

AI Data Classification of Atmospheric Ducts for Advanced Radar Modeling

Charlie Paterson^{1,*}, Rebecca Caves², Chloe Peet², William Dawber²

¹University of Portsmouth, School of Mathematics and Physics, Portsmouth, PO1 3FX, United Kingdom

²QinetiQ, Portsmouth, PO6 3RU, United Kingdom

*Corresponding author: cpaterson128@gmail.com

Abstract - This study explores the development of neural network models for predicting atmospheric refractivity profiles from the input of sea clutter returns, with a focus on two distinct duct types: evaporation ducts and hybrid ducts. Accurate prediction of these profiles is critical for various defense and communication applications, where atmospheric refractivity significantly influences radar performance and signal propagation. To capture the differing characteristics of each duct type, three models were evaluated: two specialized models trained independently on evaporation and hybrid duct data, and a joint model trained on both. Prior to training, the dataset was balanced by down-sampling the more frequent hybrid duct cases to prevent bias. Model performance was assessed using standard regression metrics including Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2). The specialized evaporation model achieved an R^2 score of 0.9894 and an RMSE of 0.10322, while the hybrid duct model demonstrated superior precision with an R^2 of 0.9996 and RMSE of just 0.02095. The joint model maintained strong overall performance ($R^2 = 0.9965$, RMSE = 0.0591), though slight performance degradation was noted in evaporation duct cases, suggesting mild overfitting to the structurally complex hybrid ducts. Model predictions closely matched true refractivity profiles across both duct types, including cases with steep gradient transitions. These results underscore the capacity of neural networks to generalize well across atmospheric conditions. While joint models offer practical efficiency, specialized models may provide enhanced accuracy for operational scenarios where precision in refractivity prediction is critical.

Keywords - Atmospheric Refractivity, Neural Networks, Ducting, Evaporation Duct, Hybrid Duct, Machine Learning, Regression, Propagation Modeling, Refractivity Prediction, Environmental Data

1 Introduction

The performance of the maritime radar systems is heavily influenced by radio frequency propagation conditions, particularly the refractivity profile of the atmosphere. Variations in temperature, humidity, and pressure can cause changes in the refractivity of a radar wave, affecting how they propagate. One of the most significant challenges is atmospheric ducting, where radar signals bend and extend over the horizon, beyond their normal range, due to temperature inversions and humidity gradients. This can lead to false detections, misinterpretations of target locations, and unexpected blind spots, impacting maritime situational awareness and operational efficiency.

This project aims to address this issue by developing a machine learning model to predict the refractivity profile using an input of radar sea clutter returns. By leveraging complex relationships

between radar clutter data and atmospheric conditions, the model will provide real-time estimations of propagation effects, allowing operators to better anticipate changes in radar detection range. The neural network model, implemented in MATLAB, will enable near-instant "now-casting" of radar propagation, optimizing radar performance and reliability.

If successful, this approach will offer a data-driven solution for predicting radar visibility under varying atmospheric conditions, improving maritime navigation, surveillance, and defense operations. By improving the prediction of the refractivity profile, this research will enhance the accuracy and adaptability of maritime radar systems, leading to more reliable detection, tracking, and operational effectiveness.

2 Theory

2.1 Wave Propagation

To understand how radar waves interact with the atmosphere, it is important to review the physics of electromagnetic (EM) wave propagation. EM waves, like all electromagnetic radiation, travel at the speed of light in a vacuum, but their speed and direction can change in different media such as air, water vapor, or clouds. These transverse waves consist of oscillating electric and magnetic fields that are perpendicular to each other and the direction of travel. Their behavior is governed by Maxwell's equations. These are a set of four fundamental equations that describe how electric and magnetic fields behave and interact. They form the foundation of classical electromagnetism, optics, and electric circuits.

Gauss's Law for Electric Fields

$$\nabla \cdot \vec{E} = \frac{\rho}{\epsilon_0} \quad (1)$$

where:
 \vec{E} = electric field
 ρ = charge density
 ϵ_0 = permittivity of free space

Gauss's Law for Magnetic Fields

$$\nabla \cdot \vec{B} = 0 \quad (2)$$

where:
 \vec{B} = magnetic field

Faraday's Law of Induction

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \quad (3)$$

where:
 $\frac{\partial \vec{B}}{\partial t}$ = a change in magnetic field

Ampere-Maxwell Law

$$\nabla \cdot \vec{B} = \mu_0 \vec{J} + \mu_0 \epsilon_0 \frac{\partial \vec{E}}{\partial t} \quad (4)$$

where:
 \vec{J} = current density
 μ_0 = permeability of free space

Together, these equations explain how electromagnetic waves (like radar or light) are generated and propagate through space. They show that a changing electric field creates a magnetic field and a changing magnetic field induces an electric field. These interactions allow EM waves to self-propagate through a vacuum or medium.

2.1.1 Wavelength and Frequency

The relationship between wavelength (λ), frequency (f), and wave speed (v) is described by:

$$v = f\lambda \quad (5)$$

where:
 v = speed of wave propagation
 f = frequency of the wave
 λ = wavelength

In the atmosphere, the speed of the wave is slightly less than in a vacuum due to the refractive index (n) shown in the equation below:

$$v = \frac{c}{n} \quad (6)$$

where:
 v = speed of wave propagation
 c = speed of light
 n = refractive index

The refractive index indicates how much the electromagnetic wave slows down when passing through a medium compared to a vacuum. A refractive index greater than 1 means the wave travels slower in that medium than in free space.

2.2 Introduction to Atmospheric Refractivity

Radar waves don't always travel in straight lines through the atmosphere. Instead, they bend due to changes in atmospheric conditions—specifically temperature, pressure, and humidity. This bending is called refraction and it occurs from 'light rays as they pass from one medium to another' ([1]). This causes it to deviate towards or away from the normal, similar to how a straw appears bent in a glass of water. The degree of bending is governed by Snell's Law, which defines 'the angles of incidence and transmission' ([2]) when a wave encounters a boundary.

$$n = \frac{v_1}{v_2} = \frac{\sin(\theta_1)}{\sin(\theta_2)} \quad (7)$$

where:
 n = refractive index
 v_1 and v_2 = velocity of waves in first and second mediums
 θ_1 = the angle of incidence
 θ_2 = the angle of refraction

This law is essential in optics but also applies to radar waves, which can slowly bend over long distances when conditions change gradually. To better capture how the atmosphere affects radar waves, we use a quantity called modified refractivity, which incorporates temperature, pressure, humidity, and altitude into a single value using equation (8):

$$M = 77.6 \frac{p}{T} - 5.6 \frac{e}{T} + 375000 \frac{e}{T^2} + 0.157z \quad (8)$$

where:
 M = modified refractivity
 p = pressure
 T = temperature
 e = vapour pressure
 z = height

The quantity 'M' from equation (8) represents a scaled-up version of the refractivity difference. Since changes in atmospheric refractivity are typically very small, this scaling helps machine learning models detect more significant patterns. Under normal atmospheric conditions, refractivity decreases gradually with altitude, causing radar waves to bend gently and follow a predictable path. The value of 'M' changes slowly with height under normal conditions, causing radar waves to bend slightly. However, when this gradient changes rapidly, it can lead to anomalous propagation which 'occurs when the refractive index is modified by changes in temperature gradient' ([3]), leading to unexpected

radar performance. Figure (1), shows how a radar wave can go through various different types of refraction, sub-refraction, super-refraction and ducting.




Figure 1:

From ([4]) showing the different types of refraction that can occur when a radar wave is placed under anomalous propagation.

Figure (1) shows a number of different variations of bending against the Earth's surface. This project focuses on when a radar wave is bent towards the Earth's surface - ducting.

2.2.1 Ducting

Ducting occurs when atmospheric conditions create channels that trap and refract electromagnetic waves, affecting their long-distance propagation and extending the normal range. A primary cause is temperature inversions, where "cold air at the surface gets trapped under a layer of warmer air" ([5]), disrupting the normal temperature gradient and altering wave propagation. Humidity gradients, defined as the "variation in concentration of their own vapor" ([6]), also significantly affect refraction. Sudden humidity shifts, especially over oceans and coastal areas, can unpredictably extend radar ranges or bend signals.

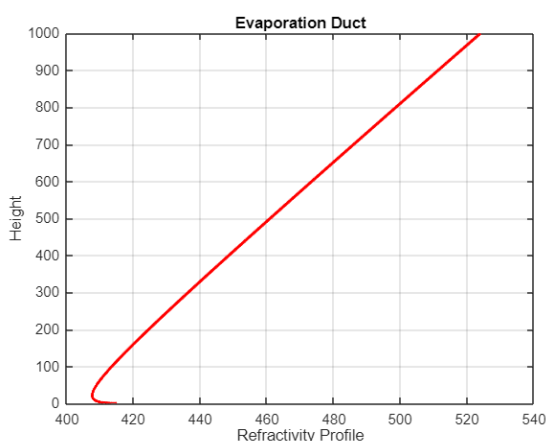


Figure 2: This illustrates the refractivity profile of an evaporation duct.

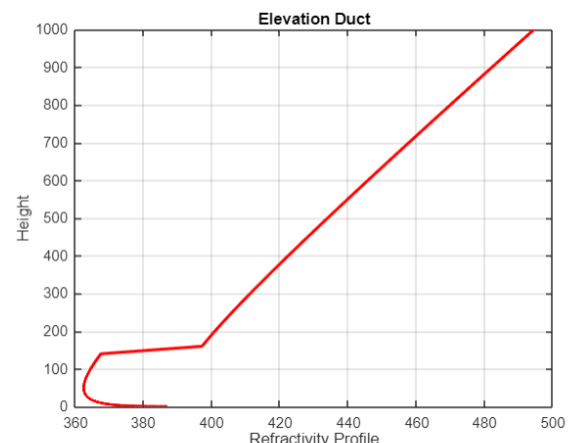


Figure 3: This illustrates the refractivity profile of a hybrid duct.

This project focuses on two duct types: evaporation ducts and hybrid ducts. Evaporation ducts, common over the sea, are "special atmospheric stratification that can affect the propagation path of electromagnetic waves at sea" ([7], p.1), shown in Figure 2. Hybrid ducts, characterized as "abnormal atmospheric refraction structure with a suspended trapped layer" ([8], p.1), combine features of evaporation and elevated ducts (Figure 3). Elevated ducts form at higher altitudes and extend wave range through atmospheric refraction driven by humidity gradient changes.

2.3 Radar Range Equation

The range at which a radar system can detect a target is given by the radar range equation:

$$R = \sqrt[4]{\frac{P_s G^2 \lambda^2 \sigma}{P_e (4\pi)^3}} \quad (9)$$

where:

R = radar range [m]
 P_s = transmitted power [W]
 G = antenna gain
 λ = wavelength [m]
 σ = radar cross section [m^2]
 P_e = received power [W]

This equation estimates the maximum range under normal conditions. However, when ducting occurs, radar signals can travel much farther than this equation predicts, making it crucial to understand atmospheric refractivity for accurate radar performance.

2.4 Introduction to Machine Learning

Machine learning (ML), a branch of AI, enables computers to learn from data and make predictions without explicit programming. Unlike traditional physics-based models, ML can capture complex patterns directly from data and adapt with new information, enabling accurate, real-time refractivity predictions that enhance radar performance in changing atmospheric conditions. In this project, a feed-forward neural network, like the one shown in figure (4), will be employed as a regression model to predict atmospheric refractivity profiles, based on the input features derived from clutter returns.




Figure 4:

Adapted from ([9]), this figure provides a visual representation of the structure of a simple neural network. It highlights the flow of data through input, hidden, and output layers, illustrating how the network processes information to make predictions.

Figure (4) works by processing input values, combining them with weights, adding a bias, and passing the result through an activation function shown in equation (10).

$$\hat{y} = f\left(\sum_{i=1}^n w_i x_i + b\right) \quad (10)$$

where:

\hat{y} = predicted output
 x_1, x_2, \dots, x_n = input variables
 w_1, w_2, \dots, w_n = weights associated with each input variable
 b = bias term
 f = activation function applied to output

The network is trained on radar data paired with known refractivity profiles, learning to map inputs to outputs. Performance is evaluated on separate test data to ensure the model generalizes well.

Once trained, the model can provide fast, accurate refractivity estimates to improve radar detection, especially under anomalous propagation conditions like ducting.

3 Material and Methods

3.1 Data

This section outlines the dataset used in this study, which includes evaporation and hybrid ducts and is vital for developing and evaluating the machine learning model. Data cleaning is essential to ensure quality, as raw data must be properly formatted for ML algorithms.

3.1.1 Data Collection

The data for this project, provided by my supervisors, consisted of two folders: one with clutter return files and one with refractivity profile files, each corresponding to specific ducts. Since both sets were initially unstructured, I first cleaned and organized the data to ensure compatibility with the machine learning workflow. This careful preparation established a reliable foundation for smooth and efficient subsequent processing.

3.1.2 Data Cleaning

After acquiring the data, I used MATLAB to clean and organize it in preparation for the machine learning workflow. The sea clutter returns and refractivity profiles were stored in separate folders, so I wrote a script to navigate through each directory, loop over all filenames, and extract the contents of each file into two separate data frames. Each data frame was constructed by iterating through the files, reading the relevant data, and appending it with the corresponding filename as a key. To merge the refractivity profiles with their associated sea clutter returns, I implemented a loop that matched entries in both data frames based on identical filenames. Once matched, I combined the two data frames into a single, unified table. This final table contained five essential columns: filename, duct strength, refractivity profile, sea clutter returns, and duct height. The resulting structure ensured that all relevant information for each case was aligned and accessible, providing a clean and consistent input format for the machine learning model.

3.1.3 Splitting

I structured the data by designating the sea clutter returns as the input variables and defining the refractivity profile as the output variable. This setup enabled the machine learning model to learn the relationship between the observed sea clutter returns and the corresponding refractivity profile. Figure (5) shows a sample of the sea clutter return data in the units of decibels.

As seen in figure (5) the strength of the sea clutter data gets smaller as the range goes further. The data for sea clutter returns was originally given to me in linear form, so to produce this graph in decibels, I had to use this equation:

$$Clutter(dB) = 20 \cdot \log_{10}(Clutter(linear)) \quad (11)$$

Before splitting the data into training and testing sets, I observed that the dataset contained significantly more samples for hybrid ducts than for evaporation ducts. To prevent the neural network from developing a bias toward the more prevalent hybrid duct type, I resized the datasets by reducing the number of hybrid duct samples to match the number of evaporation duct samples. This balancing step ensured that both duct types were equally represented during training and testing, promoting a more robust and fair learning process across different atmospheric conditions. I divided the data into training (80%), validation, and testing sets to support robust model development and evaluation. This split allowed the model to learn from a substantial portion of the data, tune its parameters, and assess its performance on unseen data. To further enhance evaluation reliability, I applied cross-validation, which uses multiple data splits to reduce over-fitting and provide a more consistent estimate of the model's generalization ability.

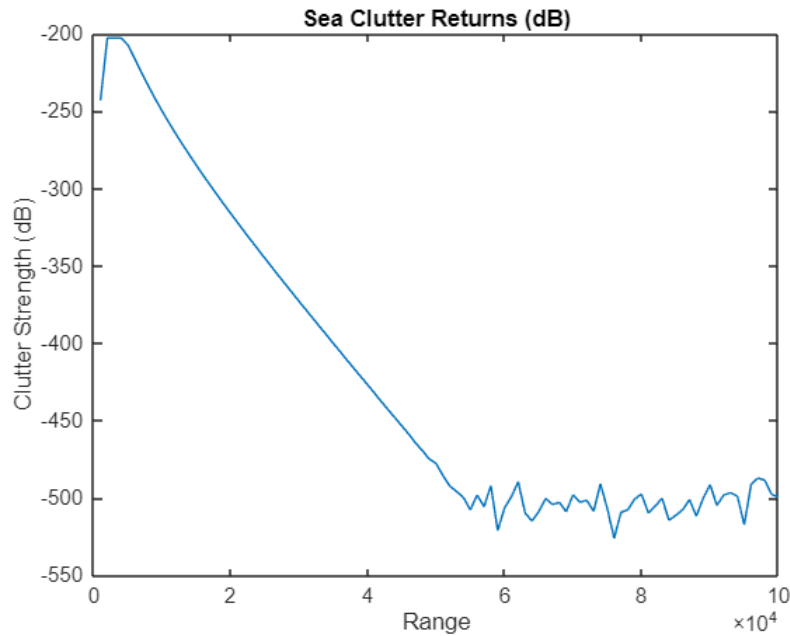


Figure 5:

This figure presents a sample from the dataset, illustrating a specific sea clutter return measured in decibels (dB).

3.1.4 Normalization

After this, I applied data pre-processing techniques to normalize the data as follows:

$$X_{norm} = \frac{X - \mu}{\sigma} \quad (12)$$

where:

X_{norm} = normalized value of the data point

X = original value of the data point

μ = mean of all data points within the feature

σ = standard deviation of all data points within the feature

f = activation function applied to output

Equation (12) illustrates z-score normalization. This technique helps reduce the impact of outliers, ensuring that extreme values do not disproportionately affect the model's performance. By normalizing the features in this way, each feature is treated with equal importance, and the model can learn more effectively from the data. This normalization step was crucial for ensuring that all features contributed equally to the model's predictions, particularly when using algorithms sensitive to the scale of input data. Once the data was normalized, I then had to transpose the data so that it fit the correct format of the MATLAB neural network as shown in Appendix C.

3.2 Neural Network

3.2.1 Neural Network Architecture

Predicting refractivity profiles from radar sea clutter data is complex due to the non-linear relationships involved. The objective is to estimate a 599-dimensional output vector from 100 input features. To tackle this, a deep feed-forward neural network was designed to capture and model these patterns effectively. The neural network consisted of 7 hidden layers where it captured complex and more abstract patterns to be able to learn through the training process. I used fully connected layers, allowing for the neurons within each fully connected layer to become connected to every neuron within the previous layer. I also used ReLU activation functions - shown in equation (13) - in between each fully connected layer as it would introduce non-linearity within the network so that it can capture more complex relationships by setting negative inputs to zero.

$$f(x) = \max(0, x) \quad (13)$$

where:
if input $x > 0$, ReLU function passes through as it is, if $x < 0$, it returns to 0.

The decreasing number of neurons across layers follows a funnel structure, which gradually reduces the data dimensionality while maintaining the most informative features.

3.2.2 Training Configuration

The training configuration defines various strategies that dynamically adjust the neural network's learning process throughout the training phase. These variations are essential for optimizing performance and ensuring efficient convergence. The neural network was trained using the Adam optimizer over 300 epochs with an initial learning rate of 5×10^{-5} . A piecewise schedule reduced the learning rate by a factor of 0.5 every 100 epochs to support better convergence. Training used mini-batches of 64 samples with data shuffled each epoch to improve generalization. L2 regularization (0.0001) was included to reduce over-fitting. Model performance was validated every 10 iterations, with early stopping triggered after 20 validation checks without improvement.

4 Results and Discussion

4.1 Training Performance

Three separate models were trained for this study to assess their individual and combined predictive capabilities: an evaporation-only model trained solely on evaporation ducts, a hybrid-only model trained only on hybrid ducts, and a joint model trained on both types.

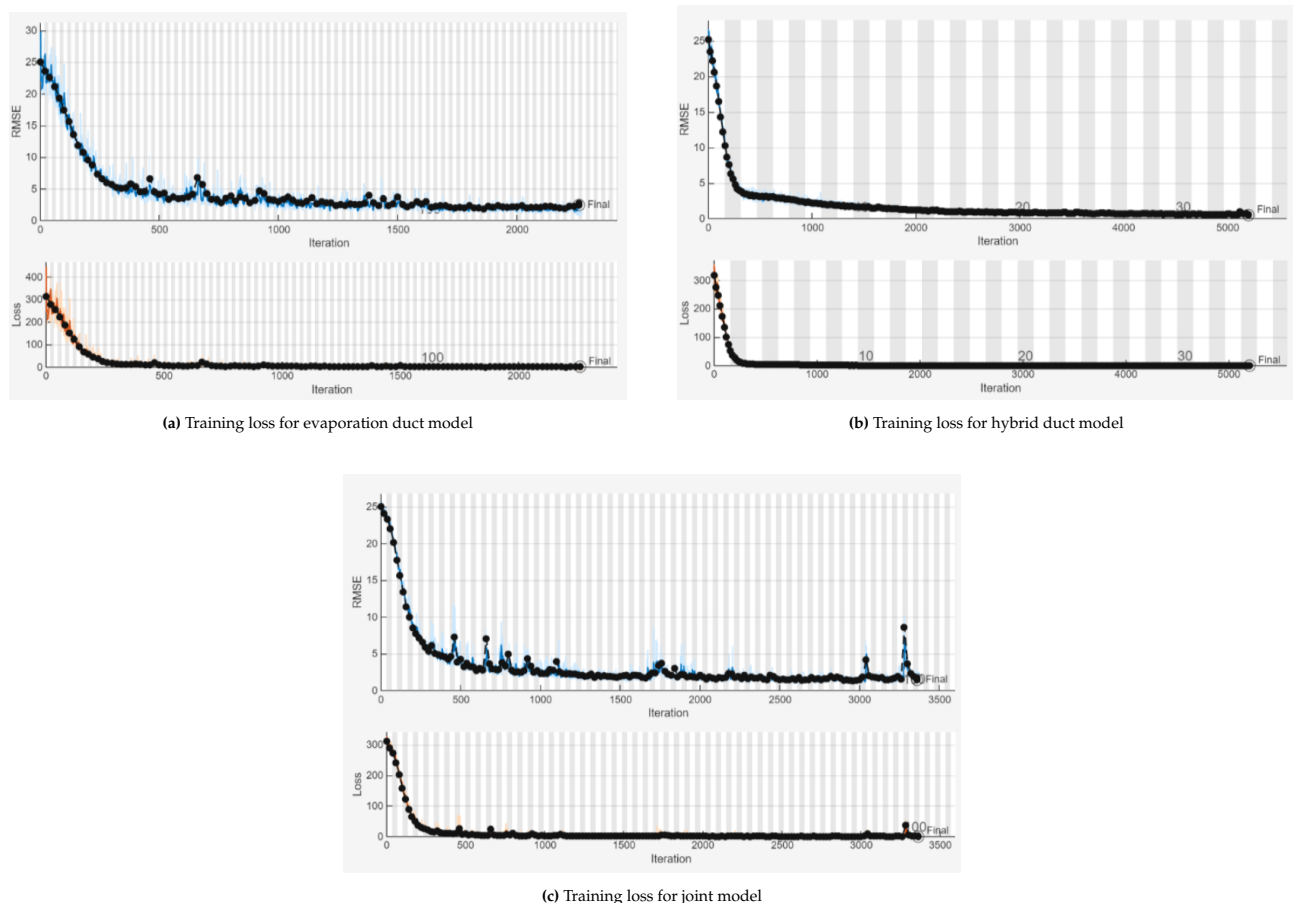


Figure 6: Training loss over time for the three models: (a) evaporation ducts, (b) hybrid ducts, and (c) joint model.

Figures (6a), (6b) and (6c) show the training progress of each model, illustrating how the loss (Mean Squared Error) and RMSE declined across epochs. The RMSE curve in the evaporation and hybrid ducts model show a smooth downward trend, indicating that the model's predictions became increasingly accurate with each epoch. Meanwhile, the training loss decreased steadily in the bottom

plot, confirming that the models were effectively minimizing the objective function. However, for the joint duct model, as seen in Figure (6c), the model shows some inconsistencies with a few peaks where the predictions diverge significantly from the actual values. These peaks indicate instances where the model struggled to adapt to the complex relationships between evaporation and hybrid ducts, resulting in inaccurate predictions for certain points.

4.2 Testing Performance

To assess model performance, four standard regression metrics were computed: MSE, MAE, RMSE, and R^2 score. The first of these was Mean Squared Error (MSE) which is calculated as shown in equation (14).

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (14)$$

where:
 n = number of data points
 y_i = actual value for i^{th} sample
 \hat{y}_i = predicted value for i^{th} sample

This shows the MSE is calculated as the quantified average squared difference between the predicted and actual refractivity values. The next metric calculated is Root Mean Squared Error (RMSE) shown in equation (15).

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (15)$$

where:
 n = number of data points
 y_i = actual value for i^{th} sample
 \hat{y}_i = predicted value for i^{th} sample

Equation (15) shows RMSE as being the square root of MSE, provides a measure of the average magnitude of prediction errors in the same units as the output variable. I also used the metric of Mean Absolute Error (MAE) to measure the performance of my model which is calculated as shown in equation (16).

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (16)$$

where:
 n = number of data points
 y_i = actual value for i^{th} sample
 \hat{y}_i = predicted value for i^{th} sample

Equation (16) shows the MAE evaluates the average absolute differences between predictions and actual values. Finally the last metric I used was the coefficient of determination (R^2); calculated as shown in equation (17).

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2} \quad (17)$$

where:
 y_i = actual value for i^{th} sample
 \hat{y}_i = predicted value for i^{th} sample
 \bar{y} = mean of the actual values

R^2 score assesses the proportion of variance in the actual output that is predictable from the inputs. A value close to 1 means the model is performing well, where a score of 0.5 means the model is performing as well as just predicting the mean of the values. A score of below 0.5 means the model is not performing very well as its predictions are worse than just predicting the mean. These metrics evaluate both the average magnitude of prediction errors and the extent to which the model explains data variance. Tables (1), (2), and (3) summarize results for the evaporation, hybrid, and joint models.

Table 1: Evaporation Only Model

Metric	Value
MSE	0.010654
RMSE	0.10322
MAE	0.0390
R ² Score	0.9894

Table 2: Hybrid Only Model

Metric	Value
MSE	0.00043891
RMSE	0.02095
MAE	0.0148
R ² Score	0.9996

Table 3: The Joint Duct Model

Metric	Value
MSE	0.0034915
RMSE	0.059089
MAE	0.0343
R ² Score	0.9965


This multi-metric approach allows for a clear comparison of model accuracy and highlights any trade-off in predictive performance when using a shared architecture across different duct types. The results in the table demonstrate that lower error values (MSE, RMSE, and MAE) indicate more accurate predictions, while a higher R² score reflects a stronger correlation between the predicted and actual values, signifying better overall model performance.

4.3 Predictions on Unseen Data


To validate generalization, each model was tested on unseen examples from their respective categories. The evaporation-only and hybrid-only models were tested on withheld cases from their specific duct types. The joint model was tested on a mixed set including both duct types.

4.3.1 Evaporation Only Model

Figures (7a) and (7b) show the correlation between the predicted values and the actual values against the height, measured in meters, for a specific evaporation duct using the evaporation duct only model. The model has never seen this specific duct before showing that it predicts the duct type very well.



(a) Predicted vs. actual refractivity for Evaporation Duct I (evaporation-only model)



(b) Predicted vs. actual refractivity for Evaporation Duct II (evaporation-only model)

Figure 7: Predicted and actual refractivity profiles for two evaporation ducts using the evaporation-only model.

4.3.2 Hybrid Only Model

Figures (8a) and (8b) show the correlation between the predicted and actual values against height, measured in meters, for a specific hybrid duct case using the hybrid-only model. This particular duct was not included in the model's training data, demonstrating that the hybrid-only model captures the duct specific cases very well.

4.3.3 Joint Duct Model

Figures (9) and (10) illustrate the correlation between the predicted and actual values as a function of height, measured in meters, for a specific evaporation duct scenario using the joint duct model. Subsequently, the model was evaluated on hybrid duct cases that were intentionally excluded from

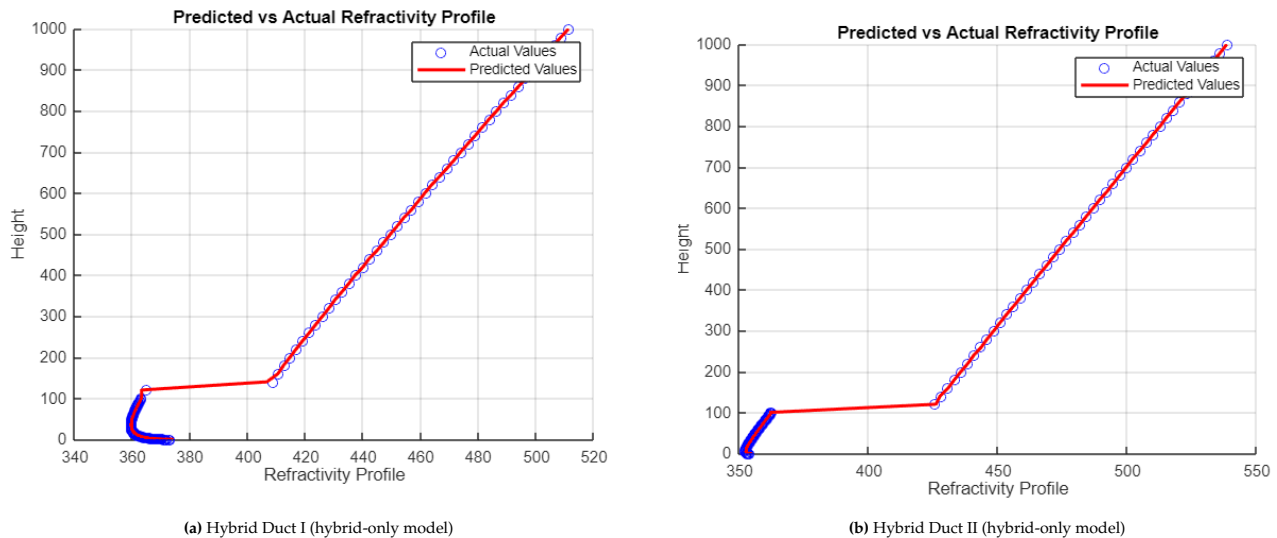


Figure 8: Predicted vs. actual refractivity profiles using the hybrid-only model.

the training phase. The results of this evaluation are presented in figures (11) and (12). These figures illustrate that the joint model was less effective at capturing the finer structural details of the ducts compared to the individual models.

4.4 Discussion

The results show that the separate neural network models for evaporation and hybrid ducts achieve near-perfect performance in predicting refractivity profiles within their respective domains. For both models, the low error metrics along with a high R^2 score underline their effectiveness in capturing the distinct structural features associated with each type of duct. These findings suggest that when trained independently, the models are able to exploit patterns that are unique to evaporation and hybrid ducts. This leads to highly accurate predictions and strong generalization on unseen data. However, when a joint model is used to predict both evaporation and hybrid ducts simultaneously, a decrease in performance occurs. While the joint model maintains similar overall error metrics, a closer examination of the prediction plots reveals subtle yet consistent distortions in evaporation duct predictions. These distortions are not present in the individual evaporation model and appear to stem from the joint model's shared internal representation. The model appears to favor the more structurally complex hybrid ducts, resulting in slight over-fitting to hybrid duct characteristics at the expense of accurately modeling the simpler, smoother evaporation ducts. This observation is further supported by the performance metrics in table (3), which show that although the joint model performs comparably in aggregate, its accuracy on evaporation ducts falls short of the dedicated evaporation model. The joint model appears to over-parametrize features more relevant to hybrid ducts, leading to localized deviations and a small but consistent drop in performance on simpler structures, a hallmark of over-fitting in mixed-domain learning scenarios. This imbalance suggests that the joint model, although theoretically appealing due to its simplicity and efficiency, may struggle to fully distinguish the unique signal characteristics of both duct types within a shared feature space. It reflects a known challenge in multi-task learning, where the learning of one task can inadvertently interfere with the learning of another. Particularly, when one is more dominant or variable in structure. Moreover, the pre-processing step of balancing the dataset before splitting was critical in enabling fair comparison and evaluation. Without this, the model would likely have developed an even stronger bias toward the more frequent duct type. Even with balanced input, the joint model's tendency to slightly skew toward hybrid ducts implies that architectural adjustments are needed. Such as introducing duct-specific sub networks or attention mechanisms might be needed to preserve class specific accuracy in future work. Another avenue worth exploring is the incorporation of metadata input features. For example surface conditions, time of day, or environmental indicators that could help the model better contextualize the input and adapt its representation accordingly. Cross-validation with real-world operational data or regional datasets could also provide a more rigorous test of generalization




Figure 9:

This figure compares the predicted and actual refractivity profiles for an evaporation duct generated by the joint model.




Figure 10:

This figure compares the predicted and actual refractivity profiles for another evaporation duct produced by the joint model.




Figure 11:

This figure compares the predicted and actual refractivity profiles for a hybrid duct generated by the joint model.




Figure 12:

This shows the comparison the predicted and actual refractivity profiles for another hybrid duct generated by the joint model.

and help assess the model's robustness across different atmospheric regimes. Ultimately, the findings suggest that while the joint model is a promising foundation, further refinement is required to achieve truly reliable, duct type specific performance.

5 Errors

Although overall model performance was strong, key error sources deserve attention, particularly when comparing the evaporation-only and joint models in specific duct cases. For an evaporation duct case, both models predict the refractivity profile with high accuracy, but minor deviations are evident, especially around sharper gradients. Figure (13a) shows the error distribution for the evaporation-only model, where small mismatches appear near surface layers. In contrast, Figure (13b) highlights the deviations of the predictions away from the line of perfect prediction.

A similar trend is observed in the hybrid duct case. As illustrated in Figures (14a) and (14b), sharp transitions in refractivity present a modeling challenge. The hybrid-only model closely aligns with

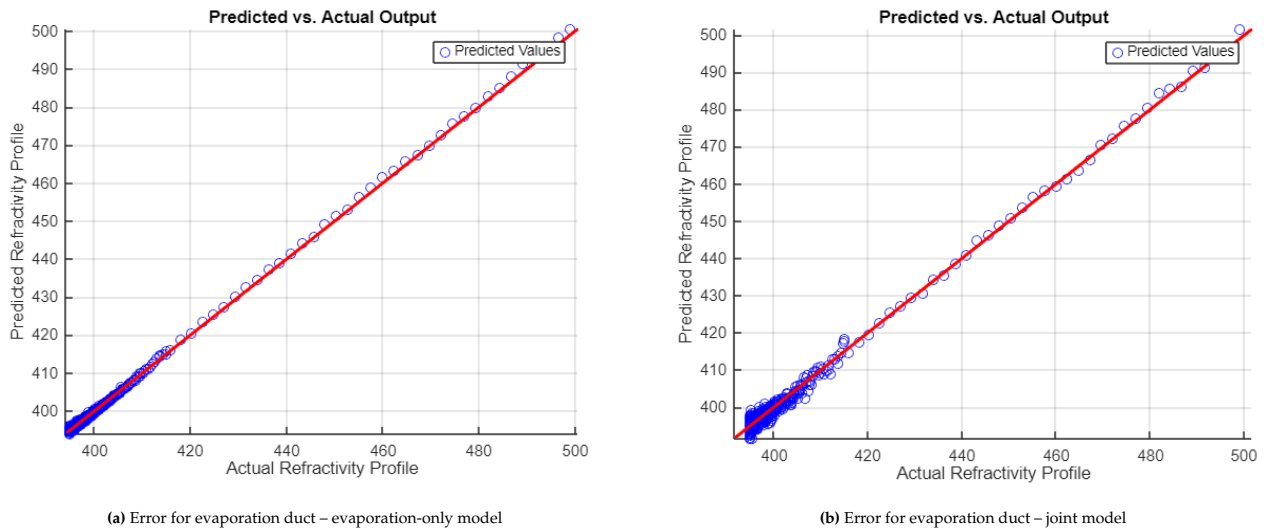


Figure 13: Comparison of prediction errors for an evaporation duct using (a) the evaporation-only model and (b) the joint model.

the actual profile, particularly capturing upper-layer gradients with high accuracy. In contrast, the joint model exhibits noticeable deviations near the lower boundary, suggesting a reduced ability to represent finer-scale features in that region.

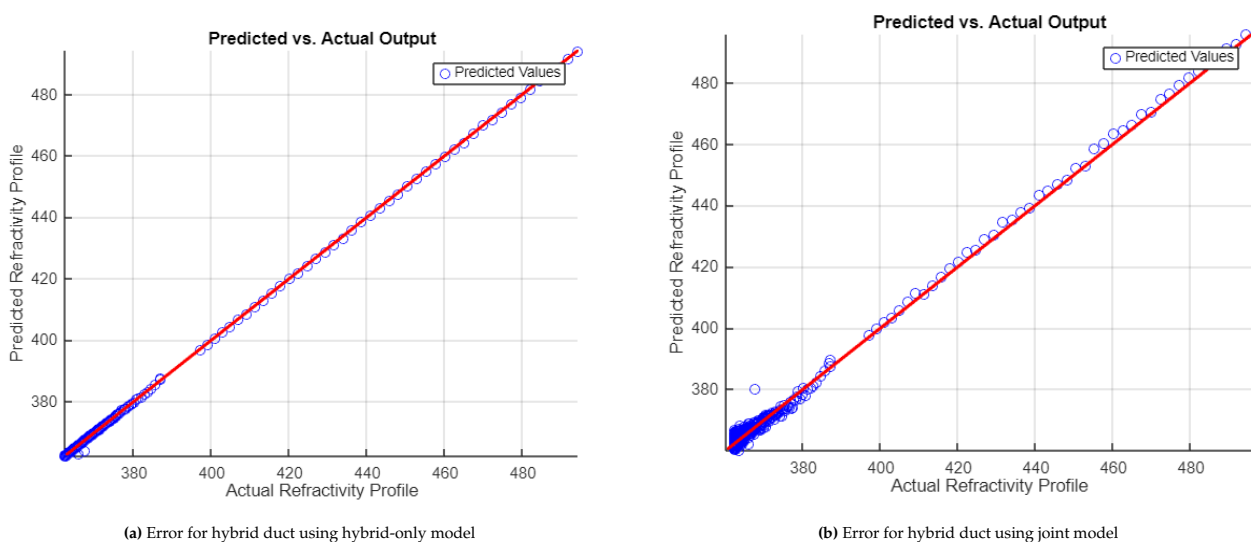


Figure 14: Error comparison for a specific hybrid duct: (a) using the hybrid-only model, and (b) using the joint model.

6 Conclusion

This study demonstrated the effectiveness of neural networks in predicting atmospheric refractivity profiles for evaporation and hybrid ducts. By comparing dedicated models with a joint model, the trade-off between specialization and generalization was explored. Individually trained models for each duct type showed excellent performance, with low errors and high R^2 scores, successfully capturing the unique structural features of each profile and generalizing well to unseen data. In contrast, the joint model, though competitive overall, showed slight degradation in accuracy—particularly for evaporation ducts. This reflects the typical challenge in multi-task learning, where shared representations can lead to interference between tasks. Balanced training data was critical in reducing bias toward the more complex and prevalent hybrid ducts. However, the joint model still leaned toward hybrid features, highlighting the need for improved modeling strategies. Future work should consider enhanced architectures, such as task-specific subnetworks or multi-branch structures, to better preserve the distinct characteristics of each duct type. Validating models with external or operational datasets will also be key to assessing their real-world reliability. This research provides a

solid foundation for further advancement in atmospheric modeling for maritime radar systems.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisors — Rebecca, Chloe, Bill, and Hooshyar — for their guidance, insight, and encouragement throughout this journey. Each of you challenged me to think deeper, aim higher, and grow not just as a student, but as a researcher. I am grateful for your time, patience, and unwavering support. To my parents and to Izzy, thank you for your constant support throughout this journey. Your encouragement has been a steady source of strength behind everything I've achieved. And finally to Scotty, whose belief in my abilities played a crucial role in restoring my confidence. I am sincerely thankful for his support. This work is a reflection not only of my efforts but of the incredible people who stood by me. I couldn't have done it without you.

References

- [1] Kirandeep Kaur and Bharat Gurnani, *Refraction of Light* (2023), <https://www.ncbi.nlm.nih.gov/books/NBK594230/>
- [2] Samudrapom Dam, *Understanding Snell's Law in Light Refraction* (2024), <https://www.azooptics.com/Article.aspx?ArticleID=2708>,
- [3] RadarTutorial, *Anomalous Propagation of Electromagnetic Waves* (n.d.), <https://www.radartutorial.eu/07.waves/wa17.en.html>
- [4] Kulesa, A.S., Barrios, A., Claverie, J., Garrett, S., Haack, T., Hacker, J.M., Hansen, H., Horgan, K., Hurtaud, Y., Lemon, C., Marshall, R., McGregor, J., McMillan, M., Periard, C., Pourret, V., Price, J., Rogers, L., Short, C., Veasey, M. and Wiss, V., 2016. The Tropical Air-Sea Propagation Study (TAPS). *Bulletin of the American Meteorological Society*. DOI: 10.1175/BAMS-D-14-00284.1.
- [5] Kieran Berryman, *What is Temperature Inversion?* (2024), <https://www2.purpleair.com/blogs/blog-home/what-is-temperature-inversion?>
- [6] Barrio-Zhang, H., Ruiz-Gutiérrez, É., Orejon, D., Wells, G. and Ledesma-Aguilar, R., 2024. Droplet motion driven by humidity gradients during evaporation and condensation. *The European Physical Journal E*, 47. DOI: <https://doi.org/10.1140/epje/s10189-024-00426-7>.
- [7] Chai, X., Li, J., Zhao, J., Hu, Y. and Zhao, X., 2022. NWPP-EDH: Numerical weather prediction products evaporation duct regional prediction model. *Frontiers in Earth Science*, 10, p.1031186. Available at: <https://www.frontiersin.org/articles/10.3389/feart.2022.1031186/full>. <https://doi.org/10.3389/feart.2022.1031186>.
- [8] Zhu, X., Yan, K., Jiang, L., Tian, L. and Tian, B., 2024. A method for elevated ducts refinement based on convolutional neural network. *Radio Science*, 59, p.1. <https://doi.org/10.1029/2023RS007789>. [Accessed 5 Jul. 2025].
- [9] Ronaghan, S., 2018. Deep Learning: Overview of Neurons and Activation Functions. [online] Available at: <https://srngn.medium.com/deep-learning-overview-of-neurons-and-activation-functions-1d98286cf1e4>.