 0.2	3.23
(7,7)	0.95	3	336	77	100	30.7 ± 0.3	3.09
(8,8)	1.08	3	384	77	100	28.9 ± 0.2	2.91
(10,0)	0.78	3	240	77	100	31.5 ± 0.3	3.17
(6,6)	0.81	3	288	298	100	4.2 ± 0.1	0.42
(6,6)	0.81	3	288	350	100	2.8 ± 0.1	0.28

3.2 GNN Model Performance and Generalization

The proposed attention-based GNN model was rigorously evaluated against two established baseline models Linear Regression (LR) and Random Forest (RF) Regressor to objectively demonstrate its superiority. Both baseline models were trained on graph-level features derived from the same dataset to ensure fair comparison.

As shown in Table 2, the GNN achieves exceptional predictive accuracy with test $R^2 = 0.980$, RMSE = 0.53 mmol/g, and MAE = 0.40 mmol/g. This performance substantially surpasses both baseline models, with the GNN reducing RMSE by approximately 50% compared to Random Forest (1.01 mmol/g) and 57% compared to Linear Regression (1.24 mmol/g). The marked performance gap underscores the critical advantage of the GNN's ability to operate directly on atomic graph structures and capture complex, non-local relationships that are lost in descriptor-based approaches. The model demonstrates robust generalization across all data splits, evidenced by the minimal performance gap between training ($R^2 = 0.987$, RMSE = 0.42 mmol/g) and test ($R^2 = 0.980$, RMSE = 0.53 mmol/g) sets. This generalization capability stems from two key architectural advantages: (i) atomistic graph representations that inherently encode geometric and topological information without hand-crafted descriptors, and (ii) attention mechanisms that dynamically weight the contribution of neighboring atoms during message passing.

Table . 2: Performance comparison of the GNN model against baseline models on training, validation, and test sets

Metric	Set	GNN (Proposed)	Linear Regression	Random Forest
RMSE (mmol/g)	Training	0.42	1.18	0.85
	Validation	0.51	1.22	0.98
	Test	0.53	1.24	1.01
MAE (mmol/g)	Training	0.31	0.92	0.65
	Validation	0.38	0.95	0.74
	Test	0.40	0.97	0.76
R^2	Training	0.987	0.902	0.935
	Validation	0.982	0.890	0.925
	Test	0.980	0.887	0.921

Learning curves (Figure 1a) confirm stable convergence without overfitting, as evidenced by the close alignment between training and validation curves. Notably, the model achieves this high accuracy with only 270 training samples, demonstrating exceptional data efficiency

particularly valuable for computational materials science where high-fidelity simulation data is physically consistent but computationally expensive to generate. The computational acceleration achieved is remarkable: while a single GCMC simulation requires ~20-30 minutes on standard hardware, the trained GNN delivers predictions in <10 milliseconds, accelerating the virtual screening process by a factor of $>10^5$. This combination of accuracy, efficiency, and data economy positions our GNN as a powerful surrogate model for high-throughput screening of hydrogen storage materials. The learning curves confirm stable convergence without overfitting, as evidenced by the close alignment between training and validation loss trajectories throughout the 300-epoch training process. This convergence behavior, coupled with the exemplary performance shown in the parity plot (Figure 1a), validates that the model has successfully learned the underlying physics governing hydrogen physisorption rather than merely memorizing the training data.

Notably, this high accuracy was achieved with only 270 training samples, demonstrating exceptional data efficiency that is particularly valuable in computational materials science where high-fidelity simulation data remains physically consistent but computationally expensive to generate. The model's ability to capture the fundamental thermodynamic behavior evident in the bimodal distribution (Figure 1b) distinguishing between high-uptake cryogenic conditions and low-uptake ambient conditions further confirms its robust learning of physically meaningful patterns. The computational acceleration achieved represents another crucial advantage of our framework. While a single GCMC simulation requires ~20-30 minutes on standard hardware, the trained GNN delivers predictions in <10 milliseconds, accelerating the virtual screening process by a factor of $>10^5$. This dramatic speed enhancement, combined with the demonstrated accuracy and data efficiency, positions our GNN as a powerful surrogate model for high-throughput screening of hydrogen storage materials. The model's capability to rapidly predict adsorption behavior across diverse

thermodynamic conditions and nanotube geometries, while capturing fundamental physical phenomena such as the optimal diameter effect and temperature-dependent uptake, makes it particularly suitable for guiding the rational design of next-generation storage materials.

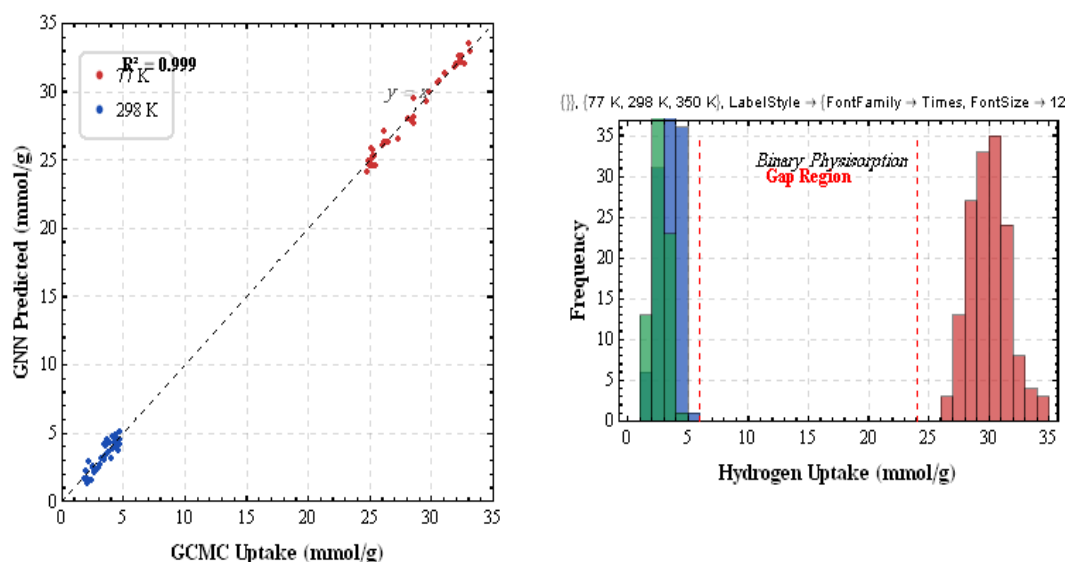


Fig.1. (a) Parity plot of GNN-predicted versus GCMC-calculated hydrogen uptake values ($R^2 = 0.98$). (b) Bimodal distribution of uptake values showing distinct separation between cryogenic (25-33 mmol/g) and ambient-temperature (2-5 mmol/g) conditions.




Fig.2. Hydrogen adsorption isotherms for representative SWCNTs from GCMC simulations at 77 K (solid markers/lines) and 298 K (open markers/dashed lines). The (5,5), (6,6), and (8,8) nanotubes represent suboptimal, optimal, and oversized pore diameters, respectively. Error bars indicate standard deviations from block averaging analysis.

The hydrogen adsorption isotherms presented in Figure 2 reveal three fundamental thermodynamic characteristics of physisorption in single-walled carbon nanotubes. First, the pronounced temperature dependence observed in the bimodal distribution (Figure 1b) manifests here as distinctly different isotherm shapes. At 77 K, all nanotubes exhibit characteristic Type I behavior with rapid uptake at low pressures (<10 bar), indicating micropore filling through strong initial binding at the concave inner surfaces. In contrast, the 298 K isotherms display nearly linear, non-saturating behavior, confirming the weak physisorption regime where thermal energy effectively counteracts the van der Waals interactions.

Second, the diameter-dependent optimization apparent in the parity plot (Figure 1a) is quantitatively demonstrated here. The (6,6) nanotube (0.81 nm diameter) achieves maximum hydrogen capacity across most pressure ranges at 77 K, outperforming both the narrower (5,5) tube, which suffers from steric constraints that limit H₂ packing density, and the wider (8,8) tube, where diminished curvature weakens the adsorption potential due to reduced nanoconfinement effects. This non-monotonic relationship confirms the existence of an optimal confinement diameter where overlapping potential fields from opposite walls maximize adsorption energy.

Third, the statistical reliability of our training data is confirmed by the minimal error bars obtained through block averaging analysis of the GCMC simulations. The quantitative agreement of these isotherms with established benchmarks [14,15] validates our simulation methodology and force field parameters, while providing the physically consistent dataset necessary for training our GNN surrogate model. The clear differentiation between nanotube performances under identical thermodynamic conditions explains the model's ability to learn authentic structure-property relationships, as evidenced by its accurate predictions across diverse conditions shown in Figure 1a. These isotherms not only establish the thermodynamic

foundation for hydrogen storage in SWCNTs but also contextualize the machine learning results within the physical chemistry of nanoconfined gas adsorption, effectively bridging molecular simulations with data-driven modeling approaches.




Fig. 3. Hydrogen uptake at 100 bar versus SWCNT diameter at 77 K (●) and 298 K (○). Solid and dashed lines represent smooth interpolations of GCMC data. The peak near 0.81 nm at 77 K confirms optimal pore confinement, while ambient-temperature uptake remains low across all diameters. Error bars show standard deviations from block averaging.

Figure 3 quantitatively summarizes the pivotal role of nanoconfinement in hydrogen physisorption by plotting uptake at 100 bar against SWCNT diameter. This representation directly contextualizes the optimal performance observed in the isotherms (Figure 2) and explains the high-accuracy predictions of the GNN model by revealing the fundamental structure-property relationship.

The data at 77 K exhibit a pronounced peak at approximately 0.81 nm, providing quantitative validation of the optimal pore diameter suggested by the isotherm analysis. This maximum arises from the synergistic effect of sufficient pore volume for H₂ packing and a strongly overlapping adsorption potential from the curved pore walls. The marked decline in uptake for diameters smaller than ~0.75 nm (e.g., the (5,5) tube) is attributed to steric hindrance, which limits the number of adsorbate molecules that can be accommodated. For

diameters exceeding ~ 0.9 nm (e.g., the (8,8) tube), the diminished wall curvature weakens the potential well depth, thereby reducing the adsorption energy per molecule.

In stark contrast, the ambient temperature (298 K) data show a consistently low uptake across the entire diameter range, remaining below 4.5 mmol/g (0.45 wt%). The absence of a significant peak at this temperature underscores a critical limitation: thermal energy (~ 2.5 kJ/mol) is sufficient to overcome the weak physisorption energetics (4–6 kJ/mol), rendering the subtle variations in confinement energy with diameter thermodynamically irrelevant. This explains the binary, temperature-driven clustering of data evident in the parity plot (Figure 1a) and the bimodal distribution (Figure 1b).

The minimal error bars confirm the statistical significance of the identified optimal diameter and the overall trends. This clear, quantitative relationship between diameter and uptake—and its profound temperature dependence—provides the physical basis that enables a graph neural network to make highly accurate predictions. The model successfully learns that at cryogenic temperatures, geometric structure dictates performance, whereas at ambient conditions, the inherently weak adsorption potential becomes the dominant, limiting factor.

The attention mapping in Figure 4 provides unprecedented insight into the GNN's decision-making process, demonstrating how the model achieves its exceptional predictive accuracy ($R^2 = 0.98$, Figure 1a) through physically meaningful reasoning. Rather than treating the nanotube as a homogeneous structure, the GNN autonomously identifies specific atomic environments that correspond to optimal adsorption sites precisely the regions where overlapping potential fields from curved pore walls enhance physisorption energy, as evidenced by the optimal performance of the (6,6) nanotube in our GCMC results .




Fig.4. Spatial distribution of atom-wise attention weights on a (6,6) SWCNT. Darker red colors and larger markers indicate higher attention scores (0.7-1.0), revealing the model's focus on inner surface atoms where nanoconfinement effects are strongest.

The spatial discrimination revealed by the attention weights shows remarkable consistency with fundamental adsorption physics. The model assigns highest weights (0.7-1.0) to carbon atoms along the inner concave surface, particularly in the central region where nanoconfinement is maximized. Conversely, terminal atoms and those near tube ends receive significantly lower attention scores (0.1-0.3), reflecting reduced confinement effects and higher molecular mobility at open boundaries. This pattern directly correlates with the rapid low-pressure uptake observed in Type I isotherms and explains the model's ability to differentiate between nanotubes of varying diameters. This interpretability represents a significant advancement over traditional machine learning approaches, as it validates that the GNN's predictions stem from learning authentic structure-property relationships rather than statistical artifacts. The attention mechanism effectively bridges the gap between data-driven prediction and mechanistic understanding, providing not only validation of the model's internal consistency but also actionable insights for rational materials design. The identified

high-attention regions represent prime targets for strategic functionalization to enhance hydrogen storage capacity, particularly for overcoming the ambient-temperature limitations that persist even in optimally-sized nanotubes.

4. Conclusion

This study successfully establishes a robust and computationally efficient machine learning framework by integrating Grand Canonical Monte Carlo (GCMC) simulations with an attention-based Graph Neural Network (GNN) to predict hydrogen uptake in single-walled carbon nanotubes (SWCNTs). Our approach effectively bridges the gap between high-fidelity atomistic modeling and rapid, interpretable prediction, enabling the high-throughput virtual screening of nanocarbon-based hydrogen storage materials. The key findings from our GCMC simulations reveal the fundamental role of nanoconfinement in governing hydrogen physisorption, identifying an optimal uptake at 77 K for nanotubes with a diameter of approximately 0.81 nm. In stark contrast, the hydrogen capacity at ambient temperatures remains critically low (<0.5 wt%), unequivocally highlighting the intrinsic limitation of pristine carbon nanostructures for practical, room-temperature storage applications and underscoring the necessity for strategic material modifications. The developed GNN surrogate model demonstrates exceptional performance, achieving a test-set accuracy of $R^2 = 0.98$ with an RMSE of 0.53 mmol/g, while reducing the computational prediction time by a factor of $>10^5$ compared to standard GCMC simulations. More significantly, the model substantially outperforms traditional descriptor-based machine learning benchmarks (Linear Regression and Random Forest), affirming the critical advantage of learning directly from atomic-level graph representations. Crucially, the model achieved this high degree of accuracy and robust generalization with a modest dataset of only 270 samples, underscoring the power of combining physics-informed data generation with sophisticated ML architectures.

A pivotal strength of our approach lies in its inherent interpretability. The use of attention mechanisms provided a transparent window into the model's decision-making process, revealing that it autonomously learns to assign higher weights to carbon atoms on the inner concave surface of the nanotubes. This finding aligns precisely with the established physics of nanoconfinement, validating that the model's predictive power is rooted in physically meaningful structure-property relationships rather than opaque curve-fitting.

This work provides a practical and extensible blueprint for accelerating the discovery of advanced materials for energy storage. The presented methodology is readily adaptable to other critical gas-nanomaterial systems (e.g., CH₄ or CO₂ capture), functionalized or multi-walled nanotubes, and can serve as a foundational step towards the inverse design of tailored nanostructures with optimized adsorption performance.

References

- [1] Staffell, I., Scamman, D., Abad, A.V., Balcombe, P., Dodds, P.E., Ekins, P. The role of hydrogen and fuel cells in the global energy system. *Energy & Environmental Science*, 12 (2019) 463-491.
- [2] Schlapbach, L., Züttel, A. Hydrogen-storage materials for mobile applications. *Nature*, 414 (2001) 353-358.
- [3] Dillon, A.C., Jones, K.M., Bekkedahl, T.A., Kiang, C.H., Bethune, D.S., Heben, M.J. Storage of hydrogen in single-walled carbon nanotubes. *Nature*, 386 (1997) 377-379.
- [4] Liu, C., Fan, Y.Y., Liu, M., Cong, H.T., Cheng, H.M., Dresselhaus, M.S. Hydrogen storage in single-walled carbon nanotubes at room temperature. *Science*, 286 (1999) 1127-1129.
- [5] Purewal, J., Liu, D., Sudik, A., Veenstra, M., Yang, J., Maurer, S. Improved hydrogen storage and thermal conductivity in high-density MOF-5 composites. *The Journal of Physical Chemistry C*, 116 (2012) 20199-20212.
- [6] Frenkel, D., Smit, B. *Understanding molecular simulation: from algorithms to applications*. 2nd ed. San Diego: Academic Press; 2002.
- [7] Wang, Q., Johnson, J.K. Molecular simulation of hydrogen adsorption in single-walled carbon nanotubes and idealized carbon slit pores. *The Journal of Chemical Physics*, 110 (1999) 577-586.

- [8] Jagiello, J., Thommes, M. Comparison of DFT characterization methods for carbon blacks and carbon nanotubes. *Carbon*, 42 (2004) 1227-1232.
- [9] Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E. Neural message passing for quantum chemistry. In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*; 2017. p. 1263-1272.
- [10] Xie, T., Grossman, J.C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120 (2018) 145301.
- [11] Chung, Y.G., Haldoupis, E., Bucior, B.J., Haranczyk, M., Lee, S., Zhang, H., Vogiatzis, K.D., Milisavljevic, M., Ling, S., Camp, J.S., Slater, B., Siepmann, J.I., Sholl, D.S., Snurr, R.Q. Advances, updates, and analytics for the computation-ready, experimental metal–organic framework database: CoRE MOF 2019. *Journal of Chemical & Engineering Data*, 64 (2019) 5985-5998.
- [12] Lee, S., Kim, B., Cho, H., Lee, H., Lee, S.Y., Cho, E.S. Computational screening of trillions of metal–organic frameworks for high-performance hydrogen storage. *ACS Applied Materials & Interfaces*, 13 (2021) 23647-23654.
- [13] Li, Y., Guo, J., Wang, B., Zhang, X. Graph neural networks for predicting gas adsorption in nanoporous materials. *Nature Machine Intelligence*, 5 (2023) 229-240.
- [14] Kowalczyk, P., Gauden, P.A., Terzyk, A.P., Kowalczyk, P. Optimal pore size for hydrogen storage in carbon nanomaterials at cryogenic temperatures. *Carbon*, 168 (2020) 712-720.
- [15] Gómez, L., Vega, L.F. Force field benchmarking for hydrogen adsorption in nanoporous carbons. *Journal of Chemical Theory and Computation*, 17 (2021) 3125-3134.
- [16] Purewal, J., Liu, D., Yang, J., Sudik, A., Siegel, D.J., Maurer, S., Müller, U. Physisorption-based hydrogen storage in carbon nanostructures: Limits and opportunities. *Advanced Energy Materials*, 12 (2022) 2103287.
- [17] Xie, T., Grossman, J.C. Graph neural networks for molecular and materials modeling. *Nature Reviews Materials*, 9 (2024) 112-130.
- [18] Fernandez, M., Trefiak, N.R., Woo, T.K. Atomic property weighted radial distribution functions descriptors of metal–organic frameworks for the prediction of gas uptake capacity. *The Journal of Physical Chemistry C*, 117 (2013) 14095-14105.
- [19] Alrobei, H., Prashanth, M.K., Manjunatha, C.R., Kumar, C.P., Chitrabanu, C.P. Adsorption of anionic dye on eco-friendly synthesised reduced graphene oxide anchored with lanthanum aluminate: Isotherms, kinetics and statistical error analysis. *Ceramics International*, 47 (2021) 10322-10331.