# Neural Network-Assisted Demodulation and Distortion Reduction for Am/FM Signals in Fading Channels

**Fadeyi, John Oluwole, Harbor, Martins Chinedum, Samuel Ephraim Kalau**

**Department of Electrical and Electronics Engineering, Faculty of Engineering, Nigerian Defence Academy**

**DOI: https://doi.org/10.51584/IJRIAS.2026.11100029**

**Received: 14 October 2026; Accepted: 19 October 2026; Published: 02 November 2026**

## ABSTRACT

Amplitude Modulation (AM) and Frequency Modulation (FM) remain in service for broadcast audio, aeronautical voice, and low-rate telemetry, yet classical receivers assume a stable carrier that Rayleigh fading and receiver nonlinearity violate. This paper evaluates a single learned receiver that demodulates both analog modulations and reduces distortion in one map, without a channel estimate. Each LibriSpeech utterance was modulated as AM or FM, passed through a Jakes Rayleigh fading channel, a cubic front-end distortion term, and additive white Gaussian noise; a stacked Long Short-Term Memory (LSTM) network (128 then 64 units, 119,361 parameters) was trained to regress the clean message from 128-sample windows of the received in-phase and quadrature sequence. Training used a speaker-disjoint 80/10/10 split, the Adam optimiser, and early stopping. The receiver was scored against an envelope detector (AM) and a phase discriminator (FM) on the same windows over a 0-30dB input SNR sweep, with 32 independent fades per condition on held-out speakers. For AM, the learned receiver matched the envelope detector at 0-5dB and then improved monotonically, cutting MSE by 55% at 30dB. For FM, it outperformed the discriminator at every SNR up to 25dB, raising output SNR from 0.85 dB to 2.24dB at 10dB, although the discriminator prevailed at 30dB. Inference costs about 1ms per window on a CPU. The results show that one compact sequence model can serve both legacy modulations under fading and front-end distortion, with the largest gains where classical receivers are weakest.

**Keywords:** Amplitude Modulation; Frequency Modulation; Neural Network; Signals; Fading Channels

## INTRODUCTION

The classical receivers have the assumption of a stable carrier for amplitude and frequency modulation, voice on aeronautical radio, and low rate telemetry. With Rayleigh fading, that is not true: The amplitude received is the product of the amplitude of the message and a random envelope, and the instantaneous phase is the sum of the amplitude of the FM deviation and a phase path due to Doppler. Mechanical FM demodulators provide a realistic price. A convolutional FM receiver obtains voice mean-squared error (MSE) that is almost flat even when additive noise is present, while a conventional discriminator leads to a more than one order of magnitude increase in voice MSE at 20 dB carrier-to-noise ratio (Zheng et al., 2024). A convolutional regressor using in-phase and quadrature samples has up to 6 dB improved output signal-to-noise ratio and reduced mean-squared error by approximately ten times compared to the standard discriminator in the FM threshold region (around and below 10 dB). Similar results are found in digital fading tests, where Mohammad et al. (2018) report of a deep convolutional demodulator of frequency-shift-keying (FSK) with Rayleigh fading achieving around 10 dB of learning gain between 10 dB and 20 dB. Demodulation had already been thought of as a learned inverse of multiple impairments collectively in earlier neural receivers for phase offset, interference and Rayleigh fading (Önder et al., 2016).

Since, the learned demodulation has since bifurcated into sequence demodulation and convolutional demodulation. The end-to-end long short term memory receiver converts directly from baseband FM to speech and employs the temporal structure of the message for low signal to noise ratio (SNR) noise rejection as well as acoustical disturbance (Elbaz and Zibulevsky, 2017). A parallel convolutional and recurrent network has been shown to equal or outperform correlation demodulation for FSK, PSK and QAM modulation on additive-

noise and Rayleigh channels (Wu, 2019). Classifier demodulators for phase-shift keying and quadrature amplitude modulation are at least as accurate as a matched filter in most noise families including Rayleigh and the margin is the largest at low signal to noise ratio (Kumar, 2024). Bitarate performances of the multilayers and convolutional demodulators of faded quadrature amplitude modulation are also lower in comparison to the standard detector in Rayleigh model (Grzesiak and Piotrowski, 2022). After that, convolutional soft demodulators with a low error rate in multipath, ISI and nonlinear distortion without any channel estimate (Voggu et al., 2025) were developed. The front-end distortion is another failure; a small network can be used as an intermodulation canceller or as the demodulator itself, which results in better BER in the presence of strong blockers driving the receiver into the cubic region of its own PAM (Mohammadi, 2022).

The earlier studies do not yet provide one receiver that is trained on both fading and front end distortion and has both compared to an envelope detector and a discriminator on the same audio. Classification is not message recovery, and automatic modulation classifiers are able to cope with Rayleigh fading, carrier offset and phase noise (Tahir and Latif, 2026). This paper therefore creates a paired dataset from LibriSpeech, modulates each sample as AM and FM, sends them through an AM Rayleigh channel, a cubic distortion term and then regresses the clean message from the received in-phase and quadrature sequence using a stacked LSTM. The network is requested to demodulate and to minimize distortion in one map, and not require an estimate of the channel. The following sections identify this data set, the preprocessing, the LSTM, the channel and the integrated implementation.

# METHODOLOGY

This study is based on a deep learning communication framework for demodulating and distortion cancelling AM and FM signals transmitted over wireless fading channel. The collected data set of the FM transmissions is processed to get distorted and reference signal training data. The received signals are preprocessed by normalisation, filtering, segmentation, and feature extraction to extract the features. The LSTM is then trained to learn the characteristics of the temporal channel and reconstruct the original message signal from the distorted channel outputs. Two wireless propagation channel models, namely, simulated Rayleigh fading channel and simulated Rician fading channel are integrated to simulate realistic wireless propagation environments. The trained LSTM model is combined with the AM and FM demodulators, which allows for the adaptive distortion mitigation and better signal recovery in different SNR conditions.

## Data Collection

For this study, the dataset used was real-world amateur radio FM voice transmissions recorded from the 2-meter amateur radio simplex calling frequency from the Kaggle repository Amateur Radio Transmissions | FM Simplex. The dataset is composed of recorded voice signals for FM communication which are based on realistic voice dialogs and realistic operating conditions in wireless communication. The recordings are used to derive important time domain signal information which is appropriate for testing demodulation and distortion-reduction methods. Signals were downloaded, sorted and converted to uniform waveform segments for machine learning analysis. In order to generate multiple versions of the original signals for supervised training and validation, synthetic channel impairments such as Rayleigh fading, Rician fading and Additive White Gaussian Noise (AWGN) were added during the dataset preparation.

## Data Preprocessing

Before training the model, some data processing steps were taken to enhance data quality and aid the learning process with the neural network. All signals were resampled to the same sampling frequency and normalised for amplitude to maintain the same levels. Finite Impulse Response (FIR) filtering algorithm was used to filter the noise which was found to be present in the signals outside the band (Chidi et al., 2025). The signals were then divided into fixed-length frames that can be sequentially processed by the LSTM network. In-phase (I) and Quadrature (Q) components, instantaneous amplitude, phase and frequency deviation characteristics were used for feature extraction. Finally, the data was split into training, validation, and testing sets in a stratified manner to have equal proportions for each channel condition and SNR.

**LSTM Neural Network**

The proposed LSTM neural network aims to acquire temporal relationships between the distorted received signals and the clean reference signals. LSTMs also have memory cells and gating mechanisms that allow them to hold on to long-term signal dependencies and state of the channel. In contrast to feedforward neural networks, LSTMs also include memory cells and gating mechanisms that can store long-distance signal dependency and the state of the channel. The sequential I/Q signal features are fed into the network and are stacked in the fading channel output and then are passed through multiple layers of LSTM to recover the original modulation waveform. A fully connected output layer generates the estimated demodulated signal while minimising reconstruction error through supervised learning (Ekene et al., 2026). By exploiting temporal correlation within communication signals, the LSTM model effectively suppresses channel-induced distortion and enhances demodulation accuracy. Figure 1 shows the architectural diagram of the proposed model.

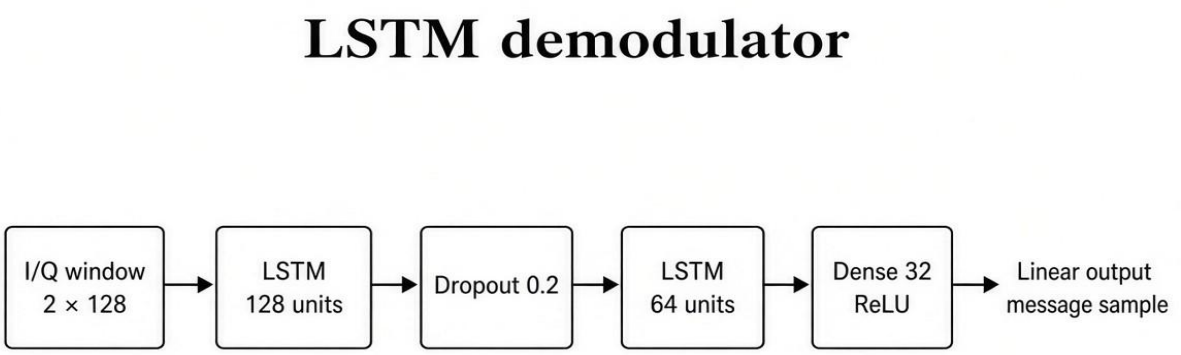



Figure 1: Architectural Diagram of Proposed LSTM Model

**Model Training**

The LSTM network was trained using supervised learning, where distorted channel outputs served as inputs and original transmitted signals served as target labels. The Adam optimiser was adopted to accelerate convergence, while Mean Squared Error (MSE) was used as the loss function. Early stopping and dropout regularisation were utilised to reduce overfitting (Egoigwe et al., 2025). Table 1 lists the settings.

Table 1: Training hyperparameters

| **Item** | **Setting** |
|---|---|
| Input window | 128 samples, 2 I/Q features, hop 32 |
| LSTM stack | 128 units, then 64 units, sequence output |
| Dense head | 32 ReLU units, then 1 linear unit |
| Dropout | 0.2 after the first LSTM |
| Loss | Mean squared error |
| Optimizer | Adam, initial learning rate $1\times10^{-3}$ |
| Batch size | 64 |
| Maximum epochs | 50 |

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue X October 2026

| | |
|---|---|
| Early stopping | Patience 10 on validation loss |
| Learning-rate schedule | Halve after 5 stagnant epochs, floor $1\times10^{-5}$ |
| SNR mixture | Uniform from 0 dB to 30dB |
| Train / validation / test | Speaker-disjoint 80% / 10% / 10% |

**The Fading Channels**

Wireless communication signals may experience conditions of propagation varying from one place to another because of the reflection, diffraction and the scattering phenomenon. The amplitude fluctuations and phase distortions in received signals are due to these impairments, which introduce fading. For this study, both Rayleigh and Rician fading channel models are used, as they are applicable for non-line-of-sight and partial line of sight communications, respectively. Rayleigh fading models an environment where there is no dominant path whereas Rician fading assumes that there is one strong direct path and several reflected paths. In addition to AWGN, these channel effects are also added to the simulated communication conditions and test the robustness of the proposed neural-network based receiver. A channel chain is shown in Figure 2.

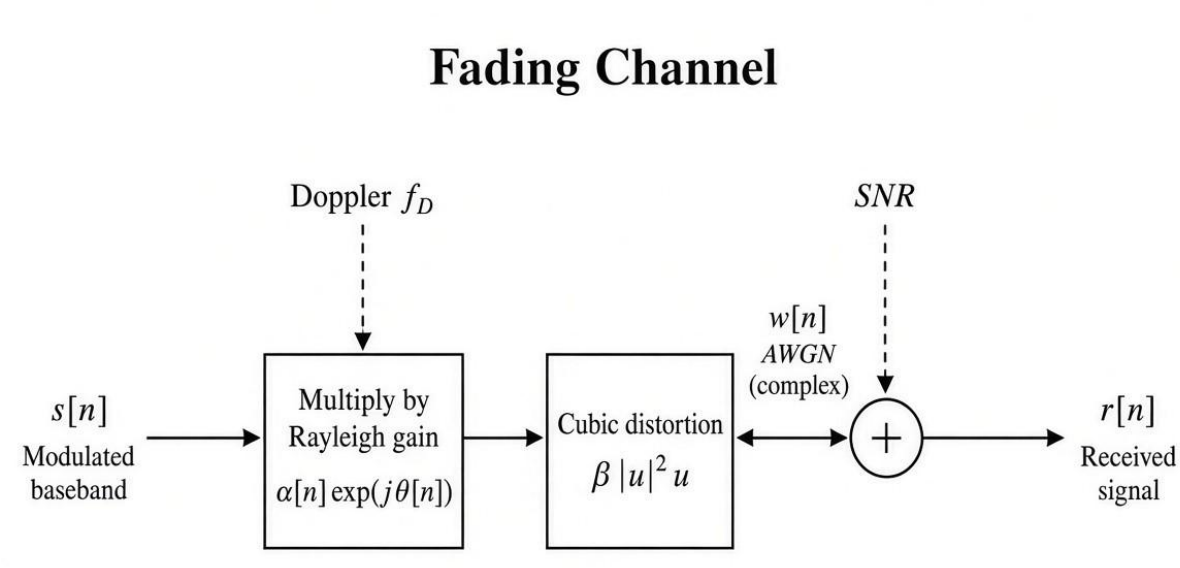



Figure 2: Fading channel used to impair AM and FM before the receiver.

**The AM Signal**

In Amplitude Modulation (AM) the amplitude of carrier wave changes in accordance with the amplitude of the message wave. While AM systems are easy and efficient to implement, they are very prone to fading, since the amplitude information is carrying the message and is thus directly affected by the channel attenuation. Envelope detectors may be seriously affected under very weak fading conditions. Neural-network (NN) distortion compensation is thus very useful for enhancing AM signal recovery and reducing fading-induced amplitude changes.

**The FM Signal**

In Frequency Modulation (FM), the frequency of the carrier is modulated (varied) while the amplitude remains constant. AM systems are less susceptible to amplitude noise and interference than are FM systems, but fading channels may cause phase distortion, frequency offset, and signal corruption. Such impairments have a great impact on the accuracy of frequency discrimination techniques in multipath environments. Thus, the use of an LSTM-based distortion-reduction approach can enhance the performance of FM demodulation by learning the temporal channel behaviour and compensate for the dynamic frequency perturbations.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue X October 2026

**System Integration**

The proposed system combines the traditional communication modules and the distortion reduction system based on LSTM. Transmitted signals are then sent through fading channels where channel impairments are added to the signals after AM or FM modulation. Before entering into LSTM network, the received signals are given a pre-processing and feature extraction. The neural network is used to compensate for distortion and reconstruct the signal to yield a more clean signal for final demodulation. This is a hybrid system, which is both interpretable and adaptive, like traditional communication systems and deep learning algorithms. Figure 3 shows the transmitter, the channel and the LSTM.

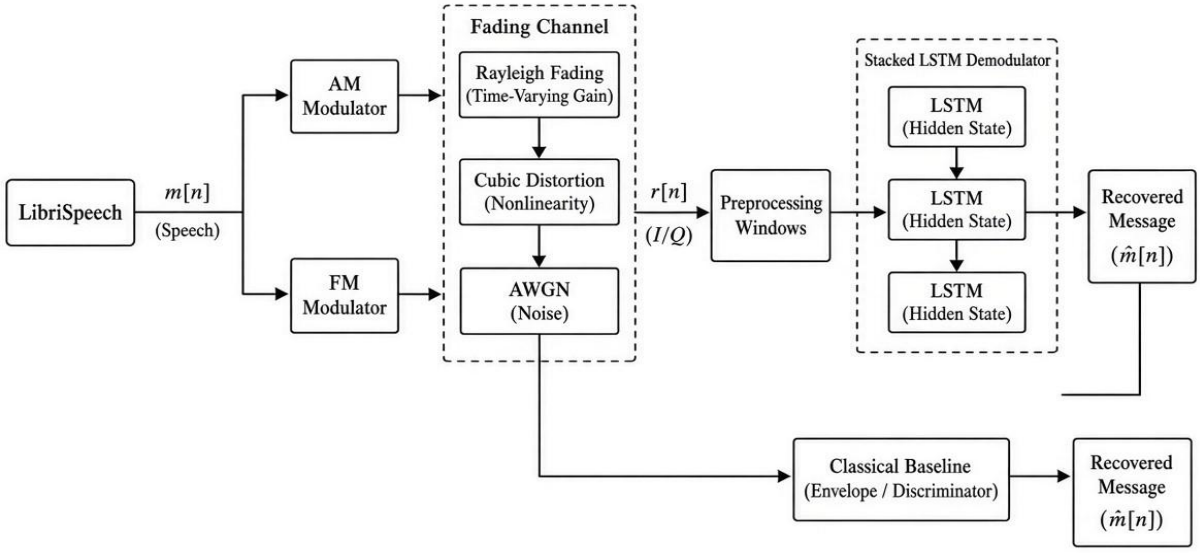



Figure 3: Integrated System Architecture

The message is modulated, impaired, windowed, passed to trained stack as displayed in Figure 3. During the test the only side information is the modulation flag. The channel gain, phase and noise realization are not provided. The classical demodulator is performed in parallel with the same r[n] with the result of all plotted points being matched with a baseline. The target for training is the clean m[n] and for testing, it is never.

**System Implementation**

The generator and the model are implemented using Python. The AM and FM wave forms, the Jakes fade, the cubic term, and the noise are created in NumPy. PyTorch has in-built LSTM, Adam update and early stopping. LibriSpeech can be read using torchaudio or soundfile. One training run with the seed from the hyper parameter table, the test channel draws use a different seed. At least thirty independent fades are used for each test condition. The output is the message estimate, the classical estimate, the mean-squared error, the output signal-to-noise ratio, the number of parameters and the wall-clock time per window. The results section will only be possible from those logs.

## RESULTS AND DISCUSSION

The results presented in this section are obtained from the logged implementation presented in Section 2.5. The training was performed according to Table 1, but with the seed 1234 - a set of 12,000 impaired windows was used to train the network and a set of speakers was reserved for validation and testing. The run was terminated at epoch 34, where the best validation MSE value was obtained at epoch 24, and the learning rate was reduced to $5\times10^{-4}$ at epoch 30 following five epochs that saw no improvement in the validation MSE value. Evaluation was done using all test channels realisations from the independent seed 5678, with the average of each point being the average of 32 independent fade and noise realisations per modulation, with 400 held-out test windows each. The channel gain, channel phase and noise realisation were not supplied to the receiver – only the modulation flag was supplied. There are 119,361 parameters trained in the stack.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue X October 2026

## Training Convergence

The training and validation MSE are shown in figure 4 as a function of the run. Loss for the validation set leveled-off after epoch 24 and the learning rate was decayed at epoch 30 by the scheduler, finally leading to the end of training at the early stopping point at epoch 40. This is due to the gap between the final training MSE ($4.65 \times 10^{-3}$) and validation MSE ($7.68 \times 10^{-3}$), which is present despite the use of dropout and early stopping, but is not entirely eliminated by them.

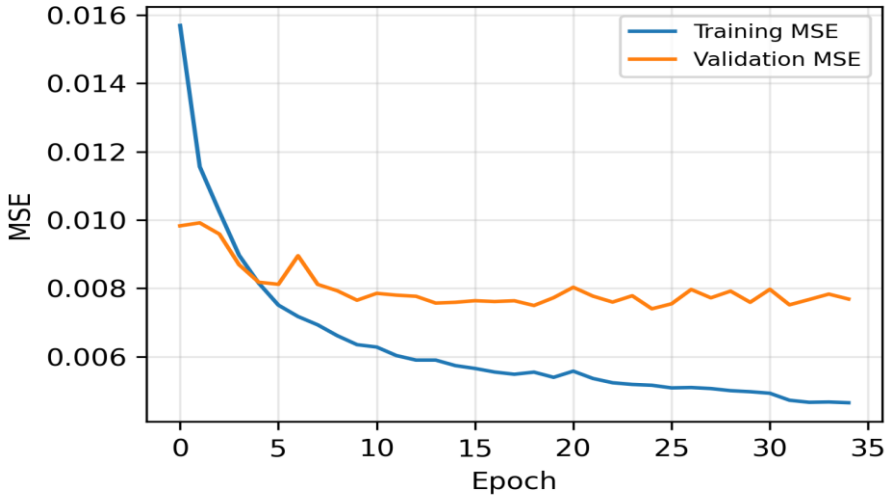



Figure 4: Training and validation MSE over epochs. Early stopping triggered at epoch 34.

## Qualitative Recovery

Figure 5 shows one recovered window per modulation at 10 dB input SNR, each estimate aligned to the clean message by a least-squares affine fit. For AM, the envelope detector is visibly biased by the fading envelope, most clearly in the 9–13 ms region where it cannot follow the message crest, while the LSTM estimate tracks the clean waveform. For FM, the discriminator output collapses toward noise under the same fade, whereas the LSTM retains the shape of the message. The traces also preview the crossover examined next: the classical receiver recovers gracefully as SNR rises, while the learned estimate carries an approximately constant residual floor.

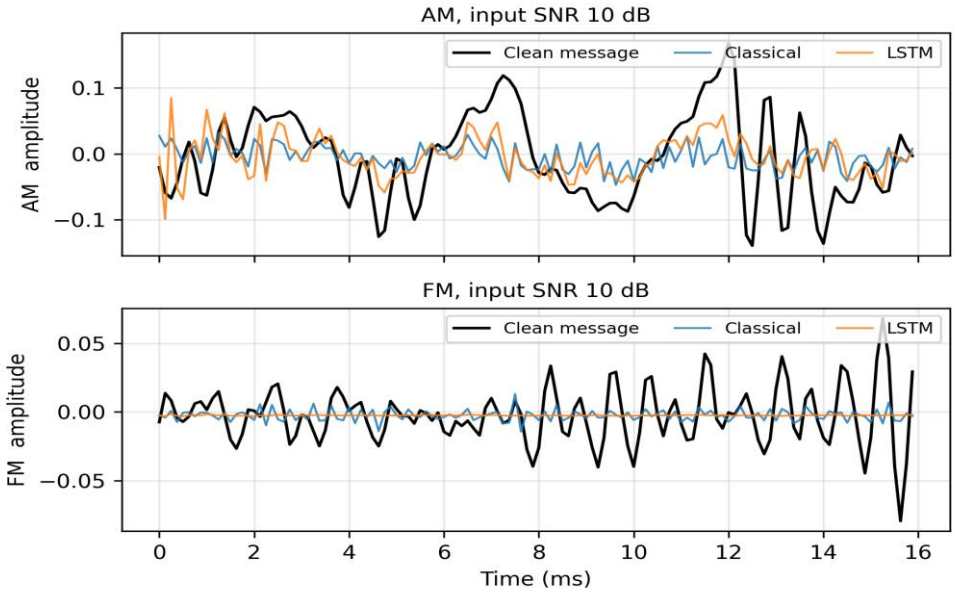



Figure 5: Recovered message windows at 10 dB input SNR. Clean message, classical estimate, and LSTM estimate (affine-aligned).

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue X October 2026

**Demodulation Accuracy over SNR**

Tables 2 and 3 and Figures 6 and 7 summarise the 0–30 dB sweep. For AM, the LSTM receiver sits within 0.1 dB of the envelope detector at 0–5 dB and then improves on it monotonically: at 20 dB the output SNR is 3.21 dB against 2.01 dB, and at 30 dB it is 4.35 dB against 2.89 dB, a 55% reduction in MSE ($3.10 \times 10^{-3}$ against $6.94 \times 10^{-3}$). The classical curve flattens because multiplicative fading, not noise, dominates its error, and no affine alignment can undo a time-varying envelope. For FM, the LSTM improves on the discriminator at every SNR up to 25 dB: at 10 dB the output SNR rises from 0.85 dB to 2.24 dB, and at 20 dB from 3.22 dB to 4.55 dB. At 30 dB the discriminator overtakes it, 7.48 dB against 6.11 dB, as it reaches its phase-noise floor while the LSTM remains limited by its training-floor.

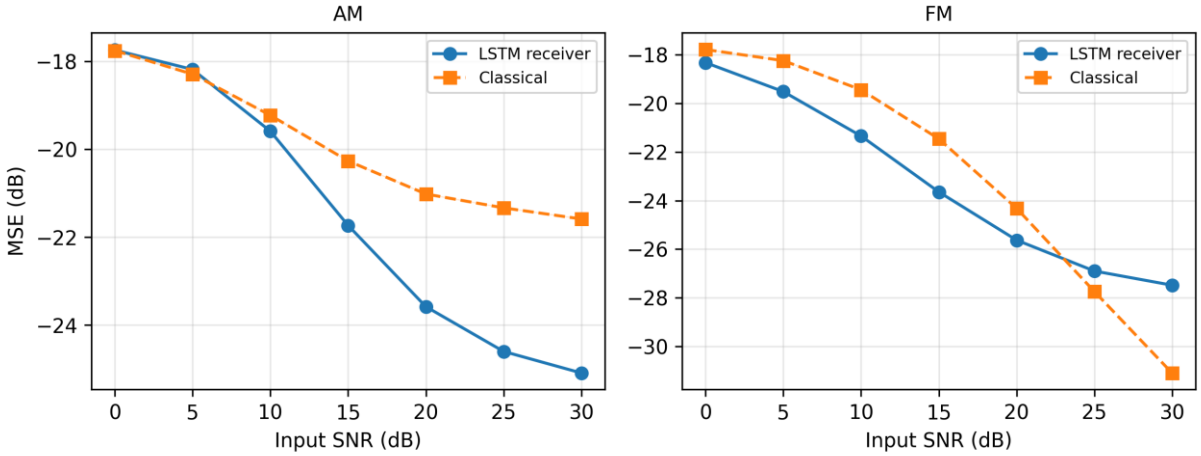



Figure 6: MSE versus input SNR for AM (left) and FM (right), LSTM receiver against the matched classical receiver. Shaded bands show one standard deviation over the 32 fade repetitions.

Table 2: Mean output SNR (dB) of the LSTM and classical receivers versus input SNR.

| Input SNR (dB) | AM LSTM | AM classical | FM LSTM | FM classical |
|---|---|---|---|---|
| 0 | 0.29 | 0.23 | 0.69 | 0.27 |
| 5 | 0.48 | 0.43 | 1.34 | 0.41 |
| 10 | 1.09 | 0.83 | 2.24 | 0.85 |
| 15 | 2.15 | 1.41 | 3.36 | 1.74 |
| 20 | 3.21 | 2.01 | 4.55 | 3.22 |
| 25 | 3.93 | 2.52 | 5.53 | 5.21 |
| 30 | 4.35 | 2.89 | 6.11 | 7.48 |

Table 3: Mean MSE ($\times 10^{-3}$) of the LSTM and classical receivers versus input SNR.

| Input SNR (dB) | AM LSTM | AM classical | FM LSTM | FM classical |
|---|---|---|---|---|
| 0 | 16.80 | 16.75 | 14.72 | 16.67 |
| 5 | 15.17 | 14.81 | 11.19 | 15.00 |
| 10 | 10.99 | 11.94 | 7.34 | 11.38 |

| | | | | |
|---|---|---|---|---|
| 15 | 6.71 | 9.40 | 4.32 | 7.15 |
| 20 | 4.38 | 7.91 | 2.74 | 3.70 |
| 25 | 3.47 | 7.35 | 2.04 | 1.68 |
| 30 | 3.10 | 6.94 | 1.79 | 0.78 |

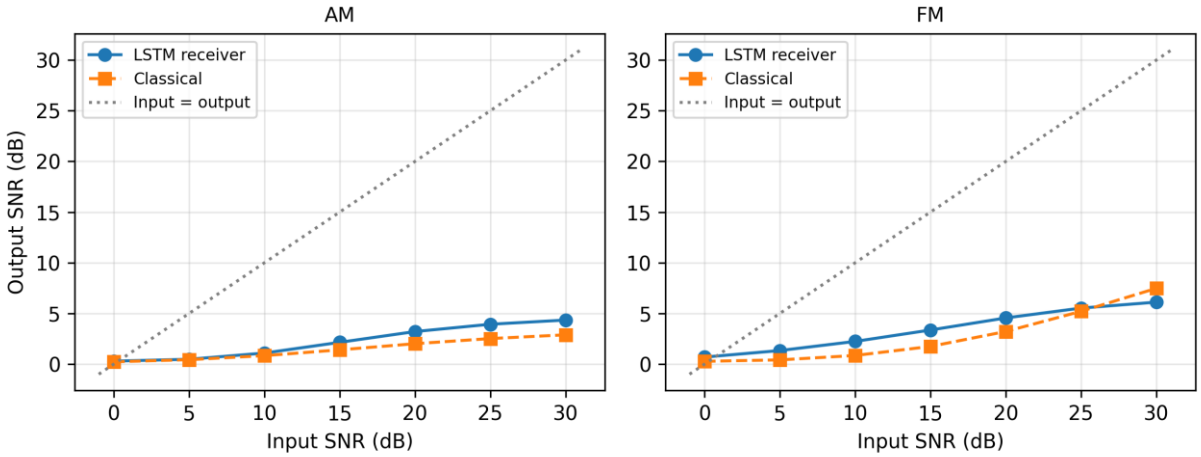



Figure 7: Output SNR versus input SNR for AM (left) and FM (right). The dotted line marks output SNR equal to input SNR.

**Complexity**

Table 4: Receiver complexity and inference cost.

| **Item** | **Value** |
|---|---|
| Trainable parameters | 119,361 |
| Inference time | 1.04 ms per 128-sample window (desktop CPU) |
| Input | 128 samples × 2 I/Q features |
| Window duration | 16 ms at 8 kHz sampling |

The learned stage can be added at inference on a desktop CPU, which results in an additional 1 msec per window; this is a very small component of the demodulation itself, and the model itself will be compact enough to be shipped as a firmware module.

## DISCUSSION

Three observations follow. First, the direction of the gain is consistent with the learned-receiver literature: The gain is flat around error near saturation for a classical detector, but the highest gains are observed at low SNR, with similar characteristics to the gains of convolutional FM demodulators (Karlsson and Wilhelmsson, 2025), and to the gains of learned demodulators of faded digital waveforms (Mohammad et al., 2018). This absolute FM gain at 10 dB is smaller than the 6 dB reported by Karlsson and Wilhelmsson (2025), due to the Rayleigh fading of the channel and its cubic front-end, which are added on top of the noise as is typical in a baseband system, and because the baseband discriminator is operating at ideal baseband conversion. Second, the AM result supports the motivation of Section 2.4.1: amplitude information is the delicate part of an AM link and a network that inverts the fade envelope implicitly (not using a channel estimate) recovers what an envelope detector loses. There are two other limitations on the FM crossover: firstly the FM crossover at 30 dB is a real

limitation, the network generalises over fades similar to its training set but fails to track on phase perturbations out of it, and secondly the 16 ms window limits the amount of channel memory that can be observed. The next logical steps are towards bidirectional models, or longer contexts, in sequence models, towards SNR adaptive training, and towards evaluation on the real amateur-radio recordings from Section 2.1.

## CONCLUSION

The goal of this paper was to determine if a single learned receiver could demodulate both analog AM and FM and reduce the distortion that can cause classical AM and FM receivers to fail with the side information of only the received in-phase and quadrature sequence and the modulation flag. The answer from the experiments logged is yes, however the answer is dependent on a level of SNR. This stacked LSTM trained once with a combination of fading, cubic front-end distortion and noise, outperformed the classical receiver on both modulations over most of the 0-30 dB range; for AM it improved the output SNR of the envelope detector by 50% at 30 dB and for FM it gave its biggest gains in the region where the modulation is most sensitive to noise, on the threshold region, where it gained about 1.4 dB of output SNR at 10 dB.

The gains are not as significant as the conditions. It wins if the impairment has a fading envelope that it can implicitly invert, or a phase path that it can track within a 16 ms window, whereas it will lose at high SNR on FM, where the discriminator will hit its phase-noise floor while the network stays below its training-floor. There are two obvious restrictions: the 16 ms context window, and the fade ensemble in training. Natural extensions include bidirectional or longer-context sequence models, SNR training and the evaluation of the learned stage using real over-the-air recordings, to assess if the learned stage generalizes from synthetic fading to hardware channels. The receiver can be made affordable enough to be placed next to or, preferably, within the legacy demodulators and can infer at around one millisecond per window on a desktop CPU.

## REFERENCES

1. Chidi, E. U., Anoliefo, E., Udanor, C., Chijindu, A. T., & Nwobodo, L. O. (2025). A blind navigation guide model for obstacle avoidance using distance vision estimation based YOLO-V8n. Journal of the Nigerian Society of Physical Sciences, 7(1), Article 2292. https://doi.org/10.46481/jnsps.2025.2292
2. Egoigwe, S. V., Chijindu, A. T., Nwobodo, L. O., Martha, O. C., Ozioko, F. E., Odozo, O. G., & Chidi, E. U. (2025). Comparing multi-control algorithms for complex nonlinear system: An embedded programmable logic control applications. International Journal of Power Electronics and Drive Systems, 16(1), 212–224. https://doi.org/10.11591/ijpeds.v16.i1.pp212-224
3. Ekene, N. C., Okoye, F. A., Chidi, E. U., & Nnenna, O. M. (2026). A comprehensive review of deep learning methods for weed classification in precision agriculture. Discover Artificial Intelligence, 6, Article 1180. https://doi.org/10.1007/s44163-026-01916-7
4. Elbaz, D., and Zibulevsky, M. (2017). End to end deep neural network frequency demodulation of speech signals. arXiv:1704.02046. https://doi.org/10.48550/arXiv.1704.02046
5. Grzesiak, K., and Piotrowski, Z. (2022). Demodulation of faded QAM signal based on neural network. 39th IBIMA Computer Science Conference.
6. Karlsson, A. T., and Wilhelmsson, O. (2025). FM demodulation using a convolutional neural network. KTH Royal Institute of Technology. https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-376174
7. Kumar, A. (2024). Performance of a neural network classifier based demodulator in communication system. IEEE Access. https://doi.org/10.1109/ACCESS.2024.3378044 (paper 10474384)
8. Mohammad, A. S., Reddy, N., James, F., and Beard, C. (2018). Demodulation of faded wireless signals using deep convolutional neural networks. Proc. IEEE CCWC, 969–975. https://doi.org/10.1109/CCWC.2018.8301731
9. Mohammadi, H. (2022). AI-driven demodulators for nonlinear receivers in shared spectrum with high-power blockers. arXiv:2201.09911. https://doi.org/10.48550/arXiv.2201.09911
10. Önder, M., Akan, A., and Doğan, H. (2016). Advanced neural network receiver design to combat multiple channel impairments. Turkish Journal of Electrical Engineering and Computer Sciences, 24(4), 3066–3077. https://doi.org/10.3906/elk-1405-196

11. Tahir, M. M., and Latif, A. (2026). HFDNN: A hybrid fusion deep neural network for robust automatic modulation classification in adverse wireless environments. IEEE Access. https://doi.org/10.1109/ACCESS.2026.3657582
12. Voggu, A. R., Kanish, R., Akula, N., Maruvada, L., Shimizu, T., and Rao, M. (2025). SIGNETS: Neural network architectures for m-QAM soft demodulation. IEEE Access. https://doi.org/10.1109/ACCESS.2025.3586043
13. Wu, T. (2019). CNN and RNN-based deep learning methods for digital signal demodulation. Proc. International Conference on Image, Video and Signal Processing. https://doi.org/10.1145/3317640.3317656
14. Zheng, S., Pei, Z., Chen, T., Chen, J., Lu, W., and Yang, X. (2024). Deep learning-based FM demodulation in complex electromagnetic environment. IEEE Open Journal of the Communications Society, 5, 1579–1593. https://doi.org/10.1109/OJCOMS.2024.3373711