Abel Kempynck
abel.kempynck@gmail.com
LinkedIn profile
+32 472491381

Forecasting DFR prices using a hybrid statistical and machine learning framework

UK Electrical Grid Modelling Hackathon

Group 33

07/04/2025

Contents

1 Background Research

This section explores the fundamental dynamics of Dynamic Frequency Response (DFR) services, focusing on the distinctions between high and low services, market influences, and their correlation with day-ahead prices. This background research was essential for understanding the complexity of DFR markets and informed the modelling choices in this study.

Differentiating high and low DFR Services

DFR services are critical for maintaining grid stability by responding to deviations in system frequency. The high and low variants of these services, such as Dynamic Containment Low (DCL) and Dynamic Containment High (DCH), are designed to address different frequency challenges:

- **DCL -** service responds to large downward frequency deviations, specifically up to -0.8 Hz. It requires significant energy availability, making it suitable for battery energy storage systems (BESS) that can provide substantial volumes. DCL is characterized by less frequent activations, which is advantageous for battery longevity due to reduced cycling and degradation [1]. The lower utilization rate of DCL results in easier compliance and lower risk, making it a safer and more straightforward option for delivery. Consequently, DCL has the highest traded volume among DFR services. Non-ABSVD (Applicable Balancing Services Volume Data) units are particularly active in DCL as they receive payments for discharging energy, unlike ABSVD units, which do not receive such compensation.
- **DCH -** addresses large upward frequency deviations, up to +0.5 Hz, and has lower volume requirements compared to DCL. However, it can be triggered more frequently during sudden increases in grid frequency, such as those caused by a sudden loss of demand, when the system is long. DCH sees significant activity from ABSVD units, which can accept negative prices because they are not required to pay for the energy they take off the system when contracted into Dx high services. These units later profit from selling energy in wholesale markets, whereas non-ABSVD units must pay for the energy they import.

Market dynamics and the impact of changes

The introduction of the Enduring Auction Capability (EAC) marked a significant shift in the DFR market dynamics. EAC allows for the co-optimization of multiple DFR services and the inclusion of negative pricing, fundamentally altering market behavior. This change represents a clear trend break, suggesting that excluding pre-EAC data from modelling efforts could enhance predictive accuracy by focusing on the more relevant post-EAC market conditions.

Correlation between day-ahead and DFR prices

The relationship between day-ahead (wholesale) prices and DFR prices is multifaceted, driven by broader market interactions and opportunity costs:

- **Market interaction and signals -** Although units committed to DFR contracts cannot directly participate in the day-ahead market, the overall supply and demand dynamics influence both markets. High demand periods in the wholesale market can lead to elevated prices, which in turn affect bidding strategies in DFR auctions as participants anticipate similar trends.
- **Opportunity cost -** The opportunity cost of not participating in the wholesale market influences DFR bidding behavior. During periods of high wholesale prices, DFR providers may bid higher in DFR auctions to compensate for the lost opportunity of selling in the wholesale market. However, you don't need to commit your full capacity to one market, allowing you to hedge your bets.

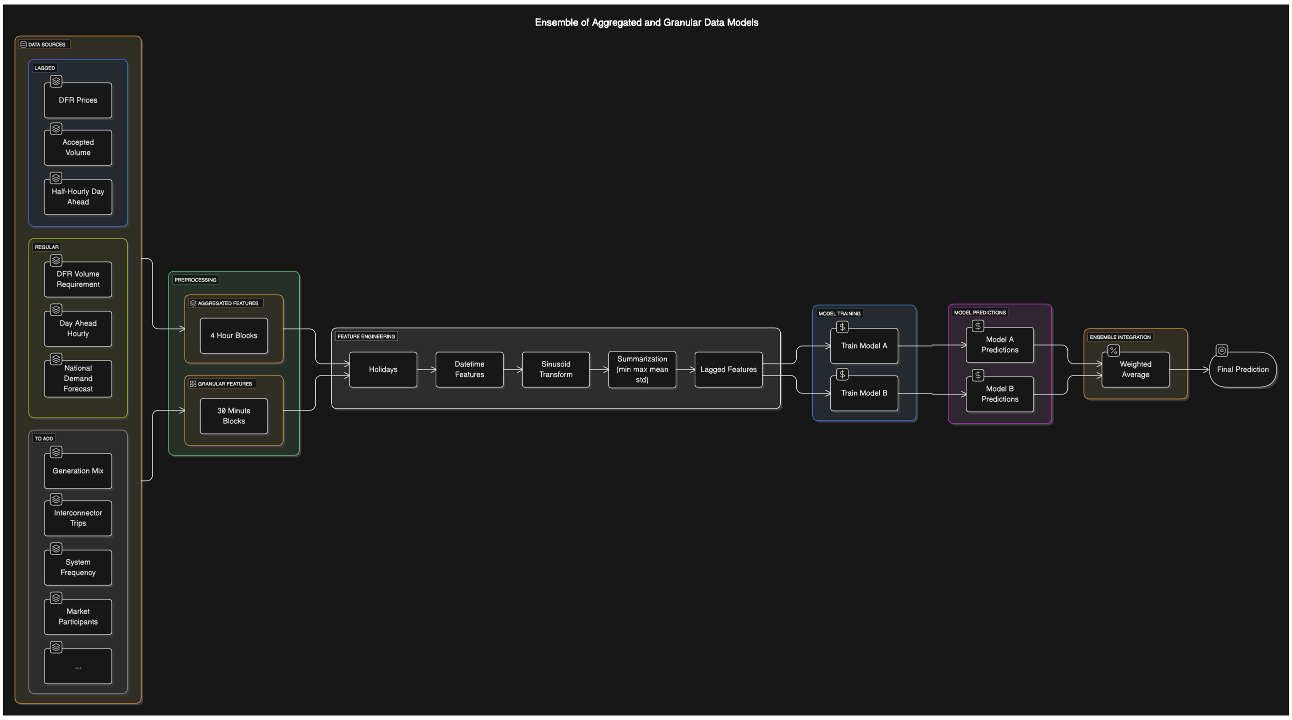

Figure 1: Model pipeline diagram

2 The Solution

General

This study aims to predict the DCL price, Y, for a 24-hour horizon in the future, structured into six 4-hour EFA blocks. The predictive forecasting model f estimates $\hat{Y} = f(X)$, where $X = \{X_1, X_2, X_3\}$ is a comprehensive feature set:

- X_1: Lagged DFR prices and feature-engineered attributes.
- X_2: Day-ahead hourly prices and related features.
- X_3: Lagged day-ahead half-hourly prices and related features.

The model f minimizes the loss function $L(Y, \hat{Y})$, to accurately forecast the DCL price.

Model pipeline

The model pipeline, illustrated in Figure 1, visualizes the idea on how to forecast the DCL price. This section will highlight the major components and provide the rationale behind the approach taken in this study.

Data sources

This study includes all available datasets provided at the start of the hackathon.

Additionally, upon examining the various auction closings and openings, an opportunity was identified to enhance forecast accuracy. Half-hourly day ahead auction results are unavailable at forecast time, necessitating the use of lagged variables. The market's 48 half-hour blocks align with conventional hours, unlike DFR EFA blocks, allowing the first two blocks to use actual values, with subsequent blocks relying on lagged data and thus highlighting a potential area for improvement given significant BESS participation.

Training

In the development of the forecasting framework, a dual-model approach was employed to capture both the macro and micro dynamics influencing DFR prices. As illustrated in Figure 1, the model pipeline is designed to integrate both the aggregated 4-hour EFA data and granular day-ahead data perspectives.

Model A

This model was trained using aggregated day-ahead hourly and half-hourly data, consolidated into 4-hour EFA blocks. The objective was to discern broader market trends and patterns that align with the temporal structure of DFR service procurement. By aggregating the data, Model A is designed to identify overarching market behaviours that influence DFR pricing over extended intervals, providing insights into long-term price determinants.

Model A employs an ARIMA statistical model with exogenous variables and L1 (Lasso) regularization. The model was trained on data spanning the first 22 months (1 year and 10 months) and evaluated on the subsequent two months. Missing values within the training and test datasets were imputed using the mean, followed by standardization.

Model B

This model was trained on granular half-hourly data basis, with the DFR data expanded to match this finer temporal resolution. Model B focuses on capturing short-term market fluctuations and detailed patterns that occur within shorter time frames. This granular approach enables the model to respond to rapid changes in market conditions, offering a nuanced understanding of the factors influencing DFR prices on a more immediate basis.

Model B utilizes a LightGBM algorithm, incorporating L1 (Lasso) and cross-validation with three folds for regularization. This model was trained on the same 22-month dataset as Model A and evaluated over the final two months of 2024. Predictions were initially generated at a 30-minute granularity and subsequently averaged over eight timeslots to align with the 4-hour EFA block time dimension. Performance was assessed using the Mean Absolute Percentage Error (MAPE) metric.

Weighted average

To leverage the strengths of both classes of models (statistical and machine learning) and hedge against model bias, a weighted average ensemble approach was adopted. Predictions from Models A and B were combined using a weighted average, where Model B, operating at a 30-minute granularity, was assigned a weight four times that of Model A. Specifically, the weights were set as 4/5 for Model B and 1/5 for Model A.

Naive benchmark

To contextualize the performance of the ensemble model, a naive forecaster was developed as a benchmark. This simple model uses the previous day's prices as the forecast for the next day, serving as a baseline to evaluate the improvements offered by the more sophisticated modelling approaches.

Feature overview

The study has carefully selected and engineered several features to capture the nuanced dynamics of the energy market, see Figure 2. Below, the study highlights a few critical features and explain the rationale behind their inclusion:

- **Summarized** - Summarization of trading day (Hourly day ahead) or previous trading day (DFR and Half-hourly day ahead). It includes minimum, maximum, mean, and standard deviation of all settlement periods of the trading day or previous trading day. Logic: Traders will be thinking

Feature Source[1]	Model A	Model B	Lagged Indicator	Lags	Summarized	Next/Previous EFA Block
DFR Prices	✔	✔	✔	**1-to 7-day**	✔	✔
DFR Accepted volume	✔	✔	✔	1-to 7-day		
Day ahead half-hourly prices	✔	✔	✔	1-and 2-day		
Day ahead hourly prices	✔	✔	✔	1 (+ original not lagged)	✔	
Volume Requirements Forecast	✔	✔				
National Demand Forecast	✔	✔				
Additional						
Datetime - Day, Month, Year - Hour, halfhour - Holidays	✔	✔				
Sinusoid - Week in month - Day in week - Hour of day - Half-hour of day - EFA_block	✔	✔				
Total features	112	137				

[1]Due to the application of L1 regularization (Lasso), not all features may be included in the final model, as some coefficients may be shrunk to zero.

Figure 2: Feature overview

about opportunity costs when setting their DFR prices, e.g., *"what's the highest export price I'm missing out on?"*

- **Sinusoid** - Represents seasonal patterns and cyclical effects in some features that are highly cyclical, e.g., hour of day, efa_block, week in month… These transformations enable the model to recognize the proximity between certain values such as hour 23 and hour 0. Additionally, encoding these features is crucial as it can significantly affect the convergence rate of the algorithm and the model's performance.
- **Next/Previous EFA Block** - Traders must anticipate the costs of ramping up or down their energy production or consumption to meet the required response levels, i.e., Response Energy Volume (REV). This involves considering the price dynamics of both the preceding and subsequent EFA blocks.

Evaluation

In assessing the performance of the forecasting models, the study compared the Mean Squared Error (MSE) of the ensemble model against a naive predictor. The ensemble model achieved an MSE of approximately 19.7. In contrast, the naive predictor yielded an MSE of about 20.1. This comparison indicates that the ensemble model provides a roughly 2% improvement in predictive accuracy over the naive approach which will ultimately lead to more efficient resource allocation and greater profits to BESS asset owners.

3 Next Steps

In the pursuit of refining my Dx price forecasting model, the following section outlines a few next steps which I believe have great potential to enhance predictive accuracy and model robustness. If I were given more time (and I'm very happy to continue working on this 😉), I would prioritize my efforts as follows (1: top priority):

1. Transforming dataset structure

To improve my forecasting model, I propose transforming the dataset from its current "wide" format to a "long" format. This transformation will increase the number of data points available for analysis, which is crucial given the limited historical data spanning only two years. Furthermore, the introduction of the EAC in November 2023 marked a clear trend break, and I hypothesize that this requires an exclusion of pre-November 2023 data due to its fundamentally different market dynamics. By using the `pandas.melt()` function, we can unpivot the DataFrame, converting Dx price columns into rows, thereby expanding the dataset size by a factor of six.

2. Integrating wind generation forecasts

Wind generation forecasts are pivotal in anticipating grid stability challenges due to the intermittent nature of wind energy. Accurate wind forecasts enable better planning for frequency regulation services, as they directly impact the need for Dx ancillary services. By incorporating wind generation forecasts into my model, I can more effectively predict fluctuations in supply and the corresponding demand for DFR services, thereby enhancing the model's responsiveness to real-time grid conditions.

3. Calculating residual demand and monitoring EAC auction changes

Residual demand, calculated as the difference between total demand and wind generation forecasts, represents the net load that must be met by dispatchable generation and ancillary services. I believe this derived feature is crucial for planning frequency response needs, as it reflects the true variability that the grid must manage. Additionally, monitoring changes in forecasts between the day-ahead and EAC auctions will provide insights into updated expectations of supply and demand, influencing the pricing and requirement for DFR services.

4. Modelling interconnector trips

Interconnector trips have a substantial impact on grid frequency, often causing deviations close to the statutory limits and necessitating the deployment of ancillary services such as DC. Recent articles by Charlotte Johnson on LinkedIn highlight the significance of these events in frequency regulation [2] [3] [4]. To model this phenomenon, I propose using a Poisson distribution to represent the days since the last interconnector trip.

Conclusion

The solution presented in this report serves as a proof of concept for a dual-model approach to forecasting DFR prices. Initial evaluations suggest that this ensemble methodology slightly outperforms a naive forecasting model, demonstrating the potential of integrating both aggregated and granular data perspectives.

I'm enthusiastic about the opportunity to further develop and refine this model, exploring my proposed next steps to enhance its predictive capabilities. I hope that the insights and methodologies presented in this report contribute to the ongoing efforts to optimize grid stability and support the integration of renewable energy sources.

I'm grateful for the opportunity to participate and contribute to this important field, and I wish to extend my sincere gratitude to LCP Delta for providing such an engaging and thought-provoking challenge!

References

[1] S. M. Malhotra, *Bess degradation*, 2025. [Online]. Available: https://www.linkedin.com/posts/shivam-malhotra_bess-fleet-degradation-of-cycles-vs-max-activity-7302674586301521921-Xafn/?utm_source=share&utm_medium=member_desktop&rcm=ACoAABWlKa8BtpOv9qyq9IDL08yU2bMyj72uNNw.

[2] C. Johnson, *Batteries and frequency activity*, 2024. [Online]. Available: https://www.linkedin.com/posts/charlotte-johnson-358790b3_batteries-frequency-activity-7291024395349688320-Gq9P?utm_source=share&utm_medium=member_desktop&rcm=ACoAABWlKa8BtpOv9qyq9IDL08yU2bMyj72uNNw.

[3] C. Johnson, *Frequency response and interconnector trips*, 2024. [Online]. Available: https://www.linkedin.com/posts/charlotte-johnson-358790b3_batteries-frequency-activity-7249392999476920320-3P35?utm_source=share&utm_medium=member_desktop&rcm=ACoAABWlKa8BtpOv9qyq9IDL08yU2bMyj72uNNw.

[4] C. Johnson, *Frequency response and interconnector trips*, 2024. [Online]. Available: https://www.linkedin.com/posts/charlotte-johnson-358790b3_batteries-frequency-activity-7211325347231862784-Gtet?utm_source=share&utm_medium=member_desktop&rcm=ACoAABWlKa8BtpOv9qyq9IDL08yU2bMyj72uNNw.