remote sensing

MDPI

Article

Physically Consistent SAR Image Generation for Unseen Aspect Angles via Attributed Scattering Center Evolution

Zihao Jiang [1], Chao Liu [1,2,*], Zhangzeyu Xing [1], Yamei Wang [1] and Shuwen Xiao [1]

[1] Institute of Systems Engineering, Academy of Military Sciences (AMS), People's Liberation Army of China (PLA), Beijing 100082, China; nwpujzh@mail.nwpu.edu.cn (Z.J.); wlcxzzy@163.com (Z.X.); wangyamei2010@126.com (Y.W.); xiaosw18@163.com (S.X.)
[2] Department of Electronic Engineering, Tsinghua University, Beijing 100084, China
* Correspondence: chao-liu23@mails.tsinghua.edu.cn

Highlights

What are the main findings?

- A physically consistent SAR image generation framework is proposed for unseen aspect angles by introducing an ASC-inspired sparse scattering-structure prior.
- Joint target scattering and background statistical consistency constraints improve the physical realism of generated SAR images under limited-view conditions.

What are the implications of the main findings?

- Unseen-angle SAR generation should be addressed as a physically constrained scattering evolution problem rather than only a conditional image synthesis task.
- The proposed framework provides a physically interpretable solution for SAR view completion under sparse angular observations.

Abstract

Synthetic aperture radar (SAR) target images are highly sensitive to aspect angle, while practical data acquisition usually provides only sparse observations over limited viewpoints. This leads to severe data scarcity at unseen aspect angles and makes cross-angle generation prone to scattering-structure distortion and background statistical mismatch. Existing SAR image generation methods either focus on distribution matching without sufficiently exploiting scattering-related structural cues, or emphasize angle conditioning while failing to jointly preserve physically plausible dominant scattering-response variations and realistic background speckle statistics at unseen aspect angles. To address this issue, we propose a physically consistent framework for SAR image generation at unseen aspect angles. The proposed method introduces an ASC-inspired sparse scattering-structure prior to approximate the dominant scattering responses in the SAR image plane. Rather than performing full parametric ASC inversion, this prior serves as a differentiable and angle-aware structural proxy that guides the generator toward synthesizing SAR images with structurally plausible scattering layouts. In addition, a dual-consistency scheme is introduced to jointly enforce target-region scattering consistency and background-region statistical consistency, thereby improving the physical realism of the generated results in both the target and background regions. Extensive experiments under strict unseen-angle interpolation and hold-out protocols demonstrate that the proposed method consistently outperforms representative baselines in image fidelity, target-region scattering consistency, background statistical consistency, and angle-conditioned consistency. Further visualization and ablation studies verify the critical role of the ASC-inspired sparse scattering-structure prior in physically consistent SAR view completion.

Check for updates

Academic Editors: Shi Chen, Yicong Zhou, Jiaqi Ma, Jiang He, Kui Jiang and Xiangyong Cao

Received: 11 May 2026
Revised: 24 June 2026
Accepted: 1 July 2026
Published: 7 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) license.

Keywords: synthetic aperture radar; unseen-angle generation; attributed scattering centers; physically constrained modeling; background statistical consistency

1. Introduction

Synthetic aperture radar (SAR) has become an indispensable sensing modality for all-weather and day-and-night observation and plays an important role in target reconnaissance, disaster monitoring, and automatic interpretation because of its robustness to illumination and atmospheric conditions [1]. With the rapid development of deep learning, SAR image understanding has gradually evolved from hand-crafted feature engineering to data-driven representation learning, leading to substantial progress in target recognition, detection, and scene analysis [2–4]. Nevertheless, the effectiveness of such methods still relies heavily on the scale, diversity, and coverage of training data [5]. This issue is particularly critical for SAR target imagery since the image appearance of a target varies significantly with the observation geometry, especially the aspect angle, due to the view-dependent nature of electromagnetic scattering [6–8]. In practical applications, however, densely sampled multi-angle SAR observations are difficult to obtain because of acquisition cost, platform constraints, and limited observation opportunities. As a result, currently available datasets often suffer from insufficient angular coverage, restricted scene diversity, or limited target categories, resulting in a severe shortage of samples at unseen aspect angles for model training and evaluation [8,9]. This mismatch between the strong angular sensitivity of SAR imagery and the sparse angular coverage of available data has become a fundamental bottleneck for robust SAR interpretation and physically reliable cross-angle image generation.

This limitation is not merely a matter of insufficient sample quantity but reflects a more fundamental challenge arising from the SAR imaging mechanism itself [10]. Unlike natural images, SAR target images are formed by coherent electromagnetic scattering, and their variation across aspect angles corresponds to changes in dominant scattering structures rather than simple geometric or textural transformations [11,12]. Consequently, generating SAR images at unseen aspect angles is not a trivial interpolation problem in image space [13]. The key requirement is that the scattering structure should evolve in a physically plausible manner as the viewing angle changes [14]. Under sparse angular coverage and limited-data conditions, however, purely data-driven generative models tend to learn global intensity statistics or weakly angle-related texture cues, instead of capturing the underlying evolution of target scattering responses [15,16]. Representative adversarial generation frameworks, such as DCGAN [17], ACGAN [18], and StyleGAN2 [19], provide useful image-level synthesis baselines, but they do not explicitly model SAR-specific scattering-structure variation or background statistical consistency. As a result, generated images may appear visually similar to real SAR data at the distribution level, while still exhibiting unreasonable local bright responses, unstable angle controllability, or mismatched background speckle characteristics [20]. Such a discrepancy between visual realism and physical credibility is particularly problematic for SAR image generation because the value of generated samples depends not only on whether they look realistic but also on whether they remain consistent with the scattering mechanism across viewpoints [21,22]. This observation suggests that physically reliable unseen-angle SAR generation requires more than conventional conditional image synthesis [23].

To alleviate the lack of physical interpretability in purely data-driven generation, an emerging line of research has attempted to incorporate physics-related priors into SAR image synthesis and augmentation [24]. Early efforts mainly relied on electromagnetic

simulation and model-based signal generation, which can reproduce scattering characteristics with high physical fidelity but usually depend on accurate target models and incur high computational cost [25]. More recently, several studies have introduced physical information into deep generative frameworks by embedding scattering-center-related features, electromagnetic scattering constraints, or physics-inspired regularization terms [26,27]. These methods have shown that incorporating physically meaningful cues can improve training stability and generation realism, especially under limited-data conditions. However, most existing physics-driven approaches still use physical information as a static prior, an auxiliary feature branch, or a weak regularizer imposed on the final output, rather than constructing an angle-aware structural prior to describe the variation of dominant scattering responses with aspect angle. As a result, although such methods may enhance overall physical plausibility, they remain insufficient for unseen-angle generation, where the key issue is not merely satisfying certain physical statistics but also ensuring that dominant scattering responses vary in an angle-consistent and structurally plausible manner across viewpoints. This limitation indicates that physically informed SAR generation needs to move beyond static physics injection toward angle-aware modeling of dominant scattering-structure variations [28].

Beyond static physical-prior injection, recent studies have further explored 3D-aware neural rendering paradigms for SAR image generation and reconstruction. Differentiable SAR rendering reformulates the SAR image formation process in a differentiable manner, allowing image-domain reconstruction errors to be back-propagated to target geometry and scattering-related attributes [29]. Following the development of neural radiance fields, SAR-NeRF-type methods model multi-view SAR observations through implicit neural fields and differentiable projection processes, while Radar Fields extend the radiance-field formulation to SAR image collections [30,31]. More recently, Gaussian-splatting-based SAR methods have introduced explicit Gaussian primitives and SAR-specific differentiable rasterization for SAR image rendering and target reconstruction [32]. These methods are theoretically attractive because they provide a more explicit connection between multi-view observations, geometry-aware representations, and SAR image formation. However, they usually rely on relatively dense multi-view observations, accurate imaging geometry, or optimization of scene-level 3D representations, which may introduce additional data requirements and computational costs. Therefore, despite their strong potential for geometry-consistent SAR rendering, lightweight generation under sparse target views remains an important and complementary problem.

In addition to the above physics-informed and 3D-aware rendering paradigms, another important line of research has focused on cross-angle or angle-controllable SAR image synthesis by explicitly introducing aspect-angle information into deep generative models [33]. Representative studies have incorporated azimuth-angle conditions into adversarial or diffusion-based frameworks through conditional vectors, auxiliary angle constraints, controllable latent modulation, or intermediate angle-interpolation mechanisms, thereby improving angular controllability in generated SAR images [34,35]. These methods have demonstrated that aspect angle is not merely an auxiliary label but a key control variable for view-dependent SAR generation. Nevertheless, in most existing approaches, angle information is still primarily treated as a low-dimensional condition used to steer feature synthesis in latent space or image space. Although such designs may enhance viewpoint control and continuity, they do not necessarily ensure that dominant scattering structures evolve in a physically meaningful manner as the observation angle changes. In particular, when the training data provide only sparse angular coverage, the generator may still satisfy the angle condition through superficial appearance adaptation rather than learning the underlying evolution of scattering structures. Moreover, background

statistics are often neglected in angle-controllable generation, even though unrealistic speckle characteristics can significantly affect the physical credibility of generated SAR imagery [20]. Therefore, current angle-aware generation methods still fall short of jointly modeling aspect-angle control, scattering-structure evolution, and background statistical consistency for unseen-angle SAR generation.

Despite these advances, a critical gap remains in unseen-angle SAR generation. Existing studies have improved either the physical plausibility of generated SAR images by incorporating scattering-related priors and physics-inspired constraints or the controllability of viewpoint synthesis by explicitly encoding aspect-angle information into generative models [13,15,22,26]. However, they have not yet systematically addressed how to jointly model aspect-angle conditioning, dominant scattering-structure variation, and SAR background statistics within a unified framework [20]. For unseen-angle generation, it is insufficient to merely enhance image fidelity or improve viewpoint controllability because neither of them directly guarantees consistency with the underlying scattering mechanism as the observation angle changes. Under sparse angular coverage, a reliable generator should not only respond to the target aspect label but also preserve structurally coherent variations of dominant scattering responses while maintaining realistic speckle characteristics in the background region. Therefore, unseen-angle SAR generation is more appropriately formulated as a physically constrained view-completion problem rather than a conventional conditional image synthesis task [14]. This observation further indicates the need for an ASC-inspired structural proxy that can link aspect-angle variation with dominant scattering-structure variation, while simultaneously supporting the preservation of background statistical consistency.

Motivated by the above observations, we developed a physically consistent framework for SAR image generation at unseen aspect angles by introducing an ASC-inspired sparse scattering-structure prior. Classical attributed scattering center models [36,37] provide a compact parametric description of electromagnetic scattering responses. In this work, we do not aim to perform full parametric ASC inversion. Instead, we adopted the ASC concept as a physical inspiration and constructed a differentiable image-domain structural proxy to approximate the dominant scattering responses of SAR targets. In this way, angle control is no longer imposed only through low-dimensional conditional encoding but is also linked to an angle-aware sparse structural prior. Based on this representation, the proposed framework guides the generator to synthesize SAR images whose scattering structures remain physically plausible under viewpoint variation. Meanwhile, to improve the overall physical credibility of the generated results, the framework jointly considers target-region structure and background-region statistics rather than focusing exclusively on target appearance. Accordingly, unseen-angle SAR generation is formulated as a joint modeling problem involving aspect-angle conditioning, scattering-structure evolution, and background statistical consistency. This design provides a mechanism-aware solution for physically consistent SAR view completion under sparse angular observations. The main contributions of this work are summarized as follows:

1. We propose a physically consistent framework for unseen-angle SAR generation, enabling structurally plausible view completion under sparse angular observations.
2. We introduce ASC-inspired generator-side structural guidance together with target-region ASC-aware discrimination and background statistical consistency, improving both target-region structural plausibility and SAR background realism.
3. We validate the proposed method through comparative and ablation experiments using multiple evaluation metrics, further supporting its physical consistency and interpretability.

2. Proposed Methods

2.1. Problem Formulation

This work considers SAR image generation at unseen aspect angles under a fixed imaging configuration. Specifically, all samples are target chips acquired under a single band, a single polarization, and a fixed depression angle, so that the dominant source of appearance variation is the aspect angle rather than changes in sensing modality or imaging geometry. Let the training set be defined as

$$\mathcal{D}_{\mathrm{tr}} = \{(x_i, y_i, \theta_i)\}_{i=1}^{N} \tag{1}$$

where $x_i \in \mathbb{R}^{H \times W}$ denotes a SAR target image, $y_i \in \{1, 2, \ldots, C\}$ is the target-category label, and $\theta_i \in \Theta_{\mathrm{tr}} \subset [0, 2\pi)$ is the corresponding aspect angle. In practice, Θ_{tr} usually covers only a sparse subset of the full angular range, which makes image generation at unseen aspect angles particularly challenging.

Given a target category y, a latent variable $z \sim \mathcal{N}(0, I)$, and an unseen aspect angle $\theta^* \in \Theta_{\mathrm{u}}$, where

$$\Theta_{\mathrm{u}} = [0, 2\pi) \setminus \Theta_{\mathrm{tr}} \tag{2}$$

the goal of this work was to learn a conditional generator G that synthesizes a SAR image

$$\hat{x} = G(z, y, \phi(\theta^*)) \tag{3}$$

where $\phi(\theta^*)$ denotes the aspect-angle encoding. To preserve the periodicity and continuity of angular variables, the aspect angle is encoded as

$$\phi(\theta^*) = [\sin\theta^*, \cos\theta^*] \tag{4}$$

Unlike conventional conditional image synthesis, unseen-angle SAR generation should satisfy not only visual realism but also physical consistency. Let Ω_{t} and Ω_{b} denote the target region and the background region of a SAR image, respectively. Then, for a generated sample $\hat{x}$, the desired solution should satisfy the following three requirements:

1. **Angle-conditioned consistency**: $\hat{x}$ should be consistent with the target category y and the required aspect angle θ^*, while remaining faithful to the distribution of real SAR imagery.
2. **Target-region structural consistency**: the dominant scattering responses within Ω_{t} should vary in an angle-consistent and structurally plausible manner as the aspect angle changes, but without assuming strictly smooth trajectories of individual scattering centers.
3. **Background-region statistical consistency**: the speckle characteristics within Ω_{b} should remain statistically compatible with those of real SAR backgrounds.

Accordingly, unseen-angle SAR generation is formulated in this work as a physically constrained view-completion problem rather than a conventional conditional image synthesis problem. The following subsections present how this objective is realized through an ASC-inspired sparse scattering-structure prior and target–background dual physical consistency constraints.

2.2. Overall Framework

The overall framework of the proposed method is illustrated in Figure 1. Building upon the problem formulation in Section 2.1, the proposed approach addresses unseen-angle SAR generation through a unified, physically consistent framework composed of three tightly coupled components: ASC-inspired sparse structural guidance, target-region

scattering consistency modeling, and background-region statistical consistency modeling. These components are designed to jointly ensure that the generated SAR image is not only visually realistic but also physically plausible in both target structure and background statistics.

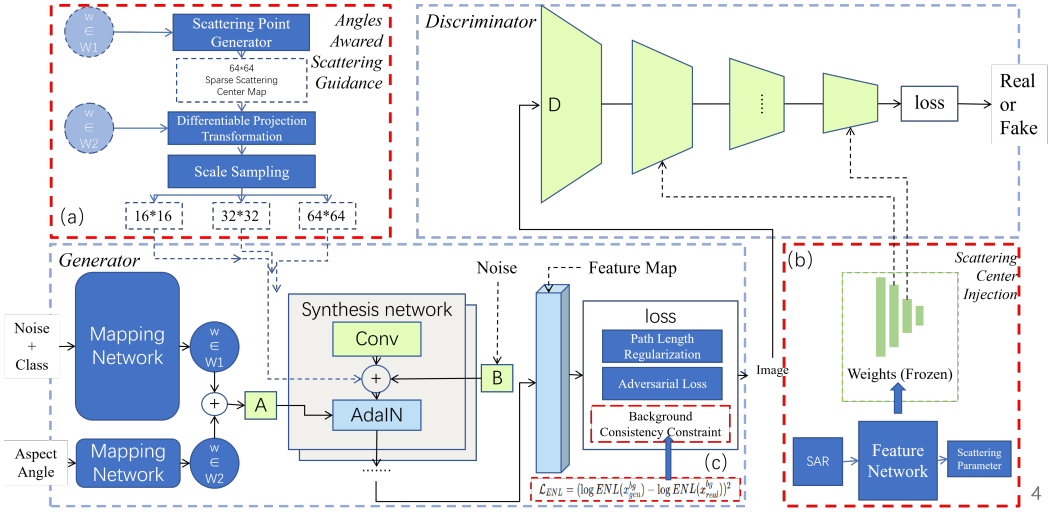


Figure 1. Overall architecture of the proposed framework. The framework consists of three tightly coupled components for physically consistent SAR image generation at unseen aspect angles: (**a**) ASC-inspired sparse structural guidance, which generates multi-scale sparse response maps to guide the synthesis of structurally plausible target scattering layouts; (**b**) target-region scattering consistency modeling, which injects frozen scattering priors into the discriminator to enforce structural plausibility of the generated target; and (**c**) background-region statistical consistency modeling, which imposes speckle statistics constraints to ensure realistic background texture. Together, these components are designed to improve both visual realism and physical consistency in terms of target structure and background statistics.

Given a latent code z, a target label y, and an aspect-angle encoding $\phi(\theta)$, the generator produces a SAR image $\hat{x}$ conditioned on both semantic and physical cues. Different from conventional conditional generators that rely solely on low-dimensional category and angle embeddings, the proposed framework introduces an ASC-inspired sparse structural branch to provide an angle-aware structural prior associated with aspect-angle variation. Let S_θ denote the angle-aligned structural representation inferred from the ASC-inspired sparse structural prior module. Then, the generation process can be written at a high level as

$$\hat{x} = G(z, y, \phi(\theta), S_\theta) \tag{5}$$

where S_θ serves as a physically meaningful guidance signal that constrains the scattering structure synthesized by the generator under viewpoint variation.

To ensure physical consistency in the target region, the discriminator is further enhanced with scattering-center-aware structural discrimination. Instead of judging realism only from image appearance statistics, the discriminator is encouraged to assess whether the generated target region exhibits physically reasonable scattering characteristics. In parallel, to suppress unrealistic background artifacts caused by adversarial learning, a background statistical consistency constraint is imposed on the generated image. This regularization explicitly aligns the statistical properties of the generated background with those of real SAR imagery, thereby improving the overall physical credibility of the synthesized results.

From a system perspective, the proposed framework can be understood as a joint modeling strategy that couples angle-conditioned generation with physically meaningful structural evolution. The ASC-inspired sparse structural prior provides structural guidance to the generator, the target-region consistency mechanism supplies physically informed

discrimination during adversarial learning, and the background statistical constraint regularizes speckle-related statistics outside the target region. Accordingly, the entire framework can be summarized as the following optimization problem:

$$\min_{G}\max_{D}\ \mathcal{L}_{\mathrm{adv}}+\lambda_{\mathrm{t}}\mathcal{L}_{\mathrm{tgt}}+\lambda_{\mathrm{b}}\mathcal{L}_{\mathrm{bg}} \tag{6}$$

where $\mathcal{L}_{\mathrm{adv}}$ denotes the adversarial objective, $\mathcal{L}_{\mathrm{tgt}}$ represents the target-region physical consistency term, and $\mathcal{L}_{\mathrm{bg}}$ denotes the background statistical consistency term. The detailed formulations of these components are presented in the following subsections.

Overall, the proposed framework does not treat unseen-angle SAR generation as a purely appearance-driven conditional synthesis problem. Instead, it formulates the task as a physically constrained view-completion problem in which aspect-angle control, scattering-structure evolution, and target–background consistency are modeled within a unified generative process.

2.3. ASC-Inspired Sparse Scattering-Structure Prior

To connect aspect-angle variation with target-region structural generation, we introduce an ASC-inspired sparse scattering-structure prior to provide the generator with angle-aware structural guidance. Different from conventional conditional generation, where aspect angle is injected only as a low-dimensional control variable, the proposed module transforms aspect-angle information into a sparse structural representation that varies with the viewpoint condition. In this way, the generator is guided not only by semantic conditions but also by a physically inspired image-domain structural proxy for dominant scattering-response variation.

It should be clarified that the proposed module does not aim to perform full parametric ASC inversion in the strict electromagnetic sense. Classical ASC models usually describe target scattering using parameters such as location, amplitude, phase, orientation, length, and frequency-dependent scattering behavior. In contrast, the present work adopts the ASC concept as a physical inspiration for generative modeling. Specifically, we construct a sparse image-domain structural proxy composed of a limited number of dominant response components, where each component is characterized by its spatial location, response strength, and spatial spread. This representation is therefore referred to as an ASC-inspired sparse scattering-structure prior. Its purpose is not to recover complete electromagnetic ASC parameters but to provide a differentiable and angle-aware structural guidance signal for unseen-angle SAR image generation.

2.3.1. ASC-Inspired Sparse Structural Representation

Let

$$w=M(z,y,\phi(\theta)) \tag{7}$$

denote the intermediate latent code produced by the mapping network $M(\cdot)$, where z is the latent variable, y is the target label, and $\phi(\theta)$ is the aspect-angle encoding. Based on w, a lightweight ASC-inspired structural head $H(\cdot)$ predicts a set of sparse structural attributes:

$$\mathcal{A}=\{(u_k,v_k,a_k,\sigma_k)\}_{k=1}^{K} \tag{8}$$

where $(u_k,v_k)\in[-1,1]^2$ denotes the normalized spatial location of the k-th scattering component, a_k denotes its response strength, and σ_k controls the spatial spread of the corresponding structural response. Here, K is a predefined small constant that limits the structural branch to representing only dominant scattering patterns rather than full image textures.

The predicted sparse attributes are then rasterized into a structural map S_0 through a differentiable Gaussian aggregation process:

$$S_0(p) = \sum_{k=1}^{K} a_k \exp\left(-\frac{\|p - c_k\|_2^2}{2\sigma_k^2}\right), \qquad c_k = (u_k, v_k) \tag{9}$$

where p denotes a spatial coordinate in the normalized image plane. The resulting map S_0 provides a compact image-domain structural description of dominant scattering responses. Since the branch is restricted to a sparse set of response components, it is prevented from degenerating into a secondary image generator and instead serves as a compact ASC-inspired structural prior. The schematic diagram of the implementation is presented in Figure 2.

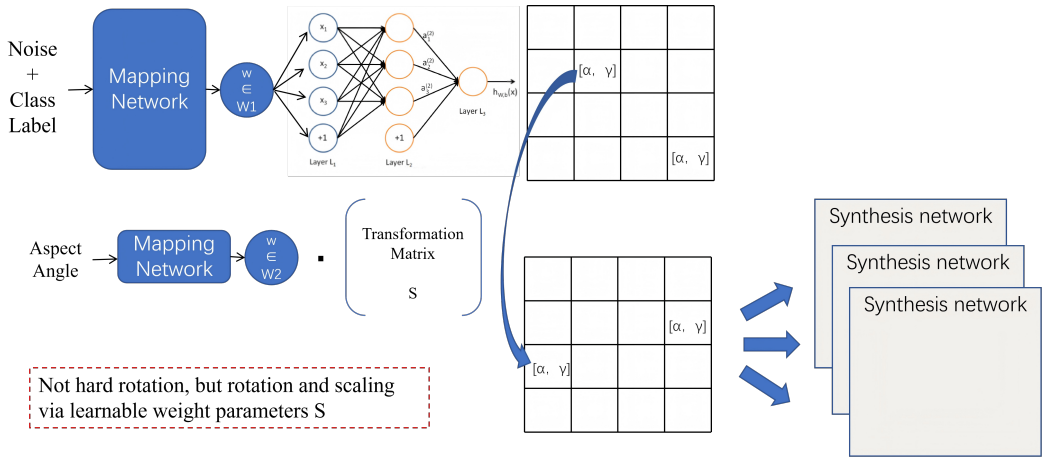


Figure 2. ASC-inspired sparse scattering-structure prior module.

2.3.2. Differentiable Angle-Aligned Evolution

The structural map S_0 itself does not yet encode explicit viewpoint variation. To model how target scattering structures evolve with aspect angle, we apply an angle-aligned differentiable transformation to the sparse representation. Specifically, the viewpoint-dependent transformation is written as

$$A(\theta) = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix} \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \tag{10}$$

where s_x and s_y are learnable scale factors used to capture anisotropic geometric variation in the image plane. The transformed coordinate of each structural component is then given by

$$\tilde{c}_k(\theta) = A(\theta) c_k \tag{11}$$

Based on the transformed coordinates, the angle-aligned structural map S_θ is obtained as

$$S_\theta(p) = \sum_{k=1}^{K} a_k \exp\left(-\frac{\|p - \tilde{c}_k(\theta)\|_2^2}{2\sigma_k^2}\right) \tag{12}$$

Equivalently, the same process can be implemented through differentiable resampling on S_0, which preserves end-to-end trainability. Through this transformation, the aspect angle is no longer treated as an abstract conditional code but is explicitly converted into a structured variation of the sparse scattering layout. As a result, the generator is encouraged to learn viewpoint-dependent structural evolution rather than relying on superficial appearance adaptation.

2.3.3. Multi-Scale Structural Guidance to the Generator

To effectively inject the ASC-inspired structural prior into the synthesis process, the angle-aligned structural map S_θ is further propagated to multiple scales of the generator. Let $\{F^{(l)}\}_{l=1}^{L}$ denote the intermediate feature maps of the generator at different synthesis stages. We first construct a multi-scale pyramid of the structural prior:

$$S_\theta^{(l)} = D_l(S_\theta), \qquad l = 1, 2, \ldots, L \tag{13}$$

where $D_l(\cdot)$ denotes the downsampling operator corresponding to the spatial resolution of the l-th stage. Each scale-specific structural prior is then projected into the feature space by a lightweight transformation $P_l(\cdot)$ and injected into the generator features via

$$\tilde{F}^{(l)} = F^{(l)} + \gamma_l P_l\left(S_\theta^{(l)}\right) \tag{14}$$

where γ_l is a learnable modulation coefficient.

This multi-scale guidance mechanism ensures that the same angle-aware structural prior constrains the generator throughout the hierarchical synthesis process, rather than affecting only a single layer. Consequently, aspect-angle control is introduced not only at the conditional-input level but also through angle-aware feature modulation. From an optimization perspective, the adversarial gradients can be propagated back through the entire ASC-inspired structural branch, so that the predicted sparse structure is progressively adjusted toward viewpoint-consistent scattering layouts. Therefore, the ASC-inspired structural prior acts as a bridge between aspect-angle conditioning and physically plausible structural generation, forming the basis of the proposed physically consistent view-completion framework.

2.4. Target–Background Dual Physical Consistency

The structural prior introduced in Section 2.3 constrains the generator to produce angle-aware scattering layouts, but physically consistent SAR generation also requires explicit supervision on the generated target and background regions. To this end, we introduce a dual physical consistency strategy that treats the target region and the background region separately. The target region is constrained by ASC-aware structural discrimination, while the background region is regularized by statistical consistency of SAR speckle. Let $M_\mathrm{t} \in \{0,1\}^{H\times W}$ denote the binary mask of the target region and $M_\mathrm{b} = 1 - M_\mathrm{t}$ denote the corresponding background mask. Then the masked target and background components of an image x can be written as

$$x_\mathrm{t} = M_\mathrm{t} \odot x, \qquad x_\mathrm{b} = M_\mathrm{b} \odot x \tag{15}$$

where $\odot$ denotes element-wise multiplication. The same notation is used for a generated image $\hat{x}$.

2.4.1. Target-Region Scattering Consistency via ASC-Aware Discrimination

The target region of a SAR image contains the dominant scattering responses of the object of interest. Therefore, physical consistency in the target region should not be judged solely by image appearance but also by whether the scattering-related structure is compatible with real SAR targets. To achieve this, we introduce an ASC-aware target-consistency mechanism into the discriminator.

Let $E_\mathrm{asc}(\cdot)$ denote a pretrained ASC extractor, which maps an input target region to a scattering-aware structural representation:

$$A(x_t) = E_{asc}(x_t) \tag{16}$$

The extractor is kept frozen during adversarial training, so that it serves as a stable physical reference rather than being absorbed into the generator–discriminator game. Based on the pretrained ASC branch, the discriminator is biased toward physically meaningful structural cues at intermediate layers. Let $W^{(l)}$ denote the convolution kernel of the l-th discriminator block and let $W_{asc}^{(l)}$ denote the corresponding ASC-related kernel projected from the pretrained extractor. We adopt a parameter-guided injection strategy:

$$\widetilde{W}^{(l)} = W^{(l)} + \alpha_l \, T_l\left(W_{asc}^{(l)}\right) \tag{17}$$

where $T_l(\cdot)$ is a channel-alignment operator and α_l is a learnable injection coefficient. Accordingly, the discriminator response at the l-th block becomes

$$F_D^{(l)}(x_t) = \sigma\left(\widetilde{W}^{(l)} * F_D^{(l-1)}(x_t) + b^{(l)}\right) \tag{18}$$

where $F_D^{(l-1)}(\cdot)$ denotes the input feature map of the current block, $b^{(l)}$ is the bias term, $*$ denotes convolution, and $\sigma(\cdot)$ is the nonlinear activation.

With the ASC-aware discriminator, the target-region consistency objective is defined as

$$\mathcal{L}_{tgt} = \mathbb{E}_{x\sim p_{data}}[\log D_t(x_t)] + \mathbb{E}_{z,y,\theta}[\log(1 - D_t(\hat{x}_t))] \tag{19}$$

where $D_t(\cdot)$ denotes the target-consistency discriminator branch. This objective encourages the discriminator to penalize generated target regions that may appear realistic in image space but are inconsistent with scattering-aware structural patterns in ASC space.

2.4.2. Background Statistical Consistency via ENL Regularization

Although the target region is the primary focus of SAR generation, unrealistic background statistics can also noticeably degrade the physical credibility of generated images. In target-chip SAR data, background speckle-like fluctuations are closely related to the coherent imaging mechanism and signal-dependent multiplicative speckle in SAR imagery [38]. Therefore, instead of imposing structural constraints on the background, we regularize its statistical properties.

For a masked background region x_b, its mean and variance are computed as

$$\mu(x_b) = \frac{1}{|\Omega_b|} \sum_{p\in\Omega_b} x(p), \tag{20}$$

$$\sigma^2(x_b) = \frac{1}{|\Omega_b|} \sum_{p\in\Omega_b} (x(p) - \mu(x_b))^2 \tag{21}$$

where Ω_b is the set of background pixels. Based on these statistics, the equivalent number of looks (ENL) is defined as

$$\mathrm{ENL}(x_b) = \frac{\mu^2(x_b)}{\sigma^2(x_b) + \varepsilon} \tag{22}$$

where ε is a small constant for numerical stability.

The background statistical consistency loss is then written as

$$\mathcal{L}_{bg} = \mathbb{E}[|\mathrm{ENL}(x_b) - \mathrm{ENL}(\hat{x}_b)|] \tag{23}$$

This regularization explicitly aligns the background speckle statistics of generated SAR images with those of real SAR data, while avoiding direct interference with target-region structural modeling.

Taken together, the target-region structural consistency term and the background-region statistical consistency term form a dual physical consistency mechanism. The former focuses on whether dominant target scattering responses are structurally plausible, while the latter constrains whether the generated background preserves realistic SAR-specific statistics. These two terms complement the ASC-inspired structural prior in Section 2.3 and will be jointly integrated into the overall training objective in Section 2.5.

2.5. Training Objective and Optimization

Based on the framework introduced in Sections 2.3 and 2.4, the proposed method is trained by jointly optimizing the generator-side ASC-inspired structural guidance, the target-region structural consistency, and the background-region statistical consistency. Let G denote the generator, D denote the image-level discriminator, and D_t denote the target-consistency discriminator branch. The pretrained ASC extractor E_{asc} remains frozen during training and is therefore not updated by gradient descent.

2.5.1. Adversarial Objective

At the image level, the generator is trained to synthesize realistic SAR images conditioned on the target label and aspect angle. Let

$$c = (y, \phi(\theta)) \tag{24}$$

denote the joint condition composed of the target label and the aspect-angle encoding. Then the generated image is written as

$$\hat{x} = G(z, c) \tag{25}$$

The image-level adversarial loss for the discriminator is defined as

$$\mathcal{L}_{adv}^{D} = -\mathbb{E}_{(x,c)\sim p_{data}}[\log D(x, c)] - \mathbb{E}_{z\sim p_z, c\sim p_c}[\log(1 - D(\hat{x}, c))] \tag{26}$$

And the corresponding adversarial objective for the generator is

$$\mathcal{L}_{adv}^{G} = -\mathbb{E}_{z\sim p_z, c\sim p_c}[\log D(\hat{x}, c)] \tag{27}$$

Unlike a conventional conditional generator, the proposed generator contains the ASC-inspired structural branch described in Section 2.3. This branch is trained jointly with the generator through back-propagation of the adversarial and structural consistency gradients. Therefore, the predicted sparse structural prior is progressively adjusted to provide angle-aware guidance for dominant scattering-layout generation.

2.5.2. Target-Region Physical Consistency Objective

To further constrain the physical plausibility of the generated target region, the ASC-aware target-consistency discriminator D_t is introduced as described in Section 2.4. The corresponding discriminator loss is

$$\mathcal{L}_{tgt}^{D} = -\mathbb{E}_{x\sim p_{data}}[\log D_t(x_t)] - \mathbb{E}_{z\sim p_z, c\sim p_c}[\log(1 - D_t(\hat{x}_t))] \tag{28}$$

while the generator-side objective is

$$\mathcal{L}_{tgt}^{G} = -\mathbb{E}_{z\sim p_z, c\sim p_c}[\log D_t(\hat{x}_t)] \tag{29}$$

This term encourages the generator to produce target regions whose dominant scattering structures are compatible with the structural patterns encoded by the frozen ASC-aware discriminator.

2.5.3. Background Statistical Consistency Objective

For the background region, the ENL-based consistency term introduced in Section 2.4 is directly applied to the generator:

$$\mathcal{L}_{\mathrm{bg}}^{G} = \mathbb{E}_{(x,c)\sim p_{\mathrm{data}}, z\sim p_z}[|\mathrm{ENL}(x_{\mathrm{b}}) - \mathrm{ENL}(\hat{x}_{\mathrm{b}})|] \tag{30}$$

Since this term regularizes only the background region, it suppresses unrealistic speckle deviations without interfering with target-region structural modeling.

2.5.4. Overall Objective and Optimization Strategy

The final generator objective combines the image-level adversarial loss, the target-region physical consistency loss, and the background statistical consistency loss:

$$\mathcal{L}_G = \mathcal{L}_{\mathrm{adv}}^{G} + \lambda_{\mathrm{t}}\mathcal{L}_{\mathrm{tgt}}^{G} + \lambda_{\mathrm{b}}\mathcal{L}_{\mathrm{bg}}^{G} \tag{31}$$

where λ_{t} and λ_{b} are trade-off coefficients controlling the relative strengths of target-region and background-region physical constraints.

The discriminator objective is defined as

$$\mathcal{L}_D = \mathcal{L}_{\mathrm{adv}}^{D} + \lambda_{\mathrm{t}}\mathcal{L}_{\mathrm{tgt}}^{D} \tag{32}$$

Accordingly, the full optimization problem can be written as

$$\min_{G}\max_{D,D_{\mathrm{t}}} \mathcal{L}_{\mathrm{adv}} + \lambda_{\mathrm{t}}\mathcal{L}_{\mathrm{tgt}} + \lambda_{\mathrm{b}}\mathcal{L}_{\mathrm{bg}} \tag{33}$$

where $\mathcal{L}_{\mathrm{adv}}$, $\mathcal{L}_{\mathrm{tgt}}$, and $\mathcal{L}_{\mathrm{bg}}$ are shorthand notations for the corresponding generator–discriminator game defined above.

During training, the generator and discriminator are updated alternately using the Adam optimizer. The ASC extractor remains fixed throughout optimization, ensuring that the structural cues used for target-region discrimination act as a stable physical reference. Through this joint optimization strategy, the proposed framework learns to synthesize SAR images that are not only visually realistic but also physically consistent in both target scattering structure and background statistics under unseen aspect-angle conditions.

3. Experiments and Results

3.1. Experimental Settings

3.1.1. Dataset Description

The experiments were conducted on a cross-angle SAR target generation dataset constructed from the original MSTAR complex data [39]. To isolate the influence of aspect-angle variation and avoid mixing it with changes in imaging geometry, all experiments were performed under a fixed imaging configuration, including a single frequency band, a single polarization, and a fixed depression angle. Following the standard target-chip setting, we selected five representative ground targets, namely, 2S1, BRDM2, D7, T62, and ZIL131.

To reduce the influence of irrelevant imaging variations, all samples were center-cropped around the target region, resized to 128×128, and linearly normalized to the range $[-1, 1]$. The final dataset contained 72 aspect-angle samples for each category, covering

the full azimuth range from 0° to 355° with an angular interval of 5°. In total, 360 labeled SAR target chips were used for the subsequent unseen-angle generation experiments. A summary of the dataset configuration is provided in Table 1.

Table 1. Dataset summary used in the unseen-angle SAR generation experiments.

Class	Number of Chips	Aspect Range	Angular Interval
2S1	72	0°–355°	5°
BRDM2	72	0°–355°	5°
D7	72	0°–355°	5°
T62	72	0°–355°	5°
ZIL131	72	0°–355°	5°
Total	360	0°–355°	5°

3.1.2. Unseen-Angle Protocols

To evaluate the proposed method under realistic limited-view conditions, we designed two complementary unseen-angle protocols, namely, an interpolation protocol and a hold-out protocol.

Interpolation protocol. In this setting, the training set contained sparsely sampled aspect angles, while the test set consisted of intermediate angles that were not directly observed during training. Specifically, the training angle set was defined as

$$\Theta_{\mathrm{tr}}^{\mathrm{int}} = \{0^\circ, 10^\circ, 20^\circ, \ldots, 350^\circ\} \tag{34}$$

and the unseen-angle test set was defined as

$$\Theta_{\mathrm{te}}^{\mathrm{int}} = \{5^\circ, 15^\circ, 25^\circ, \ldots, 355^\circ\} \tag{35}$$

Under this protocol, each class contained 36 training samples and 36 test samples, resulting in 180 training samples and 180 unseen-angle test samples in total.

Hold-out protocol. In this setting, a continuous aspect-angle sector was entirely excluded from training and used only for testing to evaluate the extrapolation capability of the model in a strictly unseen angular range. The hold-out sector was chosen as

$$\Theta_{\mathrm{te}}^{\mathrm{hold}} = \{90^\circ, 95^\circ, 100^\circ, \ldots, 145^\circ\} \tag{36}$$

and the corresponding training angle set was

$$\Theta_{\mathrm{tr}}^{\mathrm{hold}} = [0^\circ, 355^\circ] \setminus \Theta_{\mathrm{te}}^{\mathrm{hold}} \tag{37}$$

This protocol yielded 300 training samples and 60 hold-out test samples in total. Compared with interpolation, the hold-out protocol was more challenging because it required the generator to synthesize physically plausible SAR images in an entirely unobserved angular sector. The detailed split statistics are summarized in Table 2.

Table 2. Summary of unseen-angle evaluation protocols.

Protocol	Training Angles	Test Angles	Samples (Train/Test)
Interpolation	0°, 10°, . . . , 350°	5°, 15°, . . . , 355°	180/180
Hold-out	all except 90°–145°	90°–145°	300/60

3.1.3. Implementation Details

All experiments were implemented in PyTorch 1.7.1 and conducted on two NVIDIA GeForce RTX 2080Ti GPUs, each with 11 GB of memory. The official StyleGAN2 implementation was adopted for baseline and as the backbone of the proposed generator [19]. Since limited training data can easily lead to discriminator overfitting in GAN training, limited-data GAN training studies also provide an important methodological reference for this setting [40]. All compared methods were trained under the same data split, image resolution, and training budget for a fair comparison. Unless otherwise stated, the training budget was set to 1500 kimg and the batch size was fixed to 32. The generator and discriminator were optimized using Adam with an initial learning rate of 2×10^{-3}, $\beta_1 = 0.0$, and $\beta_2 = 0.99$.

For the proposed method, the number of sparse dominant response components in the ASC-inspired structural branch were set to $K = 16$, and the structural prior was injected into four synthesis stages of the generator in a multi-scale manner. The pretrained ASC extractor used in the target-consistency discriminator branch was frozen during adversarial training. The trade-off coefficients in Equation (31) were empirically set to $\lambda_t = 1.0$ and $\lambda_b = 0.1$, and the numerical stability constant in the ENL computation was fixed to $\varepsilon = 10^{-6}$. All hyperparameters were selected using a small validation subset within the training split and then kept unchanged across all comparison experiments.

For the target–background consistency constraints, the target and background masks were generated using a fixed center-square strategy. Since the MSTAR target chips were cropped and centered during preprocessing, the target-dominated region was generally located near the image center. For an input image of size $H \times W$, the target mask M_t was defined as a square region centered at the image center, with side lengths of $0.4H$ and $0.4W$, respectively. The background mask was obtained as the complement of the target mask, i.e., $M_b = 1 - M_t$. The same fixed masks were used for all real and generated images. We adopted this simple strategy because different target categories may have different spatial extents and scattering distributions, while adaptive thresholding or connected-component-based segmentation can be affected by speckle fluctuations, fragmented scattering responses, and strong local scatterers. The fixed central mask therefore provides a reproducible and consistent target/background partition while avoiding additional instability introduced by mask estimation.

3.1.4. Evaluation Metrics

To comprehensively evaluate the proposed method from both image-generation and physical-consistency perspectives, we employed four groups of metrics, including image fidelity, target-region physical consistency, background statistical consistency, and angle-condition consistency. The corresponding metrics are summarized in Table 3.

Table 3. Evaluation metrics used in the experiments. Upward/downward arrows indicate that higher/lower values are better, respectively.

Metric Group	Metric	Desired Trend
Image fidelity	FID	↓
Image fidelity	SSIM	↑
Target-region physical consistency	AFS	↑
Background statistical consistency	ΔENL	↓
Background statistical consistency	BVE	↓
Angle-condition consistency	CMAE	↓

Image fidelity. We used the Fréchet Inception Distance (FID) [41] and the Structural Similarity Index Measure (SSIM) [42] to evaluate the overall similarity between generated

SAR images and real SAR images. The FID measured the distance between the feature distributions of real and generated images:

$$\mathrm{FID} = \|\mu_r - \mu_g\|_2^2 + \mathrm{Tr}\left(\Sigma_r + \Sigma_g - 2\left(\Sigma_r\Sigma_g\right)^{1/2}\right) \tag{38}$$

where μ_r and Σ_r denote the mean vector and covariance matrix of the real-image features, respectively, and μ_g and Σ_g denote those of the generated-image features. Lower FID indicated that the generated images were closer to the real-image distribution in the feature space. The SSIM evaluated structural similarity between a generated image $\hat{x}$ and its corresponding real image x:

$$\mathrm{SSIM}(x, \hat{x}) = \frac{(2\mu_x\mu_{\hat{x}} + C_1)(2\sigma_{x\hat{x}} + C_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + C_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + C_2)} \tag{39}$$

where μ_x and $\mu_{\hat{x}}$ are the mean intensities of x and $\hat{x}$, σ_x^2 and $\sigma_{\hat{x}}^2$ are the corresponding variances, $\sigma_{x\hat{x}}$ is the covariance between them, and C_1 and C_2 are small constants for numerical stability. Higher SSIM indicated better structural similarity.

Target-region physical consistency. To quantify the structural agreement between the generated target region and the real target region in an ASC-aware feature space, we defined the ASC-aware Feature Similarity (AFS) as

$$\mathrm{AFS} = \frac{1}{N}\sum_{i=1}^{N}\frac{\langle E_{\mathrm{asc}}(x_{i,\mathrm{t}}),\, E_{\mathrm{asc}}(\hat{x}_{i,\mathrm{t}})\rangle}{\|E_{\mathrm{asc}}(x_{i,\mathrm{t}})\|_2\|E_{\mathrm{asc}}(\hat{x}_{i,\mathrm{t}})\|_2} \tag{40}$$

where $E_{\mathrm{asc}}(\cdot)$ denotes the pretrained and frozen ASC extractor; $x_{i,\mathrm{t}}$ and $\hat{x}_{i,\mathrm{t}}$ denote the target-region patches of the i-th real and generated SAR images, respectively; and N is the number of evaluated samples. A higher AFS indicated better agreement in dominant target scattering structures.

Background statistical consistency. Since the background of target-chip SAR images was mainly characterized by speckle-like statistics, we used two complementary metrics to evaluate background realism [43]. The first metric was the mean absolute ENL difference:

$$\Delta\mathrm{ENL} = \frac{1}{N}\sum_{i=1}^{N}\left|\mathrm{ENL}(x_{i,\mathrm{b}}) - \mathrm{ENL}(\hat{x}_{i,\mathrm{b}})\right| \tag{41}$$

where $x_{i,\mathrm{b}}$ and $\hat{x}_{i,\mathrm{b}}$ denote the background regions of the i-th real and generated images, respectively. The second metric was the Background Variance Error (BVE), which followed the variance-based statistical assessment used in SAR image quality evaluation and further restricted it to the background region [44],

$$\mathrm{BVE} = \frac{1}{N}\sum_{i=1}^{N}\left|\sigma^2(x_{i,\mathrm{b}}) - \sigma^2(\hat{x}_{i,\mathrm{b}})\right| \tag{42}$$

where $\sigma^2(\cdot)$ denotes the variance of the corresponding background region. Lower ΔENL and lower BVE indicated better consistency with the statistical characteristics of real SAR backgrounds.

Angle-conditioned consistency. To evaluate whether the generated image was consistent with the required aspect angle, we further computed the Circular Mean Absolute Error (CMAE) [45] using an independently pretrained aspect-angle estimator $R_\theta(\cdot)$. The estimator was trained only on real SAR images and was kept frozen during evaluation. Specifically, $R_\theta(\cdot)$ was implemented as a lightweight CNN that took a single-channel SAR image resized to 128×128 as input. It consisted of four convolutional blocks, each containing a

3×3 convolution, batch normalization, ReLU activation, and max pooling, with channel numbers of 16, 32, 64, and 128, respectively. A global average pooling layer and two fully connected layers were then used to produce a two-dimensional output corresponding to the sine and cosine components of the aspect angle:

$$\mathcal{L}_\theta = \frac{1}{M}\sum_{j=1}^{M}\left[\left(\hat{s}_j - \sin\theta_j\right)^2 + \left(\hat{c}_j - \cos\theta_j\right)^2\right] \tag{43}$$

where M denotes the number of training samples in a mini-batch and $(\hat{s}_j, \hat{c}_j) = R_\theta(x_j)$ denotes the predicted sine and cosine components of the aspect angle. No generated images were used to train or fine-tune $R_\theta(\cdot)$.

$$\mathrm{CMAE} = \frac{1}{N}\sum_{i=1}^{N}\min\left(\left|\hat{\theta}_i - \theta_i\right|, 360^\circ - \left|\hat{\theta}_i - \theta_i\right|\right) \tag{44}$$

where $\hat{x}_i$ is the generated image conditioned on the target aspect angle θ_i and $\hat{\theta}_i$ is the aspect angle estimated from $\hat{x}_i$. The circular distance was used to avoid artificial discontinuity at the $0^\circ/360^\circ$ boundary. Lower CMAE indicated better angle-conditioned consistency.

These settings ensured that the proposed method was evaluated not only in terms of visual image quality but also in terms of target-region structural plausibility, background statistical realism, and aspect-angle controllability under strictly defined unseen-angle scenarios.

3.2. Comparison Experiment

In this section, we describe comparative experiments using different generative models, including two representative GAN baselines in SAR generation (DCGAN [17] and ACGAN [18]), a strong conditional generative backbone (StyleGAN2 [19]), a recent physics-aware SAR generation method (Φ-GAN [27]), and the proposed method. All models were trained and evaluated under the same data splits, training budget, and evaluation metrics for a fair comparison.

3.2.1. Quantitative Comparison Under the Interpolation Protocol

We first evaluated the proposed method under the interpolation protocol, where the model was trained on sparsely sampled aspect angles and tested on intermediate unseen angles. This setting examined whether the generator could synthesize physically plausible SAR images when the required aspect angle lay between observed training views. Table 4 reports the quantitative comparison with representative baseline methods.

Table 4. Quantitative comparison under the interpolation protocol. Bold font indicates the best value.

Method	FID ↓	SSIM ↑	AFS ↑	ΔENL ↓	BVE ↓	CMAE ↓
DCGAN	284.73	0.214	0.521	1.87	0.0215	24.8
ACGAN	231.42	0.219	0.548	1.63	0.0189	21.4
StyleGAN2	165.95	0.223	0.612	1.21	0.0142	16.7
Φ-GAN	84.96	0.356	0.742	0.58	0.0087	9.4
Ours	**72.51**	**0.361**	**0.801**	**0.43**	**0.0081**	**6.8**

As shown in Table 4, conventional adversarial baselines such as DCGAN and ACGAN exhibited limited performance under the interpolation protocol. Although these models could generate coarse target appearances, they showed relatively large distribution mismatch, weak structural consistency in the ASC-aware feature space, and poor aspect-angle controllability. This indicates that directly treating unseen-angle SAR generation as a

generic conditional synthesis task is insufficient when the angular coverage of the training data is sparse.

Compared with these early baselines, StyleGAN2 achieved clear improvements in image fidelity and angle-condition consistency, suggesting that stronger generative backbones and adaptive data augmentation are beneficial for cross-angle SAR image synthesis. However, its performance in AFS, ΔENL, and BVE remained clearly inferior to that of the physics-aware methods, implying that appearance-level quality improvement alone does not guarantee physically plausible target scattering structures or realistic SAR background statistics.

The comparison between Φ-GAN and the proposed method further highlights the advantage of the proposed design. Although Φ-GAN already introduced physics-inspired constraints and achieved competitive FID and SSIM values, our method consistently performed better across all metrics. In particular, the proposed framework improved AFS from 0.742 to 0.801, reduced ΔENL from 0.58 to 0.43, and decreased CMAE from 9.4° to 6.8°. These improvements indicate that introducing an ASC-inspired sparse scattering-structure prior and jointly constraining target-region structure and background-region statistics are effective for interpolation-based unseen-angle SAR generation.

Overall, the interpolation results show that the proposed method achieves favorable FID and SSIM scores, while more importantly improving target-region scattering consistency, background statistical realism, and aspect-angle controllability. These results provide quantitative evidence that the proposed physically consistent framework is more suitable for unseen-angle SAR view completion than generative models with insufficient physical constraints.

3.2.2. Quantitative Comparison Under the Hold-Out Protocol

We next evaluated the proposed method under the hold-out protocol, where a continuous aspect-angle sector was completely excluded from training and used only for testing. Compared with the interpolation protocol, this setting was substantially more challenging because the generator needed to synthesize SAR images in an entirely unobserved angular range, rather than merely interpolating between neighboring training views. Therefore, the hold-out protocol provided a more stringent test of whether the model learned physically meaningful aspect-dependent structural evolution. The quantitative comparison is reported in Table 5.

Table 5. Quantitative comparison under the hold-out protocol. Bold font indicates the best value.

Method	FID ↓	SSIM ↑	AFS ↑	ΔENL ↓	BVE ↓	CMAE ↓
DCGAN	318.45	0.156	0.487	2.03	0.0241	31.6
ACGAN	276.18	0.193	0.513	1.79	0.0214	27.9
StyleGAN2	186.72	0.224	0.576	1.36	0.0165	21.8
Φ-GAN	143.57	0.261	0.631	1.08	0.0137	17.4
Ours	**132.93**	**0.308**	**0.701**	**0.81**	**0.0102**	**15.6**

As expected, all methods exhibited performance degradation under the hold-out protocol compared with the interpolation protocol. This trend confirms that unseen-angle SAR generation becomes considerably more difficult when the test angles lie in a completely unobserved angular sector. In particular, the performance drop of DCGAN, ACGAN, and the StyleGAN2-based baselines was especially pronounced in FID, AFS, and CMAE, indicating that purely appearance-driven or weakly conditioned generative models struggle to maintain structurally coherent target scattering responses when required to extrapolate beyond the observed angular support.

The comparison also reveals that physics-aware methods are generally more robust than purely data-driven baselines in this challenging setting. Φ-GAN still outperformed the standard GAN baselines, suggesting that the introduction of physics-inspired constraints alleviates, to some extent, the structural inconsistency caused by sparse angular observations. However, its performance remained clearly inferior to that of the proposed method across all metrics. These results indicate that the proposed framework is able to preserve not only image-level realism but also target-region structural plausibility and background statistical consistency in a completely unseen angular range.

More importantly, the improvement under the hold-out protocol is more indicative than that under interpolation. While interpolation performance may partly benefit from smooth transitions between neighboring training angles, successful generation under a hold-out angular sector requires the model to infer unobserved structural evolution patterns from limited angular evidence. The superior performance of the proposed method therefore suggests that the ASC-inspired sparse scattering-structure prior provides a stronger structural inductive bias for aspect-dependent generation than conventional conditional angle encoding alone. At the same time, the background statistical consistency constraint helps suppress the accumulation of unrealistic speckle artifacts, which becomes increasingly important when generation is performed far from the observed training angles.

Overall, the hold-out results provide stronger evidence for the effectiveness of the proposed method in physically consistent unseen-angle SAR generation. The proposed framework generalizes better to unobserved angular sectors than representative baselines, demonstrating that jointly modeling scattering-structure evolution and target–background physical consistency is critical for reliable SAR view completion under sparse angular observations.

3.3. Physical Consistency Evaluation

Although the quantitative comparisons in Section 3.2 already show the superiority of the proposed method in image fidelity and unseen-angle generalization, these metrics alone are insufficient to fully explain why the generated SAR images are more physically plausible. Since the proposed framework introduces an ASC-inspired sparse scattering-structure prior and enforces target–background dual physical consistency, it was necessary to further evaluate the generated results from the perspectives of target-region structural plausibility and background-region statistical realism.

3.3.1. Target-Region Scattering Consistency

We first analyzed the physical consistency of the generated target region by focusing on whether the dominant scattering structures remained compatible with those of real SAR targets. For this purpose, we report the AFS under the hold-out protocol in Table 6. The comparison includes a strong appearance-driven baseline StyleGAN2, a representative physics-aware baseline Φ-GAN, and the proposed method.

Table 6. Target-region scattering consistency under the hold-out protocol. Bold font indicates the best value.

Class	StyleGAN2	Φ-GAN	Ours
2S1	0.595	0.644	**0.709**
BRDM2	0.565	0.618	**0.693**
D7	0.583	0.629	**0.701**
T62	0.541	0.636	**0.712**
ZIL131	0.592	0.628	**0.688**
Average	0.576	0.631	**0.701**

As shown in Table 6, the proposed method consistently achieved the highest AFS values across all target categories. The improvement over StyleGAN2 indicates that appearance-level optimization alone is insufficient for preserving target-region structural plausibility in unseen-angle generation. Compared with Φ-GAN, the proposed method still yielded clear gains, suggesting that the ASC-inspired sparse scattering-structure prior provides a stronger structural guidance signal than simply introducing generic physics-inspired constraints.

To provide a more intuitive physical interpretation, we further visualized the target-region scattering consistency using scattering-center residual heatmaps. It should be clarified that the extracted points were not regarded as complete electromagnetic ASC parameters but as image-domain dominant scattering-response candidates obtained from SAR amplitude images. This design was motivated by the fact that strong target scattering mechanisms are usually manifested as localized high-intensity responses in SAR amplitude images. Therefore, the spatial distribution and relative strength of local amplitude peaks can provide a useful approximation for analyzing dominant target scattering layouts.

Specifically, for each real or generated target patch, we first applied mild denoising to suppress isolated speckle fluctuations and then detected local maxima within the target region. Non-maximum suppression was used to avoid selecting multiple neighboring pixels from the same bright response. The remaining candidates were ranked by their amplitude values, and the top-K responses were retained as dominant scattering-center candidates for visualization.

Let

$$\mathcal{S} = \{(u_k, v_k, a_k)\}_{k=1}^{K} \tag{45}$$

denote the extracted scattering-center candidates, where (u_k, v_k) and a_k represent the spatial location and amplitude of the k-th dominant scattering center, respectively. Based on these candidates, the corresponding scattering-intensity map was reconstructed through Gaussian rasterization:

$$I^{\mathrm{SC}}(p) = \sum_{k=1}^{K} a_k \exp\left(-\frac{\|p - c_k\|_2^2}{2\sigma^2}\right), \qquad c_k = (u_k, v_k) \tag{46}$$

where p denotes a pixel coordinate and σ controls the spatial spread of each dominant scattering response in the image plane.

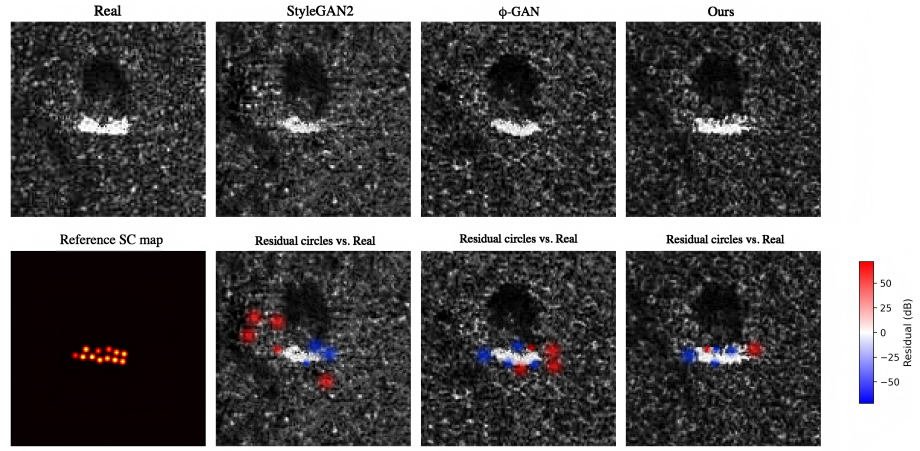
This procedure provided a compact image-domain visualization of dominant target scattering responses and was used only for structural consistency analysis.

Based on the reconstructed scattering-intensity maps of the generated and real target patches, we defined the dB-domain residual heatmap as

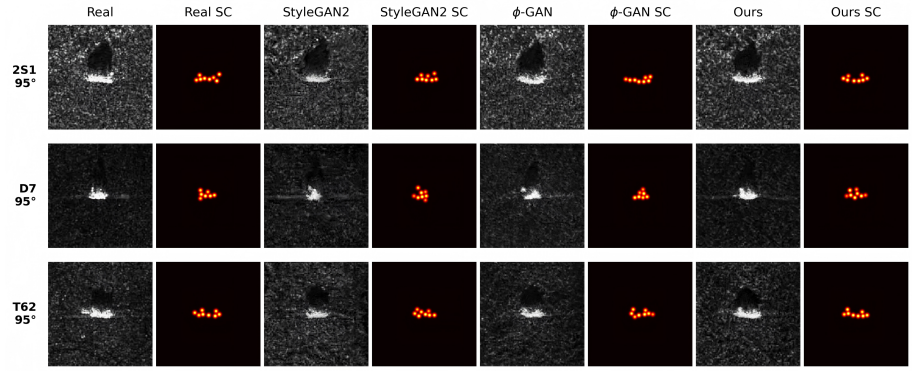
$$R(p) = 10\log_{10}\left(I_{\mathrm{gen}}^{\mathrm{SC}}(p) + \epsilon\right) - 10\log_{10}\left(I_{\mathrm{real}}^{\mathrm{SC}}(p) + \epsilon\right) \tag{47}$$

where $I_{\mathrm{gen}}^{\mathrm{SC}}$ and $I_{\mathrm{real}}^{\mathrm{SC}}$ denote the scattering-intensity maps reconstructed from the generated and real target patches, respectively, and ϵ is a small constant for numerical stability. In the visualization, red indicates overestimated scattering intensity, blue indicates underestimated scattering intensity, and white denotes high consistency with the real target.

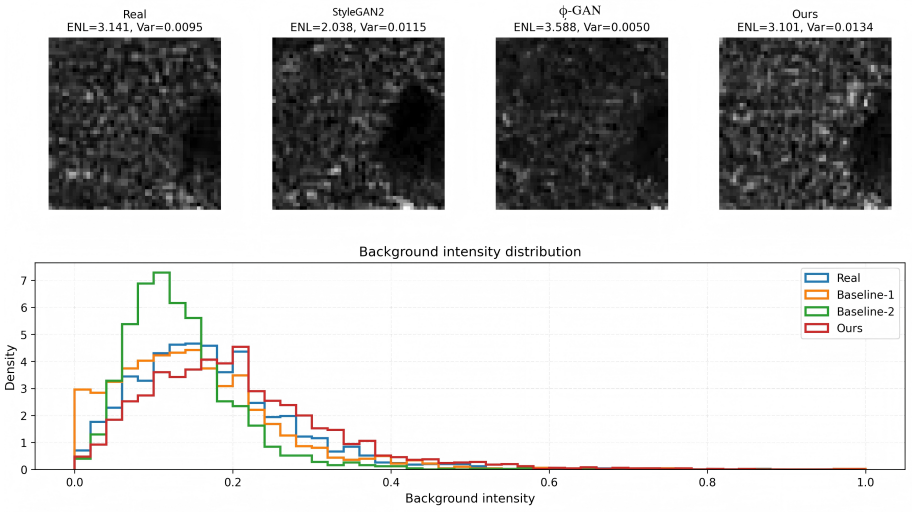
Figure 3a presents representative target-region patches together with their dominant scattering-center residual heatmaps. Compared with the baseline methods, the proposed method exhibited a more homogeneous near-white residual distribution, indicating smaller spatial discrepancies in scattering intensity with respect to the real target. In contrast, the baseline methods tended to produce more clustered red and blue regions, which suggests local overestimation or underestimation of dominant scattering responses. This observation is consistent with the improved AFS values in Table 6 and further indicates that the proposed method better preserves the physically plausible layout of dominant scattering structures under unseen-angle conditions.



(**a**) Target-region scattering consistency



(**b**) Comparison of sparse scattering structures



(**c**) Background statistical results

Figure 3. Physical consistency visualization of generated SAR images. (**a**) Target-region scattering consistency. The first row displays real and generated target patches. The second row shows the reference scattering-center map and residual heatmaps (dB domain), with red/blue indicating over/underestimation and white indicating high consistency. (**b**) Comparison of sparse scattering structures. Comparison of generated SAR images and their extracted sparse scattering structures at representative unseen azimuths under the hold-out protocol. (**c**) Background statistical results. Representative background patches of the real and generated SAR images are shown together with the corresponding background-intensity histograms for visual comparison.

To further provide a direct comparison of sparse scattering structures at unseen azimuths, Figure 3b shows representative results under the hold-out protocol. For each selected target class, one SAR image at an unseen azimuth from the retained angular sector was selected, together with its extracted sparse scattering structure. For comparison, the corresponding SAR images generated by StyleGAN2, Φ-GAN, and the proposed method, as well as their extracted sparse scattering structures, are also presented. It can be observed that the sparse scattering structures extracted from the images generated by the proposed method were generally more consistent with those extracted from the measured SAR images in terms of the spatial layout and relative intensity distribution of dominant responses. In contrast, the baseline methods more often exhibited missing responses, misplaced bright components, or structurally inconsistent scattering layouts. These results provide more direct evidence that the proposed scattering-structure guidance is effective in preserving target-region structural consistency at unseen azimuths under the hold-out setting.

3.3.2. Background Statistical Consistency

We next analyzed the realism of the generated SAR background. Since the background of target-chip SAR images is dominated by speckle-like fluctuations, its realism is primarily reflected by statistical compatibility rather than semantic structure. Table 7 reports the background consistency results of the three representative methods under both interpolation and hold-out protocols, using ΔENL and BVE as evaluation metrics.

Table 7. Background statistical consistency under different unseen-angle protocols. Bold font indicates the best value.

Protocol	Method	ΔENL	BVE
Interpolation	StyleGAN2	1.21	0.0142
	Φ-GAN	0.58	0.0087
	Ours	**0.43**	**0.0081**
Hold-out	StyleGAN2	1.36	0.0165
	Φ-GAN	1.08	0.0137
	Ours	**0.81**	**0.0102**

Several observations can be drawn from Table 7. The background statistics of all methods degraded under the hold-out protocol, as expected, because the generator was required to synthesize SAR images in a fully unseen angular range. Moreover, the proposed method consistently produced the smallest ΔENL and BVE under both protocols, indicating that the ENL-based background regularization effectively suppressed unrealistic speckle deviations.

This result can be further understood from the nature of the background constraint. Unlike the target region, where the dominant scattering responses were strongly aspect-dependent, the background region in the adopted target-chip setting was mainly characterized by low-order speckle-like statistics. Since all samples were acquired under the same frequency band, polarization, depression angle, and preprocessing protocol, the background statistics learned from the observed azimuths could still provide useful regularization for unseen azimuths. In this sense, the relatively low ΔENL and BVE values under the hold-out protocol indicate statistical generalization of the learned background distribution, rather than reconstruction of the specific background texture at a retained azimuth. During training, samples from the retained angular sector were excluded as a whole; their background regions were involved only when the final post-training evaluation metrics were computed.

The visual comparison in Figure 3c provides further evidence for the above quantitative results. Compared with StyleGAN2 and Φ-GAN, the proposed method produced background patches whose speckle granularity, local intensity fluctuation, and overall texture appearance were visually closer to those of real SAR images. Meanwhile, the corresponding background-intensity histograms show that the proposed method better matched the distribution of real background responses, with smaller deviation in both concentration and spread. These observations are consistent with the lower ΔENL and BVE values reported in Table 7, and further confirm that the target–background dual consistency design improves not only target-region structural plausibility but also the statistical realism of SAR-specific backgrounds.

Overall, the results in this subsection demonstrate that the proposed framework achieves physical consistency from two complementary perspectives. In the target region, it better preserves aspect-dependent scattering-structure evolution; in the background region, it more faithfully reproduces the speckle-like statistical characteristics of real SAR imagery. This explains why the proposed method consistently outperformed representative baselines under both interpolation and hold-out settings.

3.4. Qualitative Results and Visualization

In addition to the quantitative results, qualitative visualization was essential for understanding how the proposed framework improved unseen-angle SAR generation. Since the core objective of this work was not only to enhance image fidelity but also to preserve structurally plausible dominant scattering-response variation and realistic SAR-specific background statistics, we further analyzed the generated results from three complementary perspectives: representative unseen-angle samples, continuous angle-sequence visualization, and local target–background comparisons. For visual comparison, we adopted ACGAN and StyleGAN2 as baseline methods for comparison.

To avoid ambiguity, we explicitly distinguish the visualization protocols used in Figures 4 and 5. Figure 4 presents representative interpolation-based unseen-angle samples selected from Θ_{te}^{int}, whereas Figure 5 presents continuous angle-sequence visualization across the hold-out sector, including both angles inside Θ_{te}^{hold} and neighboring angular contexts outside the held-out range.

3.4.1. Representative Unseen-Angle Samples

Figure 4 presents representative interpolation-based unseen-angle generation results for the five target categories, including 2S1, BRDM2, D7, T62, and ZIL131. The displayed samples were selected from the interpolation test set Θ_{te}^{int} and corresponded to representative unseen aspect angles of 355°, 65°, 145°, 285°, and 215°, respectively. These aspect angles were not observed during training but lay between neighboring training angles. For each category, the generated SAR images were compared with the corresponding real SAR images at the same target aspect angle, and the same category–angle pairs were aligned across all methods for a fair visual comparison.

Several observations can be made from Figure 4. First, the proposed method preserved the global target shape and dominant bright responses more faithfully than the baseline methods. In contrast, the appearance-driven baselines tended to produce either overly smooth structures or spatially unstable bright regions, especially at challenging unseen angles. Second, the proposed method exhibited better consistency between the generated target appearance and the required aspect angle, whereas the baseline methods sometimes produced targets with ambiguous orientation or incomplete structural transition. Third, the generated background of the proposed method was visually closer to real SAR im-

agery, with more natural speckle granularity and fewer unrealistic artifacts around the target boundary.

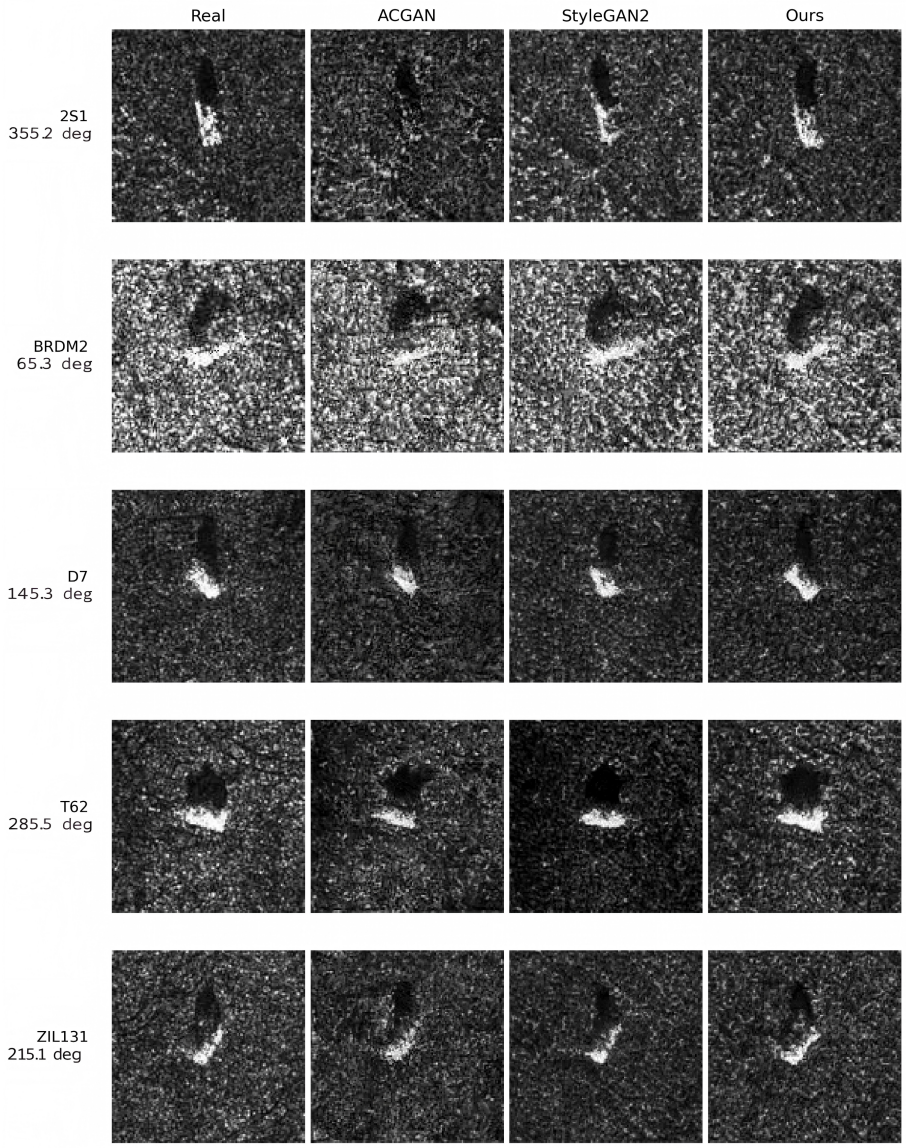


Figure 4. Representative interpolation-based unseen-angle samples for MSTAR. The visualized aspect angles were selected from the interpolation test set $\Theta_{\text{te}}^{\text{int}}$, including 355°, 65°, 145°, 285°, and 215°, which were unseen during training and lay between observed training angles.

These visual observations are consistent with the quantitative improvements reported in Sections 3.2 and 3.3. In particular, the better target-region structural plausibility and background statistical realism of the proposed method could be directly observed, which supports the argument that physically consistent unseen-angle SAR generation requires more than appearance matching alone.

It should also be noted that the proposed method did not completely recover all local scattering details within the target region. Some fine-scale bright responses still deviated from the corresponding real SAR images, especially under challenging unseen-angle conditions. Therefore, the qualitative results should be interpreted as an improvement in the overall layout and plausibility of dominant scattering responses, rather than as exact point-wise reconstruction of the target scattering structure.

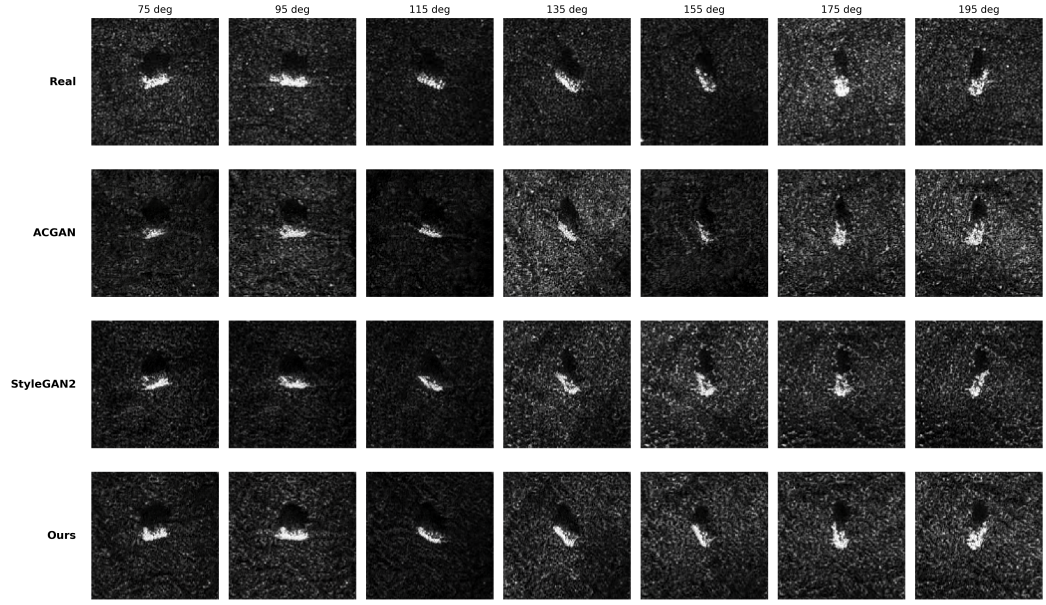


Figure 5. Continuous angle-sequence visualization across the hold-out sector for MSTAR. The sequence spanned 75°–195° with a 20° interval, covering the held-out angular sector 90°–145° and its neighboring angle ranges.

3.4.2. Continuous Angle-Sequence Visualization

Figure 5 further presents continuous angle-sequence visualization across the hold-out sector. The displayed aspect angles spanned 75°–195° with a 20° interval, namely, 75°, 95°, 115°, 135°, 155°, 175°, and 195°. Among them, 95°, 115°, and 135° fell inside the hold-out sector $\Theta_{\mathrm{te}}^{\mathrm{hold}} = \{90°, 95°, 100°, \ldots, 145°\}$, while 75°, 155°, 175°, and 195° provided neighboring angular contexts outside the held-out range. This visualization was used to complement the hold-out quantitative results by showing how the generated target scattering responses varied when the viewpoint entered, passed through, and left the held-out angular sector.

It should be noted that this visualization is not intended to demonstrate strictly smooth trajectories of identical scattering centers. In real SAR imagery, dominant scattering responses may appear, disappear, merge, split, or shift abruptly due to aspect-dependent visibility, occlusion, and electromagnetic scattering mechanisms. Therefore, this experiment focused on whether the generated SAR images exhibited plausible aspect-dependent scattering-response variations, including local response migration, intensity fluctuation, emergence, and disappearance.

As shown in Figure 5, the proposed method generated angle sequences with noticeably better structural continuity than the baseline methods. The dominant bright responses exhibited angle-dependent translation, merging, weakening, or disappearance as the aspect angle changed, which is consistent with the expected behavior of view-dependent scattering structures. In comparison, the baseline methods often exhibited abrupt changes in local bright responses, inconsistent structural deformation, or repeated texture patterns across neighboring angles. Such artifacts indicate that these methods may rely more on appearance interpolation than on physically meaningful structural evolution.

The continuous angle-sequence visualization further indicates that the proposed method encourages more plausible aspect-dependent variations of dominant scattering responses. However, this does not mean that each local scattering component was exactly recovered or continuously tracked across angles. As shown in Figure 5, local bright responses could still differ from the real target scattering details. The main advantage of the proposed method lies in producing more stable and angle-consistent dominant

scattering layouts, while fine-grained local scattering reconstruction remains challenging under sparse angular observations.

The continuous angle-sequence visualization provides qualitative evidence for the effectiveness of the ASC-inspired sparse structural prior. By introducing an angle-aligned structural prior, the proposed framework better captures the variation trend of dominant target scattering responses, rather than only producing locally plausible images at isolated angles.

3.4.3. Local Target–Background Comparisons

To more clearly demonstrate the effect of the proposed target–background dual consistency design, Figure 6 further presents magnified local comparisons of target regions and background regions. For each method, a representative generated image was selected and compared with the corresponding real SAR image using aligned local patches.

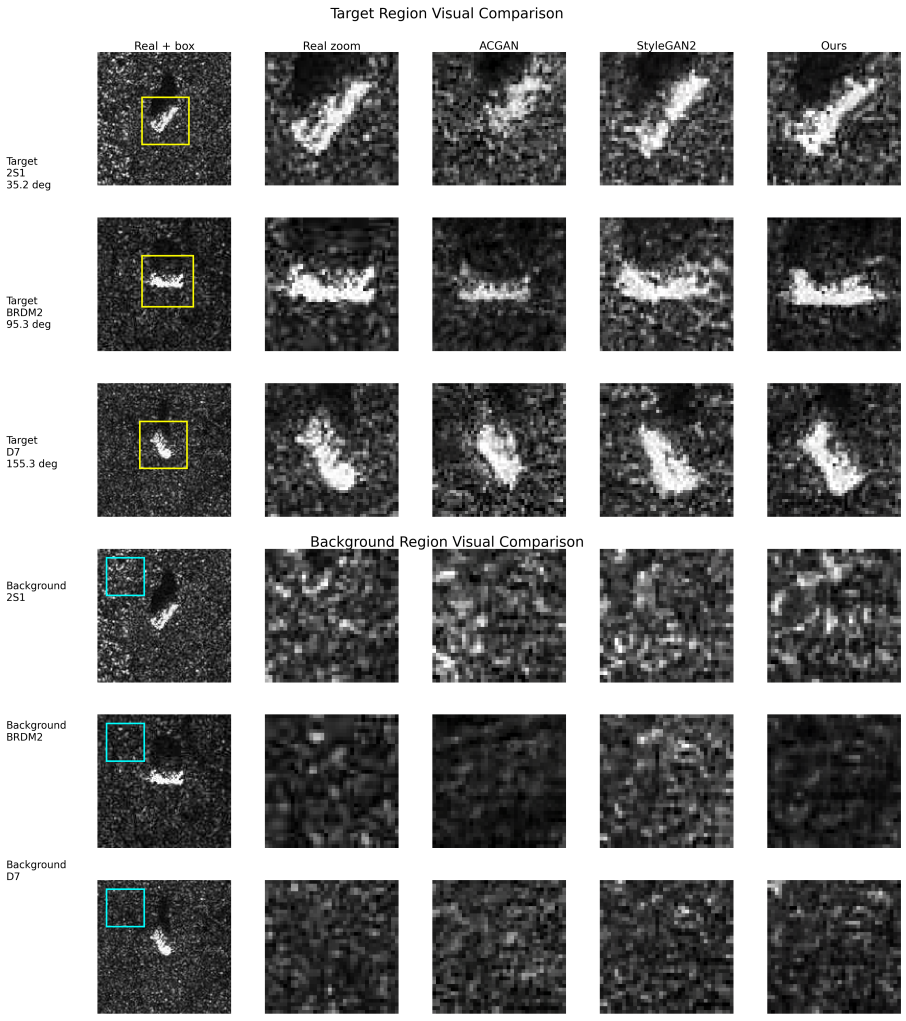


Figure 6. Comparison of target and background local patches.

For the target region, the local comparisons show that the proposed method produced dominant bright responses with sharper localization and more coherent spatial arrangement. Compared with the baselines, the generated target patches were less likely to contain blurred scattering blobs, misplaced bright structures, or locally broken outlines. This observation is consistent with the improved AFS values in Section 3.4.1 and suggests that the ASC-guided structural prior effectively regularizes the evolution of dominant scattering patterns.

For the background region, the enlarged patches show that the proposed method better reproduced the fine-grained speckle characteristics of real SAR backgrounds. In particular, the intensity fluctuation, granularity, and local statistical variation were more realistic than those generated by baselines. Therefore, the qualitative comparisons further confirm that the proposed framework improves both target-region structural plausibility and background-region statistical realism.

Overall, the visual results described in this subsection complement the quantitative evaluations by showing that the proposed method not only improves the numerical metrics but also generates SAR images with more coherent aspect-dependent structural evolution and more realistic SAR-specific background characteristics.

3.5. Ablation Studies

To further investigate the contribution of each component in the proposed framework, we conducted a systematic ablation study under the interpolation protocol. The ablation started from the conditional StyleGAN2 baseline and progressively introduced three key components: ASC-inspired structural guidance in the generator, target-region physical consistency in the discriminator, and background statistical consistency regularization. For clarity, these three components are denoted as E, T, and B, respectively.

The results in Table 8 reveal several important observations. First, the ASC-inspired structural guidance (E) contributed the most significant improvement over the baseline. Once this module was introduced, the FID dropped from 165.95 to 83.20, while the AFS and CMAE were simultaneously improved by a large margin. This indicates that explicitly converting aspect-angle information into an evolving sparse structural prior is the most critical factor for unseen-angle SAR generation. In other words, the generator benefits substantially when the angle condition is transformed from a low-dimensional code into a physically meaningful structural guidance signal.

Table 8. Ablation study of the proposed method under the interpolation protocol. Bold font indicates the best value.

Method	*E*	*T*	*B*	FID ↓	SSIM ↑	AFS ↑	ΔENL ↓	BVE ↓	CMAE ↓
Baseline				165.95	0.223	0.612	1.21	0.0142	16.7
Baseline + *E*	✓			83.20	0.344	0.754	0.63	0.0109	8.2
Baseline + *T*		✓		138.64	0.271	0.658	1.02	0.0121	14.5
Baseline + *B*			✓	109.43	0.301	0.631	0.49	0.0114	13.8
Baseline + *E* + *T*	✓	✓		83.09	0.351	0.772	0.55	0.0103	7.9
Baseline + *E* + *B*	✓		✓	80.65	**0.363**	0.763	**0.42**	0.0086	7.3
Baseline + *T* + *B*		✓	✓	108.95	0.309	0.671	0.46	0.0097	12.9
Ours (*E* + *T* + *B*)	✓	✓	✓	**72.51**	0.361	**0.801**	0.43	**0.0081**	**6.8**

Second, the target-region physical consistency term (T) also improved the results, but its individual contribution was more moderate than that of the ASC-inspired structural guidance. When applied alone, T yielded noticeable gains in FID, SSIM, and AFS, showing that ASC-inspired structural discrimination helped the model reject target regions with unrealistic scattering-response layouts. However, its improvement was not as strong as that of E, suggesting that discriminator-side physical supervision alone is insufficient to fully reshape the generative process without explicit generator-side structural guidance.

Third, the background statistical consistency term (B) had a particularly clear effect on ΔENL and BVE, while also improving the overall FID and SSIM. This confirms that background regularization is important for suppressing unrealistic speckle artifacts and aligning the generated background with SAR-specific statistical properties. Compared with

T, the isolated effect of *B* is more evident in the background-related metrics than in the target-region metric AFS, which is consistent with its intended role.

More importantly, the interactions among the three components were strongly complementary. The combination of *E* and *T* yielded better target-region consistency than either module alone, indicating that generator-side structural guidance and discriminator-side physical supervision work synergistically. Similarly, the combination of *E* and *B* produced substantial gains in both target-related and background-related metrics, which shows that physically plausible target synthesis and realistic SAR background modeling should be addressed jointly rather than separately. The complete model achieved the best overall performance, especially in FID, AFS, BVE, and CMAE. The slightly better SSIM and ΔENL obtained by baseline + *E* + *B* suggest that background regularization can strongly improve background-related statistics, but the full model provides a better balance between structural and statistical consistency.

Overall, the ablation study confirms that the improvement does not come from a single dominant module but from the collaboration of three complementary designs. ASC-inspired structural guidance improves aspect-dependent structural generation, target-region physical consistency enhances scattering-aware discrimination, and background statistical consistency stabilizes SAR-specific background realism. This joint effect is essential for physically consistent unseen-angle SAR view completion.

3.6. Recognition-Oriented Validation of Generated Samples

In addition to image-level and physical-consistency evaluations, we further conducted a recognition-oriented validation experiment to examine whether the generated SAR images could provide useful training information for downstream target recognition. The purpose of this experiment was not to optimize the recognition model itself but to evaluate the practical utility of the generated samples when they were used as auxiliary training data.

A ResNet-18 classifier was adopted as a fixed recognition network. The classifier was trained under three different data configurations: real-only training, generated-only training, and real data augmented with the generated samples produced by the proposed method. For the real-only setting, the classifier was trained using real SAR images of the five target categories at a depression angle of 15°, with aspect angles sampled at an interval of 5°. For the generated-only setting, we used the generated SAR images of the corresponding targets at the same depression angle of 15°, also with aspect angles sampled every 5°. For the augmented setting, the real training data were combined with the generated samples under the same angular sampling density. In all settings, the classifier was evaluated on real SAR test images at a depression angle of 15°, and the test set was not used during either image generation or classifier training.

As shown in Table 9, the classifier trained only with generated samples achieved an accuracy of 79.63% on real SAR test images. This result was lower than the real-only baseline of 85.13%, indicating that the generated images still could not completely replace real SAR data and that a certain domain gap remained between generated and real samples. Nevertheless, when the generated samples were combined with real training data, the recognition accuracy increased to 88.29%, which was 3.16 percentage points higher than the real-only baseline. This improvement demonstrates that the proposed generated images provide complementary aspect-angle-related information for classifier training.

These results suggest that the generated samples provide complementary information for recognition when used together with real training data, although a domain gap between generated and real samples still remains.

Table 9. Recognition-oriented validation of generated SAR samples using ResNet-18.

Training Data for the Recognition Model	Accuracy (%)
Real-only	85.13
Generated-only (ours)	79.63
Real + ours generated	88.29

4. Discussion

4.1. Why ASC-Inspired Structural Priors Matter for Unseen Angles

The experimental results consistently indicate that the main difficulty of unseen-angle SAR generation does not lie merely in limited sample size, but more fundamentally in how to model aspect-dependent structural variation. This point is particularly supported by the comparison between the interpolation and hold-out protocols. Under interpolation, a generator may still benefit from relatively smooth transitions between neighboring observed angles. However, under the hold-out protocol, the model is required to synthesize SAR images in a completely unobserved angular sector, where the generation quality depends much more strongly on whether the underlying evolution of dominant scattering structures has been captured rather than whether local textures have been memorized.

From this perspective, the advantage of the proposed method comes primarily from the introduction of an ASC-inspired sparse scattering-structure prior. Instead of treating the aspect angle as a low-dimensional control code that modulates the generator only implicitly, the proposed framework links aspect-angle variation with an angle-aware sparse structural proxy. As a result, the generator is encouraged to align viewpoint control with dominant scattering-layout generation. This explains why the proposed method achieves consistently larger gains in AFS and CMAE than in purely appearance-level metrics alone. In other words, the key contribution of the ASC-inspired structural prior is not simply to improve image realism but to guide the generated target structures toward more plausible and angle-consistent scattering-response layouts.

It is also important to clarify that the proposed structural prior does not impose a strictly smooth evolution assumption on individual scattering centers. Due to aspect-dependent visibility and electromagnetic scattering mechanisms, local dominant responses in SAR images may undergo abrupt changes even under small azimuth variations. Therefore, the term "structural variation" in this work refers to angle-consistent changes of dominant image-domain scattering layouts, rather than continuous tracking of identical physical scattering centers.

Another important implication is that physically meaningful intermediate representations are particularly valuable for extrapolative generation. When the test angle lies outside the observed angular support, the model can no longer rely on local interpolation in image space. Instead, it must infer how the target scattering structure should transform under a new viewpoint. The superior performance of the proposed method in the hold-out setting therefore suggests that the ASC-based representation provides a more robust inductive bias for unseen-angle SAR generation than conventional conditional synthesis alone.

4.2. Role of Target–Background Dual Physical Consistency

The results also show that physically consistent SAR generation should not be judged solely by target appearance. In target-chip SAR imagery, the generated image contains two physically different components: the target region, whose realism depends on scattering-structure plausibility, and the background region, whose realism is mainly reflected by speckle-like statistical characteristics. This distinction motivates the target–background dual consistency design in the proposed framework.

For the target region, the ASC-aware discriminator provides a structural consistency constraint that complements generator-side ASC-inspired structural guidance. The generator-side guidance encourages the formation of angle-dependent structural priors, while the discriminator-side consistency term further suppresses target patterns that may appear visually plausible but are inconsistent with scattering-aware structural cues. The ablation results show that this target-region consistency term alone is not as effective as the ASC-inspired structural guidance, but it becomes substantially more useful when combined with the generator-side structural prior. This suggests that physically meaningful target-region supervision is most effective when the generative process has already been structurally organized.

For the background region, the ENL-based regularization improves the statistical realism of SAR-specific speckle patterns. Although background regularization does not directly control the dominant target scattering responses, it plays an important stabilizing role in the overall generation process. In the absence of such a constraint, the adversarial objective may allow the model to produce locally unrealistic speckle granularity or intensity fluctuations, which degrades the physical credibility of the generated image even when the target structure appears reasonable. The proposed background consistency term therefore prevents the generator from sacrificing SAR-specific statistical realism for local appearance quality.

Another practical issue is the influence of pixel-level registration during training. Although SAR target chips are usually cropped, centered, and resized before being used for model training, such preprocessing only provides coarse spatial normalization and does not guarantee exact pixel-level correspondence of individual scattering responses. In SAR images, local dominant scattering responses may shift, appear, disappear, or fluctuate in intensity with small changes in aspect angle. Therefore, if the training objective relies heavily on strict pixel-wise reconstruction or registration, physically plausible local scattering variations may be over-penalized, which can lead to blurred or distorted target structures. This is one reason why target-region consistency in this work was evaluated and constrained mainly in an ASC-inspired structural feature space. Compared with direct pixel-level matching, the proposed feature-level constraint focuses more on the compatibility of dominant scattering-response layouts and is less sensitive to small local misregistration.

Taken together, the target-region structural consistency and the background-region statistical consistency are complementary rather than independent. The former improves the physical plausibility of dominant target scattering structures, whereas the latter regularizes SAR-specific background realism. Their joint effect explains why the complete model achieves stronger and more stable improvements than any single-component variant.

4.3. Comparison with Existing Paradigms

The discussion above also helps clarify the difference between the proposed method and existing SAR generation paradigms. On the one hand, physics-driven generative methods have shown that introducing scattering-related priors or physics-inspired constraints can improve physical plausibility under limited-data conditions. However, many of these methods still use physics mainly as a static prior, an auxiliary regularization term, or a post hoc consistency constraint. As a result, they may improve the realism of generated SAR images without explicitly addressing how target structures should evolve as the aspect angle changes.

On the other hand, angle-controllable generative methods explicitly encode viewpoint information and therefore improve the controllability of cross-angle synthesis. Nevertheless, these methods often treat the aspect angle primarily as a conditional variable in latent

space or image space. While such designs may enhance continuity across neighboring viewpoints, they do not necessarily guarantee physically plausible and angle-consistent structural variation, especially when the angular coverage of the training data is sparse or incomplete.

A related but different paradigm is 3D-aware SAR neural rendering, including SAR-NeRF and SAR Gaussian-splatting methods. Compared with ordinary 2D generative models, these methods have stronger potential for multi-view-consistent rendering because they explicitly optimize geometry-aware neural fields or Gaussian primitives. However, this advantage is usually obtained under a different problem setting, where relatively dense multi-view observations, accurate imaging geometry, and often target- or scene-specific optimization are available. In the present sparse target-chip generation setting, directly applying such methods would either require additional geometric supervision beyond the adopted protocol or make the 3D representation severely under-constrained. Therefore, they are not used as direct baselines in this study; instead, they are regarded as complementary geometry-aware approaches for future extensions.

In contrast, the proposed framework connects physical structural priors and angle-conditioned generation through a unified mechanism-aware design. It neither relies solely on static physics injection nor treats aspect angle as a purely abstract control code. Instead, it introduces an ASC-inspired sparse scattering-structure prior as an intermediate structural proxy, and further enforces physical consistency in both the target region and the background region. Different from 3D-aware neural rendering methods that aim to optimize geometry-aware target or scene representations, the proposed framework focuses on image-domain structural guidance under sparse 2D target-chip observations. From this viewpoint, the proposed method can be regarded as a physically constrained unseen-angle view-completion framework, rather than a conventional conditional SAR image generation model. This distinction is important because it shifts the problem formulation from appearance-oriented synthesis to structure-aware and mechanism-aware generation.

4.4. Limitations and Future Work

Despite the encouraging results, the proposed method still has several limitations. First, the current study was conducted under a relatively controlled imaging setting, including a single band, a single polarization, a fixed depression angle, and target-chip SAR data. This setting was intentionally adopted to isolate aspect-angle-induced structural variation from other imaging-geometry factors. Generating SAR images across different incidence or depression angles is indeed more meaningful for broader applications, but it also introduces additional changes in shadowing, layover, foreshortening, scattering strength, target projection, and background statistics. These factors are coupled with the aspect-angle variation and cannot be fully addressed by simply adding another angle label to the current framework. Therefore, the generalization of the proposed method to multi-band, multi-polarization, variable-depression-angle, or more complex scene-level SAR generation remains an important direction for future work.

Second, the ASC-inspired structural prior adopts a simplified sparse image-domain representation that focuses on dominant scattering responses. While this design is beneficial for interpretability and training stability, it is not equivalent to full parametric ASC inversion and does not fully capture all possible scattering mechanisms, especially for targets with highly distributed or strongly coupled scattering patterns. Although the proposed ASC-inspired structural prior improves the plausibility of dominant scattering layouts, it cannot guarantee exact recovery of all local scattering responses. This limitation is particularly evident for fine-scale target structures, where small changes in aspect angle may lead to significant variations in local bright responses. Therefore, the proposed method should be

regarded as improving structural consistency at the dominant-response level, rather than fully reconstructing the complete target scattering mechanism. Similarly, the ENL-based background regularization mainly constrains low-order background statistics. Although it is effective for target-chip SAR backgrounds, more complex backgrounds may require richer statistical or semantic modeling.

Third, the current evaluation still depends on an independently pretrained ASC extractor and a separate aspect-angle estimator. While these modules provide useful proxies for physical consistency and angle-condition consistency, they may also introduce estimation bias. A more complete evaluation framework that combines image fidelity, physical consistency, and mechanism-level validation would further strengthen the interpretability of SAR generation methods.

Future work may proceed in several directions. One promising direction is to extend the current 2D physically constrained generation framework toward 3D-aware or geometry-aware SAR generation, so that aspect-angle variation can be modeled more naturally under broader viewing conditions. Another direction is to incorporate richer physical priors, such as multi-polarization scattering behavior or frequency-dependent responses, into the generative process. In addition, constructing larger and more diverse SAR benchmarks with explicit angular and physical annotations would be highly valuable for evaluating physically consistent unseen-angle generation in more realistic scenarios.

5. Conclusions

This paper addressed the problem of SAR image generation at unseen aspect angles under sparse angular observations. Different from conventional conditional SAR image synthesis methods that mainly rely on appearance matching or low-dimensional angle encoding, the proposed approach reformulated unseen-angle SAR generation as a physically constrained view-completion problem. To this end, a physically consistent generation framework was developed by introducing an ASC-inspired sparse scattering-structure prior to guide aspect-dependent target structural generation. On this basis, a target–background dual physical consistency design was further incorporated, where ASC-aware structural discrimination was used to constrain target-region scattering plausibility, and ENL-based statistical regularization was used to preserve SAR-specific background realism.

Extensive experiments under both interpolation and hold-out protocols demonstrated that the proposed method consistently outperformed representative baseline methods in image fidelity, target-region structural consistency, background statistical consistency, and angle-condition consistency. The quantitative and qualitative results jointly showed that introducing an ASC-inspired sparse scattering-structure prior is effective for unseen-angle SAR generation, especially when the angular coverage of training data is sparse and the target structure must be inferred in a physically meaningful manner beyond the observed viewpoints. The ablation study further verified that the final performance gain arises from the coordinated effect of generator-side structural guidance, target-region physical consistency, and background statistical consistency.

Overall, this work suggests that physically consistent unseen-angle SAR generation should not be treated merely as a generic conditional image synthesis task. Instead, it should be approached as a physics-aware generation problem in which aspect-angle control, dominant scattering-structure variation, and SAR-specific background statistics are jointly modeled. This study provides a physically interpretable solution for SAR view completion under limited-view conditions and offers a useful perspective for future research on physics-constrained SAR image generation. In future work, the proposed framework can be extended toward more challenging settings, such as multi-polarization,

variable depression angles, and geometry-aware SAR generation under more complex scene conditions.

Author Contributions: Conceptualization, Z.J. and C.L.; Methodology, Z.J.; Software, Z.J. and Z.X.; Validation, Z.J.; Formal Analysis, Z.J. and Y.W.; Investigation, Z.J. and Z.X.; Resources, Y.W. and S.X.; Data Curation, Z.X.; Writing—Original Draft Preparation, Z.J.; Writing—Review and Editing, Z.J. and C.L.; Supervision, C.L., Y.W., and S.X.; Project Administration, Z.X. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by a research project of the Military Academy of China (classified, no public grant number available). The APC was funded by Chao Liu.

Data Availability Statement: The data presented in this study are available upon request from the corresponding author.

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

References

1. Zhang, T.; Zeng, T.; Zhang, X. Synthetic Aperture Radar (SAR) Meets Deep Learning. *Remote Sens.* **2023**, *15*, 303. [CrossRef]
2. Li, J.; Yu, Z.; Yu, L.; Cheng, P.; Chen, J.; Chi, C. A Comprehensive Survey on SAR ATR in Deep-Learning Era. *Remote Sens.* **2023**, *15*, 1454. [CrossRef]
3. Li, J.; Xu, C.; Su, H.; Gao, L.; Wang, T. Deep Learning for SAR Ship Detection: Past, Present and Future. *Remote Sens.* **2022**, *14*, 2712. [CrossRef]
4. Li, W.; Yang, W.; Hou, Y.; Liu, L.; Liu, Y.; Li, X. SARATR-X: Toward Building a Foundation Model for SAR Target Recognition. *IEEE Trans. Image Process.* **2025**, *34*, 869–884. [CrossRef] [PubMed]
5. Dong, G.; Song, Y. SAR Target Augmentation and Recognition via Cross-Domain Reconstruction. *Pattern Recognit.* **2025**, *159*, 111117. [CrossRef]
6. Zhang, L.; Leng, X.; Feng, S.; Ma, X.; Ji, K.; Kuang, G.; Liu, L. Azimuth-Aware Discriminative Representation Learning for Semi-Supervised Few-Shot SAR Vehicle Recognition. *Remote Sens.* **2023**, *15*, 331. [CrossRef]
7. Zhao, Y.; Zhao, L.; Zhang, S.; Ji, K.; Kuang, G.; Liu, L. Azimuth-Aware Subspace Classifier for Few-Shot Class-Incremental SAR ATR. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 1–20. [CrossRef]
8. Zhang, L.; Leng, X.; Feng, S.; Ma, X.; Ji, K.; Kuang, G.; Liu, L. Optimal Azimuth Angle Selection for Limited SAR Vehicle Target Recognition. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *128*, 103707. [CrossRef]
9. Liu, Y.; Li, W.; Liu, L.; Zhou, J.; Peng, B.; Song, Y.; Xiong, X.; Yang, W.; Liu, T.; Liu, Z.; et al. ATRNet-STAR: A Large Dataset and Benchmark Towards Remote Sensing Object Recognition in the Wild. *IEEE Trans. Pattern Anal. Mach. Intell.* **2026**, 48, 6735–6753. [CrossRef] [PubMed]
10. Huang, Z.; Zhang, X.; Tang, Z.; Xu, F.; Datcu, M.; Han, J. Generative Artificial Intelligence Meets Synthetic Aperture Radar: A Survey. *IEEE Geosci. Remote Sens. Mag.* **2026**, *14*, 6–48. [CrossRef]
11. Yue, D.X.; Xu, F.; Frery, A.C.; Jin, Y.Q. Synthetic Aperture Radar Image Statistical Modeling: Part One—Single-Pixel Statistical Models. *IEEE Geosci. Remote Sens. Mag.* **2021**, *9*, 82–114. [CrossRef]
12. Zhu, R.; Teng, F.; Hong, W. Analysis and Modeling of Statistical Distribution Characteristics for Multi-Aspect SAR Images. *Remote Sens.* **2025**, *17*, 1295. [CrossRef]
13. Sun, X.; Li, X.; Xiang, D.; Hu, C. SAR Vehicle Image Generation with Integrated Deep Imaging Geometric Information. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *132*, 104028. [CrossRef]
14. Liu, Z.; Niu, S.; Qiu, X.; Peng, L.; Shang, Y.; Zhong, L.; Ding, C. A Differentiable Method for Novel View SAR Image Generation via 3D Gaussian Splatting. *ISPRS J. Photogramm. Remote Sens.* **2026**, *231*, 167–195. [CrossRef]
15. Wang, C.; Pei, J.; Liu, X.; Huang, Y.; Mao, D.; Zhang, Y.; Yang, J. SAR Target Image Generation Method Using Azimuth-Controllable Generative Adversarial Network. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2022**, *15*, 9381–9397. [CrossRef]
16. Liu, M.; Wang, H.; Chen, S.; Tao, M.; Wei, J. A Two-Stage SAR Image Generation Algorithm Based on GAN with Reinforced Constraint Filtering and Compensation Techniques. *Remote Sens.* **2024**, *16*, 1963. [CrossRef]
17. Radford, A.; Metz, L.; Chintala, S. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2016.
18. Odena, A.; Olah, C.; Shlens, J. Conditional Image Synthesis with Auxiliary Classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2017; Volume 70, pp. 2642–2651.

19. Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; Aila, T. Analyzing and Improving the Image Quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*; IEEE: Piscataway, NJ, USA, 2020; pp. 8110–8119. [CrossRef]
20. Yue, D.X.; Xu, F.; Frery, A.C.; Jin, Y.Q. Synthetic Aperture Radar Image Statistical Modeling: Part Two—Spatial Correlation Models and Simulation. *IEEE Geosci. Remote Sens. Mag.* **2021**, *9*, 115–138. [CrossRef]
21. Xiang, D.; Liu, Y.; Cheng, J.; Lu, X.; Xie, Y.; Guan, D. SAR Target Recognition with Image Generation and Azimuth Angle Feature Constraints. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2025**, *18*, 18561–18580. [CrossRef]
22. Cui, Y.; Liu, Z.; Ruan, L.; Sheng, B.; Wang, N.; Xiao, X.; Bian, X. An Azimuth-Continuously Controllable SAR Image Generation Algorithm Based on GAN. *Remote Sens.* **2025**, *17*, 3763. [CrossRef]
23. Yang, R.; Wang, B.; Lai, T.; Huang, H. Angle-Controllable SAR Image Generation and Target Recognition via StyleGAN2. *Remote Sens.* **2025**, *17*, 3478. [CrossRef]
24. Remusati, H.; Le Caillec, J.M.; Schneider, J.Y.; Petit-Frère, J.; Merlet, T. Generative Adversarial Networks for SAR Automatic Target Recognition and Classification Models Enhanced Explainability: Perspectives and Challenges. *Remote Sens.* **2024**, *16*, 2569. [CrossRef]
25. Li, J.; Zhu, G.; Hou, C.; Zhang, W.; Du, K.; Cheng, C.; Wu, K. Ray-Tracing-Assisted SAR Image Simulation under Range Doppler Imaging Geometry. *Electronics* **2024**, *13*, 3591. [CrossRef]
26. Feng, S.; Fu, X.; Feng, Y.; Lv, X. Single-Scene SAR Image Data Augmentation Based on SBR and GAN for Target Recognition. *Remote Sens.* **2024**, *16*, 4427. [CrossRef]
27. Zhang, X.; Zhuang, Y.; Guo, Q.; Yang, H.; Qian, X.; Cheng, G.; Han, J.; Huang, Z. Φ-GAN: Physics-Inspired GAN for Generating SAR Images Under Limited Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*; IEEE: Piscataway, NJ, USA, 2025; pp. 29075–29085.
28. Liu, Z.; Niu, S.; Qiu, X.; Peng, L.; Ding, C. A Method for Mapping Strong Scattering Information in SAR Images to 3D Target Geometry Using a Customized Differentiable SAR Simulator. *J. Remote Sens.* **2026**, *6*, 1030. [CrossRef]
29. Fu, S.; Xu, F. Differentiable SAR Renderer and Image-Based Target Reconstruction. *IEEE Trans. Image Process.* **2022**, *31*, 6679–6693. [CrossRef] [PubMed]
30. Lei, Z.; Xu, F.; Wei, J.; Cai, F.; Wang, F.; Jin, Y.Q. SAR-NeRF: Neural Radiance Fields for Synthetic Aperture Radar Multiview Representation. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 1–15. [CrossRef]
31. Ehret, T.; Mar'i, R.; Derksen, D.; Gasnier, N.; Facciolo, G. Radar Fields: An Extension of Radiance Fields to SAR. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*; IEEE: Piscataway, NJ, USA, 2024; pp. 4249–4258.
32. Li, A.; Lei, Z.; Wei, J.; Xu, F. SAR-GS: Gaussian Splatting Based SAR Images Rendering and Target Reconstruction. *arXiv* **2025**, arXiv:2506.21633.
33. Wang, X.; Hui, B.; Wang, W.; Guo, P.; Ding, L.; Lin, H. Angle-Controllable SAR Image Generation for Target Recognition with Few Samples. *Remote Sens.* **2025**, *17*, 1206. [CrossRef]
34. Guo, S.; Chen, T.; Wang, P.; Liu, H.; Chen, W.; Yan, J. Label-Aspect-Conditioned Diffusion Model for SAR Image Generation and Imbalanced Target Recognition. *Appl. Soft Comput.* **2026**, *197*, 115146. [CrossRef]
35. Liu, Y.; Ma, L.; Chen, Z.; He, Z. Azimuth-Transfer Based SAR Image Generation with Few Samples. *Geocarto Int.* **2026**, *41*, 2618322. [CrossRef]
36. Potter, L.C.; Moses, R.L. Attributed Scattering Centers for SAR ATR. *IEEE Trans. Image Process.* **1997**, *6*, 79–91. [CrossRef] [PubMed]
37. Fan, J.; Tomas, A. Target Reconstruction Based on Attributed Scattering Centers with Application to Robust SAR ATR. *Remote Sens.* **2018**, *10*, 655. [CrossRef]
38. Lee, J.S. Speckle Analysis and Smoothing of Synthetic Aperture Radar Images. *Comput. Graph. Image Process.* **1981**, *17*, 24–32. [CrossRef]
39. Keydel, E.R.; Lee, S.W.; Moore, J.T. MSTAR Extended Operating Conditions: A Tutorial. In *Proceedings of the Algorithms for Synthetic Aperture Radar Imagery III*; Proceedings of SPIE; Zelnio, E.G., Douglass, R.J., Eds.; SPIE: Bellingham, WA, USA, 1996; Volume 2757, pp. 228–242. [CrossRef]
40. Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; Aila, T. Training Generative Adversarial Networks with Limited Data. In *Proceedings of the Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2020; Volume 33, pp. 12104–12114.
41. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Proceedings of the Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2017; Volume 30, pp. 6626–6637.
42. Wang, Z.; Bovik, A.C.; Sheikh, H.R.; Simoncelli, E.P. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Trans. Image Process.* **2004**, *13*, 600–612. [CrossRef] [PubMed]
43. Shen, P.; Wang, C.; Fu, H.; Zhu, J.; Hu, C. Estimation of Equivalent Number of Looks in Time-Series Pol(In)SAR Data. *Remote Sens.* **2020**, *12*, 2715. [CrossRef]

44. Yu, Z.; Dong, G.; Liu, H. SAR Image Quality Assessment: From Sample-Wise to Class-Wise. *Remote Sens.* **2023**, *15*, 2110. [CrossRef]
45. Oh, J.; Kim, M. PeaceGAN: A GAN-Based Multi-Task Learning Method for SAR Target Image Generation with a Pose Estimator and an Auxiliary Classifier. *Remote Sens.* **2021**, *13*, 3939. [CrossRef]

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.