electronics

MDPI

Article

Controlled PCA-Based Generation of Synthetic Object Images for Satellite SAR Data Augmentation †

Jolanta Chmielińska [1,*], Jacek Jakubowski [1,*], Kinga Karwowska [2] and Damian Wierzbicki [2]

[1] Faculty of Electronics, Military University of Technology, 00-908 Warsaw, Poland
[2] Faculty of Civil Engineering and Geodesy, Military University of Technology, 00-908 Warsaw, Poland; kinga.karwowska@wat.edu.pl (K.K.); damian.wierzbicki@wat.edu.pl (D.W.)
* Correspondence: jolanta.pacan@wat.edu.pl (J.C.); jacek.jakubowski@wat.edu.pl (J.J.)
† This paper is an extended version of the work presented at the 2026 Progress in Applied Electrical Engineering Conference (PAEE): J. Jakubowski, J. Chmielińska, K. Karwowska and D. Wierzbicki, Using PCA Transformation to Generate Object Images in Satellite SAR Scenes, https://doi.org/10.1109/PAEE71887.2026.11660700.

Abstract

Nowadays, synthetic aperture radars are widely used on various satellite imaging platforms. The high availability of SAR imagery makes it very popular in various automatic object detection systems. However, the visual interpretation of SAR data is difficult even for experts, which poses the primary challenge in preparing the data required for training and evaluating deep learning models. Therefore, using data augmentation techniques to expand the original training dataset is a popular approach. In this paper, a method of synthetic SAR object image generation is proposed using Principal Component Analysis. The presented technique is based on principal component perturbations and can be used for data augmentation to train object detection models. The numerical experiments performed with the YOLO model showed that the proposed method improved AP50 on the test set for transport aircraft detection, with statistical significance confirmed by bootstrap analysis.

Keywords: principal component analysis; data augmentation; SAR imagery; objects detection; YOLO model

1. Introduction

Synthetic Aperture Radar (SAR) imagery represents a valuable source of information for satellite Earth observation applications. In contrast to optical sensors, SAR systems operate as active microwave sensing instruments, allowing data acquisition regardless of daylight availability and under diverse atmospheric conditions. Contemporary high-resolution SAR systems support the analysis of various objects and infrastructures related to transportation networks, maritime operations, port facilities, airports, environmental monitoring, and disaster management [1].

The increasing availability of SAR imagery has rendered manual analysis insufficient, thereby enhancing the importance of automatic object detection methods based on deep neural networks. In the context of image analysis, this task can be performed using object detectors originally developed for conventional computer vision applications, such as the YOLO family of models. The effectiveness of these approaches, however, largely depends on the availability of extensive and diverse annotated training datasets. For SAR

Academic Editors: Aili Wang and Haibin Wu

Received: 15 September 2026
Revised: 19 October 2026
Accepted: 29 October 2026
Published: 4 November 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.

imagery, this challenge is particularly pronounced. Radar images differ substantially from optical imagery (Figure 1), as the appearance of observed objects is affected by factors including imaging geometry, electromagnetic scattering mechanisms, and speckle noise inherent to coherent imaging systems [2]. Consequently, SAR scene interpretation is more demanding, and the generation of reliable annotations frequently requires expert knowledge. This hinders the development of large annotated datasets, despite the fact that deep learning models typically benefit from large numbers of diverse training samples. The scarcity of annotated SAR data has therefore been identified as a major challenge in the literature on data augmentation for radar object recognition [3].

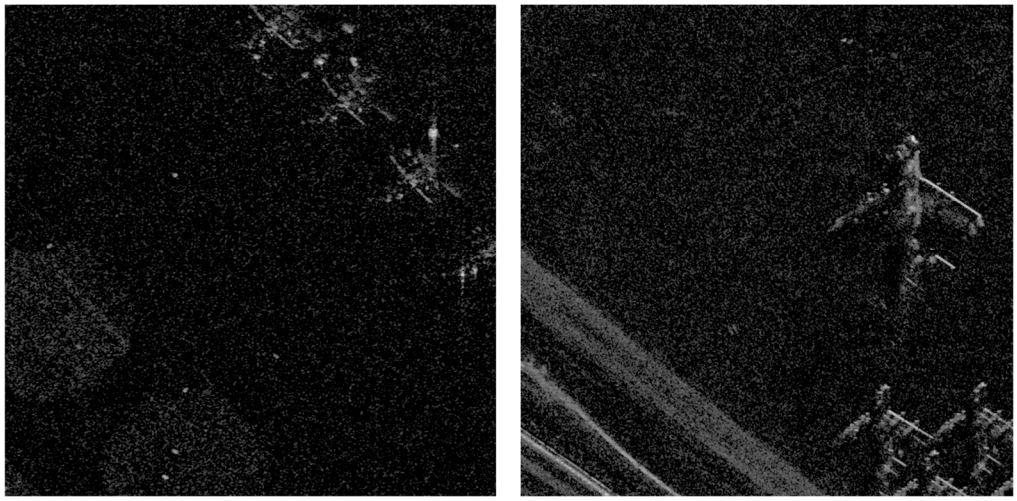


Figure 1. Examples of SAR images including transport aircraft.

A natural approach to partially addressing this problem is data augmentation. In conventional image processing, augmentation primarily involves geometric and arithmetic transformations, such as rotation, reflection, translation, scaling, cropping, as well as adjustments to brightness and contrast. These techniques are widely employed during the training of deep learning models for computer vision and often constitute an integral part of their preprocessing pipeline. Their primary objective is to increase the diversity of the training dataset without the need to acquire additional manually annotated samples [4].

However, in the case of SAR imagery, an additional limitation arises. Not every mathematically valid image transformation results in a sample that could realistically be produced through the radar imaging process. This stems from the fact that a SAR image is the outcome of a coherent scattering and radar signal processing mechanism, meaning that its statistical properties and spatial structures are subject to physical constraints that are absent in conventional optical imagery. Consequently, it is possible to generate synthetic data that are numerically correct yet unsuitable for training an object detector. An extreme example would be a transformation that reverses intensity relationships within the image, a result that cannot be associated with any realistic electromagnetic scattering mechanism responsible for the formation of SAR imagery. From the perspective of data augmentation, it is therefore important not only to increase the number of training samples but also to control the extent to which the generated samples preserve the characteristics of real SAR images.

In our previous work [5], we proposed generating new object images by perturbing principal components and subsequently reconstructing the images using the PCA matrix derived from real SAR data. This approach is attractive because it relies on a basis constructed directly from genuine images belonging to the class under investigation. However, the mere fact that an image is reconstructed from the PCA matrix extracted from real data does not guarantee the physical plausibility of every possible combination of

https://doi.org/10.3390/electronics15215125

principal components. As the magnitude of the perturbation increases, the diversity of the generated images also increases; however, so does the risk of deviating from structures characteristic of real SAR observations. From an application perspective, the suitability of a particular perturbation range could be evaluated by generating an augmented dataset, training an object detector, and assessing the resulting detection performance. However, such an approach would necessitate repeated detector training for multiple parameter settings and would therefore entail substantial computational cost. For this reason, the present work, extending the methodology proposed in [5], explores the use of a computationally efficient surrogate measure that can characterize the degree of modification of generated images before undertaking the computationally intensive detector training process.

To this end, the Structural Similarity Index Measure (SSIM) [6] was adopted as a surrogate metric. SSIM was originally developed to quantify the structural similarity between an image and a reference image by jointly evaluating characteristics such as structural content, contrast, and luminance. In the context of the present work, SSIM is not interpreted as a measure of the physical realism or correctness of SAR imagery. Rather, it serves as a computationally efficient indicator of the extent to which a generated sample departs from the original source image.

A characteristic trade-off emerges in this context. Very high SSIM values indicate that the generated samples closely resemble the original images, potentially limiting the effectiveness of augmentation due to insufficient variability. In contrast, lower SSIM values correspond to greater sample diversity, but may also indicate substantial alterations of the original image structure, increasing the risk of generating samples that are no longer representative of genuine SAR observations and are therefore of limited utility for detector training. Accordingly, the aim of this work is to identify a range of structural similarity values that achieve a favorable compromise between preserving the characteristics of real SAR data and enhancing the diversity of the training dataset.

Issues related to air traffic, which has been steadily increasing in density in recent years, are highly popular in the scientific literature [7,8]. Therefore, the proposed procedure was validated using transport aircraft as a representative class of objects commonly observed in satellite SAR imagery. This class was selected as a case study to evaluate both the process of generating new object images and their subsequent usefulness as augmentation data. The quantitative assessment of the generated samples was performed using a YOLO nano detector trained on datasets augmented with images produced under different PCA perturbation parameter configurations. The overall workflow of the proposed methodology is presented in Figure 2.

The main contributions of the present work are as follows. Building upon and extending the methodology introduced in [5], we further develop an approach for generating SAR object images through perturbation of principal components. In addition, we analyze the relationship between perturbation parameters and the resulting distribution of SSIM values, investigate the use of SSIM as a computationally efficient criterion for selecting appropriate perturbation levels, and validate the selected similarity ranges through YOLO detector training. Finally, the statistical significance of the observed improvements in the fundamental AP50 detection metric is assessed using bootstrap analysis. All these aspects are entirely novel in this paper in comparison with [5], where only the concept of generating new SAR images was presented.

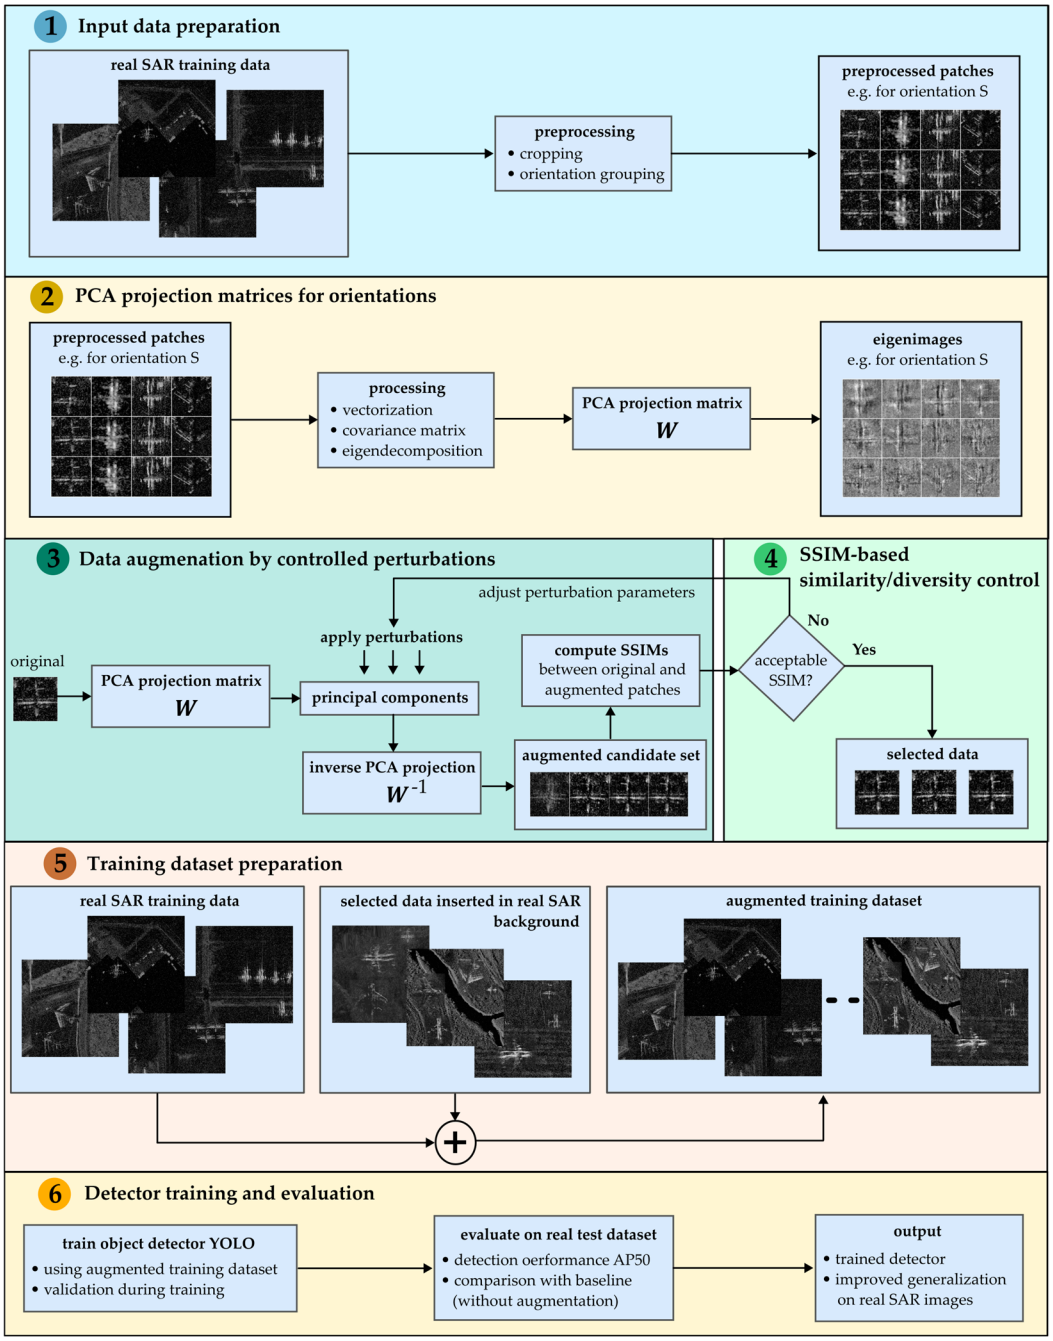


Figure 2. Overall workflow of the proposed methodology.

2. Related Work

Principal Component Analysis (PCA) is widely used across numerous application domains. Owing to its ability to extract informative features, it has been extensively employed in a variety of classification tasks, particularly those related to image recognition. In particular, in [3], PCA was applied to object recognition in SAR imagery.

A less common, yet well-established, application of PCA is data augmentation. The authors of [9] proposed a method for generating additional data vectors, referred to as PCA-distilled vectors, derived from the principal components of the training dataset and subsequently incorporated into the data used for classifier training.

Their results demonstrated improvements in selected classification performance metrics across several public datasets and multiple machine learning models. However, this approach was developed for general vector data and does not address the problem considered in the present study, namely the generation of SAR images and the control of

similarity between synthetic samples and real observations. Similarly, the authors of [10] employed principal component perturbation to modify RGB channel intensity values, thereby introducing random variations in image coloration for neural network training.

In the literature dedicated specifically to the generation of additional SAR data, several approaches can be identified, differing in their input data requirements, the representation in which new samples are generated, the amount of training data required, and the degree of control over the synthesis process. Rather than discussing this extensive research area in detail, this section focuses on selected, representative, and relatively recent methods that explicitly generate additional SAR samples and evaluate their usefulness in recognition tasks.

One group of methods comprises generative approaches based on image translation between the optical and radar domains. A broader review of generative methods used for such domain transformation is provided in [11], including GAN-based and diffusion-based approaches. They are capable of producing large quantities of synthetic data; however, the generated images are primarily derived from statistical mappings between domains rather than from direct modeling of the SAR image formation process. One approach using optical images to generate SAR data is the CDA-GAN method presented in [12]. The authors proposed an optical-to-SAR translation architecture based on the CycleGAN framework, extended with a cross-domain attention mechanism, a cross-domain multi-scale feature fusion module, and an additional reference loss. The cross-domain attention mechanism was designed to preserve the semantic information of the optical image while introducing features characteristic of the SAR domain, whereas multi-scale feature fusion supported the preservation of object geometry by combining local and global information at different representation levels. The method enabled the generation of additional SAR samples depicting ship scenes for the detection task. In an experiment with the YOLOv7 detector, an AP of 70.30% was achieved, compared with 62.84% for the standard CycleGAN. Training the translation model required access to both real optical and SAR images; however, these datasets did not need to be paired.

An alternative approach to generating additional synthetic SAR images, which does not require access to optical data, was presented in [13]. The authors used an adversarial autoencoder (AAE) to generate additional SAR images of ground vehicles, mainly military ones, for data augmentation in a classification task using A-ConvNet. A characteristic feature of the proposed model was learning a representation of the relationship between the target class and orientation and the corresponding SAR image, which subsequently enabled the generation of new samples for specified intermediate orientations by modifying explicitly defined attributes in the learned representation space. The experiments were conducted on the MSTAR dataset, containing image chips of the same target classes acquired over a wide range of aspect angles, while only 90 samples selected from a larger dataset were used to train the AAE. An additional requirement of the method was the preparation of pixel-level semantic maps distinguishing the target, radar shadow, and background. The use of data generated by the AAE increased the A-ConvNet classification accuracy from 87.59% to 93.36% in the best-performing variant with 25 additional samples generated for each augmentation case.

A mechanism for generating new SAR images, conceptually similar to that presented in this work, was proposed in [14]. The authors developed a modified RC-CapsNet architecture, based on the classical CapsNet, to obtain a deterministic and interpretable feature representation of cropped SAR images containing ground vehicles. The resulting representation was then further decoupled using an attention mechanism, allowing features related to amplitude, elevation angle, azimuth angle, and target shape to be distinguished. Selected groups of features were perturbed, and the modified representation was subsequently transformed by a reconstruction network into new synthetic SAR images. The

effectiveness of this method relied on the availability of a dataset containing observations of the same target classes under different imaging geometries, which enabled the learning of attribute representations related, among others, to elevation and azimuth angles. However, the authors limited the validation of the generated data to qualitative image assessment and the analysis of selected image quality metrics, without conducting an experiment to evaluate the effect of augmentation on the performance of a downstream target detection or classification task.

The most complete solution in terms of validating a SAR data generation method for object detection was presented in [15], which employed the open-source Stable Diffusion 1.5 model, pre-trained on a large-scale image dataset. The authors used the HRSID dataset, containing SAR images of ships, to fine-tune the model to the SAR domain. A ControlNet was then attached to the resulting model and trained using pairs consisting of a SAR image and the corresponding spatial representation of bounding boxes. The resulting model enabled the generation of new SAR scenes while controlling the spatial arrangement of objects according to the specified locations and sizes. However, the authors did not investigate whether the method provides control over orientation, preserves the diversity of object types present in the training data, or whether the generated images reproduce the physical properties of SAR data. The usefulness of the synthetic data was instead evaluated directly in an object detection task using Mask R-CNN. For the HRSID dataset without additional augmentation, an mAP50 of 0.787 was obtained; applying a simple horizontal flip increased this value to 0.825, while using synthetic data corresponding to 40% of the amount of real data yielded the highest mAP50 of 0.831.

The above review was limited to methods using SAR amplitude images, consistent with the scope of the data analyzed in this study. A detailed comparison of the discussed approaches with the proposed PCA-based method is provided after presenting the method and the experimental results.

Another line of research focuses on physics-based SAR data augmentation methods that exploit complex-valued data and information related to the radar image formation process. In [16], target images corresponding to new viewing angles were synthesized by exploiting the sparsity of dominant scattering centers and performing interpolation in the phase-history domain. This approach accounts for variations in the SAR target response associated with changes in observation geometry and therefore provides a closer link to the SAR image formation process than purely image- or feature-domain transformations. However, it requires access to complex-valued SAR data containing phase information, which may not be available in standard amplitude-only image products.

3. Materials and Methods

Principal Component Analysis [17] enables the transformation of multivariable data to a new set of uncorrelated variables, called principal components, which are ordered by decreasing variance. This means that the initial principal components capture most of the information from the original dataset. PCA is a linear transformation, and it can be described by the formula:

$$\boldsymbol{y} = \boldsymbol{W}(\boldsymbol{x} - \boldsymbol{\mu}), \tag{1}$$

where $\boldsymbol{y}$—the resulting vector of principal components, $\boldsymbol{W}$—the transformation matrix of PCA coefficients, $\boldsymbol{x}$—the original data vector, $\boldsymbol{\mu}$—the mean vector of the original data matrix.

The matrix $\mathbf{W}$ is computed by the decomposition of the covariance matrix of the multivariate data into the eigenvalues and the corresponding eigenvectors. The eigenvectors, ordered by decreasing eigenvalues, form the transformation matrix, which can include all eigenvectors or a reduced subset of them. Omitting the eigenvectors

corresponding to the smallest eigenvalues from the transformation matrix enables dimensionality reduction without a significant loss of information. In such a case, the dimensionality of the PCA space is lower than that of the original data. PCA is a reversible transformation. Thus, it is possible to reconstruct the original data vector from the principal components using the transposed PCA matrix according to the formula:

$$\boldsymbol{x}_{\mathrm{rec}} = \boldsymbol{\mu} + \boldsymbol{W}^{\mathrm{T}}\boldsymbol{y}. \tag{2}$$

If the complete projection matrix is utilized, the reconstructed vector $\boldsymbol{x}_{\mathrm{rec}}$ perfectly matches the original data vector $\mathbf{x}$. Using a reduced transformation matrix results in reconstructing the data vector with a certain error.

Applying PCA for image data processing requires unifying the image size to an established spatial resolution $m \times n$ followed by image vectorization. After this procedure, each image is presented as a data vector of length $L = m \cdot n$. This results in an immense covariance matrix that must be decomposed into eigenvalues and eigenvectors. In practice, the number of processed images (N) is often much smaller than the image vector length (L). In such cases, there are only $N - 1$ meaningful eigenvectors, which can be determined by decomposing the $N \times N$ inner product matrix instead of the full covariance matrix. This approach was presented in [18] and is widely used in image processing tasks because it simplifies computations. Therefore, this methodology was adopted by the authors of this paper.

The proposed method of synthetic image generation using Principal Component Analysis is a two-stage process that must be preceded by the computation of the PCA projection matrix based on genuine SAR object images. In the first stage, a region of interest containing a target object is projected onto the PCA subspace. In the next stage, the obtained principal components are modified, and the inverse transformation is performed for image reconstruction while retaining all modified components. As a result, a new object image is obtained that differs from the original but preserves the distinct properties of SAR imagery. SAR object images of transport aircraft were selected from the training dataset by an expert and categorized into eight subsets containing objects with similar spatial orientations. Each subset was assigned to a specific object orientation, and their labeling followed standard compass directions, namely N, NE, E, SE, S, SW, W, and NW. The images for each category were assigned based on a general visual inspection, without determining precise object orientations or establishing strict boundary criteria. The number of images in each subset varied between 17 and 32. These subsets of transport aircraft images were used to compute the projection PCA matrices for generating synthetic images presenting objects with various spatial orientations. The use of separate subsets for various object orientations was motivated by the sensitivity of Principal Component Analysis to object rotation.

The adaptation of PCA transformation to the issue of producing new synthetic SAR images of objects involves the formal application of reconstruction according to Equation (2), but after introducing perturbations to the vector $\boldsymbol{y}$. The k-th entry of the modified principal component vector can be expressed by the following formula:

$$y_{k,\mathrm{mod}} = y_k + m_k \cdot v_{\mathrm{rand}}(\beta), \tag{3}$$

where y_k—the k-th element of the original principal component vector, m_k—the variable determining the modification factor, v_{rand}—a random number taken from a fixed distribution in the interval $(-\beta,\ \beta)$. In this paper, the uniform distribution was chosen due to the lack of a priori knowledge regarding preferred values of the modification factor, ensuring that all values within the specified range are equally probable. The perturbation determined by (3) can follow different strategies, depending on how the numerical values for the variable m_k are adopted.

The first method is based on the eigenvalues that determine the variances of the respective principal components and therefore take non-negative values. In this case, the modifications of the k-th element of the vector $\boldsymbol{y}$ can be described with the equation:

$$y_{k,\mathrm{mod}} = y_k + \sqrt{\lambda_k} \cdot v_{\mathrm{rand}}(\beta), \tag{4}$$

where λ_k is the eigenvalue representing the variance of the k-th principal component. In this method, the individual principal components are modified to a degree depending on their variances, which testify to their importance in the reconstruction. The greater the eigenvalue corresponding to a given principal component, the greater its potential modification. This enables the generation of new data from the same distribution as the original SAR data.

In the second method, the multiplicative modifications of principal components are made according to the equation:

$$y_{k,\mathrm{mod}} = y_k + y_k \cdot v_{\mathrm{rand}}(\beta). \tag{5}$$

This strategy is inspired by the mechanisms of speckle noise formation in SAR images. Each principal component changes by a certain percentage of its original value. This means that the larger coefficients might change more significantly than the smaller ones. This reduces the risk that insignificant components, often containing noise or small artifacts, will be over-amplified. Particularly, coefficients with zero values do not change at all, which means that this method does not activate PCA directions that are absent in the original image. A synthetic ROI therefore remains close to its original representation.

Both presented above strategies of new data generation can be implemented with restrictions on modifying the first main components that are responsible for the geometry of the object. By making changes only to principal components with smaller variances, it is possible to preserve the shape of the object in a more faithful way while affecting the texture of the image, which in SAR imagery is associated with the granularity resulting from speckle noise.

The presented procedure of synthetic image generation using PCA transformation is performed on ROIs, not the whole images. To produce a new training sample, four synthetic ROIs were pasted into a background image. This number was selected based on the median object count per image within the training set. A soft alpha blending technique was used to avoid sharp transitions between the ROI and the background. The ROI insertion procedure followed the methodology described in [19] using a Gaussian function instead of a linear function for the transition mask. The standard deviation of the Gaussian function was scaled according to the shorter edge of the pasted image, equal to 0.15 times its length, with its maximum value strictly capped at 35. Figure 3 presents an example of a background image and an augmented image generated by combining it with synthetic ROIs.

In the presented research, a database of image radar data from the satellites of the Capella Space constellation was used. Large-scale SAR imagery was divided into 640 × 640 pixel tiles, which is a default image size for YOLO models. Transport aircraft locations in images were annotated manually by human experts using rectangular bounding boxes, whose parameters were stored in additional text files. The processed images also included other aircraft and elements of airport infrastructure. The whole dataset consisted of 467 images containing objects of interest and 640 background images, which were used for the generation of synthetic images. All background images originated from the same large-scale SAR imagery that provided the object-bearing images. No additional selection procedure was applied. The set of images containing transport aircraft was split randomly into three disjoint subsets. A total of 326 images were used for training and 50 images for validation of YOLO-based transport aircraft detectors. The test set consisted of 91 images.

The fixed data split was caused by the proposed method of synthetic ROI generation. Each subset can include images from the same SAR scene due to the presence of various transport aircraft types across different large-format SAR images.

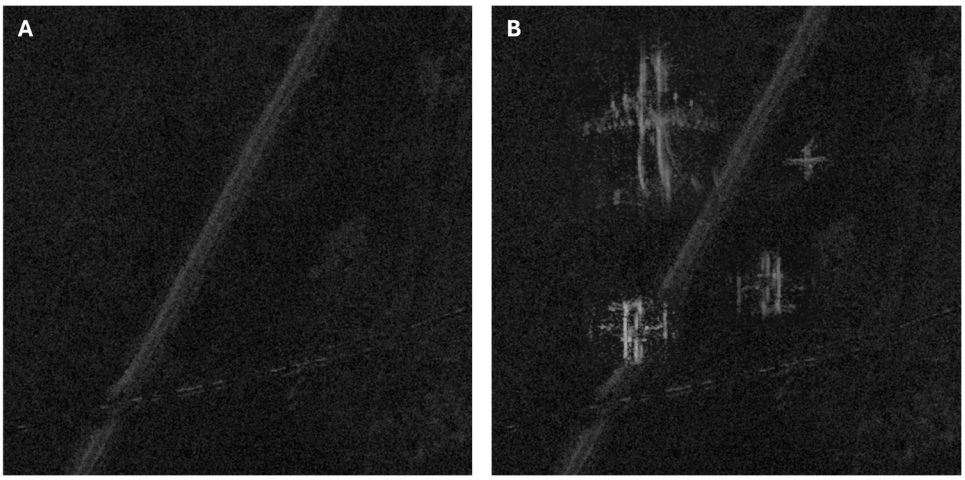


Figure 3. Example of an empty background image (**A**) and an augmented image (**B**) generated by combining the background with synthetic ROIs.

4. Results

As previously noted, the Structural Similarity Index Measure (SSIM) was used to assess the level of image modification introduced by each principal component perturbation method. The values of the selected statistical parameters of SSIM obtained for 2400 synthetic ROIs and various values of β using both proposed strategies are presented in Tables 1 and 2.

Table 1. Statistical parameters of SSIM obtained for the eigenvalue-scaled additive perturbation method.

β	SSIM			
	Min	**Max**	**Mean**	**Std**
0.05	0.764	0.977	0.925	0.030
0.1	0.660	0.970	0.869	0.049
0.15	0.551	0.958	0.806	0.065
0.2	0.463	0.924	0.740	0.077
0.3	0.325	0.904	0.617	0.094
0.4	0.160	0.826	0.514	0.105
0.5	−0.013	0.788	0.432	0.110

Table 2. Statistical parameters of SSIM obtained for the multiplicative perturbation method.

β	SSIM			
	Min	**Max**	**Mean**	**Std**
0.1	0.827	0.984	0.949	0.020
0.2	0.753	0.985	0.935	0.029
0.3	0.668	0.982	0.916	0.039
0.4	0.623	0.972	0.896	0.049
0.5	0.544	0.972	0.873	0.058
0.6	0.527	0.969	0.849	0.067
0.8	0.448	0.961	0.799	0.081
1	0.321	0.948	0.750	0.101
1.5	−0.014	0.962	0.618	0.142

It can be noticed that the higher the value of β, the greater the variability of SSIM and the lower the mean value of this measure for both methods of principal component perturbation. It means that for small values of β all synthetic images are quite similar to their original ROIs, but for relatively large values of β some of the generated images are significantly different from the originals. However, in the first method, the lower value of β is needed to obtain a certain level of modifications in comparison to the second method. It can be noted that for the highest β values, negative values of SSIM were observed in both methods. Mathematically, the SSIM index operates within the [−1, 1] range rather than being limited to positive values. A negative SSIM score indicates that a synthetic image and its original image became locally anti-correlated under the highest β values. This occurs when the modification severely inverts the local structural and luminance characteristics relative to the reference image, resulting in a negative cross-covariance. The question is how different the synthetic images should be from the original ROIs to obtain improvements in the results of the object detection procedure. It seems obvious that they should be neither nearly identical nor extremely different. For both strategies of principal component perturbations, mean SSIM values of approximately 0.75 and 0.87 were achieved; hence, the corresponding β values were selected for further investigation. The examples of synthetic images generated for the selected values of β using both methods are presented in Figures 4–7.

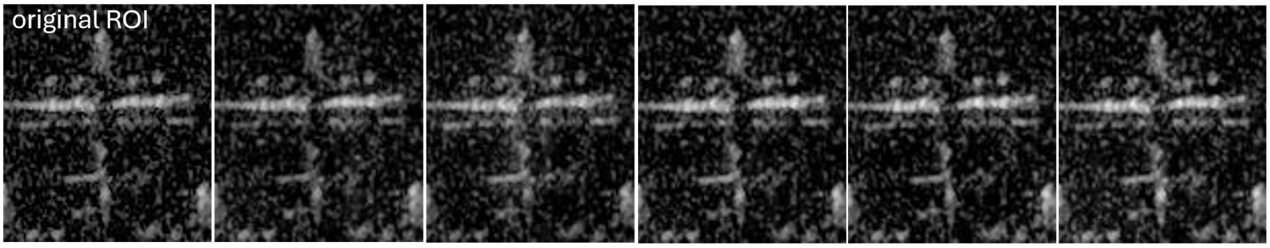


Figure 4. Original ROI and five examples of synthetic images generated with the eigenvalue-scaled additive perturbation method for β = 0.1.

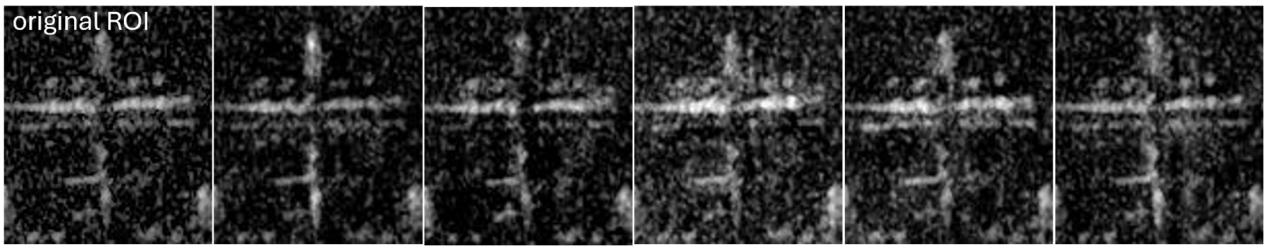


Figure 5. Original ROI and five examples of synthetic images generated with the eigenvalue-scaled additive perturbation method for β = 0.2.

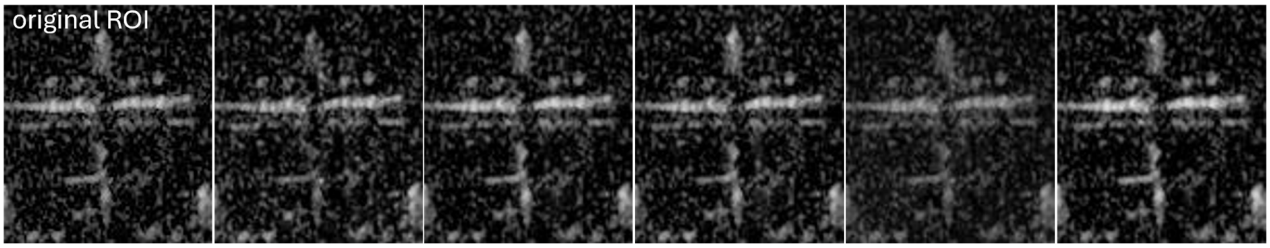


Figure 6. Original ROI and five examples of synthetic images generated with the multiplicative perturbation method for β = 0.5.

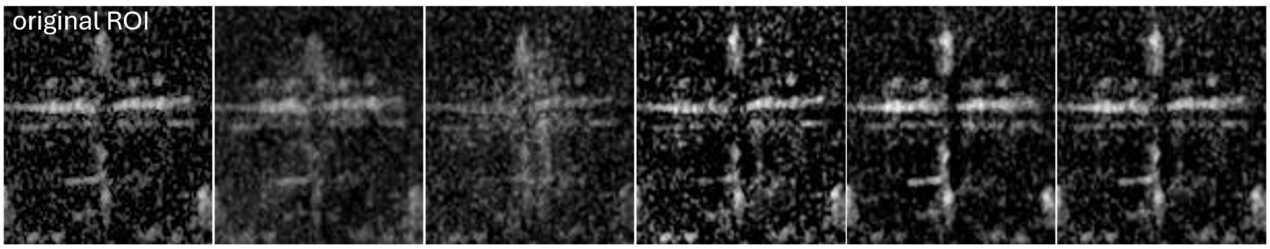


Figure 7. Original ROI and five examples of synthetic images generated with the multiplicative perturbation method for $\beta = 1$.

For the chosen values of β the masking process of the initial principal components was applied. Using transposed data while performing PCA transformation, the number of principal components is equal to the number of processed images. This results in different numbers of principal components for images presenting transport aircraft in different orientations. Therefore, the masking process includes a specific percentage of all principal components, not a fixed number. The selected statistical parameters of SSIM obtained for both discussed methods in masked versions are presented in Tables 3 and 4.

Table 3. Statistical parameters of SSIM obtained for the eigenvalue-scaled additive perturbation method in the masked version.

β	Masking	SSIM			
		Min	Max	Mean	Std
0.1	none	0.660	0.970	0.869	0.049
	20%	0.778	0.974	0.900	0.035
	50%	0.828	0.978	0.924	0.028
	80%	0.872	0.983	0.944	0.021
0.2	none	0.463	0.924	0.740	0.077
	20%	0.605	0.947	0.803	0.062
	50%	0.689	0.961	0.862	0.049
	80%	0.771	0.977	0.915	0.034

Table 4. Statistical parameters of SSIM obtained for the multiplicative perturbation method in the masked version.

β	Masking	SSIM			
		Min	Max	Mean	Std
0.5	none	0.544	0.972	0.873	0.058
	20%	0.695	0.981	0.902	0.043
	50%	0.693	0.983	0.926	0.039
	80%	0.785	0.984	0.946	0.029
1	none	0.321	0.948	0.750	0.101
	20%	0.347	0.971	0.806	0.086
	50%	0.475	0.980	0.865	0.084
	80%	0.507	0.984	0.923	0.065

The results obtained are consistent with expectations. The larger the number of masked initial principal components, the smaller the variability of SSIM and the greater the mean value of this measure for both methods of principal component perturbation, which means that synthetic ROIs are more similar to their origins.

In the next step, numerical experiments with the YOLO8 model in nano version were performed to assess the possibility of using the proposed methods of synthetic image

generation in the data augmentation procedure for training object detectors in SAR images. Each time, the training process of the YOLO model was performed using the default settings. To compare the performance of object detectors, an average precision at an intersection over union threshold of 0.5 (often denoted as AP50) was used. When computing AP50, a detection is counted as correct if its bounding box overlaps with the actual object location with an intersection over union (often denoted as IoU) of at least 0.5.

The reference model of the YOLO-based transport aircraft detector was trained using only genuine SAR images and the default training settings. In this case, the value of AP50 obtained for the test data was equal to 0.888. In further experiments, augmented images accounted for 20% of the training dataset. The validation and test datasets included only genuine SAR images. Each time, the procedure of generating synthetic images and then the training and testing process of the YOLO model were performed five times. The selected statistical parameters of SSIM for the test data are presented in Table 5.

Table 5. Statistical parameters of AP50 obtained for the test set.

Eigenvalue-Scaled Additive Perturbations				
β	**Masking**	**AP50 for the Test Set**		
		Mean	**Std**	**Max**
0.1	none	0.891	0.020	0.914
	20%	0.866	0.020	0.885
	50%	0.897	0.009	0.906
	80%	0.873	0.027	0.899
0.2	none	0.846	0.024	0.874
	20%	0.880	0.034	0.919
	50%	0.884	0.017	0.904
	80%	0.886	0.013	0.904
Multiplicative Perturbations				
β	**Masking**	**AP50 for the Test Set**		
		Mean	**Std**	**Max**
0.5	none	0.885	0.006	0.892
	20%	0.862	0.027	0.892
	50%	0.884	0.017	0.904
	80%	0.880	0.010	0.891
1	none	0.870	0.033	0.902
	20%	0.892	0.018	0.907
	50%	0.877	0.013	0.892
	80%	0.894	0.008	0.906

Taking into account the maximum values of AP50 obtained for the test set, better results were achieved using the eigenvalue-scaled additive perturbation method of principal component modifications. The highest value of AP50 on the test set is equal to 0.919 and was obtained for β = 0.2 with 20% of the initial principal components staying unchanged, and this model was selected as the best one. A slightly worse result (AP50 = 0.914) was obtained for a smaller value of β equal to 0.1 without masking. It can be supposed that the masking procedure is not necessary for relatively low values of β. However, applying this procedure with larger values of β helps to obtain synthetic ROIs that better convey the characteristics of the genuine SAR images, which results in better performance of the YOLO-based transport aircraft detector. The precision-recall curves for the reference model and the best model are presented in Figure 8. A value of AP50 is computed as the area under a PR curve. The increase in AP50 for the best model in comparison with the reference model is equal to 3.1 percentage points. Besides AP50, the

AP50-95 metric was also evaluated, reaching 0.593 for the best model on the test set. Other popular evaluation metrics such as precision and recall are highly dependent on the confidence threshold. For example, using a confidence threshold of 0.5 for the best model on the test set, the precision was 0.904 and the recall was 0.813.

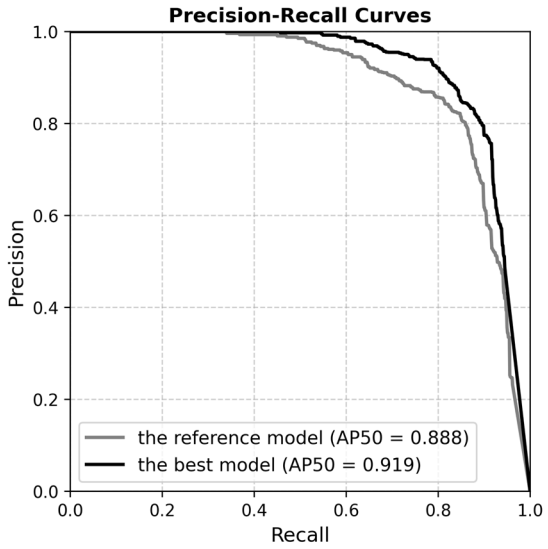


Figure 8. Comparison of PR curves for the reference model and the best YOLO-based transport aircraft detector.

To verify the validity of the proposed method of data augmentation, an ablation study was conducted. First, the training dataset was expanded using background images with the original ROIs of transport aircraft from training images (without PCA-based modifications) inserted with and without alpha blending. In both cases, the detector underperformed compared with the reference model; the values of AP50 for the test data were 0.828 and 0.840, respectively. The observed degradation of the detector's performance can be caused by duplicated (or even multiplied) ROIs or a background distribution mismatch. To exclude the second factor, an additional experiment was conducted by adding empty background images to the training dataset. In this case, the AP50 score for the test set was 0.900, which is higher than that for the reference level. This suggests that background image selection is not the reason for the observed degradation when the detector is trained using background images with the original ROIs as augmentation data.

5. Discussion

In the presented research, a fixed data split was used. Therefore, it is possible that the observed increase in AP50 for the test data using the proposed method of augmentation is highly dependent on the specific set of test images. To verify the statistical significance of the observed increase in AP50 for the test set, the stratified bootstrap procedure was performed [20]. This method enables statistical analysis of the whole population using only a finite sample even if the distribution of the processed data is unknown. In the standard bootstrap procedure, the samples are drawn with replacement from the whole dataset (in this research—from the test set). This means that in a single bootstrap replicate, a specific test image might occur multiple times, while another might not occur at all. Unfortunately, the distribution of the number of transport aircraft in the test images is highly uneven. There are images with only one and more than twenty objects of interest. This can result in either overrepresentation or underrepresentation of target objects in certain bootstrap samples. Therefore, the stratified version of the bootstrap procedure was used. For this purpose, the test images were grouped into five subsets. In each subset, there were

images with similar numbers of transport aircraft. In the stratified bootstrap method, sampling with replacement is performed within each subset. This prevents significant variation in the number of objects between samples.

The statistical analysis was performed using 1000 bootstrap replicates. This means that the evaluation procedures of the best model of the transport aircraft detector trained with augmented data and the reference model were repeated 1000 times. Each time, the values of AP50 were computed for both models using a single bootstrap sample of test images. The further statistical analysis was based on the difference between the values of AP50 obtained for the best model of transport aircraft detector trained with augmented data ($AP50_{aug}$) and for the reference model ($AP50_{ref}$)—Equation (6):

$$\Delta AP50 = AP50_{aug} - AP50_{ref}. \quad (6)$$

The positive value of this difference means that the proposed method of augmentation yields an improvement in transport aircraft detection in SAR images. The nonpositive value of $\Delta AP50$ indicates no improvement or even degradation of the object detection model. The histogram illustrating the bootstrap distribution of $\Delta AP50$ is presented in Figure 9. The vertical dotted line corresponds to the 0.05 quantile (denoted by x_α), which determines the lower bound of the one-sided 95% confidence interval. The 0.05 quantile is equal to 0.0129, which is a positive value. Therefore, it can be inferred that the presented method of augmentation yields an improvement in the performance of the YOLO-based detector of transport aircraft in SAR images.

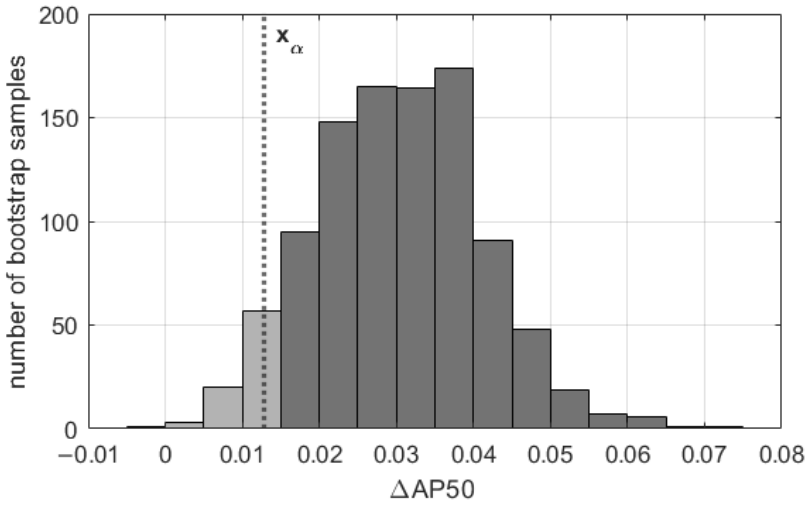


Figure 9. Distribution of ΔAP50 obtained from 1000 stratified bootstrap samples. The vertical dotted line corresponds to the 0.05 quantile, which determines the lower bound of the one-sided 95% confidence interval. Light gray bars indicate the histogram bins that fall outside the confidence interval, while dark gray bars represent those contained within it.

Besides the 0.05 quantile, some more statistical parameters of the $\Delta AP50$ bootstrap distribution were estimated. Their values are presented in Table 6. It can be noticed that the values of the median and the mean of $\Delta AP50$ are quite similar and consistent with the increase in AP50 observed in Figure 7, which was equal to 3.1 percentage points. Moreover, the bootstrap probability of nonpositive values of $\Delta AP50$ is very low and equals 0.1%.

Table 6. Statistical parameters of ΔAP50 estimated by the stratified bootstrap method.

Parameter	Value
0.05 quantile of $\Delta AP50$	0.0129
median of $\Delta AP50$	0.0305

mean of $\Delta AP50$	0.0304
$P(\Delta AP50 \leq 0)$	0.10%

The improvement in AP50 achieved in this study, amounting to several percentage points, is comparable in magnitude to the performance gains reported in other recent studies on synthetic SAR data generation discussed in Section 2. However, a direct comparison of the improvements in individual evaluation metrics is limited by differences in the datasets used, object classes, downstream tasks, and evaluation measures. Therefore, Table 7 provides a broader comparison of the proposed method with selected recent approaches, considering not only the resulting performance improvement in detection or classification tasks, but also the input data requirements, the need to train additional generative models, the computational complexity of the generation process, the degree of control over the generated samples, and the approach used to assess the quality of synthetic data.

Table 7. Comparison of the proposed method with selected recent approaches to synthetic SAR data generation.

Reference	**[12] Wu, B. et al., 2024**	**[13] Song, Q. et al., 2022**	**[14] Jin, R. et al., 2024**	**[15] Han, D. et al., 2025**	**This paper**
Separate generative model training	yes (deep GAN training)	yes (adversarial autoencoder training)	yes (RC-CapsNet + reconstruction network training)	yes (Stable Diffusion fine-tuning + ControlNet training)	no (PCA basis computations only)
Training data requirement	large set of scenes from two different domains (optical and SAR)	cropped SAR object images covering multiple viewing orientations	cropped SAR object images acquired for multiple azimuth and elevation angles	large set of SAR scenes with object annotations	limited set of cropped SAR object images
Generation mechanism	optical-to-SAR image translation	conditional generation from latent vector, class, orientation and semantic map	perturbation of learned attribute features followed by neural reconstruction	latent diffusion conditioned spatially by ControlNet	perturbation of principal components followed by inverse PCA transformation
Orientation handling of generated data	inherited from the optical source image	explicit orientation conditioning	indirect orientation control via learned latent features	no explicit orientation control	explicit grouping into 8 predefined orientation classes
Generated data unit	complete SAR scene with annotation	object SAR image (ROI)	object SAR image (ROI)	complete SAR scene with annotation	object ROI inserted into real SAR background
Target objects	ships	ground vehicles (military)	ground vehicles (civilian)	ships	transport aircraft
Main source of complexity	two imaging domains and deep translation model	rich conditional information and semantic-map preparation	rich multi-geometry dataset and learned feature decomposition	large annotated SAR dataset and complex pretrained generative architecture	manual target selection and orientation grouping rather than complex model training
Downstream validation	object detection by YOLOv7	classification by A-ConvNet	not applicable	object detection by Mask R-CNN	object detection by YOLOv8

Reported downstream effect	AP improvement (best 0.703)	accuracy improvement (best 0.934)	not evaluated	mAP50 improvement (best 0.831)	AP50 improvement (best 0.919)
Synthetic data assessment	visual and direct downstream detection	similarity assessment using NCC, MGSM, NGSC	quality assessment using PSNR, RL, DR	visual and direct downstream detection	SSIM-based similarity/diversity control

As shown in Table 7, the proposed PCA-based approach presented in this paper differs from the analyzed methods primarily in its relatively low data requirements and the absence of the need to train a separate, complex generative model. This simplification, however, comes at the cost of partially manual input data preparation, including the selection of representative object images and their assignment to predefined orientation groups. At the same time, this data organization provides direct control over the set of object variants used to construct the individual PCA bases and, when combined with explicit perturbation of the principal components, allows the degree of diversity in the generated samples to be controlled. An additional element of control is provided by the use of SSIM to assess the trade-off between the similarity of the generated samples to real data and the degree of introduced variation.

6. Conclusions

The presented results indicate that the proposed PCA-based method for the synthetic generation of SAR object images can be effectively used as a data augmentation technique for training object detectors. New images are generated by perturbing principal components and subsequently reconstructing the image using the PCA projection matrix derived from real SAR imagery. As a result, the generated samples are constructed within a subspace spanned by eigenvectors extracted from real SAR images and retain some of their statistical properties, despite the fact that the method does not directly model the physical radar image formation process. At the same time, it should be emphasized that reconstruction based on real eigenvectors alone does not guarantee the physical validity of every generated sample, particularly for large perturbation magnitudes.

For this reason, the proposed procedure employs the Structural Similarity Index Measure (SSIM) as a surrogate, computationally inexpensive indicator of the degree of modification of synthetic images relative to the source data. SSIM was not treated as a measure of the physical validity of SAR imagery; rather, it was used as a tool for supporting the selection of suitable parameter ranges for the generation process prior to the computationally intensive detector training stage. Two levels of average structural similarity, corresponding to SSIM values of approximately 0.75 and 0.87, were selected for further investigation. For both perturbation strategies considered, namely eigenvalue-scaled additive perturbation and multiplicative perturbation, an additional masking mechanism was introduced by preserving 20%, 50%, or 80% of the principal components associated with the largest eigenvalues. As expected, increasing the proportion of frozen coefficients resulted in greater similarity between the generated images and their corresponding source samples.

The resulting synthetic data variants were then independently used to augment the training set of a YOLO detector, while the same real-world test set was maintained across all experiments. This made it possible to evaluate which combination of perturbation strategy, structural similarity level, and principal component masking yielded the most effective augmentation data. The best performance was achieved using eigenvalue-scaled additive perturbation with 20% of the initial principal components kept unchanged. Extending the training dataset with 20% additional images containing synthetic objects embedded into real SAR backgrounds resulted in an AP50 increase of 3.1 percentage points

compared with the baseline model trained exclusively on real data. The statistical significance of this improvement observed on the test set under a fixed split was confirmed using a stratified bootstrap procedure.

The obtained results therefore suggest that SSIM can serve as a practical indicator for supporting the calibration of synthetic data generation processes. This does not imply, however, that higher SSIM values are always preferable or that a particular SSIM value guarantees the physical realism of a generated sample. Excessively high similarity may limit the diversity introduced by augmentation, whereas excessively low similarity may correspond to overly strong image modifications. Rather, the results indicate that SSIM can be used to narrow the parameter space of the generator and reduce the number of computationally expensive experiments requiring detector retraining.

The present study considered a single object class, namely transport aircraft. Therefore, the identified SSIM ranges and perturbation parameters should not be regarded as universally applicable to all SAR data classes, as the relationship between structural similarity and augmentation effectiveness may depend on the characteristics of a particular object class. In addition, the proposed method remains sensitive to object orientation. Future work should therefore include an evaluation of the procedure on other SAR datasets and object classes, an investigation of the transferability of SSIM ranges associated with beneficial augmentation effects, and the development of a version of the method that is less dependent on object orientation.

Author Contributions: Conceptualization, J.J. and J.C.; methodology, J.J. and J.C.; software J.J. and J.C.; validation, J.J.; formal analysis, J.C.; investigation, J.C.; data curation, K.K.; writing—original draft preparation, J.C. and J.J.; writing—review and editing, J.J., K.K. and D.W.; visualization, J.C.; supervision, J.J. and D.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Military University of Technology, Warsaw, Poland, co-financed by the Faculty of Electronics (grant no. EL_UGBWEL_12022026_09) and the Faculty of Civil Engineering and Geodesy (grant no. 531-000105-W400-22). This publication was co-financed from the budget funds awarded by the Polish Ministry of Science and Higher Education under the Vectors of Science Programme, Project No WNK/SP/0411/2025/01.

Data Availability Statement: The satellite imagery used in this study was obtained from commercial data providers under licensing terms that do not permit redistribution of the original data. Therefore, the datasets supporting the findings of this study cannot be made publicly available by the authors. Access to the original satellite imagery may be obtained directly from the respective data providers, subject to their licensing conditions.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

SAR	Synthetic Aperture Radar
PCA	Principal Component Analysis
ROI	Region Of Interest
AP50	Average Precision at intersection-over-union threshold of 0.5
IoU	Intersection over Union
AAE	Adversarial Autoencoder
GAN	Generative Adversarial Network
YOLO	You Only Look Once
SSIM	Structural Similarity Index Measure

References

1. Ramdani, F. An extensive review of SAR remote sensing in mode transportation studies: Major findings and future recommendations. *F1000Research* **2025**, *14*, 322. https://doi.org/10.12688/f1000research.160735.1.
2. Moreira, A.; Prats-Iraola, P.; Younis, M.; Krieger, G.; Hajnsek, I.; Papathanassiou, K.P. A tutorial on synthetic aperture radar. *IEEE Geosci. Remote Sens. Mag.* **2013**, *1*, 6–43. https://doi.org/10.1109/MGRS.2013.2248301.
3. Zhao, H. Target recognition of SAR images based on SVM and KSRC. *Comput. Intell. Neurosci.* **2021**, *2021*, 4322678. https://doi.org/10.1155/2021/4322678.
4. Yang, S.; Xiao, W.; Zhang, M.; Guo, S.; Zhao, J.; Shen, F. Image Data Augmentation for Deep Learning: A Survey. *arXiv* **2023**, arXiv:2204.08610v2. https://doi.org/10.48550/arXiv.2204.08610.
5. Jakubowski, J.; Chmielińska, J.; Karwowska, K.; Wierzbicki, D. Using PCA Transformation to Generate Object Images in Satellite SAR Scenes. In Proceedings of the 2026 Progress in Applied Electrical Engineering (PAEE), Kościelisko, Poland, 22–26 June 2026. https://doi.org/10.1109/PAEE71887.2026.11660700.
6. 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. https://doi.org/10.1109/TIP.2003.819861.
7. Deng, W.; Li, K.; Zhao, H. A Flight Arrival Time Prediction Method Based on Cluster Clustering-Based Modular with Deep Neural Network. *IEEE Trans. Intell. Transp. Syst.* **2023**, *25*, 6238–6247. https://doi.org/10.1109/TITS.2023.3338251.
8. Huang, C.; Peng, Y.; Deng, W. A dendrite net learning multi-objective artificial bee colony algorithm for UAV. *Appl. Soft Comput.* **2026**, *189*, 114449. https://doi.org/10.1016/j.asoc.2025.114449.
9. Shahnewaz, T.; Sirakov, N. Training Data Augmentation with Data Distilled by Principal Component Analysis. *Electronics* **2024**, *13*, 282. https://doi.org/10.3390/electronics13020282.
10. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*; Pereira, F., Burges, C.J., Bottou, L., Weinberger, K.Q., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2012; Volume 25, pp. 1097–1105. https://doi.org/10.1145/3065386.
11. Wang, Z.; Zhang, Z.; Shan, X.; Wei, H.-A.; Tang, P. Generative models for SAR–optical image translation: A systematic review. *Int. J. Appl. Earth Obs. Geoinf.* **2025**, *146*, 105009. https://doi.org/10.1016/j.jag.2025.105009.
12. Wu, B.; Wang, H.; Zhang, C.; Chen, J. Optical-to-SAR translation based on CDA-GAN for High-Quality Training sample generation for ship detection in SAR amplitude images. *Remote Sens.* **2024**, *16*, 3001. https://doi.org/10.3390/rs16163001.
13. Song, Q.; Xu, F.; Zhu, X.X.; Jin, Y.-Q. Learning to Generate SAR Images with Adversarial Autoencoder. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5210015. https://doi.org/10.1109/TGRS.2021.3086817.
14. Jin, R.; Cheng, J.; Wang, W.; Zhang, H.; Zhang, J. Attribute Feature Perturbation-Based Augmentation of SAR Target Data. *Sensors* **2024**, *24*, 5006. https://doi.org/10.3390/s24155006.
15. Han, D.; Kim, Y.; Park, J.; Yoon, D.; Jeon, M. A Diffusion-Based Data Augmentation Method for SAR Object Detection. *J. Korean Inst. Inf. Technol.* **2025**, *23*, 149–155. https://doi.org/10.14801/jkiit.2025.23.1.149.
16. Geng, Z.; Xu, Y.; Wang, B.-N.; Yu, X.; Zhu, D.-Y.; Zhang, G. Target Recognition in SAR Images by Deep Learning with Training Data Augmentation. *Sensors* **2023**, *23*, 941. https://doi.org/10.3390/s23020941.
17. Jolliffe, I.T. *Principal Component Analysis*, 2nd ed.; Springer-Verlag: New York, NY, USA, 2002.
18. Turk, M.; Pentland, A. Eigenfaces for Recognition. *J. Cogn. Neurosci.* **1991**, *3*, 71–86. https://doi.org/10.1162/jocn.1991.3.1.71.
19. Jakubowski, J.; Chmielińska, J.; Karwowska, K.; Wierzbicki, D. Hybrid Data Augmentation Using Principal Component Analysis in Object Detection on Satellite SAR Imagery. *Metrol. Meas. Syst.* 2026, *accepted.*
20. Efron, B.; Tibshirani, R.J. *An Introduction to the Bootstrap*; Chapman & Hall/CRC: New York, NY, USA, 1993.

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.