Public

Inertia Forecasting with Regional Extrapolation

NIA2_NGESO048

Final Report

NESO
National Energy System Operator

Contents

1 Executive Summary

While GE Effective Inertia solution has been operationally running since 2021 at NESO, the solution at this stage only provides inertia metering and forecasting estimates for Scotland area, due to a lack of sufficient PMU monitoring equipment installed in the south of the GB area.

This innovation project was undertaken by GE Vernova, as the supplier of the Effective Inertia solution, to explore the possibilities of using the inertia forecasting model built with Scotland data to predict the inertia for England and Wales, and thus for the entire Great Britain (GB) grid.

The initial step of the project involved re-training the original Machine Learning model currently used for Scotland using normalized predictors, to establish a baseline for future comparison with other models investigated in the project.

Neural network models have been initially studied with results compared to the baseline I model. These initial results showed comparable and slight improvement compared to the baseline (Figure 1), indicating that these models could be considered to enhance the generalization power of the model in an attempt to extrapolate inertia forecasting to other regions in the overall GB area.

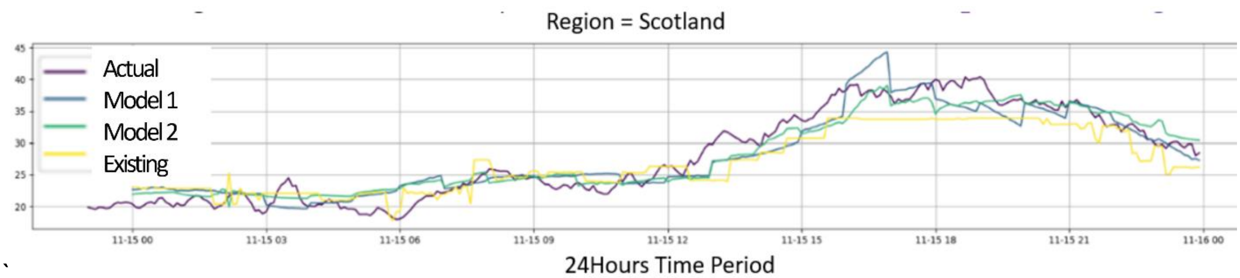

Figure 1 Comparison of Scotland Inertia Predictions produced by Model 1, Model 2 and the existing model against the actual values. X-axis shows 24hours period, 15 Nov 2023

A further model Model 4 has also been studied for the sole purpose of improving inertia forecasting for Scotland area, as this model requires past inertia values as an input to the model and cannot be considered for the extension to the GB area. The study demonstrated significant improvement compared to the baselining (Figure 2), achieving over 70% improvement in RMSE, MAE, and MAPE metrics.

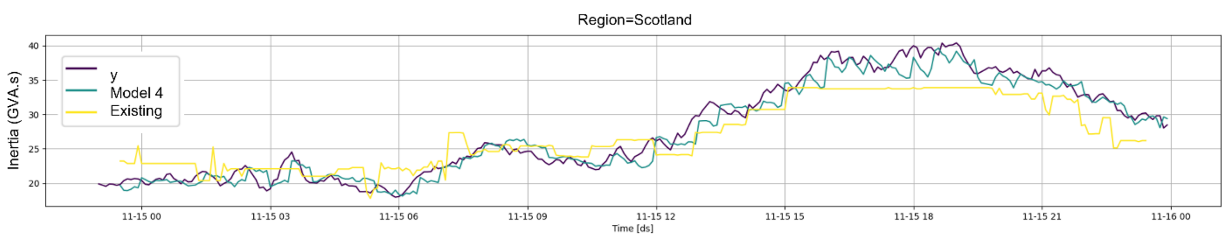

Figure 2 Comparison of Scotland Inertia Predictions produced by Model 4 and the existing ML model against the actual values (marked "y"). X-axis shows 24hours period, 15 Nov 2023

The final step of the study consisted of extending the neural network models built using Scotland data to extend to the other areas of GB (North, South East and South West).

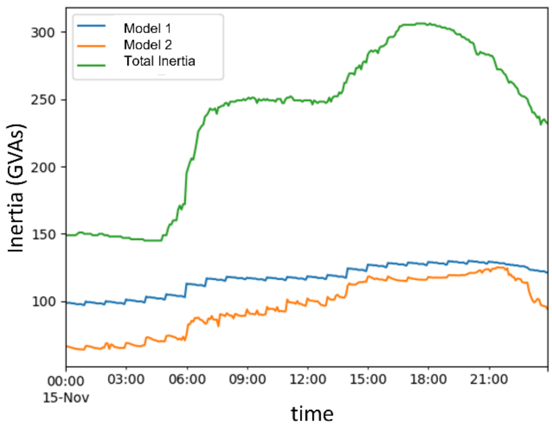

Figure 3 Total GB Inertia : NESO Linear relationship Demand/SumRotating Inertia estimation versus Model 1 /Model 2 total prediction

The findings underscore the challenges of domain adaptation and suggest advanced techniques that could potentially bridge the gap between regional data and GB-wide forecasting accuracy.

It highlighted the complexity of forecasting inertia across diverse regions with varying energy consumption, generation, and grid dynamics. Despite the limitations encountered, the insights gained from this project provide a solid foundation for future work.

For future work, it is recommend to particularly focus on acquiring a more extensive and diverse dataset, incorporating inertia data from other regions, and exploring advanced domain adaptation techniques. By addressing these areas, there is a strong potential to develop a robust, GB-wide inertia forecasting model.

2 Introduction

The GE Effective Inertia solution has been running operationally since 2021 at NESO. The solution provides real-time continuous monitoring of effective inertia metering based on PMU measurements, with Phasor Measurement Units (PMU) currently available for Scotland area only. The solution also applies machine learning techniques, leveraging predictors such as demand, sum of rotating inertia, wind generation, and solar generation to provide an effective inertia forecast for the same area. This established a reliable framework for understanding and predicting effective inertia in Scotland.

Since PMU measurements are not yet fully available in other areas of Great Britain (GB), it is currently not possible to obtain a comprehensive understanding of inertia metering and forecasting behaviour across the entire GB system.

The objective of this consultancy project is to explore the possibilities of using the inertia forecasting model built with Scotland data to predict the inertia for England and Wales, and thus for the entire Great Britain (GB) grid. The project scope also encompassed evaluating alternative approaches and possible enhancements to the existing machine learning algorithm.

Two key datasets have been used for this study:

- A set of predictors for the four GB regions defined as Scotland, North, South East and South West
- Corresponding period of metered inertia data specifically for Scotland area.

The report begins by detailing the re-training process conducted with the Machine Learning model currently used to establish a baseline for the Scotland area. Subsequently, Neural Network models are studied and evaluated to assess their generalization to other areas.

3 Data Description

This section provides a detailed description of the datasets used in this consultancy project, including the predictors dataset and the Scotland inertia dataset. These datasets were combined to create a single comprehensive dataset that served as the basis for model training and evaluation.

3.1 Predictors Dataset

The predictors dataset is composed of three separate files: DemandLive.csv, GenerationLive.csv, and RotatingInertiaLive.csv. These files contain time-series data for four regions: South-West, South-East, North, and Scotland. The data covers two time periods:

- First Set: From 13 Oct 2023 01:00 to 11 Jan 2024 11:05
- Second Set: From 05 Jan 2024 06:25 to 05 Apr 2024 04:55

The second set of files was provided in April 2024 while the consultancy project was underway and has not been included in this initial draft report as metering inertia data available in NESO Prod server could not be access in time for the analysis.

Each file contains the following predictors:

- DemandLive.csv: Electricity demand in megawatts (MW) for each region.
- GenerationLive.csv: Wind and solar generation in MW for each region.
- RotatingInertiaLive.csv: Sum of rotating inertia in MW for each region.

To better understand the differences and distributions of the predictors across the four regions, box plots presented in *Figure 1* report minimum, first quartile, median, third quartile and maximum for each predictor and each region. Values that are more than 1.5 times the interquartile range away from the box are outliers and shown as circles.

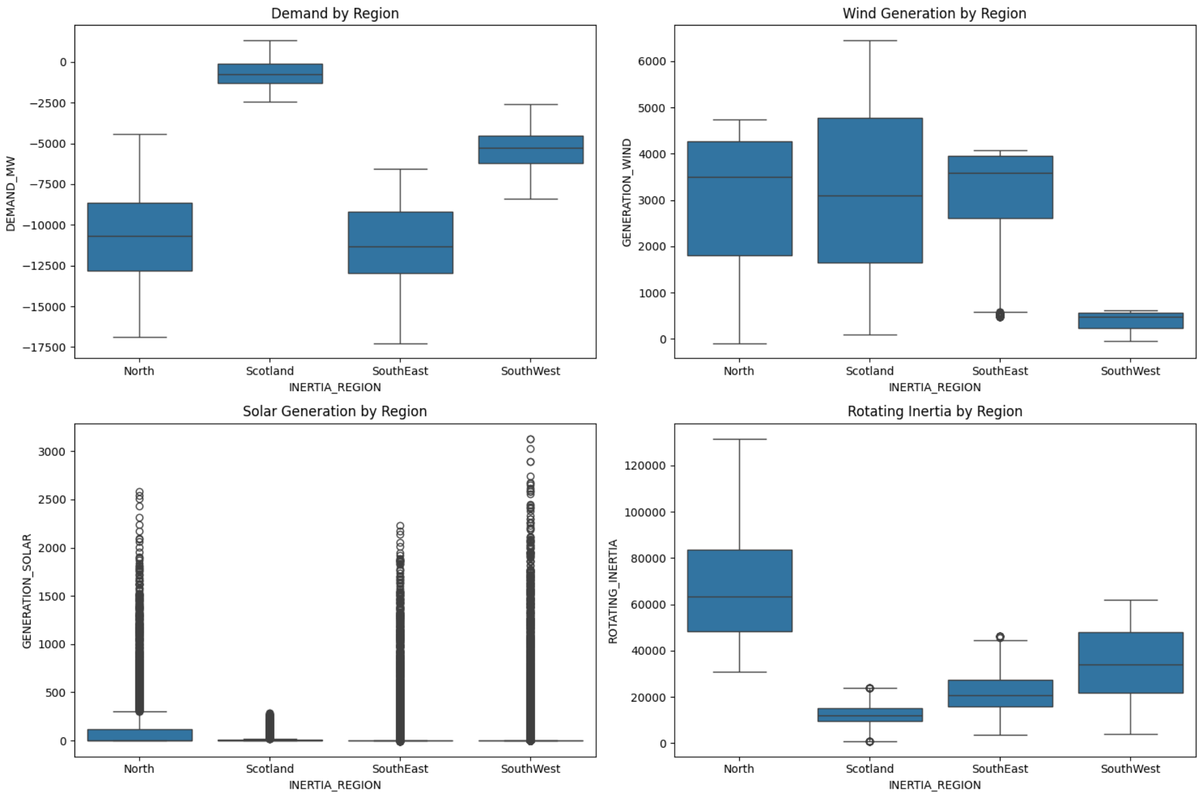

Figure 4: Comparative Analysis of Predictors Across Regions illustrated as a Box Plot

3.1.1 Detailed Observations

Demand (DEMAND_MW):
The demand varies significantly across regions, with the North and South-East regions generally showing higher demand compared to Scotland and South-West.

Wind Generation (GENERATION_WIND):
Wind generation is more variable in Scotland compared to North and South-East regions. The South-West region shows significantly lower wind generation compared to the other regions.

Solar Generation (GENERATION_SOLAR):
Solar generation is minimal or negative in some regions, possibly indicating data recording or measurement issues.

Rotating Inertia (ROTATING_INERTIA):
Rotating inertia values are highest in the North and South-West regions, indicating larger contributions to grid stability from these regions. Scotland and South-East show lower rotating inertia values.

3.2 Scotland Metered Inertia Dataset

Scotland metered effective inertia dataset was extracted from the production server to match the timespan of the 4 regions predictors dataset provided by NESO. The metered inertia values serve as the target variable for model training and validation.

To further understand the relationships between the predictors and the metered inertia for Scotland, a correlation matrix is presented in *Figure 2*, showing the correlation of each separate predictors from all regions against Scotland's inertia.

- A correlation index closer to +1 (reported in red shades) indicates strongest positive correlation (variables increase or decrease in parallel),
- A correlation index closer to -1 (reported in blue shades) indicates strongest negative correlation (as one variable increases, the other decreases or vice versa)
- A correlation index closer to 0 (reported in grey shades) indicates very poor or no correlation.

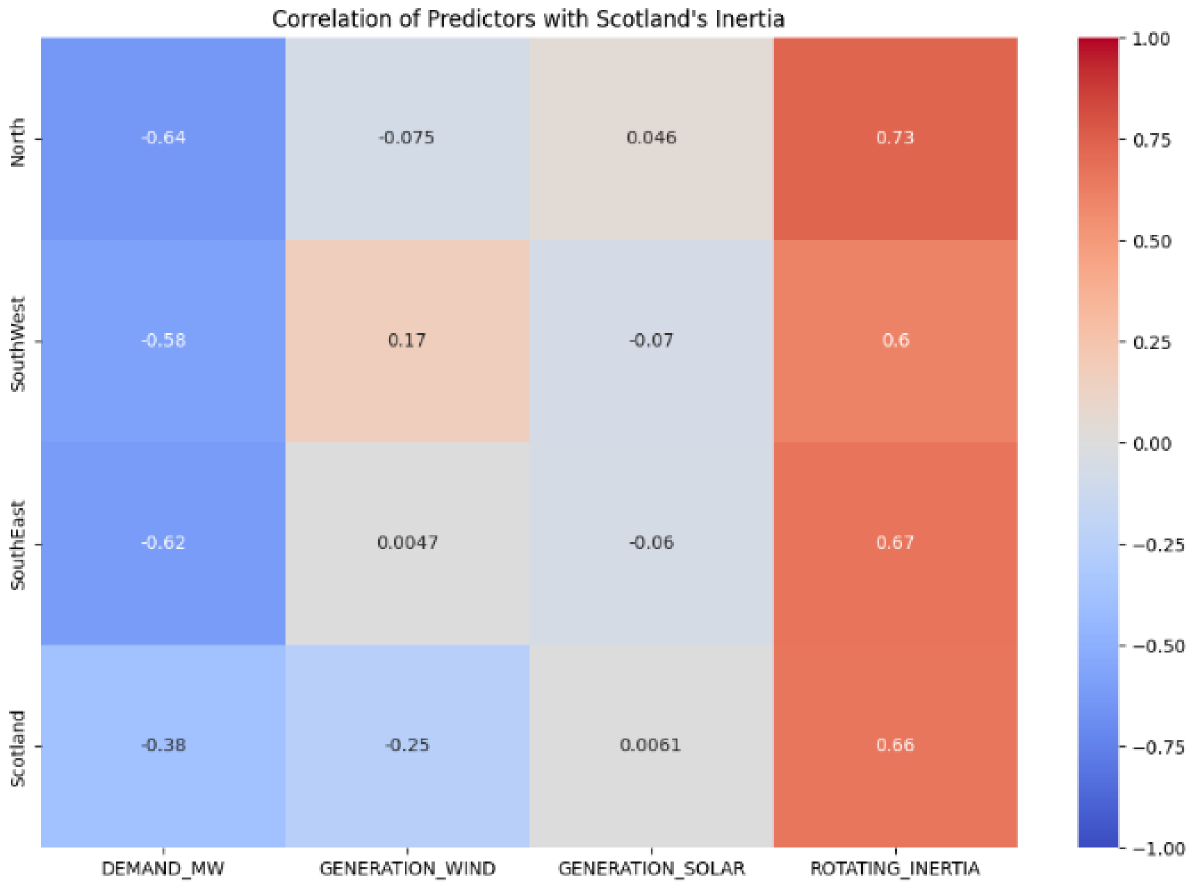

Figure 5: Correlation Heat Map of Region Predictors with Scotland Inertia

The correlation matrix provides several insights:

- Some predictors from other regions show better correlation with Scotland inertia than the predictors from Scotland itself. For example, the rotating inertia in the North region reports the highest correlation with Scotland's inertia (0.734) compared to Scotland's own rotating inertia (0.659).
- Th result suggests that there may be underlying dynamics between regions that influence inertia in ways not immediately apparent from localized data alone.
- While these correlations do not directly confirm the ability to predict inertia for regions individually, they provide a better sense of the interconnected dynamics within the grid.

These observations highlight the complexity of the task and underscore the need for advanced methods to accurately forecast inertia across different regions. Further investigation into these dynamics could provide valuable insights for improving inertia forecasting models.

4 Baseline Method

In this section, we describe the baseline method employed to assess the feasibility of using the existing Machine Learning model, built with Scotland data, to predict the inertia for other regions within the GB system.

4.1 Existing ML Model

Redacted

4.2 Normalization of Predictors

To adapt the existing ML model for our new objective, we normalized the predictors for each region using min-max scaling. This scaling method transforms the data to a fixed range, typically [0, 1], which helps in aligning the distributions of different regions. The normalization ensures that all predictors are in the same range, facilitating a more uniform input space for the model. This is particularly important given the differences in data distributions across regions, as highlighted by the earlier box plots in *Figure 1, Section 0*.

Figure 6 compares the unscaled inputs to the scaled inputs, illustrating the effect of normalization:

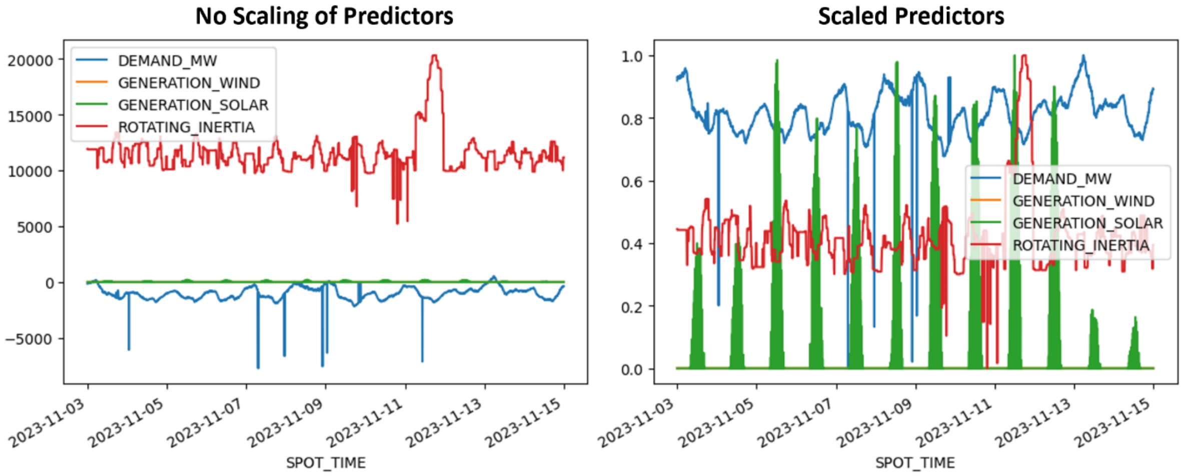

Figure 6: Unscaled (left) and Scaled Predictors (right) - Predictors of Scotland shown as an example

4.3 Extrapolation Results with Scaled Data

After normalizing the predictors, the model has been re-trained using the scaled inputs. The objective was to establish a baseline method that treats all regions data uniformly. Once the model trained, it was fed with the data of other regions. In Figure 7, left plot, we can see that this approach did not yield satisfactory results, as the model struggled to generalize beyond the region it was trained on.

To address this, further normalization of the outputs (Scotland inertia values) using min-max scaling and re-training of the model with both scaled inputs and outputs was performed. This approach aimed to provide a more balanced training process by ensuring both inputs and outputs were within a consistent range. This second trained model result presented in Figure 7, right plot also led to a saturation of the outputs when applied to the other regions. This tends to indicate that the model could not effectively learn the relationships between other regions predictors and the Scotland inertia values once all normalized.

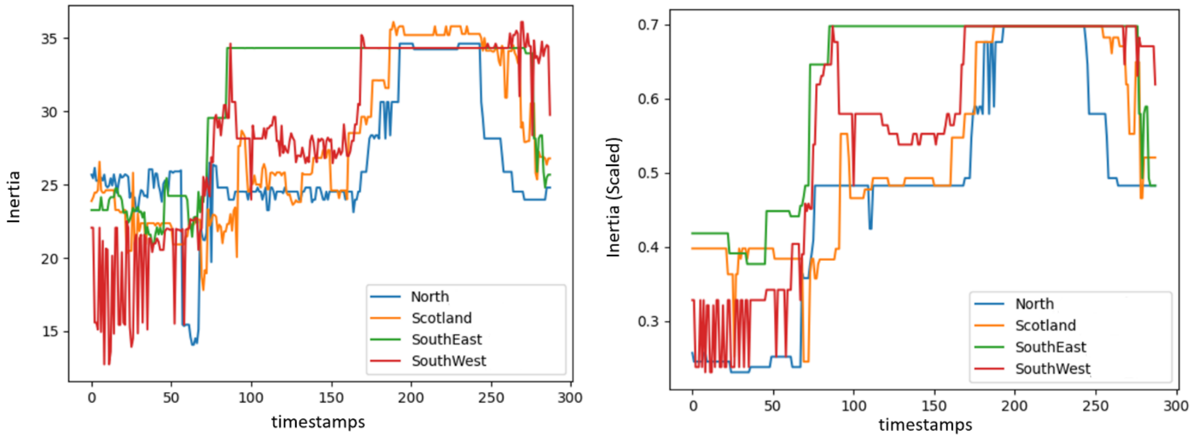

Figure 7: Inertia Extrapolation Results using the original ML model with scaled inputs (left) and scaled inputs & outputs (right) for 15 Nov 2023

The plot clearly illustrates the saturation effect observed when both inputs and outputs are scaled. The model inability to produce accurate forecasts for regions other than Scotland underscores the limitations of using the existing model for this task. This baseline method sets the stage for exploring more advanced neural network models, which are discussed in the subsequent sections.

5 Neural Network Models

Redacted

6 Enhancing Scotland Inertia Forecasting

Motivated by the improvements observed with new approaches over the existing model, we further explored advanced architectures to enhance the accuracy of inertia forecasting for Scotland. This led us to train a new model ML Model 4.

6.1 Model 4 Description

Redacted

6.2 Key Differences and Enhancements

Compared to our previous models, Model 4 incorporates several key enhancements.

Redacted

6.3 Results

The Model demonstrated remarkable improvements in forecasting accuracy compared to the existing model. Table 1 below summarizes the performances using the metrics RMSE (Root Mean Squared Error), MAE (Mean Absolute Error), and MAPE (Mean Absolute Percentage Error).

Metric	Model 4	Existing	Improvement Ratio (Existing:Model 4)	Percentage Improvement
RMSE	0.77	3.04	3.95	74.67%
MAE	0.62	2.54	4.10	75.59%
MAPE	0.023	0.088	3.83	73.86%

Table 1 Performance comparison between Model 4 and Existing model over a 24hour period for Scotland area

The results illustrate that model 4 significantly outperforms the original model across all metrics, achieving over 70% improvement in RMSE, MAE, and MAPE.

The plot below further compares the predictions of the Model 4 and the existing model against the actual inertia values for a 24-hour period. Model 4 predictions closely follow the actual values, while the existing model shows larger deviations.

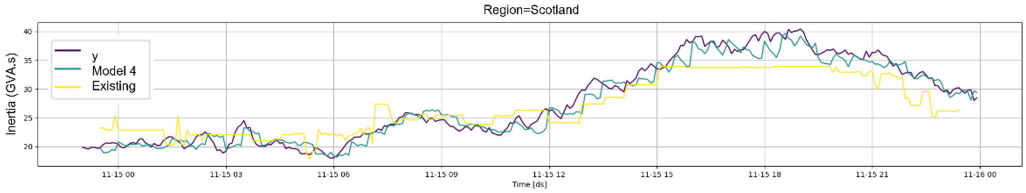

Figure 8 Comparison of Scotland Inertia Predictions produced and Existing model against the actual values (marked "y"). X-axis shows 24hours period, 15 Nov 2023

6.4 Limitations for Extrapolation

Despite its superior performance when applied for Scotland area, Model 4 is not suitable for extrapolation to other regions due to two primary reasons: Redacted.

In conclusion, Model 4 tends to show significant enhancement of inertia forecasting for Scotland area. Its design and normalization techniques however limit its use for extrapolation to other regions.

7 Extrapolation to Other Regions

After training and validating the models on the Scotland inertia data, the next step was to evaluate their performance on South-West, South-East, and North regions. The goal was to determine if the models could generalize from Scotland data the inertia prediction for these regions, thereby contributing to an overall GB inertia forecast.

7.1 Methodology

To assess the model's generalization capabilities, we used the trained models (Model 1, 2) to run inference on the predictors for each of the other regions. The input features for these regions were normalized using the same min-max scaling and standard normalization techniques applied during training as detailed in Section 4.2.

7.2 Results

The results are presented in the figure below showing the predicted inertia values from the models for all four regions. Both models produced very similar outputs for each region, indicating a lack of differentiation in the forecasts. This suggests that the models struggled to capture the unique characteristics and patterns of each region inertia dynamics.

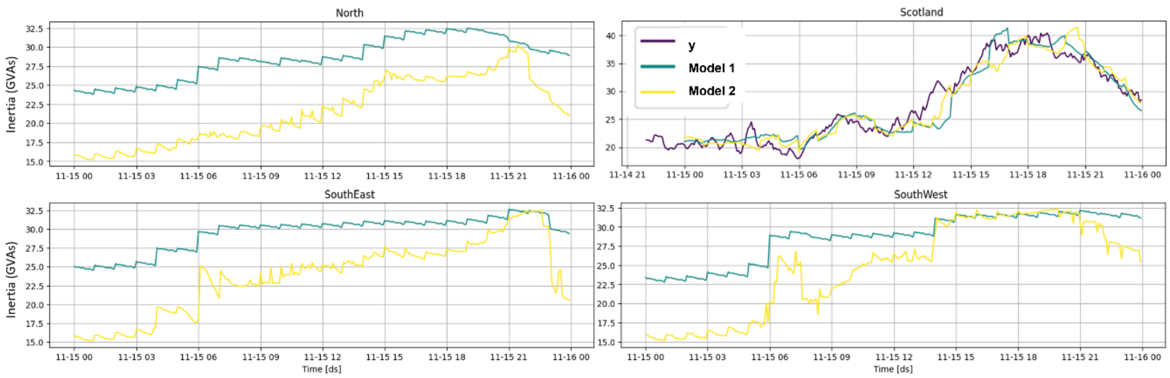

Figure 9: Extrapolation Results of Inertia Prediction on North, Southeast and Southwest, including Scotland predictions for comparison (y label) – 24 hours period 15Nov2023.

7.3 Analysis

The poor performance of the models when applied to regions other than Scotland can be attributed to several factors:

1. **Regional Differences:** Each region has unique characteristics in terms of energy consumption, generation, and grid dynamics, which were not adequately captured by the models trained solely on Scotland data.

2. **Lack of Regional Data:** The models were trained exclusively on Scotland inertia data, which does not encompass the full range of variability present in other regions. This limited training data can result in models that are not robust enough to generalize across different regions.

3. **Normalization and Scaling:** While normalization and scaling help in bringing the data to a common range, they do not address the underlying differences in the data distributions and patterns between regions. This can lead to models that perform similarly across regions without capturing their unique characteristics.

7.4 Comparison to GB Inertia

To further assess the performance, the predicted inertia values from all four regions have been summed to estimate the total inertia for the entire GB area. This total is compared to the total inertia provided by NESO (Figure 10). NESO inertia estimates are calculated using the following linear equation, with a ScalingFactor varying slightly depending on the period:

$$Total\ Inertia\ = TotalRotatingInertia + ScalingFactor\ \times Total\ Demand$$

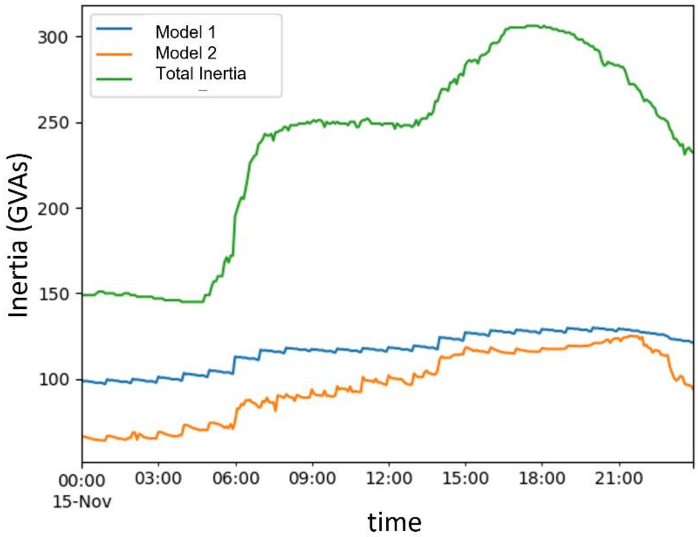

Figure 10:Total GB Inertia : NESO estimation versus Model 1 / Model 2 total prediction

The plot shows a significant discrepancy between the summed predicted total inertia and the provided total inertia. This discrepancy underscores the models inability to generalize from the Scotland data to other regions effectively. The provided total inertia, calculated by linearly scaling the total demand and Sum of Rotating Inertia, serves as a rough estimate rather than an actual measured value. Despite this, the large deviation of our predictions from this estimate indicates that our models are not capturing the underlying dynamics of inertia across the different regions of the UK.

8 Recommendations

Based on our findings, we provide the following recommendations to improve the inertia forecasting models for the entire UK grid. These recommendations are categorized into two main areas: Data and Model Training.

8.1 Data

Extended Dataset
Having access to a more extensive dataset, ideally spanning at least 12 to 15 months, would be beneficial. This extended dataset should cover a full year of training data and additional data for validation. Data collected over different seasons and varying grid conditions can help the model learn more comprehensive patterns and improve its prediction accuracy. This broader dataset will allow the model to better understand the underlying dynamics of the grid under various conditions.

Inertia Data for Other Regions
Even a small amount of inertia data for regions other than Scotland can be highly beneficial. Access to this data would make the problem more manageable and facilitate the use of transfer learning techniques. Transfer learning allows models trained on one region (e.g., Scotland) to be fine-tuned on data from other regions, improving their ability to generalize and adapt to the unique characteristics of each region.

Expanded List of Predictors
The current set of predictors may not fully capture the unique characteristics of each region. Increasing the list of predictors to include additional variables. These additional predictors could help the model learn features that are more representative of each region specific dynamics, leading to better forecasting performance.

8.2 Model Training

Custom Loss Function
Leveraging domain knowledge, such as the relationship between total inertia and total demand (e.g., total inertia = SumRotatingInertia + ScalingFactor * total demand), can be used to create a custom loss function. This custom loss function can guide the model to produce values outside the range it has seen during training, improving its ability to generalize and better match the provided total inertia estimates.
Domain Adaptation via Adversarial Learning

Domain adaptation is a field of machine learning that focuses on adapting models trained in one domain (source domain) to perform well in another domain (target domain). In the context of our problem, the source domain is Scotland, and the target domains are the other regions in the GB.

Our task can be framed as an *unsupervised domain adaptation* task because we lack inertia data for the target regions. The goal is to adapt the model to perform well in the target regions without having explicit inertia labels for those regions.

If a small amount of inertia data for the target regions becomes available, the task transitions into a semi-supervised domain adaptation task. This is more manageable as the model can leverage the small amount of labeled data in the target domain to improve its performance.

One effective approach for domain adaptation is adversarial learning. This technique involves training a model to perform well in the source domain while simultaneously training a discriminator to distinguish between the source and target domains. The goal is to make the model features indistinguishable between the two domains, thereby improving its generalization capabilities. An analogy for this process is a student (feature extractor) trying to learn concepts that are universally applicable, while a teacher (domain discriminator) tries to identify which concepts are specific to a particular context. The student improves by focusing on concepts that are harder for the teacher to classify, thereby learning more generalizable features.

While adversarial learning can significantly improve the model ability to generalize to new regions, it also comes with challenges. Implementing and productionizing such techniques can be complex and may require substantial computational resources and expertise.

9 Conclusion

This consultancy project explored the possibilities for extending inertia forecasting capabilities from a region-specific model for Scotland to a comprehensive model applicable across the entire GB grid.

The initial step in this project was to meticulously retrain the existing model to provide a baseline for model comparisons. The report then presented the exploration of advanced neural network architectures with the aim to overcome the challenges of domain adaptation and enhance the generalization power of our models.

The study of model 4 demonstrated significant improvement in the accuracy of inertia forecasting for Scotland. The exploration of the Model 4 model incorporates previous inertia values as input model and employs advanced normalization techniques . It showed substantial improvements over the traditional existing model, achieving over 70% enhancement in key performance metrics. While this model could be considered as an enhancement for the area inertia forecasting, this technique cannot be considered for the extension of the model to other areas due to the added previous inertia values input.

The report finally presented results of extending the models previously selected to other GB areas. The overall GB inertia extended prediction results were compared with estimates provided by NESO, calculated using a linear relationship between total sum of rotating inertia and demand. The results did not show successful performances of the extrapolation of the Scottish inertia model to the other GB areas when compared to the expected overall GB inertia provided by NESO (linear combination of total Demand and Sum of Rotating Inertia).

The findings highlighted the complexity of forecasting inertia across diverse regions with varying energy consumption, generation, and grid dynamics. Despite the limitations encountered, the insights gained from this project provide a solid foundation for future work.

For future work, it is recommend to particularly focus on acquiring a more extensive and diverse dataset, incorporating inertia data from other regions, and exploring advanced domain adaptation techniques. By addressing these areas, there is a strong potential to develop a robust, GB-wide inertia forecasting model.