

Benchmarking Deep Learning Classifiers for SAR Automatic Target Recognition

Jacob Fein-Ashley*, Tian Ye*, Rajgopal Kannan[†], Viktor Prasanna*, Carl Busart[†]
*University of Southern California [†]DEVCOM Army Research Lab
*{feinashl, tye69227, prasanna}@usc.edu [†]{rajgopal.kannan.civ, carl.e.busart.civ}@army.mil

Abstract—Synthetic Aperture Radar (SAR) Automatic Target Recognition (ATR) is a key technique of remote-sensing image recognition, which can be supported by deep neural networks. The existing works of SAR ATR mostly focus on improving the accuracy of the target recognition while ignoring the system’s performance in terms of speed and storage, which is critical to real-world applications of SAR ATR. For decision-makers aiming to identify a proper deep learning model to deploy in a SAR ATR system, it is important to understand the performance of different candidate deep learning models and determine the best model accordingly. This paper comprehensively benchmarks several advanced deep learning models for SAR ATR with multiple distinct SAR imagery datasets. Specifically, we train and test five SAR image classifiers based on Residual Neural Networks (ResNet18, ResNet34, ResNet50), Graph Neural Network (GNN), and Vision Transformer for Small-Sized Datasets (SS-ViT). We select three datasets (MSTAR, GBSAR, and SynthWakeSAR) that offer heterogeneity. We evaluate and compare the five classifiers concerning their classification accuracy, runtime performance in terms of inference throughput, and analytical performance in terms of number of parameters, number of layers, model size and number of operations. Experimental results show that the GNN classifier outperforms with respect to throughput and latency. However, it is also shown that no clear model winner emerges from all of our chosen metrics and a “one model rules all” case is doubtful in the domain of SAR ATR.

Index Terms—Graph Neural Network, Synthetic Aperture Radar, Automatic Target Recognition, Benchmarking

I. INTRODUCTION

Synthetic Aperture Radar (SAR) Automatic Target Recognition (ATR) is a key technique in remote sensing to recognize on-ground targets with a wide range of applications in real-world scenarios. A SAR system sends radar signals to the ground from the air, processes the reflections, generates high-resolution grayscale SAR images, and passes them to ATR systems for real-time object recognition.

SAR ATR systems must provide efficient and accurate SAR image classification in real-world applications such as military, agriculture [1], [2], and geology. Deep learning techniques have been exploited in ATR systems to classify SAR images. Convolutional Neural Networks (CNNs), which have shown great success in optical image classification [3], are also widely used as SAR image classifiers in ATR systems. They benefit from their strong capability in extracting features of SAR images. Recently, a Graph Neural Network (GNN) model was proposed in our previous paper [4] for SAR image classification. An input SAR image is converted to a grid graph where every pixel as a vertex is connected with its

eight neighbors; the SAR image classification is interpreted as a graph classification task. The GNN model was trained on the Moving and Stationary Target Recognition (MSTAR [5], [6]) dataset and achieved a state-of-the-art accuracy of 99.1%.

Although existing CNNs and GNNs have shown high accuracy in the public MSTAR dataset, SAR ATR still faces challenges in their performance in various aspects, such as latency, throughput, and memory storage. The difference in performance will affect the decision to deploy SAR ATR in real-world scenarios. For example, in one case, when fast target recognition is critical, SAR ATR should have low latency and high throughput. In contrast, in another case, the model size is more essential when memory is a limited resource in the system. Therefore, it is necessary to perform a comprehensive benchmarking to compare the performance of different deep learning models for SAR ATR so that a proper model can be chosen to deploy according to the requirements and constraints of the application.

In this paper, we benchmark deep learning-based classifiers for SAR ATR. We evaluate the performance of SAR image classifiers in terms of their classification accuracy, runtime performance, and analytical performance. The runtime version includes inference latency and throughput that represents the speed of the ATR system. The analytical performance has the size of the model, the number of parameters, the number of layers, and the number of operations, representing resource consumption and efficiency. In this empirical study, we evaluate and compare three classes of typical deep learning models: Residual Neural Network (ResNet) [7], Graph Neural Network (GNN) [4], and a variant of the Vision Transformer (ViT) [8]. The ViT (SS-ViT) variant allows for higher accuracy on a smaller dataset, proposed in [9]. To cover more real-world scenarios, we select three distinct datasets with significant diversity as summarized in Table I. The three datasets are Mov-

TABLE I: Selected SAR Imagery Datasets

	MSTAR [6]	GBSAR [10]	SynthWakeSAR [11]
SAR device	Airborne-based	Ground-based	Airborne-based
Target object	Military vehicles	Ceramic objects	Ship wakes
Data type	Images	Raw radar data	Images
Training size	27,000	5,147	36,864
Testing size	2,425	1,287	9,216
# classes	10	7	10

ing and Stationary Target Recognition (MSTAR) [6], Ground-Based SAR Data Obtained With Different Combinations of

Bandwidth and Step Size (GBSAR) [10], and a dataset containing various simulated ship wakes (SynthWakeSAR) [11]. Section III describes more details about the datasets.

Our main contributions are summarized as follows:

- We conduct a benchmarking study that compares the classification accuracy, runtime performance, and analytical performance of SAR ATR based on three types of deep learning models (ResNet, GNN, and SS-ViT) with three distinct SAR imagery datasets (MSTAR, GBSAR, and SynthWakeSAR).
- Results tend to be mixed with respect to our chosen metrics. We observe that the GNN we proposed earlier [4] wins in terms of the number of multiply/accumulate operations, the number of model layers (these metrics are explained in Section III), and accuracy for the MSTAR dataset. The SS-ViT outperforms in the number of model parameters and size but underperforms in accuracy. We notice that the ResNet architectures contain an enormous parameter and the number of MAC operations but are more accurate than the other models in the GBSAR and SynthWakeSAR datasets.
- This work finds that, like many machine learning problems, there is still “no free lunch” when it comes to a single model outperforming others in terms of all performance metrics. Additionally, we find that a more thorough assessment of deep learning models for SAR ATR requires evaluating the models across a variety of datasets. For this purpose, it is imperative to develop more public SAR datasets.

II. BACKGROUND AND RELATED WORK

A. SAR and SAR ATR

Synthetic Aperture Radar (SAR) imagery is a radar imaging technique with wide applications in military [12], climate studies [13], marine surveillance [14], and agriculture [1], [2]. A SAR system, typically equipped on satellites or aircraft, transmits radar signals toward the earth’s surface. After receiving reflections from the earth, the system generates a grayscale image depending on the backscatter intensity of the reflections. Compared to optical sensors, radar signals are independent of lighting and weather conditions, such as clouds. As a result, SAR imagery can obtain high-resolution images of the earth in all weather states and at all times of day and night.

SAR Automatic Target Recognition (ATR) automatically identifies the objects in a SAR image, which can be interpreted as a task of image classification. Different from classifying optical images, SAR image classification has several challenges: (1) SAR images are single-channel grayscale images representing the intensity of backscattered radar signal, which carries less information than RGB images with three channels of colors. (2) SAR images of the same object can appear differently when they are imaged with different **capturing parameters** such as radar frequency, azimuth, incidence angle, and resolutions. For example, incidence angles will affect the appearance of shadow regions in SAR images and make the

brightness of a surface different. (3) Large-scale public labeled datasets of SAR images are hard to access [15] and therefore it is difficult to develop effective classifiers for unseen SAR images. Despite these challenges, many classifiers based on deep learning models that have been proven successful in optical image classification can still be applied to the domain of SAR ATR.

B. Deep Learning Models for SAR ATR

Convolutional Neural Networks (CNNs) are a class of most typical models for image classifications. They apply a series of convolutional kernels to local regions of input images, which leads to a strong ability to capture local features. Residual Networks (ResNets) are a type of CNNs that are usually easier to train and show higher accuracy than standard CNNs. There are several versions of ResNets with different numbers of layers. We consider three versions with fewest layers (ResNet18, ResNet34 and ResNet50) in this paper because more layers will not further improve the accuracy in our experiments.

Unlike CNNs, SAR ATR based on Graph Neural Networks (GNNs) interprets target recognition as a graph classification task. In our earlier work [4], we proposed a SAR ATR based on GNN. A given SAR image can be converted into a graph: Each pixel becomes a vertex of the graph with the pixel value as the vertex attribute, and a non-attributed edge connects every pair of neighbor pixels. Figure 1 shows the architecture

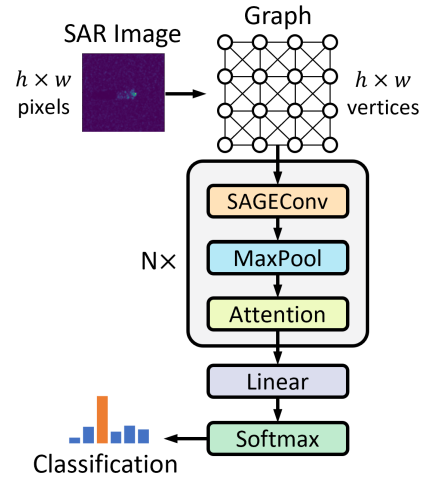


Fig. 1: Architecture of GNN-based SAR image classifier

of the GNN-based SAR image classifier [4]. N sequential groups of feature extractors generate an embedding for the graph, followed by a linear layer (“Linear”) and softmax layer (“Softmax”) that outputs the classification result. Each group of feature extractors consists of a GraphSAGE convolution operator [16] (“SAGEConv”) that generates embeddings for each vertex by aggregating the features of neighbors vertices, a max pooling layer (“MaxPool”) to reduce the dimension of the embeddings, and a layer of attention mechanism [17] (“Attention”) to guide the classifier to focus on the most important parts of the embeddings. In the experiments of our

previous paper [4], the GNN classifier obtains an accuracy of over 99% on the MSTAR dataset. Compared with CNNs, GNNs have lower computational costs while still achieving a similar level of classification accuracy.

Vision Transformer (SS-ViT [8], [9]) is a recently proposed technique that takes a different approach from CNNs and GNNs. The previous two classes of models interpret the input image as two-dimensional pixels: CNNs treat the input image as a matrix of pixel values, and GNNs build the graph by mapping each pixel as a vertex. Instead, SS-ViT divides an image into a sequence of small patches and then feeds the patches into a transformer model. The transformer can understand the relative importance of the patches and the relationship between patches. Therefore, SS-ViT has a better global understanding of the input image. Compared to CNNs and GNNs, the limitations of SS-ViT include its high requirements on the amount of training data and computational resources. To remedy this, [9] proposed a shift patch tokenization (SPT) and locality self-attention (LSA) alteration to the originally proposed SS-ViT. The change in the SS-ViT structure will allow us to use recent transformer machine-learning models on smaller datasets.

III. BENCHMARKING DEEP LEARNING MODELS FOR SAR ATR

Our motivations for a benchmark of deep learning models for heterogeneous SAR ATR datasets are two-fold: (1) With the recent release of SAR ATR datasets for classification purposes, we seek to study whether different datasets will influence the performance of deep learning models. (2) As many deep learning models are developed, we intend to compare them concerning multiple metrics, which will help practitioners to determine what model to use in a SAR ATR system.

In the space of SAR ATR deep learning, the MSTAR dataset is widely used. However, only relying on a single dataset is insufficient to assess the performance of the deep learning models on unseen SAR data. We hypothesize that deep learning models will depend on a mixture of SAR capturing methods, parameters, azimuth angles, and elevation criteria. We source from the limited supply of publicly available SAR datasets. From these datasets, we show the performance of a recently constructed GNN, ResNet18, ResNet34, ResNet50, and SS-ViT. The SS-ViT and ResNet models are widely available and often used for high-performing image classification tasks. Benchmarking performance in this variety of models for heterogeneous datasets is a vast frontier, as the number of publicly available datasets is limited. Exploring this frontier may expose the importance and sensitivity of SAR capturing parameters in the deep learning space.

A. Datasets

We obtain a variety of datasets to provide a varied perspective on azimuth, elevation, and range-to-velocity ratio, which are often factors in SAR image degradation. To retain consistency, all processed datasets require a resizing method to (88×88) pixels [18]. SAR ATR datasets are seldom

released to the public due to data sensitivity and acquisition difficulty. Thus, measuring popular deep learning models on SAR datasets with differing properties respective to azimuth, elevation, and SAR measuring device parameters is critical [19]. Details of the selected datasets and their sample images are as follows:

1. **MSTAR** [6] is a widely known SAR dataset that is publicly available and captures images at varying elevations and azimuth angles with 10 classes.




Fig. 2: Sample MSTAR Dataset Images

Concerning capturing parameters, all are constant except for elevations, which range from $14^\circ - 18^\circ$ [20].

2. **SynthWakeSAR** [11] contains synthetic (simulated) SAR images of wakes from 10 visible ships, consisting of normalized radar cross-sections (NRCS) with tilt and hydrodynamic modulations and velocity bunching. SynthWakeSAR releases distributions of noise-free images, noisy images with simulated speckles, and despeckled SAR images. For our purposes, we train our models on denoised images, as recommended by the authors of SynthWakeSAR [11]. The authors of this dataset do not state their capturing parameters for this dataset but mention that each scene remains the same size.




Fig. 3: Sample SynthWakeSAR Dataset Images

3. **GBSAR** [10] captures ceramic cups with rubber objects, each captured with different step sizes and bandwidths. GBSAR contains 7 classes. Note that GBSAR is a dataset of raw SAR data instead of reconstructed images. As suggested by the authors of GBSAR, we implement end-to-end learning by directly feeding the raw SAR data into the models. The capturing parameters of GBSAR vary, with step sizes $\in \{0.5, 1, 1.5, 3, 3.5\}$ cm, and three bandwidth cases $\in \{300, 600, 1300\}$ MHz.




Fig. 4: Sample SynthWakeSAR Dataset Images

B. Metrics

In addition to selecting datasets containing mixed capturing parameters, it is important to choose metrics that fairly compare results concerning runtime, accuracy, and practicality. The particular metrics we consider are as follows:

1. **Throughput and latency** allow for comparing the inference time across various models. Throughput is the maximum number of input instances a model can process per time unit. For our purposes, we define throughput as

$$\text{Throughput} = \frac{\text{Total number of images processed}}{\text{Total inference time}}$$

Thus, it is natural that the throughput units are in images processed per time unit. We also define **latency** as the average time of completion for a classification to be completed by a model.

2. Additionally, we use the **accuracy** metric to assess how well each classifier performs during prediction.
3. Other metrics we consider are analytical performance, including the **number of parameters (# Parameters)**, **model size, layer count (# Layers)**, and **number of multiply/accumulate operations (# MACs)**. In other words, suppose accumulator a counts an operation of arbitrary $b, c \in \mathbb{R}$. We count the number of operations as $a \leftarrow a + (b \times c)$. The layer count metric is an important factor of latency. Increasing the number of layers will also increase the latency of an inference of a model. The goal of an effective machine learning model is to maximize throughput while minimizing the number of MACs and a number of layers, in our case.

In summary, accuracy represents the ability of a model to recognize objects. Throughput, model size, # Parameters, and # MACs allow us to assess the practicality of each model. # Layers are one factor that can affect the inference latency and the potential capacity for hardware acceleration.

C. Deep Learning Models

A direct comparison of state-of-the-art models is conducted by comparing our GNN architecture to ResNet18, ResNet34, ResNet50, and SS-ViT. Baseline models are selected by heuristic performance on popular image classification datasets. ResNet architectures and SS-ViT are state-of-the-art models with well-known high accuracy for standard image classification tasks. Details of the models are described in Section II-B.

IV. EXPERIMENTS

A. Experimental Setup

We implement the training and testing of the three types of deep learning models using the PyTorch 1.11.0 [21] library on a state-of-the-art server with CUDA 11.8 powered by an NVIDIA RTX A6000 GPU. The implementation of ResNet18, ResNet34, ResNet50 and SS-ViT are based on open-source repositories [22], [23]. We use the PyTorch-OpCounter [24] library to count # Parameters and # MACs. The datasets, metrics, and deep learning models are summarized in the previous section.

B. Performance Comparison

We first compare the analytical performance using # MACs, # Parameters, model size, and # Layers for each model. As shown in Table II, among the studied deep learning models, GNN has the fewest layers and fewest operators in one forward propagation. Therefore, it is easier for potential hardware accelerators such as FPGAs and ASICs to improve the performance. The SS-ViT has the fewest parameters and smallest model size that better fits systems when memory is a critical resource.

TABLE II: Analytical Performance

Model	# MACs	# Parameters	Model Size (Mb)	# Layers
ResNet18	1.87×10^{10}	1.12×10^7	42.67	53
ResNet34	3.92×10^{10}	2.13×10^7	81.26	93
ResNet50	4.41×10^{10}	2.35×10^7	89.93	127
SS-ViT	1.70×10^{10}	5.15×10^5	2.12	247
GNN	8.15×10^7	1.27×10^6	5.08	36

Additionally, we measure the various metrics for classification accuracy and throughput for each dataset mentioned as shown in Tables III, IV, V. In most cases, our GNN outperforms in terms of throughput and latency with the exception that ResNet18 achieves better throughput on the SynthWakeSAR dataset. Each dataset yields a distinct deep learning model as the most accurate. No single model achieves better accuracy than others across all the three datasets.

TABLE III: MSTAR Performance

Model	Accuracy	Throughput (imgs/ms)	Latency (ms)
ResNet18	98.47%	5.91×10^2	1.87×10^{-3}
ResNet34	98.64%	4.96×10^2	1.94×10^{-3}
ResNet50	90.34%	4.50×10^2	2.12×10^{-3}
SS-ViT	95.61%	4.84×10^2	2.07×10^{-3}
GNN	99.09%	8.56×10^2	1.16×10^{-3}

TABLE IV: Denoised SynthWakeSAR Performance

Model	Accuracy	Throughput (imgs/ms)	Latency (ms)
ResNet18	90.30%	1.82×10^3	5.56×10^{-4}
ResNet34	92.14%	1.64×10^3	6.12×10^{-4}
ResNet50	92.42%	1.54×10^3	7.80×10^{-4}
SS-ViT	87.98%	0.67×10^3	14.66×10^{-4}
GNN	91.15%	1.75×10^3	4.55×10^{-4}

TABLE V: GBSAR Performance

Model	Accuracy	Throughput (imgs/ms)	Latency (ms)
ResNet18	99.30%	3.02×10^2	3.29×10^{-3}
ResNet34	99.99%	2.98×10^2	3.33×10^{-3}
ResNet50	99.53%	2.84×10^2	3.62×10^{-3}
SS-ViT	99.04%	3.59×10^2	1.84×10^{-3}
GNN	98.67%	5.43×10^2	1.79×10^{-3}

C. Discussions

Insight 1: No single model can outperform across all provided metrics according to our empirical measurement.

Regarding accuracy, the ResNet architectures outperform the SS-ViT and GNN with the SynthWakeSAR and GBSAR dataset cases. Our GNN classifier transcends the other models with respect to throughput and latency across all datasets. Considering the results of the inference time, model complexity, size, and performance in the machine-learning application is essential. Thus, large learning models like ResNet may not be practical. The fastest model concerning inference time in our set of models and datasets is the GNN.

Insight 2: We recommend that the practitioner thoroughly analyze their dataset before deploying deep learning models for SAR ATR usage. Due to the sensitivity of SAR capturing equipment, having a “single model rule all” case is unrealistic. As empirically evaluated in this paper, it is imperative to consider the implications of deploying an enormous model that delivers only marginally better accuracy. We leave it to the practitioner to perform a cost-benefit analysis when deploying deep learning models for SAR usage.

Insight 3: With the limited public release of SAR imaging data, predicting how models will perform on datasets is difficult. Realistically, providing SAR images with azimuth angles, elevation, and SAR capturing parameters may assist models to become more robust. Heuristically, it is seen that, although simple, the GBSAR dataset contains mixed parameters concerning bandwidth and step size. All models perform quite well on this dataset, so obtaining datasets with varying parameters, such as GBSAR, may be necessary to strengthen future deep learning models.

Insight 4: As simulated datasets become more prevalent, as seen in SynthWakeSAR, future study areas may include synthesizing SAR ATR datasets with various capturing parameters, e.g., azimuth angles and elevations. Synthesizing a dataset with diverse parameters, speckling, rotation, and other mentioned capturing methods would help create models invariant to rotation, speckling, and noisy imaging, a prevalent topic in the SAR ATR realm.

V. CONCLUSION

This work benchmarked the performance of widely popular deep learning image classification models with respect to various accuracy and computational performance metrics. We highlighted the necessity of data evaluation from the practitioner to create a robust model. We found that with the models we had experimented with, the practitioner must make informed decisions when selecting a model. The practitioner must perform a well-informed cost-benefit analysis with computational complexity and accuracy in mind.

The field of SAR ATR is exceedingly dependent on the MSTAR dataset. We believe heterogeneous datasets will be important in advancing the deep learning SAR ATR realm and future SAR image classification models.

We will consider the following future directions to improve further the robustness of deep learning models for SAR ATR. On the one hand, deep learning models should be designed to remain unaffected by variations in SAR capturing parameters. For example, the models should recognize objects in SAR

images from various azimuth angles and elevations. Also, as SAR images inevitably contain “salt and pepper” like noise spreading all over the images, deep learning models should be invariant to the speckle. On the other hand, deep learning models should be robust against perturbations from potential adversaries. In our paper [25], we have studied the vulnerability of our GNN model in the face of adversarial attacks. The GNN model’s accuracy degrades when the SAR images’ pixel values are slightly manipulated. This shows the necessity of more robust deep learning models for SAR ATR. In the future, by studying more advanced adversarial perturbations and developing defensive techniques accordingly, the robustness of deep learning models will be significantly improved. Additionally, this experiment was run using only one state-of-the-art GPU. Further studies may include experimenting with different GPUs and train-to-test ratios.

ACKNOWLEDGEMENT

Distribution Statement A: Approved for public release. Distribution is unlimited.

This work is supported by the DEVCOM Army Research Lab (ARL) under grant W911NF2220159.

REFERENCES

- [1] P. Zhan, W. Zhu, and N. Li, “An automated rice mapping method based on flooding signals in synthetic aperture radar time series,” *Remote Sensing of Environment*, vol. 252, p. 112112, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0034425720304855>
- [2] L. Landuyt, A. Van Wesemael, G. J.-P. Schumann, R. Hostache, N. E. C. Verhoest, and F. M. B. Van Coillie, “Flood mapping based on synthetic aperture radar: An assessment of established approaches,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 2, pp. 722–739, 2019.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf
- [4] B. Zhang, R. Kannan, V. Prasanna, and C. Busart, “Accurate, low-latency, efficient sar automatic target recognition on fpga,” in *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*. Los Alamitos, CA, USA: IEEE Computer Society, sep 2022, pp. 1–8. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/FPL57034.2022.00013>
- [5] T. D. Ross, S. W. Worrell, V. J. Velten, J. C. Mossing, and M. L. Bryant, “Standard SAR ATR evaluation experiments using the MSTAR public release data set,” in *Algorithms for Synthetic Aperture Radar Imagery V*, E. G. Zelnio, Ed., vol. 3370, International Society for Optics and Photonics, SPIE, 1998, pp. 566 – 573. [Online]. Available: <https://doi.org/10.1117/12.321859>
- [6] E. R. Keydel, S. W. Lee, and J. T. Moore, “MSTAR extended operating conditions: a tutorial,” in *Algorithms for Synthetic Aperture Radar Imagery III*, vol. 2757, International Society for Optics and Photonics, SPIE, 1996, pp. 228 – 242.
- [7] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [8] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [9] S. H. Lee, S. Lee, and B. C. Song, “Vision transformer for small-size datasets,” *CoRR*, vol. abs/2112.13492, 2021. [Online]. Available: <https://arxiv.org/abs/2112.13492>

- [10] F. Turčinović, "Ground based sar data obtained with different combinations of bandwidth and step size," May 2023. [Online]. Available: <https://data.mendeley.com/datasets/vpmnxw5nng/1>
- [11] I. G. Rizaev and A. Achim, "Synthwakesar: A synthetic sar dataset for deep learning classification of ships at sea," *Remote Sensing*, vol. 14, no. 16, p. 3999, Aug 2022. [Online]. Available: <http://dx.doi.org/10.3390/rs14163999>
- [12] G. G. Schaber, "Sar studies in the yuma desert, arizona: sand penetration, geology, and the detection of military ordnance debris," *Remote sensing of environment*, vol. 67, no. 3, pp. 320–347, 1999.
- [13] B. Bechtel, L. See, G. Mills, and M. Foley, "Classification of local climate zones using sar and multispectral data in an arid environment," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 9, no. 7, pp. 3097–3105, 2016.
- [14] T. Zhang, X. Zhang, J. Shi, and S. Wei, "Hyperli-net: A hyper-light deep learning network for high-accurate and high-speed ship detection from synthetic aperture radar imagery," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 167, pp. 123–153, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271620301404>
- [15] M. Rostami, S. Kolouri, E. Eaton, and K. Kim, "Deep transfer learning for few-shot sar image classification," *Remote Sensing*, vol. 11, no. 11, 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/11/1374>
- [16] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 1025–1035.
- [17] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "Cbam: Convolutional block attention module," in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham: Springer International Publishing, 2018, pp. 3–19.
- [18] S. Chen, H. Wang, F. Xu, and Y.-Q. Jin, "Target classification using the deep convolutional networks for sar images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 8, pp. 4806–4817, 2016.
- [19] P. Zhao, L. Huang, Y. Xin, J. Guo, and Z. Pan, "Multi-aspect sar target recognition based on prototypical network with a small number of training samples," Jun 2021. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8271429/>
- [20] B. Lewis, M. Levy, J. Nehrbass, T. Scarnati, E. Sudkamp, and E. Zelnio. [Online]. Available: <https://apps.dtic.mil/sti/pdfs/ADA364024.pdf>
- [21] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [22] TorchVision, "Torchvision: Pytorch's computer vision library," <https://github.com/pytorch/vision>, 2016.
- [23] Lucidrains, "Vision transformer - pytorch," <https://github.com/lucidrains/vit-pytorch>, 2022.
- [24] Lyken17, "Lyken17/pytorch-opcounter: Count the macs / flops of your pytorch model." [Online]. Available: <https://github.com/Lyken17/pytorch-OpCounter>
- [25] T. Ye, R. Kannan, V. Prasanna, and C. Busart, "Adversarial attack on GNN-based SAR image classifier," in *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications V*, L. Solomon and P. J. Schwartz, Eds., vol. 12538, International Society for Optics and Photonics. SPIE, 2023, p. 1253812. [Online]. Available: <https://doi.org/10.1117/12.2666030>