

Neural Data Augmentation Techniques for Time Series Data and its Benefits

Anindya Sarkar
IIT Hyderabad, INDIA
anindyasarkar.ece@gmail.com

Anirudh Sunder Raj
IIT Madras, INDIA
cs17b003@cse.iitm.ac.in

Raghu Sesha Iyengar
IIT Hyderabad, INDIA
raghu.sesha@gmail.com

Abstract—Exploring adversarial attacks and studying their effects on machine learning algorithms has been of interest to researchers. Deep neural networks working with time series data have received lesser interest compared to their image counterparts in this context. In a recent finding, it has been revealed that current state-of-the-art deep learning time series classifiers are vulnerable to adversarial attacks. In this paper, we introduce neural data augmentation techniques and show that classifier trained with such augmented data obtains state-of-the-art classification accuracy as well as adversarial accuracy against Fast Gradient Sign Method (FGSM) and Basic Iterative Method (BIM) on various time series benchmarks.

Index Terms—time series classification, adversarial training, gradient based adversarial attacks

I. INTRODUCTION AND RELATED WORK

Deep Neural Networks have displayed impressive results on many machine learning tasks on image ([1]–[5]), natural language processing ([6]–[9]) and time series classification ([10]–[13]). However, their fragility to small adversarial perturbations is a matter of concern for researchers. A targeted black-box attack on neural networks for image classification was formulated as an optimization problem in [14], and further improved in [15]. A targeted l_0 norm attack discussed in [16], aims to minimize the number of modified pixels in an image to cause mis-classification as a particular target class. Adversarial attacks intended to lower reliability of neural networks are also explored. Of these, gradient based l_∞ norm attacks such as [17] and [18] are very popular. Various techniques to understand and mitigate effects of adversarial perturbations have also been studied [19], [21], [22], [23], [20], [47]. An excellent review of adversarial attacks on machine learning systems can be found in [26]. The reliability and security concerns raised by adversarial attacks have been one of the main reasons for deep neural networks not yet becoming popular with safety critical applications where the cost of failure is high.

Time series data is omnipresent and classification tasks on time series data finds its applications in health care ([27]), power consumption monitoring ([28]), food safety ([29], [30]), security ([31]) etc. Current state-of-the-art deep neural networks can achieve impressive performance at classifying time series data ([10]–[12]) on various datasets ([32], [33]). However, these networks suffer in the same way to adversarial

inputs as their image counterparts. A recent finding ([34]) shows that vulnerability of state-of-the-art time series classification networks to simple adversarial attacks, such as Fast Gradient Sign Method (FGSM) and Basic Iterative Method (BIM), bring back the focus on building more robust time series classifiers.

FGSM [17] uses the gradient of the loss with respect to the input data, then adjusts the input data to maximize the loss. This can be summarised using the following expression:

$$z_{adv} = z + \epsilon * \text{sign}(\nabla_z J(\theta, z, y))$$

where original input data and its corresponding output label are represented as z and y respectively, θ denotes parameters of the classifier, $J(\theta, z, y)$ denotes the loss or objective function and the obtained adversarial data is represented as z_{adv} . Note that ϵ is used as a multiplier to ensure small perturbations. BIM [18] extends FGSM by applying it iteratively with a small step size and clip the obtained time series elements after each step to ensure that they are in an ϵ -neighborhood of the original input data.

It has been shown ([35], [36]) that data augmentation increases the size and diversity of the training set resulting in improved classification accuracy on time series data. Further, in order to achieve adversarial robustness, the classifiers should be robust to noise-corrupted data ([37]). In this paper, we propose data augmentation techniques for time series data, which helps in improving the robustness of the classifier against adversarial attacks. Our contributions can be summarized as follows:

- 1) We propose **Input gradient based Data Augmentation method**. We show that classifier trained with this augmented dataset achieves state-of-the-art classification accuracy on adversarially perturbed UCR time series datasets using FGSM and BIM attacks.
- 2) We also propose **Output gradient based Data Augmentation method**. We show that classifier trained with this augmented dataset further improve the state-of-the-art classification accuracy on adversarially perturbed UCR time series datasets using FGSM and BIM attacks.
- 3) Additionally, We propose a **spectral density based Data Augmentation method**. Our experimental finding suggests that, combination of Output gradient based

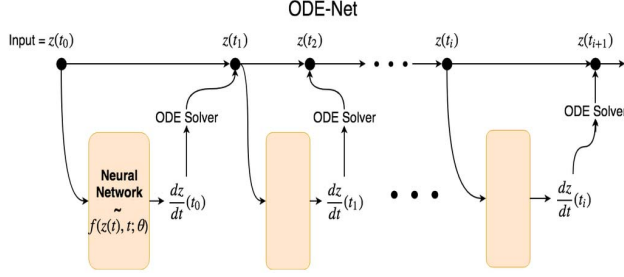


Fig. 1. ODENet building blocks. Parameters of ODENet are shared across times.

Data Augmentation method and Feature Similarity based regularized training method achieves the best classification accuracy on standard adversarial attacks such as FGSM and BIM across different UCR time series dataset.

For models to be useful in the real world, they need to be both accurate on a held-out set of time series data, which we refer to as *clean accuracy*, and robust on corrupted time series, which we refer to as *robustness*. It is believed that there exists a fundamental trade-off between the two ([38]). Our observation has been that though the trade-off exists, it is possible to build robust systems with very minimal or no drop in clean accuracy.

II. GRADIENT BASED AUGMENTATION TECHNIQUES

In the time series classification paradigm, neural networks are maximum likelihood estimators. The neural network tends to learn the input features that are important for classification, even if those features look incomprehensible to humans. It has been shown that simple gradient based perturbations in the input signal can cause a trained network to misclassify ([17], [18]). We first assume a threat model (black-box or white box) and generate adversarial samples based on that model. We augment original train data with adversarial samples and then train a reference neural network with this data. We evaluate the trained network on standard adversarial attacks on time series data ([34]).

A. Input gradient based Data Augmentation and Training method

STEP-1: Neural networks were used to model Ordinary Differential Equations(ODE) ([39]). We define a time-series $z(t)$ as a sequence of 1-D values $z(t_0), z(t_1), \dots, z(t_n)$. We model the continuous dynamics of the input time series using an ODE specified by a neural network(f), parameterized by θ :

$$\frac{dz(t)}{dt} = f(z(t), t, \theta) \quad (1)$$

We refer such neural network as **ODENet** as shown in fig.1. We use ODE solver as a sequential generative model to predict

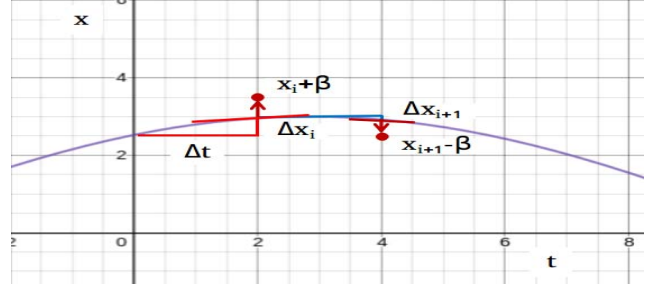


Fig. 2. Input gradient based adversarial sample generation. Each point in the timeseries is perturbed by a random number in range $[0, \epsilon]$. The direction of change is determined by the local gradient of x .

successive samples of the time-series. More specifically, we compute $z(t_i)$ from known $z(t_{i-1})$ as:

$$z(t_i) = z(t_{i-1}) + \int_{t_{i-1}}^{t_i} f(z(t), t, \theta) dt \quad (2)$$

We predict the values of $z(t_1), z(t_2), \dots, z(t_n)$ by iteratively following 2. The value predicted at time step t_i is treated as the initial value for the computation of $z(t_{i+1})$. We minimize MSE loss between predicted and true sequence, in order to find the optimum parameter set of ODENet, θ_{opt} , which can capture the continuous dynamics of $z(t)$ at any given point as follows:

$$\frac{dz(t_i)}{dt} = f(z(t_i), t_i, \theta_{opt}) \quad (3)$$

Adam optimizer with a learning rate of 0.0003 and weight decay factor of 10^{-3} was used for training the ODENet.

STEP-2: Once the ODENet is trained, we use it to compute the gradient of a timeseries at any given time point using eq. 3. we add carefully crafted perturbation guided by the direction of gradient at each time point, which allows to preserve the overall structure of the original input time series.

Accordingly the augmented samples are generated using the below equation:

$$z_{aug}(t_i) = z(t_i) + clamp\left(\frac{dz(t_i)}{dt}\right) \quad (4)$$

$$clamp(z) = \begin{cases} z & |z| < \beta \\ \beta & z > \beta \\ -\beta & z < -\beta \end{cases} \quad (5)$$

$z_{aug}(t_i)$ and $z(t_i)$ are respectively augmented timeseries and original timeseries at timestep t_i . β is a small positive constant. Note that by perturbing the samples as given in equation (4), the change in magnitude of each sample in the augmented timeseries has an upper limit of β compared to the original timeseries. This inherently sets an upper limit on the change in gradient to $2 * \beta / \delta t$ as shown in fig 2. Figure 3(a),(b) shows example timeseries data augmented using this technique. We term this augmentation method as **In-Clamp-Grad**.

Training Objective of Classification Network: We define a multiclass classification setup where the input-label pairs

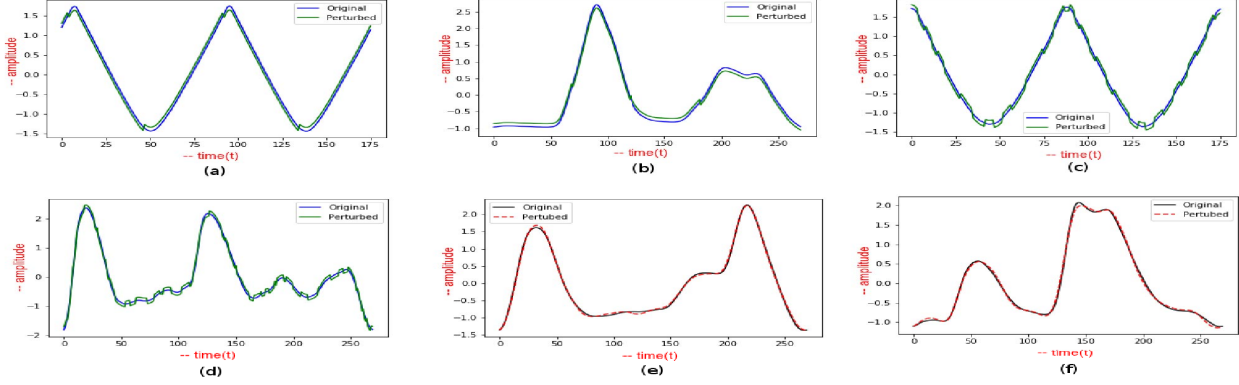


Fig. 3. Two randomly selected samples and their corresponding adversarial samples obtained using (a),(b) *Input gradient*, (c),(d) *Output gradient*, (e),(f) *Spectral Density* based adversarial sample generation techniques. (a),(c) and (e) are from Adiac time series dataset while (b),(d) and (f) from 50words dataset.

$(z, y) \in (\mathcal{X} \times \mathcal{Y})$ are sampled from data distribution D . A neural network is trained as a classifier whose goal is to predict the class label y for a given input z . The training objective of the classifier is to find parameters ϕ that minimize the multiclass cross entropy loss:

$$\min_{\phi} \mathbb{E}_{(z_{aug}, y)} [\text{CELoss}(g_{\phi}(z_{aug}, y))] \quad (6)$$

Here, g_{ϕ} represents the neural network classifier with parameters ϕ . We augment the input dataset with the perturbed examples obtained using equation (4). The labels for the perturbed examples are set to be same as the original example from which it was obtained. Thus, training the neural network classifier with this augmented dataset follows 6. We used Adam optimizer with a learning rate of 0.0002 with a weight decay of 10^{-3} for all the experiments involving training the classifier network in this paper. We denote (z_{aug}, y) as input-label pairs of augmented dataset.

We also experimented with a slight modification of equation (4) as given below:

$$z_{aug}(t_i) = z(t_i) + \epsilon * \text{sign}\left(\frac{dz(t_i)}{dt}\right) \quad (7)$$

Here, ϵ is a random value between 0.0 and 0.33. We term this augmentation method as **In-Sign-Grad**.

B. Output gradient based Data Augmentation Method

In this case, we assume the adversary has knowledge of the model and its parameters. We perturb the input samples based on the derivatives of classifier output w.r.t input. If $g(\cdot)$ is the transformation applied by the neural network and z is the input timeseries data, then the derivative of classifier output w.r.t the time variable is governed by the following chain rule:

$$\frac{dy}{dt} = \frac{dg(z)}{dz} \frac{dz}{dt} \quad (8)$$

As done in section II-A, $\frac{dz}{dt}$ is computed using an ODENet. $\frac{dg(z)}{dz}$ is computed using the automatic differentiation in Py-

torch. Using the derivative of classifier output w.r.t the time variable, we define a perturbed time series as below:

$$z_{aug}(t) = z(t) + \text{sign}\left(\frac{dg(z)}{dz}\right) \text{abs}\left(\text{clamp}\left(\frac{dz}{dt}\right)\right) \quad (9)$$

where, $\text{sign}\left(\frac{dg(z)}{dz}\right)$ is an indication of the direction of maximum change in classifier output w.r.t input. For cross entropy loss, changing the input along this direction maximizes the loss term. $\frac{dz}{dt}$ is the gradient of input w.r.t time. The function clamp is defined in equation (5). Thus, the direction of perturbation is defined by $\left(\frac{dg(z)}{dz}\right)$, while the magnitude of perturbation is defined by $\frac{dz}{dt}$ and it is limited by a small positive number (β), which helps in preserving the overall structure of original time series. As done in section II-A, we augmented the input dataset with perturbed samples and trained a neural network for classification. The trained neural network achieves state-of-the-art results on different UCR time series test datasets with FGSM and BIM perturbations. Figure 3(c),(d) shows example timeseries data augmented using this technique. We term this augmentation method as **Out-Sign-Grad**.

Note that in figure 3, sample generated using output gradient based technique has much more fluctuation compared to sample generated using input gradient based technique. This is intuitive, because in Out-Sign-Grad technique, at each time point, perturbation direction is guided by gradient of classifier output w.r.t input at that point. Hence, such perturbation can behave in arbitrary manner and have no relation with original time series. But in In-Sign-Grad technique, the perturbation direction is guided by the “learned” gradient of the original input time series at that time point. As a result, input gradient based augmented time series preserves the shape of original input time series better compared to output gradient based augmented time series.

To validate the significance of $\text{sign}\left(\frac{dg(z)}{dz}\right)$ in the generation of perturbed samples, we replaced this with random sign for each time step.

$$z_{aug}(t) = z(t) + \text{random}([-1, +1]) \text{abs}(\text{clamp}(\frac{dz}{dt})) \quad (10)$$

A neural network trained with such augmented data, which we termed as **Rand-Grad**, was found to provide lesser accuracy compared to the network trained using the augmented data obtained using **Out-Sign-Grad**.

C. Spectral density based Data Augmentation method

In the techniques described in previous sections, we introduced perturbations at each timestep of the input timeseries. The perturbations added were functions of derivatives of inputs w.r.t the time and/or derivatives of classifier output w.r.t the input. In this section, we describe the perturbations added based on the frequency domain characteristics of the input signals. The energy of a signal is given by Parseval's theorem:

$$E(z) = \sum_{n=0}^{N-1} |z[n]|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |Z[k]|^2 \quad (11)$$

Where, $Z[k]$ is the Discrete Fourier Transform of $z[n]$, both of length N .

We now obtain a new perturbed Discrete Fourier Transform $Z_{aug}[k]$ as described below. The sequence $Z[k]$ is sorted in descending order to obtain Z_{sorted} .

$$Z_{sorted} = \text{sort}(|Z[k]|) \quad (12)$$

We then find the index C such that:

$$\sum_{k=0}^C |Z_{sorted}[k]|^2 \geq 0.9 * E(z) > \sum_{k=0}^{C-1} |Z_{sorted}[k]|^2 \quad (13)$$

We apply perturbations to the frequency components in the sorted sequence ($Z_{sorted}[k]$) which lie above the index C . Let Z_1, Z_2 and Z_3 be three consecutive frequency components in $Z_{sorted}[k]$ which lie above the index C , while Z_{1p}, Z_{2p} and Z_{3p} be the corresponding perturbed components. Then, the perturbed frequency domain components are obtained as:

$$Z_{1p} = Z_1 + \frac{Z_2}{4}, Z_{2p} = Z_2 + \frac{Z_2}{2}, Z_{3p} = Z_3 + \frac{Z_2}{4} \quad (14)$$

We apply such perturbation on upto 75% of the components which lie above the index C . All components below index C remain same. The perturbed time domain timeseries is obtained by taking inverse DFT of the perturbed frequency domain components. Thus, we only redistribute energy in the frequency components forming less than 10% of the energy of the signal. Figure 3(e),(f) shows example timeseries data augmented using this technique. We term such augmentation method as **Spec-Den**.

III. EXPERIMENTS AND RESULTS

A. Network Architecture

1) *ODENet*: The timeseries dynamics of the input are modelled using differential equations. A neural network is trained to compute the gradients of input timeseries at any given timestep (eq. 3). The neural network $f(\cdot)$ consists of two

hidden layers of 25 neurons. Each hidden layer is followed by batch normalization and ELU activation. The value of the input timeseries at a given time ($z(t_i)$) and the timestep (t_i) form the input to the neural network. The network predicts the gradient of input time series at timestep (t_i).

2) *Classification Network*: The neural network for time-series classification follows the ResNet architecture as defined in [34], which is depicted in fig 4. The input to this network is a time series of length T . The output of the network is a probability distribution over the K possible classes in the dataset. The network consists of 9 convolutional layers, grouped into three residual blocks. Each layer is followed by a Rectified Linear (ReLU) activation and batch normalization. This is followed by a global average pooling layer and a softmax classification layer. We retain the same network architecture and train it with the dataset augmented with our proposed methods. This allows us to compare the effectiveness of our techniques with the baseline in [34].

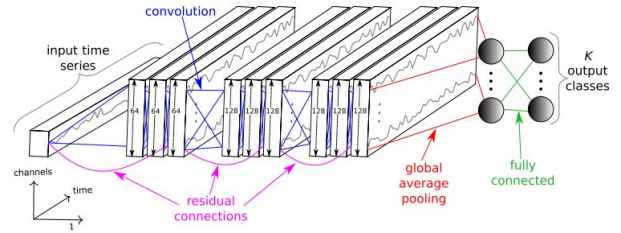


Fig. 4. Reference ResNet architecture of the neural network used for all experiments in this paper (from [34])

B. Datasets

We use the UCR timeseries dataset ([32]) for all the experiments described in this paper. Table I summarizes the datasets used, along with the baseline true classification accuracy of the ResNet classifier, baseline adversarial accuracy measured on the test data perturbed using FGSM ([17]) and BIM ([18]) attack on the ResNet classifier (as reported in [34]).

Using the techniques described in section II, we augment the datasets listed in table I. We use the augmented dataset to train the baseline classification network defined in section III-A2. The timeseries signals of all the datasets used in our experiments displayed an amplitude range close to 6 units. The value of β (equation 5) was chosen to be 0.33 and ϵ (equation 7) was a random number in range $[0.0, 0.33]$ for all the experiments described in this paper. Note that, We choose the value of β and ϵ in such a way, so as to limit the perturbation range to be almost $\frac{1}{10}$ th of the amplitude range of the original time series.

C. Additional Technique

In addition to the techniques already described for adversarial sample generation, we experimented with following regularization technique to improve the classification accuracy. We describe this below:

TABLE I
COMPARISON OF TRUE, ADVERSARIAL CLASSIFICATION ACCURACY ON NETWORKS WHERE TRAINING DATA IS AUGMENTED USING TECHNIQUES DESCRIBED IN SECTION II AND BASELINE TRUE, ADVERSARIAL ACCURACY

Dataset	Baseline			Input Gradient Method			Output Gradient Method			Spectral Density		
	No Augment			In-Clamp/In-sign-Grad			Out-Sign-Grad			Spec-Den		
	True	FGSM	BIM	True	FGSM	BIM	True	FGSM	BIM	True	FGSM	BIM
CricketX	79.00	35.4	20.8	79.49	47.95	38.72	79.49	72.31	72.05	79.23	46.67	37.44
CricketZ	81.5	27.7	16.2	82.31	43.33	34.87	81.79	42.82	37.18	81.03	40.26	33.08
50Words	73.2	17.1	8.8	74.73	38.46	29.23	76.04	67.25	52.97	74.07	25.5	12.09
UWaveGestureLibrary_X	78.0	32.1	11.1	78.28	46.76	20.91	79.23	71.19	55.05	77.11	40.90	19.88
InsectWingbeatSound	50.6	17.7	15.7	50.81	25.81	21.72	53.48	39.14	38.13	51.16	25.20	20.05
Adiac	83.1	3.1	1.5	83.38	6.39	3.84	83.63	45.27	25.83	83.38	5.63	4.09
TwoLeadECG	100	5.3	0.4	100	14.05	7.64	100	61.98	31.69	100	14.72	7.20
UWaveGestureLibrary_Y	66.7	27.7	14.9	67.03	36.24	17.64	67.06	57.96	47.35	66.14	35.87	18.01

TABLE II
ADVERSARIAL ACCURACY ON NETWORKS TRAINED USING ADDITIONAL FEAT-SIM REGULARIZER.

Dataset	(A): Input Gradient Method + Feat-Sim						(B): Out-Sign-Grad+ Feat-Sim		
	In-Clamp-Grad+ FS			In-Sign-Grad+ FS			Out-Sign-Grad+ FS		
	True	FGSM	BIM	True	FGSM	BIM	True	FGSM	BIM
CricketX	79.23	47.18	37.44	79.23	51.28	42.05	79.49	75.90	76.41
CricketZ	81.54	40.00	31.03	82.31	45.90	36.15	81.73	66.41	66.15
50Words	74.07	40.66	30.33	73.20	34.73	19.34	76.04	69.23	55.82
UWaveGestureLibraryX	78.42	51.90	28.25	78.03	47.74	24.06	78.53	72.47	62.62
InsectWingBeatSound	49.85	26.41	22.42	49.29	28.23	23.38	51.87	43.74	42.17
Adiac	79.28	7.16	5.09	82.61	6.65	4.86	83.38	46.29	28.64
TwoLeadECG	100	38.81	29.06	100	29.24	21.60	100	73.40	30.38
UWaveGestureLibraryY	66.86	36.71	20.69	66.95	41.54	25.43	67.07	59.35	58.15

1) *Feature Similarity*: Distance metric for large margin nearest neighbour classification was introduced in [46] and later used for face recognition and clustering in [48]. Let z_c be a time series sample of a specific class c . Let z_p be the augmented sample obtained by perturbing z_c using techniques described in section II. Let $z_{c'}^i$ be a randomly chosen time series sample of any other class $c' \neq c$. The augmented sample obtained by perturbing $z_{c'}^i$ is given by z_p^i . $g(z) \in \mathbb{R}$ represents the d-dimensional embedding of the timeseries z obtained by the neural network. Then, the loss function that is being minimized has an additional term (along with cross entropy) given by

$$L_{FS} = \|g(z_c) - g(z_p)\|_2^2 + \alpha - \|g(z_c) - g(z_p^i)\|_2^2 \quad (15)$$

where α is a margin that is enforced between positive and negative pairs. We also ensured that the embedding to lives on the d-dimensional hypersphere, i.e. $\|g(z)\|^2 = 1$. We term this regularized training strategy as **Feat-Sim(FS)**.

D. Results

In this section, we describe and summarize the classification accuracy observed on the reference ResNet based classification network (section III-A2), where the training dataset is augmented using techniques summarized in section II.

Comparison with baseline: Table I reports the comparison between baseline accuracy and the best accuracy achieved using our following proposed augmentation methods: Output Gradient based Augmentation, Input Gradient based augmentation and spectral Density based augmentation. We observe that even the Input Gradient based augmentation methods

(equations 4, 7) significantly improve the adversarial classification accuracies (FGSM and BIM). For every dataset, we highlight the Input Gradient Based technique that provided highest adversarial classification accuracy with **blue** and the Output Gradient based technique that provided highest adversarial classification accuracy with **green**. In both cases, the true accuracy was either same or better than the Baseline true classification accuracy of reference network. Reference network trained with augmented data obtained using Output Gradient based methods performed much better than the ones trained with data obtained using Input Gradient based method. This made intuitive sense because output gradient based method use the gradient of output of the network w.r.t input variable while input gradient based techniques use no knowledge of the network.

Comparison between different Input and Output Gradient based Augmentation methods: Table III in Appendix, we provide detailed comparison between different Input and Output Gradient based Augmentation methods. We observe that different input gradient based techniques were found to provide best adversarial classification accuracies for different datasets. This too made intuitive sense because these techniques are based on the gradients of the input time series itself, and hence it was unlikely that any one technique would give best results for all datasets. We also observe that true and adversarial accuracies achieved using Out-Sign-Grad method is much higher compared to Rand-Grad method. This result justify the significance of $sign \frac{dg(z)}{dz}$ term in equation 9.

Adversarial accuracy improvement with additional tech-

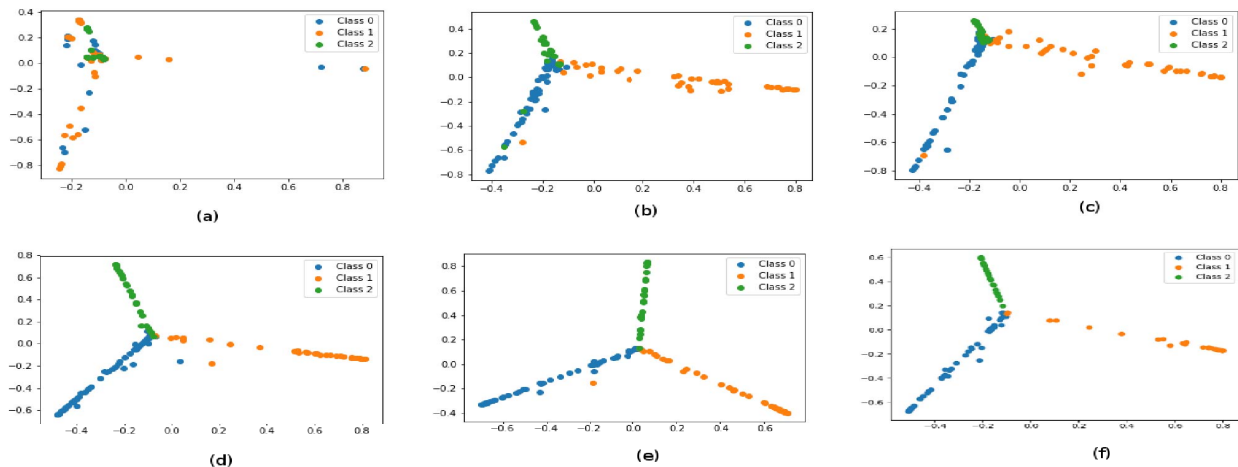


Fig. 5. PCA plots of latent representation of randomly selected samples from the 50Words dataset for network trained with various methods. (a) reference network, (b) In-Clamp-Grad, (c) In-Sign-Grad, (d) Spec-Den (e) Out-Sign-Grad, (f) Out-Sign-Grad + Feat-Sim

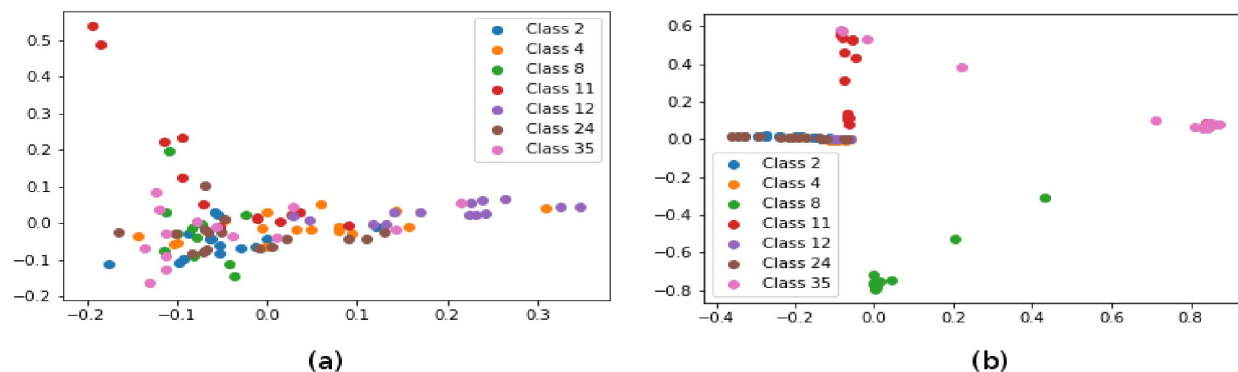


Fig. 6. PCA plots of latent representation of randomly selected samples from the Adiac dataset for (a) reference network and (b) Out-Sign-Grad + Feat-Sim

nique: Table II: We also observed that there was significant improvement in the adversarial classification accuracy when we employed additional Feat-Sim regularizer (described in section III-C) while training the reference classification network. These are summarized in table II. It can be seen that a network trained with the combination of Out-Sign-Grad(Augmentation) and Feat-Sim(Regularizer) (refer table II) provides the best adversarial classification accuracies along with significant improvements in true accuracies for all cases.

Validation In this section, we describe an independent technique that we used to validate the augmentation methods described in this paper.

a) **Principal Component Analysis:** In order to visualize the effect of augmentation methods along with additional regularizer technique described in section III-C on the latent representation of timeseries samples, we obtained PCA plots of the latent representation of randomly selected samples from the 50Words dataset. This is shown in figure 5. Here,

we can observe that the clustering of samples are better in (b) and (c) (networks trained with augmented samples obtained using input gradient based methods) compared to (a) (reference network). However, (e) and (f) (networks trained with augmented samples obtained using output gradient based methods) progressively improve the clustering, especially for the samples that lie at the intersection of clusters formed by different classes. This is inline with the results observed in table I and II.

IV. DISCUSSION

In this paper, we have introduced input gradient and output gradient based augmentation methods. We showed that networks trained using these augmented samples were more robust against standard adversarial attacks without compromising the true accuracy. We have also introduced spectral density based augmentation technique which also improved the robustness of the reference classification network. We have

also demonstrated that networks trained with the augmented data are able to cluster the samples in latent space much better than the reference network. However, in (fig 6) we have also observed that the samples belonging to certain classes are still inseparable in latent space. We would like to investigate this further and build more robust time series classifiers in our future work. We hope our findings should encourage the researchers to construct more efficient adversarial attack strategy for time series classifiers and should also open the door for new research in the area of adversarial robustness of time series classifier.

REFERENCES

- [1] Redmon, J., & Farhadi, A. (2016). YOLO9000: Better, Faster, Stronger. Retrieved from <http://arxiv.org/abs/1612.08242>
- [2] Szegedy, C., Ioffe, S., Vanhoucke, V., & Alemi, A. (2016). Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. Retrieved from <http://arxiv.org/abs/1602.07261>
- [3] He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask R-CNN. Retrieved from <http://arxiv.org/abs/1703.06870>
- [4] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. Retrieved from <http://arxiv.org/abs/1905.11946>
- [5] Zhao, Z.-Q., Zheng, P., Xu, S.-T., & Wu, X. (2019). Object Detection With Deep Learning: A Review. *IEEE Transactions on Neural Networks and Learning Systems*, 1–21. <https://doi.org/10.1109/TNNLS.2018.2876865>
- [6] Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to Sequence Learning with Neural Networks. Retrieved from <http://arxiv.org/abs/1409.3215>
- [7] Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. Retrieved from <http://arxiv.org/abs/1301.3781>
- [8] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Retrieved from <http://arxiv.org/abs/1810.04805>
- [9] Otter, D. W., Medina, J. R., & Kalita, J. K. (2018). A Survey of the Usages of Deep Learning in Natural Language Processing. Retrieved from <http://arxiv.org/abs/1807.10854>
- [10] Wang, Z., Yan, W., & Oates, T. (2017). Time series classification from scratch with deep neural networks: A strong baseline. *Proceedings of the International Joint Conference on Neural Networks, 2017–May*, 1578–1585. <https://doi.org/10.1109/IJCNN.2017.7966039>
- [11] Karim, F., Majumdar, S., Darabi, H., & Harford, S. (2019). Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116, 237–245. <https://doi.org/10.1016/j.neunet.2019.04.014>
- [12] Karim, F., Majumdar, S., & Darabi, H. (2019). Insights into LSTM Fully Convolutional Networks for Time Series Classification, 1–7. Retrieved from <http://arxiv.org/abs/1902.10756>
- [13] Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2019). Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery*, 33(4), 917–963. <https://doi.org/10.1007/s10618-019-00619-1>
- [14] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks, 1–10. Retrieved from <http://arxiv.org/abs/1312.6199>
- [15] Carlini, N., & Wagner, D. (2017). Towards Evaluating the Robustness of Neural Networks. *Proceedings - IEEE Symposium on Security and Privacy*, 39–57. <https://doi.org/10.1109/SP.2017.49>
- [16] Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. (2016). The limitations of deep learning in adversarial settings. *Proceedings - 2016 IEEE European Symposium on Security and Privacy, EURO S and P 2016*, 372–387. <https://doi.org/10.1109/EuroSP.2016.36>
- [17] Goodfellow, I. J., Shlens, J., & Szegedy, C. (2014). Explaining and Harnessing Adversarial Examples, 1–11. Retrieved from <http://arxiv.org/abs/1412.6572>
- [18] Kurakin, A., Goodfellow, I., & Bengio, S. (2016). Adversarial examples in the physical world, (c), 1–14. Retrieved from <http://arxiv.org/abs/1607.02533>
- [19] Azulay, A. & Weiss, Y. Why do deep convolutional networks generalize so poorly to small image transformations? *CoRR*, abs/1805.12177, 2018. URL <http://arxiv.org/abs/1805.12177>
- [20] Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2017). Towards Deep Learning Models Resistant to Adversarial Attacks, 1–27. Retrieved from <http://arxiv.org/abs/1706.06083>
- [21] Raghunathan, A., Steinhardt, J., & Liang, P. (2018). Certified Defenses against Adversarial Examples, 1–15. Retrieved from <http://arxiv.org/abs/1801.09344>
- [22] Wong, E., & Kolter, J. Z. (2017). Provable defenses against adversarial examples via the convex outer adversarial polytope. Retrieved from <http://arxiv.org/abs/1711.00851>
- [23] Papernot, N., McDaniel, P., Wu, X., Jha, S., & Swami, A. (2016). Distillation as a Defense to Adversarial Perturbations Against Deep Neural Networks. *Proceedings - 2016 IEEE Symposium on Security and Privacy, SP 2016*, 582–597. <https://doi.org/10.1109/SP.2016.41>
- [24] Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., & McDaniel, P. (2017). Ensemble Adversarial Training: Attacks and Defenses, 1–20. Retrieved from <http://arxiv.org/abs/1705.07204>
- [25] Lu, J., Issaranoon, T., & Forsyth, D. (2017). SafetyNet: Detecting and Rejecting Adversarial Examples Robustly. *Proceedings of the IEEE International Conference on Computer Vision, 2017–October*, 446–454. <https://doi.org/10.1109/ICCV.2017.56>
- [26] Yevgeniy Vorobeychik, Murat Kantarcioglu (2018). *Adversarial Machine Learning*. In *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers. isbn:9781681733968
- [27] S. M. Abdelfattah, G. M. Abdelrahman, & M. Wang, “Augmenting the size of EEG datasets using generative adversarial networks,” in *International Joint Conference on Neural Networks*, 2018, pp. 1–6.
- [28] Z. Zheng, Y. Yang, X. Niu, H. Dai, and Y. Zhou, “Wide and deep convolutional neural networks for electricity-theft detection to secure smart grids,” *IEEE Transactions on Industrial Informatics*, vol. 14, no. 4, pp. 1606–1615, 2018.
- [29] R. Briandet, E. K. Kemsley, & R. H. Wilson, “Discrimination of Arabica and Robusta in instant coffee by Fourier transform infrared spectroscopy and chemometrics,” *Journal of agricultural and food chemistry*, vol. 44, no. 1, pp. 170–174, 1996
- [30] A. Nawrocka and J. Lamorska, “Determination of food quality by using spectroscopic methods,” in *Advances in agrophysical research*, 2013.
- [31] S. Tobiyama, Y. Yamaguchi, H. Shimada, T. Ikuse, & T. Yagi, “Malware detection with deep neural network using process behavior,” in *IEEE Annual Computer Software and Applications Conference*, 2016, pp. 577–582.
- [32] Y. Chen, E. Keogh, B. Hu, N. Begum, A. Bagnall, A. Mueen, & G. Batista, “The UCR time series classification archive,” 2015.
- [33] Bagnall, A., Dau, H. A., Lines, J., Flynn, M., Large, J., Bostrom, A., ... Keogh, E. (2018). The UEA multivariate time series classification archive, 2018, 1–36. Retrieved from <http://arxiv.org/abs/1811.00075>
- [34] Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2019). Adversarial Attacks on Deep Neural Networks for Time Series Classification. Retrieved from <http://arxiv.org/abs/1903.07054>
- [35] Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2018). Data augmentation using synthetic data for time series classification with deep residual networks. Retrieved from <http://arxiv.org/abs/1808.02455>
- [36] Forestier, G., Petitjean, F., Dau, H. A., Webb, G. I., & Keogh, E. (2017). Generating synthetic time series to augment sparse datasets. In V. Raghavan, S. Aluru, G. Karypis, L. Miele, & X. Wu (Eds.), *Proceedings: 17th IEEE International Conference on Data Mining (pp. 865-870)*. Piscataway NJ USA: IEEE, Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/ICDM.2017.106>
- [37] Ford, N., Gilmer, J., Carlini, N., & Cubuk, D. (2019). Adversarial Examples Are a Natural Consequence of Test Error in Noise. Retrieved from <http://arxiv.org/abs/1901.10513>
- [38] Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., & Madry, A. (2018). Robustness May Be at Odds with Accuracy, 1–24. Retrieved from <http://arxiv.org/abs/1805.12152>
- [39] Chen, R. T. Q., Rubanova, Y., Bettencourt, J., & Duvenaud, D. (2018). Neural Ordinary Differential Equations, (NeurIPS). Retrieved from <http://arxiv.org/abs/1806.07366>
- [40] Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L. & Lerer, A. Automatic differentiation in pytorch. 2017.

- [41] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems 25, 2012, pp. 1097–1105.
- [42] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in International Conference on Machine Learning, 2015, pp. 448–456
- [43] Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). CoRR, abs/1511.07289, 2016.
- [44] Buckman, J., Roy, A., Raffel, C., and Goodfellow, I. Thermometer encoding: One hot way to resist adversarial examples. 2018. URL <https://openreview.net/pdf?id=S18Su-CW>
- [45] Yang, S., Luo, P., Loy, C. C., Shum, K. W., & Tang, X. Deep representation learning with target coding. In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15, pp. 3848–3854. AAAI Press, 2015. ISBN 0-262-51129-0. URL <http://dl.acm.org/citation.cfm?id=2888116.2888250>.
- [46] K. Q. Weinberger, J. Blitzer, and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. In NIPS. MIT Press, 2006. 2, 3
- [47] Sarkar, A., Gupta, N. K., and Iyengar, R. (2019). Enforcing Linearity in DNN succours Robustness and Adversarial Image Generation. <http://arxiv.org/abs/1910.08108>
- [48] Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A Unified Embedding for Face Recognition and Clustering. <https://doi.org/10.1109/CVPR.2015.7298682>

V. APPENDIX

A. Detailed results of different augmentation methods

Table III depicts detailed results of different input and output gradient based augmentation techniques as described in section II-A and II-B.

TABLE III
ADVERSARIAL CLASSIFICATION ACCURACY ON NETWORK WHERE TRAINING DATA IS AUGMENTED USING TECHNIQUES DESCRIBED IN SECTION II-A,II-B

Dataset	Input Gradient Methods						Output Gradient Methods					
	In-Clamp-Grad			In-Sign-Grad			Out-Sign-Grad			Rand-Grad*		
	True	FGSM	BIM	True	FGSM	BIM	True	FGSM	BIM	True	FGSM	BIM
CricketX	79.74	44.36	31.28	79.49	47.95	38.72	79.49	72.31	72.05	78.21	41.54	31.03
CricketZ	82.31	40.00	30.00	82.31	43.33	34.87	81.79	42.82	37.18	76.92	34.87	27.95
50Words	74.73	38.46	29.23	74.29	33.19	16.70	76.04	67.25	52.97	74.29	33.63	26.59
UWaveGestureLibrary_X	78.03	46.18	21.11	78.28	46.76	20.91	79.23	71.19	55.05	76.33	56.53	34.09
InsectWingbeatSound	50.81	25.81	21.72	49.29	27.27	23.23	53.48	39.14	38.13	48.44	26.06	23.79
Adiac	83.38	6.39	3.84	82.10	6.14	4.86	83.63	45.27	25.83	81.07	17.14	11.76
TwoLeadECG	100	13.61	6.85	100	14.05	7.64	100	61.98	31.69	99.74	17.38	6.85
UWaveGestureLibrary_Y	67.0	35.32	18.31	67.03	36.24	17.64	67.06	57.96	47.35	65.01	35.21	17.61

* Not a whitebox technique, but used here for comparison

B. Bound on energy difference between original and augmented sample augmented using spectral density based augmentation method

if $Z_1, Z_2 \in R$ with $Z_1, Z_2 > 0$ and $|Z_1| > |Z_2|$, then

$$|Z_1 + \delta|^2 > |Z_2 + \delta|^2 \quad (16)$$

Thus, if Z_1, Z_2, Z_3 are the magnitudes of three consecutive frequency components from Z_{sorted} , then the perturbed components are given by equation 14. The energy difference in these components introduced due to this operation is given by:

$$E_\delta = (|Z_{1p}|^2 + |Z_{2p}|^2 + |Z_{3p}|^2) - (|Z_1|^2 + |Z_2|^2 + |Z_3|^2) \quad (17)$$

To maximize the difference between Z_1 and Z_{1p} , the component added to Z_1 (i.e., $Z_2/4$) should be maximum. But since $Z_1 \geq Z_2$, $Z_2 = Z_1$ provides the maximum difference between Z_1 and Z_{1p} . Along the same lines, we find that E_δ is maximized when $Z_1 = Z_2 = Z_3$.

$$\max(|Z_{1p}^2 - Z_1^2|) = (Z_1 + \frac{Z_1}{4})^2 - Z_1^2 \quad (18)$$

$$\max(|Z_{2p}^2 - Z_2^2|) = (Z_1 - \frac{Z_1}{2})^2 - Z_1^2 \quad (19)$$

$$\max(|Z_{3p}^2 - Z_3^2|) = (Z_1 + \frac{Z_1}{4})^2 - Z_1^2 \quad (20)$$

Or,

$$\max(E_\delta) = 3.375Z_1^2 - 3Z_1^2 \quad (21)$$

$$\frac{\max(E_\delta)}{E_{original}} = \frac{3.375Z_1^2 - 3Z_1^2}{3Z_1^2} = 0.125 \quad (22)$$

Where $E_{original}$ is the energy of the three signal components that are perturbed. Thus, the maximum difference in energy after the perturbation (equation 14) is 12.5% of the original energy present in the three components that are perturbed.

We have assumed (equation 13) that the perturbed components form 10% of the signal energy. Hence, $\max(E_\delta)$ forms 1.25% of the energy of entire signal.