DEEP LEARNING-BASED MODULATION CLASSIFICATION USING SYNTHETIC AND OVER-THE-AIR SDR SIGNALS

Hovannes Kulhandjian*, Elizabeth Batz*, Michel Kulhandjian†
*Department of Electrical and Computer, California State University, Fresno, Fresno, CA 93740
†Department of Electrical and Computer Engineering, Rice University, Houston, TX, 77005
emails: hkulhandjian@csufresno.edu, ebatz32@mail.fresnostate.edu, mkk6@buffalo.edu
Faculty Advisor: Hovannes Kulhandjian

ABSTRACT

This work presents a convolutional neural network (CNN) for automatic classification of digital modulation schemes using signals received via software-defined radios (SDRs). A synthetic dataset was created in MATLAB for BPSK, 8-PSK, 16-PSK, QAM, 16-QAM, and 64-QAM, with 1,500 messages per scheme at five SNR levels and added phase noise. A 12-layer CNN trained on this dataset achieved 97% accuracy in classifying modulation types. To evaluate real-world performance, over-the-air signals were captured and used for validation, yielding classification accuracies ranging from 72% to 91%. While performance on live signals showed variability, the results indicate strong potential for generalization with further refinement. Enhancing the synthetic dataset with additional channel impairments may improve model robustness and real-world applicability. This research demonstrates the viability of using deep learning for signal classification in intelligent communication systems.

Index Terms: Modulation classification, deep learning, convolutional neural networks, software-defined radio, constellation images, over-the-air evaluation, domain adaptation.

INTRODUCTION

Automatic modulation recognition (AMR) is a foundational capability for modern wireless communication systems with applications in spectrum monitoring, electronic warfare, cognitive radio, and resilient link adaptation. In contested or congested environments, reliably identifying the underlying modulation scheme enables adaptive reconfiguration, interference mitigation, and enhanced situational awareness. Conventional AMR approaches rely on expert-crafted features, such as higher-order cumulants, cyclostationary statistics, or likelihood-based tests—followed by classical classifiers including support vector machines (SVMs) [1]. While effective in controlled conditions, these methods exhibit sharp performance degradation under low signal-to-noise ratio (SNR), time-varying channels, or hardware nonidealities.

Recent advances in deep learning (DL) have reshaped many inference domains by enabling discriminative feature learning directly from data. Convolutional neural networks (CNNs), recurrent neural networks (RNNs), and deep belief networks (DBNs) have demonstrated state-of-the-art accuracy in areas such as vision, radar sensing, and sonar classification [2, 3, 4, 5]. Extending these architectures to communications inference problems, and AMR in particular, has shown significant promise. Deep learning–based AMR frameworks bypass hand-crafted feature engineering by operating directly on raw in-phase/quadrature (I/Q) samples or constellation representations, thereby capturing robust discriminative patterns under diverse channel conditions [6, 7, 8, 9].

In this work, we propose and evaluate a compact 12-layer CNN framework for adaptive detection and classification of common digital modulation schemes (BPSK, QPSK, 8-PSK, 16-PSK, 4-QAM, 16-QAM, and 64-QAM). We curate a synthetic dataset that incorporates realistic impairments such as phase noise, oscillator drift, and variable SNR levels to improve generalization, and we validate transferability through over-the-air (OTA) testing using software-defined radios (SDRs). Our contributions are threefold: (i) a CNN-based architecture tailored for constellation imagery, (ii) a dataset generation pipeline that introduces channel impairments for improved robustness, and (iii) an OTA validation study that quantifies the synthetic-to-real performance gap and explores mitigation strategies such as domain randomization and mixed-domain training.

RELATED WORK

A. *Classical AMR Approaches*

Early AMR research focused on signal statistics and handcrafted features. High-order cumulants were widely employed in combination with SVMs for robust classification under Gaussian noise [1]. Cyclostationary feature extraction and likelihood-based hypothesis testing were also explored for specific modulation formats. While these methods provided high accuracy under nominal conditions, they often failed in dynamic environments characterized by frequency offsets, multipath fading, or colored noise.

B. *Machine Learning–Based AMR*

The limitations of feature-engineered approaches motivated the use of machine learning algorithms that can learn decision boundaries from examples. Early works applied SVMs, decision trees, or k-nearest neighbors to handcrafted features [1]. More sophisticated architectures such as DBNs [3] and RNNs [7] were introduced to capture nonlinear and temporal dependencies in modulation signals. While effective in low-dimensional scenarios, these models remained limited in scalability and often required extensive feature preprocessing.

C. *Deep Learning for Modulation Classification*

Convolutional neural networks emerged as a compelling solution for AMR by capturing geometric and local structural patterns in I/Q signals or constellation maps [8]. O'Shea et al. [8] demonstrated the feasibility of CNN-based AMR with competitive performance across a wide range of SNRs. Subsequent work investigated CNN variants for radar waveform recognition [6], distortion correction [7], and signal demodulation [4]. Practical toolkits such as MathWorks' Deep Learning Toolbox [9] and hardware platforms like the ADALM-PLUTO SDR [10] have further accelerated development and reproducibility.

Despite these advances, a persistent challenge is the domain gap between synthetic training data and OTA deployment. Factors such as oscillator drift, carrier frequency offset (CFO), synchronization errors, and nonlinear front-ends erode model generalization when trained exclusively on clean datasets. Recent efforts to address this gap include domain randomization (adding impairments during training), transfer learning on small OTA datasets, and hybrid training strategies [7], [11], [12]. Xu and Darwazeh [12] demonstrated the utility of deep learning for OTA non-orthogonal signal classification, emphasizing the importance of incorporating real-world variability.

D. Broader Applications of DL in Wireless and Sensing

Deep learning–based classification has also been applied to adjacent domains. Kulhandjian et al. demonstrated its use in sonar-based underwater activity recognition [5] and sign language gesture recognition using Doppler radar [13], highlighting its adaptability to noisy sensing environments. Visualization tools such as t-SNE [14] have been applied for high-dimensional separability analysis, further assisting in the design of impairment-aware training datasets. These cross-domain successes reinforce the potential of CNN-based AMR systems to generalize under practical impairments.

E. Positioning of This Work

Positioned within this trajectory, our work builds on CNN-based AMR research by introducing a robust dataset generation pipeline with channel impairments and evaluating OTA transferability using SDRs. Unlike prior studies limited to simulations or narrow datasets, we integrate practical impairments, principal-component analysis (PCA) diagnostics, and live hardware validation. The goal is to provide a reproducible framework for bridging the synthetic-to-real domain gap and advancing deployable AMR systems.

PROJECT DESIGN OVERVIEW

F. Principal Component Analysis (PCA)

PCA is a widely used technique in multivariate data analysis that reduces the dimensionality of datasets while preserving the most significant sources of variation. In this work, PCA is applied to constellation diagrams generated from both synthetically created signals and OTA captures using SDRs.

PCA works by transforming the dataset into a new coordinate system where the axes, referred to as principal components, are ordered by the amount of variance they capture. The first principal component contains the greatest variance, followed by the second, and so forth. By retaining only the leading components, redundant and correlated features are removed while the most discriminative information is preserved.

The steps of PCA include: (i) normalizing the dataset, (ii) computing the covariance matrix, (iii) extracting eigenvalues and eigenvectors of the covariance matrix, (iv) selecting a reduced set of principal components, and (v) projecting the original dataset onto the new feature space. The eigenvalues quantify the proportion of variance captured by each component, with the cumulative set representing 100% of the variation.

In the context of modulation classification, PCA serves two purposes: visualization of feature separability in lower dimensions and preprocessing to improve machine learning model efficiency. By reducing dimensionality while retaining core features, PCA can assist in making modulation classification more robust against noise and impairments.

G. Machine Learning Configuration

For the modulation classification task, a CNN was designed and implemented. CNNs are a specialized class of deep learning models particularly effective in extracting spatial features from structured inputs, such as constellation maps or spectrograms. Their architecture mimics hierarchical feature extraction similar to the connectivity patterns observed in the human brain.

The CNN architecture consists of multiple types of layers, each performing a unique role:

- Convolutional layers apply filters to extract local features from the input constellation images.
- Pooling layers perform down-sampling to reduce spatial dimensions and parameters while preserving essential patterns.
- Fully connected layers aggregate high-level features to form class-specific representations.
- SoftMax output layer normalizes probabilities and performs final classification among the candidate modulation schemes.

Figure 1 illustrates the overall CNN architecture applied for feature extraction.

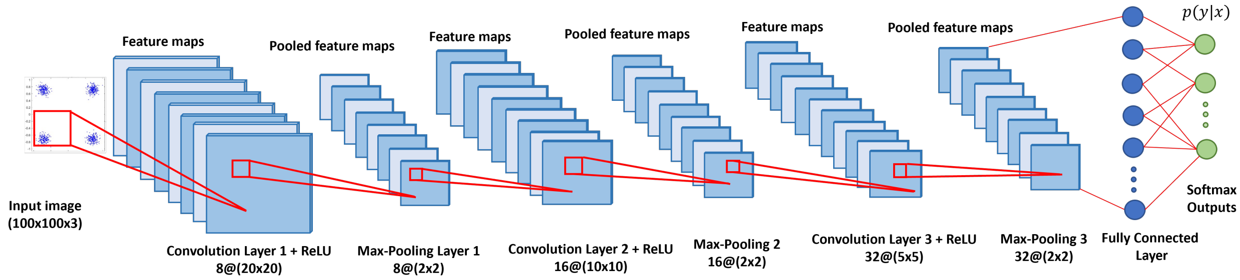

Figure 1: CNN Architecture.

Figure 2 shows the detailed 12-layer CNN configuration customized for modulation recognition.

PROJECT IMPLEMENTATION AND RESULTS

H. Dataset Generation

The training dataset for the modulation classification task was synthetically generated to ensure robustness and applicability to real-world scenarios. One limitation of purely OTA collection is the inherent variability of wireless channels, where noise, interference, and environmental dynamics

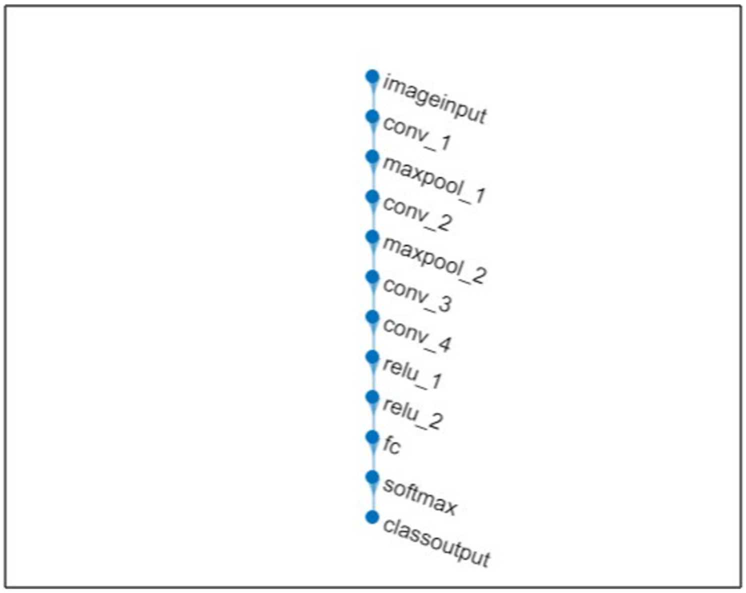

Figure 2: Custom 12 - Layer Network.

fluctuate on an hourly or daily basis. By generating the dataset synthetically, these conditions could be controlled, and realistic impairments such as variable SNR and phase noise were introduced deliberately. A total of 1,500 randomly generated messages of 1,000 bits each were encoded for every modulation scheme under consideration (BPSK, 8-PSK, 16-PSK, 4-QAM, 16-QAM, and 64-QAM). Five distinct SNR levels were incorporated to diversify training conditions, providing the network with exposure to noisy as well as clean signal environments. To ensure sufficient coverage, at least 1,000 labeled samples per class were maintained.
Figure 3 presents example constellation diagrams across different modulation formats.
Figure 4 demonstrates how constellation clarity degrades under lower SNR levels.

I. Software-Defined Radio Configuration

While synthetic data provides controlled conditions, validation on real-world OTA signals is necessary to assess transferability. For this purpose, the Analog Devices ADALM-PLUTO SDR was used to transmit and receive modulated signals at 915 MHz. The signals, once captured, were converted into constellation diagrams and compared against the synthetic dataset for validation.

OTA signals exhibited additional impairments not fully represented in the synthetic dataset, including synchronization drift, frequency offset, and phase distortion. These impairments highlighted the differences between simulated and real-world datasets.
Figure 5 shows the ADALM-PLUTO SDR used in OTA experiments.
Figure 6 depicts a received QPSK signal, highlighting synchronization improvements.

J. Principal Component Analysis Results

To further analyze the separability of modulation classes, PCA was applied to the dataset. PCA projected the high-dimensional constellation features into lower dimensions, revealing dis-

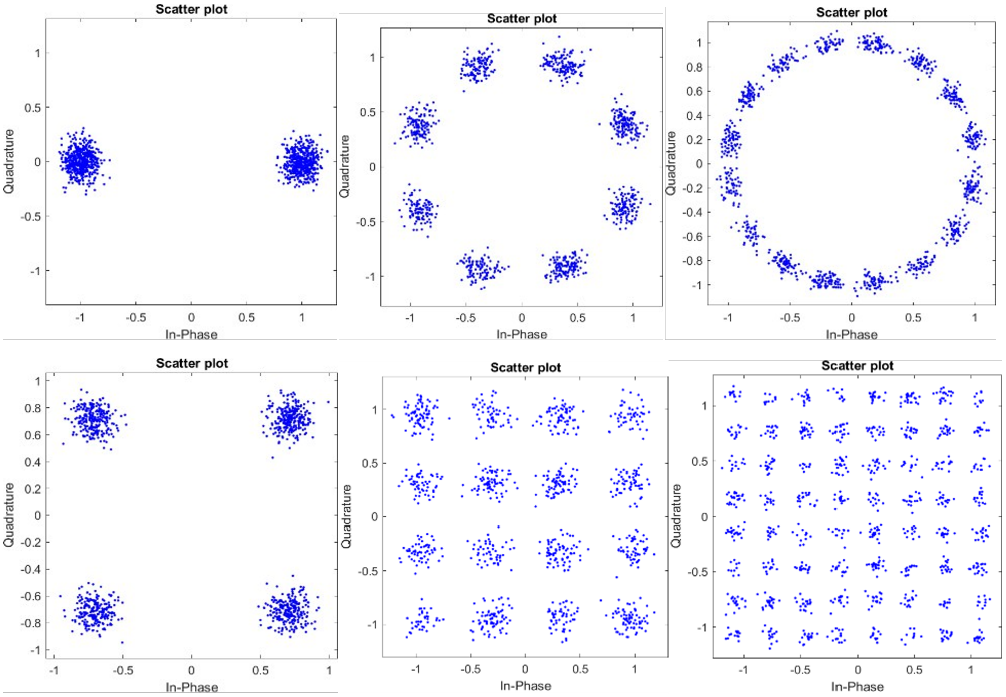

Figure 3: Constellation diagram of modulation schemes.

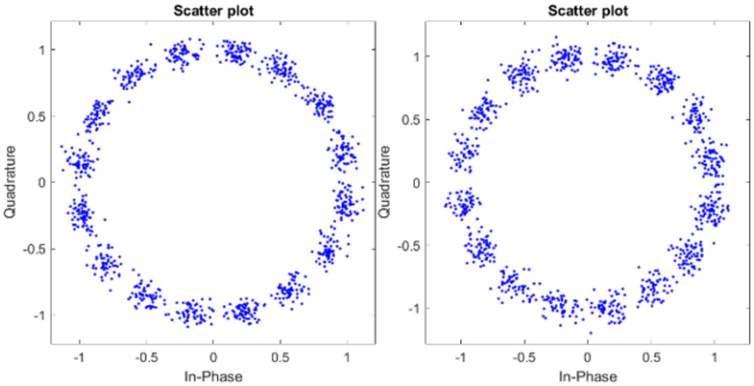

Figure 4: Sample constellation diagrams across variable SNR levels (22 dB, 20 dB, 18 dB).

tinct clusters for each modulation type. Only the first twelve components were retained, as higher-order components exhibited linear dependence and added little discriminative power.

The first two principal components provided nearly linearly separable clusters, confirming their usefulness for classification. The PCA results also suggested that kernel-based classifiers, such as SVMs, could potentially achieve strong baseline performance.
Figure 7 illustrates PCA visualization using three leading components.
Figure 8 shows nearly separable clusters using the first two components.
Figure 9 highlights separability using components 2 and 3.

Figure 5: Analog Devices ADALM-PLUTO SDR used in OTA experiments.

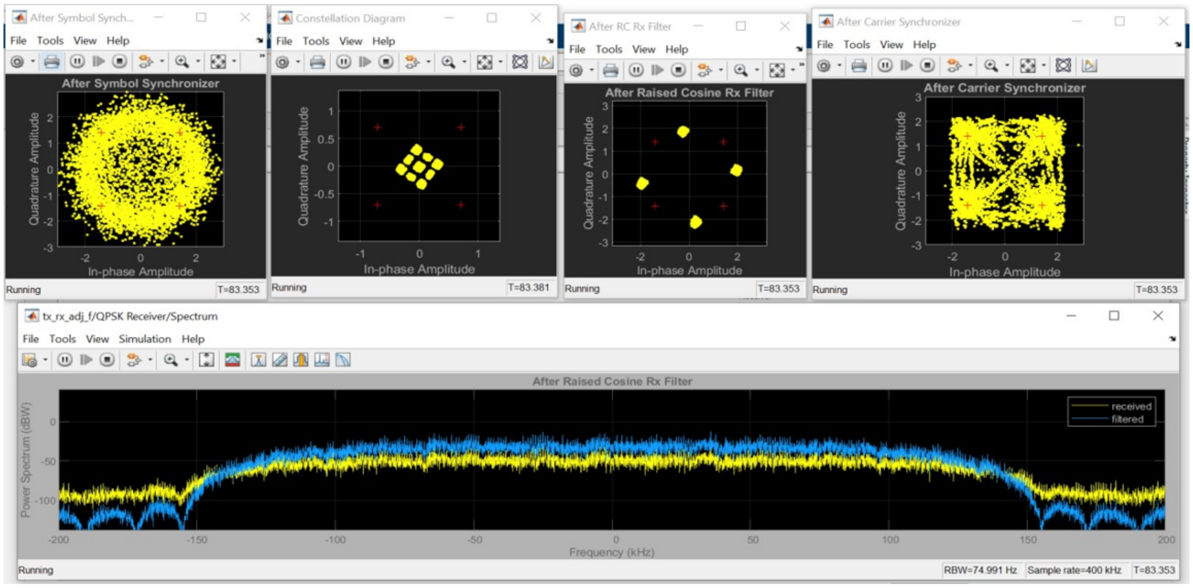

Figure 6: Example of received QPSK signal showing synchronization process.

K. Machine Learning Results

The custom 12-layer CNN was trained using an 80/20 train-validation split with multiple optimizers, including Adam, RMSProp, and stochastic gradient descent with momentum (SGDM). Initial experiments with a learning rate of 0.01 led to overfitting, with validation accuracy reaching 100% within the first epoch. Adjusting the learning rate to 3×10^{-4} reduced overfitting and produced smoother training curves. On purely synthetic data, the CNN achieved 94.48% classification accuracy using the SGDM optimizer, validating the architecture's ability to learn discriminative modulation features. Figure 10 shows the training and validation performance on the synthetic dataset.

When evaluated on OTA data, the same CNN achieved variable results between 64% and 91%

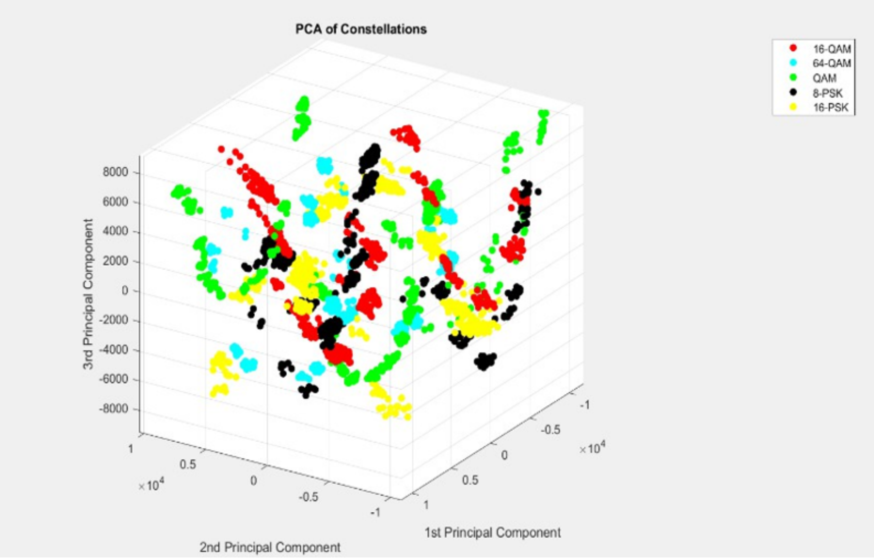

Figure 7: PCA visualization using components 1, 2, and 3.

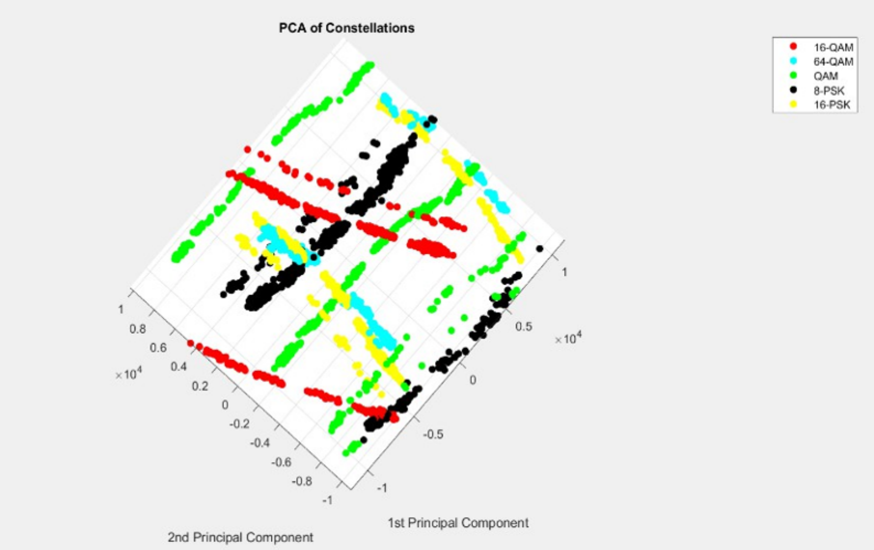

Figure 8: PCA visualization using components 1 and 2.

accuracy depending on the modulation type and channel conditions. Training required up to 20 epochs for OTA validation, compared to only 5 epochs for synthetic datasets. The observed variability was attributed to synchronization errors, frequency offsets, and unmodeled environmental impairments.

Figure 11 illustrates OTA validation curves, highlighting variability due to channel impairments.

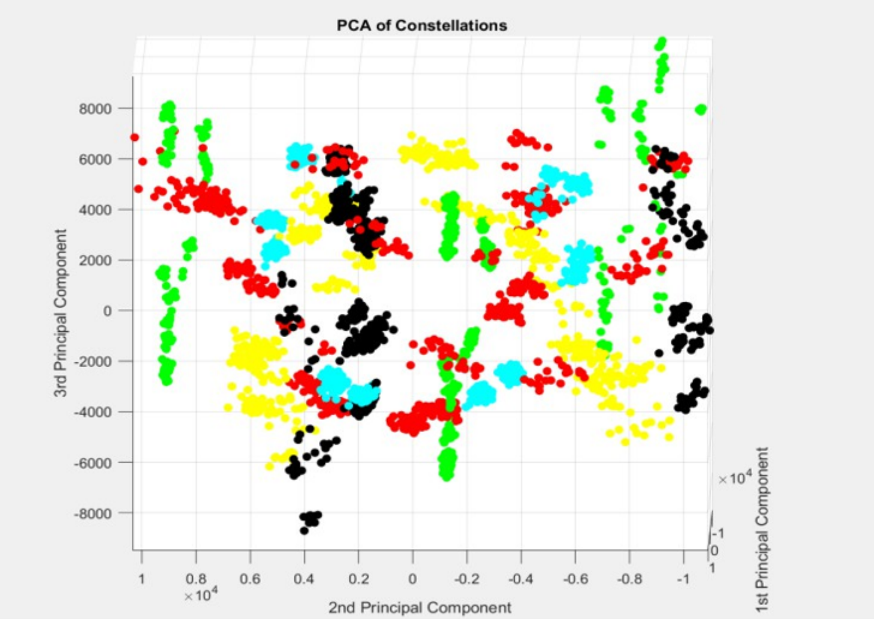

Figure 9: PCA visualization using components 2 and 3.

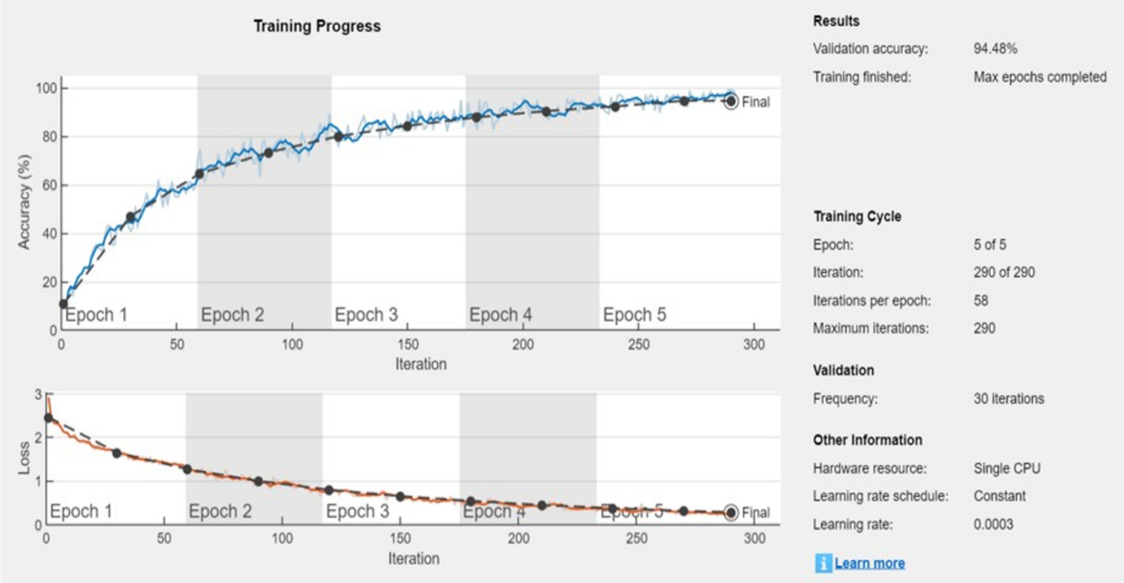

Results

Validation accuracy: 94.48%

Training finished: Max epochs completed

Training Cycle

Epoch: 5 of 5

Iteration: 290 of 290

Iterations per epoch: 58

Maximum iterations: 290

Validation

Frequency: 30 iterations

Other Information

Hardware resource: Single CPU

Learning rate schedule: Constant

Learning rate: 0.0003

Figure 10: Training and validation performance on the synthetic dataset.

CONCLUSION

This work presented the design, training, and evaluation of a deep learning-based framework for automatic modulation recognition (AMR) using constellation representations of synthetic and over-the-air (OTA) signals. A custom 12-layer convolutional neural network achieved up to 94.48% accuracy on a synthetically generated dataset that included multiple modulation schemes and varying SNR levels. These results demonstrate the potential of convolutional architectures to

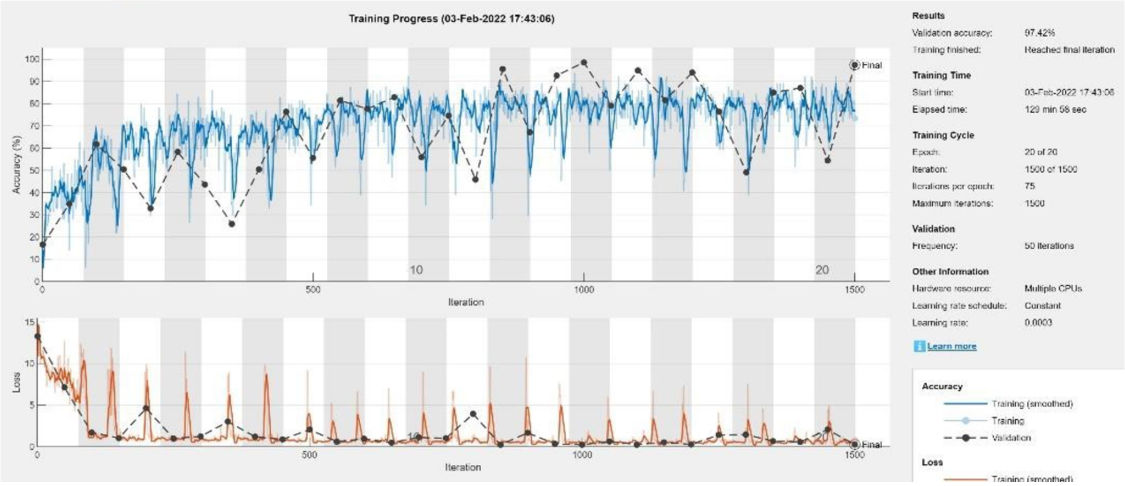

Results
Validation accuracy: 97.42%
Training finished: Reached final iteration

Training Time
Start time: 03-Feb-2022 17:43:06
Elapsed time: 129 min 58 sec

Training Cycle
Epoch: 20 of 20
Iteration: 1500 of 1500
Iterations per epoch: 75
Maximum iterations: 1500

Validation
Frequency: 50 iterations

Other Information
Hardware resource: Multiple CPUs
Learning rate schedule: Constant
Learning rate: 0.0003

Figure 11: Over-the-air validation results showing variability between training and validation curves.

extract discriminative features directly from raw signal constellations.

When validated on OTA signals collected with software-defined radios, the model exhibited classification accuracies ranging between 64% and 91%. The observed performance gap highlights the challenge of domain transfer from synthetic datasets to real-world environments, where additional impairments such as synchronization errors, oscillator drift, and channel variability play a critical role. Despite this gap, the results confirm that CNN-based modulation classification is feasible under practical conditions and can form the basis for robust, adaptive communication systems.

Future improvements will focus on enriching the synthetic dataset with a broader set of impairments, expanding OTA experiments across diverse environments, and exploring hybrid training strategies that combine simulated and real-world samples. These enhancements are expected to reduce the synthetic-to-real performance gap and strengthen the generalization of deep learning-based AMR. Ultimately, this line of research contributes to enabling resilient, adaptive, and spectrum-efficient wireless communication systems in contested and dynamic environments.

REFERENCES

[1] L. Wang and Y. Ren, “Recognition of digital modulation signals based on high order cumulants and support vector machines,” in *Proc. ISECS Int. Colloq. on Computing, Communication, Control, and Management*, vol. 4, (Sanya, China), pp. 271–274, Aug. 2009.

[2] G. J. Mendis, J. Wei, and A. Madanayake, “Deep learning-based automated modulation classification for cognitive radio,” in *Proc. IEEE Int. Conf. on Communication Systems (ICCS)*, (Shenzhen, China), pp. 1–6, 2016.

[3] M. Liang, Y. Yang, and W. Hongjun, “Dbn based automatic modulation recognition for ultra-low snr rfid signals,” in *Proc. 35th Chinese Control Conf. (CCC)*, (Chengdu, China), pp. 7054–7057, 2016.

[4] H. Wang, Z. Wu, S. Ma, S. Lu, H. Zhang, G. Ding, and S. Li, “Deep learning for signal demodulation in physical layer wireless communications: Prototype platform, open dataset, and analytics,” *IEEE Access*, vol. 7, pp. 30792–30801, 2019.

[5] H. Kulhandjian, N. Ramachandran, M. Kulhandjian, and C. D’Amours, “Human activity classification in underwater using sonar and deep learning,” in *Proc. ACM Int. Conf. on Underwater Networks and Systems (WUWNet)*, (Atlanta, GA), pp. 1–5, Oct. 2019.

[6] C. Wang, J. Wang, and X. Zhang, “Automatic radar waveform recognition based on time-frequency analysis and convolutional neural network,” in *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, (New Orleans, LA), pp. 2437–2441, 2017.

[7] K. Yashashwi, A. Sethi, and P. Chaporkar, “A learnable distortion correction module for modulation recognition,” in *Proc. IEEE Int. Conf. on Communications (ICC)*, (Kansas City, MO), pp. 1–6, 2018.

[8] T. J. O’Shea, J. Corgan, and T. C. Clancy, “Convolutional radio modulation recognition networks,” in *Proc. Int. Conf. on Engineering Applications of Neural Networks (EANN)*, (Athens, Greece), pp. 213–226, 2016.

[9] MathWorks, “Modulation classification with deep learning.” MathWorks Documentation, 2019.

[10] Analog Devices, “Adalm-pluto sdr active learning module.” Product Highlight, July 2019.

[11] C. Gravelle and R. Zhou, “Sdr demonstration of signal classification in real-time using deep learning,” tech. rep., Technical Report, 2019.

[12] T. Xu and I. Darwazeh, “Deep learning for over-the-air non-orthogonal signal classification,” *arXiv preprint arXiv:2201.12345*, 2022.

[13] H. Kulhandjian, P. Sharma, M. Kulhandjian, and C. D’Amours, “Sign language gesture recognition using doppler radar and deep learning,” in *Proc. IEEE GLOBECOM Workshop on Machine Learning for Wireless Communications*, (Waikoloa, HI), Dec. 2019.

[14] MathWorks, “Visualize high-dimensional data using t-sne.” MathWorks, 2018.