A Conditional Generative Model for Predicting Material Microstructures from Processing Methods

Akshay Iyer*
Northwestern University
Evanston, IL 60208
akshayiyer2021@u.northwestern.edu

Biswadip Dey
Siemens Corporate Technology
Princeton, NJ 08536
biswadip.dey@siemens.com

Arindam Dasgupta
Siemens Corporate Technology
Princeton, NJ 08536
arindam.dasgupta@siemens.com

Wei Chen
Northwestern University
Evanston, IL 60208
weichen@northwestern.edu

Amit Chakraborty
Siemens Corporate Technology
Princeton, NJ 08536
amit.chakraborty@siemens.com

Abstract

Microstructures of a material form the bridge linking processing conditions - which can be controlled, to the material property - which is the primary interest in engineering applications. Thus a critical task in material design is establishing the processing-structure relationship, which requires domain expertise and techniques that can model the high-dimensional material microstructure. This work proposes a deep learning based approach that models the processing-structure relationship as a conditional image synthesis problem. In particular, we develop an auxiliary classifier Wasserstein GAN with gradient penalty (ACWGAN-GP) to synthesize microstructures under a given processing condition. This approach is free of feature engineering, requires modest domain knowledge and is applicable to a wide range of material systems. We demonstrate this approach using the ultra high carbon steel (UHCS) database, where each microstructure is annotated with a label describing the cooling method it was subjected to. Our results show that ACWGAN-GP can synthesize high-quality multiphase microstructures for a given cooling method.

1 Introduction

Delineating *Processing–Structure–Property* relationships constitutes a major focus in the design of advanced material systems [Olson, 1997]. While analytical and statistical methods have been successfully used for design of certain materials [Florescu et al., 2009, Fullwood et al., 2010, Lee et al., 2017], the underlying assumptions on homogeneity and isotropy limit their generalizability and transferability to other material systems. To address these challenges, machine learning and data-driven techniques have piqued interest in the material science community. Microstructure reconstruction, by allowing an effective method to understand the high dimensional microstructure space, plays a critical role in computational material design. Prior work along this line has used deep learning to predict material property from microstructure [Cecen et al., 2018, Cang et al., 2018,

*The work was done during this author's internship at Siemens Corporate Technology.

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

Yang et al., 2018], reconstruct statistically equivalent microstructures [Li et al., 2018], and synthesize microstructures with desired properties [Yang et al., 2018, Cang et al., 2018].

Generative models, such as variational autoencoders (VAE) [Kingma and Ba, 2014] and generative adversarial networks (GAN) [Goodfellow et al., 2014], are key enablers of deep learning based microstructure reconstruction. Cang et al. [2018] used VAEs to synthesize two-phase microstructures and demonstrated that convolutional networks can be used for material property prediction. Yang et al. [2018] used deep convolutional GAN to synthesize microstructures and transfer learning to improve structure-property predictions. Both of these works augmented the generative model loss function with style transfer and mode collapse losses. Singh et al. [2018] leveraged WGAN-GP and used generative invariance checker & discriminator concurrently to generate two-phase microstructures.

However, this line of work on using generative models for microstructure reconstruction has two limitations. First, previous works have focused on two-phase microstructures, while many material systems comprise multiphase microstructures. Characterization and reconstruction of multiphase materials have been studied scarcely, especially due to the fact that evaluating higher order correlation for multiphase materials is challenging. To the best of our knowledge, transfer learning technique by Li et al. [2018] is the only method that has been used to reconstruct multiphase microstructures. This method uses first few convolutional layers of a pre-trained VGG16 [Simonyan and Zisserman, 2014] network to minimize difference in Gram Matrices of the target and reconstructed image. Although accurate, this method can only reconstruct images that match a single target microstructure and cannot model a distribution in the way generative models do. Second, previous works do not account for the influence of processing conditions on microstructure. This is a key aspect in material design since we strive to not only design an optimal microstructure but also identify the processing conditions necessary to manufacture it.

We address these two challenges by developing an auxiliary classifier Wasserstein GAN with gradient penalty (ACWGAN-GP) to synthesize multiphase alloy microstructures from user defined processing methods and demonstrate this approach using the Ultra High Carbon Steel Database [DeCost et al., 2017]. Modelling this dataset is extremely challenging owing to the multiphase, heterogeneous microstructures it contains. The key contributions of this work are:

- We demonstrate that GANs can synthesize multiphase microstructure images.
- We demonstrate that ACWGAN-GP enables us to condition the generator on a critical processing condition, namely the cooling method.
- In the absence of standardized tools to characterize multiphase microstructures and evaluate quality of synthesized images, we propose using a VGG16-based feature extractor and t-SNE to validate the proposed approach.

2 Learning Framework

2.1 Auxiliary Classifier Wasserstein GAN with Gradient Penalty (ACWGAN-GP)

To learn underlying latent distributions from high-dimensional data, such as images, GANs formulate the learning problem as a two-player zero-sum game between a generator (which tries to generate synthetic images indistinguishable from the real ones) and a discriminator (which tries to distinguish whether an image is real or has been synthesized). Furthermore, by imposing additional structure into the GAN latent space, conditional generative models provide an efficient means to better control features in synthesized samples. Conditional GAN [Mirza and Osindero, 2014], by providing side information (e.g. class labels) to both generator and discriminator, proposes an implementation of this approach and subsequently improves visual quality and diversity of the synthesized images. A later work by Odena et al. [2017] introduces the auxiliary classifier GAN (AC-GAN) architecture which, in addition to using class labels for synthesizing class conditional image samples, also includes a classifier which predicts class labels for images. This current work expands upon this architecture to synthesize microstructures from given values of processing parameters.

After GAN was introduced by Goodfellow et al. [2014], several improvements have been proposed to achieve stable training and fast convergence. Wasserstein GAN [Arjovsky et al., 2017], by using earth-mover distance (Wasserstein-1 metric) as a geometrically meaningful measure of mismatch between probability distributions, introduced a loss function that correlates with the quality of synthesized images. A later work by Gulrajani et al. [2017] has shown that penalizing the L_2-norm of discriminator gradient, in stead of constraining the discriminator weights, provides a better

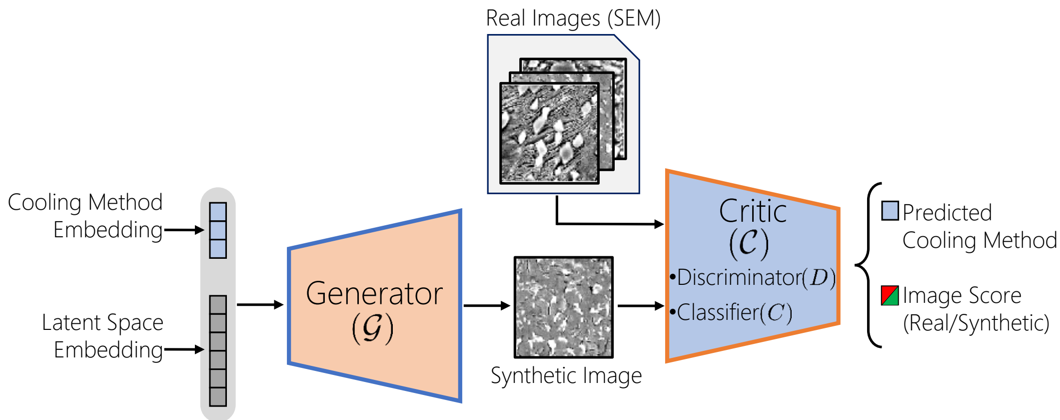

Figure 1: *ACWGAN-GP framework for learning conditional generative models to predict multi-phase microstructures from cooling methods.*

alternative to enforce the Lipschitz constraint. This change in the discriminator training results in faster convergence as well as overall improvement in the quality of synthesized images.

In this work, to synthesize alloy microstructures conditioned on processing methods, we develop a hybrid framework that uses label conditioning suggested in AC-GAN architecture but leverages Wasserstein-1 metric to design loss functions. Additionally, we use gradient penalty instead of weight clipping for discriminator training. In particular, the generator ($\mathcal{G}$) and the critic ($\mathcal{C}$) in this ACWGAN-GP framework tries to minimize the following loss functions

$$L_{\mathcal{G}} = \underbrace{-\mathop{\mathbb{E}}_{\tilde{\mathbf{x}}\sim\mathbb{P}_G}\left[D(\tilde{\mathbf{x}})\right]}_{\text{Generator loss}} + \underbrace{\lambda_2 \mathop{\mathbb{E}}_{\tilde{\mathbf{x}}\sim\mathbb{P}_G}\left[-\log\left(\text{Prob}\left[C(\tilde{\mathbf{x}}) = c_{\tilde{\mathbf{x}}}\right]\right)\right]}_{\text{Sparse cross-entropy over synthetic images}}$$

and

$$L_{\mathcal{C}} = \underbrace{\mathop{\mathbb{E}}_{\tilde{\mathbf{x}}\sim\mathbb{P}_G}\left[D(\tilde{\mathbf{x}})\right] - \mathop{\mathbb{E}}_{\mathbf{x}\sim\mathbb{P}_r}\left[D(\mathbf{x})\right]}_{\text{Discriminator loss}} + \lambda_1 \underbrace{\mathop{\mathbb{E}}_{\hat{\mathbf{x}}\sim\mathbb{P}_{\hat{x}}}\left[(\|\nabla_{\hat{x}} D(\hat{\mathbf{x}})\|_2 - 1)^2\right]}_{\text{Gradient-penalty}} + \lambda_2 \underbrace{\mathbb{E}\left[-\log\left(\text{Prob}\left[C(x) = c_x\right]\right)\right]}_{\text{Sparse Cross-entropy over all images}},$$

where, $\lambda_1 = 10$, $\lambda_2 = 1$, c_x is the true label associated with an image (both real and synthesized) x, and $\mathbb{P}_r$ and $\mathbb{P}_G$ represent the distributions of real and synthesized images, respectively. Moreover, $\mathbb{P}_{\hat{x}}$ denote the distribution of samples which have been drawn uniformly along straight lines between pairs of images sampled from $\mathbb{P}_r$ and $\mathbb{P}_G$.

2.2 Network architecture and training

In this work, we represent the cooling method via a 20-dimensional embedding vector (learned during training) and concatenate it with a 100-dimensional Gaussian noise vector. This 120-dimensional combined vector is then supplied to the generator $\mathcal{G}$. The first layer of $\mathcal{G}$ is a Fully connected network with 1024 neurons, followed by a dropout layer with rate = 0.25. This is followed by three Upsampling-Convolution-Leaky ReLU blocks and a final Upsampling-Convolution block with `tanh`-activation. The critic $\mathcal{C}$ is an approximate mirror image of the generator with four Convolution-Leaky ReLU blocks. These operations extract a 1024-dimensional feature vector which is then passed through a dropout layer with rate = 0.25 and subsequently used by two separate fully connected layers to determine the image score and the corresponding cooling method.

3 Experiment

3.1 The Ultra High Carbon Steel DataBase (UHCSDB)

UHCSDB is a collection of 961 microstructures obtained from Scanning Electron Microscopy (SEM) of samples with identical composition but subjected to varied heat treatments. The variation in heat treatment influences the microstructure and relevant properties. SEM microstructures curated

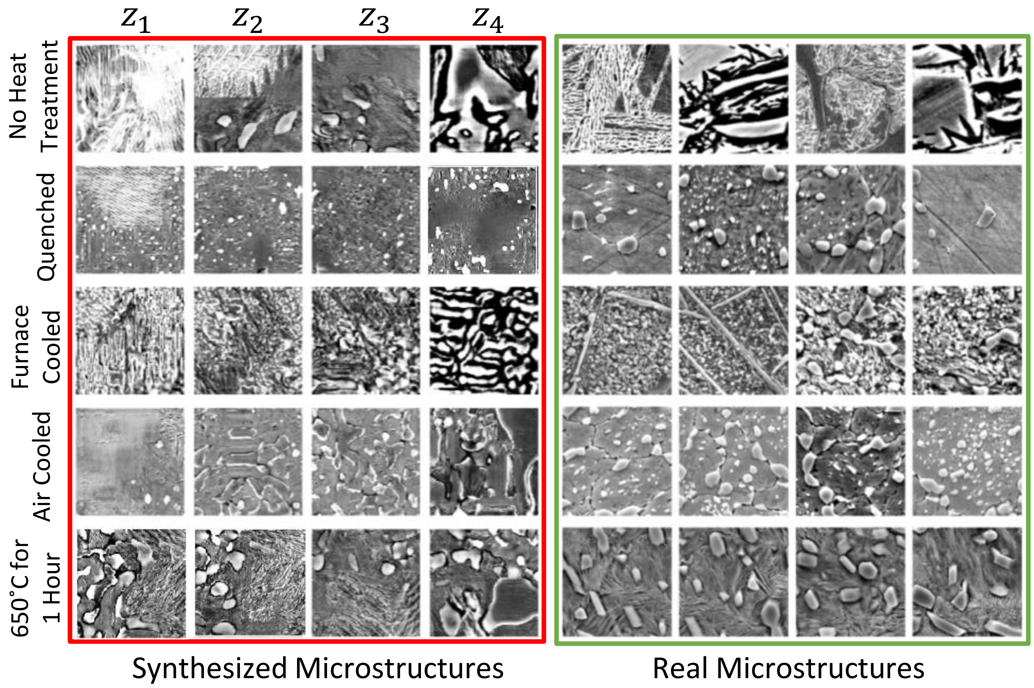

Figure 2: *Visual comparison of synthesized microstructures (after 6000 epochs) with the real ones.*

in UHCSDB correspond to 5 different cooling methods (*no heat treatment*, *quenching*, *furnace cooling*, *air cooling* and *constant heating at* $650°C$ *for 1 Hour*). Within this database, we focus on microstructures captured at a magnification of 10.3 pix/micron. To train our ACWGAN-GP model to synthesize microstructures of size 128x128 pixels, we created a dataset containing cropped images from the original 172 SEM microstructures in UHCSDB. We observed imbalance in the number of images as well as levels of variation in microconstituents among the different cooling methods. Consequently, we adjust the number of cropped segments attained from the SEM microstructures in such a way that cooling methods with higher microconstituent variation have more images. Cropping yields a dataset of 7865 images (*no heat treatment*-2850; *quenching*-1400; *furnace cooling*-1365; *air cooling*-1150; *constant heating at* $650°C$ *for 1 Hour*-1100). Finally, to build the training dataset, we augment this further by rotating the cropped images by 90, 180 and 270 degrees.

3.2 Results

We train the ACWGAN-GP for 6000 epochs with a batch size of 64 and use Adam Optimizer [Kingma and Ba, 2014] with a learning rate of 5×10^{-5}, while keeping $\beta_1 = 0.5$ & $\beta_2 = 0.9$. Also, $\mathcal{C}$:$\mathcal{G}$ training ratio is maintained at 5:1. Figure 2 shows a few representative images synthesized by ACWGAN-GP along with real microstructures for visual comparison. Each column of synthesized images correspond to the same Gaussian noise vector z_i while each row corresponds to a specific cooling method. Although this figure highlights the visual resemblance between real and synthesized microstructures, we present rigorous evaluation methods in what follows.

Furthermore, to perform a quantitative evaluation of the synthesized microstructures, we randomly select 4000 pairs of real and synthesized images and compare their 2-point spatial correlations [Yeong and Torquato, 1998]. Figure 3(a) shows a good match between the correlation values of the real and the synthesized microstructures. However, for complex, multiphase microstructures, this is only a necessary condition, not sufficient. We provide an alternative evaluation which first extracts appropriate feature vectors using a pre-trained VGG16 and then uses t-SNE based dimensionality reduction to project these vectors onto a 2-dimensional space. We applied this procedure to 4000 pairs of real & synthesized images over the course of training to obtain Figs. 3(b-d). After 1 epoch, the generated microstructures do not resemble the training dataset which leads to high dissimilarity in features extracted from VGG16. Consequently, tSNE places the real and generated images in different

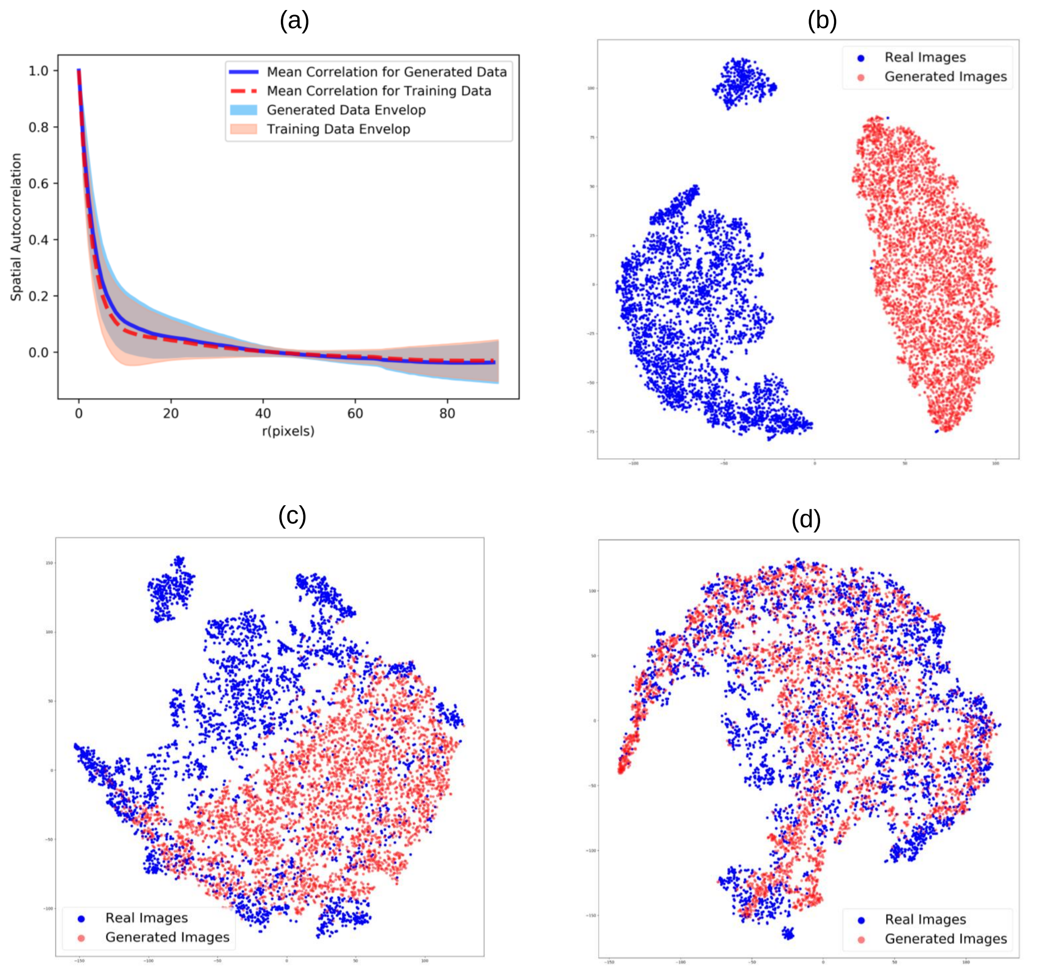

Figure 3: *Validation of the trained ACWGAN-GP model using (a) 2-point spatial correlation and (b-d) 2-dimensional projection of a t-SNE based embedding associated with real and synthesized microstructures after (b) 1, (c) 10 & (d) 6000 epochs, respectively.*

regions(Fig. 3b). However, the quality of generated microstructures improves as training proceeds, and a well-trained model shows significant overlap between real and generated microstructures (as shown in Fig. 3d), indicating visual similarities between them.

4 Conclusion

To the best of our knowledge, this is the first attempt to model processing-structure relationship as a conditional image synthesis problem. To accomplish this objective we have utilized the ACWGAN-GP framework, which inherits training stability of WGAN-GP and conditional image generation of ACGAN. We use this framework for reconstructing multiphase microstructures from UHCSDB and demonstrate its capability in synthesizing high quality microstructures from a given cooling method. In effect, ACWGAN-GP establishes the Process-Structure linkage which represents one-half of the *Processing–Structure–Property* relationship needed for design of advanced material system. In the future, we will introduce additional processing parameters, such as annealing temperature and time, to achieve tighter control and better insight over the synthesized microstructures.

References

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.

Ruijin Cang, Hechao Li, Hope Yao, Yang Jiao, and Yi Ren. Improving direct physical properties prediction of heterogeneous materials from imaging data via convolutional neural network and a morphology-aware generative model. *Computational Materials Science*, 150:212–221, 2018.

Ahmet Cecen, Hanjun Dai, Yuksel C Yabansu, Surya R Kalidindi, and Le Song. Material structure-property linkages using three-dimensional convolutional neural networks. *Acta Materialia*, 146: 76–84, 2018.

Brian L DeCost, Matthew D Hecht, Toby Francis, Bryan A Webler, Yoosuf N Picard, and Elizabeth A Holm. UHCSDB: UltraHigh carbon steel micrograph database. *Integrating Materials and Manufacturing Innovation*, 6(2):197–205, 2017.

Marian Florescu, Salvatore Torquato, and Paul J Steinhardt. Designer disordered materials with large, complete photonic band gaps. *Proceedings of the National Academy of Sciences*, 106(49): 20658–20663, 2009.

David T Fullwood, Stephen R Niezgoda, Brent L Adams, and Surya R Kalidindi. Microstructure sensitive design for performance optimization. *Progress in Materials Science*, 55(6):477–562, 2010.

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.

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.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Won-Kyu Lee, Shuangcheng Yu, Clifford J Engel, Thaddeus Reese, Dongjoon Rhee, Wei Chen, and Teri W Odom. Concurrent design of quasi-random photonic nanostructures. *Proceedings of the National Academy of Sciences*, 114(33):8734–8739, 2017.

Xiaolin Li, Yichi Zhang, He Zhao, Craig Burkhart, L Catherine Brinson, and Wei Chen. A transfer learning approach for microstructure reconstruction and structure-property predictions. *Nature Scientific Reports*, 8, 2018.

Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2642–2651, 2017.

Gregory B Olson. Computational design of hierarchically structured materials. *Science*, 277(5330): 1237–1242, 1997.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Rahul Singh, Viraj Shah, Balaji Pokuri, Soumik Sarkar, Baskar Ganapathysubramanian, and Chinmay Hegde. Physics-aware deep generative models for creating synthetic microstructures. *arXiv:1811.09669*, 2018.

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):111416, 2018.

CLY Yeong and Salvatore Torquato. Reconstructing random media. *Physical Review E*, 57(1):495, 1998.