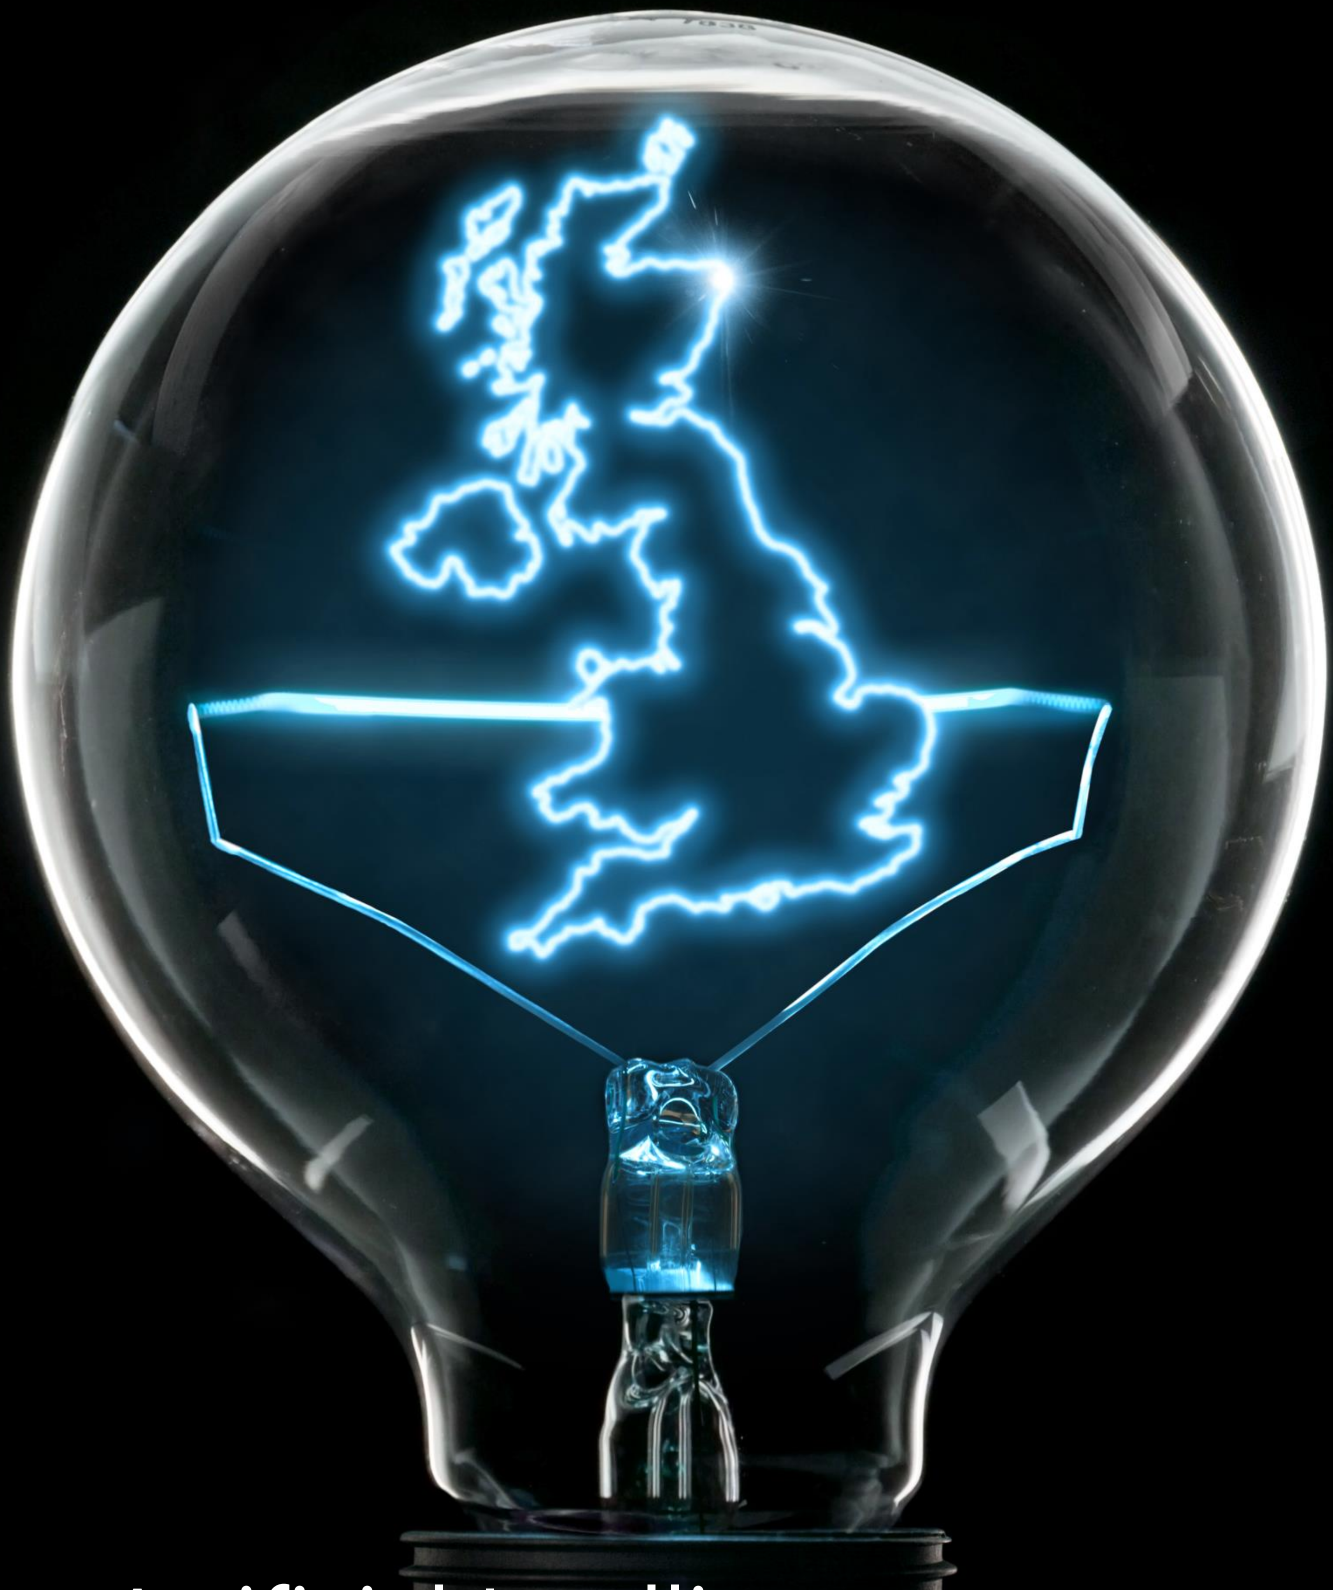

Using Artificial Intelligence to Design GB's Future Power Generation Mix

ISSUE 1

DECEMBER 2024

FRAZER-NASH CONSULTANCY
A KBR COMPANY

PREFACE

This article aims to highlight the potential for Artificial Intelligence (AI) technologies (specifically, reinforcement learning) to improve energy system planning and policy making.

Planning the future electricity system for Great Britain is a complex task due to the scale of the system and the many uncertainties that surround its future use. As the system grows, while also transitioning to be less carbon intensive, many choices will need to be made by Government, the system operators, asset owners and investors. These choices will affect the carbon emissions of the system, the costs and the security of supply for consumers.

We have developed a prototype AI system that can support decision makers in optimising their choices during the ongoing energy system transformation. This was achieved using open energy data and a simple bespoke energy system dispatch model. Whilst the results of this exercise are hopefully useful for demonstrating the potential for AI techniques within energy system planning, they have only been validated for demonstration purposes.

National energy system planning and policy design require a significant amount of effort annually by the National Energy System Operator (NESO) and the Department for Energy Security and Net Zero (DESNZ). We hope that in the future, the application of data science techniques will enable these teams to explore policy options that are more robust to alternative future scenarios. This will help to reduce system emissions, reduce costs and increase resilience for GB energy consumers.

FRAZER-NASH CONSULTANCY
A KBR COMPANY

Frazer-Nash Consultancy are a systems, engineering and technology company based in the UK and Australia.

We want to do things that matter, working on innovative technology solutions to help make lives safe, secure, sustainable and affordable.

We work with power generators, asset owners, system operators, government and academia to solve some of the industry's hardest technical and strategic challenges.

You can find out more about our recent projects and services by visiting our website.

www.fnc.co.uk

Taking the politics out of power generation strategy

How would a rational AI algorithm choose to optimise the path to a net zero electricity grid?

Following the open data publication of the Digest of UK Energy Statistics 2024 (DUKES) [1], we ran an experiment to gain some new insights into GB's power generation options, using a systems thinking approach and AI modelling techniques.

The future demand for electricity in the UK is highly uncertain. Not knowing the future demand makes it harder to make robust policy decisions on the best balance of power generation technologies that should supply our future grid.

Any future power mix needs to ensure security of supply, i.e. it needs to be capable of meeting demand at all times of the day, whatever the weather, or plant maintenance requirements. Alongside this uncertainty sits the imperative to decarbonise electricity generation, which currently accounts for around 15% of UK carbon emissions. Finally, given the above, we want to ensure that the solution provides the best value for money to consumers. This challenge is present internationally and is known throughout the world as the Energy Trilemma.

> *"The Energy Trilemma – how to balance security of supply, emissions, and cost."*

We used a systems modelling approach to explore how the UK power industry could securely decarbonise and at what cost. Applying AI techniques to the problem helped us to quickly explore a range of decarbonising strategies and optimise the balance between investment cost, security of supply and overall emissions for potential future demand scenarios. Our aims were to:

1. Test the usefulness of AI on a complex policy challenge.
2. Calculate some "no regret" options to inform power generation policy.

In other words, what generation technologies should we encourage in the next 5-10 years that we won't regret 15-20 years later. Here's what we found:

- AI, and probabilistic modelling techniques, can be applied to complex and uncertain problems like the Energy Trilemma to generate robust sets of strategic options. The AI technique applied here generates new solutions quickly, which could allow more scenarios to be explored, new technologies tested, or policy options analysed faster.
- The AI model is successful at reducing the carbon emissions of the grid mix, whilst the costs of electricity show a slight reduction in current generation costs. It achieves this whilst matching demand for each year modelled.
- The best options over the period to 2050 focus on increasing the capacity of nuclear, wind and solar power. Reducing the generating capacity of fossil fuels is feasible while still meeting demand. The AI determined that building more nuclear capacity than current plans can mitigate the risk of uncertain future demand.

Exploring the current energy landscape

313 TWh, the yearly average electricity generation over the last five years. 129 TWh of this came from renewable sources, according to the latest published yearly accounts from the Department for Energy Security and Net Zero (DESNZ) [1].

Extra capacity is needed to meet peaks in demand.

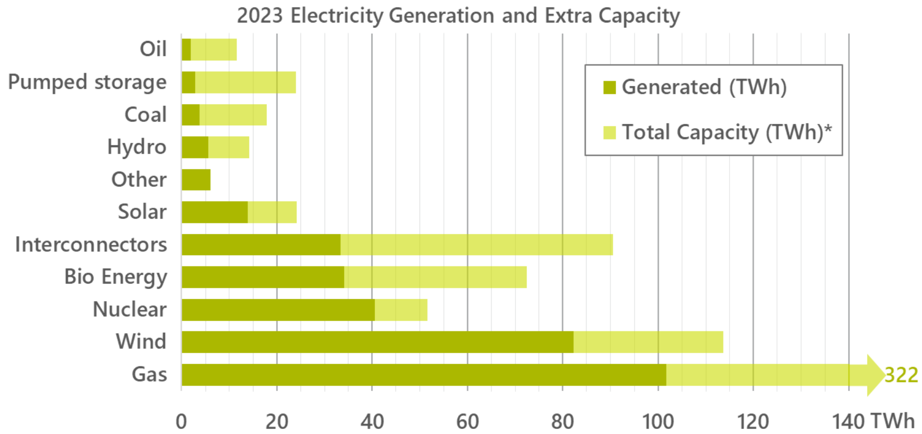

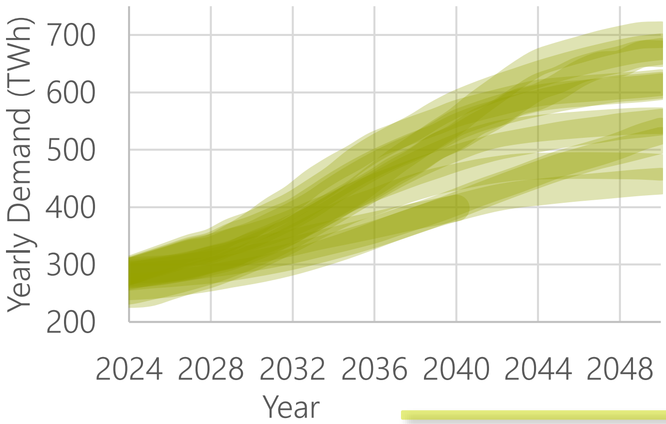

What will yearly demand look like in the future?

NESO's Future Energy Scenarios [2] suggest electricity demands are expected to increase and could feasibly double compared to 2023 demand. Demand pathways from NESO, Climate Change Committee [3] and DESNZ [4], all shown here, highlight the uncertainty in future demand.

The costs of building and deploying different power generation technologies on a future grid are highly uncertain and difficult to predict. Since our analysis is only intended to test the usefulness of AI in developing solutions to policy challenges, we have taken a simplified approach to capturing the costs of different technologies.

In our analysis we used the cost per MWh over the lifetime of each power generation technology (the Levelised Cost of Electricity, LCOE). This does not consider the additional grid connection, reinforcement, curtailment or balancing costs required for some technologies over others, but does allow for high-level, long-term comparisons to be made on construction and generation potential. There is variability in the costs and emissions of technologies, but for our purposes we have assumed the values below (from [5] and [6]) for new projects starting in 2025.

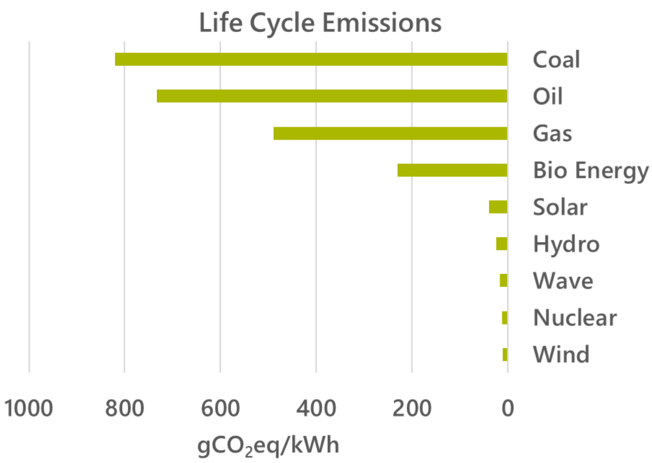

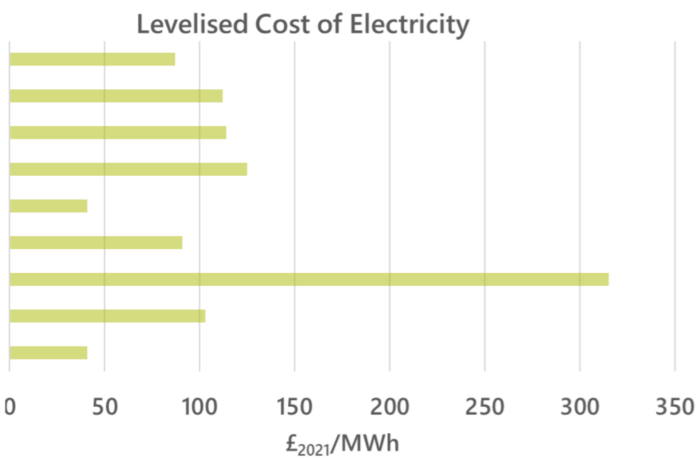

*Total capacities of solar, wind and wave are the derated capacities reported by DESNZ to allow for a reasonable comparison of total capacity between generation fuel sources.

How do you design a system which ensures security of supply when you don't know how much supply will be needed? How do you guarantee that such a system also minimises emissions and delivers electricity at the best value for money to consumers?

An excellent way to better understand, predict, manage or design a complex system is to build a computer model of it. With a model, we can test generation technology strategies before deploying them to understand the likely costs and emissions. Then we can run computer simulations of the strategy repeatedly with different electricity demand scenarios to make sure that it is robust to variations in future demand profiles.

In this experiment we let AI loose on our electricity system model to plan technology investment policies. We gave it the goal of teaching itself to make 'least regret' decisions for a robust future electricity system. This AI approach has speed and flexibility advantages over some of the established energy policy modelling approaches.

Unlocking the value in the data

To make good decisions on energy policy we must first have a means to test options before we commit to them. Computer models of the energy system provide the ability to do that.

Various models for the energy system exist already for use in policy decisions, for example the UK TIMES model [7], ESME [8], as well as circa 75 others [9]. Many of these models are detailed whole energy systems models that consider the many microeconomic factors at play, future technology options and multiple energy vectors in order to answer specific types of questions relating to energy policy.

Given the work already done on energy system modelling, we did not want to completely reinvent the wheel, but we needed to make sure we had a wheel that was the right size to bolt on to our AI decision engine and lightweight enough to get us to our policy destination quickly.

So we built a new custom model of the GB electricity market which enabled us to do that.

At Frazer-Nash we take a systems approach to building models and tackling complex problems (like solving the energy trilemma). This involves four main steps:

- Understand the purpose and parts. What question is the model trying to solve and what are the important factors in the wider system to consider when answering that question?
- Understand the relationships and the whole. How do the parts relate and interact to form a coherent whole system?
- Experiment and optimise. Using our understanding of the system, redesign it, make predictions with it and use data to optimise it.
- Test and validate. Ensure that the models address the original question, are coherent, credible, robust and fit for purpose.

The purpose of our model is to experiment with the capacities of power generation technologies to develop a capacity mix that will be robust to uncertain future demand at the lowest cost and lowest emissions.

The parts

As the saying goes "*All models are wrong, but some are useful*", it is important to realise that our model of the electricity system will not be perfect.

When defining a model to generate insights for a decision, one of the most important questions to understand is how much detail should we include, and how much detail is unnecessary?

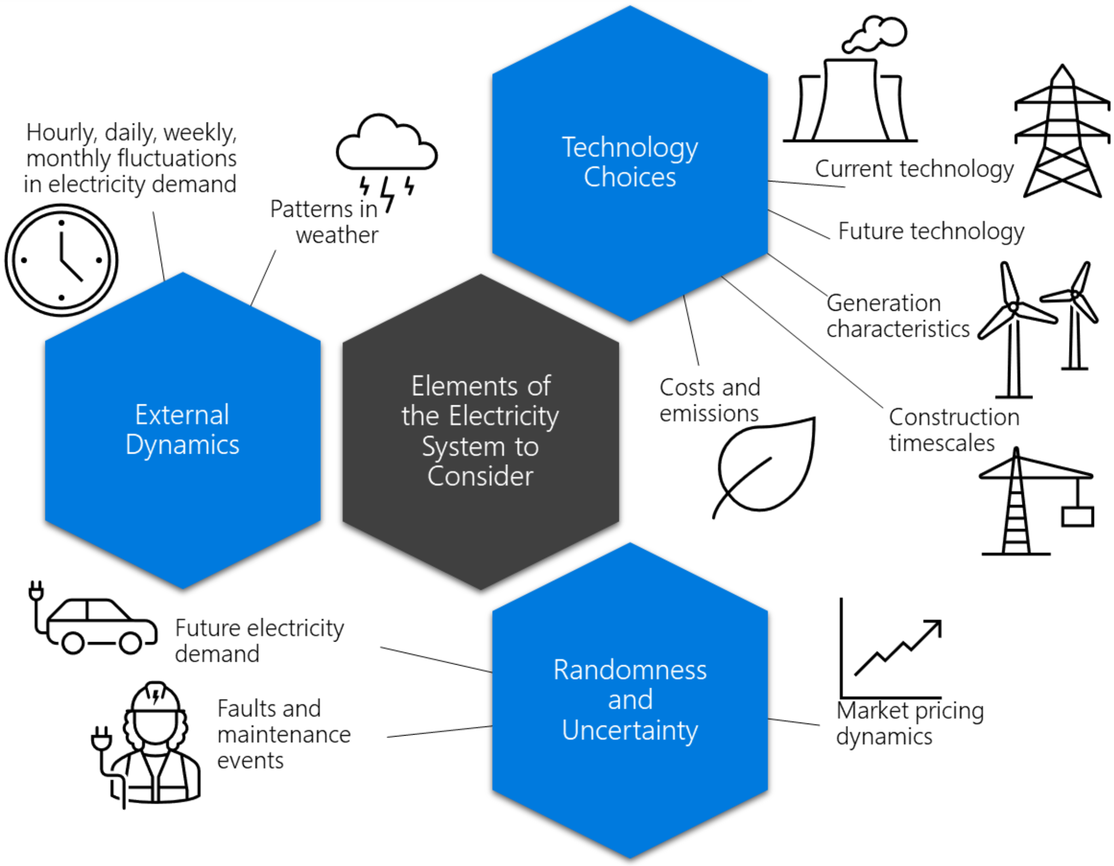

The following principles guided the inclusion of system parts in our model:

- We included a clustered mix of proven technology options:
  - Gas (natural gas)
  - Wind (onshore and offshore)
  - Nuclear (GW scale fission)
  - Bio energy
  - Interconnectors
  - Solar
  - Hydro and wave energy
  - Coal
  - Pumped storage
  - Oil
- We excluded any whole system effects from the provision of heat or future technologies still in development, or where LCOE and emissions data was unavailable. This results in recommendations that can be improved upon if, or when, technologies under development become available (e.g., new large-scale long-term battery storage technologies or advanced nuclear).
- Given the trade-off between speed and calculation complexity, we used a proximate model for the electricity market pricing mechanisms. Testing this demonstrated no significant loss in accuracy, this allowed us to run simulations of future electricity demand scenarios quickly, giving more time to train and apply our AI.

To bring the model to life we need to program **the relationships** between system parts that allow us to simulate future energy scenarios.

The model was made deliberately fast and light so that we could quickly test policy ideas and optimise a strategy for the energy trilemma. As such, the model we developed simply follows a time-stepping routine to calculate the costs and emissions associated with electricity supplied from the selected technologies on an hourly basis from now to 2050. This included factors such as intermittency of renewables based on weather patterns and other technologies based on historical maintenance requirements.

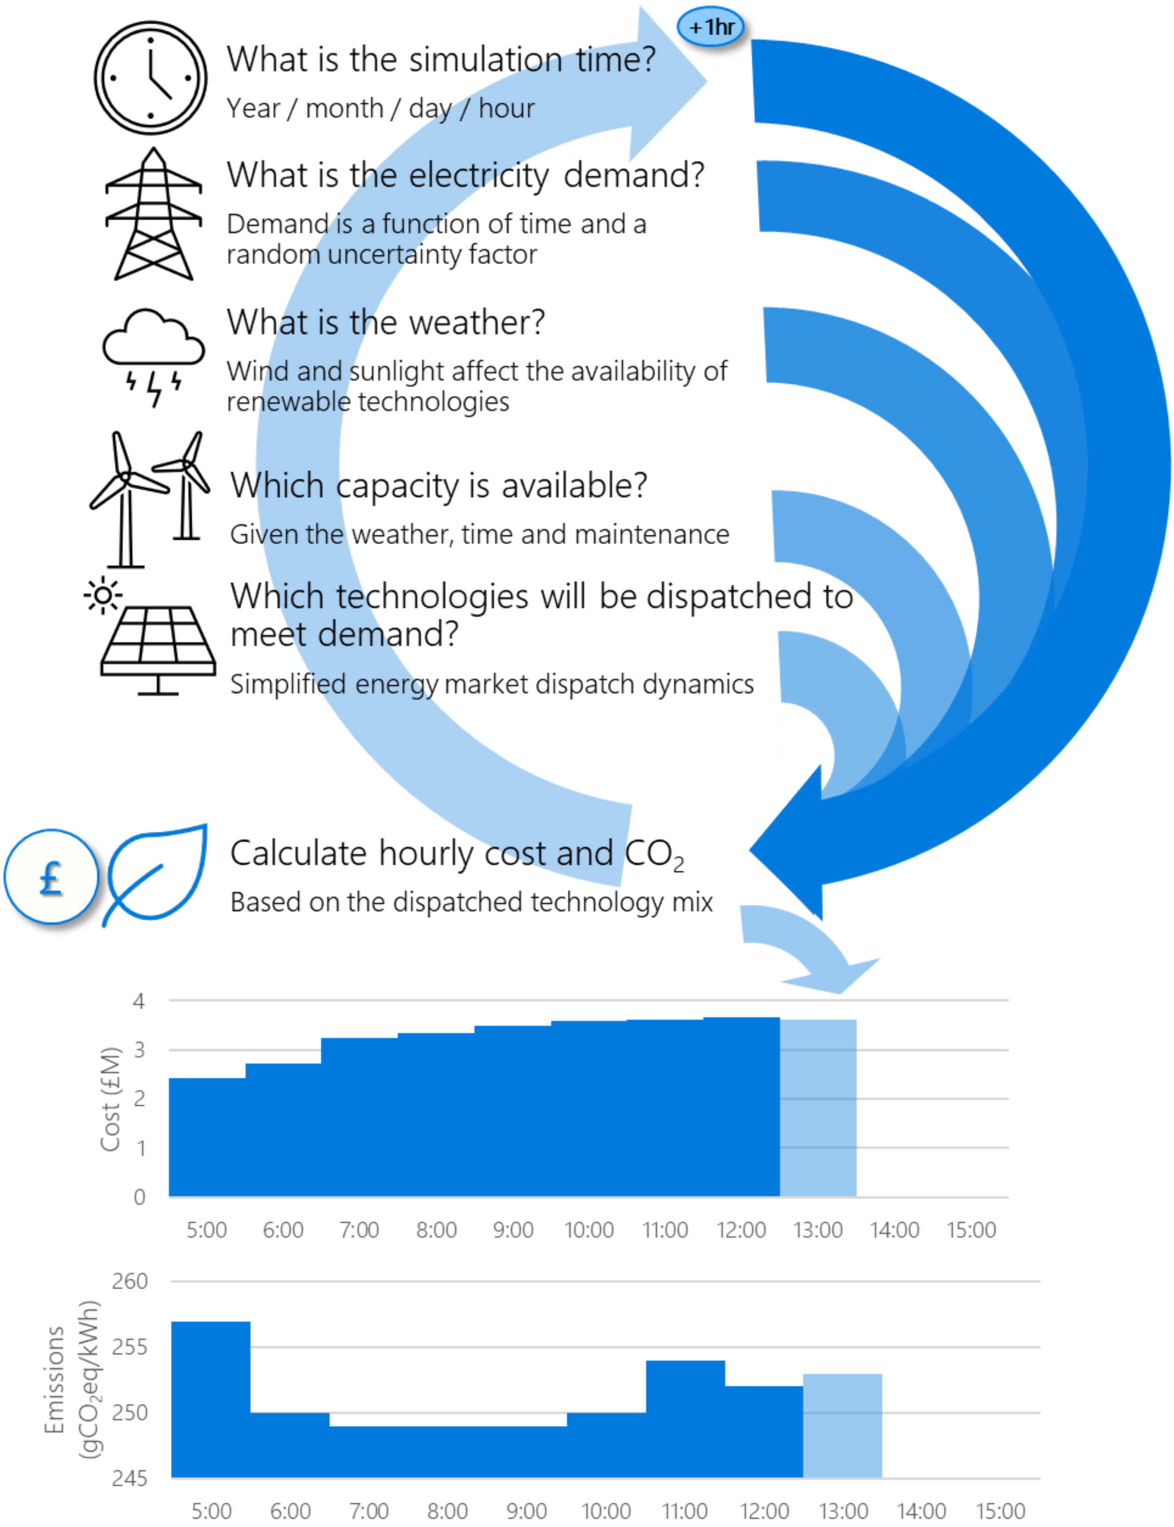

The importance of timing

Why time-step on an hourly basis? Timescales are key considerations for models of the energy system in GB. Uncertain multi-decadal changes in demand are at the core of this problem, but unless we model the hourly fluctuations in demand, we cannot accurately capture the peaks and troughs. So hourly time-stepping to understand the total capacity needed on the grid. Running hourly time-steps did not limit our model, since calculating 30 years' worth of dispatch takes seconds.

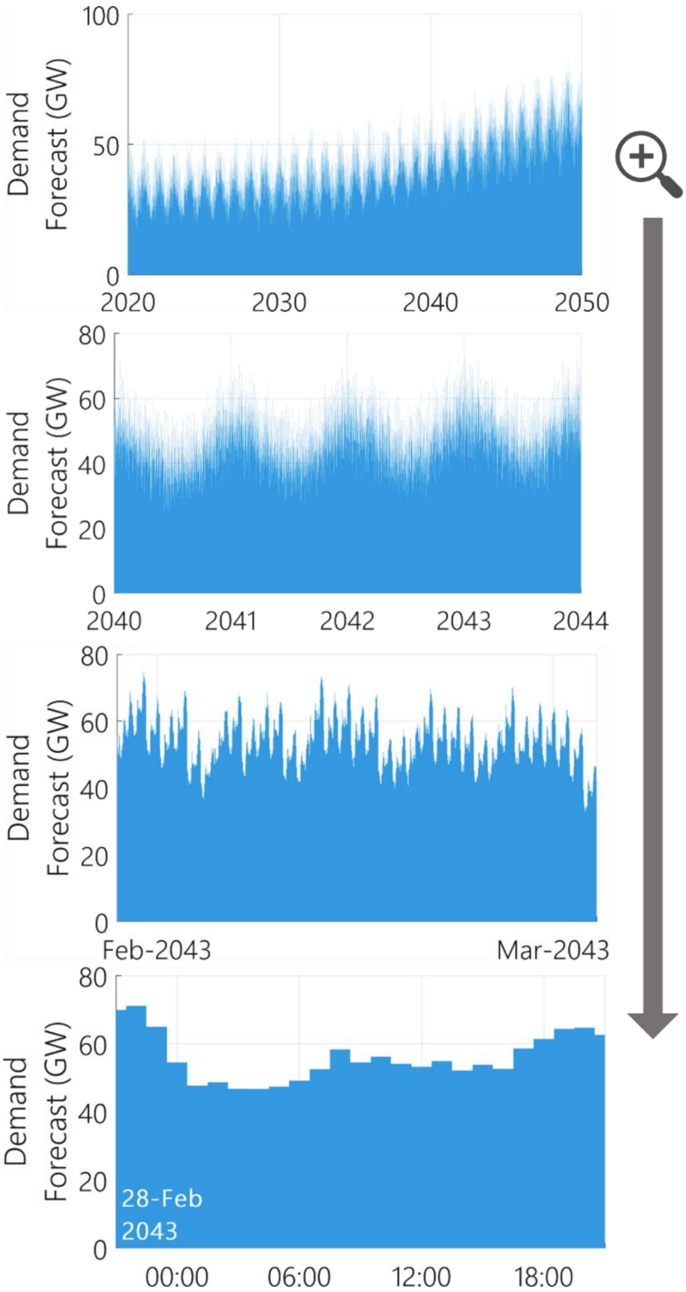

Experimentation

For most GB power system requirement projections, the currently installed generation capacity will not be enough to meet future demand. This leads to the need to decide which technologies should be built to meet the gap in capacity as demand rises. Our model allows us to experiment with designing future generation mixes. We can compare these options against conceivable future demand profiles to see how they hold up against uncertainty.

Testing

A range of model tests and validation checks were carried out on our simple economic dispatch algorithm, including the use of historical demand data as an input to predict which technologies the model would dispatch compared to real dispatch data. Our dispatch model matched historical profiles well.

Optimising the path to net zero

Whilst manually exploring technology investment policies with the model can be interesting, we can apply mathematical optimisation techniques to our model to produce a substantiated 'best' policy.

If we knew exactly what the future demand was going to be, this would be easy to do using mathematical optimisation, or even brute force – we could just run lots of combinations of technology investment choices and choose the best. In reality, we have to contend with long-term uncertainty in demand alongside hourly uncertainty in technology availability. This means there are tricky timing issues to contend with.

Given all the future uncertainty, the key question becomes

> *"How much of each technology should we build in the next 5-10 years?"*

If we don't start building soon it will be too late to meet increases in demand. If we build too soon, we might build too much and demand might never materialise. If we build the wrong mix of technologies at the wrong time then we could inflate the cost of generation, reduce the security of supply, or not meet our decarbonisation goals.

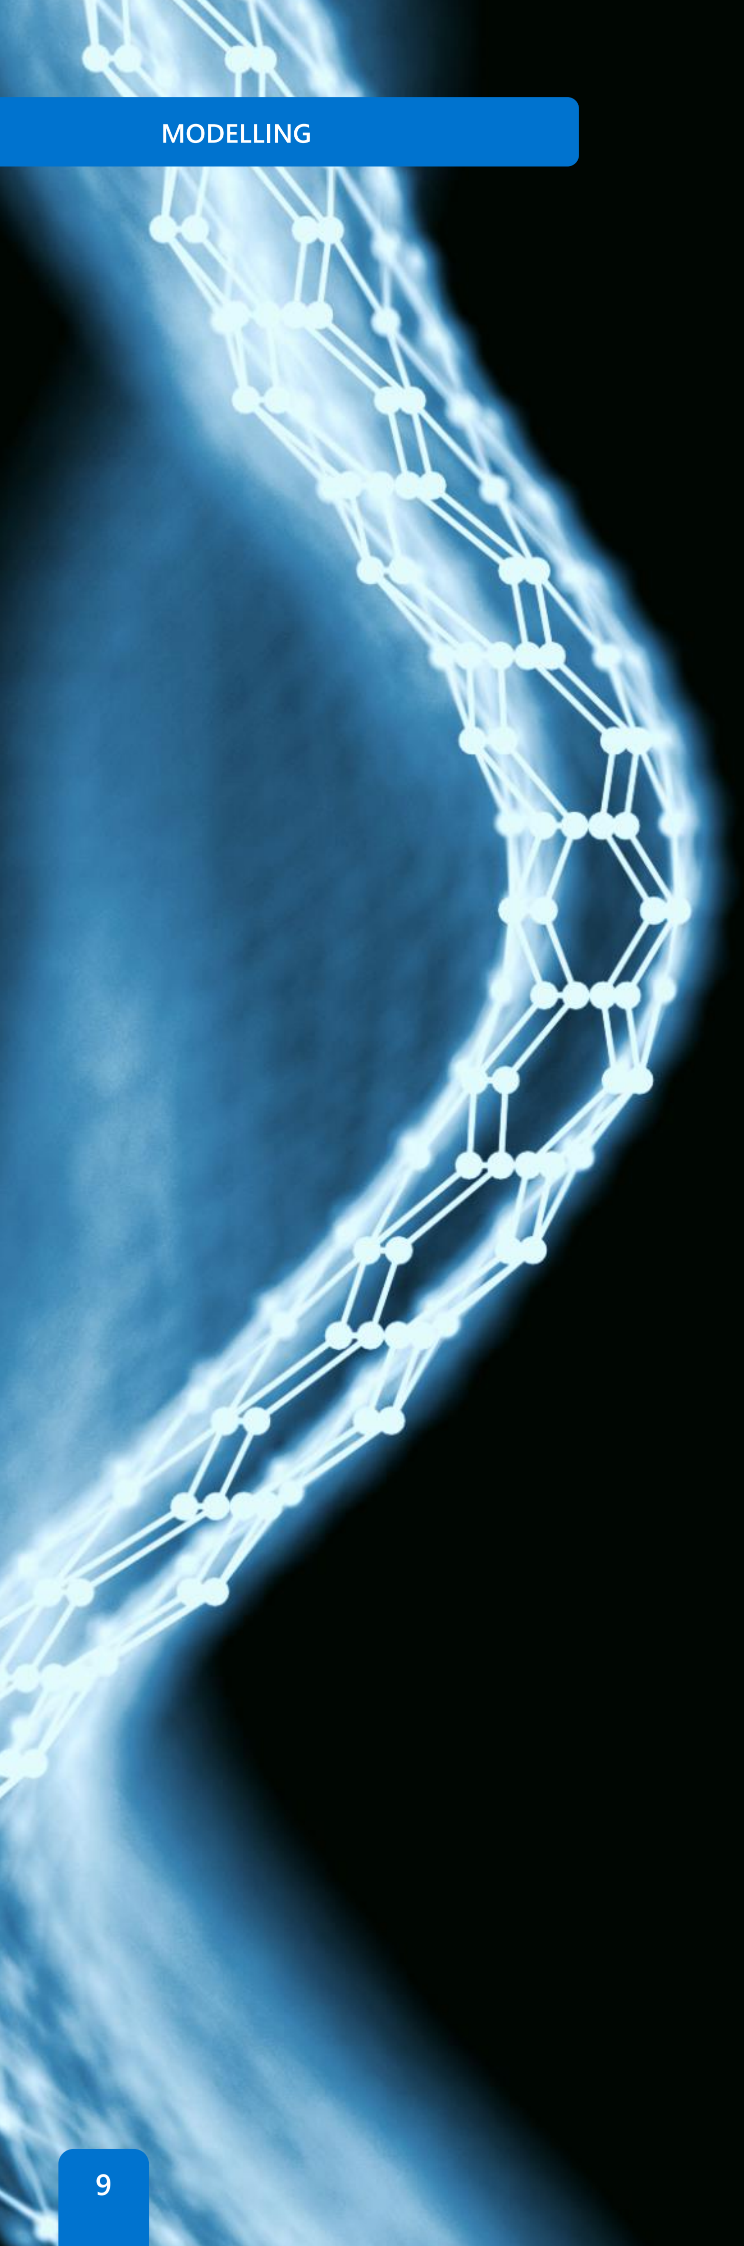

Applying Artificial Intelligence

There are plenty of traditional optimisation techniques that could crunch out a good solution for what the power mix should be to at some time in the future. But traditional optimisation can be slow to update when a new system configuration or demand scenario needs testing. We also want to know what sequence of decisions to make in the short term, given a large amount of uncertainty about the future.

AI can help us to navigate this very complex multi-choice landscape. Through a technique called reinforcement learning we created an AI agent that learnt how to make good long-term decisions in the power technology investment game.

Using reinforcement learning gave us the ability to generate sequences of sensible power technology investment decisions for different starting conditions or points in the timeline.

If you haven’t heard of reinforcement learning before, you may have read some of the headlines the technique has generated in the pursuit for artificial general intelligence that mimics human intelligence. Deepmind’s AlphaGo [10] brought the technique into the limelight by using it to beat the world champions at the boardgame Go. Subsequent versions have proven themselves on more complex multi-player computer games [11].

Deepmind have published their reinforcement learning algorithms and we are using similar algorithms here to learn to play our power system model as if it were a computer game.

How does Reinforcement Learning work?

Reinforcement learning teaches an AI 'agent' how to succeed in its environment using a carrot and stick approach. The agent observes its environment and makes choices about what to do next, initially through trial and error, but later developing a 'policy' to maximise long-term reward. In our case the agent observes the electricity demand and technology mix every year and makes choices for the following year about how much to invest in each generation technology. It receives rewards for desirable outcomes (low emissions) and a punishment for undesirable outcomes (high costs or unmet demand).

The agent develops its policies to achieve the greatest reward by simulating the years 2024-2050 millions of times with different randomised conditions, slowly generating an internal map of which decisions maximise reward in certain situations.

For a complex problem like this, the map will not be perfect, but a well-trained reinforcement learning agent is able to learn from the past, observe the present and create the desired future through intelligent actions.

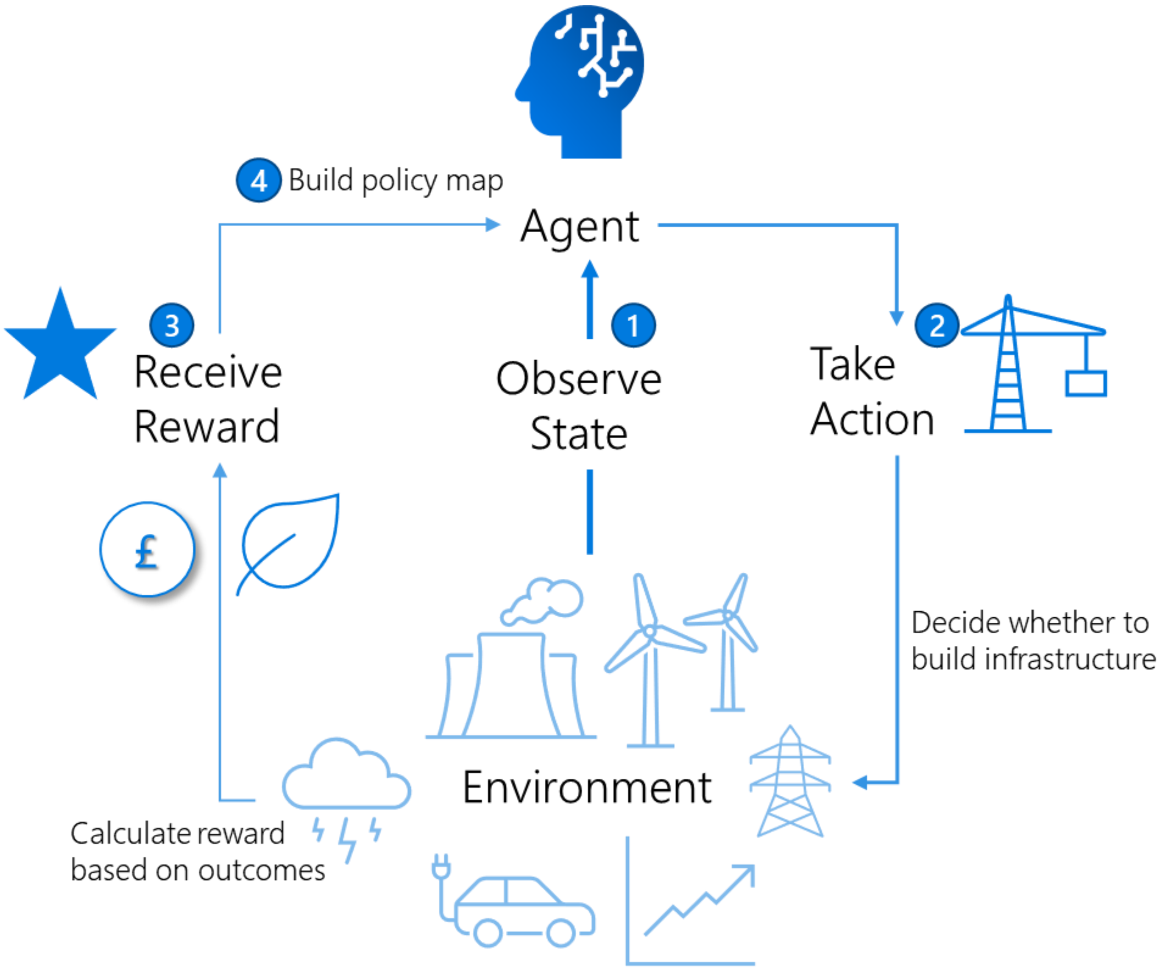

During training, the agent will repeat steps 1-4 above millions of times to build a map of successful policies for different states. It must balance the trade-off between exploration of new states and exploitation of the best states to fine tune its policies. Different reinforcement learning algorithms take their own approach to this challenge.

Once adequately trained, the agent can be deployed in an environment with new states and it will make near optimised decisions based on its policies to generate a high reward. In some systems it can continue learning in the new environment, or it can simply be deployed to make good decisions following steps 1 and 2 repeatedly.

According to our AI, how much of each technology should we build in the next 5-10 years?

The AI agent has been trained to make low cost, low carbon, high security-of-supply electricity capacity improvement decisions based on future demand scenarios. Whilst the GB electricity system is currently built to allow for 3 hours a year of Loss of Load Expectation (LOLE), we heavily penalised our AI agent for allowing any blackouts. This resulted in the agent slightly overbuilding capacity to avoid this outcome, meaning that the solutions were highly robust to any security of supply issues.

After training was complete (a few day's computing time) the agent was deployed to make capacity building decisions for a range of scenarios. It took less than a second to decide the best power generation strategy for the different demand scenarios we tested.

By the year 2030, the AI chose to increase the capacity for wind, solar, nuclear and hydropower to ensure there was always enough capacity available to meet demand. While wind energy saw the largest increase in available GW capacity (an extra 50GW), nuclear energy saw the largest percentage increase by roughly doubling the current capacity*.

This trend of capacity growth continued beyond 2030, with our AI agent choosing how to proportionally scale up the capacities of these technologies, which were preferred due to their low LCOE and emissions, while still being able to meet a variety of possible demand, weather and maintenance scenarios.

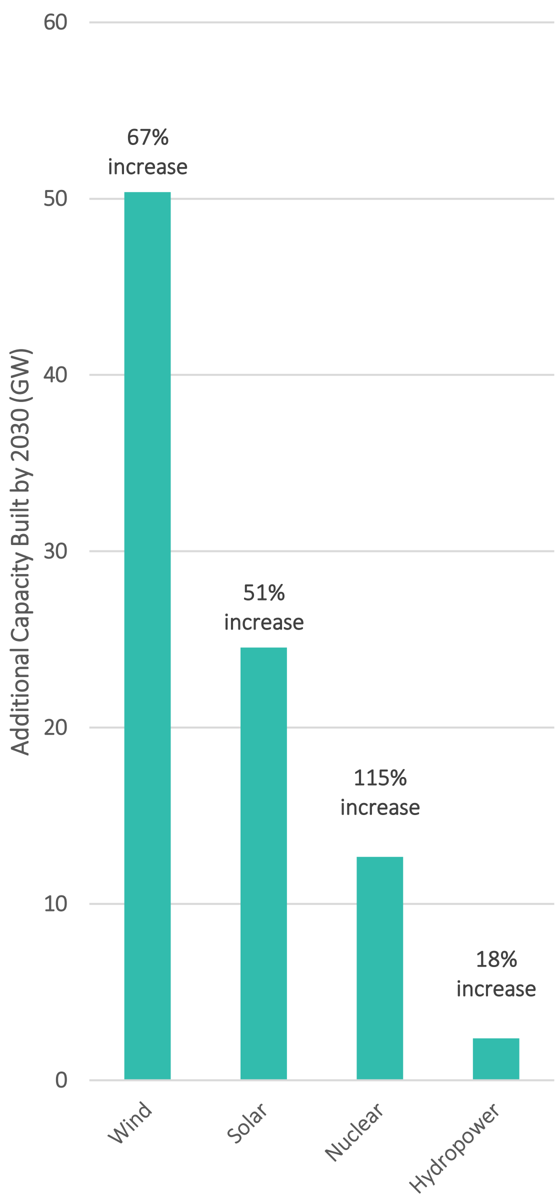

*Unfortunately this result does not mean that our AI agent found a magic shortcut for the long timescales required to build new nuclear power stations in the UK. In reality, new nuclear stations take a relatively long time to build and start generating. However, to keep our capacity construction model simple and consistent across large and small scale generation technologies, we assumed a linear build rate. So if we were to build a 3GW nuclear station by 2031, our model assumes incorrectly that we will have 2.5GW available by 2030.

We tested these capacities by running a range of future demand scenarios through our dispatch model to see what the total cost of generation and emissions would be. Since the costs and emissions are proportional to electricity demand, we calculated the cost and carbon savings compared to a *business-as-usual* counterfactual scenario, where all the current generation technologies increased in their current proportions to meet consumer demand.

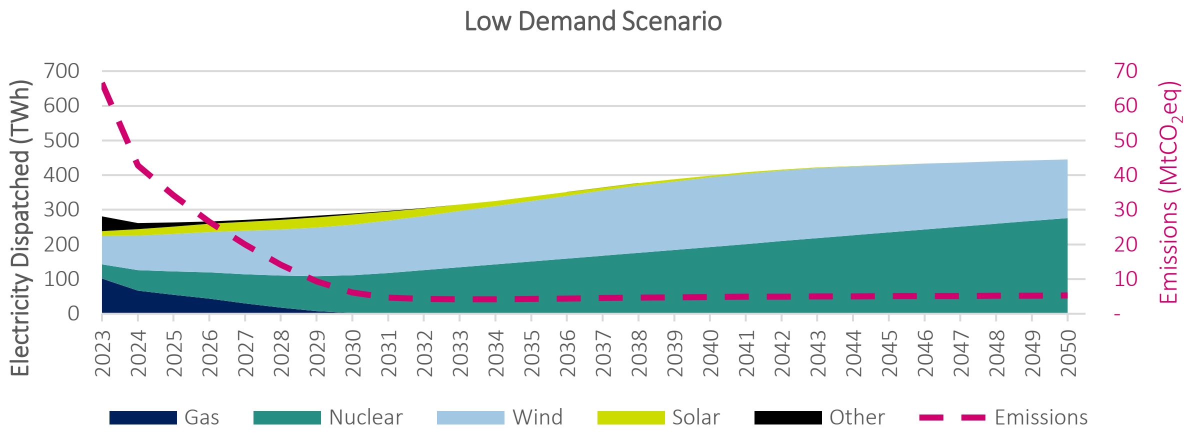

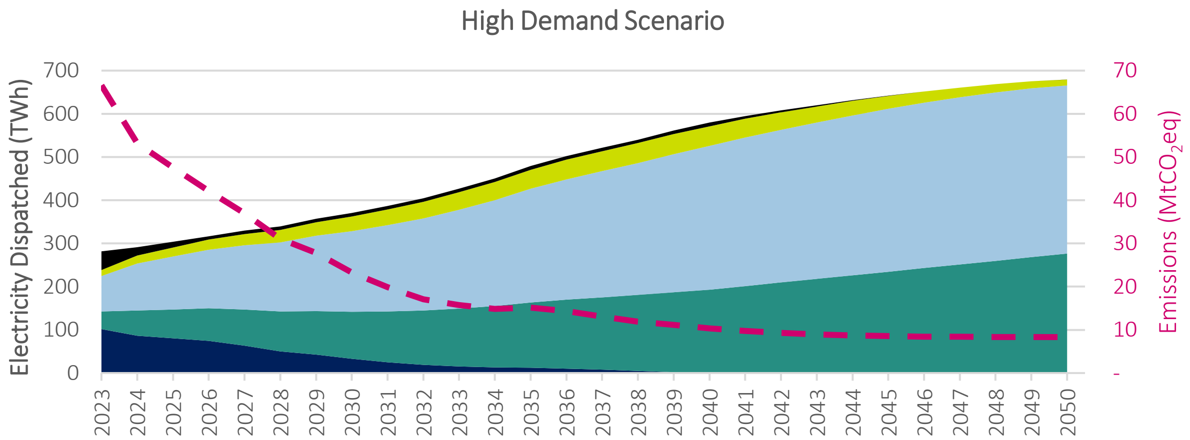

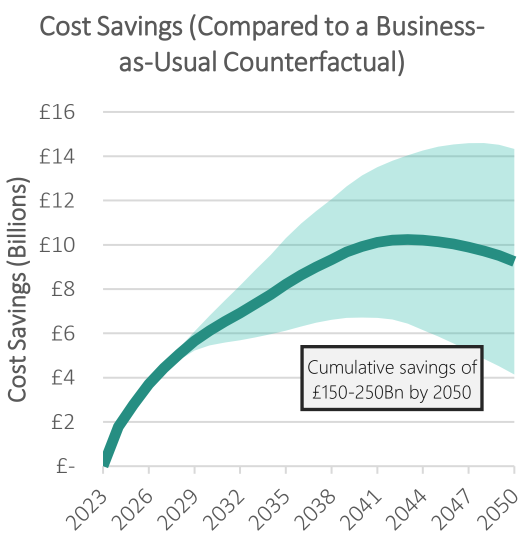

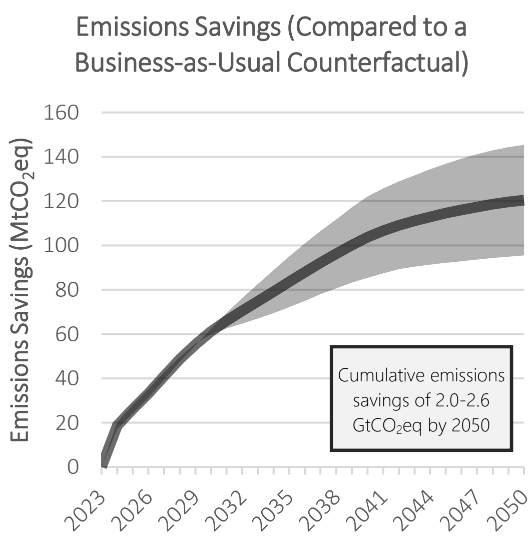

Is this solution feasible?

By 2050 the AI agent was able to design and build a generation technology mix that achieved between £150Bn and £250Bn (roughly 20%) cumulative savings compared to a business-as-usual counterfactual*. It achieved these savings while rapidly decarbonising the electricity generation market. This was primarily due to the reduction in the use of gas-powered technologies, which have relatively high emissions. These were replaced with wind power, nuclear, solar and a few other low carbon technologies.

In our current electricity transmission system, gas power plants provide useful frequency response and other ancillary services. These considerations were not included in this version of our AI system and a future grid design like the AI solution here would need to find a solution for these services with low cost and low carbon technologies.

The Government's *Net Zero Growth Plan* published last year [12] outlines plans for 50GW of offshore wind capacity by 2030, 70GW of solar by 2035 and 24GW of nuclear by 2050, which aligns reasonably well with the AI power generation plan. However, our AI agent preferred to build around 60% more nuclear capacity by 2050 than the Government's *Civil Nuclear: Roadmap to 2050* [13], while also including more hydropower and marine energy than has been announced in Government plans to date.

The AI agent came to a similar conclusion to NESO in their recently published their *Clean Power 2030* advice to government [14], where they advised rapid growth in wind and solar, with some small increases in long term energy storage such as pumped hydropower. However, the AI agent concluded that nuclear power would provide strong cost and emission benefits and should also be prioritised for rapid development alongside these other technologies to contribute further capacity in the 2030s.

The AI agent's strategy to build more nuclear capacity mitigates the risk of uncertain future demand. If demand is lower than anticipated, it is straight forward to build less wind and solar capacity if we have a reasonable nuclear baseload. However, if demand is higher than anticipated it will be challenging to quickly build more nuclear capacity due to the longer construction timescales.

* based on dispatched electricity LCOE

Background art generated using Leonardo AI "An AI vision of the future energy landscape"

How to get the benefits?

Our AI approach to power generation capacity building concluded that we could build the capacity needed to meet future demand at a lower generation cost, and emitting less carbon, by focussing policy efforts on enabling nuclear, wind and solar power capacities. The scale of capacity needed by the 2030s is high for all future energy demand scenarios, so we will need to deploy a mix of technologies quickly if we are going to meet net zero targets.

The value of our AI approach over more traditional energy system optimisation models will be realised when planners need to quickly understand what an optimised system could look like for different assumed future scenarios. For example, what if the demand pattern changed in 2035 due to a new vehicle technology? What if nuclear fusion or space-based solar power are deployable at grid scale capacities in 2040? What if we decide to relax or tighten up our resilience standards?

Numerous assumptions were used to test the feasibility of this approach that will now need to be addressed in future development. The options below outline the next steps to developing the capability for use in energy policy and elsewhere.

Requirements

Capture clear user requirements to guide future development, ensuring any model, software and interfaces can be planned for and engineered.

Enable technology costs and performance to evolve over time

Alongside a "proven technologies" pathway, calculate a pathway that highlights the benefits to be gained by R&D investment, lowering the cost of developing technologies in the future.

Techno-economics

Replace LCOE with a full economic cost model for the build, operation and decommissioning stages. Include the sunk costs of current infrastructure.

Spatial Planning

Account for the geospatial challenges of locational power generation. Include the costs and losses associated with electricity transmission and distribution.

Test and Validate

Benchmark the dispatch model and solutions against those developed using established models like UK TIMES or ESME

More technologies

We deliberately limited the technology mix available to allow us to develop the AI capability more easily. Adding a fuller suite of power generation and storage technology options would increase the usefulness of the outputs.

More short-term accuracy of technology build rates

Simulate the planning and construction timescales and their uncertainties to capture the risk of construction delays.

AI personalities

Train multiple agents with different preferences to aid decision making. A combination of agents would provide policy makers with multiple perspectives on different challenges. For example:

- **Cost Agent** – priorities system costs
- **Emission Agent** – prioritises carbon reduction
- **Technology Agent** – prioritises R&D and export opportunities
- **Defence Agent** – prioritises system resilience to geopolitical shocks

Whole System

Develop the approach for a whole energy system including electricity, natural gas and hydrogen to support whole system planning re-optimisation challenges.

Summary

This analysis set out to understand how a rational AI algorithm might choose to optimise the path to a net zero electricity system in Great Britain. Our aim is to highlight the potential for techniques such as reinforcement learning to enhance our energy system planning processes and policy decision making more generally.

There are an extremely large number of future scenarios for the GB energy system in the next few decades. Variability in weather patterns, longer term climate change, changing demand profiles, interconnectedness with Europe, new technologies, carbon emission and cost volatility; these are just some of the uncertainties that planners need to take account of to design a resilient and low-cost net zero energy system.

The reinforcement learning AI approach allows for robust decision making in the face of these uncertainties. The AI agent calculated that our immediate priorities should be to be build a large amount of extra capacity of wind, nuclear, solar and a few other low carbon technologies. While we don't know exactly what electricity demand will be in the future, the AI agent based its decisions on a broad range of possible scenarios to ensure the decision is robust to different future eventualities.

The benefit of the AI approach over traditional optimisation techniques is the speed at which you can deploy the agent once trained. Rather than slowly re-optimising for every new 'what-if' scenario you want to assess, the agent can use its training to highlight a sensible approach almost instantly.

Reinforcement learning also has the potential to develop new solution options by exploring strategies that have not previously been considered. Our agent chose to build more nuclear capacity than current government plans. This mitigates the risk of uncertain future demand levels, bearing in mind the relative timescales for building new nuclear capacity compared to solar and wind.

From our experience, we believe there are untapped potential benefits to deploying reinforcement learning in energy planning and other policy challenges. We have listed some examples below.

This analysis was only possible due to the sharing of open data sources like DUKES and the NESO Future Energy Scenarios data workbook. This open energy data is valuable for industry and academia to enable innovative research into energy policy and planning approaches.

Future Uses of Reinforcement Learning

Energy Industry

- Network infrastructure planning – how much transmission and distribution infrastructure should be built? Where? When? How does that answer change under different energy technology future scenarios?
- Investors – how much of a given technology should be deployed amongst the future grid under different demand scenarios and alternative technology mixes? What is the rate of return confidence level?
- Resilience planning – how do we optimise the design of the grid so that it is robust to external shocks at given resilience standard level? Where are the vulnerabilities and how do we mitigate them?

Other policy challenges:

- What workforce skills and numbers are needed to maximise future economic performance?
- What level of investment in the healthcare system would provide the best balance of cost and performance for a given level of service or budget?
- What transport infrastructure should we build to best meet the needs of future transport users on roads, rail and air?
- What business and R&D investments should we prioritise under different long-term geopolitical scenarios to maximise economic performance?

This analysis was developed by the Frazer-Nash Consultancy Strategic Modelling group. We specialise in solving policy and strategy problems with the help of advanced modelling and decision science approaches.

For more information about our services or this project please contact the authors:

Andy Moore
Group Leader
Strategic Modelling
a.moore@fnc.co.uk

Chris Williams
Senior Engineer
Strategic Modelling

Daniel Kent
Consultant
Strategic Modelling

Michael Graham
Senior Consultant
Strategic Modelling

References:

[1] Digest of UK Energy Statistics (DUKES) 2024, Department for Energy Security and Net Zero, https://www.gov.uk/government/statistics/digest-of-uk-energy-statistics-dukes-2024 (2024).
[2] Future Energy Scenarios: ESO Pathways to Net Zero, https://www.neso.energy/publications/future-energy-scenarios-fes (2024).
[3] Sixth Carbon Budget, Climate Change Committee, https://www.theccc.org.uk/publication/sixth-carbon-budget/ (2020).
[4] Energy and Emissions Projections, Department for Energy Security and Net Zero, https://www.gov.uk/government/collections/energy-and-emissions-projections (2023).
[5] Electricity Generation Costs 2023, Department for Energy Security and Net Zero, https://www.gov.uk/government/publications/electricity-generation-costs-2023 (2023).
[6] Schlömer et al, Annex III: Technology-specific cost and performance parameters. In: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change Cambridge University Press (2014).
[7] Hannah E. Daly, Brigit Fais, UK TIMES Model Overview, UCL Energy Institute (2014).
[8] Chris Heaton, Modelling Low-Carbon Energy System Designs with the ETI ESME Model, ETI (2014).
[9] Neil Strachan, Pei-Hao Li, Energy Modelling in the UK, Briefing Paper 1: The Modelling Landscape, UKERC (2021).
[10] David Silver et al, Mastering the Game of Go with Deep Neural Networks and Tree Search, Nature 529, 484–489 (2016).
[11] Deepmind, Mastering the Game of Stratego with Model-Free Multiagent Reinforcement Learning, https://arxiv.org/pdf/2206.15378 (2022)
[12] Powering Up Britain: Net Zero Growth Plan, Policy Paper, Department for Energy Security and Net Zero, https://www.gov.uk/government/publications/powering-up-britain/powering-up-britain-net-zero-growth-plan (2023).
[13] Civil Nuclear: Roadmap to 2050, Policy Paper, Department for Energy Security and Net Zero, https://www.gov.uk/government/publications/civil-nuclear-roadmap-to-2050 (2023).
[14] Clean Power 2030, Advice on achieving clean power for Great Britain by 2030, NESO, https://www.neso.energy/publications/clean-power-2030 (2024)

FRAZER-NASH CONSULTANCY
A KBR COMPANY

www.fnc.co.uk