Let's Not Forget About Machine Learning

Forecasting Beyond the AI Buzz

Andrew Akman

Principal Data Scientist - Operational Forecasting

Australian Energy Market Operator

June 2025

Agenda

- Introducing AEMO & Operational Forecasting
- Unpacking the AI Umbrella
- Our Machine Learning Journey
  Charting Our Machine Learning Journey Through Demand Forecasting
- Looking Elsewhere – New & Emerging Opportunities in Machine Learning & AI

Introducing AEMO & Operational Forecasting

About AEMO

- AEMO is a member-based, not-for-profit organisation.
- We are the independent energy market and system operator for the National Electricity Market (NEM) and the WA Wholesale Electricity Market (WEM), and system planner for the NEM.
- We also operate retail and wholesale gas markets across south-eastern Australia and Victoria's gas pipeline grid.

AEMO SERVICES

AEMO Services is an independent subsidiary of AEMO, established in 2021 to enable the transparent provision of advisory and energy services to National Electricity Market jurisdictions.

Electricity

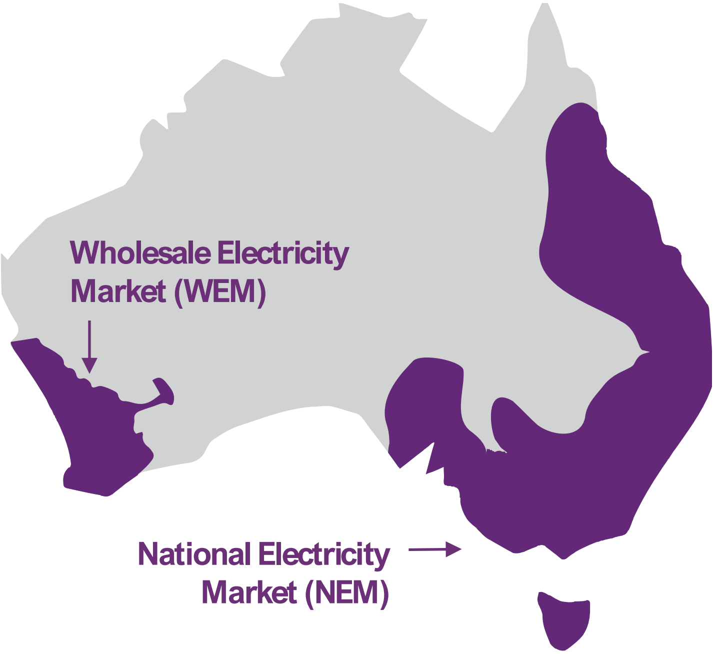

Gas

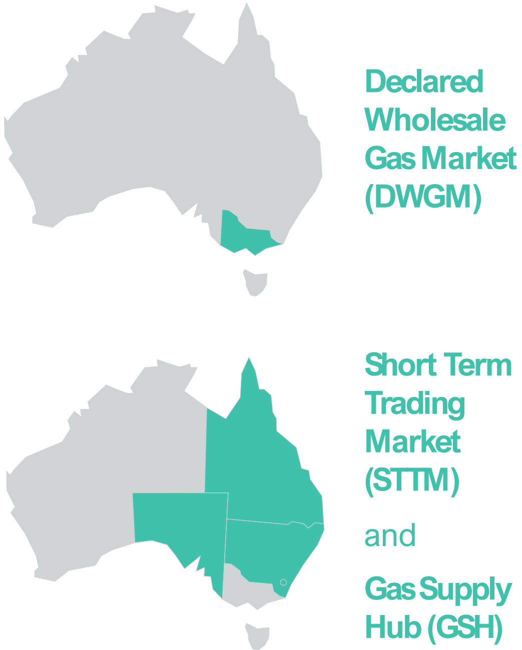

Operational Forecasting – Our Forecasts

System	Why it matters	Horizon
PASA Network outages Generator outages **Demand and rooftop PV forecasts** **Wind and solar forecasts** **Forecast uncertainty**	**Operational planning** to maintain a secure and reliable power system	1 week ahead: PD & ST PASA
Pre-dispatch Outages Bids and capacity available **Demand and rooftop PV forecasts** **Wind and solar forecasts**	**Contributing to market efficiency** by providing the best information for participants decision making	1 day ahead: Pre-dispatch 1 hour ahead: 5-min Pre-dispatch
Dispatch Bids and capacity available **Demand forecasts** **Wind and solar dispatch forecasts**	**Control of the power system** through dispatch and frequency control	5-min ahead: Dispatch

AEMO produces ~3 million point-forecasts a day

Impact of The Energy Transition on Operational Forecasting

AEMO's 2024 ISP step change scenario

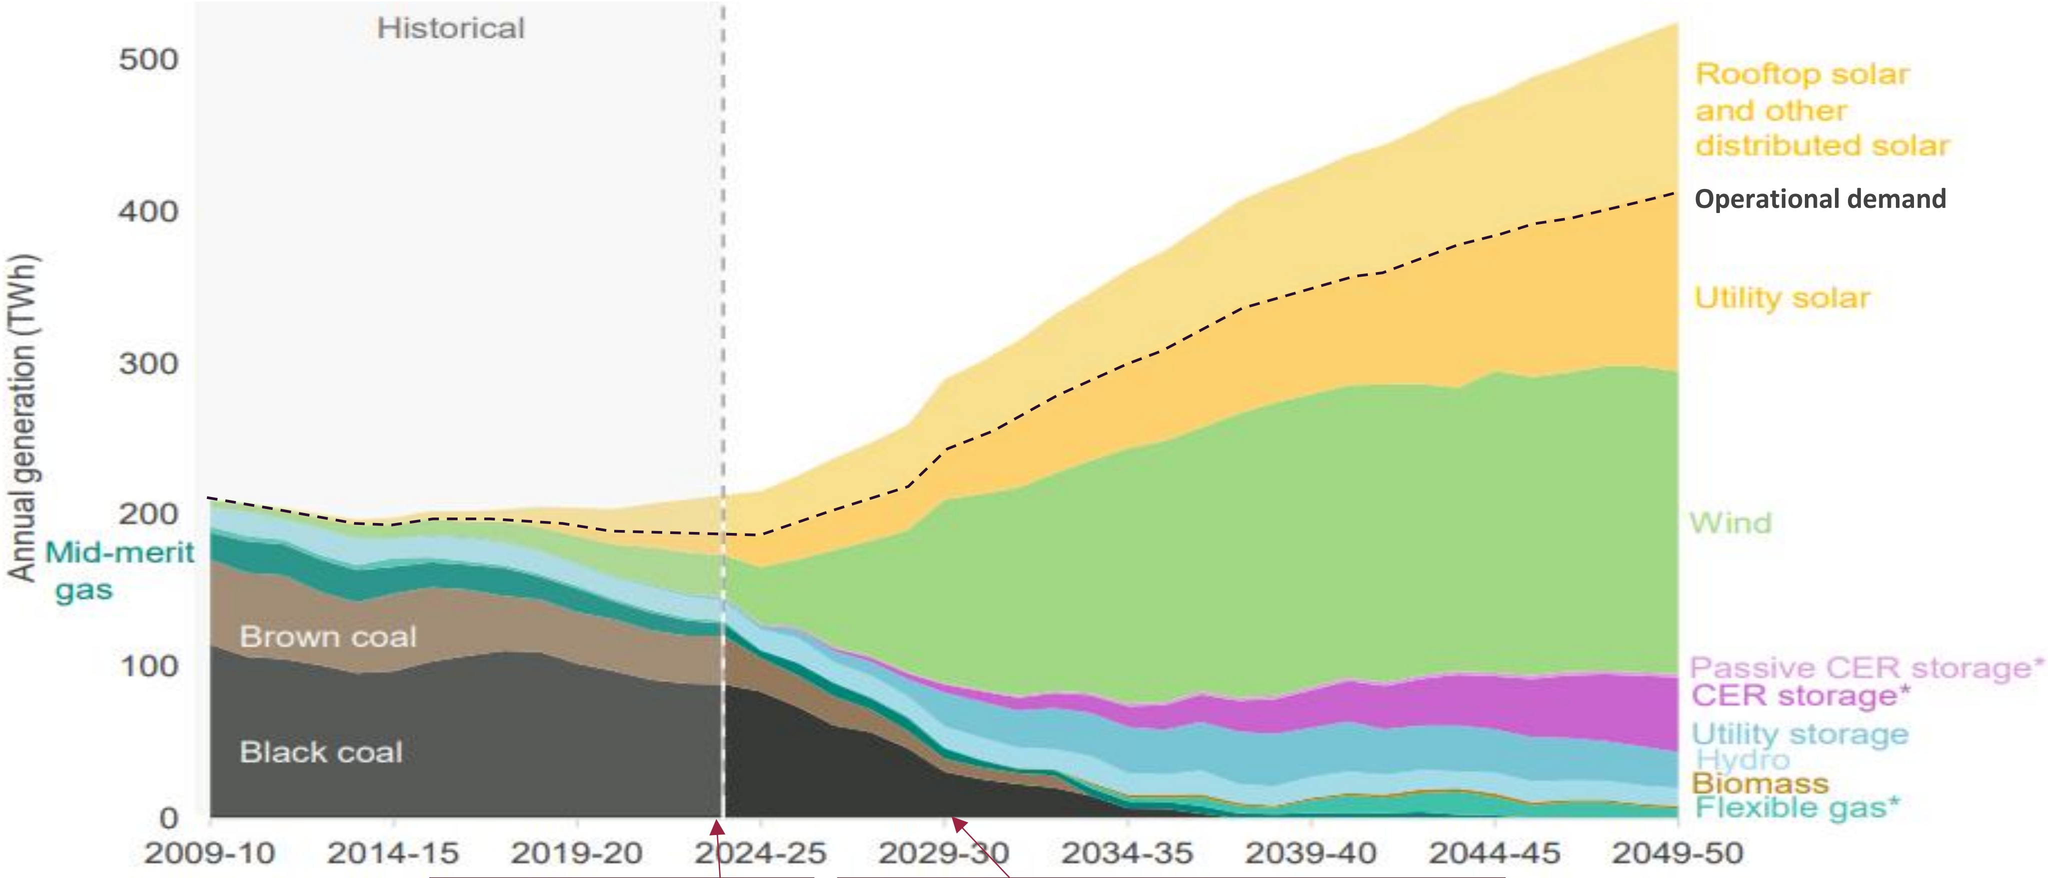

Today, around 28% of operational demand is met by large-scale VRE

By 2030, close to 60% will be met by large-scale VRE, increasing the reliance on VRE forecasts

"Flexible gas" includes gas-powered generation and potential hydrogen capacity.
"CER storage" means consumer energy resources such as batteries and EVs.

Unpacking the AI Umbrella

The AI Umbrella

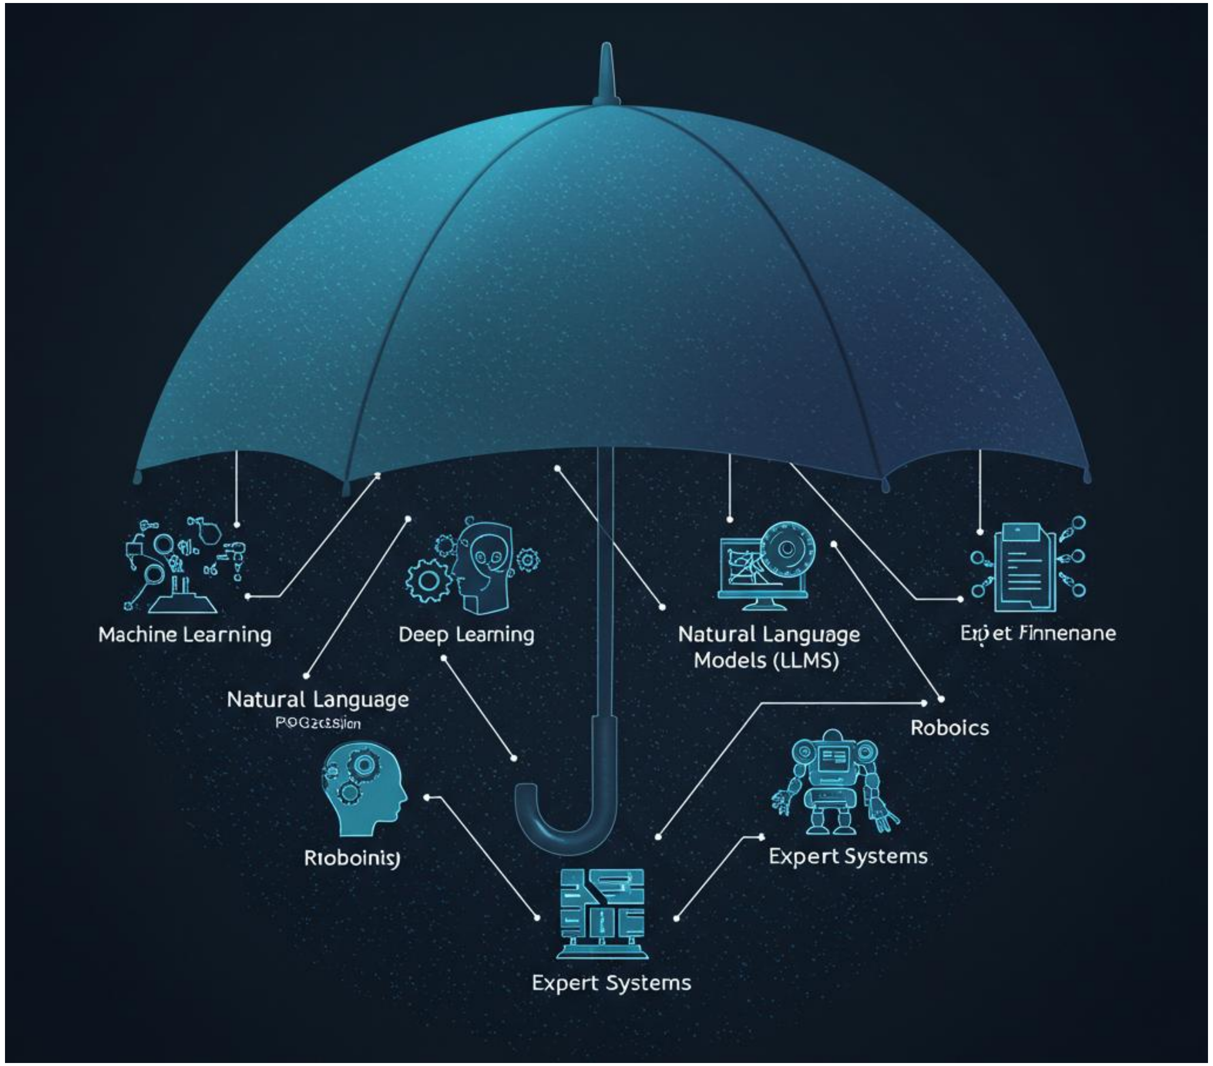

Artificial Intelligence (AI) is a broad and encompassing field, often misunderstood as solely referring to the latest advancements like large language models. In reality, AI serves as an umbrella term, covering a diverse range of disciplines and technologies.

- **Machine Learning (ML)**: The core of AI's practical applications, where systems learn from data without explicit programming.
  - Includes: Supervised Learning, Unsupervised Learning, Reinforcement Learning.
- **Deep Learning (DL)**: A subset of Machine Learning using multi-layered neural networks, excelling in complex pattern recognition (e.g., image and speech).
- **Natural Language Processing (NLP)**: Enables computers to understand, interpret, and generate human language.
- **Large Language Models (LLMs)**: A prominent subset of NLP and Deep Learning, focused on generating and understanding human-like text at scale.
- **Computer Vision (CV)**: Allows machines to "see" and interpret visual information from the world.
- **Robotics**: Deals with the design, construction, operation, and use of robots.
- **Expert Systems**: Early AI systems designed to mimic human decision-making in specific domains using rule-based reasoning.

Our Machine Learning Journey

Charting Our Machine Learning Journey Through Demand Forecasting

Demand Forecasting for the NEM

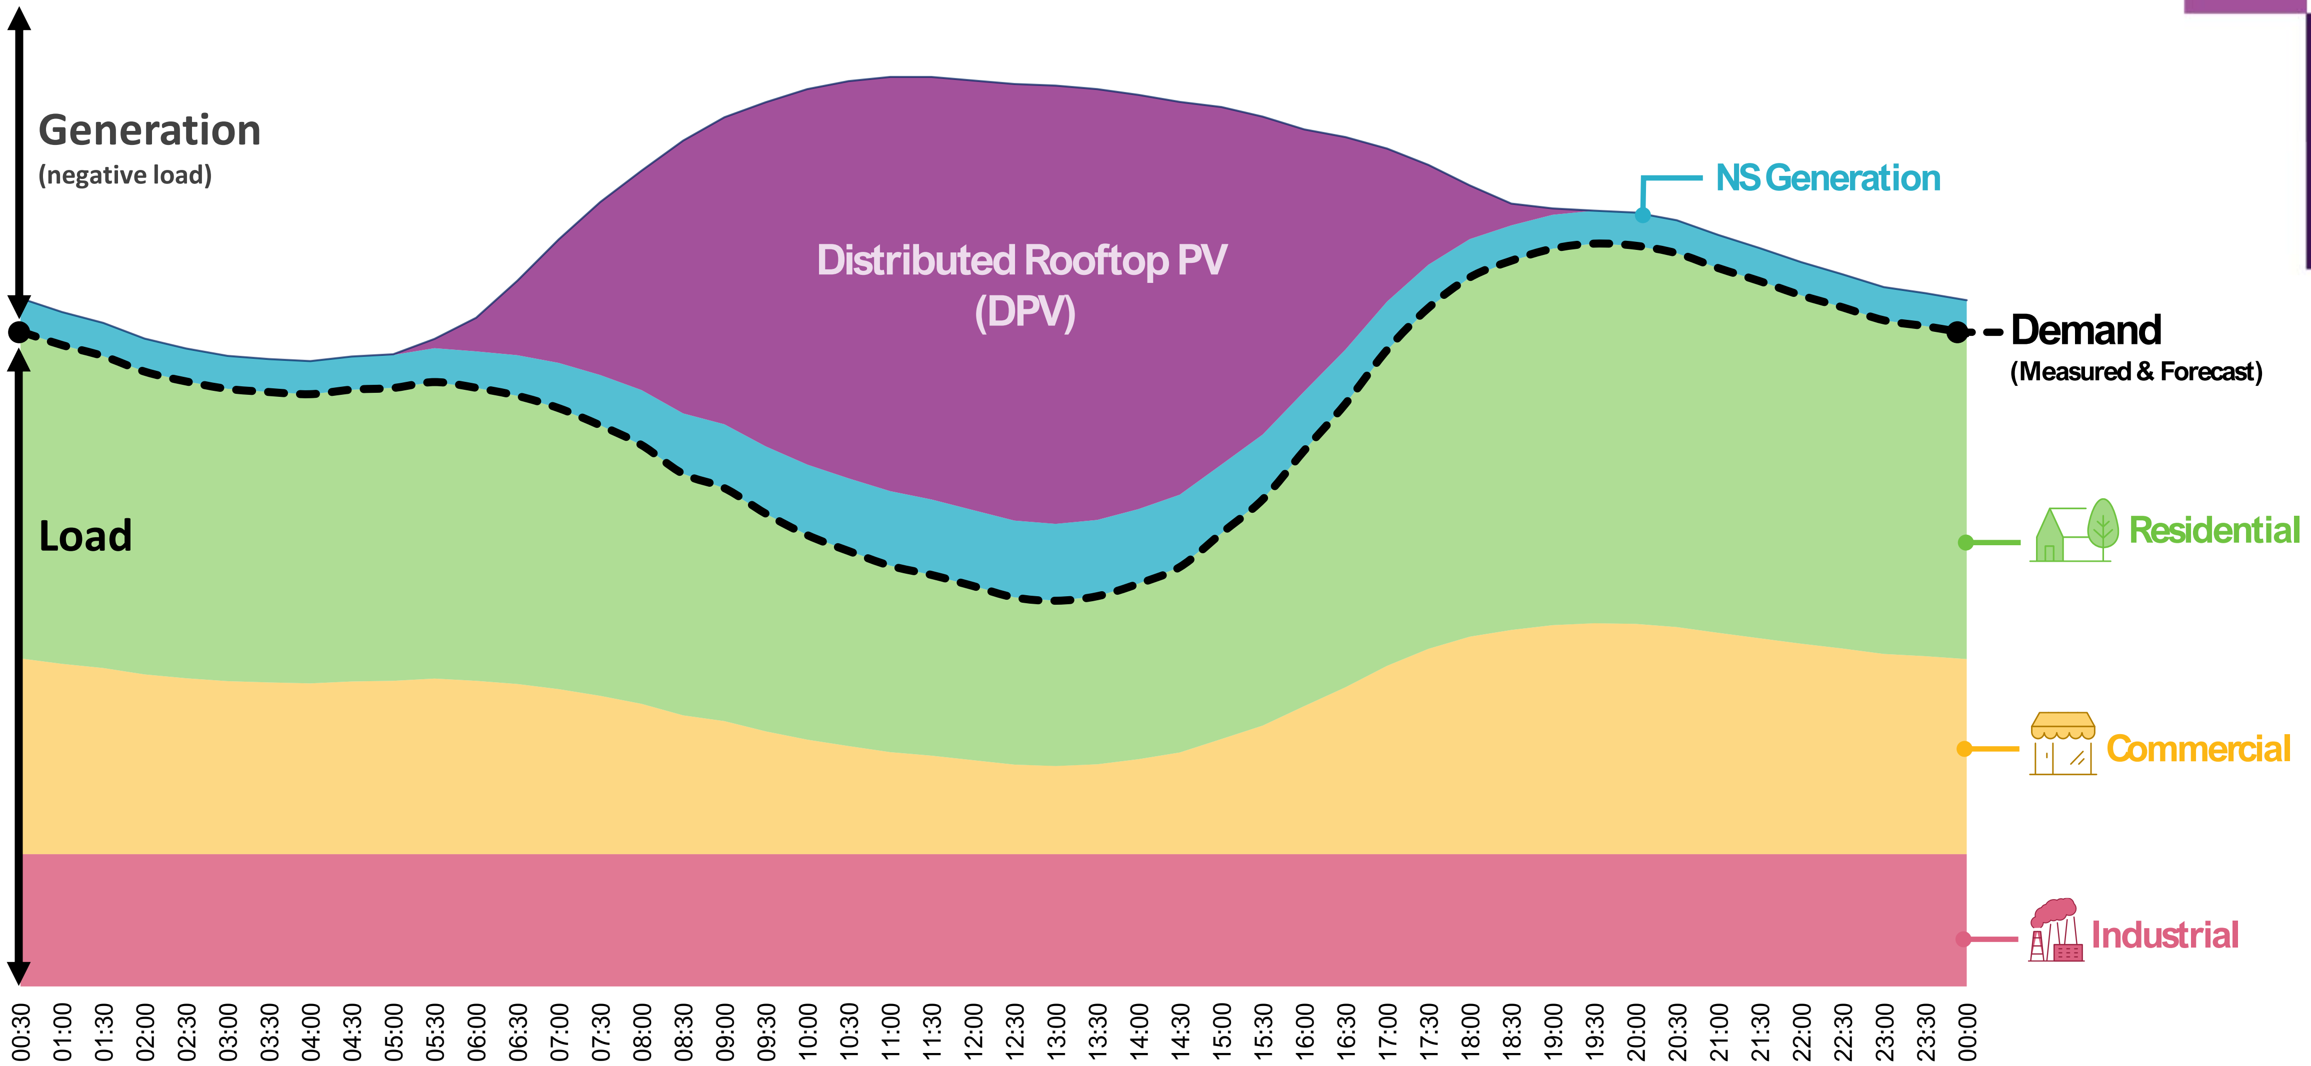

Operational Demand in the NEM

The demand of each NEM region is a mix of residential, commercial, and industrial.

SA, most weather sensitive region

- Low amount of industrial load
- High proportion of residential load
- Max DPV penetration >100%

QLD, diverse climate

- Highest amount of industrial load ~1GW

NSW, highest demand region

- Largest population
- High industrial load
- Highest DPV capacity

TAS, industrial

- Demand is dominated by large industrial loads
- Low DPV penetration

VIC, largely residential

- High proportion of residential demand due to low industrial activity
- Highest historical residential gas consumption

Our Foundation:

Linear Models & The Power of Simplicity

Demand Forecasting System (DFS)

Operational Forecasting use statistical models in the Demand Forecasting System (DFS) to produce demand forecasts for each region of the NEM. The forecast variables used by the models are designed to best predict how demand will change under certain conditions.

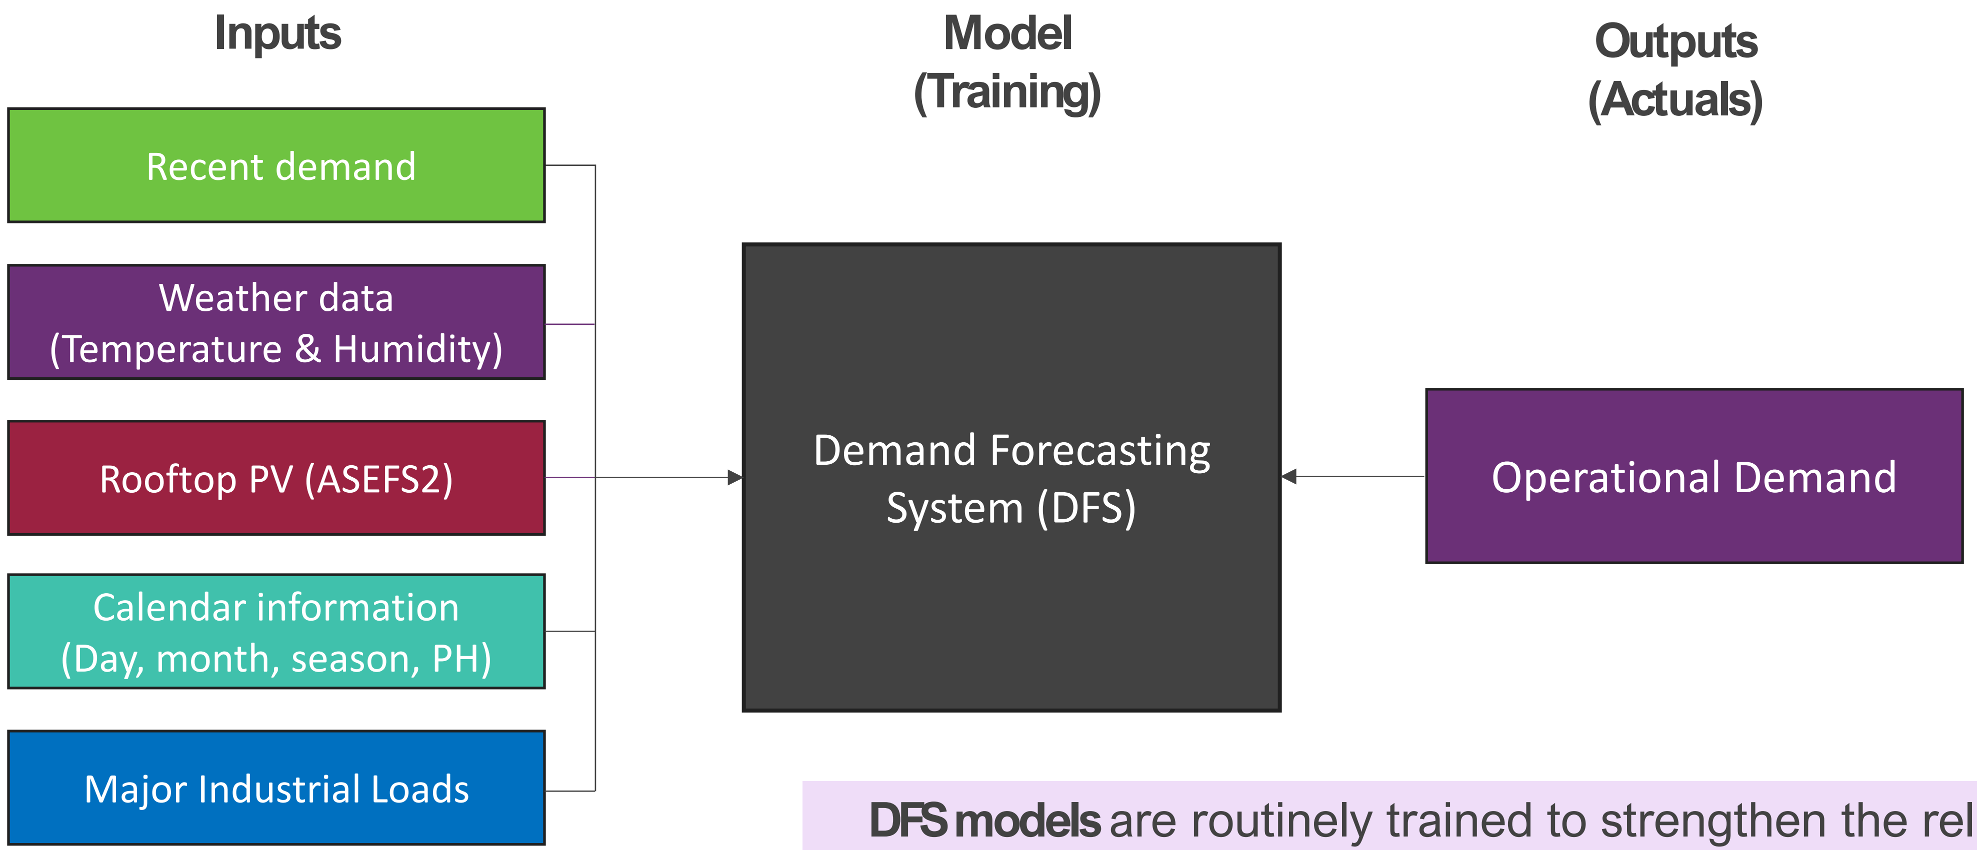

DFS models are routinely trained to strengthen the relationship between explanatory variables and target operational demand.

Linear Models and the Strength of Simplicity

- Our journey with Machine Learning began with a pragmatic approach. For years, our primary market-facing forecasting model has been, at its heart, a sophisticated linear regression model. It is incredibly effective.

- The value here wasn't just in its performance, but in its inherent characteristics:

  Simplicity: Linear models are straightforward to understand.

  Transparency: Their predictions are highly interpretable, allowing us to easily explain why a particular forecast was made.

- This simplicity meant easy management and clear insights. However, relying on a single, complex model, especially one implemented in third-party software, came with its limitations. **We lacked the flexibility to rapidly expand our model suite or adapt to new challenges without significant overhead.**

Boosting Ahead:

Embracing Open-Source Ensemble & Machine Learning Models

Where Boosted Models Take the Lead

- Recognizing the need for greater agility and a more robust modeling toolkit, we began to adopt open-source Machine Learning libraries. This allowed us to build a richer, more diverse suite of models.

- Specifically, we leaned into ***boosted tree models*** like XGBoost and LightGBM. These techniques empowered us with an ensemble of information, combining the strengths of many individual models to create a more powerful and nuanced prediction. We also began to apply ML techniques to develop probabilistic models and quantify uncertainty margins, moving beyond single-point forecasts.

- **Providing a consensus of different forecasts for Analysts revealed a more diverse outlook useful for situational awareness and preparedness**

- **Machine Learning Models also allowed Analysts to overwrite or refine forecasts in key situations where they know the Machine Learning models offer better accuracy & reliability (such as Public Holidays)**

The Power of Consensus Forecasts

There is no one-size-fits-all Approach

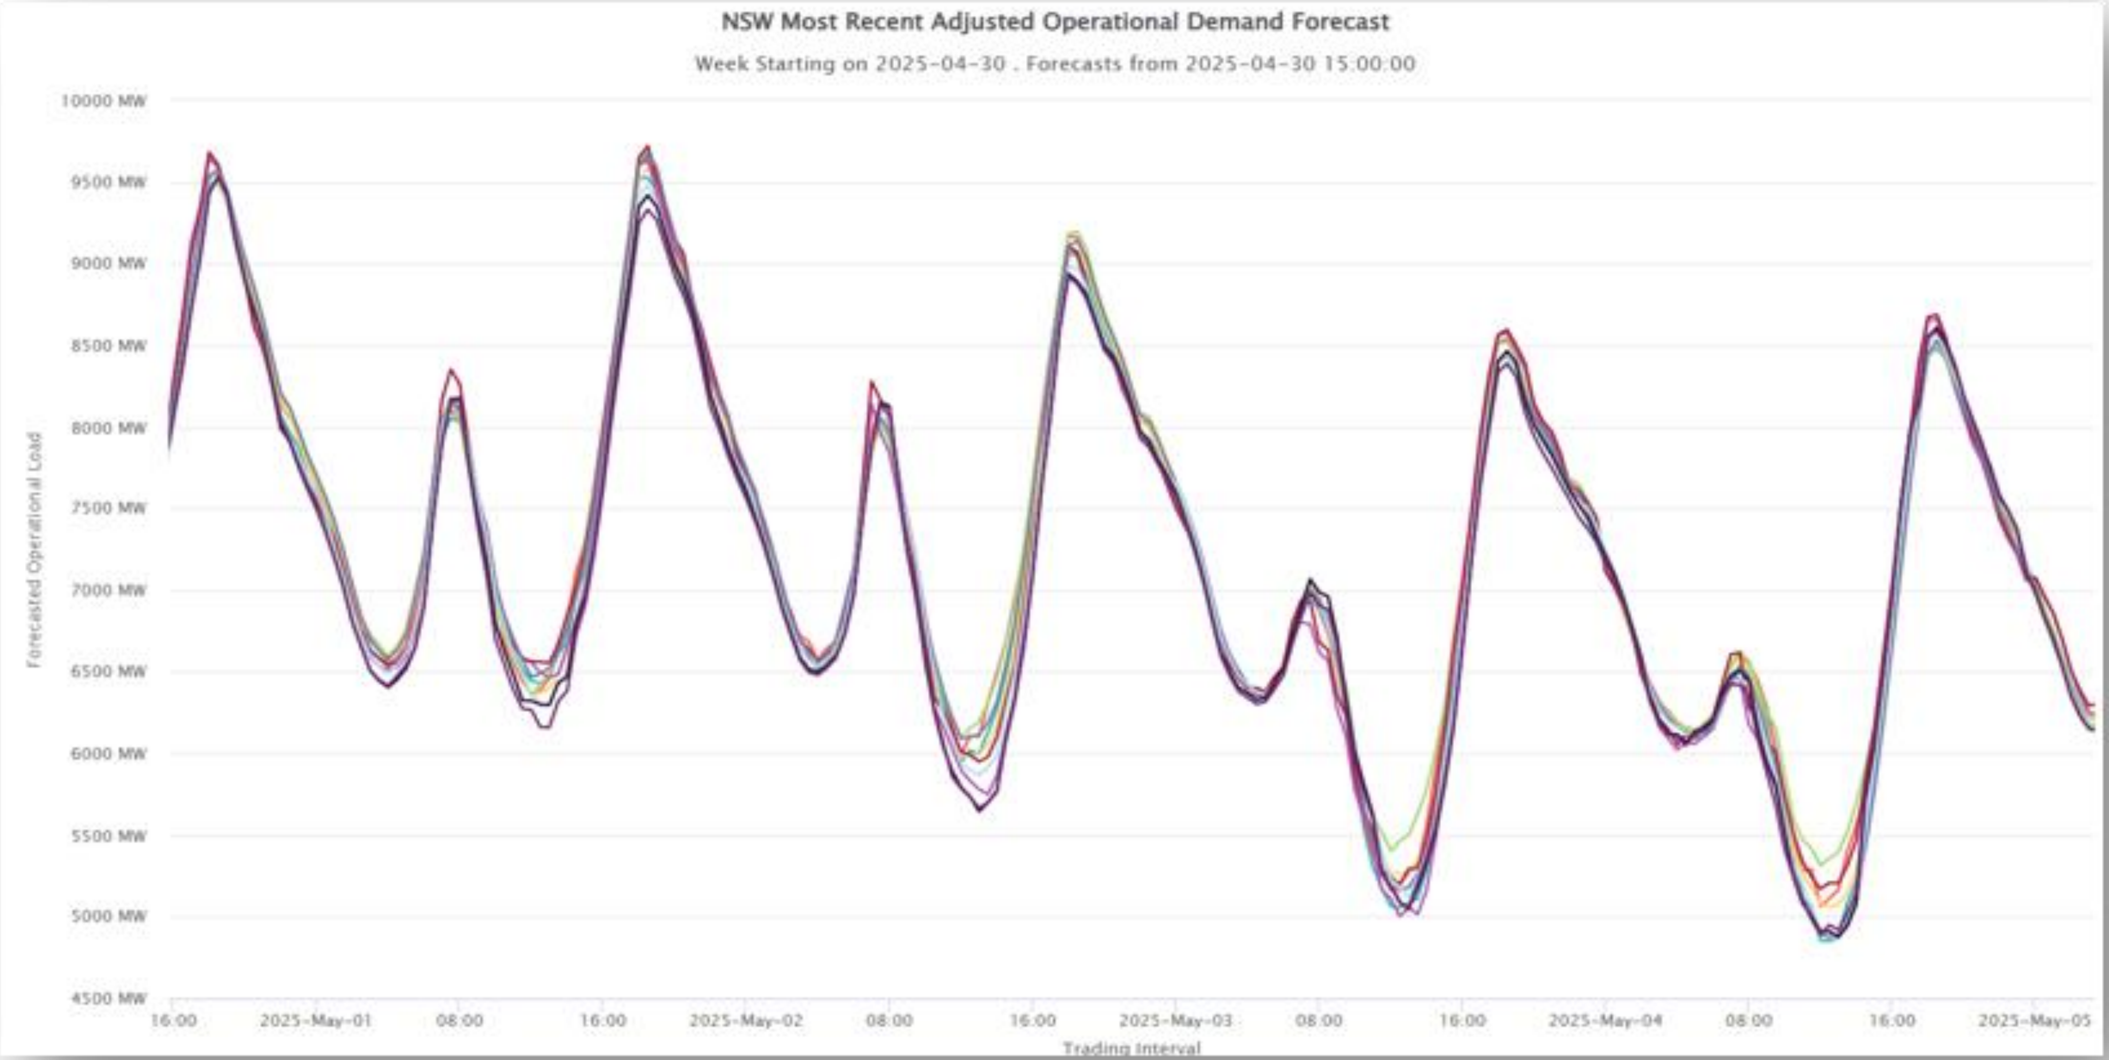

The Power of Consensus Forecasts
There is no one-size-fits-all Approach

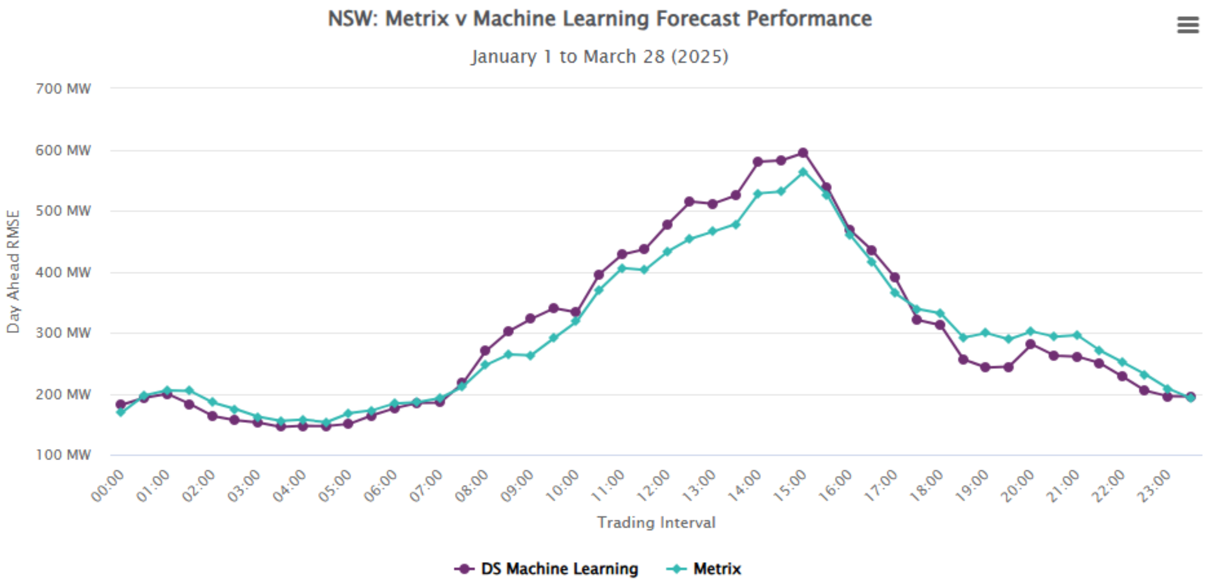

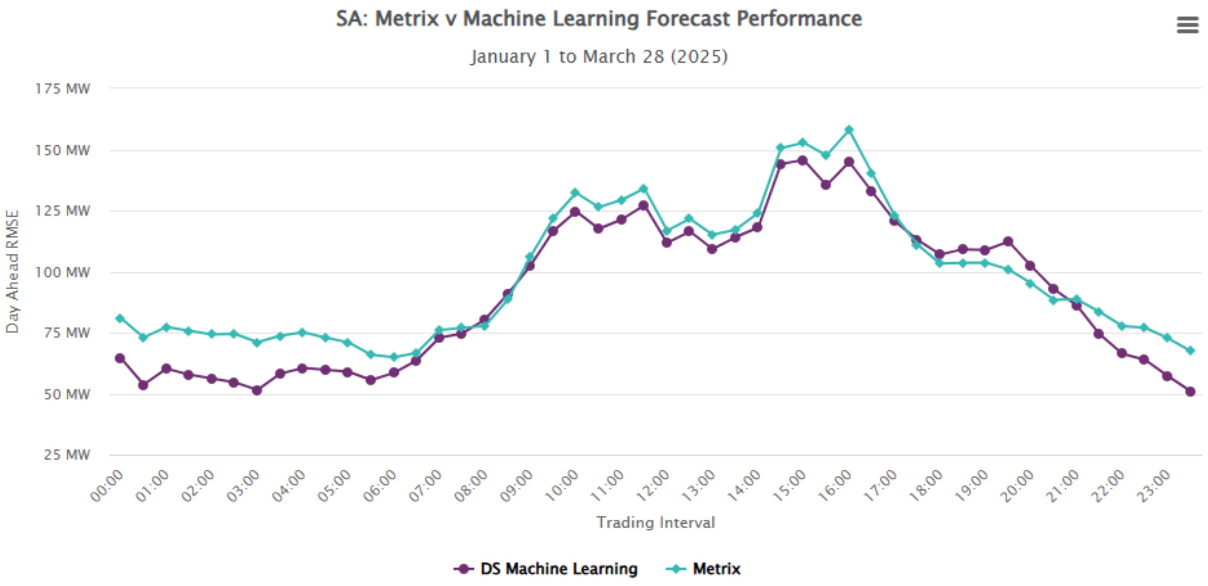

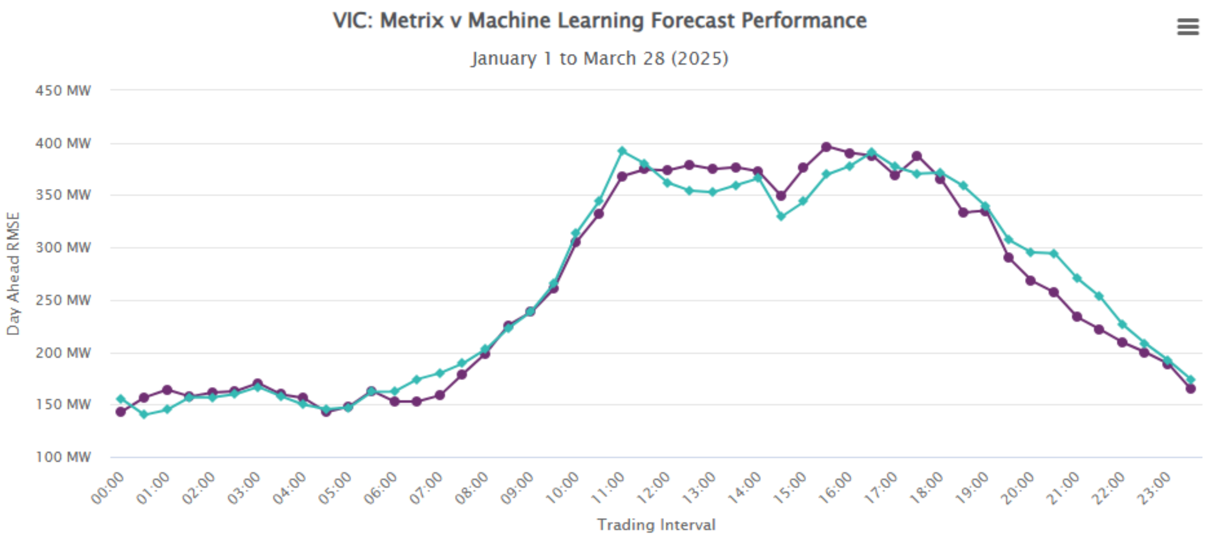

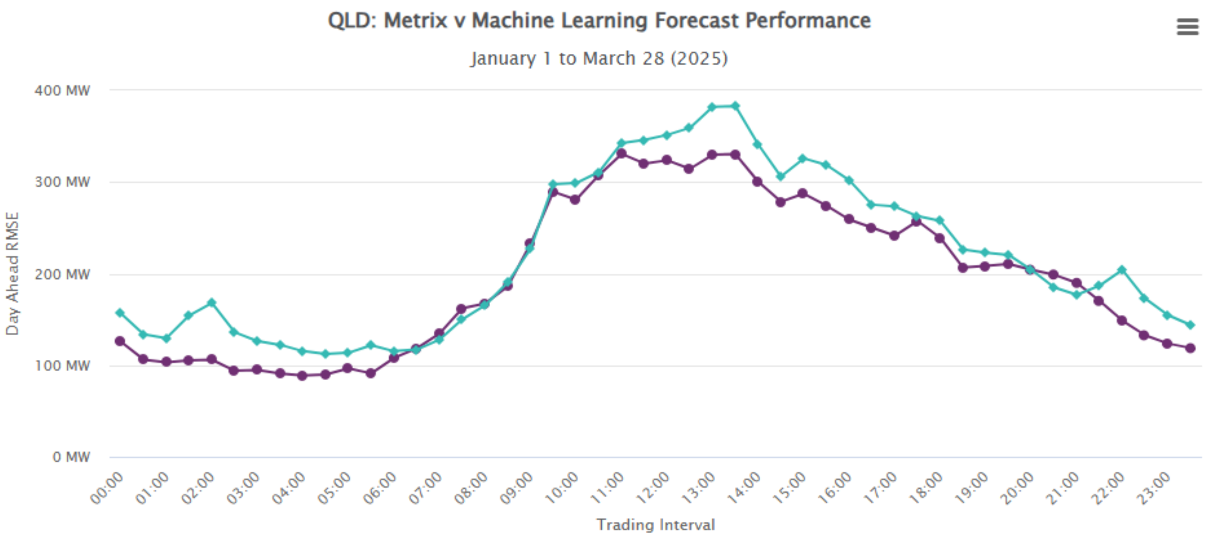

Limitations of Boosted Models

Even these advanced models had a key limitation: **their inability to extrapolate.**

Boosted Tree Models can struggle to forecast new minimums or maximums outside the range of their training data.

If new input data contained values unseen before, our models couldn't reliably predict outcomes far beyond their learned boundaries.

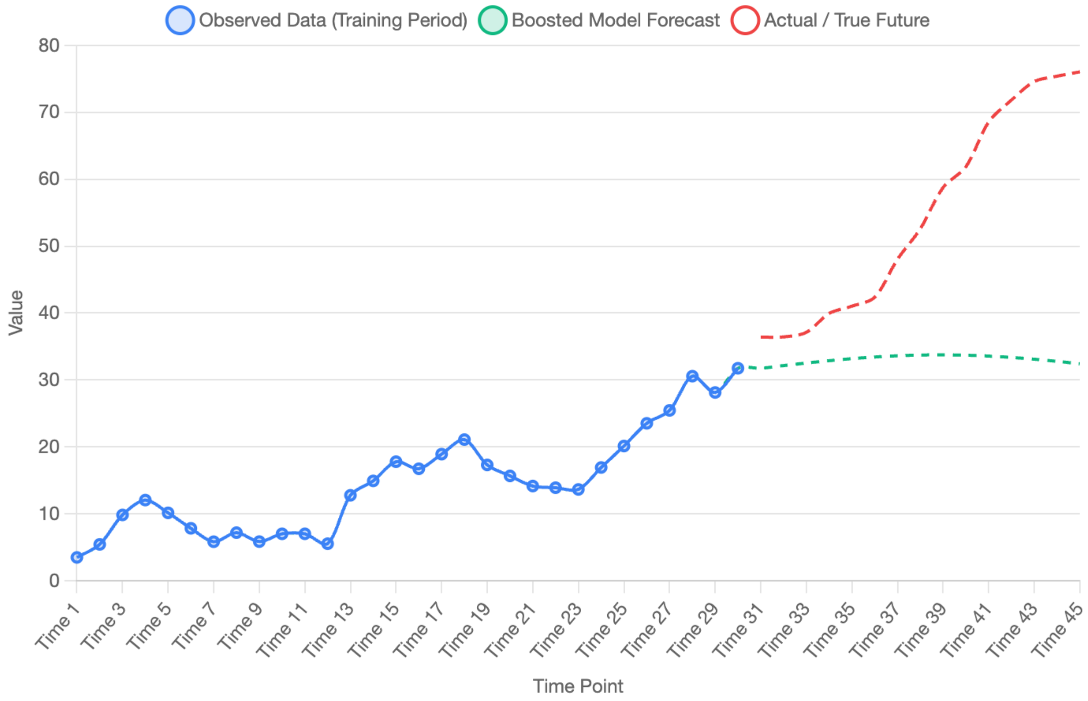

Expanding Our Horizons:

Deep Learning for Extrapolation

Deep Learning and Neural Network approaches

- Deep Learning and Neural Network models represent a significant leap forward. Their intricate architectures allow them to capture highly complex patterns and, crucially, possess a much greater capacity for extrapolation.

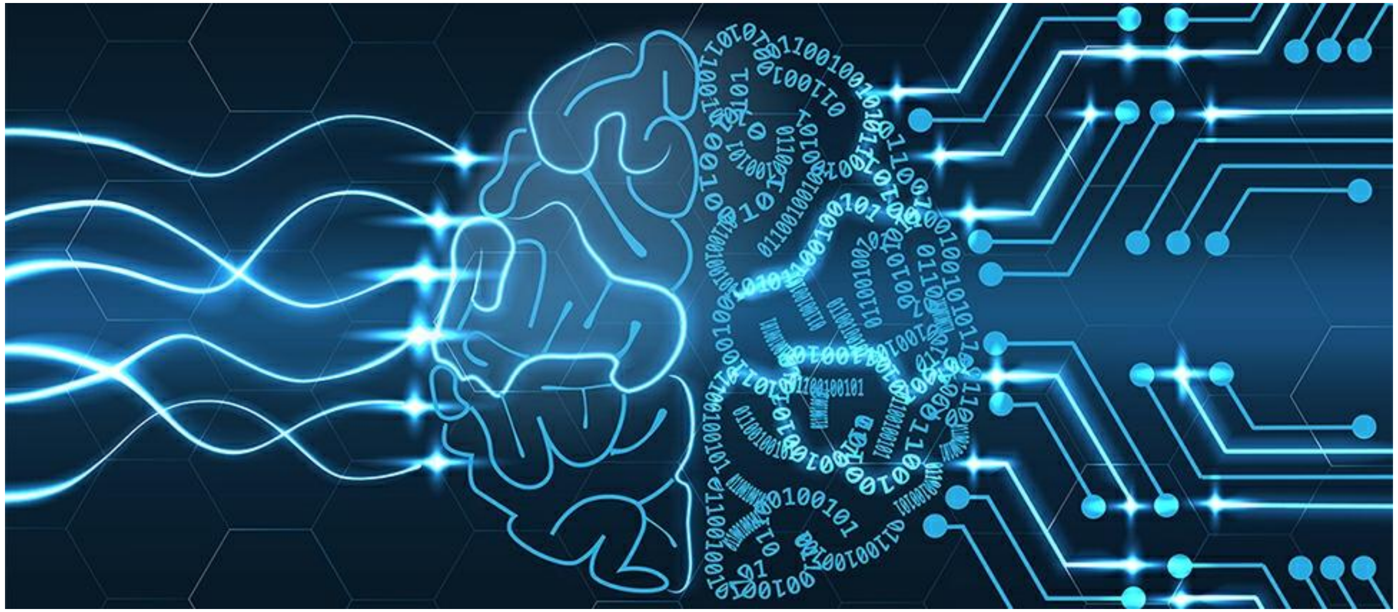

- This means we no longer have to compromise between achieving high accuracy and the ability to forecast truly novel scenarios.

Benefits of Neural Networks

- Specialised Deep Learning & Neural Network models such as LSTM show great potential for sequence/time-series forecasting.
- Deep learning models offer great performance across the board, and their inherent properties enable them to predict beyond the historical data range, providing valuable insights into unprecedented situations.
- Automatic feature extraction happens during training process, no fuss.

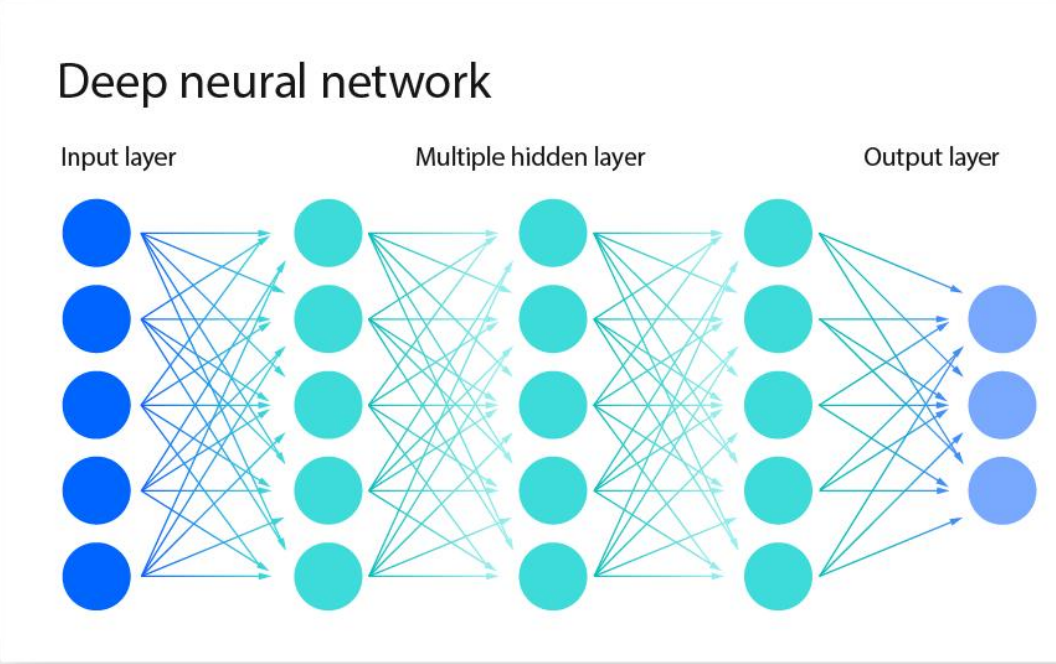

A Neural Network Case Study – South Australia Day Ahead Forecasts (Dec 21 & 22)

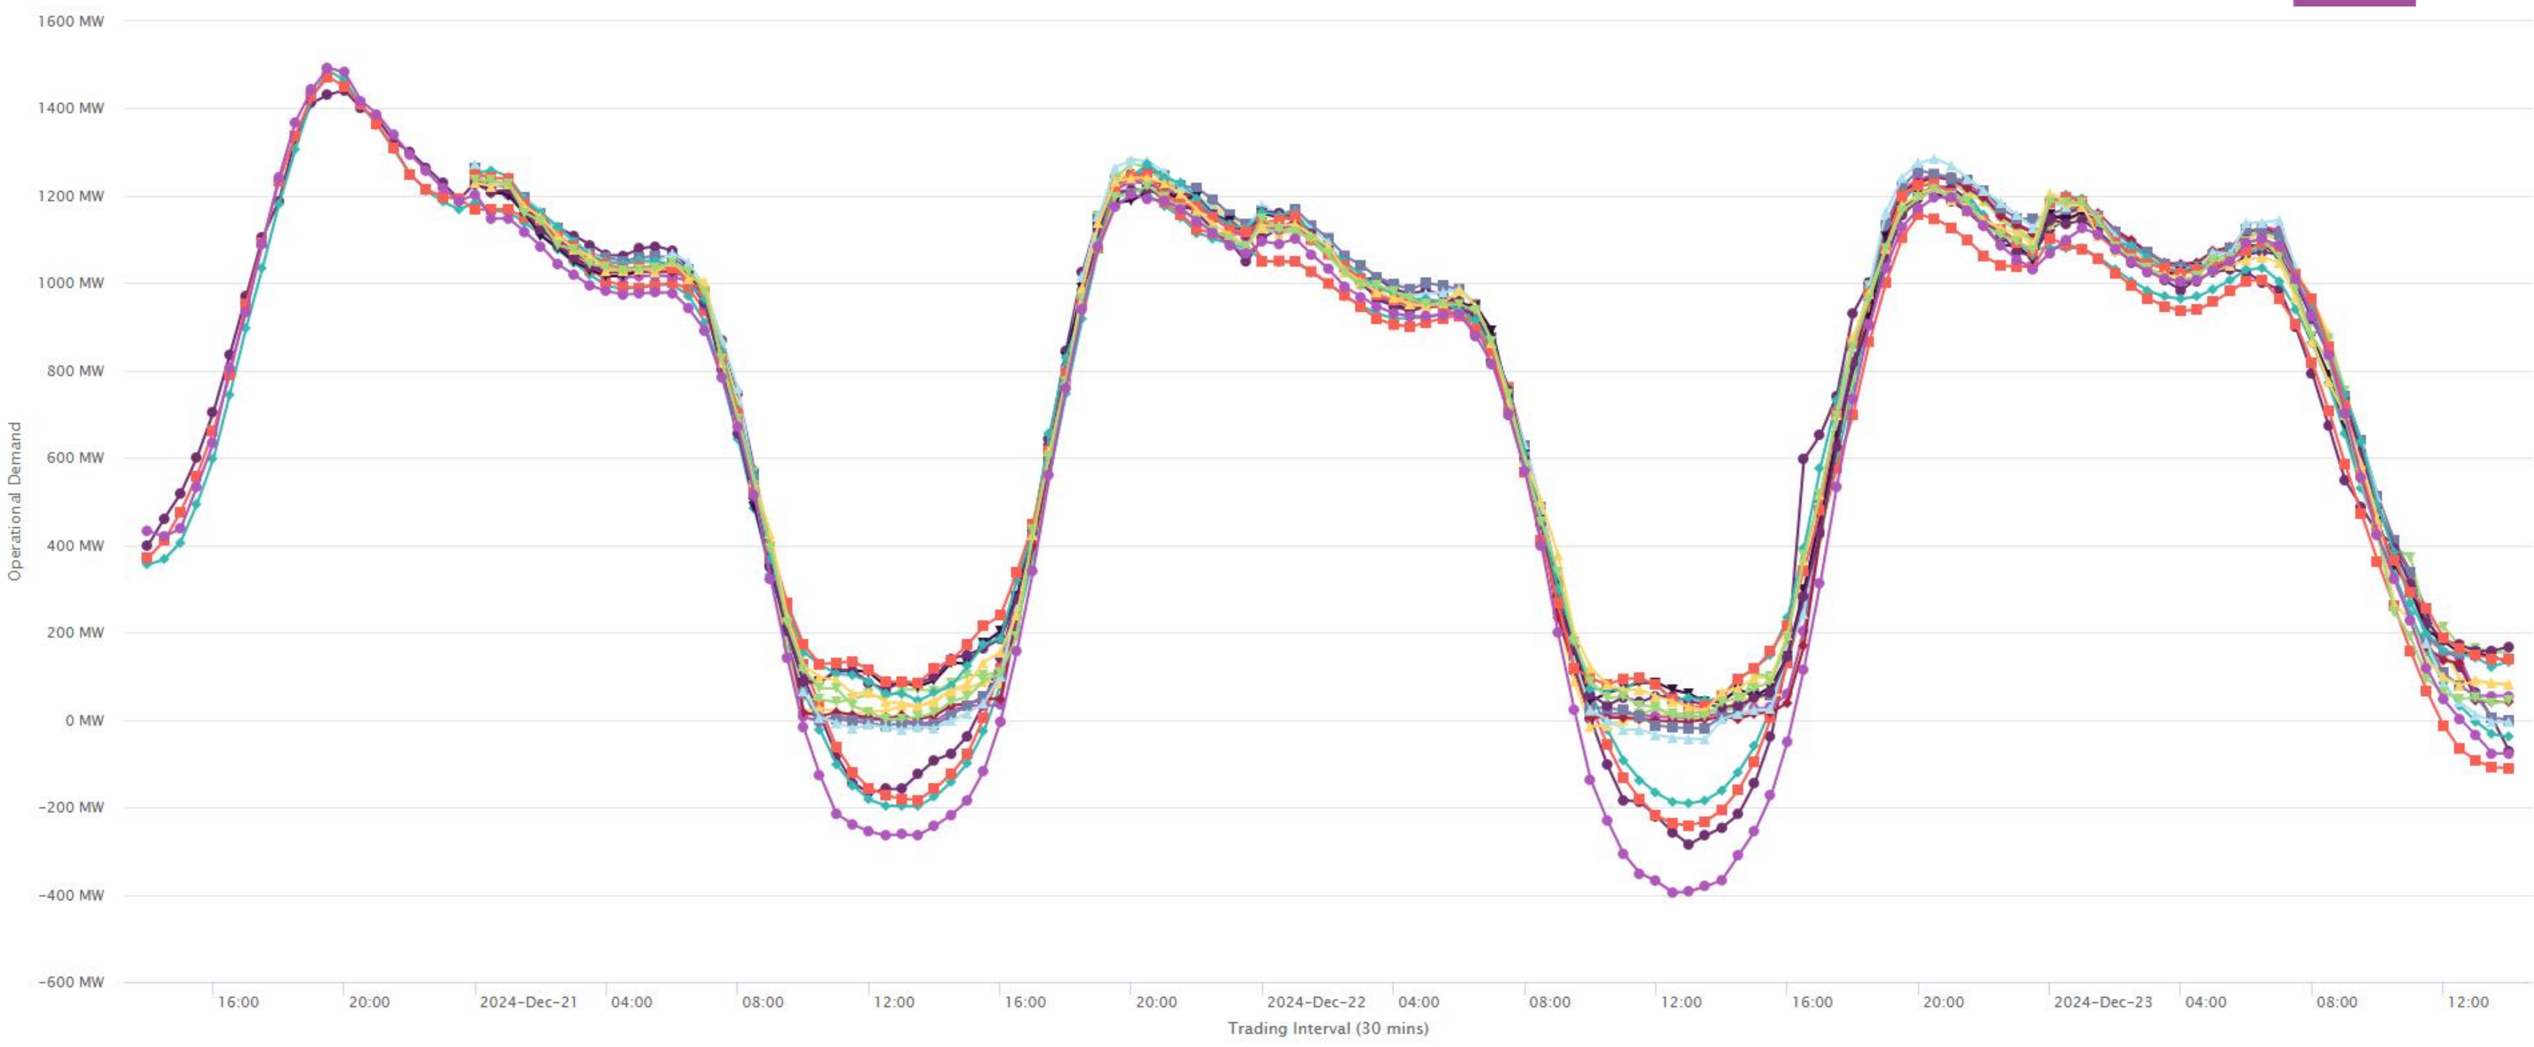

A Neural Network Case Study – South Australia Day Ahead Forecasts (Dec 21 & 22)

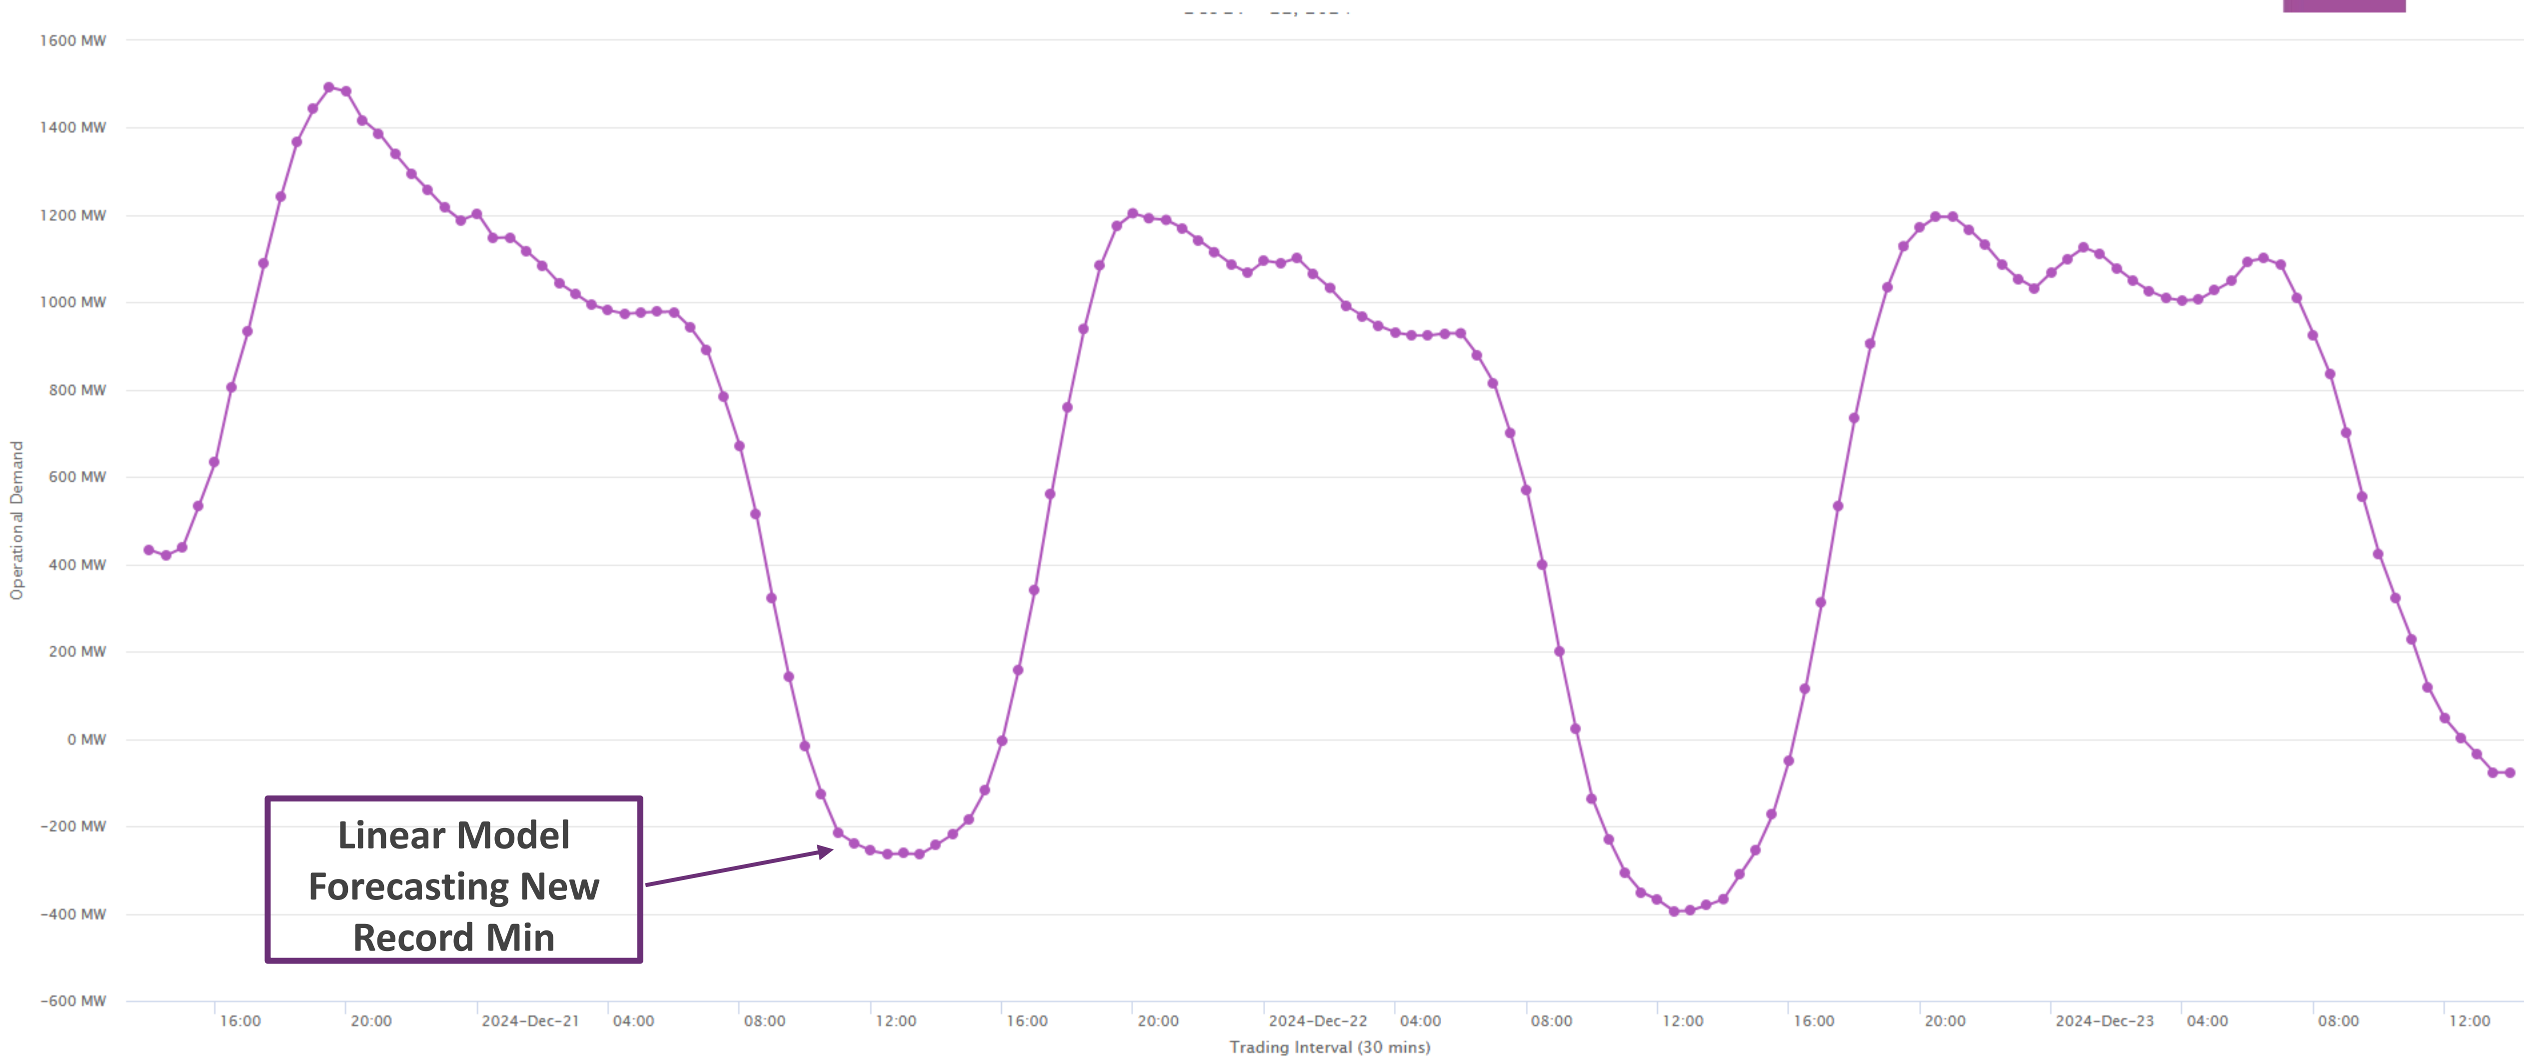

A Neural Network Case Study – South Australia Day Ahead Forecasts (Dec 21 & 22)

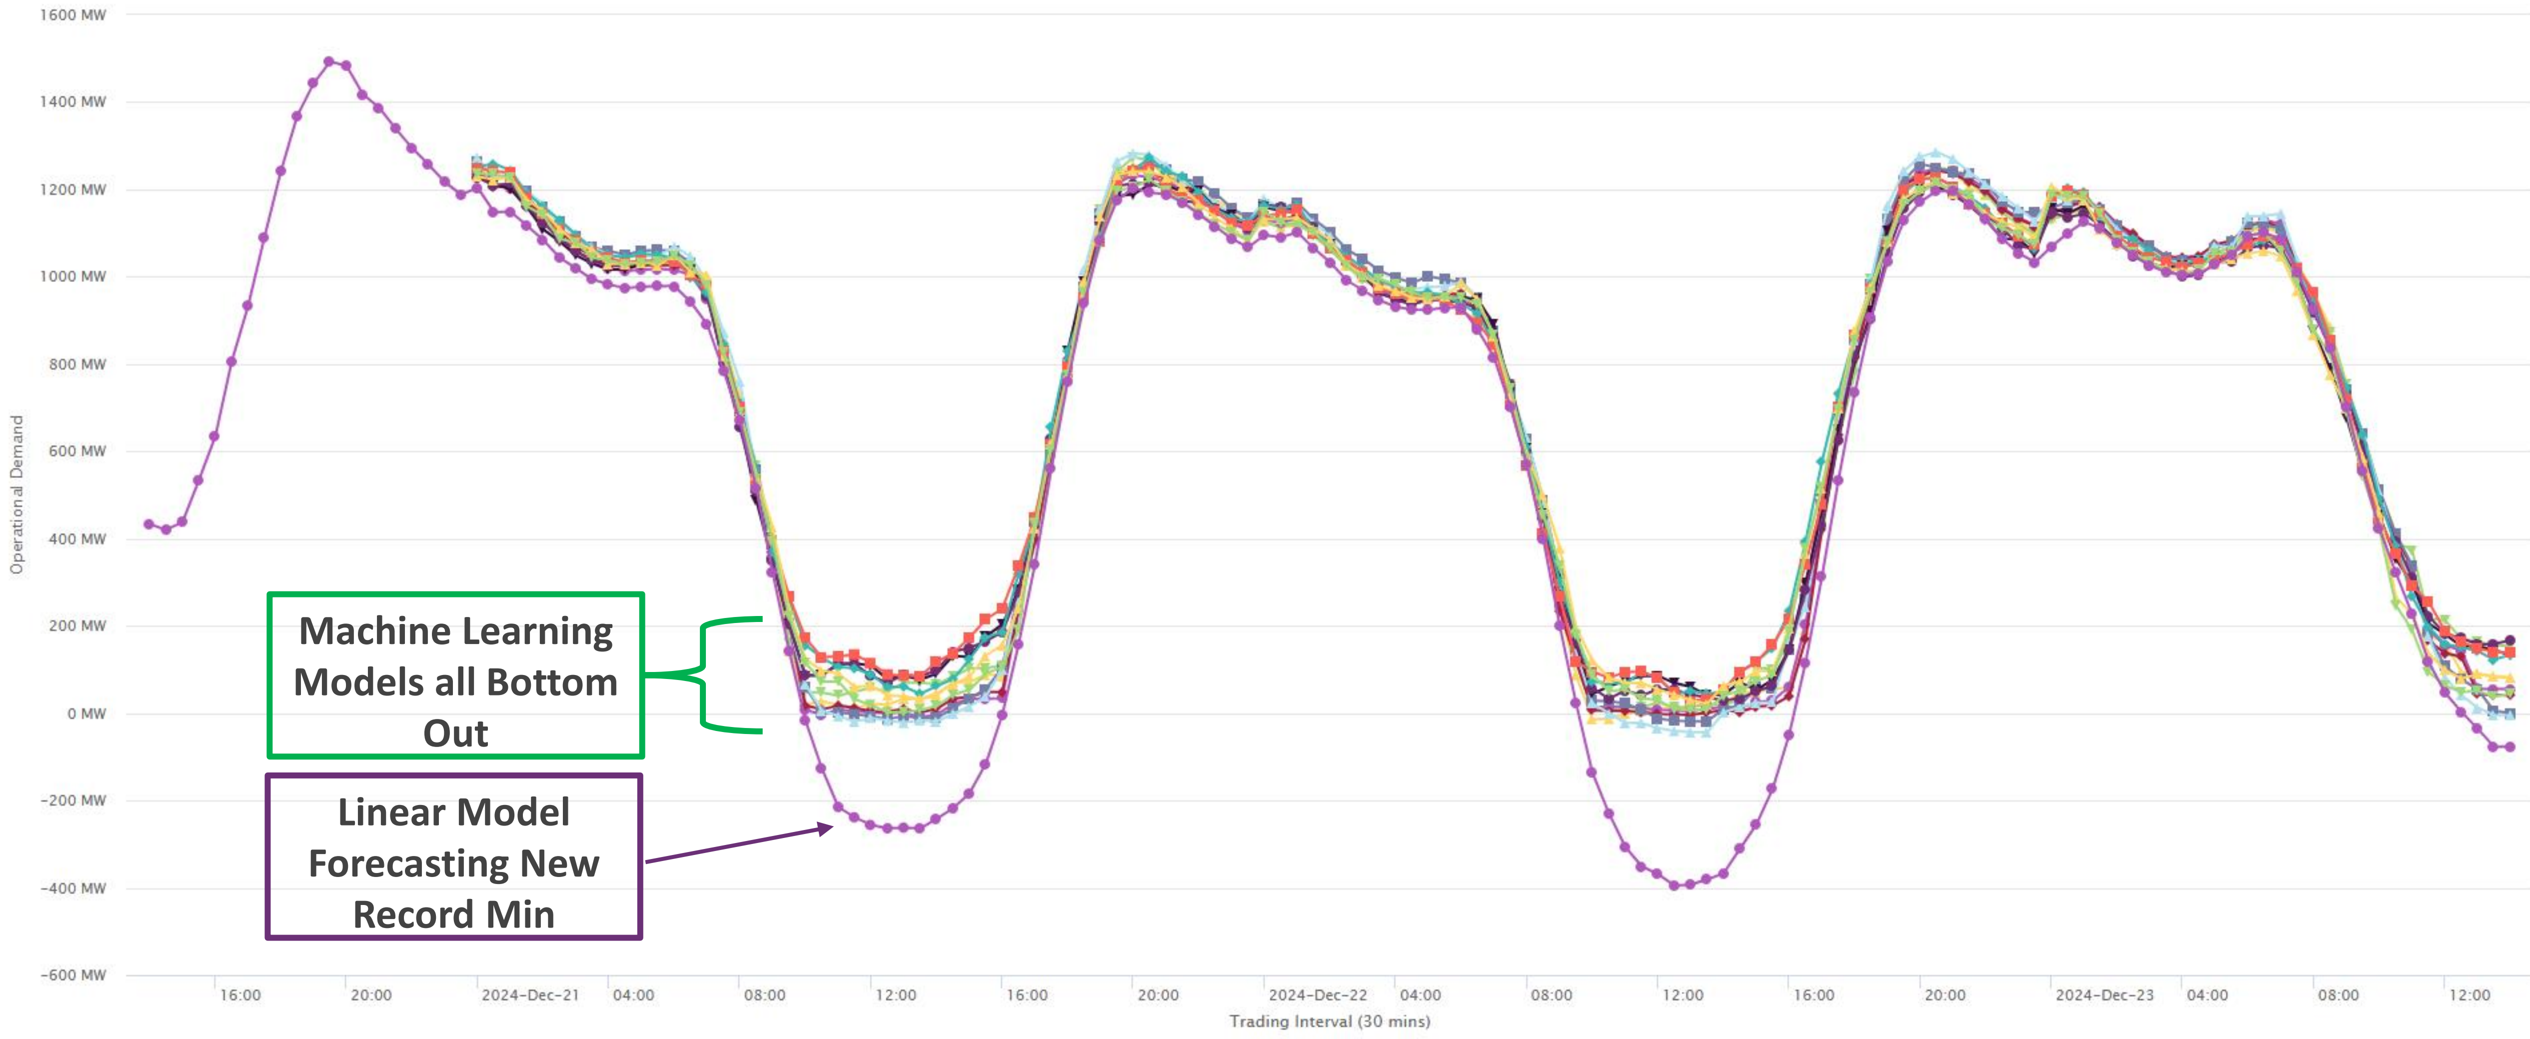

A Neural Network Case Study – South Australia Day Ahead Forecasts (Dec 21 & 22)

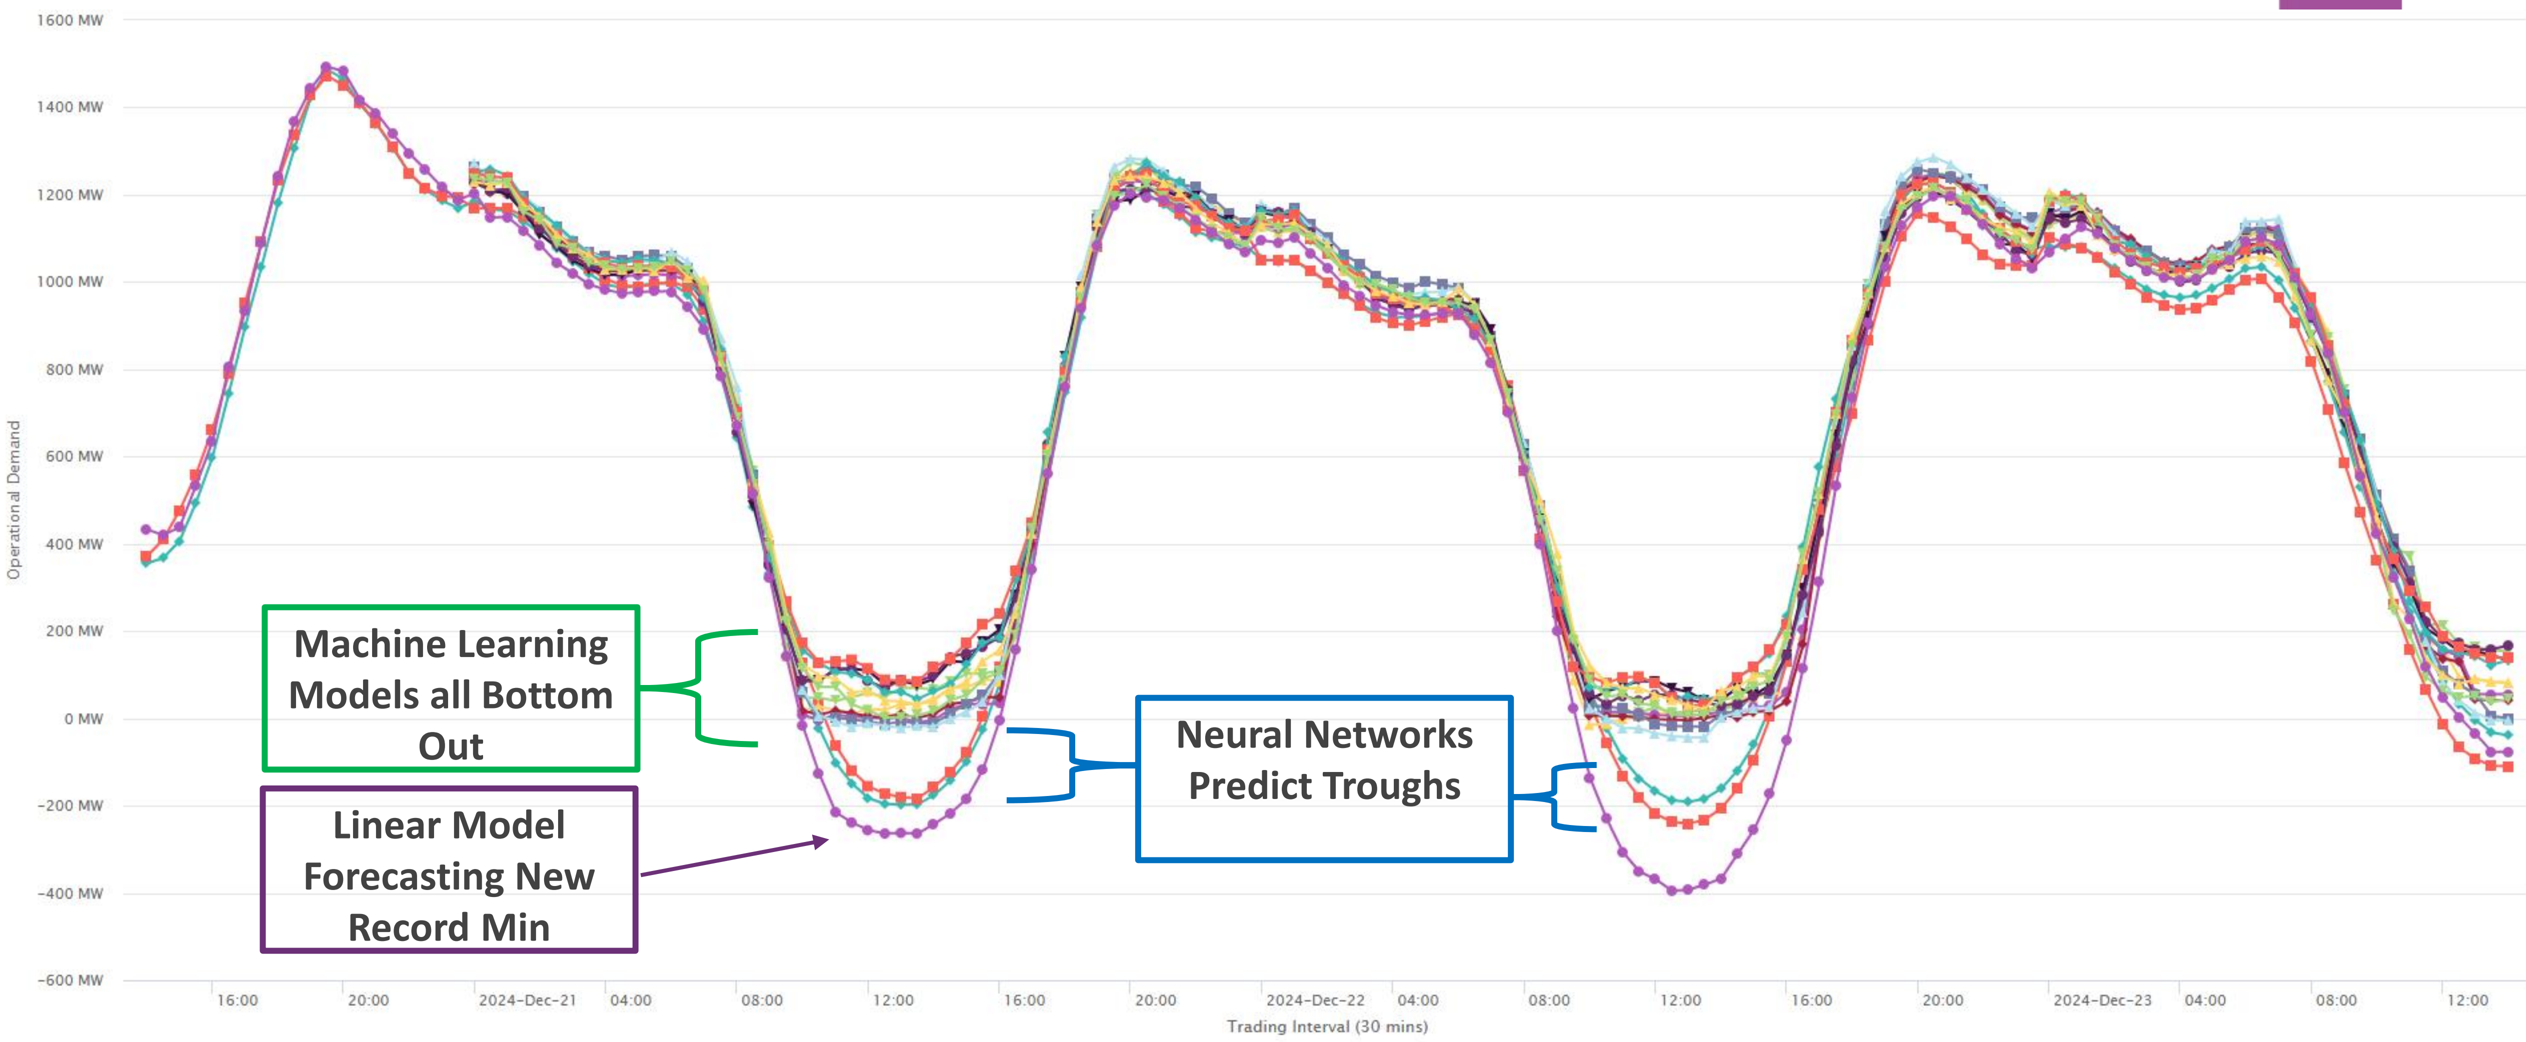

A Neural Network Case Study – South Australia Day Ahead Forecasts (Dec 21 & 22)

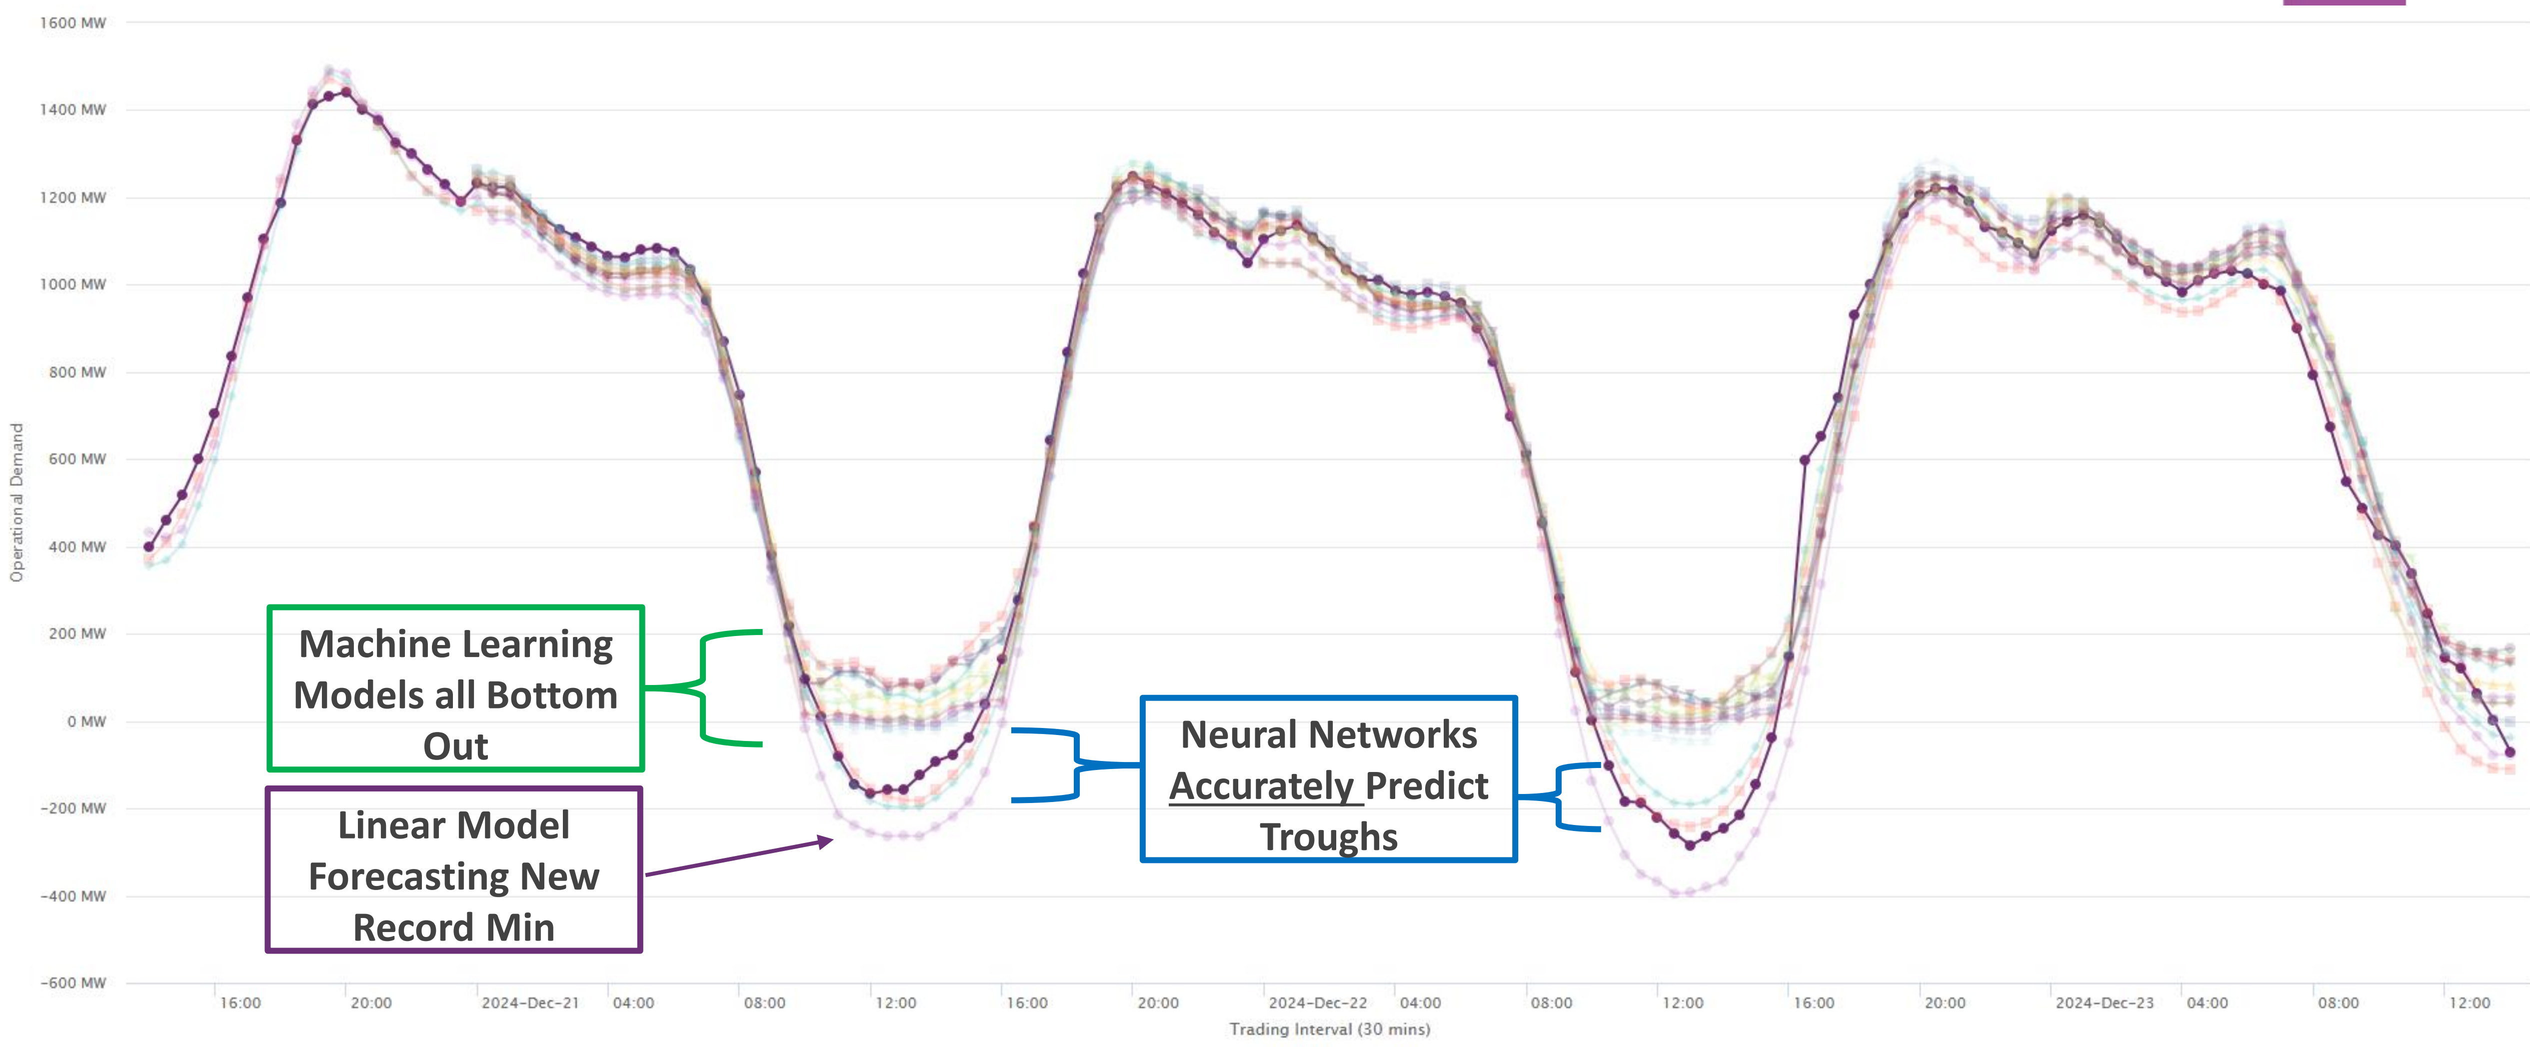

The Next Frontier:

Embracing The Black Box with Automated Machine Learning (AutoML)

Automated Machine Learning (AutoML)

- Looking ahead, our focus is actively shifting towards AutoML tools.
- While the term "black box" might traditionally be seen as a negative, we are now embracing it for specific applications. AutoML platforms allow us to automate the entire model development process, from data preprocessing and feature engineering to model selection and hyperparameter tuning.
- **The goal here is not to replace human expertise, but to rapidly identify and deploy the best-performing solutions for given problems, allowing our data scientists to focus on higher-level strategic challenges and model interpretation.**

Model Development Lifecycle

Current Machine Learning / AI lifecycle and process

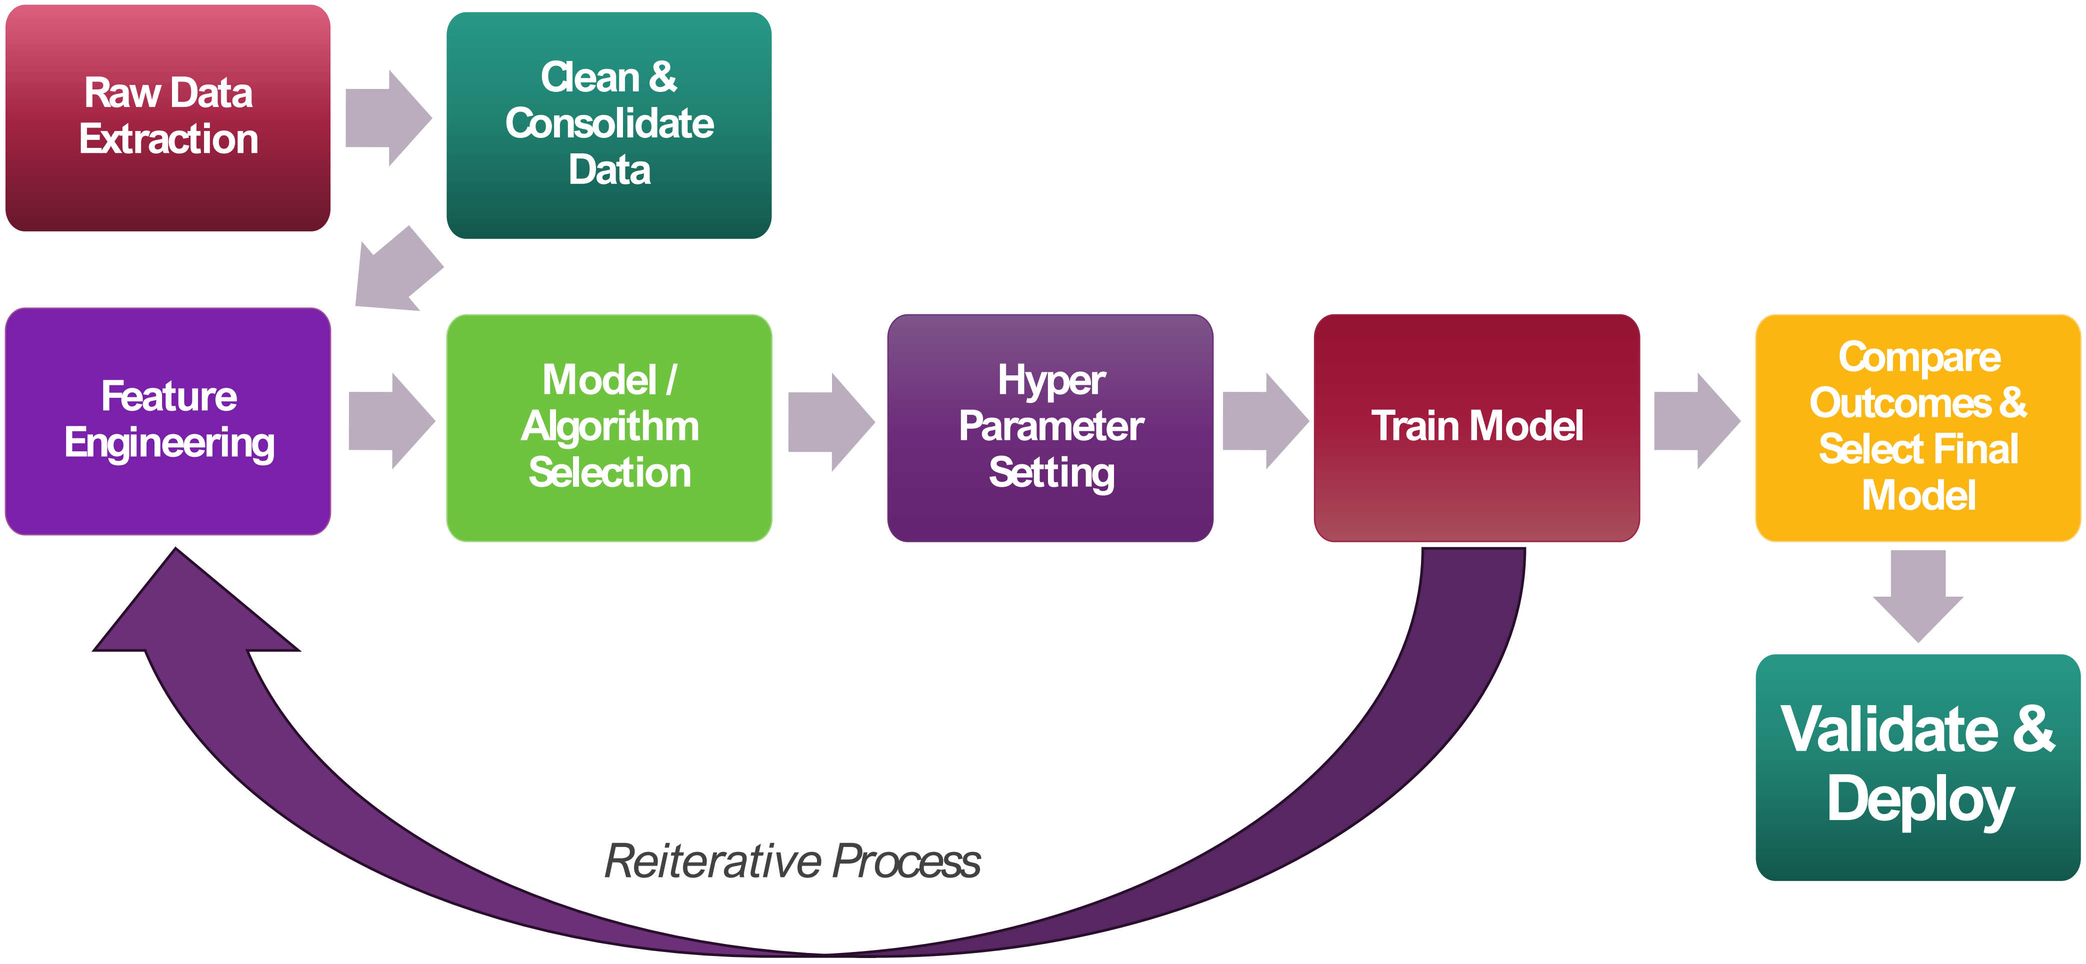

Model Development Lifecycle

Automated Machine Learning (AutoML) Lifecycle

Raw Data Extraction

Clean & Consolidate Data

AutoML Engine

(Feature Engineering / Model Selection / Hyperparameter Tuning / Model Training now automated)

- Enables faster, easier build and deployment of models
- Reduces resources & lowers barrier to entry for model development.
- Comprehensive explainability toolkit
- Scalable and adaptable to many use cases

Validate & Deploy

Looking Elsewhere – New & Emerging Opportunities in Machine Learning & AI

Recent Advances in Machine Learning

AEMO

- In the last 3-5 years have seen a shift from meticulously handcrafted time series models to more generalized, automated, and powerfully capable models driven by advancements in deep learning (especially Transformers) and the emergence of "foundation models" like TimeGPT.

- These new tools are making high-accuracy time series forecasting more accessible, efficient, and applicable in diverse, real-world organizational settings, particularly where data volume, complexity, or privacy concerns were once major barriers.

Recent Advances in Machine Learning

Advanced Deep Learning

How they're new/different: They overcome limitations of older RNNs (like the vanishing gradient problem) allowing them to "remember" patterns over much longer sequences. This is crucial for time series with complex seasonality or long-term trends.

Practical Use: Widely used for electricity demand forecasting (e.g., predicting energy load based on historical consumption, temperature, and day of the week), product demand prediction in retail, and traffic flow forecasting.

Transfer Learning

How it's new/different: It involves taking a pre-trained model (like TimeGPT, or a model trained on a large dataset from a similar domain) and fine-tuning it with a smaller, specific dataset. This is highly beneficial when you have limited historical data for a new forecasting problem.

Practical Use: A new startup entering the energy market can leverage a model pre-trained on vast public energy consumption data and then fine-tune it with their specific, albeit limited, customer data to quickly get accurate demand forecasts.

Recent Advances in Machine Learning

Reinforcement Learning (RL)

How it's new/different: RL agents can be designed to learn optimal strategies for combining different forecasting models (ensemble learning) or even for dynamically augmenting data, especially in "few-shot" learning scenarios where data is scarce. This means the RL model can learn how to forecast better given different situations.

Practical Use: For time series forecasting, combining multiple models (like a linear model, a boosted tree, and a deep learning model) often yields better results than a single model. RL can learn to dynamically adjust the weights given to each base model in the ensemble based on the current time series characteristics, adapting to non-stationary data distributions.

Deep Reinforcement Learning (The combination of RL with Deep Neural Networks)

How it's new/different: DRL allows the agent to learn directly from raw time series data (e.g., sequences of sensor readings, financial ticks) without needing extensive manual feature engineering. Models like LSTMs, GRUs, and even Transformers are used within the RL framework to process the time series input and inform the agent's decisions.

Practical Use: In Energy Management & Grid Control, RL can optimize energy scheduling and distribution in smart grids. An RL agent can learn to dispatch power from various sources (solar, wind, battery storage, traditional plants) in real-time, balancing supply and demand, minimizing costs, and ensuring grid stability given fluctuating renewable energy inputs and demand forecasts.

GraphCast: AI model for faster and more accurate global weather forecasting

GraphCast is an AI model built on Graph Neural Networks (GNNs). Instead of solving intricate physics equations, it learns patterns directly from decades of historical weather data (like temperature, wind speed, and pressure). It takes the current global weather state, along with the state from a few hours prior, and predicts how the weather will evolve in the future. This process can then be "rolled forward" to generate forecasts for up to 10 days ahead.

In essence, GraphCast leverages the power of AI to learn from vast amounts of historical data, providing faster, more accurate, and more energy-efficient weather forecasts that have significant practical implications for industries relying on timely and precise weather information, from disaster preparedness to renewable energy management.

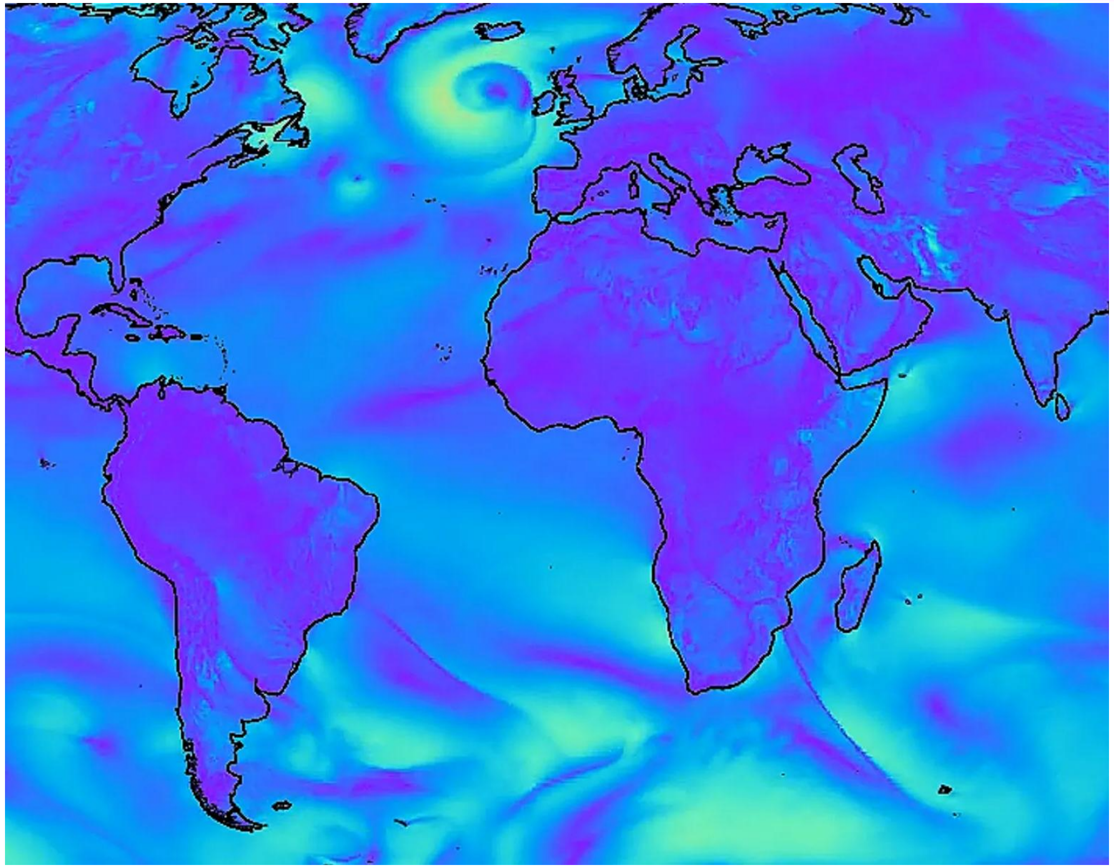

GraphCast: Value and Significance

- **Speed & Efficiency:** GraphCast can generate a 10-day global weather forecast in under a minute using significantly less computational power (e.g., a single AI accelerator chip) compared to the hours and massive supercomputers required for traditional methods. This also translates to being approximately 1,000 times more energy-efficient.
- **Accuracy:** It has demonstrated unprecedented accuracy, often outperforming the industry's gold-standard traditional forecasting systems (like the European Centre for Medium-Range Weather Forecasts' HRES) on a large majority of weather variables and lead times.
- **Earlier Warnings for Extreme Weather:** GraphCast excels at predicting severe weather events, such as the tracks of tropical cyclones and the onset of extreme temperatures, earlier than conventional models. This provides crucial extra lead time for communities and emergency services to prepare.
- **Accessibility & Democratization:** Its reduced computational requirements mean that high-quality weather forecasting could become more accessible to a wider range of organizations and regions globally, including those without access to multi-million dollar supercomputing infrastructure.
- **Complementary Tool:** While powerful, it's generally seen as a complement to, rather than a full replacement for, traditional physics-based models. Meteorologists can use AI forecasts like GraphCast as an additional, fast, and accurate source of information to enhance their overall predictions and provide better insights.
- **Open Source:** Google DeepMind has made GraphCast open-source, encouraging further research, development, and adaptation by the scientific and meteorological communities worldwide.

Our journey is one of continuous evolution, driven by the desire to leverage the most effective tools to provide the most accurate and insightful forecasts for our market

Thank You

Contact Information

Andrew.Akman@aemo.com.au

Operational.Forecasting@aemo.com.au

AUSTRALIAN ENERGY MARKET OPERATOR

For more information visit

aemo.com.au

Appendix

National Electricity Market (NEM)

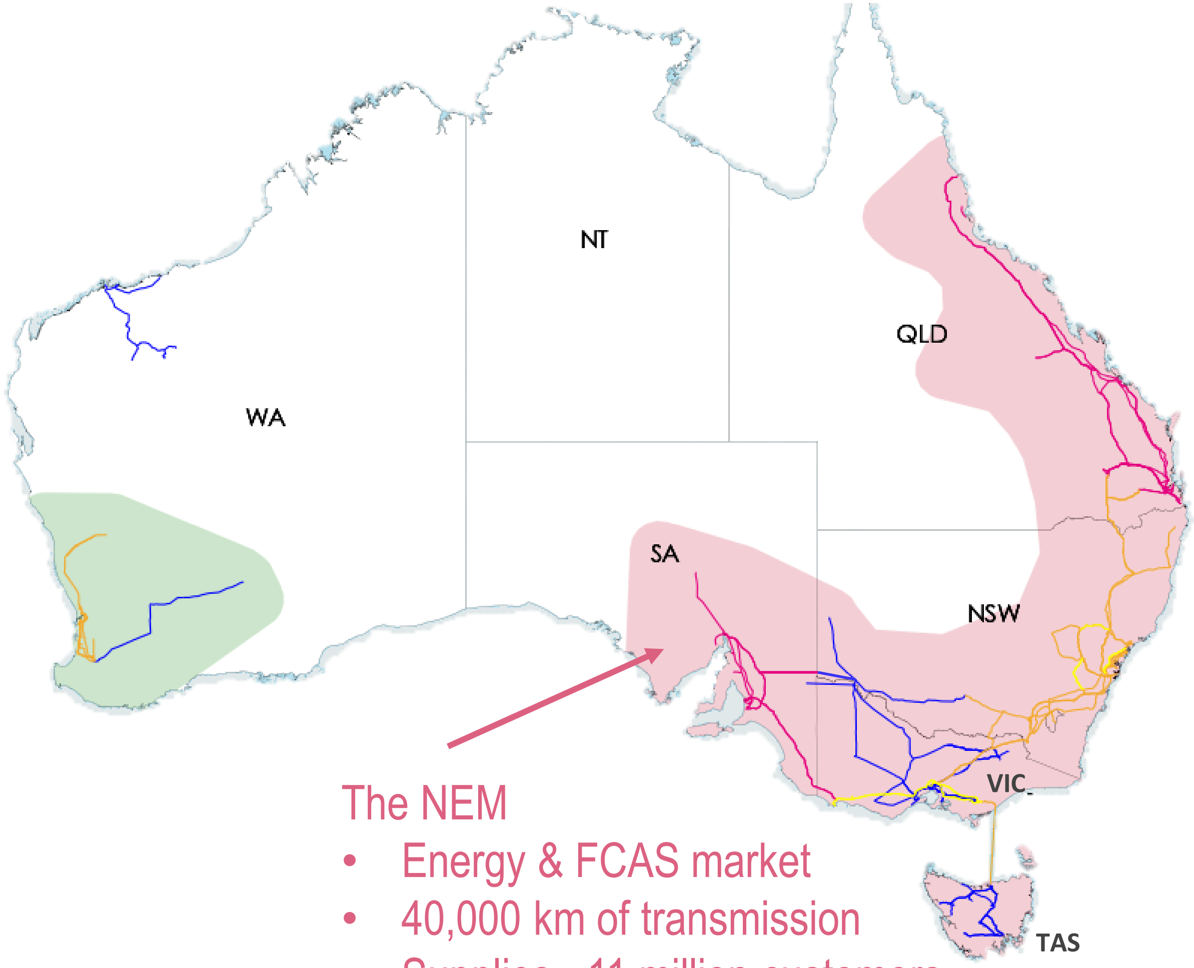

The NEM
- Energy & FCAS market
- 40,000 km of transmission
- Supplies ~11 million customers
- 71 GW of generation capacity

Population **23.5 million**

NEM Max Demand **35,796 MW**
NEM Min Demand **10,073 MW**
Wind Capacity **13,613 MW**
Solar Capacity **11,637 MW**
Rooftop Solar Capacity **22,780 MW**

Approximately **3.5 million** homes have a solar system

Managing Operational Forecasts

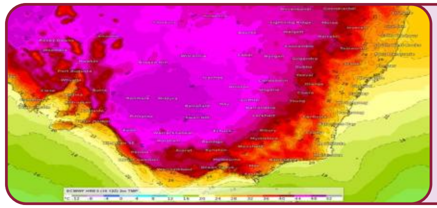

Check the weather, demand, wind & solar for the day

Weather risks and possible scenarios are assessed by analysts, with information provided by meteorologists and data from our data science team. The demand, wind and solar forecasts are reviewed by the analysts.

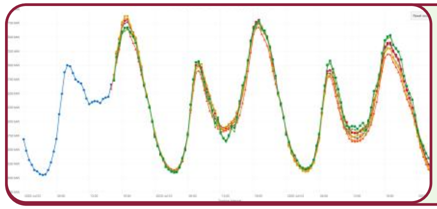

Translate the weather forecasts into power system impacts

How the weather will impact demand, wind and solar is interpreted by the analysts, and any severe weather risks which could damage assets or limit generation is assessed.

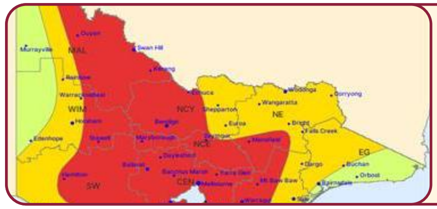

Communicate the forecast & associated risks

The forecast outcomes and weather-related risks are communicated to the wider business including our real time control room through a variety of methods.

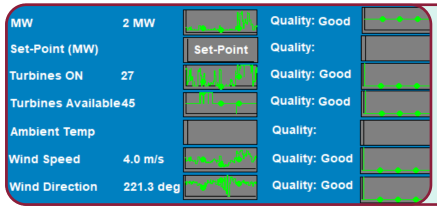

Assist in the management of wind and solar farms, help manage the demand forecasts, monitor the current situation

Wind and solar forecast issues are investigated, and generation risks are communicated. Demand forecasts are continually monitored and managed with our real time control room.

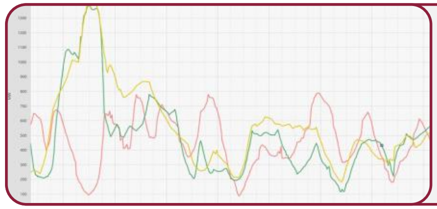

Review the outcomes of the day, adjust for tomorrow

Any demand forecast adjustments are made and managed, and the forecast models are continually improved and monitored.

Wider Operational Forecasting Functions

- Situational awareness for control room and wider AEMO
- Supporting AEMO's onboarding and connections process
- Provide input to Market Reform initiatives
- Reporting and documentation
- Ongoing improvement of tools and processes

Demand Forecasting for the NEM

New demand forecasts are produced:

- Every 5 minutes for Dispatch & 5-minute Pre-Dispatch, covering one hour at 5-minute resolution.
- Every 30 minutes for Pre-Dispatch & STPASA, covering eight calendar days at a half-hourly resolution.

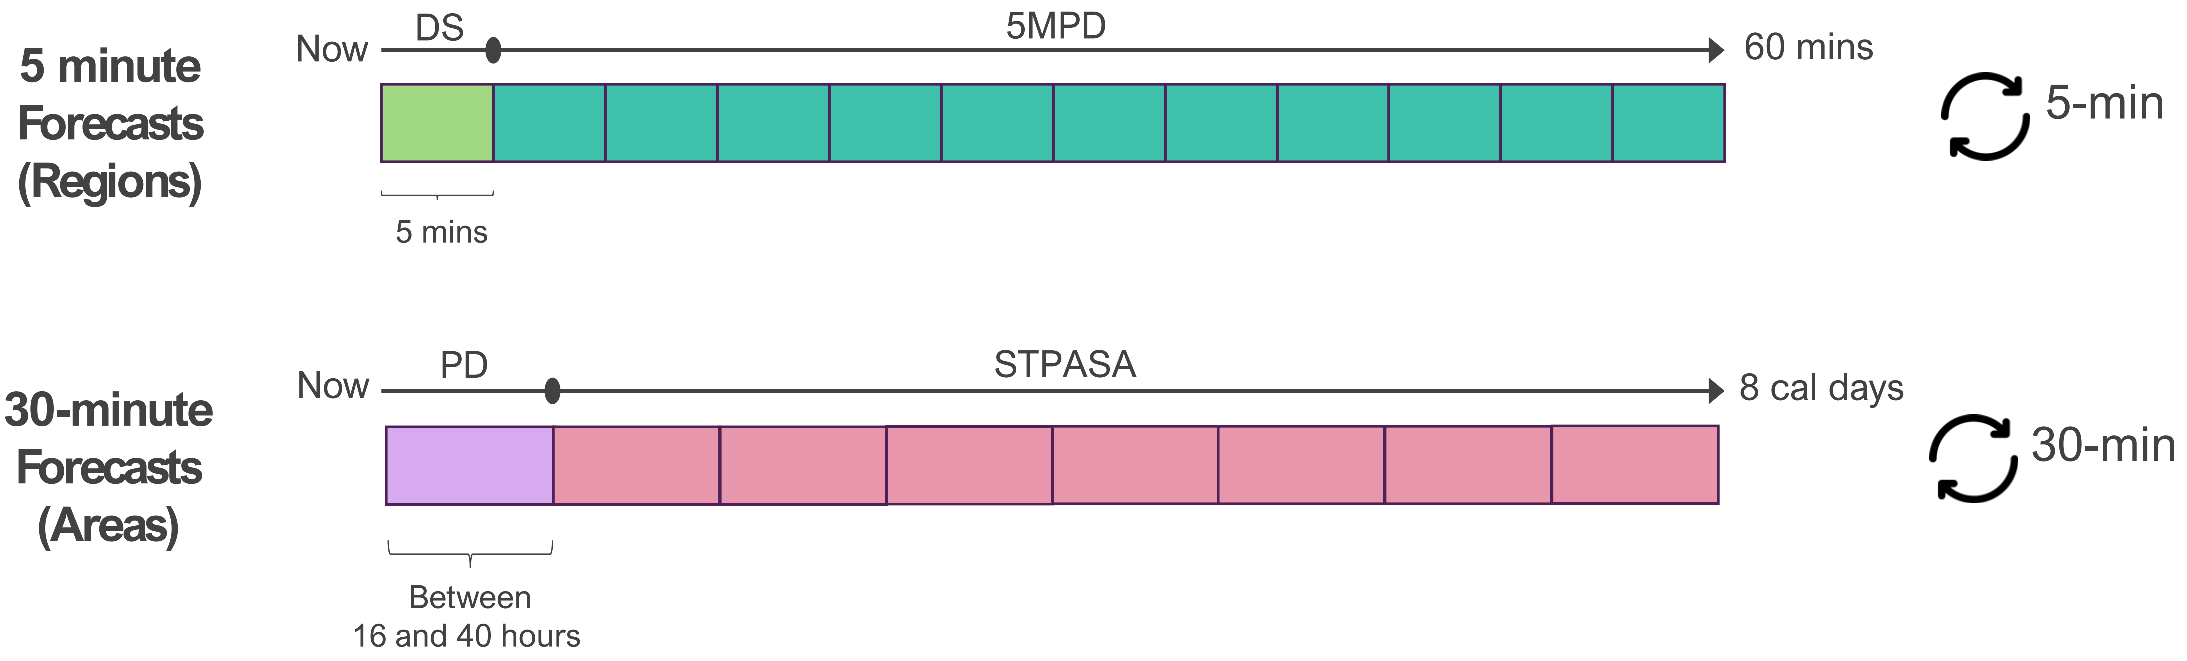

Source: [Spot Market Operations Timetable]

Demand Forecasting for the NEM

The five NEM **regions** are modelled as eight demand forecast **areas**

Modelled as Single Areas

New South Wales
(NSW)

Victoria
(VIC)

South Australia
(SA)

Modelled as Multiple Areas

Northern,
Central,
Southern
(QLD)

Northern,
Southern,
+ 4 Major Industrials
(TAS)

Demand Sources in the NEM

End-use electricity demand can be divided into three major sources; **residential, commercial** and **industrial**. Each with different sensitivities and each NEM region has a different make-up of the three. It is important to understand these differences to accurately forecast regional demand.

<table>
<tr><th>Sensitivity</th><th colspan="2">Residential</th><th>Commercial</th><th colspan="2">Industrial</th></tr>
<tr><td>Weather & DPV</td><td colspan="2">High</td><td>Moderate</td><td colspan="2">Low</td></tr>
<tr><td>Day of week</td><td colspan="2">Moderate</td><td>High</td><td colspan="2">Low</td></tr>
<tr><td>Holidays</td><td colspan="2">Moderate</td><td>High</td><td colspan="2">Low</td></tr>
<tr><td>Electricity Price</td><td>Low</td><td>Moderate</td><td>Moderate</td><td>Moderate</td><td>High</td></tr>
<tr><td>Biggest Impact</td><td colspan="2">Temperature
Rooftop PV, batteries</td><td>Public Holidays</td><td colspan="2">Maintenance
Prices</td></tr>
</table>

An Ensemble of Forecasting Information

Ensemble

- *A group of items viewed as a whole rather than individually.*
- *A group of similar systems, or different states of the same system, often considered statistically.*

Operational Forecasting

- *Multiple weather providers for energy forecasting*
- *Several scenarios of possible weather from each provider*
- *A suite of weather tools for situational awareness & monitoring*
- *A consensus of demand forecasts for assessment and validation*

Why AutoML

Enables faster and easier build and deployment of energy supply and demand models

- The Operational Forecasting Team currently deploys hundreds of machine learning (ML) models that generate forecasts in real-time. These models are maintained in-house through a time-consuming process.

- We aim to leverage AutoML to streamline our MLOps, accelerate our development cycles, and continually push the boundaries of what our forecasting capabilities can achieve.

- Moving forward, new uses cases will also require more scalable solutions.