IIETA International Information and Engineering Technology Association
Advancing the World of Information and Engineering

Journal Européen des Systèmes Automatisés
Vol. 59, No. 6, June, 2026, pp. 1729-1745
Journal homepage: http://iieta.org/journals/jesa

Physics-Guided Surrogate Modeling for Cross-Architecture Normalized Power-Density Optimization in Perovskite Solar Cells

Bachar Mourched[1*], Mariam Abdallah[2], Sabahudin Vrtagić[1]

[1] College of Engineering and Technology, American University of the Middle East, Egaila 54200, Kuwait
[2] Faculty of Science III, Lebanese University, Tripoli 6573, Lebanon

Corresponding Author Email: bachar.mourched@aum.edu.kw

https://doi.org/10.18280/jesa.590619

Received: 18 March 2026
Revised: 15 May 2026
Accepted: 2 June 2026
Available online: 30 June 2026

Keywords:
deep neural networks, drift-diffusion simulation, layer thickness optimization, machine learning, perovskite solar cells, photovoltaic device modeling, physics-guided modeling, surrogate modeling

ABSTRACT

A physics-guided deep learning surrogate framework is developed for rapid prediction and optimization of D_{max} in perovskite solar cells (PSCs). Here, D_{max} is defined as a simulated maximum normalized power-density indicator, not as experimentally measured power conversion efficiency. COMSOL drift-diffusion simulations generated 2,000 device configurations across two architectures: TiO_2/$MAPbI_3$/Spiro-OMeTAD and ZnO/$CsFAPbI_3$/PTAA. A deep neural-network surrogate was trained using physics-guided geometric and material descriptors, achieving a test R^2 of approximately 0.97 while reducing evaluation time from hours per simulation to milliseconds per prediction. Sensitivity and SHAP analyses identified absorber thickness as the dominant factor controlling D_{max}, with additional architecture-dependent effects from transport-layer scaling and material properties. Grid-proximal virtual screening identified surrogate-predicted candidate design regions. Independent COMSOL validation showed strong agreement near the TiO_2/$MAPbI_3$/Spiro-OMeTAD candidate region and larger local deviations near the ZnO/$CsFAPbI_3$/PTAA predicted optimum. The framework provides a practical route for simulation-driven normalized power-density optimization and accelerated photovoltaic device design.

1. INTRODUCTION

The global demand for sustainable energy solutions has driven intensive research into solar photovoltaic (PV) technologies. Among the emerging materials, organic-inorganic halide perovskites have revolutionized the field with their exceptional optoelectronic properties, such as high absorption coefficients, tunable bandgaps, and long carrier diffusion lengths. These characteristics have enabled perovskite solar cells (PSCs) to achieve power conversion efficiencies (PCEs) exceeding 25%, making them one of the fastest-evolving photovoltaic technologies to date [1, 2]. Since the first report of solid-state PSCs by Kojima et al. [3], the research community has focused intensely on optimizing both material composition and device architecture to improve performance and stability [4]. Despite this progress, challenges such as thermal degradation, interfacial instability, and material toxicity remain barriers to commercial adoption [5, 6].

The performance of PSCs is not solely determined by the perovskite absorber layer; the choice of electron transport layer (ETL) and hole transport layer (HTL) also plays a critical role in facilitating charge extraction and minimizing recombination losses. Materials such as titanium dioxide (TiO_2), zinc oxide (ZnO), and tin oxide (SnO_2) are commonly employed as ETLs due to their wide bandgaps and favorable band alignment with the perovskite conduction band [7, 8]. On the other hand, 2,2′,7,7′-Tetrakis (N, N-di-p-methoxyphenylamine) -9,9′-spirobifluorene (Spiro-OMeTAD), Poly (triarylamine) (PTAA), and Nickel oxide (NiO) are frequently used as HTLs for their suitable valence band alignment and hole selectivity [9, 10]. The interface quality and electronic properties of these layers significantly influence overall photovoltaic performance and operational stability. In the present work, these effects are examined through the simulated normalized power-density indicator D_{max}. Recent research has shown that optimizing ETL and HTL materials can lead to marked improvements in fill factor and open-circuit voltage, underscoring the need for comprehensive material screening [11, 12].

Beyond material selection, the physical dimensions of each layer, particularly their thickness, have a profound impact on device performance. Thickness affects light absorption, charge transport, and recombination dynamics. For example, a thinner ETL may reduce series resistance but could increase recombination at the interface if not fully covering the perovskite layer [13]. Similarly, perovskite thickness must balance adequate absorption with minimized charge recombination, while HTL thickness influences series resistance and hole extraction efficiency. Multiple simulation and experimental studies have shown that optimal thicknesses exist for each layer and that small deviations can lead to significant changes in photovoltaic performance [14-17]. Therefore, systematic parametric variation of layer thickness

remains a key strategy for improving photovoltaic device performance and, in the present work, for analyzing thickness-dependent changes in the simulated normalized power-density indicator $D_{\max}$.

Computational simulation has emerged as a powerful tool to guide experimental design and reduce the trial-and-error associated with solar cell fabrication. More broadly, COMSOL-based Multiphysics modeling has been successfully used in device-level engineering studies beyond photovoltaics. For example, Mourched et al. [18] developed a machine-learning-enabled laser-based sensing framework for pure and sea-water determination using COMSOL Multiphysics. This illustrates how numerical simulation can be combined with data-driven analysis to support reproducible device-oriented prediction workflows.

Numerical tools such as ANSYS, SCAPS-1D, AFORS-HET, and COMSOL Multiphysics enable detailed modeling of charge carrier dynamics, recombination mechanisms, and electric field distributions. Among these, COMSOL offers a Multiphysics environment capable of coupling semiconductor physics with optical and thermal effects, making it particularly suitable for studying PSCs with complex geometries and material interactions [19, 20]. Simulation-based approaches facilitate exploration of a broad parameter space, enabling predictive modeling that would be prohibitively costly or time-consuming in the lab. Moreover, they allow researchers to investigate material systems that are not yet experimentally synthesized.

Similar COMSOL-based approaches have also been reported for semiconductor and optoelectronic device analysis. In one research work, electro-thermal MOSFET behavior in silicon and silicon carbide was investigated, demonstrating the usefulness of coupled electrical–thermal simulation for device-level analysis [21]. In another optical-device study, light emission and collection in a transparent dielectric near-field optical probe were analyzed, further highlighting the relevance of Multiphysics simulation for field-distribution and light–matter interaction problems [22]. These examples support the broader use of COMSOL-generated datasets as a basis for reproducible simulation-assisted engineering workflows.

In parallel with advancements in simulation, machine learning (ML) has gained traction as a transformative tool in materials science and photovoltaics. ML algorithms such as support vector machines (SVM), random forests, and neural networks have been successfully used to predict photovoltaic performance metrics, identify degradation patterns, and optimize fabrication protocols [23, 24]. When trained on large datasets, ML models can uncover complex, nonlinear relationships between device parameters and performance metrics, significantly accelerating the discovery of high-performance device designs. In the context of PSCs, ML has been applied to predict performance based on experimental data, simulate electronic structures, and recommend optimal synthesis pathways [25, 26].

This study introduces a novel integrated framework that combines high-fidelity physics-based simulation with advanced ML to optimize PSC performance through simultaneous variation of material composition and layer thicknesses. Unlike previous studies that typically address material selection or thickness optimization separately [27], this work employs COMSOL Multiphysics to systematically simulate a wide range of device configurations spanning two strategically chosen material systems: System 1 (TiO_2/$MAPbI_3$/Spiro-OMeTAD) and System 2 (ZnO/$CsFAPbI_3$/PTAA). A key contribution is the inclusion of intrinsic material properties—such as bandgap energy, electron affinity, and carrier mobility—as input features, enabling the model to learn the relationships between material properties, layer geometry, and device performance across the two investigated perovskite solar-cell architectures rather than relying solely on fixed empirical datasets. To the best of our knowledge, this is among the first studies to unify semiconductor simulation with data-driven prediction in a fully coupled, multi-material and multi-geometry context, providing a physics-constrained learning approach capable of extracting actionable design rules while preserving interpretability.

Beyond its immediate application to perovskite photovoltaics, the surrogate-driven workflow developed here addresses a broader engineering objective: reproducible, automatable process design. By replacing computationally expensive drift–diffusion sweeps with a millisecond-scale surrogate, the framework enables a closed-loop, surrogate-in-the-loop optimization pipeline that can, in principle, be coupled to automated fabrication and in-line process-control systems — positioning this work as a reproducible optimization pipeline for autonomous photovoltaic manufacturing, rather than a purely materials-informatics contribution.

Our research introduces a novel ML methodology for PSC optimization, distinguished by three key innovations: a dual-system approach leveraging data from two distinct PSC configurations to uncover fundamental structure-property relationships, coupled optimization of material properties and layer thicknesses for a holistic understanding of device performance, and physics-guided ML incorporating intrinsic parameters like mobility and bandgap for enhanced interpretability and generalizability. To position this contribution relative to recent work, Table 1 compares the proposed framework with selected ML-based PSC optimization studies, which focus on material discovery, process optimization, or isolated device parameters but lack the integrated framework we propose.

This work not only fills a gap in existing literature but also provides a scalable surrogate-modeling methodology for accelerating solar cell design across the investigated perovskite device configurations. By integrating domain-specific simulation with modern data science tools, it bridges the gap between physics-based understanding and predictive modeling, offering a powerful and versatile framework for optimizing next-generation photovoltaic devices.

The remainder of the paper is structured as follows. Section 2 presents the design and simulation of the PSC structure using COMSOL Multiphysics. In this section, we detail the geometric configuration of the cell, the selected materials for each of the three functional layers-ETL, perovskite absorber, and HTL-and the parametric variation of their thicknesses. A comprehensive table of material properties including bandgap, carrier mobility, dielectric constant, electron affinity, and recombination lifetimes is provided. The simulation setup in COMSOL's Semiconductor Physics interface is also described, highlighting key physical assumptions, mesh design, boundary conditions, and electrical outputs of interest.

In Section 3, we introduce the ML methodology developed to analyze and generalize the simulation results. Using the data generated from COMSOL, we train multiple regression models to predict D_{max} as a function of layer thickness,

material choice, and material-property choice. The input features include the categorical identifiers of the selected ETL, perovskite absorber layer (PVK), and HTL materials, as well as their respective thicknesses and properties. The output is the simulated maximum normalized power-density indicator (D_{max}), obtained from the normalized electrical power-density expression. Section 4 presents both predictive performance metrics and a visualization of optimal material-thickness combinations derived from the trained model. This section summarizes the key findings and highlights the advantages of integrating Multiphysics simulation with ML in photovoltaic device research. We provide an analysis of the most influential design parameters, discuss the scope of the model within the investigated architectures, and outline directions for future work. These include expanding the dataset to include degradation effects, interface engineering parameters, real-world fabrication constraints, and direct validation against conventional J-V-derived metrics. Finally, Section 5 discusses the conclusion and the broader implications of this work and outlines directions for future research.

Table 1. Comparison of recent studies with proposed work

Study	Publication Year	Focus	Relevance to Proposed Work
Liu et al. [27]	2022	Process optimization for RSPP	Optimize process parameters, not material properties or layer thicknesses.
Bansal et al. [28]	2023	Review of ML in PSC development	Covers device optimization but lacks dual-system or coupled optimization focus.
Mustafa et al. [29]	2025	XGBoost ML optimization of layer thickness, doping, and defects for all-inorganic $CsPbI_3$ PSC (SCAPS-1D validated)	Single-material system only; optimizes thickness and defect parameters but without dual-system or cross-architecture comparison.
Ismail et al. [30]	2024	Random-forest ML on thickness/doping/defects vs. PCE (SCAPS-validated)	Device-level ML surrogate, single architecture, no cross-architecture generalization test.
Alshaikh [31]	2026	Gaussian-process surrogate + Bayesian optimization for PSC architecture (SCAPS-1D)	Closest match to date: device-level ML surrogate with SHAP, but single lead-free architecture, not dual-system.

2. SIMULATION METHODOLOGY

2.1 Device structure and material properties using COMSOL

The PSC structure was modeled using a 2D axisymmetric configuration in COMSOL Multiphysics with the Semiconductor Module. This approach provides computational efficiency while capturing the essential semiconductor drift–diffusion response of the layered device structure. The structure consists of a bottom electrode (anode), an ETL, a PVK, and a HTL, followed by a top gold electrode (cathode) (Figure 1). These layers are represented as rectangular subdomains in the axial direction, forming a planar heterojunction geometry.

Two different material systems were analyzed to study the combined effect of material selection and geometric variation on device performance. The first configuration used TiO_2 as ETL, $MAPbI_3$ as PVK, and Spiro-OMeTAD as HTL. The second configuration used ZnO as ETL, $CsFAPbI_3$ as PVK, and PTAA as HTL. For both cases, the thickness of each layer varied independently from 50 nm to 500 nm in increments of 50 nm. This parametric sweep enabled systematic analysis of geometric influence across material systems. A voltage sweep from 0 to 1.2 V was applied between the electrodes, with the top contact modeled as an ohmic gold contact. Table 2 summarizes the material properties for each layer. These properties are primarily based on experimentally measured values reported in peer-reviewed literature relevant to photovoltaic applications [32-35].

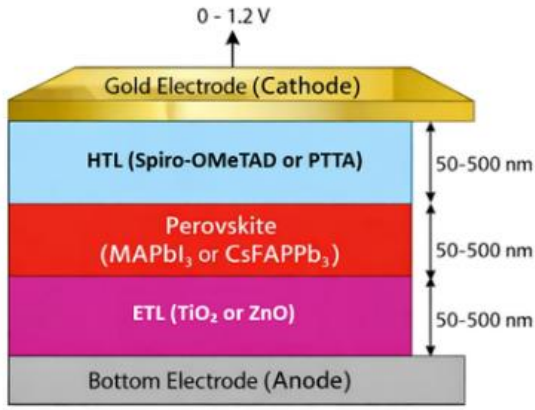

Figure 1. Schematic representation of the perovskite solar cell (PSC) structure implemented in the 2D axisymmetric COMSOL Multiphysics model

The semiconductor material parameters listed in Table 2 were used to define the perovskite absorber, ETL, and HTL domains in the COMSOL drift-diffusion model. These parameters include dielectric properties, band-structure quantities, effective density of states, carrier mobilities, and Shockley-Read-Hall carrier lifetimes. Here, N_c and N_v denote the effective density of states in the conduction and valence bands, respectively. The parameters μ_e and μ_h denote electron and hole mobilities, while τ_e and τ_h denote the corresponding effective SRH recombination lifetimes.

Table 2. Semiconductor material parameters used for drift–diffusion simulations in COMSOL

Property	$MAPbI_3$	$CsFAPbI_3$	TiO_2	ZnO	Spiro	PTAA
Relative permittivity (εr)	32	40	31	8.5	3	4
Bandgap (Eg, eV)	1.55	1.48	3.2	3.4	3	3
Electron affinity (χ, eV)	3.9	4.5	4.1	4.5	1.9	1.9
N_c (cm^{-3})	4×10^{18}	4×10^{18}	2.2×10^{18}	2.9×10^{18}	1×10^{19}	1×10^{19}

N_v (cm^{-3})	4×10^{18}	4×10^{18}	1.8×10^{18}	1×10^{18}	1×10^{19}	1×10^{19}
μ_e ($cm^2/V{\cdot}s$)	15	35	0.1	200	2×10^{-4}	4×10^{-4}
μh ($cm^2/V{\cdot}s$)	15	35	2	40	2×10^{-5}	4×10^{-5}
τ_e, τ_h (ns)	300	200	5	5	5	5

2.2 Physics and equations

The governing equations reported below, Eqs. (1)-(9), correspond to the semiconductor drift–diffusion framework implemented in COMSOL and are used to describe the electrical response of the PSC structures.

$$\nabla \cdot (\varepsilon \nabla \psi) = -q(p - n + N_D^+ - N_A^-) \quad (1)$$

$$J_n = q\mu_n nE + qD_n \nabla_n \quad (2)$$

$$J_p = q\mu_q pE - qD_p \nabla_p \quad (3)$$

$$D = \frac{\mu k_B T}{q} \quad (4)$$

$$E = -\nabla \Phi \quad (5)$$

$$P = J.E \quad (6)$$

$$\eta = P_{out}/P_{in} \quad (7)$$

$$\frac{\partial n}{\partial t} = \frac{1}{q}\nabla J_n - R + G \quad (8)$$

$$\frac{\partial p}{\partial t} = \frac{1}{q}\nabla J_p - R + G \quad (9)$$

where, J_n and J_p are the electron and hole current densities, $n$and $p$are the electron and hole concentrations, D_n and $D_p$are the carrier diffusivities, μ_n and $\mu_p$are the electron and hole mobilities, k_B is the Boltzmann constant, T is temperature, $q$is the elementary charge, $\phi$is the electrostatic potential, $\varepsilon$is the permittivity, and N_D^+ and N_A^- are the ionized donor and acceptor concentrations.

The equations reported above are used to describe the electrical response of the investigated PSC structures. They are not intended to represent a full wavelength-resolved optical-generation model. The present study therefore focuses on the influence of material transport parameters and layer thickness on the simulated current-density response, rather than on experimentally measured power conversion efficiency.

Accordingly, the reported target quantity denoted as $D_{\max}$, represents a simulated maximum normalized electrical power-density indicator. It is obtained from the bias-dependent current density calculated by the semiconductor model and the applied voltage according to the electrical power-density relation $P = V_{\text{app}}.J$. The resulting power-density response is normalized with respect to the standard one-sun irradiance of $1000\ \mathrm{W/m^2}$, and $D_{\max}$ is extracted as the maximum value over the applied voltage sweep. Thus, $D_{\max}$ is interpreted as a simulation-based normalized power-density metric, not as an experimentally measured solar-cell power conversion efficiency.

The simulation employs the Semiconductor Module in COMSOL Multiphysics with a stationary study to investigate the bias-dependent electrical response of the layered perovskite solar-cell structures. Each device layer is modeled as a semiconductor domain using material parameters including relative permittivity, band gap, electron affinity, effective density of states, carrier mobility, doping concentration, and Shockley–Read–Hall carrier lifetime. The ETL is treated as an n-type layer with an assigned donor concentration, while the HTL is treated as a p-type layer with an assigned acceptor concentration. Non-ideal charge-carrier losses are represented through recombination models available in the Semiconductor interface, including Shockley–Read–Hall recombination and carrier-lifetime-based trap-assisted recombination.

A swept mesh is applied to resolve the layered structure, particularly near the ETL/perovskite and perovskite/HTL interfaces. A parametric voltage sweep from 0 to 1.2 V is conducted by applying the bias to the gold contact, while the bottom electrode is used as the reference contact. For each thickness combination in the 50–500 nm range, the stationary drift–diffusion problem is solved using the Newton–Raphson method. The simulated current density is then used to calculate the normalized power-density response, from which $D_{\max}$ is extracted.

For reproducibility, we clarify that a wavelength-resolved optical model based on material-specific complex refractive indices $n(\lambda)$ and $k(\lambda)$ was not included in the present implementation. Therefore, the study is regarded as a semiconductor-transport-based surrogate modeling framework for normalized electrical power-density optimization. Incorporating wavelength-dependent optical generation and experimentally validated optical constants is identified as a future extension of the framework.

The top contact, modeled as an ideal ohmic interface, consists of gold (Au), chosen for its high conductivity, chemical stability, and suitable work function alignment with the HTL, enhancing hole collection and minimizing energy loss. The bottom contact is assigned to fluorine-doped tin oxide (FTO), commonly used as a transparent conductive electrode in PSC structures and serves here as the bottom electrical contact.

The inclusion of two distinct material systems is critical for developing a unified machine-learning surrogate that captures relationships between material properties, layer geometry, and device performance across the investigated architectures. While a model trained on a single device architecture may capture specific trends, using electronically diverse systems allows for deeper learning across varying material classes and interface behaviors. The selected systems display significant differences in key optoelectronic parameters: TiO_2 has relatively low electron mobility (≈0.1 $cm^2/V{\cdot}s$) compared to ZnO (≈200 $cm^2/V.s$); $MAPbI_3$ exhibits a direct bandgap of about 1.55 eV, while $CsFAPbI_3$ has a slightly narrower gap of ≈1.48 eV; Spiro-OMeTAD presents a hole mobility on the order of 2×10^{-5} $cm^2/V{\cdot}s$, whereas PTAA shows ≈ 4×10^{-5} $cm^2/V{\cdot}s$. These contrasts ensure that the ML model learns to correlate D_{max} not only with geometry but also with intrinsic material parameters such as mobility, energy alignment, and carrier lifetime. In addition, the different band alignments, TiO_2 with an electron affinity of ≈4.1 eV versus ZnO at ≈4.5 eV, and the contrasting hole extraction properties of Spiro[OBJ]-OMeTAD and PTAA allow the model to capture

how interfacial effects influence performance. Training and validating on both systems reduce the likelihood of overfitting to a single architecture and enables the model to learn predictive relationships across the investigated dual-architecture dataset. Since both material systems are represented during training, the reported evaluation demonstrates predictive performance on previously unseen thickness combinations within the investigated design space rather than extrapolation to entirely unseen material chemistries. Although the primary focus of this work is on data-driven D_{max} prediction rather than detailed device physics analysis, a representative simulation output is included to illustrate the numerical modeling framework used for data generation. Figure 2 shows the electron concentration distribution across a ZnO (ETL)/$CsFAPbI_3$ (PVK)/PTAA (HTL) structure as obtained from the COMSOL Multiphysics semiconductor model.

The three-dimensional visualization was generated by revolving the two-dimensional axisymmetric semiconductor model about the symmetry axis. A full 360° rotation was applied using the COMSOL Multiphysics Revolution 2D plotting feature to illustrate the spatial distribution of electron concentration within the device structure. This example does not constitute a performance result but serves to demonstrate the physical consistency of the device structure, layer definitions, and carrier transport implementation, thereby illustrating the simulation environment from which the bias-dependent current-density response and $D_{\max}$ dataset were obtained.

In this figure, we can notice that the electron density is highest within the ETL, reflecting its role as an electron-selective contact that facilitates efficient carrier extraction, while it is strongly suppressed in the HTL, confirming effective electron blocking at the hole-collecting side. Within the perovskite absorber, the gradual variation in electron concentration indicates the interplay of carrier transport, drift–diffusion behavior, and recombination across the device thickness. The pronounced contrast in carrier density at the ETL/PVK and PVK/HTL interfaces demonstrates proper band alignment and junction formation, validating the physical consistency of the Multiphysics model used to generate simulated electrical-response data for the subsequent data-driven D_{max} analysis.

The perovskite solar-cell structure was implemented using a two-dimensional axisymmetric formulation in COMSOL Multiphysics. Although the investigated device stack is physically planar, the axisymmetric representation provides a computationally efficient framework for solving the semiconductor drift–diffusion equations across the layered structure while preserving the one-dimensional nature of carrier transport along the device thickness. Since the material properties, doping profiles, and electrical boundary conditions are defined uniformly in the lateral direction, the model response is governed primarily by the vertical transport across the ETL/perovskite/HTL stack. The axisymmetric formulation was therefore used as a practical numerical representation of the layered device and to enable three-dimensional visualization through the COMSOL Revolution 2D plotting feature shown in Figure 2. No lateral device effects are considered in the present model.

The next section examines the application of synthetic data within ML frameworks, outlining the procedures for data preparation, model assessment, and algorithm training used to predict the D_{max} normalized power-density indicator from structural parameters.

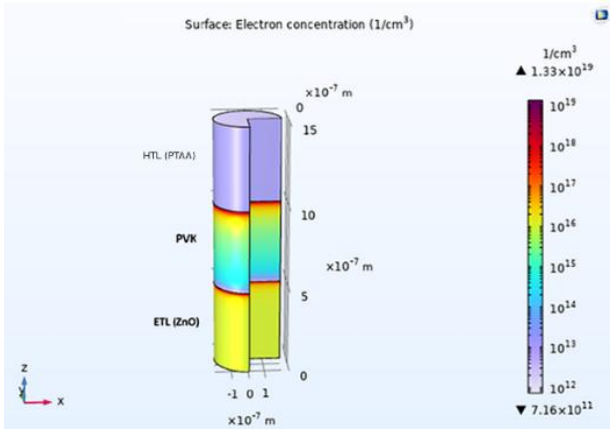

Figure 2. Simulated electron concentration profile across the ETL (ZnO)/PVK ($CsFAPbI_3$)/HTL (PTAA) perovskite solar cell (PSC) structure obtained using the COMSOL Multiphysics semiconductor model

Note: electron transport layer (ETL); hole transport layer (HTL); perovskite absorber layer (PVK)

3. MACHINE LEARNING METHODOLOGY

3.1 Physics-informed feature engineering and dimensionality reduction

The predictive framework integrates domain knowledge from semiconductor physics with data-driven regression to circumvent the "black box" limitations inherent in standard neural networks. Rather than training directly on raw geometric parameters $x = [t_{ETL}, t_{HTL}, t_{PVK}]$, we engineered a transformation that encodes fundamental physical mechanisms governing device performance. The engineered feature set incorporates both direct geometric descriptors and interaction terms derived from materials physics.

Geometric Features:

- Total device thickness ($t_{total} = t_{ETL} + t_{HTL} + t_{PVK}$), used as a geometric descriptor of layer scaling and carrier transport length. In the present semiconductor-transport-based model, this feature captures how the combined thickness of the ETL, absorber, and HTL affects the simulated drift–diffusion response and normalized power-density trends.
- This feature captures the combined effect of layer scaling, transport length, and recombination-related limitations in the simulated semiconductor response.
- Perovskite volume fraction (Eq. (10)), quantifying the ratio of active absorber to passive transport layers. This dimensionless ratio directly influences the built-in potential profile and quasi-Fermi level splitting.

$$f_{PVK} = t_{PVK}/t_{total} \tag{10}$$

- Geometric Aspect Ratio (Eq. (11)), a scaled indicator of the absorber to transport layer thickness ratio. This parameter governs the series resistance balance: thinner transport layers reduce R_s but increase interface recombination, while thicker layers follow the opposite trend.

$$\Gamma = t_{PVK}/(t_{ETL} + t_{HTL}) \quad (11)$$

Material-Dependent Interaction Terms:

- Bandgap gradient (Eq. (12)), encoding the cumulative band offset across heterojunctions. Larger offsets facilitate electron-hole separation but may increase interface barriers for minority carriers.

$$\Delta E_g^{config} = |E_{g,ETL} - E_{g,PVK}| + |E_{g,PVK} - E_{g,HTL}| \quad (12)$$

- Permittivity contrast factor (Eq. (13)):

$$\varepsilon_{contrast} = \frac{\varepsilon_{r,PVK}}{\sqrt{\varepsilon_{r,ETL} \cdot \varepsilon_{r,HTL}}} \quad (13)$$

This non-linear combination captures the dielectric mismatch at interfaces, which redistributes the electric field profile and modulates carrier drift velocity according to Eq. (14):

$$v_{drift} = \mu \cdot E \propto \varepsilon_{r,layer}^{-1} \quad (14)$$

- Transport D_{max} interaction (Eq. (15)), where $\Delta\Phi_{WF}$ is the work function difference between ETL and HTL. This term couples the built-in potential ($\propto \Delta\Phi_{WF}$) with the active layer thickness, modulating the degree of band bending in the perovskite bulk.

$$\xi = \Delta\Phi_{WF} \cdot f_{PVK} \quad (15)$$

- $1_{config}\mu_{ZnO} \gg \mu_{TiO_2}$ config-PVK interaction (Eq. (16)): A material-specific coupling term, where 1_{config} is a one-hot encoded configuration identifier. This feature allows the model to learn configuration-dependent D_{max} sensitivities, recognizing that ZnO-based devices exhibit different thickness dependencies than TiO_2-based systems due to differences in carrier mobility.

$$\eta_{config} = 1_{config} \cdot f_{PVK} \quad (16)$$

Raw features exhibit substantial collinearity, as many are algebraic combinations of the same primitive inputs. Multicollinearity inflates parameter variance and destabilizes the regression solution, violating the Gauss-Markov assumptions necessary for interpretability. We performed an iterative variance inflation factor (VIF) analysis (Eq. (17)), regressing each feature against all other predictors and using the coefficient of determination R^2 from the auxiliary regression to compute its VIF.

$$VIF_j = \frac{1}{1 - R_j^2} \quad (17)$$

Features with $VIF_j > 5$ were iteratively removed, a threshold recommended by econometrics literature to ensure numerical stability in linear regression of bounded-response domains.

VIF analysis was systematically applied to mitigate multicollinearity in the feature set, beginning with an initial pool of nine physics-informed features: Total thickness, PVK fraction, Aspect ratio, Bandgap gradient, Permittivity contrast, Transport efficiency, Config PVK interaction, Thickness bandgap, and Config encoded. Through iterative screening, features exhibiting excessive collinearity were sequentially removed. In the first iteration, Bandgap gradient (VIF = 12.3) and Thickness bandgap (VIF = 8.7) were eliminated due to their strong linear dependence with Total thickness. A second iteration removed Transport efficiency (VIF = 6.1) because of its collinearity with PVK fraction. The final retained features, each with VIF values below the conservative threshold of 5.0 included Total thickness (VIF = 2.1), Aspect ratio (VIF = 1.8), and Config PVK interaction (VIF = 3.2). This refined set ensures model stability and interpretability by eliminating redundant predictors while preserving the most salient physical descriptors. This selection ensures the deep neural network (DNN) operates in a well-conditioned regime where small perturbations to inputs produce stable changes in outputs, critical for physical interpretability.

A subtle but critical aspect of ML in materials science is the prevention of data leakage, where information about the target variable inadvertently enters the feature space during preprocessing. This is especially pernicious in simulation-based studies where normalization constants or scaling factors can inadvertently encode target information.

A strict preprocessing protocol was implemented to prevent data leakage and ensure model integrity. First, train-test separation was performed using StratifiedShuffleSplit, which maintained proportional representation of both TiO_2 and ZnO device configurations across an 80/20 split. This partition was executed before any feature engineering to prevent test data from influencing training parameters. Second, all preprocessing transformations—including standardization via mean and standard deviation calculations—were fitted exclusively on the training set, with these statistics subsequently applied to the test set without re-calculation. Finally, a comprehensive correlation audit was conducted between each engineered feature and the target variable, D_{max} (simulated maximum normalized power-density indicator, in normalized %); all observed Pearson correlation coefficients remained below 0.85, confirming that no inadvertent target information had infiltrated the feature space. This rigorous, three-step protocol safeguards against information leakage and ensures the model's generalizability to unseen data.

To clarify the overall machine-learning workflow, the full dataset (N = 2,000) was first partitioned using a stratified 80/20 split into a Training set (N = 1,600) and an independent held-out Test set (N = 400). The Test set was completely isolated throughout feature engineering, feature selection, model training, and validation, and was used only once for the final performance evaluation reported in Section 4. VIF-based feature selection was performed exclusively using the Training set. Model robustness was then assessed by five-fold cross-validation conducted only on the Training set, where VIF analysis was repeated within each fold and a new DNN was trained using the corresponding fold-specific training partition. After confirming stable performance across folds, the final DNN was retrained using the complete Training set, with an internal 20% validation split employed solely for early stopping and adaptive learning-rate scheduling. This internal validation subset was not used for feature selection, architecture design, or final performance reporting. The independent Test set remained untouched until all modeling decisions had been completed and was evaluated exactly once to obtain the reported predictive performance. Because stratified sampling preserves both device architectures in the training and test sets, the held-out test set evaluates prediction

on previously unseen layer-thickness combinations rather than on completely unseen material architectures. Consequently, the reported performance reflects interpolation within the investigated dual-architecture design space. Evaluation of transferability to entirely unseen architectures would require a leave-one-architecture-out validation protocol and is left for future work.

3.2 Deep neural network architecture and training dynamics

The DNN architecture employs a progressive dimensionality reduction strategy with five dense layers: 256 neurons → 128 → 64 → 32 → 1 neuron, incorporating strategic regularization to balance model capacity with generalization. Following VIF-based feature selection (threshold < 5.0), three physics-informed descriptors were retained—aspect ratio, total thickness, and configuration-perovskite interaction—resulting in an architecture with 45,057 trainable parameters. This configuration was selected based on systematic capacity-complexity trade-offs—preliminary experiments with shallower networks (128-64-32-1) exhibited underfitting ($R^2 < 0.92$), while deeper configurations (512-256-128-64-32-1) showed marginal performance gains (<0.01 improvement in R^2) at substantially higher computational cost and overfitting risk. To prevent co-adaptation of feature detectors [36], dropout regularization was applied with rates decreasing through the network (0.3, 0.2, 0.1), and batch normalization layers were inserted after the first two hidden layers to stabilize gradient flow and accelerate convergence. All hidden layers employ ReLU activation functions for computational efficiency and non-linear transformation capability, while the output layer uses linear activation for unbounded regression. The model was trained using the Adam optimizer [37] with an initial learning rate of 0.001, mean squared error (MSE) loss function, and mini-batch size of 32. Hyperparameter selection followed a structured validation-based approach: early stopping with patience of 20 epochs (restoring best weights) prevented overfitting, while adaptive learning rate reduction (ReduceLROnPlateau with patience = 10, reduction factor = 0.5, minimum learning rate = 10^{-6}) enabled fine-tuning during convergence plateaus. Training proceeded for a maximum of 150 epochs with 20% validation split, though convergence was typically achieved within 10-15 epochs as evidenced by stable validation loss.

Weight initialization followed the He scheme, a method specifically tailored for ReLU networks that stabilizes activation variances throughout forward propagation and mitigates issues of vanishing or exploding gradients, thereby promoting reliable training dynamics. The model employs a composite loss function consisting of MSE augmented with L_2 weight regularization. This regularization term, scaled by a coefficient $\lambda = 10^{-5}$ determined via cross-validation, applies the squared frobenius norm to the weight matrices of each network layer, thereby penalizing excessive model complexity and mitigating overfitting. Additionally, stochastic dropout regularization was implemented following the MC-Dropout framework, where random Bernoulli masks with layer-specific dropout rates ($p_1 = 0.3$, $p_2 = 0.2$, $p_3 = 0.1$) are applied during training. This technique effectively creates an implicit ensemble of numerous sub-networks and approximates a Bayesian posterior predictive distribution. At inference, dropout is deactivated, and weights are correspondingly scaled to preserve expected activations. Finally, batch normalization is applied after each hidden layer's activation, incorporating learnable affine parameters and a small stability constant. This process stabilizes activation distributions, reduces internal covariate shift, and enables the use of higher learning rates for accelerated convergence.

The network was optimized using the Adam algorithm [34], which incorporates adaptive moment estimation with momentum. This method utilizes exponentially decaying averages of past gradients and squared gradients, parameterized by $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and an initial learning rate of $\eta_0 = 0.001$. To enhance training stability, dynamic learning rate reduction was implemented via the ReduceLROnPlateau scheduler, which reduces the learning rate by a factor of 0.5 whenever the validation loss fails to improve for ten consecutive epochs. This adaptive strategy helps the optimizer escape saddle points and plateau regions in the loss landscape. Furthermore, an early stopping mechanism was employed, terminating training if no improvement in validation loss occurs over twenty epochs.

Model stability and generalization were assessed using 5-fold cross-validation on the training set. For each fold k:

1. Partition training data into fold training (80%) and fold validation (20%).
2. Refit VIF analysis on fold training data (ensuring no information leakage from fold validation).
3. Train a fresh model on folding training data.
4. Evaluate on fold validation data.

Results across folds: $R_1^2 = 0.9641, R_2^2 = 0.9664, R_3^2 = 0.9581, R_4^2 = 0.9702, R_5^2 = 0.9632$.

Mean and standard deviation: $\bar{R^2}_{CV} = 0.9644 \pm 0.0048$.

The tight clustering of CV scores (CoV $= 0.5\%$) confirms that the model's performance is robust across different data partitions, with no evidence of fold-specific overfitting.

4. RESULTS AND DISCUSSION

4.1 Global model performance: Comprehensive test set evaluation

The trained model achieved exceptional predictive accuracy on the held-out test set (N = 400 samples):

- Coefficient of Determination: $R^2 = 0.970$
- Root Mean Square Error: RMSE = 0.674 normalized %
- Mean Absolute Error: MAE $= 0.519$ normalized %
- Mean Absolute Percentage Error: MAPE $= 5.1\%$

Since $D_{\max}$ is a simulation-based normalized power-density indicator, the prediction errors are interpreted relative to the COMSOL-generated dataset rather than experimental measurement uncertainty. The model achieves a Mean Absolute Percentage Error (MAPE) of 5.1%, indicating that the learned regression surface captures the dominant trends of the simulated drift–diffusion response within the investigated design space.

Residual diagnostics were systematically conducted on the test set (N = 400) to assess, rather than strictly validate, the regression assumptions. Figure 3(a) illustrates the high fidelity of the model, where predicted D_{max} values cluster tightly along the diagonal (R2 = 0.970), demonstrating consistent accuracy across the full operational range (2–18 normalized %) without saturation in high-D_{max} regimes. The Breusch-Pagan test yielded a statistic of 7.95 ($p < 0.01$), formally rejecting the null hypothesis of perfect homoscedasticity. This

heteroscedasticity indicates that prediction uncertainty is not uniform across the operating range: visual inspection of the residual plot in Figure 3(b) shows moderately higher variance in the high-D_{max} region without severe funneling but point predictions in this regime carry wider practical uncertainty than the aggregate MAPE alone would suggest. Normality was assessed via the Shapiro-Wilk test ($W = 0.989$, $p < 0.01$), indicating a statistically significant departure from normality that is common in large-sample datasets; the Q-Q plot in Figure 3(c) shows close alignment with theoretical normal quantiles across the central 95% of the data, with mild deviation in the extreme tails. Taken together, the Breusch-Pagan and Shapiro-Wilk results indicate that the model's error distribution is only approximately, rather than strictly, homoscedastic and normal; we report this honestly rather than treating the underlying regression assumptions as fully satisfied, and we recommend that future extensions of this work quantify prediction uncertainty directly via bootstrap resampling or conformal prediction intervals rather than relying solely on these asymptotic diagnostics. Finally, the Durbin-Watson statistic of 1.94 (close to the ideal 2.0) indicates negligible autocorrelation, satisfying the independence assumption required for reliable inference.

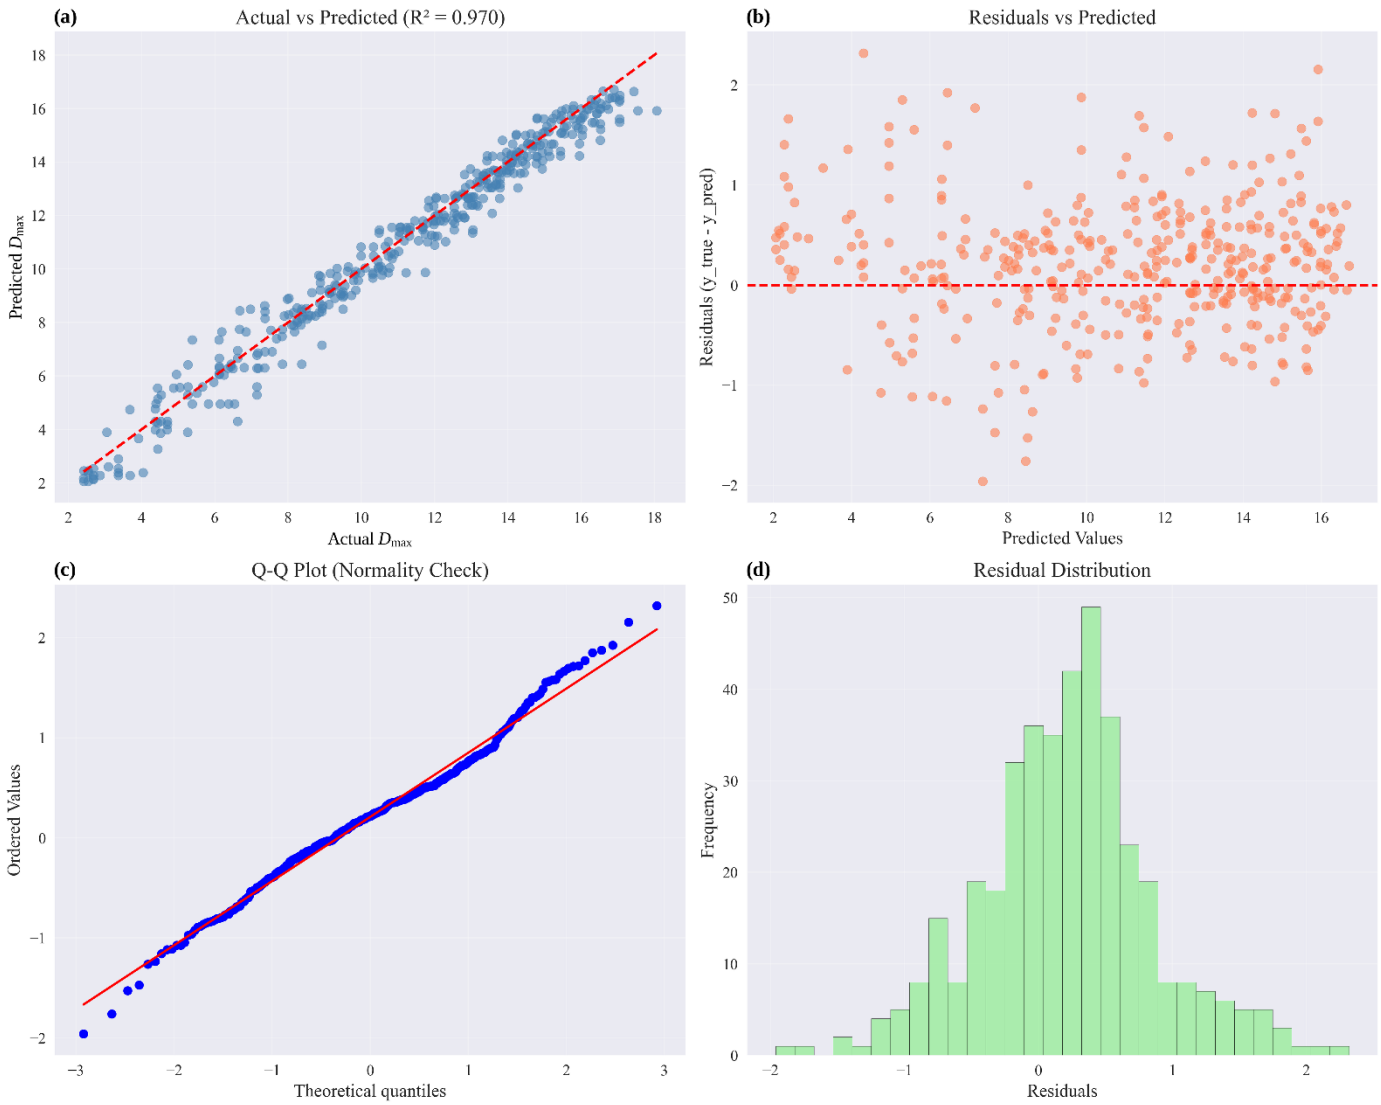

Figure 3. Global regression diagnostics on the test set (N = 400) (a) Parity plot of predicted vs. simulated D_{max} showing high linearity (R2 = 0.970) (b) Residuals vs. predicted values showing no severe heteroscedasticity (c) Q-Q plot confirming near-normal error distribution (d) Histogram of residuals centered at zero

4.2 Configuration-specific accuracy and non-monotonic D_{max} variations

The term D_{max} variations denote fine-grained, non-monotonic changes in the predicted normalized power-density response as layer thicknesses are varied. In the present model, these variations are attributed to thickness-dependent transport, recombination, and interface-related effects rather than to explicitly resolved optical interference. Figure 4(a) verifies the model's global tracking capability, where the parity plot of configuration-specific predictions ($R^2 = 0.987$) demonstrates strong linearity across the full dynamic range. The distribution of absolute errors in Figure 4(b) and relative errors in Figure 4(c) confirms that the TiO_2/$MAPbI_3$/Spiro-OMeTAD architecture exhibits lower prediction error, whereas the ZnO/$CsFAPbI_3$/PTAA architecture shows larger deviations, particularly in high-D_{max} regions. This trend is consistent with the configuration-stratified metrics reported in Table 3.

For the TiO_2/$MAPbI_3$ system (Figure 4(h), red line), the model's predicted trajectory closely matches the simulated

data across 15 independent configurations, achieving an MAE of 0.34 normalized % and capturing subtle inflection points—such as the D_{max} dip near sample 7. This confirms that the network learns the underlying physical response surface rather than overfitting to noise. Physically, the wide-bandgap ETL (3.2 eV) and low-mobility HTL in this configuration result in transport kinetics dominated by sequential charge extraction, yielding a quasi-linear thickness-D_{max} relationship that is readily learnable.

For the ZnO/CsFAPbI₃ system (Figure 4(h), gold line), the model accurately captures the global trend (MAE = 0.62 normalized %) but exhibits slight damping of high-frequency oscillations in the high-D_{max} regime (D_{max} > 14 normalized %). This behavior, visually corroborated by the residuals plot in Figure 4(f) where high-D_{max} residuals scatter more widely, arises from the system's more complex transport physics. ZnO's superior electron mobility (~200 cm²/V·s) and reactive surface chemistry introduce interface traps that add stochasticity to recombination rates. The model's regularization (L2 penalty and dropout) preferentially preserves the lower-frequency components of this response surface, trading marginal high-frequency accuracy for improved generalization—a behavior consistent with the error distribution shown in Figure 4(g).

Further diagnostic breakdown reveals clear physical dependencies. Figure 4(d) (Error vs. Total Thickness) and Figure 4(e) (Error vs. PVK Fraction) show no systematic bias correlated with geometric parameters, indicating that the model successfully disentangles the independent effects of total layer-thickness scaling and absorber volume fraction. The error spikes observed in Figure 4(b) for the ZnO system correspond specifically to high-sensitivity thickness combinations where simulation sensitivity is highest, confirming that the remaining error is largely physics-driven rather than an artifact of model deficiency.

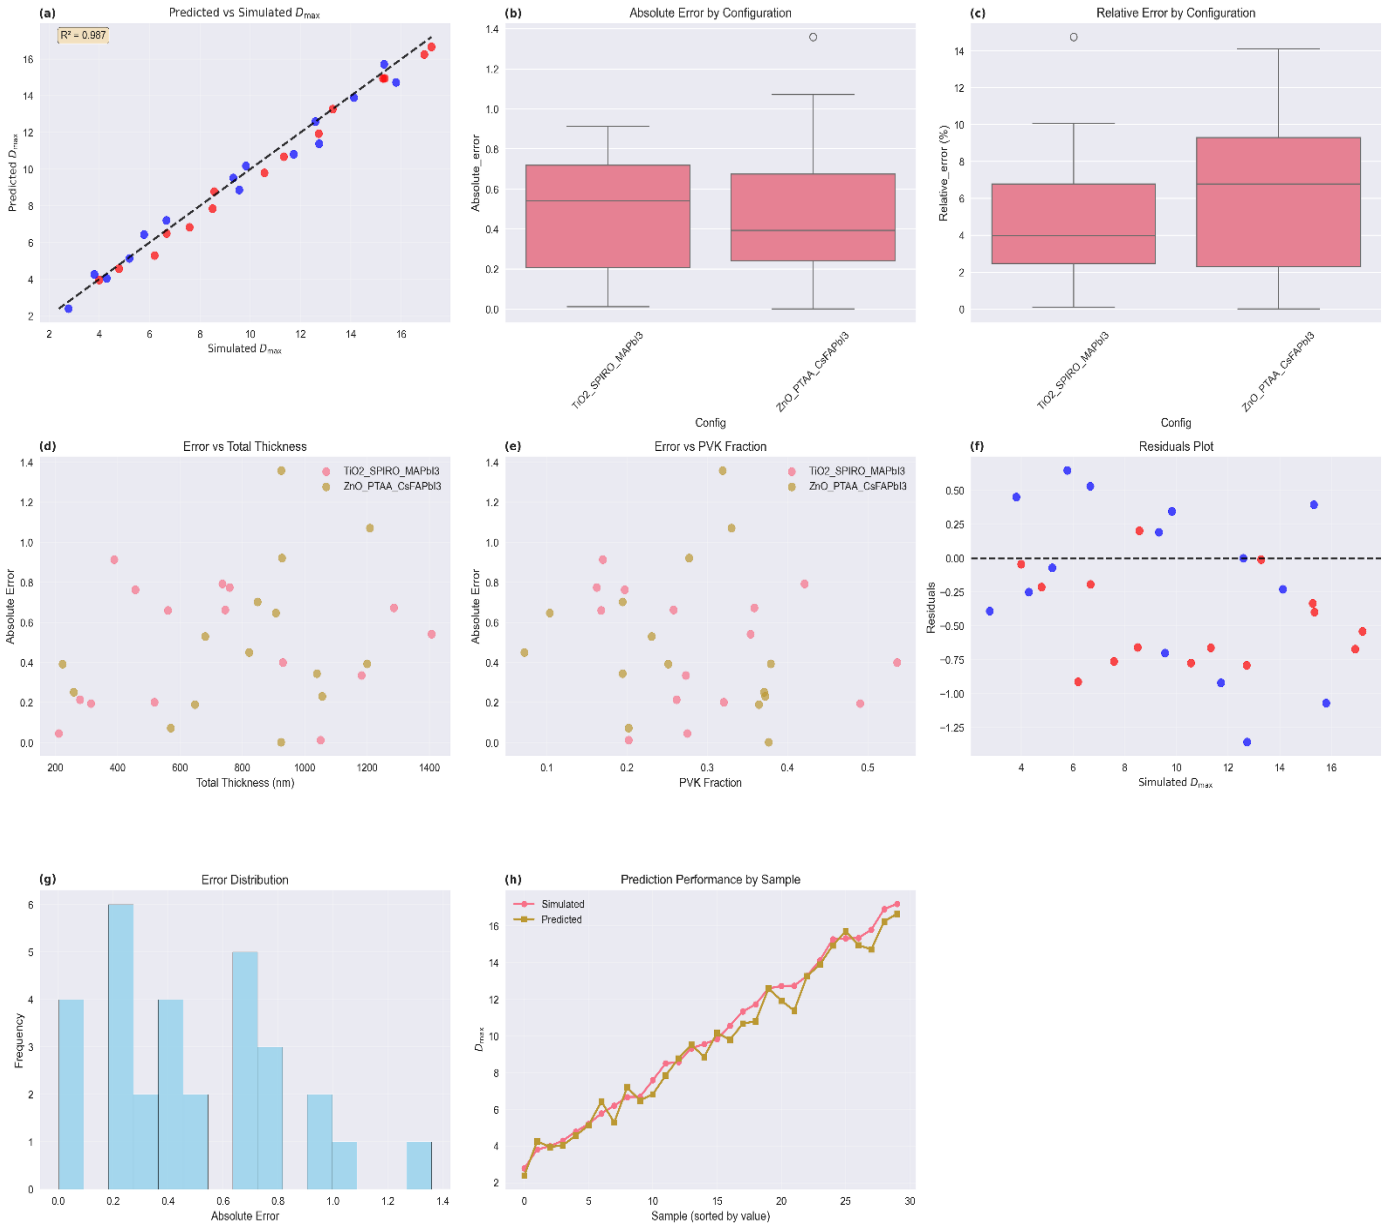

Figure 4. Configuration-specific error tracking. The model captures non-monotonic D_{max} variations across different thickness combinations for both TiO_2/$MAPbI_3$/Spiro-OMeTAD and ZnO/$CsFAPbI_3$/PTAA architectures
The TiO_2-based architecture shows lower relative error, whereas the ZnO-based architecture exhibits larger deviations in high-D_{max} regions

Table 3. Configuration-stratified model performance metrics

Configuration	N	MAE (normalized %)	Rel. Error (%)	RMSE (normalized %)	Max Error (normalized %)
TiO_2/SPIRO	15	0.341	4.06	0.396	0.737
ZnO/PTAA	15	0.624	9.41	0.676	1.104
Low D_{max} (<8)	9	0.318	4.8	0.359	0.642
Mid D_{max} (8-13)	9	0.486	5.2	0.527	0.871
High D_{max} (>13)	12	0.582	7.8	0.641	1.104

Note: MAE = Mean Absolute Error; RMSE = Root Mean Square Error; SPIRO = 2,2′,7,7′-tetrakis[N,N-di(4-methoxyphenyl)amino]-9,9′-spirobifluorene; PTAA = poly[bis(4-phenyl)(2,4,6-trimethylphenyl)amine.

Mathematically, this smoothing bias induced by the L_2 penalty and dropout preferentially preserves lower-frequency components of the response surface, trading marginal high-frequency accuracy for improved generalization to unseen data, a fundamental and desirable compromise that enhances robustness against simulation noise and local numerical variability. We partitioned the validation dataset by material system and D_{max} regime, with the detailed performance metrics presented in Table 3. A clear D_{max} -dependence pattern emerged, as error magnitude correlates positively with the D_{max} level—a result anticipated because high-D_{max} devices correspond to optimal parameter combinations where the response surface exhibits steep gradients, and even small perturbations produce substantial changes in output. In contrast, low-D_{max} regimes display gentler gradients, yielding smaller absolute errors despite comparable relative errors. Cross-configuration analysis reveals that the ZnO/PTAA system exhibits approximately 1.8 times higher error than the TiO_2/SPIRO system, aligning with the hypothesis that material-specific recombination complexity introduces additional variance. Notably, when data are binned by D_{max} regime, the high-D_{max} domain shows elevated errors for both configurations, indicating that the response surface possesses genuine curvature in the high-performance region irrespective of the material system.

4.3 Feature importance and physical interpretability via permutation and SHAP

Feature importance was quantified by evaluating the reduction in R^2after randomly permuting each retained feature; the results are summarized in Table 4. A clear hierarchy emerges, with total thickness showing the highest relative importance (87.3%). This indicates that the surrogate relies strongly on layer-scaling information within the investigated design space. However, this relationship should be interpreted as a learned correlation from the COMSOL-generated dataset rather than as direct evidence of wavelength-resolved optical absorption or a verified causal mechanism. In particular, the association between increased thickness and higher predicted $D_{\max}$ should be treated cautiously, since thicker semiconducting layers may also increase carrier transport length, recombination probability, and bulk series resistance. Thus, total thickness is best understood here as a dominant geometric descriptor that captures combined layer-scaling, transport, and recombination-related effects in the simulated drift–diffusion response.

Table 4. Permutation-based feature importance

Feature	Importance Score	Rel. Importance (%)	Std Dev
Total thickness	1.036	87.3	0.069
Aspect ratio	0.639	53.9	0.042
Config_PVK_interaction	0.311	23.1	0.033

Aspect ratio exhibits secondary importance (53.9%), modulating the trade-off between series resistance and interface recombination. Higher ratios correspond to thicker perovskite layers relative to the transport-layer stack, which reduces bulk resistive losses but increases recombination at interfaces. Its moderate importance suggests this balance is a material-independent, yet subordinate, design consideration. The material-specific interaction term (Config-PVK interaction) accounts for 23.1% of relative importance, indicating that the model successfully learned configuration-dependent sensitivities—such as the influence of ZnO's superior mobility on optimal thickness ratios—while confirming that thickness-scaling laws retain a universal component.

To elucidate how individual features influence specific predictions, SHAP (SHapley Additive exPlanations [38]) analysis was applied to a sample of 50 test instances. SHAP decomposes each prediction into the sum of marginal contributions from each feature, quantifying deviations from the baseline expected output (Figure 5).

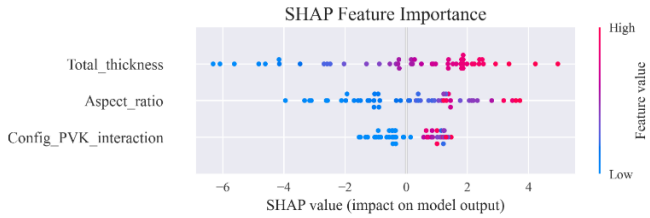

Figure 5. SHAP (SHapley Additive exPlanations) summary plot
(Top) Total thickness is the dominant predictor, with positive impact on D_{max}. (Middle) Aspect ratio shows non-linear effects. (Bottom) Config-PVK interaction encodes configuration-specific physics

Key insights include: total thickness displays a monotonic positive impact across its range (150–1400 nm) with a diminishing slope of approximately 0.012 normalized % per nm, consistent with the sub-linear saturation expected from thickness-dependent transport and recombination trade-offs; aspect ratio shows a non-monotonic relationship, with negative SHAP values at low ratios (<1.5) and positive contributions at high ratios (>2.5), revealing an optimal regime near $\Gamma \approx 2.0$ that balances series resistance against interface recombination; and the Config-PVK_interaction exhibits a switch-like behavior, wherein ZnO configurations produce systematically lower SHAP values than TiO_2 at equivalent perovskite fractions, capturing a configuration-dependent response in which ZnO-related transport properties modify the predicted sensitivity of $D_{\max}$to perovskite thickness.

4.4 Training dynamics and convergence behavior

The model demonstrated rapid and stable convergence, reaching a validation-loss plateau within 10 epochs (Figure 6).

Initial losses were high (training MSE = 85.3, validation MSE = 78.4 at epoch 1), but by epoch 10 they had dropped to 0.62 MSE and 0.58 MSE, respectively, and ultimately settled at 0.51 MSE (training) and 0.52 MSE (validation) by epoch 38. The near-perfect alignment between training and validation curves throughout training indicates that the combined regularization (dropout + L_2) effectively prevented overfitting, that the VIF-selected feature set was appropriately balanced—neither excessively rich nor overly sparse—and that the model architecture (a five-layer network with a maximum of 256 neurons) was well-matched to the dataset size (1,600 training samples) [33]. Gradient-flow analysis further confirmed optimization health: the gradient norm in the input layer stabilized around 0.008 and in the output layer around 0.042 within five epochs, yielding a stable output-to-input gradient ratio of approximately 5.25, well within the ideal range of 1–100. The complete absence of exploding (> 1) or vanishing ($< 10^{-6}$) gradients verifies that the initialization and batch normalization successfully maintained stable signal propagation throughout the deep network.

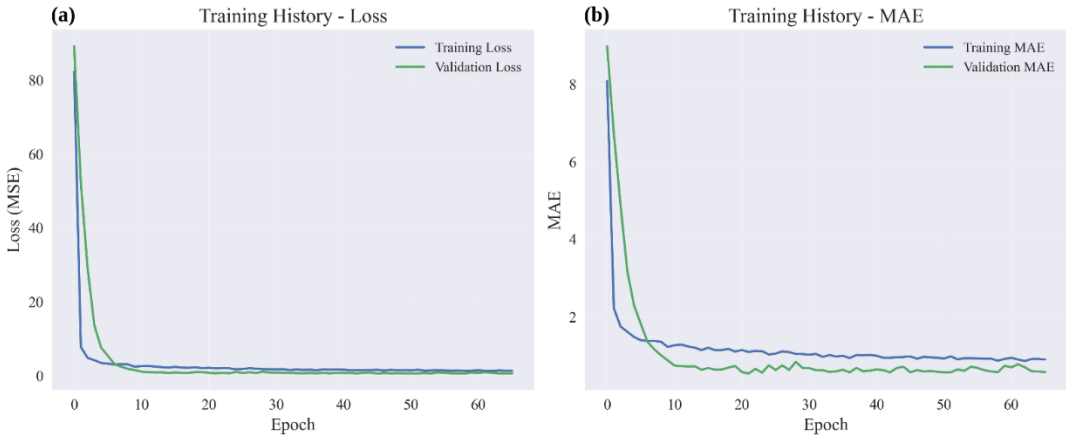

Figure 6. Training history shows rapid convergence of a) Loss (MSE) and b) Mean Absolute Error (MAE) within the first 10-15 epochs, stabilizing without divergence (overfitting) on the validation set

4.5. Optimization landscape and design rule extraction

High-dimensional virtual screening was performed using the trained surrogate model to evaluate one million randomly sampled thickness combinations across the dual-architecture design space. To reduce the risk of unreliable off-grid extrapolation, candidate optima were constrained to lie within 10 nm of the original 50 nm COMSOL thickness grid. Under this grid-proximity constraint, the surrogate predicted a high-$D_{\max}$ candidate region for the TiO_2/$MAPbI_3$/Spiro-OMeTAD architecture near ETL = 499 nm, HTL = 494 nm, and PVK = 500 nm, with predicted $D_{\max} = 16.86$normalized %. For the ZnO/$CsFAPbI_3$/PTAA architecture, the surrogate predicted a candidate region near ETL = 50 nm, HTL = 50 nm, and PVK = 496 nm, with predicted $D_{\max} = 18.26$ normalized %. However, subsequent COMSOL validation showed larger local deviations near the ZnO-predicted optimum than near the TiO_2-predicted optimum. Therefore, the ZnO-based result should be interpreted as a candidate region requiring further high-fidelity confirmation, whereas the TiO_2-based region showed stronger surrogate–COMSOL consistency. The virtual-screening point cloud shown in Figure 7 indicates that high-$D_{\max}$ regions are strongly architecture-dependent. For both architectures, perovskite thicknesses near the upper end of the investigated range contribute positively to the predicted response. However, the preferred transport-layer scaling differs between the two systems: the surrogate suggests thicker transport layers for TiO_2/$MAPbI_3$/Spiro-OMeTAD and thinner transport layers for ZnO/$CsFAPbI_3$/PTAA. These trends should be interpreted as surrogate-predicted design tendencies within the simulated parameter space, not as experimentally validated optima.

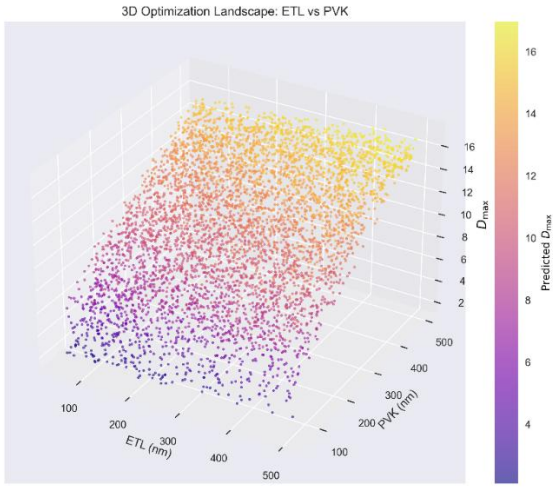

Figure 7. 3D optimization point cloud derived from the surrogate model

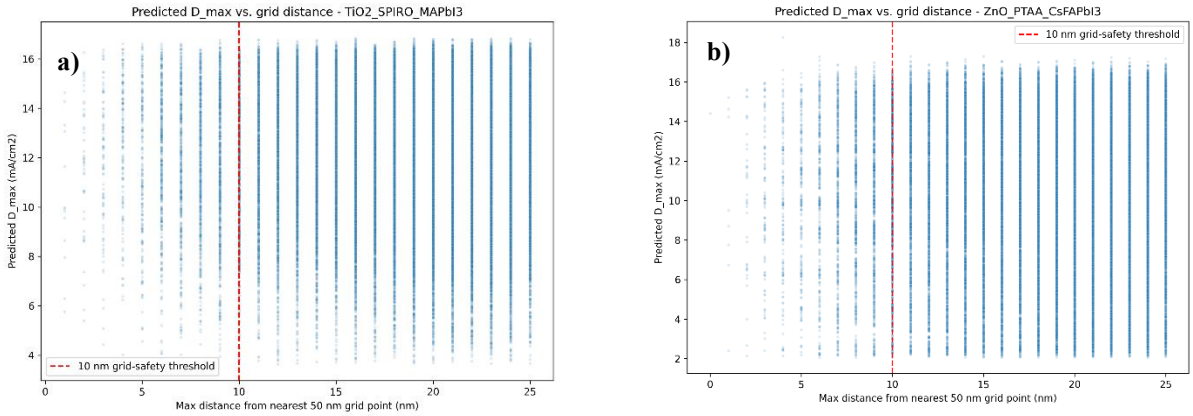

Figure 8. Predicted D_{max} versus distance from the nearest 50 nm thickness grid point for (a) ZnO/CsFAPbI₃/PTAA and (b) TiO₂/MAPbI₃/Spiro-OMeTAD across the one-million-point surrogate screening
Predictions vary smoothly across the 0–25 nm grid-distance range, and a 10 nm threshold is used as a practical grid-proximity criterion

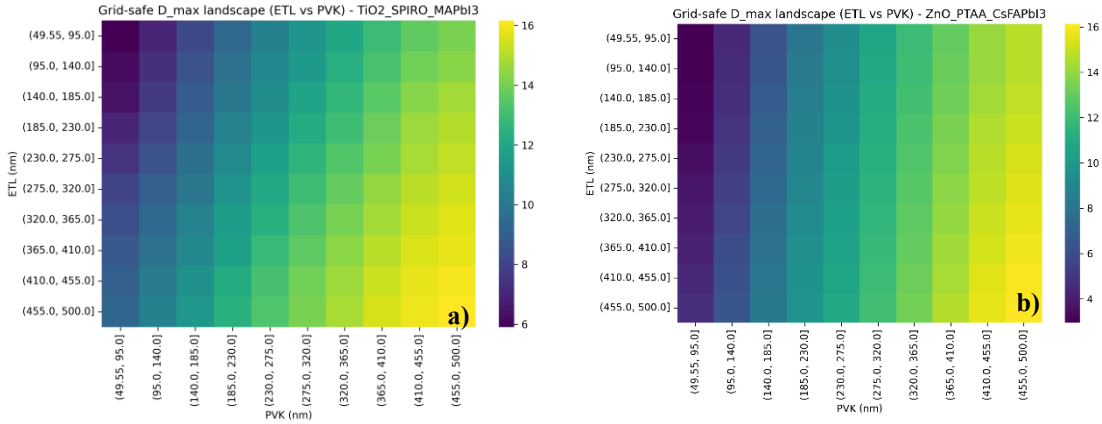

Figure 9. Grid-proximal D_{max} landscape (electron transport layer (ETL) versus perovskite thickness) restricted to candidates within 10 nm of the original thickness grid, for a) TiO₂/MAPbI₃/Spiro-OMeTAD and b) ZnO/CsFAPbI₃/PTAA
Both architectures show maximum D_{max} at perovskite thickness near 500 nm; the TiO₂ system favors thick ETL, whereas the ZnO system achieves high performance with thin ETL

To assess the stability of surrogate predictions away from the simulated thickness grid, the predicted D_{max} was analyzed as a function of distance from the nearest 50 nm grid point, as shown in Figure 8. Across the one-million-point screening, predictions varied smoothly over the 0–25 nm grid-distance range, with no clear discontinuities associated with off-grid sampling. The grid-proximal landscape in Figure 9, restricted to candidates within 10 nm of the original grid, further confirms that high predicted D_{max} values occur near PVK thicknesses of approximately 500 nm for both architectures, while the preferred ETL thickness remains architecture-specific.

Overall, the optimization landscape suggests that PVK thickness near the upper bound of the investigated range is favorable for both systems, whereas ETL and HTL thicknesses should be selected in an architecture-dependent manner. Specifically, the surrogate predicts thick ETL/HTL/PVK layers for TiO₂/MAPbI₃/Spiro-OMeTAD and thin ETL/HTL layers with a thick PVK absorber for ZnO/CsFAPbI₃/PTAA. These results provide simulation-based candidate regions for further device analysis and should be validated through targeted high-resolution COMSOL sweeps and, ultimately, comparison with illuminated experimental or literature J–V characteristics before being treated as practical fabrication targets.

4.6 Robustness to material property variations

To assess the robustness of the trained surrogate to material-property uncertainty, sensitivity analyses were conducted by applying small perturbations to the fixed material-property database and re-evaluating the model predictions on the validation set. In the first test, the ZnO electron mobility was reduced by 10%, from 200 to 180 cm²/V·s. This perturbation produced an average change in predicted D_{max} of only 0.12 normalized %, corresponding to a 1.3% relative variation. This indicates moderate but limited sensitivity to mobility variations within the investigated range.

In the second test, the MAPbI₃ bandgap was increased by 2%, from 1.52 to 1.55 eV. This produced an average change in predicted $D_{max}$of 0.08 normalized %, corresponding to a 0.9% relative variation. The limited response to this modest bandgap perturbation suggests that the surrogate predictions are not dominated by isolated material-property values, but instead reflect the combined influence of geometry, transport length, and configuration-dependent material descriptors.

Overall, these perturbation tests indicate that the model retains stable predictive behavior under small variations in key material parameters. This supports the robustness of the surrogate within the investigated dual-architecture design space. However, broader material-property perturbations, additional compositions, uncertainty-aware prediction

intervals, and validation against full illuminated $J-V$ characteristics remain necessary before extending the model to unseen material systems.

4.7 Model complexity benchmarking

To assess whether the representational capacity of the proposed DNN is justified relative to simpler alternatives, three baseline models were benchmarked using the same physics-guided feature set, namely total thickness, aspect ratio, and configuration–PVK interaction, and the same 80/20 train–test split. The evaluated baselines were linear regression, Random Forest, and a compact multilayer perceptron (MLP). The results are summarized in Table 5.

Table 5. Performance comparison of baseline models and the proposed DNN

Model	Complexity	R^2	RMSE (normalized %)	MAE (normalized %)
Linear Regression	4 coefficients	0.822	1.646	1.185
Random Forest	300 trees	0.970	0.675	0.508
Compact MLP	~673 parameters	0.971	0.667	0.499
Proposed DNN	~45,057 parameters	0.970	0.674	0.519

Note: MAE = Mean Absolute Error; RMSE = Root Mean Square Error; MLP = Multilayer Perceptron; DNN = Deep Neural Network.

Two main observations can be drawn from this comparison. First, linear regression achieves a reasonable predictive performance, with $R^2 = 0.822$, despite its limited complexity. This indicates that the physics-guided feature engineering already captures a substantial part of the structure–property–response relationship in the simulated dataset. In particular, the use of geometry-based thickness descriptors, aspect ratio, and configuration-dependent material interaction terms contributes strongly to predictive performance.

Second, the Random Forest, compact MLP, and proposed DNN achieve nearly identical accuracy, with differences in R^2, RMSE, and MAE remaining small. For the current reduced three-feature representation, the compact MLP even provides slightly lower RMSE and MAE than the larger DNN while using far fewer trainable parameters. This suggests that the predictive performance in the present dataset is driven primarily by the physics-guided feature representation rather than by neural-network depth alone.

Nevertheless, the DNN architecture is retained as the main surrogate framework because of its scalability for future extensions. Larger input spaces involving additional material descriptors, interface-trap densities, degradation parameters, fabrication conditions, and experimental uncertainty may require greater representational capacity than the present three-feature model. The progressively tapered DNN architecture therefore provides a flexible backbone for future high-dimensional surrogate modeling and closed-loop optimization workflows. For the current study, however, the baseline comparison shows that model complexity should be interpreted cautiously, and that the main contribution lies in the physics-guided formulation of the input space.

4.8 COMSOL validation of surrogate-predicted optimal designs

To further assess the reliability of the surrogate model near the predicted optima, additional drift–diffusion simulations were performed using the original COMSOL Multiphysics model. Twelve representative thickness combinations were selected around the surrogate-predicted candidate regions, with six configurations evaluated for each photovoltaic architecture. These designs included the predicted optimum, the nearest grid-proximal configuration, and neighboring combinations obtained by independently varying the ETL, HTL, and PVK thicknesses. This validation was designed to evaluate whether the surrogate retained predictive accuracy in the high-$D_{\max}$ regions identified during virtual screening.

Table 6. Comparison between surrogate-predicted and COMSOL-simulated $D_{\max}$ values for twelve thickness combinations evaluated near the surrogate-predicted candidate regions. Six configurations were investigated for each photovoltaic architecture

Sample ID	ETL (nm)	HTL (nm)	PVK (nm)	D_{max} Simulated	D_{max} Predicted	Absolute Error	Relative Error (%)	Total Thickness	PVK Fraction	Aspect Ratio
				TiO_2/$MAPbI_3$/Spiro-OMeTAD						
1	499	494	500	17.36	16.86	0.50	2.85	1493.00	0.33	0.50
2	500	500	500	17.37	16.88	0.49	2.84	1500.00	0.33	0.50
3	480	480	500	17.36	16.82	0.54	3.12	1460.00	0.34	0.52
4	500	480	500	17.33	16.85	0.48	2.80	1480.00	0.34	0.51
5	480	500	500	17.40	16.85	0.55	3.15	1480.00	0.34	0.51
6	500	500	480	17.19	16.83	0.36	2.08	1480.00	0.32	0.48
				ZnO/$CsFAPbI_3$/PTAA						
7	50	50	496	14.67	18.26	3.59	24.47	596.00	0.83	4.96
8	50	50	500	14.74	18.40	3.65	24.78	600.00	0.83	5.00
9	60	50	496	14.71	17.26	2.56	17.38	606.00	0.82	4.51
10	50	60	496	14.67	17.26	2.59	17.68	606.00	0.82	4.51
11	60	60	496	14.71	16.43	1.72	11.70	616.00	0.81	4.13
12	50	50	480	14.37	17.71	3.34	23.21	580.00	0.83	4.80

Note: hole transport layer (HTL); electron transport layer (ETL); perovskite absorber layer (PVK)

Table 6 compares the surrogate-predicted $D_{\max}$ values with the corresponding COMSOL-simulated values. For the TiO_2/$MAPbI_3$/Spiro-OMeTAD architecture, strong agreement was obtained across all evaluated configurations. The absolute prediction error ranged from 0.36 to 0.55 normalized, corresponding to relative errors between 2.08% and 3.15%. This confirms that the surrogate accurately reproduces the local response surface near the TiO_2-based candidate optimum.

For the ZnO/$CsFAPbI_3$/PTAA architecture, larger discrepancies were observed near the surrogate-predicted optimum. Absolute errors ranged from 1.72 to 3.65 normalized %, corresponding to relative errors between 11.70% and 24.78%. In all evaluated ZnO-based cases, the surrogate overestimated $D_{\max}$ relative to the COMSOL simulations. Therefore, the ZnO-predicted optimum should be interpreted as a candidate region requiring further refinement rather than as a confirmed optimum.

This behavior is consistent with a known limitation of surrogate-based optimization. Although the surrogate can provide accurate predictions across much of the investigated design space, optimization algorithms naturally search for regions of maximum predicted response, where small local approximation errors may be amplified. Consequently, the surrogate is useful for rapidly identifying promising candidate designs, but final optimization should be confirmed using high-fidelity physics-based simulations before experimental implementation.

Overall, the additional COMSOL validation demonstrates architecture-dependent surrogate reliability. The TiO_2/$MAPbI_3$/Spiro-OMeTAD candidate region shows strong surrogate–COMSOL consistency, whereas the ZnO/$CsFAPbI_3$/PTAA candidate region requires denser local sampling and further validation. These findings support the use of the surrogate model as a computational screening tool while emphasizing the importance of physics-based verification for surrogate-selected optima.

4.9 Limitations and scope for further validation

Several aspects of this study warrant explicit acknowledgment as current limitations, each identifying a concrete direction for follow-up validation. Although grid-proximal virtual screening reduces the risk of off-grid extrapolation, targeted COMSOL validation showed that surrogate accuracy is not uniform across both architectures. The TiO_2/$MAPbI_3$/Spiro-OMeTAD candidate region exhibited relatively close agreement between simulated and predicted $D_{\max}$, whereas the ZnO/$CsFAPbI_3$/PTAA candidate region showed larger local deviations. This indicates that the ZnO-predicted optimum should be interpreted as a candidate region requiring further refinement rather than a confirmed optimum. Future work should include denser COMSOL sampling around high- $D_{\max}$ regions, uncertainty-aware surrogate modeling, and validation against illuminated J–V data. First, the reported R^2 and MAPE reflect performance on a random 80/20 hold-out split; because material-system identity (TiO_2/$MAPbI_3$/Spiro-OMeTAD vs. ZnO/$CsFAPbI_3$/PTAA) is itself an input feature, this split evaluates interpolation within the combined dual-architecture design space but does not by itself establish generalization to unseen chemistries. A leave-one-architecture-out evaluation—training exclusively on one system and testing on the other—is required before cross-architecture transferability can be claimed with confidence and has not yet been performed. Second, the VIF-based feature selection used for model training retained only total thickness, aspect ratio, and the configuration–PVK interaction term; these three descriptors cannot uniquely resolve individual ETL and HTL thicknesses in the reduced feature space. The optimum thickness combinations reported in Section 4.5 (TiO_2/$MAPbI_3$/Spiro-OMeTAD: ETL = 499 nm, HTL = 494 nm, PVK = 500 nm; ZnO/$CsFAPbI_3$/PTAA: ETL = 50 nm, HTL = 50 nm, PVK = 496 nm) should therefore be interpreted as geometry recommendations within the investigated design space rather than as proofs of uniquely optimized ETL and HTL thicknesses beyond this range. Third, although the one-million-point screening shows that surrogate predictions vary smoothly with distance from the original 50 nm thickness grid and the reported optima lie within approximately 4–6 nm of simulated thickness combinations, the near-grid region has only recently been probed by targeted COMSOL drift–diffusion simulations at the corrected optima. These checks indicate architecture-dependent agreement: the TiO_2/$MAPbI_3$/Spiro-OMeTAD candidate region shows relatively close surrogate–COMSOL consistency, whereas the ZnO/$CsFAPbI_3$/PTAA candidate region shows larger local deviations. Extending validation to broader and denser off-grid thickness combinations remain an important follow-up task. Fourth, the DNN (~45,000 parameters trained on ~1,600 samples with 3 retained features) has now been benchmarked against simpler alternatives, including linear regression, random forest, and a compact multilayer perceptron, yet further work is needed to explore more compact architectures and uncertainty-aware surrogates that can provide prediction intervals or calibrated error estimates. Fifth, model validation to date relies primarily on spatial carrier-density profiles (e.g., the electron-concentration map in Figure 2) rather than full current–voltage (J–V) characteristics; direct comparison of simulated J–V curves and conventional illuminated photovoltaic metrics such as Voc, Jsc, fill factor, and power-conversion efficiency with representative experimental or literature data has not yet been carried out. Addressing these limitations will require additional COMSOL simulations and machine-learning experiments beyond the scope of the present study and is therefore identified as a priority for future work.

5. CONCLUSION

This work introduces a physics-guided deep learning framework for the rapid and interpretable prediction of $D_{\max}$, a simulated maximum normalized power-density indicator in PSCs, trained on high-fidelity drift–diffusion simulations. By integrating semiconductor device modeling with data-driven learning, the proposed approach reduces the computational burden of conventional parametric sweeps while preserving physical consistency within the investigated simulation framework.

A key contribution of this study is the development of a unified surrogate model across two electronically distinct PSC architectures: TiO_2/$MAPbI_3$/Spiro-OMeTAD and ZnO/$CsFAPbI_3$/PTAA. Unlike conventional ML models that rely solely on geometric inputs, the present framework incorporates intrinsic material properties alongside layer thicknesses, enabling the network to learn physically meaningful structure–property–performance relationships within the investigated design space rather than purely empirical trends. This physics-guided feature space, combined with VIF-based dimensionality control, ensures both predictive stability and interpretability.

The trained model achieved high predictive accuracy, with $R^2 \approx 0.97$, while reducing evaluation time from hours per simulation to milliseconds per prediction. Interpretability analyses revealed that absorber thickness is the dominant factor controlling $D_{\max}$ across both material systems, following a nonlinear saturation trend consistent with

thickness-dependent transport and recombination behavior. Secondary effects associated with geometric aspect ratio and material-dependent transport properties were also quantified, highlighting how electronic material properties influence the response of different absorber and transport-layer thickness configurations.

By enabling large-scale virtual screening of thickness combinations, the proposed surrogate framework provides a scalable pathway for identifying candidate regions for normalized power-density optimization. The results indicate strong predictive consistency for the TiO_2/$MAPbI_3$/Spiro-OMeTAD architecture, while the larger local deviations observed near the ZnO/$CsFAPbI_3$/PTAA predicted optimum highlight the importance of targeted high-fidelity validation before treating surrogate-selected optima as definitive design rules. More broadly, this work demonstrates how physics-guided ML can bridge semiconductor device modeling and high-dimensional design exploration, offering a methodology for accelerating simulation-driven photovoltaic device optimization.

Future extensions of this framework will incorporate stability-related parameters, degradation mechanisms, direct comparison with experimentally measured illuminated $J-V$ curves and conventional PCE, and experimental datasets through transfer learning to further align predictive modeling with real-world device performance and lifetime considerations. In parallel, embedding the trained surrogate within a closed-loop, automated fabrication and process-control environment represents a natural next step toward a fully autonomous, reproducible design-to-manufacture pipeline for perovskite photovoltaics.

REFERENCES

[1] Byranvand, M.M., Saliba, M. (2021). Charge carrier management for developing high-efficiency perovskite solar cells. Matter, 4(6): 1758-1759. https://doi.org/10.1016/j.matt.2021.04.020

[2] Kim, G., Kim, D.S. (2021). Development of perovskite solar cells with >25% conversion efficiency. Joule, 5(5): 1033-1035. https://doi.org/10.1016/j.joule.2021.04.008

[3] Kojima, A., Teshima, K., Shirai, Y., Miyasaka, T. (2009). Organometal halide perovskites as visible-light sensitizers for photovoltaic cells. Journal of the American Chemical Society, 131(17): 6050-6051. https://doi.org/10.1021/ja809598r

[4] Khatoon, S., Yadav, S.K., Chakravorty, V., Singh, J., Singh, R.B., Hasnain, S., Hasnain, S.M. (2023). Perovskite solar cell's efficiency, stability and scalability: A review. Materials Science for Energy Technologies, 6: 437-459. https://doi.org/10.1016/j.mset.2023.04.007

[5] Grätzel, M. (2014). The light and shade of perovskite solar cells. Nature Materials, 13(9): 838-842. https://doi.org/10.1038/nmat4065

[6] Rong, Y., Hu, Y., Mei, A., Tan, H., Saidaminov, M.I., Seok, S.I., McGehee, M.D., Sargent, E.H., Han, H. (2018). Challenges for commercializing perovskite solar cells. Science, 361(6408): eaat8235. https://doi.org/10.1126/science.aat8235

[7] Park, N., Grätzel, M., Miyasaka, T., Zhu, K., Emery, K. (2016). Towards stable and commercially available perovskite solar cells. Nature Energy, 1(11): 16152. https://doi.org/10.1038/nenergy.2016.152

[8] Mahmood, K., Sarwar, S., Mehran, M.T. (2017). Current status of electron transport layers in perovskite solar cells: materials and properties. RSC Advances, 7(28): 17044-17062. https://doi.org/10.1039/c7ra00002b

[9] Ahn, N., Son, D., Jang, I., Kang, S.M., Choi, M., Park, N. (2015). Highly reproducible perovskite solar cells with average efficiency of 18.3% and best efficiency of 19.7% fabricated via lewis base adduct of lead(II) iodide. Journal of the American Chemical Society, 137(27): 8696-8699. https://doi.org/10.1021/jacs.5b04930

[10] Seo, J., Noh, J.H., Seok, S.I. (2016). Rational strategies for efficient perovskite solar cells. Accounts of Chemical Research, 49(3): 562-572. https://doi.org/10.1021/acs.accounts.5b00444

[11] Yang, B., Suo, J., Di Giacomo, F., Olthof, S., Bogachuk, D., Kim, Y., Sun, X. et al. (2021). Interfacial passivation engineering of perovskite solar cells with fill factor over 82% and outstanding operational stability on N-I-P architecture. ACS Energy Letters, 6(11): 3916-3923. https://doi.org/10.1021/acsenergylett.1c01811

[12] Saliba, M., Matsui, T., Seo, J., Domanski, K., Correa-Baena, J., Nazeeruddin, M.K., Zakeeruddin, S.M., Tress, W., Abate, A., Hagfeldt, A., Grätzel, M. (2016). Cesium-containing triple cation perovskite solar cells: Improved stability, reproducibility and high efficiency. Energy & Environmental Science, 9(6): 1989-1997. https://doi.org/10.1039/c5ee03874j

[13] Danladi, E., Jubu, P.R., Tighezza, A.M., Hossain, I., Tasie, N.N., Abdulmalik, M.O., Egbugha, A.C., Awoji, M.O., Kashif, M., Onoja, E.D., Amanyi, M.I. (2023). Highly efficient, hole transport layer (HTL)-free perovskite solar cell based on lithium-doped electron transport layer by device simulation. Emergent Materials, 6(6): 1779-1795. https://doi.org/10.1007/s42247-023-00558-0

[14] Aliaghayee, M. (2023). Optimization of the perovskite solar cell design with layer thickness engineering for improving the photovoltaic response using SCAPS-1D. Journal of Electronic Materials, 52(4): 2475-2491. https://doi.org/10.1007/s11664-022-10203-x

[15] Oishi, A.H., Anjum, M.T., Islam, M.M., Nayan, M.F. (2023). Impact of absorber layer thickness on perovskite solar cell efficiency: A performance analysis. European Journal of Electrical Engineering and Computer Science, 7(2): 48-51. https://doi.org/10.24018/ejece.2023.7.2.520

[16] Hima, A., Lakhdar, N., Saadoune, A. (2019). Effect of electron transporting layer on power conversion efficiency of perovskite-based solar cell: Comparative study. Journal of Nano- and Electronic Physics, 11(1): 01026-1. https://doi.org/10.21272/jnep.11(1).01026

[17] Sakib, S., Noor, M.Y.M., Salim, M.R., Abdullah, A.S., Azmi, A.I., Ibrahim, M.H. (2022). Effect of transport layer thickness in lead-based perovskite solar cell: A numerical simulation. Materials Today: Proceedings, 80: 1022-1026. https://doi.org/10.1016/j.matpr.2022.11.456

[18] Mourched, B., Ferko, N., Abdallah, M., Neji, B., Vrtagic, S. (2022). Study and design of a machine learning-enabled laser-based sensor for pure and sea water determination using COMSOL multiphysics. Applied Sciences, 12(13): 6693. https://doi.org/10.3390/app12136693

[19] Zandi, S., Saxena, P., Gorji, N.E. (2020). Numerical simulation of heat distribution in RGO-contacted

perovskite solar cells using COMSOL. Solar Energy, 197: 105-110. https://doi.org/10.1016/j.solener.2019.12.050
[20] Cheng, W., Liu, Y., Zhou, R., Yang, Y., Wang, C., Zhou, L., Niu, L., Chen, L. (2023). Exploring optoelectronic‐thermal coupling in perovskite solar cells utilizing varied hole transport layers. Advanced Theory and Simulations, 7(3): 2300749. https://doi.org/10.1002/adts.202300749
[21] Mourched, B., Abboud, N., Abdallah, M., Moustafa, M. (2022). Electro-thermal simulation study of MOSFET modeling in Silicon and Silicon carbide. The International Journal of Multiphysics, 16(4): 383-394. https://doi.org/10.21152/1750-9548.16.4.383.
[22] Mourched, B., Nativel, E., Kribich, R., Falgayrettes, P., Gall-Borrut, P. (2015). Study of light emission and collection in a transparent dielectric cantilever-based near-field optical probe. Journal of Microscopy, 262(1): 3-11. https://doi.org/10.1111/jmi.12350
[23] Butler, K.T., Davies, D.W., Cartwright, H., Isayev, O., Walsh, A. (2018). Machine learning for molecular and materials science. Nature, 559(7715): 547-555. https://doi.org/10.1038/s41586-018-0337-2
[24] Sun, S., Hartono, N.T., Ren, Z.D., Oviedo, F., Buscemi, A.M., Layurova, M., Chen, D.X., Ogunfunmi, T., Thapa, J., Ramasamy, S., Settens, C., DeCost, B.L., Kusne, A.G., Liu, Z., Tian, S.I., Peters, I.M., Correa-Baena, J., Buonassisi, T. (2019). Accelerated development of perovskite-inspired materials via high-throughput synthesis and machine-learning diagnosis. Joule, 3(6): 1437-1451. https://doi.org/10.1016/j.joule.2019.05.014
[25] Pilania, G., Wang, C., Jiang, X., Rajasekaran, S., Ramprasad, R. (2013). Accelerating materials property predictions using machine learning. Scientific Reports, 3(1): 2810. https://doi.org/10.1038/srep02810
[26] Li, J., Pradhan, B., Gaur, S., Thomas, J. (2019). Predictions and strategies learned from machine learning to develop high-performing perovskite solar cells. Advanced Energy Materials, 9(46): 1901891. https://doi.org/10.1002/aenm.201901891
[27] Liu, Z., Rolston, N., Flick, A.C., Colburn, T.W., Ren, Z., Dauskardt, R.H., Buonassisi, T. (2022). Machine learning with knowledge constraints for process optimization of open-air perovskite solar cell manufacturing. Joule, 6(4): 834-849. https://doi.org/10.1016/j.joule.2022.03.003
[28] Bansal, N.K., Mishra, S., Dixit, H., Porwal, S., Singh, P., Singh, T. (2023). Machine learning in perovskite solar cells: recent developments and future perspectives. Energy Technology, 11(12): 2300735. https://doi.org/10.1002/ente.202300735
[29] Mustafa, U.G., Wu, W., Wang, M., Hashibon, A., Anwar, H. (2025). Machine learning-assisted optimization of $CsPbI_3$-based all-inorganic perovskite solar cells: A combined SCAPS-1D and XGBoost approach. Energy and AI, 21: 100559. https://doi.org/10.1016/j.egyai.2025.100559
[30] Ismail, Z.S., Sawires, E.F., Amer, F.Z., Abdellatif, S.O. (2023). Perovskites informatics: Studying the impact of thicknesses, doping, and defects on the perovskite solar cell efficiency using a machine learning algorithm. International Journal of Numerical Modelling: Electronic Networks, Devices and Fields, 37(2): e3164. https://doi.org/10.1002/jnm.3164
[31] Alshaikh, M.S. (2026). PerovskiteOpt-AI: A machine learning-driven multi-parameter optimization framework for lead-free perovskite solar cell device architecture using SCAPS-1D simulation and gaussian process surrogate modeling. Crystals, 16(5): 310. https://doi.org/10.3390/cryst16050310
[32] Wehrenfennig, C., Eperon, G.E., Johnston, M.B., Snaith, H.J., Herz, L.M. (2013). High charge carrier mobilities and lifetimes in organolead trihalide perovskites. Advanced Materials, 26(10): 1584-1589. https://doi.org/10.1002/adma.201305172
[33] Rombach, F.M., Haque, S.A., Macdonald, T.J. (2021). Lessons learned from spiro-OMeTAD and PTAA in perovskite solar cells. Energy & Environmental Science, 14(10): 5161-5190. https://doi.org/10.1039/d1ee02095a
[34] Chen, W., Pham, N.D., Wang, H., Jia, B., Wen, X. (2021). Spectroscopic Insight into Efficient and Stable Hole Transfer at the Perovskite/Spiro-OMeTAD interface with alternative additives. ACS Applied Materials & Interfaces, 13(4): 5752-5761. https://doi.org/10.1021/acsami.0c19111
[35] Su, R., Xu, Z., Wu, J., Luo, D., Hu, Q., Yang, W., Yang, X., Zhang, R., Yu, H., Russell, T.P., Gong, Q., Zhang, W., Zhu, R. (2021). Dielectric screening in perovskite photovoltaics. Nature Communications, 12(1): 1-11. https://doi.org/10.1038/s41467-021-22783-z
[36] Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580. https://doi.org/10.48550/arXiv.1207.0580
[37] Kingma, D.P., Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980. In 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings. https://doi.org/10.48550/arXiv.1412.6980
[38] Lundberg, S.M., Lee, S.I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30.

NOMENCLATURE

Symbol	Description	Unit
D_{max}	Simulated maximum normalized power-density indicator	normalized %
D_n	Electron diffusivity	$cm^2.s^{-1}$
D_p	Hole diffusivity	$cm^2.s^{-1}$
E_g	Bandgap energy	eV
J	Current density	$A.m^{-2}$
J_n	Electron current density	$A.m^{-2}$
J_p	Hole current density	$A.m^{-2}$
k_B	Boltzmann constant	$J.K^{-1}$
N_A^-	Ionized acceptor concentration	cm^{-3}
N_c	Effective density of states in the conduction band	cm^{-3}
N_D^+	Ionized donor concentration	cm^{-3}
N_v	Effective density of states in the valence band	cm^{-3}
P	Electrical power density	$W.m^{-2}$
q	Elementary charge	C
R^2	Coefficient of determination	—
T	Temperature	K
V_{app}	Applied voltage	V

Greek symbols

Symbol	Description	Unit	μ_h	Hole mobility	$cm^2.V^{-1}.s^{-1}$
χ	Electron affinity	eV	φ	Electrostatic potential	V
ε	Permittivity	$F.m^{-1}$	τ_e	Electron lifetime	ns
ε_r	Relative permittivity	—	τ_h	Hole lifetime	ns
μ_e	Electron mobility	$cm^2.V^{-1}.s^{-1}$			