

Predicting Vertical Radio Refractivity Gradients Across Temperate, Subarctic, and Polar Climates Using Ensemble Machine Learning

Solomon O. Adeola

`Solomon.adeola@bowen.edu.ng`

Bowen University

Francis O. Aweda

Bowen University

Jacob A. Akinpelu

Bowen University

Saeed A. Bello

University of Ilorin

Nathaniel O. Adeniji

Bowen University

Research Article

Keywords: Radio refractivity gradient, Tropospheric propagation, Ensemble machine learning, ERA5 reanalysis, Anomalous propagation

Posted Date: September 24th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10880805/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Predicting Vertical Radio Refractivity Gradients Across Temperate, Subarctic, and Polar Climates Using Ensemble Machine Learning

Solomon O. Adeola^{a,*}, Francis O. Aweda^a, Jacob A. Akinpelu^a, Saeed A. Bello^b, Nathaniel O. Adeniji^a

^a Department of Physics, College of Agriculture, Engineering, and Science, Bowen University, Iwo, Osun State, 232102, Nigeria.

^b Department of Physics, University of Ilorin, Ilorin, Kwara State, Nigeria.

* Corresponding author. E-mail address: Solomon.adeola@bowen.edu.ng

Abstract

Existing machine-learning studies of lower-tropospheric radio refractivity gradients have been predominantly limited to basic surface thermodynamic predictors and single-climate environments, leaving the influence of atmospheric dynamical and cloud microphysical variables across contrasting climatic zones largely unexamined. This research examined refractivity-gradient prediction across selected Temperate, Cold/Subarctic, and Polar regions to fill this gap. Long-term ERA5 reanalysis data (January 1991–April 2025) were collected from 14 stations representing these contrasting climates. Fourteen predictors, including thermodynamic, dynamical, and cloud microphysical variables, were used as inputs. Six machine learning models, LightGBM, Random Forest, XGBoost, SVR, MLPRegressor, and AdaBoost, were evaluated alongside a Ridge regression baseline with a 60:20:20 data split, using feature selection and hyperparameter tuning on the validation set. Ensemble tree-based models consistently outperformed the baseline. LightGBM achieved an R^2 of 0.877 and an RMSE of 1.700 N-units/km in the Temperate zone, an R^2 of 0.929 with an RMSE of 1.023 N-units/km in the Cold/Subarctic zone, and an R^2 of 0.814 with an RMSE of 0.395 N-units/km in the Polar zone. Across all zones, predictors such as specific humidity, potential vorticity, divergence, and geopotential were identified as most relevant. Applying ensemble machine learning to ERA5 reanalysis data effectively captures refractivity-gradient variability across different climates. These results are based on a random data split and have not been validated against independent radiosonde profiles or chronological test sets. Future studies should include leave one station out, chronological validation and independent radiosonde measurements.

Keywords: Radio refractivity gradient, Tropospheric propagation, Ensemble machine learning, ERA5 reanalysis, Anomalous propagation

1. INTRODUCTION

Radio wave propagation is mostly determined by the properties of the lower atmosphere, particularly the troposphere, the atmospheric layer nearest to human habitation [1, 2]. Electromagnetic waves, as they traverse this medium, undergo continuous scattering, absorption, reflection, and refraction due to spatial and temporal fluctuations in atmospheric characteristics. The extent of this bending is measured by radio refractivity, which is significantly influenced by absolute air temperature, water vapour pressure, and atmospheric pressure [3]. In normal atmospheric conditions, these thermodynamic factors decrease with altitude, causing radio signals to bend slightly downward toward the Earth's surface [4]. Nonetheless, when the vertical refractivity gradient diverges markedly, severe anomalous propagation conditions such as sub-refraction, super-refraction, and air ducting arise. These anomalies function as dielectric waveguides, resulting in significant signal attenuation, multipath interference, radar clutter, and communication disruptions [5].

Propagation models based on standard atmospheric assumptions, including the conventional 4/3 k-factor, have been developed and validated primarily in temperate environments, where conditions such as those studied by Norin [6] are representative.

In temperate mid-latitude regions, the refractivity profile is highly variable and significantly affected by seasonal weather fronts, summer sea-breeze circulations, and nighttime radiative cooling. As a result, maritime and continental temperate regions often encounter heightened super-refractive layers and significant atmospheric evaporation ducting, especially during summer periods [7].

As communication systems and radar networks extend into higher latitudes, it becomes increasingly evident that the physical determinants of radio refractivity undergo substantial changes across climatic environments. In high-latitude regions, the troposphere is markedly shallower, and anomalous propagation is influenced more strongly by severe cold conditions and temperature inversions than by the intense moisture gradients that dominate tropical environments. In extreme polar environments, characterised by severely cold winters and minimal precipitation, stable standard refraction predominates for much of the year, with vertical refractivity gradients generally fluctuating between -79 and -59 N-units/km [5]. In contrast, in Subarctic marine and continental regions, intense radiative cooling during prolonged cold nights frequently induces strong temperature inversions. These temperature-driven atmospheric dynamics contribute to a high occurrence of sub-refractive anomalies, which have historically been reported to occur during up to 80% of summer nights in some Subarctic environments [8]. Despite these substantial propagation challenges, many conventional operational systems continue to rely on simplified atmospheric assumptions developed primarily from temperate observations, thereby contributing to range-height inaccuracies and coverage deficiencies in climatically contrasting regions.

These strong climatic contrasts demonstrate that refractivity-gradient behaviour cannot be adequately represented using uniform atmospheric assumptions across all propagation environments. Conventional empirical propagation frameworks, including standard effective-earth-radius approximations, were largely developed using observations from limited climatic settings and therefore often struggle to represent the nonlinear variability observed across dynamically different atmospheric regimes. As communication and

radar systems increasingly operate across diverse geographical and climatic conditions, there is growing need for predictive approaches capable of adapting to complex thermodynamic, dynamical, and microphysical atmospheric interactions [5-7, 9].

To mitigate the limitations of traditional empirical models, researchers have increasingly adopted soft-computing and machine learning approaches for modelling complex atmospheric datasets. Traditional predictive tasks in radio meteorology have often relied on Simple Linear Regression (SLR) or Multiple Linear Regression (MLR); however, these models are constrained by assumptions of linearity, which may reduce predictive accuracy when modelling highly nonlinear atmospheric interactions [2]. In moisture-laden or dynamically variable environments, linear models often fail to adequately represent the coupled interactions between thermodynamic and atmospheric dynamical processes. For example, in the tropical monsoon climate of Calabar, Nigeria, simple linear regression revealed that relative humidity had a stronger linear relationship with refractivity, with an R^2 value of 0.62, compared to ambient temperature, which had an R^2 value of 0.19, or atmospheric pressure, which had an R^2 value of 0.02 [10]. Similar limitations have been reported in coastal temperate environments. Ranran et al. [11] demonstrated that multiple linear regression produced a relatively high Root Mean Square Error (RMSE) of 57.2597 when forecasting refractivity gradients within the lowest 65 m over Qingdao, China. However, the same study also reported that linear regression performed comparatively better in the moisture-limited plateau environment of Lhasa, where weaker humidity variability reduced atmospheric complexity [11].

To address these limitations, recent studies have increasingly explored machine learning architectures capable of modelling nonlinear atmospheric relationships. Unlike conventional statistical approaches, machine learning methods can capture complex interactions involving moisture variability, atmospheric dynamics, cloud microphysics, and thermal structure. Consequently, neural-network, decision-tree, and ensemble-learning approaches have increasingly been applied to refractivity and atmospheric-propagation studies across different climatic environments.

A Deep Neural Network utilising Long Short-Term Memory (LSTM) was developed to forecast time-series air refractivity in temperate Asia using 11,942 meteorological observations from Osan, South Korea [9]. The model achieved a Mean Absolute Percentage Error (MAPE) of 2.75%, outperforming the conventional ARIMA model, which produced an error rate of 7.61% [9]. In the marine temperate environment of the French North Atlantic Coast, Claverie and Motsch [7] applied both a Decision Tree Regressor (DT-20) and a Neural Network (MLP-500) to estimate vertical refractivity profiles associated with evaporation ducts. Their results showed that the Decision Tree model achieved stronger predictive performance, while the Neural Network occasionally generated physically unrealistic negative duct heights. Similarly, Jicha et al. [12] applied an Artificial Neural Network with a 20-10-1 architecture in the continental temperate climate of the Czech Republic and reported that models trained solely on standard surface meteorological variables achieved only moderate predictive success, indicating that conventional thermodynamic variables alone may be insufficient for fully representing propagation dynamics.

Despite the increasing application of machine learning in radio meteorology, several important gaps remain [13, 14]. Most existing studies restrict their predictor inputs to basic surface thermodynamic variables while neglecting atmospheric dynamical and cloud-microphysical processes [1-2, 15]. Furthermore, advanced architectures such as ANN, LSTM, GRU, and RVM have been applied primarily in single-climate environments without cross-climate comparative evaluation [7, 9, 11-12, 16]. Comparatively few studies have examined explainable ensemble machine learning for refractivity-gradient prediction in high-latitude environments, where atmospheric thermodynamic structure differs substantially from tropical and temperate conditions.

Consequently, this study investigates ensemble and nonlinear machine-learning approaches for predicting vertical radio refractivity gradients across selected Temperate, Cold/Subarctic, and Polar environments, using long-term ERA5 reanalysis data and 14 atmospheric predictors drawn from thermodynamic, dynamical, and cloud-microphysical processes.

2. METHODOLOGY

2.1 Study Area

The study comprised 14 meteorological stations grouped into three climatically distinct environments: the Temperate zone, the Cold/Subarctic zone, and the Polar zone, as shown in Figure 1.




Figure 1: Map showing the study area: Temperate, Cold/Subarctic and Polar zones.

The climate-zone classification adopted in this study was based on a combination of Köppen–Geiger climatic characteristics [17], geographical latitude, station elevation, and the dominant lower-tropospheric atmospheric behaviour relevant to radio propagation. Temperate stations such as Bordeaux, Copenhagen, New York, and Denver were assigned to the Temperate class on the basis of their strong seasonal thermodynamic variability and mid-latitude atmospheric dynamics [7]. High-latitude stations such as Anchorage, Helsinki, Kebnekaise, and Galdhøpiggen were assigned to the Cold/Subarctic class because of their persistent cold lower-tropospheric structure, enhanced atmospheric stability, and inversion-prone conditions associated with subarctic propagation environments [5, 8]. Courmayeur was also assigned to the Cold/Subarctic class, as this station is situated at high elevation, where persistently cold, moisture-limited, and stable lower-tropospheric conditions produce refractivity-gradient characteristics consistent with subarctic propagation behaviour, despite its classical low-altitude Köppen–Geiger classification [18]. The Polar class comprised Vostok and McMurdo, two Antarctic stations characterised by highly stable atmospheric conditions and negligible atmospheric moisture typical of polar propagation environments [5].

2.1.1 The Temperate Climate Zone

The Temperate dataset comprises 6 stations located across North America, Europe, and Asia. In these mid-latitude environments, the tropospheric refractivity profile is highly variable and primarily governed by seasonal thermodynamic cycles, including warm-season moisture loading, temperature inversions associated with nocturnal radiative cooling, and the passage of synoptic weather systems[7]. These atmospheric conditions are associated with enhanced super-refractive layers and anomalous propagation events, particularly during summer months when moisture gradients in the lower troposphere are strongest [7]. The geographical coordinates of the selected Temperate stations are presented in Table 1.

Table 1. Temperate Climate Stations

Station	Country	Latitude (°N)	Longitude (°E)
Bordeaux	France	44.84	-0.58
Copenhagen	Denmark	55.68	12.57
New York	USA	40.71	-74.01
Atlanta	USA	33.75	-84.39
Denver	USA	39.74	-104.99
Lhasa	China	29.65	91.10

2.1.2 The Cold/Subarctic Climate Zone

The Cold/Subarctic dataset comprises 6 stations situated in high-latitude regions across Norway, Sweden, Finland, Alaska, and Mongolia, as well as high-altitude environments in Italy and China. Radio wave propagation in these environments is strongly influenced by intense radiative cooling during prolonged cold periods, which induces deep temperature inversions and suppresses vertical moisture gradients. These temperature-driven atmospheric dynamics are associated with a high prevalence of anomalous propagation, particularly sub-refractive conditions, which have been observed on up to 80% of summer nights in some subarctic environments [8]. The geographical coordinates of the selected Cold/Subarctic stations are presented in Table 2.

Table 2. Cold/Subarctic Climate Stations

Station	Country	Latitude (°N)	Longitude (°E)
Galdhøpiggen	Norway	61.63	8.31
Kebnekaise	Sweden	67.90	18.50
Helsinki	Finland	60.17	24.94
Anchorage	Alaska	61.22	-149.90
Ulaanbaatar	Mongolia	47.92	106.92
Courmayeur	Italy	45.79	6.97

2.1.3 The Polar Climate Zone

The Polar dataset comprises 2 stations located in Antarctica. The troposphere in these environments is significantly shallower than at mid-latitudes and is characterised by severely cold winters and negligible precipitation [5]. The very low absolute moisture content of the polar atmosphere means that highly stable standard refraction is the prevailing condition for the majority of the year, with vertical refractivity gradients typically ranging between -79 and -59 N-units/km [5]. Anomalous propagation events are comparatively infrequent in this zone relative to the Cold/Subarctic and Temperate zones [5]. The geographical coordinates of the selected Polar stations are presented in Table 3. The Polar climatic group includes two Antarctic stations (Vostok and McMurdo).

Table 3. Polar Climate Stations

Station	Country	Latitude (°N)	Longitude (°E)
Vostok	Antarctica	-78.47	106.83
McMurdo	Antarctica	-77.85	166.67

2.2 Data Source and Study Period

Atmospheric data were obtained from the ERA5 reanalysis dataset produced by the European Centre for Medium-Range Weather Forecasts (ECMWF), which provides high-resolution, globally consistent meteorological fields at multiple pressure levels [19]. The dataset spanned January 1991 to April 2025, providing a 34-year record sufficient for characterising both seasonal and interannual refractivity-gradient variability.

Monthly averaged pressure-level data were extracted for 14 stations representing the Temperate, Cold/Subarctic, and Polar climate zones. The analysis focused on the lower troposphere, specifically the 1000–850 hPa pressure layer, which encompasses the atmospheric depth most relevant to surface-based radio wave propagation and refractivity-gradient assessment [4-5]. Station data were extracted using a nearest-grid-point approach, in which the ERA5 grid cell with the minimum squared distance from each station's coordinates was identified and all predictor variables were extracted from that single grid cell at each pressure level and time step. This approach was applied consistently across all stations and time steps throughout the study period.

2.3 Computation of Refractivity Gradient

Atmospheric radio refractivity (N) was computed using the Smith–Weintraub refractivity formulation [1-2, 20]:

$$N = 77.6 \left(\frac{P}{T} \right) + 3.73 \times 10^5 \left(\frac{e}{T^2} \right) \quad (1)$$

where N is radio refractivity in N-units, P is atmospheric pressure in hPa, T is temperature in Kelvin, and e is water vapour pressure in hPa.

Water vapour pressure e was derived from ERA5 relative humidity values using [1-2, 20]:

$$e = \frac{RH}{100} \times e_s \quad (2)$$

where RH is relative humidity in percent and e_s is saturation vapour pressure.

The saturation vapour pressure was calculated using the Tetens-based formulation commonly applied in atmospheric and refractivity studies [2, 21]:

$$e_s = 6.112 \exp \left[\frac{17.67(T - 273.15)}{T - 29.65} \right] \quad (3)$$

where T is temperature in Kelvin and e_s is in hPa.

The vertical refractivity gradient was subsequently calculated between the 1000 hPa and 850 hPa pressure levels using [4-5]:

$$\frac{dN}{dz} = \frac{N_{850} - N_{1000}}{Z_{850} - Z_{1000}} \quad (4)$$

where N_{850} and N_{1000} are refractivity values at 850 hPa and 1000 hPa, while Z_{850} and Z_{1000} are geopotential heights in km at the same pressure levels.

The ERA5 monthly averaged pressure-level variables used for refractivity-gradient computation included temperature and relative humidity at multiple pressure levels. Pressure-level altitude values were approximated using standard atmospheric pressure-height relationships, where each ERA5 pressure level was assigned a representative geopotential height based on standard atmospheric conditions [20]. For example, the 1000 hPa and 850 hPa levels were assigned approximate heights of 0.11 km and 1.46 km respectively.

For elevated stations such as Lhasa and Courmayeur, the 1000 hPa pressure level may intersect or fall below the local terrain surface because ERA5 atmospheric variables are provided on fixed pressure levels [19]. Consequently, refractivity-gradient results from these stations were interpreted as representative large-scale lower-tropospheric estimates rather than exact near-surface atmospheric conditions.

2.4 Atmospheric Predictors

Fourteen atmospheric predictors were extracted from ERA5 pressure-level data to represent the thermodynamic, dynamical, and cloud microphysical processes governing lower-tropospheric refractivity variability. Most existing refractivity studies restrict their predictor inputs to surface meteorological variables (temperature, pressure, and relative humidity [1]) without incorporating broader atmospheric dynamical or microphysical variables. The present study extended this predictor space to include variables representing all three process categories, as described below. The use of ERA5 for regional refractivity-gradient investigations across different climatic environments has been demonstrated in recent studies [13, 22].

2.4.1 Thermodynamic and Chemical Variables

(1) Specific humidity ($\text{kg}\cdot\text{kg}^{-1}$): The ratio of the mass of water vapour to the total mass of an air parcel [21]. Specific humidity is a fundamental meteorological variable used to compute the water vapour pressure (e) and the wet component of atmospheric radio refractivity [1, 23].

(2) Ozone mass mixing ratio ($\text{kg}\cdot\text{kg}^{-1}$): The ratio of the mass of ozone gas to the total mass of dry air in a given atmospheric parcel [21]. Ozone mass mixing ratio was included as a candidate predictor because it reflects large-scale atmospheric transport, vertical mixing, and dynamical variability within the troposphere and lower stratosphere [24-25]. Although ozone does not enter the Smith–Weintraub refractivity equation directly, variations in its distribution are associated with changes in atmospheric stability, temperature structure, and circulation patterns that may indirectly relate to refractivity-gradient variability. Ozone was therefore subjected to feature selection without any prior assumption of predictive relevance.

2.4.2 Atmospheric Dynamical Variables

These variables govern atmospheric motion, turbulence, and wind structure, all of which influence refractivity variability in dynamically active environments [22].

(1) Potential vorticity ($\text{K}\cdot\text{m}^2\cdot\text{kg}^{-1}\cdot\text{s}^{-1}$): A conserved dynamic and thermodynamic quantity that combines the rotational component of the wind field with the vertical temperature stratification of the atmosphere, widely used to track air mass movement and tropopause dynamics [26].

(2) Divergence (s^{-1}): The rate at which an air mass spreads horizontally, driving vertical motion through mass continuity [26].

(3) Geopotential ($\text{m}^2\cdot\text{s}^{-2}$): The potential energy per unit mass at a given point in the atmosphere relative to the Earth's surface [26]. Geopotential height at specific pressure levels has been used to characterise spatial variations in refractivity and atmospheric ducting [23].

(4) Vertical velocity ($\text{Pa}\cdot\text{s}^{-1}$): The rate of change of pressure with time, representing the speed and direction of vertical air motion in the atmosphere [26].

(5) U-component of wind ($\text{m}\cdot\text{s}^{-1}$): The zonal (east–west) component of the horizontal wind vector [26]. Horizontal wind speed and direction influence boundary-layer moisture transport and ducting variability [7, 15].

(6) V-component of wind ($\text{m}\cdot\text{s}^{-1}$): The meridional (north–south) component of the horizontal wind vector [26].

(7) Vorticity (relative) (s^{-1}): A measure of the local rotation of the atmospheric wind field relative to the Earth's surface [26].

2.4.3 Cloud Microphysical Variables

(1) Fraction of cloud cover (dimensionless, 0–1): The proportion of a given model grid box covered by cloud [21].

(2) Specific cloud liquid water content ($\text{kg}\cdot\text{kg}^{-1}$): The mass of suspended liquid water droplets per unit mass of air [21]. Liquid water content contributes to atmospheric refractivity through cloud absorption and scattering effects, referred to as cloud refractivity [23].

(3) Specific rain water content ($\text{kg}\cdot\text{kg}^{-1}$): The mass of precipitating liquid water per unit mass of air [21]. Elevated precipitation rates are associated with signal attenuation and variations in the atmospheric refractive index [1, 15].

(4) Specific snow water content ($\text{kg}\cdot\text{kg}^{-1}$): The mass of precipitating frozen water per unit mass of air [21].

(5) Specific cloud ice water content ($\text{kg}\cdot\text{kg}^{-1}$): The mass of suspended ice crystals per unit mass of air [21]. Ice water content contributes to cloud refractivity in a manner analogous to liquid water content and is relevant to the modelling of propagation depth and signal fading in cold-cloud environments [23].

2.5 Data Preprocessing

Data preprocessing was conducted to ensure consistency and quality across all stations and predictor variables. Stations were standardised to a consistent naming convention to ensure accurate climate-zone mapping, and all predictor variables were converted to numeric format. Records containing non-numeric values were identified as missing values during this conversion. A missing-value check was subsequently conducted across all stations and variables; no missing values were detected across the final set of modelling variables, and no imputation was required.

A physical validity check was performed for variables with bounded or non-negative physical ranges. Fraction of cloud cover was verified to lie within the 0–1 range, and specific humidity and ozone mass mixing ratio were checked for the absence of negative values. No physically invalid values were detected.

Outliers were assessed using the interquartile range (IQR) method for all 14 predictor variables. Statistical outliers were identified but were not removed, as many represented physically plausible atmospheric extremes, particularly for variables such as potential vorticity, divergence, vertical velocity, wind components, and cloud microphysical variables. The outlier analysis was therefore used as a diagnostic quality-control step rather than as a basis for data deletion.

Predictor variables were standardised using zero-mean unit-variance scaling within the modelling pipeline. Standardisation was applied consistently across all models and was particularly important for scale-sensitive approaches such as Ridge Regression, Support Vector Regression (SVR), and MLPRegressor. Although tree-based models do not require standardisation, a consistent preprocessing pipeline was applied across all models to ensure comparability.

The dataset was partitioned into training, validation, and test subsets using a 60:20:20 random split, as summarised in Table 4. Twenty percent of the data was reserved as an independent test set prior to any model development. The remaining 80% was divided into training (60%) and validation (20%) subsets. Feature selection and hyperparameter tuning were performed exclusively within the training-validation framework, while the independent test set remained unseen until final evaluation. A fixed random seed was used throughout to ensure reproducibility. It is acknowledged that random rather than chronological splitting may introduce temporal leakage, as monthly records from different years may appear across subsets; a chronological split would provide a more stringent test of temporal generalisation.

Table 4. Sample distribution across climate zones and data subsets. Each sample represents a station–month observation partitioned into training (60%), validation (20%), and test (20%) subsets.

Climate Zone	Stations	Total Samples	Training (60%)	Validation (20%)	Test (20%)
Temperate	6	2,472	1,484	494	494
Cold/Subarctic	6	2,472	1,484	494	494
Polar	2	824	494	165	165

2.6 Feature Selection

Feature selection was performed separately for each climate zone and machine-learning model to identify the most relevant subset of atmospheric predictors for refractivity-gradient prediction. The initial candidate pool comprised all 14 atmospheric variables described in Section 2.3.

For tree-based models evaluated within a cross-validation framework, Sequential Forward Selection (SFS) was implemented within the training dataset using 3-fold cross-validation. Feature subsets ranging from 2 to $n-1$ variables were evaluated iteratively, and the subset producing the lowest mean cross-validation RMSE was retained as the optimal combination.

For non-tree models evaluated using a validation-based framework, features were added progressively one at a time, with the feature producing the greatest reduction in validation RMSE retained at each step. The optimal subset was defined as the combination yielding the minimum validation RMSE across all tested subset sizes.

Feature selection and subsequent hyperparameter tuning were conducted exclusively within the training-validation framework, ensuring that the independent test set remained completely unseen until final performance evaluation. This ensured that reported test RMSE, MAE, and R^2 values reflect held-out test performance within the same station-month distribution; it is acknowledged, however, that the random split does not guarantee strict temporal generalisation, and these results should be interpreted accordingly. A top-subset feature frequency analysis was also conducted by counting how often each predictor appeared among the five best-performing feature subsets ranked by validation RMSE. This analysis identified the most consistently selected atmospheric variables across climate zones and models.

2.7 Machine Learning Models

Six regression models were implemented: Light Gradient Boosting Machine (LightGBM), Random Forest Regressor, Extreme Gradient Boosting (XGBoost), Support Vector Regression (SVR), Multi-Layer Perceptron Regressor (MLPRegressor), and Adaptive Boosting (AdaBoost). Ridge regression was additionally included as a regularised linear baseline to provide a benchmark for comparison with the nonlinear approaches. These models were selected to represent a range of learning paradigms, including ensemble tree-based methods, kernel-based regression, neural networks, and linear regression.

2.8 Hyperparameter Tuning

Model hyperparameters were optimised using grid search based on validation RMSE, ensuring that tuning decisions remained independent of the test dataset. The searched hyperparameter ranges and the optimal configurations identified for each model and climate zone are reported in the results tables (Tables 6, 8, and 10).

2.9 Model Evaluation Metrics

Model performance was evaluated using Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Coefficient of Determination (R^2). These metrics were computed on the independent 20% hold-out test dataset, which remained unseen during all stages of model development and selection, ensuring that the reported performance values reflect held-out test accuracy within the same station-month distribution rather than strict out-of-sample temporal generalisation.

Bootstrap confidence intervals were estimated for the independent test-set performance metrics of each regression model across all three climate zones. For each model, 1,000 bootstrap samples were generated by randomly resampling the test observations with replacement, and RMSE and R^2 were computed for each resample. The 2.5th and 97.5th percentiles of the resulting distributions were used to construct 95% confidence intervals, providing an indication of the stability of the reported performance metrics.

2.10 Model Interpretation

Feature importance was assessed using SHAP (SHAPley Additive Explanations) for tree-based models (LightGBM, Random Forest, XGBoost) and permutation importance for non-tree models (SVR, MLPRegressor, AdaBoost). SHAP values quantify the average marginal contribution of each feature to individual predictions, whereas permutation importance measures the increase in prediction error when a feature's values are randomly shuffled. These two metrics operate on fundamentally different scales and are not numerically comparable. Feature-importance results were therefore interpreted separately within each model group, and cross-model comparisons were restricted to the qualitative direction of influence specifically, which variables ranked highest within each method rather than to the magnitude of the reported scores. This approach enabled the identification of atmospheric variables most consistently associated with refractivity-gradient variability across climate zones.

3. RESULTS AND DISCUSSION




Figure 2. Monthly variation of meteorological variables (dashed, left axis) and the 1000–850 hPa vertical refractivity gradient (solid, right axis) across the Temperate, Sub-Arctic, and Polar climate zones: (a) temperature, (b) relative humidity, and (c) water vapour pressure. The dotted grey line marks the ITU-R standard gradient of -40 N/km separating standard from super-refractive conditions. Meteorological variables are shown at 850 hPa and the refractivity gradient is computed between the 1000 and 850 hPa levels, based on ERA5 monthly means (January 1991–April 2025).

Figure 2 presents the mean monthly cycles of the principal meteorological drivers together with the resulting 1000–850 hPa vertical refractivity gradient for the three climatic zones examined in this study. In each panel, the dashed curves show a meteorological variable on the left axis and the solid curves show the refractivity gradient on the right axis, so that the coupling between the atmospheric state and gradient behaviour can be read directly. The dotted grey line denotes the ITU-R reference gradient of -40 N/km, which separates standard refraction from the super-refractive regime.

Temperature (Figure 2a). All three zones followed a clear seasonal temperature cycle whose amplitude decreased with latitude. In the Temperate zone the 850 hPa temperature ranged from about 276.2 K in January to 287.2 K in July, and in the Sub-Arctic zone from about 267.2 K in January to 284.7 K in July. The Polar zone was both the coldest and the most weakly forced, with a reversed (Southern-Hemisphere) cycle peaking near 259.2 K in January and falling to about 242.7 K in August. The refractivity gradient responded inversely to this thermal forcing: as temperature rose toward the summer peak the gradient in the Temperate and Sub-Arctic zones steepened, reaching its most negative monthly values of -46.7 and -43.3 N/km in August before relaxing to its least negative values of -38.4 N/km (March) and -34.5 N/km (February), respectively. The gradient minimum therefore lagged the July temperature and moisture peaks by about one month. The Temperate gradient fell below the ITU-R -40 N/km super-refraction threshold in seven of the twelve months and the Sub-Arctic gradient in four, whereas the Polar gradient remained almost invariant (mean -31.3 N/km, seasonal range only 1.7 N/km) and never crossed the -40 N/km line, its slight steepening to -32.5 N/km occurring in January during the austral summer.

Relative humidity (Figure 2b). Relative humidity varied over only a modest range in each zone (approximately 59–68% in the Temperate zone, 64–76% in the Sub-Arctic zone, and 75–83% in the Polar zone) and did not co-vary consistently with the refractivity gradient across zones. Most tellingly, the Polar zone maintained the highest relative humidity of the three zones yet exhibited the least negative and least variable gradient. This dissociation confirms that relative humidity alone is a poor indicator of gradient behaviour, because it expresses moisture relative to the saturation capacity of the air rather than the absolute water vapour content that governs the wet term of refractivity.

Water vapour pressure (Figure 2c). The clearest and most consistent coupling was between water vapour pressure and the refractivity gradient. Absolute moisture rose sharply into July, peaking near 11.4 hPa in the Temperate zone and 10.4 hPa in the Sub-Arctic zone, and the gradient reached its most negative value about a month later, in August, indicating that the accumulation of lower-tropospheric moisture drives the seasonal steepening. In the Polar zone water vapour pressure remained very low throughout the year (mean 0.73 hPa, maximum only 1.6 hPa in January), consistent with the small, nearly constant Polar gradient. This correspondence identifies absolute moisture content as the dominant control on the seasonal cycle of the vertical refractivity gradient.

Overall, Figure 2 shows that the seasonal steepening of the refractivity gradient, and the associated excursion into super-refractive conditions in the Temperate (seven months) and Sub-Arctic (four months) zones, is driven primarily by the summer increase in absolute water vapour pressure, reinforced by rising temperature, with the gradient minimum lagging the driver peak by roughly one month. Relative humidity is only weakly related to the gradient. The Polar zone behaves as a limiting case in which persistently cold, dry air suppresses both the magnitude and the seasonal variability of the gradient, holding it near standard refraction throughout the year. These physically consistent relationships are reflected in the refractivity-gradient behaviour modelled in the following sections and in the predictor-importance results, where absolute-moisture and thermodynamic variables carry the greatest predictive weight.

Figure 3: Monthly predicted vs observed refractivity gradient (N/km) for the Temperate climate zone across models, based on monthly averages from ERA5 data (January 1991–April 2025).

3.1. Temperate





Figure 4: Predicted vs observed refractivity gradient (N/km) scatter plots for the Temperate climate zone across models, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).



Predicted vs Actual - Temperate - Ridge




Figure 5: Monthly predicted versus observed refractivity gradient (N/km) and predicted versus actual scatter plot for Ridge regression in the Temperate climate zone, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

The Temperate climate zone exhibited the largest seasonal variability in refractivity gradients among the three investigated climatic environments, with monthly gradients ranging between approximately -38 and -46 N-units/km (Figure 3). The physical basis for this seasonal range was evident in Figure 2. Although relative humidity showed a modest inverse seasonal pattern, water vapour pressure underwent a more dramatic increase from approximately 7–8 hPa in winter to 20–21 hPa in summer (Figure 2c), consistent with absolute moisture content being the primary source of seasonal refractivity variability rather than relative humidity. This pattern was consistent with the Clausius–Clapeyron relationship, whereby the warmer summer atmosphere held approximately two to three times more absolute water vapour than the winter atmosphere [27], amplifying the wet term of the Smith–Weintraub equation (Equation 1). These findings were consistent with Gao et al. [28], who reported that refractivity gradients are more sensitive to moisture gradients than to temperature gradients during the warm season.

All models have their minimum in July, the same month as the actual refractivity gradient values. The ensemble tree-based models, particularly LightGBM, Random Forest, and XGBoost, reproduced the seasonal evolution of refractivity gradients with strong agreement (Figure 3), consistent with the known capacity of ensemble learners to represent nonlinear predictor interactions [29–31]. AdaBoost showed the largest deviations, consistent with its reliance on shallow weak learners less suited to high-dimensional nonlinear predictor spaces [30]. Scatter plots (Figure 4) confirmed strong agreement between predicted and observed values across the gradient range of approximately -55 to -30 N-units/km, with LightGBM, XGBoost, and Random Forest exhibiting the tightest clustering along the 1:1 reference line.

Figure 5 shows the performance of Ridge regression in the Temperate zone. As shown in Figure 5(a), Ridge reproduced the broad seasonal direction but produced systematically compressed predictions, underestimating the most negative gradients during summer and overestimating the least negative gradients during winter. The scatter plot in Figure 5(b) further illustrated this compression, with predicted values clustering toward the centre of the observed range rather than tracking the full extent of gradient variability. These results were consistent with the lower performance metrics reported for Ridge regression, with an RMSE of 2.852 N-units/km and R^2 of 0.654 (Tables 5a and 5b), indicating that linear modelling was less suited to representing the full range of refractivity-gradient variability in the Temperate environment.

Top-Subset Feature Frequency by Model - Temperate




Figure 6: Top-subset feature frequency across regression models in the Temperate climate zone, showing the frequency of features appearing in the five best-performing subsets based on validation RMSE using ERA5 data (January 1991–April 2025).




Figure 7: Feature importance across regression models in the Temperate climate zone derived from SHAP (tree-based models) and permutation importance (SVR, MLPRegressor, AdaBoost), evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

Figure 6 showed broadly distributed feature-selection patterns across all models in the Temperate zone, with most predictors appearing in 3 to 5 of the top-performing subsets. In the tree-based models, vorticity (relative) and specific snow water content reached the maximum frequency of 5 in Random Forest, while V component of wind, specific humidity, specific rain water content, and specific snow water content reached frequency 5 in XGBoost, and V component of wind and specific cloud ice water content reached frequency 5 in LightGBM. Specific humidity appeared consistently across all models but did not universally reach the maximum frequency in every tree-based model, indicating that while it was an important predictor, the Temperate zone required a broad and distributed predictor set rather than a single dominant variable. Among the non-tree models, specific humidity reached frequency 5 in AdaBoost, while potential vorticity was the most frequently selected variable in MLPRegressor, and vertical velocity reached frequency 5 in SVR. Ozone mass mixing ratio appeared with low frequency across most models, consistent with its marginal predictive relevance observed across all three zones.

Figure 7 presented the feature-importance results for each model group. Among the tree-based models, SHAP values indicated that specific snow water content was the highest-ranked predictor in Random Forest, with a mean SHAP value of around 1.4 to 1.5, followed by geopotential (around 1.0 to 1.1) and specific humidity (around 0.7 to 0.8). In XGBoost, geopotential ranked highest at around 1.2, followed by specific snow water content (around 1.0) and specific humidity (around 0.6 to 0.7). LightGBM assigned the highest importance to geopotential (around 1.3 to 1.4), with potential vorticity as a secondary contributor at around 0.7 and additional contributions from V component of wind and specific cloud ice water content. The prominence of specific snow water content and geopotential in the tree-based SHAP profiles distinguished the Temperate zone from the Cold/Subarctic and Polar zones, where dynamical variables and specific humidity more clearly led the importance rankings.

Among the non-tree models, specific humidity was the dominant predictor across all three models. In AdaBoost, specific humidity reached a permutation importance of around 2.5, substantially exceeding geopotential at around 0.5. In SVR, specific humidity dominated at around 3.0, with fraction of cloud cover as a notable secondary contributor at around 0.8 to 0.9. MLPRegressor showed the widest importance distribution of the non-tree group, with specific humidity leading at around 4.5 but potential vorticity carrying substantial secondary importance at around 3.0, followed by divergence. The permutation importance

values for non-tree models are not numerically comparable to the SHAP values reported for tree-based models and are interpreted within each model group only.

Overall, these results indicated that refractivity-gradient variability in the Temperate zone was associated with a broader and more distributed set of predictors than in the other two zones. No single variable dominated consistently across all models, with specific snow water content, geopotential, and specific humidity emerging as the most prominent predictors in the tree-based group, while specific humidity led the non-tree group but was accompanied by meaningful contributions from potential vorticity, divergence, and fraction of cloud cover.

Table 5a. Feature-selection and hyperparameter-tuning performance metrics for regression models in the Temperate climate zone based on ERA5 data (January 1991–April 2025).

Model	Selected Feature Count	Feature Selection RMSE	Full 14 RMSE	Best RMSE After Tuning
LightGBM	13	1.812	1.843	1.812
Random Forest	12	1.845	1.849	1.845
XGBoost	11	1.878	1.937	1.826
MLPRegressor	7	1.993	2.092	1.980
SVR	11	2.012	2.017	1.853
AdaBoost	9	2.429	2.480	2.439
Ridge	13	2.863	2.870	2.863

Table 5b. Independent test-set performance and bootstrap-derived 95% confidence intervals for regression models in the Temperate climate zone. Confidence intervals were estimated using 1,000 bootstrap resamples of the test predictions.

Model	Test RMSE (95% CI)	Test MAE	Test R ² (95% CI)
LightGBM	1.700 (1.581–1.825)	1.269	0.877 (0.851–0.896)
SVR	1.814 (1.657–1.962)	1.333	0.860 (0.831–0.884)
Random Forest	1.836 (1.710–1.985)	1.345	0.857 (0.828–0.879)
XGBoost	1.853 (1.732–1.996)	1.409	0.854 (0.825–0.877)
MLPRegressor	2.048 (1.888–2.207)	1.540	0.822 (0.786–0.851)
AdaBoost	2.418 (2.270–2.570)	1.965	0.751 (0.712–0.778)
Ridge	2.852 (2.685–3.019)	2.313	0.654 (0.604–0.694)

Tables 5a and 5b present the feature-selection, tuning, and independent test-set performance of all regression models in the Temperate climate zone.

Feature selection reduced validation RMSE relative to full-feature models across most approaches. LightGBM retained 13 of the 14 available predictors, achieving a feature-selection RMSE of 1.812 compared to a full-feature RMSE of 1.843, a marginal improvement reflecting the limited redundancy in the Temperate predictor set. Hyperparameter tuning produced no further change for LightGBM, indicating that the feature-selected configuration was already well-optimised. XGBoost showed the largest improvement from tuning, with RMSE decreasing from 1.878 to 1.826. SVR benefited most from tuning relative to its feature-selection baseline, with RMSE declining from 2.012 to 1.853, and subsequently generalised well to the test set with an RMSE of 1.814, performing slightly better than its tuned validation result. AdaBoost showed an unusual pattern in which the full-feature RMSE of 2.480 was slightly higher than the feature-selection RMSE of 2.429, indicating that feature reduction was beneficial but tuning did not fully recover the improvement. Ridge retained 13 predictors, the same count as LightGBM, yet showed no improvement across any stage, with RMSE remaining near 2.863 throughout. This contrast highlighted that the number of retained predictors alone did not determine predictive performance, but rather the model's capacity to exploit nonlinear predictor interactions.

On the independent test set, LightGBM achieved the lowest RMSE of 1.700 N-units/km, the lowest MAE of 1.269, and the highest R² of 0.877, supported by narrow bootstrap confidence intervals for RMSE (1.581–1.825) and R² (0.851–0.896), indicating stable performance across resampling experiments. SVR, Random Forest, and XGBoost produced closely comparable results, with R² values of 0.860, 0.857, and 0.854 and RMSE values of 1.814, 1.836, and 1.853 N-units/km respectively. The substantial overlap among their bootstrap confidence intervals confirmed that performance differences among these three models were modest. MLPRegressor recorded an RMSE of 2.048 N-units/km and an R² of 0.822, with a slight deterioration relative to its tuned validation RMSE of 1.980, suggesting mild overfitting despite its compact seven-predictor feature set. AdaBoost and Ridge showed the weakest predictive performance, with Ridge recording an RMSE of 2.852 N-units/km and R² of 0.654.

Overall, these results indicated that LightGBM, SVR, Random Forest, and XGBoost all provided reliable refractivity-gradient predictions in the Temperate zone, with the large performance gap between these models and Ridge regression consistent with the limitations of linear modelling in this dynamically variable environment.

Table 6. Selected features and optimal hyperparameter configurations for regression models in the Temperate climate zone based on ERA5 data (January 1991–April 2025), with model performance evaluated on the 20% hold-out test dataset.

Model	Selected Features	Best Params
-------	-------------------	-------------

LightGBM	Specific humidity, Divergence, Geopotential, Specific cloud ice water content, Specific rain water content, U component of wind, V component of wind, Potential vorticity, Fraction of cloud cover, Vertical velocity, Specific cloud liquid water content, Vorticity (relative), Ozone mass mixing ratio	{'regressor__learning_rate': 0.1, 'regressor__n_estimators': 100, 'regressor__num_leaves': 31}
SVR	Specific humidity, Divergence, Geopotential, V component of wind, Fraction of cloud cover, Vertical velocity, Specific rain water content, Specific cloud ice water content, U component of wind, Specific snow water content, Potential vorticity	{'regressor__C': 10, 'regressor__gamma': 'scale', 'regressor__kernel': 'rbf'}
Random Forest	Specific humidity, Divergence, Geopotential, V component of wind, Specific cloud ice water content, Potential vorticity, Specific rain water content, U component of wind, Vorticity (relative), Fraction of cloud cover, Vertical velocity, Specific snow water content	{'regressor__max_depth': None, 'regressor__n_estimators': 100}
XGBoost	Specific humidity, Divergence, Geopotential, Potential vorticity, V component of wind, Specific cloud liquid water content, U component of wind, Specific snow water content, Specific rain water content, Vorticity (relative), Specific cloud ice water content	{'regressor__learning_rate': 0.1, 'regressor__max_depth': 6, 'regressor__n_estimators': 200}
Model	Selected Features	Best Params
MLPRegressor	Specific humidity, Vertical velocity, Fraction of cloud cover, Divergence, V component of wind, Ozone mass mixing ratio, Potential vorticity	{'regressor__alpha': 0.0001, 'regressor__hidden_layer_sizes': (100, 50), 'regressor__learning_rate_init': 0.001}
AdaBoost	Specific humidity, Divergence, Geopotential, V component of wind, Potential vorticity, Vertical velocity, Vorticity (relative), Ozone mass mixing ratio, Fraction of cloud cover	{'regressor__learning_rate': 1.0, 'regressor__n_estimators': 100}

Table 6 complemented the performance patterns observed in Table 5a. A notable contrast was observed between LightGBM and Ridge regression, both of which retained 13 predictors. Despite using an identical feature count, LightGBM achieved a test RMSE of 1.700 N-units/km and R^2 of 0.877 while Ridge recorded an RMSE of 2.852 N-units/km and R^2 of 0.654, illustrating that feature count alone did not determine predictive performance.

MLPRegressor retained the smallest subset of seven variables. This compact selection was associated with a modest tuning improvement in Table 5a, from 1.993 to 1.980, suggesting that the reduced feature set constrained the gains available from hyperparameter optimisation.

The hyperparameter configurations varied across models. XGBoost used 200 estimators while LightGBM and Random Forest used 100. This larger ensemble in XGBoost was consistent with the tuning improvement observed in Table 5a, where XGBoost showed the largest reduction from feature-selection RMSE of 1.878 to tuned RMSE of 1.826 among the tree-based models. LightGBM and Random Forest showed no improvement from tuning, with Best RMSE After Tuning matching their feature-selection RMSE exactly, indicating that their initial configurations were already well-suited to the Temperate predictor space. SVR used an RBF kernel with a regularisation parameter of 10 and showed the largest overall tuning improvement among all models, with RMSE declining from 2.012 to 1.853. AdaBoost used a learning rate of 1.0 and 100 estimators, which may have contributed to its comparatively weaker test generalisation. Ridge used a regularisation parameter of 0.1, the simplest configuration of all models, consistent with its substantially weaker predictive performance across all evaluation stages in Table 5a.

3.2. Cold/Subarctic

Monthly Predicted vs Observed by Model - Cold/Sub-Arctic




Figure 8: Monthly predicted vs observed refractivity gradient (N/km) for the Cold/Subarctic climate zone across models, based on monthly averages from ERA5 data (January 1991–April 2025).

Predicted vs Observed by Model - Cold/Sub-Arctic




Figure 9: Predicted vs observed refractivity gradient (N/km) scatter plots for the Cold/Subarctic climate zone across models, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).




Figure 10: Monthly predicted versus observed refractivity gradient (N/km) and predicted versus actual scatter plot for Ridge regression in the Cold/Subarctic climate zone, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

Figure 8 presents the mean monthly predicted and observed refractivity gradients for the Cold/Subarctic climate zone. The seasonal cycle is well-defined but more constrained than in the Temperate zone, consistent with the lower peak water vapour pressure in this region as shown in Figure 2c. The least negative gradients occur in January and February, coinciding with the winter temperature minimum and suppressed water vapour pressure. Despite relatively high relative humidity during this period, low absolute temperatures limit actual atmospheric moisture content, confirming that relative humidity alone is a poor indicator of refractivity-gradient strength at subarctic latitudes.

A notable distinction from the Temperate zone is the timing of the summer minimum. In the Temperate zone, the most negative gradients occur in July and are reproduced accurately by all six models. In the Cold/Subarctic zone, however, the most negative actual gradient occurs in August. None of the six models captured this August minimum. All models predicted the lowest gradient in July, with August as the next most negative month. This systematic one-month shift suggests the models respond primarily to the temperature and vapour pressure maxima in July rather than to the actual atmospheric moisture conditions that drive the gradient minimum later in the season.

LightGBM and XGBoost show the closest overall agreement with actual values, capturing the descent toward the summer trough and the autumn recovery with low deviation. Random Forest performs comparably with minor underestimation of trough depth. SVR and MLPRegressor track the seasonal cycle adequately but exhibit greater smoothing around the summer minimum. AdaBoost produces the largest departure, with a shallower seasonal curve that underestimates the depth of the summer minimum and the steepness of the spring descent.

Figure 9 presents the scatter plots of predicted versus actual refractivity gradient values for all six models on the 20% hold-out test dataset. A strong positive linear relationship is evident across all models, with points distributed closely around the 1:1 reference line. LightGBM, XGBoost, and Random Forest display the tightest clustering, indicating strong predictive accuracy across the full gradient range. SVR and MLPRegressor perform well but show slightly increased spread at more negative values. AdaBoost exhibits the widest dispersion, consistent with the seasonal underestimation identified in Figure 8.

Figure 10a shows the monthly predicted versus actual refractivity gradients for Ridge regression. Ridge reproduced the general seasonal direction but agreement breaks down at the spring transition from April through June, where predicted values are considerably more negative than actual. This differs from the Temperate zone where Ridge errors were distributed more broadly across seasons. In the Cold/Subarctic zone, errors are concentrated at the point of spring moisture accumulation, where the nonlinear rate of boundary-layer moistening cannot be represented by a linear predictor-response relationship. Figure 10b confirms these limitations at the individual observation level, with a noticeably wider point cloud than the ensemble models in Figure 9 and systematic departures concentrated at the most negative end of the distribution. These results are consistent with the weaker test-set performance of Ridge regression reported in Table 7b and reinforce the conclusion that linear modelling is insufficient to capture the nonlinear moisture-driven dynamics governing refractivity-gradient variability in this climate zone.

Top-Subset Feature Frequency by Model - Cold/Sub-Arctic




Figure 11: Top-subset feature frequency across regression models in the Cold/Subarctic climate zone, showing the frequency of features appearing in the five best-performing subsets based on validation RMSE using ERA5 data (January 1991–April 2025)




Figure 12: Feature importance across regression models in the Cold/Subarctic climate zone derived from SHAP (tree-based models) and permutation importance (SVR, MLPRegressor, AdaBoost), based on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

Figure 11 showed the top-subset feature selection frequency across all six models in the Cold/Subarctic climate zone. Among the tree-based models, specific cloud ice water content reached the maximum frequency of 5 in LightGBM, alongside V component of wind and vorticity relative. In XGBoost, specific humidity reached frequency 5, with V component of wind, divergence, and specific snow water content also prominently selected. Random Forest showed a more distributed pattern, with specific cloud ice water content, specific humidity, U component of wind, and potential vorticity all appearing with high frequency. Among the non-tree models, vorticity relative reached maximum frequency in AdaBoost, with geopotential, divergence, and fraction of cloud cover also consistently retained. SVR selected specific humidity and vorticity relative at maximum frequency, with specific cloud ice water content, potential vorticity, and fraction of cloud cover as prominent secondary selections. MLPRegressor showed specific cloud liquid water content and specific humidity reaching the highest frequencies, with geopotential, fraction of cloud cover, and divergence also regularly selected. Specific snow water content appeared with moderate to high frequency across several models, reflecting the predictive relevance of cold-region microphysical processes in this zone. Compared to the Temperate zone, specific snow water content showed reduced prominence in the tree-based importance rankings, where it led SHAP values in the Temperate zone but contributed minimally across all three tree-based models in the Cold/Subarctic zone, with dynamical variables such as potential vorticity and geopotential taking a more dominant role.

Figure 12 showed the feature importance results for each model group in the Cold/Subarctic zone. Among the tree-based models, SHAP values indicated that potential vorticity was the leading contributor in both Random Forest and LightGBM, with fraction of cloud cover and geopotential as prominent secondary contributors in Random Forest, and geopotential and U component of wind in LightGBM. In XGBoost, geopotential ranked highest with the largest mean SHAP value among all tree-based models, followed by potential vorticity and fraction of cloud cover as strong secondary predictors. Among the non-tree models, permutation importance identified specific humidity as the dominant predictor across all three models, with a substantially larger importance value than all other variables. Potential vorticity appeared as a secondary contributor in AdaBoost, while specific cloud ice water content and vorticity relative ranked prominently in SVR. In MLPRegressor, divergence and specific cloud liquid water content contributed alongside specific humidity as secondary predictors. The contrast between dynamical variables ranking highest under

SHAP and specific humidity dominating under permutation importance reflected methodological differences between the two approaches rather than a physical contradiction. Both groups qualitatively agreed that geopotential, potential vorticity, specific humidity, and cloud-related variables were the most consistently influential predictors of refractivity-gradient variability in the Cold/Subarctic zone.

Table 7a. Feature-selection and hyperparameter-tuning performance metrics for regression models in the Cold/Subarctic climate zone based on ERA5 data (January 1991–April 2025).

Model	Selected Feature Count	Feature Selection RMSE	Full 14 RMSE	Best RMSE After Tuning
LightGBM	12	0.945	0.968	0.944
XGBoost	9	0.968	1.021	0.997
Random Forest	9	1.001	1.041	0.999
SVR	12	1.088	1.085	0.987
MLPRegressor	12	1.095	1.151	1.066
AdaBoost	11	1.317	1.409	1.327
Ridge	13	1.387	1.387	1.387

Table 7b. Independent test-set performance and bootstrap-derived 95% confidence intervals for regression models in the Cold/Subarctic climate zone.

Model	Test RMSE (95% CI)	Test MAE	Test R ² (95% CI)
LightGBM	1.023 (0.908–1.137)	0.706	0.929 (0.908–0.946)
XGBoost	1.061 (0.947–1.177)	0.741	0.923 (0.900–0.941)
Random Forest	1.077 (0.960–1.194)	0.736	0.921 (0.897–0.939)
SVR	1.097 (0.939–1.265)	0.709	0.918 (0.887–0.942)
MLPRegressor	1.104 (0.974–1.224)	0.765	0.917 (0.892–0.937)
AdaBoost	1.413 (1.298–1.527)	1.093	0.864 (0.835–0.886)
Ridge	1.554 (1.419–1.678)	1.167	0.835 (0.794–0.867)

Table 8. Selected features and optimal hyperparameter configurations for regression models in the Cold/Subarctic climate zone based on ERA5 data (January 1991–April 2025), with model performance evaluated on the 20% hold-out test dataset.

Model	Selected Features	Best Params
LightGBM	Specific humidity, Potential vorticity, Fraction of cloud cover, U component of wind, Geopotential, Specific snow water content, Divergence, Specific rain water content, Ozone mass mixing ratio, Vorticity (relative), Specific cloud ice water content, Vertical velocity	{'regressor__learning_rate': 0.1, 'regressor__n_estimators': 200, 'regressor__num_leaves': 31}
XGBoost	Specific humidity, Potential vorticity, Fraction of cloud cover, U component of wind, Vertical velocity, Geopotential, Vorticity (relative), Specific rain water content, Ozone mass mixing ratio	{'regressor__learning_rate': 0.1, 'regressor__max_depth': 6, 'regressor__n_estimators': 200}
Random Forest	Specific humidity, Potential vorticity, Fraction of cloud cover, U component of wind, Geopotential, Divergence, Specific snow water content, Specific cloud ice water content, Ozone mass mixing ratio	{'regressor__max_depth': None, 'regressor__n_estimators': 200}
SVR	Specific humidity, Specific cloud liquid water content, U component of wind, Vorticity (relative), Geopotential, Divergence, Fraction of cloud cover, Specific snow water content, Vertical velocity, Potential vorticity, Specific rain water content, Specific cloud ice water content	{'regressor__C': 10, 'regressor__gamma': 'scale', 'regressor__kernel': 'rbf'}
MLPRegressor	Specific humidity, Divergence, Specific cloud liquid water content, U component of wind, Geopotential, Fraction of cloud cover, Vorticity (relative), Vertical velocity, Specific snow water content, Potential vorticity, Specific rain water content, V component of wind	{'regressor__alpha': 0.0001, 'regressor__hidden_layer_sizes': (100, 50), 'regressor__learning_rate_init': 0.001}
AdaBoost	Specific humidity, Fraction of cloud cover, Potential vorticity, U component of wind, Vorticity (relative), Ozone mass mixing ratio, V component of wind, Specific cloud ice water content, Divergence, Geopotential, Specific rain water content	{'regressor__learning_rate': 1.0, 'regressor__n_estimators': 100}

Tables 7a and 7b present the feature-selection, tuning, and independent test-set performance of all regression models in the Cold/Subarctic climate zone.

Feature selection reduced validation RMSE relative to full-feature models across most approaches. LightGBM achieved the lowest feature-selection RMSE of 0.945, with negligible change after tuning (0.944), indicating that the selected predictor subset adequately captured the relevant predictive information. XGBoost and Random Forest showed improvements from feature

selection, with RMSE decreasing from 1.021 to 0.968 and 1.041 to 1.001 respectively. SVR benefited most from hyperparameter tuning, with RMSE declining from 1.088 to 0.987. Ridge showed no improvement at any stage, remaining at approximately 1.387 throughout.

On the independent test set, LightGBM recorded the lowest RMSE of 1.023 N-units/km, lowest MAE of 0.706, and highest R^2 of 0.929. XGBoost, Random Forest, SVR, and MLPRegressor produced closely comparable results, with R^2 values of 0.923, 0.921, 0.918, and 0.917 respectively. The substantial overlap among their bootstrap confidence intervals, for example R^2 intervals of 0.908–0.946 for LightGBM and 0.887–0.942 for SVR, indicated that performance differences among the five leading models were modest. Notably, SVR recorded a lower MAE of 0.709 than both Random Forest (0.736) and XGBoost (0.741) despite a slightly higher RMSE, suggesting its errors were generally small but included occasional larger deviations. AdaBoost and Ridge showed the weakest performance, with RMSE values of 1.413 and 1.554 N-units/km and R^2 values of 0.864 and 0.835 respectively.

Table 8 shows that specific humidity, potential vorticity, geopotential, fraction of cloud cover, divergence, and wind components were consistently retained across models, with specific snow water content also selected by several models, reflecting the predictive relevance of cold-region microphysical variables. LightGBM and XGBoost required 200 estimators compared to 100 in the Temperate zone, and Random Forest similarly employed 200 trees, reflecting the increased model capacity needed to represent the more stable but dynamically complex Cold/Subarctic atmospheric environment.

Overall, these results indicated that nonlinear ensemble approaches provided the strongest predictive performance for refractivity-gradient estimation in the Cold/Subarctic zone, with the consistent selection of both dynamical and thermodynamic predictors distinguishing this zone from the moisture-dominated Temperate environment.

3.3. Polar Zone

Monthly Predicted vs Observed by Model - Polar




Figure 13: Aggregated monthly predicted versus observed refractivity-gradient behaviour for the for the Polar dataset comprising two Antarctic stations, based on ERA5 monthly averages from January 1991 to April 2025.

Predicted vs Observed by Model - Polar




Figure 14: Predicted vs observed refractivity gradient (N/km) scatter plots for the Polar climate zone across models, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).




Figure 15: Monthly predicted versus observed refractivity gradient (N/km) and predicted versus actual scatter plot for Ridge regression in the Polar climate zone, evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

The Polar climate zone presented a fundamentally different predictive challenge from the Temperate and Cold/Subarctic zones. The observed refractivity gradient varied within a considerably narrower seasonal range, with monthly values spanning approximately -79 to -59 N-units/km and a seasonal amplitude of approximately 20 N-units/km, substantially smaller than those

observed in the warmer zones. Within this compressed range, the seasonal cycle was nonetheless well-defined, with the least negative gradients occurring in February and November and the strongest negative values during July–August, driven by the modest but consistent increase in water vapour pressure during the Antarctic summer (Figure 2c).

The narrower target range meant that even small absolute prediction errors represented a proportionally larger fraction of the total observed variability than in the other zones. This context is important when interpreting the lower absolute RMSE values reported for this zone; they reflected the limited variability of the Polar dataset rather than inherently stronger predictive skill. The R^2 values, which account for the variance structure of the target, therefore provided a more informative basis for cross-zone comparison than RMSE alone.

LightGBM, Random Forest, XGBoost, and SVR all tracked the observed seasonal cycle with reasonable fidelity (Figure 13), reproducing both the winter minima in gradient intensity and the summer deepening. However, even the leading models showed occasional underestimation of gradient magnitude during October–November, a period characterised by the rapid atmospheric transition into the Antarctic polar night. AdaBoost and MLPRegressor produced notably smoother seasonal profiles, underestimating the strongest negative gradients during late summer. In the Polar zone this smoothing behaviour was more consequential than in the other zones precisely because the observable seasonal range was already small; flattening the seasonal curve here effectively removed a substantial portion of the real atmospheric signal.

Scatter plots (Figure 14) confirmed the tighter clustering of LightGBM, Random Forest, XGBoost, and SVR around the 1:1 reference line, with departures concentrated at the more negative end of the gradient distribution. Ridge regression (Figure 15) underestimated both the February–November peaks and the July–August minimum, compressing the predicted seasonal range further than any other model and producing the widest scatter dispersion, consistent with its inability to represent the nonlinear predictor–response structure present even within the Polar dataset's constrained variability range.




Figure 16: Top-subset feature frequency across regression models in the Polar climate zone, showing the frequency of features appearing in the five best-performing subsets based on validation RMSE using ERA5 data (January 1991–April 2025).




Figure 17: Feature importance across regression models in the Polar climate zone derived from SHAP (tree-based models) and permutation importance (SVR, MLPRegressor, AdaBoost), evaluated on the 20% hold-out test dataset using ERA5 data (January 1991–April 2025).

Figures 16 and 17 present the feature-selection frequency and feature-importance results for the Polar climate zone. The results reveal a consistent dominance of moisture-related variables, particularly specific humidity, together with atmospheric circulation indicators such as divergence, potential vorticity, wind components, and cloud-related parameters.

Figure 16 showed a broader and more evenly distributed feature-selection pattern across models. Specific humidity reached the maximum frequency of 5 across most models but was not uniquely dominant in the selection stage, as specific cloud liquid water content, U and V wind components, potential vorticity, divergence, and fraction of cloud cover also achieved frequencies of 4 to 5 across multiple models. Geopotential appeared with moderate frequency in LightGBM and XGBoost but less consistently in Random Forest and the non-tree models. Ozone mass mixing ratio and specific rain water content appeared infrequently across most models, with frequencies of 1 to 2, indicating limited independent predictive relevance in the Polar environment.

Figure 17 provided a more nuanced picture of predictor influence within each model. Among the tree-based models, specific humidity was the highest-ranked predictor across all three, with mean SHAP values of around 0.6 in Random Forest, 0.8 in XGBoost, and 0.5 in LightGBM. The margin of dominance varied across models. In Random Forest, divergence (around 0.35) and potential vorticity (around 0.3) were notable secondary contributors. In XGBoost, divergence (around 0.45) and potential vorticity (around 0.3) similarly occupied prominent secondary positions. LightGBM presented the most distributed importance profile of the three tree models, with geopotential (around 0.5) approaching specific humidity (around 0.5) as a co-leading predictor, and divergence, U wind, and V wind contributing at intermediate levels. This suggested that large-scale atmospheric structure carried comparatively more predictive relevance in LightGBM's representation of the Polar environment than in the other tree-based models.

Among the non-tree models, specific humidity was again the dominant predictor across AdaBoost, SVR, and MLPRegressor. In AdaBoost, the margin was clear, with specific humidity (around 0.5) substantially exceeding potential vorticity and geopotential. In SVR, specific humidity (around 1.0) was the leading predictor but specific cloud liquid water content (around 0.6) and V wind (around 0.35) were notable secondary contributors. MLPRegressor showed the most compressed importance distribution of the non-tree group, with specific humidity (around 1.45) leading but specific cloud liquid water content (around 1.2) and divergence

(around 1.0) also carrying substantial predictive weight, suggesting that this model relied on a broader combination of moisture and dynamical variables than the other non-tree approaches.

Comparing across zones, the Polar results showed a clearer and more consistent dominance of specific humidity than either the Temperate or Cold/Subarctic zones, but the secondary predictors, particularly divergence, potential vorticity, wind components, and cloud liquid water content, remained meaningful contributors rather than negligible ones. The prominence of cloud liquid water content as a secondary predictor in SVR and MLPRegressor was more pronounced in the Polar zone than in the other two zones, a pattern potentially reflecting the sensitivity of these model architectures to cloud-phase information under persistently cold atmospheric conditions.

Across all three climate zones, ozone mass mixing ratio consistently ranked among the lowest-importance predictors and appeared infrequently in the top-performing feature subsets. In the Temperate zone, ozone reached a selection frequency of only 1 in most models and was absent from the upper importance rankings in both SHAP and permutation importance analyses. In the Cold/Subarctic zone, ozone showed slightly higher selection frequencies in AdaBoost and SVR, reaching around 3 to 4, but contributed negligible importance values in the tree-based SHAP profiles. In the Polar zone, ozone appeared with frequencies of 1 to 2 across the non-tree models and was effectively absent from all importance rankings. This consistent pattern across all three zones and both importance methods indicated that ozone mass mixing ratio carried limited independent predictive information for refractivity-gradient estimation within this dataset. Its inclusion in the candidate predictor pool was not detrimental to model performance, as the feature-selection procedure correctly marginalised it in most cases, but it did not emerge as a meaningful contributor to predictive accuracy in any of the three climatic environments investigated.

Table 9a. Feature-selection and hyperparameter-tuning performance metrics for regression models in the Polar climate zone based on ERA5 data (January 1991–April 2025).

Model	Selected Feature Count	Feature Selection RMSE	Full 14 RMSE	Best RMSE After Tuning
LightGBM	11	0.444	0.478	0.444
Random Forest	7	0.425	0.455	0.425
XGBoost	10	0.455	0.472	0.466
SVR	10	0.497	0.555	0.450
MLPRegressor	4	0.498	0.537	0.503
AdaBoost	8	0.782	1.202	0.669
Ridge	13	0.684	0.687	0.682

Table 9b. Independent test-set performance and bootstrap-derived 95% confidence intervals for regression models in the Polar climate zone.

Model	Test RMSE (95% CI)	Test MAE	Test R ² (95% CI)
Random Forest	0.390 (0.329–0.452)	0.388	0.819 (0.746–0.877)
LightGBM	0.395 (0.324–0.468)	0.375	0.814 (0.746–0.874)
SVR	0.433 (0.366–0.509)	0.422	0.777 (0.682–0.839)
XGBoost	0.409 (0.345–0.469)	0.451	0.800 (0.714–0.861)
AdaBoost	0.483 (0.410–0.551)	0.554	0.721 (0.615–0.802)
Model	Test RMSE (95% CI)	Test MAE	Test R ² (95% CI)
MLPRegressor	0.612 (0.521–0.718)	0.582	0.553 (0.413–0.657)
Ridge	0.646 (0.576–0.717)	0.615	0.502 (0.358–0.604)

Table 9a showed that feature selection reduced validation RMSE relative to full-feature models across most approaches. Random Forest achieved the lowest feature-selection RMSE of 0.425 while retaining only 7 predictors, whereas LightGBM reached a comparable validation RMSE of 0.444 using 11 predictors, indicating that a more compact predictor subset was sufficient for effective gradient estimation by Random Forest in the Polar zone. SVR showed the largest improvement from tuning, with RMSE declining from 0.497 to 0.450, while AdaBoost benefited substantially from both feature selection and tuning, with RMSE falling from a full-feature value of 1.202 to 0.669 after optimisation. Ridge showed negligible improvement across all stages, with RMSE remaining between 0.682 and 0.687, consistent with the limited predictive capacity of linear modelling in this environment.

On the independent test set, Random Forest recorded the lowest RMSE of 0.390 N-units/km and the highest R² of 0.819, followed closely by LightGBM with an RMSE of 0.395 and R² of 0.814 (Table 9b). This represented a reversal of the ranking observed in the Temperate and Cold/Subarctic zones, where LightGBM consistently led. The narrow margin between the two models, together with the extensive overlap in their bootstrap confidence intervals (R² of 0.746–0.877 for Random Forest and 0.746–0.874 for LightGBM, RMSE of 0.329–0.452 and 0.324–0.468 respectively), indicated that this difference was not practically meaningful and should not be interpreted as a definitive performance advantage for Random Forest in this zone. XGBoost also produced competitive results with an RMSE of 0.409 N-units/km and R² of 0.800, with overlapping confidence intervals confirming the broadly comparable predictive performance of the three leading ensemble models.

SVR recorded an RMSE of 0.433 N-units/km and R² of 0.777, performing creditably within the Polar dataset. AdaBoost showed weaker predictive performance with an RMSE of 0.483 N-units/km and R² of 0.721. The most notable result in this zone was the performance of MLPRegressor, which recorded an RMSE of 0.612 and R² of 0.553, a substantially weaker result than in the Temperate zone having an R² of 0.822 and Cold/Subarctic zone having an R² of 0.917. This deterioration was likely associated with the limited size of the Polar training sample, which comprised only 2 stations, reducing the volume of data available for the

neural network to learn stable predictor-response relationships. Ridge recorded the lowest R^2 of 0.502, consistent with its performance pattern across all three zones.

Table 10. Selected features and optimal hyperparameter configurations for regression models in the Polar climate zone based on ERA5 data (January 1991–April 2025), with model performance evaluated on the 20% hold-out test dataset.

Model	Selected Features	Best Parameters
LightGBM	Specific humidity, Divergence, Ozone mass mixing ratio, Fraction of cloud cover, Potential vorticity, U component of wind, Specific rain water content, Specific snow water content, Specific cloud ice water content, V component of wind, Geopotential	{'regressor__learning_rate': 0.1, 'regressor__n_estimators': 100, 'regressor__num_leaves': 31}
Random Forest	Specific humidity, Divergence, Potential vorticity, U component of wind, Fraction of cloud cover, V component of wind, Specific snow water content	{'regressor__max_depth': None, 'regressor__n_estimators': 100}
XGBoost	Specific humidity, Divergence, Ozone mass mixing ratio, Potential vorticity, U component of wind, Vorticity (relative), Fraction of cloud cover, Specific cloud ice water content, Specific cloud liquid water content, Specific rain water content	{'regressor__learning_rate': 0.1, 'regressor__max_depth': 6, 'regressor__n_estimators': 200}
SVR	Specific humidity, Divergence, Fraction of cloud cover, Geopotential, Specific snow water content, U component of wind, V component of wind, Specific cloud ice water content, Specific cloud liquid water content, Specific rain water content	{'regressor__C': 10, 'regressor__gamma': 'scale', 'regressor__kernel': 'rbf'}
MLPRegressor	Specific humidity, Specific snow water content, Specific rain water content, Specific cloud ice water content	{'regressor__alpha': 0.0001, 'regressor__hidden_layer_sizes': (100, 50), 'regressor__learning_rate_init': 0.001}
AdaBoost	Specific humidity, Divergence, V component of wind, Geopotential, Potential vorticity, U component of wind, Fraction of cloud cover, Specific cloud liquid water content	{'regressor__learning_rate': 1.0, 'regressor__n_estimators': 100}
Ridge	Specific humidity, U component of wind, V component of wind, Specific cloud liquid water content, Divergence, Specific cloud ice water content, Fraction of cloud cover, Potential vorticity, Ozone mass mixing ratio, Specific snow water content, Geopotential, Vertical velocity, Specific rain water content	{'regressor__alpha': 0.1}

Table 10 showed that specific humidity was retained in every optimal feature subset across all models, indicating that it carried consistent predictive relevance for refractivity-gradient estimation in the Polar zone. Divergence, U and V wind components, and fraction of cloud cover were also repeatedly selected across most models, indicating that these variables contributed consistent predictive information across model architectures.

Potential vorticity appeared in the optimal feature sets of LightGBM, Random Forest, XGBoost, AdaBoost, and Ridge, suggesting that variables associated with atmospheric stability and circulation contributed useful predictive information. Geopotential was retained by LightGBM, SVR, AdaBoost, and Ridge, indicating that large-scale vertical atmospheric structure remained a relevant predictor in this zone. Ozone mass mixing ratio was selected by LightGBM, XGBoost, and Ridge, consistent with the pattern observed across all three zones in which ozone appeared in some feature subsets but contributed negligible importance values, with its predictive relevance likely reflecting statistical covariance with dynamical variables rather than independent refractivity-related information.

Cloud microphysical variables were widely represented across the selected feature sets. Specific cloud ice water content, specific cloud liquid water content, specific rain water content, and specific snow water content were retained by several models, particularly LightGBM, XGBoost, SVR, and Ridge. Their repeated selection indicated that cloud microphysical variables carried useful predictive information within the Polar dataset. The selection of snow and cloud ice water content variables was consistent with the prevalence of frozen hydrometeors in Polar atmospheric conditions.

Differences in selected feature counts illustrated variations in model complexity across architectures. Random Forest achieved its strongest validation performance using only 7 predictors, while LightGBM retained 11 and Ridge retained 13. MLPRegressor selected the smallest subset, consistent with its reliance on a compact set of moisture and hydrometeor variables. These differences in feature count were consistent with the known architectural differences among the models, with tree-based ensemble methods generally benefiting from broader predictor sets and the neural network architecture relying on a more concentrated subset. Overall, refractivity-gradient predictions in the Polar zone were associated with a combination of moisture, dynamical, and cloud microphysical predictors rather than any single predictor, a pattern broadly consistent across all three climatic environments investigated in this study.

4. CONCLUSION

This study developed and evaluated a machine learning framework for predicting lower-tropospheric vertical radio refractivity gradients across selected Temperate, Cold/Subarctic, and Polar climatic environments, using long-term ERA5 reanalysis data spanning January 1991 to April 2025 and a set of 14 thermodynamic, dynamical, and cloud microphysical predictors drawn from a 14-station, multi-climate sample. The results indicate that nonlinear ensemble tree-based architectures achieved the highest predictive accuracy across all three climate zones, outperforming simpler boosting approaches and the linear Ridge regression

baseline. LightGBM performed best in the Temperate and Cold/Subarctic zones, while Random Forest marginally led in the Polar zone, though overlapping bootstrap confidence intervals indicate no decisive practical difference among the leading ensemble models. LightGBM explained approximately 88% of refractivity-gradient variability in the Temperate zone, with a high R^2 of 0.877 and an RMSE of 1.700 N-units/km. Its best performance was in the Cold/Subarctic zone, showing an R^2 of 0.929 and an RMSE of 1.023 N-units/km. In the Polar zone, it captured the refractivity-gradient behaviour with an R^2 of 0.814 and an RMSE of 0.395 N-units/km. It should be noted that the lower absolute RMSE observed in the Polar zone partly reflects the narrower range of the target variable in that zone rather than a straightforward improvement in predictive skill relative to the other climate zones. Random Forest and XGBoost produced broadly comparable results across all zones, with bootstrap confidence intervals that overlapped substantially, indicating that the performance advantage of LightGBM over these models is modest and should not be treated as definitive. AdaBoost and Ridge regression consistently recorded lower accuracy, consistent with the limitations of shallow boosting and linear modelling for high-dimensional nonlinear atmospheric regression tasks.

These findings are broadly consistent with recent work in radio meteorological machine learning. The tendency for ensemble tree-based methods to outperform both linear baselines and simpler neural architectures aligns with results reported by Onawumi et al. [32], who found that Random Forest and LightGBM outperformed LSTM and GRU models which yielded negative R^2 values in ERA5-based refractivity gradient prediction across West Africa, and with Claverie and Motsch [7], who found that a decision tree regressor outperformed a neural network for vertical refractivity profile estimation in maritime environments. Yang et al. [9] demonstrated that an LSTM-based model achieved a MAPE of 2.75% for time-series refractivity prediction, outperforming ARIMA, supporting the broader finding that nonlinear architectures represent atmospheric refractivity more effectively than conventional statistical approaches. The present study extends these comparisons to high-latitude and cold-climate environments where such evaluations have been limited [35].

The feature-selection and feature-importance analyses identified cross-climate differences in the predictors most consistently associated with refractivity-gradient variability. In the Temperate zone, specific humidity, geopotential, potential vorticity, and cloud microphysical variables were most frequently selected. In the Cold/Subarctic zone, dynamical variables particularly potential vorticity and geopotential ranked most prominently under SHAP analysis for tree-based models, while specific humidity was the leading predictor under permutation importance for non-tree models. In the Polar zone, specific humidity was the most consistently selected and highest-ranked predictor across all models and both importance methods. It is important to note that SHAP values and permutation importance scores operate on different scales and cannot be compared numerically across model groups; these results therefore reflect the qualitative direction of predictor influence within each method rather than any causal attribution.

The framework has not been validated against independent radiosonde profiles or operational propagation measurements, and its practical applicability to radar or communication system design cannot be established from the present results alone. The dataset was split randomly rather than chronologically, which means temporal ordering was not preserved and the reported test performance reflects interpolation within the same station-month distribution rather than out-of-sample temporal generalisation; a chronological or leave-one-station-out split would provide a more stringent test.

Consequently, the reported test metrics should not be interpreted as evidence of generalisation to unseen time periods or locations, and the results are best understood as preliminary ERA5-based predictive evidence rather than a demonstration of operational readiness.

Notwithstanding these limitations, the results indicate that ensemble tree-based machine learning applied to ERA5 pressure-level predictors can capture a substantial proportion of refractivity-gradient variability across climatically contrasting environments. The finding that specific humidity, potential vorticity, divergence, and geopotential are consistently associated with refractivity-gradient variability across all three zones suggests that a common core predictor set may underpin ERA5-based gradient estimation in different selected climatic settings, with zone-specific microphysical variables providing supplementary predictive information. These findings may support future work on reanalysis-based refractivity-gradient modelling and anomalous propagation studies in high-latitude environments [34] where in-situ radiosonde observations are sparse, though further validation against independent observational data would be needed before any operational application could be considered.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Solomon O. Adeola: Conceptualization, Methodology, Software, Formal analysis, Data curation, Visualization, Writing - original draft. **Francis O. Aweda:** Supervision, Writing - review & editing. **Jacob A. Akinpelu:** Supervision, Methodology, Validation, Writing - review & editing. **Saeed A. Bello:** Formal analysis, Writing - review & editing. **Nathaniel O. Adeniji:** Writing - review & editing.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the European Centre for Medium-Range Weather Forecasts (ECMWF) for making the ERA5 daily mean reanalysis data available for this study.

FUNDING

This research received no external funding.

DATA AVAILABILITY STATEMENT

Data used in this study can be obtained from the following link: <https://cds.climate.copernicus.eu/datasets/reanalysis-ERA5-pressure-levels?tab=download>.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE MANUSCRIPT PREPARATION PROCESS

During the preparation of this work the authors used a generative AI-based language assistant in order to check grammar and spelling and to improve the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

- [1] Javeed, S., Alimgeer, K.S., Javed, W., Atif, M., Uddin, M. (2018). A modified artificial neural network based prediction technique for tropospheric radio refractivity. *PLoS ONE*, 13(3): e0192069. <https://doi.org/10.1371/journal.pone.0192069>
- [2] Ashidi, A.G. (2024). Autoregressive modelling of tropospheric radio refractivity over selected locations in tropical Nigeria using artificial neural network. *Earth Science Informatics*, 17: 5401-5416. <https://doi.org/10.1007/s12145-024-01489-y>
- [3] Muhammed, K.D., Igwe, K.C., Ezenwora, J.A. (2022). Computation of radio refractivity using machine learning techniques. *Proceedings of International Conference on Information Systems and Emerging Technologies*. <https://doi.org/10.2139/ssrn.4331592>
- [4] Oyewole, O., Sirola, C., Pandit, A., and Ayo, S. (2025). The study of altitudinal refractivity and refractivity gradient variations by weather parameters in some selected locations in Nigeria. *Physical Science International Journal*, 29(3), 42–49. <https://doi.org/10.9734/psij/2025/v29i3882>
- [5] Bettouche, Y., Agba, B.L., Kouki, A.B., Obeidat, H., Alabdullah, A., Abdussalam, F., Ghauri, S., Abd-Alhameed, R. (2020). Estimation and analysis of the radio refractivity, its gradient and the geoclimatic factor in Arctic regions. *Progress in Electromagnetics Research M*, 92: 181-192. <https://doi.org/10.2528/PIERM20021906>
- [6] Norin, L. (2023). Observations of anomalous propagation over waters near Sweden. *Atmospheric Measurement Techniques*, 16(7), 1789–1801. <https://doi.org/10.5194/amt-16-1789-2023>
- [7] Claverie, J., Motsch, J. (2022). Applying machine learning techniques to compute vertical refractivity profiles in maritime environments. *Proceedings of the 16th European Conference on Antennas and Propagation (EuCAP 2022)*. <https://doi.org/10.23919/EuCAP53622.2022.9769254>
- [8] Skura, J.P. (1987). Worldwide anomalous refraction and its effect on electromagnetic wave propagation. *Johns Hopkins APL Technical Digest*, 8(4). Available: <https://secwww.jhuapl.edu/techdigest/content/techdigest/pdf/V08-N04/08-04-Skura.pdf>
- [9] Yang, J., Heo, J., Kim, J., Park, Y., Choo, H., Park, Y.B. (2023). Deep neural network-based time series atmospheric refractivity prediction model using meteorological observation data. *Journal of Korean Institute of Electromagnetic Engineering and Science*, 34(11): 860-863. <https://doi.org/10.5515/KJKIEES.2023.34.11.860>
- [10] Agbo, E.P., Ettah, E.B., Eno, E.E. (2020). The impacts of meteorological parameters on the seasonal, monthly and annual variation of radio refractivity. *International Journal of Physics*, 8(6): 198-205. <https://doi.org/10.1007/s12648-020-01711-9>
- [11] Ranran, H., Leke, L., Zhenwei, Z., Shuji, S. (2021). Regression analysis of refractivity gradient in the lowest 65 m near the ground based on relevance vector machine algorithm. *Proceedings of 2021 13th International Symposium on Antennas, Propagation and EM Theory*. <https://doi.org/10.1109/ISAPE54070.2021.9753320>
- [12] Jicha, O., Pechac, P., Kvicera, V., Grabner, M. (2013). Estimation of the radio refractivity gradient from diffraction loss measurements. *IEEE Transactions on Geoscience and Remote Sensing*, 51(1): 12-18. <https://doi.org/10.1109/TGRS.2012.2199995>
- [13] Emmanuel, I., Ojo, O.S., Abe, O.E., Adedayo, K.D. (2020). Geostatistical distribution of vertical refractivity gradient over Nigeria. *Radio Science*, 55: e2020RS007109. <https://doi.org/10.1029/2020RS007109>
- [14] Fashade, O.O., Akinwumi, S.A., Omotosho, T.V., Arijaje, T.E., Ayo-Akanbi, O.A., Ogundolie, O.I., Ometan, O.O., Adewusi, O.M. (2023). Evaluation of refractivity gradient and k-factor within the lower troposphere of Maiduguri and Enugu under two climatic zones in Nigeria. *WSEAS Transactions on Communications*, 22. <https://doi.org/10.37394/232030.2023.2.1>
- [15] Adebayo, S., Aweda, F.O., Ojedokun, I.A., Olapade, O.T. (2022). Refractive index perception and prediction of radio wave through recursive neural networks using meteorological data parameters. *International Journal of Engineering*, 35(4): 810-818. <https://doi.org/10.5829/ije.2022.35.04a.21>
- [16] Liu, X., et al. (2021). Tropospheric Refractivity Profile Estimation by GNSS Measurement at China Big-Triangle Points. *Atmosphere*, 12(11), 1468. <https://doi.org/10.3390/atmos12111468>

- [17] Beck, H.E., McVicar, T.R., Vergopolan, N., et al. (2023). High-resolution (1 km) Köppen-Geiger maps for 1901–2099 based on constrained CMIP6 projections. *Scientific Data*, 10, 724. <https://doi.org/10.1038/s41597-023-02549-6>
- [18] Pepin, N.C., Arnone, E., Gobiet, A., Haslinger, K., Kotlarski, S., Notarnicola, C., Palazzi, E., Seibert, P., Serafin, S., Schöner, W., Terzago, S., Thornton, J.M., Vuille, M., and Adler, C. (2022). Climate changes and their elevational patterns in the mountains of the world. *Reviews of Geophysics*, 60(1), e2020RG000730. <https://doi.org/10.1029/2020RG000730>
- [19] Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Thépaut, J.N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049. <https://doi.org/10.1002/qj.3803>
- [20] International Telecommunication Union. (2019). The Radio Refractive Index: Its Formula and Refractivity Data. Recommendation ITU-R P.453-14. Available: <https://www.itu.int/rec/R-REC-P.453>
- [21] Wallace, J.M., Hobbs, P.V. (2006). *Atmospheric Science: An Introductory Survey*. 2nd ed. Elsevier Academic Press. <https://doi.org/10.1016/C2009-0-00034-8>
- [22] Usman, S., Osuagwu, F., Akinbolati, A., Owolabi, A., Akusu, C.O. (2024). Exploring vertical gradients of radio refractivity and their significance for radio wave propagation in Abuja, Jos and Makurdi of North-Central Nigeria. *Science World Journal*, 19(1): 58–63. <https://doi.org/10.4314/swj.v19i1.8>
- [23] Xu, X., Zou, X. (2022). COSMIC-2 RO profile ending at PBL top with strong vertical gradient of refractivity. *Remote Sensing*, 14(9): 2189. <https://doi.org/10.3390/rs14092189>
- [24] Millán, L.F., Hoor, P., Hegglin, M.I., Manney, G.L., Boenisch, H., Jeffery, P., Kunkel, D., Petropavlovskikh, I., Ye, H., Leblanc, T., Walker, K. (2024). Exploring ozone variability in the upper troposphere and lower stratosphere using dynamical coordinates. *Atmospheric Chemistry and Physics*, 24: 7927–7962. <https://doi.org/10.5194/acp-24-7927-2024>
- [25] Wang, M., and Fu, Q. (2023). Stratosphere-troposphere exchanges of air mass and ozone concentrations from ERA5 and MERRA2: Annual-mean climatology, seasonal cycle, and interannual variability. *Journal of Geophysical Research: Atmospheres*, 128, e2023JD039270. <https://doi.org/10.1029/2023JD039270>
- [26] Holton, J.R. (2004). *An Introduction to Dynamic Meteorology*. 4th ed. Elsevier Academic Press.
- [27] IPCC. (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press. <https://doi.org/10.1017/9781009157896>
- [28] Gao, J., Brewster, K., Xue, M. (2008). Variation of radio refractivity with respect to moisture and temperature and influence on radar ray path. *Advances in Atmospheric Sciences*, 25(6): 1098–1106. <https://doi.org/10.1007/s00376-008-1098-x>
- [29] Grinsztajn, L., Oyallon, E., and Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, 35, 507–520. <https://arxiv.org/abs/2207.08815>
- [30] Alkhayyal, M.A., and Mostafa, A.M. (2025). Enhancing LoRaWAN performance using boosting machine learning algorithms under environmental variations. *Sensors*, 25(13), 4101. <https://doi.org/10.3390/s25134101>
- [31] Bochenek, B., and Ustrnul, Z. (2022). Machine learning in weather prediction and climate analyses: applications and perspectives. *Atmosphere*, 13(2), 180. <https://doi.org/10.3390/atmos13020180>
- [32] Onawumi, K.C., Adediji, A.T., Ogunjo, S.T. (2026). Comparative analysis of machine learning models for predicting refractivity gradients in signal propagation. *Current Journal of Applied Science and Technology*, 45(3): 7–20. <https://doi.org/10.9734/cjast/2026/v45i34671>
- [33] Soci, C., Hersbach, H., Simmons, A., et al. (2024). The ERA5 global reanalysis from 1940 to 2022. *Quarterly Journal of the Royal Meteorological Society*, 150(764), 4014–4048. <https://doi.org/10.1002/qj.4803>
- [34] Qiu, Z.J., Zhang, C., Wang, B., Hu, T., Zou, J., Li, Z.Q., Chen, S.Z., and Wu, S. (2023). Analysis of the accuracy of using ERA5 reanalysis data for diagnosis of evaporation ducts in the East China Sea. *Frontiers in Marine Science*, 9, 1108600. <https://doi.org/10.3389/fmars.2022.1108600>
- [35] Haji-Aghajany, S., Amerian, Y., Verhagen, S., and Rohm, W. (2024). Machine learning-based wet refractivity prediction through GNSS troposphere tomography for ensemble troposphere conditions forecasting. *IEEE Transactions on Geoscience and Remote Sensing*, 62, 4106118. <https://doi.org/10.1109/TGRS.2024.3417487>

Nomenclature

Symbol	Description
N	Atmospheric radio refractivity, N-units (dimensionless)
G	Vertical refractivity gradient, $N - \text{units km}^{-1}$
P	Atmospheric pressure, hPa
T	Absolute temperature, K
e	Water vapour pressure, hPa
H	Geopotential height, km
N_{1000}	Refractivity at 1000 hPa level, N-units
N_{850}	Refractivity at 850 hPa level, N-units
H_{1000}	Geopotential height at 1000 hPa, km
H_{850}	Geopotential height at 850 hPa, km
RMSE	Root Mean Square Error, N-units km^{-1}
MAE	Mean Absolute Error, N-units km^{-1}
R^2	Coefficient of determination, Dimensionless

Atmospheric Predictor Variables

Symbol	Description
q	Specific humidity, kg kg^{-1}
ω	Vertical velocity (pressure coordinate), Pa s^{-1}
u	Zonal wind component, m s^{-1}
v	Meridional wind component, m s^{-1}
ζ	Relative vorticity, s^{-1}
$\nabla \cdot \vec{V}$	Horizontal divergence, s^{-1}
PV	Potential vorticity, $\text{Km}^2\text{kg}^{-1} \text{s}^{-1}$
Φ	Geopotential, $\text{m}^2 \text{s}^{-2}$
C_f	Fraction of cloud cover, Dimensionless
q_l	Specific cloud liquid water content, kg kg^{-1}
q_r	Specific rain water content, kg kg^{-1}
q_s	Specific snow water content, kg kg^{-1}
q_i	Specific cloud ice water content, kg kg^{-1}
O_3	Ozone mass mixing ratio, kg kg^{-1}

Model Abbreviations

Symbol	Description
LightGBM	Light Gradient Boosting Machine
RF	Random Forest
XGBoost	Extreme Gradient Boosting
SVR	Support Vector Regression
MLP	Multi-Layer Perceptron
AdaBoost	Adaptive Boosting