# Fast Inverse Design of Microstructures via Generative Invariance Networks

**Xian Yeow Lee[1], Joshua R. Waite[1], Chih-Hsuan Yang[1], Balaji Sesha Sarath Pokuri[1], Ameya Joshi[2], Aditya Balu[1], Chinmay Hegde[2], Baskar Ganapathysubramanian[1+], and Soumik Sarkar[1+]**

[1]Department of Mechanical Engineering, Iowa State University, Ames, IA, 50011, USA.
[2]Tandon School of Engineering, New York University, Brooklyn, NY, 11201, USA.
[+]Corresponding authors. Email: baskarg@iastate.edu, soumiks@iastate.edu

## ABSTRACT

The problem of efficient design of material microstructures exhibiting desired properties spans a variety of engineering and science applications. An ability to rapidly generate microstructures that exhibit user-specified property distributions transforms the iterative process of traditional microstructure-sensitive design. We reformulate the microstructure design process as a *constrained* Generative Adversarial Network (GAN). This approach explicitly encodes invariance constraints within a GAN to generate two-phase morphologies for photovoltaic applications obeying design specifications: specifically, various short circuit current density and fill-factor combinations. Such invariance constraints can be represented by deep learning-based surrogates of full physics models mapping microstructure to photovoltaic properties. To circumvent data generation bottlenecks, we utilize a multi-fidelity surrogate that reduces the requirements of expensive labels by 5X. Our approach enables fast generation of microstructures (in ≈190ms) with user-defined properties. Such physics-aware data-driven methods for inverse design problems are expected to democratize and accelerate the field of microstructure-sensitive design.

## Introduction

Advances in manufacturing (additive manufacturing, 3D printing, layer-by-layer deposition, real time control) allow us to precisely tailor the spatial distribution (i.e. microstructure) of materials. This opens up the possibility of microstructure-sensitive design, which involves identification of optimal material configurations that produce a desired property. Microstructure-sensitive design can impact a diverse array of applications ranging from membranes design to enhance water reclamation, battery electrode design to improve energy transport, and organic electronics active layer design to improve energy harvesting or sensing[1,2]. The systematic creation of *fast methods* for microstructure-sensitive design is, however, a challenging problem, usually due to the complexity of the 'forward model' that maps the microstructure to property. The availability of a fast inverse design framework will transform the field of microstructure-sensitive design, and significantly impact how we harvest, store and distribute energy and mass.

Over the past decade, a wide variety of approaches have been explored for solving the inverse microstructure design problem[3–5]. Traditionally, iterative optimization approaches have been the most popular strategy to search for microstructures that yield desirable characteristics. These approaches are typically time-consuming, computationally expensive, and often require manual supervision by domain experts. Furthermore, such approaches often lack the ability to generalize to new design constraints and require repeated exploration of the design space for each new design constraint or user choice. More generally, optimization based approaches are susceptible to challenges arising from (a) the combinatorial explosion of plausible microstructures, and (b) the computational complexity of function evaluation, i.e., solving the forward problem, especially for complex multi-physics problems. Thus, conventional optimization strategies based on using the full-physics forward model is a very challenging proposition, with limited deployment by groups with the skill set to use, and dedicated access to large computational clusters. As an example, current microstructure optimization approaches which rely on multiple forward-model PDE calculations for the discovery of a single optimal (inverse) design may typically require up to 160,000 CPU-hours for a given chemical system[3]. The push to democratize microstructure-sensitive design led to efforts focused on relaxing the challenges described above. Some approaches relied on developing cheaper, but less accurate surrogates for the forward model (example, using graphs[3]), while other approaches made the problem computationally tractable by severely constraining the design space (allowing only specific parameterized shapes)[4].

In this context, recent advances in the field of deep learning and scientific machine learning show promise for solving inverse design problems[6–12]. Particularly promising are Generative Adversarial Networks (GANs)[13], a class of generative deep learning models. Given a set of data, these generative models are capable of learning the underlying data distribution to generate new, realistic-looking samples. In the context of engineered systems, GANs have been successfully applied to areas such as differential equations[14], system-modelling[15,16], and material and drug discovery[17–19].

Generative models can be trained to reconstruct realistic looking microstructures[17,20,21]. The challenge is to train them to reconstruct microstructures that satisfy a user-defined set of properties – or more generally, satisfying a set of constraints. Recently, a modified version of generative models called Invariance Networks (InvNets), have been proposed to enable imposition of explicit constraints on the model outputs[22]. The formulation of InvNet, which allows constraints to be defined independently, provides great flexibility in terms of incorporating domain knowledge or user specifications, cast as constraints. In this work, we formulate the microstructure-sensitive design problem into that of training an InvNet with physics-based constraints. We deploy this framework to generate candidate two-phase microstructures/morphologies for organic photovoltaic (OPV) applications due to this microstructure's potential aspect in addressing a broad range of problems elaborated below. Flexible, light-weight, and wearable electronics and solar cells made from organic components provide a promising solution to wide array of societal needs. The potential application of these devices range from sensing (for precision and personalized medicine) to ambient energy harvesting (indoor solar cells) and energy efficient lighting and electronics. For example, the newest generation of small molecule acceptors have pushed single-junction organic photovoltaic (OPV) efficiencies over 14% and tandem efficiencies over 17%, potentially revolutionizing cheap, flexible, and green energy harvesting. A large body of work has demonstrated that the morphology in the active layer of OPV devices is key to enabling high-performance devices[23–25]. Thus, controlling the morphology in the active layer of these devices continues to be crucial for maximizing performance. Tremendous advances in self-assembly of flexible polymers suggest remarkable control of hierarchical structure, but the impact on high-performance organic electronics has been limited. This is because, despite the importance of active layer mesoscale morphology to OPV device performance, it remains a challenge to identify "ideal" microstructures that maximize power conversion efficiencies. A key question is whether multiple 'families' of optimal morphology exist, and whether these morphological characteristics depend on material specific parameters such as electron mobility, exciton diffusion length and biomolecular recombination (i.e. the molecular chemistry). Thus, a rapid, physics-aware microstructure design strategy will enable practitioners to systematically explore and unravel questions of how materials limitations affect optimal morphological features, thereby accelerate materials design leading to high-performance devices.

Here, we train an InvNet to generate microstructures that simultaneously obey multiple constraints, specifically a user defined short-circuit current density $J_{sc}$ and fill-factor $FF$. These two properties characterize the current-voltage performance of an OPV (for a given material system). The short circuit current density, $J_{sc}$, represents the maximum amount of current per unit area that can be drawn across a solar cell(when the applied voltage is zero). Meanwhile, the fill-factor $FF$ denotes the maximum amount of power that can be supplied by the solar cell as a ratio of peak theoretical power. These properties are intimately (and non-trivially) influenced by morphology[26]. We first demonstrate that a deep neural network (DNN) can be trained as a surrogate model to accurately predict the values of $J_{sc}$ and $FF$ given a two-phase microstructure. This requires a substantial amount of full fidelity training data, and only serves as a baseline forward model surrogate. Next, we propose a multi-fidelity neural network that achieves a similar predictive accuracy while utilizing a small fraction of full fidelity labels alongside low-fidelity labels. We then formulate an InvNet based inverse design framework where these surrogate (baseline as well as multi-fidelity) models are used as invariances to generate morphologies that satisfy user-design specifications.

# Results

We begin with a brief overview of our proposed methodology of using InvNets for fast generation of targeted two-phase morphologies. Figure 1(a) illustrates the overall InvNet framework with a Wasserstein-GAN (WGAN)[27] formulation. The WGAN model architecture ensures that the distribution of generated morphologies matches the true data distribution. Design specifications are enforced via an explicit invariance constraint, whereby the invariance loss is computed using the surrogate physics model represented by a deep neural network. This invariance loss ensure that the generator produces morphologies that satisfies the invariance constraint. Figure 1(b) and (c) shows two alternatives of surrogate models: a high-fidelity convolutional neural network (CNN) trained with a large amount of expensive labels obtained from high-fidelity simulations, and a multi-fidelity network trained with a mixture of high- and low-fidelity labels. The multi-fidelity network is trained on low-fidelity but computationally cheap labels alongside a fraction of high-fidelity labels to reduce the overall computation cost.

We present our results in the following order. First, we validate our CNN surrogate model for accurately predicting photovoltaic properties of a given morphology, followed by comparative assessment of prediction results of the multi-fidelity surrogate model. We then provide illustrative results of the microstructures generated by InvNet using both the high-fidelity network as well as the multi-fidelity network.

[Figure 1 about here.]

## High-fidelity short circuit current and fill-factor estimation

In this section, we validate the approach of utilizing a deep neural network as a surrogate of the physics-based forward model. This surrogate model also serves as the invariance constraint within the InvNet framework. As deep neural networks are known to be powerful function approximators with fast prediction times[28], we hypothesize that they are suitable candidates for representing the physics-based forward models.

We previously showed that a CNN can accurately classify microstructures in binned classes of $J_{sc}$[5]. Here, extend the idea to train a regressor that is capable of predicting $J_{sc}$ and $FF$ as continuous values. We train a CNN-based high-fidelity regressor $R_{HF}$ on a dataset of two-phase morphology images with high-fidelity simulated $J_{sc}$ values ranging from 0 to 7 $mA/cm^2$ and $FF$ values ranging from 0.4 to 0.8 (see Methods and Supplemental Materials for details on full physics simulations and material properties). The output of this surrogate model is a vector which consists of the estimated $J_{sc}$ and $FF$ values. Training this model yields a $R^2$ statistic of 0.994 for the estimations of $J_{sc}$ and $R^2$ statistic of 0.928 for estimations of $FF$ as seen in bottom two scatter plots of Figures 2(a). This suggests that the surrogate model is capable of estimating the properties with sufficiently high accuracy. The top two plots in Figure 2(a) shows the histograms of absolute error for both $J_{sc}$ and $FF$ respectively. The error distribution of $J_{sc}$ has a mean of $0.002 \pm 0.08$ $mA/cm^2$ while the error distribution of $FF$ have a mean magnitude of 7.68E-5 $\pm$ 0.02.

[Figure 2 about here.]

## Reducing data cost with multi-fidelity labels

An expected bottleneck of training a surrogate model on high-fidelity labels is the challenge of initial data generation, which may be computationally expensive. Various works in literature have previously explored the idea of leveraging both high- and low-fidelity data to accelerate computational models[29–32]. We exploit a similar idea by proposing a multi-fidelity surrogate model to circumvent the challenge of generating computationally expensive high-fidelity labels.

Having demonstrated that the high-fidelity surrogate is capable of estimating $J_{sc}$ and $FF$, we next provide a preliminary overview of the computation of low-fidelity labels, followed by results of estimating such low-fidelity labels via another neural network surrogate. Then, we present the results of the multi-fidelity surrogate model which alleviates the requirement of a large-labelled dataset required for training. We achieve this by using computationally inexpensive low-fidelity labels and a fraction of expensive high-fidelity labels to train the multi-fidelity surrogate model.

We use our prior work where a mechanistic consideration of the photo physics as three distinct processes (absorption and generation of excitons; exciton diffusion and dissociation; charge transport and collection) allowed identification of three morphology descriptors that together showed high correlation with $J_{sc}$[33,34]. These descriptors are computed by representing the two-phase morphology as a weighted-graph and evaluating standard graph measures (like connected components and path lengths). Since the complexity of graph-based problems and corresponding algorithms are well understood, these descriptors are computationally inexpensive to compute. These morphology descriptors thus provide a low fidelity link between morphology and performance. We next describe how using only a small fraction of simulated high-fidelity labels ($J_{sc}$ and $FF$) along with information from low-fidelity labels (morphology descriptors), we train a surrogate model that has the similar predictive performance as a model trained purely using high-fidelity labels. Training this multi-fidelity surrogate required availability of differentiable low-fidelity descriptors, which we get by training another neural network, $R_g$, that maps a morphology to the low-fidelity descriptors (see SI for results, network and training details)

### *Multi-fidelity short circuit current and fill-factor estimation*

We train a multi-fidelity model that estimates the magnitudes of $J_{sc}$ and $FF$ using a limited amount of high-fidelity labels with the help of the low-fidelity descriptors. This multi-fidelity network consists of the low-fidelity network, $R_g$ and a separate shared-embedding network, as illustrated in Figure 1(c). The purpose of the shared embedding network is to learn additional features that are useful in estimating the properties which were not captured by the low-fidelity model. In our experiments, we used only 20% of randomly sampled high-fidelity labels to train the multi-fidelity network, which resulted in a $R^2$ of 0.989 and 0.894 for $J_{sc}$ and $FF$ respectively. The absolute error distributions of the $J_{sc}$ and $FF$ predictions have a mean of $-0.009 \pm 0.12$ $mA/cm^2$ and $-0.003 \pm 0.02$ respectively. Figure 2(b) shows the scatter plots of the properties estimated by the multi-fidelity model against the ground truth values as well as the distribution of errors. As seen in Figure 2(a) and (b), the $R^2$ of the $R_{HF}$ and $R_{MF}$ models are similar although the label requirements of the multi-fidelity model is reduced by 80%. We stress that while the low-fidelity network was trained using the entire dataset, the multi-fidelity model was only trained with 20% of the high-fidelity labels, which are significantly more expensive to generate (e.g., evaluating the $J_{sc}$ and $FF$ of one morphology needs about 1 cpu-hr, whereas the low fidelity metrics can be computed in less than a minute). Hence, by using the multi-fidelity network, we alleviate the problem of requiring a large labelled dataset to train a surrogate physics model as the invariance constraint evaluator in the InvNet.

## Targeted Microstructure Generation

[Figure 3 about here.]

We present the results of generating targeted morphologies that are tailored to design specifications using our proposed *InvNet* with multi-fidelity surrogate model framework. In Figure 3(a), we show samples of microstructures generated with InvNet for different design specifications. In the top row, we show examples of morphologies with low $J_{sc}$ values and high $FF$ values. As we traverse down the rows of Figure 3(a), the specified $J_{sc}$ values are increased while the $FF$ values are decreased. It is observed that the InvNet-trained generator is able to generate a variety of candidate microstructures with different morphologies given the same design specifications. This signifies that the generator has learnt the underlying distribution of the actual data and no mode collapse occurred during training which can result in only similar morphologies being generated. This also anecdotally validates a hypothesis in the OPV community that there exist multiple families of morphologies that produce identical performance.

To further verify that the generated morphologies satisfy the imposed design constraints, we generated an additional 1000 morphologies for different ranges of $J_{sc}$ and $FF$ values and compared the estimated properties of these morphologies with the actual design specifications. The values of these estimated properties and design specifications are plotted as densities and shown in Figure 3(b). We observe that the specified values and generated values for both $J_{sc}$ and $FF$ have highly overlapping densities. These overlapping densities show that generator is capable of creating morphologies that satisfy the imposed design specifications, hence enabling targeted design of candidate two-phases microstructures.

Nonetheless, we observe that there are situations where the generated morphologies do not adhere to the design specifications, as seen in the first row of Figure 3(b), where the density of generated morphologies (in solid green) had a range of $J_{sc}$ values that are higher than the specified range of $J_{sc}$ values (in dotted blue). Since the proposed framework is fundamentally data-driven, we hypothesize that this failure mode was caused by an imbalanced dataset where samples from the low $J_{sc}$ and high $FF$ regions might be sparse. To confirm this hypothesis, we visualize the training data distribution in Figure 3(c). Based on the visualization of the joint density, we observe that there are indeed very few samples in the top left region, where morphologies have a low $J_{sc}$ and high $FF$ values. However, it is interesting to recognize that even when the generator fails to generate morphologies with specified $J_{sc}$ in such sparse training data regions, the rank order of the morphologies' $J_{sc}$ are still preserved. Instead of generating morphologies with random $J_{sc}$s', the generated morphologies defaulted to morphologies with low $J_{sc}$ and high $FF$ values which are well supported with data.

## Comparing high-fidelity and multi-fidelity InvNets

Next, we provide qualitative results to compare the effects of using the high-fidelity, $R_{HF}$, and multi-fidelity $R_{MF}$ surrogate model as the invariance constraint evaluator in InvNet framework. In Figure 2, we have shown that the performances of the high- and multi-fidelity surrogate models are comparable. Moreover, we are also interested in investigating if the higher variance of the multi-fidelity surrogate will compound and affect the results of the generated morphologies. To study this, we trained InvNet with the same network architecture and replaced the $R_{MF}$ with $R_{HF}$. We illustrate the results from both methods in Figure 4. In terms of the generated morphologies, we do not observe any significant difference between the two methods. Both the high- and multi-fidelity InvNets are capable of generating microstructures of varying morphologies without signs of mode collapse. However, the density plots which are used to validate the constraint invariances reveal two interesting observations.

First, we observe that the high-fidelity InvNet is more capable of generating low $J_{sc}$/high $FF$ morphologies in comparison with the multi-fidelity InvNet. This is evident in the first row, where the density of morphologies generated by high-fidelity InvNet has a higher overlapping area with the design specifications as compared to the density of morphologies created by multi-fidelity InvNet. We attribute this to the fact that $R_{HF}$ was exposed to a much larger and diverse set of morphologies as compared to $R_{MF}$, which results in the high-fidelity InvNet being able to learn the underlying structure of the low $J_{sc}$/high $FF$ morphologies better when training for the invariance. Thus, this suggests we can expect the performance of high-fidelity InvNet to be more robust and consistent when queried in regions where training data is sparser.

The second interesting observation we make is that the high-fidelity InvNet also tends to generate morphologies that are a little more biased in terms of the $FF$. This can be observed in the second, third, and fourth rows where the densities of high-fidelity $FF$ are slightly shifted from the $FF$ design specifications. Referring back to Figure 3(c), we observe that the marginal density of $FF$ data is highly skewed towards the lower regions. Therefore, it is possible that by training $R_{HF}$ on the entire high-fidelity dataset and subsequently using it as the invariance constraint evaluator to train InvNet does result in generated morphologies that are more biased in terms of the design specifications. This highlights the importance of having a balanced dataset when using our proposed framework for morphology generation.

[Figure 4 about here.]

### Efficiency of neural-network based methods versus physics-based models

In Table 1, we compare the wall-clock running times of our proposed neural-network based methods with physics-based methods for a few different scenarios. All timings were performed on the same platform using a NVIDIA Titan RTX GPU and averaged across 100 function evaluations. In the first two columns, we show the average computation times for evaluating the $J_{sc}$ and $FF$ properties of a given morphology. We observe that both multi- and high-fidelity methods are several orders of magnitude faster than a high-fidelity physics simulation. A second advantage is that with the surrogate models, only one evaluation is required to estimate both $J_{sc}$ and $FF$ simultaneously. In comparison, performing the physics simulation requires separate individual evaluations for $J_{sc}$ and $FF$. Comparing the multi-fidelity surrogate model $R_{MF}$ with the high-fidelity surrogate model $R_{HF}$, we note that $R_{HF}$ is an order of magnitude faster than $R_{MF}$. However, training $R_{HF}$ comes at the cost of requiring a large dataset with high-fidelity labels. On the other hand, $R_{MF}$ requires a smaller amount of high-fidelity labels, but requires training a more complex model architecture, which increases computation time. Hence, we view the benefits of each method as a trade-off between availability of data with computation time.

In the third column, we show the total time required to train InvNet for 1E5 epochs. We observe that the high-fidelity InvNet is $\approx$ 3X faster than multi-fidelity InvNet, which is expected since the training of InvNet is dependent on the surrogate model to compute the invariance loss. We also include an estimate of the time required to train the InvNet if we were to replace the invariance constraint evaluator with an actual physics-based model to compute the invariance loss. As observed, training such an InvNet will require $\approx$ 60k hours, which is not tractable in compared to using a neural network-based surrogate model.

Last but not least, we provide the morphology generation time for a single morphology. Since the process of generating a morphology using InvNet during inference is independent of surrogate model, there is no significant difference time difference between using the high-fidelity versus multi-fidelity InvNet. In summary, we conclude that there is no significant difference in terms of the querying a trained high-fidelity versus multi-fidelity InvNet to generate targeted morphologies. Instead, the deciding factor of which model to apply depends on the availability of high-fidelity labels or computation resources. The high-fidelity InvNet framework is faster to train but requires a large dataset of high-fidelity labels to pre-train the surrogate model. Conversely, the multi-fidelity InvNet model requires less high-fidelity labels but requires a more complex network architecture which results in longer training times.

[Table 1 about here.]

## Discussion

The ability to rapidly synthesize targeted microstructure designs is essential in a broad range of scientific and engineering applications. We propose a data-efficient generative framework (InvNet) that casts user-specifications as explicit invariance constraints to generate candidate two-phase microstructures that adheres to design specifications. While recent works with similar objectives have proposed frameworks that demonstrated promising results[12,22], we highlight that those approaches is not capable of solving our specific application in a tractable manner. This is particularly due to the extremely long and expensive computation required to evaluate the constraints, which is a common bottleneck in the community. Hence, to remedy this challenge, we leverage neural network-based surrogates for the purpose of fast constraint evaluation. Using a surrogate, our framework addresses the challenge of expensive constraint evaluation while simultaneously circumventing the need of having a differentiable and explicit, closed-form expression of the constraints. Combining these advantages, we believe that our method results in a far more general-purpose framework that is applicable to a wider range of inverse design problems. Additionally, we have also supplemented our surrogate-based generative framework with a multi-fidelity approach to improve the data requirements of the model. This multi-fidelity approach reduces foreseeable expensive label generation procedures, which is an obstacle that is not present in inverse design problems where design constraints can be tractably computed. For further discussion on the motivation of our framework, we defer reader to the Methods section. From our experimentation, our results illustrate that neural networks are capable of being accurate surrogates of expensive full-physics simulations and the InvNets trained with multi-fidelity surrogates are capable of generating various candidate morphologies which caters to both $J_{sc}$ and $FF$ specifications. Furthermore, comparing the results of InvNet trained with high-fidelity and multi-fidelity surrogates reveals no significant differences in performance, thus reinforcing the fact of data-efficiency benefits of using the multi-fidelity surrogate. A wall-clock comparison of training times reveals that a trade-off exists between the high-fidelity and multi-fidelity modes, with the multi-fidelity version of the surrogate and InvNet having longer training times.

While we have demonstrated our proposed framework through the lens of a material microstructure design problem that uses a data-driven surrogate, we emphasize that our InvNet framework is certainly not limited to purely data-driven surrogate approaches. Since the invariance constraint of InvNet is explicit, it can be easily replaced or combined with other data-free approaches. In this regard, a key future direction is to develop InvNets that explicitly incorporate complex physics/domain knowledge in a computationally tractable manner. This approach will significantly reduce the dependency of the proposed framework on data availability and extend the capability of the framework to extrapolate beyond the support of data. Other promising directions include extending the current framework to generate morphologies with more than two phases as well as validating the generalizability of the framework on a dataset with more than two target properties. To conclude, our vision is that the computational tools developed in this paper will serve to democratize and accelerate the area of microstructure-sensitive design.

# Methods

## Description of two-phase morphology microstructures

**Microstructures:** We use a large dataset of microstructure images arising from solving the Cahn-Hilliard (CH) equation with varying initial conditions. The Cahn-Hilliard equation[35] describes phase separation occurring in a binary alloy under thermal annealing. It tracks the evolution of local volume fraction of each phase, in the presence of spatial gradients in chemical potential of the system. Hence, in the time evolution process, one first observes an initial rapid separation of the well-mixed system into its constituent phases, followed by slow coarsening of the respective domains. Thus, the microstructures generated will have lower energy compared to those at the beginning stages of the simulation, according to the second law of thermodynamics. Image data arising from the simulations provide a rich dataset for design of microstructures. Specifically, the morphologies obtained through the simulation will be similar to the morphologies in a real active layer of organic photovoltaic cells[5]. We use an in-house solver for generating the microstructure images.

**Photophysics Annotation of Microstructures:** Each of the morphologies is virtually interrogated to extract its current-voltage characteristics, by solving a morphology aware (i.e. spatially heterogeneous) photophysics device model. We deploy a validated, in-house software that uses a finite element based solution strategy for solving the photophysics device model. The photophysics model is described by the steady state *excitonic drift diffusion (XDD) equations*. The XDD equations are a set of four tightly coupled partial differential equations that model the optoelectronic physics of energy harvesting in organic photovoltaic devices. The photophysics consists of the following stages:

- Incident solar radiation causes the generation of energetically active electron-hole pairs, called excitons (denoted by X), in the donor regions of the microstructure. These excitons diffuse across the microstructure and have a finite lifetime before becoming ground state electron-hole pairs;
- Excitons that diffuse and reach the donor-acceptor interface undergo dissociation into electrons (denoted by n) and holes (denoted by p) at the donor-acceptor interface. The dissociation mechanism is material and field dependent (denoted by D);
- These generated charges (n,p) traverse the microstructure and reach their corresponding electrodes (cathode and anode) to produce a current. Two mechanisms are responsible for driving carrier transport or current flow. First, the drift, which is caused by the presence of an electric field (denoted as the gradient of the potential, $\nabla\varphi$, and second, the diffusion, which is caused by a spatial gradient of electron or hole concentration;
- The distribution of electrons and holes in the microstructure interacts with the applied voltage and influences the electrostatic potential $\varphi$ across the microstructure. Finally, electrons and holes can recombine (denoted by R) to create excitons

The photophysics described above is encoded as the exciton drift diffusion (XDD) equations[26].

$$\nabla.J_n - R + D = 0 \tag{1}$$
$$-\nabla.J_p - R + D = 0 \tag{2}$$
$$\nabla.(\varepsilon_r \varepsilon_0 \nabla\varphi) = q(n-p) \tag{3}$$
$$-\nabla.(V_t \mu_x \nabla X) - fD_{[\nabla\varphi,X]} - R_{[x]} = -G - R_{[n,p]} \tag{4}$$

Here, $X$, $n$, $p$ represent the exciton, electron and hole distributions respectively. $\varphi$ represents the electric potential. $q$ represents the elementary charge. $V_t$ represents the thermal voltage. $\varepsilon$ is the dielectric constant in the donor and recipient materials. $\mu_{n/p/X}$ are the mobilities of electron/hole/exciton respectively. The current densities $J_n$ and $J_p$ are given by the constitutive equations

$$J_n = -qn\mu_n\nabla\varphi + qV_t\mu_n\nabla n \tag{5}$$
$$J_p = -qp\mu_p\nabla\varphi - qV_t\mu_p\nabla p \tag{6}$$

These set of high-dimensional, complex PDEs are solved to get the performance of the solar cell device, which is charecterized by the short-circuit current, and fill factor.

## Framework development

In this section, we discuss the motivation of creating a data-driven framework capable of generating microstructures with various targeted morphologies while adhering to design specifications. Previous studies have demonstrated that InvNets can effectively generate novel two-phase microstructures that satisfy explicit constraints such volume fractions and domain size and also generate poly-crystalline microstructures (a discrete-valued generation problem) by relaxing the generation problem to a probabilistic assignment problem[22]. However, we consider a couple of drawbacks of the existing InvNet in terms of scalability. Previously, it has only been shown that InvNet worked with explicitly defined constraints or invariances. Nevertheless, existing works have not addressed what happens when the invariances cannot be explicitly defined. Additionally, evaluation of the invariances are often computationally expensive and time-consuming. For example, in the case of our application, evaluating $J_{sc}$ and $FF$ of a given morphology involves solving a set of differential equations that can take up to approximately several hours.

As such, this limitation hampers the scalability of training InvNet. Hence, we represent such invariance constraints, which cannot be explicitly expressed or are too computationally expensive with a deep neural network. By representing the invariance with a deep neural network surrogate, the evaluation of the invariance constraints can be significantly accelerated since the forward evaluation of a neural network is fast once the model is trained. Utilizing a neural network surrogate also has the benefit of not requiring the invariances, such as the equations governing a physical system, to be explicitly known as long as labeled data are available to train the surrogate model. Also, a neural network representation of the invariance simplifies the training of InvNet. During training of the InvNet, the parameters of the entire model are optimized by utilizing gradient information from the invariance loss function. Since neural network models are differentiable, gradient information with respect to the invariance loss can be easily computed using modern deep learning libraries with automatic-differentiation capabilities. In comparison, using other forms of explicit invariances will necessitate the constraints to be differentiable, and the gradients will have to be calculated separately.

Nonetheless, as alluded above, representing the invariance with a deep neural network does result in a second drawback, which is the availability of labeled data. In the context of our application, creating a labeled data set of morphologies with corresponding $J$ and $FF$ values is computationally expensive, and defeats the goal of avoiding costly physics-based simulations. This second drawback motivates the development the multi-fidelity surrogate which alleviates the problem of generating expensive labels.

## Training details

**High-fidelity surrogate model**: To improve the robustness of the surrogate model, we first performed standard image augmentation techniques, image rotation and flipping, which resulted in an augmented dataset of $\approx 38k$ images of augmented morphologies. To ensure a stable training process, we also scaled the labels of $J_{sc}$ and $FF$ to belong in the same numerical range. Following standard practices, we partitioned 80% of the data as training data and reserved 20% of data as a test data. Since the task of the surrogate model is to essentially perform a multi-target regression, the loss function of the regressor is formulated as:

$$L_{R_{HF}} = \|R_{HF_{\phi,J_{sc}}}(I) - J_{sc}\|_2^2 + \|R_{HF_{\phi,FF}}(I) - FF\|_2^2 \tag{7}$$

where $R_{HF}$ denotes the high-fidelity surrogate model, parameterized by parameters $\phi$, $I$ is the input image of the microstructure and $J_{sc}$ and $FF$ are the true label values. The high-fidelity model architecture we used is a sequential model which consists of two convolution layers, each followed by batch normalization layer, ReLU activation, and a max pooling layer. Two dense layers were used after the two convolution blocks along with dropout layers to avoid over-fitting during training. The output of $R_{HF}$ is a vector of two values that corresponds to the estimated $J_{sc}$ and $FF$ values. The model was trained using Adam optimizer with a learning rate of 3E-4 for 25 epochs. Additionally, we also investigated network architectures with separate final layers that do not share parameters. We observed no significant improvement in prediction accuracy while the cost of computational memory requirement was increased.

**Multi-fidelity surrogate model**: Before describing the training details, we briefly justify the need to replace the graph-based computation of low-fidelity descriptors with another neural network surrogate, $R_g$ in the multi-fidelity model. While multi-fidelity frameworks are effective in reducing the requirement of expensive labels[32], they are currently not tractable for application as an invariance constraint in InvNets. This is because updating the generator's parameters in InvNet requires the gradient computation of the invariance-loss function. However, graph-based methods used to compute the low-fidelity descriptors are often non-differentiable. Therefore, optimizing the parameters of the generator via conventional back-propagation becomes a non-trivial problem. Additionally, evaluating the low-fidelity descriptors using previously proposed graph-based method requires that the generated images be converted into nodes and edges on-the-fly during training, which incurs additional computational cost and time. Hence, a neural network surrogate which is differentiable and can directly evaluate graph features of morphologies in the pixel domain circumvents both of these challenges.

As illustrated in Figure 1(c), the multi-fidelity network encompasses both low-fidelity network (described in SI) and a shared-embedding network. The purpose of the shared-embedding network is to learn additional features that are not already captured by the low-fidelity network for estimating $J_{sc}$ and $FF$. During training of the multi-fidelity network, the low-fidelity network predicts the low-fidelity descriptors of a given microstructure, which are combined with the image embeddings from the shared embedding network. These two vectors are then passed through a dense layer to estimate $J_{sc}$ and $FF$. As we are only using a limited amount of high-fidelity labels, it is possible that training the multi-fidelity network might lead to a biased model due to label imbalance. To avoid such issues, we constructed the following weighted loss function with empirically-determined scaling constants that balances the errors between the estimations of $J_{sc}$ and $FF$. $R_{MF}$ denotes the multi-fidelity surrogate model where $\chi$ and $\omega$ represents the parameters of the shared-embedding network and low-fidelity network respectively.

$$L_{R_{MF}} = L_{J_{sc}} + L_{FF} \tag{8}$$

$$L_{J_{sc}} = \lambda_1 (J_{sc}^2 + J_{sc}) \| R_{MF_{\chi,\omega,J_{sc}}}(I) - J_{sc} \|_2^2 \tag{9}$$

$$L_{FF} = \lambda_2 \| R_{MF_{\chi,\omega,FF}}(I) - FF \|_2^2 \tag{10}$$

with $\lambda_1$ and $\lambda_2$ heuristically set to 0.008 and 0.0005, respectively. We highlight that in principle, the weights of the low-fidelity network $R_g$ are already trained and can be frozen. Nevertheless, in practice, we find that allowing the weights of the low-fidelity network to optimized alongside the entire network does result in a slightly better estimations. To train the multi-fidelity network, we used SGD optimizer with a learning rate of 1E-3 and trained the network for 100 epochs.

**Generator and Discriminator of InvNet**: In this section, we provide the training description of InvNet with the multi-fidelity surrogate model as the invariance constraint evaluator. Since the main modification that we've proposed occurs in the invariance constraint, the formulation of InvNet's loss function remains as

$$L_{InvNet} = L_G(\theta, \psi) + L_I(\theta) \tag{11}$$

where $L_G$ denotes the standard loss function of the WGAN, with $\theta$ being the parameters of the generator, $\psi$ being the parameters of the discriminator. Both the generator and discriminator are also represented using deep neural networks. The invariance loss $L_I$ is expressed as:

$$L_I = \| R_{MF_{J_{sc}}}(G_\theta(z)) - R_{MF_{J_{sc}}}(I) \|_2^2 + \| R_{MF_{FF}}(G_\theta(z)) - R_{MF_{FF}}(I) \|_2^2 \tag{12}$$

with $G_\theta$ denoting the generator, $z$ denoting a latent vector sampled from a uniform distribution, $G_\theta(z)$, denoting the image of generated morphology and $I$ denoting a real morphology sampled from the dataset. During training, the weights of the surrogate physics model, $\phi$ are kept frozen, and $R_{MF}$ acts purely as an invariance constraint evaluator that estimates the morphological properties of the generated microstructures. Only the parameters of the discriminator $\psi$, and generator $\theta$ are optimized.

To train the InvNet, we instantiate the generator with an architecture that consists of one dense layer, five residual blocks with skipped connections, and one convolution layer. Each residual block is made up of two batch-normalization layers and two convolution layers with up-sampling operations. ReLU activation functions were used after every layer, except for the last convolution layer. We used the sigmoid activation function on the output of the convolution layer to generate 128 x 128 images of microstructures. The Discriminator network consists of one convolution layer, four residual blocks, and a dense layer. The residual blocks are similar to the blocks used in the Generator, with the exception that the convolution layers are paired with down-sampling operations and layer-normalization is used instead of batch-normalization. As we've chosen to use the WGAN-GP[36] variant of GAN, the output of the discriminator is a single scalar value estimating the Wasserstein distance between the distributions of generated and real microstructures. To compute the invariance loss, we use the multi-fidelity surrogate model $R_{MF}$ to ensure that generated morphologies had properties that are similar to the properties of the real morphologies. Both the Generator and Discriminator are trained alternatively using Adam optimizer with a learning rate of 3E-4 for 1E5 epochs. We include specific details of network layers we used in the generator, discriminator and multi-fidelity network in the Supplementary Materials. Note that in the methodology presented, we have described the InvNet framework using the multi-fidelity surrogate, $R_{MF}$ as the invariance constraint evaluator. We highlight that the methodology for training the framework using high-fidelity network is exactly the same, with only $R_{HF}$ replacing $R_{MF}$.

## References


1. Anthony, J. E. Addressing challenges. *Nat. materials* **13**, 773–775 (2014).
2. Yao, H.-B. *et al.* A flexible and highly pressure-sensitive graphene–polyurethane sponge based on fractured microstructure design. *Adv. Mater.* **25**, 6692–6698 (2013).
3. Du, P., Zebrowski, A., Zola, J., Ganapathysubramanian, B. & Wodo, O. Microstructure design using graphs. *npj Comput. Mater.* **4**, 1–7 (2018).
4. Noruzi, R., Ghadai, S., Bingol, O. R., Krishnamurthy, A. & Ganapathysubramanian, B. Nurbs-based microstructure design for organic photovoltaics. *Comput. Des.* **118**, 102771 (2020).
5. Pokuri, B. S. S., Ghosal, S., Kokate, A., Sarkar, S. & Ganapathysubramanian, B. Interpretable deep learning for guided microstructure-property explorations in photovoltaics. *npj Comput. Mater.* **5**, 1–11 (2019).
6. Stoecklein, D., Lore, K. G., Davies, M., Sarkar, S. & Ganapathysubramanian, B. Deep learning for flow sculpting: Insights into efficient learning using scientific simulation data. *Sci. reports* **7**, 46368 (2017).
7. Lore, K. G., Stoecklein, D., Davies, M., Ganapathysubramanian, B. & Sarkar, S. A deep learning framework for causal shape transformation. *Neural Networks* **98**, 305–317 (2018).
8. Bock, F. E. *et al.* A review of the application of machine learning and data mining approaches in continuum materials mechanics. *Front. Mater.* **6**, 110 (2019).
9. Lee, X. Y., Balu, A., Stoecklein, D., Ganapathysubramanian, B. & Sarkar, S. A case study of deep reinforcement learning for engineering design: Application to microfluidic devices for flow sculpting. *J. Mech. Des.* **141** (2019).
10. Balu, A. *et al.* A deep learning framework for design and analysis of surgical bioprosthetic heart valves. *Sci. reports* **9**, 1–12 (2019).
11. Geng, Y., van Anders, G., Dodd, P. M., Dshemuchadse, J. & Glotzer, S. C. Engineering entropy for the inverse design of colloidal crystals from hard shapes. *Sci. advances* **5**, eaaw0514 (2019).
12. Kim, B., Lee, S. & Kim, J. Inverse design of porous materials using artificial neural networks. *Sci. advances* **6**, eaax9324 (2020).
13. Goodfellow, I. *et al.* Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680 (2014).
14. Yang, L., Zhang, D. & Karniadakis, G. E. Physics-informed generative adversarial networks for stochastic differential equations. *SIAM J. on Sci. Comput.* **42**, A292–A317 (2020).
15. de Oliveira, L., Paganini, M. & Nachman, B. Learning particle physics by example: location-aware generative adversarial networks for physics synthesis. *Comput. Softw. for Big Sci.* **1**, 4 (2017).
16. Wu, J.-L. *et al.* Enforcing statistical constraints in generative adversarial networks for modeling chaotic dynamical systems. *J. Comput. Phys.* **406**, 109209 (2020).
17. Yang, Z. *et al.* Microstructural materials design via deep adversarial learning methodology. *J. Mech. Des.* **140** (2018).
18. Sanchez-Lengeling, B. & Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **361**, 360–365 (2018).
19. Prykhodko, O. *et al.* A de novo molecular generation method using latent vector based generative adversarial network. *J. Cheminformatics* **11**, 74 (2019).
20. Fokina, D., Muravleva, E., Ovchinnikov, G. & Oseledets, I. Microstructure synthesis using style-based generative adversarial networks. *Phys. Rev. E* **101**, 043308, DOI: 10.1103/PhysRevE.101.043308 (2020).
21. Gayon-Lombardo, A., Mosser, L., Brandon, N. P. & Cooper, S. J. Pores for thought: generative adversarial networks for stochastic reconstruction of 3d multi-phase electrode microstructures with periodic boundaries. *npj Comput. Mater.* **6**, 1–11 (2020).
22. Joshi, A. *et al.* Invnet: Encoding geometric and statistical invariances in deep generative models. In *AAAI 2020* (2020).
23. Yu, G., Gao, J., Hummelen, J. C., Wudl, F. & Heeger, A. J. Polymer photovoltaic cells: enhanced efficiencies via a network of internal donor-acceptor heterojunctions. *Science* **270**, 1789–1791 (1995).
24. Shaheen, S. E. *et al.* 2.5% efficient organic plastic solar cells. *Appl. Phys. Lett.* **78**, 841–843 (2001).
25. Guerrero, A. & Garcia-Belmonte, G. Recent advances to understand morphology stability of organic photovoltaics. *Nano-micro letters* **9**, 10 (2017).

26. Kodali, H. K. & Ganapathysubramanian, B. Computer simulation of heterogeneous polymer photovoltaic devices. *Model. Simul. Mater. Sci. Eng.* **20**, 035015 (2012).

27. Arjovsky, M., Chintala, S. & Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 214–223 (2017).

28. Csáji, B. C. *et al.* Approximation with artificial neural networks. *Fac. Sci. Etvs Lornd Univ. Hungary* **24**, 7 (2001).

29. Sarkar, S. *et al.* Multifidelity and multiscale bayesian framework for high-dimensional engineering design and calibration. *J. Mech. Des.* **141** (2019).

30. Costabal, F. S., Perdikaris, P., Kuhl, E. & Hurtado, D. E. Multi-fidelity classification using gaussian processes: accelerating the prediction of large-scale computational models. *Comput. Methods Appl. Mech. Eng.* **357**, 112602 (2019).

31. Babaee, H., Bastidas, C., DeFilippo, M., Chryssostomidis, C. & Karniadakis, G. A multifidelity framework and uncertainty quantification for sea surface temperature in the massachusetts and cape cod bays. *Earth Space Sci.* **7**, e2019EA000954 (2020).

32. Yang, C.-H. *et al.* Multi-fidelity machine learning models for structure-property mapping (2020). ArXiv.

33. Wodo, O., Zola, J., Pokuri, B. S. S., Du, P. & Ganapathysubramanian, B. Automated, high throughput exploration of process–structure–property relationships using the mapreduce paradigm. *Mater. discovery* **1**, 21–28 (2015).

34. Wodo, O., Tirthapura, S., Chaudhary, S. & Ganapathysubramanian, B. Computational characterization of bulk heterojunction nanomorphology. *J. Appl. Phys.* **112**, 064316 (2012).

35. Cahn, J. W. & Hilliard, J. E. Free energy of a nonuniform system. I. interfacial free energy. *The J. chemical physics* **28**, 258–267 (1958).

36. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. & Courville, A. C. Improved training of wasserstein gans. In *Adv. Neural Inf. Proc. Sys. (NeurIPS)*, 5767–5777 (2017).

## Acknowledgements

### General

We gratefully thank Dr. Adarsh Krishnamurthy for fruitful discussions and positive feedbacks.

### Funding

This work was supported by the ARPA-E DIFFERENTIATE program under grant DE-AR0001215. BG, C-HY, and BP were supported in part by DoD MURI 6119-ISU-ONR-2453. CH and AJ were supported in part by NSF grants 2005804 and 1815101. Computing support from XSEDE and Iowa State University is gratefully acknowledged.

### Author contributions

CH, BG, SS initiated the project; BP, BG planned and generated dataset. XYL, AJ, CH, BG, SS designed the ML framework; XYL, JW, CH-Y performed the training; XYL, AJ and AB analyzed the data; all authors contributed to writing the manuscript.

### Competing interests

The authors declare no competing interests.

### Data and materials availability

The datasets generated and/or analysed during the current study will be available upon acceptance of the paper.

## List of Figures

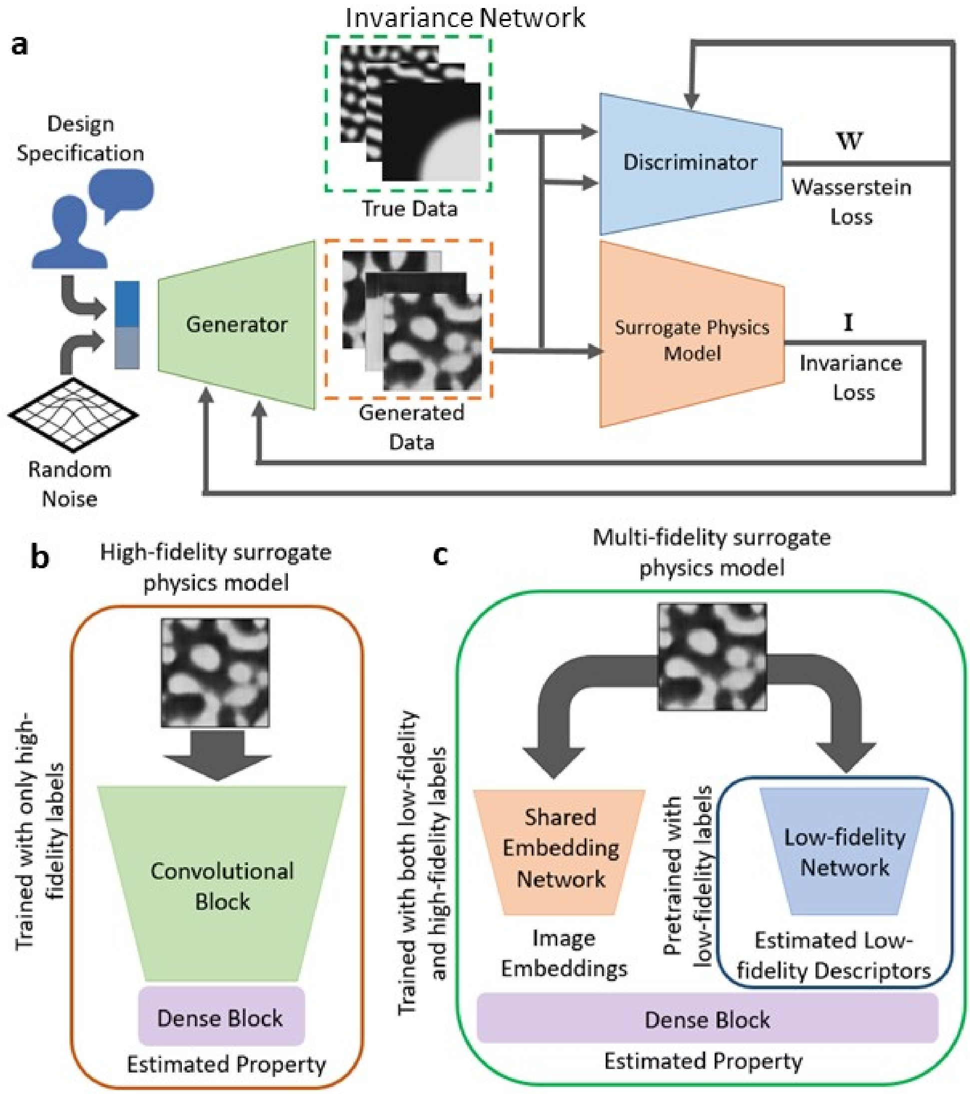


**Figure 1. Overview of our proposed framework for targeted microstructure generation based on user design specifications.** **(a)** Illustration of InvNet architecture, which utilizes a modified GAN to learn the underlying data distribution to generate new morphologies. We propose the use of a surrogate physics model, represented by a deep neural network, to enforce an invariance constraint by evaluating the properties of generated microstructures to ensure that user design specifications are satisfied. **(b)** A high-fidelity CNN is trained to predict the properties of the microstructures, which is used as the surrogate physics model in the InvNet. **(c)** To reduce the requirements of expensive, high-fidelity labels to train the surrogate model, we propose a multi-fidelity network which attains the same predictive accuracy as training the network on high-fidelity data by combining information from cheap, low-fidelity labels and a fraction of high-fidelity labels.

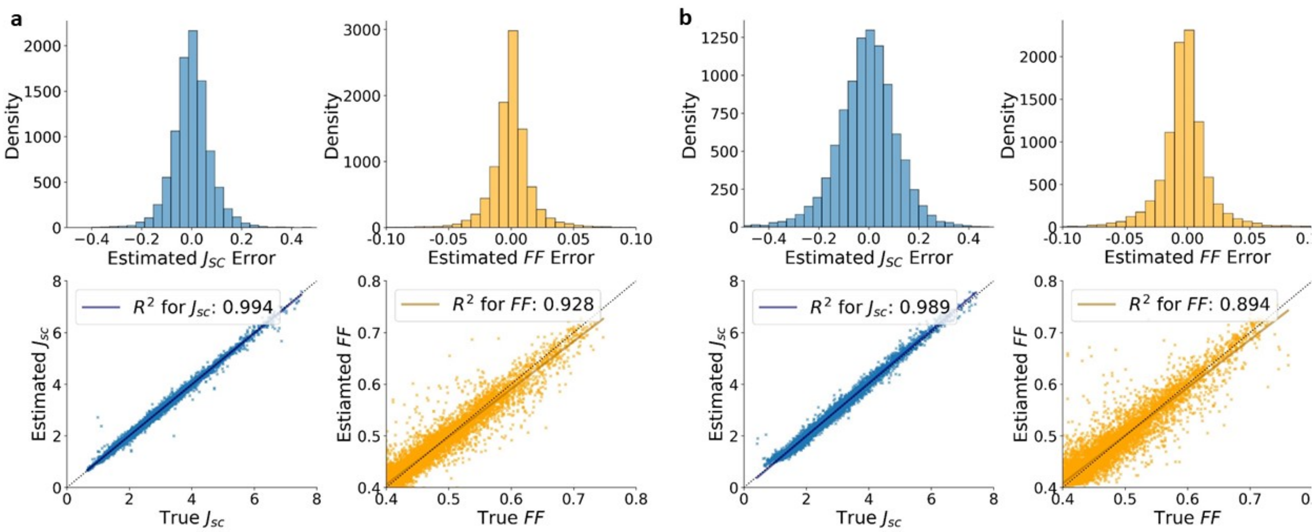


**Figure 2. Results of high-fidelity and multi-fidelity surrogate models.** (**a**) Left figures summarize the distribution of errors for both $J_{sc}$ and $FF$ estimation using the high-fidelity surrogate physics model. Bottom plots visualize the correlation plot of the estimated properties with respect to the ground truth values. In both cases, the predicted values have high correlation coefficients, $R^2$ values of greater than 0.9. (**b**) Summary of error distributions for $J_{sc}$ and $FF$ estimation using the multi-fidelity surrogate model which was trained with only 20% of high-fidelity labels. We observe that while there is slight drop in $R^2$ and increase in variance, there is a huge marginal gain in terms of decreasing the amount of expensive simulations required to generate the high-fidelity labels.

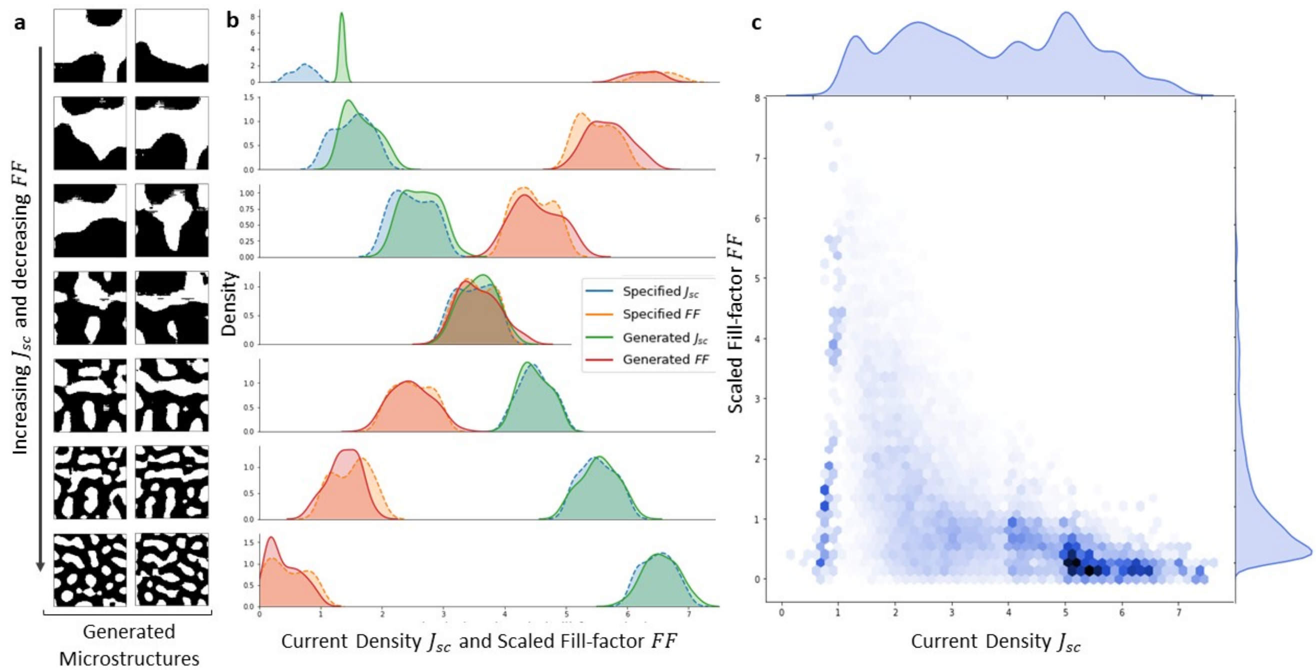


**Figure 3. Results of targeted microstructure design using multi-fidelity InvNet. (a)** Examples of morphologies generated by InvNet for the specified $J_{sc}$ and $FF$ ranges shown on the right densities. **(b)** Densities of estimated $J_{sc}$ and $FF$ from generated morphologies compared with a range of respective design specifications for 1000 samples. Observe that the densities of the design specifications and generated morphologies properties in the mid- and high-ranges (rows 2 to 7) are highly overlapping, signifying that the invariances are satisfied. In contrast, the densities at the region of low $J_{sc}$ are more deviated, signifying a more biased model at the region where the training data is sparse. **(c)** Visualization of joint and marginal densities of training data for both $J_{sc}$ and $FF$. Notice that the marginal density of $J_{sc}$ labels is relatively well balanced, while the marginal density of $FF$ is extremely skewed, resulting in sparser data around certain regions.

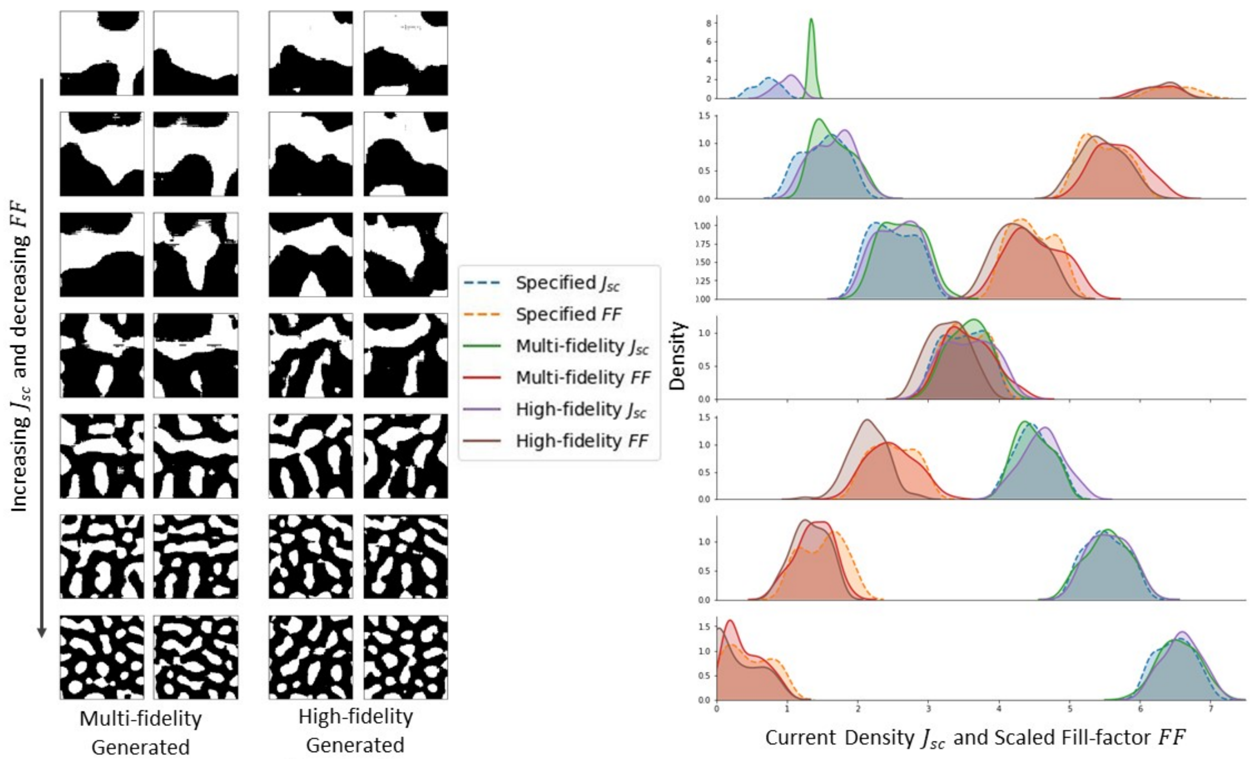


**Figure 4. Qualitative comparison of morphologies generated by the high-fidelity InvNet vs multi-fidelity InvNet.** Visually, we observe that both models are capable of generating varying morphologies which follows a similar trend as we varied the design specifications. Looking at the densities of property invariances, we observe that the high-fidelity InvNet performs slightly better than multi-fidelity InvNet by generating morphologies which are closer to design specifications in the low $J_{sc}$ high $FF$ regions where training data is sparse. However, the high-fidelity InvNet also tend to generate morphologies which are slightly biased in terms of the $FF$, as observed in rows 3, 4 and 5.

# List of Tables

| | $J_{sc}$ Evaluation | $FF$ Evaluation | InvNet Training | Morphology Generation |
|---|---|---|---|---|
| High-Fidelity | 5.9 ms | 5.9 ms | 5.8 hr | 191.0 ms |
| Multi-Fidelity | 55.3 ms | 55.3 ms | 18.7 hr | 192.0 ms |
| Physics-Model | 9.0 min | 72.0 min | 60,017.0 hr* | N/A |

**Table 1. Comparison of average computation times of neural network-based methods vs physics-based methods for different processes.** $J_{sc}$ and $FF$ columns denotes the time required to evaluate the corresponding properties given a morphology. InvNet training times are based on our training scheme of 1E5 epochs. *Physics model-based InvNet training is based on an estimate if the invariance loss were to be computed using high-fidelity physics simulation. Morphology Generation column denotes the time required for a trained InvNet to generate a single morphology given design specification values of $J_{sc}$ and $FF$.