The Influence of Weather on Fashion Retail: Developing a Sustainable and Interpretable Forecasting Model

MSc Research Project
MSc in Artificial Intelligence for Business

César Antonio Ledesma González
Student ID: 23425016

School of Computing
National College of Ireland

Supervisor: Victor del Rosal

National College of Ireland
Project Submission Sheet

Student Name:	César Antonio Ledesma González
Student ID:	23425016
Programme:	MSC AI for Business — **Year:** 2024
Module:	MSCAIBUS
Lecturer:	Victor del Rosal
Submission Due Date:	11/08/2025
Project Title:	"The Influence of Weather on Fashion Retail: Developing a Sustainable and Interpretable Forecasting Model"
Word Count:	4800

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the references section. Students are encouraged to use the Harvard Referencing Standard supplied by the Library. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action. Students may be required to undergo a viva (oral examination) if there is suspicion about the validity of their submitted work.

Signature: César Antonio Ledesma González

Date: 11/08/2025

PLEASE READ THE FOLLOWING INSTRUCTIONS:

1. Please attach a completed copy of this sheet to each project (including multiple copies).
2. Projects should be submitted to your Programme Coordinator.
3. **You must ensure that you retain a HARD COPY of ALL projects**, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer. Please do not bind projects or place in covers unless specifically requested.
4. You must ensure that all projects are submitted to your Programme Coordinator on or before the required submission date. **Late submissions will incur penalties.**
5. All projects must be submitted and passed to successfully complete the year. **Any project/assignment not submitted will be marked as a fail.**

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

The Influence of Weather on Fashion Retail: Developing a Sustainable and Interpretable Forecasting Model

Ledesma César Antonio (MsC)[1]

[1] National College of Ireland, Mayor Street, IFSC, Dublin 1, Ireland.
x23425016@student.ncirl.ie

Abstract. This study explores how weather affects sales across U.S. fashion retailer, but more than that, it is an attempt to build a sustainable, transparent forecasting architecture that balances predictive performance and ethical considerations. Led by a focus on actionable forecasting and environmental protection, the analysis integrates and refines sales transactions with weather records, applying a Linear Regression, XGBoost and Facebook Prophet, and then adding interpretability via SHAP. Based on the findings, the most streamlined approach, Linear Regression achieves the highest coefficient of determination (R^2= 0.61), underscoring the comparative significance of carefully curated input features and intuitive modelling over more intricate, complex systems. Weather data showed only modest predictive lifts; in contrast, attributes reflecting product hierarchy and calendar structure proved more impactful. A distinctive contribution is the embedding of CO_2 emissions forecasting into model performance metrics, filling a conspicuous gap in the evolving literature on sustainable AI models. The prototype functions simultaneously as a forecasting engine for the fashion domain and as a modest step toward AI systems that are accountable, transparent, and ecologically minded. The work is premised on the conviction that data-driven analysis can advance corporate objectives while respecting environmental constraints, honouring the minimal ethical duty of practitioners to reduce the adverse impacts of the technologies they design and deploy.

Keywords – *Sales Forecasting, Prophet, Explainable AI, SHAP, Fashion Retail, Responsible AI.*

1. Introduction

Fashion has always captured my interest as an industry of art, culture, and society. With this project, I aimed to combine an appreciation for fashion with a technological approach to predict market demand. Throughout the process, I became more committed to not only building a model but making sure it was interpretable, actionable, sustainable, and useful for business in the real world, as a complement to their operations. Understanding now how external factors like weather influence our consumption and finding a way to align operational performance with eco-friendliness was one way to address this problematic. Creativity, ethical concern and data science guided choices across the entire project. This work became a space where I could merge those creative insights with a structure responsible use of AI, not to replace human decisions, but to inform more wisely.

The initial methodological framework proposed in Practicum 1 aimed at developing a hybrid, interpretable forecasting model for the fashion retail domain. It integrated product visual characteristics obtained through a pre-trained Convolutional Neural Networks (CNNs) with contextual variables, for instance, weather patterns. The pipeline model was based on Facebook Prophet, a framework with a modular approach, and combined with SHAP (SHapley Additive Explanations) to ensure evaluative model transparency and stakeholder interpretability.

This research seeks to examine the impact of weather conditions on fashion retail sales in the United States with the goal of developing a sustainable and interpretable forecasting model. The primary contribution of this research is the creation, deployment and implementation of a streamlined forecasting framework that integrates weather data with sales data, using forecasting models such as Linear Regression, XGBoost, and the Facebook Prophet. A minor contribution of this research is the incorporation of sustainability metrics such as computational cost and estimated CO_2 emissions into the model evaluation pipeline. Building this, I wasn't just testing algorithms; I was also testing my belief that technology can reflect our values if we desire it.

This paper discusses the evolution of modelling strategy beginning with a multi-modal approach, and the structure is structured as follows: in Section 2 related work in weather-influenced retail forecasting, interpretable machine learning, and sustainable AI practices. The research methodology, data sources, and tools are discussed in Section 3. Section 4 discusses the design components and the evolution of the proposed model. The implementation of this research is discussed in Section 5. Section 6 presents and discusses the evaluation results, performance metrics, alongside environmental impact considerations. Section 7 concludes the research, discusses key findings and proposes future work directions.

2. Related Work

Understanding demand in the fashion retail industry is a complex challenge due to short product life cycles, seasonality, and rapidly shifting consumer behaviour. Traditional time-series models such as ARIMA and SARIMA have been used fundamentally in sales forecasting (3), despite the fact that their structural inflexibility increasingly limits their use in today's context of rapid changes in demand, especially in the fashion retail domain, where items have very short life cycles. These models tend to fail in the non-linear demand and to incorporate external variables like weather, holidays, etc., which can contribute to an insightful effect on consumer purchasing behaviour patterns and on companies' logistics operations.

Recent studies explored the application of machine learning models and Deep Learning techniques to tackle these limitations. Nguyen et al. (2021) show how external factors, such as weather and meteorological conditions, can be effectively incorporated into an LSTM-based forecasting model to improve its capability of anomaly detection in retail sales. In a similar way, Li et al. (2024) proposed a multi-modal model that incorporates product images and textual features with temporal data to predict new product sales. While these models showed great promise, they are computationally intensive and frequently not interpretable.

Giri and Chen (2022) have demonstrated that visual product features have been shown to play a significant role and influencing consumer preferences in fashion. Their work explores the use of CNNs for extracting visual attributes such as textures and colour for forecasting demand. Despite these approaches achieving a high level of accuracy, their computational demands may present substantial barriers to implementation, especially for small and medium-sized enterprises (SMEs). Additionally, the majority of popular visual forecasting models do not consider contextual variables like weather conditions or holidays, which reduces and limits their practical applicability and relevance from the business perspective in a competitive market.

Explainable AI (XAI) has emerged as a response to the "black-box" nature of complex models. SHAP, according to Lundberg & Lee (2017), provide model-agnostic explanations by quantifying the contribution of each feature to individual predictors. As a result, it is considered a valuable and reliable tool in time-series and regression contexts. Prophet, a model developed by Facebook, achieves a balance between simplicity and interpretability while allowing the incorporation of external regressors. Studies such as Negre et al. (2024) have demonstrated the Prophet's superiority over SARIMA in the context of handling missing data and seasonal trends in retail forecasting.

Samek et al. (2019) emphasise the growing demand for transparency in AI applications, especially in commercial environments where trust and accountability are paramount. In contrast, certain techniques like Grad-CAM (Selvaraju et al., 2017), though valuable for visual explanations in image-based tasks, are underutilised in demand forecasting.

At the same time, an increasing growing number of publications have emphasised the need for sustainable AI. The issue of energy consumption and CO_2 emissions from model training is becoming a significant consideration. The integration of environmental impact assessments into modelling pipelines represents a recent and still emerging area of research (Schwartz et al. 2020), with the potential to align AI research with broader sustainability goals.

While CNN-based and transformer-based models are highly accurate across forecasting tasks, their high computational requirements limit their practical deployment in low-resource environments. Consequently, intuitive forecasting strategies that deliberately sacrifice computational demand for ecological and ethical sustainability have not received sufficient methodological attention. Even though the relevance of such trade-offs is gaining recognition, the field currently lacks validated, scalable alternatives that are both technically sound and amenable to routine application across diverse operational contexts. This research aims to reduce the gap by combining environmental responsibility with explainable and low-cost computing techniques.

3. Methodology

The research methodology consists of six main steps named *data collection, data cleaning & preprocessing, feature engineering, model training & tuning, explainability & interpretation, evaluation & sustainability* analysis and results as shown in Fig. 1. These steps are organised in a sequential data science workflow designed to ensure methodological clarity, computational efficiency, and sustainability impact tracking.

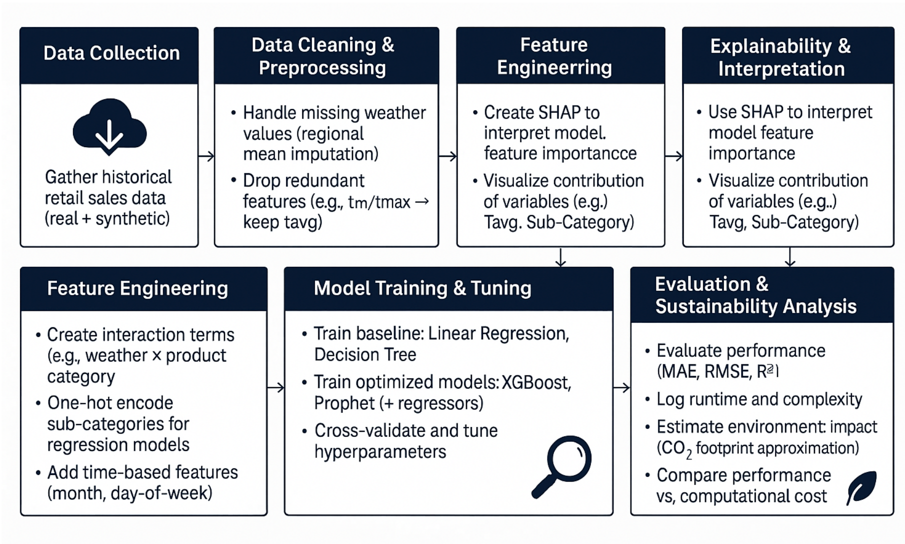

Fig. 1. Research Methodology

The first step, *Data collection,* involved gathering two core datasets. The first comprised real and synthetically generated retail sales transactions from a US-based company, categorised by product type, sub-category and region (23). The second dataset obtained from a public APIs, Metostat and NOAA, contained weather conditions such as temperature, precipitation, and wind speed. Both datasets were combined using date and geographic location as join keys, ensuring no loss of information, which would otherwise affect or bias subsequent modelling.

The second step is *Data Cleaning and Pre-processing.* This step includes activities to prepare data to make it consistent, ready and significant for analysis. This phase involved imputing missing weather values with region-based averages and removing hyper-correlated variables (replacing *tmin* and *tmax* in favour *tavg*). Categorical parameters like region and product sub-category were encoded using one-hot and label encoding strategies, where applicable.

The third step, *Feature Engineering,* involves enriching the dataset with additional metrics. Interaction features between weather variables and product categories were created to evaluate potential relationships and improve model performance. Temporal indicators like month, day-of-week, and holiday flags were also added to capture seasonal or cyclical behaviours within the sales data. To improve model performance and interpretability, the feature space was optimised to reduce redundancy while preserving variance.

The fourth step, *Model training and tuning*, a couple of algorithms were implemented and refined. Baseline models such as Linear Regression and Decision Tree Regressor were first used to establish a performance benchmark. These were followed by more robust algorithms, including XGBoost (2) and Facebook Prophet, the latter configured both with and without external weather regressors. K-Fold and Time Series Cross-Validation were applied to tune hyperparameters to ensure generalisation and mitigate overfitting.

In the fifth step, *Explainability and interpretation*, the focus shifted to providing clarity on the reasoning behind model decisions after training. Model explainability was provided through prediction attribution using SHAP. This resulted in identifying some of the variables driving model performance, for instance, average temperature and specific product sub-categories. Seasonal trends within and between categories were further analysed through clustering.

Finally, the sixth step, Evaluation *and sustainability analysis*, model performance was evaluated using traditional metrics such as MAE, RMSE, and R^2. Additionally, the computational cost of each modelling pipeline was documented through runtime and resource logging. These values were translated into estimated CO_2 impact using emission factor benchmarks aligning with this methodology, Sustainable AI practices to contribute to achieving the UN Sustainable Development Goals, such as ISO 14064 -1 and GHG Protocol.

3.3 Tools and Technologies Used

A critical factor for this methodological approach was the balance of real-world applicability alongside interpretability and computational efficiency (17). As the research progressed, it became clearer that despite more complex models like Deep Learning models providing higher predictive performance, they were not transparent and required significant computational resources. Given the limited dataset size and the environmental concerns associated with large-scale model training, the decision was made to prioritise simpler, interpretable models supported by robust feature engineering. These tools struck the balance between depth and simplicity, which is where the integration of explainable AI techniques like SHAP and forecasting tools like Prophet, without sacrificing the accessibility of the system for small and medium-sized fashion retailers and preserving interpretability.

Furthermore, aligning model development with sustainability criteria added an evaluative layer, ensuring the final models met the accuracy criteria but were also responsible for the design and deployment.

4. Design Specification and Model Evolution

This step illustrates the starting iterations of the forecasting model from the initial multi-modal framework to the final lightweight solution applied in this study. It also includes design considerations that shaped and enriched model architecture and explains the rationality behind the methodological decisions made during the research process. The architecture summarised in Fig. 2 reflects both the evolution of the model and its final structure.

The model is organised into related blocks: a data preparation pipeline, the forecasting engine, and the interpretability-enhancing elements. The forecasting engine, which forms the core of the framework, supports structured tabular data inputs and includes multiple algorithms ranging from classical regressors to Prophet for time series models. Complementary to that, a SHAP analysis module was added to provide model explanations.

Each component of the system supports a specific function: the ***data preparation module*** handles merging, cleaning, and encoding; ***the model block*** executes training and forecasting; and the ***interpretation module*** visualises the impact of weather and category-level variables on forecast outputs. These are discussed in the subsections 4.1 to 4.4.

This section describes how the model evolved, starting from the initial multi-modal concept to the final, lightweight and sustainable version applied in this project. It includes design considerations that shaped model architecture and explains the rationale behind methodological decisions made during the research process.

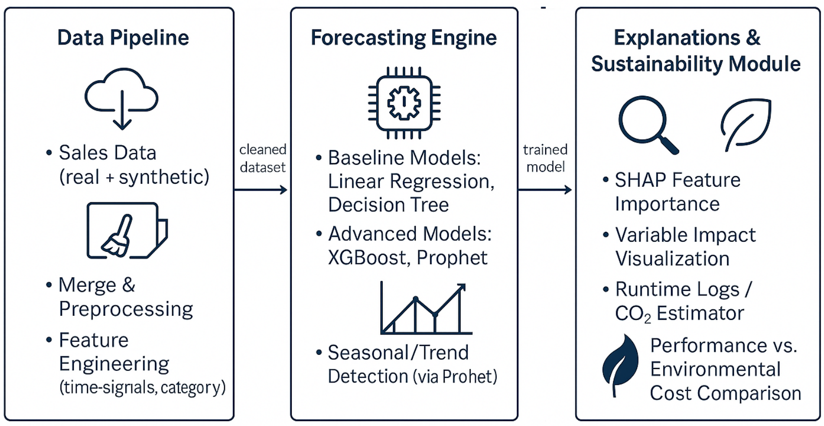

Fig 2. Forecasting and Modelling Architecture

4.1 Final Model Architecture Design

The final model framework focuses on integrating structured numerical and categorical data: average daily temperature, precipitation, product category, sub-category, and temporal flags (e.g. month, weekday, holiday). These variables were fed into three different types of models: baseline regressors (Linear Regression and Decision Tree), advanced models (XGBoost), and a time series model (Prophet with and without regressors) to evaluate the performance and significance.

The Prophet model was maintained in the design due to its ability to handle and capture seasonality and its compatibility with SHAP for post-interpretation. All models were tuned and evaluated not only based on their predictive accuracy but also on their inference time and estimated computational cost, following the best practices within the industry and relevant research.

4.2 Design Considerations and Ethical Alignment

Throughout the model design process, key considerations include interpretability, model simplicity, environmental cost, and practical deployment. By selecting SHAP-compatible models and avoiding black-box algorithms, the research promoted transparency and facilitated stakeholder trust. Model selection was guided by both performance metrics (MAE, RMSE, R^2) and the estimated CO_2 footprint derived from runtime and hardware usage.

This approach reflects a personal commitment to ethical and sustainable practices in AI real-world scenarios, providing a replicable and responsible framework for forecasting in the fashion retail industry. This tailored forecasting framework has been designed with ease of use in mind for small to medium retailers and removes barriers to the utilisation of predictive analytics and looking to be aligned with the UN Sustainable Development Goals.

Moreover, the project identifies a possibility of bias within the system, even in a "non-biased" area such as sales integrated with weather data. To mitigate these biases, the data was purposefully bounded to omit personal, demographic attributes such as age, gender, race or income. Although this approach reduces the ability to provide tailored experiences, it strengthens the ethical position of avoiding biased and discriminatory social inferences. Future replication of this framework could be enhanced by fairness criteria to assess bias and representation of the estimates across different locations, categories, and demographic variables.

The goals and criteria for this project seek to integrate and balance model performance, interpretability, environmental awareness, and fairness, setting a precedent for inclusive, transparent and sustainable AI solutions in the fashion retail domain. Any potential bias found must be mitigated for reproducibility.

5. Implementation

5.1 Tools and Environment Setup

The forecasting system was developed using Python 3.7.17 within Visual Studio Code environments. Core libraries included Pandas and Numpy for data manipulation, *scikit-learn* for classical models, *XGBoost* for gradient boosting, and *Prophet* for time series forecasting. *SHAP* was used for model interpretation. Data preprocessing, visualisation, and runtime logging were managed using *seaborn, matplotlib,* and *Python's* built-in time and resources modules. All models were trained and evaluated on a local CPU setup to support low-resource use cases.

5.2 Data Integration Workflow

Sales and weather data were loaded separately from CSV and JSON formats. Precipitation and wind variables were added to the weather data, which was extracted using the Metostat API. The records were aligned daily and combined based on location. Region-level averages were used to fill in missing values during integration. When required, one-hot encoding was applied to encode categorical variables, and date-time tools were used to programmatically build temporal aspects such as month, weekday, and holiday flags for seasonal patterns.

5.3 Outputs and Deliverables

The last stage of the project was for producing a detailed output set, which incorporated an assessment and implementation of the forecasting framework. These outputs made a significant contribution towards narrative, within the parameters of the study, replication and iterative purposes.

Codebase: A complex forecasting system was designed from a collection of Python scripts and Visual Studio Code notebooks, which were constructed in a modular and reusable way. The complete machine learning pipeline is documented and covered by these scripts, starting from raw data ingestion and data cleaning, all the way to feature engineering, model training, achieved through methodical file organisation and comprehensive inline comments. The codebase is designed to allow for model configuration changes and enable model testing, resulting in facilitating quick testing and different configurations.

Data Transformations: that stage includes a variety of multi-state processes, including the transformation of unstructured datasets into structured and pre-processed datasets, which are ready for modelling. Some of the processes include meteorological data gaps fulfilled through region-based averaging, temporal feature (month, weekday, or holiday) creation, and one-hot encoding for categorical features such as product sub-category. These tasks were automated for both training and test datasets through script solutions.

Model Artefacts: The models were explored using *.PKL* for regressors and *.JSON* for Prophet configurations. This allows loading model states for future testing or deployment. Each model artefact captures a specific experiment configuration (e.g. *Prophet* with and without regressors, *XGBoost* with engineered features), which supports and allows a critical comparison across modelling strategies. Additionally, these artefacts facilitate downstream integrations into dashboards or inventory tools.

Visual Outputs: Multiple visualisations were produced to enhance model understanding, data attributes, and the forecasting performance evaluation. These include the SHAP summary plots (Fig. 9), featuring importance scores (Fig. 6), time series Prophet forecasts showing and illustrating seasonal patterns (Fig. 5), bar plots of MAE/RMSE scores across different models, and clustered scatter plots grouped by product subcategories (Fig. 10). The visual outputs are shown as follows.

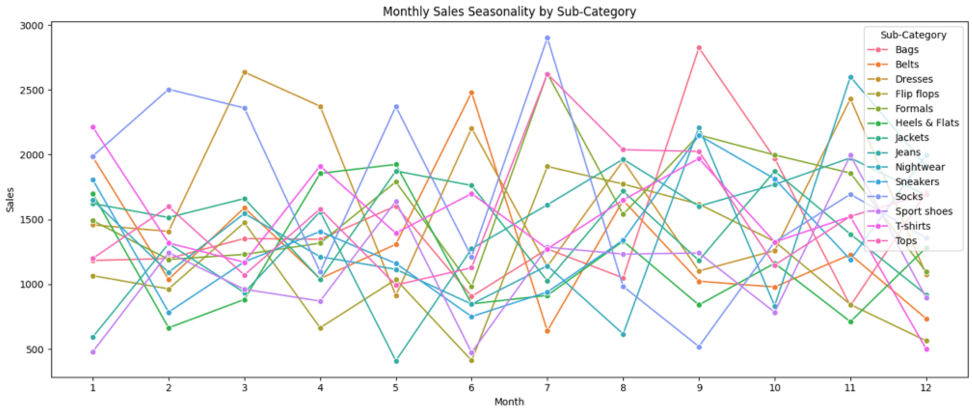

Fig 3. Sub-category Sales on Time Series

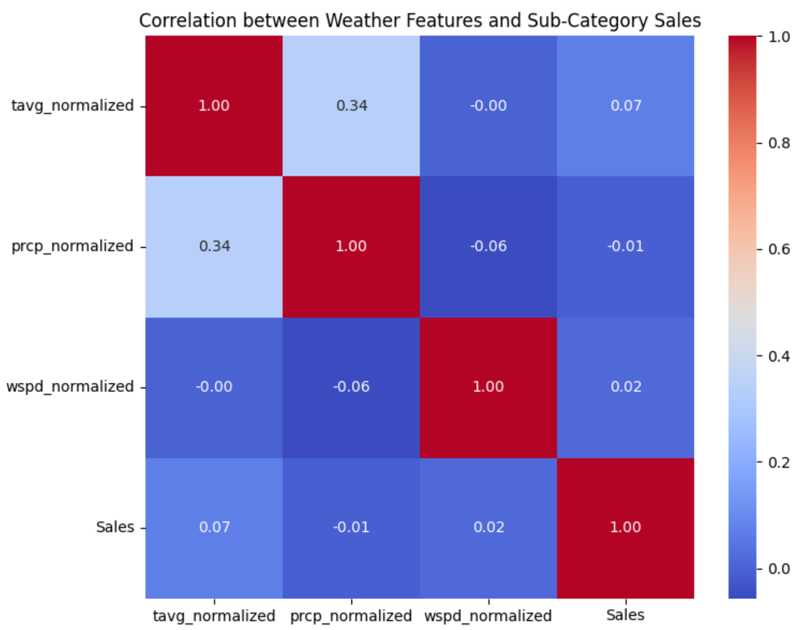

Fig 4. Correlation Heatmap on Weather Features

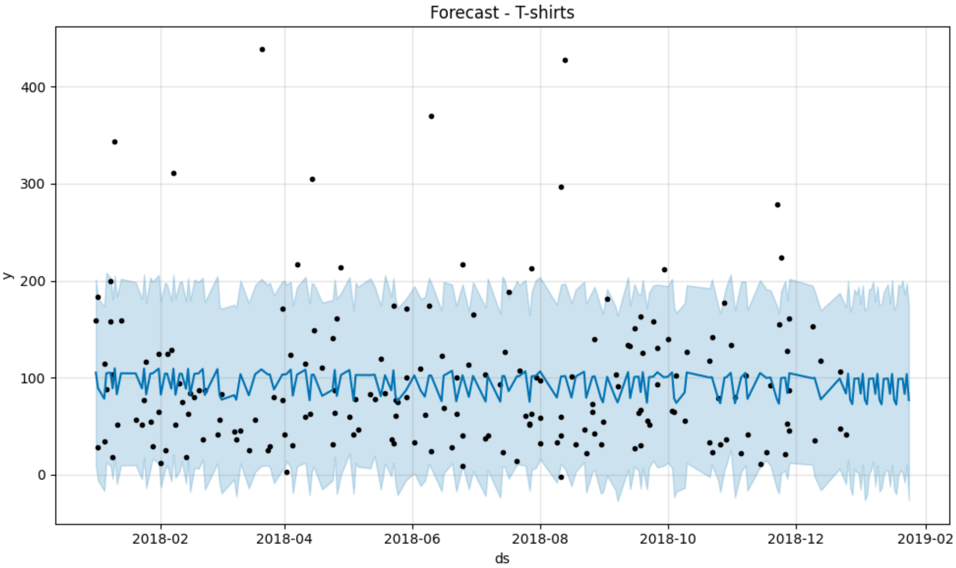

Fig 5. Prophet Forecast of a High-Volume Subcategory Item

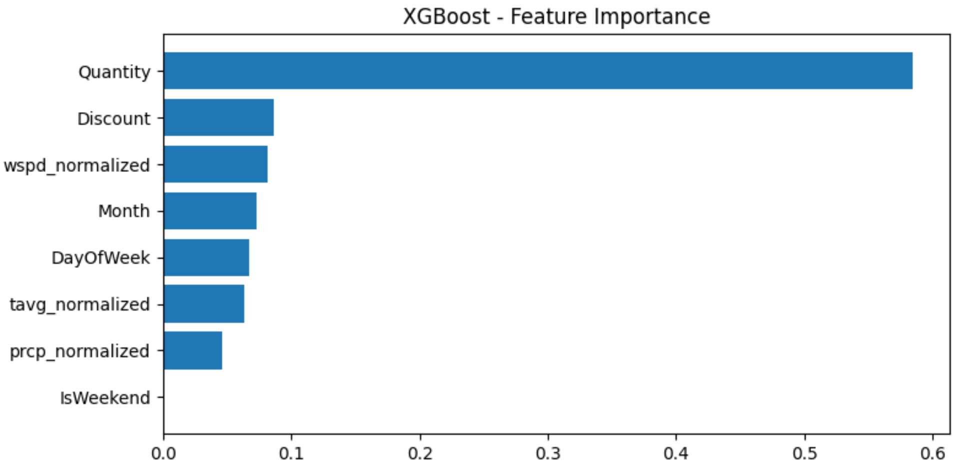

Fig. 6. XGBoost Feature Importance

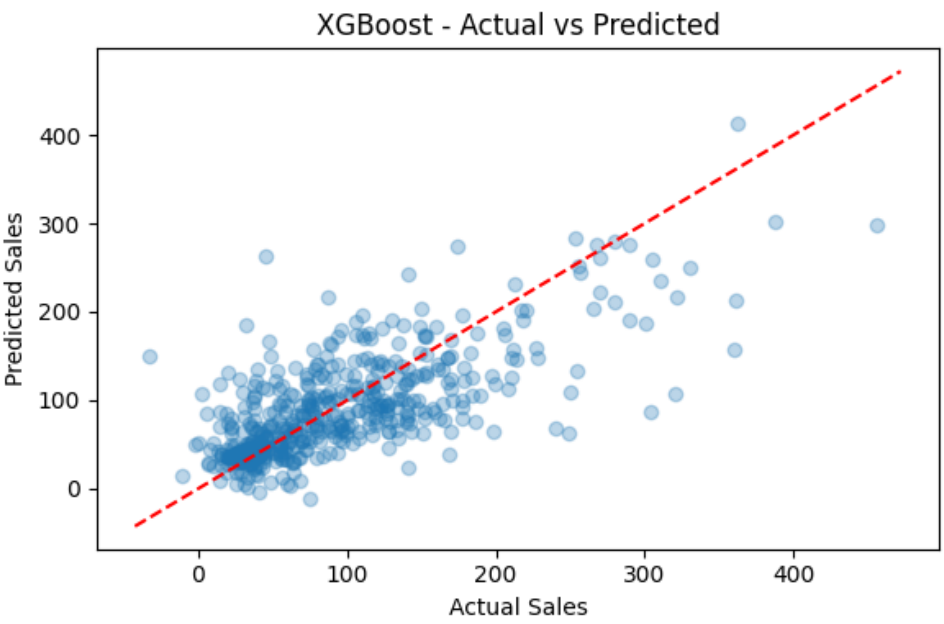

Fig 7. Actual Sales vs. XGBoost prediction

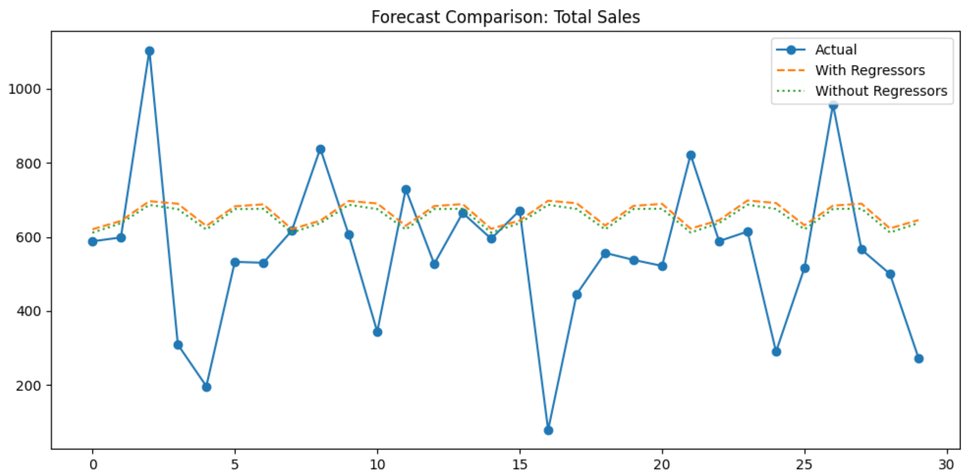

Fig 8. Forecast with and without Regressors

5.4 Runtime and Sustainability Logs

Part of the environmental evaluation for the forecasting models, dedicated resource monitoring tools were deployed to track time, CPU workload, and memory usage for each configuration. In the lack of direct CO_2 emissions monitoring equipment, operational logs provided means to indirectly estimate energy consumption through the following formula:

$$COe = Runtime(h)\ x\ \frac{Power\ Draw(X)}{1000}\ x\ PUE\ x\ Grid\ Emission\ Factor\ (kg\ COe/kWh)$$

This formula is obtained from the Green Algorithms framework (Lannelongue et al. 2021), which provides a well-known method for estimating the computational carbon emissions. It also supports concepts from the ISO-14064-1 (9) standard and the GH Protocol, which prioritise transparency, traceability and the use of verifiable activity data for the calculations based on indirect emissions.

6. Evaluation

This research project aims to design and deploy a forecasting framework and based on the initial research goals, to construct a lightweight, interpretable, and environmentally sustainable model for predicting fashion retail sales based on meteorological and categorical external factors. The results focused on prediction accuracy, variable relevance, identified temporal trends, and a life-cycle carbon impact.

6.1 Feature Importance and SHAP Interpretation

SHAP analysis identified *Sub-Category* encoding, *tavg*, and *Holiday* variables as the leading determinants of forecasting across all architectures (Fig. 9). The influences of product specification dynamics were most pronounced for *Jeans, T-shirts, and Sneakers,* which are consistently remarked as dominant predictors. Both XGBoost and the Decision Tree implementations demonstrated that categorical encodings and lagged sales features outperformed any extrinsic meteorological predictors. This finding underscores the notion that localised retail demand patterns tend to exhibit temporal invariance across product lines, diminishing the marginal impact of varying external climatic conditions. In other words, the result suggests that customers' shopping behaviour for different fashion items tends to be consistent over time, even with weather variances.

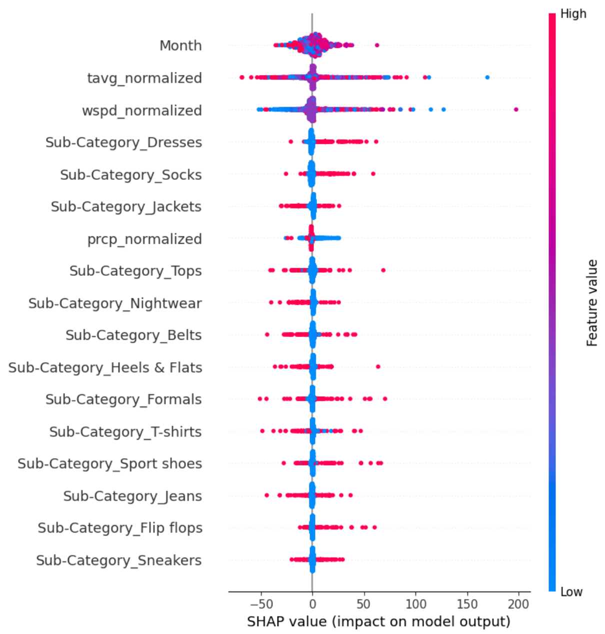

Fig 9. Shap Summary Plot

Prophet models captured seasonal demand oscillations with high fidelity. T-shirts and Sneakers registered pronounced surges in the warm months, whereas Nightwear and Jackets attained their maxima during colder intervals. K-Means clustering segregated the sub-catalogue into three seasonally coherent groups, substantiating the efficacy of segment-specific modelling. Subcategories such as Belts, Socks and Heels (Fig. 10), by contrast, exhibited relatively flat temporal profiles, suggesting that the predictive gains from adopting elaborate modelling frameworks in these cases are marginal at best.

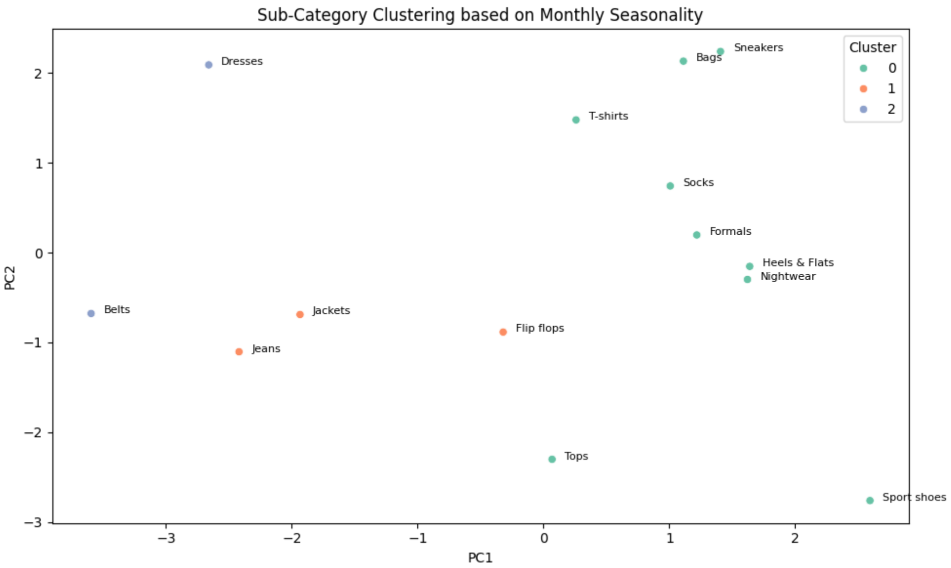

Fig 10. Sub-Category Seasonality Clusters

Initially, the weather seemed to present a potential level for enhancing the demand forecast, yet it failed to substantially elevate model accuracy. Analysis of weather's influence across various categories revealed weak explanatory power, with R^2 values rarely exceeding 0.10. Only isolated gains emerged for Bags, Flipflops, and Socks (Fig. 10), where incorporation of weather data produced a marginal reduction in error, confined to fluctuations of +/-1 RMSE unit.

These restricted gains imply that the national meteorological averages may obscure more salient microclimatic influences. Furthermore, an underappreciated explanatory space remains, as critical covariates like discounts, promotions, or pricing were omitted from the specification.

6.2 Model Performance

A comparison of six predictive approaches was performed. The Linear Regression model provided the strongest performance, recording an MSE (Mean Absolute Error) of 31.74 units, an RMSE (Root Mean Square Error) of 44.57 units, and a Coefficient of determination of 0.61. The XGBoost implementation showed a lower R^2 of 0.50, while the Decision Tree exhibited excessive variance, evident from the minimal R^2 of 0.14. The Prophet framework, tested in configurations with and without the meteorological predictors, successfully captured a large-scale seasonal effect; however, even with the incorporation of weather, it contributed little to the overall accuracy gain.

Model	**MAE**	**RSME**	**R^2**
Linear Regression	31.74	44.57	0.61
Decision Tree	46.09	66.08	0.14
XGBoost	35.75	50.39	0.50
Prophet (No Regressors)	~50-70	Varies	N/A
Prophet (Weather Regressors)	±1 RMSE	~0.0-0.7	Mixed results, minor gains

Table 1. Model Performance

Three different models were evaluated: Linear Regression, Decision Tree, XGBoost and time series models. Linear Regression resulted as the best performer, achieving an R^2 score of 0.61, MAE of 31.74, and RMSE of 44.57. The model's simplicity worked to its advantage, and it was able to outperform all complex models due to feature engineering and the dataset's relatively low-noise structure. Furthermore, it offered full explainability, the lowest training time, and was suitable for low-resource environments.

Decision Tree, on the other hand, underperformed with and R^2 of 0.14. The dominant feature appears to violate the general picture since trees are known to perform well at capturing non-linear interactions, and the small dataset with limited variability in some features contributed to high overfitting. This result raises the risk of using flexible models on modest datasets without ensemble techniques (13).

A moderate performance was reported by XGBoost (R^2=0.50). This model demonstrates a balance between strength and complexity. While it offered more robust generalisation than Decision Trees, it required significantly longer training time and was less interpretable without SHAP post-processing.

Within the time series category, the Facebook Prophet model offered valuable insights about seasonality. The model captured expected retail cycles, including January slowdowns and summer peaks, particularly for categories like T-shirts, Nightwear, and Sneakers. Prophet's performance differed with the addition of regressors. With weather variables, it showed marginal performance changes, with an average RMSE fluctuation of ±1 across runs. This suggests that weather had a limited influence on short-term sales fluctuations in the dataset, and this finding was later confirmed via SHAP analysis.

Prophet's modularity and interpretability were strong advantages. On the other hand, it was less accurate than the regression-based models estimating absolute values, which makes it more useful as a trend detector than a sales predictor in this context.

Overall, the evaluation confirmed that model performance was more influenced by internal features like sub-category and temporal signals, rather than external variables like weather. This matches the SHAP findings and reinforces the conclusion that data quality and contextual are more relevant features than model sophistication.

6.3 Sustainability Assessment

Quantitative chronometry for training and inference was collected from all models. Experiments employed a local CPU, with GPU resources withheld. CO_2 emissions were approximated on execution time and energy footprints, leading to the inference that sparser architectures, such as Linear Regression, yield the best performance and sustainability.

Model	Avg. Training Time (sec)	Energy Impact
Linear Regression	~1.2	Very low footprint
Decision Tree	~1.6	Minimal complexity
XGBoost	~4.7	Moderate compute, no GPU needed
Prophet (no regressors)	~2.1 per sub-model	Efficient.
Prophet (regressors)	~2.3 per sub-model	Slight increase from feature overhead
SHAP	~6.5 (incl. plots generation)	Interpretability costs but has low energy

Table 2. Models, Runtime and Energetic Impact

The computational demands of this study, centred on training and validating forecasting algorithms within a CPU, produced an estimated carbon footprint of 0.65 kg of CO_2 equivalent for each complete model execution. This estimation was calculated using the standard European grid emission factor of 0.231 kg CO_2 e/kWh (7)(14). These emission figures are the energy required to operate a 10W LED lamp for 32 hours. Despite the relatively low absolute value, the footprint underscores the necessity of integrating energy efficiency and carbon assessment into the model design process, particularly when the forecasting systems are planned for extensive and large-scale real-world deployment(16) (17) (18).

While the absolute CO_2 emissions per model run were low, the emissions become more significant considering frequent retraining or multi-region deployments. As the study highlights, models such as Linear Regression not only provided strong predictive accuracy but also minimised energy consumption when compared with more complex models. These observations demonstrate the need to shift the focus of model evaluation alongside predictive accuracy to include the environmental impact and computation cost.

Additionally, the study supports the principles of responsible AI as defined by the EU with Green AI policies and the IEEE with sustainable computing frameworks. Future work could explore the possibility of real-time emissions tracking and reporting through an API or any other available tool.

7. Conclusion and Future Work

This research aimed to quantify how external variables, such as weather variables, affect fashion performance in the United States while constructing a sustainable forecasting model that focuses on both interpretability and ethical decision support. The proposed architectures are lightweight and combine structured and meteorological data with transparent techniques named Linear Regression, XGBoost, and Facebook Prophet, augmented by SHAP to facilitate model interpretability.

Experimental results indicate that Linear Regression exceeded performance benchmarks, achieving R^2=0.61 and full transparency in coefficient interpretation. Prophet successfully delineated seasonal demand patterns, particularly for T-shirts, Sneakers, and Nightwear sub-categories. SHAP analysis showed explainability metrics, particularly product sub-category and calendar period, dominated predictive power, while weather inputs exhibited uniformly weak correlation. Subsequent clustering analysis confirmed robust seasonal cycles across subcategories, a finding that allows a tighter segmentation for inventory allocations and tailored marketing strategies.

A significant advancement presented in this project is the integration of green gas emission modelling, specifically CO_2, into the evaluative framework for predictive systems. By analysing runtime and computational resources, the study estimates the latent energy costs of model execution, and this introduces a quantitative axis for differentiation that extends beyond predictive precision and computational latency. This methodological development fills a gap in the current literature and is consistent with the emerging discipline of energy-conscious data-driven science, making the findings especially relevant for startups and resource-constrained practitioners.

The academic contribution is accompanied by a personal motivation to promote and contribute to achieving the United Nations Sustainable Development Goals, with explicit focus on responsible consumption (SDG 12), Climate Action (SDG 13), and inclusive, Impactful Technology (SDG 9). The endeavour stands as a technical milestone and a measured contribution to the cultivation of lean, transparent, and ethically aware Artificial Intelligence that mitigates ecological tension while addressing pressing societal demands. Looking back, this project showed me how, from every modelling choice, the features I engineered to the hardware I used, it is a reflection of values as much as technique.

Research future trajectories that follow this work could integrate market-variable features, such as pricing and promotional cadences, enhance the spatial resolution of ambient data, and extend the modelling architecture to multi-regional demand projections. Furthermore, supplementary CO_2

quantification could be achieved by deploying public application programming interfaces or dedicated sensing arrays, thereby tightening the precision of the ecological footprint assessment.

In the context of expanding digital ecosystems and growing environmental pressure, the following research introduces a flexible, ethically informed AI framework designed to harmonise ecological stewardship with technological advancement and managerial discipline. While intentionally focused, the discussion captures the baseline ethical duty incumbent upon engineers to reduce the planetary costs incurred by our creations. The outcome is not merely a suite of algorithms but an invitation to design systems that respect and operate with the ecological thresholds that sustain human and environmental health. I have learned that great and modern models are not just built with code, they are shaped by the questions we ask, the impact and benefit we consider, and the kind of world we want to help create together.

References

[1] A. Shehabi et al., United States Data Centre Energy Usage Report, Lawrence Berkeley National Laboratory, U.S. Department of Energy, Jun. 2016. [Online]. Available: https://eta.lbl.gov/publications/united-states-data-center-energy

[2] Chen, T. and Guestrin, C. (2016) ‘XGBoost: A scalable tree boosting system’, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794. Available: https://arxiv.org/abs/1603.02754

[3] Chollet, F. (2021) *Deep Learning with Python*. 2nd edn. Manning Publications.

[4] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for deep learning in NLP,” in Proc. 57th Annu. Meeting Assoc. Comput. Linguistics (ACL), 2019, pp. 3645–3650. [Online]. Available: https://doi.org/10.1609/aaai.v34i09.7123

[5] European Environment Agency, “CO_2 emission intensity of electricity generation,” EEA, 2023. [Online]. Available: https://www.eea.europa.eu/en/analysis/maps-and-charts/co2-emission-intensity-15

[6] Giri, C. and Chen, Y., 2022. Deep learning for demand forecasting in the fashion and apparel retail industry. Forecasting, 4(2), pp.565–581. Available: https://doi.org/10.3390/forecast4020031

[7] Greenhouse Gas Protocol, A Corporate Accounting and Reporting Standard, Revised Edition, World Resources Institute and World Business Council for Sustainable Development, 2004. [Online]. Available: https://ghgprotocol.org/corporate-standard

[8] H. D. Nguyen, K. P. Tran, S. Thomassey, and M. Hamad, ‘Forecasting and Anomaly Detection approaches using LSTM and LSTM Autoencoder techniques with the applications in supply chain management’, International Journal of Information Management, vol.57, p. 102282, Apr. 2021, Available: https://doi.org/10.1016/j.ijinfomgt.2020.102282

[9] International Organisation for Standardisation, ISO 14064-1:2018 Greenhouse gases – Part 1: Specification with guidance at the organisation level for quantification and reporting of greenhouse gas emissions and removals, Geneva, Switzerland: ISO, 2018. [Online]. Available: https://www.iso.org/standard/66453.html

[10] J. Schmidhuber, “Deep Learning in Fashion Retail,” Journal of Retail Analytics, vol. 17, no. 2, pp. 23–31, 2022.

[11] L Li, X., Shen, J., Wang, D., Lu, W. and Chen, Y., 2024. Multi-modal transform-based fusion model for New product sales forecasting. Engineering Applications of Artificial Intelligence, 133, 108606. [Online]. Available https://doi.org/10.1016/j.engappai.2024.108606

[12] L. Lannelongue, J. Grealey, and M. Inouye, “Green algorithms: Quantifying the carbon footprint of computation,” Patterns, vol. 2, no. 3, p. 100274, Mar. 2021. [Online]. Available: https://doi.org/10.1002/advs.202100707

[13] L. Rokach and O. Maimon, "Top-down induction of decision tree classifiers - a survey," in IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 35, no. 4, pp. 476-487, Nov. 2005, https://doi.org/10.1109/TSMCC.2004.843247

[14] M. Samek et al., “Explainable AI: Interpreting, Explaining and Visualising Deep Learning,” Springer, 2019. https://link.springer.com/book/10.1007/978-3-030-28954-6

[15] M. Strubell, A. Ganesh, and A. McCallum, “Energy and Policy Considerations for Deep Learning in NLP,” ACL, pp. 3645–3650, 2019. Available: https://doi.org/10.1609/aaai.v34i09.7123

[16] P. Negre, R. S. Alonso, J. Prieto, Ó. García, and L. de-la-Fuente-Valentín, “Prediction of footwear demand using Prophet and SARIMA,” Expert Systems with Applications, vol. 255, Part B, Article no. 124512, 2024. [Online]. Available: https://doi.org/10.1016/j.eswa.2024.124512

[17] R. Schwartz et al., “Green AI,” Commun. ACM, vol. 63, no. 12, pp. 54–63, 2020.

[18] S. Anitha and R. Neelakandan, “A Demand Forecasting Model Leveraging Machine Learning to Decode Customer Preferences for New Fashion Products,” Complexity, vol. 2024, Article ID 8425058, 10 pages, 2024. [Online]. Available: https://doi.org/10.1155/2024/8425058

[19] S. M. Lundberg and S. Lee, “A Unified Approach to Interpreting Model Predictions,” Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017. Available: https://dl.acm.org/doi/10.5555/3295222.3295230

[20] S. Taylor and B. Letham, “Forecasting at Scale,” The American Statistician, vol. 72, no. 1, pp. 37–45, 2018. Available: https://doi.org/10.1080/00031305.2017.1380080

[21] Selvaraju, R.R., Cogswell, M., Das, A. et al. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localisation. Int J Comput Vis 128, 336–359 (2020). https://doi.org/10.1007/s11263-019-01228-7

[22] United Nations, “Transforming our World: The 2030 Agenda for Sustainable Development,” 2018. Available: https://www.undp.org/ukraine/publications/transforming-our-world-2030-agenda-sustainable-development

[23] “H&M Sales data,” Kaggle, Feb. 04, 2022. https://www.kaggle.com/datasets/tulasiram574/hm-sales-data