NIA Project Close Down Report Document

Notes on Completion: Please refer to the appropriate NIA Governance Document to assist in the completion of this form. The full completed submission should not exceed 6 pages in total. Network Licensees must publish the required Project Progress information on the Smarter Networks Portal by 31st July 2014 and each year thereafter. The Network Licensee(s) must publish Project Progress information for each NIA Project that has developed new learning in the preceding relevant year.

Date of Submission

Jul 2024

Project Reference Number

NIA2_NGESO003

Project Progress

Project Title

Probabilistic Machine Learning Solution for Dynamic Reserve Setting

Project Reference Number

NIA2_NGESO003

Funding Licensee(s)

NG ESO - National Grid ESO

Project Start Date

May 2021

Project Duration

1 year and 0 months

Nominated Project Contact(s)

Gabriel Griffin-Booth

Scope

Dynamic Reserve Level Setting Approach

- The different types of reserve setting processes should be integrated to ensure that the total reserve held is representative of NGESO's risk appetite.

- In this project, we expect to apply the approach summarised below to both the basic reserve (possibly including interconnectors) and the reserve for renewable energy sources and we expect that our focus will be on 4–24-hour lead times.

Dynamic Reserve Setting Solution Design

This project will create a proof of concept for a DRS solution with the following features:

- Data pipeline which automates the extraction, cleaning, and preparation of raw data into storage.

- Probabilistic ML (machine learning) model which makes use of predictor variables in that data (e.g. temperature and wind forecast quantiles, and generation mix) to create more accurate predictions (and prediction intervals) of forecast errors and therefore set reserve levels, which better reflect NGESO's risk appetite.

- A database and dashboard for the display of the results of the ML model.

- Automated upload of the ML model results to control room systems.

- Retraining of the ML model to enable a cycle of continuous learning using new data about recent system conditions and forecast errors.

Objectives

- Probabilistic ML model which makes use of predictor variables in that data (e.g. temperature and wind forecast quantiles, and generation mix) to create more accurate predictions (and prediction intervals) of forecast errors and therefore set reserve levels, which better reflect NGESO's risk appetite.

- Retraining of the ML model to enable a cycle of continuous learning using new data about recent system conditions and forecast errors.

Success Criteria

Optimising reserve levels and reducing balancing costs is a key challenge for the NGESO as we move towards our net zero targets and a more volatile system. Uncertainty is more difficult to predict in the short term and we are seeking to optimise reserve levels in the most efficient way.

The end result would be better optimised levels, avoiding 'overholding' and better value for end consumer, and avoiding 'underholding' in risky and uncertain periods.

Performance Compared to the Original Project Aims, Objectives and Success Criteria

National Grid Electricity System Operator ("NGESO") has endeavoured to prepare the published report ("Report") in respect of Probabilistic Machine Learning Solution for Dynamic Reserve Setting, NIA2_NGESO003 ("Project") in a manner which is, as far as possible, objective, using information collected and compiled by NG and its Project partners ("Publishers"). Any intellectual property rights developed in the course of the Project and used in the Report shall be owned by the Publishers (as agreed between NG and the Project partners).

Project Overview

Reserves is a term used to describe spare MWs which are synced ahead of time and are used to correct for changes to the volume of demand and supply between a forecast and real-time. Currently, the volume of reserves which is scheduled on the system is determined by a calculation which is performed bi-annually, once at BST clock-change and once at GMT clock change. This project aimed to develop a probabilistic machine learning model for allocating reserve dynamically based on the latest real-time and forecast system data. Reserve levels must be sized efficiently such that there is always an appropriate volume of reserve on the system. In this definition, appropriate is defined such that there is always enough reserve to operate the system according to ESOs risk appetite whilst also minimising excess reserve holding where possible, thus minimising costs.

The project specifically looked at applying ML techniques to the problem of forecasting operational reserves. One key requirement from the control room for any reserve's recommendation methodology is that there must be some accompanying explanation for why the model has suggested the reserve levels. Therefore, the Dynamic Reserve Setting project aimed to deliver a set of explainable ML models which use the latest real-time and forecast data to recommend reserve which meets ESOs risk appetite.

The work for the Dynamic Reserve Setting project can be broken down into three parts: creating the robust data processing pipelines on ESO systems, improving the models built in the PoC phase of the project and evaluating them against current ESO methods, and building a dashboard for end users to use the model's outputs.

Part 1: Data Pipelines

In this phase of the project up to date data was supplied to the Smith Institute and processed. The data was downloaded from multiple ESO systems, each data source was cleaned to remove any outliers and erroneous data points, missing data was handled appropriately, every data source was aligned temporally and spatially, and finally input into a MySQL database. In parallel with this a PostgreSQL database was built on ESO systems and pipelines were built to feed this database with processed data directly from the source systems.

Part 2: Modelling

Modelling improvements began alongside updating the data as historic data was available from the PoCt phase of the project. These improvements took the form of:

- updating the target variable
- using different methods of feature selection
- building multiple models for different lead times
- hyperparameter tuning.

Updating the target variable: reserves are used to provide resilience against all sources of error and uncertainty which affect system operation at some lead time. To know how much reserve to hold we must first be able to quantify the uncertainty in system operation at

any given lead time. In general, there are four contributions to system level uncertainty which we can easily measure: solar forecast error, wind forecast error, demand forecast error, and power plant unreliability (either changes to Physical Notifications by real-time or under, or over, delivery in real-time or both). Previously, power plant uncertainty was not considered as the data to remove trips from the calculation was not available, however trips have now been removed. Real-time trading data has also been used to remove any trades from power plant uncertainty as this capacity cannot be used for reserve if it has been traded.

Different methods of feature selection: Initially the model for the PoC was built using a bottom-up approach. This consisted of manually selecting features that Subject Matter Experts (SMEs) highlighted as impacting the errors, these included solar radiance for solar forecast error and wind speed for wind forecast error. Once a core set of features was established, features could be added to see their impact on the model and following this process, the model was refined. Top-down modelling was then used to refine the model further. Using this approach, cross correlations between all features are calculated and any with high cross correlation are removed in order to reduce the number of features which describe the same thing. All features are then used and ranked based on their contribution to the result, the features that are lesser used are then removed until all features pass selection criteria and the results compared.

Building multiple models for different lead times: The above process was repeated several times to create different models. It was noticed that different models were needed for different needs for reserve setting, especially with regards to the balancing reserve service. A longer-term model was required to inform the market of the general direction of the market 5 days ahead, this was tuned to only use temporal features which have the largest impact on reserve. For balancing reserve, a model was tuned to be run at 6am every day to give results for the day ahead auction, as the auction takes place early in the morning. In addition, two more models were trained, one at 11am to provide the day ahead view and one that is tuned for 6 hours ahead of real-time to allow control room operators to plan around contingency reserve. As before, SHAP values are provided with all models to explain how the algorithm came to it's conclusion and which features contributed the most.

Hyperparameter tuning: Hyperparameter tuning is the method used to fine tune a model once all other major features have been established. This was done with the use of automatic tuning utilities initially to narrow down the solution, once these gave an acceptable result manual intervention was used to refine them fully.

Part 3: Dashboard

A dashboard has been built in PowerBI, a business supported dashboard software. The dashboard displays the outputs of all the models to users in real-time and allows them to select a model, a percentage expectation of shortfall (both positive and negative) and a time horizon. There are multiple options to adjust how the results are displayed, switching between graph and table view, cardinal points and settlement periods and an option to overlay the national demand forecast. The dashboard shows the SHAP values for each point in time, allowing the user to see which features contributed to the reserve number. This can be shown as on a per feature basis or aggregating features into groups such as solar or wind features. End users were liaised with at various points in this process to help refine the dashboard and gather feedback.

Results

The dynamic reserve setting model has been shown to outperform the existing approach to setting reserves as it is better able to flex reserve based on real-time data and short-term forecasts. Overall, the volume of reserve recommended is 380MW lower than the existing approach for upwards reserve and performs equally for downwards reserve; however, in any given settlement period the reserve recommendations better follow the ESO's risk appetite.

A well-structured, documented, and production-ready suite of software for training and running the probabilistic ML models for setting both upwards and downwards reserve levels was delivered. The approach adopted and the quality of the deliverables produced in the course of the project has ensured the models can be implemented into the business quickly and easily. The pipelines and the database built allow easy integration into ESO systems. Further areas have been identified that the project can help with including validating improvements to demand forecasts (as an improved demand forecast should lower the reserves held).

Project Review

- The Dynamic Reserve Setting project can be considered a success for the following reasons.

- It demonstrated that it is possible to build and deploy an explainable probabilistic machine learning model which can be used to set operational reserves. This model makes use of the latest real-time and forecast data so that the reserve recommendations are demonstrably better at meeting the company risk appetite when compared with the current approach for setting reserves.

- It delivered a software suite which is compliant with internal IT systems and is deployable and fully enables both real-time operation and periodic re-training of the ML models

- It delivered extract, transform, load pipelines for transforming the data held within ESO systems into a cleaned format stored in a PostgreSQL database for implementation into the ESO.

Project Dissemination

- https://www.current-news.co.uk/news/national-grid-eso-eyes-dynamic-day-ahead-reserve-setting-in-smith-institute-project
- https://www.nationalgrideso.com/news/national-grid-eso-and-smith-institute-begin-industry-pioneering-ynamic-reserve-setting-drs

- https://www.smithinst.co.uk/insights/national-grid-eso-and-smith-institute-begin-pioneering-drs-project/
- The Dynamic Reserve Setting project was presented at the 2021 Energy Networks Innovation Conference.
- Dynamic Reserve Setting innovation project seeks to reduce balancing costs | ESO (nationalgrideso.com)
- Data IQ Shortlist Awards 2023

All project reports are available on the Smarter Networks Portal.

Required Modifications to the Planned Approach During the Course of the Project

Changes to programme

The project was extended from the end of financial year 2022 until June 2023 as the following work packages were identified:

- Completing outstanding changes and feature requests, including documentation and dashboard updates as required, following NG feedback
- Support for, and contributing to, integration and end-to-end testing
- Further enhancements and improvements, including:
-Unique models per (some) lead times
-Short term forecasting
-Initial exploration of what could be done to incorporate free reserve

These work packages were identified towards the end of the project as a value add, the feedback received was broadly positive. It was decided it was best to leverage the existing partner's expertise for testing and additional development as they were best placed to deliver the enhancements that were desired.

Changes to cost

The additional work packages described above cost an additional £417k.
Next steps are currently being scoped and will include:
The introduction of solar PV data from the Solar PV Nowcasting NIA project which has cost £425k, this was a change request to the initial project extending the scope to include the work listed in the OCF phase listed below.

A follow up project has been approved internally to include adding regional functionality to DRS to suggest in which region of the UK reserves will be needed.

Lessons Learnt for Future Projects

A key lesson learnt was that liaising with control room staff to gather requirements and feedback is difficult. Control staff work on a shift pattern whereas the rest of the members of the project team worked office hours. This was exacerbated by the fact that control room staff are often too busy to join meetings when they are working in the control room. For this reason, it was determined that a single point of contact with control staff should be established, this took the form of a control engineer whom had been seconded to an office based team. This member of staff worked one shift a week across all shifts and so was able to gather feedback from control staff and deliver feedback to the project.

A further lesson learnt was how the control staff interact with the data that is provided to them. How the data is presented on a dashboard is directly related to how useful it is in a decision-making scenario. There are many different opinions amongst control staff so many options to customise the data display were given to make the dashboard have as much utility as possible. Many lessons have been learnt about the factors that effect reserve levels; tests were performed with various different data sources and the results measured. Some tests were successful such as using perfect knowledge of system imbalance increasing both risk performance and number of MW saved and some were unsuccessful producing no benefit i.e. using European weather data.

Note: The following sections are only required for those projects which have been completed since 1st April 2013, or since the previous Project Progress information was reported.

The Outcomes of the Project

The Dynamic Reserve Setting has achieved the following:

- Built an ETL pipelines for cleaning data from ESO systems to be used in reserves forecasting
- Built a database to store the cleaned data and forecast predictions on ESO systems

- Delivered a software suite which can run and re-train several probabilistic ML model for reserves levels

- Demonstrated that Machine Learning can be used to make better predictions of reserves when compared to the current approach

- Successfully used SHAP values to provide explanations for each prediction made by the

- Machine Learning model and these explanations were well received by the control room engineers

- Created a dashboard to display the models outputs to be used by control engineers to inform their decision making about reserves in real time.

Data Access

Details on how network or consumption data arising in the course of NIA funded projects can be requested by interested parties, and the terms on which such data will be made available by National Grid can be found in our publicly available "Data sharing policy related to NIC/NIA projects" and www.nationalgrideso.com/innovation.

National Grid Electricity System Operator already publishes much of the data arising from our NIC/NIA projects at www.smarternetworks.org. You may wish to check this website before making an application under this policy, in case the data which you are seeking has already been published.

Foreground IPR

The final report from Phase 1 of the project is published on the Smarter Networks Portal. Further reports will be added when available. A final report is currently in production and will be published shortly on the Smarter Networks Portal.

Foreground IPR generated in the course of the project to date:

Produced in DRS Phase 1:

- Code to process relevant data into a format appropriate for ESO GB-wide reserve setting modelling
- Code to train and execute probabilistic machine learning models for use with ESO GB-wide reserve setting
- Code to generate explanations for ESO GB-wide reserve recommendations generated from reserve setting models
- Definitions of errors used to define appropriate measures of ESO GB-wide energy reserve holdings, as-well as code to compute these errors.

Produced in DRS Phase 2:

- Code to perform feature selection for ESO GB-wide reserve setting models
- Code to perform expanding window cross validation to assess ESO GB-wide reserve setting model performance
- A methodology & code to map a given reserve holding (in MWs) to a given risk appetite (quantile) for ESO GB-wide dynamic reserve setting
- A dashboard for visualisation of ESO GB-wide reserve recommendations & associated explanations
- Updated code for data-processing of data required for the ESO GB-wide dynamic reserve setting approach (incorporation of generator trips, logging, and the code updates on ESO's side)
- A methodology & code to assess drift in features & targets variables for the ESO GB-wide reserve setting model
- Updates to the definitions of errors used to define appropriate measures of required GB-wide energy reserve holdings for ESO, as-well as code to compute these errors.
- A database managing relationships between datasets used for ESO GB-wide dynamic reserve setting, as-well as the ESO GB-wide database setup code
- Documentation of the modelling and data-pipelines implemented to perform ESO GB-wide dynamic reserve setting
- A methodology & code to identify data drift over time in use of the ESO GB-wide dynamic reserve setting model

Produced in OCF Phase:

- Code to process solar forecast data from OCF
- Code to process other features such as system imbalance and European weather data

- Code to automatically generate features from these additional data sources
- Methodology to compare one version of a model against a baseline model and evaluate feature selection
- A solar forecast model extended to 48 hours ahead

Planned Implementation

A dashboard has been created and training has been given to end users. The system has been trialled initially awaiting full productionisation for operation in shadow mode. The trial has allowed us to set our default risk level for the control room at 99%. In the first one off trial in December 2023 we saved 1000MW of reserve over the 3B cardinal point. Subsequently the longer term trial caused no disruption to the control room.

Source data for DRS is being ingested into the data analytics platform as an enduring solution for the productionisation of DRS which will be hosted in the advanced analytics environments. Once this has happened this will feed into the SPICE system for use by the control room.

An extension to the project was undertaken with Open Climate Fix to see how an improved solar forecast would improve the model. Other experiments were also done to test the impact of specific features, such as adding European weather data or system imbalance. Using OCF's best probabilistic forecast PVNet an additional 100MW were saved with negligible impact on risk performance. This will be incorporated into our BAU process after the initial productionisation of DRS has been implemented.

Other Comments

The Project outcomes and results contain confidential information and intellectual property rights that cannot be disclosed in this Report due to their proprietary nature. Should the viewer of this Report ("Viewer") require further details this may be provided on a case by case basis following consultation of all Publishers. In the event such further information is provided each and any Publisher that owns such confidential information or intellectual property rights shall be entitled to request the Viewer enter into terms that govern the sharing of such confidential information and/ or intellectual property rights including where appropriate formal licence terms or confidentiality provisions. Dependent upon the nature of such request the Publishers may be entitled to request a fee from the Viewer in respect of such confidential information or intellectual property rights.

Standards Documents