

AlloyGAN: Domain-Promptable Generative Adversarial Network for Generating Aluminum Alloy Microstructures

Biao Yin^{*1}, Yangyang Fan^{*2}, Nicholas Josselyn¹, Elke Rundensteiner^{1,3}

¹Department of Data Science, Worcester Polytechnic Institute, Worcester, MA, USA

²Chief Material Scientist, DeepAlloy, Portland, Oregon, USA

³Department of Computer Science, Worcester Polytechnic Institute, Worcester, MA, USA

Abstract—The global metal market, expected to exceed \$18.5 trillion by 2030, faces costly inefficiencies from defects in alloy manufacturing. Although microstructure analysis has improved alloy performance, current numerical models struggle to accurately simulate solidification. In this research, we thus introduce AlloyGAN – the first domain-driven Conditional Generative Adversarial Network (cGAN) involving domain prior for generating alloy microstructures of previously not considered chemical and manufactural compositions. AlloyGAN improves cGAN process by involving prior factors from solidification reaction to generate scientifically valid images of alloy microstructure given basic alloy manufacturing compositions. It achieves a faster and equally accurate alternative to traditional material science methods for assessing alloy microstructures. We contribute (1) a novel AlloyGAN design for rapid alloy optimization; (2) unique methods that inject prior knowledge of the chemical reaction into cGAN-based models; and (3) metrics from machine learning and chemistry for generation evaluation. Our approach highlights the promise of GAN-based models in the scientific discovery of materials. AlloyGAN has successfully transitioned into an AIGC startup with a core focus on model-generated metallography. We open its interactive demo at: <https://deepalloy.com/>

I. INTRODUCTION

Background. The global metal market, crucial for industries like construction and aerospace, hinges on alloys for their reliable properties. With a projected worth of \$18.5 trillion by 2030, alloy manufacturing is a key market driver [1].

Motivation. Alloy development is crucial yet expensive, with a high rate of product rejections, which must be re-melted and re-cast due to various defects. Annually, the industry sees tens of millions of tons of metal casting products fall into this cycle [2]. Thus, efficient and reliable simulation models for accelerating scientific alloy discovery and manufacturing will bring immense economic and environmental benefits.

Challenges. Traditional numerical methods have struggled with the complex solidification process in alloy formation, characterized by vast nonlinear chemical and physical interactions [3]–[5]. This complexity challenges conventional mathematical modeling attempting to accurately simulate the final microstructure based on the basic alloy compositions like the process shown in Figure 1. These models are intricate,

^{*}These two authors contribute equally to this work.




Figure 1: AlloyGAN: Promptable generation algorithm serving a novel scientific alloy discovery platform built by an AIGC startup. Red arrows indicate model training, purple for model inference, black solid for material science verification, and black dash for data collection.

computationally demanding, and require specific knowledge of material science, limiting their accessibility to the broader research community [6]–[9].

Problem Definition. The problem thus is the urgent need for more efficient and accurate methods to generate microstructure images of metal alloys, based on initial conditions like chemical composition and manufacturing setting. Training from limited data, these methods should accurately model complicated chemical reactions while ensuring scientific validity and less computational complexity for practical application.

State-of-the-Art. Deep learning techniques like VAEs, GANs, and Diffusions are being explored for microstructural analysis of materials in an early stage [10]–[13]. Their application remains limited however, particularly in generating scientifically valid images based on basic alloy compositions. The existing deep learning methods haven't fully addressed the complexities inherent in material science [14]–[20]. This

void emphasizes the need for innovative solutions capable of dealing with the complex nonlinear dynamics of the alloy microstructure formation and delivering robust generation capabilities to accelerate scientific alloy discovery.

Proposed Method. To bridge this gap, we propose AlloyGAN, a ground-breaking approach leveraging the power of deep learning to create scientifically valid alloy microstructure images from basic alloy compositions (See Figure 1). By integrating domain-specific priors from solidification reactions, AlloyGAN, as a chemically-constrained cGAN architecture, conditions the generation process related to the basic alloy compositions, thereby successfully simulating the complex solidification process to verify alloy properties.

Findings and Impact. Our research results bear witness to the *transformative potential of AlloyGAN* in the process of accelerating the scientific discovery of alloys. Over standard cGAN, AlloyGAN demonstrates its ability to quantify the impact of individual chemical elements while generating the microstructure’s metallography of previously not-considered alloy compositions. In evaluating AlloyGAN, we employ metrics from the machine learning community, such as FID [21], across alloys seen and unseen in the training process. In addition, we also design domain-specific measures confirming its ability to accurately depict various chemical properties that determine any castable alloy microstructure, such as 1) micrograph evolution with Si content, 2) the effect of Cooling Rate, and 3) the effect of modification from Strontium and Phospho process. This makes AlloyGAN a promising tool for expedited research in material science.

Contributions. Our proposed AlloyGAN model introduces significant advancements to alloy manufacturing including:

- **Interdisciplinary Innovation for Material Discovery.** By merging insights from material science with generative models, we significantly improve the efficiency of generating alloy microstructure conditioned with basic alloy compositions on limited domain data – thereby providing a practical alloy discovery platform (Figure 1).
- **Novel Methodologies Injecting Prior Chemical Knowledge to cGAN:** We offer new methodologies that bridge standard cGAN and understanding of the chemical reactions in the alloy solidification process for generating scientifically valid images to reflect alloy microstructure, especially on alloys previously not considered (Figure 3).
- **Domain-Specific Evaluation Metrics:** Not only machine learning metrics, we employ a unique set of domain-specific metrics for assessing the chemical properties of the generated alloy microstructure images, enriching the field with robust evaluation tools.
- **Accessibility and Real-world Impact:** Established by our collaborative startup from the two co-first authors, AlloyGAN is freely accessible via an interactive demo at <https://deepalum.com/>. Our findings may inspire further research in advanced deep-learning models for material innovation and thus catalyze efficiency and sustainability practices in the alloy manufacturing industry.

II. RELATED WORK

Conventional Alloy Microstructure Modelling Approach.

Traditional alloy microstructure modeling demands significant resources and expertise in solving intricate Partial Differential Equations (PDEs) over analytical models and Finite Element Method (FEM) [3]–[5], [22], [23]. Phase Field modeling requires extensive computation – 768 GPUs for 12 days to compute a voxel cell [6], while the Monte Carlo method struggles with aligning simulation and physical time [7], [8]. Cellular Automation (CA) assists in simulating recrystallization and solidification [7], [8] but faces heavy domain challenges with complicated metallurgical processes and transition rules [9]. Despite their mathematical and physical clarity, these conventional methods may fall short in providing prompt, scientifically robust decisions in large-scale alloy discovery.

Material Generation using Deep Learning Approach. Deep learning methods have emerged to explore possible solutions. Notably, GANs [24]–[28], variational autoencoders (VAEs) [29], [30], and convolutional neural networks (CNNs) [10]–[13], [31] have been applied to exploring microstructure segmentation [10], classification [11], and impurity inspection [12]. Using a modified GAN architecture, Tang et al. [14] modeled aluminum microstructures according to the intensity of the laser. GANs have been applied for microstructural image augmentation, such as using a GAN to generate microstructures and predicting heat treatment conditions using a random forest [15]–[18], [32]. Alternatively, our GAN structure utilizes fundamental alloy compositions as conditions to prompt microstructure image generation, thereby facilitating the process of scientific alloy discovery.

III. ALLOY MICROSTRUCTURE DATASET

Dataset and Standard of Scientific Alloy Discovery.

From standard material science experiments collecting alloy microfractography [33], our dataset has 4200 microstructure binary images conditioned on 42 manufactured and labeled alloys. Each alloy has 100 sample images. Each image is 128 by 128 pixels in size. Example images are shown in Figure 2. Each alloy is designed with 6 chemical conditioning features and manufactured under 2 environmental conditioning features – cooling rate and material modification status called Sr/P. In material science, the microstructure of any alloy will




Figure 2: Example images from different alloys in our real-world alloy microstructure image dataset used to train and evaluate AlloyGAN vs. Standard cGAN

be stable if considering these 8 conditioning features [33]. Table I shows a summary of the alloy labels according to the conditioning features. The Alloy ID identifies a designed but not manufactured alloy from different chemical conditions which are signified by the 2nd to 7th columns – the weight percentage (wt.) of Silicon (Si), Iron (Fe), Copper (Cu), Manganese (Mn), and Nickel (Ni). The Cooling Rate indicates the cooling rate of a designed alloy when it solidifies in Kelvin per second (K/s). The last column indicates if the alloy is processed by a fixed amount of Phosphorus or Strontium (Sr/P), additions into the alloy which significantly change the microstructures of the aluminum alloys in complex chemical reactions. An alloy with its composition ID can be processed or not in our dataset. For Si content, the value range is usually less than 22 wt.; Cu contents range less than 4.5 wt.; contents of other chemical elements are not higher than 2.5 wt. The cooling rate value varies between the magnitude of 0.1 K/s to 100 K/s. It is challenging to train an ideal generative model with continuous prompts on the chemical conditions using this real-world domain image dataset due to the availability of limited images and alloys. This issue is typical since casting an alloy sample is labor and time-intensive and domain experts are required for safety control and scientific validation.

The Crucial Role of Silicon in Aluminum-based Alloys. Silicon is crucial for casting alloys, aiding in defect-free solidification while reacting with aluminum. Its unique expansion offsets aluminum’s shrinkage, exhibiting dark particle flaws shown in Figure 2. On the other hand, aluminum displays a bright contrast and its form resembles ‘fingers’, a characteristic known as dendrites. Furthermore, silicon enhances the fluidity of aluminum melts, enabling them to fill complex geometries in casting products such as motor engines.

Data Splits for Training and Evaluating AlloyGAN. To test the generative performance on materials not included in the initial training set, we will remove a subset of the alloys in the dataset from the training process and instead use them as the test set for the alloy discovery. We randomly select 8 of 42 labeled alloys and filter out all the images of these alloys in our dataset to establish the test set of alloys not seen during training. For the remaining alloys utilized for training, we randomly split the images representing each given alloy 80% and 20% as train and validation sets.

IV. THE METHODOLOGY: ALLOYGAN

Our proposed method, AlloyGAN, builds upon the concepts of conditional Generative Adversarial Networks (cGANs), augmented with unique adaptations that incorporate prior knowledge from solidification reaction factors to generate scientifically valid alloy microstructure images. Illustrated in Figure 3, both the AlloyGAN Generator and Discriminator are equipped with a core module where prior chemical reaction factors can be applied to the basic compositions of alloy. This injects prior knowledge with domain-specific conditions into the Generator and Discriminator pair – providing the capability of AlloyGAN for generating the alloy microstructure images with rich informative chemical properties.

Alloy ID	Si	Fe	Cu	Mn	Mg	Ni	Cooling Rate	Sr/P
A356	7	0.5	0.25	0.35	0.3	0	2.5, 10, 57, 143	Yes, No
A360	9.5	0.6	0.1	0.05	0.5	0	2.5, 10, 57, 143	No
A369	11.5	1	0.5	0.25	0.4	0.05	2.5, 10, 57, 143	Yes, No
A339	12	1.2	2	0.5	1	1	2.5, 10, 57, 143	No
A393	22	1.3	0.9	0.1	1	2.3	2.5, 10, 57, 143	Yes, No
A355	5	0.65	1.25	0.55	0.5	0	2.5, 10, 57, 143	No
A308	5.5	0.8	4.5	0.5	0.1	0	2.5, 10, 57, 143	No
A319	6	1	4	0.4	0.55	0.35	2.5, 10, 57, 143	No
A332	9.5	0.9	3	0.5	2.1	0.5	2.5, 10, 57, 143	No

Table I: Summary of labels: Chemical composition with experienced manufacturing environments that together determine Alloy microstructure in our dataset.




Figure 3: AlloyGAN: Pink-colored module illustrates how we incorporate the prompted inputs with the domain prior knowledge for cGAN enhancements for this domain-specific task. The resulting conditions are concatenated with the respective images as additional channels in the workflow. Conv indicates Convolutional layers, while Norm indicates Max-Min normalization on features.

A. The AlloyGAN Generator

The AlloyGAN Generator concatenates two different types of inputs: a random noise tensor and a condition tensor. The noise vector introduces a degree of randomness into the generation process following a standard normal distribution. The condition tensor receives prompted user inputs of basic alloy compositions containing the amount of chemical elements and the manufacturing settings. Before concatenating, these conditions will be processed with prior deterministic factors from the chemical reaction of the alloy solidification process. To produce the desired realistic images, the processed conditions are randomly selected and concatenated with the noise tensor before passing through several transposed convolutional layers. This ensures the prompted conditions upon priors to better fool the Discriminator so that the generated images can be discriminative according to the conditions upon priors. The convolutional layers include batch normalization and leaky ReLU activation functions to stabilize the training process and prevent overfitting. The output of these layers is a 2-dimensional tensor in the same shape as our target binary microstructure image.

B. The AlloyGAN Discriminator

Illustrated in Figure 4, the Discriminator of AlloyGAN serves as a classifier that tries to distinguish between real and generated data. It takes in both real images with their corresponding conditions upon priors and the generated images produced by the Generator and assigns a probability that a generated image can be predicted as a real image.

C. Prior from Chemical Reaction in Alloy Solidification

The incorporation of targeted domain-specific prior knowledge into the prompted conditions is what differentiates AlloyGAN from a standard cGAN. Illustrated in the pink subnet in Figure 3, the integration of the quantifiable prior with the prompted alloy composition conditions is formulated as a concatenation of two types of normalized features:

- 1) **Features of Chemical Elements:** This feature subset accounts for the influence of each chemical element on the generated alloy by multiplying its quantity with a conversion factor in Table II. This factor represents the expected impact of each element on the solidification process critical for the resulting alloy properties. The features were then normalized with a unified scale from 0 to 1 for training stability using max-min normalization.
- 2) **Features of Manufacturing Environments:** This feature subset represents the conditions under which the alloy is manufactured. Key factors like the cooling rate and Sr/P modification are taken into account. Similar to the chemical features, these manufacturing features also undergo a max-min normalization process. This step accounts for potential variations in manufacturing environments.

The above intricate formulations of the prompted alloy compositions with the chemical reaction prior factors capture domain-specific alloy chemical and manufacturing conditions that are critical for AlloyGAN to generate microstructure images taking the complex interplay among these conditions into consideration. This is why conditions not previously in our input data set could be inferred and thus meaningfully generated by our trained AlloyGAN with limited data.

Domain-driven Chemical Element Normalization. The standard cGAN individually normalizes each raw input feature shown in Table I. However, this approach does not consider the relative contributions of each element to the alloy. For example, the Si element greatly changes the alloy microstructures, whereas other elements have less impact. This may be because their content is relatively low (less than 4.5 wt.) compared with Si. Moreover, the chemical reaction outcomes of these elements with aluminum need also be considered. Table II shows the outcome of chemical reactions of each element to aluminum and the factor of converting a unit of weight percentage of elements to the atomic percentage of outcomes. The conversion factor embodies an indicator of the impact of elements on the microstructure metallograph. This is what we thus propose to use as a prior to prompted chemical composition conditions.

Element	Si	Fe	Cu	Mn	Mg	Ni
Outcome	Si	Al3Fe	Al2Cu	Al12Mn	Mg2Si	Al3Ni
Factor	0.036	0.018	0.016	0.018	0.021	0.017

Table II: Summary of chemical reaction outcome and factors in alloy manufacturing: a factor indicates the conversion factor of a unit of weight percentage of the Element to the amount of atomic percentage (at.) of the outcome – formulating the domain prior onto prompted conditions in AlloyGAN.

D. Training Process

AlloyGAN is trained using a two-step iterative process. In the first step, the generator creates a batch of synthetic microstructure images with the processed prior. The later are then evaluated by the discriminator. The generator’s weights are updated to minimize the difference between the discriminator’s output to distinguish the synthetic images from the real images. In the second step, the discriminator’s weights are updated using a batch of real microstructure images from the training alloys and the synthetic images from the first step. The goal of this step is to classify real versus generated images from random noise with conditions normalized with the prior.

This process is repeated for several epochs until the generator can produce synthetic microstructure images that the discriminator can no longer distinguish from real ones. The resulting trained model is then ready to generate new microstructure images for any given alloy composition — including those compositions that had previously not been unseen in the training process.

AlloyGAN training models a two-player minimax game with prior chemical reactions on the prompted alloy compositions for a robust, accurate, and fast generation. The Generator (G) tries to minimize this objective against an adversary D that tries to maximize it, i.e., $\min_G \max_D V(D, G)$. The prompted alloy compositions are represented in a conditioning vector y that is concatenated with G (as y_1 after *Conv*) or D (as y_2 after *Conv*) illustrated in Figure 3 and formulated as:

$$y = \text{Concat}(\text{Norm}(\mathbf{w} * \mathbf{I}_1), \text{Norm}(\mathbf{I}_2)) \quad (1)$$

where Norm denotes min-max normalization, $\mathbf{I}_1$ User Input 1 to prompt the wt.% content of each chemical element, $\mathbf{I}_2$ User Input 2 to prompt manufacturing environments, and $\mathbf{w}$ Domain Prior from chemical reaction factors. The loss function of training AlloyGAN is:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x, y_1 \sim p_{\text{data}}(x, y_1)} [\log D(x|y_1)] + \mathbb{E}_{z \sim p_z(z), y_2 \sim p_{\text{data}}(y_2)} [\log(1 - D(G(z|y_2)))] \quad (2)$$

In Eq. (1), x denotes the sample image, and y_1 the conditioning vector capturing the domain prior illustrated as “Chemical Reaction Factors” in Figure 3. z is a noise vector. $D(x|y_1)$ corresponds to the discriminator’s estimate of the probability that real alloy microstructure image x , given y_1 . $G(z|y_2)$ is the fake alloy microstructure image by the generator, given z and y_2 . The integration of the domain prior factors, $\mathbf{w}$, provides chances to further distinguish the importance of chemical composition prompts, $\mathbf{I}_1$, in AlloyGAN, thereby enhancing cGAN to generate applicable images with robust performance.




Figure 4: Example generated images in the training process. Top row shows standard cGAN generations. Bottom row shows AlloyGAN generations. Each column shows the last batch of the epoch 1st, 20th, 40th, 60th, 80th, 90th, and 100th. The random seed is the same for both models. Upon same hyperparameters, including the random seed, AlloyGAN outperforms cGAN to generate informative microstructure images of alloys.

E. Evaluation Metrics

We evaluate AlloyGAN and cGAN performance in this study using FID [21], a widely-used machine learning metric. FID evaluates the performance of a generative model by looking at how close the generated data is to the real data in the embedding space of an object classification model (commonly using the InceptionV3 model in particular) [21]. We calculate FID between real and generated images on our test set so that we know how the model generates the microstructure images of previously not-trained alloys. To build the embedding space required for the calculation, we not only utilize the pre-trained InceptionV3 network upon the ImageNet dataset as a general space but also finetuned the network according to our domain train dataset as a domain-specific space for the evaluation.

We note however that general evaluation metrics in machine learning, such as, FID, cannot measure the chemical properties of generated alloy microstructure images for scientific use. We thus also utilize three domain-science metrics to evaluate whether AlloyGAN creates valid images that are feasible for fast material verification in alloy discovery.

Micrograph Evolution with Si Content. This metric measures whether the generated images show the trend that Si content in metallography increases with the addition of Si wt.% in alloys. The Si content in metallography can be quantifiable by calculating the area of fractions (black area after binarizing the image).

Effect of Cooling Rate on the Secondary Dendritic Arm Spacing. The theory of computational material science shows that there must be a correlation between the cooling rate and the Secondary Dendritic Arm Spacing in the images satisfying the following formula:

$$\lambda = k(CR)^{-n} \quad (n > 0), \quad (3)$$

where λ represents the secondary dendritic arm spacing measured by the distance between the neighborhood white areas separated by a black arm in the generated image and CR represents the cooling rate (K/s) in the manufacturing

environment. For scientifically valid microstructure images of an alloy, there must be a strong linear relationship between $(CR)^{-n}$ and λ with a constant n and a coefficient k .

Modification of Strontium and Phosphor on the Alloy Microstructure. This metric measures whether Primary Al dendrites are branched after the modification. Sr/P is added to aluminum in very small amounts (<0.05 wt.%), yet they have a significant impact on the microstructures of aluminum alloys. Sr and P atoms act as "poisons" to the solidification of the Si phase, causing it to form more branched and rounded structures. Aluminum alloys with Sr/P modification typically exhibit microstructure images with dispersed dark particles. In contrast, alloys without Sr/P modification display larger and segregated dark Si particles.

V. EXPERIMENTAL STUDY

To compare AlloyGAN and standard cGAN performances, we evaluate their generated images on the same training hyperparameters default from standard cGAN. The latent layer dimension was set to $n_z=100$, allowing for a diverse range of features to be generated in the Generators. The Adam optimizer was utilized for both the generator and discriminator, with a beta value of 0.5. The training was conducted using a batch size of 64, a learning rate of 0.0002, and lasted for 100 epochs which are all recommended settings in standard cGAN. According to our dataset, the images used for training and testing are of size 128x128 pixels, with a single channel ($n_c=1$). We use NVIDIA-A100 GPU with the same global random seed for all the GAN experiments. Moreover, in order to further tune and test the model performance, we tune AlloyGAN vs. standard cGAN by a combination of the epoch length (ep) from a set of $\{100, 1000\}$ and latent layer dimension (n_z) from a set of $\{30, 100\}$.

AlloyGAN vs. cGAN. Figure 4 shows that AlloyGAN generates informative images that effectively represent the microstructure of the training alloys, unlike standard cGAN. We evaluate both models via machine learning metrics in Table III or material science metrics in Figures 5, 6, and 7.

FID from Pretrained Inception Network ↓	ep = 100, nz = 30	ep = 100, nz = 100	ep = 1000, nz = 30	ep = 1000, nz = 100
AlloyGAN Overall test performance	804.08	746.15	1394.04	1550.11
cGAN Overall test performance	1948.17	1398.75	1455.63	1298.18
AlloyGAN Alloy-wise test performance	1022.15 ± 251.85	1001.22 ± 184.40	1650.96 ± 403.76	1769.98 ± 169.39
cGAN Alloy-wise test performance	2150.85 ± 226.71	1634.39 ± 231.64	1578.84 ± 288.56	1602.19 ± 151.22
FID from Finetuned Inception Network ↓	ep = 100, nz = 30	ep = 100, nz = 100	ep = 1000, nz = 30	ep = 1000, nz = 100
AlloyGAN Overall test performance	14.88	14.02	26.35	39.84
cGAN Overall test performance	68.01	38.00	38.18	46.82
AlloyGAN Alloy-wise test performance	29.83 ± 20.24	33.55 ± 22.32	43.35 ± 16.58	54.22 ± 22.80
cGAN Alloy-wise test performance	68.69 ± 21.61	43.51 ± 20.04	50.30 ± 10.90	60.49 ± 26.95

Table III: We report one FID score over all test images as the Overall calculation method and mean $\pm$ standard deviation across different test alloys as the Alloy-wise calculation method. The FID score by the finetuned inception network shows AlloyGAN outperforms standard cGAN in every hyperparameter setting. The best AlloyGAN Alloy-wise test performance is under a small epoch with less latent dimension (ep = 100, nz = 30). The best AlloyGAN Alloy-wise test performance is under cGAN default hyperparameter setting (ep = 100, nz = 100). The best cGAN Overall or Alloy-wise test performance is under cGAN default hyperparameter setting (ep = 100, nz = 100).

In Table III, the smaller FID scores, the better the generation images perform as real images. Since the inception network is pre-trained on ImageNet dataset which is distinctive from our domain microstructure dataset, we finetuned this network by optimizing the alloy classification according to the train and validation microstructure images. Using the finetuned inception network as the feature extractor to calculate FID scores, AlloyGAN outperforms standard cGAN on the test set with alloys previously unseen during the training process. The result remains consistent when calculating the score for each alloy and then aggregating the scores through methods like taking the mean or standard deviation. No matter the pretrained or fine-tuned networks as feature extractors, as a general metric, FID cannot evaluate whether the generated images are scientifically valid for alloy discovery. We thus also employ domain-science metrics to evaluate alloy foundational chemical properties of the generated images from the conditions of unseen alloys on our test set. **1. Micrograph Evolution with Si Content:** we evaluate whether generated images of AlloyGAN are able to reflect the proposed effect of Si content. Shown in Figure 5, generated images using AlloyGAN perform a clear trend of the area of fractions when Si content increases. This illustrates AlloyGAN has the capability to create valid microstructure images reflecting the Si effects, especially with alloys that previously had not been manufactured. **2. Effect of Cooling Rate:** we then evaluate whether AlloyGAN is able to handle the proposed effect of the Cooling Rate. Shown in Figure 6, its generated images exhibit a strong correlation between Cooling Rate and SDRS. This illustrates AlloyGAN has the capability to create valid microstructure images reflecting the cooling rate effects. This is of particular interest for generating alloys that previously had not been manufactured before. **3. Effect of Sr/P Modification:** Figure 7 shows the model successfully predicts the effects of Sr/P modification on A356 and A393 alloys. This illustrates that AlloyGAN has the capability to create valid microstructure images reflecting the Sr/P modification, especially with never-manufactured alloys.

VI. ALLOYGAN WEBSITE

AlloyGAN is a product from a data science project team formulated by machine learning researchers, software engi-




Figure 5: Area fraction of Si phase vs. Si amount in created images: Si wt. content (a) 4%, (b) 9%, (c) 15%, and (d) 22%.




Figure 6: Effect of Cooling Rate: Red bars illustrate how we measure SDRS in each generated image in terms of cooling rates and manufactured alloys. R^2 indicates the fitness of linearity based on samples measured in each example alloy.




Figure 7: Example generations given different Sr/P conditions: (a) A356 alloys with Sr/P Non-Modified and (b) A356 alloys with Sr/P Modified; (c) A393 alloys with Sr/P Non-Modified and (b) A393 alloys with Sr/P Modified. All alloys are under a Cooling Rate of 10 K/S.

neers, and domain material scientists. We open a publicly accessible website to serve AlloyGAN interactive demo at <https://deepalum.com/>. The website is built by our collaborative startup. Users can create scientifically valid images given their text prompt to determining expected Alloy compositions within 1 second. Enhancements of AlloyGAN to support different types of materials continue to drive up its value to the material science community and customers.

VII. CONCLUSION

AlloyGAN generates microstructures of aluminum alloys under promptable conditions, taking into account chemical elements, manufacturing environments, as well as fundamental chemical reaction knowledge. Distinct from standard cGAN outputs, our model produces informative images of alloy microstructures reflecting crucial chemical properties. AlloyGAN achieves scientifically valid results that rival traditional computational material science methods in accuracy, however significantly reducing time and dependency on complex domain knowledge. With the launch of AlloyGAN, we open a path to efficient deep learning applications in material verification and scientific discovery in the field.

VIII. ACKNOWLEDGEMENT

We thank DeepAlloy (<https://deepalloy.com/>) and NSF NRT-CEDAR 2021871 for sponsoring this research.

REFERENCES

- [1] Allied Market Research, "Metal & metal manufactured products market 2021-2030," <https://www.alliedmarketresearch.com>, 2022, accessed: April, 2022.
- [2] A. S. H. Makhlof and M. Aliofkhazraei, *Handbook of materials failure analysis with case studies from the aerospace and automotive industries*. Butterworth-Heinemann, 2015.
- [3] J. Eiken, "Dendritic growth texture evolution in mg-based alloys investigated by phase-field simulation," *International Journal of Cast Metals Research*, vol. 22, no. 1-4, pp. 86-89, 2009.
- [4] I. Steinbach and M. Apel, "The influence of lattice strain on pearlite formation in fe-c," *Acta Materialia*, vol. 55, pp. 4817-4822, 2007.
- [5] N. Zhou, C. Shen, P. Sarosi, M. Mills, T. Pollock, and Y. Wang, " γ rafting in single crystal blade alloys: a simulation study," *Materials Science and Technology*, vol. 25, no. 2, pp. 205-212, 2009.
- [6] T. Takaki, T. Shimokawabe, M. Ohno, A. Yamanaka, and T. Aoki, "Unexpected selection of growing dendrites by very-large-scale phase-field simulation," *Journal of Crystal Growth*, vol. 382, pp. 21-25, 2013.
- [7] M. Anderson, G. Grest, and D. Srolovitz, "Grain growth in three dimensions: a lattice model," *Scr. Metall.*, vol. 19, no. 2, 1985.
- [8] A. Rollett, D. J. Srolovitz, R. Doherty, and M. Anderson, "Computer simulation of recrystallization in non-uniformly deformed metals," *Acta Metallurgica*, vol. 37, no. 2, pp. 627-639, 1989.
- [9] J. Zhang, X. Li, D. Xu, and R. Yang, "Recent progress in the simulation of microstructure evolution in titanium alloys," *Progress in Natural Science: Materials International*, vol. 29, no. 3, pp. 295-304, 2019.
- [10] E. A. Holm, R. Cohn, N. Gao, A. R. Kitahara, T. P. Matson, B. Lei, and S. R. Yarasi, "Overview: Computer vision and machine learning for microstructural characterization and analysis," *Metallurgical and Materials Transactions A*, vol. 51, pp. 5985-5999, 2020.
- [11] S. M. Azimi, D. Britz, M. Engstler, M. Fritz, and F. Mücklich, "Advanced steel microstructural classification by deep learning methods," *Scientific reports*, vol. 8, no. 1, p. 2128, 2018.
- [12] I. Y. Moon, H. W. Lee, S.-J. Kim, Y.-S. Oh, J. Jung, and S.-H. Kang, "Analysis of the region of interest according to cnn structure in hierarchical pattern surface inspection using cam," *Materials*, vol. 14, no. 9, p. 2095, 2021.
- [13] Y. Kim, H. K. Park, J. Jung, P. Asghari-Rad, S. Lee, J. Y. Kim, H. G. Jung, and H. S. Kim, "Exploration of optimal microstructure and mechanical properties in continuous microstructure space using a variational autoencoder," *Materials & Design*, vol. 202, p. 109544, 2021.
- [14] J. Tang, X. Geng, D. Li, Y. Shi, J. Tong, H. Xiao, and F. Peng, "Machine learning-based microstructure prediction during laser sintering of alumina," *Scientific Reports*, vol. 11, no. 1, pp. 1-10, 2021.
- [15] W. Ma, E. J. Kautz, A. Baskaran, A. Chowdhury, V. Joshi, B. Yener, and D. J. Lewis, "Image-driven discriminative and generative machine learning algorithms for establishing microstructure-processing relationships," *Journal of Applied Physics*, vol. 128, no. 13, p. 134901, 2020.
- [16] K.-H. Lee and G. J. Yun, "Microstructure reconstruction using diffusion-based generative models," *Mechanics of Advanced Materials and Structures*, pp. 1-19, 2023.
- [17] X. Li, Y. Zhang, H. Zhao, C. Burkhart, L. C. Brinson, and W. Chen, "A transfer learning approach for microstructure reconstruction and structure-property predictions," *Scientific reports*, 2018.
- [18] Z. Yang, X. Li, L. Catherine Brinson, A. N. Choudhary, W. Chen, and A. Agrawal, "Microstructural materials design via deep adversarial learning methodology," *Journal of Mechanical Design*, 2018.
- [19] Z. Yang, S. Li, S. Li, J. Yang, and D. Liu, "A two-step data augmentation method based on generative adversarial network for hardness prediction of high entropy alloy," *Computational Materials Science*, vol. 220, p. 112064, 2023.
- [20] I. Ferreira, L. Ochoa, and A. Koeshidayatullah, "On the generation of realistic synthetic petrographic datasets using a style-based gan," *Scientific Reports*, vol. 12, no. 1, p. 12845, 2022.
- [21] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *Advances in neural information processing systems*, vol. 30, 2017.
- [22] R. Kobayashi, "Modeling and numerical simulations of dendritic crystal growth," *Physica D: Nonlinear Phenomena*, vol. 63, no. 3-4, pp. 410-423, 1993.
- [23] Y. Li and L. Chen, "Temperature-strain phase diagram for ba ti o 3 thin films," *Applied physics letters*, vol. 88, no. 7, p. 072905, 2006.
- [24] A. Iyer, B. Dey, A. Dasgupta, W. Chen, and A. Chakraborty, "A conditional generative model for predicting material microstructures from processing methods," *arXiv preprint arXiv:1910.02133*, 2019.
- [25] J. Jung, J. Na, H. K. Park, J. M. Park, G. Kim, S. Lee, and H. S. Kim, "Super-resolving material microstructure image via deep learning for microstructure characterization and mechanical behavior analysis," *npj Computational Materials*, vol. 7, no. 1, p. 96, 2021.
- [26] J.-W. Lee, N. H. Goo, W. B. Park, M. Pyo, and K.-S. Sohn, "Virtual microstructure design for steels using generative adversarial networks," *Engineering Reports*, vol. 3, no. 1, p. e12274, 2021.
- [27] S. Oh, H. K. Kim, T.-E. Jeong, D.-H. Kam, and H. Ki, "Deep-learning-based predictive architectures for self-piercing riveting process," *IEEE Access*, vol. 8, pp. 116254-116267, 2020.
- [28] L. Ning, Z. Cai, Y. Liu, and W. Wang, "Conditional generative adversarial network driven approach for direct prediction of thermal stress based on two-phase material sem images," *Ceramics International*, vol. 47, no. 24, pp. 34115-34126, 2021.
- [29] S. Sundar and V. Sundararaghavan, "Database development and exploration of process-microstructure relationships using variational autoencoders," *Materials Today Communications*, vol. 25, p. 101201, 2020.
- [30] J. Jung, J. I. Yoon, H. K. Park, H. Jo, and H. S. Kim, "Microstructure design using machine learning generated low dimensional and continuous design space," *Materialia*, vol. 11, p. 100690, 2020.
- [31] J. Na, G. Kim, S.-H. Kang, S.-J. Kim, and S. Lee, "Deep learning-based discriminative refocusing of scanning electron microscopy images for materials science," *Acta Materialia*, vol. 214, p. 116987, 2021.
- [32] M. Conti, R. Di Pietro, L. V. Mancini, and A. Mei, "(old) distributed data source verification in wireless sensor networks," *Inf. Fusion*, vol. 10, no. 4, pp. 342-353, 2009.
- [33] M. Warmuzek, *Aluminum-silicon casting alloys: an atlas of microfractographs*. ASM international, 2004.