

# Deep Convolutional Neural Networks as a Method to Classify Rotating Objects based on Monostatic Radar Cross Section

 ISSN 1751-8644  
 doi: 0000000000  
 www.ietdl.org

 Eric Wengrowski<sup>†,1</sup>, Matthew Purri<sup>†,1,\*</sup>, Kristin Dana<sup>1</sup>, Andrew Huston<sup>2</sup>
<sup>†</sup>Equally contributing co-primary authors

<sup>1</sup>Rutgers University, Department of Electrical and Computer Engineering, New Brunswick, USA

<sup>2</sup>Lockheed Martin, Radar Group, Moorestown, USA

\*E-mail: matthew.purri@rutgers.edu

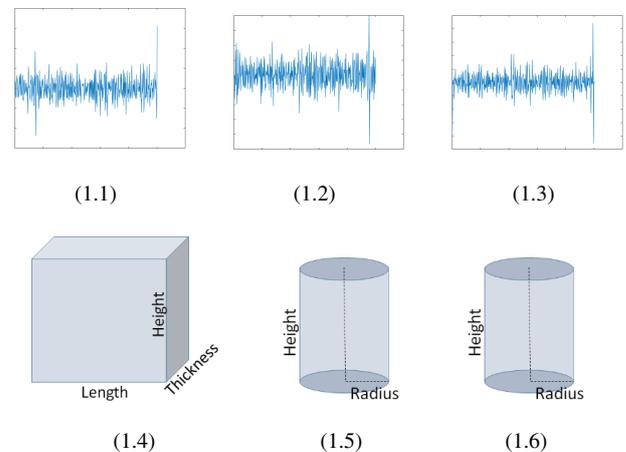
**Abstract:** Radar systems emit a time-varying signal and measure the response of a radar-reflecting surface. In the case of narrowband, monostatic radar signal domain, all spatial information is projected into a Radar Cross Section (RCS) scalar. We address the challenging problem of determining shape class using monostatic RCS estimates collected as a time series from a rotating object tumbling with unknown motion parameters under detectability limitations and signal noise. Previous shape classification methods have relied on image-like synthetic aperture radar (SAR) or multistatic (multiview) radar configurations with known geometry. Convolutional neural networks (CNNs) have revolutionized learning tasks in the computer vision domain by leveraging images and video rich with high-resolution 2D or 3D spatial information. We show that a feed-forward CNN can be trained to successfully classify object shape using only noisy monostatic RCS signals with unknown motion. We construct datasets containing over 100,000 simulated RCS signals belonging to different shape classes. We introduce deep neural network architectures that produce 2% classification error on testing data. We also introduce a refinement network that transforms simulated signals to appear more realistic and improve training utility. The results are a pioneering step toward the recognition of more complex targets using narrowband, monostatic radar.

## 1 INTRODUCTION

When illuminated with a narrowband radar signal, an object reflects incident energy and the reflectance depends on the object's geometry and material properties. The amount of energy that is reflected directly back toward the source of illumination is a function of its monostatic RCS (Radar Cross Section). As an object changes orientation, the RCS changes as well. We wish to classify the 3D shape of objects based only on a time series of monostatic RCS as the object moves according to force-free rigid body motion. Our set of target objects includes right circular cones, right circular cylinders, rectangular planes, spheroids, and trapezoidal prisms. The target object set varies in size with respect to a geometric parameter for each class (e.g. radius and height variation for cylinders). The chosen geometric properties in the test set are selected by radar wavelength so that each object is modeled as a Perfect Electrical Conductor (PEC). Labelled data, i.e. RCS of known objects, are required to train and test our supervised classifier. We create a large dataset of geometric objects and their corresponding RCS time-series signals.

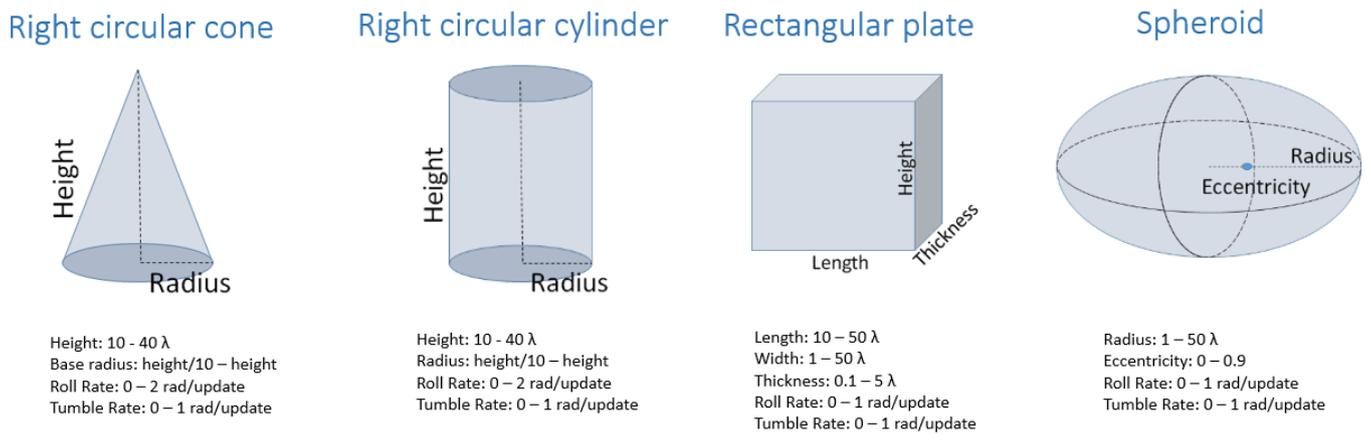
To simulate real-world conditions, the input signals for testing are corrupted by Gaussian noise and *Swerling dropout*. The Swerling Model [1] is a standard method for determining the detectability of an object based on SNR and waveform characteristics. The instantaneous probability of detecting each object at a given time is explicitly included in order to make the performance closer to real world operation. If the Signal to Noise Ratio (SNR) at a given time point is too small, a real-world radar system may be unable to separate the object from noise and will therefore be unable to detect the object and estimate its RCS.

A subset of the generated signals are used to train a feed-forward convolutional neural network classifier. We employ an end-to-end learning architecture, where signal features and the classifier are jointly solved for. The inputs are a series of RCS samples over time as the object rotates through free space. These objects belong to one of four shape families, illustrated in Figure 2. When the rotation is simple and follows a known path (as shown in Figure 5, top row), the



**Fig. 1:** Our goal is to correctly predict object shape family from a noisy monostatic RCS signal. RCS is highly sensitive to motion, and the rotation rates and viewing angles are unknown to the classifier. For example, objects may be rotating very fast or very slow about multiple axes. These signals contain added white Gaussian noise and a Swerling detection model, where the probability of detection is smaller for lower RCS values results in missing data points. A convolutional neural network (CNN) is used to *learn* the separating features that accurately recognize each object class overcoming the challenge of noisy data, missing data and unknown trajectories.

problem is trivial. However, the problem becomes substantially more difficult when the motion parameters are unknown (see Figure 5, bottom row).



**Fig. 2:** The four shape families correspond to four target classes in our classifier. Each shape class has a range of geometric parameters and motion parameters. The parameter ranges are listed under each shape.  $\lambda$  is wavelength of the incident radar signal.

In this work, we successfully classify the shape family for rotating objects with unknown roll rates, tumble rates, and unknown initial orientations. We train deep neural network classifiers that return the probability of each signal belonging to each shape family. The deep learning training and testing is implemented using PyTorch, a machine learning and optimization library for the Python programming language [2]. The SVM and Decision Tree algorithms are implemented using the SciPy library for the Python programming language [3]. To our knowledge, our methods are the first application of deep learning for object shape classification using monostatic radar signals.

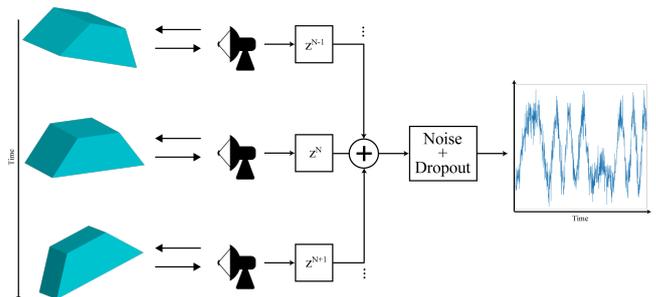
## 2 RELATED WORK

Producing an accurate representation of a target object’s narrowband monostatic RCS is a challenging problem. Radar specific properties such as wavelength and sampling rate, as well as object-specific properties such as surface material, shape, and motion may dramatically influence the resulting RCS time series. In this application, the objects under investigation are geometrically simple, convex shapes with uniform material construction. The incident energy wave is assumed to be a simple plane wave. The environment is not modeled, except for the addition of Gaussian noise. Due to these constraints, the physical optics (PO) approximation is appropriate to produce realistic returns. Open source RCS signal generation tools such as the Matlab toolbox POFacets are readily available [4] and have been used to approximate RCS of aircraft models [5].

A powerful new class of supervised machine learning algorithms called *convolutional neural networks* (CNNs) leverage optimization to learn complex latent features for robust classification. This family of algorithms is called *deep learning* when networks contain many convolutional layers. In 2012, a convolutional neural network significantly outperformed all other algorithms on the object classification dataset ImageNet [6] and CNNs have become the algorithm of choice for image recognition in computer vision [7–11].

Traditional neural networks have been used for radar classification tasks for decades, often derived from architectures developed for speech recognition such as the time-delay neural network [12, 13]. Early work on neural networks for processing radar signals were applied to identifying the number and type of radar emitters in a simulated multisource environment [14]. Pulse-train radar signal classification and source identification remains a topic of active research [15, 16]. Another recent challenge for neural networks in radar is the identification of radar jamming signals [17, 18]. Traditional neural networks have been applied to: SAR imagery for ground terrain classification [19] and crop classification [20]; microwave radar for classifying pedestrians and vehicles [21]; doppler radar for identify human breathing [22]; ground penetrating radar for the classification of geological structures [23]; forward scattering radar for identifying very small marine targets [24].

While traditional neural networks have been used widely in radar classification tasks, modern deep learning and CNNs are beginning to take hold in recent applications [25–29]. The success of the 2D CNNs on standard color images has translated well into radar applications. While most deep learning networks are designed for 2D imagery and can be directly applied to radar-based imagery, however, the RCS time series signals in our work are one-dimensional signals. In fields such as natural language processing [30] and medical applications [31], 1D CNNs have provided successful classification. In this work, we leverage successful deep networks for 2D image recognition, but adapt the networks to the 1D monostatic RCS signals.



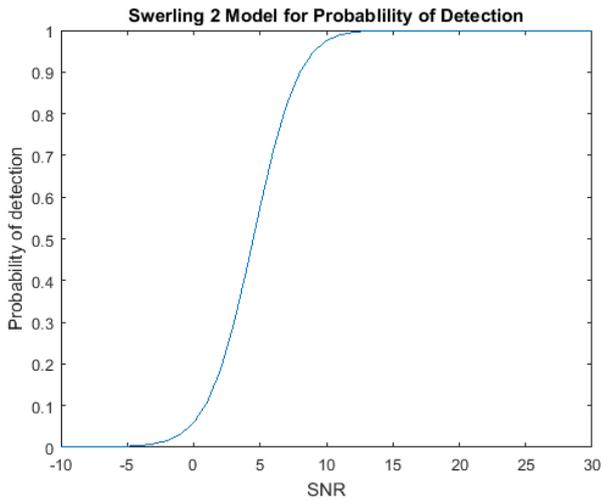
**Fig. 3:** The POFacets library was used to generate RCS signals from geometric shape models. Generalized Euler motion, additive Gaussian noise, and Swerling 2 dropout are then incorporated to generate the final signal.

Multi-static radar systems utilize a set of receivers and transmitters to create multiple 1D RCS signals of a target object. In prior work, multistatic RCS signals are classified individually using CNNs [25, 32] and the average of multiple CNNs [33] for multi-static contextual target signatures. The monostatic system addressed in our work contains a single collocated receiver-transmitter pair, compared to multistatic systems which have one or more spatially separated receivers and transmitters. The classification problem of monostatic RCS signals is particularly challenging since the signals do not contain contextual information from multiple sources.

## 3 GENERATING RCS SIGNALS

The first step in RCS classification is generating 3D models of our target objects. The parameters of these objects are listed in Figure 2. 128 geometric models were generated, each corresponding to one of four shape classes in the primary experiments. For each of the 3D models, POFacets is used to generate narrowband monostatic RCS values. In the case of monostatic radar, we assume that the radar

source and receiver are at the same location. The radar frequency is kept constant. It is important to note that in the physical optics model, RCS behavior depends only on the size of the object in wavelengths. Thus we can arbitrarily set the chosen frequency to  $0.3\text{GHz}$  while preserving the general behavior of any wavelength. Since the 3D model parameters are scaled by wavelength, this allowed for unit shape size parameters. POFacets is used to generate narrowband monostatic RCS responses, sensitive to object rotation parameterized by  $\theta$  and  $\phi$ . The mapping is done by specifying an angular sweep from  $0^\circ$  to  $180^\circ$  at high sampling intervals ( $0.1^\circ$ ). Symmetry about the shapes allows us to simulate to a maximum rotation of  $180^\circ$ .



**Fig. 4:** Swerling detectability is an important parameter in our model. As the RCS SNR decreased, so does the probability of detection. According to the above graph, SNRs of  $25\text{dB}$  and  $15\text{dB}$  provide almost no dropped measurements. But for  $\text{SNR} = 5\text{dB}$ , the probability of detection drops significantly, to roughly 50%. The RCS measurements with the lowest magnitude have a greater likelihood of being dropped to 0. Although Swerling dropout did have a major effect on our results, it often preserves larger RCS values in the time series signal, and the larger RCS values are expected to play a more substantial role in feature selection.

### 3.1 Generalized Euler Motion

Once an RCS map had been generated, a motion path is drawn over the surface and the map is interpolated. The target objects are assigned tumble, roll, and initial rotation angle. The initial conditions are then propagated following the physics of rigid body motion in the presence of no external forces (free motion). A quaternion model is used to generate the motion path parameterized by  $\theta$  and  $\phi$  over the precomputed 2D RCS map. The roll and tumble parameters are bound by the values described in Figure 2. For each shape class, the center of mass and moment of inertia are calculated and used for the simulation of realistic, geometry-dependent object motion.

### 3.2 Randomizations in Motion Parameters

It would be relatively easy to classify RCS signals from objects at integer-valued roll, tumble, and viewing angle. To make the problem more realistic and challenging, randomizations were applied to the values of each parameter. A random variable  $x$  with  $\mu = 1$  and  $\sigma = 0.5$  was multiplied with the viewing angle ( $\theta$  and  $\phi$ ), tumble rate, and rotation rate for each signal. The random variation allows for the construction of a database where the same 2D RCS map could be used to generate multiple signals. The ability to scale motion parameters with random jitter allowed the creation a nearly equal number of signals between the four classes, even though there were

more 3D models created for plates. CNN performance is generally improved when there are equal number of training examples in all classes.

### 3.3 Update Rate, Swerling, Gaussian Noise, Gradients, and Pyramids

A realistic radar model has a finite update rate. The number of samples as an object rotates are related to the update rate (in Hz) and the rotation rates (in radians/second). In this study the kinematic bounds of the objects are defined in radians/update, thus the performance of a highly sampled signal that rotates quickly is the same as if it were rotating more slowly with a corresponding decrease in radar update rate. The motion parameters are specified in radians per update. The radar update rate is arbitrarily set to 1 Hz. To simulate realistic distortions of each RCS value, Gaussian noise and a Swerling detectability model are incorporated into each RCS signal. The addition of Gaussian noise transforms the RCS from a truth value to an estimate. The specific parameters can be found in Table 1.

To summarize, the objects under test have complex motion with tumble, roll, and variable viewing angles, yielding complex time series of RCS estimates. The signals are noisy and have missing data points. Each RCS signal dataset contains variable values for each of the aforementioned parameters. Therefore, the same classifier is expected to correctly label RCS signals from objects moving at highly varied speeds in highly varied motion paths with different amounts of noise.

**Table 1** Generation parameters for A4 and B4 datasets

Parameters	A4	B4
Number of classes	4	4
Tumble rate (rad/sec*max)	0.015, 0.1, 0.5	0.015, 0.1, 0.5, 1
Roll rate (rad/sec*max)	0.015, 0.1, 0.5	0.015, 0.1, 0.5, 1
Signal to noise ratio (dB)	25, 15	25, 15, 5
Viewing vector angle (degrees)	0, 20	0, 20, 40, 60
Swerling model	2	2
Probability of false alarm	0.0001	0.0001
Number of pulses	10	10
Signal length (samples)	501	501
Number of signals	121,320	363,960

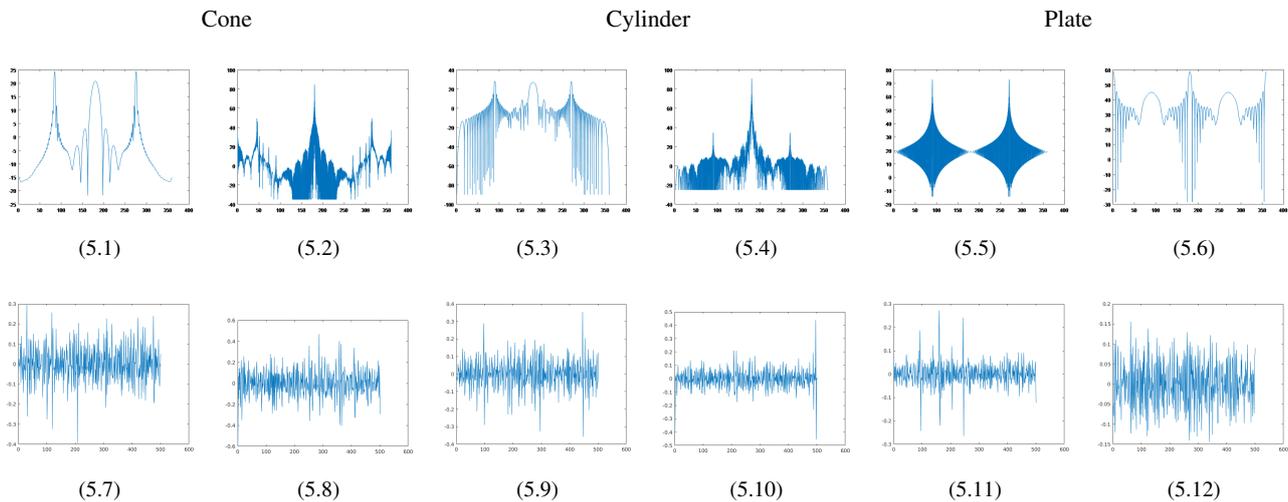
## 4 EXPERIMENTS

Two datasets are created using the methods described. One is used for training, and the other for evaluation/testing. The parameters used to create these dataset are listed in Table 1. The datasets in this paper are named A4 and B4 respectively because they both contain four classes but have different parameter values.

All experiments were run on a Ubuntu 16.04 machine with 32GB of RAM, a Xeon E5-1620 v4 @ 3.5GHz x 8 CPU, a Samsung 860 EVO SSD, and a Nvidia Titan X (Maxwell edition) GPU. The PyTorch and SciPy library versions used for training and evaluation are 0.1 and 1.1 respectively.

### 4.1 Residual Network

Our 1D residual network architecture is inspired from He et al. [10]. Two-dimensional  $3 \times 3$  convolutional filters were replaced by one-dimensional  $3 \times 1$ ,  $5 \times 1$ , and  $7 \times 1$  filters, but the original block module structure and skip connections are maintained. See Figure 6 for a detailed view of the 18-layer network architecture. The residual network was run over 30 epochs and updated using the Adam [34] optimizer with a learning rate of 0.001. Unlike the original implementation of ResNet, batch normalization is done during training to avoid overfitting. The batch size for training is 128 signals for all models except for the 152-layer residual network due to GPU memory constraints and is instead run with a batch size of 32 signals. The learning rate is decayed by 70% if the current validation accuracy does not improve compared to the average of the previous five



**Fig. 5:** There is tremendous variation among the cone, cylinder, and plate RCS signals on the top row. Those signals have rotation about a fixed axis at a relatively slow speed and zero noise. The bottom row features 2 more realistic cone, cylinder, and plate RCS signals. The salient features present in the top examples are now gone.

**Table 2** The number of each respective model in the A5 dataset.

	Cone	Cylinder	Plate	Sphere	Trapezoidal Prism
Train (#)	24,201	25,189	29,041	19,311	2,258
Test (#)	1,871	2,038	2,214	1,623	254
Train (%)	24.2	25.2	29.0	19.3	2.3
Test (%)	23.4	25.4	27.3	20.3	3.2
Models	11	12	30	5	1

validation accuracies. The network with the lowest validation error is saved and used to evaluate the test data. The 18-layer residual network requires five minutes to train while the 152-layer residual network requires nearly three hours to train. The time required to evaluate a signal with the listed hardware is on the order of tens of microseconds, allowing real time signal classification.

#### 4.2 Expanding the A4 Dataset

In secondary tests we expand our four class dataset to include a new trapezoidal prism class. We augment the dataset to answer the question of how our model performance would be affected by the addition of a smaller class of signals. This object is selected such that it closely resembles one of the original classes, i.e. the plate class. One trapezoidal prism class model was created. The new dataset distribution is recorded in Table 2. The number of signals for the new class is significantly lower than the other classes. We call this dataset A5 because it contains the same motion parameters as A4 but has an extra shape class.

#### 4.3 Siamese Network

Our initial hypothesis was that our residual network would misclassify signals belonging to the class with the fewest instances, confusing them with one of the larger classes. If we assume one class will be confused, the loss function will be minimized by misclassifying signals in the smallest class. In order to test our hypothesis, we compare the performance of the residual network with a siamese network. A siamese network consists of two feature extractor modules, each outputting a lower dimensional, compared to the original input, feature vector. The goal of our siamese network is to cluster signals from the same class in close proximity while moving signals from different classes farther apart in feature space. This network is chosen such that the smaller class is less likely to be grouped with another class. The feature extractor modules share the same parameter so that the output vectors can be compared symmetrically. The 18-layer residual networks are used as the feature extractors in the

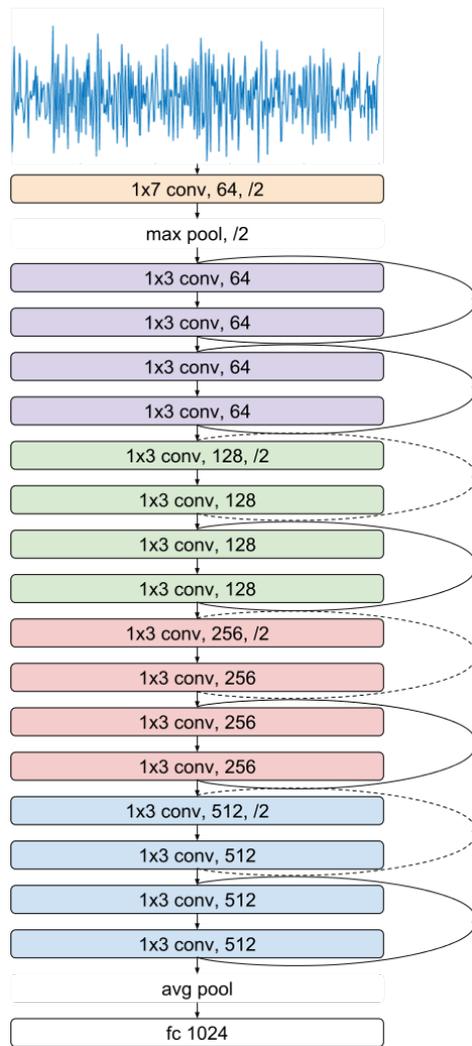
siamese architecture. As with our other trained CNNs, the siamese network is trained using the Adam optimizer with batch sizes of 128 signals for 30 epochs. The learning rate was also initialized and adjusted congruently. The comparator or loss function requires a margin hyperparameter to separate signals of different classes:

$$L = \sum_i^N y^i \cdot \|x_1^i - x_2^i\|_2^2 + (1 - y^i) \cdot \max(0, m - \|x_1^i - x_2^i\|_2^2) \quad (1)$$

The loss function encourages signals in feature space synthesized from the same type of model to converge while forcing signals in feature space belonging to different models farther apart. A CNN generates a fixed length feature representation of the input signal from learned feature extractors. The similarity between feature representations of two signals,  $x_1$  and  $x_2$ , is measured with the  $L_2$  distance metric. The binary label  $y = 1$  if the signals are from the same shape primitive model then, and  $y = 0$  if the signals are not from the same primitive. Signals from the same shape primitives are forced closer in feature space. Whereas, signals from different shape primitives are forced apart if the distance between the feature representations are closer than the margin  $m$ . Since the network requires two signals, evaluation is computed by measuring the similarity between a test signal and a set of signals from the training dataset. Several methods were attempted as classifiers but ultimately a nearest neighbor classifier performed with the greatest accuracy. An input signal first passes through the feature extractor network to produce the corresponding test signal feature vector. The test feature vector is compared to a set of training feature vectors. The most similar feature vector to the test feature vector assigns its label to the test vector. Other methods such as a  $k$ -nearest neighbor with  $k > 1$  and a support vector machine (SVM) were also used but did not perform as well.

#### 4.4 Robustness Test

In order for the classifier to be utilized in real-world applications, it must make accurate predictions on signals with previously unseen distortions. Signal distortions such as occlusion, saturation, and clutter can affect monostatic RCS signals. Occlusion, in this work, is defined as zeroing a subset of a signal's RCS values. Clutter is defined as random amplitude spikes at random locations within a signal. Saturation or clipping is a hard cutoff at a set threshold that limits a signal's amplitude. Subsampling is the removal of a random contiguous section of a signal. Occlusion differs from subsampling because occluded signals have the same number of samples after the distortion is applied unlike signal subsampling. As a robustness

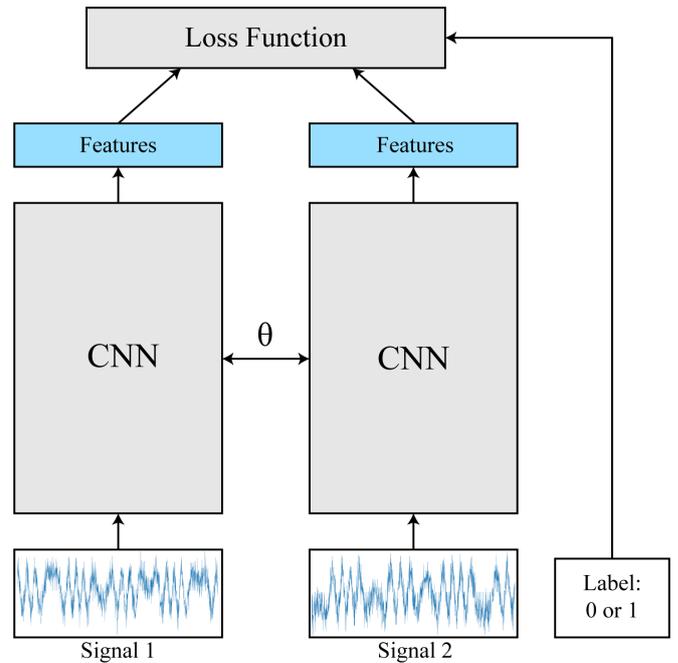


**Fig. 6:** An 18-layer convolutional network is trained to analyze a noisy RCS signal. The architecture is strongly inspired by ResNet [10]. Skip connections are shown as curved arrows. Unlike ResNet, batch normalization is incorporated into the model.

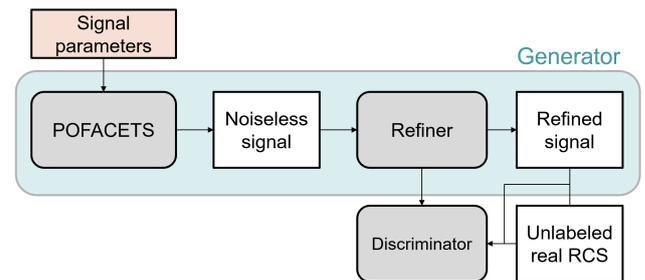
test, the network is trained on dataset A4 which only contains signals distorted by noise and Swerling dropout. The trained network then evaluates a test set of the A4 dataset that is distorted by one of the previously mentioned distortions. The degree of distortion is varied in each test, e.g. the test signals are saturated to 75% of their maximum amplitude. The residual architecture can receive signals of various dimensions as its input because of an average pooling layer before the end of the feature extractor module. Subsampling is implemented by circularly shifting the signal by a random integer and then setting the last  $n$  elements to zero.

#### 4.5 Refiner Network

This section is inspired by the work done by Shrivastava et al. [35], where the authors train a refiner network to make generated images appear more realistic. This network resembles a generative adversarial network (GAN) [8] where a generating network tries to create “realistic” data and a discriminator network decides whether the data is real or fake. The generator network iteratively improves the generated image while the discriminator network learns to more accurately discern the real and fake data apart. Instead of generating data from a noise distribution, as with the classic GAN example, a refiner network converts simulated data into data that more resembles the realistic data. In this work we use a refiner network to make our simulated RCS signals look like simulated signals with added



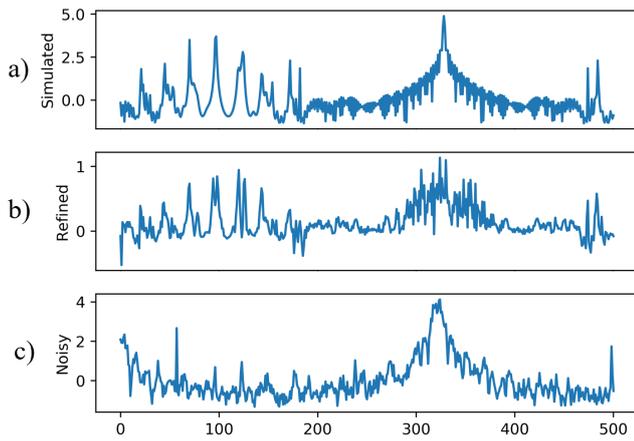
**Fig. 7:** Two signals are fed into two CNNs with shared parameters. The output feature vectors are compared via the Siamese network loss function 1. The target label is equal to one if the two signals belong to the same class and zero otherwise.



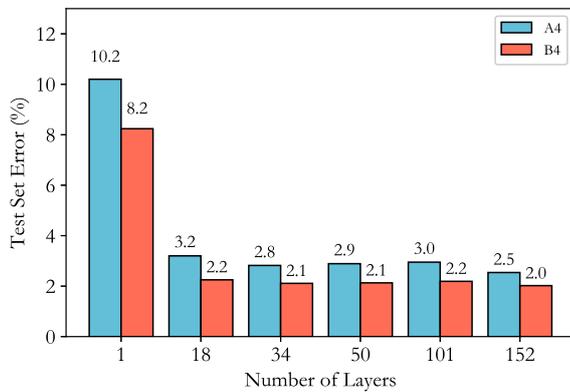
**Fig. 8:** The refiner network and the discriminator work in a similar adversarial manner as a generative adversarial network. The refiner optimizes the simulated signals to look more like the unlabeled realistic data while the discriminator tries to distinguish the difference between the refined and realistic signals.

noise. The refiner network maintains the structure of our signal while adding features to make it appear more like the signals with noise. The parameters used for the simulated dataset are similar to A4 dataset except that no noise is added to the signal and rotation and roll rates are decreased.

The refiner network is a 3-layer CNN that takes a simulated signal as input and outputs a refined signal of the same size. The discriminator network is a 5-layer CNN that receives the refined signal as input and outputs a vector probability map. The probability map determines which parts of the input signal appear realistic to the discriminator. The refiner and discriminator networks have separate loss functions and are trained iteratively. The refiner network’s loss function is a combination of the distance between the input signal and the generated signal and the likelihood that the discriminator believes that the refined signal is real. The discriminator network’s loss function is a combination of the likelihood that the discriminator believes that the refined signal is real and the likelihood that the discriminator is unsure that the real data is real. Both networks are trained for 50 epochs with the Adam optimizer. For each epoch the refiner network is trained twice while the discriminator is only trained once.



**Fig. 9:** The result of the refiner network is shown above. The refiner network takes the signal in (a) as input and returns the signal in (b). It learns to make this transformation by observing signals with noise like the signal in (c).



**Fig. 10:** Several residual networks of different lengths are evaluated on both the A4 and B4 datasets. As the number of layers in the architecture increases the test error on either dataset decreases but only achieves marginal improvement past a depth of 18 layers.

## 5 RESULTS & DISCUSSION

In this section we explore the performance of our trained CNNs on our generated datasets. We also compare different architecture performance using an augmented dataset, investigate the robustness of our classifier, and explore improving our simulated data post-generation.

### 5.1 Classification on A4 and B4 Datasets

Several residual networks with layer depths shown in Figure 10 are trained as described in the experiments section, on both A4 and B4 datasets. Best performance is achieved using the 152-layer residual networks, with classification error scores of 2.5% and 2.0% on datasets A4 and B4 respectively, as shown in Figure 10. While the general trend implies that deeper networks perform better, this is not always true. The 101-layer network performs slightly worse than the 50-layer and 152-layer networks for both the A4 dataset (2.9% vs. 3.0% vs 2.5% for A4) and the B4 dataset (2.1% vs. 2.2% vs. 2.0%). Since all of these networks were trained with the same data, hyperparameters, and appropriately scaled architecture for the given depths, it is difficult to explain this fluctuation in test performance. Test performance saturates for the 18-layer network, and performance changes only slightly for larger networks. As network size

**Table 3** Accuracy performance of support vector machine (SVM), decision tree (DT), single-layer convolutional neural network (CNN), and residual network (RN) algorithms on the A4 dataset. The leftmost column represents the signal features that were used by each classification algorithm. Common signal statistics (SS) represents feature vectors comprised of the mean, standard deviation, and extremum of a signal. Transform representations (TR) represent feature vectors comprised of coefficients from the Fourier and Wavelet transforms of a signal. Since convolutional neural networks learn a feature representation, only the original signals are used as input.

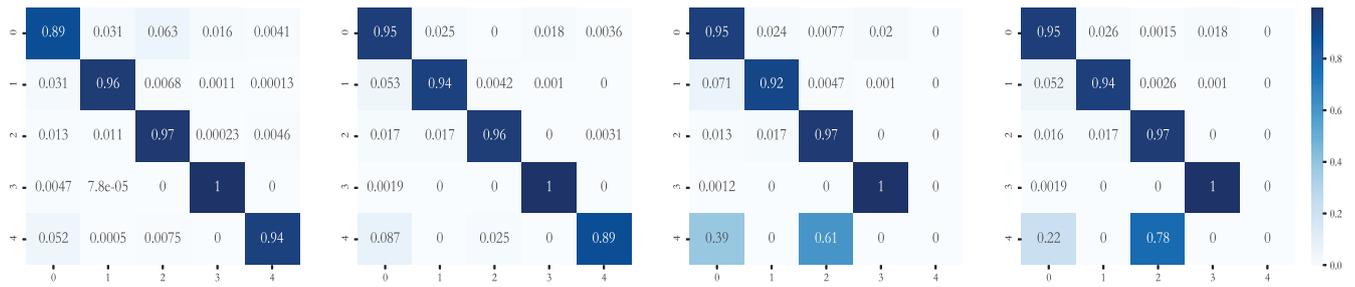
	SVM	DT	1L-CNN	RN18	RN152
Original	0.557	0.485	0.898	0.968	0.975
SS	0.863	0.849	-	-	-
TR	0.707	0.607	-	-	-
TR+SS	0.742	0.828	-	-	-

increases, so does the ability to learn more complex features. But larger networks also have a propensity to overfit if the dataset used for training is not sufficiently large and representative of the distribution of each class. When overfitting occurs, training accuracy will continue to improve while test accuracy continues to degrade. Since Figure 10 features test error, and the A4 and B4 datasets are sufficiently large, the networks are likely not overfit, but at saturation for test accuracy given the complexity of useful signal features. Likely, the small deviations in test performance stem from each network converging on different local minima in the optimization plane. Initial conditions and when training is stopped may have effects on which minima a network is likely to converge on.

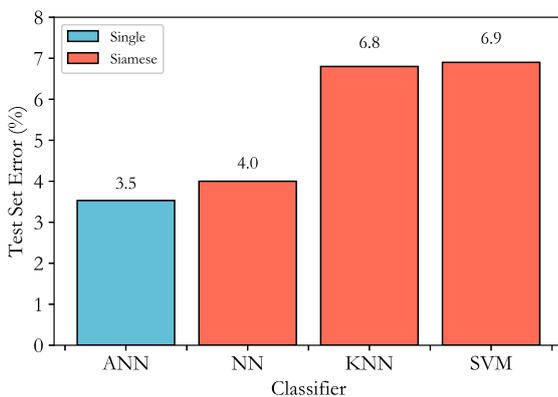
Models trained on the B4 dataset perform better than models trained on the A4 dataset across all network depths. As a baseline, a neural network and non-residual convolutional neural network were trained and evaluated on the A4 dataset with the corresponding test errors, 29.5% and 6.1%. The neural network contains six layers, dropout, and non-linear layers. Increasing the number of layers in the neural network did not significantly improve results. The non-residual convolutional neural network contained 18 layers and is trained with the same training parameters described in the experiments section. When the number of layers in the non-residual convolutional network was increased, performance plateaued and then began to degrade.

Our classification results for the residual networks may appear counter-intuitive at first glance, since CNNs typically perform worse on datasets that have more variation. Datasets with more variation are simply more difficult to learn because the CNN will have to learn specific filters to deal with that variation. Not only does the B4 dataset contain more signals but it contains faster roll and tumble rates. The faster roll and tumble rates for our signals actually increases the amount of information per sample because the models we use to generate our signals have large distinct edges and smooth surfaces. If instead the models used had rough surfaces and less distinct edges, information would be lost by increasing the roll and tumble rates. The B4 dataset also contains signals with lower SNR rates and more varied viewing angles, which decrease the amount of information within the signals. Regardless of the size of the network, test performance on the B4 dataset was greater than on the A4 dataset. It was for this reason that the A4 dataset was selected to create new datasets and to further train/test our models. If a more difficult dataset is used, then there will be a clearer distinction between the results of more advanced networks.

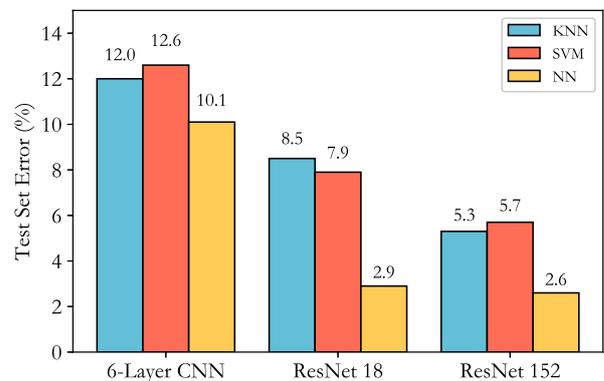
In addition to neural networks, we assess the performance of other machine learning classification algorithms such as support vector machines (SVM) and decision trees (DT) on the A4 dataset. The SVM algorithm utilizes the radial basis function kernel with a gamma value equal to reciprocal of the number of input features. Multiple one-against-one classifiers are aggregated to form the final SVM classifier. As for the DT, the Gini criterion is used to measure the quality of the split in the tree and decision nodes are randomly chosen to be further split. The minimum number of samples to be a leaf node is set to five, and the minimum number of samples required to split a decision node is two. Unlike deep learning algorithms, features must be manually crafted for the SVM and DT classifiers to



**Fig. 11:** The confusion matrices for all siamese networks and the single residual network. Confusion matrices starting from the left to the right belong to the single network, the siamese network with nearest neighbor, siamese network with k nearest neighbor, and siamese network with support vector machine. The classes are enumerated as (0) cone, (1) cylinder, (2) plate, (3) spheroid, and (4) trapezoidal prism.



**Fig. 12:** A single residual network’s performance on the A5 test dataset is compared to the performance of three Siamese networks with various classification layers. ANN stands for artificial neural network, NN is nearest neighbor, kNN is k nearest neighbor, and SVM stands for support vector machine.



**Fig. 13:** Three different classifier modules are compared after a CNN feature extractor of varied depths. The nearest neighbor classifier achieves the highest overall accuracy consistently across all architectures tested.

attain optimal performance. For comparison the SVM and DT classifiers are trained and evaluated on the complete length signals from the A4 dataset and achieve accuracies of 55.7% and 48.5% respectively, as shown in Table 3. Common signal statistics (SS) such as minimum and maximum are combined with low order cumulants [36] to form a representation of the RCS signals. This representation improves upon the previous the accuracy of the classifiers to 86.3% and 84.9%. Following the work of Byl [37] and Zhang [38], more complex descriptive features such as Fourier Transform frequency responses and Wavelet Transform coefficients are used to represent the signals. Specifically, the Fast Fourier Transform generates frequency coefficients and the Discrete Wavelet Transform (DWT) symmetrically pads signals during the transform in order to avoid inaccurate calculation of the DWT. The first 50 coefficients from each transform are concatenated to form the feature vector representation. This method, which we call transform representations (TR), is combined with the SS features to achieve accuracies of 74.2% and 82.8%. For reference our one layer CNN (1L-CNN) has a test accuracy of 89.8% on the A4 dataset, Figure 10.

## 5.2 Classification on A5 Dataset

The siamese network structure has been used on a variety of tasks such as signature matching and facial identification with high performance [39, 40]. This type of network performs most effectively when the number of classes in a dataset is large and the number of data per class is relatively low. The architecture’s unique comparator function forces input from the same class to cluster in high dimensional space and input from different classes to be farther apart in

high dimensional space. The loss function for a typical CNN classifier is the negative log likelihood function which does not contain any constraint on how far apart the output vectors of the feature extractor module are. The A5 dataset contains the same set of parameters as A4 but includes an additional geometric model of trapezoidal prism. The additional class contains only one model and makes up a small portion of the total signals in the A5 dataset.

The A5 dataset is a superset of the A4 dataset, but augmented with an additional and easily-confused shape class. The results of this experiment are shown in Table 4. The single residual network outperforms all types of the siamese networks in terms of overall accuracy as shown in Figure 12. Initially it appears that the lack of clustering term in the objective function does not reduce performance on the A5 dataset, however the CNN could maintain high accuracy even while misclassifying all of the signals in the newest class. To further investigate this result the precision, recall, and the F1-score of each class is calculated and shown in Table 4. The siamese networks with the k-nearest neighbor and support vector machine classifiers misclassified the trapezoidal prism class in every case. The single residual network and the siamese network with the nearest neighbor classifier were both able to correctly classify the trapezoidal prism class a majority of the time.

In Table 4 we can see that the F1-score for the trapezoidal class is greater in the single network section than the siamese network section. Overall the average F1-score across classes is 0.948 and 0.956 for the single network and siamese network respectively. If we weigh the F1-score by the number of signals per class there is an even larger difference in performance. The weighted F1-score of

**Table 4** Accuracy performance comparison between a single residual network and a siamese network with a NN classifier on the A5 dataset

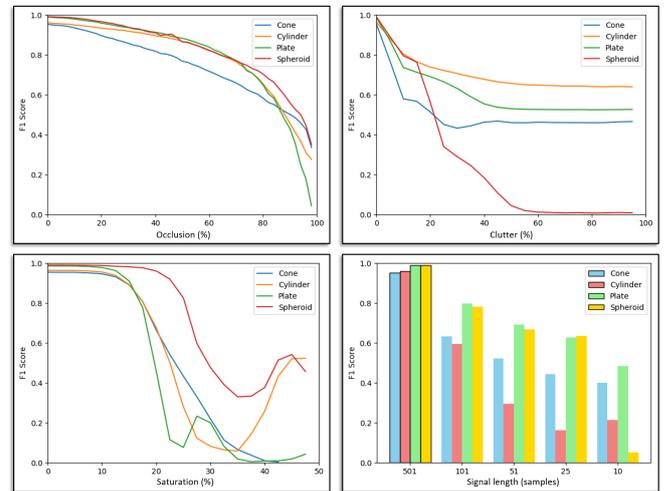
Class	Single Network			Siamese Network+NN		
	Precision	Recall	F1-Score	Precision	Recall	F1-Score
Cone	0.94	0.89	0.91	0.92	0.95	0.94
Cylinder	0.96	0.96	0.96	0.95	0.94	0.95
Plate	0.94	0.97	0.95	0.99	0.96	0.98
Spheroid	0.98	1.00	0.99	0.98	1.00	0.99
Trapezoidal Prism	0.93	0.94	0.93	0.95	0.89	0.92

the single network and siamese network are 0.947 and 0.959 respectively. It appears that the single network showed high performance on the trapezoidal prism class because it misclassified more of the signals in the cone class. The siamese network with the nearest neighbor classifier performs well because the feature extractor module is better able to separate the clusters for each class. Intuitively we expect the k nearest neighbor and support vector machine classifiers to outperform the nearest neighbor classifier, but our results in Figure 13 suggest otherwise. The dimensionality of the output vector from the feature extractor module may be a potential reason that the nearest neighbor classifier performs better. As the number of dimensions increase, the k nearest neighbor algorithm tends to perform worse due to the increasing space in between points.

### 5.3 Robustness Metric Performance

A CNN classifier's ability to handle noisy input data can be evaluated in multiple ways, such as testing on a novel set of data with distortions seen in the training data or testing on a novel set of data with distortions unseen in the training data. Monostatic radar signals can have a variety of distortions in real applications such as signal occlusion, clutter, sensor saturation, subsampling, or a combination of several. Since generating a dataset with every combination of signal distortions is unwieldy, we instead decide to evaluate our system's robustness to distortions by evaluating our model on data with distortions not seen in the training data. The results shown in Figure 14 are the F1-score per class from a single 18-layer residual network. However, the robustness results for networks with more layers is nearly identical and not presented. The evaluation set was generated via the method described in the experiments section.

The network performs remarkably well on signals that have been occluded by even 75% of the total signal, even though no dropout layers are used to train the model. Occlusion may not affect our network significantly because the rotation rates used in our dataset generation are relatively large and occasionally the shape model is rotating several times within the full window of sampling. Even if the signal is occluded significantly, some signals with high rotation rates may contain enough information for classification. However signals generated with slower rotation rate parameters do not appear to complete rotations multiple times within a full window. For these cases the CNN is able to discern the object within a limited viewing window. The CNN is however very sensitive to signal clutter, accuracy-per-class drops as soon as clutter is introduced. Clutter in this work is the addition of random peaks in a signal and CNNs are sensitive to slight distortions to input data. This distortion is similar to the distortion created by adversarial attacks such as FGSM [41], except that we are adding distortions with random amplitudes at random locations. Most CNNs are not robust to adversarial attacks and it appears that clutter approximates an adversarial attack in this domain. The CNN is resilient to signal saturation up to roughly 15%, then performance decreases significantly soon after. Signals with heavy saturation begin to appear indistinguishable from each other, and the filters that the CNN uses to detect features cannot distinguish between each class. The rise in F1-score of some of the classes seems to be an artifact of the dataset instead of a feature of the network. The final distortion is subsampling the input signal. This measure is similar to the occlusion distortion but the number of total samples in the signal do not change in the occlusion distortion. The results of subsampling show that the CNN can use signals with lengths as small as 25 samples as input and achieve a reasonable F1-score. The performance halves when input size is 5% of its



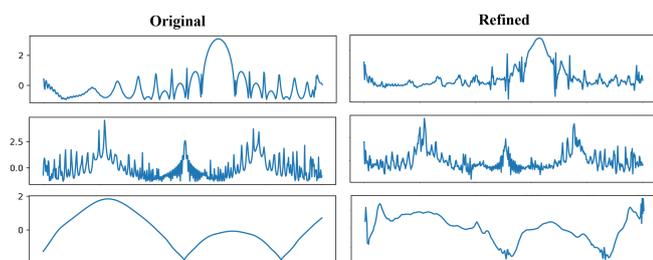
**Fig. 14:** The single 18 layered residual network's robustness performance is shown for several novel distortions. This benchmark is a way to compare a network robustness to realistic signal distortions found RCS systems. Signal occlusion, clutter, saturation, and subsampling are the realistic distortions used for this benchmark.

original length. The siamese network evaluated with the robustness metric is not included because the previously mentioned siamese testing method compares an input signal to a subset of the training data. Since the training data does not contain the distortions of the evaluation data, unsurprisingly, the siamese network performs very poorly.

### 5.4 Classification on Refined Dataset

In order to compare the difference between the simulated dataset and the refined dataset we train separate three layered convolutional neural networks. The network's performance was evaluated by classifying simulated signals with added white Gaussian noise. The simulated signals with added noise were also used as "real" data in the refiner network training. Overall the model's performance on the evaluation dataset is greater when the model is trained using the refined dataset by 3.5%. The accuracy of the network trained on the simulated subset A4 dataset is 86.7%, while the accuracy of the network trained on the refined dataset was 90.2%.

No simulator can perfectly model the all of the nuances and variables that are required to create real data. Therefore training a CNN on simulated data typically does not perform well on real data. This does not mean that networks should be trained with only real data because representative real data is difficult and expensive to obtain. Real data is also potentially biased in terms of only representing certain occurrences and typically few variables are able to be controlled when creating real datasets. Simulated data is useful because very large datasets can be generated easily. Adjustments can be made one variable at a time and all parameters used to create that data is known at every timestep. The generative CNN called the refiner network described in Section 4 makes simulated data appear more like real data, shown in Figure 9. Using the refined data to train a small network on a subset of our A4 dataset results in a 3.5% accuracy improvement over training using the equivalent simulated data. For that test the only "realistic" feature added to the "real" data was



**Fig. 15:** Some examples of signals pre and post refinement. The structure of the signal is maintained but pseudo noise is added to the original signal from the refiner network.

noise. In Figure 9 we see that the refined signal seemingly adds noise to the simulated signal but maintains the structural elements of the signal.

## 6 Conclusion

To the best of our knowledge, we are the first to train convolutional neural networks to classify object shape from monostatic radar signals. We expand upon the MATLAB library POFacets to generate large datasets with a variety of selected parameters. Realistic motion, added noise, and Swerling dropout enhance the initial simulation generation. Utilizing the latest in deep learning architecture we create a 1D residual network capable of achieving test error results as low as 2-2.5% on our generated datasets. Our A4 dataset is augmented with an additional test and then evaluated with a siamese network architecture. The siamese CNN does perform as well in terms of accuracy but surpasses the performance of the single residual network in terms of average F1-score. The robustness of our CNN is then evaluated on signals with previously unseen realistic distortions. The single residual network performs well on signals with occlusion and subsampling but performs poorly on signals with clutter and saturation. We explored increasing the quality of the simulated signals using a state of the art refiner network. Deep learning models trained on the refined signals outperform models trained on the original simulated data.

## ACKNOWLEDGMENT

This project was funded by Lockheed Martin. We would like to thank Rowland Escritor and Kevin Vance for their substantial efforts coordinating and facilitating this project.

## 7 References

- 1 Swerling, P.: 'Probability of detection for fluctuating targets', *IRE Transactions on Information Theory*, 1960, pp. 269-308
- 2 Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., et al.: 'Automatic differentiation in pytorch', 2017
- 3 Jones, E., Oliphant, T., Peterson, P.: '{SciPy}: open source scientific tools for {Python}', 2014
- 4 Chatzigeorgiadis, F.: 'Development of code for a physical optics radar cross section prediction and analysis application'. Monterey California. Naval Postgraduate School, 2004
- 5 Touzopoulos, P., Boviatsis, D., Zikidis, K.C.: '3d modelling of potential targets for the purpose of radar cross section (rcs) prediction: Based on 2d images and open source data', *Military Technologies (ICMT)*, 2017 International Conference on, IEEE, 2017, pp. 636-642
- 6 Krizhevsky, A., Sutskever, I., Hinton, G.E.: 'Imagenet classification with deep convolutional neural networks', *Advances in neural information processing systems*, 2012, pp. 1097-1105
- 7 Simonyan, K., Zisserman, A.: 'Very deep convolutional networks for large-scale image recognition', *arXiv preprint arXiv:14091556*, 2014
- 8 Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al.: 'Generative adversarial nets', *Advances in neural information processing systems*, 2014, pp. 2672-2680
- 9 Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., et al.: 'Going deeper with convolutions', *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1-9

- 10 He, K., Zhang, X., Ren, S., Sun, J.: 'Deep residual learning for image recognition', *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770-778
- 11 Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: 'Densely connected convolutional networks', *arXiv preprint arXiv:160806993*, 2016
- 12 Kouemou, G.: 'Radar target classification technologies'. In: *Radar Technology*. (InTech, 2010).
- 13 Lang, K.J., Waibel, A.H., Hinton, G.E.: 'A time-delay neural network architecture for isolated word recognition', *Neural networks*, 1990, pp. 23-43
- 14 Anderson, J.A., Gately, M.T., Penz, P.A., Collins, D.R.: 'Radar signal categorization using a neural network', *Proceedings of the IEEE*, 1990, pp. 1646-1657
- 15 Jordanov, I., Petrov, N.: 'Sets with incomplete and missing datann radar signal classification', *Neural Networks (IJCNN)*, 2014 International Joint Conference on, IEEE, 2014, pp. 218-224
- 16 Jordanov, I., Petrov, N., Petrozziello, A.: 'Supervised radar signal classification', *Neural Networks (IJCNN)*, 2016 International Joint Conference on, IEEE, 2016, pp. 1464-1471
- 17 Soto, A., Mendoza, A., Flores, B.C.: 'Optimization of neural network architecture for classification of radar jamming fm signals', *Radar Sensor Technology XXI*, International Society for Optics; Photonics, 2017, p. 101881H
- 18 Mendoza, A., Soto, A., Flores, B.C.: 'Classification of radar jammer fm signals using a neural network', *Radar Sensor Technology XXI*, International Society for Optics; Photonics, 2017, p. 101881G
- 19 Hara, Y., Atkins, R.G., Yueh, S.H., Shin, R.T., Kong, J.A.: 'Application of neural networks to radar image classification', *IEEE Transactions on Geoscience and Remote Sensing*, 1994, pp. 100-109
- 20 Zhang, Y., Wu, L.: 'Crop classification by forward neural network with adaptive chaotic particle swarm optimization', *Sensors*, 2011, pp. 4721-4743
- 21 Park, S., Hwang, J.P., Kim, E., Lee, H., Jung, H.G.: 'A neural network approach to target classification for active safety system using microwave radar', *Expert Systems with Applications*, 2010, pp. 2340-2346
- 22 Rahman, A., Yavari, E., Lubecke, V.M., Lubecke, O.B.: 'Noncontact doppler radar unique identification system using neural network classifier on life signs', *Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireless)*, 2016 IEEE Topical Conference on, IEEE, 2016, pp. 46-48
- 23 Szymczyk, P., Szymczyk, M.: 'Classification of geological structure using ground penetrating radar and laplace transform artificial neural networks', *Neurocomputing*, 2015, pp. 354-362
- 24 Kabakchiev, C., Behar, V., Garvanov, I., Kabakchieva, D., Rohling, H.: 'Detection, parametric imaging and classification of very small marine targets emerged in heavy sea clutter utilizing gps-based forward scattering radar', *Acoustics, Speech and Signal Processing (ICASSP)*, 2014 IEEE International Conference on, IEEE, 2014, pp. 793-797
- 25 Lundén, J., Koivunen, V.: 'Deep learning for hrrp-based target recognition in multistatic radar systems', *Radar Conference (RadarConf)*, 2016 IEEE, IEEE, 2016, pp. 1-6
- 26 Morgan, D.A.: 'Deep convolutional neural networks for atr from sar imagery', *Proceedings of the Algorithms for Synthetic Aperture Radar Imagery XXII, Baltimore, MD, USA*, 2015, p. 94750F
- 27 Kim, Y., Moon, T.: 'Human detection and activity classification based on micro-doppler signatures using deep convolutional neural networks', *IEEE Geoscience and Remote Sensing Letters*, 2016, pp. 8-12
- 28 Gong, M., Zhao, J., Liu, J., Miao, Q., Jiao, L.: 'Change detection in synthetic aperture radar images based on deep neural networks', *IEEE transactions on neural networks and learning systems*, 2016, pp. 125-138
- 29 Mason, E., Yonel, B., Yazici, B.: 'Deep learning for radar', *Radar Conference (RadarConf)*, 2017 IEEE, IEEE, 2017, pp. 1703-1708
- 30 Zhang, X., LeCun, Y.: 'Text understanding from scratch', *arXiv preprint arXiv:150201710*, 2015
- 31 Kiranyaz, S., Ince, T., Gabbouj, M.: 'Real-time patient-specific ecg classification by 1-d convolutional neural networks', *IEEE Transactions on Biomedical Engineering*, 2016, pp. 664-675
- 32 Stinco, P., Greco, M.S., Gini, F., LaManna, M.: 'Non-cooperative target recognition in multistatic radar systems', *IET Radar, Sonar & Navigation*, 2013, pp. 396-405
- 33 Mathews, Z., Quiriconi, L., Böniger, U., Schüpbach, C., Weber, P.: 'Learning multi-static contextual target signatures', *Radar Conference (RadarConf)*, 2017 IEEE, IEEE, 2017, pp. 1568-1572
- 34 Kingma, D.P., Ba, J.: 'Adam: A method for stochastic optimization', *arXiv preprint arXiv:1412.6980*, 2014
- 35 Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: 'Learning from simulated and unsupervised images through adversarial training', *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, p. 6
- 36 Xin, Z., Ying, W., Bin, Y.: 'Signal classification method based on support vector machine and high-order cumulants', *Wireless Sensor Network*, 2010, p. 48
- 37 Byl, M.F., Demers, J.T., Rietman, E.A.: 'Using a kernel adaptor for object classification with rcs data', *arXiv preprint arXiv:10055337*, 2010
- 38 Zhang, L., Zhou, W., Jiao, L.: 'Wavelet support vector machine', *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 2004, pp. 34-39
- 39 Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: 'Signature verification using a "siamese" time delay neural network', *Advances in Neural Information Processing Systems*, 1994, pp. 737-744
- 40 Sun, Y., Wang, X., Tang, X.: 'Deep learning face representation from predicting 10,000 classes', *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1891-1898
- 41 Goodfellow, I.J., Shlens, J., Szegedy, C.: 'Explaining and harnessing adversarial examples', *arXiv preprint arXiv:1412.6572*, 2014