

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

Dec 2024

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

1 year and 11 months

Nominated Project Contact(s)

Lyndon Ruff

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 NWP; 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 NGENSO project.

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

(Standard ESO Text – do not amend)

National Grid Electricity System Operator ("NGESO") has endeavoured to prepare the published report ("Report") in respect of Solar PV Nowcasting, NIA2_NGENSO002, NIA2_NGENSO0020 ("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).

The Report provided is for information only and viewers of the Report should not place any reliance on any of the contents of this Report including (without limitation) any data, recommendations or conclusions and should take all appropriate steps to verify this information before acting upon it and rely on their own information. None of the Publishers nor its affiliated companies make any representations nor give any warranties or undertakings in relation to the content of the Report in relation to the quality, accuracy, completeness or fitness for purpose of such content. To the fullest extent permitted by law, the Publishers shall not be liable howsoever arising (including negligence) in respect of or in relation to any reliance on information contained in the Report

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 ESO model. At the end of WP1, the national solar generation forecast was 2.8 times better than the existing ESO PV forecast (for forecasts up to two hours ahead). The ESO'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 ESO requirements. A session was held in June 2022 with relevant ESO stakeholders to demonstrate the first UI prototype of the service.

As a result of the further ML research in WP2, a better performing model has been developed that has high potential to outperform the current ESO PV forecast. More research in this area will be conducted in WP3.

WP3 –

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 detail.

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 to the project costs.

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

The conversations on integrating the developed PV forecast into the ESO systems were kicked-off in the final months of the extension. Further conversations are ongoing to add a final extension to the project to investigate the feasibility of including a further dataset into the probabilistic forecasts.

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 to 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 six GPUs 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 utmost importance

Merging the code after frontend testing proved to be time-consuming and it is something 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 the 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.

Further 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 National Grid ESO

Particularly high turnover at National Grid ESO 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 allowed 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.

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 outcomes of the project to date are (2024 progress):

- 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 out of these is the achievement of the target set by NG-ESO of a 20% reduction in MAE error. This is extremely large in renewable forecasting, and is the result of numerous machine learning improvements.

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 NG-ESO, 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 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

All reports and materials will be uploaded to the Smarter Networks Portal in the course of the project, including two reports produced as part of project extensions:

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)