Marshall Plan Scholarship

Final Report

Nikolaus Houben

September 10, 2023

Contents

1 Executive Summary

In an era where climate change looms large, the transition towards renewable energy sources is more than just a trend—it's a necessity. Central to this is the electricity grid, a foundational pillar of modern society, that must adapt and evolve to handle the dynamics of renewable integration, decentralization, and the increasing demand. As the world shifts, there is an opportunity to harness technological advancements, particularly the capabilities of artificial intelligence (AI), to navigate the complex terrains of the modern energy landscape.

The spotlight in recent advancements is energy forecasting, an arena where AI, and in particular, supervised learning methods, have shown immense potential. The intricacy of predicting energy demand and supply is accentuated with the integration of variable renewable sources, and this is where advanced AI techniques have made groundbreaking progress. The projects undertaken during the research visit to the Berkeley Lab, made possible by the support of the Marshall Plan Foundation, delved deeply into this domain, focusing on multi-scale load forecasting and net load forecasting.

The multi-scale load forecasting project aimed to understand and predict energy demand at different scales—from the expansive transmission grids spanning vast regions to the individual nuances of building-level consumption. Recognizing that each scale presents unique challenges, the project's approach integrated granular data sources with advanced machine learning models. This holistic view ensures that forecasting remains consistent and reliable across various levels, from high-voltage transmissions to localized distribution networks.

The net load forecasting project was initiated to address the complexities arising from the increasing integration of renewables. With solar and wind energy's inherent variability, predicting the net load—the difference between forecasted load and renewable generation—becomes pivotal. This endeavor sought to develop models that can cater to this nuanced requirement, ensuring that the grid remains balanced and resilient even with the fluctuating nature of renewable energy sources.

However, while the potential of AI in energy forecasting is immense, the path is laden with challenges. Data accessibility remains a pressing issue, with the need for high-fidelity, consistent datasets often clashing with proprietary and security concerns. Additionally, ensuring true generalization across different energy time series, given the myriad of local influencing factors, is a task that demands further innovation.

In summing up, the global energy landscape is at a crossroads, with AI set to redefine its trajectory. The projects undertaken at the Berkeley Lab signify an important step in this direction, addressing key challenges and harnessing technological advancements to pave the way for a more sustainable, efficient, and resilient energy future. This report encapsulates these endeavors, offering insights into the progress made and setting the stage for what lies ahead.

2 Introduction

2.1 Electrification to Mitigate Climate Change

In our epoch, termed the Anthropocene by many, humans have become a driving force in shaping the Earth's environment. The urgency to mitigate climate change, a colossal challenge resulting from human activities, primarily the emission of greenhouse gases, has been widely recognized. Governments, industries, and societies worldwide are engaged in a multidimensional effort to mitigate these effects, with the transition to sustainable energy at the center of these endeavors. At this nexus of humanity's response to climate change lies the evolution and restructuring of electricity grids to support an increasingly electrified world.

2.2 Electrification and the Evolution of Grids

The shift towards sustainable energy sources like wind and solar, in response to the climate crisis, mandates a transformation of our electricity grids. These inherently variable and decentralized sources of energy introduce unprecedented challenges in grid management. The traditional electricity grids, designed for centralized and predictable energy generation, are now being re-envisioned to accommodate a new era of energy. The integration of renewable energy, storage technologies, and the looming prospect of electrified transport has exponentially increased the complexity of grid operation. It's within this evolving tapestry of energy systems that the intricate dance of supply and demand unfolds. Electrification, while essential in our fight against climate change, places an enormous responsibility on ensuring stability, reliability, and efficiency in our grids.

2.3 Multi-Scale Nature of Energy Systems

Furthermore, the energy system exhibits a multi-scale nature, ranging from the high-voltage transmission grid to regional distribution networks, down to individual buildings and homes. Each level presents its unique challenges and opportunities, and the pursuit of an optimized energy system necessitates understanding and harmonizing these various scales. In this intricate setup, forecasting — especially energy forecasting — becomes a cornerstone of reliable grid operation. Recognizing patterns, anticipating demand and supply, and making informed decisions are pivotal in an environment where renewable energy, with its inherent variability, is playing an ever-increasing role.

2.4 Digitization and the Advent of AI in Grid Operation

The advent of digitization has brought with it a trove of data from smart meters, grid sensors, and a multitude of IoT devices. This rich data landscape is both a challenge and an opportunity. The vast amount of data can be overwhelming, but with the right tools, it offers a chance to revolutionize grid operations. Artificial

Intelligence (AI), and in particular, (semi-)supervised learning, have emerged as potent tools to harness this data for actionable insights. Their ability to capture long-term dependencies and intricate patterns has particular relevance in energy forecasting. However, these algorithms have specific challenges, such as high data requirements and accessibility concerns. Moreover, true generalization between time series for forecasting and control remains a central problem.

In the ensuing pages of this report cover the scientific advancements achieved during my research visit at the Berkeley Lab. This transformative endeavor, made possible by the generous sponsorship of the Marshall Plan Foundation, has allowed me to deepen my knowledge of AI-driven methods for energy forecasting. These methods, while heralding great potential, come replete with their unique set of challenges. It is by navigating some of these complexities that we have contributed to lay the groundwork for shaping a more sustainable and reliable energy future. Through these accomplishments, I hope to demonstrate the invaluable impact of the Foundation's support.

3 Research Background

The rapid integration of artificial intelligence (AI) into various sectors has revolutionized the way we perceive and handle data, leading to substantial advancements in predictive accuracy and operational efficiency. One of the sectors most impacted by these technological evolutions is energy systems, where the infusion of AI techniques has the potential to reshape traditional paradigms.

3.1 Artificial Intelligence in Energy Systems

Machine learning, at its core, is akin to teaching computers to learn from experience. Imagine trying to teach a child to recognize fruits by showing them pictures of apples, bananas, and cherries while naming each one. The more pictures and names you show, the better the child becomes at identifying each fruit on their own. Similarly, machine learning involves feeding the computer, or more precisely, a model, with a plethora of data and letting it discern patterns within this data.

The past few years have witnessed the meteoric rise of large language models within the realm of machine learning. These models, built on vast datasets of textual information, are not just about understanding language but are exemplars of how intricate patterns can be recognized and utilized by machine algorithms. Just as the human brain can grasp and generate complex language structures from exposure to linguistic inputs, these models leverage billions of sentences to generate, understand, and even transform human-like text. This evolution towards more sophisticated and data-intensive models underpins the broader shift in machine learning, where capacity, complexity, and capability are continually being expanded.

One primary method within machine learning is supervised learning. Much like our fruit-naming exercise, supervised learning involves providing the model with both the input (e.g., historical weather data) and the correct output (e.g., past electricity consumption). By repeatedly exposing the model to various inputs and their corresponding outputs, it learns to associate patterns in the input data with the correct output. This setup, where past data guides the prediction of future outcomes, is particularly suited for time series forecasting, such as predicting future energy demand based on past consumption and conditions. It's within this context of supervised learning that the energy sector has witnessed transformative shifts.

Energy systems, being the backbone of modern society, have seen a shift with the incorporation of renewable energy sources, decentralization, and digitalization. These changes have made the system increasingly complex, emphasizing the need for sophisticated techniques to manage and optimize its operations. AI offers unparalleled precision in forecasting, real-time monitoring, anomaly detection, and thereby control of the grid. Among the various algorithms introduced to the energy sector, transformers, originally designed for natural language processing tasks, have demonstrated their utility in understanding temporal dependencies in energy time series, making them particularly suitable for forecasting.

3.2 Challenges and the Nuances of Data

However, the journey to incorporate AI in energy systems encounters its share of roadblocks. Foremost among these is the issue of data accessibility. Owing to the proprietary nature of energy data and understandable security reservations, procuring the requisite datasets for modeling is often an uphill battle. Furthermore, the very nature of the energy domain, with its emphasis on precision and reliability, demands data of impeccable quality and consistency.

Compounding these challenges is the quest for genuine generalization across time series. Achieving accurate energy forecasting isn't simply about predicting unseen future data points but also about ensuring that these predictions hold water across diverse regions, grids, and scenarios. This becomes particularly daunting when considering local variables like unique weather conditions, varying consumption behaviors, and the specificities of regional infrastructures.

3.3 The Multi-scale Nature of Energy Systems

Zooming out, our energy landscape reveals its multi-scale essence, cascading from vast transmission grids to singular buildings:

- **Transmission Grid**: Here, the game is all about seamless electricity conveyance across large distances, with an emphasis on anticipating demand and supply shifts, especially considering the unpredictable nature of renewables.

- **Distribution Grid**: At this juncture, the lens sharpens to address local demand-supply equilibriums, assimilate decentralized energy contributions, and maintain voltage stability.
- **Building Level**: On this micro-scale, the goal pivots to refining consumption behaviors, weaving in smart technologies, and potentially feeding surplus energy back into the overarching grid.

Such a multi-tiered structure underlines the indispensability of an encompassing AI strategy. It underscores the need for a symphony of operations, harmonizing decisions made on a macro scale with those on the micro.

3.4 The Significance of the Upcoming Projects

With a comprehensive understanding of the energy systems' challenges and intricacies, the impending projects encapsulated in this report gain heightened relevance. By delving into forecasting nuances, harnessing the power of AI paradigms like transformers, and unraveling the multi-scale dynamics, these projects aim to contribute to the path forward for the energy transition. Their overarching goal is not merely technological elevation but translating these advancements into user-friendly code repositories, distributing tools to improve the resilience and efficiency of our future energy systems.
In culmination, while the energy sector grapples with the monumental task of addressing climate change, AI stands ready to offer its toolkit, promising innovations that could shape our energy future. In the following project descriptions, we will delve deeper into the practical applications and innovations developed during my research visit at the Berkeley Lab.

4 Projects Undertaken

The following projects were undertaken during the course of my study visit at the Berkeley Lab:

1. **Multi-Scale Electricity Load Forecasting**
   - **Description:** This project aimed to deliver a comprehensive comparison of accessible machine learning algorithms, systematically evaluated across a multitude of electricity load time series, and spanning diverse levels of aggregation.
   - **Outcome:** Journal Paper (expected submission October, 2023)
   - **Collaborators:** Han Li, Miguel Heleno, Tianzhen Hong
2. **Net Load Forecasting**
   - **Description:** Accurate net load forecasting relies on the availability of high-quality data from various sources, including weather forecasts, historical and/or real-time data from smart meters, and PV system specification data. The challenge of data collection can make it expensive or even impossible, making it imperative for grid operators and energy retailers to utilize the data effi-

ciently. This project presents a method for day-ahead net load forecasting that is robust to varying data availability.

- **Outcome:** IEW Presentation and Conference Paper (presented June 16th, 2023)
- **Collaborators:** Lennard Visser, Wilfried van Sark

In the subsequent sections, we will dive into the motivation, methods, and results associated with each project, offering a comprehensive understanding of the undertakings and their significance in the greater context of the energy transition.

4.1 Multi-Scale Electricity Load Forecasting (Main Project)

In the early stages of this project, the intricacies of the load forecasting landscape really struck us. The electricity grid, dynamic and continually evolving, had grown not just in size but also in complexity. With the constant integration of countless new assets into the grid, from large-scale power plants to small decentralized generators, it became clear that predicting electricity usage wasn't just a task—it was a grand challenge. This myriad of interconnections and dependencies posed numerous questions and demanded intricate solutions. Among them were:

1. How can we efficiently manage the *peak load*, especially when it varies across different temporal scales?
2. Are the existing forecasting methods myopic, focusing too much on specific spatial scales?
3. What is the ideal algorithmic approach for a sprawling metropolitan city compared to a more sporadic rural landscape?
4. How does the temporal granularity of data (hourly vs. minute-wise) affect the choice of forecasting models?
5. With the ever-growing integration of renewables, how can forecasting be tuned to adapt to this intermittency, especially across different scales?

As we delved into existing studies, another realization came to the forefront: while there were numerous research efforts centered around load forecasting, many were tunnel-visioned, concentrating either on macro or micro-levels, but seldom bridging the two. The lack of a holistic view across different spatial scales, from large grids covering entire regions to individual homes and buildings, became evident. There was an essential need for a comprehensive approach—one that could seamlessly span from the vast to the granular while retaining a high degree of forecasting accuracy. And thus, with these gaps and challenges as a backdrop, the project was born out of a pursuit to offer solutions that were both encompassing and tailored to each scale's unique requirements.

4.1.1 Setting the Research Objective

The initial phase of any research endeavor is pivotal, setting the trajectory for what is to come. As we delved into the preliminary research, the contours of the problem space became sharper. The path was clear: We wanted to bridge the evident gap in the literature. It was not merely about plugging a hole but rather about charting new territories in the field of load forecasting. The project's mission was to comprehensively compare available machine learning algorithms, assessing them across multiple electricity load time series and aggregation levels. We also sought to understand how each algorithm performs under varying conditions and challenges, from data scarcity to noise and from rapid demand fluctuations to seasonal patterns. Our aim was not only to identify the top performers but to delve into the underlying reasons behind their success or lack thereof.

4.1.2 Understanding the Literature

Diving deep into the existing literature is akin to exploring a vast forest; one encounters familiar territories, but there are always uncharted lands waiting to be discovered. As we ventured into the existing body of knowledge, two categories stood out: spatial (aggregation level) and temporal (forecast horizon) aspects of electricity load forecasting. The richness and depth of existing studies were commendable. Yet, upon closer examination, it was evident that many researches had focused on one of the two. While many studies offered insights for specific aggregation levels, a holistic cross-scale evaluation was largely missing. Moreover, fewer ventured into the economic implications of forecast accuracies. We realized the importance of an overarching framework that could collate these disparate pieces of research into a coherent narrative, shedding light on not just the "what" but also the "why" and "how" behind various forecasting successes and challenges.

4.1.3 Original Contributions

Through this research, we aimed to bring something novel to the table:

- A thorough analysis of popular machine learning forecasting techniques across 15 distinct electricity load time series.
- A comparison that juxtaposed tree-based and neural network-based forecasting methods, taking into consideration seasonal variations.
- The introduction of the *neat load error* metric, an innovative lens to discern the tangible effects of forecast inaccuracies.

4.1.4 Methodological Insights

Throughout our work, we employed a range of algorithms that not only served our immediate forecasting objectives but also epitomized significant milestones in the history of supervised learning for time series forecasting. As we journeyed through this domain, it became evident how the progression and sophistication of these tools have reshaped the landscape of predictive analytics. Each algorithm we used offers a unique perspective, highlighting the importance of leveraging diverse techniques to stay at the forefront of this evolving field.

- **Linear Regression**
  - A linear approach to establish the relationship between dependent and independent variables.
- **Random Forest [Bre01]**
  - An ensemble of decision trees, used for classification and regression.
- **XGBoost [CG16]**
  - An optimized gradient boosting library.
- **LightGBM [Ke+17]**
  - A gradient boosting framework that uses tree-based algorithms.
- **Gated Recurrent Unit (GRU) [Chu+14]**
  - A type of recurrent neural network that is especially effective for sequences, such as time series data.
- **N-BEATS [oreshkin_n-beats_2019]**
  - A deep learning model that provides interpretable time series forecasts using neural basis expansion.
- **Temporal Fusion Transformer by [Lim+21]**
  - An attention-based deep learning model for time series forecasting.

Machine learning libraries, such as scikit-learn and pytorch, have democratized access to a plethora of sophisticated forecasting tools. From the user-friendly interfaces to their comprehensive documentation, these libraries have become indispensable to both neophytes and experts in the field. The diversity of available algorithms, ranging from traditional techniques like Linear Regression to the cutting-edge ones like the Temporal Fusion Transformer, is indicative of the depth and breadth of methods at our disposal. However, this vastness can be overwhelming, which is why it is paramount to understand the underlying mechanics and principles of each method. By applying these methodologies, we aim to sift through the noise and highlight the distinctive attributes and advantages of each, positioning them within the broader forecasting landscape.

4.1.5 Measuring Success: The Net Load Error (NLE) Metric

Forecasting success in the energy domain demands a blend of theoretical accuracy and practical applicability. The traditional metrics such as RMSE, MAE, and MBE, offer a quantitative means to gauge forecast accuracy.

However, the real challenge is understanding how these forecast errors translate to tangible operational costs in a realistic scenario.

The *Net Load Error* (NLE) was conceived to address this challenge. At its core, NLE seeks to answer a pivotal question: Given perfect market signal foresight, what additional costs will arise from the inaccuracies in a load forecast? In simpler terms, how do the operational costs differ when utilizing forecasts versus the actual ground truth for battery-operated electrical systems?

Calculating the Net Load Error (NLE):

The calculation of NLE is structured as follows:

Scaling the Time Series The real-time data, or ground truth, undergoes scaling using the MinMax formulation. This scaling ensures compatibility with the BESS parameters, specifically its capacity and C-rate.

Price Signals Formulation Price signals are constructed by perturbing the lagging scaled ground truth time series with a noise term drawn from a Gaussian distribution with mean = 0.0 and standard deviation = 0.1.

Operational Costs Computation Utilizing the Model Predictive Control (MPC) setup, the system calculates optimal setpoints for the BESS's state-of-charge at each timestep. This approach modulates the actual load profile, leading to a residual load that is subsequently evaluated against the cost structure employed during optimization.

Cost Function The crux of the setpoint calculation at each timestep hinges on a cost function. This function is divided into two components: The first, a two-segment time-dependent rate, assesses the system's ability to adapt to energy price fluctuations. The inherent tier structure captures the repercussions of forecast errors, especially during peak load times. The second component introduces a monthly demand charge, linked to the highest observed peak in that month. This evaluates the system's prowess in forecasting and subsequently curtailing peak load.

System Constraints Various system constraints ensure the model's robustness and practicality. These include energy balance equations, battery state-of-charge bounds, and power charge/discharge limits. They play a vital role in ensuring that the operational scenario simulated by the MPC mirrors real-world battery-operated electrical systems.

Parameterization: A fair assessment mandates that the operational costs derived are consistent across varying load time series. To achieve this, a systematic parameterization of prices and BESS based on the load time series is crucial. Furthermore, all load time series (including forecasts and ground truth) are normalized within a range of [0, 1].

In summary, the *Net Load Error* metric bridges the gap between abstract forecast errors and tangible operational consequences. The heart of NLE lies in this comparison: The difference in operational costs between the sub-optimal residual load (stemming from forecast-based decisions) and the optimal one (informed by the ground truth). By offering a holistic view of the repercussions of forecasting inaccuracies in a real-world context, NLE emerges as a key metric, ensuring that our assessment remains grounded in both theory and practice.

4.1.6 Preliminary Results

As we began collecting, exploring and cleaning our datasets, and subjecting them to various forecasting models, intriguing patterns emerged. A consistent observation was the stellar performance of tree-based methods, namely Random Forest, XGBoost, and LightGBM. These methods, backed by robust statistical algorithms, frequently overshadowed their neural network counterparts in terms of conventional statistical error metrics. The unbiased nature of errors produced by these tree-based models renders them invaluable across myriad forecasting horizons and scenarios.

Conversely, when we shifted our gaze towards neural networks, the picture became more nuanced. While they occasionally lagged behind on some traditional metrics, they showcased commendable prowess when assessed using the more innovative, application-driven error metric, NLE. This divergence underscores the pivotal role of context. Tree-based models, with their robustness, are unparalleled for general forecasting. However, when the spotlight is on application-centric scenarios, where nuanced intricacies matter, neural networks ascend to the fore.

4.1.7 Conclusions

Forecasting, a domain that marries the age-old art of prediction with modern computation, has witnessed substantial growth and diversification, as evidenced by our exploration of various methodologies and metrics. As we navigate through the myriad of available algorithms, certain patterns emerge that offer valuable insights into the application and efficacy of these techniques.

At the heart of our findings is the notable prominence of tree-based methods, which include Random Forest, XGBoost, and LightGBM. Their inherent simplicity, combined with remarkable efficiency, sets them apart as reliable workhorses in the forecasting arena. The nature of their construction, rooted in the division of data into decision trees, allows for an intuitive understanding and interpretation of results. Their unbiased error distribution further endears them to researchers and practitioners, making them an optimal choice for those who prioritize accuracy using traditional statistical error metrics.

However, the narrative takes a nuanced turn as we delve deeper into the intricacies of application-specific requirements. This is where neural networks, specifically models like the Temporal Fusion Transformer, come into play. While they demand greater computational resources and often involve more intricate setups, their flexibility in capturing non-linear relationships and their adaptability to various data structures makes them indispensable. When assessed using our innovative *Net Load Error* (NLE) metric, which focuses on the practical implications of forecasting errors, neural networks often showcased superior performance. Their ability to gauge and account for intricate patterns in data, which might elude more traditional methods, emphasizes their value in scenarios where the context and nuanced application specifics are paramount.

In conclusion, the journey of forecasting is not a one-size-fits-all endeavor. The choice of methodology hinges significantly on the end goals. For those prioritizing computational efficiency and accuracy based on traditional metrics, tree-based methods stand out as the stalwarts. Conversely, when the scenario demands a deeper understanding of data intricacies and the consideration of application-specific metrics, neural networks rise to the challenge, bridging the gap between raw computational forecasting and real-world application nuances. As we continue to advance in this domain, it's essential to remain adaptable, understanding that the optimal forecasting method is often a dynamic choice, contingent upon the evolving requirements and constraints of each unique situation.

4.1.8 Next Steps and Key Learnings

The analysis has provided valuable insights, but further work is needed. The next stages will focus on a detailed evaluation using the *Net Load Error* metric, aiming to get a clearer picture of the intricacies of forecasting performance. Following this, the final draft of the journal article will be composed, encapsulating all the findings.

From this project, we learned:

- It is essential to choose the method that aligns best with the specific forecasting challenge.
- Comprehensive error analysis can reveal more than just raw performance metrics.
- Continuous research and re-evaluation are crucial to achieve the best forecasting results.

4.1.9 Limitations and Reflection on the Project

Embarking on this multiscale forecasting project presented both a series of challenges and rewarding moments of clarity. Foremost, working across 15 distinct electricity load time series brought with it a challenge in data harmonization. Each dataset, while individually consistent, posed its own unique intricacies and patterns. Understanding and pre-processing this vast and varied dataset was the first hurdle we encountered.

The decision to juxtapose tree-based methods with neural networks was not without its trials. While both methodologies have their merits, ensuring consistency in the evaluation framework for both presented technical intricacies. To facilitate a fair comparison, particularly across seasonal variations, required meticulous data handling and model tuning. Often, this involved returning to the drawing board to tweak our methodology. However, one of the most challenging aspects was crafting a training and evaluation framework robust enough to handle inhomogeneous datasets. We felt the weight of responsibility, given that our findings could influence future research directions. The triumph, though, was the integration of Weights & Biases, which allowed us to develop an end-to-end solution. The process was demanding, but the results were clear: a robust framework that didn't just serve this project but became a beacon for subsequent projects. Today, the pride is evident, as our solution (see Figure 1) has been adopted for diverse projects at the Berkeley Lab, testament to the enduring value of our endeavor.

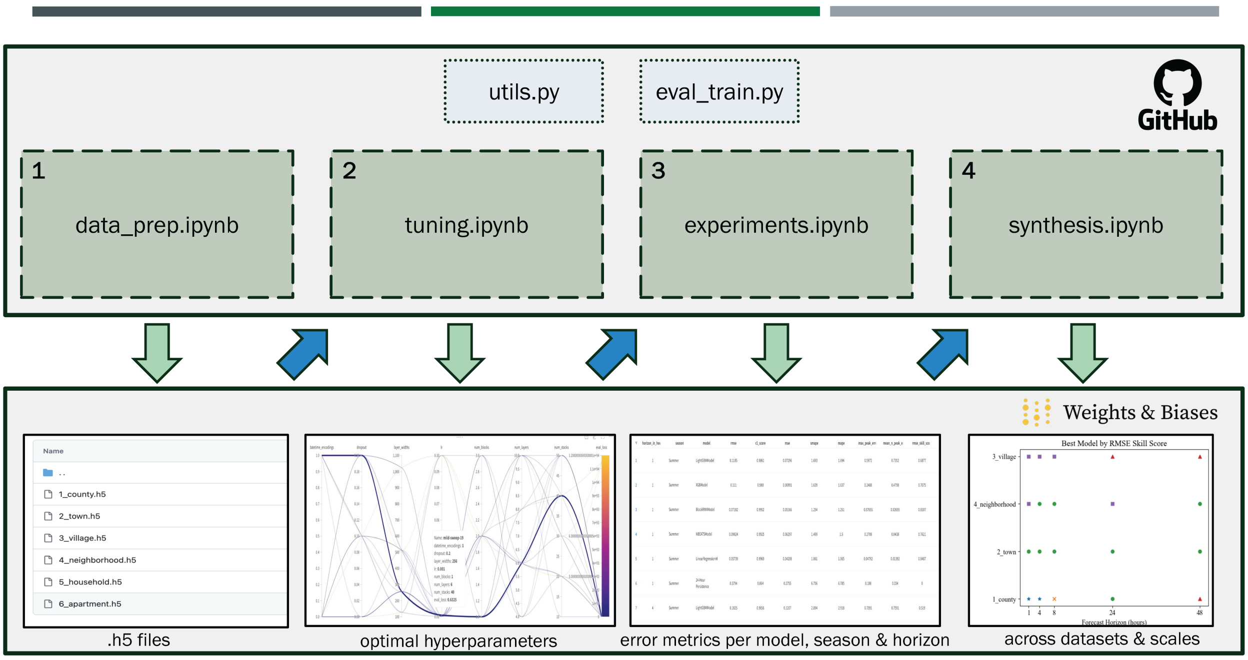

Figure 1: Wattcast Repository Structure: End-to-End Forecasting Pipeline for Model Comparison

Yet, no research journey is devoid of introspection. The introduction of the *neat load error* metric was a pivotal moment, one that shifted our perspective on forecast inaccuracies. It reminded us that while numerical accuracy is vital, the real-world implications of our models, the tangible effects, are of utmost importance. This metric has now set a precedent, pushing us and others in the field to consider the broader implications of forecasting errors.

In hindsight, while the path was strewn with challenges, they were the catalysts for innovation. They pushed us to think laterally, seek novel solutions, and in the process, contribute meaningful advancements to the realm of multiscale forecasting. The journey underscored the importance of resilience, collaboration, and the pursuit of real-world applicability in research.

4.2 Net Load Forecasting in the Age of Decentralized Energy Resources

4.2.1 Setting the Research Objective

The surge in the adoption of roof-top solar photovoltaics (PV) signifies a monumental shift in our energy paradigm. With their widespread acceptance, grid operators and energy retailers are posed with the challenge of effectively integrating them. Historically, the electricity grid was designed with centralized generation in mind, where large power plants produced electricity and transmitted it over long distances. However, the growth of DERs introduces a two-way flow of electricity, challenging the conventional operation of power networks.

Herein lies the prominence of net load, a metric representing the electrical load after accounting for generation from distributed energy resources (DERs). This metric, often dubbed the "duck curve" due to its characteristic shape in some regions, presents new operational challenges such as over-generation during solar peaks and steep ramps in the evening when the sun sets and electricity demand rises. Addressing these challenges necessitates accurate forecasting.

The aim of this research is to present a day-ahead net load forecasting technique that is not only accurate but also efficient in using varying availability of data. We believe that by offering a robust solution to predict these fluctuations, we can assist grid operators in maintaining system reliability while maximizing the benefits of renewable integrations.

4.2.2 Understanding the Literature

While the world transitions towards decentralized energy sources, there remains ambiguity regarding the ideal method for day-ahead net load forecasting. The dynamic nature of renewable generation, influenced by a myriad of factors ranging from weather conditions to local shading, makes predicting the net load a complex endeavor.

The current literature is deficient in discussions about timesteps below the 5-minute interval, which are crucial in capturing the rapid fluctuations, especially with increasing penetration of renewables. These shorter intervals are vital for efficient grid operations, ensuring that supply and demand are balanced in real-time. Furthermore, there's a noticeable gap in in-depth explorations on the effects of data availability scenarios. Many regions, while witnessing an influx of DERs, may not have extensive historical datasets, making it crucial to devise techniques that perform well under data constraints.

Additionally, the existing literature often gravitates towards particular methods, leaving a wide swath of algorithms and techniques underexplored. This lack of comprehensive exploration curtails the potential for discovery of novel or more efficient approaches tailored to specific challenges associated with decentralized energy resource integration.

4.2.3 Original Contributions

This research goes beyond traditional approaches to fill the gaps left by the existing body of work. Our contributions include:

1. A groundbreaking multi-step ahead net load forecasting method, bridging a physical PV model and a transformer neural network.
2. The novel introduction of time series patching for improved data efficiency with high-frequency smart meter data.
3. A holistic comparison of additive methods against their direct and integrated counterparts.
4. A comprehensive study on the impact of data availability on forecasting, with special emphasis on meter and PV system metadata.
5. Translating error analyses into pragmatic recommendations for energy retailers.

4.2.4 Methodological Insights

Central to this research is an indirect, additive net load forecasting technique. This method blends the strengths of a physical PV model with the capabilities of a transformer neural network. The dual approach enables indirect forecasting by adding the outputs of a data-driven machine learning model to that of a physical model. A cornerstone of this strategy is the physical model, which forecasts for individual PV systems using specific metadata, encompassing factors like system tilt, azimuth angle, capacity, and more. This multi-faceted approach is further elucidated in Figure 2.

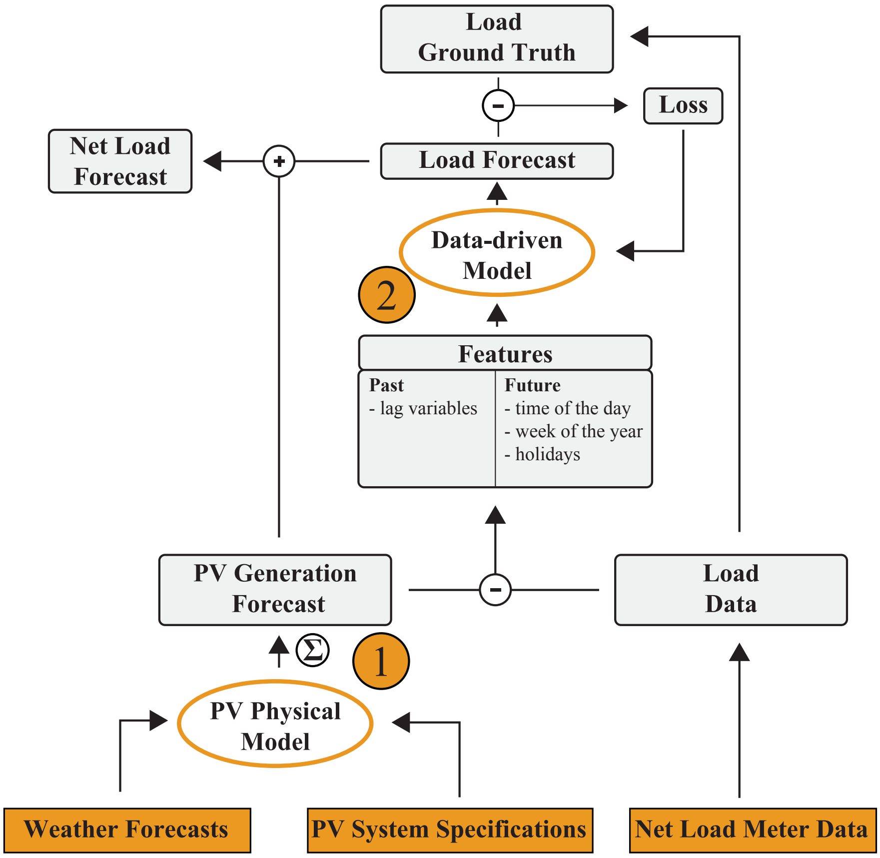

Figure 2: Model Framework – Bottom (1): Weather forecasts (irradiance) and PV specification data are used as inputs to the physical PV model to produce a PV generation forecast per system. Top (2): The PV generation forecast is aggregated and subtracted from the net load meter data. Features for the past and future horizon are extracted from the load data and a data-driven model is trained by back-propagating the loss through the neural network. The loss is calculated as the *mean squared error* of the forecast and the ground truth of the measurements. Once trained the load forecasts are added to the corresponding PV Generation Forecast to retrieve the Net Load Forecast

4.2.5 Measuring Success

The essence of our research lies not only in developing a robust forecasting methodology but also in rigorously assessing its viability and effectiveness. The ultimate goal of our forecasting technique is its practical application, ensuring a more balanced and cost-effective integration of distributed energy resources into the grid. To this end, we established a multifaceted framework for evaluation that bridges theoretical accuracy with tangible real-world impacts.

Central to our assessment strategy is the calculation of imbalance costs, particularly in the context of the Dutch energy system. The imbalance costs give us an insight into the potential financial ramifications of forecasting inaccuracies. In essence, when energy is either over-supplied or under-supplied relative to the forecasted demand, imbalance costs are incurred. These discrepancies between the forecasted and actual net loads can result in significant financial penalties for grid operators, reinforcing the importance of an accurate forecasting mechanism.

In addition to the direct financial metrics, it was crucial to understand the statistical robustness of our model. Traditional statistical error metrics, such as the Root Mean Square Error (RMSE) and Mean Square Error (MSE), provide an analytical lens into the method's accuracy. While these metrics are conventionally used to gauge the deviation of predicted values from the actual measurements, their relevance in our study extends further. By analyzing these metrics in conjunction with the imbalance costs, we aimed to establish a direct link between theoretical accuracy and its financial implications. Furthermore, to ascertain the broader implications of our methodology, we compared these metrics against established benchmarks and state-of-the-art methods in the realm of net load forecasting.

4.2.6 Results

Embarking on a rigorous assessment of our proposed methodology yielded interesting outcomes, increasing our confidence in the research's robustness and practicality. Our primary metric of success, the imbalance costs within the Dutch energy system's context, provided a real-world lens through which we could discern the potential financial consequences of our forecasting mechanism.

Upon extensive evaluation, our model emerged as a frontrunner in minimizing these costs. Specifically, when juxtaposed against other prevalent forecasting techniques, our methodology exhibited a substantial reduction in the associated financial implications. This directly translated to significant potential savings for grid operators and energy retailers, showcasing the practical value of our research.

Further delving into the statistical robustness of our method, the RMSE and MSE values obtained offered a promising picture. Notably, our technique's RMSE value was consistently lower than other established forecasting methods by a margin that signifies not just statistical significance but also real-world applicability. The MSE values further reinforced this, offering a testament to the method's consistent accuracy across different scenarios and test cases.

But beyond the numbers and metrics, what truly set our methodology apart was its adaptability. The innovative combination of a physical PV model with a transformer neural network allowed for a degree of flexibility in catering to different data scenarios and grid conditions. This adaptability ensures that our model remains relevant and effective even as the dynamics of decentralized energy integration evolve.

In conclusion, the results obtained through our rigorous evaluation framework underscore the efficacy of our research. By bridging the gap between theoretical accuracy and practical financial implications, our study offers a promising avenue for the future of net load forecasting in an increasingly decentralized energy landscape.

4.2.7 Conclusions

1. Our innovative approach to net load forecasting surpasses both "integrated" and "direct" methods consistently for the day-ahead forecasting horizon.
2. Distinguishing our method from others, it remains viable and robust, even when faced with limited training data sets.
3. By implementing patching, we've been able to successfully incorporate high-resolution data. This inclusion not only boosts performance but also enables the forecasting model to consider longer context lengths.

4.2.8 Next Steps and Key Learnings

Our journey in net load forecasting, while successful, has shed light on potential areas for further exploration and refinement.

1. There's significant potential in aligning our forecasts with imbalance price expectations. Doing so could make our forecasting results even more relevant from a financial perspective.
2. Broadening our horizons by inviting more community participants could enrich our data sets and present newfound insights.
3. Employing real-time weather and irradiance forecasts will likely offer a genuine reflection of our model's applicability and accuracy in real-world scenarios. This pivot could further cement our method's reputation as a go-to solution in net load forecasting.

4.2.9 Limitations and Reflection on the Project

Embarking on this journey of advancing the frontiers of net load forecasting was no smooth sail. The initial excitement of devising a novel method was soon met with the sobering reality of research complexities and collaborative intricacies.

Data Acquisition Hurdles One of the primary challenges was data acquisition. Suitable and relevant data sets are the lifeblood of any forecasting model, and our project was no exception. Our initial surveys of available datasets turned out to be either insufficient in granularity or not entirely relevant to the specific Dutch energy context we were focusing on. Collaborating with our esteemed counterparts in Utrecht seemed promising, given their repository of data that seemed ideal for our needs. However, this wasn't without its share of hurdles. Convincing them about the potential impact of our project, and ensuring that data sharing adhered to privacy and other regulatory standards, took longer than anticipated. The process involved a maze of NDAs, data-sharing agreements, and multiple rounds of negotiations.

Crafting the Novel Method The task of formulating a unique method brought its own set of challenges. While the theoretical understanding was clear, translating it into a practical, workable model took countless iterations, brainstorming sessions, and testing. It was a humbling experience, reminding us that innovation often demands perseverance and a willingness to embrace failure. There were moments of doubt, where we questioned the feasibility of our approach, especially when initial results were not as promising.

Silver Linings However, every challenge we faced turned into a learning experience. The difficulties in data acquisition underscored the importance of collaboration and the value of open data sharing in the research community. Crafting our method, despite its challenges, strengthened our resolve and honed our problem-solving skills. It reiterated the age-old adage – 'Rome wasn't built in a day.' Research, especially one aiming to push the boundaries, is a marathon, not a sprint.
In conclusion, the challenges we faced, while taxing, enriched our research journey, instilling in us lessons that extend beyond just this project.

5 Personal and Professional Development

Throughout my academic trajectory, particularly during my time at TU Wien and my transformative journey to Berkeley, my professional and personal boundaries expanded in ways we hadn't previously envisioned. The challenges faced, experiences gathered, and connections formed laid the foundation for this exponential growth.
From a professional perspective, the two major projects—multiscale forecasting and net load forecasting—have been monumental milestones. They underscored the importance of resilience in the face of challenges. The multiscale forecasting project, for instance, reminded me of the intricacies of handling diverse datasets, juxtaposing various machine learning methodologies, and thinking critically about real-world impacts. Delving into the depths of neural networks, tree-based methods, and crafting metrics like the *neat load error* enhanced

my skill set and fortified my expertise in energy informatics.

The subsequent net load forecasting project was equally, if not more, enlightening. Beyond the technical challenges, this project introduced me to the complexities of collaboration. Negotiating data access, understanding contractual intricacies like NDAs, and coordinating with partners in Utrecht were lessons in patience, persistence, and diplomacy. These experiences greatly expanded my ability to navigate real-world research hurdles.

Yet, it wasn't just the research complexities that facilitated growth. Being at the heart of academic excellence in Berkeley and engaging with leading minds pushed me to consistently elevate my standards. Mentorship from figures like Dr. Tianzhen Hong, collaborations with peers like Han Li and Miguel Heleno, and the camaraderie of fellow scholars all played pivotal roles in shaping my professional journey.

On a personal level, the journey was equally transformative. Living in Berkeley, a melting pot of cultures and ideas, reshaped my perspectives. Interacting with fellow scholars, exploring the city, and forging bonds with visitors provided a depth to my personal development. These experiences broadened my horizons, teaching me the importance of openness, adaptability, and embracing diversity. The invitation to Han Li's barbecue in Oakland, for instance, wasn't just about food; it was a moment of cultural immersion and deepening friendships.

Reflecting upon these experiences, it's evident that my time in Berkeley was much more than just an academic endeavor. It was a journey of self-discovery, professional refinement, and building enduring relationships. While the academic accolades and successful projects stand as testament to the professional growth, the personal narratives, cultural experiences, and lasting connections are treasures that have enriched my life in immeasurable ways.

6 Recommendations for Future Scholars

Pursuing research at an esteemed institution like Berkeley Lab, particularly within the realm of energy informatics, presents a plethora of opportunities and challenges. Drawing from my experiences and the insights from our discussed projects, we offer the following recommendations for aspiring scholars:

1. **Embrace Collaboration Early On**: The energy informatics domain is expansive and multidisciplinary. Engaging with experts, as we did with Han Li and Miguel Heleno, can open doors to fresh perspectives and methodologies. This interdisciplinary approach was pivotal in our multiscale forecasting project, highlighting the importance of diverse knowledge bases.
2. **Data Accessibility and Integrity**: One of the critical hurdles we faced was securing suitable data, especially from collaborators in different regions. Always be prepared to navigate administrative challenges, including understanding NDAs. Furthermore, ensure the integrity and quality of your data, as

it forms the bedrock of machine learning applications in energy systems.

3. **Adopt Flexible Frameworks**: Our multiscale forecasting project underlined the importance of a flexible training and evaluation framework, especially when dealing with inhomogeneous datasets. Platforms like weights and biases can provide an end-to-end solution, streamlining the research process.
4. **Stay Updated and Innovate**: The field of machine learning is ever-evolving. During our net load forecasting project, we realized the value of creating novel methods and metrics (like the *neat load error*). Stay updated with the latest advancements and don't hesitate to introduce innovations tailored to energy system challenges.
5. **Reflect on the Real-world Impact**: Beyond technical accuracy, understand the tangible impacts of forecast inaccuracies, especially in the energy sector. Tools like the *neat load error* metric can offer insights into the real-world ramifications of your findings.
6. **Plan for the Future**: As highlighted in our net load forecasting project, always keep an eye on future enhancements. Whether it's optimizing forecasts for imbalance price expectations or integrating real weather data, foreseeing future steps can guide the direction of your current work.
7. **Immerse in the Cultural and Academic Fabric**: Berkeley's vibrant academic and cultural landscape is a treasure trove for personal and professional growth. Engage in intellectual dialogues, attend conferences like the International Energy Workshop, and explore the city's rich tapestry. These experiences, beyond enriching your research, will significantly enhance your overall journey.
8. **Nurture Relationships Beyond the Lab**: Building on personal experiences, such as the barbecue at Han Li's place, always remember that the relationships you forge extend beyond professional boundaries. These connections can often lead to future collaborations, provide emotional support during challenging phases, and add depth to your overall experience.

In conclusion, the journey of research in energy informatics at Berkeley is a blend of technical challenges, innovative thinking, and nurturing relationships. By embracing collaboration, staying updated, and integrating with the vibrant Berkeley culture, future scholars can maximize their growth and contributions in this ever-evolving domain.

7 Conclusion

As the curtain falls on our research visit to Berkeley Lab, one is reminded of the dynamism and transformative potential of the global energy landscape. The urgency brought about by climate change, combined with the technological strides in artificial intelligence (AI) and machine learning (ML), underscores the vitality of our undertaking. The fusion of these disciplines, as evidenced by our endeavours in multi-scale load forecasting and net load forecasting, presents a beacon of hope and innovation in these changing times.

The multi-scale load forecasting project, which sought to fathom the intricate layers of energy demand, has yielded significant results. With an exhaustive analysis of diverse machine learning forecasting techniques across 15 distinct electricity load time series, we've shone a light on the vast expanse of this domain. Our efforts in contrasting tree-based and neural network-based forecasting methods against the backdrop of seasonal variations have offered profound insights, shaping our understanding of the field. Moreover, the introduction of the *neat load error* metric serves as a testament to our commitment to both innovation and practical relevance, providing a fresh perspective on quantifying and addressing forecast inaccuracies.

The net load forecasting project embarked on a journey to decipher the challenges entailed by the increasing integration of renewables into our grids. Our multi-step ahead net load forecasting methodology bridges a physical PV model and a transformer neural network, paving the way for more accurate and responsive predictions. Coupled with the pioneering concept of time series patching, we've not only optimized data efficiency but also set a benchmark for high-frequency smart meter data utilization. Furthermore, our holistic examination of additive methods, comprehensive studies on data availability, and actionable insights for energy retailers present a well-rounded, impactful contribution to the domain.

However, as is the case with all pioneering endeavors, the path was not devoid of challenges. Issues like data accessibility and the ever-looming quest for true generalization across different energy time series reiterate that while we have made significant strides, there is yet much ground to cover.

In retrospect, our stay at the Berkeley Lab, facilitated by the gracious support of the Marshall Plan Foundation, has been nothing short of transformative. It has not only expanded our horizons but also solidified our commitment to driving change in the energy sector through the power of AI methods.

As we conclude our time at the Berkeley Lab, it is with a spirit of gratitude, determination, and anticipation. The insights gleaned and the contributions made are but stepping stones towards a larger goal. There's a palpable sense of excitement for what lies ahead, especially as we channel our learnings into the formulation of comprehensive journal papers. This mission, while rooted in our past endeavors, is very much about charting the path forward – a future where our collective efforts contribute to a more sustainable, efficient, and resilient energy world.

In the grand tapestry of the energy transition, our projects form crucial threads. They not only validate the role of AI and ML in shaping the future but also emphasize the collaborative spirit required to navigate the complexities of the modern energy landscape. With an unwavering commitment, we look forward to furthering this work, contributing to the global dialogue, and, most importantly, making a difference.

8 Acknowledgements

Acknowledging the efforts and contributions of individuals and teams is an integral part of any research endeavor. My journey at the Berkeley Lab was made possible and incredibly enriching due to the wisdom, expertise, and guidance of many, and it's imperative that I extend my deepest gratitude to them.

First and foremost, I am profoundly grateful to Tianzhen Hong, Han Li, and Miguel Heleno from the Berkeley Lab. Their insights, unwavering support, and dedication played a pivotal role in molding my projects and research trajectory. A special mention and heartfelt thanks to Michael Stadler, who not only introduced me to the Building Technology & Urban System Division at the Berkeley Lab but also laid the foundation for this invaluable collaboration.

Back at the Energy Economics Group, my home institute, I am thankful to Hans Auer, Amela Ajanovic, and Reinhard Haas. Their continuous encouragement and endorsement were the cornerstones that strengthened my resolve and vision throughout this research journey. Their unwavering belief in my endeavors has been a constant source of inspiration.

Lastly, I cannot overlook the pivotal roles played by Diana Tsenova (TU Wien) and Claudia Kraif (Austrian Marshall Plan Foundation). Their exceptional organisational acumen ensured a seamless experience regarding the scholarship.

In conclusion, this research is not just the result of my academic pursuit but also a reflection of the collaborative spirit, guidance, and support of many. My gratitude for their contributions is great, and it is their support that has rendered this journey truly unforgettable.

References

[Bre01] Leo Breiman. “Random forests”. In: *Machine learning* 45 (2001), pp. 5–32. issn: 0885-6125. *Publisher: Springer*.

[Chu+14] Junyoung Chung et al. “Empirical evaluation of gated recurrent neural networks on sequence modeling”. In: *arXiv preprint arXiv:1412.3555* (2014).

[CG16] Tianqi Chen and Carlos Guestrin. “Xgboost: A scalable tree boosting system”. In: 2016, pp. 785–794.

[Ke+17] Guolin Ke et al. “Lightgbm: A highly efficient gradient boosting decision tree”. In: *Advances in neural information processing systems* 30 (2017).

[Lim+21] Bryan Lim et al. “Temporal fusion transformers for interpretable multi-horizon time series forecasting”. In: *International Journal of Forecasting* 37.4 (2021), pp. 1748–1764. issn: 0169-2070. *Publisher: Elsevier*.

[Ore+21] Boris N. Oreshkin et al. “N-BEATS neural network for mid-term electricity load forecasting”. In: *Applied Energy* 293 (2021), p. 116918. issn: 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2021.116918. url: https://www.sciencedirect.com/science/article/pii/S0306261921003986.