
Generative Design of Material Microstructures for Organic Solar Cells using Diffusion Models

Ethan Herron

Xian Yeow Lee

Aditya Balu

Balaji Sesha Sarath Pokuri

Baskar Ganapathysubramanian

Soumik Sarkar

Adarsh Krishnamurthy

Iowa State University

{edherron, xylee, baditya, baskarg, soumiks, adarsh}@iastate.edu

Abstract

Score-based methods, particularly denoising diffusion probabilistic models (DDPMs), have demonstrated impressive improvements to state-of-the-art generative modeling. Due to their impressive ability to sample from complex distributions, DDPM models and related variants, all broadly categorized under diffusion models, apply to various applications. In this work, we compare the performance of a diffusion model with a Wasserstein Generative Adversarial Network in generating two-phase microstructures of photovoltaic cells. We demonstrate the diffusion model’s performance improvements in generating realistic-looking microstructures and its ability to cover several modes of the target distribution.

1 Introduction

The microstructure—spatial distribution of electron donor and acceptor domains—plays an important role in determining the photocurrent in thin film organic photovoltaics (OPVs). Optimizing the microstructure can lead to higher photocurrent generation and is an active area of experimental research. A framework to reliably generate microstructure-sensitive materials would have a nontrivial impact on various energy sectors. Designing flexible and lightweight OPV materials would enable customized electronics and solar cells. These highly customized materials would have applications in personalized medicine, energy harvesting, and energy-efficient electronics.

In this work, we focus on designing two-phase microstructures for organic solar cells, but the broader implication is the design of material microstructures with the assistance of machine learning models, particularly deep neural networks. Machine learning (ML) models have become a popular method for driving materials design because they are fast during inference and nearly match the performance of their numerical and computationally intensive counterparts. Designing material microstructures requires using generative modeling algorithms, a broad category encompassing any method that transforms noise into a sample that matches some target distribution. In this work, we train a diffusion model to sample a random normal distribution and map it to a sample from the target distribution. The target distribution we have defined is a dataset of 2D images representative of the microstructures of two-phase morphologies for photovoltaic cells. The goal of this work is to highlight the improved performance of diffusion models for generating material microstructures compared to previous methods.

Background and Related Work

Generative modeling for material design has been well-studied [1–4], but typically from the lens of feasibility. Most of these works use Generative Adversarial Networks [5] at their core and either aim to strictly generate plausible structures or from a more controllable perspective where the




Figure 1: The inverse diffusion trajectory taken by the trained diffusion model to produce 5 different sample microstructures.

microstructures generated satisfy some set of user-defined characteristics. In either case, the results are impressive but are hindered by choice of a generative modeling algorithm.

Generative Adversarial Networks (GANs) are a type of generative modeling algorithm in which two different neural networks are pitted against each other in a zero-sum game. One neural network, the generator, is tasked with generating samples that appear to be from the target distribution. The other network, the discriminator, is tasked with measuring the likelihood that a sample is from the target distribution. This amounts to optimizing each network to beat the other at their respective task. This dynamic between two highly nonlinear actors often leads to instability. GANs have been shown to produce exceptionally high-quality results but are equally notorious for being difficult to train as well as suffering from mode collapse. In this work, we implement a Wasserstein Generative Adversarial Network (WGAN) [6]. The discriminator in a WGAN measures the Wasserstein distance instead of the Jensen-Shannon divergence between two probability distributions.

Diffusion models [7–9] are a recent advancement in generative modeling algorithms. Diffusion models learn the inverse of a diffusion process to transform random noise into a sample from an arbitrary target distribution. These models are optimized to predict the noise added at each time step of a predefined and analytic diffusion process. The inverse diffusion process is then carried out by starting with pure white noise and progressively removing the predicted noise for a predefined number of iterations.

2 Methods

Diffusion models are composed of two distinct parts: the forward and backward diffusion processes. The forward diffusion process is simply the iterative addition of Gaussian noise to a sample from some target distribution. This forward process may be carried out an arbitrary number of times, creating a set of progressively noisier and noisier samples of the original data sample. The following Markov chain defines this process:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}) \quad (1)$$

where $\mathbf{x}_0$ is the sample from some target distribution $q(\mathbf{x})$, and a variance schedule defined as $\{\beta_t \in (0, 1)\}_{t=1}^T$. The reverse diffusion process would then iteratively remove the noise added during the forward diffusion process, i.e., $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$. In practice, it isn't feasible to sample $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ since we do not know the entire distribution, therefore we approximate the conditional probabilities via a neural network, $G_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$, where the parameters of G , θ , are updated via gradient-based optimization. In this work, training a diffusion model consists of training a neural network to predict the noise added at the given time step. The objective function is then:

$$\|z - G_\theta(\mathbf{x}_t, t)\|^2 = \|z - G_\theta(\sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{(1 - \alpha_t)}z, t)\|^2 \quad (2)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and $z \sim \mathcal{N}(0, \mathbf{I})$. In summary, the network is given an image with some level of white noise, dependent on the time step, and predicts the noise applied at that time step. The mean squared error is then used to compare the predicted noise at that time step with randomly sampled white noise.

3 Experiment Details

The following section outlines the data used during experimentation and the network architectures of each neural network. The training and sampling process is outlined for the diffusion model, as well as the training and inference times of the WGAN compared to the diffusion model.

The entire dataset of two-phase microstructures is created by generating roughly 25,000 microstructures by solving the Cahn-Hilliard (CH) equation [10, 11] with randomly sampled initial conditions. The CH equation describes the phase separation in a binary alloy under thermal annealing. The trajectory taken by the CH equation alters the volume fraction of each phase according to the spatial gradients in the chemical potential of the system. The final two-phase microstructures generated by the CH solver will be comparable to real organic photovoltaic cells. To increase the size of the dataset, we implement horizontal and vertical flips of each sample in the initial dataset. The final dataset is then composed of nearly 80,000 different microstructures; all represented using 128×128 pixel binary images. All evaluations of the proposed diffusion model and the baseline WGAN were carried out on the test dataset, which we define to be 15% of the total dataset.

The diffusion model utilizes a typical U-Net architecture with three blocks of downsampling and upsampling. Three different two-layer MLPs are used to map scalar values, indicative of the current time step to embeddings. These embeddings are injected into each of their respective upsampling blocks. In total, the U-Net consists of 18,237,185 parameters. During training, the Adam optimizer [12] was used with a learning rate of $1e-4$, with a linear learning rate decay scheme. The forward diffusion process used to inject noise into the true data distribution comprised 1,000 time steps. Because of this, the diffusion model is trained to predict the noise over 1,000 different time steps. It is important to note that the model does not directly reverse the 1,000 diffusion steps iteratively during training. Training is structured such that the loss is computed for only a single time step of each sample in the batch. These time steps are different for each batch and are randomly sampled from a uniform distribution. The diffusion model is given an initial image of pure white noise during inference. The diffusion model is then run for 1,000 time steps, predicting the noise to be removed at each time step, finally converging to the final microstructure. This process is visualized in Figure 1.

The diffusion model was trained to randomly sample from the distribution of all potential microstructures defined by the CH numerical solver. To quantitatively evaluate the final optimized distribution we trained a surrogate model [13] to approximate the solution to the *excitonic drift-diffusion equations* (XDD) using an in-house numerical method [14]. These equations provide the current-voltage characteristic of a given microstructure. This surrogate model takes a microstructure as input and predicts the short-circuit current as well as the fill factor. We trained the surrogate model to achieve r^2 values of 0.993 and 0.939 for the short-circuit current and fill factor, respectively, on the test dataset. To assess the performance of the proposed diffusion model against the baseline WGAN, as well as the ground truth microstructures, we compare each generated microstructure by approximating the current-voltage characteristics with the surrogate model. Additionally, we have compared the P1 and P2 correlations amongst the three types of generated microstructures. The P1 correlation is the volume fraction of each material in the microstructure, which amounts to computing the mean of each image. The P2 correlation [15] between material microstructures is also used to measure the differences between the three distributions. See Figure 2 and Figure 3 for reference.




Figure 2: A comparison of the P1 and P2 correlations of the generated microstructures.




Figure 3: A comparison of predicted short-circuit current and fill factor values, respectively.




Figure 4: Generated microstructures.

4 Results and Discussion

During experimentation, we were able to demonstrate the increased capacity to model the desired target distribution of potential microstructures with the use of a diffusion model. Qualitatively, [Figure 4](#) shows the diffusion model’s increased performance in generating microstructures similar to those generated by solving the CH equation. The baseline WGAN-generated microstructures are representative of the target distribution but lack clarity as well as smooth boundaries. In [Figure 2](#), we quantitatively showcase the diffusion model’s improved ability to sample the target distribution of the microstructures. In addition, in [Figure 3](#), we show the diffusion model’s ability to also match the output performance characteristics of the generated microstructures better with the baseline as compared to the baseline WGAN model. In future works, we look to incorporate a conditioned diffusion model to generate microstructures with user-defined short-circuit current and fill factor attributes. Previous work [\[16\]](#) has demonstrated this task can be accomplished using invariance-driven GANs. Given this work, we hypothesize we can further improve upon the previous works using diffusion models and extend that work by conditioning the model on additional attributes such as material type and other current-voltage characteristics.

5 Broader Impacts

The learning capacity of generative algorithms showcased by diffusion models does raise concern for the potential of malicious applications. In this case, we do not see any reason for concern with this application of diffusion models, as the barrier to entry for manufacturing these microstructures is exceptionally high. Additionally, we cannot immediately recognize any malicious activities one could conduct using customized material microstructures.

6 Conclusions

The potential impacts of customizable material microstructures demand robust and fast generative design methods. This work shows that Diffusion models are thus far the highest-performing and a reliable option to generate high-fidelity microstructures. This is compared to previous methods, which were susceptible to mode collapse, and often failed to represent the breadth of the target distribution. This work demonstrated the advantages of using diffusion models for generating material microstructure designs and provides a footing for additional work in this direction.

7 Acknowledgments

This work was supported in part by the National Science Foundation under grants LEAP-HI-2053760 and ARPA-E DIFFERENTIATE-DE-AR0001215.

References

- [1] Baekjun Kim, Sangwon Lee, and Jihan Kim. Inverse design of porous materials using artificial neural networks. *Science advances*, 6(1):eaax9324, 2020.
- [2] Zijiang Yang, Xiaolin Li, L Catherine Brinson, Alok N Choudhary, Wei Chen, and Ankit Agrawal. Microstructural materials design via deep adversarial learning methodology. *Journal of Mechanical Design*, 140(11), 2018.
- [3] Andrea Gayon-Lombardo, Lukas Mosser, Nigel P Brandon, and Samuel J Cooper. Pores for thought: generative adversarial networks for stochastic reconstruction of 3d multi-phase electrode microstructures with periodic boundaries. *npj Computational Materials*, 6(1):1–11, 2020.
- [4] Daria Fokina, Ekaterina Muravleva, George Ovchinnikov, and Ivan Oseledets. Microstructure synthesis using style-based generative adversarial networks. *Phys. Rev. E*, 101:043308, Apr 2020. doi: 10.1103/PhysRevE.101.043308. URL <https://link.aps.org/doi/10.1103/PhysRevE.101.043308>.
- [5] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. URL <https://arxiv.org/abs/1406.2661>.
- [6] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan, 2017. URL <https://arxiv.org/abs/1701.07875>.
- [7] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- [8] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.
- [9] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models, 2021. URL <https://arxiv.org/abs/2102.09672>.
- [10] John W. Cahn and John E. Hilliard. Free energy of a nonuniform system. i. interfacial free energy. *The Journal of Chemical Physics*, 28(2):258–267, 1958. doi: 10.1063/1.1744102. URL <https://doi.org/10.1063/1.1744102>.
- [11] Balaji Sesha Sarath Pokuri, Viraj Shah, Ameya Joshi, Chinmay Hegde, Soumik Sarkar, and Baskar Ganapathysubramanian. Binary 2D morphologies of polymer phase separation, February 2019. URL <https://doi.org/10.5281/zenodo.2580293>.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Chih-Hsuan Yang, Balaji Sesha Sarath Pokuri, Sangeeth Balakrishnan, Chinmay Hedge, Soumik Sarkar, and Baskar Ganapathysubramanian. Multi-fidelity machine learning models for structure-property mapping, 2020. ArXiv.
- [14] Balaji. pbssarath/excitonic_drift_diffusion: v1.0, February 2021. URL <https://doi.org/10.5281/zenodo.4552777>.
- [15] Surya R. Kalidindi. 3-Statistical quantification of material structure. In Surya R. Kalidindi, editor, *Hierarchical Materials Informatics*, pages 75–110. Butterworth-Heinemann, Boston, 2015. ISBN 978-0-12-410394-8.
- [16] Xian Yeow Lee, Joshua R. Waite, Chih-Hsuan Yang, Balaji Sesha Sarath Pokuri, Ameya Joshi, Aditya Balu, Chinmay Hegde, Baskar Ganapathysubramanian, and Soumik Sarkar. Fast inverse design of microstructures via generative invariance networks. *Nature Computational Science*, 2020.