Small Methods

small methods
www.small-methods.com

RESEARCH ARTICLE OPEN ACCESS

Minimal Input, Maximum Insight: Inferring Material Parameters From Solar Cell JV Curves Alone

Cai Williams[1] | Chen Wang[1] | Alexander Ehm[1] | Dietrich R. T. Zahn[1] | Maria Saladina[1] | Carsten Deibel[1] | Roderick C. I. Mackenzie[2]

[1]Institut für Physik, Technische Universität Chemnitz, Chemnitz, Germany | [2]Department of Engineering, Durham University, Durham, UK

Correspondence: Roderick C. I. Mackenzie (roderick.mackenzie@durham.ac.uk)

Received: 10 February 2026 | **Revised:** 15 July 2026 | **Accepted:** 20 July 2026

Keywords: artificial neural network | carrier lifetime | computer science | electron mobility | fabrication | organic solar cell | photovoltaic system | random walk | solar cell | transistor

ABSTRACT

A key challenge in the development of materials for the next generation of solar cells, sensors and transistors is linking macroscopic device performance to underlying microscopic properties. For years, fabrication of devices has been faster than our ability to characterize them. This has led to a random walk of material development, with new materials being proposed faster than our understanding. We present two neural network-based methods for extracting key material parameters, including charge carrier mobility and trap state density, in optoelectronic devices such as solar cells. Our methods require only a single measured light current-voltage curve and modest computational resources, making our approach applicable in even minimally equipped laboratories. Unlike traditional machine learning models, our methods place the final material values in a non-Gaussian likelihood distribution, allowing confidence assessment of each predicted parameter.We demonstrate these techniques on freshly fabricated PM6:Y12 and PM6:BTP-eC9 organic solar cells, and then track a single PM6:BTP-eC9 device as it degrades in air, recovering the evolution of carrier lifetime, mobility and shunt resistance. This approach enables rapid, low-cost extraction of key material parameters from simple JV measurements alone, providing a practical route to accelerate optimisation of next-generation solar-energy materials and devices.

1 | Introduction

In recent years, there has been an explosive growth in the development of novel materials for optoelectronic devices, including organic polymers [1–3], small molecules [4], and perovskites [5–7]. Today, thin film analogues based on disordered materials exist for nearly all conventional inorganic devices, including organic LEDs [8], organic [9–11] solar cells, perovskite solar cells [12–14], transistors [15, 16], and sensors [17, 18]. Some of these technologies, most notably OLED displays, have seen significant commercial success. Nevertheless, progress across the field has been relatively slow. For instance, more than a decade after their modern rediscovery [19], perovskite solar cells have yet to achieve broad market adoption, and organic photovoltaics, despite their promise, remain largely confined to the lab. In order to see commercial success the learning rate of such systems must be increased. Accelerating this learning rate is essential for advancing renewable-energy technologies, where device optimisation cycles remain slow and costly.

If one takes the perovskite material system as an example, a cursory search on any on-line archive will reveal tens of thousands of articles published over the last 10 years, each reporting some improvement over the state-of-the-art. However, in most cases, characterization is limited to reporting of the overall device efficiencies (PCEs) based on current density–voltage (JV) curves,

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.
© 2026 The Author(s). *Small* Methods published by Wiley-VCH GmbH

yielding the solar cell parameters short-circuit current density (J_{sc}), fill factor (FF) and open-circuit voltage (V_{oc}). Often, this is supplemented by basic measurements such as external quantum efficiency or absorption. However, researchers rarely employ and analyse techniques that truly clarify why changes in fabrication conditions or formulation/molecular structure affected device performance in terms of microscopic electrical parameters, such as trap densities, Urbach tail slopes, charge carrier mobilities, and recombination rates. Although many microscopic factors influence device performance, their effects ultimately manifest in macroscopic transport parameters such as charge carrier mobility and recombination lifetime, whose product determines the carrier diffusion length and strongly influences charge collection. Within organic photovoltaics, generally speaking, when a chemist changes the length of an aliphatic side chain on a novel acceptor molecule, it is common to report the change in efficiency, but not a detailed analysis as to what this did to the charge carrier dynamics and the core reasons for changes in efficiencies. This lack of closed loop feedback has turned device and material optimization into a random walk that spans decades. Consequently, it can take tens of years to turn a promising new material system into a commercially viable product; this is time we do not have in the context of the CO_2/energy crisis [20].

Techniques capable of probing complex physical processes, such as charge transport and recombination, remain underutilized due to their time-intensive nature and reliance on specialized equipment. To use these techniques, one often needs high speed pulsed lasers, fast electronics, and dedicated dark-room facilities. Consequently, access to these methods is often confined to a small number of well-resourced laboratories. Furthermore, expertise is compartmentalized across subdisciplines of materials science. For instance, techniques for analysing charge carrier dynamics are more likely to be found within condensed matter physics or physical chemistry groups, whereas the synthesis and processing of novel materials typically occur in synthetic chemistry laboratories. While cross-disciplinary collaboration is common in the field, this separation of skills and infrastructure makes it impractical to apply comprehensive characterization to every device that is fabricated. The rate at which new materials and devices are fabricated continues to accelerate, outpacing the ability to meaningfully characterize them. Devices that underperform are frequently discarded without any detailed analysis, leaving the reasons for failure unknown.

One way to address this challenge is to apply modern machine learning (ML) techniques [21–25], with the aim of extracting key device parameters that are difficult to extract, such as mobilities and trap densities, from rudimentary measurement techniques. In this work, we combine state-of-the-art ML techniques based on residual blocks [22] which are used in most modern Large Language Models (LLMs) [26–28], with synthetic data sets and experimental results. We demonstrate a ML framework that is capable of extracting key material parameters from single standard light JV measurement and, most importantly, places reasonable bounds on the accuracy of the results, without the need for frequency-domain characterisation.

It is the overarching aim of this paper to develop a method that can extract as much information as possible from a functioning device using the simplest possible measurements. We want to enable any laboratory that fabricates devices to be able to gain detailed information on their devices quickly and at low cost. To achieve this, we will focus on small datasets derived from minimal, low-complexity simulations and experiments. We prioritize steady-state measurements and simulations, deliberately avoiding frequency and time-domain techniques due to their higher experimental and computational complexity. Throughout this work, the more demanding techniques of IMPS, IMVS and SCLC are used exclusively to validate the machine-learning predictions against independently measured values. They are not required in order to apply the method, and are precisely the class of measurement that our approach is intended to make unnecessary.

ML models are often criticized for producing predictions without clear justification, therefore incorporating a measure of confidence is essential. The method presented will quantify uncertainty in the extracted values. We introduce two key innovations: First, instead of predicting absolute values of the parameter of interest, the model is trained to learn the difference between a well-characterized reference device and a device under test. Second, this relative prediction enables a fully probabilistic output, where material parameters are represented as likelihood distributions. The form of the distribution is not assumed (unlike Bayesian based methods). We validate our approach on a series of PM6:BTP-eC9 and PM6:Y12 organic solar cells, including both freshly fabricated and aged devices, demonstrating that the method captures performance evolution over time. The acceptor molecules chosen represent the state-of-the-art for both chlorinated and fluorinated organic acceptor molecules. While this study focuses on organic photovoltaics, the approach is broadly applicable to a wide range of semiconductor systems and other classes of devices, offering a simple and accessible pathway toward interpretable, quantitative characterization. There is nothing preventing the proposed method from being applied to other material systems and device structures.

2 | Experimental

Devices were fabricated on pre-cleaned patterned ITO substrates with a 37 nm ZnO nanoparticle layer. PM6:BTP-eC9 and PM6:Y12 active layers (1:1.2 donor:acceptor ratio, 12 mg/mL in chloroform) were spin-coated at 45°C and annealed at 100°C, yielding 50–60 nm thick films. An ethanol-based PEDOT:F layer [29] was deposited on top, followed by thermal evaporation of 150 nm thick Ag electrode to define a 0.04 cm^2 active area. All fabrication steps beyond ZnO deposition were conducted in a nitrogen glove box. See the supplementary information for further detail. The overall device structure was ITO/ZnO/active layer/PEDOT:F/Ag. This architecture was chosen because it provides stable devices and is compatible with scalable fabrication approaches such as printed organic photovoltaics. The device structure, and chemical structure of the donor and acceptors used can be seen in Figure 1. Both devices were measured in the as-fabricated (pristine) state. The PM6:BTP-eC9 device shown here is the same device that is subsequently followed through the degradation study of Section VII, where its pristine state corresponds to the first measurement point.

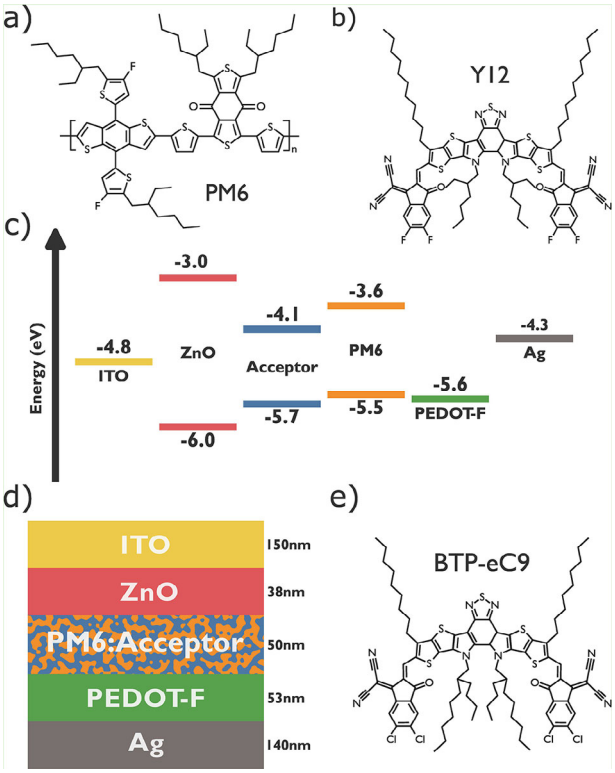


FIGURE 1 | (a) The molecular structure of organic donor PM6. (b) The molecular structure of organic acceptor Y12. (c) The energy level alignment diagram referenced to vacuum of both devices. (d) The layer structure and thicknesses of the device within simulation. (e) The molecular structure of organic acceptor BTP-eC9.

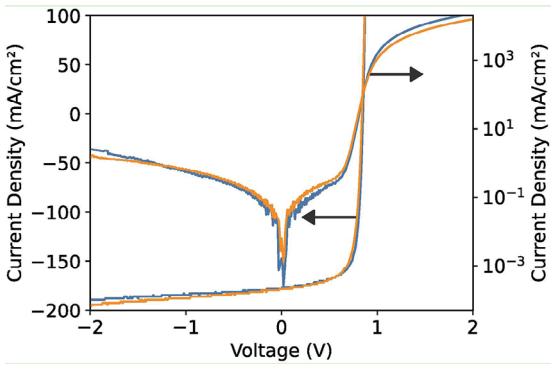


FIGURE 2 | As-fabricated light and dark JV curves for the PM6:Y12 (blue), and the PM6:BTP-eC9 (orange) device. The Light JV curves are indicated to the left, and dark JV curves indicated to the right.

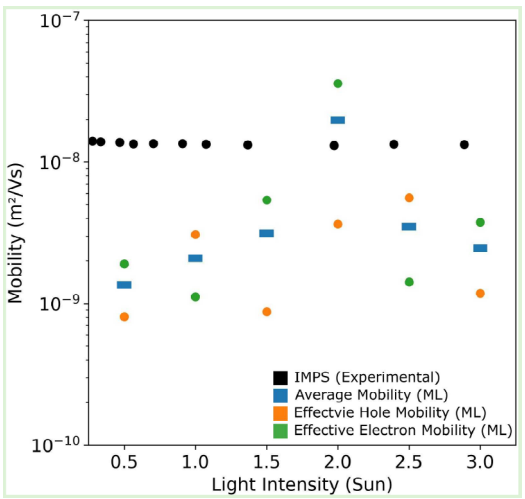


FIGURE 3 | Experimental and predicted charge carrier mobilities. The black dots represent the values measured by IMPS, while other dots show ML-predicted mobilities for electrons, holes and their mean. The predicted mobilities are within an order of magnitude of the experimental values and remain largely independent of light intensity. The model struggles to differentiate between electron and hole mobilities, likely due to their close values. Note that the predicted values represent effective mobilities, averaged over both free and trapped charge carriers.

For prediction, JV measurements were performed under AM1.5 illumination from a Wavelabs LS-2 solar simulator or Omicron A350 laser (515 nm), with experiments conducted in nitrogen atmosphere. For validation, Intensity-modulated photovoltage (IMVS) and photocurrent (IMPS) spectroscopy used modulated laser illumination and a Zurich MFLI lock-in amplifier, maintaining small-signal excitation with 10% modulation depth. Space charge limited current (SCLC) devices were fabricated and mobility was extracted from hole- and electron-only devices using the Mott-Gurney law, with tailored layer stacks for selective carrier injection. Full details of the characterization techniques can be found in the Supporting Information. It should be noted that IMPS, IMVS and SCLC are all relatively complex measurements not available in all labs. Therefore, they are only used in this article for validation of the ML method and are not required to apply the ML method presented later.

Figure 2 shows the light and dark JV curves for the PM6:Y12 and PM6:BTP-eC9 device. The PM6:Y12 and PM6:BTP-eC9 devices had a PCE of 11.0% and 10.7% respectively, this difference is due to the slightly higher V_{oc} of the PM6:Y12 device. The PM6:Y12 device was further characterised using IMPS.

3 | Dataset Generation

Modern neural networks require a significant amount of training data to produce reliable predictions, for most tasks they need to see between 10 000 and 100 000 examples. In practical terms, if one were trying to differentiate between cat and dog pictures, one would need a set of images already labelled by a human with the words cat or dog to train the network. In our case, producing an experimental data set on such a magnitude would be very challenging, as we would need to perform complex experiments on a large number of devices. Furthermore, many experimental techniques used to measure microscopic device quantities, such as trap density and mobility, often struggle to produce accurate values due to the variation of these parameters with charge carrier density [30, 31]. Thus, were we to attempt to generate such a data set based on experimental data, we would risk learning incorrect values. For these two reasons we used a synthetic data set generated by our drift–diffusion model that includes a distribution of trap states. It should be highlighted here that, because we are modelling disordered semiconductors, it is essential to use a model that includes carrier transport; carrier

trapping, and recombination through trap states [32]. If one were to attempt to use a model without trap states, one would risk getting the carrier density in the device wrong, and thus the recombination rates/mobility values. This is discussed at length in the Supporting Information.

To generate the data set, we first built the device structure of interest in our drift–diffusion simulation OghmaNano [33, 34]. We inputted the device geometry and also imported experimental complex refractive index (n, k) data for each layer (listed and explained in Supporting Information: Section S1). Using this base device simulation, a set of 1 000–10 000 child simulations with an identical structure were made. For each of these devices the electrical parameters such as free carrier mobilities, trap densities, Urbach tail slopes, capture cross sections, shunt and series resistances were randomly varied within defined bounds. A log scale was used where appropriate. Each of these simulations was then run to produce a light JV curve. This left us with a data set of light JV curves, and the exact electrical parameters used to produce them. Around 20% of this data was held back for validation. We also removed devices with an efficiency below 8% percent as devices of low performance otherwise become dominant within the datasets increasing the difficulty of training. This process could be completed in a couple of hours on a modern PC. If the method is to be applied to devices with a different geometry or thickness, a new training dataset corresponding to the experimental structure can simply be generated. Alternatively, datasets containing multiple device thicknesses can be used, with the thickness provided as an additional input to the neural network. We have tested this approach previously and found that it works well, although it requires a larger dataset and increases model complexity.

4 | A Simple Neural Network for Point Prediction

We first use a deep neural network in the most simple and classical way that we can [35, 36], to baseline more advanced methods described later. This can be seen in Figure 4a, with a JV curve on the left hand side of the image and a predicted value on the output. The neural network, once trained, acts as a transform between the input data and the predicted output. In this case, the network is extracting mobility from a JV curve, but the parameter could just as well be any material property, such as trap density. The actual size of the network is much larger than depicted in the figure and is detailed in the Supporting Information. Each node of the network is made up of an artificial neuron, schematically shown in Figure 4b. Each neuron takes a series of inputs, first multiplying them by the weight w_n. The sum of these values is then applied to activation function allowing the output of the neuron to be non-linear, and when combined together universal function approximators. In this work, we use the Sigmoid Linear Unit (SiLU) activation function [23], as it is widely used in large models such as LLMs. Initially, the network is assigned a set of random weights. It is then trained by presenting known input data (e.g., a JV curve) and comparing the network's output (e.g., mobility) to the known target value used to generate that input. The difference between the predicted and true values defines an error function, which is used to update the weights through backpropagation [21, 37]. Repeating this process thousands of times over the training

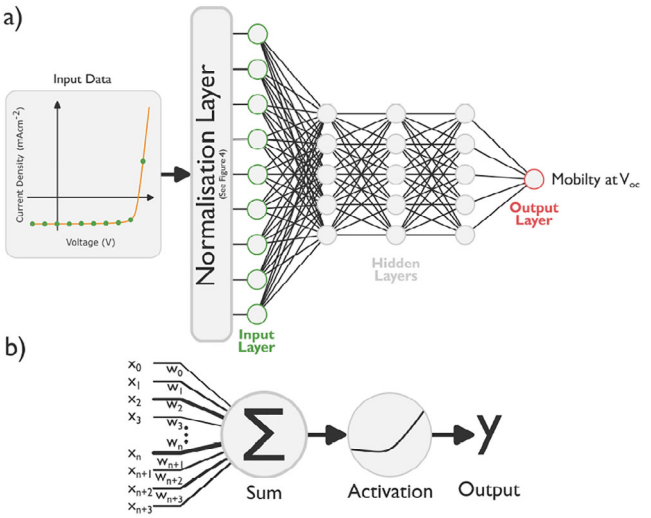


FIGURE 4 | (a) A neural network consisting of an input layer, hidden layers and an output layer. The JV curve of interest is placed on the input by discretizing it at the points indicated by the green dots. The network then predicts the mobility that was used to generate the JV curve. (b) An expanded view of a single neuron, where inputs are weighted, summed, and passed through an activation function that determines the neuron's output.

dataset yields a network capable of predicting the desired material parameter.

Key to having a well performing network that takes note of all parts of the data is to correctly normalize the input data. If one considers the current density of a dark JV curve (e.g., Figure 2), the values at 0.1 V are up to 4 orders of magnitude smaller than those at around 1.5 V. Thus, feeding such a curve into a neural network would cause the high-current values to dominate, while lower-current values would be effectively ignored. To mitigate this imbalance, the data is normalized, as shown in Figure 5.

The input data can be seen in Figure 5a,b the green dots highlight specific voltages at which the current density values are sampled from the JV curve, while red/blue bars indicate the maximum and minimum range across the whole dataset at these voltages. Figure 5c presents the final normalized input to the neural network. Note that normalization is performed relative to all data points in the dataset at a given voltage; not relative to the rest of the curve. This ensures that all input values fall between 0.0 and 1.0. The same normalization approach is applied to the output data, and logarithmic scaling is used where appropriate.

Once trained, the network can be tested on the 20% of the data set that was not used for training. Figure 6 shows the results from these tests plotted as so called confusion matrices. Here we show confusion matrices for the mobility at the maximum power point μ_{pmax}, charge carrier lifetime at open-circuit voltage τ_{voc}, and electron Urbach energy E_U^e. To generate these graphs, the normalized values for each JV curve are placed on the inputs of the network, the true value of mobility (or other parameter) used to generate the curve are placed on the x-axis of the plot, and the predicted value from the network is then placed on the y-axis. For a perfect prediction, the plot should be a straight line.

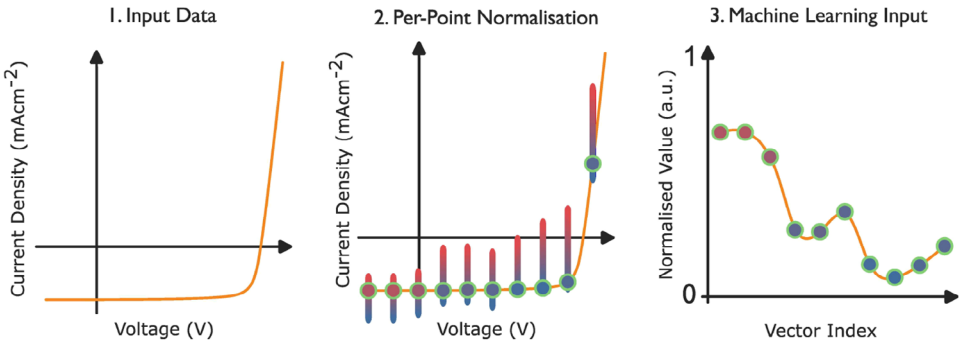


FIGURE 5 | Normalisation process of input data. (a) Experimental current–voltage data. (b) Green dots represent voltages at which the current values are extracted, with red/blue bars showing the maximum/minimum range of the values across the entire dataset at each corresponding voltage. (c) The normalized input to the neural network. Note the points are normalized relative to all data points in the dataset at a given voltage, not relative to the rest of the curve.

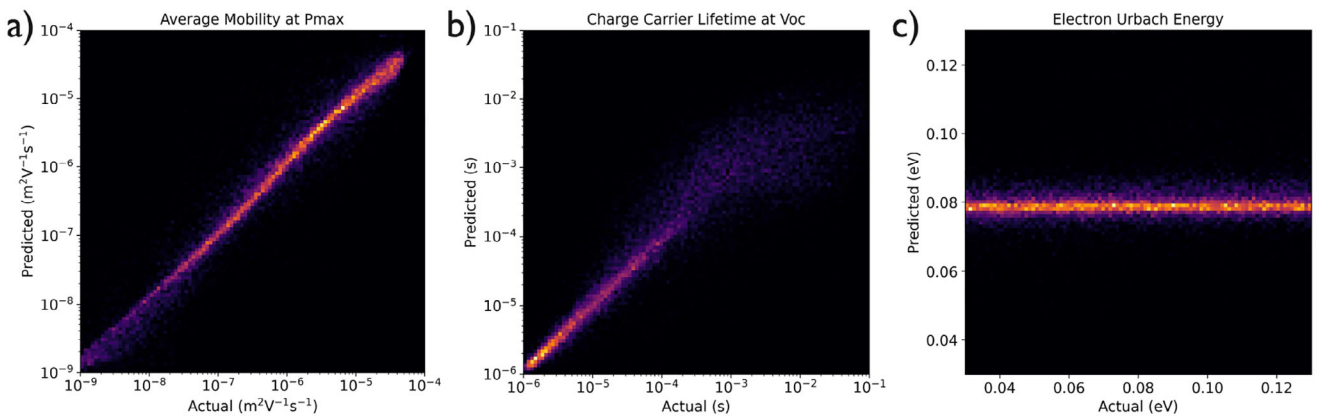


FIGURE 6 | Confusion matrices for (a) average carrier mobility at the maximum power point μ_{pmax}, (b) charge carrier lifetime at open-circuit voltage τ_{voc}, and (c) electron Urbach energy E_U^e. Average carrier mobility at the maximum power point is straightforward to predict. Charge carrier lifetime at V_{oc} is also reliably predicted, provided sufficient training data is available. However, Urbach tail slope is very difficult for a simple feed forward neural network to extract, simply because there is not much information about the Urbach tail slope in a single light JV curve.

The fewer off diagonal elements that exist, the better the ML system performs.

The three parameters presented here were chosen because they display very different learning properties (plots of other parameters are available in the SI). From Figure 6a, it is evident that the model accurately predicts the average charge carrier mobility at the maximum power point (μ_{pmax}), as all points fall on or near the diagonal. The model also performs well in predicting τ^e_{pmax} for values well below 1×10^{-3} s, but struggles above this threshold due to insufficient training data: very few devices in the training set have recombination lifetimes as long as 1 ms, which would correspond to exceptionally, and perhaps unrealistically, high-performance devices. In contrast, the Urbach energy for electrons (E_U^e) cannot be predicted by the model at this stage, simply because there is not a lot of information about the tail slope embedded within the IV quadrant of the JV curve and what information exists in the curve is hard to extract. Later we discuss a more sensitive method, that shows more success.

Once the network was properly trained and tested, we applied it to the experimental JV curves from Figure 2, extracting electron and hole mobilities as a function of light intensity. The resulting average mobility values, shown in Figure 3, are in close agreement with those obtained via IMPS. Notably, all values were inferred from a single light JV curve measured at 1 Sun.

However, we cannot assign error bars to these values. The best we can do is look at the confusion matrix to assess how well the network predicts μ_{Jsc} and thereby determine whether the results are trustworthy. A closer look at the graph reveals that the predicted electron and hole mobilities alternate, with electron mobility sometimes exceeding hole mobility and vice versa. This indicates that the model struggles to consistently distinguish between the two charge carrier species, likely because their mobilities are close in magnitude. Since the network only provides point predictions, no further information can be extracted at this stage. The issue of trust and uncertainty in these predictions is addressed in the following two sections.

5 | The Differential Objective Method

The method described in the previous section provides estimates of material parameters, and when combined with the confusion matrices, it offers a reasonable indication of how much

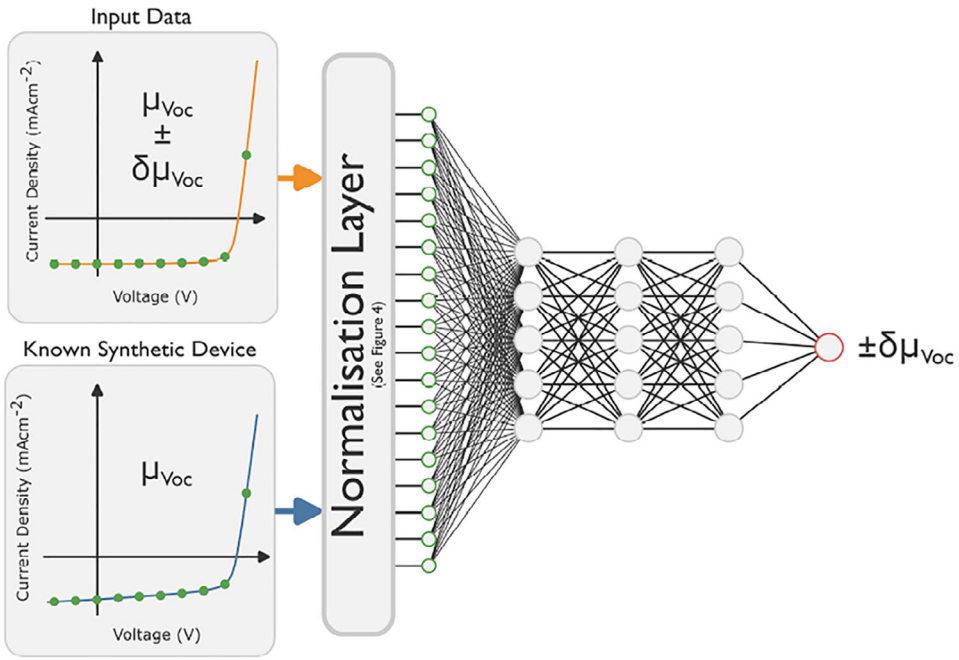


FIGURE 7 | Application of the difference method: The network receives a "known curve" with a known parameter (e.g., μ_{voc}) and an "unknown curve." Instead of predicting an absolute value, it estimates how much higher or lower the parameter value is in the unknown curve relative to the known one, with an output $\pm\delta\mu_{voc}$. Repeating this over many known curves yields a distribution of parameter values predicted by the network.

confidence one can place in each prediction. However, challenges arise when an experimental input lies outside the distribution of the training data, for instance, an s-shaped JV curve [38] may produce predictions with no meaningful confidence, as it deviates significantly from what the model has learned. Here, we introduce a novel approach that outputs a probability distribution representing the confidence in each predicted parameter. Importantly, unlike Bayesian methods, our technique does not require assuming a prior distribution for the calculation.

The method is depicted in Figure 7. Instead of taking a single JV curve and asking the model to predict a single parameter, we present the network with two JV curves. One curve comes from the simulated dataset, where the parameters used to generate it are known (μ_{voc} in the example shown in Figure 7), and is referred to as the "known curve." The other curve, for which the generating parameters are unknown, is referred to as the "unknown curve." Rather than asking the model to predict a single value of mobility (or another parameter), we ask it to estimate how much higher or lower it believes the mobility in the "unknown curve" is relative to the "known curve," that is, $\pm\delta\mu_{voc}$. By adding $\pm\delta\mu_{voc}$ to μ_{voc}, we obtain the estimated mobility for the "unknown curve." Repeating this process over many known curves allows us to obtain a distribution of predictions, and thus we can quantify how reliable the prediction is.

This is shown in Figure 8, where we present histograms indicating where the model thinks the most likely material parameters lie for the experimental PM6:Y12 device. Four histograms are shown: average electron/hole mobility at V_{oc}, charge carrier lifetime at V_{oc}, electron trap density, and average electron/hole mobility at the maximum power point. For this device, the mobility at V_{oc} can be predicted with reasonable confidence within two orders of magnitude. Experimental SCLC mobility values are also included as vertical lines; orange for hole mobility, green for electron mobility, and black for the average. The distribution centres around 1×10^{-7} m^2 V^{-1} s^{-1}, which compares well to the experimental SCLC values. As always, however it should be noted that experimental values of mobility in disordered materials are always difficult to accurately measure [39]. With that in mind, the predictions and measurements show very good agreement. In contrast to the charge carrier mobility, the distribution for charge carrier lifetime in Figure 8b is broader, but has a peak near 5×10^{-7} s. That value coincides with the lifetime measured by IMVS, which is 1.3×10^{-6} s under 1 Sun illumination. However, again it is worth noting that in disordered systems different measurement techniques often extract slightly different values for constants due to the dependence of most parameters on charge density.

Figure 8c demonstrates that the electron trap density is, as expected, more difficult to extract as shown by the broad distribution generated. Nevertheless we can state that the value probably lies around 1×10^{28} m^{-3}, which is the peak of the distribution. The average electron/hole mobility at the maximum power point is shown in Figure 8d, and we can say with confidence that it is very likely to be 1×10^{-8} m^2 V^{-1} s^{-1} due to the sharp peak in the probability distribution. The SCLC lines are also plotted for comparison. It is interesting to note, when comparing Figure 8a,d, that the distribution of mobilities in the latter is shifted to lower values than in the former. This is expected, as the carrier density is higher at V_{oc}, resulting in a higher average charge carrier mobility.

The method described above has three distinct advantages: Firstly, one can visually assess where the value of interest is most likely to lie. Secondly, our method does not assume a prior distribution, unlike many approaches based on Bayesian

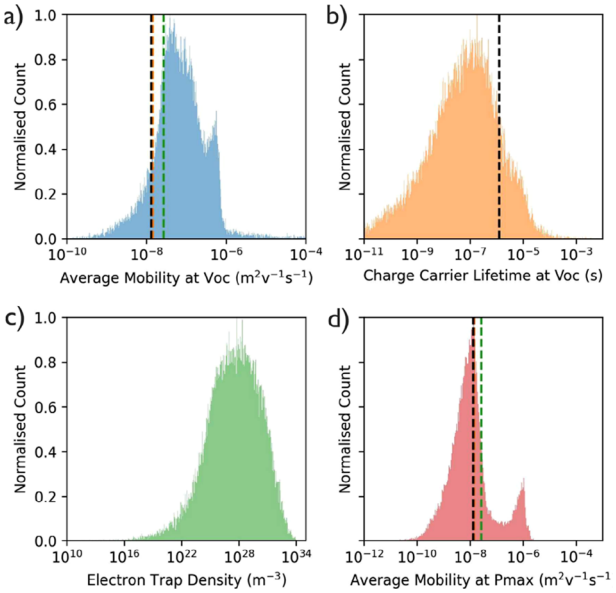


FIGURE 8 | Difference network predictions for: (a) Average Mobility at V_{oc}, (b) charge carrier lifetime at V_{oc}, (c) electron trap density, and (d) average mobility at the maximum power point. Some parameters can be predicted more accurately than others, each showing a unique, non-Gaussian distribution. In panels a and c, dashed lines indicate experimentally determined mobility values measurements: SCLC electron-only device (green), SCLC hole-only device (orange), and IMPS (black). The black line in panel b represents IMVS measured charge carrier recombination time.

statistics; for example, a Gaussian distribution would not capture the shapes shown in Figure 8a,d. Thirdly, when it comes to training, it can dramatically reduce the required dataset size. In a dataset of 1 000 curves, we can train against all permutations of the curves, yielding an effective dataset size of 999 000. This not only decreases data generation time but also boosts the quality of the trained network.

6 | The Differential Objective Method With Residual Blocks

In deep neural networks such as ours, training becomes more difficult as the network depth increases. This is because adjusting a weight in the centre of the network often has little impact on the quality of the prediction, making it hard to determine which way to tune a weight [22, 40]. The overall effect of this to limit how deep a network can be before additional layers do not help with the task at hand. In the previous example we found that having more than four hidden layers did not add much to the predictive capabilities of the model. Another issue with deep networks is that neurons deeper in the network often lose sight of the input data and are only working on a transformed version of it. As a result, they may struggle to extract meaningful information. A solution to both these issues, proposed by He et al. [22], is the residual block, shown in Figure 9b. The core idea of this block is that some fraction of the input data always bypasses the neurons and is routed directly toward the output. This allows deeper neurons to retain access to the information available at shallower layers of the network. This strategy is used to great effect by many modern LLMs [26–28]. Figure 9b presents the structure of our difference network with residual blocks included. The residual block allows us to use a network that is much deeper. For comparison, in the pervious section our network was 4 layers deep with 226 049 trainable parameters, while the network in this section was 32 layers deep with 1 007 233 trainable parameters. It should be noted that this is still small in comparison with most LLMs, which typically have have 1–1.8 trillion weights.

Results from the residual block approach are shown in Figure 10. Overall, we found that residual networks are better at extracting hard to learn parameters, such as Urbach energy, trap density and also energetically averaged charge carrier mobility but due to their larger size, longer training times were needed. Generally in terms of RMS error, residual block networks outperformed the networks in the previous section by a factor of between 2 and 1000, this can be seen in Table S5.

Figure 10a,b displays the energetically averaged electron and hole mobility at J_{sc}. The distribution of predictions for electron mobility is very sharp, indicating strong predictive performance, while the distribution for holes is more square-shaped, suggesting greater difficulty in predicting the value. This ties in well with the results in Figure 3, where the model struggles to understand whether electron or hole mobility is higher. It is also worth noting that both distributions fall within roughly the same range, which is supported by previous work [41]. As in the previous figures experimental lines are also shown. It should be emphasized that disentangling electron/hole mobility, rather than stating an average value is a difficult task even a well equipped lab.

In the previous section we showed that trap density (see Figure 8c) was hard to extract and the network predicted a broad distribution of $1 \times 10^{16} - 1 \times 10^{34}$ m^{-3}. This distribution has been recalculated in Figure 10c,d for both electrons and holes. It can be seen that the distributions have narrowed, with peaks at approximately 1×10^{26} m^{-3}. These are reasonable parameters when one considers that each molecule or polymer subunit capable of trapping charge might occupy a volume of $V = (1 \times 10^{-9})^3$ m^{-3}, and this would yield an expected trap density of around 1×10^{27} m^{-3} - were all sites efficiently connected to conducting pathways.

Figure 10e plots a confusion matrix of predicted and extracted Urbach energy. If one compares this to the previous confusion matrix shown in Figure 6c, it can be seen that there is a considerable improvement, with the network being able to identify the $x = y$ trend. Although, the result is far from perfect, it should be taken in the context that in a single light JV curve there is very limited information about Urbach energy. Indeed to-date there are no reliable analytical methods, that have been proposed to extract Urbach energy from a single light JV curve.

7 | Application to Degradation

We now apply the method developed above to a degradation study of PM6:BTP-eC9 device. Degradation studies provide a good testing ground for the method because the solar cell changes over time and there is no simple way to pause the process to perform complex charge carrier measurements. The

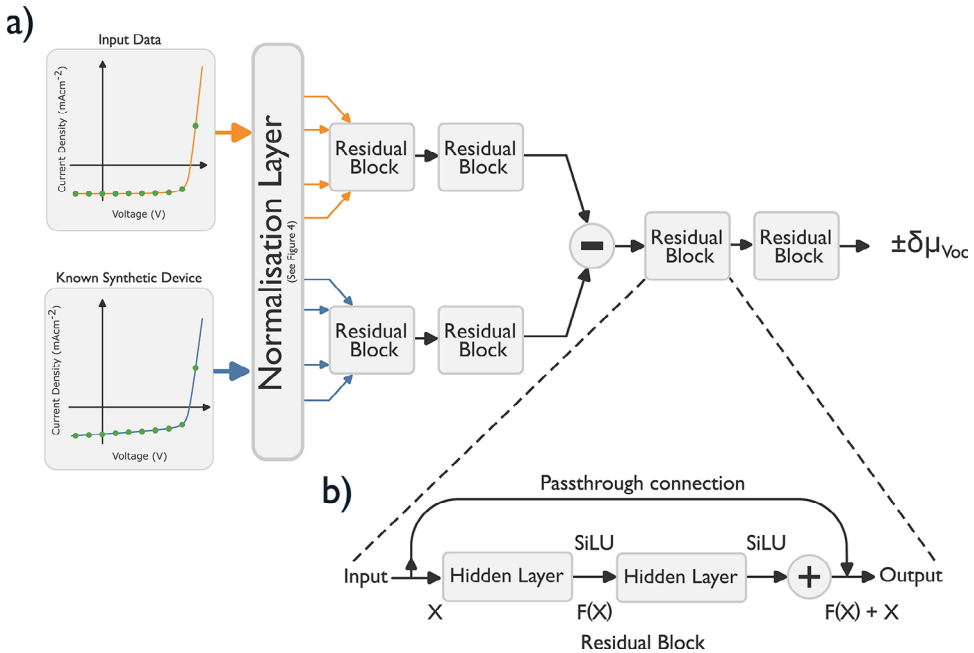


FIGURE 9 | Implementation of the difference method with residual blocks. (a) Schematic of the implemented network; the actual network includes additional residual blocks beyond those depicted. (b) A residual block showing how the input is added to the output.

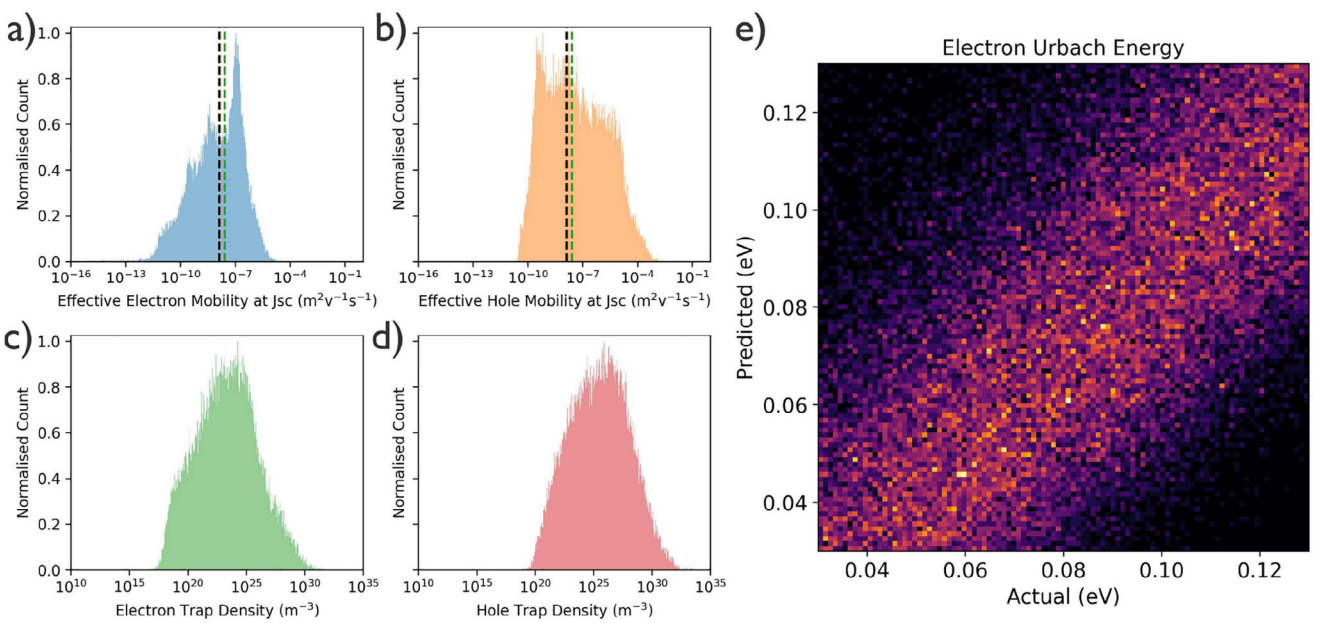


FIGURE 10 | Residual network predictions for: (a) Effective electron mobility at J_{sc}, (b) effective hole mobility at J_{sc}, (c) electron trap density, and (d) Hole trap density. In panels a and b, dashed lines represent experimentally determined mobility values from SCLC measurements: electron-only device (green), hole-only device (orange), and IMPS mobility mean (black). (e) A confusion matrix for electron Urbach energy prediction.

experiments can also produce large volumes of data that are time-consuming to interpret using standard fitting techniques. We aged the solar cell in ambient air at V_{oc} under illumination from a 1000 W m^{-2} Xenon Lamp, in accordance with the ISOS-L1 [42] testing protocol. To increase confidence in the reported results and separate light soaking effects, due to the ZnO electrode [43, 44], from true degradation, the device was aged in stages. After fabrication, it was left in the glove box to stabilise for 20 h, followed by 4 hours of exposure to light and air. The device was then returned to the glove box for two days to allow light soaking effects to subside. This was followed by another 8 h of light exposure in air, and another rest period in the glove box. Periodic JV measurements were taken throughout the entire process.

If one examines the line in Figure 11a representing PCE, the effect of this staggered degradation becomes apparent. Between

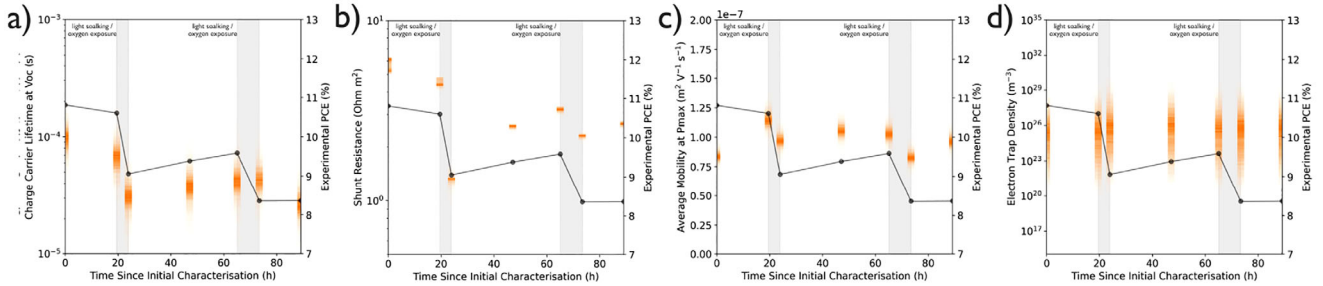


FIGURE 11 | Residual network predictions upon degrading PM6:BTP-ec9 devices for: (a) Charge carrier lifetime at V_{oc}, (b) shunt resistance, (c) average mobility at P_{max}; and (d) electron trap density. In each figure the measured experimental PCE during degradation is shown as a reference for overall device condition.

the first two black dots, there is a gradual reduction in PCE, likely indicating device stabilization after fabrication due to residual solvent evaporation or molecular relaxation. During the first ageing step, a sharp drop in efficiency occurs, potentially due to oxygen reacting with donor polymers and moisture ingress. After the device is returned to the glove box for 20 h, the efficiency increases by a few percent, followed by a rapid decline in PCE during the next ageing step. This behaviour suggests that the device undergoes initial burn-in, partial recovery in an inert atmosphere, and then accelerated degradation upon further exposure.

The left-hand axis of Figure 11a shows the predicted carrier lifetime, with the orange spots representing the probability distributions of the predicted values. It can be seen that the predicted recombination rates follow the trend of the PCE, mirroring degradation and recovery as it happens. Figure 11b,c,d plots the shunt resistance, effective electron mobility at P_{max} and electron trap density. Shunt resistance and effective mobility follow the trends of PCE, while average mobility and as far as we can tell trap density do not change.

The reduction in shunt resistance is likely due to moisture-induced leakage pathways. While, the decline in charge carrier lifetime probably reflects the formation of non-radiative recombination centres, such as oxidised species or defect states, either within the bulk or at critical interfaces. In contrast, the stability of charge carrier mobility suggests that long-range transport remains largely unaffected, potentially due to preserved molecular packing. The invariance in trap density further implies that deep trap states are not forming in significant numbers to be detected or the newly introduced traps are too shallow to be seen.

8 | Conclusion

We present two practical methods for extracting key material parameters, such as charge carrier mobility and recombination constants, from experimental JV measurements. By employing a differential learning framework, our approach expresses material properties as probability distributions rather than point estimates, enabling direct quantification of uncertainty. Performance is enhanced through the use of residual blocks, which improve the accuracy and robustness of the deep neural network. The approach presented here enables laboratories to extract physically meaningful material parameters and associated uncertainties directly from standard JV measurements, without requiring specialised characterisation techniques.

Our method requires significantly less training data, making it computationally efficient and accessible to experimental groups with limited resources. We validate the approach on both fresh PM6:Y12 and aged PM6:BTP-eC9 organic solar cells, demonstrating its ability to yield meaningful physical insights into device behaviour. These techniques represent a step toward integrating ML into routine materials characterization workflows. By lowering computational and experimental barriers, they open opportunities for rapid feedback during device optimization. Because the method relies only on standard JV measurements and modest computation, it provides an accessible pathway for accelerating materials discovery and device optimization across solar-energy research.

Acknowledgments

We thank the Deutsche Forschungsgemeinschaft (DFG) for funding this work (Research Unit FOR 5387 POPULAR, project no. 461909888).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

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

References

1. E. Mazzolini, R. A. Pacalaj, Y. Fu, et al., "Pathways to Upscaling Highly Efficient Organic Solar Cells Using Green Solvents: A Study on Device Photophysics in the Transition From Lab-to-Fab," *Advanced Science* 11, no. 31 (2024): 2402637.

2. S. Simotko, Y. Yadav, and G. L. Frey, "A Blending Approach for Dual Surface and Bulk Functionality in Organic Transistors," *Advanced Electronic Materials* 11 (2025): 2500085.

3. R. Chen, N. Liang, and T. Zhai, "Dual-Color Emissive OLED with Orthogonal Polarization Modes," *Nature Communications* 15 (2024): 1331.

4. D. Li, Q. Lian, T. Du, R. Ma, H. Liu, Q. Liang, Y. Han, G. Mi, O. Peng, G. Zhang, et al., "Co-Adsorbed Self-Assembled Monolayer Enables High-Performance Perovskite and Organic Solar Cells," *Nature Communications* 15 (2024): 7605.

5. S. Qiu, M. Majewski, L. Dong, et al., "Over One-Micron-Thick Void-Free Perovskite Layers Enable Highly Efficient and Fully Printed Solar Cells," *Energy & Environmental Science* 18 (2025): 5926–5939.

6. M. Li, Y. Yang, Z. Kuang, et al., "Acceleration of Radiative Recombination for Efficient Perovskite LEDs," *Nature* 630 (2024): 631.

7. J. Pang, H. Wu, H. Li, T. Jin, J. Tang, and G. Niu, "Reconfigurable Perovskite X-Ray Detector for Intelligent Imaging," *Nature Communications* 15 (2024): 1769.

8. X.-C. Fan, X. Tang, T.-Y. Zhang, et al., "Stable Narrowband Blue OLEDs by Modulating Frontier Molecular Orbital Levels," *Nature Communications* 16 (2025): 4936, https://doi.org/10.1038/s41467-025-60172-y.

9. Y. Bae, S. A. Park, S. Kim, et al., "Boosting the Mechanical Stability and Power Output of Intrinsically Stretchable Organic Photovoltaics with Stretchable Electron Transporting Layer," *Advanced Energy Materials* 15 (2025): 2405217.

10. S. Zahra, D. H. Ryu, J.-W. Ha, et al., "Trifluorophenyl Side Chain Engineering in Benzotriazole-Core Acceptors for High-Performance Organic Photovoltaics," *Energy & Environmental Materials* 8 (2025): e12875.

11. J.-D. Chen, H. Ren, F.-M. Xie, et al., "Harnessing Plasmon-Exciton Energy Exchange for Flexible Organic Solar Cells With Efficiency of 19.5%," *Nature Communications* 16 (2025): 3829.

12. Z. Li, C. Jia, Z. Wan, et al., "Boosting Mechanical Durability under High Humidity by Bioinspired Multisite Polymer for High-Efficiency Flexible Perovskite Solar Cells," *Nature Communications* 16 (2025): 1771.

13. Z. Wei, Q. Zhou, X. Niu, et al., "Surpassing 90% Shockley–Queisser V Oc Limit in 1.79 eV Wide-Bandgap Perovskite Solar Cells Using Bromine-Substituted Self-Assembled Monolayers," *Energy & Environmental Science* 18 (2025): 1847.

14. B. Ding, Y. Ding, J. Peng, et al., "Dopant-Additive Synergism Enhances Perovskite Solar Modules," *Nature* 628 (2024): 299.

15. F. Reginato, E. Lunedei, S. Mattiello, et al., "Improved Charge Recombination Efficiency in Organic Light-Emitting Transistors via Luminescent Radicals," *Advanced Functional Materials* 35 (2025): 2411845.

16. L. Wang, Q. Song, F. Pei, et al., "Strain Modulation for Light-Stable N– I– P Perovskite/Silicon Tandem Solar Cells," *Advanced Materials* 34 (2022): 2201315.

17. Y. Zheng, Z. Zhan, Q. Chen, et al., "Highly Sensitive Perovskite Photoplethysmography Sensor for Blood Glucose Sensing Using Machine Learning Techniques," *Advanced Science* 11 (2024): 2405681.

18. C. Shin, W. Jeong, E. Darici Lee, et al., "Molecular Engineering of Benzobisoxazole-Based Conjugated Polymers for High-Performance Organic Photodetectors and Fingerprint Image Sensors," *Energy & Environmental Materials* 8, no. 1 (2025): e12806.

19. M. M. Lee, J. Teuscher, T. Miyasaka, T. N. Murakami, and H. J. Snaith, "Efficient Hybrid Solar Cells Based on Meso-Superstructured Organometal Halide Perovskites," *Science* 338 (2012): 643.

20. P. Friedlingstein, M. O'Sullivan, M. Jones, et al., *Global Carbon Budget 2023*. Earth System Science Data, https://essd.Copernicus.org/articles/15/5301, (2023).

21. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," preprint, *arXiv*, January 30, 2017, https://doi.org/10.48550/arXiv.1412.6980.

22. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016): 770–778.

23. S. Elfwing, E. Uchibe, and K. Doya, "Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning," *Neural Networks* 107 (2018): 3.

24. R. Cao, C.-R. Zhang, X.-M. Liu, et al., "Molecular Design of Organic Photovoltaic Donors and Non-Fullerene Acceptors: A Combined Machine Learning and Genetic Algorithm Approach," *Journal of Materials Chemistry C* 13 (2025): 12150.

25. E. Kim, P. Hartnagel, B. Urbano, L. Christen, and T. Kirchartz, "Inferring Material Parameters from Current-Voltage Curves in Organic Solar Cells via Neural-Network-Based Surrogate Models," *Solar RRL* 9 (2025): e202500648, https://doi.org/10.1002/solr.202500648.

26. J. Achiam, S. Adler, S. Agarwal, et al., "GPT-4 Technical Report," preprint, *arXiv*, March 4, 2024, https://doi.org/10.48550/arXiv.2303.08774.

27. G. Team, R. Anil, S. Borgeaud, et al., "Gemini: A Family of Highly Capable Multimodal Models," preprint, *arXiv* May 9, 2025, https://doi.org/10.48550/arXiv.2312.11805.

28. A. Grattafiori, A. Dubey, A. Jauhri, et al., "The Llama 3 Herd of Models," preprint, *arXiv*, November 23, 2024, https://doi.org/10.48550/arXiv.2407.21783.

29. Y. Jiang, X. Dong, L. Sun, et al., "An Alcohol-Dispersed Conducting Polymer Complex for Fully Printable Organic Solar Cells With Improved Stability," *Nature Energy* 7 (2022): 352.

30. M. Hußner, C. Deibel, and R. C. MacKenzie, "The Physical Meaning of Time-Delayed Collection Field Transients From Disordered Devices," *Journal of Applied Physics* 135, no. 12 (2024): 125502.

31. T. Kirchartz, "Challenges and Opportunities for the Characterization of Electronic Properties in Halide Perovskite Solar Cells," *Chemical Science* 16, no. 19 (2025): 8153–8195.

32. W. Shockley and W. Read Jr, "Statistics of the Recombinations of Holes and Electrons," *Physical Review* 87 (1952): 835.

33. Y. Gao, R. Mackenzie, Y. Liu, B. Xu, P. H. Van Loosdrecht, and W. Tian, "Engineering Ultra Long Charge Carrier Lifetimes in Organic Electronic Devices at Room Temperature," *Advanced Materials Interfaces* 2, no. 4 (2015): 1400555.

34. R. MacKenzie, V. Balderrama, S. Schmeisser, et al., "Loss Mechanisms in High Efficiency Polymer Solar Cells," *Advanced Energy Materials* 6 (2016): 1501742.

35. N. Majeed, M. Saladina, M. Krompiec, S. Greedy, C. Deibel, and R. C. MacKenzie, "Using Deep Machine Learning to Understand the Physical Performance Bottlenecks in Novel Thin-Film Solar Cells," *Advanced Functional Materials* 30 (2020): 1907259.

36. M. Hußner, R. A. Pacalaj, G. Olaf Müller-Dieckert, et al., "Machine Learning for Ultra High Throughput Screening of Organic Solar Cells: Solving the Needle in the Haystack Problem," *Advanced Energy Materials* 14 (2024): 2303000.

37. D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning Representations by Back-Propagating Errors," *Nature* 323 (1986): 533.

38. A. Wagenpfahl, D. Rauh, M. Binder, C. Deibel, and V. Dyakonov, "S-Shaped Current-Voltage Characteristics of Organic Solar Devices," *Physical Review B* 82 (2010): 115306.

39. R. Hanfland, M. A. Fischer, W. Brütting, U. Würfel, and R. C. I. MacKenzie, "The Physical Meaning of Charge Extraction by Linearly Increasing Voltage Transients from Organic Solar Cells," *Applied Physics Letters* 103 (2013): 063904.

40. Y. Bengio, P. Simard, and P. Frasconi, "Learning Long-Term Dependencies with Gradient Descent Is Difficult," *IEEE Transactions on Neural Networks* 5 (1994): 157.

41. J. Vollbrecht, N. Tokmoldin, B. Sun, et al., "On the Relationship of the Effective Mobility and Photoconductance Mobility in Organic Solar Cells," *Energy Advances* 2 (2023): 1390.

42. M. O. Reese, S. A. Gevorgyan, M. Jørgensen, et al., "Consensus Stability Testing Protocols for Organic Photovoltaic Materials and Devices," *Solar Energy Materials and Solar Cells* 95 (2011): 1253.

43. J. Kim, G. Kim, Y. Choi, J. Lee, S. Heum Park, and K. Lee, "Light-Soaking Issue in Polymer Solar Cells: Photoinduced Energy Level Alignment at the Sol-Gel Processed Metal Oxide and Indium Tin Oxide Interface," *Journal of Applied Physics* 111, no. 11 (2012): 114511.

44. M. Günther, S. Lotfi, S. S. Rivas, et al., "The Neglected Influence of Zinc Oxide Light-Soaking on Stability Measurements of Inverted Organic Solar Cells," *Advanced Functional Materials* 33 (2023): 2209768.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Supporting File: smtd70921-sup-0001-SuppMat.pdf.