

Article

Prediction of H₂–CNT Interaction Energies on a Chiral (2,1) Carbon Nanotube Using Multilayer Perceptrons

Luis Josimar Vences Reynoso ¹, Roberto Alejo Eleuterio ¹, Everardo Efrén Granda Gutiérrez ²,
Daniel Villanueva Vázquez ¹, Juan Horacio Pacheco Sánchez ¹, Allan A. Flores Fuentes ²,
and Federico Del Razo López ^{1,*}

¹ Instituto Tecnológico de Toluca, Tecnológico Nacional de México, Av. Tecnológico s/n, Col. Agrícola, Metepec, Estado de México 52149, Mexico; lvencesr@toluca.tecnm.mx (L.J.V.R.); ralejoe@toluca.tecnm.mx (R.A.E.); daniel.vv@toluca.tecnm.mx (D.V.V.); jpachecos@toluca.tecnm.mx (J.H.P.S.)

² Centro Universitario UAEM Atlacomulco, Universidad Autónoma del Estado de México, Km. 60 Carretera Toluca-Atlacomulco, Atlacomulco, Estado de México 50450, Mexico; eegrandag@uaemex.mx (E.E.G.G.); aafloresf@uaemex.mx (A.A.F.F.)

* Correspondence: fdelrazol@toluca.tecnm.mx

Abstract

Accurate estimation of molecule–nanotube interaction energies is critical for the computational screening of carbon-based materials for hydrogen storage; however, density functional theory (DFT) calculations remain computationally expensive for extensive configurational sampling. In this work, we develop a multilayer perceptron (MLP) surrogate model to predict H₂–CNT interaction energies, represented by E_{ads} , for H₂ interactions with a chiral (2,1) carbon nanotube. A curated dataset comprising 696 configurations was generated using DMol³ (BIOVIA Materials Studio), varying intermolecular distance, molecular orientation, and interaction site across three regions: internal cavity, edges, and external surface. The proposed MLP architecture (64–32–1) incorporates GELU activation functions, L2 regularization, and dropout to improve generalization. The model achieves coefficients of determination in the range $R^2 = 0.90$ – 0.96 across all interaction regions, with particularly strong performance at the nanotube edges ($R^2 = 0.9358$, $MSE = 0.046$ eV²), as well as on the external surface ($R^2 = 0.9625$, $MSE = 0.574$ eV²) and within the internal cavity ($R^2 = 0.9051$, $MSE = 1.506$ eV²). The original E_{ads} distribution had a mean of 4.0955 eV and a sample standard deviation of 4.3189 eV. The elevated energy values observed in the internal cavity (up to 12 eV) are consistent with steric repulsion induced by geometric confinement rather than predictive artifacts. The trained MLP showed close agreement with DFT-derived trends, enabling exploration of interaction-energy landscapes spanning both attractive and repulsive regimes. These results indicate that MLP-based models trained on diverse configurational datasets provide a computationally efficient alternative for screening carbon nanostructures in hydrogen storage applications, without substantially compromising accuracy relative to first-principles methods.



Academic Editor: Sergei Manzhos

Received: 18 May 2026

Revised: 10 June 2026

Accepted: 25 June 2026

Published: 27 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: carbon nanotubes; multilayer perceptron; interaction energy; hydrogen storage; density functional theory

1. Introduction

In modern materials science, computational simulations have become a fundamental tool for the discovery and design of advanced functional materials [1]. In hydrogen storage research, carbon nanostructures remain attractive candidates because they combine low

density, high surface-to-volume ratio, chemical stability, and tunable surface geometry [2,3]. For H₂–CNT systems, the local molecule–nanotube interaction energy is a key quantity for identifying accessible interaction sites, repulsive regions, and configurational barriers associated with molecular uptake and release [4,5].

Density functional theory (DFT) provides a quantum-mechanical route for evaluating interaction energies from electronic-structure calculations. However, systematic exploration of intermolecular distances, molecular orientations, and region-specific interaction sites requires many individual DFT evaluations. The computational cost of this procedure increases rapidly with system size and configurational complexity [6], which limits high-throughput screening and fine-grained mapping of the H₂–CNT potential-energy landscape. This limitation motivates the use of machine-learning models that can approximate DFT-derived energy trends after being trained on representative configurational datasets.

Artificial neural networks, particularly multilayer perceptrons (MLPs), have been widely used as nonlinear regression models for physicochemical property prediction when suitable training data are available [7–9]. Nevertheless, many previous machine-learning studies related to hydrogen storage have focused on global adsorption capacities, averaged energetic quantities, or descriptor-based screening of large material families [10,11]. In contrast, fewer studies have addressed the local dependence of H₂–CNT interaction energy on explicit atomic coordinates, molecular orientation, and region-specific interaction sites. This gap is relevant because an atom-coordinate-based model can map internal-cavity, edge, and external-surface configurations, including both attractive and strongly repulsive regimes.

In the present work, the computed quantity is denoted as E_{ads} to preserve consistency with the total-energy expression and dataset notation. Nevertheless, E_{ads} is interpreted here as a generalized interaction-energy quantity rather than as an energy restricted only to stable adsorption states. Accordingly, negative or lower-energy values describe attractive physisorption-like configurations, whereas positive values describe repulsive or endothermic configurations that form part of the same H₂–CNT potential-energy landscape.

It is also important to emphasize that the chiral CNT(2,1) considered in this work is not proposed as a technologically optimized hydrogen storage material. Owing to its very small diameter, this finite nanotube model is used as a compact proof-of-concept system for evaluating whether an atom-coordinate-based MLP can learn local DFT-derived H₂–CNT interaction-energy trends. Therefore, the scope of the present study is methodological rather than a direct assessment of practical storage performance in experimentally optimized CNTs.

The main contributions of this work are summarized as follows:

- A DFT-based configurational dataset was constructed for a finite chiral CNT(2,1) interacting with molecular H₂ across internal-cavity, edge, and external-surface regions.
- An MLP surrogate model was implemented using explicit atomic coordinates, molecular orientation, and interaction distance as input information to approximate DFT-derived H₂–CNT interaction energies.
- The predictive performance of the model was evaluated by region, allowing for a comparison of attractive, weakly interacting, and repulsive configurational regimes.

2. Related Work

Machine learning offers tools for exploring high-dimensional material spaces and property prediction. Recent reviews assert that these computational methods enhance the synthesis and characterization of CNTs and help to expedite the advancement of energy-related applications [12].

Specialized neural network architectures that can describe carbon–hydrogen systems with high accuracy have been created. These architectures can be used on different structures and make it possible to run high-throughput simulations [13]. Additionally, the integration of *ab initio* methodologies, including DFT, with explainability frameworks such as SHAP (Shapley Additive Explanations) has facilitated the identification of the key structural characteristics that influence H₂ adsorption in doped CNTs [14].

ANNs have been employed to enhance geometric calculations and predict properties related to adsorption processes in applications of CNTs and associated materials. Vences-Reynoso et al. [15] showed that MLP, BiLSTM, and one-dimensional CNN architectures can pre-optimize CNT geometries, cutting DFT computational overhead by 39% to 90%. Also, research into porous carbon materials has produced interpretable random forest models utilizing SHAP to forecast H₂ adsorption at 77 K, extending beyond pure nanotubes. These studies show that pressure and BET surface area are the most important factors, while micropore volume and oxygen content are less important [16]. In addition, Thanh et al. [10] used random forest models that were improved using evolutionary algorithms (PSO, GWO, GA, and DA) to predict H₂ uptake across a range of carbon-based adsorbents. The coefficients of determination (R^2) for the training and test sets were 0.98 and 0.91, respectively.

The potential of ML to model structure–function relationships extends to Metal–Organic Frameworks (MOFs). For example, Altintas et al. [17] demonstrated how simulations combined with ML can unravel structure–property relationships in MOFs for gas storage and separation applications. Similarly, Kuriakose et al. [11] introduced AABBA vector representations to train ANNs on 11,500 simulated structures, predicting gravimetric and volumetric H₂ storage capacities with descriptor selection assisted by gradient boosting techniques.

In recent years, significant progress has demonstrated the value of ML in hydrogen storage research. Nevertheless, a thorough analysis of the literature uncovers two principal methodological biases. First, there is a tendency to forecast global adsorption capacities, including gravimetric uptake or system-averaged energies, without explicitly modeling the local dependence of interaction energy on specific atomic configurations [10]. Second, many studies depend on overarching structural descriptors from extensive datasets, frequently lacking a direct connection to atomic coordinates or the exact geometric configurations of the system [11].

Few studies have attempted to predict interaction energies at the molecular level for H₂–CNT systems. Feedforward networks and random forest models have been utilized to interpolate experimental or simulated adsorption data [16], yet these methodologies typically do not fully represent the explicit dependence of interaction energy on atomic coordinates. For example, the research conducted by Kronberg et al. [14] represents an advancement by forecasting energies calculated by DFT and analyzing them using SHAP; nonetheless, it is limited to N-doped CNTs and specific electronic descriptors. On the other hand, the work of Faraji and Liu [13] on neural networks for C–H systems shows that it is possible to reproduce formation energies and geometries across different systems when trained on DFT data. This suggests that a well-calibrated MLP architecture could be used to predict interaction energies in other molecular systems.

In this context, we introduce an MLP-based approach to forecast interaction energy in the H₂–CNT(2,1) system. This study differs from previous research because it focuses on a small-diameter chiral (2,1) CNT, which enables us to describe electronic and geometric interactions with H₂ at higher atomic-level resolution. The MLP is trained to approximate the interaction-energy profile from atomic coordinates, interatomic distances, and relevant angles. To accomplish this aim, a reference dataset of configurations calculated using DFT

with DMol³ (Materials Studio) was created for training and testing the model. Consequently, this study aims to map the energetic landscape of the H₂–CNT interaction and evaluate the model's generalization capability across different interaction regions, providing a predictive method for adsorption studies in chiral nanotubes.

3. Theoretical Foundations

The development of a robust model for predicting interaction energy requires an understanding of both the physical substrate governing the interaction and the mathematical architecture used for function approximation. This section therefore establishes the theoretical framework of the study by outlining the structural and electronic properties of chiral carbon nanotubes that modulate hydrogen adsorption, as well as the operational principles of multilayer perceptrons as universal approximators.

3.1. Carbon Nanotubes: Structure and Adsorption Mechanisms

CNTs are nanomaterials formed by the seamless rolling of graphene sheets into cylindrical geometries with diameters typically ranging from sub-nanometer to several tens of nanometers [18]. The graphene lattice consists of carbon (C) atoms in *sp*² hybridization, arranged in a hexagonal honeycomb structure where each atom forms three σ -bonds with its neighbors (C–C), leaving a delocalized π -electron cloud perpendicular to the surface [19]. This electronic structure confers remarkable mechanical strength, chemical inertness, and a polarizable surface that facilitates non-covalent interactions with adsorbates.

The topology of a CNT is defined by the chiral vector $\mathbf{C}_h$ in Equation (1), which specifies the rolling direction of the graphene sheet.

$$\mathbf{C}_h = n \mathbf{a}_1 + m \mathbf{a}_2, \quad (1)$$

where (n, m) are integer chiral indices and $\{\mathbf{a}_1, \mathbf{a}_2\}$ denote the primitive lattice vectors of graphene.

From $\mathbf{C}_h$, two structural parameters are derived—the tube diameter $d_t = |\mathbf{C}_h|/\pi$ and the chiral angle $\theta \in [0^\circ, 30^\circ]$ —which determine the electronic band structure and surface curvature [20]. Based on the (n, m) indices, CNTs are classified into three families: *armchair* ($n = m$), *zigzag* ($m = 0$), and *chiral* ($n \neq m \neq 0$), where each family exhibits distinct electronic and adsorption properties.

For hydrogen storage applications, the interaction between H₂ and CNTs is typically governed by physisorption, where binding energies arise mainly from van der Waals (dispersion) forces, with a smaller contribution from electrostatic interactions with the π -electron system [21]. Reported adsorption energies for H₂ on pristine carbon surfaces generally fall within the range of -0.5 to -1.5 kcal/mol, reflecting the weak and reversible nature of physisorption [2]. Surface curvature can modulate these interactions: increased curvature in small-diameter tubes may enhance local polarizability and electric field gradients, potentially strengthening adsorbate-substrate coupling [3]. Conversely, strong confinement, as in the (2, 1) CNT considered here (with an estimated diameter of ≈ 2.1 Å), can lead to steric repulsion when the adsorbate kinetic diameter (2.89 Å for H₂) exceeds the available cavity space. In such cases, positive (endothermic) interaction energies are expected, reflecting penetration barriers rather than stable binding.

In this work, the chiral (2, 1) CNT is used to probe the limits of geometric confinement in adsorption phenomena. Its small diameter accentuates curvature effects, while its low symmetry introduces nonequivalent adsorption sites (e.g., internal cavity, rim regions, and external hexagonal rings), which represent a nontrivial challenge for modeling approaches. Thus, capturing the energetic landscape requires a method capable of resolving atomistic-level dependencies, for which MLP models provide a suitable framework.

3.2. Multilayer Perceptrons

An MLP is a canonical feedforward neural network architecture composed of successive layers of fully connected neurons, each applying a nonlinear activation function to an affine transformation of its inputs [22].

MLPs with at least one hidden layer and non-polynomial activation functions satisfy the universal approximation theorem, guaranteeing the capacity to approximate any continuous function on a compact domain to arbitrary precision [23,24]. This property, combined with efficient gradient-based optimization, has established MLPs as versatile tools for regression tasks in computational materials science, particularly in the modeling of physicochemical properties [8].

Let $\mathbf{x} \in \mathbb{R}^d$ denote a feature vector (input vector) encoding the geometric configuration of a H₂–CNT system. In a hidden layer, the pre-activation of neuron j is computed as in Equation (2):

$$s_j = \sum_{i=1}^n w_{ji} x_i + b_j, \quad (2)$$

For a network with L hidden layers, the forward pass is defined recursively as depicted in Equation (3):

$$\mathbf{h}^{(\ell)} = \phi(\mathbf{W}^{(\ell)} \mathbf{h}^{(\ell-1)} + \mathbf{b}^{(\ell)}), \quad \ell = 1, \dots, L, \quad \mathbf{h}^{(0)} = \mathbf{x}, \quad (3)$$

where $\mathbf{W}^{(\ell)}$ and $\mathbf{b}^{(\ell)}$ are the weight matrix and bias vector of layer ℓ , and $\phi(\cdot)$ is a nonlinear activation function, typically Rectified Linear Unit (ReLU), Gaussian Error Linear Unit (GELU), or Swish (an extension of Sigmoid Linear Unit, SiLU) [25,26].

Modern implementations often use GELU [25], defined as Equation (4):

$$\phi_{\text{GELU}}(z) = z \cdot \Phi(z), \quad (4)$$

$$\Phi(z) = \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{z}{\sqrt{2}}\right) \right], \quad (5)$$

The GELU activation provides smooth, non-monotonic transitions that improve gradient flow and numerical stability during optimization compared to traditional ReLU activations [26]. The final output $\hat{y} \in \mathbb{R}$, representing the predicted interaction-energy quantity E_{ads} , is obtained via a linear projection of the last hidden layer.

Model parameters $\Theta = \{\mathbf{W}^{(\ell)}, \mathbf{b}^{(\ell)}\}_{\ell=1}^L$ are optimized by minimizing a differentiable loss function $\mathcal{L}(\hat{y}, y)$ over a training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$. Additionally, the mean squared error (MSE) is commonly employed, as specified in Equation (6):

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2. \quad (6)$$

Optimization proceeds via stochastic gradient descent (SGD) variants such as Adam [27], which adaptively scale learning rates per parameter to accelerate convergence.

Generalization performance is the ability of an MLP to predict on unseen configurations. It depends on three factors: (i) model capacity, controlled by network depth and width, which must be balanced against the risk of overfitting given finite data; (ii) input representation, where the choice of descriptors (e.g., raw Cartesian coordinates versus rotationally invariant features) determines whether the model must learn physical symmetries from data or can exploit them by design; and (iii) regularization, including techniques such as L_2 weight decay, dropout, and early stopping, which constrain the hypothesis space to favor solutions that generalize beyond the training distribution [28].

For interaction-energy prediction, the MLP offers distinct advantages: it can capture complex, nonlinear dependencies between atomic positions and interaction energies without requiring explicit physical equations; it scales efficiently to high-dimensional input spaces; and, once trained, it can evaluate new configurations rapidly, enabling rapid screening of configurational spaces that would be prohibitive with *ab initio* methods.

However, its success depends on the quality and representativeness of the training data, as well as validation processes to ensure that learned patterns reflect genuine physical trends rather than artifacts of sampling or numerical noise.

4. Experimental Design

This study follows a three-stage workflow to generate a high-fidelity dataset of hydrogen interaction energies on a chiral carbon nanotube and to train a multilayer perceptron model capable of accurate prediction from geometric descriptors.

4.1. Phase I: Ab Initio Simulations

A (2,1) chiral carbon nanotube comprising 28 carbon atoms was selected to balance computational tractability with sufficient structural complexity for probing confinement effects.

The CNT(2,1) model was selected because its small size allows the generation of multiple DFT reference configurations at a manageable computational cost. This choice enables a controlled methodological evaluation of the MLP surrogate model. However, the results should be interpreted as specific to this finite chiral nanotube model and should not be generalized directly to larger-diameter CNTs or experimentally optimized hydrogen storage systems.

The finite CNT(2,1) model was considered in its pristine, non-functionalized form in order to keep the simulation protocol controlled and computationally tractable. The configurational sampling was designed to explore the molecular H₂–CNT interaction-energy landscape, mainly associated with physisorption-like and repulsive regimes, rather than dissociative chemisorption processes.

In this sense, H₂ was treated as a molecular adsorbate throughout the simulations, and no explicit H–H bond dissociation or C–H bond formation pathway was imposed. Although molecular hydrogen is a closed-shell species and H–H bond cleavage is not expected to occur spontaneously under the present sampling protocol, the non-passivated terminal carbon atoms may still affect the local electronic structure and the computed interaction energies near the tube edges.

Therefore, edge-region results should be interpreted within the limits of the finite pristine and unpassivated CNT model used in this proof-of-concept study.

All electronic structure calculations were performed using the DMOL³ module in BIOVIA Materials Studio[®], employing the GGA-PBE exchange-correlation functional and a double numerical basis set with polarization functions. To ensure self-consistent field (SCF) convergence, electronic smearing of 0.005 Ha was applied, which smooths discontinuities near the Fermi level.

The isolated CNT geometry was first optimized, and its total energy, $E_{\text{tot}}(\text{CNT})$, was subsequently calculated using the same DFT settings. Similarly, the isolated H₂ molecule was independently optimized, and its total energy, $E_{\text{tot}}(\text{H}_2)$, was then computed. After obtaining these reference energies, each combined H₂–CNT configuration was evaluated through a single-point DFT calculation to obtain $E_{\text{tot}}(\text{CNT} + \text{H}_2)$. Finally, these three energy terms were substituted into the supermolecular expression used to compute the generalized interaction-energy quantity E_{ads} .

Then, the interaction-energy quantity E_{ads} was computed via the standard supermolecular approach in Equation (7):

$$E_{\text{ads}} = E_{\text{tot}}(\text{CNT} + \text{H}_2) - E_{\text{tot}}(\text{CNT}) - E_{\text{tot}}(\text{H}_2), \quad (7)$$

where E_{ads} is used in this work as a generalized interaction-energy quantity derived from total-energy differences. Under this definition, negative or lower-energy values describe attractive interaction regimes associated with physisorption, whereas positive values correspond to repulsive or endothermic configurations. Therefore, the term E_{ads} does not exclusively denote stable adsorption states, but rather the complete H_2 –CNT interaction-energy landscape sampled in this study.

To make the calculation procedure explicit, Figure 1 shows the representative external CNT surface configuration used for the energy calculation example, while Table 1 presents the corresponding total-energy components for a fixed molecular orientation of $\alpha = 135^\circ$, where the centroid–centroid distance was varied in 0.5 Å increments. The total-energy components originally obtained from the single-point DFT evaluations were converted to eV for consistency with the reported target variable. The resulting positive E_{ads} values decrease as the H_2 –CNT separation increases, illustrating the transition from a strongly repulsive close-contact configuration toward a weak-interaction regime.




Figure 1. Representative external CNT surface configuration used to illustrate the calculation of E_{ads} . The sampled interaction site on the CNT surface for an H_2 molecule with molecular orientation $\alpha = 135^\circ$, where the centroid–centroid distance was varied to obtain the total-energy values reported in Table 1.

Table 1. Representative calculation of E_{ads} for an external CNT surface configuration at $\alpha = 135^\circ$. All energy terms are reported in eV after conversion from the original total-energy output.

α ($^\circ$)	r (Å)	$E_{\text{tot}} \text{ H}_2\text{–CNT}$	$E_{\text{tot}} \text{ CNT}$	$E_{\text{tot}} \text{ H}_2$	E_{ads}	Interpretation
135	0.0	−157.9818	−164.8210	−4.4690	11.3175	Strong repulsion
135	0.5	−159.8599	−164.8210	−4.4690	9.4379	Strong repulsion
135	1.0	−164.8985	−164.8210	−4.4690	4.3951	Repulsion
135	1.5	−167.2893	−164.8210	−4.4690	2.0023	Moderate repulsion
135	2.0	−168.5197	−164.8210	−4.4690	0.7709	Weak repulsion
135	2.5	−168.9849	−164.8210	−4.4690	0.3054	Weak interaction
135	3.0	−169.1517	−164.8210	−4.4690	0.1385	Weak interaction

To sample the interaction landscape, the relative configuration space was partitioned into three distinct regions of the CNT: internal cavity, tube edges, and external surface. This procedure made it possible to capture the energetic variability associated with the different interaction sites.

For the internal cavity, the H_2 centroid was aligned with the CNT axis, as illustrated in Figure 2. For orientations $\alpha \in \{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$ in the transverse plane, the center-to-center distance r was varied from 0 to 7.0 Å in 0.5 Å increments. The system energy was calculated for each (r, α) combination.




Figure 2. H_2 –CNT interaction within the internal cavity of the nanotube, illustrating the relative configuration employed in the sampling procedure.

For the case of tube edges, H_2 was positioned relative to local centroids defined at each terminal carbon atom, using the same (r, α) grid as above to probe rim-specific interactions. The open ends of the nanotube were considered as active interaction zones between the H_2 molecule and the terminal carbon atoms of the CNT, as shown in Figure 3.




Figure 3. H_2 –CNT interaction at the nanotube edges, showing the sampling configurations employed to compute interaction energies.

Finally, for the external surface zone represented in Figure 4, the interaction sites were aligned to the geometric centers of individual hexagonal rings along the tube wall, again sampled over the full (r, α) parameter space.




Figure 4. H_2 –CNT interaction on the external surface of the nanotube, exhibiting the sampling sites located on the hexagonal rings.

Following this procedure, 696 curated configurations were obtained, where total energies were computed and Equation (7) evaluated to obtain E_{ads} for each case.

4.2. Phase II: Dataset

Each one of the $N = 696$ configurations was encoded as a 31×3 matrix of Cartesian coordinates (x, y, z) of the atoms, as follows:

- Rows 1–28: CNT carbon atoms.
- Rows 29–30: H₂ hydrogen atoms.
- Row 31: Metadata, where the first column is α (degrees), the second column is r (Å), and the third column is E_{ads} (eV).

This yields 93 scalar quantities per sample: 92 input variables and one target value. The scalar output corresponds to a supervised regression problem based on a vectorial representation of the configuration $(\mathbf{r}_{\text{CNT}}, \mathbf{r}_{\text{H}_2}, \alpha, r)$, where the dataset is organized into blocks of 31 rows per sample. For each sample i , the input vector $\mathbf{x}$ is defined as Equation (8):

$$\mathbf{x}_i = [\mathbf{r}_{\text{CNT}}^{(i)} \in \mathbb{R}^{84}; \mathbf{r}_{\text{H}_2}^{(i)} \in \mathbb{R}^6; \alpha^{(i)}; r^{(i)}] \in \mathbb{R}^{92}, \quad (8)$$

with the target value $y_i = E_{\text{ads}}^{(i)} \in \mathbb{R}$.

The input variables used by the MLP were organized according to their physical meaning, as summarized in Table 2. The Cartesian coordinates of the CNT atoms encode the local geometry of the finite nanotube, while the Cartesian coordinates of the H₂ molecule define its position and orientation relative to the carbon framework. The centroid–centroid distance and molecular orientation angle correspond to the explicit geometric parameters varied during the sampling protocol.

Table 2. Input variables and target variable used in the MLP-based regression workflow.

Variable Group	Number of Variables	Physical Interpretation
CNT Cartesian coordinates	84	Local atomic geometry of the finite CNT model
H ₂ Cartesian coordinates	6	Molecular position and orientation
Orientation angle, α	1	Relative molecular orientation during sampling
Centroid–centroid distance, r	1	Main separation descriptor controlling interaction strength
Target variable, E_{ads}	1	Generalized H ₂ –CNT interaction-energy quantity

The input variables do not explicitly encode the electronic density, molecular orbitals, exchange–correlation terms, or individual quantum-mechanical energy contributions computed during the DFT procedure. These physical and chemical effects are instead implicitly contained in the reference target values generated with DMol³ using the GGA–PBE functional. Therefore, by training on atomic coordinates, centroid–centroid distance, molecular orientation angle, and the corresponding DFT-computed E_{ads} values, the MLP learns a supervised numerical relationship between the sampled H₂–CNT geometry and the resulting interaction-energy response. Thus, the model should be interpreted as an approximation to the DFT-derived interaction-energy trends within the sampled configurational domain, rather than as an explicit quantum-mechanical solver or a transferable interatomic potential.

From a physical perspective, the centroid–centroid distance r is expected to strongly influence E_{ads} , because it controls the transition between short-range repulsive interactions and longer-range attractive regimes. The molecular orientation angle α modifies the relative exposure of the H₂ molecule to the CNT wall, edge, or internal cavity. In addition, the Cartesian coordinates of the CNT and H₂ atoms define the local geometric environment and allow the model to distinguish between internal-cavity, edge, and external-surface configurations. Although a formal feature-importance analysis such as SHAP or permutation importance was not performed in this work, this geometric interpretation supports the physical relevance of the descriptors used as input variables.

For reproducibility, the original target-energy distribution was characterized before model training. The final curated dataset contained 696 H₂–CNT configurations. The generalized interaction-energy variable, denoted as E_{ads} , exhibited a mean value of 4.0955 eV and a sample standard deviation of 4.3189 eV. The minimum and maximum values were

0.1161 eV and 12.9977 eV, respectively. The complete dataset was divided into 487 training configurations and 209 internal test configurations.

Raw Cartesian coordinates were used as direct input features because the objective was to evaluate a simple atom-coordinate-based MLP surrogate under a controlled configurational sampling protocol. This representation preserves explicit atomic-position information for each H₂-CNT configuration. However, raw Cartesian coordinates are not intrinsically invariant to global translations, rotations, or permutations of equivalent atoms. Therefore, the trained MLP should be interpreted as a dataset-specific regression model rather than as a transferable interatomic potential.

Afterwards, the full dataset was standardized using Gaussian normalization to homogenize feature scales, as expressed in Equation (9):

$$\tilde{\mathbf{x}} = \frac{\mathbf{x} - \mu_X}{\sigma_X}, \quad \tilde{y} = \frac{y - \mu_y}{\sigma_y}, \quad (9)$$

with statistics $(\mu_X, \sigma_X, \mu_y, \sigma_y)$ computed over the entire set prior to any train/validation split, and all numerical operations employed double-precision floating-point arithmetic.

4.3. Phase III: MLP Architecture and Training Protocol

The predictive task was framed as scalar regression: given the vectorized atomic coordinates and metadata $\mathbf{x}_i \in \mathbb{R}^{92}$, predict $y_i = E_{\text{ads}}^{(i)}$. A compact feedforward architecture (64–32–1) was adopted, featuring GELU activations and regularization to mitigate overfitting from spatial correlations in the input. Specifically, each dense layer included L_2 weight decay ($\lambda = 10^{-4}$) and a dropout of $p = 0.15$. Table 3 summarizes the MLP architecture.

Table 3. Topology of the multilayer perceptron utilized in this study.

Layer	Dimension/Parameters	Details
Input	$\tilde{\mathbf{x}}$	Input vector (92 features)
Dense 1	64 neurons	GELU activation, L2 ($\lambda = 10^{-4}$)
Dropout 1	$p = 0.15$	Regularization
Dense 2	32 neurons	GELU activation, L2 ($\lambda = 10^{-4}$)
Dropout 2	$p = 0.15$	Regularization
Output	$\tilde{y}$ (1 neuron)	Prediction of E_{ads}

The model was optimized using the Adam algorithm ($\eta = 5 \times 10^{-4}$) with mean squared error loss on the normalized target $\tilde{y}$. Training used mini-batches of size 16 for up to 2000 epochs, with early stopping (patience = 50) to prevent overtraining and to restore the weights corresponding to the best validation performance. A randomly held-out validation subset (10% of samples) was used. Reproducibility was ensured by fixing global random seeds (NumPy/TensorFlow, SEED = 42).

No exhaustive hyperparameter optimization or formal ablation study was performed in this work. The 64–32–1 architecture was selected as a compact baseline model after preliminary tests aimed at balancing predictive performance, training stability, and architectural simplicity were completed. The use of GELU activation functions, L2 regularization, dropout, and early stopping was intended to reduce overfitting while preserving a simple feedforward structure suitable for evaluating the feasibility of the atom-coordinate-based regression approach. Future work should include systematic ablation studies to quantify the individual contribution of the input descriptors, regularization strategies, activation functions, and network depth to the final predictive performance.

Performance was quantified via the coefficient of determination, depicted in Equation (10), with an acceptance threshold of $R^2 \geq 0.85$ to validate physical fidelity and generalization capability that reproduces the energetic variability of the H₂–CNT (2,1) interaction governing the adsorption process.

$$R^2 = 1 - \frac{\sum_{i=1}^N (\hat{y}_i - y_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}, \quad (10)$$

The workflow described above establishes a reproducible pipeline for generating interaction-energy data and training an MLP model. This methodology ensures that learned patterns reflect physical dependencies by sampling configurational space across distinct interaction regions.

5. Results

The following sections present the performance of the trained model, evaluating its ability to generalize to unseen configurations and to reproduce the energetic trends governing the H₂–CNT (2,1) interaction across the internal cavity, edges, and external surface.

5.1. Predictive Performance of the MLP Model

The predictive capability of the trained multilayer perceptron was evaluated using a held-out test set comprising previously unseen configurations sampled from the internal cavity, edges, and external surface of the CNT. The dataset consisted of $N = 696$ curated configurations, partitioned into training and testing subsets.

The internal split contained 487 configurations for training and 209 configurations for testing. On this internal test set, the trained MLP achieved $R^2 = 0.8690$ and $\text{MSE} = 2.6398 \text{ eV}^2$, confirming that the model reproduced the main DFT-derived interaction-energy trends on unseen samples from the same configurational distribution.

Model performance was quantified using the coefficient of determination (R^2) to assess trend fidelity and the MSE to evaluate absolute accuracy in squared energy units (eV^2).

The regional performance plots were generated after completing the MLP training process. The trained model was saved and subsequently reloaded together with the corresponding input and output scalers. Independent test files extracted from the original configurational database were then used to evaluate the model after training. These test files contained only the input information required for inference, namely the Cartesian coordinates (x, y, z) of the CNT atoms, the Cartesian coordinates of the H₂ molecule, the centroid–centroid distance, and the molecular orientation angle. The trained MLP then predicted the corresponding generalized interaction-energy quantity, E_{ads} , which was compared against the DFT reference values computed with the DMol³ module. Therefore, the curves shown in this section correspond to post-training inference tests with the saved model and not to the training-loss or fitting process itself.

Table 4 summarizes the performance metrics across the three interaction regions. The model demonstrates robust generalization capability, achieving R^2 values exceeding 0.90 in all regions. However, a distinct trade-off is observed between correlation strength and absolute error magnitude: while the model captures the energetic ordering of configurations (high R^2), the MSE varies significantly depending on the energy scale of the specific region.

Table 4. Performance of the MLP model for each interaction region.

Region	Test Set Size	R^2	MSE (eV ²)
Internal cavity of the CNT	15	0.9051	1.505950
CNT edges (rim regions)	4	0.9358	0.046259
CNT surface (ring sites)	7	0.9625	0.574098

5.2. Analysis of Interaction Energy by Region

Figure 5 illustrates the performance of the model within the internal cavity of the nanotube. This region exhibits the highest energy magnitudes, ranging from approximately 1 to 12 eV, indicative of strong steric repulsion due to the confinement of the H_2 molecule within the narrow tube diameter.




Figure 5. Post-training prediction test for the internal cavity of the CNT. The trained and reloaded MLP model predicts the generalized interaction-energy quantity E_{ads} from atomic coordinates, interaction distance, and molecular orientation angle; results are compared against DFT reference values.

The model achieves an R^2 of 0.9051, confirming that it can learn the repulsive trend as the molecule approaches the tube walls. However, the MSE is notably higher in this region (1.506 eV²) compared to the external surfaces. This elevated error suggests that while the network captures the qualitative physics of confinement, it exhibits slight deviations in predicting the steepness of the repulsive wall. For high-throughput screening purposes, this level of accuracy is sufficient to identify forbidden configurations, though absolute energy values in highly repulsive regimes may require further calibration.

In the edge-region test, the tube rims display the highest absolute accuracy, with the lowest MSE of 0.046 eV² and an R^2 of 0.9358 (Figure 6). The interaction energies in this region are confined to a lower range (approximately 0 to 2.5 eV), corresponding to interactions with terminal carbon atoms. The agreement between predicted and reference values indicates that the local geometric environment at the edges is well represented by the input descriptors. This suggests that edge sites, often critical for catalytic activity or functionalization, can be modeled with high precision using this MLP architecture.

Finally, the external-surface test yielded the highest correlation ($R^2 = 0.9625$), as shown in Figure 7. The energy profile spans from near 0 to 10 eV, capturing the transition from repulsive close contact to favorable physisorption distances. Although the MSE (0.574 eV²) is higher than that of the edges, it is proportionally lower relative to the total energy range compared to the internal cavity. The model effectively preserves the ranking

of interaction sites across the hexagonal rings, validating its utility for mapping energetic landscapes on the outer tube surface where practical hydrogen storage interactions occur.




Figure 6. Post-training prediction test for the CNT edge region. The comparison between DFT reference values and MLP-predicted E_{ads} demonstrates the model response after loading the saved trained network and evaluating independent test configurations.




Figure 7. Post-training prediction test for the external CNT surface. The trained MLP model predicts the generalized interaction-energy quantity E_{ads} for independent surface configurations using only structural input variables.

5.3. Summary of Findings

The results show that the multilayer perceptron is effective for approximating *ab initio* energy trends in this system. Overall, the model reproduces the intrinsic physical behavior of H_2 –CNT interactions, achieving coefficients of determination of at least 0.90 across all regions, which supports its use for configurational screening. However, the prediction error increases with the magnitude of the interaction energy. In particular, larger errors are observed within the internal cavity, where steep repulsive potentials and confined geometries are more difficult to model. In contrast, lower errors occur near the edges, where the energy landscape is smoother. In addition, the model provides substantial

computational advantages by enabling a rapid evaluation of interaction energies while reducing reliance on repeated density functional theory calculations within the sampled configurational domain.

6. Discussion

The predictive performance of the proposed MLP model, characterized by R^2 exceeding 0.90 across all interaction regions, aligns with recent advances in ML applied to materials discovery. While studies such as those by Thanh et al. [10] and Kronberg et al. [14] have successfully employed ensemble methods and explainability techniques to predict global uptake or averaged energies, the present work distinguishes itself by resolving the local energetic landscape at the atomic level.

The ability of the MLP to reproduce *ab initio* trends validates the utility of feedforward architectures as surrogate models for nanostructured systems, corroborating the findings of Vences-Reynos et al. [15] regarding the computational efficiency of MLPs in CNT-related tasks. However, beyond mere predictive accuracy, the physical interpretation of the energy magnitudes and the composition of the training dataset warrant critical discussion in the context of existing theoretical literature.

A notable feature of the results is the range of interaction energies observed, particularly within the internal cavity of the (2,1) CNT, where values reach up to approximately 12 eV. While typical physisorption energies for H_2 on carbon surfaces are often cited in the range of 0.1–0.5 eV, the higher magnitudes obtained in this study reflect the specific geometric constraints of the system rather than computational artifacts.

Literature on confined molecular hydrogen indicates that in small diameter tubes, the interaction energy becomes highly sensitive to radius and orientation, often resulting in strongly repulsive interactions when the molecule is squeezed against the wall [2]. For very narrow tubes, such as the (2,1) chirality studied here, the kinetic diameter of H_2 exceeds the available cavity space, leading to a dominance of Pauli repulsion over dispersion forces [29,30]. Consequently, the positive energy values observed in the internal cavity represent the repulsive wall of the potential-energy surface, a physical reality confirmed by DFT studies on confined H_2 which show that binding can become unfavorable in the smallest radii [31,32]. Thus, the model is not predicting erroneous binding energies but rather mapping the steric barriers that define the accessibility of the nanotube interior.

The inclusion of these configurations in the training dataset is a methodological choice supported by recent work on interatomic potentials. Montes De Oca Zapiain et al. [33] argue that ML models achieve greater transferability and accuracy when trained on diverse datasets that encompass stable equilibrium states and high-energy, non-equilibrium configurations. For distances from 0 to 7 Å, including the short range repulsive regime, the MLP learns the full shape of the interaction curve rather than merely interpolating around the energy minimum.

This approach aligns with the findings of Costanzo et al. [34], who demonstrated that a correct description of H_2 adsorption on graphene and nanotubes requires analyzing the energy as a continuous function of distance, capturing both the attractive physisorption well and the short range repulsion caused by electronic cloud overlap. Similarly, Lugo et al. [35] emphasized that interaction barriers and repulsive regions are integral to the PES, particularly in defective or confined systems, and ignoring them limits the model's physical validity.

Furthermore, the selection of configurations where H_2 is positioned inside the tube serves as a stress test for the model's generalization capability. Cheng et al. [32] noted that inside adsorption in narrow CNTs is more likely to expose steric limitations compared to outside adsorption. Thus, the MLP demonstrates robustness against configurational out-

liers when it can predict energies in the repulsive regime (as evidenced by the $R^2 = 0.9051$ for the internal cavity). This supports the argument that a well-calibrated model should identify favorable adsorption sites and correctly penalize physically inaccessible configurations with high energy values.

The ability to distinguish between favorable external sites and repulsive internal cavities is helpful for applications involving gas separation or storage, where understanding exclusion volumes is as important as identifying binding sites.

Another limitation of the present model is that the finite CNT was not terminally passivated. Dangling bonds at the tube ends may introduce local electronic effects that are not representative of ideal infinite nanotubes or hydrogen-passivated finite fragments. Moreover, this study was restricted to molecular H_2 interacting with a pristine CNT model; dissociative adsorption, H–H bond cleavage, C–H bond formation, and chemical functionalization of the nanotube edges were not explicitly considered. Future work should compare pristine, passivated, defective, and functionalized CNT models to quantify how terminal states and possible chemisorption pathways modify the H_2 –CNT interaction-energy landscape.

Another methodological limitation is the use of raw Cartesian coordinates as input variables. Although this representation is straightforward and preserves the explicit geometry of each configuration, it does not encode rotational, translational, or permutational invariance. Consequently, the present MLP should not be interpreted as a transferable interatomic potential, but as a supervised regression model valid within the coordinate convention and configurational protocol used to generate the training data. Future work should evaluate invariant geometric descriptors to reduce the dependence of the model on the selected coordinate frame and to improve transferability within broader configurational domains.

Limitations and Future Work

The present study has a specific methodological scope that should be considered when interpreting the results. First, the model was developed using a finite chiral CNT(2,1) system; therefore, the learned trends should not be directly generalized to larger-diameter CNTs, other chiralities, periodic nanotube models, or experimentally optimized hydrogen storage materials without additional training and validation. Second, the use of a small fixed chirality was intentional for a proof-of-concept study because it allowed for a systematic generation of DFT reference configurations at a manageable computational cost, but it also restricted the application domain of the trained MLP.

A second limitation is related to the finite pristine and unpassivated representation of the nanotube. As discussed above, non-passivated terminal carbon atoms may influence local interaction energies near the tube edges. Consequently, edge-region predictions should be interpreted as specific to the finite CNT model used in this work. Future studies should compare pristine, passivated, defective, and functionalized CNT models in order to evaluate how terminal states, defects, and chemical modifications affect the H_2 –CNT interaction-energy landscape.

A third limitation concerns the input representation and model-selection strategy. The MLP uses raw Cartesian coordinates, the centroid–centroid distance, and the molecular orientation angle as input variables. Although this representation is useful for evaluating the feasibility of an atom-coordinate-based surrogate model, it does not provide intrinsic invariance to global rotations, translations, or permutations of equivalent atoms. In addition, no exhaustive hyperparameter optimization or formal ablation study was performed. Future work should evaluate invariant geometric descriptors, additional CNT chiralities and

diameters, larger test sets, systematic ablation studies, and quantitative feature-importance analyses to improve transferability, interpretability, and robustness.

7. Conclusions

This study demonstrated that a compact multilayer perceptron can reproduce DFT-derived H₂–CNT interaction-energy trends from an atom-coordinate-based representation. The final dataset comprised 696 H₂–CNT configurations. Each configuration was encoded using 92 input variables, including 84 Cartesian coordinates for the CNT atoms, 6 Cartesian coordinates for the H₂ molecule, the molecular orientation angle, and the centroid–centroid distance. The corresponding target variable, denoted as E_{ads} , was treated as a generalized interaction-energy quantity. Its original distribution exhibited a mean value of 4.0955 eV, a sample standard deviation of 4.3189 eV, a minimum value of 0.1161 eV, and a maximum value of 12.9977 eV.

The MLP architecture consisted of two hidden dense layers with 64 and 32 neurons, GELU activation functions, L2 regularization with $\lambda = 10^{-4}$, dropout with $p = 0.15$, and a single linear output neuron. The model was trained with the Adam optimizer using a learning rate of 5×10^{-4} , mini-batches of 16 samples, mean squared error loss, a 10% validation split from the training set, and early stopping. From the 696 curated configurations, 487 were used for training and 209 were reserved for internal testing. On the internal test set, the model achieved $R^2 = 0.8690$ and $\text{MSE} = 2.6398 \text{ eV}^2$.

Region-specific external evaluations further confirmed the ability of the MLP to reproduce DFT-computed trends across different H₂–CNT interaction environments. For the internal cavity, the model achieved $R^2 = 0.9051$ and $\text{MSE} = 1.5060 \text{ eV}^2$ using 15 test configurations. For the CNT edge region, the model obtained $R^2 = 0.9358$ and $\text{MSE} = 0.0463 \text{ eV}^2$ using 4 test configurations. For the external surface, the best trend agreement was observed, with $R^2 = 0.9625$ and $\text{MSE} = 0.5741 \text{ eV}^2$ using 7 test configurations. These results show that the model captures the main interaction-energy profiles in the three evaluated regions, while larger errors in the internal cavity are consistent with the steeper repulsive wall associated with geometric confinement.

The results support the use of MLP-based surrogate models as computationally efficient tools for the rapid evaluation of H₂–CNT configurations within the sampled domain. Nevertheless, the present model should be interpreted as a dataset-specific regression model rather than as a transferable interatomic potential. Its scope is limited by the finite pristine and unpassivated CNT(2,1) representation and the direct use of raw Cartesian coordinates. Future work should extend the approach to larger-diameter CNTs, additional chiralities, passivated and functionalized nanotube models, larger external test sets, systematic ablation studies, quantitative feature-importance analyses, and invariant geometric descriptors capable of improving transferability, interpretability, and physical consistency.

Author Contributions: Conceptualization, L.J.V.R. and R.A.E.; methodology, F.D.R.L. and J.H.P.S.; software, L.J.V.R. and F.D.R.L.; validation, E.E.G.G. and R.A.E.; formal analysis, R.A.E., D.V.V. and A.A.F.F.; investigation, L.J.V.R., J.H.P.S. and A.A.F.F.; data curation, L.J.V.R. and D.V.V.; visualization, L.J.V.R. and E.E.G.G.; writing—original draft preparation, L.J.V.R. and F.D.R.L.; writing—review and editing, E.E.G.G.; supervision, J.H.P.S. and R.A.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the Tecnológico Nacional de México, under Project No. 25653.26-P, and by the Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI), Mexico, through Scholarship Grant No. 845048.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: The authors thank the Instituto Tecnológico de Toluca, Tecnológico Nacional de México, for their support.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Le, T.H.; Kim, M.P.; Park, C.H.; Tran, Q.N. Recent Developments in Materials for Physical Hydrogen Storage: A Review. *Materials* **2024**, *17*, 666. <https://doi.org/10.3390/ma17030666>.
2. Lyu, J.; Kudiiarov, V.; Lider, A. An Overview of the Recent Progress in Modifications of Carbon Nanotubes for Hydrogen Adsorption. *Nanomaterials* **2020**, *10*, 255. <https://doi.org/10.3390/nano10020255>.
3. Baghai, B.; Ketabi, S. Hydrogen storage efficiency of Fe doped carbon nanotubes: Molecular simulation study. *RSC Adv.* **2024**, *14*, 9763–9780. <https://doi.org/10.1039/D3RA08382A>.
4. Araujo, R.B.; Rodrigues, G.L.S.; dos Santos, E.C.; Pettersson, L.G.M. Adsorption energies on transition metal surfaces: Towards an accurate and balanced description. *Nat. Commun.* **2022**, *13*, 6853. <https://doi.org/10.1038/s41467-022-34507-y>.
5. Pablo Garcia, S.; Morandi, S.; Vargas Hernandez, R.A.; Jorner, K.; Ivković, Z.; López, N.; Aspuru-Guzik, A. Fast evaluation of the adsorption energy of organic molecules on metals via graph neural networks. *Nat. Comput. Sci.* **2023**, *3*, 433–442. <https://doi.org/10.1038/s43588-023-00437-y>.
6. Orban, B.; Holtzl, T. Acetylene and Ethylene Adsorption during Floating Fe Catalyst Formation at the Onset of Carbon Nanotube Growth and the Effect of Sulfur Poisoning: A DFT Study. *Inorg. Chem.* **2024**, *63*, 13624–13635. <https://doi.org/10.1021/acs.inorgchem.4c01830>.
7. Bahmanzadgan, F.; Ghaemi, A.; Qasemnazhand, M.; Molaei, M. Simulation of gas adsorption on single-walled carbon nanotubes. *Sci. Rep.* **2025**, *15*, 15595. <https://doi.org/10.1038/s41598-025-99988-5>.
8. Zhang, Q.; Zeng, R.; Lu, Y.; Zhang, J.; Zhou, W.; Yu, J. Machine learning-based prediction of the adsorption energy of CO on boron-doped graphene. *New J. Chem.* **2022**, *46*, 10451–10457. <https://doi.org/10.1039/D2NJ01278B>.
9. Jiang, D.; Wu, Z.; Hsieh, C.Y.; Chen, G.; Liao, B.; Wang, Z.; Shen, C.; Cao, D.; Wu, J.; Hou, T. Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models. *J. Cheminform.* **2021**, *13*, 12. <https://doi.org/10.1186/s13321-020-00479-8>.
10. Thanh, H.V.; Ebrahimnia Taremsari, S.; Ranjbar, B.; Mashhadimoslem, H.; Rahimi, E.; Rahimi, M.; Elkamel, A. Hydrogen Storage on Porous Carbon Adsorbents: Rediscovery by Nature-Derived Algorithms in Random Forest Machine Learning Model. *Energies* **2023**, *16*, 2348. <https://doi.org/10.3390/en16052348>.
11. Kuriakose, N.; Morán-González, L.; Kneiding, H.; Betten, J.E.; Balcells, D. Machine Learning the Hydrogen Adsorption Capacity of Metal Organic Frameworks. *ChemRxiv* **2025**. <https://doi.org/10.26434/chemrxiv-2025-mccz2>.
12. Chen, G.; Tang, D.M. Machine Learning as a “Catalyst” for Advancements in Carbon Nanotube Research. *Nanomaterials* **2024**, *14*, 1688. <https://doi.org/doi:10.3390/nano14211688>.
13. Faraji, S.; Liu, M. Transferable machine learning interatomic potential for carbon hydrogen systems. *Phys. Chem. Chem. Phys.* **2024**, *26*, 22346–22358. <https://doi.org/10.1039/D4CP02300E>.
14. Kronberg, R.; Lappalainen, H.; Laasonen, K. Hydrogen Adsorption on Defective Nitrogen-Doped Carbon Nanotubes Explained via Machine Learning Augmented DFT Calculations and Game-Theoretic Feature Attributions. *J. Phys. Chem. C* **2021**, *125*, 15918–15933. <https://doi.org/10.1021/acs.jpcc.1c03858>.
15. Vences-Reynos, L.J.; Villanueva-Vasquez, D.; Alejo-Eleuterio, R.; Del Razo-López, F.; Martínez-Gallegos, S.M.; Granda-Gutiérrez, E.E. Evaluation of Neural Networks for Improved Computational Cost in Carbon Nanotubes Geometric Optimization. *Modelling* **2025**, *6*, 36. <https://doi.org/10.3390/modelling6020036>.
16. Maulana Kusdhany, M.I.; Lyth, S.M. New insights into hydrogen uptake on porous carbon materials via explainable machine learning. *Carbon* **2021**, *179*, 190–201. <https://doi.org/10.1016/j.carbon.2021.04.036>.
17. Altintas, C.; Altundal, O.F.; Keskin, S.; Yildirim, R. Machine Learning Meets with Metal Organic Frameworks for Gas Storage and Separation. *J. Chem. Inf. Model.* **2021**, *61*, 2131–2146. <https://doi.org/10.1021/acs.jcim.1c00191>.
18. Iijima, S. Helical microtubules of graphitic carbon. *Nature* **1991**, *354*, 56–58. <https://doi.org/10.1038/354056a0>.
19. Geim, A.K.; Novoselov, K.S. The rise of graphene. *Nat. Mater.* **2007**, *6*, 183–191. <https://doi.org/10.1038/nmat1849>.
20. Mattia, D.; Gogotsi, Y. Review: Static and dynamic behavior of liquids inside carbon nanotubes. *Microfluid. Nanofluid.* **2008**, *5*, 289–305. <https://doi.org/10.1007/s10404-008-0293-5>.
21. Wang, S.; Zhang, Z.; Tan, Y.; Liang, K.; Zhang, S. Review on the characteristics of existing hydrogen energy storage technologies. *Energy Sources Part A Recovery Util. Environ. Eff.* **2023**, *45*, 985–1006. <https://doi.org/10.1080/15567036.2023.2175938>.
22. Shwartz-Ziv, R.; Armon, A. Tabular data: Deep learning is not all you need. *Inf. Fusion* **2022**, *81*, 84–90. <https://doi.org/10.1016/j.inffus.2021.11.011>.

23. Mehta, P.; Bukov, M.; Wang, C.H.; Day, A.G.; Richardson, C.; Fisher, C.K.; Schwab, D.J. A high-bias, low-variance introduction to Machine Learning for physicists. *Phys. Rep.* **2019**, *810*, 1–124. <https://doi.org/10.1016/j.physrep.2019.03.001>.
24. Carleo, G.; Cirac, I.; Cranmer, K.; Daudet, L.; Schuld, M.; Tishby, N.; Vogt-Maranto, L.; Zdeborová, L. Machine learning and the physical sciences. *Rev. Mod. Phys.* **2019**, *91*, 045002. <https://doi.org/10.1103/RevModPhys.91.045002>.
25. Hendrycks, D.; Gimpel, K. Gaussian Error Linear Units (GELUs). *arXiv* **2016**. arXiv:1606.08415. <https://doi.org/10.48550/ARXIV.1606.08415>.
26. Ramachandran, P.; Zoph, B.; Le, Q.V. Searching for Activation Functions. *arXiv* **2017**. arXiv:1710.05941. <https://doi.org/10.48550/arXiv.1710.05941>.
27. Loshchilov, I.; Hutter, F. Decoupled Weight Decay Regularization. In Proceedings of the 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, 6–9 May 2019. <https://doi.org/10.48550/arXiv.1711.05101>.
28. Wang, B.; Xia, H.; Nguyen, T.; Osher, S. How does momentum benefit deep neural networks architecture design? A few case studies. *Res. Math. Sci.* **2022**, *9*, 57. <https://doi.org/10.1007/s40687-022-00352-0>.
29. Petrushenko, I. DFT Calculations of Hydrogen Adsorption inside Single-Walled Carbon Nanotubes. *Adv. Mater. Sci. Eng.* **2018**, *2018*, 9876015. <https://doi.org/10.1155/2018/9876015>.
30. Rehak, F.; Piccini, G.; Alessio, M.; Sauer, J. Including Dispersion in Density Functional Theory for Adsorption on Flat Oxide Surfaces, in Metal–Organic Frameworks and in Acidic Zeolites. *Phys. Chem. Chem. Phys.* **2020**, *22*, 7577–7585. <https://doi.org/10.1039/D0CP00394H>.
31. Gtari, W.; Tangour, B. A Theoretical Study of the Dihydrogen Molecule Confined inside Carbon Nanotubes. *Int. J. Quantum Chem.* **2013**, *113*, 2397–2404. <https://doi.org/10.1002/qua.24474>.
32. Cheng, H.; Cooper, A.; Pez, G.; Kostov, M.; Piotrowski, P.; Stuart, S. Molecular Dynamics Simulations on the Effects of Diameter and Chirality on Hydrogen Adsorption in Single Walled Carbon Nanotubes. *J. Phys. Chem. B* **2005**, *109*, 3780–3786. <https://doi.org/10.1021/jp045358m>.
33. Montes De Oca Zapiain, D.; Wood, M.; Lubbers, N.; Pereyra, C.; Thompson, A.; Perez, D. Training Data Selection for Accuracy and Transferability of Interatomic Potentials. *npj Comput. Mater.* **2022**, *8*, 189. <https://doi.org/10.1038/s41524-022-00872-x>.
34. Costanzo, F.; Silvestrelli, P.; Ancilotto, F. Physisorption, Diffusion, and Chemisorption Pathways of H₂ Molecule on Graphene and on (2,2) Carbon Nanotube by First Principles Calculations. *J. Chem. Theory Comput.* **2012**, *8*, 1288–1294. <https://doi.org/10.1021/ct300143a>.
35. Lugo, G.; Cuesta, I.; Sánchez Marín, J.; Sánchez De Merás, A. MP2 Study of Physisorption of Molecular Hydrogen onto Defective Nanotubes: Cooperative Effect in Stone–Wales Defects. *J. Phys. Chem. A* **2016**, *120*, 4951–4960. <https://doi.org/10.1021/acs.jpca.5b12589>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.