Physics-inspired data augmentation for SAR ATR: a new approach to tackle the synthetic-to-measured Domain Gap

Elisa Delhommé
Thales LAS France SAS
Elancourt, France
0009-0007-9501-2117

Héloïse Remusati
Thales LAS France SAS
Elancourt, France
0000-0002-5528-5491

Caroline Lesueur
Thales LAS France SAS
Elancourt, France
0009-0000-8557-7180

Jacques Petit-Frère
Thales LAS France SAS
Elancourt, France
0009-0008-8384-5581

***Abstract*—Synthetic Aperture Radar (SAR) imaging supports applications from environmental monitoring to defence. For Automatic Target Recognition (ATR), Deep Learning (DL) delivers strong results but needs large datasets, and SAR data is scarce due to cost and confidentiality. A common workaround is training on synthetic data generated by simulators and Computer-Aided Design (CAD) models, but these simplify complex electromagnetic effects, creating a domain shift between training (synthetic) and test (measured) domains. Although Data Augmentation (DA) is used to improve representativeness and robustness, many methods lack semantic, physics-informed changes to improve recognition performance. In this paper, we propose a physics-based DA to address the Synthetic-to-Measured (S2M) gap, first validating physical parameter extraction from measured images and then leveraging this knowledge to improve ATR. Training solely on synthetic data, our approach achieves 70.97% accuracy.**

***Index Terms*—ATR, classification, synthetic data, MOCEM, SAR, data augmentation, ASC, deep learning, MSTAR.**

I. Introduction

Synthetic Aperture Radar (SAR) Automatic Target Recognition (ATR) is often used in many real-world applications where object identification is required via a classification system. Deep learning-based SAR ATR has achieved increasingly high performance in the past few years, particularly on the open Moving and Stationary Target Acquisition and Recognition (MSTAR) dataset under standard operating conditions (SOCs). As a matter of fact, deep neural networks have proven indispensable for tackling the complex task of classifying SAR images due to the multiple electromagnetic effects they exhibit.

Most of the studies generally show cases where the test conditions are close to the learning conditions using measured data from neighbouring depression angles (15° versus 17° most of the time). This case, although effective, is rarely adapted to real operational conditions where exact target information is rare or even unavailable, and where learning on synthetic data is recommended (or even required). However, synthetic data are based on simulators that cannot accurately reproduce all the effects present in the measured data and cannot cover all the variability found under real operating conditions. As a result, models trained on these data often have limited generalisation and robustness potential in the real world. This is why data augmentation (DA) techniques are essential, whether to bring synthetic and real data closer together, to present numerous variations to the model, or simply to expand the training database. SAR images although have unique physical properties that set them apart from optical images, for which most data augmentation techniques in the literature are designed and which do not always make physical sense for this type of data (for example, rotations or affine transformations). Techniques designed specifically for the optical domain may raise questions, such as:

When does a target cease to be the same? While this question is more obvious in optics, it is more difficult to answer when dealing with a SAR image to which transformations are to be applied. Some transformations, such as affine transformations, can change the target shape.

How to ensure that the applied transformations are sufficient to cover the target domain? This kind of augmentation techniques can improve the extraction of features of interest for classification and provide generalisation ability for data exhibiting the same type of variability. However, this does not guarantee robustness to other variations such as pixel level modification or target state variation (open doors...).

This paper therefore introduces a data augmentation technique that addresses three issues:

- Improving recognition performance while learning only from synthetic data,
- Providing a new data augmentation technique for SAR images,
- Introducing an artificial augmentation that resorts to physical mechanisms and that is interpretable.

The proposed SAR augmentation process draws on physical knowledge based on scattering mechanisms. The objective of this method is to confront synthetic training databases with real measurements in order to achieve better performance and robustness of SAR ATR algorithms. We believe that the use of physics-based transformations can improve the guarantee that the algorithm will work in real-world conditions and cover the domain of use, by providing interpretable augmented samples. The paper is presented as follows. Section II introduces background and related works. Section III brings in our approach

and presents the methodology adopted for carrying out this work. Section IV shows the results of the data analysis and details the results obtained for SAR ATR on the MSTAR dataset. Section V addresses the limitations of the proposed method. Finally, Section VI concludes the paper and discusses future work.

II. Related works

When training machine learning algorithms, it is typically assumed that the data distribution of the training and test sets is consistent. Nevertheless, this assumption often fails to hold in practical applications and in particular when training is performed on synthetic data [3]. Indeed, even though simulation models are particularly efficient in providing physically accurate SAR images, they are only simplified representations of the complex mechanisms that form the signature of targets in images [2], leading to residual discrepancies between measured and synthetic images. For example, Fig. 1 illustrates a visible domain gap between synthetic and real images.

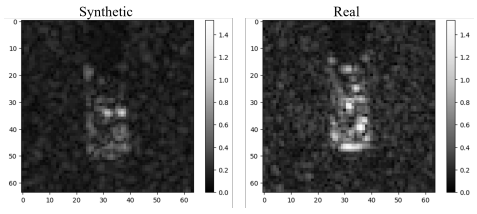

Fig. 1. Domain Gap Example: 2S1 target (azimuth 358°). Images are displayed with the same colour range with the QPM LUT.

Domain generalisation and domain adaptation are areas of research aiming at alleviating this issue, often referred to as domain shift or gap. To do so, several strategies have been proposed in the literature to make models able to generalise, to an unseen domain (also called target domain – measured samples in our case), using a known domain (also called source domain – synthetic samples in our case). They can be primarily categorised into three types [3], being domain alignment, meta-learning, and data augmentation, but we are going to focus on the last one in this paper. More specifically, data augmentation is the process of artificially generating new data from existing data.

Particularly for the SAR S2M domain gap, we can name two different (non-exhaustive) data augmentation strategies:

Image-level augmentation. According to [3], it creates new examples similar to the original images but does not explicitly focus on the concept of domain by performing a series of transformations or noise addition to the original images in the training set. For example, in [5], Inkawhich et al. tested and compared several data augmentation techniques such as rotation and the addition of Gaussian noise and analysed their impact on saliency and feature-space representation. Similarly, in his PhD work, Denton [6] proposed an alpha blending approach to generate new data points by combining SAMPLE extracted targets with MSTAR clutter samples.

Domain-level augmentation. According to [3], this type of data augmentation method tries to increase the breadth of the training domains and to cover the unseen target domain as comprehensively as possible by generating a large number of new samples with diverse distributions. For example, the authors of [4] directly worked on synthetic and measured images in order to bring them closer. They observed that the synthetic and measured images in the SAMPLE dataset are linearly separable and utilise a linear content erasure method (LCDE) in order to transform the images and eliminate this separability. Differently, Camus et al. in [2] proposed a domain randomisation technique, in order to introduce randomness and variations into samples within a simulated environment, combined with adversarial training. More recently, they proposed in [12] an even more advanced image generation pipeline, by combining their ADASCA block, presented in [2], with two other kinds of domain-level augmentation: semantic and radar augmentations using, respectively, their ADAMO and M3D Exploit tools.

Similarly to our approach, these different works try to bridge the domain gap between synthetic and measured SAR images. Nevertheless, we realise that most research papers:

1) either try to characterise and bridge the domain gap by using learnt transformations (for example [4]), but do not provide physical comprehension of them (for instance, synthetic and measured samples exhibit specific target signature differences),
2) or produce new samples without any guidance, hoping that the augmented database will cover the target domain as we assume that more diverse data can help reduce generalization error (for example [12]).

As a consequence, and as stated in [3], we believe that addressing the following question is of utmost importance: "which data augmentation method is most useful in the current task and why". Indeed, currently, most data augmentation methods create samples that are not interpretable, such as adversarial perturbation and adding noise, or without being able to justify why they will be beneficial. With our approach, we propose to resort to physical knowledge provided by Attributed Scattering Centres (ASC) in order to both bridge the gap between synthetic and measured samples and physically characterise it. We believe that our approach, combined with other DA techniques, such as the one of SCALIAN DS [12], can improve even more recognition performance.

By passing, we found some papers in the literature that also used attributed scattering centres for data augmentation (for instance [7], [8]), but they only consider the measured-measured configuration and do not deal with the synthetic-measured gap. To our knowledge, this is the first work to attempt to benefit from ASC to characterise and bridge the synthetic-to-measured domain gap.

III. Methodology

Our approach uses a physical modelling of SAR images, provided by the Attributed Scattering Centres Model. Thanks to it, it is possible to describe both synthetic and measured

images and to use this physical knowledge to generate augmented samples. Even though similar to the pipeline proposed by SCALIAN DS in [12], our modelling and proposed DA technique differs on two aspects: physical parameters can be estimated from both synthetic and measured images, without being adherent to MOCEM or CAD models, and can thus provide an interpretable feature space shared between the two domains, providing hints about their discrepancies.

A. Our approach

1) Attributed Scattering Centres (ASC) model: Proposed by Gerry et al. [9] in 1997, the Attributed Scattering Centres (ASC) model is based on the geometrical theory of diffraction (GTD) and physical optics (PO) and enables one to accurately model the scattering of a target. This model deals with both localized (the scatterer appears to exist at a single point in space) and distributed (the scatterer, in the imaging plane, appears as a finite, non-zero-length current distribution) scattering mechanisms, and characterises them with a set of several parameters corresponding to:

- Frequency and aspect dependence,
- Physical attributes (such as the structure, location, orientation, geometry, size).

Formally, the ASCM describes the total scattered field as a function of frequency f and aspect angle ϕ as:

$$E(f,\phi;\Theta_N) = \sum_{i=1}^{N} E_i(f,\phi;\theta_i) \tag{1}$$

Where $\Theta_N = \{\theta_i | \theta_i = [A_i, x_i, y_i, \alpha_i, \gamma_i, L_i, \bar{\phi}_i], 1 \leq i \leq N\}$ is the parameter set of N individual scatterers and

$$\begin{aligned} E_i(f,\phi;\theta_i) =& A_i \cdot (j\frac{f}{f_c})^{\alpha_i} \\ & \cdot \exp(-j\frac{4\pi f}{c}(x_i\cos\phi + y_i\sin\phi)) \\ & \cdot \text{sinc}(2\pi f c L_i \sin(\phi - \bar{\phi}_i)) \\ & \cdot \exp(-2\pi f \gamma_i \sin\phi) \end{aligned} \tag{2}$$

With each term $E_i(f,\phi;\theta_i)$ representing the backscatter from a single scattering mechanism, f_c the centre frequency of radar wave, c the velocity of light.

The parameter set of the ith scatterer has physical interpretations related to the location and the geometry of the scatterer. Specifically, A_i represents the relative amplitude of the measured field ($A_i \in \mathbb{C}^2$), x_i and y_i correspond, respectively, to range and cross-range locations (in meters), $\alpha_i \in \{-1, -0.5, 0, 0.5, 1\}$ models the frequency dependence, γ_i describes the mild aspect dependence of localized scattering centre cross section, L_i models the length of the scattering centre and $\bar{\phi}_i$ models the orientation angle with respect to the broadside.

This model has the advantage of being able to represent several types of scattering centres depending on the different values the parameters take, as presented in Table I.

TABLE I
GEOMETRIC SCATTERING TYPES DIFFERENTIATED BY FREQUENCY AND ASPECT DEPENDENCE.

	Geometric Scattering Type	α	γ	L	ϕ
Localized	Trihedral	1	> 0	0	0
	Top Hat	0.5			
	Sphere	0			
	Corner Diffraction	-1			
Distributed	Dihedral	1	0	> 0	$\neq 0$
	Cylinder	0.5			
	Edge Broadside	0			
	Edge diffraction	-0.5			

2) Extracting ASC parameters: Given a SAR image $D(x,y)$, the objective is to find the set of parameters $\Theta_N = [\theta_0, ..., \theta_N]$ that best fit the N scattering mechanisms present in the object that is imaged.

To do so, we resort to the minimisation of a cost function that represents the difference between the model and the actual response in the image domain, as scattering centre responses are isolated in the image domain when data is gathered at high frequencies. Consequently, our problem can be expressed as:

$$\hat{\Theta}_N = \arg\min_{\Theta_N} \left\| D(x,y) - \tilde{D}(x,y;\Theta_N) \right\|_2 \tag{3}$$

where Θ_N are the scattering parameters, $D(x,y)$ the SAR image and $D(x,y) - \tilde{D}(x,y;\Theta_N)$ the reconstructed image via the ASC model.

First, as with every numerical optimisation process, we initialise a set of values before starting the iterations. Next, we use a gradient-based optimisation to numerically obtain the set of parameter values that minimise the cost function. Choosing the order of the problem, here N, that corresponds to the number of scatterers, is really difficult. For simplicity, we fix it to a upper bound, the value of which we will examine in IV-B1b. It is worth noting that this number of scatterers will change, depending on the considered target, azimuth angle and database.

3) Reconstructing images from ASC parameters: Given the set of ASC parameters, it is next possible to generate the SAR signal in the frequency-aspect angle domain using the equation 1. In order to retrieve the Cartesian domain data, it is then needed to resample to a uniform grid on coordinates (f_x, f_y) expressed as $f_x = f\ \cos(\phi)$ and $f_y = f\ \sin(\phi)$.

This resampling gives us an $M \times P$ array depending on f_x and f_y in the frequency domain. In order to be comparable with SAR images, we need to convert these data into the image domain. To do so, the data are first multiplied by a Taylor window (with a -35 dB side lobe level), and is next zero-padded to a new size of $M_z \times P_z$ where $M_z = 1.5M$ and $P_z = 1.5N$. Finally, we get the SAR image $D(x,y)$ thanks to a two-dimensional inverse Fourier (2D-IFFT). It is thus possible to compare the images generated with the ASCM with measured samples.

4) Using ASC for data augmentation:

a) Training pipeline: In the context of SAR ATR, our approach uses the ASC model to obtain a representation in

the ASC domain of sample images. In this way, the data augmentation process can be carried out in a domain ancillary to SAR images, which is more physical and allows for better control of the transformations performed.

The training pipeline relies on the process presented in Fig. 2:

1) Starting from a synthetic training dataset, the ASC parameters are extracted on each image separately, obtaining an ASC training dataset.
2) Augmentations and transformations are then optionally performed on the ASC parameters (red block).
3) In order to train the model of interest on SAR images, the reconstruction step is necessary to convert the ASC parameters back into SAR images.

We then obtain a synthetic reconstructed training dataset. An example of reconstruction is shown in Fig. 3. In this new dataset, only the targets are reconstructed (Step 3), resulting in a zero background, which is not consistent with the real test data set.

For a matter a time, we first overcome this issue and compensate for reconstruction defects by averaging each reconstructed sample with its original version (Step 4). We emphasise the fact that it is only a temporary measure until we have implemented a better ASC parameter estimation algorithm and found a process to add background information in reconstructed images. Finally, as is traditionally the case in deep learning processes, the image is standardised: the model is trained with the Quarter Power Magnitude (QPM) LUT on the images, that are then normalised and cropped to a 64×64 format. These images are then used to train our model, which is a ResNet18. For information, we trained it for 500 epochs, with early stopping (that happened at epoch 257), with batches of 32 samples. We used an Adam optimiser with a learning rate of 0.001.

b) Transformation types: As mentioned earlier, our data augmentation method do not apply directly to the training samples, in this case SAR images, but to the ASC parameters. All augmentation types are then performed on ASC parameters data type, before the reconstruction, i.e. between stages 2 and 3 in Fig. 3. Three families of transformation are used:

- Adding Gaussian noise to the parameters, on the reflectors position or amplitude for example.
- Adding new reflectors:
  - defined (type of effects and amplitude) randomly or similar to the target.
  - placed randomly or on the target area.
- Removing random reflectors.

Fig. 4 illustrates the different types of transformation in the resulting SAR reconstruction (Step 3) in Fig. 3). The SAR reconstruction from ASC parameters is also performed on the original sample without any transformation for comparison purposes.

In this paper, we then refer to data augmentation techniques applied directly to SAR images as *image augmentations* and those applied in the domain of ASC parameters as *ASC augmentations*.

B. Study of the ASC modelling

Our first experiments aim at evaluating the quality of our ASC parameter extraction algorithm and to analyse the sets of parameters extracted from our datasets.

1) Quality evaluation of the ASC parameters extraction algorithm: The purpose of this study is to analyse the impact of ASC extraction hyper-parameters on the reconstructions, and thus select the most effective set of values for the experiments. The extraction of ASC parameters is considered effective if the reconstruction of the SAR image is close to the original image. The impact of both the number of iterations and the number of scattering centres will be evaluated.

2) Analysis of estimated ASC parameters: Analysing the ASC parameters not only allows trends in the data to be identified, but also any gaps in the data to be determined. To do this, several aspects are studied:

- The number of parameters overall and per effect according to the type of target, but also the angular sectors,
- The comparison of the parameters obtained from real and synthetic data,
- The effectiveness of extraction on both types of data.

3) Metrics: To perform the experiments presented in the previous section, we will use the following metrics.

a) Image comparison: For comparing an original image I_O and its reconstruction I_R, we use the Image Correlation Coefficient metric (ICC):

$$ICC(I_{\mathrm{O}}, I_{\mathrm{R}}) = \frac{\sum_m \sum_n \left[I_{\mathrm{O}}(m,n) - \bar{I}_{\mathrm{O}}\right]\left[I_{\mathrm{R}}(m,n) - \bar{I}_{\mathrm{R}}\right]}{\sqrt{\sum_m \sum_n \left[I_{\mathrm{O}}(m,n) - \bar{I}_{\mathrm{O}}\right]^2 \sum_m \sum_n \left[I_{\mathrm{R}}(m,n) - \bar{I}_{\mathrm{R}}\right]^2}} \quad (4)$$

We also computed other metrics but, due to space considerations, we only report ICC values that appear to be fairly aligned with the human eye.

b) ASC parameters analysis: **Heatmaps.** Heatmaps are used to show data depending on two independent variables as a colour coded image plot. In this case, they represent a number of elements associated with a colour bar.
Violin plots. Violin plots are used to compare data distributions in a similar way to box plots, but with the addition of probability density information.

C. ASC data augmentation study for SAR ATR

The focus of the following experiments is on the benefits of our DA technique regarding recognition performance. In order to quantify the contribution of the proposed method for SAR ATR, several test configurations are carried out. We gradually add different ASC augmentation techniques, starting from baselines. All tested configurations are listed in Table II. Each of the three baselines has its own objective:

- The *Baseline* allows to define a starting point for performance on original SAR images, without resorting to any data augmentation.

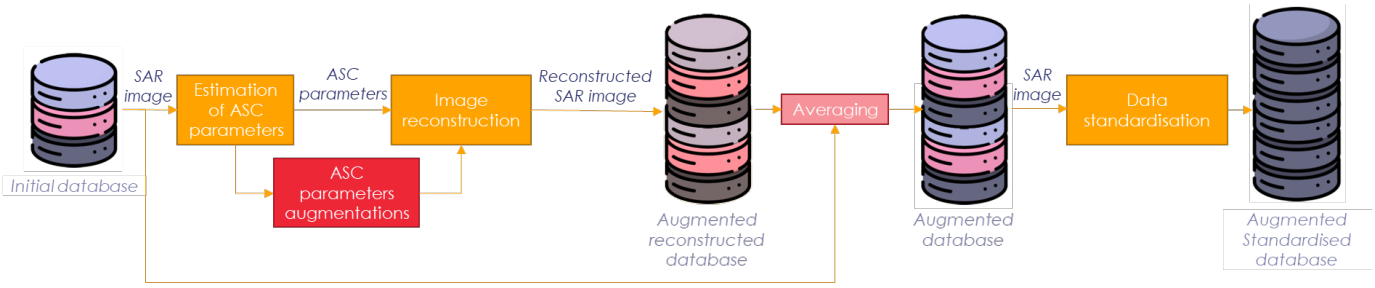

Fig. 2. Our data augmentation pipeline (ASC augmentation).

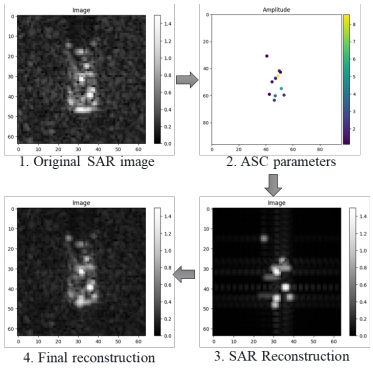

Fig. 3. Example of an ASC extraction and reconstruction on a 2S1 real target. Images are displayed with the same colour range with the QPM LUT.

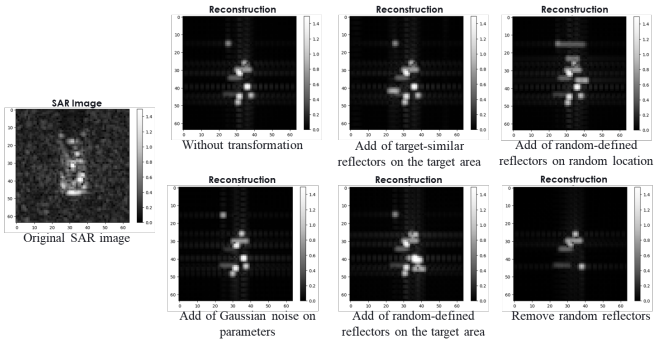

Fig. 4. The different transformations effects on the reconstructed SAR image on a real 2S1 example.
Images are displayed with the same colour range with the QPM LUT.

- The *Augmented Baseline* allows to quantify the final gain obtained by adding traditional image data augmentation techniques.
- The *Baseline ASC* represents a point of comparison for each ASC-augmented test training via SAR reconstructed images, without resorting to any data augmentation.

In order to confront our work to the state of the art, some test configurations use traditional SAR image augmentation techniques, inspired from and detailed in [1]: Gaussian noise addition, colour jittering and random area erasing. We also add random high pixels dropout and value shifting. What's more, all configurations use circular shifts (x and y offsets) on the synthetic training images (original and reconstructed) for the significant performance gain it provides, as mentioned in [2]. The study is evaluated thanks to traditional classification metrics Accuracy and F1-score.

IV. Experimental results

A. Data under study

1) Synthetic MOCEM MSTAR: The training set is based on a synthetic database generated using the MOCEM software, which is a CAD-based SAR imaging simulator developed by SCALIAN DS for the French MoD (DGA) [11]. Synthetic samples were directly provided by the French MoD who reproduced the MSTAR dataset. They used one CAD model per MSTAR class and the radar parameters provided by the MSTAR metadata. They ran parametric simulations for the depression angle of 15° and generated images at every 1° azimuth for the full [0°, 360°[range. Thus, they did not consider exactly the same azimuth angles than MSTAR. These 3600 images were randomly split into train/validation sets using a 75%/25% repartition for the first two baselines. For the *Baseline ASC* and TSM1 to TSM10, the corresponding ASC-based reconstructions are added to these splits, leading to a total of 7200 images. Special attention is paid to ensuring that both the original image and its reconstruction are included in either the training set or the validation set. Addition is considered instead of replacement to still take benefit from the initial physical information provided by MOCEM.

2) MSTAR: The test set is the measured MSTAR dataset, collected by the Sandia National Laboratory SAR sensor platform and sponsored by Defence Advanced Research Projects Agency and Air Force Research Laboratory [10]. The MSTAR dataset allows the benchmark for experiments related to problems dealing with SAR images. Each MSTAR file has the same structure: it includes experimental conditions (such as azimuth angles, depression angles, and target classes, among others), and the amplitude and phase information. The target images were captured by X-band SAR sensor and have a resolution of 0.3×0.3m. The test dataset consists of the central-cropped 64 × 64-sized 2425 SAR target images under a depression angle of 15°. This dataset contains 10 target classes that are composed of one bulldozer (D7), one truck (ZIL131), one air defence unit (ZSU234), one rocket launcher (2S1), two tanks

TABLE II
TESTING SYNTHETIC-TO-MEASURED CONFIGURATIONS

Test name	**Training dataset** (number of images)	**Augmentations**		
		ID	**Images**	**ASC**
Baseline	S[2] (3600)	-	-	-
Augmented Baseline	S (3600)	A0	Traditional SAR images augmentations[1]	-
Baseline ASC	S + R[3] (7200)	-	-	-
TSM1	S + R (7200)	A1	-	Gaussian noise on reflectors amplitude.
TSM2	S + R (7200)	A2	-	Gaussian noise on reflectors position (x and y)
TSM3	S + R (7200)	A3	-	A1 + A2
TSM4	S + R (7200)	A4	-	Adding reflectors (randomly defined and similar to the target)
TSM5	S + R (7200)	A5	-	Random removing of reflectors.
TSM6	S + R (7200)	A6	-	A3 + A4
TSM7	S + R (7200)	A7	-	A3 + A5
TSM8	S + R (7200)	A8	-	A4 + A5
TSM9	S + R (7200)	A9	-	A6 + A5
TSM10	S + R (7200)	A10	A0	A9

[1]The data augmentation methods used are detailed in III-C.
[2]S = Original synthetic dataset.
[3]R = Reconstructed synthetic dataset.

(T62, T72), and four armoured personnel carriers (BMP2, BTR60, BTR70, BRDM2).

B. Study of the ASC modelling

1) Quality evaluation of ASC extraction algorithm:

a) Impact of the number of iterations: Analysing the effectiveness of ASC parameters extraction according to different iteration values allows us to choose the optimal value for generating training data. Fig. 5 presents the ICC metrics resulting from all 2S1 synthetic SAR images compared with their pairwise reconstruction. The left plot shows sorted ICC values obtained from all pairs for different number of iterations. The right heatmap presents the same results according to azimuthal sectors. Both graphs allow us to conclude that reconstruction improves as the number of iterations increases. However, after 100 iterations, the gain on the ICC metric weakens, which is why a value of 200 seems to strike a balance between quality and generation speed.

It is also very interesting to note that ICC performs much better for certain angular ranges, particularly around 90° and 270°, which correspond to cases where the target is facing forwards or backwards, regardless of the number of iterations.

b) Impact of the number of scattering centres: To conclude on the effect of the number of scattering centres considered in the algorithm, we propose to calculate the ICC between the reconstructed image and the original image of the synthetic data for all targets. Histograms for 5, 25, 50 and 100 considered reflectors are displayed Fig. 6. Five points are clearly insufficient to reconstruct a high-quality image and a real improvement is observed from 50 points onwards. Using 100 reflectors reduces the risk of obtaining data with a low ICC, but considering the increasing calculation time with the number of reflectors, a value of 50 reflectors will be retained for the rest of the experiments.

2) Analysis of estimated ASC parameters:

a) Stability of extraction according to classes and angular sectors: To observe trends in ASC parameters in terms of number of points between different targets, but also according

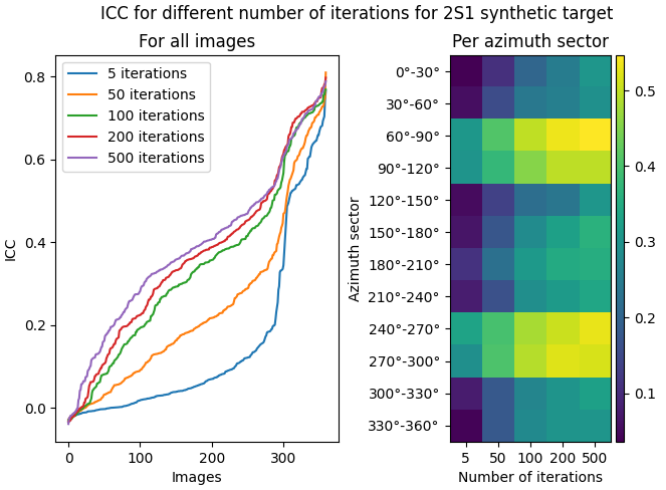

Fig. 5. ICC analysis between the original SAR image and the image reconstructed from the ASC parameters extracted according to 5, 50, 100, 200 and 500 iterations. The extraction is performed with 50 reflectors. On the left are the ICC curves sorted for all images, and on the right is the heatmap by angular sector for the different iteration values.

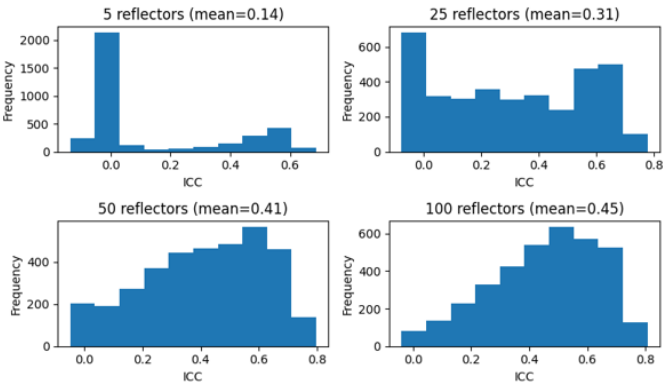

Fig. 6. ICC distribution between reconstructed and synthetic original images for 5, 25, 50 and 100 scattering scatters. The extraction is performed with 50 reflectors and 200 iterations.

to the angular sector considered, the Fig. 7. presents heat maps for synthetic and real datasets. All data is obtained by extracting 50 reflectors, but some are null. Here, we consider the number of points to be the number of non-null points among the 50.

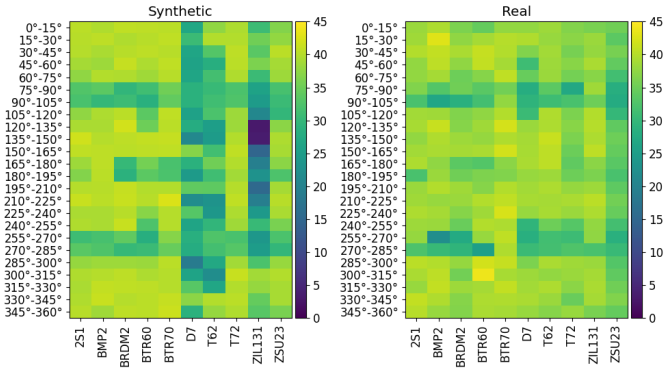

Fig. 7. Average number of reflectors per angular sector for each target for synthetic and real datasets.

Two main findings emerge from this study. Firstly, the study reveals a similar trend between synthetic and real datasets according to angular sectors around azimuths 90° and 270° which have fewer reflectors. This corresponds to cases where the target is facing forwards and backwards. This is not a problem, as it has been observed Fig. 5 that reconstruction is better at these angles.

Secondly, some synthetic targets show a significant gap with measured data, which manifests itself in a lack of reflectors. In particular, the D7 target has fewer reflectors in the synthetic case, which is consistent with its smaller size compared to other targets, but this trend, is in fact, absent in the real case. This is also the case for T62 target, even though it is similar to the T72. Also, the ZIL131 seems to be pathological for azimuths between 120° and 150° as there are practically no extracted reflectors.

The study can be refined by looking at the number of reflectors per radar effect in the same cases. Fig. 8 presenting the average number of effects for each angular sector for all targets combined shows that trihedral and corner-diffraction effects are predominant, while cylinder effects are rarely detected.

Fig. 9 focuses on the ZIL131 pathological case via a synthetic versus measured violin plot confronted to the 2S1 case. The 2S1 distributions are more symmetrical between synthetic and real cases, as shown by the EMD coefficient, which shows greater similarity than the ZIL131, which highlights significant asymmetry.

For further analysis, Fig. 10 reveals that pathological cases generally happen when a pixel or a small group of pixels is significantly brighter than other pixels of the target, which end up being ignored in the ASC extraction.

b) Extraction efficiency between real and synthetic data: In order to study the differences in extraction between synthetic and real data, we display Fig. 11 showing the ICC

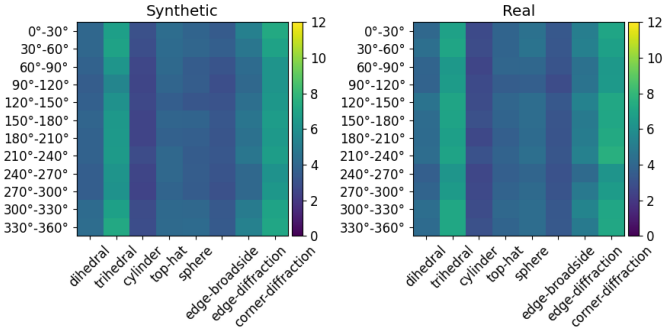

Fig. 8. Average number of reflectors for each radar effect.

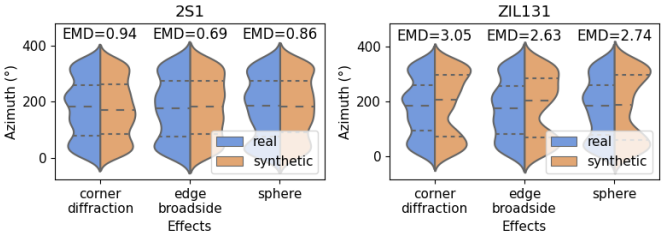

Fig. 9. Violin plot for the 2S1 and ZIL131 effects: synthetic vs real. For visualisation purposes, we only display three effects. The Earth Mover Distance (EMD) coefficient is also displayed for each effect. This calculates the difference between the real and synthetic distribution pairs. It can also be seen as the amount of change required for the synthetic distribution to equal the real one.

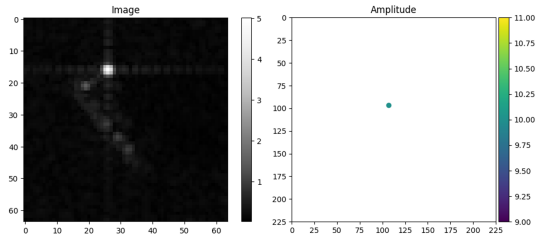

Fig. 10. Example of a pathologic case of the ZIL131 target, azimuth 125°. SAR Image is dispalyed on the left, while ASC parameters are shown on the right.

distribution via histograms for four different targets. Overall, better reconstructions are possible based on real data. This fact is somewhat reassuring, as it shows that ASC extraction reflects physical reality.

C. ASC data augmentation study for SAR ATR

Training configurations are completed as described in Table II. For all cases, the trained ATR is then tested on the same MSTAR real set and the results are provided in Table III with accuracy and F1-score.

Training the ATR model on reconstructed ASC dataset set results (Baseline ASC) in an accuracy of 55.88%, while adding ASC augmentations process (TSM9) ends up with 60.86%. Adding ASC augmentation process confers an improvement almost identical to adding image augmentation, which gives

TABLE III
EXPERIMENTAL RESULTS

Test name	Baseline	Augmented Baseline	Baseline ASC	TSM1	TSM2	TSM3	TSM4	TSM5	TSM6	TSM7	TSM8	TSM9	**TSM10**
Test Accuracy	53.25	59.58	55.88	54.63	58.8	51.67	62.21	54.3	60.81	56.62	52.78	60.86	**70.97**
F1-Score	51.24	58.5	55.39	53.45	57.79	49.93	60.24	53.07	59.79	54.88	51.49	59.83	**70.16**

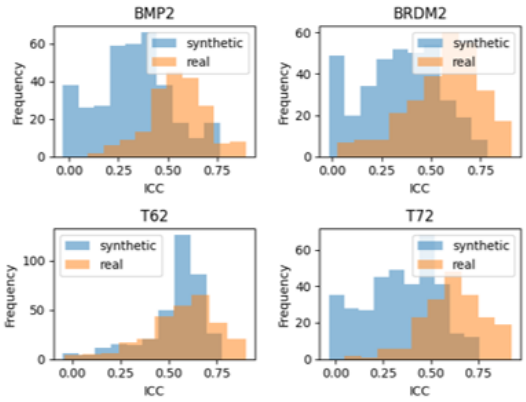

Fig. 11. ICC distribution: synthetic vs real (50 reflectors, 200 iterations).

a 59.58% accuracy. However, combining the two methods (TSM10) gives an accuracy of 70.97%.

According to the TSM4 test, adding new reflectors near the target seems to have a positive effect on performance. TSM6 and TSM9 tests confirm this trend, as they are also based on reflectors addition augmentations. This result could be explained by the fact that synthetic and measured samples may exhibit different target signatures, with scatterers amplitude and position slightly shifted.

We present our method as a complement to existing data augmentation techniques and compare it to those described in the literature. For example, Baffour et al. [1] use traditional augmentation and report 92% accuracy on MSTAR, but their results rely on the flawed SAMPLE dataset [2]. The approach by Camus et al. [2] is most comparable to ours: with their augmentation pipeline and MOCEM dataset, they reach 75% accuracy. The observed differences stem from the use of broader augmentation (affecting both target and clutter, versus only the target in our ASC-based DA), and more techniques employed (bagging, adversarial training, Test-Time Domain Augmentation). In order to quantify the isolated contribution of the data augmentation method, we can compare our TSM9, which achieves 60.86% accuracy, with their test incorporating the ResNet architecture and domain randomisation (both tests also involve random shifting), which achieves 50.48% accuracy. This result highlights the promise and robustness of our approach, which is independent of MOCEM and induces more localized changes.

V. CONCLUSION AND PERSPECTIVES

In this paper, we propose a data augmentation technique based on physical knowledge about SAR images to bridge the gap between synthetic and measured data. First, we presented the Attributed Scattering Centres Model, which is a physical model aiming at accurately describing the target information in SAR images. Next, we described our data augmentation method using this physical knowledge. In fact, having a set of ASC parameters, it is next possible to generate new samples by slightly perturbing them. These augmented samples can then be used to train models. Given the MSTAR and the corresponding synthetic datasets, we analyse different augmentation configurations and their impact on recognition performance. We showed that our proposed data augmentation pipeline can help gain up to 15.09% in terms of accuracy (see TSM10 vs Baseline ASC in Table III). Moreover, we identified that adding scattering mechanisms near the target is the most prolific augmentation, being able to give physical clues about the SAR S2M domain gap.

As a consequence, we believe that this approach shows potential for the future. Forthcoming work could involve the several following aspects.

Our ASC parameter estimation algorithm is not yet fully operational as it can be seen in Fig. 6. Some samples are not correctly reconstructed, leading to incomplete target information for some classes and azimuth angles, and thus deteriorating the representativeness of the dataset. As a consequence, we believe that our data augmentation pipeline and its corresponding results are not at the maximum of their potential. These moderate results can be explained by different reasons, for example by our way of initialising parameters. On the other hand, we need to further investigate and characterise the S2M domain gap. In fact, this paper shows some preliminary results that enabled us to test our data augmentation process. However, as explained earlier, we truly believe that we can also unveil which data augmentation method is most useful in the current task and why with our DA process. For example, we would like to be able to provide explanations such as: "we identified that for these measured and synthetic datasets, the rear of the 2S1 target is different because it shows dihedral instead of trihedral effects. As a consequence, we can guarantee domain generalisation by using the corresponding data augmentation during training." In other words, we would like to both characterise physically why synthetic and measured samples are different, and bridge the domain by using the most useful and justified data augmentation technique.

ACKNOWLEDGEMENT

We deeply thank the French MoD (DGA) for providing the synthetic dataset and reviewing this article. We also address our gratitude to SCALIAN DS for our constructive discussions.

REFERENCES

[1] A. A. Baffour, I. Osei Agyemang, I. Adjei-Mensah and R. E. Nuhoho, "Towards Fully Synthetic Training: Exploring Data Augmentations for Synthetic-to-Measured SAR in Automatic Target Recognition", 2025, IEEE. [Online]. Available: https://ieeexplore.ieee.org/document/11105313,

[2] B. Camus, C. Le Barbu, E. Monteux, "Robust SAR ATR on MSTAR with Deep Learning Models trained on Full Synthetic MOCEM data", 2022, arXiv:2206.07352. [Online]. Available: https://arxiv.org/abs/2206.07352

[3] Y. Zhong, W. Zhou, and Z. Wang, 'A Survey of Data Augmentation in Domain Generalization', Neural Process Lett, vol. 57, no. 2, p. 34, Mar. 2025, doi: 10.1007/s11063-025-11747-9.

[4] M. Scherreik and K. Grubaugh, 'Linear separability of the SAR domain gap', in Algorithms for Synthetic Aperture Radar Imagery XXXII, E. Zelnio and F. D. Garber, Eds., Orlando, United States: SPIE, May 2025, p. 32. doi: 10.1117/12.3067969.

[5] N. Inkawhich et al., 'Bridging a Gap in SAR-ATR: Training on Fully Synthetic and Testing on Measured Data', IEEE J. Sel. Top. Appl. Earth Observations Remote Sensing, vol. 14, pp. 2942–2955, 2021, doi: 10.1109/JSTARS.2021.3059991.

[6] A. W. Denton, 'DATA AUGMENTATION FOR SYNTHETIC APERTURE RADAR USING ALPHA BLENDING AND DEEP LAYER TRAINING'.

[7] J. Lv and Y. Liu, 'Data Augmentation Based on Attributed Scattering Centers to Train Robust CNN for SAR ATR', IEEE Access, vol. 7, pp. 25459–25473, 2019, doi: 10.1109/ACCESS.2019.2900522.

[8] B. Ding, G. Wen, X. Huang, C. Ma, and X. Yang, 'Data Augmentation by Multilevel Reconstruction Using Attributed Scattering Center for SAR Target Recognition', IEEE Geosci. Remote Sensing Lett., vol. 14, no. 6, pp. 979–983, Jun. 2017, doi: 10.1109/LGRS.2017.2692386.

[9] M. J. Gerry, "Two-dimensional inverse scattering based on the GTD model /," phdthesis, The Ohio State University, 1997. [Online]. Available: http://rave.ohiolink.edu/etdc/view?acc_num=osu1487946103567201.

[10] MSTAR database. Available online: https://www.sdms.afrl.af.mil/index.php?collection=mstar.

[11] C. COCHIN, P. POULIGUEN, B. DELAHAYE, D. l. HELLARD, P. GOSSELIN and F. AUBINEAU, "MOCEM - An 'all in one' tool to simulate SAR image," 7th European Conference on Synthetic Aperture Radar, Friedrichshafen, Germany, 2008, pp. 1-4. keywords: Solid modeling;Radar imaging;Synthetic aperture radar;Radar cross section;Computational modeling;Image resolution,

[12] B. Camus, S. Ds, J.-C. Louvigné, and C. Saleun, "Génération massive d'images SAR synthétiques pour l'IA à fins de classification automatique (ATR)," presented at ENVIREM 2025, Palaiseau, France, 2025.