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.

NIA Project Annual Progress Report Document

Date of Submission

Jun 2025

Project Reference Number

NIA2_NGESO002

Project Progress

Project Title

Solar PV Nowcasting

Project Reference Number

NIA2_NGESO002

Funding Licensee(s)

NESO - National Energy System Operator

Project Start Date

September 2021

Project Duration

3 years and 8 months

Nominated Project Contact(s)

innovation@neso.energy

Scope

The project is split into three different work packages:

WP1 - Design
Research & develop the use of machine learning & satellite images to nowcast PV power generation at GSP-level. Research will be conducted in close collaboration with academia & industry.

WP1 will include the following:

- Use historical data to train machine learning models.
- Evaluate forecast skill using historical data.
- Compare against ESO's current approach; and ESO's approach + gridded NWPs; and CM SAF.
- ML model output: Probabilistic predictions for total solar PV power generation for each GSP at 5-minute intervals.
- The system will also estimate PV outturn 'now' (situational awareness).
- Calibrate forecasts in near-real-time using live PV power data.
- Static designs for user-interface (UI)

WP2 – Prototype Development

Research & develop a prototype of real-time PV nowcasting system. Research how to present nowcasts to Control Room users via an interactive web user-interface. Build a prototype of API & web UI & nowcasting engine capable of running in real-time.

WP2 will include:

- Develop operational requirements with Control Room engineers
- Research ways to run the nowcasting ML models every 5 minutes for all of the 1 million PV systems in the UK.
- Develop a suite of PV nowcast performance validation metrics
- Build a functional prototype of nowcasting system & web user interface.
- Expose PV nowcasts via the API for PEF implemented in work package 1.
- The UI will give an overview of the country's UK PV fleet and also allow the user to drill into details.

WP3 – Prototype Demonstration
Evolve nowcasting system through multiple rounds of user feedback. Quantify impact on grid balancing.

WP3 will include:

- Feedback, evolve, develop with Control Room engineers to learn what they do and don't like about the prototype PV nowcasting system.
- Quantify the effect of PV nowcasting on balancing costs & CO2 emissions (building on the Control REACT NIA project).
- Explore using probabilistic PV nowcasts to dynamically set reserves (working with the Dynamic Reserve NIA project).
- Explore ways to measure users' interactions with the PV nowcasting system.
- Research ways to blend PV Live with PEF's half-hourly PV forecasts with OCF's 5-minutely PV nowcasts.

Objectives

The project proposed to complete the following objectives:

1. Research report comparing the performance of our 'Deep Learning' nowcasting system against other forecasting techniques (including ESO's current approach).
2. Cost estimates for running business-as-usual (BaU).
3. Prototype web UI & PV nowcasting service running in real-time. Including a suite of validation metrics.
4. Report on the feasibility of running ML-powered PV nowcasts in real-time.
5. Developed and socialized agile/CI/CD methodology.
6. Workshop on implementing satellite-powered nowcasts BaU.
7. Workshop on agile/CI/CD in critical energy applications.
8. Joint proposal with ESO for candidate adjacent project with planning on cadence and target level of agile/CI/CD.

Success Criteria

Project success will be defined by:

1. Demonstration of improved PV nowcasting skill compared to ESO's existing PV forecasting approach.
2. Delivery of a prototype of PV nowcasting service.
3. Measured change in balancing costs & CO2 emissions. User satisfaction. Identified next steps to apply agile/CI/CD to a future NGESO project.

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

The project is performing well against its main objectives and success criteria.

WP1 – Design

Throughout WP1, we conducted Machine Learning (ML) research, involving thousands of experiments, and tested four different models which have all outperformed the existing NESO model. At the end of WP1, the national solar generation forecast was 2.8 times better than the existing NESO PV forecast (for forecasts up to two hours ahead). The NESO's existing national solar PV forecasts have a Mean Absolute Error (MAE) of 650 MW. The best national PV forecasts developed in the course of the project to date have a MAE of 233.

WP2– Prototype Development
During WP2, we continued ML research, started UX & UI architecture development as well as open-source community engagement to deliver a prototype of the PV Nowcasting service. Early in WP2, we engaged with key stakeholders and end users of the forecasts to make sure the service meets specific NESO requirements. A session was held in June 2022 with relevant NESO stakeholders to demonstrate the first UI prototype of the service. As a result of further ML research in WP2, a better performing model has been developed that has high potential to outperform the current NESO PV forecast. More research in this area will be conducted in WP3.

WP3 – Prototype Demonstration
In WP3, we started working on creating national forecasting models that should significantly improve the accuracy of our forecast. We also continued work on productionising the service. Due to PVLive updating its GSP regions, we spent a reasonable amount of time making the GSP data updates. Further work on developing a UI and engaging with users resulted in the successful deployment of the first version of the forecast for use by the control room.

WP1 (extension) - Adding probabilistic forecasts
In WP1 of the project extension, we completed our work on the national model and delivered a user-consumable probabilistic forecast. We back tested the probabilistic forecast with a train and test set to check the results are not over fit, and the live results will have the same accuracy. We allow the user to access the probabilistic results through the API and the UIAPP we have made

WP4 – Ensemble Model Development
Design Throughout WP4, we conducted literature review of the approaches to combining ensemble weather forecasts to generate solar PV forecasts. We created a pipeline for ensemble weather data download and transformation into a format that can further be used in the existing ML pipeline, as well as the creation of an archive. We adjusted the existing ML model setup for work with ensemble weather data and running inference on the collected data to obtain initial results and produce visualization of solar PV forecasts at Grid Supply Point (GSP) level using ensemble weather data.
Work has commenced on generating solar PV forecasts at GSP level using ensemble weather data and aggregating these forecasts to produce a national-level solar PV forecast. Exploratory efforts are underway to incorporate recent ensemble weather observations as additional features in the deep learning model, which will improve model responsiveness to rapidly changing weather conditions and enhance short-term forecast accuracy. We will further evaluate the accuracy of solar PV forecasts using ensemble weather data by comparing it against existing solar forecasts output. The next session is scheduled in June 2025 to review ongoing model developments and gather feedback on forecasting outputs/accuracy.

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

As an R&D project, some modifications to how we carry out ML research were envisaged from the start, see the ‘Lessons Learnt’ section below for details.

The project was extended by 15 months to deliver the model’s probabilistic forecast results, developing the machine research further and focusing on model precision through adding more features to the forecast. The project extension is targeting a further 20% improvement in the MAE across horizons up to 4 hours, and 10% for day ahead horizons. This extension included an additional £355,000 in the project costs.

The project was extended to assess the impact of adding improved forecasts to the Dynamic Reserve Setting (DRS) model. This extension did not affect the project timeline but added an additional £80,000 to the project costs.

The project was extended further to investigate the use of NWP ensembles to comprehensively account for both weather and ML model uncertainty in probabilistic outputs. This further extension added 7 months to the project timeline and an additional £103,500 to the project costs.

The conversations on integrating the developed PV forecast into NESO systems are ongoing. During the course of the latest extension, the initial timeline was extended by 5 months due to delays in acquiring weather ensemble data from external sources. The original plan involved exploring weather data from two different sources into the modeling framework. However, in light of time constraints and challenges encountered in data acquisition, the approach was revised to utilise only one weather data source.

This approach will not affect project deliverables as we are still exploring methods of ensemble-based forecasting. No modifications were made to the project scope or key objectives.

Lessons Learnt for Future Projects

In the course of W1 and WP2, the project identified the following lessons:

Never underestimate the importance of cleaning up and checking data in advance:
Several approaches to loading data were tried, from on-the-fly to pre-preparing, and instituted automatic and visual tests of the data to ensure the project was always lining up the various data sources correctly.

Having infrastructure as code allows the main production service to run uninterrupted:
Having code to easily instantiate infrastructure is very useful to the efficient management of environments to ensure the project could bring the algorithm into productive use. The Terraform software tool was used which makes spinning up (and down) environments very easy and repeatable. Being able to spin up new environments allowed the project to test new features in development environments while allowing the main production to keep on running uninterrupted.

Using Microservices to “start simple and iterate” accelerates development:
Using a microservice architecture allowed the project to upgrade individual components as we see benefit in improving them, independently of changing other components’ behaviour. This has been very useful when building out the prototype service, as it has allowed the project team to start with a simple architecture - even a trivial forecast model - and iteratively improve the functionality in the components. For example, first starting out with one PV provider of data has allowed the project to get a prototype working, and in WP3 we will expand onboard an additional PV provider.

Data processing may take longer than expected:
While it was initially planned to extend our dataset back to 2016 for all data sources during WP2, it turned out that data processing takes much longer than expected. This does not have a direct impact on project deliverables but is something to consider in further ML research.

Data validation is important:
For both ML training and inference, using clear and simple data validation builds trust in the data. This helps build a reliable production system and keeps software bugs at a minimum.

Engaging specialist UX/UI skills is important:
By acknowledging that UX and UI design is a specialised area and incorporating those skills, a UI has been developed which will be easier to use and convey information effectively. This will be validated over WP3 through working with the end users.

Building our own hardware demonstrates value for money but may pose other challenges for a small team:
Two computers have been built during the project with a total of sixGPUs and it is estimated that using on-premises hardware instead of the cloud for data-heavy & GPU heavy machine learning R&D can significantly reduce the direct costs. However, the time it would require for a small team to put together all the components is significant (approx. 25 days for one person in total). While the total costs would still be lower, appropriate resource planning should be considered if planning hardware upgrades in the future.

In the course of WP3 and WP1 (extension), the project identified the following lessons:

Merging the code right away when performing frontend testing is of upmost importance:
Merging the code after frontend testing proved to be time-consuming, and it is important to consider when performing tests.

Large Machine Learning models are harder to productionise:
Large Machine Learning models proved to be difficult to productionise and the size of the model makes it difficult to use. Going forward, we need to investigate further how to deploy large models.

Machine Learning training always takes longer than expected:
Even with an already made model, making datapipes to work correctly takes time. It is important to always have enough time allocated when planning ML training activities.

Security and authentication is hard:
Ensuring robust authentication/security measures are in place is harder than we envisaged. It may be easier to implement packages already built or contract third party providers to support the process.

Separate National Forecast Model: PVLive estimate of National Solar generation does not equal the sum of PV Live's GSP generation estimate. This motivates us to build a separate National forecast, compared to adding up our GSP forecast.

Investment is needed to take open-source contributions to the next level:
Time and resources are needed to engage with open-source contributors and develop an active community. We may want to consider hiring an additional resource to support this activity.

Lessons from WPS (extension) were as follows:

Expensive cloud machines storage disk left idle:
We use some GCP/AWS machines for R&D, and we often pause the machine when they are not in immediate use. This is because some GPU machines can cost significant amounts per hour. It was discovered that costs still accrue for the disk (storage) of the paused machines. Balancing the pausing of the machines with the ability to start them up quickly versus starting a cloud machine from scratch each time has no golden rule, but it is useful to be aware of.

Challenging to maintain active communication with NESO:
Particularly high turnover at NESO forecasting team has affected communication on the project. This has evolved over the duration of the project, and more active and easier communication has been observed in the latter phase.

Reproducibility on cloud vs local servers:
When results differ between the cloud and local services, it can be tricky to determine the cause. Verbose logging, saving intermediate data, and maintaining consistent package versions and setup on both machines helped. One particular bug involved results differing when multiple CPUs were used locally, but only one CPU was used in the cloud.

Protection of production data:
Two environments in the cloud, "development" and "production," are maintained to protect the "production" data. This setup allows developers to access the "development" environment, where changes do not affect the live service. Although maintaining two environments increases costs, it is considered worthwhile.

Probabilistic forecasts:
Some unreleased open-source packages were used to implement one of the probabilistic forecasts. The advantage of using this code before its release is noted, but it also means more thorough checks for bugs are required, which can take more time.

Further lessons from WPS (extension) of NWP ensemble model development were as follows:

Clear project goals are important from the outset:

Clarifying project goals early, even in research projects with high uncertainty, prevents scope creep and wasted effort. This ensures everyone is aligned with what success looks like from the outset.

Schedule enough time for obtaining data:
Obtaining ECMWF EPS data proved challenging due to its substantial size (over 100GB per day) and the considerable time it takes to download, which can range from over 60 hours for a month's archive to as much as 120 hours depending on the archival system's workload. It is therefore important to schedule enough time to obtain data in the early stages of project planning.

Optimizing Data for Efficiency:
Streamlining the number of variables used in the model can significantly cut down on data download times and storage requirements. A study on feature importance pinpointed 6 out of 12 variables that collectively hold most of the information PVNet needs for forecasting, enabling a halving of both download time and storage by providing regular deterministic data for the remaining variables during inference.

The Importance of Visualisation for Ensemble Data:
The multifaceted nature of ensemble data makes it difficult to quickly grasp. Various visualisation techniques are essential for comprehending general trends, spread, and areas of uncertainty.

National ensemble-based forecasts are achievable within reasonable timeframes:
Running inference for 50 forecast versions for a single Grid Supply Point (GSP) takes approximately 5 seconds, following an initial setup period of under 6 minutes. This suggests that a national ensemble-based forecast could be produced in around 25 minutes.

Investing in foundational tools streamlines the entire research and development process:
Making sure enough time is allocated to develop the right tools, such as ocf-data-sampler or new training models, accelerates future development and improves overall efficiency.

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

Fully operational PV Nowcasting service running on two ML models:

PVNet for 0-6 hours
Blend of PVNet and National_xg for 6-8 hours, and National_xq from 8 hours and beyond.

Accuracy improvement over the previous OCF model by approximately 30% for the GSP and National forecasts (4-8 hours), resulting in forecasts approximately 40% more accurate than the BMRS model and over 40% against the PEF forecast (for 0-8 hours).

Probabilistic forecasts for all horizons:
Backtest runs for DRS project
UI including a new Delta view, dashboard view and probabilistic display
UI speedup, with query times from 20s down to <1 second

The most significant of these is the achievement of the target set by NESO of a 20% reduction in MAE error. This is extremely large in renewable forecasting and is the result of numerous machine learning improvements.

NWP ensembles:
Literature Review of NWP ensembles to identify common trends and finalise the methodologies for Solar PV forecasts creation using ensemble weather data.

Visualisations to highlight the general trend of the overall ensemble prediction and a probabilistic representation of the ensemble spread.

The most important outcome of the latest extension is in development now, which is to forecast Solar PV using ensemble weather data and evaluate the accuracy against existing solar PV forecasts.

Lastly, the forecast from Open Climate Fix is delivered completely open and documented. The resilience was significantly increased over the project duration, resulting in over 99.5% availability. This resilience is implementable by NESO, with all the infrastructure constructed in code to allow replicability

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 NESO can be found in our publicly available "Data sharing policy related to NIA projects (and formerly NIC)" and [Innovation | National Energy System Operator](#).

National Energy System Operator already publishes much of the data arising from our 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

- NIA2_NGESO002_Feasibility of Running ML-powered PV Nowcasts in Real Time
- NIA2_NGESO002_Research Report: PV Nowcasting Using Deep Learning_WP1-1
- NIA2_NGESO002_Open-Source Research Code and Results_WP1-2
- NIA2_NGESO002_Operational Pilot Service API_WP1-4

- NIA2_NGESO002_Progress Report_WP1
- NIA2_NGESO002_Progress Report_ WP2
- NIA2_NGESO002_Progress Report_WP3
- NIA2_NGESO002_Progress Report_ WP1 (extension)
- NIA2_NGESO002_Progress Report_WP2 (extension)
- NIA2_NGESO002_Progress Report_WP4 (extension)

ll reports and materials will be uploaded to the Smarter Networks Portal in the course of the project.