

Article

Time-Dependent Feature Importance of Source Intensity and Meteorological Variables in Simulation of Air Pollutant Concentrations

Yuval ¹, Yoav Levi ² , Pavel Khain ²  and David M. Broday ^{1,*} 

¹ Faculty of Civil and Environmental Engineering, Technion, Israel Institute of Technology, Haifa 3200003, Israel; lavuy@technion.ac.il

² Israel Meteorological Service, P.O. Box 25, Bet Dagan 5025001, Israel; leviyo@ims.gov.il (Y.L.); khainp@ims.gov.il (P.K.)

* Correspondence: dbroday@technion.ac.il

Abstract

Disentangling the relative roles of emissions and atmospheric processes in controlling air-pollutant concentrations remains a central challenge in air-quality management. Although machine-learning (ML) models can accurately predict pollutant concentrations, the temporal evolution of the importance of individual drivers is often difficult to interpret. Here, we introduce a framework for reconstructing time-resolved feature importance (FI) in ML air-quality models. Hourly NO₂ and PM_{2.5} concentrations were simulated across clusters of observations, defined along concentration trajectories in a state-space spanned by source intensity and meteorological variables. Within each cluster, predictor importance is quantified and mapped back onto the corresponding time points, yielding continuous FI time series for all predictors. The framework is demonstrated using observations from the nationwide air-quality network in Israel, together with traffic-related source indicators and meteorological parameters. The dominant drivers differ markedly between the two pollutants: NO₂ variability is primarily associated with local emissions, mechanical transport, and turbulent mixing, whereas PM_{2.5} variability reflects predictors that are related to nucleation, coagulation, hygroscopic growth, long-range transport, and chemical transformation. The feature importance exhibits pronounced seasonal, regional, and diurnal variability, including modulation around traffic rush hours. These results demonstrate the value of time-resolved interpretability for diagnosing drivers of air-pollutant variability and improving the representation of processes in statistical air-quality models.

Keywords: environmental informatics; feature importance; machine learning; meteorology-air pollution relations; Path Accumulated Local Effects; state space



Academic Editor: Junwen Liu

Received: 14 June 2026

Revised: 6 July 2026

Accepted: 16 July 2026

Published: 21 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Ambient air pollution is a major environmental and public health concern, with well-documented links to a wide range of adverse health outcomes [1–3]. In urban environments, pollutant concentrations are governed by a complex interplay between emission sources and meteorological processes, including mechanical transport, turbulent mixing, and chemical transformation [4–7]. Quantifying the relative contributions of these drivers and how they vary in time remains a central challenge in air-quality research and management.

To address the limitations of sparse observations and process-based modeling, statistical and machine-learning approaches are widely used to estimate spatiotemporal pollutant

concentrations [8–14]. While such models often achieve high predictive accuracy, they typically provide limited insight into the underlying drivers of variability. Improving the interpretability of these models—particularly in a manner consistent with physical understanding—is therefore an important objective in environmental modeling.

Feature importance (FI) methods offer one pathway towards interpreting complex models by quantifying the contribution of individual predictors to the model output [15]. This approach has been applied in air-quality studies across a range of spatial scales and predictor sets, including land-use, meteorological, and satellite-derived variables [11,16–19]. However, existing applications exhibit several important limitations. First, FI is typically assessed using data at a coarse temporal resolution, thereby neglecting intra-daily variability despite its known importance. Second, most studies rely on a single global FI estimate, implicitly assuming stationarity in the relationships between predictors and pollutant concentrations. To the best of our knowledge, none of the air pollution informatics studies to date addressed the FI's temporal and spatiotemporal variability. Third, commonly used FI methods, e.g., the Mean Decrease in Impurity [20] and the simple permutation method that is implemented in the Python (version 3.14) scikit-learn package, do not adequately account for predictor inter-correlation, which is ubiquitous in environmental data and can bias importance estimates and their interpretation [21–23].

Addressing these limitations requires methodological advances that enable high temporal resolution of feature importance attribution, which is robust in the presence of correlated predictor variables. In this study, we introduce a novel framework for feature importance assessment that explicitly accounts for correlated predictors and resolves temporal variability. The approach combines state-space clustering with machine-learning-based attribution to estimate FI within dynamically defined regimes, which are subsequently mapped back onto the time domain to yield continuous, time-resolved FI series. This enables the identification of regime-dependent and non-stationary relationships between predictors and pollutant concentrations.

The framework is demonstrated using NO₂ and PM_{2.5} hourly observations from a nationwide monitoring network in Israel, together with traffic-related and meteorological predictors. By resolving the temporal evolution of predictor importance, the proposed approach provides insights into the dynamics of drivers of air-pollutant variability. More broadly, it offers a transferable methodology for interpreting complex environmental models in the presence of correlated inputs and non-stationary system behavior.

2. Materials and Methods

2.1. Meteorological Data

Meteorological predictors were derived from the limited-area configuration of the ICON (ICOsahedral Nonhydrostatic) model (ICON-LAM), operated by the Israel Meteorological Service within the framework of the Consortium for Small-Scale Modeling (COSMO; [24]). To ensure temporal consistency, a retrospective simulation was performed using a single ICON-LAM model l version for the period from 1 January 2019 to 31 December 2022. The model was initialized daily at 00:00 UTC, and hourly fields were constructed from forecasts with lead times of 6–30 h. Simulations were conducted at ~2.5 km horizontal resolution with 65 vertical levels, using boundary conditions that were updated every three hours, obtained from the ECMWF Integrated Forecasting System (IFS, Cycle 45r1 to Cycle 47r3) [25].

A subset of 11 meteorological variables was selected from a larger pool of model outputs, and additional derived variables relevant to pollutant dispersion were constructed (Table S1 in the Supplementary Material). Only data from the lowest four model levels (up to 122 m) were retained, since Yuval et al. [4] found that the lowest model levels' data were

sufficient to characterize near-surface stratification in the study area. Each hourly time point in the study period was assigned forecast data with the shortest forecast lead time. The data domain spans $34\text{--}36^\circ$ east and $29.5\text{--}33.5^\circ$ north, encompassing 13,041 grid points. This domain covers the whole of Israel and the surrounding areas of the Mediterranean, Lebanon, Syria, Jordan, the Palestinian Authority, and Egypt (Figure 1). The validity of the ICON output has been established by numerous past studies (e.g., [26,27]). We tested the surface ICON data at the study area against observations from 147 automatic meteorological stations throughout Israel. The model's upper atmosphere output was evaluated using twice-daily radiosonde data with 10 m vertical spacing launched at the IMS headquarters in Bet Dagan, central Israel.




Figure 1. Study area map, ICON grid points, and the locations of the air quality monitoring stations, color-coded by region. Map lines depict the study area and do not necessarily delineate accepted national boundaries.

2.2. Air Pollution Data

Raw half-hourly air pollutant observations were obtained from the Israeli Ministry of Environmental Protection website, from which we extracted the concentrations of NO_2 and $\text{PM}_{2.5}$. The data underwent a quality-assurance and quality-control process using a suite of dedicated computer codes developed in-house, including outlier identification, temporal statistics, consistency checks based on spatiotemporal correlations of observation time series, etc. (for more details, see Section S3 in the Supplementary Material). Data

identified as clearly erroneous was removed from the database. Only data series with more than 90% completeness during the study period were retained. The machine-learning and FI methods we employed do not require strictly regularly spaced data series, so missing values were not filled to avoid imputation errors.

Locations of the air quality monitoring (AQM) stations are shown in Figure 1. Figure S1 specifies the pollutants observed by each of the AQM stations. The dataset includes 85 AQM stations. Seven of the stations are located within train and bus stations and are not fully exposed to outdoor road-traffic emissions and meteorological conditions. Therefore, their data were used only for preliminary qualitative comparisons with the fully outdoor stations.

2.3. Traffic Volume Data

The predominant source of surface-level primary pollutant concentrations in Israel is road traffic [28]. Road-traffic volumes were used in this study for estimating combustion-related traffic emissions of NO_x and PM_{2.5}. The traffic volumes were obtained from three sources and were combined into a single spatiotemporal varying emissions proxy.

The first source is the Israeli National Traffic Model, which calculates 30 traffic attributes for five daytime periods during working days across 17,207 segments of the major roads throughout Israel. It uses a traffic-demand matrix, based on 2018–2019 cellular-location data from the two leading Israeli mobile-phone providers. The model reports traffic volumes on working days at the following time periods: 06:00–09:00, 09:00–12:00, 12:00–15:00, 15:00–18:00, and 18:00–21:00. We used only the private-car and truck volume attributes and the locations and lengths of the road segments. Bus volumes were not available and were assumed to be proportional to car volumes.

The second source comprises traffic volumes of private cars, trucks, and buses derived from vehicle-fleet GPS signals from 2012, purchased from Decell Inc. (Tel Aviv, Israel). These data were provided on a 250 × 250 m² spatial grid covering all the AQM-station locations, for 11 daily time periods, and for working days, Fridays, and Saturdays, separately. Bus volumes in the Decell's dataset were found in an earlier study to be unreliable and were therefore not used [29].

The third source contained routine direct traffic-volume observations collected by the Israeli Central Bureau of Statistics during 2016–2019, as well as observations made by private operators of toll roads. These observations were made using the pneumatic-tube method, which counts the number of vehicle axles passing a segment per hour. The data are therefore not vehicle-type specific and are reported as private-car equivalents (PCE; axle count divided by two). Due to their limited spatiotemporal nature, these records were used for validating the combined National Traffic Model and Decell data traffic volume estimates. Specifically, the traffic volumes we used were synthesized from the three traffic volume sources as follows. Decell's data served as the basic traffic volumes due to their higher temporal resolution. They were modified by the more recent National Traffic Model data and then corrected by regressing them on the direct traffic volume observations. A detailed assessment of the three datasets is given in Section S3 of the Supplementary Material.

2.4. Initial Data Processing

2.4.1. Spatial and Temporal Grid Adjustments

All the AQM station locations were transformed to the WGS84 coordinate system. Each AQM station location was then associated with the meteorological time series calculated at its nearest ICON grid point. While this choice may not be optimal in every instance, given the comparatively fine spatial resolution of the model grid (Figure 1) and the fact that most AQM stations are situated in areas of modest topographical variability, it is likely to be

appropriate in most cases. Hourly means of the raw half-hourly air-quality observations were taken to match the ICON model's hourly temporal resolution.

2.4.2. Traffic Emission Impact Indicators

Since the National Traffic Model and Decell data correspond to specific time periods of the day, they were first interpolated to provide 24 h daily cycles. The hourly daily cycles were expanded to time series by assigning each hour in the study period the appropriate traffic volume data from the daily cycles.

The major road-network segments in Israel vary in length from about 40 m to more than 19 km. We divided all segments into 100 m-long sub-segments (with shorter segments left unchanged) and assigned them to their midpoint proxies for traffic emissions derived from the synthesized traffic-volume data. The traffic attribute used as an emission proxy is vehicle-kilometer traveled (VKT) per hour [veh km h^{-1}] — the total distance that vehicles (of each type) travel within each road sub-segment per hour in both directions.

The direct impact of a pollution source on a given location is negligible unless the airflow transports the emissions towards it. We therefore quantified the impact of traffic on the concentrations observed at the AQM stations, accounting for the local wind direction. Two measures were calculated for each vehicle type, for each observation location (i.e., AQM station), at each hour of the study period. (1) The VKT at the road sub-segment closest to the pollutant observation location and lying within 15° of the upwind ICON wind direction. This measure represents the local emissions impact. While the choice of $\pm 15^\circ$ is rather arbitrary, it is based on our domain knowledge and is physically sound. Using other angles between $\pm 10^\circ$ and $\pm 30^\circ$ did not significantly change the accumulated VKT that impacted each grid cell under varied scenarios. (2) The total VKT at all road sub-segments, which, given the local ICON wind speed and direction, could influence the concentrations at the observation location during that hour. This measure represents the impact of background emissions. It was calculated as the sum of the VKT at all sub-segments within a sector of $\pm 15^\circ$ from the local ICON wind direction and within a radius equal to the distance that the local airflow could travel in one hour, given the local ICON wind speed.

2.5. Data Space Partition Using Cluster Analysis

The traffic and meteorological variables associated with each AQM station location formed a state space—a space spanned by variables. Following Yuval et al. [30], the state space of each AQM station was partitioned using cluster analysis, whereby each point on the emission–dispersion curve was assigned to a specific cluster, with the clustering performed independently for each AQM station. Consequently, both the number of clusters and their sizes vary between stations. The predictor variables used for this partitioning are listed in Table S1. Prior to clustering, the meteorological predictors' time series were detrended, i.e., their daily and seasonal cycles were removed, and the variables were standardized by subtracting the mean and dividing by the standard deviation. Partitioning was carried out using the *K*-means algorithm — an unsupervised machine-learning method that assigns data points to clusters such that the total sum of squared distances between points and their respective centroids is minimized.

K-means is efficient and well-suited to large datasets [31], but its implementation faces two main challenges. First, it is required to determine, a priori, the optimal number of clusters, *K*. Second, the final clustering is sensitive to the initial choice of centroids. These issues are interrelated, since stable clustering for a given number of clusters, *K*, is required for determining the optimal *K* reliably. The methodology we applied to address these challenges and to carry out the partitioning is described in detail in Yuval et al. [30].

2.6. Minimizing Intercorrelation Between Predictor Variables

Many of the predictor variables are highly correlated (Figure S2a), which is known to impair model training by destabilizing parameter estimation and promoting overfitting to shared noise. Correlation might affect FI inference even more severely, as it can obscure the attribution of predictive contributions to individual variables. Our choice of primary FI method was designed to mitigate these effects (Section 2.8). However, reducing predictor correlation remains preferable whenever possible. We therefore retained, for both the ML modeling and the FI inference, only a subset of predictors that captures the dominant dispersion mechanisms while avoiding the strongest correlations: 10 m vertical velocity (W_v), 2 m temperature (T), 2 m relative humidity (RH), temperature lapse rate (Γ), 10 m horizontal wind speed (W_h), and the background and local car indicators (Car_{BG} and Car_{Loc} , respectively). Figure S2b demonstrates a clear reduction in inter-correlation, from the range $[-0.51, 0.93]$ for the full predictor set to $[-0.4, 0.57]$ within this selected subset.

2.7. Feed-Forward Neural Network

A two-layer feed-forward neural network (FFNN) was employed as the machine-learning model to simulate pollutant observations within the data clusters. The FFNN models were trained separately for each cluster at each AQM station location using a four-fold cross-testing procedure to assess the modeling performance. An additional internal cross-validation step was performed within each training fold to optimize the model hyperparameters.

Machine learning modeling of cluster data can be generally expressed by

$$\hat{Y}_i = \mathcal{F}(X_{ij}, \mathbf{P}), \quad i = 1, 2, \dots, N, \quad j = 1, 2, \dots, M \quad (1)$$

where $\hat{Y}_i$ simulates the true predictand value Y_i using $\mathcal{F}$, M is the number of predictor variables, and N is the number of data points in the cluster. Each row of the matrix X_{ij} is denoted as X_i and forms a data point in the M -dimensional state space defined by the M predictor variables. Each of the columns of X_{ij} is a predictor data series referred to as X_j . The predictor variables used for the modeling and subsequently for the FI analysis are noted in Table S1. In Equation (1), $\mathbf{P}$ is a vector of model parameters that can be determined by minimizing the cost function.

$$\Phi = \sum_{i=1}^N |Y_i - \hat{Y}_i(X_{ij}, \mathbf{P})|^2 \quad (2)$$

The process of determining the parameters $\mathbf{P}$ is termed in the machine learning jargon as model training. Assessment of the model performance evaluates the trained model using data not used in the training. We refer to this process as model testing. We used a four-fold cross-testing procedure, which is described in detail in the Supplementary Material.

The two-layer FFNN is formulated as

$$\hat{Y}_i = \sum_{k=1}^H w_k^{(2)} \cdot \mathcal{T} \left(\sum_{j=1}^M w_{kj}^{(1)} X_{ij} + b_k^{(1)} \right) + b^{(2)}, \quad i = 1, 2, \dots, N \quad (3)$$

where $w^{(1)} \in \mathbb{R}^{H \times M}$, $w^{(2)} \in \mathbb{R}^{H \times 1}$, $b^{(1)} \in \mathbb{R}^{1 \times H}$, and $b^{(2)}$ contain the parameters to be determined by minimizing the cost function in Equation (2). Here, $\mathcal{T}$ is a nonlinear activation function, which renders the FFNN a nonlinear mapping of the input X_{ij} . The value of H is usually referred to as the number of hidden neurons of the neural network. Larger H corresponds to a more complex model but also to a higher number of model parameters. The latter poses stricter requirements on the training to avoid model overfitting

to the observations. With a proper choice of $\mathcal{T}$ and sufficiently large H , the FFNN model can be made highly flexible. In the following, we use the tansig (hyperbolic tangent sigmoid) activation function

$$\mathcal{T}(x) = \tan \operatorname{sig}(x) = \frac{2}{1 + e^{-2x}} - 1 \quad (4)$$

and $H = 5$. Larger networks would substantially increase the number of free parameters and the risk of overfitting without a clear benefit for the objective of this study, which is an interpretable regression model, not necessarily the best predicting one. The tansig function has smooth, bounded derivatives, which enable numerical stability in the cost function minimization process. Its output values are in the range $[-1, 1]$, which naturally centers the transformed data around zero, often leading to faster convergence when using gradient-based optimization for relatively small networks like in our case. With $H = 5$ and $M = 9$ predictors, the number of free parameters was 56. This number of model parameters and the tansig activation yield sufficient flexibility to capture complex nonlinear relationships while minimizing the risk of overfitting the data in the small data clusters. The FFNN training, validation, and testing were carried out using Matlab built-in functions [32] and additional Matlab software (Version R2025b) developed in-house.

2.8. Feature Importance in Machine Learning Regression Models

The goal of FI analysis is to assign each predictor a quantitative measure of its contribution to the model output. A wide range of methods has been developed to interpret machine-learning models by quantifying the contributions of individual predictors (i.e., features in the ML jargon). Among the most widely used methods are the SHapley Additive exPlanations (SHAP; [33]), which attributes feature importance using concepts from cooperative game theory, and Integrated Gradients [34], a gradient-based attribution method that was originally developed for deep neural networks. Here, we employed the recently introduced Path Accumulated Local Effects method (PALE; [23], which was specifically designed to perform robustly in the presence of correlated predictors. For comparison, we also applied the established permutation-based FI approach with conditional permutation sampling, which relies on a different conceptual framework and computational procedure [21,22]. Using two complementary FI methods allows us to assess the robustness of the inferred predictor contributions. All algorithms were implemented in-house using Matlab [32].

2.8.1. Path Accumulated Local Effects Feature Importance

The Accumulated Local Effects (ALE) method was introduced by Apley and Zhu [35] as a visual tool for interpreting machine learning models. Zhu and Apley [23] subsequently extended the ALE framework to define Path Accumulated Local Effects (PALE) as a measure of predictor importance. PALE quantifies both the direct contribution of a predictor, termed the main effect (PALE-main), and its total effect (PALE-total), which incorporates contributions arising from interactions with other predictors. The fundamental idea underlying ALE is to quantify the average impact of a predictor on the model output within narrow intervals of the predictor's domain (local effects). The remaining predictors are implicitly sampled from their conditional distribution given the predictor of interest. As a result, ALE avoids extrapolation into regions of the predictor space unsupported by the data and is robust in the presence of correlated predictors [35]. Owing to this ability to mitigate the adverse effects of linear dependence between predictors, PALE was adopted as the primary measure of feature importance in this study.

The PALE-main effect can be obtained directly from the ALE function. For a given cluster, the trained model is evaluated twice using only the data points that fall within each interval of the value range of the examined predictor: (1) with all values of the

examined predictor X_j set to the lower interval boundary $x_j^{(l)}$, and (2) with all values of X_j set to the upper interval boundary, $x_j^{(u)}$. The local effects of predictor X_j are defined as the differences in predicted values between the two model evaluations while keeping all other predictors unchanged. Specifically, for an observation i belonging to bin k within the interval $[x_{j,k}^{(l)}, x_{j,k}^{(u)}]$, the local effect is

$$\Delta_{j,k}^{(i)} = \mathcal{F}(X_1^{(i)}, \dots, X_{j-1}^{(i)}, x_{j,k}^{(u)}, x_{j+1}^{(i)}, \dots, X_M^{(i)}) - \mathcal{F}(X_1^{(i)}, \dots, X_{j-1}^{(i)}, x_{j,k}^{(l)}, x_{j+1}^{(i)}, \dots, X_M^{(i)}) \quad (5)$$

where $\mathcal{F}$ denotes the trained model. The mean local effect in the k th bin is

$$\text{LE}_{j,k}([x_{j,k}^{(l)}, x_{j,k}^{(u)}]) = \frac{1}{Q} \sum_{i=1}^Q \Delta_{j,k}^{(i)} \quad (6)$$

where Q is its number of data points in the interval $[x_{j,k}^{(l)}, x_{j,k}^{(u)}]$.

The range of predictor X_j is partitioned into K consecutive, non-overlapping intervals $[x_{j,1}^{(l)}, x_{j,1}^{(u)}], \dots, [x_{j,K}^{(l)}, x_{j,K}^{(u)}]$, ordered such that $x_{j,1}^{(l)} < x_{j,1}^{(u)} = x_{j,2}^{(l)} < \dots < x_{j,K}^{(u)}$. Suppose a given value x_j belongs to the k th interval, $x_j \in [x_{j,k}^{(l)}, x_{j,k}^{(u)}]$ for some $k \in \{1, \dots, K\}$. The Accumulated Local Effects function for predictor X_j is then defined as the integration along the mean path of the local effects, which for the discrete intervals is given by

$$\text{ALE}_j(x_j) = \sum_{i=1}^{K-1} \text{LE}_{j,i}([x_{j,i}^{(l)}, x_{j,i}^{(u)}]), \quad (7)$$

with the convention that $\text{ALE}_j(x_j) = 0$ if x_j falls in the first interval ($k = 1$).

The Accumulated Local Effects function, $\text{ALE}_j(x_j)$, is centered by subtracting its mean over the observed data distribution of X_j ,

$$\tilde{\text{ALE}}_j(x_j) = \text{ALE}_j(x_j) - \frac{1}{N} \sum_{n=1}^N \text{ALE}_j(x_j^n), \quad (8)$$

where N is the total number of data points in the data cluster and x_j^n denotes the n th value of predictor X_j . The ALE functions are centered to ensure that the resulting feature importance reflects only fluctuations around the average effect, independent of arbitrary offsets introduced by discretization or finite-sample effects. Examples of centered ALE functions are shown in Figure S3 for the NO_2 data of cluster 25 of AQM station KV4, which is located in the GUD region of central Israel (Figure 1).

The PALE-main feature importance for X_j is defined as the variance of the centered ALE function,

$$\tilde{\text{FI}}_j^{\text{main}} = \text{Var}[\tilde{\text{ALE}}_j]. \quad (9)$$

This quantity measures the contribution to model output explained by the X_j predictor in isolation. To compare FIs across different pollutants and clusters, we considered the relative feature importance of predictor X_j , which is computed as

$$\text{FI}_j^{\text{main}} = \frac{\tilde{\text{FI}}_j^{\text{main}}}{\sum_{k=1}^M \tilde{\text{FI}}_k^{\text{main}}}, \quad (10)$$

where $\text{FI}_j^{\text{main}}$ is the relative PALE-main, considering only the main effects.

In addition to the main effect, the PALE framework also defines a total-effect feature importance, which captures both the direct contribution of predictor X_j and its contribution through interactions with all other predictors. The total feature importance, $\tilde{\text{FI}}_j^{\text{tot}}$, is obtained by randomly sampling one value from the set $\Delta_{j,k}^{(i)}$, $i = 1, 2, \dots, Q$ in each interval and constructing from these values a path along which the accumulated sum ALE function is computed. Repeating this process multiple times and averaging the variances of the path-specific ALEs yields the PALE-total $\tilde{\text{FI}}_j^{\text{tot}}$. The number of repetitions, B , should be large enough for the results to be stable. We found that $B = 100$ was usually sufficient and used $B = 200$ in all cases. A rigorous description of the PALE-total procedure is provided in the Supplementary Material. The relative PALE-total value is computed analogously to Equation (10). Note that, considering interactions, the relative PALE-main due to predictor X_j should be

$$\text{FI}_j^{\text{main,int}} = \frac{\tilde{\text{FI}}_j^{\text{main}}}{\sum_{k=1}^M \tilde{\text{FI}}_k^{\text{tot}}}, \quad (11)$$

and the corresponding relative importance of the interactions is thus $\text{FI}_j^{\text{tot}} - \text{FI}_j^{\text{main,int}}$.

The PALE-main and PALE-total results of this study closely matched in most instances. Consequently, the analysis and discussion of this work focus on the main-effect feature importance. The interaction effects are not negligible, though (Figure S4 and Section S7 in the Supplementary Material), but their thorough study and impact are deferred to future studies.

2.8.2. Permutation Feature Importance

Permutation feature importance quantifies the contribution of a predictor variable by measuring the increase in the model's prediction error after the information contained in that variable is removed. This is achieved by permuting the predictor's values, thereby breaking their association with the predictand data points. We define the prediction error as the sum of the absolute differences as

$$\varepsilon = \sum_{i=1}^N |Y_i - \hat{Y}_i|. \quad (12)$$

Alternative measures, such as the root mean square error (RMSE) in Equation (2), can also be employed. However, the mean absolute error is a more robust measure of predictive performance since it is less sensitive to outliers compared to the RMSE. This is especially important to prevent bias in the modeling of data in small clusters.

The baseline prediction error is obtained by evaluating the model without permuting any predictors. Subsequently, each predictor is permuted in turn, the model is re-evaluated, and a new error is computed. The difference between the baseline error and the permuted error defines the feature importance of predictor X_j ,

$$\tilde{\text{FI}}_j^{\text{perm}} = \varepsilon(Y_i, X_{ij}) - \varepsilon(Y_i, X_{ij}^{(j,\pi)}), \quad (13)$$

where $X_{ij}^{(j,\pi)}$ denotes the set of predictors with the j th predictor permuted. After repeating this procedure for all the M predictors, the relative importance $\text{FI}_j^{\text{perm}}$ of predictor X_j can be computed analogously to Equation (10).

In the presence of inter-variable dependencies (e.g., collinearity among predictors), random shuffling of a predictor's marginal distribution, $\mathbb{P}(X_j)$, may generate unrealistic value combinations. For instance, associating zero nighttime global solar radiation with summer

mid-day temperatures or with strongly positive temperature lapse rates. Such combinations cause the model to extrapolate beyond the natural predictand data domain [21–23]. Moreover, marginal distribution shuffling eliminates information shared with other predictors, which can lead to overestimation of the predictor’s importance [23]. To mitigate this, the permutation should sample from the conditional distribution $\mathbb{P}(X_j|X_{-j})$, where X_{-j} denotes the set of all predictors except X_j [21]. Our conditional permutation procedure, described in the Supplementary Material, prevents out-of-domain extrapolation [21,22]. An alternative approach for handling multicollinearity is the multi-pass permutation FI method [36]. However, we found the conditional permutation strategy to be more intuitive and computationally efficient.

3. Results

The clustering procedure has been applied to the 85 AQM stations’ pollutant records [29], yielding between 21 and 33 clusters per AQM station (see Figure S5) with each cluster containing at least 200 data points. Smaller clusters were excluded, as they did not contain a sufficient number of observations to permit reliable model training and performance evaluation. Typically, the excluded clusters accounted for less than 1.5% of the total data points at each AQM station. These small clusters were mostly associated with downdraught, high cloud cover, low solar radiation, high relative humidity, and precipitation, representing scarce precipitation events in the study area. To illustrate the characteristics of the retained clusters, Table S2 presents the median predictor values for the 28 clusters extracted from the emission–dispersion (i.e., meteorological and emissions proxy) data at AQM station KV4, which is located in close proximity to one of the major longitudinal highways in the densely urbanized GUD region of central Israel (Figure 1). The corresponding median NO₂ and PM_{2.5} concentrations observed at this station (and for which the clustering process was blind) are also reported. Each cluster represents a distinct emission-dispersion scenario, resulting in substantial variability in pollutant concentrations, with more than sevenfold and threefold differences in the median cluster concentrations of NO₂ and PM_{2.5}, respectively.

The PALE and permutation FI methods were implemented using the AQM station and cluster-specific, two-layer, feed-forward neural network (FFNN) model simulations of the concentrations of the two pollutants across all the clusters and all the AQM stations. Figure 2 and Figure S6 synthesize the results from all the AQM stations by presenting summary statistics of the weighted mean of the relative PALE and permutation FIs for each predictor, respectively. The weighted mean, $^{WF}FI_j$, of the relative FI of the j th predictor was computed using the cluster sample sizes as weights,

$$^{WF}FI_j = \frac{1}{\sum_{k=1}^L N_k} \sum_{k=1}^L N_k \cdot FI_j^k \quad (14)$$

where FI_j^k is the relative FI (either PALE or permutation) of the j th predictor in the k th cluster of the AQM station, N_k is its number of data points, and L denotes the total number of clusters.

The statistics shown for $^{WF}FI^{\text{main}}$ in Figure 2 and the corresponding statistics for $^{WF}FI^{\text{perm}}$ (Figure S6) are very similar, with nearly identical relative magnitudes of the medians and widths of the interquartile ranges (IQRs). Figure 2 and Figure S6 indicate that the most dominant factors in determining NO₂ levels are the local traffic indicator Car_{Loc} , which represents local emissions, and meteorological variables that promote pollutant dispersion through mechanical transport and turbulent mixing, such as Γ , W_h , and W_v . A different pattern is observed for the FIs of PM_{2.5}, for which T exhibits by far the largest

impact, followed by RH and Γ . The influence of T and RH on $PM_{2.5}$ concentrations arises from physical processes, such as nucleation, coagulation, and hygroscopic growth, as well as from chemical reactions that result in the formation of secondary organic aerosols (SOA) and secondary inorganic aerosols (SIAs). Regarding $PM_{2.5}$, the impacts of both traffic emission indicators (Car_{BG} and Car_{Loc}) are smaller than those of any of the meteorological variables, suggesting that ambient $PM_{2.5}$ in the study area is predominantly driven by long-range transport (LRT) rather than local traffic emissions. This result is in general agreement with the findings of Belachsen and Broday [19,37] that the contribution of traffic-related particulate emissions is smaller than that of long-range transport of dust and secondary aerosols (SOA and SIA) — about 46% vs. 54%, respectively, in the summer months (JJA) and about 40% vs. 46%, respectively, in the non-summer months.




Figure 2. Boxplots of the weighted mean relative feature importance $^{WF}FI^{main}$ of PALE-main per predictor variable for NO_2 and $PM_{2.5}$.

Figure 3 shows the median and interquartile range (IQR) of the nighttime temporal variability of FI^{main} , calculated from the FFNN simulations of NO_2 and $PM_{2.5}$ in the MHD region along Israel's southern coast (AQM stations marked by red circles in Figure 1). Figure 4 presents the corresponding plots for daytime conditions. The analogous figures for FI^{perm} (Figures S7 and S8) convey similar information. The temporal curves used to generate these figures were obtained by assigning each cluster's FI values to the cluster's data points corresponding to observation times. The curves were subsequently smoothed using a running-mean operator with a one-month window. The MHD region encompasses a mixture of urban and rural environments and hosts several power plants and other large industrial facilities. It also contains the largest regional AQM network in the study area, with 21 NO_2 and 10 $PM_{2.5}$ AQM stations (Figure 1 and Figure S1). Nighttime periods at each location were defined as those with zero global solar radiation, and all remaining time points were classified as daytime.

Some of the FIs shown in Figures 3 and 4 exhibit pronounced seasonal variability. In the case of NO_2 , particularly large seasonal variations are observed in the nighttime impacts of the relative humidity, RH , and the temperature lapse rate, Γ (Figure 3 and Table 1). For $PM_{2.5}$, the largest seasonal variability is seen in the daytime impact of background traffic, Car_{BG} (Figure 4 and Table 1). The fluctuations of the nighttime FI curves are, in all cases, more pronounced for NO_2 than for $PM_{2.5}$. However, during daytime, there is no clear pollutant-dependent trend (Table 1).




Figure 3. Nighttime temporal variation in the median and interquartile range of FI^{main} for NO_2 and $PM_{2.5}$ in the region MHD.




Figure 4. Daytime temporal variation in the median and interquartile range of FI^{main} for NO_2 and $PM_{2.5}$ in the region MHD.

Table 1. Coefficient of variation (the ratio of standard deviation to the mean, %) of the median curves in Figure 3 (nighttime) and Figure 4 (daytime).

Predictor	NO ₂ , Night	NO ₂ , Day	PM _{2.5} , Night	PM _{2.5} , Day
Γ	24	7	14	12
W_h	17	16	11	10
W_v	10	9	8	17
T	13	6	8	9
RH	26	11	19	5
Car_{BG}	19	10	14	21
Car_{Loc}	15	9	8	13

Nighttime and daytime dynamics differ in several cases. For example, in the winter, FI^{main} for NO₂ is associated with wind speed (W_h) peaks at nighttime, whereas in the summer, the corresponding FI^{main} for NO₂ is associated with wind speed (W_h) peaks at daytime. The opposite behavior is observed for RH , whose impact on NO₂ peaks during nighttime in the summer and during daytime in the winter. Nonetheless, in other cases, the impact curves for both NO₂ and PM_{2.5} exhibit similar seasonal variability. For example, the nighttime impact of RH exhibits clear seasonal variability, peaking in summer for both pollutants. On the other hand, the nighttime impact of T and the daytime impact of W_v for the two pollutants are clearly in opposing phases.

Differences in the FI magnitudes between the two pollutants, previously noted in Figure 2, are also apparent in Figures 3 and 4. For instance, the FIs of T and RH for PM_{2.5} are substantially larger throughout the year and exhibit relatively wider IQRs, indicating greater inter-AQM station variability. By contrast, the mean impacts of Γ , W_h , W_v , and Car_{Loc} are larger for NO₂, although the differences range from very large to almost negligible, depending on the season.

Many of the temporal trends observed in Figures 3 and 4 for the MHD AQM stations are also evident in other regions; however, clear regional differences are apparent. Figures 5 and 6 show, respectively, the medians of the nighttime and daytime temporal variability of FI^{main} , calculated from the FFNN simulations of NO₂ in the SHC, GUD, and MHD regions. The colors of the FI curves correspond to those of the AQM station markers for each region in Figure 1. Additional smoothing was applied to the curves to enhance the visibility of regional differences. Figures S9 and S10 correspond to Figures 5 and 6, respectively, and present similar results for FI^{perm} .

In most cases, similar seasonal variability is observed across all regions, albeit with differing magnitudes. For example, nighttime variability of the Γ and W_v FIs is evident in SHC, but its magnitude is markedly smaller than in the other two regions. By contrast, the magnitude of the W_h FI in SHC is more pronounced. The nighttime FI^{main} of T and RH exhibit clear seasonal cycles, peaking in summer in all regions. The corresponding daytime variability, however, is complex and differs between regions. The importance of the source indicators Car_{BG} and Car_{Loc} also varies seasonally, with substantially larger magnitudes and greater variability in SHC. Interestingly, the impacts of the traffic source proxies in SHC are clearly out of phase — the importance of Car_{BG} peaks in winter during nighttime and in summer during the daytime — with exactly opposite impacts for Car_{Loc} . Much weaker variability in the impacts of the traffic indicators is observed in the GUD and MHD regions.




Figure 5. Nighttime temporal variation in the median FI^{main} for NO_2 in regions SHC, GUD, and MHD (for geographic reference, see Figure 1).




Figure 6. Daytime temporal variation in the median FI^{main} for NO_2 in regions SHC, GUD, and MHD (for geographic reference, see Figure 1).

An additional aspect of the temporal variability of the FIs is their daily cycles. Figure 7 presents the median daily cycles of FI^{main} and their IQRs for both NO_2 and $PM_{2.5}$ in the

MHD region. The corresponding plots for FI^{perm} are shown in Figure S11. The fluctuations of the daily cycles are similar in magnitude to those of the seasonal cycles (Tables 1 and 2). The impacts of the lapse rate, Γ , and the wind speed, W_h , on NO_2 appear to change abruptly during the morning and evening rush hours, but they change in opposite directions: the impact of Γ is elevated during rush hours, whereas that of W_h is reduced. Other variables also appear to exhibit changes in their impact during traffic rush hours, but these effects are minor. The impact of the temperature, T , on $PM_{2.5}$ concentrations is larger during the daytime, while the corresponding impact of relative humidity, RH , is slightly larger at night. The impact of the local traffic indicator, Car_{Loc} , is consistently larger for NO_2 than for $PM_{2.5}$ throughout the day, although no pronounced daily cycle is observed.




Figure 7. Daily cycle of the median and interquartile range of FI^{main} for NO_2 and $PM_{2.5}$ in the region MHD.

Table 2. Coefficient of variation (the ratio of standard deviation to the mean, %) of the median curves in Figure 7.

Predictor	NO_2	$PM_{2.5}$
Γ	22	12
W_h	14	8
W_v	14	13
T	16	23
RH	16	16
Car_{BG}	12	12
Car_{Loc}	10	10

4. Discussion

This study introduces a framework for time-resolved feature importance (FI) assessment that is robust to correlated predictors and applicable to complex environmental systems. By integrating state-space clustering with machine-learning-based attribution, the approach enables the estimation of continuous FI series that capture non-stationary and regime-dependent relationships between predictors (i.e., features) and pollutant concentrations. The application to NO_2 and $\text{PM}_{2.5}$ across Israel demonstrates both the methodological capabilities of the framework and its potential to yield physically interpretable insights. More broadly, the results highlight the value of combining regime identification with model interpretation, which enables better uncovering of relationships that would otherwise be obscured by temporal averaging.

Our main FI tool, the PALE method, quantifies the effect of a predictor on the model output, whereas the permutation-based approach assesses its impact on predictive performance under predictor perturbation. The consistency observed between PALE and permutation-based FIs, despite their distinct definitions of feature importance, strengthens confidence in the robustness of the results. This finding is consistent with Zhu and Apley [23], who showed that the two approaches converge under moderate predictor inter-correlation.

The results reveal clear contrasts in the dominant drivers of the two pollutants. NO_2 variability is strongly influenced by local traffic emissions, as reflected in the pronounced importance of the local traffic indicator, Car_{Loc} , whereas the regional traffic indicator, Car_{BG} , plays only a minor role. Meteorological influences on NO_2 are primarily associated with atmospheric mixing and transport, as indicated by the importance of the temperature lapse rate, Γ , and the horizontal and vertical wind components, W_h and W_v , respectively. In contrast, $\text{PM}_{2.5}$ concentrations are governed predominantly by thermodynamic and microphysical processes. The strong influence of the temperature, T , and relative humidity, RH , reflects the importance of nucleation, coagulation, hygroscopic growth, and secondary aerosol formation. The comparatively limited contribution of traffic-related predictors suggests that in our study area, the $\text{PM}_{2.5}$ variability is shaped by processes operating over broader spatial and temporal scales than those affecting NO_2 , i.e., long-range transport (LRT).

The prominent roles observed for certain meteorological variables are consistent with previous findings [4,5,38]. Changes in meteorological conditions associated with climate change may therefore influence future pollutant concentrations. In particular, the strong dependence of $\text{PM}_{2.5}$ on temperature suggests that the pronounced regional warming trend [39,40] could alter $\text{PM}_{2.5}$ levels, but the direction of this response remains uncertain and warrants further investigation.

Air pollution processes are known to involve interactions between drivers. For example, the effect of relative humidity on ambient $\text{PM}_{2.5}$ concentrations depends on the temperature. The PALE framework is, in principle, capable of detecting such interactions. However, the present analysis indicates that interaction effects are generally small. This may reflect limited intra-cluster variability, which constrains the identification of interaction terms, or the tendency of feed-forward neural networks (particularly with correlated inputs) to approximate interaction effects through nonlinear main effects. The Supplementary Material provides an example in which synthetic correlated predictor data, simulating air pollution formation and dispersion, were designed to interact in producing a pollutant concentration series. However, the interaction contributions remain comparatively small despite being explicitly imposed. Further investigation is needed to better understand the representation of interactions within this framework.

This study has several limitations: (a) Apart from traffic-related emissions, other local emission sources were not accounted for (e.g., seaports, industry, biomass burning, domestic heating, etc.). Since, in general, such sources are either relatively small or have mostly local impact, their contributions are expected to be mostly discerned at AQM stations that are located in their vicinity. Nonetheless, not accounting for them might have biased the results, particularly for PM_{2.5}. (b) Our analysis relies on model-derived meteorological data. While previous studies [6,7,30,38] support the suitability of model-derived meteorological data for different applications, associated uncertainties remain. (c) It is noteworthy that feature importance reflects the internal structure of the statistical model and does not establish causality. The present analysis is therefore exploratory, and the inferred mechanisms should be regarded as initial results that require further validation using independent data and complementary modeling approaches, such as ensemble or hybrid methods. (d) The present study employs a specific machine-learning model, namely a feed-forward neural network (FFNN). The extent to which the FI variation in any given variable is model-dependent was not investigated. However, in a previous study, we have shown that model performance slightly changed when the same set of predictors and dependent variables was input to either Random Forest (RF) or XGBoost regression models [41].

5. Conclusions

The framework proposed in this work offers a transferable methodology for analyzing complex environmental systems that are characterized by correlated inputs and non-stationary behavior. Specifically, the framework involves the reconstruction of time-resolved feature importance (FI) of ML regression models, developed independently for distinct clusters of ambient pollutant concentration observations. These clusters were obtained using a stabilized and generalized *K*-means partitioning algorithm that operated on a state space spanned by source intensity proxies and meteorological variables but was blind to all other information, in particular to the pollutant concentrations, date, and time. By resolving pollutant-specific feature importance in both time and space and reconstructing the complete time series of the resolved feature importance, our framework provides a more nuanced view of pollutant dynamics than conventional approaches. Using records from the nationwide air-quality network in Israel, we found that the dominant drivers of the observed NO₂ and PM_{2.5} ambient concentrations differ markedly, with NO₂ variability affected mainly by the spatial distribution and intensity of local emissions and mechanical mixing (turbulence), while PM_{2.5} variability was found to be affected mainly by particle growth and formation mechanisms. Consequently, the feature importance was found to exhibit pronounced seasonal, regional, and diurnal variation. These results demonstrate the value of identifying time-resolved drivers of air-pollutant variability, thus improving our understanding of atmospheric processes through AI- and ML-assisted air-quality models' interpretability. Future work should focus on testing and extending the framework proposed here to incorporate additional data sources and modeling strategies, with the aim of strengthening the link between data-driven interpretation and process-based understanding in air-quality research.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/atmos17070704/s1>.

Author Contributions: Conceptualization, Y. and D.M.B.; data curation, Y., Y.L. and P.K.; formal analysis, Y.; funding acquisition, D.M.B.; investigation, Y.; methodology, Y.; project administration, D.M.B.; resources, D.M.B.; software, Y., Y.L. and P.K.; supervision, D.M.B.; validation, Y.; visualization,

Y.; writing—original draft, Y.; writing—review and editing, Y., Y.L., P.K. and D.M.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by an Israel Science Foundation (ISF) research grant No. 1878/23 and an Israel Ministry of Environmental Protection (IMoEP) research grant No. 226-1-2.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Phalen, R.F.; Phalen, R.N. *Introduction to Air Pollution Science: A Public Health Perspective*, 1st ed.; Jones & Bartlett Learning: Burlington, ON, Canada, 2011.
- Cohen, A.J.; Brauer, M.; Burnett, R.; Anderson, H.R.; Frostad, J.; Estep, K.; Balakrishnan, K.; Brunekreef, B.; Dandona, L.; Dandona, R.; et al. Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: An analysis of data from the Global Burden of Diseases Study 2015. *Lancet* **2017**, *389*, 1907–1918. [[CrossRef](#)] [[PubMed](#)]
- Abbasi-Kangevari, M.; Malekpour, M.-R.; Masinaei, M.; Moghaddam, S.S.; Ghamari, S.-H.; Abbasi-Kangevari, Z.; Rezaei, N.; Rezaei, N.; Mokdad, A.H.; Naghavi, M.; et al. Effect of air pollution on disease burden, mortality, and life expectancy in North Africa and the Middle East: A systematic analysis for the Global Burden of Disease Study 2019. *Lancet Planet Health* **2023**, *7*, e358–369. [[CrossRef](#)] [[PubMed](#)]
- Yuval; Levi, Y.; Dayan, U.; Levy, I.; Broday, D.M. On the association between characteristics of the atmospheric boundary layer and air pollution concentrations. *Atmos. Res.* **2020**, *231*, 104675. [[CrossRef](#)]
- Yuval; Trischer, T.; Raz, R.; Levi, Y.; Levy, I.; Broday, D.M. Emissions vs. turbulence and atmospheric stability: A study of their relative importance in determining air pollutant concentrations. *Sci. Total Environ.* **2020**, *733*, 139300. [[CrossRef](#)] [[PubMed](#)]
- Wen, W.; Ma, X.; Tang, Y.; Wei, P.; Wang, J.; Guo, C. The impacts of meteorology on source contributions of air pollution in winter in Beijing, 2015–2017 changes. *Atmos. Pollut. Res.* **2020**, *11*, 1953–1962. [[CrossRef](#)]
- Guo, Q.; Wu, D.; Yu, C.; Wang, T.; Ji, M.; Wang, X. Impacts of meteorological parameters on the occurrence of air pollution episodes in the Sichuan basin. *J. Environ. Sci.* **2022**, *114*, 308–321. [[CrossRef](#)] [[PubMed](#)]
- Beckerman, B.S.; Jerrett, M.; Serre, M.; Martin, R.V.; Lee, S.J.; van Donkelaar, A.; Ross, Z.; Su, J.; Burnett, R.T. A hybrid approach to estimating national scale spatiotemporal variability of PM_{2.5} in the contiguous United States. *Environ. Sci. Technol.* **2013**, *47*, 7233–7241. [[CrossRef](#)] [[PubMed](#)]
- Akita, Y.; Baldasano, J.M.; Beelen, R.; Cirach, M.; de Hoogh, K.; Hoek, G.; Nieuwenhuijsen, M.; Serre, M.L.; de Nazelle, A. Large scale air pollution estimation method combining land use regression and chemical transport modeling in a geostatistical framework. *Environ. Sci. Technol.* **2014**, *48*, 4452–4459. [[CrossRef](#)] [[PubMed](#)]
- Di, Q.; Kloog, I.; Koutrakis, P.; Lyapustin, A.; Wang, Y.; Schwartz, J. Assessing PM_{2.5} exposures with high spatiotemporal resolution across the continental United States. *Environ. Sci. Technol.* **2016**, *50*, 4712–4721. [[CrossRef](#)] [[PubMed](#)]
- Hu, H.; Belle, J.H.; Meng, X.; Wildani, A.; Waller, L.A.; Strickland, M.J.; Liu, Y. Estimating PM_{2.5} concentrations in the conterminous United States using the random forest approach. *Environ. Sci. Technol.* **2017**, *51*, 6936–6944. [[CrossRef](#)] [[PubMed](#)]
- Brokamp, C.; Jandarov, R.; Hossain, M.; Ryan, P. Predicting daily urban fine particulate matter concentrations using a random forest model. *Environ. Sci. Technol.* **2018**, *52*, 4173–4179. [[CrossRef](#)] [[PubMed](#)]
- Tripathy, S.; Tunno, B.J.; Michanowicz, D.R.; Kinnee, E.; Shmool, J.L.C.; Gilooly, S.; Clougherty, J.E. Hybrid land use regression modeling for estimating spatiotemporal exposures to PM_{2.5}, BC, and metal components across a metropolitan area of complex terrain and industrial sources. *Sci. Total Environ.* **2019**, *673*, 54–63. [[CrossRef](#)] [[PubMed](#)]
- Tularam, H.; Ramsay, L.F.; Mutoo, S.; Brunekreef, B.; Meliefste, K.; de Hoogh, K.; Naidoo, R.N. A hybrid air pollution/land use regression model for predicting air pollution concentrations in Durban, South Africa. *Environ. Pollut.* **2021**, *272*, 116513. [[CrossRef](#)] [[PubMed](#)]
- Molnar, C. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. Leanpub. 2018. Available online: <https://christophm.github.io/interpretable-ml-book/> (accessed on 14 January 2026).
- Chan, K.Y.; Jian, L. Identification of significant factors for air pollution levels using a neural network-based knowledge discovery system. *Neurocomputing* **2013**, *99*, 564–569. [[CrossRef](#)]
- Adams, M.D.; Kanaroglou, P.S. Mapping real-time air pollution health risk for environmental management: Combining mobile and stationary air pollution monitoring with neural network models. *J. Environ. Manag.* **2016**, *168*, 133–141. [[CrossRef](#)] [[PubMed](#)]

18. Araki, S.; Shima, M.; Yamamoto, K. Spatiotemporal land use random forest model for estimating metropolitan NO₂ exposure in Japan. *Sci. Total Environ.* **2018**, *634*, 1269–1277. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Belachsen, I.; Broday, D.M. Decomposing PM_{2.5} concentrations in urban environments into meaningful factors 1. Separating the contribution of local anthropogenic activities from background and long-range transport. *Sci. Total Environ.* **2024**, *940*, 173749. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [\[CrossRef\]](#)
21. Strobl, C.; Boulesteix, A.L.; Kneib, T.; Augustin, T.; Zeileis, A. Conditional variable importance for random forests. *BMC Bioinform.* **2008**, *9*, 307. [\[CrossRef\]](#) [\[PubMed\]](#)
22. Hooker, G.; Mentch, L.; Zhou, S. Unrestricted permutation forces extrapolation: Variable importance requires at least one more model, or there is no free variable importance. *Stat. Comput.* **2021**, *31*, 82. [\[CrossRef\]](#)
23. Zhu, J.; Apley, D.W. Measuring variable importance via accumulated local effects. *arXiv* **2025**, arXiv:2512.21124. [\[CrossRef\]](#)
24. Baldauf, M.; Seifert, A.; Förstner, J.; Majewski, D.; Raschendorfer, M.; Reinhardt, T. Operational convective-scale numerical weather prediction with the COSMO model: Description and sensitivities. *Mon. Weather Rev.* **2011**, *139*, 3887–3905. [\[CrossRef\]](#)
25. Bauer, P.; Thorpe, A.; Brunet, G. The quiet revolution of numerical weather prediction. *Nature* **2015**, *525*, 47–55. [\[CrossRef\]](#) [\[PubMed\]](#)
26. Avgoustoglou, E.; Shtivelman, A.; Khain, P.; Marsigli, C.; Levi, Y.; Cerenzia, I. The sensitivity of the icosahedral non-hydrostatic numerical weather prediction model over Greece in reference to observations as a basis towards model tuning. *Atmosphere* **2023**, *14*, 1616. [\[CrossRef\]](#)
27. Gochakov, A.V.; Tokarev, V.M.; Kolker, A.B. Verification of vertical temperature profiles in the COSMO and ICON models with available observational data. *IOP Conf. Ser. Earth Environ. Sci.* **2022**, *1023*, 012004. [\[CrossRef\]](#)
28. Chen, S.; Yuavli; Broday, D.M. A new modeling approach for assessing the contribution of industrial and traffic emissions to ambient NO_x concentrations. *Atmos. Environ.* **2018**, *173*, 173–184. [\[CrossRef\]](#)
29. Chen, S.; Bekhor, S.; Yuavli; Broday, D.M. Aggregated GPS tracking of vehicles and its use as a proxy of traffic-related air pollution emissions. *Atmos. Environ.* **2016**, *142*, 351–359. [\[CrossRef\]](#)
30. Yuval; Levi, Y.; Khain, P.; Broday, D.M. Clustering of emission–dispersion dynamics in a state space defined by pollution sources and meteorological variables. *Sci. Total Environ.* **2025**, *995*, 180051. [\[CrossRef\]](#) [\[PubMed\]](#)
31. Ikotun, A.M.; Ezugwu, A.E.; Abualigah, L.; Abuhaija, B.; Heming, J. K-means clustering algorithms: A comprehensive review, variants analysis, and advances in the era of big data. *Inf. Sci.* **2023**, *622*, 178–210. [\[CrossRef\]](#)
32. MathWorks. *Matlab*, Version R2025b; The MathWorks Inc.: Natick, MA, USA, 2025.
33. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. *arXiv* **2017**, arXiv:1705.07874. [\[CrossRef\]](#)
34. Sundararajan, M.; Taly, A.; Yan, Q. Axiomatic attribution for deep networks. *arXiv* **2017**, arXiv:1703.01365. [\[CrossRef\]](#)
35. Apley, D.W.; Xhu, J. Visualizing the effects of predictor variables in black box supervised learning models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2020**, *82*, 1059–1086. [\[CrossRef\]](#)
36. Lakshmanan, V.; Karstens, C.; Krause, J.; Elmore, K.; Ryzhkov, A.; Berkseth, S. Which polarimetric variables are important for weather/no-weather discrimination? *J. Atmos. Ocean. Technol.* **2015**, *32*, 1209–1223. [\[CrossRef\]](#)
37. Belachsen, I.; Broday, D.M. Decomposing PM_{2.5} concentrations in urban environments into meaningful factors 2. Extracting the contribution of traffic-related exhaust emissions. *Sci. Total Environ.* **2024**, *940*, 173715. [\[CrossRef\]](#) [\[PubMed\]](#)
38. Yuval; Levi, Y.; Broday, D.M. Revealing causality in the associations between meteorological variables and air pollutant concentrations. *Environ. Pollut.* **2024**, *345*, 123526. [\[CrossRef\]](#) [\[PubMed\]](#)
39. Lionello, P.; Scarascia, L. The relation between climate change in the Mediterranean region and global warming. *Reg. Environ. Change* **2018**, *18*, 1481–1493. [\[CrossRef\]](#)
40. Shafran-Nathan, R.; Broday, D.M. Long-term recurrent exposure to excessive heat in the East Mediterranean and climate resilience development—A case study in Israel. *Reg. Environ. Change* **2024**, *24*, 153–165. [\[CrossRef\]](#)
41. Falah, S.; Banerjee, T.; Kizel, F.; Broday, D.M. Utilizing multiple satellite data products for estimation of PM_{2.5} concentrations—Accounting for aerosol properties improves the prediction. *Environ. Pollut.* **2023**, *320*, 121119. [\[CrossRef\]](#) [\[PubMed\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.