
Physics-aware Deep Generative Models for Creating Synthetic Microstructures

Rahul Singh*, Viraj Shah*, Balaji Pokuri

Iowa State University
Ames, IA 50010

rahulsn@iastate.edu, viraj@iastate.edu, balajip@iastate.edu

Soumik Sarkar, Baskar Ganapathysubramanian, Chinmay Hegde

Iowa State University
Ames, IA 50010

soumiks@iastate.edu, baskarg@iastate.edu, chinmay@iastate.edu

Abstract

A key problem in computational material science deals with understanding the effect of material distribution (i.e., microstructure) on material performance. The challenge is to synthesize microstructures, given a finite number of microstructure images, and/or some physical invariances that the microstructure exhibits. Conventional approaches are based on stochastic optimization and are computationally intensive. We introduce three generative models for the fast synthesis of binary microstructure images. The first model is a WGAN model that uses a finite number of training images to synthesize new microstructures that weakly satisfy the physical invariances respected by the original data. The second model explicitly enforces known physical invariances by replacing the traditional discriminator in a GAN with an invariance checker. Our third model combines the first two models to reconstruct microstructures that respect both explicit physics invariances as well as implicit constraints learned from the image data. We illustrate these models by reconstructing two-phase microstructures that exhibit coarsening behavior. The trained models also exhibit interesting latent variable interpolation behavior, and the results indicate considerable promise for enforcing user-defined physics constraints during microstructure synthesis.

1 Introduction

1.1 Motivation

An overarching theme of materials research is the design of material distributions (also called microstructure) so that the ensuing material exhibits tailored properties. The final property of the material system is intricately connected to the underlying microstructure. This theme encompasses several material systems including porous materials [15], steels and welds [2], composites [10], powder metallurgy [16], 3D printing [13], energy storage devices as batteries [7], and energy converting devices like bulk hetero-junction solar cells [11]. Microstructure-sensitive design has been used to tailor a wide variety of properties including strengths, heat and mass diffusivities, energy storage capacity and lifetime, and energy conversion efficiency. In microstructure-sensitive design, quantifying the effect of microstructure features on performance is critical for the efficient design of application-tailored devices.

*Equal contribution

There are extensive efforts to experimentally image the microstructure (including X-ray, optical, and electron microscopy). However, getting a complete virtual instance of a microstructure is non-trivial, as is the process of perturbing specific features of the microstructure. Thus, an entire sub-field in computational material science is devoted to the development of methods for the *simulation* of microstructures [6, 5, 12]. Here, microstructure realizations are synthesized that satisfy certain target statistical properties of the material distribution. These statistical properties could be scalars (such like total volume fraction of a material) or more complex functions (like 2-point correlations and other material statistics) [15].

1.2 Our contributions

Over the past several decades, a number of methods have been developed for microstructure simulations, including Gaussian random fields [12], optimization-based methods [19], multi-point statistics [4], and layer-by-layer reconstruction [14]. Most of these techniques use some form of optimization to refine an initial microstructure to satisfy target constraints. They are formulated as a minimization of the difference between the constraints/invariances of the simulated microstructures and the target. Consequently, these approaches are computationally intensive, requiring several compute hours even to simulate just *one* synthetic example. Such simulation methods are prohibitively slow to perform analysis for complex problems, which typically involve millions of pixels and require multiple constraints/invariances satisfaction.

In this paper, we introduce novel techniques that use generative adversarial networks (GANs) to generate microstructures. We propose and test three separate techniques:

1. The first approach uses a standard GAN architecture trained with a Wasserstein metric. We show that the generated images respect the distribution of certain physical invariances — specifically, the one-point and two-point correlations — of the training data. In this approach, we effectively let the discriminator learn the features of the data.
2. The second approach replaces the traditional discriminator with a *checker* function. This checker function is defined by the user and identifies the most physically informed features of the microstructures. There is no discriminator training involved and thus the issue of mode collapse can be averted. Here, the data requirements are minimal, and data is only used to calibrate the checker. We demonstrate how the generated images are diverse and mimic the supplied invariance metric (two point correlation curve).
3. Finally, we propose a hybrid of the above two architectures. We demonstrate its potential to simultaneously assimilate patterns both from the available data and user description. This enables the exploration and replication of non-quantifiable phenomena in the data, along with user-defined constraints.

We validate our techniques using a range of numerical experiments. We employ a 2D microstructure dataset that is generated by simulating a phase separation process of two (immiscible) components under thermal annealing [17]. Diverse microstructures can be generated, of varying domain-purity, domain size, interfacial area and (relative) volume fractions of the components. Overall, our results indicate that the models that we obtain have considerable promise to not just capture visually salient microstructure features but *also the physics* underlying the data generation.

2 Proposed techniques

In the context of this paper, we consider the underlying material to be a two-phase homogeneous, isotropic material. Our setup for statistical characterization of microstructure follows with Torquato [15]. Consider an instance of the two-phase homogeneous isotropic material within d -dimensional Euclidean space \mathbb{R}^d (where $d \in \{2, 3\}$). A phase function $\phi(\cdot)$ is used to characterize this two-phase system, defined as:

$$\phi^{(1)}(\mathbf{r}) = \begin{cases} 1, & \mathbf{r} \in V_1, \\ 0, & \mathbf{r} \in V_2, \end{cases} \quad (1)$$

where $V_1 \in \mathbb{R}^d$ is the region occupied by phase 1 and $V_2 \in \mathbb{R}^d$ is the region occupied by phase 2.

Given this microstructure defined by the phase function, ϕ , statistical characteristics can be evaluated. These include the n -point correlation functions for $n = 1, 2, 3, \dots$. For homogeneous and isotropic media, n -point correlations depend neither on the absolute positions of n -points, nor on the rotation of these spatial co-ordinates; instead, they depend only on relative displacements. The 1-point correlation function, p_1 , commonly known as *volume fraction*, is constant throughout the material. The volume fraction of phase 1, $p_1^{(1)}$, is defined as:

$$p_1^{(1)} = \mathbb{E}_{\mathbf{r}} \phi^{(1)}(\mathbf{r}).$$

The 2-point correlation is defined as:

$$p_2^{(1)}(r_{12}) = \mathbb{E}_{\mathbf{r}_1, \mathbf{r}_2} \left[\phi^{(1)}(\mathbf{r}_1) \phi^{(1)}(\mathbf{r}_2) \right].$$

The 2-point correlation is one of the most important statistical descriptors of microstructures. An alternate interpretation of 2-point correlation is the probability that two randomly chosen points \mathbf{r}_1 and \mathbf{r}_2 a certain distance apart both share the same phase.

Henceforth we omit the superscript representing the phase and subscripts representing the spatial points for simplicity, and refer to volume fraction as p_1 , and 2-point correlation as p_2 . It can be shown that $p_2(r=0) = p_1$ and $\lim_{r \rightarrow \infty} p_2(r) = p_1^2$.

We now propose three generative models capable of generating grayscale microstructure images of two-phase materials. Each model gives us a varying degree of control over the statistical invariances (p_1 and p_2) of the generated samples.

2.1 Wasserstein GAN

Generative Adversarial Networks (GAN) are powerful models that attempt to learn high-dimensional data distributions such as images [8]. The key strategy is to pose the estimation of generative model as two-player minimax game, with each player being a trainable model, referred as the generator and discriminator respectively. Both the networks are trained simultaneously in an adversarial manner. The generator is trained to generate realistic images by learning a nonlinear mapping from a low-dimensional space of latent parameters to the space of real images, while the discriminator is trained to discriminate or classify the samples generated by the generator (G) as either ‘real’ or ‘fake’. After sufficient training, the generator network is (often) able to reproduce synthetic images that closely resemble the original images. Formally, the minimax objective of GAN can be represented as:

$$\min_G \max_D \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r} [\log(D(\mathbf{x}))] + \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_g} [\log(1 - D(\tilde{\mathbf{x}}))], \quad (2)$$

where \mathbb{P}_r is the real data distribution, and \mathbb{P}_g is data distribution of images generated by generator. $\tilde{\mathbf{x}}$ is defined as $\tilde{\mathbf{x}} = G(\mathbf{z})$ with \mathbf{z} being drawn from noise distribution p , typically a Gaussian distribution. If the discriminator is trained optimally, then at each generator update the minimization of the function in Eq. 2 turns out to be the minimization of the Jensen-Shannon divergence between \mathbb{P}_r and \mathbb{P}_g .

The GAN framework proposed in [8] can produce visually appealing samples, but it usually suffers from training instability. The reason for the training difficulty, as explained in [1] is that the divergence is not continuous with respect to the parameters of the generator. An proposed improvement, known as Wasserstein GAN (WGAN), advocates using the Earth Mover Distance (also known as the Wasserstein-1 distance), which is continuous under reasonable assumptions. Moreover, the training can be stabilized via a suitable weight-clipping step in each epoch [1].

However, even the WGAN can sometimes generate poor samples and fail to converge. To tackle the issue, an alternative to weight clipping is proposed in [9], which stabilizes the GAN training by penalizing the norm of the gradient of the discriminator with respect to its input. The new objective function, which is a combination of Wasserstein distance and a gradient penalty (GP), becomes:

$$L = \underbrace{\mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_g} [D(\tilde{\mathbf{x}})] - \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r} [D(\mathbf{x})]}_{\text{Wasserstein loss}} + \underbrace{\lambda \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_{\tilde{\mathbf{x}}}} \left[(\|\nabla_{\tilde{\mathbf{x}}} D(\tilde{\mathbf{x}})\|_2 - 1)^2 \right]}_{\text{Gradient penalty}}. \quad (3)$$

Here, the samples $\tilde{\mathbf{x}}$ from the distribution $\mathbb{P}_{\tilde{\mathbf{x}}}$ are obtained by interpolating uniformly along straight lines between pairs of points sampled from the data distribution \mathbb{P}_r and the generator distribution \mathbb{P}_g .

In our first approach, we train a WGAN with gradient penalty (WGAN-GP) according to Eq. 3 as proposed in [9]. The details of training and the obtained results are in Section 3.

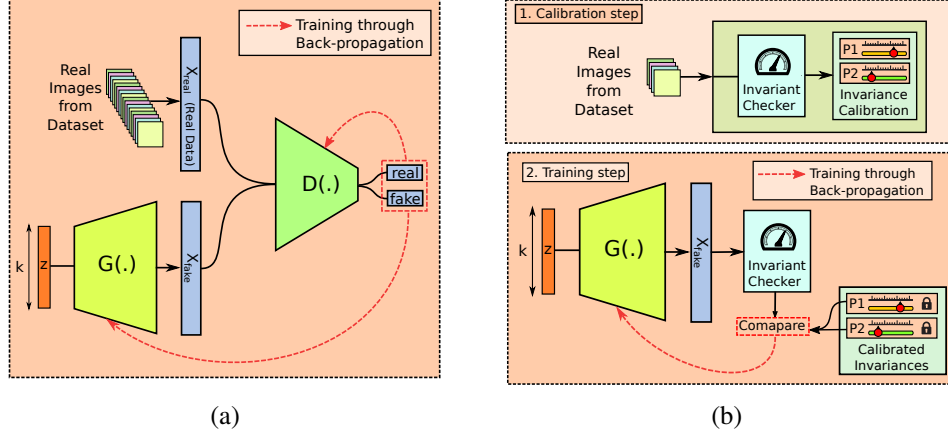


Figure 1: (a) Generative Adversarial Network model; (b) Generative Invariance Network model

2.2 Generative Invariance Network

Though our WGAN-GP model shows an ability to learn the statistical descriptors of the given microstructure morphologies to some extent, it gives very coarse control on the properties of the generated samples. The only factor determining the generated samples is the training data, making it difficult for the model to generate samples with precise values of p_1 , p_2 and other higher order statistical measures on demand. Such a need usually arises in multiple fields of material science. To resolve this, we advocate a physics-based generative model that we call the **Generative Invariance Network** (GIN).

Our proposed GIN model consists of two parts, illustrated in Fig. 1(b). The key innovation is to replace the discriminator in a traditional GAN model with a function that verifies whether the generated samples obey known physical invariances. We use such a mathematical description the *Invariance Checker* function.

The first step in training our model is the *calibration step*. Here, the invariance checker estimates, from the training data, the parameters of the invariances that need to be verified. For example, in our application, the checker calculates the target p_1 and p_2 values for all the training image, and fine-tunes (calibrates) threshold parameters p_1^* and p_2^* . The second step is standard back-propagation, where the calibrated invariance checker produces a loss value that can be used to train the weights of the generator. More specifically, the checker calculates the p_1 and p_2 values for each image generated by the generator (G), and compares it against the thresholds set in the calibration step. A mismatch in these values can be used to define the loss. In the simplest form, this can be some sort of weighted ℓ_2 -loss, defined as:

$$L_{inv} = \lambda_1 \sum_i |p_1^{(g)} - p_1^*| + \lambda_2 \sum_i \|p_2^{(g)} - p_2^*\|_2, \quad g \in \text{set of images generated by generator, } \mathcal{G}. \quad (4)$$

This loss can be back-propagated to train the generator weights. After sufficient training, generator is expected to produce images with statistical properties closely matching with the calibrated invariances. In Eq. 4, the parameters λ_1 and λ_2 can be used to trade-off the fidelity of the generated samples to the target invariances. Therefore, we obtain fine-grained control over the physical properties for the outputs, compared to traditional GANs.

The benefits of being able to incorporate physical invariances into generative model training are two-fold. First, we obtain superior control over features of the target images. Second, there is no requirement of training a full-fledged discriminator, but rather, only calibrating the parameters of the invariance-checker. That means that such generative models can be trained even with substantially less training data.

2.3 Hybrid (GAN+GIN) model

We propose a third approach that combines the best aspects of both the previous models into a Hybrid (GAN+GIN) model.

Similar to GIN, the hybrid model too proceeds in two steps: a calibration step followed by training step. Calibration step is essentially the same as the GIN, where the invariances are calibrated. However, in the training step, instead of solely relying on either discriminator or the invariance checker, we use a combination of both to train the generator. Two separate optimizers are run to train the generator, one each for minimizing the discriminator loss and the invariant loss. The optimization procedures remain same as in the cases of GAN and GIN respectively. Similar to GAN, we update the generator and discriminator alternatively, with the generator getting updated two times in every iteration: once for minimizing discriminator loss, and once for minimizing the invariance loss (4).

3 Experiments

3.1 Cahn-Hilliard Dataset

We use the Cahn-Hilliard equation [3] to generate microstructures for training and testing. Originally proposed to study phase separation in alloys, this equation can account for such phenomena in polymers, ceramics, and other material systems. We use an in-house modular finite element software [17] to solve this equation to generate time evolving microstructures. Independent solutions were obtained for a range of system parameters like volume fraction (ratio of black pixels to white pixels) and immiscibility parameters (that determine the degree of purity of domains). The simulation domains are square in shape and have physical dimensions of 100 units each. This data set is a modified version of the data used for [18].

3.2 Experiments on WGAN-GP model

3.2.1 Architecture

In all of our WGAN-GP experiments, we use the ResNet-type architecture from [9] with some modifications. In our architecture, both the generator and the discriminator are made of 4 residual blocks, each block consisting two 3×3 convolutional layers with the ReLU nonlinearity. Input to generator is 128-dimensional random Gaussian vector \mathbf{z} . In the generator, some residual blocks perform nearest neighbor upsampling, while in the case of discriminator, some residual blocks perform downsampling through average pooling. Batch normalization is applied only in the case of generator. For optimization, we use Adam with learning rate 1×10^{-4} , and batch size 32. Dimensions of the input images and output samples are $64 \times 64 \times 1$. The hyper-parameter for gradient penalty term, λ in Eq. 3 is set to 10. The number of discriminator updates per generator update is set to 5.

3.2.2 Qualitative results and comparisons

Using the Cahn-Hilliard (CH) dataset, we prepare two smaller datasets referred as CH_{p_1} and CH_{p_2} by segregating the images based on their p_1 and p_2 values respectively. The first subset CH_{p_1} is a collection of all images of CH dataset having a volume fraction (p_1) value between 0.35 to 0.45. The second subset is segregated on the basis of p_2 values of images, and contain all the images from CH dataset having 2-point correlation (p_2) value equal to 0.0625. We train 3 WGAN-GP using CH, CH_{p_1} and CH_{p_2} as the training data respectively. As these segregated datasets typically contain images with similar statistical properties (either p_1 or p_2), we can testify the ability of our model to preserve such properties by observing the p_1 or p_2 values of the images generated by these 3 networks.

To validate our model’s capability of correctly translating the statistical properties of the training data into the generated images, we provide various qualitative results for 3 WGAN-GP networks each trained on CH, CH_{p_1} and CH_{p_2} respectively. The model architecture and other hyperparameters are kept same for all 3 networks. In Fig. 2 (b), (c) and (d) we depict the samples generated by all 3 networks respectively. Qualitatively, we can conclude that these results maintain the diversity within the dataset, while preserving the uniformity in the statistical properties. We further provide the density plots/histograms of p_1 or p_2 values of the images for both real data and generated data.

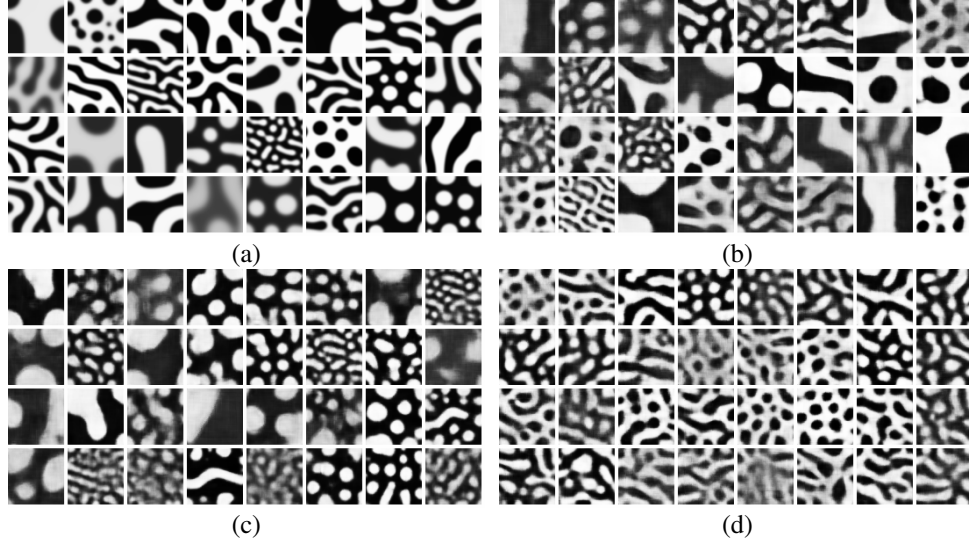


Figure 2: (a) Sample images from Cahn-Hilliard dataset; (b) samples generated by WGAN-GP trained on CH-dataset; (c) Samples generated by WGAN-GP trained over the morphologies from CH_{p1} dataset (only includes the images with volume fraction between 0.35 to 0.45); (d) Samples generated by WGAN-GP trained over the morphologies from CH_{p2} dataset (only includes the images with 2-point correlation equal to 0.0625).

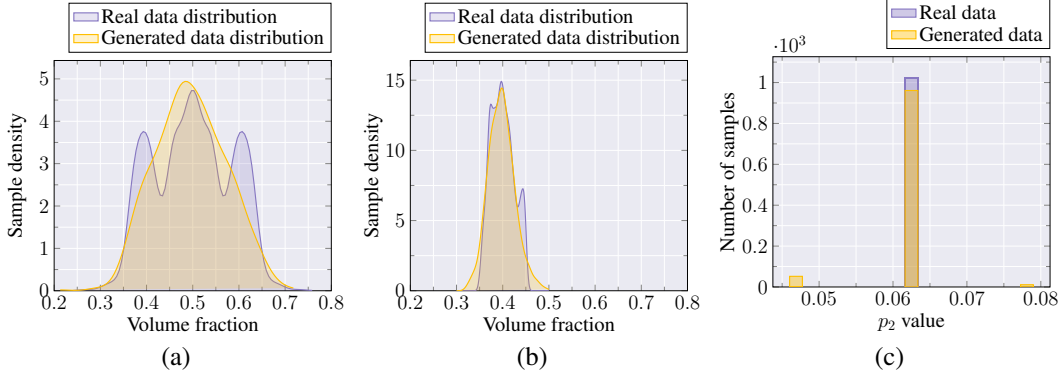


Figure 3: (a) Comparisons of the distributions of volume fractions of training dataset and that of the samples generated by WGAN-GP trained over entire CH dataset; (b) Comparisons of the distributions of volume fractions of training dataset and that of the samples generated by WGAN-GP trained over the CH_{p1} dataset; (c) Histograms of p_2 correlation values for samples from CH_{p2} dataset and samples generated by WGAN-GP trained over CH_{p2} dataset.

The striking similarities in the spread of both the density plots/histograms suggest that our network successfully reproduces the statistical properties of the real (training) images in the simulated images. In Fig. 3 (a,b), densities of p_1 value is compared between the real data and generated data for network trained on CH dataset and CH_{p1} dataset. We provide the histogram for p_2 values for models trained using CH_{p2} in Fig. 3 (c). Both the histograms closely match.

3.2.3 Interpolation results

Further, we display interesting behavior of the learned image manifold through results of interpolation over latent vectors \mathbf{z} . For the first WGAN-GP network trained over the full CH dataset, we randomly pick two different noise vectors \mathbf{z}_1 and \mathbf{z}_2 , and linearly interpolate between them to obtain 10 more

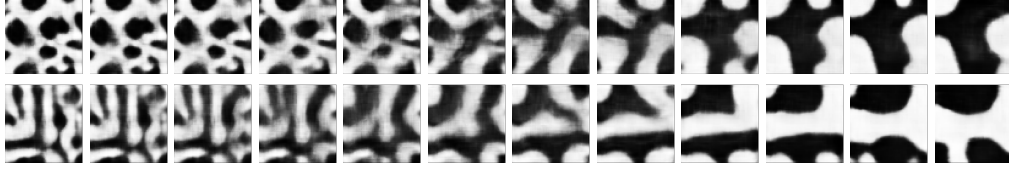


Figure 4: (a) Results of the linear interpolation over latent variable z for WGAN-GP trained over the entire CH dataset

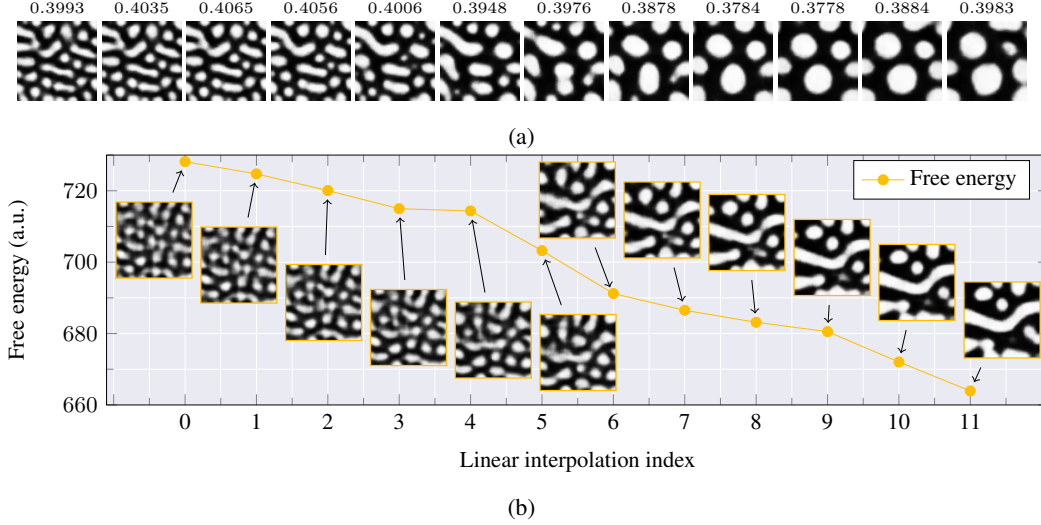


Figure 5: (a) Results of the linear interpolation over latent variable z for WGAN-GP trained over the images from CH_{p1} dataset (only includes the images with volume fraction between 0.35 to 0.45); volume fraction values are printed above each image; (b) Free energy of the interpolated morphologies (volume fraction of each image is $(0.41 \pm 5\%)$).

such noise vectors. All 12 images are plotted as in Fig. 4. It is noticeable that the image manifold appears interpretable, providing a smooth transition from one morphology to another.

For networks trained on CH_{p1} , we repeat similar experiments, except that here we make sure that the initial and final images of the interpolation have nearly the same values of p_1 (volume fraction). It is to verify the continuity of the learned manifold with respect to the volume of each material. Fig. 5 (a) displays the interpolation results for the p_1 case. It is evident that the volume fraction (p_1) values remain almost uniform throughout the interpolation. Such behavior can effectively be used to decide series of manufacturing processes for obtaining final morphology from initial morphology without adding any new material. Very interestingly, we observe that an unseen invariance of energy minimization is captured. As shown in Fig. 5 (b), the free energy of the morphologies decreases as the we move from one interpolation step to the next. In this process, the volume fraction (p_1) is also preserved within reasonable limits. This suggests that the WGAN framework is able to learn latent physical rules from the dataset.

3.3 Experiments on Generative Invariant Network (GIN)

3.3.1 Architecture

The input to the generator is a 64×64 square matrix of random noise (\mathbf{z}) bounded in $[-1, 1]$. \mathbf{z} is fed into a combination of convolution layer and max-pooling layer. We use two such combinations followed by a batch normalization layer in each combination. The first convolution layer has 32 filters with a filter size of 4×4 , while the second layer contains 16 such filters. A max-pooling layer of size 4×4 with a stride of 2×2 is applied to each filter, which reduces the size of the representation. It helps in eliminating the dependence in the representations. The output of the two convolutional

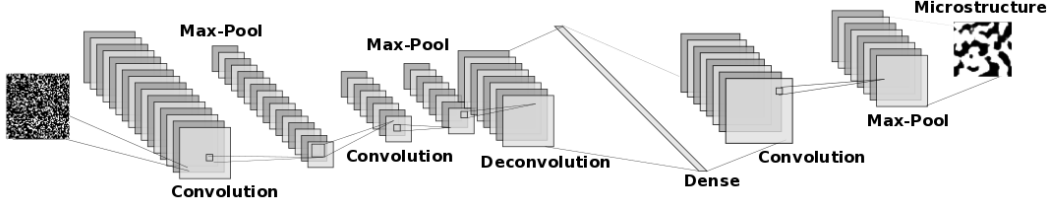


Figure 6: Architecture of the generator (G) for Generative Invariant Network

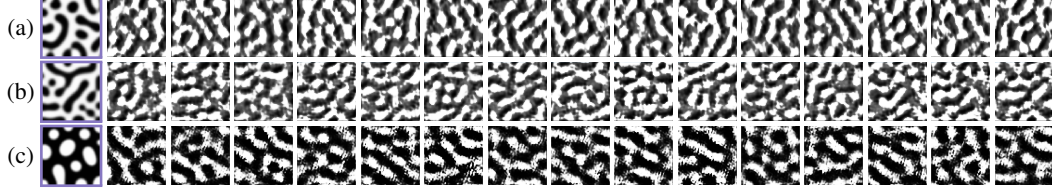


Figure 7: Images generated by GIN models; with first image in each row being the real image used for calibration. (a,b) are trained over entire p_2 curve while model in (c) used only the initial portion of p_2 curve.

layers is passed to a deconvolution layer, which has 16 filters and filter size of 4×4 . Its output is then fed to linear layer of 128×128 neurons followed by a batch normalization layer. The linear layer collects and aggregates the important information from the local features collected by the previous layers. This is further distributed by means of convolution and max pooling layers with 32 filters and a filter size of 4×4 . In the final layer, the architecture generates an image of desired size. To maintain consistency between the other models the size of the output images are kept at 64×64 . In the calibration step, a single image from the CH dataset is used for calibrating the GIN invariances. In training step, generator is trained for 15000 epochs using Adam optimizer with learning rate equal to 10^{-5} . As $p_2(r=0) = p_1$, in our training, the loss function in Eq. 4 is reduced just to the second term, *i.e.* the term corresponding to p_2 value. Further, we consider p_2^* to be a vector containing all p_2 values corresponding to the varying distance (r).

3.3.2 Qualitative results and comparisons

Synthetic microstructures are generated using random variables in space bounded in $[-1, 1]$. In Fig. 7 we show the samples generated by three GIN models, each calibrated with a different image from the CH dataset. Real image used for calibration of the invariances are depicted with blue box. Two of these GIN models (a,b) are trained by employing the entire p_2 curve of the training image as the invariance, while in the third model (c), we use only the initial part of the p_2 curve, as the initial values play a dominating role in deciding the physical and chemical properties of the material. Visually, generated images have similar domains sizes as the original image and exhibit similar spatial distribution patterns. We plot the 2-point correlation curves of the generated samples for GIN models and compare it with the 2-point correlation curve of the real image used for the calibration. Fig. 8 (a,b) depicts such curves for GIN models (a,b) trained with entire p_2 curve. Fig. 8 (c) depicts such curve for GIN model (c) trained with only the initial values of the p_2 curve. It is evident that the p_2 curves of the generated images conform to the calibrated p_2 curves closely. In es sense we capture three descriptors of the microstructure by capturing 2-point correlation. The lowest order descriptor is the volume fraction, which is equal to p_1 . In other words composition of the generated microstructures is very close to the original composition and varies in the range $p_1 \pm 5\%$. The second descriptor is the surface area, which is equal to $\frac{dp_2(0)}{dr}$ and graphically represents the slope of the curve near $r=0$. This values is also close with the value of the original image in all 3 cases.

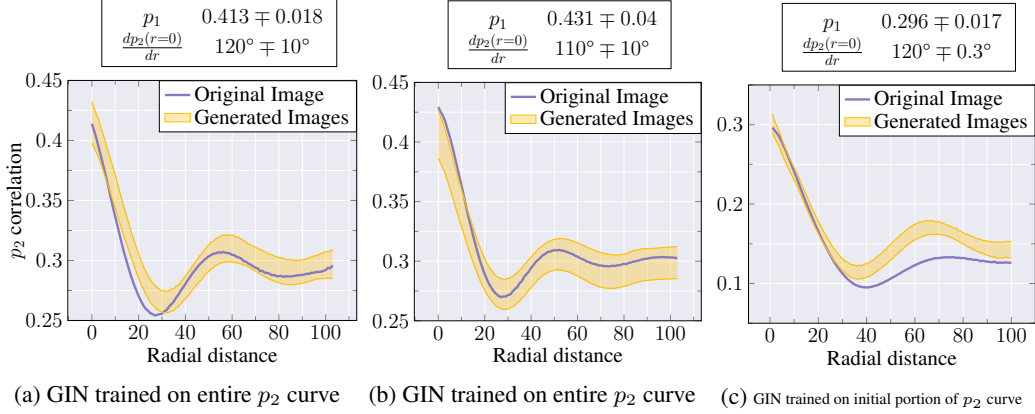


Figure 8: Comparisons of 2– point correlation curves between the images generated by GIN models and the target image.

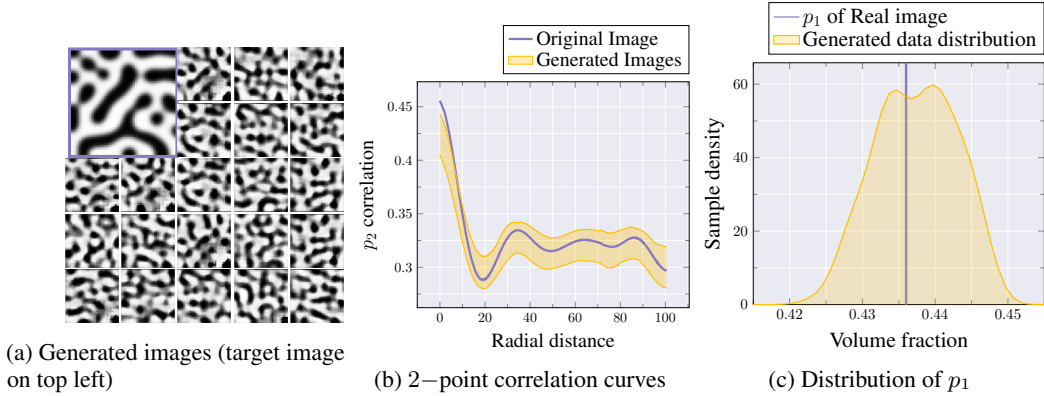


Figure 9: Comparisons of volume fraction (p_1) distribution and 2– point correlation curves between the images generated by hybrid GAN and the target image.

3.4 Experiments on Hybrid (GAN+GIN) model

3.4.1 Architecture

The hybrid (GAN+GIN) model combines both the adversarial discriminator and invariance checker in the training of the generator. The architectures of the generator (G) and the adversarial discriminator (D) are kept same as the WGAN-GP case (described in section 3.2.1). The invariant checker is calibrated in a similar fashion as the GIN case.

3.4.2 Qualitative results

The results for the GAN+GIN model are shown in Fig. 9. The first image on the top right represents the original image and all the other images are synthetically generated by the hybrid model. The similarities are visually evident. Moreover, we also observe sufficient diversity in the simulated images. We further validate the efficacy of the model via the values of statistical properties as shown in Fig. 9. We compare the value of p_1 and p_2 image descriptors of the synthetic images with the original image. The volume fraction of the original image is 0.436, and the distribution of volume fraction of synthetic images is in the range [0.418, 0.452]. Also, the distribution is symmetric around the mean and it follows a Gaussian distribution with mean at ≈ 0.436 . Similar values are observed from the p_2 curve shown in Fig. 9(b). The p_2 curves of the synthetic images closely match the p_2 curve of the original image. Thus, the hybrid approach combines advantages from both the previous models, and is easily extendable to include higher order descriptors of the microstructures.

References

- [1] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- [2] H. K. D. H. (Harshad Kumar Dharamshi Hansraj) Bhadeshia and R. W. K. (Robert William Kerr) Honeycombe. *Steels : microstructure and properties*. Butterworth-Heinemann, 2017.
- [3] 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.
- [4] Junxi Feng, Qizhi Teng, Xiaohai He, and Xiaohong Wu. Accelerating multi-point statistics reconstruction method for porous media via deep learning. *Acta Materialia*, 159:296–308, oct 2018.
- [5] Baskar Ganapathysubramanian and Nicholas Zabarar. Modeling diffusion in random heterogeneous media: Data-driven models, stochastic collocation and the variational multiscale method. *Journal of Computational Physics*, 226(1):326–353, sep 2007.
- [6] Baskar Ganapathysubramanian and Nicholas Zabarar. A non-linear dimension reduction methodology for generating data-driven stochastic input models. *Journal of Computational Physics*, 227(13):6612–6637, jun 2008.
- [7] Ramiro García-García and R. Edwin García. Microstructural effects on the average properties in porous battery electrodes. *Journal of Power Sources*, 309:11–19, mar 2016.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [9] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.
- [10] James Chen-Min Li. *Microstructure and Properties of Materials, Volume 2*. WORLD SCIENTIFIC, oct 2000.
- [11] Jonathan Rivnay, Stefan C B Mannsfeld, Chad E Miller, Alberto Salleo, and Michael F Toney. Quantitative Determination of Organic Semiconductor Microstructure from the Molecular to Device Scale. *Chemical reviews*, 2012.
- [12] Anthony P Roberts. Statistical reconstruction of three-dimensional porous media from two-dimensional images. *Physical Review E*, 56(3):3203, 1997.
- [13] Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. Microstructures to control elasticity in 3d printing. *ACM Transactions on Graphics (TOG)*, 34(4):136, 2015.
- [14] Pejman Tahmasebi, Ardeshtir Hezarkhani, and Muhammad Sahimi. Multiple-point geostatistical modeling based on the cross-correlation functions. *Computational Geosciences*, 16(3):779–797, 2012.
- [15] Salvatore Torquato. *Random heterogeneous materials: microstructure and macroscopic properties*, volume 16. Springer Science & Business Media, 2013.
- [16] Hongtao Wang, Z Zak Fang, and Pei Sun. A critical review of mechanical properties of powder metallurgy titanium. *International Journal of Powder Metallurgy (Princeton, New Jersey)*, 46(5):45–57, 2010.
- [17] Olga Wodo and Baskar Ganapathysubramanian. Modeling morphology evolution during solvent-based fabrication of organic solar cells. *Computational Materials Science*, 55:113–126, apr 2012.
- [18] Olga Wodo, Jaroslaw Zola, Balaji Sesha Sarath Pokuri, Pengfei Du, and Baskar Ganapathysubramanian. Automated, high throughput exploration of process–structure–property relationships using the MapReduce paradigm. *Materials Discovery*, 1:21–28, jan 2015.
- [19] C. L. Y. Yeong and S. Torquato. Reconstructing random media. *Phys. Rev. E*, 57:495–506, Jan 1998.