

Article

Reflectance-Consistent CycleGAN for Low-Sample Data Augmentation in Graphite Ore Grade Recognition

Caolu Liu ¹, Le Chen ¹, Xueyu Huang ^{1,*} and Binghui Wei ²

¹ School of Information Engineering, Jiangxi University of Science and Technology, Ganzhou 341000, China; 6720240900@mail.jxust.edu.cn (C.L.)

² School of Information Engineering, Gannan University of Science and Technology, Ganzhou 341000, China

* Correspondence: 6120190561@mail.jxust.edu.cn

Abstract

Accurate grade detection in graphite ore, which is a strategic and critical mineral resource, plays an important role in improving beneficiation efficiency and overall resource utilization. However, the scarcity of high-grade samples limits the performance of deep learning models in grade identification tasks. This limitation makes it difficult for models to learn stable and representative features. This paper proposes an enhanced CycleGAN-based image augmentation framework designed for graphite ore imagery. The method works within an unpaired image translation architecture. It introduces a distributed reflectance consistency loss. This loss encodes the graphite ore's typical low reflectance and high optical contrast as explicit statistical constraints. The design enforces consistency in both the intensity distribution and the textural structure of the generated images. The model further integrates a convolutional block attention module into the generator. This module helps refine feature representation under a physics-inspired heuristic. The study constructs augmented training sets using the proposed method. It then evaluates these datasets with a downstream grade classification model. Experimental results show clear improvements. The method reduces Fréchet Inception Distance by 21.9% and Kernel Inception Distance by 39.4%. It also improves peak signal-to-noise ratio by 3.3% and structural similarity index measure by 2.6% compared with the baseline CycleGAN. The classification accuracy in the grade identification task increases by about 2.3 percentage points. These results show that the proposed method improves both the perceptual quality and the statistical consistency of synthetic graphite ore images. It also helps reduce the performance drop caused by limited training data in few-shot learning conditions.

Keywords: CycleGAN; graphite ore grade; data augmentation; distributed reflectance consistency loss; few-shot learning



Received: 8 May 2026

Revised: 25 June 2026

Accepted: 30 June 2026

Published: 2 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Graphite ore is widely used in energy storage materials, conductive media, and high-end manufacturing sectors. It serves as a strategic and critical mineral resource in these fields. The grade of the graphite ore directly affects beneficiation efficiency and the overall utilization rate of the resource. Rapid and high-precision grade identification therefore plays an important role in intelligent mineral processing and refined resource development [1]. Conventional grade detection methods rely on manual sampling and laboratory chemical analysis. These methods often use instruments such as XRD, SEM-EDS, and optical microscopy. These techniques can provide high analytical accuracy. However,

they usually require long processing times and involve high operational costs. They also cannot support continuous online monitoring in real production environments [2]. The internal structure of ore bodies is often complex, and sampling practices show strong spatial variation [3]. In practical exploration, high-grade natural graphite deposits are limited and difficult to access [4]. Low-grade samples make up most of the available data. This imbalance often leads models to fit too closely to low-grade features while failing to learn high-grade characteristics during training. As a result, recognition accuracy declines. At the same time, resource conditions continue to worsen and mining environments become more challenging. Efficient and intelligent methods for graphite grade identification have therefore become an important research focus in mining engineering.

In recent years, deep learning has developed rapidly and has driven the use of data-based methods in ore grade identification. Convolutional neural networks and Transformer-based models can extract high-level semantic features from image or spectral data. These models have shown strong performance in ore classification and grade prediction tasks. Cao et al. developed a framework that combines PGNA technology with a Vision Transformer to estimate copper grade [5]. Ren et al. showed that CNN-based models can effectively analyze sensor signals in magnetite sorting tasks [6]. Other studies have explored multimodal image analysis to improve ore feature representation. Liu et al. proposed a weighted evaluation method based on sample usage frequency and distance measures. This method helps reduce estimation bias caused by uneven spatial sampling and improves ore identification performance [7]. The application of deep learning methods still depends on large amounts of labeled training data. In real mining environments, high-grade ore samples are rare [8]. This situation leads to long-tail data distributions and strong class imbalance. Models trained under such conditions are more likely to overfit and produce biased results. This issue remains a key challenge and limits the practical use of deep learning in ore grade assessment.

To address small-sample limitations, similar constraints have been reported across various vision-related tasks, such as medical imaging and scientific machine learning, where insufficient annotated data directly limits model reliability and training efficiency [9,10]. To address this issue, generative methods have been gradually introduced as an effective alternative. Generative adversarial networks, first proposed by Ian Goodfellow, learn the underlying distribution of real data through the interaction between a generator and a discriminator [11]. Building on this idea, prior works have shown that GAN-generated samples can improve classification performance and enhance model robustness in medical imaging, remote sensing, and industrial vision tasks where data acquisition is expensive and limited [12–14].

Among these models, CycleGAN has emerged as a promising solution for graphite ore image translation, mainly because it supports unsupervised mapping without paired samples. Jun-Yan Zhu and colleagues introduced cycle-consistency constraints into the framework [15], which helps maintain a stable correspondence between source and target domains. This mechanism preserves structural information during translation and reflects the idea of symmetry in feature mapping. As a result, the model shows strong adaptability in tasks that require both structural consistency and visual fidelity. This capability has already been explored in several application areas. In medical imaging, CycleGAN has been used to support classification and segmentation tasks with improved reliability [16]. For example, Chehade et al. generated artifact-free chest X-ray images, which led to better performance when combined with CBAM-enhanced DenseNet-121 models [17]. In engineering scenarios, the model has been applied to generate nighttime bridge images and to simulate underwater crack evolution [18]. Furthermore, combining CycleGAN with an optimized YOLO11 achieved a detection rate of 91.8% in crack identification tasks [19].

These studies provide a useful reference for further exploration. When transfer learning is introduced alongside CycleGAN in graphite ore data augmentation, it offers a practical way to reduce data scarcity. At the same time, this strategy helps avoid the feature distortion that often appears in conventional augmentation methods, making the overall learning process more stable and reliable. However, it should be noted that existing reviews and applications of CycleGAN-based translation have largely focused on medical and civil engineering scenarios, while critical discussions on its use for mineral and ore grade image translation remain limited. To the best of our knowledge, existing computer vision research for mineral and ore images has been largely limited to conventional downstream tasks, including ore classification, rock type recognition, conveyor belt object detection, and ore image segmentation [20–23], rather than exploring how cross-domain translation can mitigate data scarcity while preserving the material and textural characteristics of different ore grades. This gap motivates our present work, which adapts and improves CycleGAN to address these domain-specific requirements for graphite ore image translation.

Direct use of CycleGAN in graphite ore image augmentation does not always bring clear gains in downstream classification. Graphite ore images show distinct material and textural properties. They usually have low overall reflectance, while local regions present high contrast. These features play a key role in accurate grade discrimination. Conventional CycleGAN models mainly depend on pixel-level consistency constraints. This design makes it hard to keep these empirically meaningful patterns in the generated images [24]. At the same time, the original model does not include an attention mechanism. The generator therefore relies only on stacked convolutional receptive fields to process features. This limitation often leads to overly broad global changes. It can also damage or remove important structural details during generation. Existing evaluation methods add another layer of mismatch. Many studies still rely on general metrics such as Fréchet Inception Distance and peak signal-to-noise ratio. These metrics do not always align with the needs of grade recognition tasks, so they cannot fully reflect the practical value of generated samples [25]. As a result, several issues remain to be resolved. It remains necessary to integrate domain-specific priors and attention mechanisms into the model, while designing evaluation criteria that relate directly to the target task.

To deal with these problems, this paper presents a small-sample data augmentation framework for graphite ore images. The framework builds on an improved CycleGAN architecture. The main contributions are outlined below.

1. A target-specific high- and low-grade graphite ore dataset is constructed from industrial production environments. High-resolution imaging equipment captures graphite ore samples under varied illumination conditions; manual annotation is subsequently corroborated by quantitative carbon-sulfur analyzer measurements to support labeling fidelity.
2. A physics-inspired distributed reflectance consistency loss is incorporated into the CycleGAN objective. This term constrains both the mean and standard deviation of synthetic outputs to align with the statistical signature of authentic high-grade graphite ore. An exponential moving average scheme concurrently helps stabilize global distributional statistics, mitigating training fluctuations and reducing the model's tendency to overfit spurious local textural variations.
3. A convolutional block attention module is embedded within the generator's intermediate layers. Isolated deployment of CBAM yields marginal gains; however, its integration with distributed reflectance consistency loss produces a favorable synergistic effect that improves generative fidelity.
4. Assessment extends beyond conventional visual quality metrics. A downstream grade classification pipeline built upon the ConvNeXt architecture serves as a task-specific

indicator, providing a practical evaluation of the utility of augmented samples for graphite ore grade discrimination within the tested environment.

The remainder of the manuscript is structured as follows. Section 2 delineates the design rationale and implementation specifics of the improved CycleGAN variant. Section 3 describes dataset curation protocols, experimental configurations, and the performance metrics employed. Section 4 presents a comprehensive empirical evaluation, situating the proposed method against baseline models and multiple ablation configurations. Section 5 offers concluding remarks and directions for future inquiry.

2. Methods

2.1. Overall Framework

This study builds on the CycleGAN framework to handle image translation between different grades of graphite ore. The model defines two adversarial mappings. One mapping converts images from the source domain to the target domain, and the other mapping performs the reverse transformation. In this setting, low-grade graphite ore images are translated into the high-grade domain, while high-grade images are translated in the opposite direction. The model applies cycle consistency to keep structural features stable during translation. This constraint makes the training process more stable and allows the model to work without paired data. Similar designs have shown reliable results in medical imaging, remote sensing, and industrial inspection tasks [16,26]. The overall structure follows the common unpaired image-to-image translation scheme. The network includes two generators, $G_{A \rightarrow B}$ and $G_{B \rightarrow A}$, as well as two discriminators, D_A and D_B . The generator $G_{A \rightarrow B}$ maps low-grade graphite ore images into the high-grade domain, and $G_{B \rightarrow A}$ performs the reverse mapping. The model also uses a patch-based discriminator. This component focuses on local regions and strengthens the learning of texture realism and fine details [27]. This setup helps reduce the gap between generated images and real samples while keeping structural consistency.

The framework further extends CycleGAN at both the structure level and the loss level. The model inserts a Convolutional Block Attention Module into the generator to improve feature representation. Section 2.2 provides the detailed design. At the same time, the model introduces a Distribution Reflectance consistency Loss to guide the training process, and Section 2.3 explains this part. The method does not rely only on visual style transfer. It also considers the statistical appearance properties of graphite ore images. In this task, visual similarity alone cannot ensure reliable results, and statistical consistency also matters. The model therefore adds a reflectance-based constraint into the loss function. This design leads to a joint optimization of the network structure and the training objective. With these changes, the model generates images that look realistic and also follow the statistical distribution of real high-grade graphite ore. Figure 1 shows the overall framework. The adversarial loss helps match data distributions, cycle consistency keeps structural stability, and the Distribution Reflectance consistency Loss maintains reflectance consistency from a statistical viewpoint.




Figure 1. Overall architecture of the improved CycleGAN algorithm.

2.2. Generator Network Structure

The design of the generator directly affects the expressive capacity and the quality of generated images. In this study, the model uses a deep residual learning framework based on an encoder–decoder structure with a ResNet backbone. This structure places nine residual blocks in sequence. These blocks strengthen feature extraction and help keep the transformed graphite textures clear and distinguishable. The ResNet architecture was first proposed by Kaiming He and has been widely adopted in image generation tasks [28]. Later work by Justin Johnson introduced a perceptual loss framework. That method combines a feedforward convolutional network with residual blocks. It follows the same design idea and supports real-time style transfer as well as high-quality single-image super-resolution [29]. This study focuses on translating low-grade graphite images into high-grade ones, so the design of the generator $G_{A \rightarrow B}$ needs careful attention. The generator plays a central role in capturing both texture patterns and structural features. Figure 2 shows its main configuration.




Figure 2. Generator structure of the improved CycleGAN algorithm.

The generator comprises an encoding stage, a feature transformation stage containing nine residual blocks, and a decoding stage. Instance Normalization and ReLU activation are applied after every convolutional layer, with the final output layer employing a Tanh function. Regarding normalization, we follow the established convention of Instance Normalization (IN) for generative tasks [30]. In contrast to batch normalization, which aggregates statistics across the batch dimension, IN computes normalization parameters independently for each image. In the context of geological image translation, grade-specific attributes—such as metallic luster and cleavage texture—often reside in subtle local variations. IN effectively prevents cross-style contamination between different graphite ore samples within a batch, thereby safeguarding fine graphite particle details from unwanted smoothing or artifact introduction. The ResNet backbone, grounded in residual learning, further alleviates the vanishing gradient issue commonly encountered

in deep architectures [31]. In image synthesis, residual connections excel at preserving low-frequency structural contours while selectively refining high-frequency textural details. This property is particularly valuable for graphite ore imagery. While macroscopic geometric differences between graphite ore grades remain minimal, the discrimination hinges on nuanced variations in reflectance and surface texture. By maintaining structural integrity while reconstructing textural cues, the residual framework substantially enhances the perceptual fidelity of the generated outputs.

The original CycleGAN architecture lacks an explicit self-attention mechanism, limiting its capacity to fully leverage global structural context and precisely refine localized details. To mitigate these limitations, we integrate the Convolutional Block Attention Module into the generator. This mechanism enables the network to jointly attend to salient channel-wise features and spatial positions, facilitating a dynamic recalibration of the feature maps [32]. Considering the specific demands of our transformation task—where preserving structural fidelity is paramount—CBAMs are inserted exclusively into the intermediate layers of the residual stack, specifically within the 4th, 5th, and 6th blocks of the nine-block sequence. This restricted placement enhances feature expressiveness while insulating the critical structural pathways between the input and output layers from potential interference. Functioning as a lightweight, plug-and-play component, CBAM first derives attention weights along the channel axis, subsequently generates a spatial attention map, and finally performs an element-wise multiplication of these attention maps with the input feature volume. This process enables adaptive refinement and selective emphasis of salient geological features. The overall configuration of the CBAM module is depicted in Figure 3.




Figure 3. Overall structure of CBAM reproduced from [32].

Embedding the convolutional block attention module into the core of the CycleGAN generator enables adaptive, dynamic refinement of feature maps and sharpens feature selectivity throughout the transformation from low-grade to high-grade graphite ore imagery. Channel-wise attention actively suppresses background activations and noise components irrelevant to high-grade characteristics, while amplifying channels that encode the defining traits of superior graphite ore quality—thereby directing the generator’s representational focus with greater precision. In parallel, spatial attention steers the model’s modeling capacity toward locally discriminative regions of the graphite surface, confining fine-grained adjustments to those critical areas rather than imposing uniform modifications across the entire image. This targeted treatment helps avert the distortion of geological signatures that often accompanies global homogenization. The interplay of channel and spatial attention mechanisms bolsters both detail fidelity and generative quality. Consequently, the generator becomes adept at capturing the nuanced visual cues that distinguish graphite ore grades. The synthesized high-grade outputs more closely approximate real specimens in essential geological particulars, mitigating common artifacts such as image blur and

textural inconsistency while elevating both the perceptual authenticity and the geological plausibility of the generated results.

2.3. Distributed Reflectance Consistency Loss

The motivation for introducing a distributed reflectance consistency loss stems directly from the appearance and material characteristics of graphite. A well-documented geological correspondence links graphite grade to its reflectance behavior. High-grade material, rich in carbon, strongly attenuates visible light; the resulting reflectance remains low, lending the graphite ore a dark or near-black visual appearance. Moreover, the well-developed crystallinity of high-grade graphite imparts a distinct metallic luster, which translates in imagery to pronounced contrast and richly articulated textural detail. Low-grade graphite, by comparison, contains substantially less carbon, and its oxidized surfaces frequently host impurity phases such as quartz and clay minerals. These inclusions interact only weakly with visible light, yielding dark gray hues that occasionally carry brownish undertones. Incomplete crystallization and poorly expressed cleavage surfaces further manifest as diminished contrast, textural blurring, and a notable absence of specular highlights.

Conventional image generation pipelines tend to prioritize visual resemblance while overlooking such domain-specific constraints. Cyclic consistency losses, in their standard formulation, offer no mechanism to enforce these material-specific attributes. Consequently, a generated output may approximate the color distribution of the target domain yet diverge meaningfully in empirical character. For instance, synthesized images of high-grade graphite may appear excessively pale or deficient in contrast, failing to capture the defining optical signature of carbon-rich graphite ore. Such discrepancies inevitably degrade classification fidelity in downstream grade identification tasks. To address this, we propose a distributed reflectance consistency loss that anchors the generated output to the statistical profile of authentic high-grade samples. This is accomplished by jointly constraining the mean and the standard deviation of the synthesized imagery. The mean constraint governs the correctness of the global tonal balance, while the standard deviation constraint safeguards textural richness and sufficient local contrast. Operating in tandem, these two statistical bounds impose a principled regularization on the visual plausibility of the generated results, ensuring that they satisfy not only perceptual realism but also the empirical appearance expected of graphite mineralogy. Meanwhile, since all experiments are conducted under consistent and fixed lighting conditions, the global intensity mean and local standard deviation of image pixels can serve as effective and stable statistical proxies to characterize the macroscopic optical properties of graphite ore. These statistics are not intended to replace professional laboratory measurements of reflectance, but rather to capture the domain-specific visual-physics-inspired heuristic that are most critical for distinguishing between low-grade and high-grade samples. This design ensures the generated images conform to the inherent optical laws of graphite mineralogy, rather than merely learning superficial color and texture patterns from data.

Therefore, guided by physics-inspired statistical priors, we denote the synthesized high-fidelity image as I_{fake} . To characterize its photometric and textural properties, we compute the global mean intensity μ_I and standard deviation σ_I , which respectively quantify the overall brightness and the contrast texture richness of the image. Let H and W represent the image height and width, and let $I_{i,j}$ denote the intensity value at pixel location (i, j) . The corresponding mathematical definitions are given as follows:

$$\mu(I) = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W I_{i,j} \quad (1)$$

$$\sigma(I) = \sqrt{\frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W (I_{i,j} - \mu(I))^2} \quad (2)$$

The conventional CycleGAN framework is limited by GPU memory, typically requiring a batch size of one during training. Enforcing alignment between the generated synthetic graphite sample and a single randomly selected real graphite sample per iteration can induce training instability. This instability originates from the inherent microscopic heterogeneity of natural high-grade graphite ore. Despite belonging to the same high-grade category, individual specimens exhibit significant variation: some surfaces feature bright quartz veins, while others are uniformly dark and pure. If the model overfits to these stochastic inter-sample fluctuations, the training trajectory becomes unstable, preventing convergence to a robust optimum.

To prevent the model from being adversely affected by statistical fluctuations from a single random sample—such as noise, outliers, or incidental variations within a batch—during early training, which can cause oscillation and sharp jumps in the objective, this work abandons the single-sample alignment scheme. Instead, we employ an exponential moving average (EMA) strategy to stably estimate the global mean and variance throughout training [33]. We maintain two global variables, μ_{target} and σ_{target} , denoting the target mean and target standard deviation of the real high-grade graphite ore dataset, respectively. Together, they represent the expected distribution of the real high-grade dataset over the entire training cycle.

$$\overline{\mu_{new}} = m \cdot \overline{\mu_{old}} + (1 - m) \cdot \mu(I_{real}) \quad (3)$$

$$\overline{\sigma_{new}} = m \cdot \overline{\sigma_{old}} + (1 - m) \cdot \sigma(I_{real}) \quad (4)$$

Here, $\overline{\mu_{new}}$ denotes the updated target mean of pixel intensity for the high-grade graphite ore imagery, serving as the stabilized estimation reference during subsequent training. $\overline{\mu_{old}}$ represents the corresponding target mean prior to the update—namely, the historically accumulated statistic. $\mu(I_{real})$ indicates the mean intensity computed from I_{real} , the authentic high-grade sample of the current batch. Analogously, $\overline{\sigma(I_{real})}$ and $\overline{\sigma_{old}}$ designate the updated and prior target standard deviations, while $\sigma(I_{real})$ captures the standard deviation of the current real high-grade graphite ore image. The momentum coefficient m is assigned a value of 0.99 throughout the experiments reported here. Examining this update rule reveals two complementary justifications for the chosen momentum setting. Mechanistically, a coefficient of 0.99 ensures that the evolving global mean and variance incorporate only approximately one percent of the information contributed by any single incoming sample; the overwhelming majority of the weight remains anchored to historical global estimates. This conservative update strategy effectively dampens the disruptive influence of sample-level noise and statistical outliers, thereby preserving training stability. Empirically, this particular momentum value has become a widely adopted convention across image generation, style transfer, and domain adaptation literature [34–37]. While the global distribution is progressively smoothed, the global pixel mean and standard deviation of the authentic high-grade graphite dataset undergo a gradual, self-adjusting drift that tracks the underlying data manifold. The model thus optimizes under the guidance of a remarkably stable heuristic constraint, circumventing common pitfalls such as textural blurring and chromatic deviation.

The distributed reflectance consistency loss is formally defined as the L1 distance between the statistics of the generated image and the corresponding global target values:

$$L_{DRL} = \lambda_{\mu} \left| \mu(I_{fake}) - \overline{\mu_{target}} \right| + \lambda_{\sigma} \left| \sigma(I_{fake}) - \overline{\sigma_{target}} \right| \quad (5)$$

where λ_μ and λ_σ are weighting hyperparameters that balance the relative importance of the mean alignment constraint and the standard deviation alignment constraint, respectively. Conventionally, λ_μ and λ_σ are inherently fixed to 1.0 to assign equal geometric importance to both first- and second-order statistical moments, thereby avoiding redundant over-parameterization of the loss function. The two components embedded in this formulation are jointly essential, each imposing a distinct semantic constraint. Mean alignment establishes the global tonality by actively suppressing the overall intensity scale of the synthesized output, thereby eliminating the grayish-white cast characteristic of low-grade graphite ore. Relying on this term alone, however, invites a well-known pitfall: neural networks, adept at exploiting expedient shortcuts, will simply converge upon a uniform gray image whose average brightness exactly matches the target value. Within the generative modeling literature, this degenerate outcome is referred to as mean collapse. To avoid trivial solutions, this study introduces a second term based on standard deviation alignment. Specifically, standard deviation alignment refers to mathematically constraining the pixel intensity variance of the generated image. The constraint aims to match the statistical spread of real high-grade samples. Visually, this alignment actively prevents the network from taking a shortcut. It stops the network from outputting a flat and featureless image known as mean collapse. Forcing a target variance explicitly guides the model to recover critical high-frequency visual details. These details include metallic luster, crystalline graininess, and specular highlights. Such characteristics are inherent to the rough fracture surfaces of authentic high-grade graphite. In visual feature space, standard deviation is reflected as contrast. It corresponds to the crystalline graininess that can be observed by the naked eye. The model is guided to match the standard deviation of real high-grade graphite ore. This constraint encourages the network to recover high-frequency details on top of the tonal structure. These details include scaly specular reflections and rough fracture surfaces. This dual constraint on the first- and second-order moments steers the model toward learning empirically meaningful representations of graphite ore texture.

Following the integration of the distributed reflectance consistency loss, the composite objective function comprises four components: the cycle consistency loss L_{cycle} , the distributed reflectance consistency loss λ_{DRL} , the identity loss $L_{identity}$, and the adversarial loss L_{GAN} . The overall optimization target is expressed as

$$L_{total} = L_{GAN} + \lambda_{cyc}L_{cycle} + \lambda_{id}L_{identity} + \lambda_{DRL}L_{DRL} \quad (6)$$

In this experiment, we set the cycle consistency loss weight λ_{cyc} to 10.0. This value helps the network change the graphite ore grade while keeping the overall structure of the graphite intact. The identity loss weight λ_{id} is set to 5.0. All images used in this study have solid color backgrounds. Under such conditions, the generator can easily shift the color tone during unsupervised mapping. The identity loss works as an anchor here. It holds the color and texture steady so they do not drift away. The distributed reflectance consistency loss is the main contribution of this paper. Its weight λ_{DRL} is set to 0.1. A value this small means the empirical constraints act only as soft regularization. They should never dominate the training process. If we give this weight a much larger value, say close to 1.0, the network will chase a shortcut. It will try to match the mean and variance statistics perfectly while ignoring overall image fidelity. The adversarial loss would then become ineffective. The final output would exhibit simplified textures and unnatural artifacts, lacking any sense of depth. A carefully tuned weight of 0.1 avoids this problem. It nudges the brightness and contrast toward the target statistical characteristics of high-grade graphite ore, yet still preserves natural mineral textures. With this balance between adversarial pressure and heuristic guidance, the model robustly transforms samples from low-grade to high-grade appearances.

It is worth noting that although components such as CBAM, EMA, and moment distance mentioned above have been widely adopted in machine learning, their specific, hybridized synthesis under our framework represents a distinct departure from generic statistically constrained GANs or naive domain-adaptation methods. Generic methods typically fail when encountering the intense microscopic heterogeneity inherent to natural graphite ore. By anchoring the training manifold to a macro-statistical expectation via EMA and jointly regulating the tone-contrast trade-off using mineralogical priors, this composite formulation explicitly solves the dual problem of parameter oscillation and texture blurring in geological data augmentation, which constitutes the primary contribution of this work.

To explicitly clarify the execution sequence and the exact integration points of the proposed components (CBAM and DRL), Algorithm 1 summarizes the training loop for a single iteration. As detailed in Section 2.2, the CBAMs are structurally embedded within the 4th, 5th, and 6th residual blocks of the generator $G_{A \rightarrow B}$. Furthermore, the DRL target statistics are updated via an exponential moving average (EMA) prior to the loss computation. This step ensures that the statistical distribution constraint is applied stably during the backward pass. This formulation serves as a reliable heuristic regularization throughout the training stage. The variables I_{Real_A} , I_{Real_B} , I_{Fake_B} , I_{Rec_A} , I_{Fake_A} and I_{Rec_B} are explained in detail in Section 4.4 later in the manuscript.

Algorithm 1: Training loop for one iteration of the Reflectance-Consistent CycleGAN

Input: Source domain image I_{Real_A} , Target domain image I_{Real_B}
Parameters : EMA momentum $m = 0.99$, Loss weights λ_{cyc} , λ_{id} , λ_{DRL}

- 1: //Step 1: Forward Pass (Generators)
- 2: $G_A \rightarrow B(I_{Real_A}) \rightarrow I_{Fake_B}$ //CBAM dynamically refines features in ResBlocks 4, 5, 6
- 3: $G_B \rightarrow A(I_{Fake_B}) \rightarrow I_{Rec_A}$
- 4: $G_B \rightarrow A(I_{Real_B}) \rightarrow I_{Fake_A}$
- 5: $G_A \rightarrow B(I_{Fake_A}) \rightarrow I_{Rec_B}$
- 6: //Step 2: DRL Target Update (Exponential Moving Average)
- 7: Compute current batch statistics from real target: $I_{Real_B} \rightarrow \mu(I_{real}), \sigma(I_{real})$
- 8: Update global anchors (Equations (3) and (4)):
- 9: $0.99 \cdot \mu_{old} + 0.01 \cdot \mu(I_{real}) \rightarrow \mu_{target}$
- 10: $0.99 \cdot \sigma_{old} + 0.01 \cdot \sigma(I_{real}) \rightarrow \sigma_{target}$
- 11: //Step 3: Loss Computation
- 12: Compute adversarial loss LGAN using Discriminators DA,DB
- 13: Compute L_{cycle} and $L_{identity}$
- 14: Extract statistics from generated image: $I_{Fake_B} \rightarrow \mu(I_{Fake}), \sigma(I_{Fake})$
- 15: Compute Distributed Reflectance consistency Loss L_{DRL} (Equation (5)) using updated
- 16: Compute total objective L_{total} (Equation (6))
- 17: //Step 4: Backward Pass & Weights Update
- 18: Update Generators $G_A \rightarrow B, G_B \rightarrow A$ by minimizing L_{total}
- 19: Update Discriminators DA,DB by minimizing discriminator adversarial loss

3. Data and Experiment Preparation

3.1. Dataset Construction

The experimental data come from China Minmetals Graphite Industry Co., Ltd. in Hegang, China, and are collected under real industrial production conditions using available graphite ore resources. The dataset construction follows a complete workflow that includes raw graphite ore collection, multi-angle image acquisition, crushing the samples

to particle sizes below 2 mm, grinding into powder, drying at 105 °C, removing inorganic carbon through acid treatment, eliminating organic carbon by low-temperature combustion, measuring with a carbon–sulfur analyzer, and finally labeling and organizing all samples. Image acquisition uses a Canon EOS 5D Mark II digital single-lens reflex camera (Canon Inc., Tokyo, Japan), with an original resolution of 5616×3744 pixels, and a total of 3152 graphite ore images are collected. Graphite grade refers to the fixed carbon content and serves as a key indicator of quality and purity, where higher carbon content corresponds to better quality. Mineralogical analysis shows that the samples include both primary and oxidized graphite ore. Primary graphite ore mainly contains graphite, quartz, and silicate minerals, while oxidized graphite ore is often associated with minerals such as hematite and pyrite, which make its surface color more complex. Based on measured carbon content, the samples are divided into three grades: 0–10% as low grade, 10–20% as medium grade, and above 20% as high grade. In general, high-grade graphite ore appears darker and more uniform due to its higher carbon content, while low-grade graphite ore contains more light-colored minerals such as quartz and silicates, leading to a mottled surface. Oxidized graphite ore is further affected by weathering and iron-bearing minerals, which often produce reddish-brown or yellow tones. Figure 4 presents the characterized photographs of representative samples, where the first row corresponds to low-grade oxidized graphite ore, the second row to low-grade primary graphite ore, and the last row to high-grade graphite ore.



Figure 4. Images of different types of graphite ore.

This experiment selected a portion of low-grade graphite ore and a limited number of high-grade graphite ore as representative image samples. The dataset contains 1084 images of low-grade graphite ore and 337 images of high-grade graphite ore. In this study, low-grade refers to graphite ore with a grade less than 10%, while high-grade refers to graphite ore with a grade greater than 20%. The study constructs a dedicated dataset based on the collected samples. The dataset is divided into a training set and a test set at a ratio of 9:1. In the high-grade graphite training set, the study increases the sample size by capturing each graphite ore from multiple angles. This process results in 303 images of high-grade graphite ore in the training set. The low-grade graphite training set contains 976 images. The test set includes 34 high-grade and 108 low-grade independent samples. The study applies a standardized batch preprocessing procedure to all collected images. This step adjusts every image to a fixed size of 256×256 . The process ensures that the input data remain consistent and effective during model training. It also meets the input requirements

of deep learning models and aligns with practical needs in industrial production. The details of the dataset partition are shown in Table 1 below.

Table 1. Details of dataset partitioning.

	Training Set	Test Set	Total
0–10%	976	108	1084
20%+	303	34	337
Total	1279	142	1421

3.2. Experimental Environment

All experiments ran on a single cloud computing node. The node was equipped with an AMD EPYC 7642 processor, an NVIDIA RTX 3090 graphics card, 90 GB of system memory, and a 50 GB high-speed data disk. This hardware configuration fully supported the heavy gradient computations and feature iterations required by CycleGAN training. The system environment was Ubuntu 20.04 LTS. The model code was built with Python 3.8 and the PyTorch 2.0.0 framework. GPU acceleration relied on CUDA 11.8 and the driver version was 580.76.05, a combination that ensured compatibility and stability throughout the whole training process.

The training hyperparameters and execution flow were kept identical across all comparison experiments. A fixed batch size of 8 and 200 training epochs were adopted, which was sufficient for full convergence. We chose the Adam adaptive moment estimation optimizer and set its core parameters to $\beta_1 = 0.5$ and $\beta_2 = 0.999$. This setting effectively balanced momentum accumulation and adaptive learning rate and suited the non-stationary nature of adversarial training. The learning rate schedule combined an initial constant phase with a later linear decay. Starting at 0.0002, the rate was held constant for the first 100 epochs to drive the model rapidly into a near-optimal region. Over the next 100 epochs, the rate decayed linearly to zero, allowing the model to fine-tune its parameters near the optimum and avoiding oscillations. Generator and discriminator were updated synchronously with a 1:1 frequency. The discriminator was trained using a pool of 50 historical generated images. Gradient clipping was also applied to constrain the gradient range, preventing gradient explosion and ensuring stable convergence of the entire training network.

3.3. Evaluation Criteria

To evaluate the generation performance of the improved CycleGAN model in practical use in a systematic and objective way, this study adopts the Fréchet Inception Distance (FID), Kernel Inception Distance (KID), Structural Similarity Index (SSIM), Peak Signal-to-Noise Ratio (PSNR), and a downstream grade classification task as evaluation metrics. SSIM and PSNR serve as auxiliary measures, and the study computes them between the input images and the reconstructed images to assess the cycle consistency of the model. FID and KID are treated as the primary metrics, and their mathematical definitions are given as follows:

$$\text{FID} = \|\mu_r - \mu_g\|_2^2 + \text{Tr}(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2}) \quad (7)$$

$$\text{KID} = \frac{1}{n(n-1)} \sum_{i \neq j} k(x_i, x_j) + \frac{1}{m(m-1)} \sum_{i \neq j} k(y_i, y_j) - \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m k(x_i, y_j) \quad (8)$$

In FID, μ_r and μ_g denote the mean vectors of the features extracted by the InceptionV3 network from the real and generated high-grade graphite ore image sets. Σ_r and Σ_g denote the covariance matrices of the corresponding features. FID measures the distance between two multivariate Gaussian distributions. A smaller value indicates that the generated high-grade graphite ore images are closer to the real sample distribution, which

reflects a better conversion from low-grade to high-grade. In KID, n and m denote the numbers of real and generated high-grade graphite ore image samples. x_i and y_j denote the corresponding feature vectors extracted by the InceptionV3 network. The function $k(\cdot, \cdot)$ is a commonly used cubic polynomial kernel, defined as $k(u, v) = (u \cdot v + 1)^3$. KID corresponds to the squared unbiased maximum mean difference. It does not assume that the features follow a Gaussian distribution. A smaller value indicates higher consistency between the generated and real distributions.

In downstream classification tasks, this study uses five metrics, including accuracy, precision, recall, F1 score, and the area under the receiver operating characteristic curve (AUC), to evaluate the classification model. The specific formulas are given as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (9)$$

$$Precision = \frac{TP}{TP + FP} \quad (10)$$

$$Recall = \frac{TP}{TP + FN} \quad (11)$$

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (12)$$

In these formulas, TP denotes the number of samples that are actually high-grade graphite ore and are predicted as high-grade by the model. FP denotes the number of samples that are actually low-grade graphite ore but are incorrectly predicted as high-grade. FN denotes the number of samples that are actually high-grade graphite ore but are misclassified as low-grade. TN denotes the number of samples that are actually low-grade graphite ore and are correctly identified as low-grade by the model. Accuracy represents the proportion of samples that the model predicts correctly, and it reflects the overall classification performance. Precision reflects the proportion of truly high-grade samples among those predicted as high-grade, which indicates the reliability of the prediction results. Recall measures the proportion of actual high-grade samples that are successfully identified. The F1 score is the harmonic mean of precision and recall, and it provides a balanced evaluation under imbalanced data conditions. AUC measures the ability of the model to distinguish graphite ore images of different grades from a global perspective. It serves as a stable and objective evaluation metric in binary classification tasks.

4. Experiments and Discussion

4.1. Model Training

The improved CycleGAN model was trained on a customized graphite ore dataset and evaluated against the baseline model. To ensure stable training of the improved model, we adopt the EMA strategy for robust global distribution estimation, apply gradient clipping to prevent gradient explosion, use a two-phase learning rate schedule with initial constancy and subsequent linear decay, update the generator and discriminator synchronously at a 1:1 frequency, and employ a 50-sample historical image pool to smooth training dynamics.

Figure 5 shows the training process. Meanwhile, Figure 5 provides clear and direct evidence for the convergence analysis of both models. Through the loss curves, we can quantitatively compare the convergence speed, stability, and oscillation degree between the baseline model and the improved model. The loss curves present the evolution of generator and discriminator losses for both the baseline CycleGAN and the improved model across training epochs. In the early stage, both models adapted quickly to the data distribution. The generator loss decreased rapidly, and the discriminator loss decreased at the same

time. This stage shows good initial convergence. As training continued, the baseline model became gradually unbalanced. The generator loss showed stronger oscillations and even increased in the later stage. The discriminator loss kept decreasing to very low values, which can lead to vanishing gradients. The improved model shows a different pattern. The generator loss stayed at a lower level, and the convergence process remained more stable. The discriminator loss fluctuated within a reasonable range and did not drop excessively. This behavior reflects a healthier adversarial interaction.




Figure 5. Training comparison between the improved CycleGAN and baseline model. (a) Generator loss-epochs curve; (b) Discriminator loss-epochs curve.

In general, in terms of convergence performance, the baseline model gradually suffers from severe oscillation and unstable convergence as the training proceeds. The loss fluctuates violently and cannot maintain a steady downward trend, which indicates insufficient convergence robustness. In contrast, the improved model maintains continuous and smooth convergence throughout the whole training process without obvious fluctuation or rebound. Such results demonstrate that the improved model achieves more stable, robust, and sufficient convergence, which effectively addresses the convergence issues existing in the baseline CycleGAN. Overall, the improved model achieves more stable training dynamics and a more balanced adversarial state by introducing loss constraints and attention mechanisms.

4.2. Ablation Experiment

4.2.1. Module Ablation Study

We used CycleGAN as the baseline and gradually added improved modules to evaluate the contribution of each component. All tests were performed under identical experimental settings and parameter configurations, keeping variable interference strictly under control and allowing a fair assessment of different improvement measures. The evaluated configurations are listed below.

- A: Basic CycleGAN model
- B: Adding only the Convolutional Block Attention Module (CBAM)
- C: Introducing only the Distributed Reflectance Consistency Loss
- D: Adding both the Convolutional Block Attention Module and the Distributed Reflectance Consistency Loss

The ablation results are reported in Table 2. To ensure a strictly fair and deterministic comparison, all evaluation metrics were computed on the exact same evaluation set using the finalized model checkpoints from a single training run, rather than being averaged across multiple random stochastic trials. Figure 6 presents the visual outcomes on graphite ore after each module is added. The first column presents original low-grade graphite ore images, which serve as inputs for the image translation task. The last column shows real

high-grade graphite ore images from the dataset. Columns two to four illustrate the effects of different modules, which gradually adjust the surface features of input images to match the target style.

Table 2. Module ablation evaluation results. “√” indicates that the corresponding module or strategy has been integrated into the model.

Model	A	B	C	D	FID	KID	SSIM	PSNR
CycleGAN	√				112.7	0.0786	0.948	29.62
		√			111.82	0.0738	0.905	28.19
			√		104.49	0.0603	0.969	29.71
				√	87.97	0.0479	0.973	30.62




Figure 6. Visualization results of graphite ore ablation experiment. The red dotted boxes mark the key local regions for visual comparison of generation quality among different models.

The table data indicate that progressively adding different modules has a noticeable effect on model performance. The overall change moves from a distribution mismatch toward a gradual coordination between distribution and structure. The baseline CycleGAN model already provides a basic ability to convert low-grade graphite ore into high-grade domains. Yet the relatively high Fréchet Inception Distance and Kernel Inception Distance still point to a significant deviation in the overall distribution of the generated results. Introducing only the CBAM attention module lowers FID and KID slightly. This brings the model somewhat closer to the target domain distribution. SSIM and PSNR, however, drop. The attention mechanism appears to strengthen discriminative features while it weakens the preservation of the original structure. In other words, the image looks more like a high-grade graphite ore but does not necessarily represent the original graphite ore.

Introducing only the distributed reflectance consistency loss leads to a more stable improvement. FID and KID both decrease, and SSIM and PSNR improve significantly. This indicates that the constraint guides distribution alignment and also restricts structural information. It produces a more reasonable balance between visual consistency and

target statistical characteristics in the generated results. The balance is especially suitable for graphite images where texture and reflectance features dominate. When CBAM is combined with the distributed reflectance consistency loss, the model achieves its best experimental performance. FID drops to 87.97, KID drops to 0.0479, and SSIM and PSNR reach their highest values. One way to interpret this is that the attention mechanism enhances local expressive power, while the distribution constraint limits the overall offset. The two components are, to some degree, complementary. However, it is worth noting that despite these relative improvements, the remaining absolute FID and KID values suggest that minor generative artifacts persist, which can be attributed to the simplified 2D static nature of the images and the single-source dataset constraints.

Meanwhile, we take the final Model D as the baseline and analyze the ablation results from a reverse perspective to verify the necessity of each component by removing them one by one. After removing CBAM, we can see that the FID increases from 87.97 to 104.49 and the KID rises from 0.0479 to 0.0603, with obvious performance degradation, which indicates that CBAM is crucial for fine-grained feature extraction and focusing on key regions. Similarly, after removing DRL, the FID increases from 87.97 to 111.82 and the KID rises from 0.0479 to 0.0738, accompanied by reduced cross-domain consistency, proving that DRL loss acts as a critical constraint for maintaining generation quality. Therefore, the results show that removing either component will lead to a remarkable performance drop, and the synergistic effect of the two components is the core guarantee for the model to achieve the optimal performance.

This ablation experiment demonstrates a clear pattern: a single improvement often works on only one aspect. In the task of mapping low-grade graphite to high-grade graphite, constraining both local features and global distribution simultaneously is more critical. The improved final performance of the model further validates this combined strategy.

4.2.2. Hyperparameter Sensitivity Analysis

To further verify the robustness of our model regarding parameter choices, this section provides a sensitivity analysis for the distributed reflectance consistency loss weight λ_{DRL} and the Exponential Moving Average (EMA) momentum. All tests were done under the same baseline architecture.

First, we fixed the cycle consistency loss weight $\lambda_{cyc} = 10.0$ and identity loss weight $\lambda_{id} = 5.0$, while changing the value of λ_{DRL} . As shown in Table 3, when $\lambda_{DRL} = 0.01$, the physical-inspired heuristic rule acts as an insufficient constraint, failing to give enough guidance. When λ_{DRL} increases to 1.0, the network tends to chase a shortcut. It tries to match the mean and variance numbers perfectly while ignoring the adversarial loss. This makes the final generated graphite ore images exhibit simplified textures and unnatural artifacts, lacking any sense of visual depth. Consequently, the downstream classification accuracy drops noticeably. This drop occurs because the classifier cannot extract distinct texture features from these flattened surfaces. The lack of surface details directly confuses the feature recognition network. The results prove that setting $\lambda_{DRL} = 0.1$ pushes the generation metrics to the optimal state, reaching a great balance between adversarial pressure and physical-inspired heuristic guidance.

Table 3. Sensitivity analysis results of the loss weight λ_{DRL} .

λ_{DRL}	FID	KID	SSIM	PSNR
0.01	92.65	0.0572	0.915	28.19
0.1	87.97	0.0479	0.973	30.62
1.0	96.12	0.0594	0.864	21.95

Looking closely at Table 3, we can see a clear tug-of-war between our statistical constraints and the generative freedom of the GAN. When λ_{DRL} is too small, the restriction is too loose; the model basically ignores the unique macro-reflectance of high-grade graphite, leading to sub-optimal results. On the flip side, cranking the weight up to 1.0 backfires completely. The network tries to take a shortcut—it perfectly matches the pixel intensity numbers to satisfy the math formula but completely forgets how to draw the intricate, flaky details of the rock. This causes a massive collapse in structural metrics, where SSIM drops heavily to 0.864 and PSNR falls to 21.95, leaving us with flat, artificial images that lack visual depth. Setting $\lambda_{DRL} = 0.1$ hits the perfect sweet spot. It acts as a gentle guide that nudges the brightness and contrast toward a realistic look without getting in the way of the adversarial loss crafting genuine mineral textures.

Additionally, we tested the EMA momentum parameters for weight updates. As reported in Table 4, training without EMA or with a lower value causes strong oscillations because adversarial training is non-stationary. Meanwhile, a value too high leads to overly conservative updates and delays convergence. Setting it to 0.99 provides the smoothest historical tracking and the best classification performance. The summarized experimental results are shown in Table 4.

Table 4. Empirical testing results of the EMA momentum.

EMA Momentum	FID	KID	SSIM	PSNR
No EMA	95.62	0.0638	0.895	27.19
0.90	90.86	0.0515	0.955	29.23
0.99	87.97	0.0479	0.973	30.62
0.999	94.31	0.0564	0.924	28.42

The data in Table 4 highlights just how crucial training stability is when dealing with complex, natural mineral images. Without EMA, or with a low momentum value like 0.90, the model gets caught in intense gradient swings. This is caused by the highly non-stationary nature of GAN training. The generator’s weight fluctuates too wildly from batch to batch. This results in blurry rock surfaces and noisy artifacts. These issues push the FID value up to 95.62. However, pushing the momentum too close to 1, such as 0.999, makes the model overly conservative and sluggish. It relies too much on its past states and fails to adapt quickly enough to the discriminator’s real-time feedback. This slows down the convergence process. An EMA momentum of 0.99 works like an ideal shock absorber. It filters out random training spikes to keep the core features stable. At the same time, it remains nimble enough to capture the true structural layout of the graphite ore. This setting achieves the highest scores across all metrics.

4.3. Downstream Classification Tasks

To verify the practical value of the high-grade graphite images generated in this paper, we run a downstream grade classification experiment for quantitative evaluation. ConvNeXt-Base serves as the classification network. It is a modern CNN architecture proposed by Liu et al. in 2022 [38]. We load its pre-trained weights from ImageNet-22K and replace the last fully connected layer with binary heads for high grade and low grade. All input images are uniformly resized to 224×224 . Training runs for 50 rounds, the batch size is fixed at 16, and the learning rate is $1e-4$. We first pool all real low-grade and high-grade graphite images together. The dataset includes 1084 real low-grade images, 337 real high-grade images and 80 generated images. These 80 synthetic high-grade images are obtained by randomly selecting 80 real low-grade images from the training set and translating them using our proposed model. All generated images are directly added to the training set

without manual selection or filtering to ensure objective data augmentation. We then split it into training, validation and test sets at a ratio of 7:1.5:1.5. Three groups of control schemes are set up. The Baseline scheme trains the classification model on real images only. An ablation scheme trains the model on generated high-grade images together with real low-grade images. The Ours scheme takes the real high-grade data and supplements it with generated images to form an enhanced training set. The three groups share the same validation set and the same test set. The generated high-grade graphite images are added only to the high-grade portion of the training set and never appear in validation or testing, so the comparisons stay fair and credible. Table 5 shows the specific dataset split, where real means real graphite ore images, fake means generated fake high-grade images, and Low and high indicate low grade and high grade.

Table 5. Dataset partitioning details.

Experiment	Training Set (Low-Grade)	Training Set (High-Grade)	Validation Set	Test Set
Baseline	758 (real)	236 (real)	163 low + 51 high	163 low + 50 high
Ablation	758 (real)	80 (fake)	163 low + 51 high	163 low + 50 high
Ours	758 (real)	316 (236 real + 80 fake)	163 low + 51 high	163 low + 50 high

Table 6 summarizes the classification outcomes of the three schemes on the test set, and Figure 7 offers a more intuitive presentation. Trained exclusively on real data, the Baseline model attains a test-set accuracy of 97.24%, with precision, recall, and F1 score each reaching 0.9816 and an AUC of 0.9990, reflecting excellent overall classification capability. When the Ablation scheme is trained solely on generated high-grade images, it still yields an accuracy of 93.55%. This finding directly demonstrates that the improved CycleGAN can effectively capture the core visual features of high-grade graphite, and that the generated images share representational characteristics with real samples. Grade classification based only on generated samples remains stable and dependable within this experimental scope. The Ours scheme delivers the strongest performance, with a test-set accuracy of 99.54%, precision of 99.39%, and perfect recall of 100%, meaning that no high-grade sample is missed. The F1 score reaches 0.9969 and the AUC reaches 0.9999. Compared with the Baseline, the Ours scheme improves classification accuracy by roughly 2.3 percentage points, and all evaluation metrics show clear improvement.

Table 6. Results of classification experiments.

Experiment	Accuracy (%)	Precision (%)	Recall (%)	F1-Score	AUC
Ablation	93.55 ± 0.11	96.27 ± 0.05	95.09 ± 0.23	0.9568 ± 0.0012	0.9743 ± 0.0004
Baseline	97.24 ± 0.09	98.16 ± 0.07	98.16 ± 0.15	0.9816 ± 0.0010	0.9990 ± 0.0001
Ours	99.54 ± 0.06	99.39 ± 0.02	100.00 ± 0.00	0.9969 ± 0.0003	0.9999 ± 0.0000

Based on the preceding analysis, high-grade images synthesized by the improved CycleGAN model effectively enhance the performance of downstream classification models under offline evaluation conditions. To address the scarcity of authentic high-grade graphite ore samples, the generated images augment training data diversity, thereby mitigating the limited generalization capacity characteristic of small-sample scenarios. Concurrently, the ablation experiment isolates the contribution of high-fidelity synthetic samples from naive data stacking, directly validating that the features of the generated images are empirically authentic and functionally effective within the tested dataset. From an engineering perspective, although the baseline already achieves a high accuracy of 97.24%, the practical significance of the 2.3-point gain remains noteworthy. This marginal improvement corre-

sponds to a substantial reduction in the error rate within the high-accuracy zone, helping the classifier better resolve long-tail visual ambiguities. More importantly, our scheme achieves a high recall of 100% on the test set, substantially minimizing the false negative rate for high-grade graphite ores. In actual industrial sorting pipelines, high-grade graphite is exceptionally valuable and scarce; the baseline's missing rate implies that a portion of premium graphite ores might be inadvertently routed into low-grade tailings, leading to cumulative economic losses. Therefore, suppressing this critical marginal error and approaching a near-zero leakage rate for premium graphite minerals significantly strengthens the operational viability of deploying automated visual inspection systems. In conclusion, the high-grade graphite ore images generated in this study demonstrate substantial practical utility, consistently improve downstream grade classification performance, and establish a reliable data augmentation methodology for visual grade identification tasks in graphite ore analysis. Nevertheless, these promising findings are inherently bounded by laboratory conditions; due to the current lack of edge-deployment optimization, the actual robustness of this methodology under real-world, dynamic 3D industrial conveyor belt environments remains to be further validated.




Figure 7. Results of the classification experiment.

4.4. Visualization Experiment

To explicitly validate the class of the generated images, our visual inspection moves beyond general realism and specifically targets key visual attributes that define high-grade graphite. In such geological imagery studies, a successful class translation must exhibit three core attributes: a uniform pure black or dark gray tonal distribution, a distinct metallic luster indicative of crystalline structure, and the preservation of original edge contours without structural collapse. Figure 8 displays the translation results from the improved model. Each example groups six related images together. *real_A* shows a real low-grade graphite ore image. *fake_B* presents the generated high-grade version. *rec_A* gives the reconstructed low-grade image that comes from *fake_B*. The opposite direction follows the same pattern. *real_B* holds a real high-grade sample. *fake_A* shows the generated low-grade counterpart. *rec_B* completes the cycle by reconstructing the high-grade input from *fake_A*. We examine these outputs to evaluate the model. The visual inspection targets three specific aspects. We first check the quality of the generated high-grade image. We then compare both reconstructed images against their original inputs. Finally, we assess the realism of the generated low-grade sample. Real high-grade graphite ore looks quite different from low-grade material. High-grade samples usually appear pure black

or dark gray. The surface maintains a uniform and rich color. A distinct metallic luster gives it a brighter overall look. Low-grade graphite ore shows a different pattern. The color distribution looks uneven across the surface. Grayish-black or grayish-white tones dominate the visual field. The rock matrix contains many light-colored gangue minerals. These inclusions scatter light and produce a dull surface finish. Weak reflection dominates the overall texture. Only a few isolated flakes show strong local reflection. Oxidation further alters the low-grade material. The chemical weathering exposes secondary substances on the rock face. Yellow or red stains frequently appear in these weathered zones.




Figure 8. Image of the result after the improved model transformation.

The images show that the improved model maintains stable visual performance during the mapping from low-grade graphite ore to high-grade regions. The result in fake_B approaches real high-grade samples in both color distribution and surface reflection. The gray levels and metallic luster appear close to those of the target domain, which suggests effective learning at the distribution level. A closer view reveals that the surface texture of the graphite ore is largely preserved during the transformation. Traditional models often produce blurred textures or structural collapse. In contrast, the generated results here keep clear edge contours and consistent particle distribution. No obvious breaks or over-smoothing appear in the structure. The reconstruction results rec_A and rec_B support this observation. The model completes cross-domain mapping while retaining the original structural information. The cycle consistency remains reliable throughout the process. Some local differences are still visible. Certain regions show slight color shifts or enhanced contrast. These changes reflect a bias toward emphasizing target domain features. However, the overall effect remains controlled. The model does not introduce noticeable artifacts or unrealistic structures.

Synthesizing multiple groups of samples, it can be found that the above phenomena have good consistency and are not unique. This shows that the improved model has certain generalization ability under different graphite ore morphology and texture conditions. In other words, it can not only realize the transformation from low grade graphite ore image to high grade graphite ore image, but also become more stable, which is more meaningful for the actual industrial scene.

4.5. Comparison Experiments

To rigorously validate the superiority of the proposed framework Reflectance-Consistent CycleGAN, a comprehensive comparative analysis was conducted against classical data augmentation, algorithmic class-balancing strategies, and advanced unpaired generative baselines in the context of graphite ore grade recognition. To ensure fairness, all comparative methods were evaluated using the identical downstream classifier ConvNeXt-Base under uniform baseline hyperparameter settings, where the batch size was 16, with no additional optimizations or empirical adjustments. Detailed comparison results are presented in Table 7. To eliminate random errors inherent in deep learning training, all values reflecting classification accuracy and performance metrics in the table represent the average of three independent experiments.

Table 7. Results of comparisons with other models.

Experiment	Accuracy (%)	Precision (%)	Recall (%)	F1-Score	AUC
Baseline	97.24 ± 0.21	98.16 ± 0.15	98.16 ± 0.11	0.9816 ± 0.0012	0.9990 ± 0.0006
Classical Augmentation	97.85 ± 0.12	98.40 ± 0.06	98.25 ± 0.09	0.9832 ± 0.0004	0.9992 ± 0.0001
Oversampling + Class-Balanced	97.56 ± 0.08	97.10 ± 0.03	99.12 ± 0.05	0.9810 ± 0.0005	0.9991 ± 0.0004
CUT	98.62 ± 0.13	98.95 ± 0.07	99.10 ± 0.10	0.9902 ± 0.0014	0.9995 ± 0.0002
Ours	99.54 ± 0.09	99.39 ± 0.03	100.00 ± 0.00	0.9969 ± 0.0001	0.9999 ± 0.0000

In terms of implementation, the Baseline scheme directly utilizes the original unbalanced real graphite dataset for classifier training. Classical Augmentation introduces random flips and color jittering during training iterations. The oversampling + Class-Balanced strategy duplicates rare high-grade real images to achieve a one-to-one balance and modifies the loss weight. Both the CUT and Ours methods train generative networks based on the real low-grade and high-grade images, translating low-grade images into pseudo-high-grade samples to expand the training set for the downstream ConvNeXt-Base classifier. Here, Ours uniquely applies distributed reflectance consistency loss statistical constraints.

As shown in Table 7, Ours outperforms other models in terms of classification metrics, achieving an accuracy of 99.54% and an F1-Score of 0.9969. Compared with Baseline, Classical Augmentation, oversampling + Class-Balanced, and CUT, Ours improves the classification accuracy by 2.30%, 1.69%, 1.98%, and 0.92% respectively. This demonstrates its strong performance in addressing grade imbalance. Notably, CUT relies entirely on a data-driven generation paradigm. Benefiting from the introduced reflectance constraints, our method achieves further gains in precision and recall. This verifies that embedding domain-specific physics-inspired heuristic into generative data augmentation can effectively boost model performance. The intuitive comparison of all evaluation metrics across different methods is presented in Figure 9.




Figure 9. Quantitative comparison of performance metrics for various approaches.

4.6. Generalization Experiment

To evaluate the generalization ability of the proposed model beyond graphite ore, this study tests the improved CycleGAN on a public rock and mineral image dataset. The dataset is the Type of Rocks and Minerals Dataset from Kaggle [39], released by Usta et al. in 2024. It contains about 10,000 rock images. The data cover three main rock types, including igneous, metamorphic, and sedimentary rocks. The dataset also includes diverse mineral textures, bedding structures, and surface reflection characteristics. This dataset does not provide explicit grade annotations. This study therefore constructs pseudo-low-grade and pseudo-high-grade domains based on color, gloss, and texture uniformity of the mineral surface. The pseudo-low-grade samples mainly show light gray tones, mottled patterns, rough textures, and weak luster. These samples are mainly selected from sedimentary rocks and metamorphic rocks containing mottled gangue. The pseudo-high-grade samples show dark black color, uniform appearance, dense texture, and clear reflection features. These samples are selected from dark and dense metamorphic rocks. This setting simulates the visual transformation process from low-grade to high-grade minerals. This grouping principle aligns with general mineralogical observations, where dark and dense components often indicate a higher concentration of target minerals. However, this preliminary division inevitably introduces some human subjectivity.

The results are presented in Figure 10. Visual inspection shows that the transformation from low-grade to high-grade remains stable. The fake_B column shows clear darkening or

graying compared with the original images. The results differ from the input images and move closer to real high-grade samples. The reconstruction results are less satisfactory. The reconstructed images do not reach the same level as those in the graphite ore experiments. This gap indicates that the model still has room for improvement. The fake_A column shows that most generated samples are successfully mapped back to the low-grade domain. Overall, the improved CycleGAN maintains stable performance in the generalization experiment. Some limitations remain and need further refinement. Specifically, the lack of quantitative geochemical assays for the public dataset limits the absolute precision of our pseudo-grade domains. This subjective grouping constraint represents a known limitation of the current validation stage.




Figure 10. Visual comparison of other datasets.

4.7. Limitations and Future Research Directions

Although the improved CycleGAN model shows good performance in converting low-grade graphite ore images to high-grade ones, several limitations remain. First, the dataset in this study is collected under specific conditions and exclusively from a single mining area, namely the Hegang graphite mine of China Minmetals Corporation. This setting limits the geographic and geological diversity of the data. It also restricts the representation of variations under different shooting conditions, backgrounds, geological environments, and mineral compositions. It makes it challenging to fully verify the model's adaptability to graphite ore samples with different geological origins and different mineral association types. Second, the model improves performance through the attention mechanism and the distribution reflectance Consistency Loss constraint. However, the FID and KID values of the generated results remain relatively high. This result suggests that the generated images are not yet fully aligned with real high-grade graphite ore images in high-dimensional feature distribution. Methodologically, the proposed distribution reflectance Consistency loss relies on exponential moving averages of real high-grade statistics, which makes training somewhat sensitive to the initial estimate and may not adapt quickly when the target distribution is multimodal. Also, it only enforces first- and second-order moments, leaving higher-order texture features like skewness or local patterns unregulated. As for the CBAMs inserted into three residual blocks of the generator, they inevitably increase parameters and inference time. Finally, the current application scenario is simplified. The study focuses on two-dimensional static graphite ore images. It does not consider complex interferences in industrial environments, such as dust, illumination disturbance, and motion blur. In addition, the model is not optimized for lightweight deployment. This limitation makes it difficult to apply directly to real-time detection equipment in mineral processing sites.

Future work will build on the limitations discussed above and move forward along three directions: data, model design, and practical deployment. On the data side, although the underlying robustness of the algorithm has been verified through cross-dataset generalization experiments and multiple independent repeated tests, future research still needs to collect larger-scale graphite ore samples from mining areas with diverse geographical features and varied environments. This will enable further evaluation and optimization of the practical value of the proposed framework in complex industrial scenarios. For the model, the focus will shift to distribution alignment. The goal is to further narrow the gap between generated and real samples and to reduce FID and KID values. In addition, using regularization techniques such as dropout or stochastic depth can help mitigate overfitting when training on a small dataset. In terms of application, more attention will be given to deployment in real scenarios. The model will be adapted for lightweight use through techniques such as pruning and quantization of the generative network, which will further reduce parameters and inference time. These changes will support deployment on edge computing devices and help move the method from controlled experiments toward practical engineering use.

Finally, the proposed framework is mainly designed for real-time online grade monitoring. Its target application scenario is conveyor belts in industrial sorting plants. Actual industrial field environments contain various inevitable interference factors. These factors include dust coverage, uneven illumination, motion blur, and graphite ore stacking occlusion. The dynamic DRL constraint and the CBAM attention mechanism are integrated in our method. This design provides theoretical robustness against global style drifting and local texture blurring. However, we have not yet fully conducted systematic evaluations targeting these specific operational noise factors at the current stage. The next core step of our research is clear. We will deploy the model in a simulated industrial pipeline. We will

stress-test its industrial adaptability. We will also explicitly optimize its inference stability under complex multi-interference field conditions of graphite ore.

5. Conclusions

High-grade graphite samples are scarce in field exploration of the Hegang graphite mine in Heilongjiang Province. Expanding the sample size of high-grade graphite ore therefore plays an important role in improving the accuracy of the grade recognition model. Deploying the optimized grade recognition model at the mine site can drive the digital transformation of mining operations. It enables precise sorting and processing of graphite ore and raises the utilization efficiency of graphite resources. At the same time, refined and intelligent mining can significantly cut manpower and material consumption, which helps realize green and efficient mining. To address these challenges, this paper proposes an improved CycleGAN model. The model converts low-grade graphite ore images into high-grade graphite ore images. This expands the sample size of high-grade graphite ore and in turn improves the accuracy of the grade recognition model. The main contributions of this study include:

1. We built a high-quality graphite ore image dataset through a systematic data preparation process in a real industrial scenario. Graphite ore images under different lighting conditions were collected with high-resolution imaging equipment. Each image was manually labeled and its carbon and sulfur content was measured using a carbon-sulfur analyzer. The resulting dataset contains diverse samples that cover two types of graphite ore morphology and two grade levels: high grade and low grade.
2. We introduce a distributed reflectance consistency loss, based on the empirical appearance characteristics of graphite ore, into the basic CycleGAN model. Overfitting to local texture differences is avoided by dynamically maintaining and smoothly updating the global pixel mean and global standard deviation of the real high-grade graphite ore dataset using exponential moving averages. We also insert CBAMs into the 4th, 5th, and 6th intermediate layers of the 9 residual blocks in the CycleGAN generator. This brings in an attention mechanism that lets the neural network focus on important channels and locations simultaneously, dynamically adjusting the feature maps.
3. To verify the effectiveness of the generated data and to evaluate the performance impact of each module, we designed a downstream classification and recognition task that uses the ConvNeXt architecture. We then carried out several experiments. Ablation experiments evaluated the contribution of each module to model performance, and visual comparisons further demonstrated the advantages of the improved model in this task. We also ran generalization experiments on other publicly available datasets to verify the model's generalization ability and robustness.
4. Compared with the base model, the improved CycleGAN outperforms it by roughly 21.9% in FID, 39.4% in KID, 2.6% in SSIM, and 3.3% in PSNR. Adding the newly generated samples to the grade recognition model brings a 6.9% gain. These results highlight the improved model's advantage in image conversion for graphite mining and confirm that the experimental findings extend to downstream grade recognition tasks.

In summary, this work presents a low-sample data augmentation method for graphite ore image translation. Experimental results indicate that the framework can effectively improve the visual quality of the data generated. Sensitivity analysis and ablation studies further verify that each module maintains a stable contribution to model performance. The consistent trends of experimental data lend reasonable support to the practical feasibility of our mathematical constraints. Accordingly, this method can serve as a reliable tool to assist downstream grade recognition networks in intelligent mining. Furthermore, it should be noted that the primary task is restricted to low- and high-grade binary translation. We

clarify that this binary focus is a purposeful design targeted at the relative data scarcity of high-grade graphite in real-world industrial sorting. Isolating these polar domains provides a relatively high-contrast learning signal that improves the practical utility of the domain translation framework. Future work will continue to refine this method and explore additional technical improvements, with a focus on lightweight design. These efforts aim to reduce the model's reliance on a single dataset and improve its generalization across different scenarios and devices. At the same time, the image translation model will be integrated with intelligent mineral processing and online grade detection systems. The study will further extend to more complex industrial settings, including dynamic video data and hyperspectral images.

Author Contributions: Conceptualization, C.L., L.C. and X.H.; methodology, X.H. and B.W.; software, C.L. and L.C.; validation, C.L. and L.C.; resources, L.C.; data curation, C.L. and L.C.; writing—original draft preparation, C.L.; writing—review and editing, C.L.; visualization, C.L.; supervision, X.H. and B.W.; project administration, X.H. and B.W.; funding acquisition, X.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Ministry of Science and Technology of the People's Republic of China under the National Key Research and Development Program of China (Grant 2020YFB1713700).

Data Availability Statement: The raw data supporting the conclusions of this article are available from the corresponding author upon request. The Type of Rocks and Minerals dataset is available online at: <https://www.kaggle.com/datasets/kutlayusta/type-of-rocks-and-minerals-dataset> (accessed on 2 May 2026).

Acknowledgments: The author sincerely thanks Xueyu Huang for his guidance on the research direction and the provision of financial support for this work. The author also appreciates the support from China Minmetals Group (Heilongjiang) Graphite Industry Co., Ltd., which provided ore samples and research facilities. In addition, special gratitude is extended to Jun-Yan Zhu and his research team for their pioneering work on the CycleGAN algorithm, which forms the core methodological foundation for the graphite ore image augmentation research presented in this paper.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AUC	Area Under the Receiver Operating Characteristic Curve
CBAM	Convolutional Block Attention Module
CycleGAN	Cycle-Consistent Generative Adversarial Network
DRL	Distributed Reflectance Consistency Loss
EMA	Exponential Moving Average
FID	Fréchet Inception Distance
KID	Kernel Inception Distance
PSNR	Peak Signal-to-Noise Ratio
SEM-EDS	Scanning Electron Microscopy with Energy-Dispersive Spectroscopy
SSIM	Structural Similarity Index Measure
XRD	X-ray Diffractions

References

1. Keeling, J. Graphite: Properties, Uses and South Australian Resources. *MESA J.* **2017**, *84*, 28–41.
2. Qiu, Z.; Huang, X.; Li, S.; Wang, J. Stellar-YOLO: A Graphite Ore Grade Detection Method Based on Improved YOLO11. *Symmetry* **2025**, *17*, 966. [[CrossRef](#)]

3. Cao, Y.; Xing, C.-M. Textural Relationships and Mineral Compositions of Fe-Ti Oxide Ores of the Panzhihua Layered Intrusion, SW China: Insights for Magma Emplacement and Lateral Extension of Fe-Ti Oxide Ore Bodies. *J. Asian Earth Sci.* **2023**, *255*, 105772. [\[CrossRef\]](#)
4. Park, J.; Cho, S.-J.; Shin, S.; Kim, R.; Shin, D.; Shin, Y. Overview of Graphite Supply Chain and Its Challenges. *Geosci. J.* **2025**, *29*, 329–341. [\[CrossRef\]](#)
5. Cao, J.; Zhong, C.-G.; You, H.-T.; Zhang, Y.; Wang, R.-B.; Zhou, S.-M.; Qu, J.-H.; Chen, R.; Liu, S.-L. High-Precision Copper-Grade Identification via a Vision Transformer with PGNA. *Nucl. Sci. Tech.* **2025**, *36*, 120–130. [\[CrossRef\]](#)
6. Ren, Y.; Yang, Y.; Wang, J.; Pan, C.; Yuan, F.; Chen, W.; Wang, J. A Sensor-Based Magnetite Ore Sorting System Integrating Empirical Mode Decomposition and Convolutional Neural Network. *Minerals* **2026**, *16*, 210. [\[CrossRef\]](#)
7. Liu, Z.-N.; Lu, C.-L.; Tian, R.; Deng, Y.-Y.; Liu, Z.-H.; Zhang, P.-W. Evaluation Method of Ore Grade Estimation Effectiveness. *PLoS ONE* **2024**, *19*, e0309696. [\[CrossRef\]](#) [\[PubMed\]](#)
8. Jia, S.; Fu, J.; Leng, W.; Cui, S.; Zhao, B.; Meng, J. Identification and Location of High-Grade Iron Deposits in the Anshan-Benxi Area, China, Using Integrated Geophysical Techniques. *Ore Geol. Rev.* **2026**, *188*, 107056. [\[CrossRef\]](#)
9. Bowles, C.; Chen, L.; Guerrero, R.; Bentley, P.; Gunn, R.; Hammers, A.; Dickie, D.A.; Hernández, M.V.; Wardlaw, J.; Rueckert, D. GAN Augmentation: Augmenting Training Data Using Generative Adversarial Networks. *arXiv* **2018**, arXiv:1810.10863. [\[CrossRef\]](#)
10. Li, W.; Wang, C.; Guan, H.; Wang, J.; Yang, J.; Zhang, C.; Tao, D. Generative Adversarial Physics-Informed Neural Networks for Solving Forward and Inverse Problem with Small Labeled Samples. *Comput. Math. Appl.* **2025**, *183*, 98–120. [\[CrossRef\]](#)
11. Goodfellow, I.J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Networks. *arXiv* **2014**, arXiv:1406.2661. [\[CrossRef\]](#)
12. Frid-Adar, M.; Klang, E.; Amitai, M.; Goldberger, J.; Greenspan, H. Synthetic Data Augmentation Using GAN for Improved Liver Lesion Classification. *arXiv* **2018**, arXiv:1801.02385. [\[CrossRef\]](#)
13. Han, C.; Murao, K.; Noguchi, T.; Kawata, Y.; Uchiyama, F.; Rundo, L.; Nakayama, H.; Satoh, S. Learning More with Less: Conditional PGGAN-Based Data Augmentation for Brain Metastases Detection Using Highly-Rough Annotation on MR Images. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management, Beijing, China, 3–7 November 2019; pp. 119–127.
14. Oubara, A.; Wu, F.; Amamra, A.; Yang, G. Survey on Remote Sensing Data Augmentation: Advances, Challenges, and Future Perspectives. In *Advances in Computing Systems and Applications*; Senouci, M.R., Boulahia, S.Y., Benatia, M.A., Eds.; Lecture Notes in Networks and Systems; Springer International Publishing: Cham, Switzerland, 2022; Volume 513, pp. 95–104.
15. Zhu, J.-Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. *arXiv* **2020**, arXiv:1703.10593. [\[CrossRef\]](#)
16. Sandfort, V.; Yan, K.; Pickhardt, P.J.; Summers, R.M. Data Augmentation Using Generative Adversarial Networks (CycleGAN) to Improve Generalizability in CT Segmentation Tasks. *Sci. Rep.* **2019**, *9*, 16884. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Hage Chehade, A.; Abdallah, N.; Marion, J.-M.; Hatt, M.; Oueidat, M.; Chauvet, P. Advancing Chest X-Ray Diagnostics: A Novel CycleGAN-Based Preprocessing Approach for Enhanced Lung Disease Classification in Chest X-Ray14. *Comput. Methods Programs Biomed.* **2025**, *259*, 108518. [\[CrossRef\]](#) [\[PubMed\]](#)
18. Zhu, J.; Ma, L.; Ma, X.; Xiong, Z.; Wu, M.; Li, Y. Enhancing Computer Vision-Based Bridge Traffic Identification at Nighttime through CycleGAN-Enabled Data Augmentation. *J. Civ. Struct. Health Monit.* **2026**, *16*, 13–32. [\[CrossRef\]](#)
19. Wei, B.; Zheng, Z.; Hu, Y.; Yuan, D.; Chi, Y.; Qiu, C. Underwater Concrete Crack Detection of Dams via CycleGAN-Based Data and Optimized Multi-Scale YOLO11. *Constr. Build. Mater.* **2026**, *514*, 145497. [\[CrossRef\]](#)
20. Li, Z.; Zhan, Z.; Hu, J.; Yi, S.; Zhang, X.; Weng, Z.; Zhang, Z.; Ding, K. An Adaptive Generalized Regression Neural Network Approach for Ore Grade Estimation Considering Spatial Anisotropy. *Nat. Resour. Res.* **2025**, *34*, 2423–2442. [\[CrossRef\]](#)
21. Paty, S.; Kamilya, S. Identification of Different Types of Rocks and Ores Using Cellular Automata Assisted CNN Models: An Application of Mining Industry. *Nat. Comput.* **2026**, *25*, 14–33. [\[CrossRef\]](#)
22. Liu, Y.; Wang, X.; Zhang, Z.; Deng, F. LOSN: Lightweight Ore Sorting Networks for Edge Device Environment. *Eng. Appl. Artif. Intell.* **2023**, *123*, 106191. [\[CrossRef\]](#)
23. Zhao, S.; Zhan, Y.; Niu, W. EU-Net and ACFS: An Effective Method for Segmenting Ore Images Collected on-Site. *J. King Saud Univ. Comput. Inf. Sci.* **2025**, *37*, 35–50. [\[CrossRef\]](#)
24. Skandarani, Y.; Jodoin, P.-M.; Lalande, A. GANs for Medical Image Synthesis: An Empirical Study. *arXiv* **2021**, arXiv:2105.05318. [\[CrossRef\]](#)
25. Dash, A.; Panigrahi, S.; Adhikary, D.R.D.; Swarnkar, T. Data Augmentation Strategies for GAN-Based Medical Image Analysis: An Empirical Review. *Syst. Sci. Control Eng.* **2026**, *14*, 2634452. [\[CrossRef\]](#)
26. Carmona, J.M.; Cortes, J.L.; Relancio, J.J. Relativistic Deformed Kinematics from Momentum Space Geometry. *Phys. Rev. D* **2019**, *100*, 104031. [\[CrossRef\]](#)

27. Isola, P.; Zhu, J.-Y.; Zhou, T.; Efros, A.A. Image-to-Image Translation with Conditional Adversarial Networks. *arXiv* **2018**, arXiv:1611.07004. [[CrossRef](#)]
28. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. *arXiv* **2015**, arXiv:1512.03385. [[CrossRef](#)]
29. Johnson, J.; Alahi, A.; Fei-Fei, L. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. *arXiv* **2016**, arXiv:1603.08155. [[CrossRef](#)]
30. Ulyanov, D.; Vedaldi, A.; Lempitsky, V. Instance Normalization: The Missing Ingredient for Fast Stylization. *arXiv* **2017**, arXiv:1607.08022. [[CrossRef](#)]
31. Ledig, C.; Theis, L.; Huszar, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. *arXiv* **2017**, arXiv:1609.04802. [[CrossRef](#)]
32. Woo, S.; Park, J.; Lee, J.-Y.; Kweon, I.S. CBAM: Convolutional Block Attention Module. *arXiv* **2018**, arXiv:1807.06521. [[CrossRef](#)]
33. Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv* **2015**, arXiv:1502.03167. [[CrossRef](#)]
34. Tarvainen, A.; Valpola, H. Mean Teachers Are Better Role Models: Weight-Averaged Consistency Targets Improve Semi-Supervised Deep Learning Results. *arXiv* **2018**, arXiv:1703.01780. [[CrossRef](#)]
35. Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.H.; Buchatskaya, E.; Doersch, C.; Pires, B.A.; Guo, Z.D.; Azar, M.G.; et al. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. *arXiv* **2020**, arXiv:2006.07733. [[CrossRef](#)]
36. French, G.; Mackiewicz, M.; Fisher, M. Self-Ensembling for Visual Domain Adaptation. *arXiv* **2018**, arXiv:1706.05208. [[CrossRef](#)]
37. He, K.; Fan, H.; Wu, Y.; Xie, S.; Girshick, R. Momentum Contrast for Unsupervised Visual Representation Learning. *arXiv* **2020**, arXiv:1911.05722. [[CrossRef](#)]
38. Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; Xie, S. A ConvNet for the 2020s. *arXiv* **2022**, arXiv:2201.03545. [[CrossRef](#)]
39. Type of Rocks and Minerals Dataset. Available online: <https://www.kaggle.com/datasets/kutlayusta/type-of-rocks-and-minerals-dataset> (accessed on 2 May 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.