Naval Information Warfare Center PACIFIC

TECHNICAL REPORT 3302
SEPTEMBER 2023

Toward Channel Estimation with Conditional Generative Adversarial Networks

Erich C. Walter
Tanya G. Cheung
NIWC Pacific

Naval Information Warfare Center (NIWC) Pacific
San Diego, CA 92152-5001

This page is intentionally blank.

TECHNICAL REPORT 3302
SEPTEMBER 2023

Toward Channel Estimation with Conditional Generative Adversarial Networks

Erich C. Walter
Tanya G. Cheung
NIWC Pacific

Administrative Notes:
This report was approved through the Release of Scientific and Technical Information (RSTI) process in November 2022 and formally published in the Defense Technical Information Center (DTIC) in September 2023.

NIWC Pacific
San Diego, CA 92152-5001

NIWC Pacific
San Diego, California 92152-5001

P.M. McKenna, CAPT, USN
Commanding Officer

M.J. McMillan
Executive Director

ADMINISTRATIVE INFORMATION

The work described in this report was performed by the Spectrum Exploitation Branch (Code 75671) of the Cyber S&T Department, Naval Information Warfare Center (NIWC) Pacific, San Diego, CA. The Office of Naval Research (ONR) In-house Laboratory Independent Research (ILIR) program provided funding for this Basic Applied Research project.

Released by
Alex Vassel, Division Head
Information Operations Division

Under authority of
Joseph Simonetti, Department Head
Cyber/S&T Department

ACKNOWLEDGMENTS

Editor: RJP/MM

EXECUTIVE SUMMARY

OBJECTIVE

The goal of this paper is to model the complex channel effects that can occur during radio frequency (RF) transmission using Generative Adversarial Networks (GANs). Two different GANs were trained on paired and unpaired RF datasets with varying modulation schemes and channel effects to investigate the different GAN's potential, and to learn the complex channel effects that can occur during RF transmission. An expert feature-based system, GNU Radio (physics-based modelling), was used to generate the synthetic RF transmit and receive dataset pairs. After training on the synthetic data, the conditional GANs produced output that qualitatively fit the training data. This is a first step toward training a GAN that can qualitatively and quantitatively reproduce the transformation between the transmitted RF data and the received RF signal. Ultimately this approach can be applied to a paired dataset recorded in real-world conditions.

METHODS

To train a conditional GAN, it is necessary to have a representative dataset that contains a transmit and receive pair.

- Synthetic datasets were created using GNU Radio, expanding on the published modulation dataset RadioML[1].
    - This modulation dataset has eight digital modulations set at various signal-to-noise ratio (SNR) levels.
    - The examples were created with and without a channel model to simulate this transmit and receiver pair that would be fed into GAN training.
- There were two main different conditional GAN architectures that were used in this study:
    - Pix2Pix, a style transfer GAN that requires paired data.
    - Wasserstein GAN, a stochastic generator that does not require paired data.
- GAN output was visualized with In-phase and Quadrature (IQ) plots and time series plots and then qualitatively compared to the training data.

CONCLUSIONS AND RECOMMENDATIONS

While further work is needed, GANs show promise for generating RF data as well as channel effect estimation. Both GAN architectures produced raw IQ data that was qualitatively consistent with the training data for the modulation and SNR datasets. Further optimization with the use of quantitative metrics will improve this performance. The datasets used in this study contained small snippets of RF signals. Increasing signal lengths generated by the GAN to seconds will require a new GAN architecture. Current work shows qualitative imitation of synthetic data, and therefore the GAN is not better than synthetic generation at this point. The goal moving forward is to expand this study into over-the-air data collections.

This page is intentionally blank.

ACRONYMS

cGAN	Conditional Generative Adversarial Networks
c-WGAN-GP	Wasserstein GAN with Gradient Penalty
GAN	Generative Adversarial Networks
GP	Gradient Penalty
ILIR	In-house Laboratory Independent Research
IQ	In-phase and Quadrature
LSTM	Long Short-Term Memory
ML	Machine Learning
NIWC Pacific	Naval Information Warfare Center Pacific
ONR	Office of Naval Research
PDF	Probability Distribution Function
POR	Program of Record
RFR	Refractivity from Radio
RF	Radio Frequency
SNR	Signal-to-Noise Ratio
WGAN	Wasserstein GAN

CONTENTS

FIGURES

1. INTRODUCTION

1.1 PURPOSE

The purpose of this report is to model the effects of atmospheric propagation for radio frequency (RF) transmissions given a set of transmission parameters using a conditional generative adversarial network (GAN). A data-driven, machine learning approach has the potential to outperform expert feature crafted models for channel effect forecasting [3, 4]. Channel characteristics are learned from In-phase and Quadrature (IQ) data. The model predicts how new signals will likely appear at the receiver given a set of transmission conditions.

GANs have been heavily researched in the image domain over the past decade and have shown successful results that generate realistic adversarial images [2, 5]. Little research has applied GANs to synthetic RF data [3], and fewer studies have utilized real RF data collections [4].

If successful, an RF GAN would have the potential to augment existing refractivity from radio (RFR) capabilities in force level ships, as well as provide a range of complete “RF propagation prediction from radio” prediction services for forces that lack access to existing Program of Record (POR) solutions for different computational and RF sensor capabilities. The flexible approach afforded by these algorithms will blend such estimates with spectrum management, interference estimation, and pattern of life analysis.

1.1.1 Objectives

The objective of this study is to adapt a conditional GAN architecture from the image domain to estimate the channel effects from synthetic IQ data generated with GNU Radio’s dynamic channel model.

1.1.1.1 Creating Synthetic Paired RF Data

A synthetic dataset is the best method to get a quick and clean paired dataset to facilitate research on the feasibility of GANs applied to RF.

For this experiment, the synthetic RF dataset has following requirements:

- Complex IQ data
- Transmit and receive pairs
- Conditional information to feed the GAN

1.1.1.2 Training GAN Architectures

GANs were chosen based upon their successes in other domains and their different specifications and capabilities:

- Stochastic vs. deterministic
- Paired vs. unpaired datasets
- Non-linear
- Resistance to typical GAN failure modes

1.2 BACKGROUND

Currently, many different complex physics-based algorithms can be used to predict the quality of channels in RF communications with varying levels of fidelity across many parameters. RF channels are influenced by complicated, sparsely measured processes, including propagation (due to weather, platform movement, etc.), interference (due to human activity, automated processes, etc.), and receiver/transmitter hardware. A conditional generative adversarial network (cGAN) can potentially learn channel characteristics directly from IQ data, and then use that model to predict how new signals would appear at the receiver. A heuristic Machine Learning (ML) approach may be well suited to sparse measurements and latency constraints for operational RF channel forecasting. Recent ML advances have shown a data-driven approach may have better fidelity to experimental results than physics-based models [9]. Additionally, they have outperformed expert feature-based systems in many domains [10] including RF [1].

GANs can produce novel examples from a learned dataset distribution through the use of two competing networks, a generator and a discriminator. GAN training iteratively learns the distribution of the training data. The generator network produces an example to pass to the discriminator which evaluates if this newly generated example is part of the training dataset. This effort is a basic step in the direction of mapping out the equivalent possibilities using this technique with RF data.

1.2.1 GANs

GANs were first defined by Ian Goodfellow in 2014. They consist of two networks: a generator that generates new data, and a discriminator that captures the training data distribution and then estimates the probability that new samples from the generator came from the training data distribution [5]. The competing architectures are trained based on a zero-sum game, where the discriminator gets better at distinguishing between real and fake samples, while the generator is updated based on how well the fake samples fool the discriminator. Since generative models are generally unsupervised, this method allows the generator to act as a supervised model using the discriminator. This iterative training between the generator and the discriminator allows the GAN to generate new plausible examples that fit within the distribution of the training set. Figure 1 shows a diagram of the flow graph of GANs.

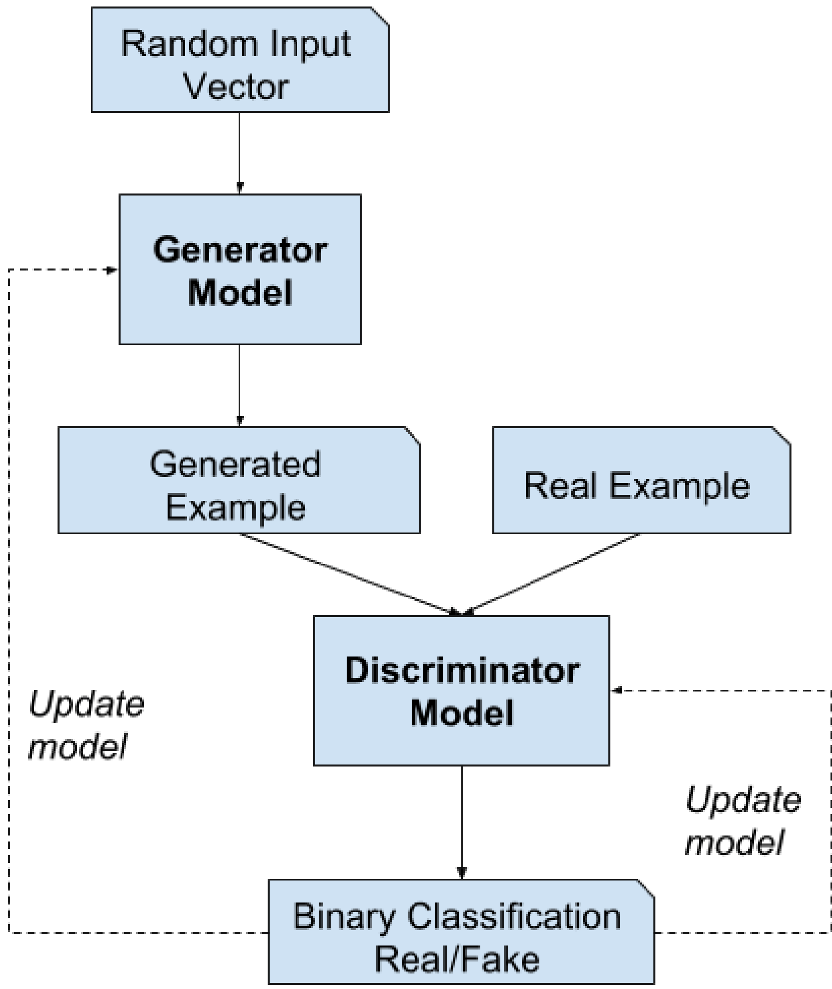

Figure 1. Flow Graph of GANs [5].

These competing architectures have been studied and expanded immensely within the image domain, and many variations of GANs were developed. One of which is the conditional GAN, where the GAN is provided with additional input to inform the output, which allowed for many new applications including but not limited to image-to-image translation and style transfer. This includes the two main architectures that were used in this study: conditional Wasserstein GANs, and Pix2Pix. Each will be explained in more detail in the Methods Section 2.

1.2.2 Literature Review

Very few papers have been published regarding GANs applied to RF datasets. One of the first papers in this area is "Approximating the Void: Learning Stochastic Channel Models from Observation with Variational Generative Adversarial Networks," which goes into detail on using GANs to jointly approximate both a channel response and a modulation scheme [3]. To more accurately estimate the probability distribution function (PDF) of the channel, they introduce a variational GAN that captures stochastic behavior. They have had some successes, but also have had some struggles to learn the stochastic processes of the channel responses accurately.

In 2020, Dorner published “Wasserstein GAN (WGAN)-based autoencoder training over-the-air,” which expands on the idea previously mentioned by using GANs to learn channel behavior but on real over-the-air measurements [4]. This uses a conditional Wasserstein GAN and embeds it into an autoencoder architecture. They explain how GANs were more successful in creating a more realistic PDF in comparison to other machine learning methods such as reinforcement learning, but GANs still have some limitations, specifically with non-stationary channels.

Both of these papers use GANs to estimate the PDF of a channel response. In the case of this study, GANs are fed raw IQ data to generate IQ that is representative of the signal after going through a channel.

2. METHODS AND MATERIALS

2.1 METHODS

Synthetic modulation data was generated using RadioML package for GNU Radio. This Python code allowed for various characteristics in the paired dataset that was generated, including number of samples, length of examples, signal-to-noise ratio (SNR), and channel effects. The codebase was used to create synthetically "transmitted" and "received" data by the inclusion of the channel models. The RadioML modulation recognition dataset source code [1] was modified to generate the synthetic datasets meeting the above requirements. Obtaining a paired dataset is necessary to establish ground truth that the algorithm can train on.

The simulated dataset is very similar to the RadioML dataset, with a 128-length complex IQ with eight digital modulations. To make it paired, the dataset was generated both with and without the simulated channel effects to mimic transmitted and received data with varying modulation and SNR.

Conditional GANs adapted for RF data:

- Wasserstein GAN with Gradient Penalty (c-WGAN-GP)
- Pix2Pix GAN

A Wasserstein GAN [6] with gradient penalty [7] was chosen as a starting point for generating RF data due to the theoretical improvements over previous GANs. The main changes in the Wasserstein GAN (WGAN) are the switch to using the Earth-Mover distance (Wasserstein) and the switch from a discriminator to a critic model. This architecture offers improved training stability and performance while being more resilient towards mode collapse, a typical GAN failure mode. This network was used to generate new RF examples from a stochastic latent space. This network can train on unpaired data.

The c-WGAN-GP was designed with the generator consisting of dense layers and a critic based upon our earlier work with RF modulation recognition classifiers (2D convolutional long short-term memory (LSTM) layers). The model consisted of seven dense layers with hyperbolic tangent activation with batch normalization followed by a reshape layer that returns the appropriate dimensions of the RF signal in the training dataset. Figure 2 shows the generator network that consisted of dense layers followed by batch normalization and tanh activation. Figure 3 shows a six-layer diagram of the discriminator architecture.

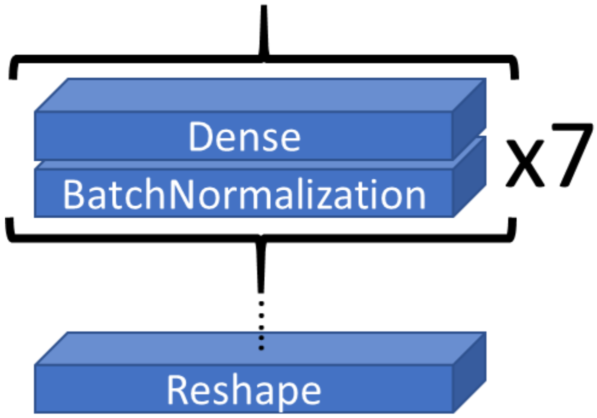

Figure 2. The generator network consisted of dense layers followed by batch normalization and tanh activation.

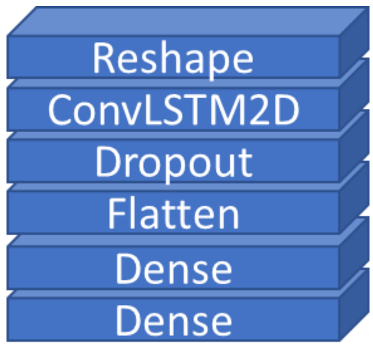

Figure 3. The discriminator architecture consisted of a dense layer followed by 2D convolutional LSTM layer, another two dense layers, and a dense layer with the number of features as the output.

The Pix2Pix network [8] is an image-to-image transfer GAN that learns the mapping from an input to an output image and the loss function needed to train the input to output mapping. This network was used to map channel effects to a high SNR RF signal. This network trains on paired data. The 2 x 128 RF signals were padded with zeros to create 128 x 128 arrays, which were then passed into the Pix2Pix input. Figure 4 shows two views of the Pix2Pix Architecture.

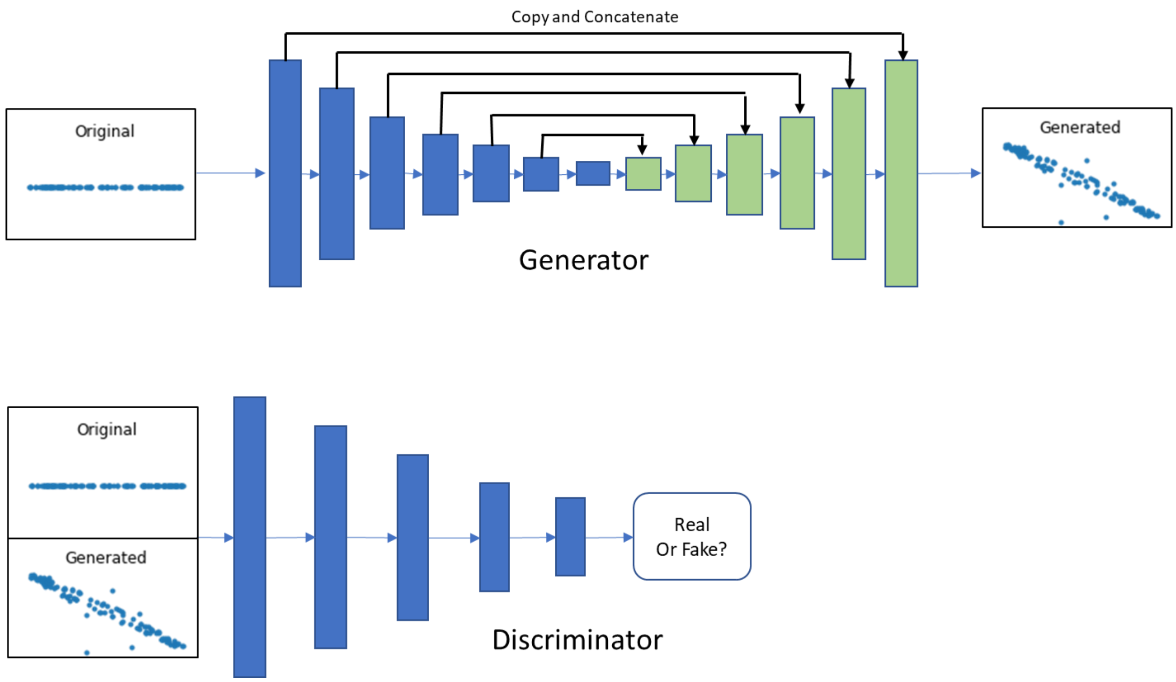

Figure 4. Pix2Pix Architecture.

The networks used were based on the Pix2Pix architecture: a U-Net Generator with 2D 2 x 2 convolutions and a stride of two. The discriminator used down sampling 2D 2 x 2 convolutions and a stride of two.

2.2 MATERIALS

All the development was done in Python on an Nvidia DGX-station equipped with four V100 GPUs. Although optimization and processing speed is not the focus of this study, the architecture development time benefits from the GPUs and are utilized to accelerate training.

This page is intentionally blank.

3. RESULTS

3.1 RESULTS

GANS were programmed to save images periodically during training. Visual analyses of these images after training were used to determine the quality of the GAN output. Both time series data and IQ plots were used to judge how similar the data was to the training data.

3.1.1 Synthetic Data

3.1.1.1 Modulation Dataset

The eight digital modulations that are used in the generated dataset are: BPSK, QPSK, 8PSK, CPFSK, GFSK, PAM4, QAM16, and QAM64. Examples of these modulations are shown in Figure 5 without any channel effects. These signals are used as our "transmit" signal.

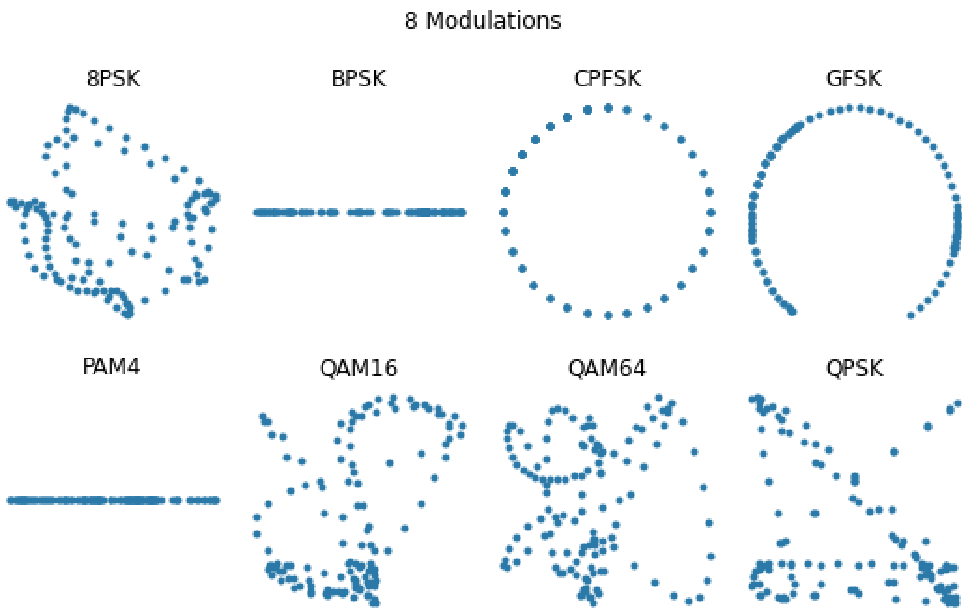

Figure 5. Examples of our eight different modulation classes used as transmit signals.

3.1.1.2 SNR Dataset

Examples of the synthetically-generated dataset are shown on the next page. With the same dynamic channel model as used in RadioML, channel effects include varying SNR with white Gaussian noise, fading channel, frequency, and sample rate offset. This is our "received" signal. Figure 6 has examples of receive signals at different SNRs.

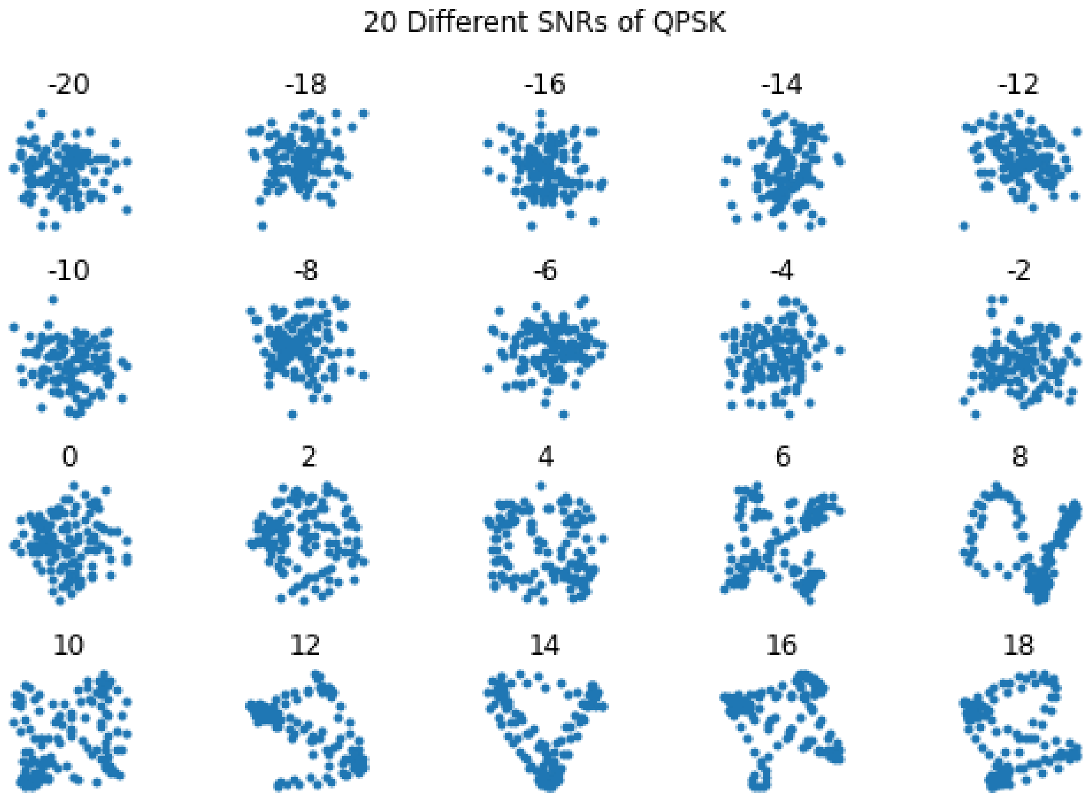

Figure 6. Examples of receive signals at different SNRs.

3.1.2 Wasserstein GAN Output

WGANs were typically trained for 60,000 epochs and up to 200,000 epochs. Training used eight classes, a latent space of 100, and a batch size of 32. Typical WGAN hyperparameters used were a learning rate of 0.0002, a beta_1 of 0, a beta_2of 0.9, and a epsilon of 1e-07. The critic was updated at a 5:1 ratio relative to the generator.

3.1.2.1 GAN-generated Modulation Data

The c-WGAN was trained on a dataset consisting of high SNR modulation data separated into eight classes based upon modulation scheme. Output from the GAN is in Figure 7. Figure 8 shows four comparative time domain plots generated for the same input.

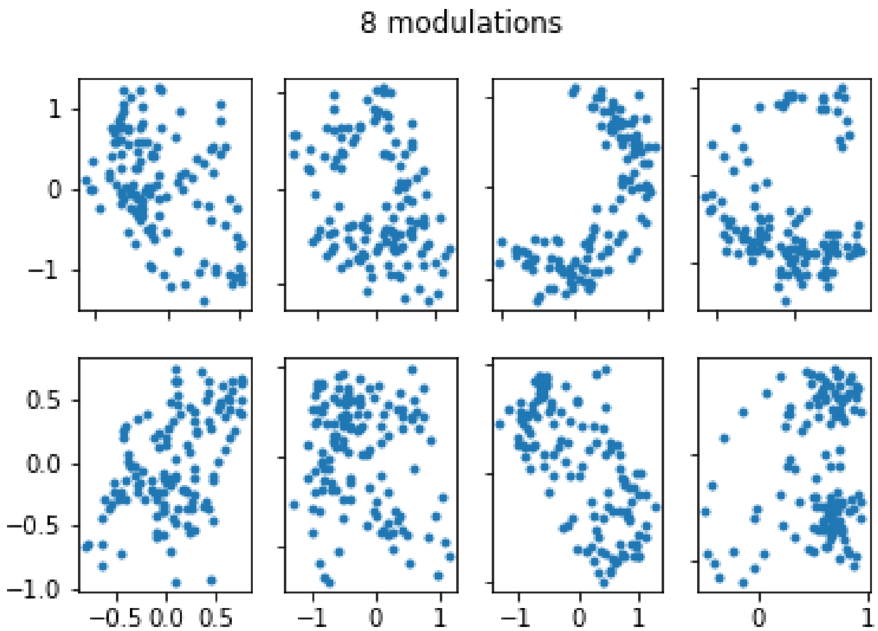

Figure 7. IQ Plots of eight different generated modulations.

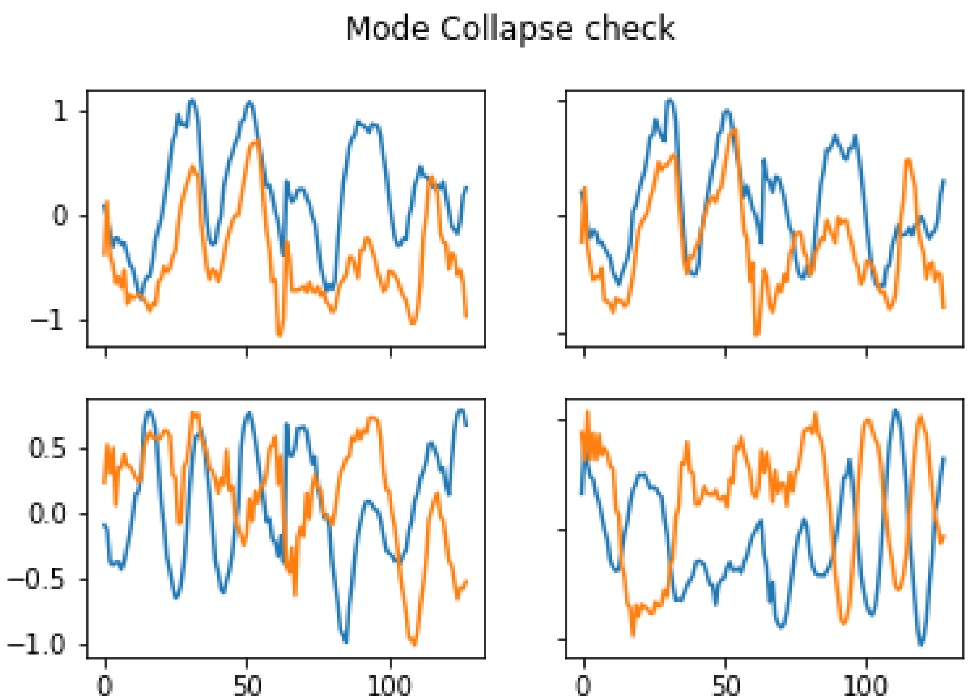

Figure 8. Time domain plots generated for the same input (one modulation class) demonstrating mode collapse resistance where blue is in-phase and orange is quadrature.

3.1.2.2 GAN-generated SNR Data

C-WGAN-GP output resulted in IQ data that qualitatively represented the training data. Figure 9 shows the GAN output for eight different SNR values for the modulation scheme QPSK.

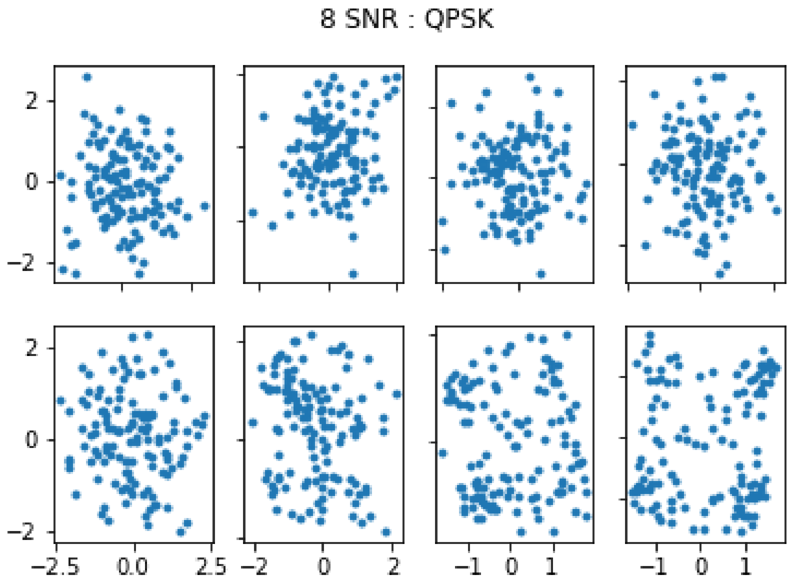

- Top row: -20, -16, -12, -8.
- Bottom row: -4, 0, 4, 8

Figure 9. GAN output for eight SNR values.

3.1.3 Pix2Pix GAN Output

The Pix2Pix GAN was trained on the paired SNR dataset with all eight modulations for 3,000 epochs with a sample interval of 200, and a batch size of 32. Figure 10 shows the outputs of the Pix2Pix GAN and shows visually how the IQ plots compare to the synthetically generated receive signals.

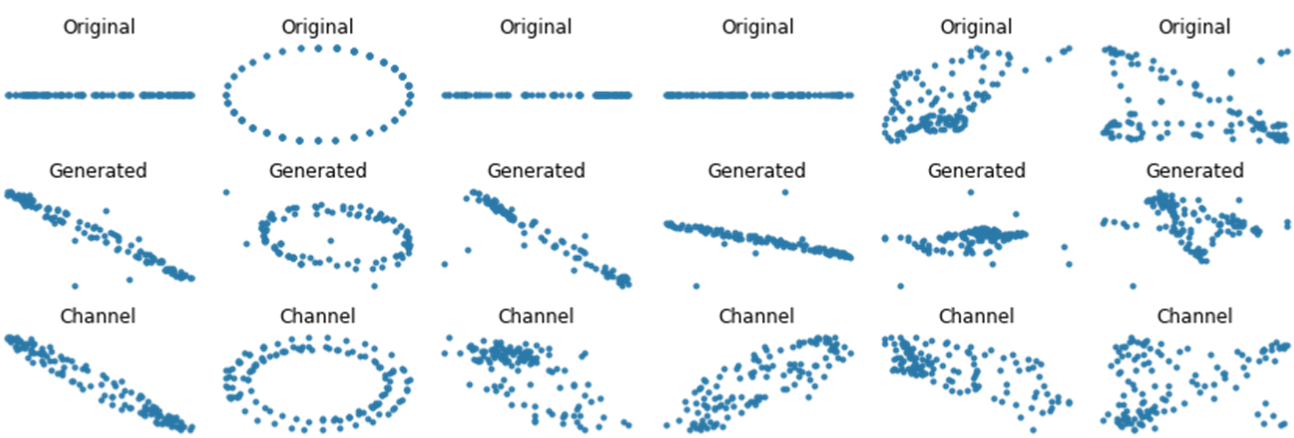

- Top row: transmit signal
- Middle row: output generated from Pix2Pix GAN for an SNR of 10
- Bottom row: 'Actual' receive signal.

Figure 10. Pix2Pix GAN IQ plots.

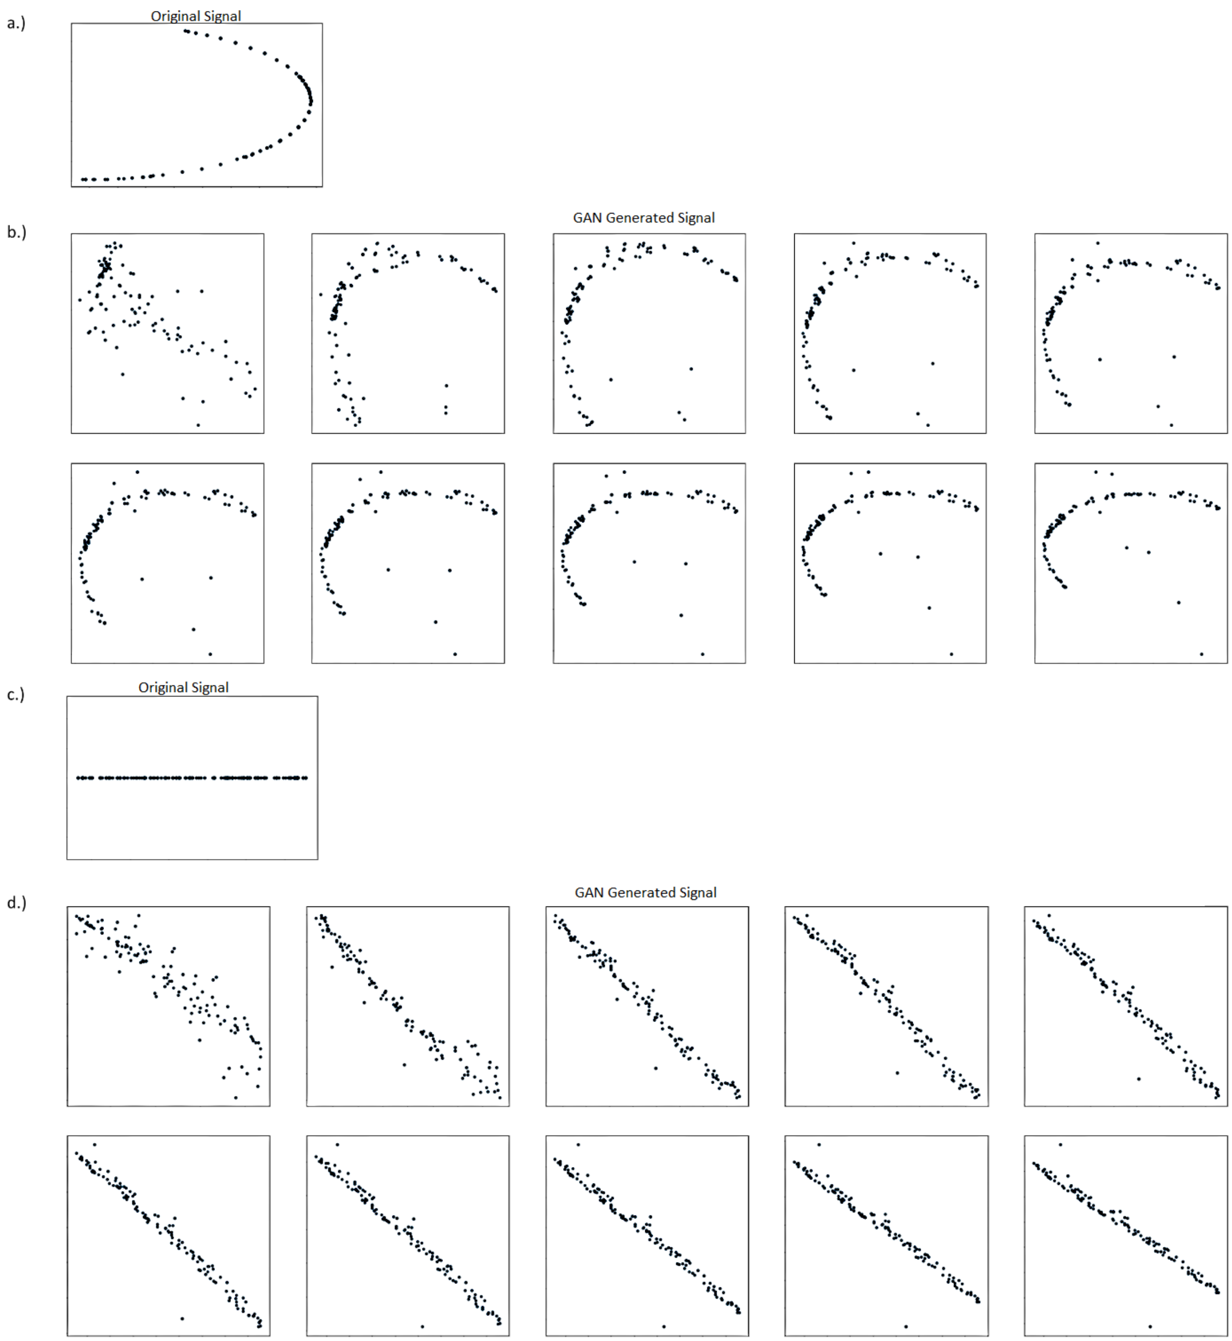

- The original signal input is shown in a) and c).
- The output generated by a trained Pix2Pix GAN is shown in b and d.

Figure 11. Two original signals and the corresponding 10 different SNR classes (0, 2, 4, 6, 8, 10, 12, 14, 16, 18) generated by Pix2Pix.

This page is intentionally blank.

4. DISCUSSION

4.1 MAJOR FINDINGS

GANs have been adapted from the image domain to the IQ domain. With this initial approach, RF signals generated from GANs seem to have qualitatively fit in the training dataset distribution of received signals, shown in both the IQ and time domain plots in the Results Section 3. One of the metrics used in this study is SNR. Because of the short signal length, it was hard to measure if the SNR was correct. Further statistical analysis also needs to be adapted to quantitatively measure GANs training progress as well as to be able to compare the output of different GANs.

4.1.1 Implications

While GANs can generate output that looks like channel effects for RF transmissions, it remains to be seen if it can stand up to rigorous statistical analysis. The relatively short sections of RF signals used also do not present a readily available way to generate a longer signal.

4.1.2 Future Research

One of the next major steps includes integrating more validation techniques that can quantitatively measure a GAN's success in RF. This includes validating if the GAN generated output signal contains the various channel effects that are generated in the synthetic IQ signals. Research on how to validate this performance, such as bit error rate, still needs to be done, but with this, it has been difficult to evaluate some of our basic metrics (e.g., SNR) due to the short examples. This implies the need of longer signals to perform some of these metrics. A metric that has been developed in the image domain is known as the inception score. The inception score is used to evaluate the performance of a GANs relative to each other and is based on the inception-v3 network.

From this initial study, there are a few of issues that need to be addressed. First is the issue of longer length signals of IQ, and second is the more complex channel domain. The current dataset only models one specific channel with its specific channel effects, but there are many other effects that do come into play for other types of data. Because of this, there are also a wide varying range of temporal effects. This implies that the input data signal of the GANs needs to be longer to capture the longer time pattern characteristics. An over-the-air dataset collection has been started here at NIWC Pacific with the corresponding metadata and is expected to be used for training GANs in the following year. With this, there are implications that there may need to be significant changes to some of the model architectures that are currently being used.

This page is intentionally blank.

5. CONCLUSIONS

Both GAN architectures tested produced RF signals qualitatively similar to the training data. The Pix2pix network suffered from mode collapse and was deterministic, as expected. The WGAN architecture was not as limited by mode collapse and could produce different output signals from the same input signal. Since these architectures are adapted from the image domain, there is still much research to be done in this area to validate the generated signals. Additionally, the next steps would include more validation techniques, as well as modifications to the architecture to be able to allow for larger and varying time scales.

This page is intentionally blank.

REFERENCES

1. O’shea, T., West, N. and Davidson, Y., 2017, May. “RadioML Dataset: Open RadioML Synthetic Benchmark Dataset,” *Github*. https://github.com/radioML/dataset
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y., 2020. “Generative adversarial networks,” *Communications of the ACM*, 63(11), pp.139–144.
3. O’Shea, T.J., Roy, T. and West, N., 2019, February. “Approximating the void: Learning stochastic channel models from observation with variational generative adversarial networks,” In *2019 International Conference on Computing, Networking and Communications (ICNC)* (pp. 681–686). IEEE
4. Dörner, S., Henninger, M., Cammerer, S. and ten Brink, S., 2020, May. “WGAN-based autoencoder training over-the-air,” *2020 IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)* (pp. 1–5). IEEE.
5. Brownlee, Jason. “A Gentle Introduction to Generative Adversarial Networks (Gans).” Machine Learning Mastery, 19 July 2019, https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans/.
6. Arjovsky, M., Chintala, S. and Bottou, L., 2017, July. “Wasserstein generative adversarial networks.” In *International conference on machine learning* (pp. 214–223). PMLR.
7. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. and Courville, A.C., 2017. “Improved training of wasserstein gans.” Advances in neural information processing systems, 30.
8. Isola, P., Zhu, J.Y., Zhou, T. and Efros, A.A., 2017. “Image-to-image translation with conditional adversarial networks.” In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1125–1134).
9. Willard, J., Jia, X., Xu, S., Steinbach, M. and Kumar, V., 2021. “Integrating scientific knowledge with machine learning for engineering and environmental systems.” *ACM Computing Surveys (CSUR)*.
10. Cuomo, S., Di Cola, V.S., Giampaolo, F., Rozza, G., Raissi, M. and Piccialli, F., 2022. “Scientific Machine Learning through Physics-Informed Neural Networks: Where we are and What's next.” *arXiv preprint arXiv:2201.05624*.

This page is intentionally blank.

INITIAL DISTRIBUTION

84310	Technical Library/Archives	(1)
71740	E. Walter	(1)
75671	T. Cheung	(1)

Defense Technical Information Center Fort Belvoir, VA 22060–6218	(1)

This page is intentionally blank.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-01-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
September 2023	Final	

4. TITLE AND SUBTITLE

Toward Channel Estimation with
Conditional Generative Adversarial Networks

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

~~5d. PROJECT NUMBER~~

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS

Erich C. Walter
Tanya G. Cheung
NIWC Pacific

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

NIWC Pacific
53560 Hull Street
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER

TR-3302

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

The Office of Naval Research In-house Laboratory
8500 Bluffstone Cove, Ste A201
Austin, TX 78759

Independent Research Program
875 N Randolph Street
Arlington, VA 22217

10. SPONSOR/MONITOR'S ACRONYM(S)

ONR and ILIR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release. Distribution is unlimited.

13. SUPPLEMENTARY NOTES

This work may be copied and disseminated without restriction.

14. ABSTRACT

The focus of this paper is to train Generative Adversarial Networks (GANs) on radio frequency (RF) datasets to learn the complex channel effects that can occur during RF transmission. Two different GANs were trained on paired and unpaired RF datasets with varying modulation schemes and channel effects to investigate the different GAN's potential to learn the complex channel effects that can occur during RF transmission. An expert feature-based system, GNU Radio (physics-based modelling), was used to generate the synthetic RF transmit and receive dataset pairs. After training on the synthetic data, the conditional GANs produced output that qualitatively fit the training data. This is a first step toward training a GAN that can qualitatively and quantitatively reproduce the transformation between the transmitted RF data and the received RF signal. Ultimately this approach can be applied to a paired dataset recorded in real-world conditions.

15. SUBJECT TERMS

Generative Adversarial Networks; published modulation dataset; RadioML; cGAN; ILIR; machine learning; Signal-to-Noise Ratio; Wasserstein GAN

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Erich C. Walter
					19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	SAR	36	619-553-9771

This page is intentionally blank.

This page is intentionally blank.

Naval Information Warfare Center
Warfare Center
PACIFIC

Naval Information Warfare Center (NIWC) Pacific
San Diego, CA 92152-5001