

CycleChemist: A Dual-Pronged Machine Learning Framework for Organic Photovoltaic Discovery

Hou Hei Lam^{*1}, Jiangjie Qiu^{*1}, Xiuyuan Hu¹, Wentao Li¹, Fankun Zeng¹, Siwei Fu¹, Hao Zhang¹, Xiaonan Wang¹

¹Tsinghua University

linhx22@mails.tsinghua.edu.cn, qjj21@mails.tsinghua.edu.cn, huxy22@mails.tsinghua.edu.cn, liwt24@mails.tsinghua.edu.cn, zfk23@mails.tsinghua.edu.cn, fsw22@mails.tsinghua.edu.cn, haozhang@tsinghua.edu.cn, wangxiaonan@tsinghua.edu.cn

Abstract

Organic photovoltaic (OPV) materials offer a promising pathway for sustainable energy generation. However, their development is hindered by the challenge of identifying high-performance donor-acceptor pairs with optimal power conversion efficiencies (PCEs). Most existing design strategies focus exclusively on either the donor or the acceptor, rather than employing a unified model capable of designing both components. In this work, we introduce a dual-pronged machine learning framework for OPV discovery, integrating predictive modeling and generative molecular design. In this study, we propose the newly curated Organic Photovoltaic Donor-Acceptor Dataset (OPV²D), the largest of its kind, comprising 2,000 experimentally characterized donor-acceptor pairs. This dataset serves as a comprehensive foundation for model training and evaluation. To enable accurate property prediction in organic photovoltaic (OPV) materials, we first introduce the Organic Photovoltaic Classifier (OPVC) to predict the likelihood that a given material exhibits OPV behavior. Complementing this, we develop a hierarchical graph neural network framework that integrates multi-task learning and cross-modal donor-acceptor interaction modeling. This framework includes the Molecular Orbital Energy Estimator (MOE²) for predicting the highest occupied molecular orbital-lowest unoccupied molecular orbital (HOMO-LUMO) energy levels, and the Photovoltaic Performance Predictor (P³) for estimating power conversion efficiency (PCE). In addition, we introduce the Material Generative Pretrained Transformer (MatGPT) to generate synthetically accessible organic semiconductors. Building on this, we propose a reinforcement learning strategy with three-objective policy optimization that guides molecular generation while preserving chemical validity. By bridging molecular representation learning with device performance prediction, our framework advances computational OPV material discovery.

Introduction

Artificial intelligence (AI) is increasingly recognized as a transformative force for advancing social good, especially in the field of renewable energy. Its integration with materials science is revolutionizing the discovery and design of functional materials through data-driven approaches that accel-

erate innovation. In the realm of sustainable energy, organic photovoltaic (OPV) materials have garnered significant interest due to their flexibility, lightweight nature, and low cost. However, despite these advantages, the development of high-performance OPV systems remains a major challenge, particularly in designing optimal donor-acceptor combinations. Traditional design methods typically rely on incremental modifications of existing molecular structures and extensive trial-and-error experimentation. This approach not only demands considerable human and material resources but also limits the exploration of the design space, making it difficult to identify potential high-performance molecules beyond current systems. Moreover, most molecular design strategies focus exclusively on either the donor or the acceptor, typically optimizing one while fully fixing a specific instance of the other. As a result, there is a lack of a generalizable model capable of designing both components in a flexible manner.

These limitations have motivated us to explore novel machine learning techniques for more efficient OPV material discovery. As shown in Figure 1, we develop CycleChemist, a dual-pronged machine learning framework that integrates predictive modeling and generative molecular design, thereby overcoming previous paradigm limitations. The primary contributions are:

- **Dataset Construction:** We curated the **Organic Photovoltaic Donor-Acceptor Dataset (OPV²D)**, consisting of 2,000 experimentally characterized donor-acceptor pairs with calibrated properties, making it one of the largest OPV datasets available.
- **OPV Property Prediction Models:** We introduced **Organic Photovoltaic Classifier (OPVC)** that predicts the likelihood of a material exhibiting OPV behavior. In addition, our multi-task Graph Neural Network (GNN) predicts HOMO-LUMO levels and power conversion efficiency (PCE). It integrates the **Molecular Orbital Energy Estimator (MOE²)** for HOMO-LUMO predictions and the **Photovoltaic Performance Predictor (P³)** for PCE estimation.
- **Generative Model Development:** We designed and developed the **Material Generative Pretrained Transformer (MatGPT)** model. By incorporating Rotary Position Embedding (RoPE), a Gated Linear Unit (GLU)

^{*}These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

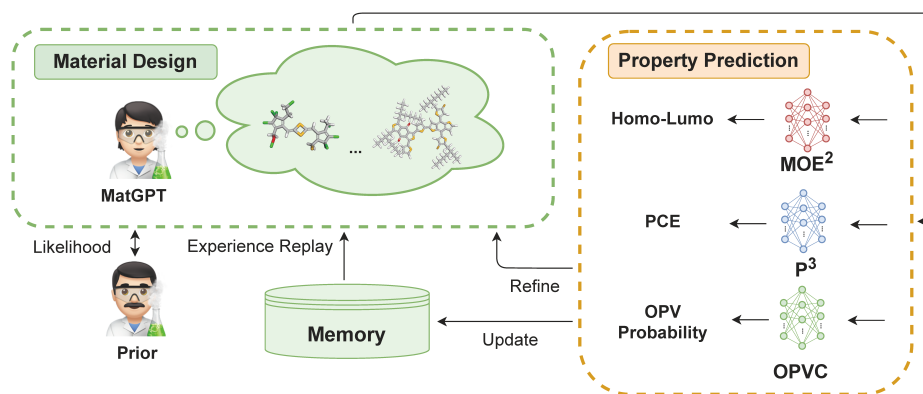


Figure 1: The dual-pronged OPV materials discovery framework of our work integrates material generation, property prediction, and reinforcement learning for iterative enhancement. **Molecular Orbital Energy Estimator (MOE²)** is pre-trained via masked language modeling (MLM) and fine-tuned for HOMO-LUMO prediction. **Photovoltaic Performance Predictor (P³)** predicts power conversion efficiency (PCE) through supervised regression. **Organic Photovoltaic Classifier (OPVC)** is a Random Forest model that predicts the probability of a material being an organic photovoltaic (OPV). **Material Generative Pretrained Transformer (MatGPT)** is pre-trained using causal language modeling (CLM) and fine-tuned with reinforcement learning (RL) for high-PCE OPV molecular generation.

variant, and a diversified sampling method, this model is able to generate valid and diverse molecules.

- **Reinforcement Learning Optimization:** A novel **three-objective RL strategy** balances theoretical performance, molecular validity, and electronic property distribution, guiding the generation of optimal OPV candidates.

Related Work

Prediction of Power Conversion Efficiency (PCE)

Current machine learning models used for predicting the power conversion efficiency (PCE) of organic photovoltaic (OPV) materials remain relatively simplistic. Limitations in both algorithm design and dataset sizes have resulted in insufficient predictive accuracy and generalizability. For example, a comprehensive survey of various embedding methods and machine learning models was conducted using 558 data points, with tree-based algorithms achieving the highest predictive correlation (approximately 0.61) (Seifrid et al. 2024). However, such a small dataset may not adequately represent mainstream OPV material systems, thereby significantly affecting generalizability. In another study, a larger dataset was compiled and a random forest model was trained, reporting an R^2 value of 0.71 (Miyake and Saeki 2021). Nonetheless, the reliance on experimental data as model input renders this approach impractical for large-scale screening studies. In (Zhang et al. 2025b), training was performed on 1,060 donor–acceptor pairs using various graph neural network algorithms, with the Graph Attention Network (GAT) achieving the best result (reported r value of 0.74). Additionally, the model presented in (Das and Mondal 2024) was trained on 300 donor–acceptor pairs and involved detailed quantum mechanical calculations for each molecule during dataset preparation, introducing new quantum mechanical descriptors. On this small dataset, the gradient boosting (GB) model achieved the best performance with

a reported r value of 0.92. However, the complexity of the quantum mechanical (QM) calculations limits the model’s potential for large-scale screening. Similarly, another study utilizing only 100 data points incorporated numerous QM descriptors and reported an r value of 0.918 (Khatua, Das, and Mondal 2024). In summary, these methods are either trained on relatively small datasets or require complex QM calculations or additional experimental parameters, all of which limit their applicability for large-scale screening and may lead to performance that is somewhat contingent on chance.

Photovoltaic Material Discovery

Previous research has predominantly focused on predicting the properties of photovoltaic materials, and in recent years, some scholars have begun to explore donor–acceptor design based on machine learning. According to the literature, methods for exploring high-efficiency OPV materials mainly include strategies such as disassembling and recombining existing molecular structures, as well as employing variational autoencoders (VAE) and genetic algorithms (GA) (Greenstein and Hutchison 2023; Sun et al. 2024; Zhang et al. 2025a). These approaches largely rely on optimizing known molecules, thereby limiting their capacity to generate and explore entirely novel molecules. Some studies do not involve molecular design at all but instead exhaustively enumerate possible donor–acceptor combinations, exploring potential high-efficiency pairs within a predefined design space (Khatua, Das, and Mondal 2024). Moreover, most related studies employ datasets comprising only 100 to 1,500 samples (Miyake and Saeki 2021; Wu et al. 2020; Greenstein and Hutchison 2023), covering relatively narrow material systems, which consequently restricts the diversity of molecular fragments that the models can explore.

Methodology

Dataset Preparation

Based on previous works (Miyake and Saeki 2021; Wu et al. 2020), we integrated and curated multiple organic photovoltaic datasets to construct the **Organic Photovoltaic Donor-Acceptor Dataset (OPV²D)**, comprising approximately 2,000 donor-acceptor pairs—one of the largest OPV datasets to date. Molecules are represented using **SMILES (Simplified Molecular Input Line Entry System)** (Weininger 1988). Duplicate entries were removed, and for pairs with varying PCEs due to fabrication differences, the highest reported value was retained to align with our model’s goal of predicting maximum achievable PCE. Data formats were standardized through manual inspection and corrections using the original literature. To ensure consistency, simplified alkyl chains in some datasets were restored to their authentic forms.

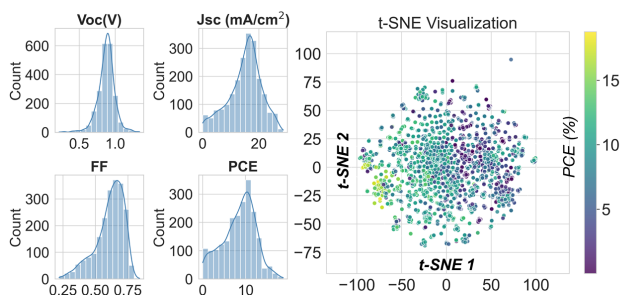


Figure 2: Dataset overview: photovoltaic parameter distributions and molecular structural diversity revealed by t-SNE visualization. The data spans a wide range of property values and structural variations.

Figure 2 illustrates the statistical distribution and structural diversity of the photovoltaic parameters in our dataset. Histograms (left) for Voc (open-circuit voltage), Jsc (short-circuit current density), FF (fill factor), and PCE reveal central tendencies and spread. The t-SNE visualization (right) maps donor-acceptor pairs into a 2D latent space, color-coded by PCE. This underscores the broad coverage of the molecular configurations and experimental outcomes in our dataset, which supports robust model generalization.

Property Prediction

The hierarchical model integrates three complementary components to predict key aspects of molecular and device performance: (1) the Organic Photovoltaic Classifier (OPVC) for identifying OPV-active molecules, (2) the Molecular Orbital Energy Estimator (MOE²) for predicting orbital energies HOMO-LUMO, and (3) the Photovoltaic Performance Predictor (P³) for estimating power conversion efficiency (PCE) of donor-acceptor pairs (see Figure 3).

The design addresses three major challenges in OPV modeling: (1) representing complex molecular structures and electronic interactions, (2) learning transferable representations that capture both local and global chemical features, and (3) modeling donor-acceptor interactions gov-

erning charge transfer and exciton dynamics. These are tackled through a combination of cheminformatics features, molecular graph encoding, multi-task pretraining, and cross-attention mechanisms, balancing physical interpretability with predictive accuracy.

OPV Classifier (OPVC) To estimate the likelihood of organic photovoltaic (OPV) activity, we train a binary Random Forest (RF) classifier (Breiman 2001). The dataset comprises 1,500 OPV-active molecules from the OPV²D dataset and 1,500 presumed-inactive molecules randomly selected from PubChem (Kim et al. 2025). The output indicates the model’s confidence that the molecule is OPV-active.

Molecular Orbital Energy Estimator (MOE²)

Graph Encoder Architecture The molecular encoder uses a three-stage GATv2-based hierarchy (Brody, Alon, and Yahav 2021) to capture local and global interactions. Feature propagation at each layer is defined as:

$$\mathbf{z}_i^{(l+1)} = \text{GATv2} \left(\mathbf{z}_i^{(l)}, \bigoplus_{j \in \mathcal{M}(i)} \psi(\mathbf{z}_j^{(l)}, \mathbf{e}_{ij}) \right), \quad (1)$$

where, $\mathbf{z}_i^{(l)}$ denotes the feature vector of node i at layer l , and $\mathcal{M}(i)$ represents its neighboring nodes. The operator \bigoplus is an attention-based message aggregation function with 8 attention heads, and $\psi(\mathbf{z}_j^{(l)}, \mathbf{e}_{ij})$ incorporates edge attributes \mathbf{e}_{ij} from molecular graph, which encode bond properties such as type, aromaticity, and conjugation. Molecular graphs are derived from SMILES using RDKit (RDKit, online), with atoms as nodes and bonds as edges. An edge-conditioned dynamic attention mechanism allows the model to distinguish between different interatomic interactions, such as covalent bonds, and π - π .

Attentional Graph Regression. To predict orbital energies, node embeddings \mathbf{Z} are pooled via an attentional aggregation layer into a graph-level representation \mathbf{h}_G . This vector is then fed into a Multi-Layer Perceptron (MLP) regressor to simultaneously output \hat{y}_{HOMO} and \hat{y}_{LUMO} .

Photovoltaic Performance Predictor (P³)

Cross-Domain Interaction. We model the donor-acceptor interplay using a dual cross-attention mechanism. By treating one molecule’s embeddings as Queries and the other’s as Keys and Values (and vice-versa), the model captures directional physical interactions:

$$\begin{aligned} \text{Attn}_D &= \text{Attention}(Q_A, K_D, V_D), \\ \text{Attn}_A &= \text{Attention}(Q_D, K_A, V_A). \end{aligned} \quad (2)$$

This formulation explicitly encodes both donor-to-acceptor charge transfer and acceptor-induced polarization effects, which are subsequently concatenated with molecular features to predict PCE via a non-linear regression head.

Nonlinear Regression Final power conversion efficiency (PCE) prediction is computed via a non-linear fully connected feedforward regression function Ψ , integrating donor

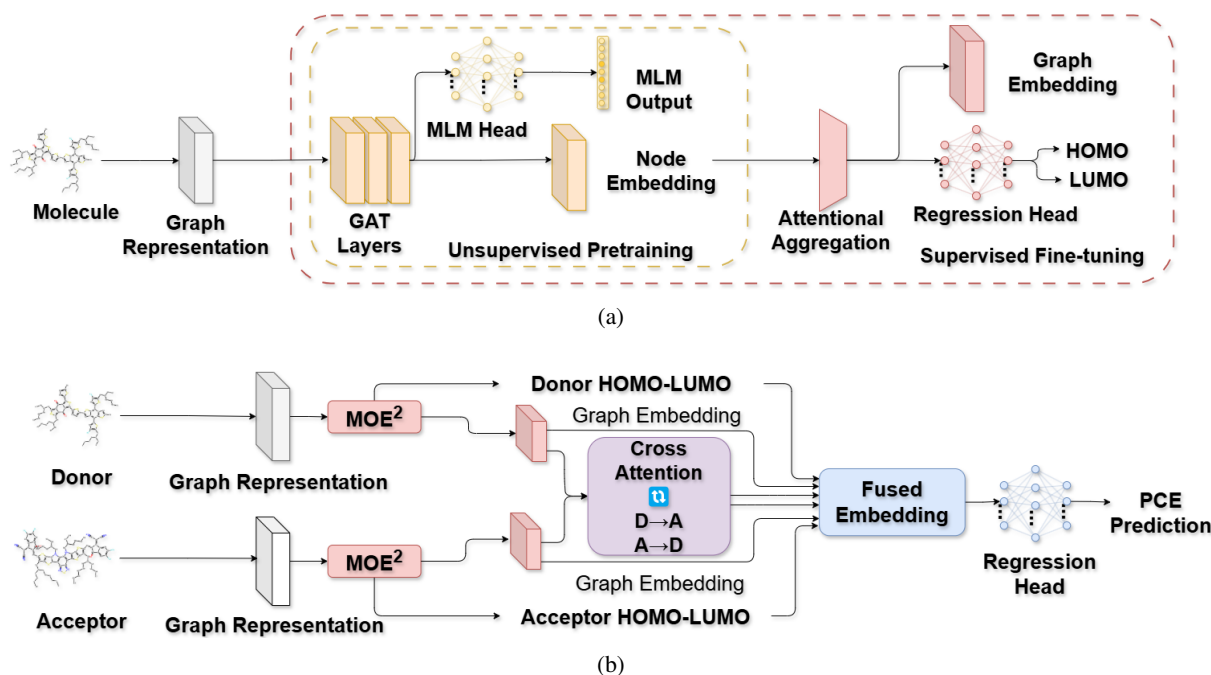


Figure 3: The architecture of (a) **MOE²** and (b) **P³** integrates hierarchical graph neural networks, multi-task learning, and cross-attention mechanisms to jointly predict molecular electronic properties (HOMO-LUMO) and photovoltaic performance (PCE), enabling efficient donor-acceptor screening for organic photovoltaics.

and acceptor features, cross-attention outputs, and predicted orbital energies through feature concatenation:

$$\hat{y}_{PCE} = \Psi([\mathbf{h}_D; \mathbf{h}_A; \text{Attn}_D; \text{Attn}_A; \hat{y}_{HOMO}^D; \hat{y}_{LUMO}^D; \hat{y}_{HOMO}^A; \hat{y}_{LUMO}^A]). \quad (3)$$

The architecture maintains physical interpretability by: (1) separating molecular and interaction features, (2) using dedicated heads for energy prediction, and (3) enabling attention weights to reflect orbital overlap. This unified design supports simultaneous molecular and device-level prediction.

MatGPT Architecture

The **Material Generative Pretrained Transformer (MatGPT)** adapts the GPT-2 architecture for molecular generation. To capture the cyclic and relative patterns in SMILES, we replace absolute positional embeddings with **Rotary Positional Encoding (RoPE)** (Su et al. 2024). Furthermore, we integrate **Gated Linear Units (GLU)** (Hendrycks and Gimpel 2016) into the feed-forward layers to improve local-global feature interaction via dynamic gating.

To mitigate mode collapse and ensure diversity, we implement a dynamic sampling strategy. We combine *top-p* filtering with temperature annealing, where the temperature $T(t)$ linearly decays from T_{\max} to T_{\min} over the generation steps. This schedule encourages exploration in early tokens while sharpening the distribution for valid syntax towards the end.

Reinforcement Learning Framework

We employ a three-objective optimization strategy to guide the MatGPT agent. The agent is initialized from a pre-trained prior and optimized via a composite loss function that balances photovoltaic performance, policy alignment, and statistical property distribution:

$$\mathcal{L}_{\text{total}} = \underbrace{\mathbb{E}_u \left[(\zeta_1 \cdot r(u) - (\log P_{\text{prior}}(u) - \log P_{\text{agent}}(u)))^2 \right]}_{\text{Photovoltaic Performance Optimization and Prior-Policy Alignment}} + \underbrace{\zeta_2 \cdot D_{\text{KL}}(\mathcal{N}_{\text{gen}} \parallel \mathcal{N}_{\text{data}})}_{\text{Property Distribution Control}}, \quad (4)$$

Here, the first term aligns the agent with the prior while maximizing the reward $r(u)$, scaled by ζ_1 . The second term, regulated by ζ_2 , minimizes the KL divergence between the distribution of predicted electronic properties (HOMO/LUMO) and the experimental reference distribution from OPV²D.

The reward function $r(u)$ is designed to maximize Power Conversion Efficiency (PCE) within the Shockley-Queisser theoretical limit (33.7%), conditioned on validity:

$$r(u) = \min(\hat{y}_{\text{PCE}}(u) \cdot P(\text{OPV} \mid \mathbf{x}), 33.7) \cdot \mathbb{I}_{\text{valid}}(u). \quad (5)$$

To stabilize training, we employ Experience Replay (Lin 1992), recirculating the top-5% high-performing unique molecules (Tanimoto similarity < 0.7) into training batches.

Experiments

Property Prediction

OPV Likelihood Prediction We trained OPVC on a dataset consists of 1,500 positive examples from the OPV²D dataset and 1,500 negative samples randomly drawn from the PubChem database (Kim et al. 2025), using 8:2 train-test split. The model was configured with 200 estimators, a maximum depth of 15, and class-balanced weighting. It achieved excellent performance, with a test accuracy of 99.16% and F1 score of 0.9916. The AUC reached 0.9985 on the test set, indicating strong generalization and robust classification performance.

HOMO-LUMO Prediction We pre-trained the graph encoder of MOE² using three tasks: molecular mask reconstruction, computed HOMO-LUMO prediction, and literature-reported HOMO-LUMO prediction. For the first two tasks, we used over 51,000 non-fullerene acceptors (NFAs) from the reported dataset (Lopez et al. 2017), and for the third task, we used the proposed OPV²D dataset. A 9:1 train-test split was employed for all tasks. The model achieved 99.997% accuracy on the reconstruction task and R^2 scores of 0.9866 and 0.9784 for the computed and literature-reported HOMO-LUMO prediction tasks, respectively.

To ensure a fair comparison in the baseline experiments, we also pre-trained other graph encoders using the same three tasks and datasets. Three architectures—GCNConv (Kipf 2016), GATv2Conv (Brody, Alon, and Yahav 2021), and GINConv (Xu et al. 2018)—were trained under identical hyperparameter settings.

PCE Prediction After pretraining, we evaluated the learned graph embeddings against traditional Morgan fingerprints on the OPV²D dataset. For each sample, the embeddings were concatenated and used as input to a range of regression models (Breiman 2001; Cortes and Vapnik 1995; Friedman 2001; Chen and Guestrin 2016; Prokhorenkova et al. 2018; Ke et al. 2017; Seeger 2004; Pedregosa et al. 2011; Duan et al. 2020). To predict power conversion efficiency (PCE). We used 5-fold cross-validation and the coefficient of determination (R^2) as the evaluation metric. Baseline results indicated that GAT-based embeddings, which incorporate edge weight information, consistently outperformed traditional approaches.

Building upon these GAT embeddings, we introduced P³, a model that integrates self-attention and cross-attention mechanisms to more effectively capture intra- and inter-molecular interactions. This enhancement led to a significant improvement in performance, achieving a peak R^2 score of 0.716. To ensure fairness and reproducibility, we applied consistent hyperparameter configurations across all models. Performance results are summarized in Table 1.

Molecular Generator Performance

To evaluate the molecular generation capability of MatGPT, we adopted the Molecular Sets (MOSES) (Polykovskiy et al. 2020) benchmarking platform with protocol adaptations for organic photovoltaic (OPV) applications. The

Model	Morgan FP	Graph Embeddings		
		GNN embedding	GAT embedding	GIN embedding
RF	0.691±0.030	0.673±0.038	0.688±0.023	<i>0.670±0.031</i>
SVM	0.577±0.039	0.585±0.040	0.611±0.044	<i>0.575±0.046</i>
GB	0.658±0.036	0.663±0.038	0.667±0.022	<i>0.640±0.030</i>
XGB	<i>0.675±0.036</i>	0.689±0.041	0.700±0.027	0.682±0.030
CatBoost	0.648±0.031	0.648±0.037	0.655±0.037	<i>0.633±0.038</i>
LGBM	0.689±0.038	0.687±0.037	0.703±0.017	<i>0.682±0.033</i>
GP	0.131±0.085	<i>-0.366±0.240</i>	0.301±0.155	-0.110±0.168
HGB	0.686±0.038	0.686±0.038	0.704±0.022	<i>0.681±0.032</i>
NGB	0.568±0.033	0.565±0.040	0.565±0.034	<i>0.535±0.041</i>
P³ (Ours)	–	–	0.716±0.040	–

Table 1: Average R^2 Scores for Models Predicting PCE using Morgan FP and Graph Neural Embeddings. Best values in bold, worst in *italic* (horizontal comparison within each row). Error values are formatted in \pm_{error} .

MOSES metrics were selected based on their ability to address three critical requirements for OPV material discovery: (1) *Validity* to ensure synthesizability of donor/acceptor candidates, (2) *Novelty* to avoid rediscovery of known low-efficiency materials, and (3) *Diversity* (*IntDiv*, *IntDiv2*) to explore broad chemical spaces for synergistic pairings. Besides vanilla GPT-2, we compared MatGPT against baseline models spanning diverse generative paradigms, including VAE (autoencoder-based) (Gómez-Bombarelli et al. 2018), ORGAN (Generative Adversarial Networks) (Guimaraes et al. 2017), a combinatorial generator (template-based) (Polykovskiy et al. 2020), and NGram (rule-based) (Brown et al. 1992).

Experimental Setup We selected PubChem (Kim et al. 2025) for pretraining due to its broader chemical diversity compared to drug-focused datasets like ZINC Clean Leads (Sterling and Irwin 2015) and ChEMBL (Gaulton et al. 2012). Domain-specific preprocessing was applied to tailor the data for OPV materials, filtering SMILES strings to 15–300 characters and excluding ionic or multi-molecule entries. From the cleaned data, 5 million molecules were sampled to train MatGPT-Large, while MatGPT, GPT-2, and other baselines used a 1 million subset. An 8:2 split was used for training and reference sets. MatGPT and GPT-2 were trained for 30 epochs on the Causal Language Modeling task (Radford et al. 2018) using a cosine annealing schedule (initial learning rate: 1e-4); other baselines followed MOSES defaults for fair comparison. Each model generated 30,000 molecules, evaluated with MOSES metrics. MatGPT used top-p sampling ($p = 0.91$) and temperature scaling (0.5–1.4) for diversity. All experiments ran on a single NVIDIA RTX 4090 GPU.

Experimental Results As shown in Figure 4, MatGPT achieves a strong balance between validity (97.4%) and novelty (99.9%) while generating 100% unique molecules, outperforming traditional baselines such as VAE (89.0% va-

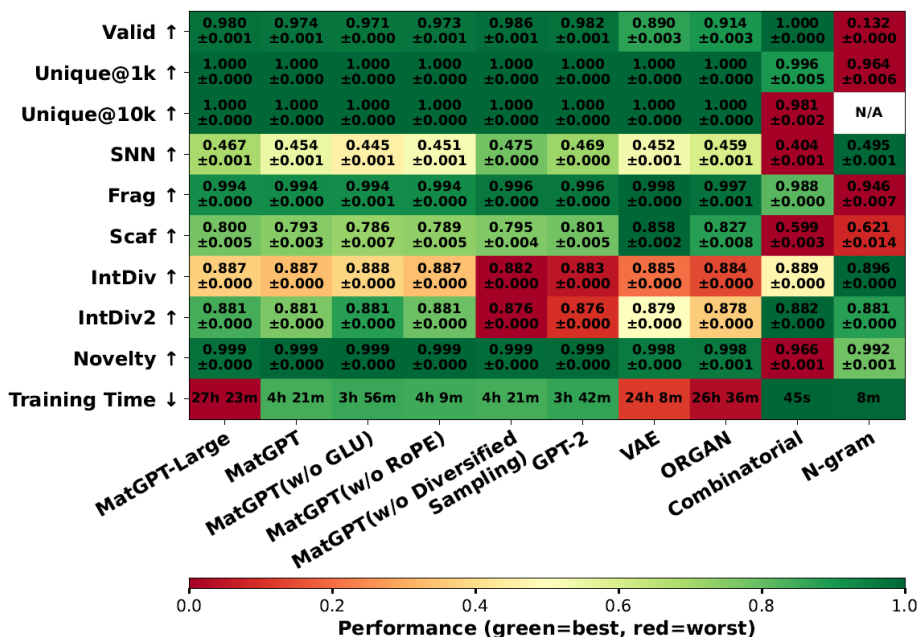


Figure 4: MOSES benchmark performance comparison. Arrows indicate optimal direction (↑=higher better, ↓=lower better). Best values in green, worst in red. Values show mean \pm std. dev. All models were trained on the 1M PubChem dataset, except for MatGPT-Large, which was trained on 5M. N-gram generated only around 3,200 valid molecules (insufficient for Unique@10k).

lidity) and ORGAN (91.4%). Training on a larger dataset, MatGPT-Large further improves validity to 98.0% without compromising diversity.

MatGPT maintains acceptable similarity to the nearest neighbor (SNN = 0.793), indicating that generated molecules stay close to the reference distribution. Its high fragment similarity (Frag = 0.994) suggests balanced representation of chemical fragments, while the scaffold similarity (Scaff = 0.793) reflects good alignment in core structural motifs. Additionally, MatGPT shows relatively high internal diversity (IntDivP) compared to other baselines, demonstrating its ability to generate a chemically diverse set of molecules.

Ablation studies highlight the importance of architectural components: removing RoPE or GLU layers slightly reduces validity, SNN, and Scaff, while disabling diversified sampling increases validity (98.6%) but lowers diversity (IntDiv2 = 0.876).

Designing Donor and Acceptor Molecules

We begin with the pre-trained MatGPT-Large model and fine-tune it on the OPV²D dataset, updating only the gated linear unit (GLU) layers while freezing all other parameters. This is followed by RL-based fine-tuning, where all model parameters remain frozen except for the GLU layers. During RL, we use the trained OPVC to estimate OPV activity, MOE² for predicting HOMO and LUMO energy levels, and P³ for predicting power conversion efficiency (PCE). The loss function includes a reward scaling coefficient $\zeta_1 = 100$. The RL fine-tuning was run for 1,000 steps with a batch size

of 256 and a learning rate of 1×10^{-4} , completing in approximately 3 hours on a single NVIDIA RTX 4090 GPU.

Figure 5 shows the mean top-10 reward in the memory over training steps for various reinforcement learning (RL) configurations. Two search settings are evaluated: fixing PTB7-Th (C49H57FO2S6) as the donor while searching for acceptors, and fixing Y6 (C82H86F4N8O2S5) as the acceptor while searching for donors. Each setting is tested without and with the Property Distribution Control term, controlled by the coefficient $\zeta_2 = 500$, to guide optimization toward OPV-favorable molecules based on OPV²D. With the Property Distribution Control term enabled, both settings consistently achieve higher mean rewards, indicating improved generation of OPV-favorable candidates.

To validate the practical effectiveness of our framework, we selected representative candidates from two design tasks for quantum chemistry analysis. All DFT and TD-DFT calculations were performed using the Gaussian 16 and GaussView 6 software package (Frisch et al. 2016; Dennington, Keith, and Millam 2019). In organic photovoltaics, two factors are critical for high performance: (1) Energy Level Alignment and (2) Spectral Complementarity. Proper alignment of the donor-acceptor frontier molecular orbitals (HOMO/LUMO) is essential to create a driving force for separating light-induced excitons into free charges. Concurrently, the donor and acceptor should absorb different regions of the solar spectrum, working in concert to achieve panchromatic light harvesting and maximize current generation. Our framework demonstrated success in optimizing both factors, as illustrated in Figure 6. The figure displays the calculated

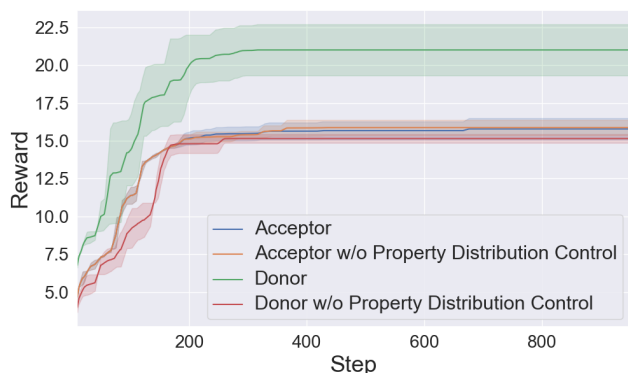


Figure 5: Mean top-10 reward in the memory during RL training.

monomer absorption spectra (Monomer, Calc.) for our representative generated molecules and their estimated solid-state film spectra (Film, Est.). Guided by experimental literature (Yuan et al. 2019; P Christopholi et al. 2024; Li et al. 2021), which typically reports a 100-200 nm redshift from solution to solid-state films due to increased molecular planarity and strong intermolecular $\pi - \pi$ interactions, we selected a representative shift of 150 nm for our estimation. For the PTB7-Th pair, our generated acceptor extends the absorption deep into the near-infrared region (700-900 nm), perfectly complementing the donor’s visible light absorption. Conversely, for the Y6 pair, our generated donor effectively fills the visible spectrum gap (400-650 nm) left by the NIR-absorbing Y6. In addition to this ideal spectral complementarity, quantum chemistry calculations confirmed that both pairs possess well-aligned frontier molecular orbital energies conducive to efficient exciton dissociation. This dual validation of optical and electronic properties underscores our framework’s ability to discover novel, high-potential donor-acceptor partners.

Discussion

The GNN-based model provides interpretable predictions by linking molecular features to photovoltaic performance, with cross-attention highlighting donor-acceptor relationships aligned with known principles. However, real-world discrepancies arise due to fabrication variability, molecular stability, and synthetic feasibility, which are not fully captured by theoretical models. While adding fabrication data or quantum-level accuracy could improve predictions, it would significantly increase complexity. To maintain efficiency, our reinforcement learning framework relies on SMILES representations for practical donor-acceptor screening. Future directions include enhancing model generalization via active learning and domain adaptation, embedding experimental constraints directly into the model, and refining the reward function to better balance competing goals like performance, synthetic accessibility, and sustainability.

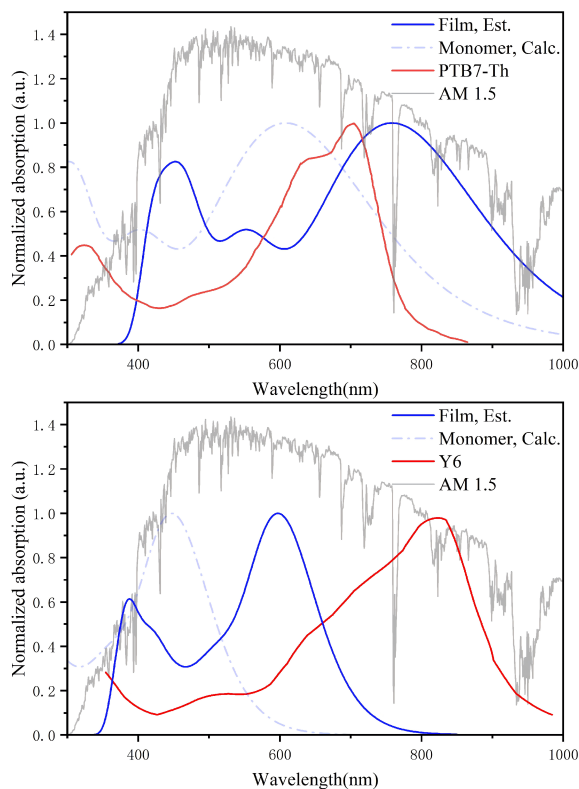


Figure 6: Spectral complementarity of generated OPV molecules. (a) Normalized absorption spectra for the donor PTB7-Th (red) and the corresponding generated acceptor. (b) Spectra for the acceptor Y6 (red) and the corresponding generated donor. In both panels, the calculated monomer absorption (Monomer, Calc.) of the generated molecule is shown alongside its estimated solid-state film spectrum (Film, Est.). The estimated film absorption successfully complements the absorption window of its partner, enabling broad coverage of the AM 1.5 solar spectrum (grey).

Conclusion

This study introduces a dual-pronged machine learning framework combining accurate performance prediction with efficient molecular generation to accelerate OPV material discovery. We constructed the large-scale OPV²D dataset and employed the OPV Classifier (OPVC) for activity screening, alongside a multi-task GNN comprising MOE² and P³ for precise prediction of frontier orbital energies and PCE. In parallel, MatGPT, fine-tuned via reinforcement learning, was used to generate candidate molecules with both high performance and validity. Experimental results demonstrate superior predictive accuracy and molecule quality, offering a powerful tool for rapid OPV screening. Future work will integrate experimental constraints and synthetic conditions to improve generalizability and real-world relevance for sustainable energy development.

Acknowledgments

This work is supported by the National Key R&D Program of China (No. 2022ZD0117501), the Scientific Research Innovation Capability Support Project for Young Faculty (ZYGXQNJSKYCXNLZCXM-E7), the Tsinghua University Initiative Scientific Research Program (Student Academic Research Advancement Program: Zhuiguang Special Project) (Grant No. 20257020001), as well as the Tsinghua University Initiative Scientific Research Program and the Carbon Neutrality and Energy System Transformation (CNEST) Program led by Tsinghua University.

References

- Breiman, L. 2001. Random forests. *Machine learning*, 45(1): 5–32.
- Brody, S.; Alon, U.; and Yahav, E. 2021. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*.
- Brown, P. F.; Della Pietra, V. J.; Desouza, P. V.; Lai, J. C.; and Mercer, R. L. 1992. Class-based n-gram models of natural language. *Computational linguistics*, 18(4): 467–480.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Cortes, C.; and Vapnik, V. 1995. Support-vector networks. *Machine learning*, 20(3): 273–297.
- Das, B.; and Mondal, A. 2024. Predictive Modeling and Design of Organic Solar Cells: A Data-Driven Approach for Material Innovation. *ACS Applied Energy Materials*, 7(20): 9349–9363. Doi: 10.1021/acs.aem.4c01847.
- Dennington, R.; Keith, T. A.; and Millam, J. M. 2019. GaussView Version 6. Semichem Inc. Shawnee Mission KS.
- Duan, T.; Anand, A.; Ding, D. Y.; Thai, K. K.; Basu, S.; Ng, A.; and Schuler, A. 2020. Ngboost: Natural gradient boosting for probabilistic prediction. In *International conference on machine learning*, 2690–2700. PMLR.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189–1232.
- Frisch, M. J.; Trucks, G. W.; Schlegel, H. B.; Scuseria, G. E.; Robb, M. A.; Cheeseman, J. R.; Scalmani, G.; Barone, V.; Petersson, G. A.; Nakatsuji, H.; Li, X.; Caricato, M.; Marenich, A. V.; Bloino, J.; Janesko, B. G.; Gomperts, R.; Mennucci, B.; Hratchian, H. P.; Ortiz, J. V.; Izmaylov, A. F.; Sonnenberg, J. L.; Williams-Young, D.; Ding, F.; Lipparini, F.; Egidi, F.; Goings, J.; Peng, B.; Petrone, A.; Henderson, T.; Ranasinghe, D.; Zakrzewski, V. G.; Gao, J.; Rega, N.; Zheng, G.; Liang, W.; Hada, M.; Ehara, M.; Toyota, K.; Fukuda, R.; Hasegawa, J.; Ishida, M.; Nakajima, T.; Honda, Y.; Kitao, O.; Nakai, H.; Vreven, T.; Throssell, K.; Montgomery, J. A., Jr.; Peralta, J. E.; Ogliaro, F.; Bearpark, M. J.; Heyd, J. J.; Brothers, E. N.; Kudin, K. N.; Staroverov, V. N.; Keith, T. A.; Kobayashi, R.; Normand, J.; Raghavachari, K.; Rendell, A. P.; Burant, J. C.; Iyengar, S. S.; Tomasi, J.; Cossi, M.; Millam, J. M.; Klene, M.; Adamo, C.; Cammi, R.; Ochterski, J. W.; Martin, R. L.; Morokuma, K.; Farkas, O.; Foresman, J. B.; and Fox, D. J. 2016. Gaussian 16 Revision C.01. Gaussian Inc. Wallingford CT.
- Gaulton, A.; Bellis, L. J.; Bento, A. P.; Chambers, J.; Davies, M.; Hersey, A.; Light, Y.; McGlinchey, S.; Michalovich, D.; Al-Lazikani, B.; et al. 2012. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic acids research*, 40(D1): D1100–D1107.
- Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2): 268–276.
- Greenstein, B. L.; and Hutchison, G. R. 2023. Screening Efficient Tandem Organic Solar Cells with Machine Learning and Genetic Algorithms. *The Journal of Physical Chemistry C*, 127(13): 6179–6191.
- Guimaraes, G. L.; Sanchez-Lengeling, B.; Outeiral, C.; Farias, P. L. C.; and Aspuru-Guzik, A. 2017. Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843*.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; and Liu, T.-Y. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Khatua, R.; Das, B.; and Mondal, A. 2024. Physics-Informed Machine Learning with Data-Driven Equations for Predicting Organic Solar Cell Performance. *ACS Applied Materials & Interfaces*, 16(42): 57467–57480. Doi: 10.1021/acsami.4c10868.
- Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; et al. 2025. PubChem 2025 update. *Nucleic acids research*, 53(D1): D1516–D1525.
- Kipf, T. 2016. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*.
- Li, C.; Yue, Q.; Wu, H.; Li, B.; Fan, H.; and Zhu, X. 2021. Small bandgap non-fullerene acceptor enables efficient PTB7-Th solar cell with near 0 eV HOMO offset. *Journal of Energy Chemistry*, 52: 60–66.
- Lin, L.-J. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3): 293–321.
- Lopez, S. A.; Sanchez-Lengeling, B.; de Goes Soares, J.; and Aspuru-Guzik, A. 2017. Design Principles and Top Non-Fullerene Acceptor Candidates for Organic Photovoltaics. *Joule*, 1(4): 857–870.
- Miyake, Y.; and Saeki, A. 2021. Machine Learning-Assisted Development of Organic Solar Cell Materials: Issues, Analyses, and Outlooks. *The Journal of Physical Chemistry Letters*, 12(51): 12391–12401.
- P Christopholi, L.; Marchiori, C. F. N.; Jalan, I.; Opitz, A.; Muntean, S. A.; and Moons, E. 2024. Role of the Solvent on the Orientation of Y-Type Acceptor Molecules in Spin-Coated Films. *The Journal of Physical Chemistry C*, 128(42): 17825–17835.

- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12: 2825–2830.
- Polykovskiy, D.; Zhebrak, A.; Sanchez-Lengeling, B.; Golo- vanov, S.; Tatanov, O.; Belyaev, S.; Kurbanov, R.; Artamonov, A.; Aladinskiy, V.; Veselov, M.; et al. 2020. Molecular sets (MOSES): a benchmarking platform for molecular generation models. *Frontiers in pharmacology*, 11: 565644.
- Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A. V.; and Gulin, A. 2018. CatBoost: unbiased boosting with categorical features. *Advances in neural information processing systems*, 31.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- RDKit, online. 2013. RDKit: Open-source cheminformatics. <http://www.rdkit.org>. [Online; accessed 11-April-2013].
- Seeger, M. 2004. Gaussian processes for machine learning. *International journal of neural systems*, 14(02): 69–106.
- Seifrid, M.; Lo, S.; Choi, D. G.; Tom, G.; Le, M. L.; Li, K.; Sankar, R.; Vuong, H.-T.; Wakidi, H.; Yi, A.; Zhu, Z.; Schopp, N.; Peng, A.; Luginbuhl, B. R.; Nguyen, T.-Q.; and Aspuru-Guzik, A. 2024. Beyond molecular structure: critically assessing machine learning for designing organic photovoltaic materials and devices. *Journal of Materials Chemistry A*, 12(24): 14540–14558.
- Sterling, T.; and Irwin, J. J. 2015. ZINC 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11): 2324–2337.
- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.
- Sun, J.; Li, D.; Zou, J.; Zhu, S.; Xu, C.; Zou, Y.; Zhang, Z.; and Lu, H. 2024. Accelerating the discovery of acceptor materials for organic solar cells by deep learning. *npj Computational Materials*, 10(1): 181.
- Weininger, D. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1): 31–36.
- Wu, Y.; Guo, J.; Sun, R.; and Min, J. 2020. Machine learning for accelerating the discovery of high-performance donor/acceptor pairs in non-fullerene organic solar cells. *npj Computational Materials*, 6(1): 120.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yuan, J.; Zhang, Y.; Zhou, L.; Zhang, G.; Yip, H.-L.; Lau, T.-K.; Lu, X.; Zhu, C.; Peng, H.; Johnson, P. A.; Leclerc, M.; Cao, Y.; Ulanski, J.; Li, Y.; and Zou, Y. 2019. Single-Junction Organic Solar Cell with over 15% Efficiency Using Fused-Ring Acceptor with Electron-Deficient Core. *Joule*, 3(4): 1140–1151.
- Zhang, S.; Li, S.; Song, S.; Zhao, Y.; Gao, L.; Chen, H.; Li, H.; and Lin, J. 2025a. Deep Learning-Assisted Design of Novel Donor–Acceptor Combinations for Organic Photo-voltaic Materials with Enhanced Efficiency. *Advanced Materials*, 37(4): 2407613.
- Zhang, W.; Zou, Y.; Wang, X.; Chen, J.; and Xu, D. 2025b. Deep learning accelerated high-throughput screening of organic solar cells. *Journal of Materials Chemistry C*, 13(10): 5295–5306.